NOAA Technical Memorandum OAR GSD-55

http://doi.org/10.7289/V5/TM-OAR-GSD-55

A Survey of Selected Aviation Weather Products

February 2012

Thomas W. Schlatter

Earth System Research Laboratory
Global System Division
Boulder, Colorado
February 2012

noaa NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION / Office of Oceanic and Atmospheric Research

NOAA Technical Memorandum OAR GSD-55

A Survey of Selected Aviation Weather Products

Thomas W. Schlatter[1]

[1]Cooperative Institute for Research in Environmental Sciences (CIRES) and
NOAA/ESRL/GSD

Acknowledgements

I especially thank John C. Osborn of ESRL/GSD for comprehensive graphics support during the entire year this report was in preparation. Special thanks also to Ann M. Reiser (ESRL), who edited the entire manuscript and linked the Table of Contents with the text. Barbara Schlatter helped prepare the list of acronyms. Jeannine Schlatter checked the reference list.

Visit to the Aviation Weather Center in Kansas City, October, 2010
Thanks to Bruce Entwistle, who hosted the visit, planned the itinerary, and spent hours explaining AWC operations. I spent time on shift with Nolan Duke at the tropical desk, Jim Roets (FA East), Gregory Harris (convective SIGMETs), and Don Baker (CCFP). All these meteorologists were generous with their time and eager to share their experiences and past presentations.

Visit to the University of Wisconsin, Madison, January, 2011
Thanks to Ralph Petersen and Wayne Feltz, who coordinated the visit and planned the itinerary.
I conducted interviews with the following on the use of satellite information for diagnosis of aviation weather hazards:
Wayne Feltz introduced the GOES-R Proving Ground concept.
Justin Sieglaff, Lee Cronce, and Kris Bedka (NASA, Hampton, VA, participated remotely) discussed convective initiation.
Jason Brunner, Scott Lindstrom, and Kris Bedka (again, participating remotely from NASA) explained detection of overshooting tops.
Mike Pavolonis and Justin Sieglaff explained an algorithm for the detection of volcanic ash and several of its properties.
Mike Pavolonis and Corey Calvert presented the GOES-R algorithm for detecting fog and low clouds.
Brad Pierce discussed his research on inferring surface visibility from GOES-R ABI data.
Tim Schmidt summarized current and expected GOES satellite capabilities.
Andy Heidinger explained how MODIS data are being used to simulate GOES-R data for the detection of cloud properties especially as they relate to icing.
Tony Wimmers explained algorithms for detecting tropopause folds, gravity wave, and mountain waves—all related to turbulence.
Chris Velden gave a status report on feature tracking by satellite for the inference of atmospheric winds.
Ralph Petersen gave an update on WVSS-II (water vapor sensing system) and explained his research on the inference of convective initiation and dissipation from satellite data.

Introduction
Jennifer Mahoney (ESRL/GSD provided verification statistics and verification reports for many products. Her group's RTVS was the source of several verification figures and tables.
Stephen Walden, Weather Systems Engineering, introduced me to a large body of regulations about flying and how they relate to the use of aviation weather information.

Turbulence

Bob Sharman (NCAR) gave the initial interview on GTG2. He helped with the GTG2 turbulence section and reviewed a draft.
Rod Frehlich (NCAR) also participated in discussions about GTG2 and provided references.
John Williams (NCAR) and Michael Istok (NWS Office of Science and Technology) provided information on the NEXRAD Turbulence Detection Algorithm.
Tom Fahey (Delta Airlines) sent information on the Delta Turbulence Plot System.

Icing
Marcia Politovich (NCAR) gave the initial interview on the icing products.
Marcia, Cory Wolff, and Tressa Fowler, also of NCAR, provided references and reviewed the draft segment on CIP/FIP.
Patrick Minnis steered me to a host of references and web sites relating to NASA LaRCs work on cloud properties and their connection with icing.
William L. Smith, Jr. (NASA LaRC), and Patrick Minnis reviewed the LaRC segment on cloud Properties.
William Moninger (ESRL/GSD) provided plots of atmospheric soundings used to infer icing.

Convection
Marilyn Wolfson (MIT Lincoln Laboratory) gave several interviews on CIWS and CoSPA, was quick to exchange emails, and provided numerous formal and informal references.
Haig Iskenderian, (MIT Lincoln Laboratory), with help from Marilyn Wolfson, reviewed many pages of manuscript pertaining to CIWS and CoSPA and offered extensive, constructive comments.
Colleen Reiche (MIT Lincoln Laboratory) reviewed the segment on CoSPA.
Greg Thompson (NCAR) and John Brown (ESRL/GSD) explained the calculation of radar reflectivity from model mixing ratios.
Keith Brewster (CAPS/University of Oklahoma) gave a phone interview and provided references on merging the CoSPA extrapolation forecast with the HRRR model forecast.
James Pinto (NCAR) helped me understand how merging VIL fields from the HRRR model and the 88D radar occurs; he reviewed this portion of the CoSPA writeup.
Steven A. Lack (ESRL/GSD and CIRES) explained verification of the CoSPA product and provided several figures.
Judy Ghirardelli (NWS/MDL) helped with a phone interview about LAMP and pointed me to many references. She and her colleague, David Rudack, read the draft description of LAMP and sent comments.
Rita Roberts, Jim Wilson, and James Pinto, NCAR, gave the initial interview on the Autonowcaster and provided many references.
Rita Roberts and Juanzhen (Jenny) Sun (NCAR), reviewed the draft description of Autonowcaster. Jenny Sun also provided information about VDRAS.
David Sharp and Matthew Volkmer, NWS/Melbourne, FL, described their experience with the Autonowcaster.
Haig Iskenderian, MIT Lincoln Labs, provided extensive documentation on SATCAST.
Justin Sieglaff, SSEC, University of Wisconsin, pointed me to information sources about UWCI.

Fog and Low Ceiling
Paul Herzegh gave two interviews on CVA and CVF and provided a number of references and figures used in the report. He also reviewed the draft description of CVA and CVF and sent comments and corrections.
Andrew Loughe (ESRL/GSD) provided a report on the critical evaluation of CVA.
Charles Kluepfel, Performance Branch, NWS, spoke with me about verification, especially of TAFS.
Mike Pavolonis (STAR/NESDIS/U. Wisc.) and Corey Calvert (CIMMS, U. Wisc.) helped me understand the GOES-R product for fog/low clouds. They pointed me to extensive formal and informal documentation. Corey and Mike reviewed my description of the GOES-R product.
Paul Herzegh (NCAR) informed me about HEMS.

Hazards at and near Airports
Roy Rasmussen (NCAR) interviewed on ground de-icing of aircraft and the WSDDM project.
Scott Landolt (NCAR) gave a very quick review and helpful comments on the segment on ground de-icing.
David Johnson (NCAR) enlightened me on lightning protection at airports.
William Mahoney and David Johnson (NCAR) sent information on LLWAS.
William Bobb (MIC, Denver CWSU) arranged for a half-day visit to his facility, including interviews and a tour of the collocated ARTCC. He also provided considerable documentation about the CWSU.
Larry Larkin (TMU, Denver ARTCC) interviewed and explained operations at the Traffic Management Unit.

Alaska and Volcanoes
Kristine A. Nelson (MIC, Anchorage CWSU) gave a personal interview and provided documentation.
Tony Hall (MIC, AAWU, Anchorage), gave briefings and arranged for access to his staff during the day shift.
Jeffrey M. Osiensky (SSD, Alaska Region, NWS) arranged my itinerary for a visit to Anchorage in January 2011, gave personal briefings, and offered extensive documentation regarding special problems faced by aviation in Alaska.
Christina A. Neal, Volcanologist (USGS, AVO, Anchorage), gave an interview and facility tour at AVO. She also provided extensive documentation on Alaska volcanoes.
I thank Gary Hufford (SSD) and Phil Dutton (Alaska Field Support Unit), both at the Alaska Region headquarters, NWS, in Anchorage for briefings and discussions.
Nathan Eckstein and Dave Hadley, forecasters at AAWU, Anchorage, took time out from their shift to discuss their jobs with me.
Barbara Stunder (ARL) helped with documentation on HYSPLIT.
Mike Pavolonis and Justin Sieglaff (SSEC, U.Wisc) provided documentation and consultation on satellite products relating to volcanic ash.

Atmospheric Modeling
Stan Benjamin and John M. Brown (ESRL/GSD) provided interviews, Powerpoint presentations, and documentation on RUC, RAP, and HRRR models. They also reviewed my draft description of these models.
Geoff DiMego (NCEP/EMC) interviewed on GFS and NAM models and provided extensive documentation. He reviewed what I wrote about these two models.

A Survey of Selected Aviation Weather Products

By Thomas W. Schlatter

February 2012

Table of Contents

1. Introduction

1.1 Purpose of the Survey

This survey was commissioned by the National Weather Service (NWS), U.S. Department of Commerce. Its purpose is to obtain a fairly comprehensive inventory of aviation weather products available at the middle of calendar year 2011. The survey will be used by the NWS to determine where gaps in service lie and how to fill those gaps through development of new or improved products. The survey is also a reference work that developers and consumers of aviation weather products should find useful.

An aviation weather product is any source of information about atmospheric conditions from the surface to the lower stratosphere, which is of interest to those who fly. Much of this information is available to the public on the Internet in the form of text, weather charts, and other graphical displays that pertain to aviation safety and the efficiency of flight operations. Some information is available only to government agencies responsible for controlling air traffic, for example, maintaining aircraft separation, routing aircraft around hazardous weather, maintaining efficient use of airspace, and expediting landings and takeoffs. Such specialized information appears on dedicated graphical display devices or, when seconds count, may even be transmitted verbally. Another large class of aviation weather products, most of it proprietary, is generated by the private sector and sold to commercial airlines and a host of general aviation (GA) customers.

1.2 Scope and Limitations

The survey is limited to aviation weather products generated in the U.S. and with geographical coverage including the contiguous states, Alaska, and nearby coastal areas.

The products must be either operational (guaranteed to be available as scheduled) or in prototype testing, which means that a representative class of intended users receives the product regularly and in real time for evaluation.

The products must be *visual*, in the form of text, graphics, or pictures. Almost all such products are on the web, but access is sometimes restricted. This excludes information sent or received via phone, hotline, or headsets used in the cockpit.

The products relate to specific aviation weather hazards:

En route hazards

- Turbulence
- Icing
- Convection

Hazards affecting airports: approaches, departures, and ground operations

- Low ceiling and visibility
- Freezing or frozen precipitation (de-icing)
- Cloud-to-ground lightning
- Strong low-level winds or wind shear

The grouping of hazards is mostly a matter of convenience, in that thunderstorms, low clouds, icing, and turbulence affect both en route aircraft as well as ground operations.

This survey emphasizes basic products, of which some are stand-alone and others are part of larger systems. Basic products receive more detailed coverage than larger systems in this survey. The scientific underpinnings of products are emphasized throughout because improvements to existing products or development of new products to fill service gaps will inevitably build on this scientific foundation.

This survey includes products approved by the Federal Aviation Administration (FAA) and the NWS—products available to the public and products whose distribution is restricted to government agencies that help ensure aviation safety and efficiency. The survey does not extend to space weather. The primary space weather hazard occurs on very high-latitude commercial flights when solar storms affect the ionosphere and knock out communications, a potentially dangerous situation. However, as far as the author knows, no air accidents have resulted from this. The survey excludes proprietary products, for example, those used by airlines that subscribe to a commercial service. However, because the private sector generates a host of aviation weather products, it receives further mention in Section 1.3.2.

1.3 How Government Regulations Affect Providers and Users of Aviation Weather Information

A few definitions will be helpful in the following subsections. *Civil aviation* refers to all non-military aviation. Civil aviation has two major categories, *scheduled air transport*, carrying passengers and/or cargo over designated routes, and *general aviation*, including all other civil flights, private or commercial, that are *un*scheduled.

Title 14 of the Code of Federal Regulations (14-CFR) contains voluminous Federal Aviation Regulations (FARs). These regulations are divided into sections, called *Parts*. The parts mentioned in this report pertain to scheduled, commercial aviation: Part 121 refers to all major and most regional commercial airlines carrying ten passengers or more. Part 135 refers to airlines carrying nine passengers or fewer on regularly scheduled flights. Such airlines are subject to the most stringent safety requirements.

Pilots fly under either *Visual Flight Rules* (VFR) or *Instrument Flight Rules* (IFR). Under VFR, a pilot relies solely on visual cues as to location, aircraft attitude, and separation from the ground and other aircraft. Typical daytime VFR requires visibility of three miles and maintaining separation from clouds, at least 1,000 ft above them, 500 ft below them, and 2,000 ft on either side. On a VFR flight, a pilot does not have to file a flight plan nor be in communication with Air Traffic Control (ATC). Under IFR rules, a pilot relies on instruments for position, attitude, and navigational information. ATC tracks IFR flights via radar and maintains aircraft separation. IFR flights in controlled airspace require a flight *clearance*, essentially permission to fly, based upon the route, speed, time in the air, speed, and altitude. From a practical standpoint, controlled U.S. airspace includes all airspace above 18,000 ft, some airspace corridors below 18,000 ft, and airspace near airports.

1.3.1 Legally defined classes of aviation weather products

The FAA defines two classes of aviation weather products, primary and supplementary. All flight-related, aviation weather decisions must be based on *primary* aviation weather products; these meet regulatory requirements. *Supplemental* aviation weather products may be used to enhance situational awareness. Aviation weather products produced by NWS are primary products unless designated as supplementary by the FAA. Any limitations imposed by the FAA on the use of a product must appear on the product label. Advisory Circular AC 00-45G, Change 1, published jointly by NWS and FAA in July 2010, describes primary and supplementary aviation weather products.

Many products developed with support of the FAA's Aviation Weather Research Program have become supplementary products; some have become primary products. The path from research to operations has, in the past, involved a formal process involving several steps:

1) Acceptance of a new weather forecast concept; sponsor agrees to fund development.
2) Initial development and testing of the new product to verify forecast quality relative to a baseline of performance measured from comparable operational weather products.

3) Scientific and technical review of the product after prototype testing and verification that it meets user needs and maintains sufficient forecast quality. At this stage, the product is still considered experimental; though it may be available on the web, it cannot be used to meet regulatory requirements. Further testing and refinement of the product leads to a second evaluation and a decision either to approve or reject transition to operations. If the transition is approved, the product moves to Step 4.
4) Routine, real-time generation of the product (designated supplementary, at least at first) by the Aviation Weather Center (AWC) in Kansas City

This process is under review, and, though the details may change, the main point is that FAA-approved products have withstood extensive and rigorous testing and scrutiny.

1.3.2 FAA-approved sources of aviation weather information

FAA Order 8900.1, Flight standards Information Management System (FIMS), contains fifteen volumes. Volume 3, *General Technical Administration*, contains 56 Chapters, of which chapter 26, "Aviation Weather Information Systems for Air Carriers," is germane to this section.

1.3.2.1 National Weather Service (NWS)

All Part 121 and Part 135 operators must rely on approved sources of weather observations at any airport where IFR departures or approaches are conducted. METAR reports, available at NWS offices and most sizeable airports, are the gold standard for surface aviation observations. The NWS also sanctions automated reports from a few other sources. Within the 48 contiguous states, all Part 121 operators, which always fly under IFR, must use weather forecasts issued by the NWS or by an EWINS provider (next section). All Part 135 operators flying under IFR must follow the same requirement. With few exceptions, aviation weather products generated by the NWS are designated primary.

1.3.2.2 Enhanced Weather Information System (EWINS)

FIMS 8900.1, Vol. 3, Chapter 26, cited earlier, states "Certificate holders [those licensed to fly certain types of aircraft] may choose to provide their own weather information system or to contract with a private company for weather information. Regardless of the weather sources used, a weather information system must provide all operationally necessary meteorological information for each phase of flight and ground operations." This directive essentially

authorizes Enhanced Weather Information Systems (EWINS). EWINS incorporate advanced technical capabilities, are approved by the FAA, and provide certificate holders with aviation weather data that permit quick, flexible, and operationally efficient responses to changing meteorological conditions. Part 121 certificate holders are especially motivated to operate their own EWINS or enter into contracts with EWINS providers in order to protect schedules and use equipment and personnel with maximum efficiency.

The requirements for an EWINS are stringent: an EWINS must provide sufficient procedures, personnel, and communication and data processing equipment to effectively obtain the necessary aviation weather data from approved sources. It must have the necessary qualified personnel, procedures, and equipment for effective analysis and evaluation of aviation weather data and of the effect of changing weather conditions on current and future operations.

EWINS personnel must be either an aviation meteorologist or a flight dispatcher with authority to make flight movement forecasts (FMFs). An FMF is a forecast prepared and issued by an EWINS containing meteorological conditions expected during departure, en route flight, and landing of a specific aircraft. An FMF is an authoritative forecast that meets legal requirements for flight planning and release by a dispatcher.

An aviation meteorologist must have a degree in meteorology, meet specific course requirements, and be certified by his or her employer as competent to perform aviation forecasting duties. Dispatchers are authorized to make FMFs after they have completed an FAA-approved initial training course in meteorology. In addition, they must complete a refresher training course at least once every two years.

Principal Operations Inspectors from the FAA may approve an EWINS following submittal of required documents, and examination of facilities and personnel. Based upon meteorological data acquired through an EWINS, aviation meteorologists are authorized to issue, update, and/or modify general forecasts of routine or adverse weather phenomena. Based upon the same information, a dispatcher with FMF authority may modify forecasts of adverse weather phenomena as they affect a particular flight.

Many major airlines operate their own EWINS or obtain contract EWINS services from the private sector.

1.3.2.3 Qualified Internet Communication Provider (QICP)

Again quoting from FIMS 8900.1, Vol. 3, Chapter 26: "If a Part 121 or Part 135 operator obtains aviation weather and Notices to Airmen (NOTAM) data via the Public Internet, the operator must use an approved Qualified Internet Communication Provider (QICP)." QICPs must meet certain eligibility standards of reliability, accessibility, and security. In particular, (1) users must be able to retrieve requested data from the provider with no outage lasting longer than 10 minutes, and no more than 30 minutes of outages in any continuous three-month period; (2) the provider must initiate transmission of data with all its users within two minutes of a data request; (3) the provider must demand user authentication and protect data from corruption, intentional or not.

Potential QICPs must submit an application to the FAA containing a service description, a security plan, a capability demonstration plan, a maintenance plan, and a warning label. If the FAA accepts these documents, the applicant must then successfully complete a demonstration of capability. The warning label is explicit that the FAA does *not* necessarily approve the quality of the vendor's weather content or data. It merely certifies that standards of reliability, accessibility and security are being met. The user assumes the entire risk when making decisions based upon information obtained from a QICP. Table 1.1, current on 3 August 2010, lists approved QICPs and gives contact information.

Table 1.1. List of approved Qualified Internet Communication Providers

	Approved QICPs	Contact Information
1.	NAS Aeronautical Information Management Enterprise System (NAIMES)	Amy Johns 202-385-8436
2.	Jeppesen	Mike Cetinich 303-328-6106
3.	WSI	Paul Devlin 978-983-6766
4.	DTC DUAT	Bill Young 856-228-3232
5.	CSC DUATS	Leon Thomas 703-818-4971
6.	Telvent DTN	Steve Madsen 952-882-4587
7.	Weathernews Americas, Inc.	Ryoichi Unten +81 43-296-8099
8.	FltPlan.com	Ken Wilson

		203-262-8900
9.	Flight Explorer	Jeffrey Krawczyk 609-601-7049
10.	Weather Decision Technologies	Matt Orr 405-579-7675 x222
11.	WeatherTap	Robert Parsons 800-337-5263
12.	Air Routing International	Greg Murray 713-430-7200
13.	Honeywell	Lisa Mueller 425-885-8788
14.	Universal Weather and Aviation, Inc.	Bert Poore 713-947-5657
15.	ARINC Direct	David Hancock 413-410-5564
16.	IPS MeteoStar, Inc.	David VandenHeuvel 303-242-5002 x313
17.	Aviation Weather Center	John Sereno 816-584-7247

1.3.3 How regulations regarding aviation weather affect pilots

This section gives three examples of regulations affecting pilots. The first is very simple: pilots must base their decisions on primary products, generated either by the NWS or an EWINS. Supplemental products may enhance awareness of the weather to be encountered during flight, but they are not the basis for decisions regarding flight safety.

The second example pertains to Instrument Flight Rules (IFR) for Part 135 aircraft (scheduled commuter aircraft carrying up to nine passengers). Here are some pertinent regulations.

- "No pilot may begin an instrument approach procedure to an airport unless that airport has a weather reporting facility operated the NWS or a source approved by the FAA *and* the latest weather report indicates that weather conditions [ceiling and visibility] are at or above the authorized IFR landing minimums for that airport." This applies to the airport of origin as well as the destination airport.
- "No person may begin an IFR flight unless the latest weather reports or forecasts—indicate that weather conditions at the estimated time of arrival at the next airport of intended landing will be at or above authorized IFR landing minimums."

- "No person may designate an alternate airport unless the weather reports or forecasts…. indicate that the weather conditions will be at or above authorized alternate airport landing minimums for that airport at the estimated time of arrival."

NWS offices issue Terminal Aerodrome Forecasts (TAFs), which include ceiling and visibility information. Often such forecasts will include conditional wording such as "occasional," "intermittent," or "chance of." The FAA has ruled that the worst conditions implied by the modifying phrases are to be controlling when deciding upon the destination or alternate airports. In other words, if the prevailing conditions at the destination airport are above minimums but intermittent conditions are below minimums, then the pilot may not take off for that location.

This situation becomes more interesting in the cases of EWINS providers. A qualified meteorologist or aircraft dispatcher may issue an FMF based on a detailed analysis of the specific flight with*out* including conditional phrases. The dispatcher may even release a flight to a destination (for which conditional remarks of an NWS forecast indicate the possibility of the destination being below minimums) when the FMF for that specific flight indicates the airport will be at or above minimums. In other words, a certificate holder may rely on an EWINS forecast even if it conflicts with an NWS forecast. Many Part 121 operators (large commercial carriers) who conduct domestic operations have obtained exemptions to release flights to destinations at which the forecast remarks include conditions below minimums. However, when this occurs, additional precautions must be followed, such as the dispatcher monitoring and advising the crew during flight and the pilot designating a second alternate airport.

The third example pertains to a product to be discussed later—the Collaborative Convective Forecast Product, which indicates areas where thunderstorms above a certain intensity are expected, the percent of area covered, and a measure of uncertainty in the forecast. The product is collaborative in that a variety of users may confer on a draft forecast and provide differing opinions before the forecast is formally issued. If Air Traffic Control (ATC) decides that predicted nearly solid coverage and severity of the thunderstorms poses an unacceptable threat and the probability of a correct forecast high, it may decide to divert high-level flights around the area, thus adding to the distance traveled and fuel consumed. Pilots must follow routes specified by ATC.

1.4 Verification

Ultimately, the value of an aviation weather product is gauged by the decisions it prompts and the benefits of those decisions in terms of lives saved, injury and property damage averted, or improved efficiency, e.g., fewer ground delays, less time in the air, and reduced fuel consumption. Because such benefits are often hard to quantify, most measures of value rely on statistical verification scores. A variety of scores is available at the website for the Real-Time Verification System (RTVS). See http://rtvs.noaa.gov/ for verification data pertaining to the main hazards considered in this report: convection, icing, turbulence, and low ceiling and visibility. In a report of this size it will only be possible to sample representative charts.

The following definitions will prove helpful in interpreting various verification measures.
YY — number of predicted events that actually occurred
YN — number of predicted events that did not occur
NY — number of events that occurred but were not predicted
NN— number of times an event was not predicted and did not occur
A two-by-two contingency table has values arranged like this:

Event Predicted?		Event observed?	
		Yes	No
	Yes	YY	YN
	No	NY	NN

Table 1.2 lists some of the skill scores computed from values in the contingency table.

Table 1.2

Standard verification measures that can be computed from the 2x2 contingency table.

Statistic	Definition	Description
PODy	$\frac{YY}{YY + NY}$	Probability of detection of "yes" observations: the fraction of "yes" observations that were forecast correctly. Best case: PODy = 1
PODn	$\frac{NN}{YN + NN}$	Probability of detection of "no" observations: The fraction of non-events that were forecast correctly. Best case: PODn = 1
FAR	$\frac{YN}{YY + YN}$	False Alarm Ratio: the fraction of "yes" forecasts that were incorrect. Best case: FAR = 0
CSI	$\frac{YY}{YY + YN + NY}$	Critical Success Index: Number of correct "yes" forecasts relative to the number of "yes" forecasts and observations. Best case: CSI = 1
TSS	PODy + PODn -1	True Skill Statistic: A measure of discrimination. Best case: TSS = 1
Bias	$\frac{YY + YN}{YY + NY}$	Number of "yes" forecasts relative to the number of "yes" observations. Best case: Bias = 1

When a binary classification system such as the one above depends upon a threshold, then it is useful to look at the *receiver operating characteristic* (ROC), a plot of PODy vs (1 – PODn), as the discrimination threshold is varied. For example, if a cloud contains liquid water, one could 1) make a binary forecast of icing (yes or no) based upon a temperature threshold, 2) count the number of hits and misses at thresholds varying from 0°C and -40°C, and 3) compute PODy and PODn for each threshold and thereby obtain points on the ROC curve.

1.5 Guide to the Remainder of this Document

Sections 2 and 3 occupy most of this report. They describe products associated with aviation weather hazards. Section 2 concentrates on en route hazards; Section 3 concentrates on hazards near the ground, with emphasis on airport operations. The division is not crisp in that many atmospheric phenomena can affect aircraft and flight operations both in the air and on the ground. Section 4 discusses the aviation weather problems specific to Alaska and the products designed to offer guidance to pilots who fly there. Section 5 summarizes the essential and pervasive role that computer prediction models play in the generation of aviation weather products.

Appendix A gives a list of acronyms. Acronyms are always defined upon first use, but Appendix A will prove useful when acronyms appear again many pages later. Appendix B, "Technical Details" is reserved for highly technical or mathematical topics. This information is essential for developers of aviation weather products but will probably be a hindrance to the general reader. The numbering system in Appendix B follows the numbering system in the main body of the text. For example, Section 2.1.5 discusses "Graphical Turbulence Guidance," a turbulence product. The mathematics behind this product is discussed in Section B2.1.5 of Appendix B.

Many figures in this document are at high resolution. If you are viewing this document in electronic format, it is helpful to know that more detail becomes visible in such figures if you zoom in.

Finally, subsections labeled as "commentary" contain opinions by the author. The opinions were solicited by the National Weather Service, which commissioned this work, but they are not necessarily shared by the National Weather Service, nor do they necessarily reflect a broad consensus.

2. En Route Hazards

This major section and the one to follow deal with aviation weather products categorized by the hazard addressed. For each product, the following questions will be answered, if appropriate and providing the answer is known:

- Who developed it?
- What is its purpose?
- How does it work? The scientific basis receives special emphasis in this survey because it provides the rationale for developing a product in the first place.
- What kinds of information are required for input?
- What does a sample product look like?
- Who generates the product?
- What resources are needed to generate the product?
- Is the product stand-alone or part of a larger system?
- Is the product operational? Real time? In prototype testing?
- Who uses the product?
- What types of operational decisions are made based on the product?
- Where is the product available?
- Have the accuracy and utility been assessed?
- Any plans for upgrades?
- Anything else one should know about the product?

2.1 Turbulence

This section deals with the kinds of atmospheric turbulence that cause a bumpy ride at altitude, in particular, clear-air turbulence (CAT), with mountain-wave turbulence as a sub-category. Convection, either weak as in small cumulus clouds, or strong as in a thunderstorm, invariably generates turbulence; it is considered separately in Section 2.4. Turbulence near the ground is covered in Section 3.3.

The *Aeronautical Information Manual* (Table 7-1-9) defines four classes of turbulence: *Light* turbulence causes slight erratic changes in the altitude or attitude (spatial orientation) of the aircraft. Passengers may feel a slight strain against seat belts or shoulder straps. Food service may be conducted and there is little or no difficulty in walking. *Moderate* turbulence causes rapid bumps or jolts without appreciable changes in aircraft altitude or attitude, but the aircraft remains in positive control. Occupants feel definite strains against seat belts and shoulder straps. Unsecured objects are dislodged. Food service and walking are difficult. *Severe*

turbulence causes large, abrupt changes in the altitude or attitude of the aircraft, and usually large variations in indicated air speed. The aircraft may be momentarily out of control. Occupants are forced violently against seat and shoulder straps. Unsecured objects are tossed about. Food service and walking are impossible. *Extreme* turbulence tosses the aircraft violently about and makes it practically impossible to control. Structural damage may occur.

2.1.1 Turbulence Observations

Those who forecast turbulence consult a variety of sources, including observations and numerical model output. Among observational sources, pilot reports (PIREPs) are paramount. Figure 2.1 plots PIREPs for the northeast U.S. The source is http://aviationweather.gov/adds/pireps/displayPireps.php .
The altitude of the turbulence report is in hundreds of feet. The symbols for turbulence intensity are given at the bottom of the figure.

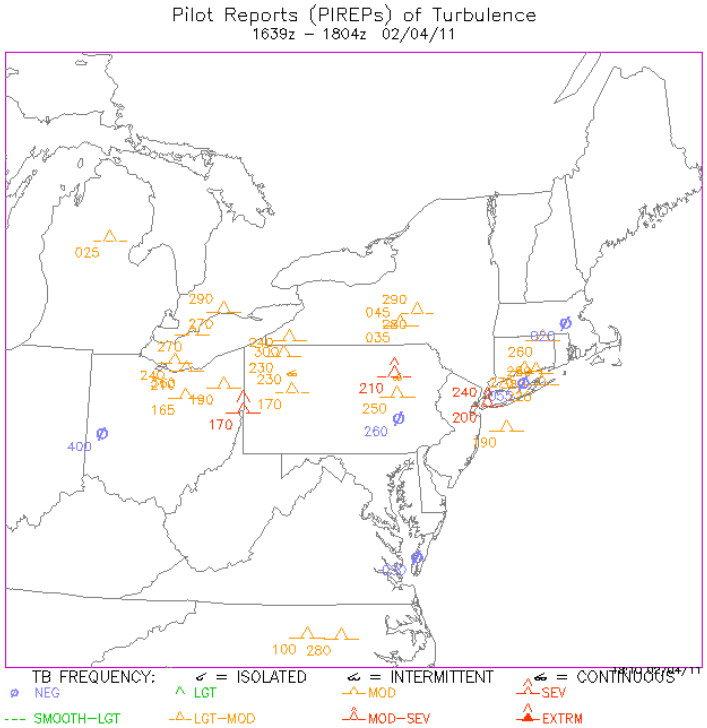

Fig. 2.1. Pilot reports of turbulence, 1639–1804 Greenwich Mean Time (GMT), 4 Feb 2011. Symbols and colors for turbulence intensity are at bottom.

A second source of turbulence observations resides within the avionics systems on selected commercial aircraft. Parameters measured by these systems permit calculation of the eddy

dissipation rate (EDR or ε), a measure of the rate at which energy associated with turbulent eddies decreases with time. The units of ε are $[(m^2s^{-2})/s]$. $\varepsilon^{1/3}$ [units: $m^{2/3}s^{-1}$] is the number usually reported; it is independent of aircraft type, which means that the same value of $\varepsilon^{1/3}$ might give a Cessna a severe jolt but a Boeing 767 only a mild bump.

EDR is calculated in two different ways (Cornman et al. 2004). The first method estimates EDR from a sequence of measured vertical accelerations (the accelerometer takes a measurement many times a second) and a model of the aircraft response to these accelerations. The second method does not require the sometimes hard-to-obtain data on aircraft response. It calculates the vertical wind from the true air speed, the rate of climb or descent, the roll and pitch angles, and the body-axis angle of attack. Mathematical operations on a time sequence of vertical winds permits an estimate of $\varepsilon^{1/3}$.

As of December 2010, 100 United Airline 757s were reporting $\varepsilon^{1/3}$ every minute during cruise. 80 Delta Airline 737s were reporting every 15 min by default, but turbulence triggers a report at any time. 10 Southwest Airline (SWA) 737-700s were testing EDR systems, but over 300 SWA aircraft will eventually carry them.

The normal range of $\varepsilon^{1/3}$ values is from zero to above 0.80. The air is considered smooth for $\varepsilon^{1/3} \leq 0.05$. Figure 2.2 is a sample plot of EDR reports obtained mostly from Illinois, Indiana, and Ohio during three hours beginning at 2100 UTC 01 Feb 2011. A powerful winter storm was crossing the Midwest states at the time. All the blue dots and squares indicate $\varepsilon^{1/3}$ values ≤ 0.05. Other colors indicate higher values and the likelihood of at least some turbulence. One sample report is indicated by the arrow. Aircraft #496 (specific tail numbers are not identified), flying at 36,000 ft on a pressure surface at 227 mb, en route from Dulles International Airport (Washington, D.C.) to Phoenix, reported an average $\varepsilon^{1/3}$ value of 0.15 and a maximum of 0.35. Like most automated aircraft reports, EDR data are proprietary to the airlines that collect them; they are available in real time only to government agencies such as NOAA in support of forecasting operations. The public may view data samples of automated aircraft reports at least 48 hours old at http://amdar.noaa.gov/demo_java/ .

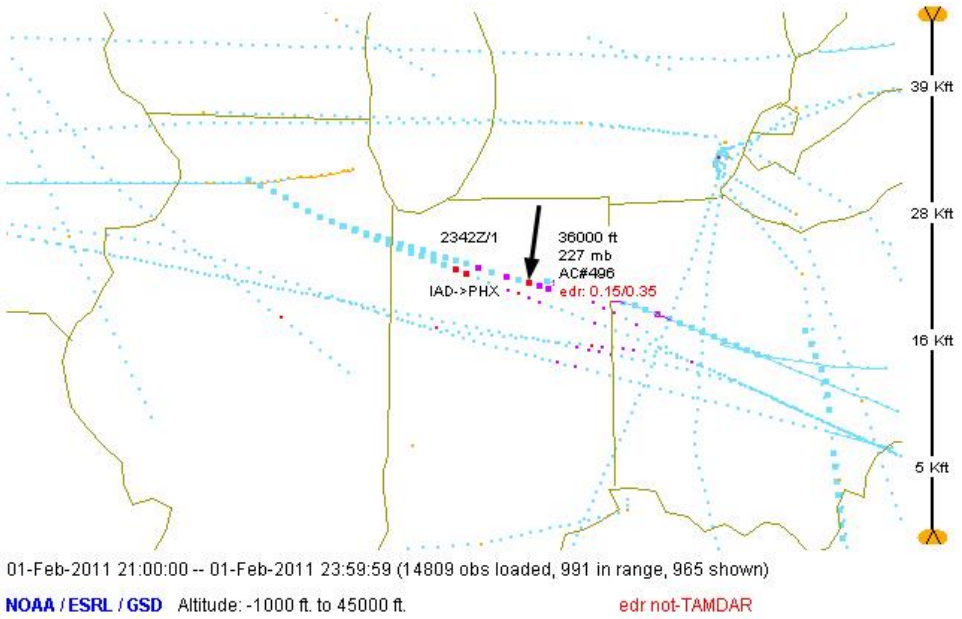

Fig. 2.2. Automated reports of eddy dissipation rates from 2100 thru 2359 GMT, 1 Feb 2011 over Illinois, Indiana, Ohio, and Michigan.

Both voice PIREPS and automated EDR reports are used to verify the accuracy of turbulence products.

2.1.2 AIRMET (Airmen's Meteorological Information)

An AIRMET is a concise description in abbreviated language of the occurrence or expected occurrence of specified en route weather phenomena which may affect the safety of aircraft operations, but at intensities lower than those which require the issuance of Significant Meteorological Information (SIGMET) (see next section). In the case of turbulence, an AIRMET is issued whenever moderate turbulence is expected. The AIRMET specifies the top and bottom altitudes of the turbulent layer. The Aviation Weather Center (AWC) in Kansas City, MO, and the Alaska Aviation Weather Unit (AAWU) in Anchorage issue AIRMETs on a six-hour schedule in both plain language and graphical form. Unscheduled updates are issued if unexpected conditions develop. An example of a plain language AIRMET appears below.

AIRMET (U.S. WEST COAST)

```
WAUS46 KKCI 032045
WA6T
SFOT WA 032045
AIRMET TANGO UPDT 4 FOR TURB VALID UNTIL 040300.
```

```
AIRMET TURB...CA ID WY NV UT CO AZ NM AND CSTL WTRS
FROM DDY TO BFF TO GLD TO 50W LBL TO 30ESE TBE TO INK TO ELP TO
50S TUS TO BZA TO 20S MZB TO 170SW MZB TO 160SW RZS TO 120SW PYE
TO MOD TO 80SW TWF TO DDY
MOD TURB BTN FL180 AND FL400. CONDS CONTG BYD 03Z THRU 09Z.
.
AIRMET TURB...WA OR ID MT WY AND CSTL WTRS
FROM 30S YQL TO 50NNW ISN TO 70NW RAP TO 70ESE DLN TO 20NE DBS
TO 30W REO TO 50ESE DSD TO 30S PDX TO 110WNW ONP TO 150W TOU TO
20ESE HUH TO 50SW YXC TO 30S YQL
MOD TURB BLW 150. CONDS CONTG BYD 03Z THRU 09Z.
.
AIRMET TURB...CA WY NV UT CO AZ NM AND CSTL WTRS
FROM 30SSE OCS TO 30NW PUB TO 20WSW LAA TO 30ESE TBE TO 60SE CME
TO ELP TO 50S TUS TO BZA TO 20S MZB TO 120WSW MZB TO 70WNW RZS
TO 30SSE OCS
MOD TURB BLW 150. CONDS CONTG BYD 03Z THRU 09Z.
```

These AIRMETs specify the boundaries of areas subject to moderate turbulence and the altitude ranges, but they are not easy to read unless one knows the word contractions, three-letter station identifiers, and their locations. The turbulent areas are even more difficult to visualize. For that purpose, AWC makes available a tool called *Graphical AIRMET* or simply *G-AIRMET* on the Web at http://aviationweather.gov/products/gairmet/ . Figure 2.3 shows the AIRMET for turbulence. Note the two types of scalloped areas. Orange boundaries enclose areas of high-level turbulence; red boundaries enclose areas of low-level turbulence. Note also the choice of valid times at top right. One can choose from among five valid times.

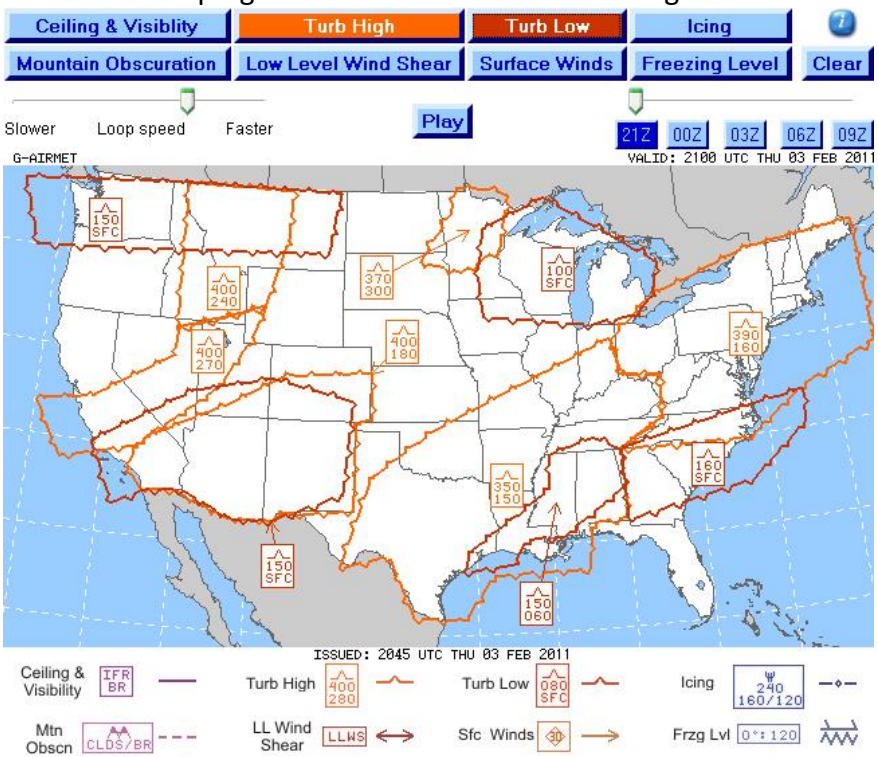

Fig. 2.3 AIRMET for turbulence issued 2045 GMT, 03 Feb 2011.

2.1.3 SIGMET (Significant Meteorological Information)

A SIGMET is an unscheduled advisory in abbreviated plain language concerning the occurrence or expected occurrence of potentially hazardous en route weather phenomena that may affect the safety of aircraft operations. They are intended for dissemination to all pilots in flight. AWC or AAWU forecasters issue a SIGMET whenever conditions meeting certain criteria are expected to affect an area of at least 3000 square miles. If the hazard covers an area less than 3000 square miles, the Center Weather Service Unit (CWSU), collocated with an Air Route Traffic Control Center (ARTCC), will issue a Center Weather Advisory (CWA – next section). SIGMETs are valid for not more than four hours at a time. AIRMETs and SIGMETs are *primary* aviation weather products. Pilots are legally required to avoid SIGMET areas for severe icing unless their aircraft is equipped to deal with it. Pilots may legally enter a SIGMET area for severe turbulence at their own risk.

A SIGMET for turbulence is issued whenever severe or extreme turbulence is reported or expected. Here is a sample of a text SIGMET issued on the basis of a pilot report.

```
WSUS01 KKCI 041543
BOSU WS 041543
SIGMET UNIFORM 2 VALID UNTIL 041943
SIGMET
NH VT MA RI CT NY NJ PA
FROM SYR TO CON TO PVD TO JFK TO 40WSW HAR TO SLT TO SYR
OCNL SEV TURB BTN FL210 AND FL270. RPTD BY CRJ AT FL250-220 NR
HAR. CONDS CONTG BYD 1943Z.
```

The corresponding graphical SIGMET in Fig. 2.4 is extracted from http://aviationweather.gov/adds/airmets (click on turbulence SIGMETS). Like many aviation weather products, this graphic is available on the Aviation Digital Data Service (ADDS) website.

The report of severe turbulence that triggered the issuance of this SIGMET is shown in Fig. 2.1. It turns out that a maximum in jet stream winds (jet max), over 150 knots at 300 mb, was passing through the area of the SIGMET, with strong vertical shear below.

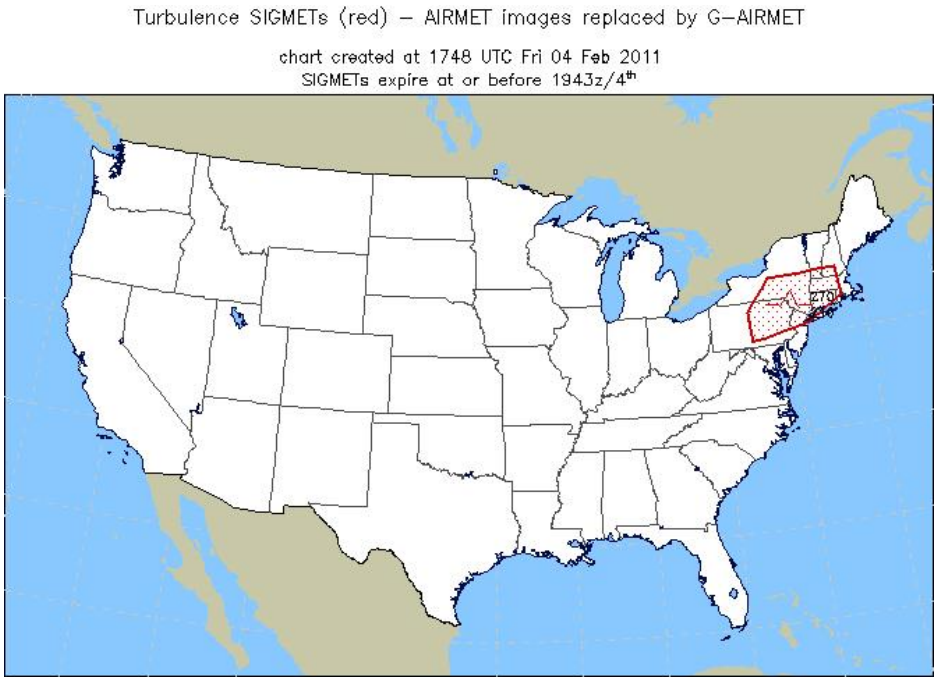

Fig. 2.4. A SIGMET for severe turbulence in the northeast U.S. issued at 1748 GMT, 04 Feb 2011.

Forecasters use many rules of thumb when issuing AIRMETS and SIGMETS for turbulence. They look for strong winds at various altitudes, for example, > 60 knots at 700 mb, >90 kt at 500 mb, and >110 kt at 400 mb. They look for layers with strong vertical shear. The Richardson Number (Ri), is a dimensionless ratio containing a measure of atmospheric stability in the numerator and of vertical shear in the denominator. When Ri <1.0 in a layer at least 5000 ft thick, turbulence is indicated. Atmospheric soundings from rawinsondes, aircraft on ascent or descent, or derived from model output, are available on the interactive display systems in all NWS offices: the Advanced Weather Information Processing System (AWIPS). The geographic location of soundings from model output is arbitrary, and many parameters derivable from the temperature, dewpoint, and wind profiles are available, including Ri and turbulence kinetic energy (TKE), a measure of the mean kinetic energy per unit mass associated with eddies in turbulent flow.

Forecasters of turbulence are interested in horizontal shear as well as vertical shear. A 90^{o} change in wind direction or more within 150 km with winds >10 knots on either side of the wind shift is a favored location for turbulence. So is a region of strong cold or warm air advection at any level where the wind is strong.

Certain features in satellite images give clues about the presence of turbulence. In water vapor images, dark or darkening bands (indicating dry air) on the poleward side of the jet stream, or sinking (dark) signatures in the lee of mountains suggest turbulent flow. On visible satellite images, wide and thick bands of cirrus oriented perpendicular to the wind direction and curved bands in other kinds of clouds indicate turbulence. If a cirrus cloud shield is associated with a straight or anticyclonically curved jet developed scallop pattern on its poleward edge, turbulence is indicated. Billow clouds, marking a series of so-called Kelvin-Helmholz waves, almost always indicate turbulence. Many other rules of thumb for turbulence, and aviation weather forecasting in general, may be found in Bader et al. 1995.

From model output, forecasters plot the time tendency of divergence (three-hour change) in units of $10^{-9}s^{-2}$ in 50-mb thick layers between 500 and 200 mb. If they see positive and negative centers close together, both with magnitude of at least 20 and values differing by ≥ 60, turbulence is likely between couplets. These limits apply for computations on a 30-km grid, well away from convection.

AWC forecasters use TKE diagnostics from the North American Mesoscale (NAM) model, mainly to make sure they're not missing any turbulence. Numerical predictions of turbulence are still in their infancy, partly because the scales of turbulence lie far below the scales resolved by any operational models. TKE diagnostics are a strong function of model resolution.

AIRMETs and SIGMETs are routinely verified by the Real Time Verification System (RTVS – Section 1.4). Figure 2.5 plots the PODy for turbulence SIGMETs issued anywhere in the lower 48 states for each month of 2010. Table 2.1, below the figure, helps in the interpretation.

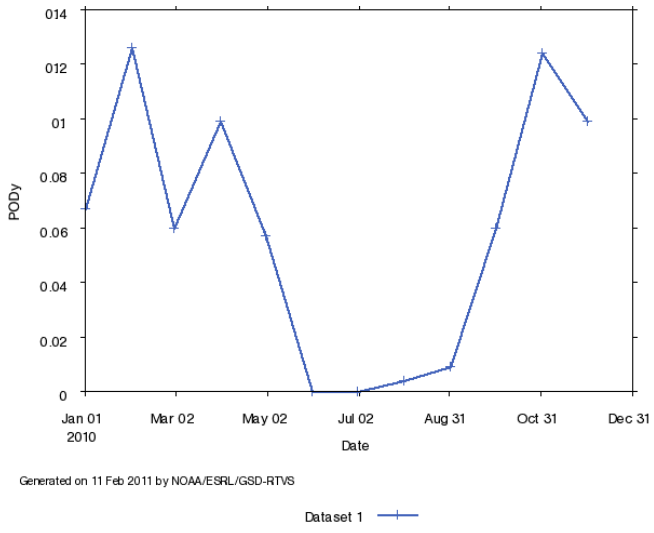

Fig. 2.5 PODy (percent of observed moderate or severe turbulence reports correctly predicted) by month for 2010. For this plot, a report of moderate or greater turbulence was deemed to verify the SIGMET.

Table 2.1

Date	YY	YN	NY	NN	PODy
2010-01-01	48	352	667	26613	0.067
2010-02-01	80	468	556	21149	0.126
2010-03-01	47	457	739	29580	0.060
2010-04-01	77	545	701	25665	0.099
2010-05-01	45	220	747	23774	0.057
2010-06-01	0	3	366	16837	0.000
2010-07-01	0	0	203	13849	0.000
2010-08-01	1	0	262	13297	0.004
2010-09-01	3	89	335	17999	0.009
2010-10-01	28	86	438	18901	0.060
2010-11-01	117	724	823	24348	0.124
2010-12-01	66	280	603	23894	0.099
Totals	512	3224	6440	255906	0.074

For Fig. 2.5 and Table 2.1, YY is defined as an aircraft report of moderate or severe turbulence within the SIGMET area and valid period. This inflates the score because SIGMETs are issued only for severe or extreme turbulence. In any case, note the following: 1) Only 3736 SIGMETS for turbulence were issued nationwide in all of 2010, of which 3224 were false alarms (no verifying aircraft reports). There is a good reason for this: once a SIGMET is issued, aircraft avoid the area. 2) On the other hand, aircraft encountered moderate or greater turbulence 6440 times in 2010, at locations where no SIGMET was in effect. This reflects the properties of turbulence—small scales, intermittency, and short lifetime—that make it challenging for humans and models to predict. 3) The vast majority of aircraft (97.4%) reported smooth air or light turbulence. 4) Many summertime encounters with moderate or greater turbulence (831) may have been associated with convection. Only four SIGMETs for turbulence were issued all summer, no doubt because the conditions that generate moderate or greater CAT are rarely present in summer.

2.1.4 Center Weather Advisory

A Center Weather Advisory (CWA) is a weather watch or warning focusing on weather hazards that affect aviation safety and traffic flow control. It is valid for up to two hours; the valid period must begin with two hours of issuance. The CWA uses abbreviations and coded language. Flight crews use CWAs to avoid adverse conditions en route and at airports. CWAs are verified with pilot reports, METARs (surface observations), lightning detection, weather radar, and flight deviation records from the ARTCC's Traffic Management Unit (TMU). To find CWAs for any air traffic region in the country, go to http://aviationweather.gov/products/cwsu/.

CWSU forecasters may issue a CWA for turbulence 1) if moderate or greater turbulence covering an area less than 3000 square miles is expected within the ARTCC region they serve; 2) if the forecaster believes that existing SIGMETs do not cover the hazard he or she perceives (in which case the CWA should be issued in coordination with AWC or AAWU; 3) to cancel an existing CWA if the hazard no longer exists.

```
FAUS21 KZLC 312010
ZLC1 CWA 312010
ZLC CWA 101 VALID UNTIL 312210
FROM LWT-SHR
AREA MOD-SEV MTN WAVE TURB 40NM WIDE. FL340-FL380. ± 15-20KT AND ± 500FT.
RPTD BY NMRS ACFT. CONDS EXP TO CONT THRU PD.
```

Translation: This is a Center Weather Advisory issued by the CWSU for the Salt Lake City air traffic control region, valid on the 31st day of the month from 2010 to 2210 UTC. From Lewistown, Montana, to Sheridan, Wyoming, expect moderate to severe mountain wave turbulence in a band 40 nautical miles wide at flight levels from 34,000 to 38,000 ft. Fluctuations in air speed of ±15–20 knots and in altitude of ±500 ft have been reported by numerous aircraft. Conditions are expected to continue through the period.

2.1.5 Graphical Turbulence Guidance

The Graphical Turbulence Guidance (GTG2) product is a fully automated, *supplemental* product available at http://aviationweather.gov/adds/turbulence/turb_nav.php. It displays turbulence intensity in three categories: none, light, and moderate or greater. It predicts clear air turbulence (CAT) but not turbulence due to mountain waves or convection. The product consists of a one-, two-, and three-hour forecast updated every hour, and a six-, nine-, and twelve-hour forecast updated every third hour. The primary input is a numerical forecast from

the Rapid Update Cycle (RUC) model. This model is operational at the Environmental Modeling Center (EMC), one of the National Centers for Environmental Prediction (NCEP). An algorithm calculates a number of turbulence-related parameters from the model output and weights these parameters to give a "GTG index" of turbulence. Currently, values from 0.0 to 0.3 indicate no CAT; values from 0.3 to 0.475 indicate light CAT; values from 0.475 to 1.0 indicate moderate or greater CAT.

One can display the GTG2 product at levels from 11,000 to 17,000 ft above sea level and flight levels (FL) from 190 and 450 (19,000 to 45,000 ft), all in increments of 2,000 ft, and for any available forecast time. The break point at 18,000 ft is chosen because aircraft are under active control above that level, and FL is determined with reference to a standard atmosphere. Below 18,000 ft, terrain height is a consideration, and so the turbulence product is referenced to geometric altitude. The user can display the maximum turbulence intensity expected anywhere in the vertical. He or she can also display GTG2 information in a vertical cross section following a flight path. The latter capability is part of the Flight Planner, a hybrid, interactive product to be discussed in Section 4.1.

Figure 2.6 is a sample GTG2 product as displayed on ADDS that shows the color-coded *maximum* turbulence intensity for *any* altitude between 10,000 and FL 450. Because this image is for the initial forecast time, available turbulence reports are plotted as well.

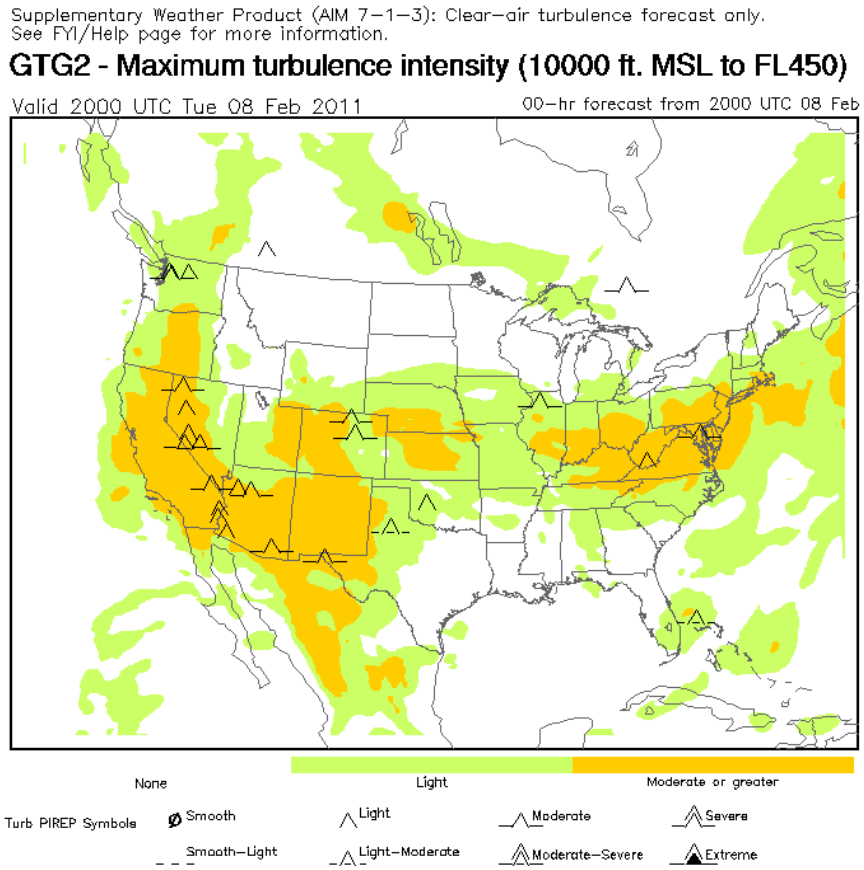

Fig. 2.6. A sample product from Graphical Turbulence Guidance showing the greatest intensity of turbulence to be expected at any level between 10,000 ft above mean sea level (MSL) and a Flight Level of 45,000 ft. Because the displayed product is for the initial time (not a forecast), available turbulence reports are plotted as an overlay.

The scientific basis for GTG2 is described in Sharman et al. 2006. The authors tested a number of algorithms for diagnosing turbulence. Many of these algorithms express mathematically what forecasters are looking for subjectively when they examine various maps and charts for the likelihood of turbulence. These algorithms are described in Appendix B, Section B2.1.5.

The algorithms generate indices which have physical units attached to them. The indices are calculated from Rapid Refresh (RAP) model output on its computational grid, then interpolated to flight levels. Threshold values corresponding to the breakpoints between turbulence categories (null-light, light-moderate, moderate-severe, and severe-extreme) are obtained for each index by extensive comparisons of index values with turbulence reports from pilots. These threshold values are then mapped to a common turbulence intensity scale, which is based on in situ EDR (units $m^{2/3} s^{-1}$) measurements. On this scale, for a medium-sized

commercial aircraft such as a B757, 0.0 corresponds to smooth (i.e. no turbulence), 0.14 corresponds to a "light" turbulence intensity threshold, 0.31 corresponds to moderate, and 0.54 corresponds to severe. The largest acceptable value on this scale is 0.8 $m^{2/3} s^{-1}$. Note that these thresholds apply to the GTG2 product to be implemented in the spring of 2012, when the RAP model becomes operational. The current thresholds, listed at the beginning of this section, apply to indices computed from the RUC model.

Each normalized index (scaled from 0.0 to 0.8) is scored as described by Sharman et al. (2006) by comparing it with pilot reports of turbulence and in situ EDR data within a three-hour time window. Turbulence reports from the vicinity of thunderstorms are not considered. On the basis of these scores, each index is assigned a weight, and a weighted sum of all the indices, separately computed for upper levels and mid levels, gives a grand Graphical Turbulence Guidance (GTG2) index, which itself is displayed on the EDR scale.

Note the following:

- The best set of turbulence diagnostics is not the same at mid levels and upper levels.
- The optimum thresholds at mid and upper levels for the same diagnostic are not necessarily the same.
- The number of pilot reports for verification is much greater at upper levels than at mid levels,
- The weights applied to the different turbulence diagnostics change with time, depending upon atmospheric conditions and the number of available pilot reports.
- New diagnostics may replace the current ones if they prove more effective in diagnosing turbulence.

The ROC curve (see again Section 1.4) in Fig. 2.7 is one way to examine the skill of the GTG2 product. (Other ways are available at http://rtvs.noaa.gov/turb/stats/.) Table 2.2 gives information about the data that support Fig. 2.7. (Note that these statistics are derived from the current, RUC-based GTG2 product). The + marks on the curve correspond to the threshold values in the first column of Table 2.2.

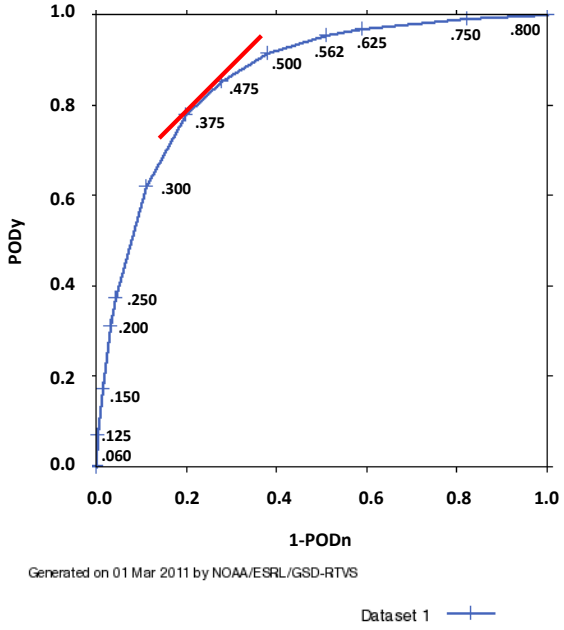

Fig. 2.7. The Receiver Operating Characteristic (ROC) curve for all 3-h predictions of moderate or greater clear air turbulence with the GTG2 turbulence algorithm in 2010. All predictions within the contiguous United States and at flight levels between 20,000 and 40,000 ft were included. The value of the GTG2 Index is plotted beside each point marked on the curve. The bold red line tangent to the ROC curve is explained in the text.

Table 2.2
Data supporting the Receiver Operating Characteristic (ROC) curve in Fig. 2.7.

Data Source	Variable	Obs Threshold	Start Date	End Date	Product Time	Lead Time	Region	Flight Level
GTG2	Index 0.0-1.0	N=no Y≥Mod	2010 Jan 01	2010 Dec 31	All	3 hours	Nat'l	20K-40K

Fcst Thresh	YY	YN	NY	NN	1-PODn	PODy
0.8	15	0	31910	24728	0.000	0.000
0.75	109	3	31816	24725	0.000	0.003
0.625	2296	111	29629	24617	0.004	0.072
0.562	5536	338	26389	24390	0.014	0.173
0.5	9985	774	21940	23954	0.031	0.313
0.475	11904	1034	20021	23694	0.042	0.373
0.375	19802	2719	12123	22009	0.110	0.620
0.3	24860	4917	7065	19811	0.199	0.779
0.25	27245	6881	4680	17847	0.278	0.853
0.2	29185	9429	2740	15299	0.381	0.914
0.15	30467	12622	1458	12106	0.510	0.954
0.125	30933	14534	992	10194	0.588	0.969
0.06	31672	20338	253	4390	0.822	0.992

Over 46,000 turbulence reports from 2010 are represented in Fig. 2.7 and Table 2.2, including reports of no turbulence. Note that a report of moderate or greater turbulence is considered a "yes" observation. Reports of light turbulence are not considered. Only three-hour forecasts are verified. For GTG2 index values of 0.2 or less, the probability of detection of moderate or greater turbulence is > 0.91, and the false alarm rate [FAR = YN/(YY+YN)] is > 0.24. The downside is that the GTG2 index is designed so that values <0.25 are meant to suggest no turbulence. Values of the GTG2 index between 0.50 and 0.75 are meant to suggest moderate turbulence. Less than 19% of all GTG2 forecasts have an index as high as 0.50, and the percentage decreases rapidly with still higher index values. Only 15 times did the 3-h predicted index reach 0.80 in 2010. Every time the index was 0.80, moderate or greater turbulence was experienced, and the false alarm rate was perfect. However, nearly 32,000 events of moderate

or greater turbulence were recorded for GTG index values *less* than 0.80, so one would definitely not want to use this value as a threshold for finding moderate or greater turbulence. Some statisticians suggest this method for determining the optimum threshold: find the GTG2 index value where a straight line at a 45° angle is tangent to the ROC curve. The red line in Fig. 2.7 is this tangent. The point of tangency is at a GTG2 index of about 0.40. Thus, if one wanted to make a binary decision about moderate or greater turbulence, one would say "no" for index values $\leq$0.40 and "yes" for values >0.40.

The FAA's Aviation Weather Research Program has funded development of the GTG product. Plans for improving GTG stretch out for more than five years. They include the detection and prediction of mountain-wave and convectively-induced turbulence; the incorporation of new data sources, namely, Doppler radar and satellites; and the generation of probabilistic forecasts. The next two sections describe products already available that rely on satellite and Doppler radar.

2.1.6 Satellite Turbulence Products

2.1.6.1 Background

GOES-R is the next in a series of Geostationary Operational Environmental Satellites, now expected to be launched in 2015. It will feature these improvements over the current GOES satellites:

- An Advanced Baseline Imager (ABI) with 16 channels (viewing in 16 different wavelength intervals). Depending on the channel, the image resolution will be 2 km or better at the subpoint as opposed to 4 km on the current GOES. The ABI will be able to scan the full earth disk in 5 min (25 min for the current GOES), the continental U.S. (CONUS) in 5 min, and designated regions every 30s. GOES R will *not* have an infrared (IR) sounding instrument.
- A lightning mapper that measures total lightning (in-cloud, cloud-to-air, and cloud-to-ground). The lightning mapper will work day and night, but detection efficiency will be better at night.

In preparation for GOES-R, satellite product developers at the Cooperative Institute for Meteorological Satellite Studies (CIMSS) in Madison, WI; the Cooperative Institute for Research in the Atmosphere (CIRA) in Fort Collins, CO; and the Short-term Prediction Research and Transition (SPoRT) Center at the National Aeronautics and Space Administration (NASA)

Marshall Space Flight Center in Huntsville, AL, and the NASA Langley Research Center (LaRC) in Hampton, VA, are participating in a program called the *GOES-R Proving Ground.* They are developing aviation weather products that will mimic those to be available from the GOES-R ABI and will be distributing them in real time to NWS offices and regional centers around the country for evaluation. The look-alike products are generated from data provided by instruments already in space, primarily MODIS, the Moderate Resolution Imaging Spectroradiometer. Flying aboard the NASA Terra and Aqua satellites, MODIS observes in 36 spectral bands in the visible and infrared with a resolution of 1 km or better. This section examines three turbulence-related products developed for the GOES-R Proving Ground.

2.1.6.2. Turbulence Associated with Tropopause Folding

The tropopause is the usually well-defined boundary between the troposphere and the stratosphere. Tropopause folding is a pronounced deformation of the tropopause when strato spheric air intrudes into the troposphere due to ageostrophic flow around the jet stream. This usually occurs in connection with a strong upper tropospheric front and frequently leads to turbulence both above and below the fold, where vertical shear is strong, as illustrated in Fig. 2.8.

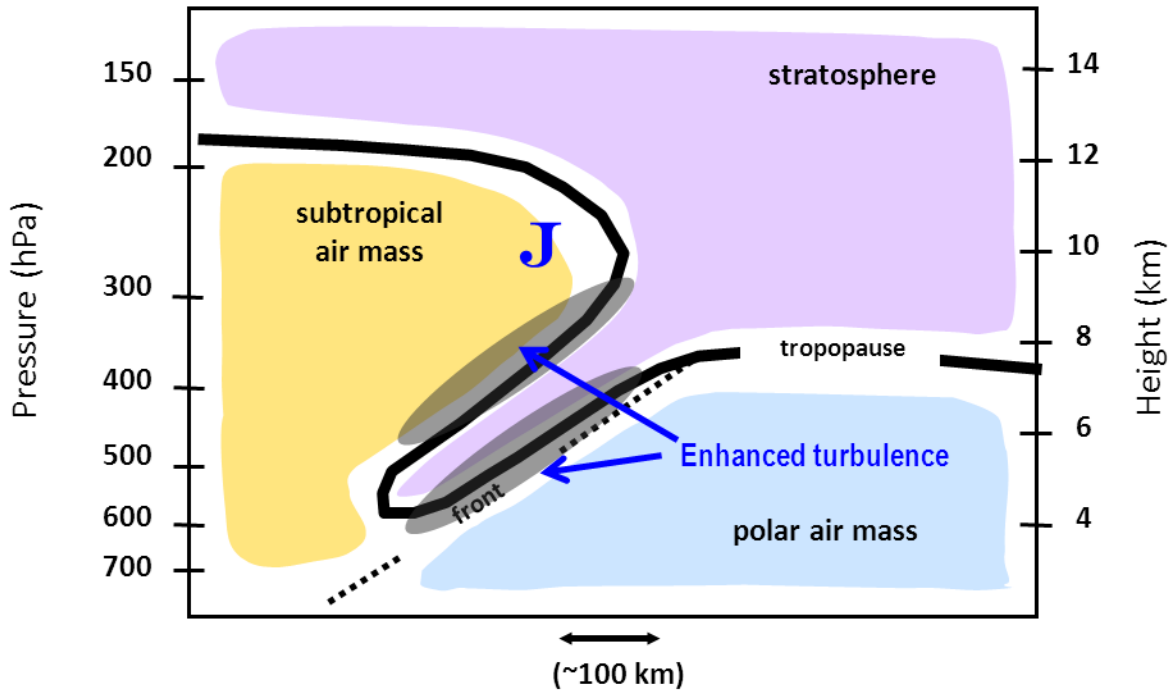

Fig. 2.8 A tropopause fold, represented by the thick, black line. J indicates the position of the jet stream. Stratospheric air is colored purple, tropospheric air, either blue or yellow. Regions where turbulence is likely are shaded in gray. Courtesy of Tony Wimmers, Space Science and Engineering Center (SSEC), University of Wisconsin at Madison.

Very dry stratospheric air descends within tropopause folds, causing a distinctive signature in water vapor images. Sharp gradients in brightness temperature occur near the jet stream and the folds. An edge-detection algorithm captures the boundaries of the dry-air intrusion, and these boundaries, in turn, enclose areas where turbulence might be encountered.
A prototype of the product appears in Fig. 2.9. The shaded areas show where tropopause folding might result in CAT. Automated EDR observations from 30 minutes either side of the product valid time are plotted. The light gray stripes indicate portions of flight paths with negligible turbulence, but the colored dots indicate light or moderate turbulence. As of late 2010, the false alarm rate was < 50%, and the PODy was near 50%.

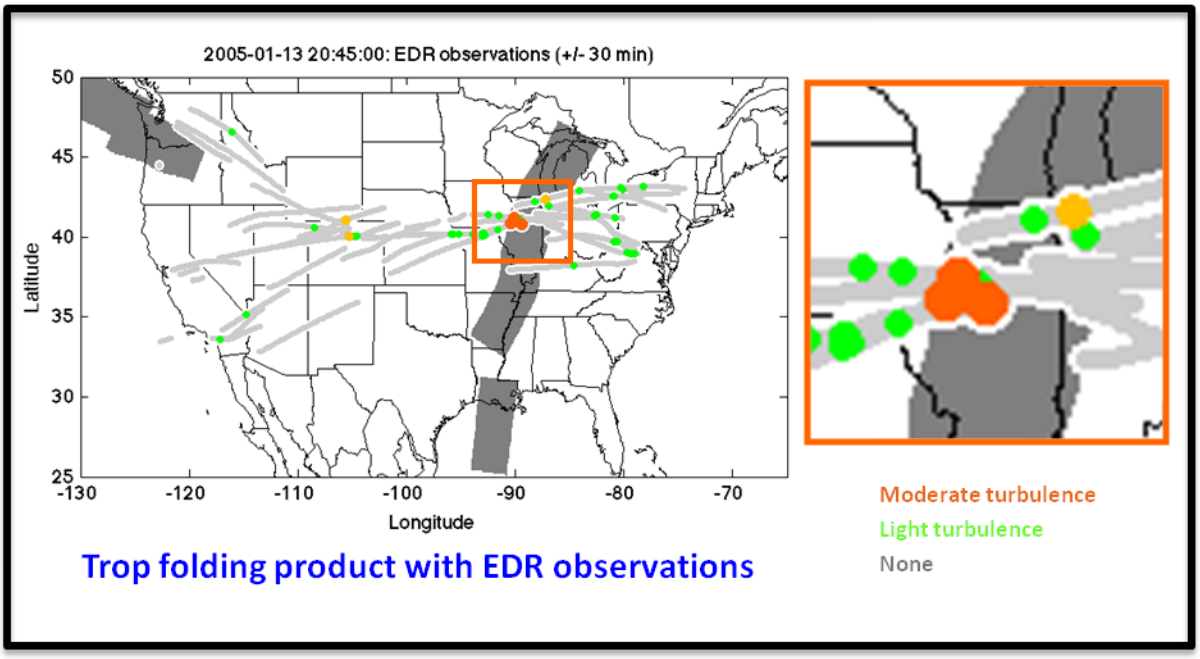

Fig. 2.9. Verification of tropopause fold product with automated aircraft reports of eddy dissipation rate (EDR) within 30 minutes either side of 2045 GMT, 13 Jan 2005. Courtesy of Tony Wimmers, Space Science and Engineering Center (SSEC), University of Wisconsin at Madison.

2.1.6.3 Convective Overshooting Tops

Vigorous thunderstorm updrafts occasionally penetrate the tropopause. This results in a brief excursion (lasting a few minutes) of a portion of the cumulonimbus cloud into the stratosphere—an overshooting top—before it sinks back into the anvil. The overshooting top perturbs the stable layer that is the tropopause and thereby generates gravity waves, which, in turn, cause turbulence nearby. Scientists at SSEC have developed a product that uses MODIS data to detect overshooting tops. Detection with MODIS data is considerably easier than with GOES-12 IR data because the latter have higher spatial resolution.

The algorithm works as follows:

1. Analyze an image in the IR window (where water vapor doesn't absorb upwelling radiation). A channel near 11-μm wavelength is the normal choice.
2. Identify pixels for which the brightness temperature (BT) is less than or equal to a) 215K *and* b) the tropopause temperature as determined from the Global Forecast System (GFS) model running at EMC.
3. Rank the pixels identified in Step 2 from lowest to highest BT. Start with the lowest BT in the list. It can be labeled a "candidate overshooting pixel." Find the second lowest BT in the list. If it has no neighbors as cold as it is within 15 km, it too can be labeled a "candidate overshooting pixel." Find the third lowest BT in the list, and continue the process. This will effectively isolate those pixels with the lowest BT in each cluster of cold pixels. The 15-km radius is chosen because overshooting tops are usually not more than 15 km in diameter.
4. For each candidate overshooting pixel, sample the surrounding anvil at a distance of 8 km in 16 directions (at 22.5^{o} intervals). This distance should be great enough as to lie outside an overshooting top.
5. If the candidate overshooting pixel has a BT at least 7.5K lower than the average BT of the surrounding anvil, then it is considered to be an "overshooting center pixel (OCP)."
6. Say that the BT of the OCP is ΔT degrees lower than the average BT of the surrounding anvil. Within a 7x7 pixel box centered on the OCP, find those pixels with BT at least $0.5\Delta T$ less than the anvil BT. These pixels, together with the overshooting center pixel, define the entire overshooting top.

Part of this process is illustrated in Fig. 2.10, courtesy of Kristopher Bedka, Science Systems and Applications Inc., at the NASA Langley Research Center. At left is a high-resolution image of some anvils. The cauliflower-like bulges are candidates for overshooting tops. Of course, this image would not be available at night, and so the algorithm uses high-resolution IR images. The center frame shows the candidate overshooting pixels and how pixels from the surrounding anvil are identified. The right frame shows not only the overshooting center pixels but

surrounding pixels that were cold enough to be included in the cluster defining the overshooting top.

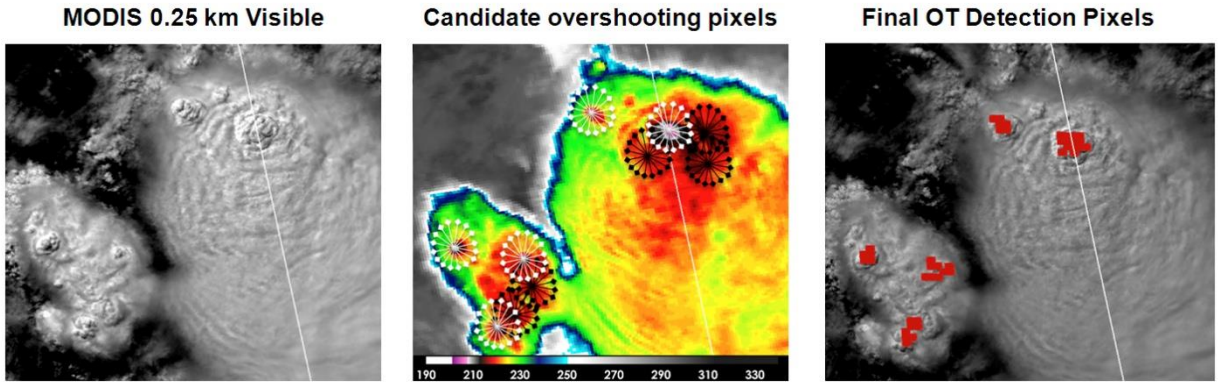

Fig. 2.10. Left: high-resolution visible image of thunderstorm tops from MODIS. Center: Checking the brightness temperatures surrounding candidate overshooting pixels. White pinwheels: candidate overshooting top (OT) significantly colder than surrounding anvil. Black pinwheels: Cloud-top brightness temperature pattern too uniform–not an OT. Right: Overshooting center pixels plus neighbors (red) sufficiently colder than the surrounding anvil to be considered part of an overshooting top. The thin diagonal line across all three images represents an overpass of the NASA CloudSat satellite, used to provide independent information for this case (not discussed further here). Figure courtesy of Kristopher Bedka, NASA Langley Research Center.

The probability of detecting an overshooting top using MODIS data to simulate the GOES-R ABI instrument is 75%. Using lower-resolution GOES-12 data, the POD is 58%. The false alarm rates are, respectively, 16% and 18%.

Bedka et al. (2010) have used EDR turbulence reports to examine the prevalence of turbulence at various distances from overshooting tops as identified in the algorithm. Figure 2.11, from Bedka et al. (2010), shows the frequency of turbulence in three categories (light or greater, moderate or greater, and severe) in the vicinity of cold thunderstorm tops, with and without overshooting turrets. For this case, note that overshooting tops were determined from GOES-12 data. At least within about 20 km of cold anvil tops, the frequency of turbulence reports in any of the three categories is higher when an overshooting top is present than when it is not. This is not surprising because a storm with an overshooting top is much more likely to be severe than one without.

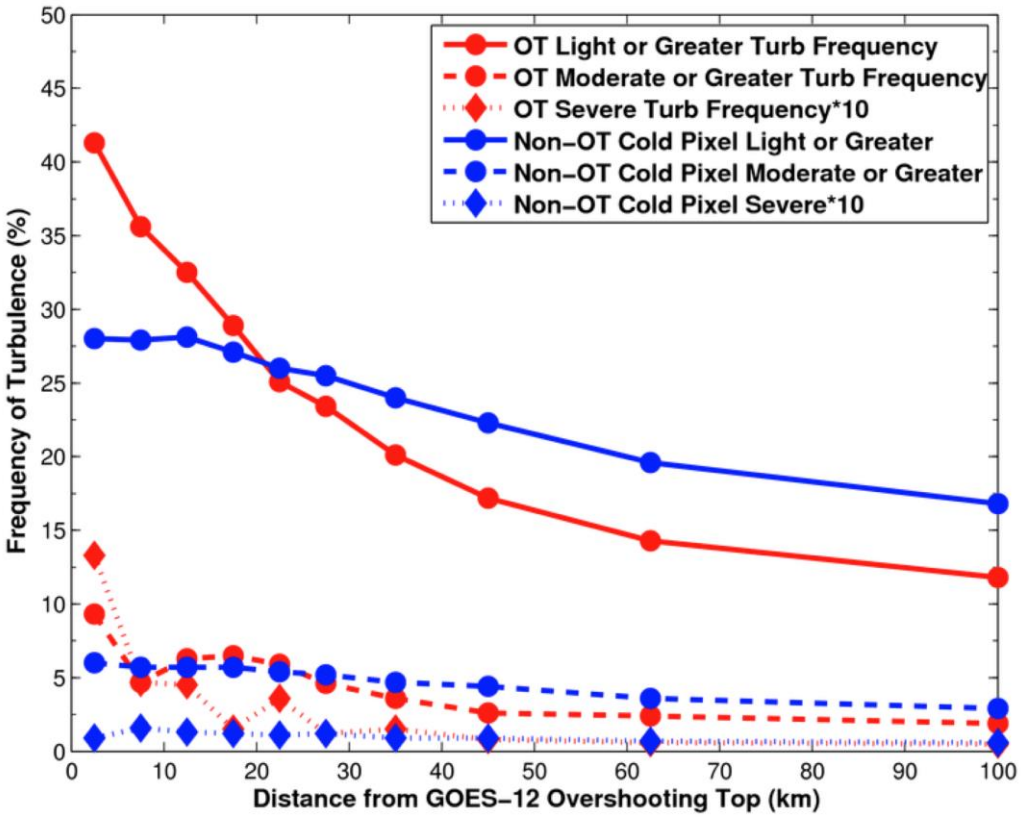

Fig. 2.11. The frequency of turbulence in three different categories inferred from EDR observations at varying distances from overshooting tops (OT) and non-overshooting cold pixels (non-OT) as identified from GOES-12 images from April to September 2005–2008 over the eastern U.S. The frequency of severe turbulence (a rare event) is multiplied by ten so that variability in the two dotted curves may be discerned. Reprinted with permission of the American Meteorological Society (AMS).

The overshooting top product, along with others, is available in real time at http://cimss.ssec.wisc.edu/snaap/convinit/quicklooks/ . Daily verification of this product occurs at http://cimss.ssec.wisc.edu/snaap/convinit/validation/ . One can choose a specific date, as in Figs. 2.12a and b, for 8 June 2010. The locations where overshooting tops were diagnosed for the 24 hours ending at 1200 UTC 09 June 2010 are shown in Fig. 2.12a. Fig. 2.12b, a plot of severe weather reports from the Storm Prediction Center, shows that most locales where overshooting tops occurred also experienced some kind of severe weather.

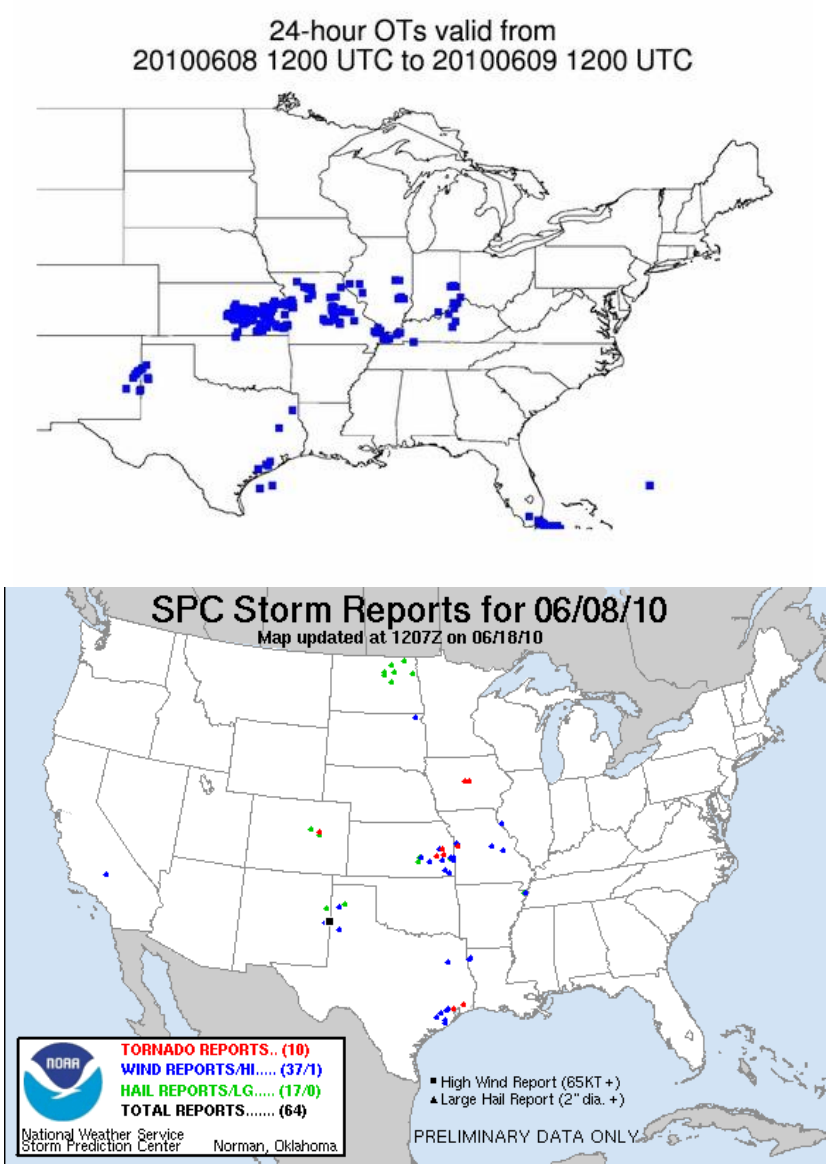

Fig. 2.12. Overshooting tops identified with MODIS satellite data (top) and their correlation with severe storm reports from NOAA's Storm Prediction Center (bottom) on 08 Jun 2010.

2.1.6.4 Mountain Wave Turbulence

This survey came across two algorithms that attempt to diagnose mountain wave turbulence. These algorithms are beyond proof of concept but not yet ready for operations. The first algorithm examines GOES-12 water vapor images for evidence of gravity waves. Gravity waves appear as regular undulations in brightness temperature, caused by the regular up down

motions in a layer of statically stable air. Cross-mountain flow is a frequent cause of gravity waves. Not all gravity waves produce turbulence, but they do affect altimeter settings of aircraft. Should the gravity waves break, for example, in the lee of mountain ranges, severe or even extreme turbulence can result.

The algorithm begins by searching for wave-like patterns in the water vapor image. It filters out wave-like patterns that do not lie across the wind, for example, clouds streets that accompany cold advection and lie parallel to the wind. Then it assigns a score from 0.0 to 2.0 to the patterns it finds. Higher scores indicate greater regularity in wave spacing and higher amplitude between peaks and troughs in brightness temperatures.

Figure 2.13 (left) shows a water vapor image over Colorado. The algorithm has outlined areas in black where it has discovered wavelike patterns that lie roughly perpendicular to the wind direction (known from independent information). At right, the outlined areas, where turbulence is possible, are color-coded according to the assigned score, ranging from 0.0 to 2.0. The higher the score, the more likely is turbulence.

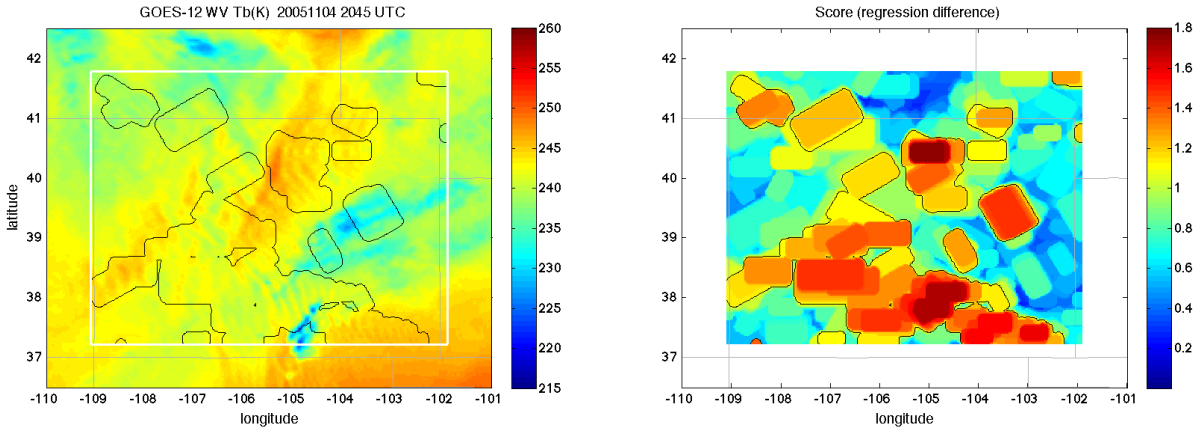

Fig. 2.13. At left, thin black contours enclose areas where wave-like patterns have been detected in the GOES-12 water vapor image. At right, regions where gravity waves have been detected are color-coded according to a score. As the score increases, turbulence becomes more likely.
Courtesy of Tony Wimmers.

Figure 2.14 shows hundreds of EDR turbulence reports from commercial aircraft superposed on the water vapor image (left) and the color-coded areas where turbulence is likely. Most turbulence reports fall within the higher risk areas (right); the correlation is encouraging but still far from perfect.

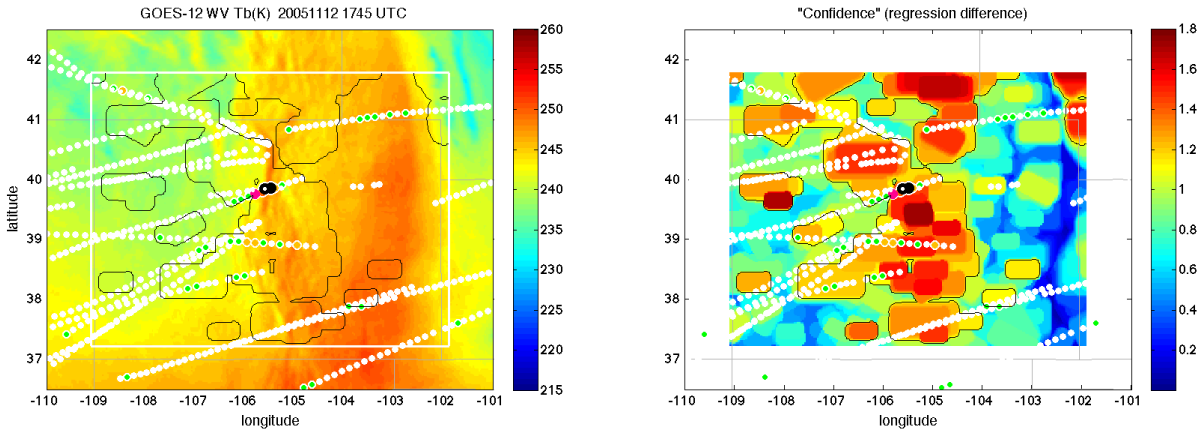

EDR turbulence intensity

Light:	.05 white, .15 green
Moderate:	.25 light orange, .35 dark orange, .45 magenta
Severe:	.55 red, ≥ .65 black

Fig. 2.14. Same as Fig. 2.13 except that automated EDR turbulence reports are superposed on the water vapor image (left) and on the diagnostic turbulence image (right). The colors of the dots indicate the observed intensity of turbulence, as indicated at bottom. The EDR values range from 0.05 to greater than 0.65. Courtesy of Tony Wimmers.

A second algorithm tries to detect mountain wave turbulence in the lee of major mountain ranges. It essentially computes the dot product between the brightness temperature (BT) gradient in a water vapor image and the terrain gradient. Mountain waves typically break just downwind of the crest of a divide in regions of strongly sinking air. With a strong west-to-east flow of stable air across the Front Range of the Rockies, the terrain elevation drops rapidly to the east, but the brightness temperature increases rapidly to the east, so a strongly negative dot product of terrain gradient and BT gradient would suggest turbulence. EDR reports indicate that mountain-wave turbulence is concentrated near the area of strongest downslope winds.

GOES-R water vapor images will have finer resolution than GOES-12 images; this should permit improved detection of gravity waves and mountain waves.

2.1.7 Turbulence Inferred from Doppler Radar Data

More than half of turbulence-related aviation accidents are associated with convective clouds, either within them or nearby. The network of more than 130 WSR-88D (Doppler) radars in the U.S. provides good but not always continuous coverage at 10,000 ft above ground and easily detects thunderstorms. The obvious utility of Doppler radar data is to identify regions of high reflectivity and/or strong rotation for the issuance of severe thunderstorm or tornado warnings. But the standard output of a Doppler radar is also useful for detecting turbulence.

Pilots routinely avoid thunderstorms because of the multiple hazards they pose, including turbulence, icing, hail, and lightning, but sometimes they experience moderate or greater turbulence in convective clouds that are not obviously thunderstorms but still have sufficient reflectivity to be detected by radar. Doppler radars provide, in addition to returned power and radial velocity, a spectrum width measurement. The latter is the standard deviation of radar velocities measured within a radar pulse volume (at a specific range, azimuth, and elevation angle) and hence gives a local measure of turbulence.

During the past decade, a NEXRAD Turbulence Detection Algorithm (NTDA) has been developed and tested at the National Center for Atmospheric Research (NCAR). (NEXRAD stands for NEXt generation RADar, the WSR-88D radar deployed by NWS). This algorithm complements the GTG product because it is specific to in-cloud turbulence, to which GTG does not apply. The best description of this algorithm is in Williams et al. (2006), which is summarized in Appendix B, Section B2.1.7.

NCAR field-tested the NTDA in 2002 using a NASA B-757 research aircraft equipped with an early version of the EDR software described in Section 2.1.1. The B-757 gathered EDR measurements in convective clouds, and NCAR compared these with EDR values computed after the fact with the NTDA software using archived 88-D data. NCAR conducted more field tests from 2005 to 2007 using sixteen 88-D radars in the upper Midwest and EDR reports from United Airlines flights in the vicinity. The outcome of these field tests was sufficiently favorable that the FAA funded further development of the NTDA with the goal of having it eventually run at the Aviation Weather Center. A few improvements in data processing and quality control have been added since the Williams et al. (2006) description discussed above, including a confidence factor related to the presence of ground clutter. NTDA is *not* currently included in the suite of NEXRAD algorithms available at NWS offices on AWIPS.

An example of the NTDA product appears in Fig. 2.15. At top is an 88-D reflectivity mosaic for the CONUS; at bottom is a corresponding mosaic of EDF values calculated with the NTDA.

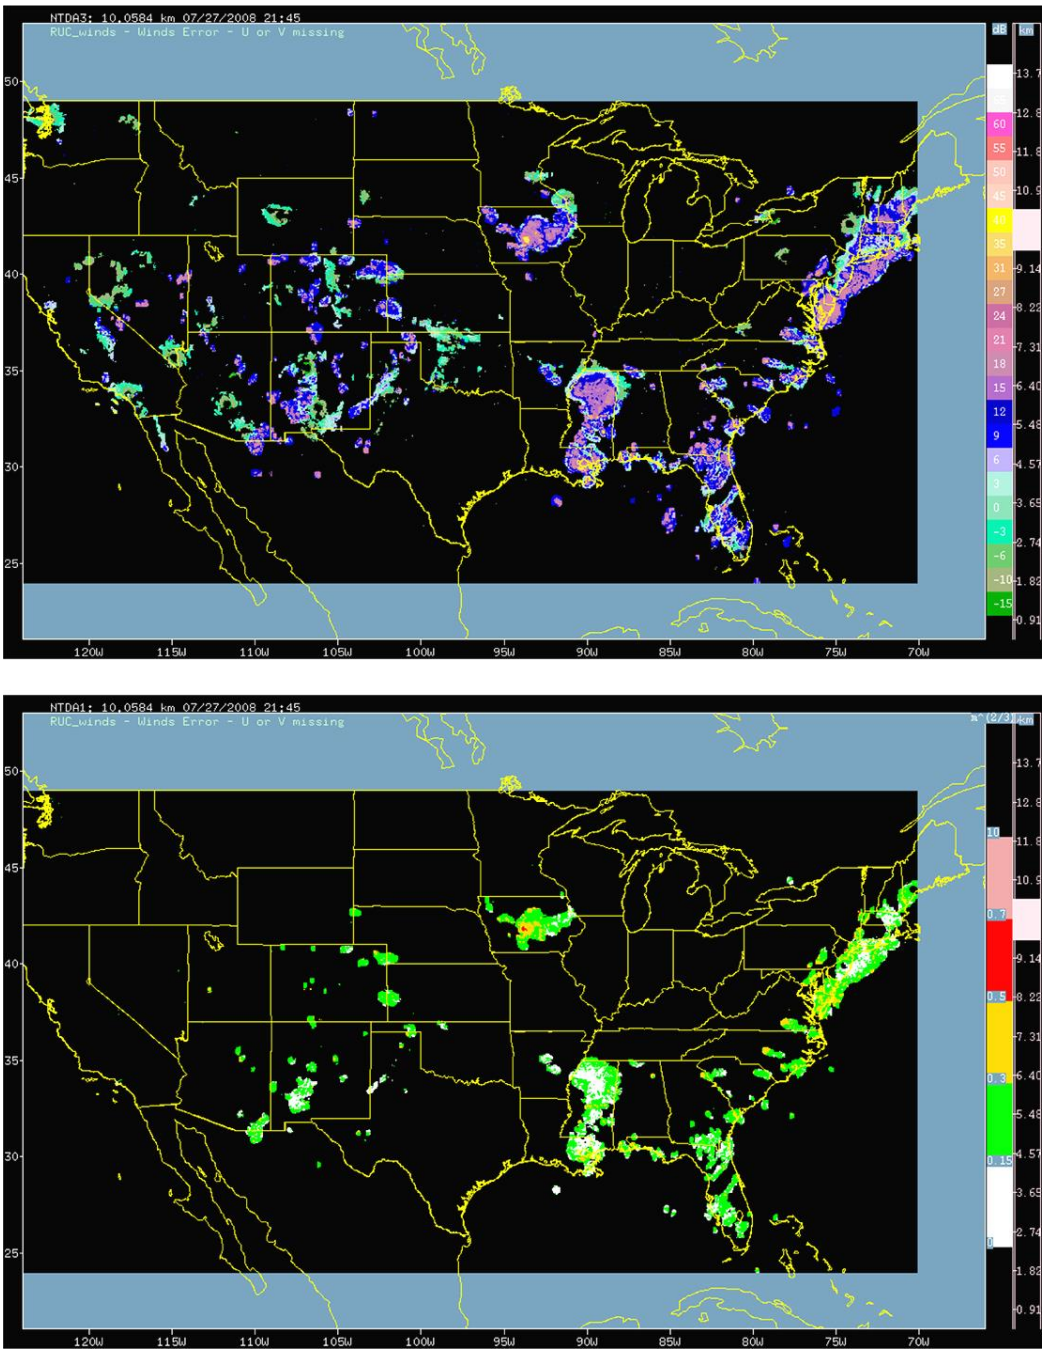

Fig. 2.15. Top: a national mosaic of reflectivity (dBZ) from WSR-88D radars valid at 10-km altitude for 27 July 2008 at 2145 GMT. The color code for reflectivity is given at right. Bottom: the corresponding EDR ($\varepsilon^{1/3}$) mosaic computed with the NDTA. The colors at right correspond to EDR values from 0.0 to 1.0. White indicates smooth air. Green, yellow, red, and coral colors indicate light, moderate, severe, and extreme turbulence, respectively. Note that moderate turbulence is indicated at some locations where reflectivity is less than 30 dBZ.

2.1.8 TP System

Originally called the Turbulence Plot System because it was designed to warn pilots, dispatchers, and other airline personnel of en route turbulence, this system now issues messages about a number of aviation hazards: clear-air turbulence, mountain wave activity, thunderstorms, tropical storms, wind shear, volcanic ash, icing, elevated levels of ozone, and geomagnetic disturbances that affect aircraft communications. Meteorologists at Delta Airlines (which now includes the former Northwest Airlines) developed and continue to generate these products. They are intended to be part of an organization's EWINS, an official FAA sanctioned program mentioned on Section 1.3.2.2. The products are mentioned here because some of them are used in ARTCCs. (They are used widely by dispatchers at several airlines.) Messages pertain to hazards that exist or have the potential to develop. They are called *advisory* if the intensity of the phenomenon lies below the threshold at which safety is threatened, *alert* if the intensity is such that avoidance is recommended, and *avoid* if avoidance is required.

2.2 Icing

This section covers ice accumulation on aircraft surfaces during flight through clouds and/or stratiform precipitation. It does not specifically cover icing in thunderstorms, which is presumed to be severe. The thunderstorm hazard is covered in Section 2.3. Aircraft de-icing on the ground prior to takeoff is covered in Section 3.2

Liquid water is common in the atmosphere at temperatures below 0^{o}. It is most common in the form of cloud droplets (diameter 10–50 µm: 1000 µm = 1 mm). Less common but more dangerous are supercooled large drops (SLDs) with diameter > 50 µm. SLDs occur in the form of freezing drizzle (diameter between 0.05 and 0.5 mm) or freezing rain (diameter > 0.5 mm). Though supercooled droplets can exist down to temperatures approaching $-40^{o}C$, icing occurs far more often between $0^{o}C$ and -20^{o} than at lower temperatures.

Icing occurs when drops of any size strike the skin of the aircraft; this is most likely on leading edge surfaces. The drops freeze following impact. Cloud droplets freeze instantly, producing a frost-like coating called rime ice. SLDs may not freeze instantly. If the supercooled liquid water content is high enough, some of the water may flow rearward on the aircraft to surfaces not protected by de-icing equipment and freeze there, where only melting or evaporation can remove it.

The accumulation of ice on airfoils can quickly lead to problems such as greatly increased drag, thereby decreasing the air speed and increasing fuel consumption, but far more serious problems can follow. The maximum lift decreases; the stall speed increases. The maneuverability, even the controllability of the aircraft can be affected.

The Aeronautical Information Manual (AIM), Part 7-1-21 defines four categories of icing that are to be used in transmitting pilot reports (PIREPs).

Trace. Ice becomes perceptible. The rate of accumulation is slightly greater than sublimation. De-icing/anti-icing equipment is not utilized unless encountered for an extended period of time (over 1 hour).

Light. The rate of accumulation (¼ inch in 15-60 minutes) may create a problem if flight is prolonged (over one hour) in this environment. Occasional use of de-icing/anti-icing equipment removes/prevents accumulation. It does not present a problem if de-icing/anti-icing equipment is used.

Moderate. The rate of accumulation (¼ inch in 5–15 minutes) is such that even short encounters become potentially hazardous and use of de-icing/anti-icing equipment or flight diversion is necessary.

Severe. The rate of accumulation (¼ inch in less than 5 minutes) is such that de-icing/anti-icing equipment fails to reduce or control the hazard. Immediate flight diversion is necessary.

In conflict with this, the FAA approved a different set of definitions in Federal Register Document No. FAA-2000-8560, Icing Terminology, issued on 7 May 2003. Note the omission of "trace" and the addition of "heavy" in the four categories.

Light. The rate of ice accumulation requires occasional cycling of manual de-icing systems[2] to minimize ice accretions on the airframe. A representative accretion rate for reference purposes is ¼ inch to one inch (0.6 to 2.5 cm) per hour[1] on the outer wing. The pilot should consider exiting the condition[3].

Moderate. The rate of ice accumulation requires frequent cycling of manual de-icing systems[2] to minimize ice accretions on the airframe. A representative accretion rate for reference purposes is 1 to 3 inches (2.5 to 7.5 cm) per hour[1] on the outer wing. The pilot should consider exiting the condition as soon as possible[3].

Heavy. The rate of ice accumulation requires maximum use of the ice protection systems to minimize ice accretions on the airframe. A representative accretion rate for reference purposes is more than 3 inches (7.5 cm) per hour[1] on the outer wing. Immediate exit from the conditions should be considered[3].

Severe. The rate of ice accumulation is such that ice protection systems fail to remove the accumulation of ice and ice accumulates in locations not normally prone to icing, such as areas aft of protected surfaces and any other areas identified by the manufacturer. Immediate exit from the condition is necessary.

Footnotes:

[1] These rates can be measured by a suitable icing rate meter.

[2] It is expected that de-icing or anti-icing systems will be activated and operated continuously in the automatic mode, if available, at the first sign of ice accumulation, or as directed in the Airplane Flight Manual. Occasional and frequent cycling refers to manually activated systems.

[3] It is assumed that the aircraft is approved to fly in the cited icing conditions. Otherwise, immediate exit from any of these intensity categories is required by regulations.

[4] Severe icing is aircraft dependent, as are the other categories of icing intensity. Severe icing may occur at any ice accumulation rate when the icing rate or ice accumulations exceed the tolerance of the aircraft. Icing certification implies an increased tolerance to icing intensities up through heavy.

Two icing products to be considered later in this section, the Current Icing Potential (CIP) and the Forecast Icing Potential (FIP), generate an icing severity index (IceSev) with several thresholds:

$$\text{IceSev} < 0.010 \text{ implies no icing.}$$
$$0.010 \leq \text{IceSev} < 0.175 \text{ implies trace icing}$$
$$0.175 \leq \text{IceSev} < 0.375 \text{ implies light icing.}$$
$$0.375 \leq \text{IceSev} < 0.700 \text{ implies moderate icing.}$$
$$0.700 \leq \text{IceSev} \text{ implies heavy icing.}$$

As of this writing, a proposal for standardization is under consideration: 1) Make the reportable icing intensities for PIREPs light, moderate, heavy, and severe. 2) For icing forecasts issued by AWC, use light, moderate, and heavy. 3) For CIP and FIP products, use light, moderate, and heavy. 4) Revise Federal Meteorological Handbook 12 to permit coding for heavy icing in PIREPs. 5) In all of this, use the definitions approved by the FAA (above) in May 2003.

2.2.1 Observations of Icing

Pilot reports of icing are available at http://aviationweather.gov/adds/icing/ . One can click on any of seven regions to see the current reports. Figure 2.16 is a sample from the North Central States. Commercial aircraft on ascent and descent normally pass quickly through the altitudes where icing is most likely to occur. At flight altitudes, outside air temperatures are usually too low for icing to occur. General aviation pilots usually fly at lower altitudes, where icing is more prevalent.

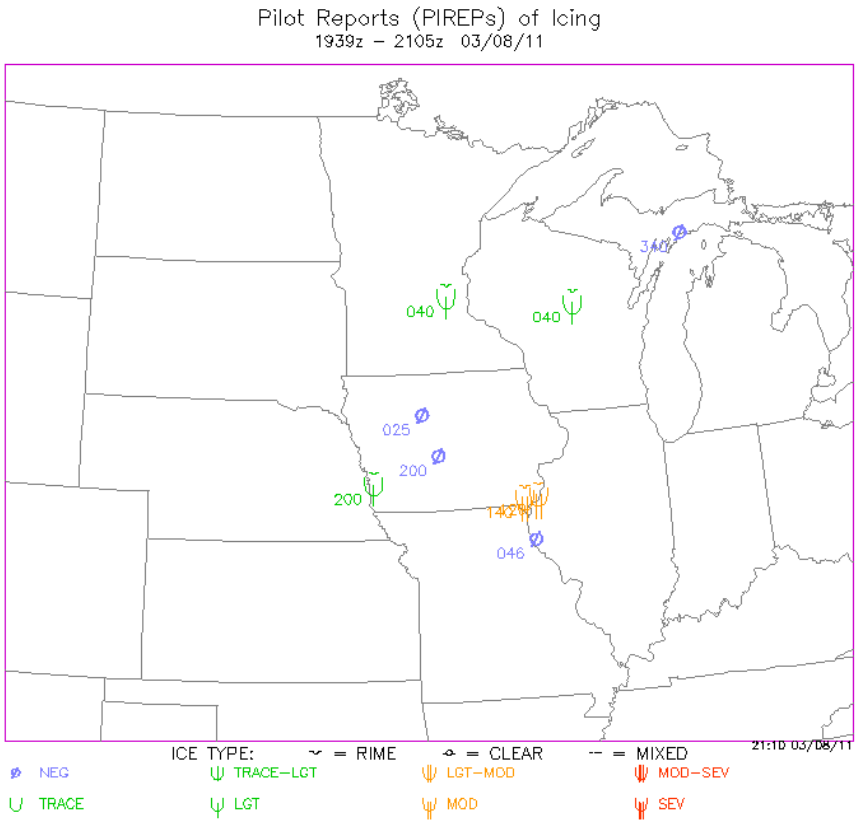

Fig. 2.16. Pilot reports of icing between 1939 and 2105 UTC 08 March 2011. The intensity of reported icing is given by the symbols and colors at bottom.

2.2.2 AIRMETs for Icing

AIRMETs for icing are issued by the AWC or the AAWU whenever moderate icing is expected. The top and bottom of each icing layer is designated along with the areal extent. Icing forecasters rely heavily on mesoscale model output from the RUC and NAM models; they check the hydrometeor fields for liquid water at temperatures below freezing. Layers of supercooled liquid water (SLW) thicker than 5000 ft catch their attention. They also look at the percentage of model levels below 30,000 ft that contain SLW. Experience indicates that icing tends to be worst near ridgetops in the Pacific Northwest. Large supercooled drops often concentrate in the compression zone where mountain blocking causes the greatest uplift. Over the Great Lakes, the worst icing occurs near cloud tops when cold air passes over warm water. Some tops of lake-effect thunder-snowstorms can lie as low as 10,000 ft. Electrically charged snowflakes can stick to wings, not just SLW.

Here is a text example of an icing AIRMET. Text is still used because many general aviation pilots do not have ready access to the Internet when they get their briefings.

```
WAUS43 KKCI 082045
CHIZ WA 082045
AIRMET ZULU UPDT 3 FOR ICE AND FRZLVL VALID UNTIL 090300
AIRMET ICE...ND SD MN
FROM 70WNW INL TO FSD TO 70SW RAP TO 50NNW ISN TO 70WNW INL
MOD ICE BLW 080. CONDS CONTG BYD 03Z THRU 09Z.
OTLK VALID 0300-0900Z
AREA 1...ICE SD NE KS MN IA
BOUNDED BY ODI-DSM-PWE-ICT-GCK-70WNW ANW-FSD-ODI
MOD ICE BLW 120. CONDS CONTG THRU 09Z.
AREA 2...ICE ND SD MN WI
BOUNDED BY 60N MOT-20NNW INL-DLH-ODI-FSD-70WNW ANW-60N MOT
MOD ICE BLW 080. CONDS CONTG THRU 09Z.
FRZLVL...RANGING FROM SFC-095 ACRS AREA
 MULT FRZLVL BLW 050 BOUNDED BY 40NE TVC-YVV-30SE ECK-DXO-
 40ESE BAE-50SSW GRB-40NE TVC
 SFC ALG LBL-40SE GCK-70WSW FOD-40E MCW-30NE BAE-20SW DXO
 040 ALG 40SSW ICT-50NNE MCI-30W GIJ-40NE FWA
 080 ALG 30NW RZC-20NE SGF-50SSE COU-20ESE LOZ-40W HMV
```

The above AIRMET was issued at 2045 GMT on 08 March 2011 and was valid until 0300 on 09 March. The first part is given in plain language: In an area bounded by points 70 miles WNW of International Falls, MN, Sioux Falls, SD, 70 miles SW of Rapid City, SD, 50 miles NNW of Williston, ND, and back to 70 miles WNW of International Falls, moderate icing is expected below 8000 ft. These conditions are expected to continue through 0900 UTC. At the end of the message is information about the freezing level.

Easier to interpret is the G-AIRMET product at http://aviationweather.gov/products/gairmet/

shown in Fig. 2.17. This icing product became operational in March 2010. Note the six areas across the country where moderate icing is expected and the upper and lower bounds for each layer in thousands of feet. A third number appearing alone gives the lower bound of a higher icing layer when two layers are present.

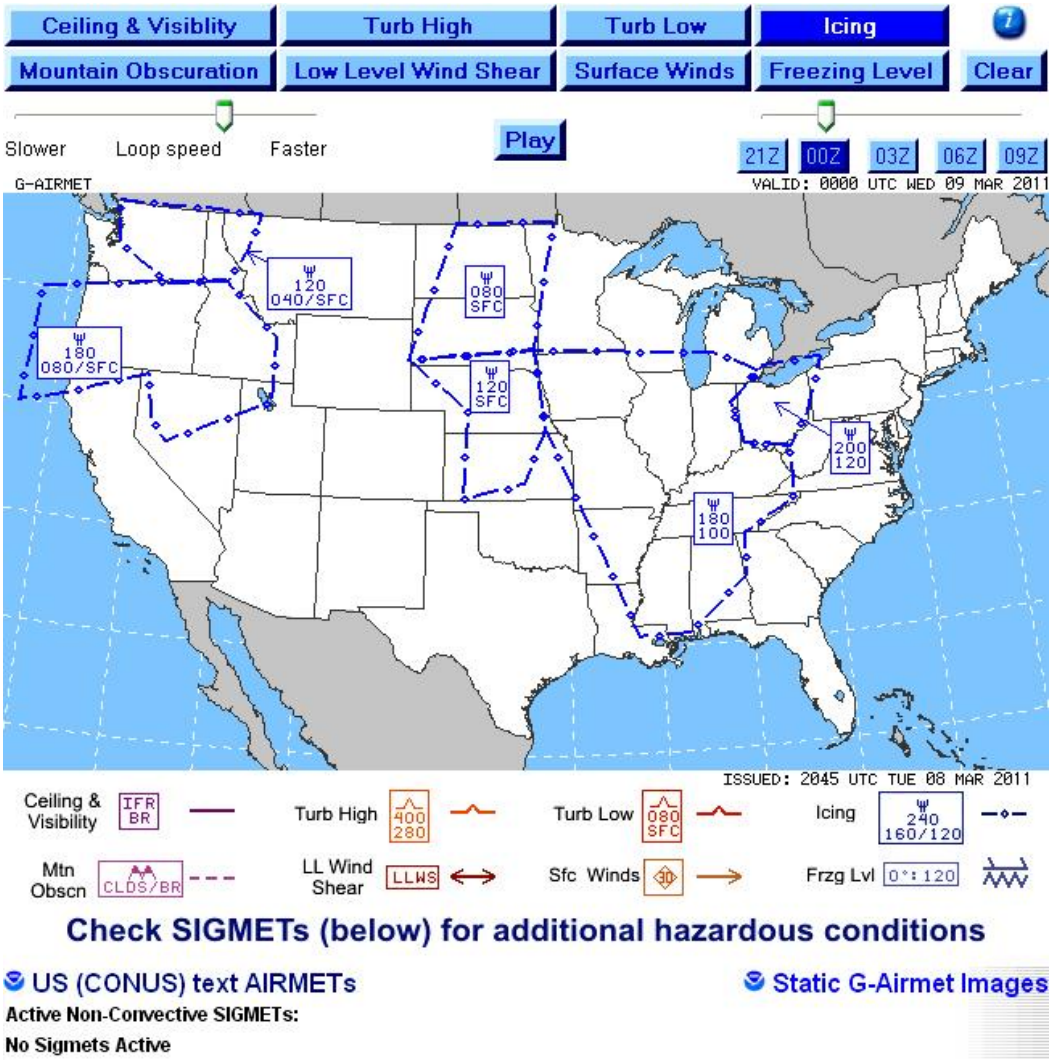

Fig. 2.17. The G-Airmet product for icing, in this case a 3-h forecast. Note that AIRMETs for hazards other than icing are available on the same web page and that one can choose from among several valid times for the forecast. One can hit the “Play” button at top middle to see the forecasts in animation.

Table 2.3 gives some idea of the accuracy of AIRMET icing forecasts during 2010. Compared with the rest of the year, not many icing reports were submitted from June through September. When icing was observed during those months, it was generally predicted less than half the time. Some of the icing reports may have been associated with thunderstorms, to which the icing AIRMET does not apply.

Table 2.3

Verification statistics for 3-h AIRMET icing forecasts for each month of 2010. A "yes" means moderate or greater icing was observed; a "no" means no icing was observed.

Date	YY	YN	NY	NN	1-PODn	PODy
2010-01-01	641	197	227	454	0.303	0.738
2010-02-01	638	251	173	344	0.422	0.787
2010-03-01	319	130	161	427	0.233	0.665
2010-04-01	235	40	125	293	0.120	0.653
2010-05-01	134	40	109	352	0.102	0.551
2010-06-01	53	8	50	207	0.037	0.515
2010-07-01	6	0	10	51	0.000	0.375
2010-08-01	8	3	43	135	0.022	0.157
2010-09-01	61	9	83	251	0.035	0.424
2010-10-01	225	61	120	212	0.223	0.652
2010-11-01	414	103	186	333	0.236	0.690
2010-12-01	594	188	194	360	0.343	0.754
Totals	3328	1030	1481	3419	0.232	0.692

During the winter months of December, January, and February, icing is correctly predicted about ¾ of the time. 1-PODn is the fraction of no-icing observations for which icing was incorrectly predicted. This fraction is less than ¼ except during the winter months.

Verification of icing forecasts is difficult because a pilot is much more likely to report icing than a lack of it. Tropospheric Airborne Meteorological DAta Reporting (TAMDAR) is a system developed by AirDat for collecting and disseminating automated aircraft reports from short-hop regional carriers, which generally fly lower and take off and land more frequently than the larger carriers. Unique with TAMDAR is an icing report, which serves the same purpose as the EDR reports of turbulence: it gives a null report of icing as readily as a positive report. So far, however, TAMDAR data have not been used routinely for verifying icing forecasts.

2.2.3 SIGMETs for Icing

SIGMETs for icing are issued only in response to a report of severe icing. They are not common. The AWC does not forecast severe icing and, because "heavy icing" is not a reportable category of aircraft icing, neither does it forecast heavy icing at the present time. This situation is expected to change soon.

2.2.4 CWAs for Icing

CWSU meteorologists issue Center Weather Advisories for icing only rarely, when either the affected area is less than 3000 square miles or the icing hazard is already apparent but AWC has not yet issued an AIRMET or SIGMET.

2.2.5 Freezing Level Product

The big-picture view of freezing level is available at http://aviationweather.gov/adds/icing/frzg_nav.php . The analysis of freezing level is available every hour from the RUC model, as are forecasts out to 12 h. Fig. 2.18 shows a sample product.

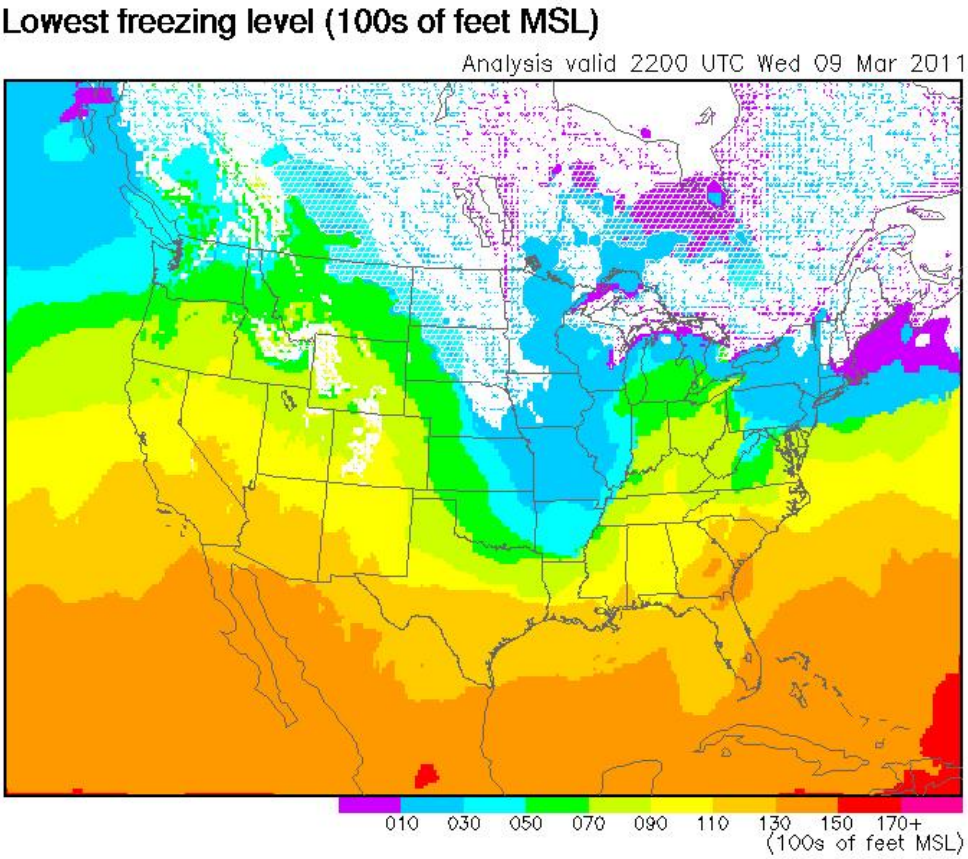

Fig. 2.18. The lowest level at which the vertical temperature profile crosses 0°C. A white background indicates that the entire profile lies below freezing. Colors indicate the height above mean sea level (MSL) in hundreds of feet of the lowest 0°C crossing. White hatching over colors indicates a surface temperature below freezing and multiple 0°C crossings above ground. The noisy, colored scratching on a white background is a defect in the product. The entire vertical column was below freezing in this area.

Another way of diagnosing the freezing level is to plot soundings available from model output, rawinsondes, or aircraft that relay temperature (sometimes moisture) data while on ascent or descent. The great majority of automated aircraft data in the U.S. come from MDCRS, the Meteorological Data Collection and Reporting System, funded jointly by the U.S. government and seven participating airlines (American, Delta, FedEx, Northwest, Southwest, United, and United Parcel Service) and operated by Aeronautical Radio, Inc. A significant fraction of reports comes from TAMDAR, mentioned above. Both MDCRS and TAMDAR reports are proprietary, that is, not available to the public, but the government does use them in generating aviation weather products.

2.2.6 Inferring Icing from Vertical Profiles of Temperature and Moisture

Sounding data show details in the vertical temperature profile not discernible in Fig. 2.18. If the sounding includes dewpoint as well as temperature, icing is possible in layers where the difference between temperature and dewpoint is less than 2°C (this suggests clouds) *and* the temperature lies between -20° and 0°C. Soundings from rawinsondes, aircraft on ascent or descent, or generated from the output of operational prediction models are available at http://rucsoundings.noaa.gov/plot_soundings.cgi . For help in plotting the soundings of choice, either observed or predicted, by source, location, date and time, go to http://amdar.noaa.gov/java/help.html and see the information under "Help for Soundings."

Figure 2.19 shows a 3-h RUC prediction of temperature, dewpoint, and wind profiles valid for Baltimore-Washington International (BWI) Airport at 1700 GMT on 10 March 2011 in a region where aircraft icing was frequent. The data are plotted on a Skew-T/Log-p diagram. At upper left is a hodograph, which traces the path the tip of the wind vector would follow from the surface up to above 12 km. Kilometers are labeled 3, 6, 9, and 12 on the hodograph curve. At right, corresponding to the hodograph, is a vertical stack of winds, plotted in the conventional way (with barbs for 5 and 10 knots and a flag for 50 knots).

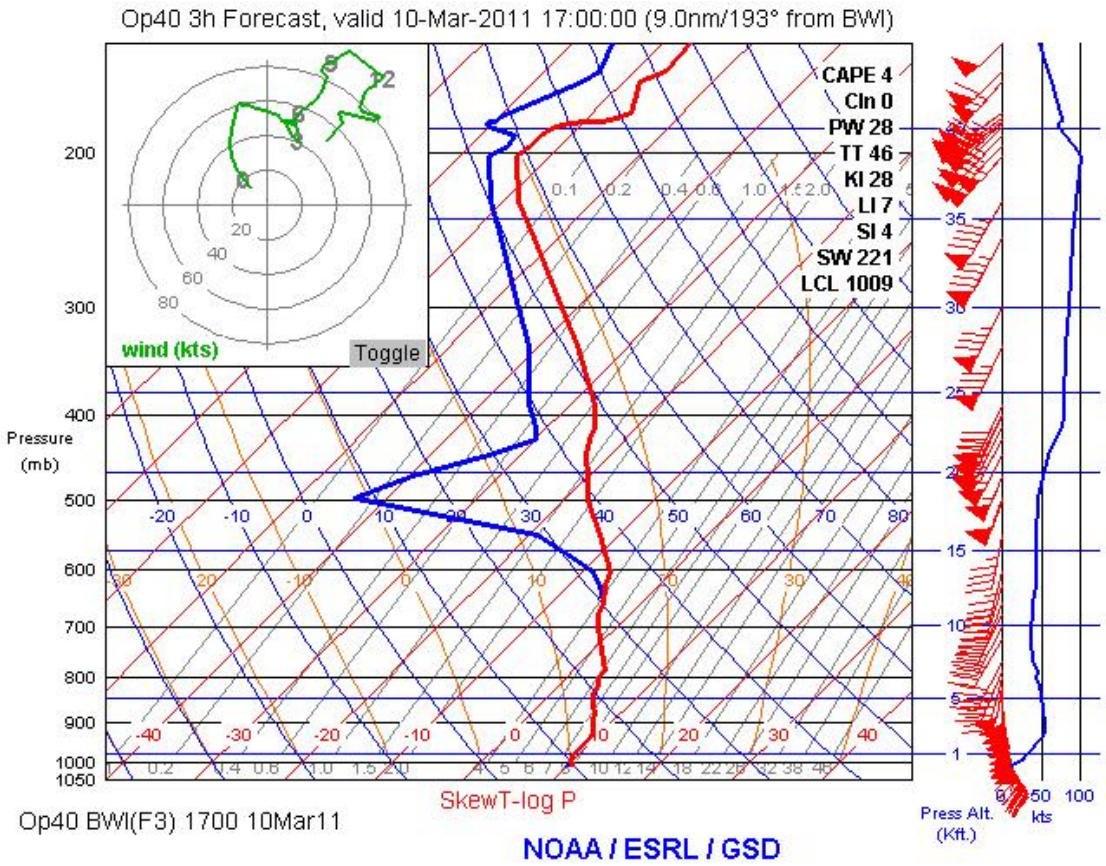

Fig. 2.19. A 3-h forecast of temperature, dewpoint, and wind profiles from the RUC model, plotted on a Skew-T/Log-p diagram and valid at Baltimore-Washington Airport at 1700 GMT 10 March 2011. Further details are in the text.

Figure 2.20 shows a sounding measured by an aircraft on descent into BWI. The aircraft landed at 1701 GMT, a minute later than the valid time of the 3-h forecast in Fig. 2.19. Temperature, moisture, and wind data are plotted beginning from a point 230 nautical miles (nm) and 49 min out, and a flight level near 35,000 ft. For distances more than 100 nm from BWI, the data are plotted in gray. Closer to BWI, the temperature is plotted in red and the dewpoint in blue. At upper left, the position of the aircraft within 100 nm is shown as it descends on an ENE heading until 20 nm out. The fraction of aircraft soundings including moisture profiles is still small but growing. A sensor called WVSS2 (Water Vapor Sensing System, Version 2) recently met World Meteorological Organization (WMO) requirements for accuracy and, as of January 2011, is deployed on 25 United Parcel Service aircraft and 9 Southwest aircraft. Together, these aircraft are producing more than 200 soundings per day.

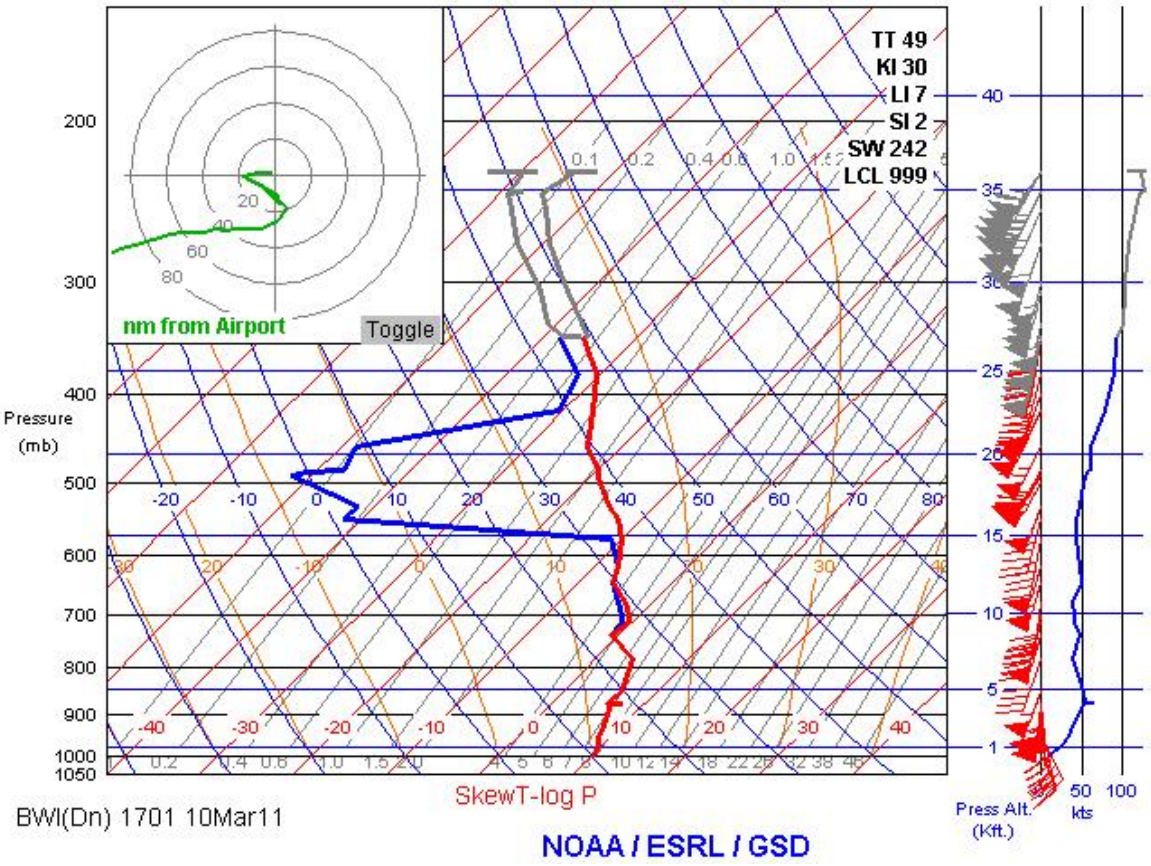

Fig. 2.20. Similar to Fig. 2.19, except this sounding came from an aircraft descending into BWI from the WSW and landing at 1701 GMT, 10 March 2011. This aircraft was equipped with WVSS2, a water vapor sensor. See text for details.

Note that both soundings are saturated or nearly so from the surface to near 600 mb, the air is dry from 580 up to 430 mb, and then more moist above that. The freezing level on both soundings is near 720 mb. Icing, if present, would occur between 720 and 580 mb.

The two soundings are superposed in Fig. 2.21. In this case, the 3-h RUC forecast below 600 mb was very good: the temperature was within a degree or so of the observed value, and the air was saturated. The forecast also captured the dry layer rather well. Above about 350 mb, comparisons are risky because the aircraft was more than 100 nm distant from BWI, where the forecast is valid. If horizontal gradients are strong, the agreement may be poor (see, for example, the wind speeds above the 400-mb level, graphed at far right. As noted earlier, AWC meteorologists use model soundings extensively when generating AIRMETs.

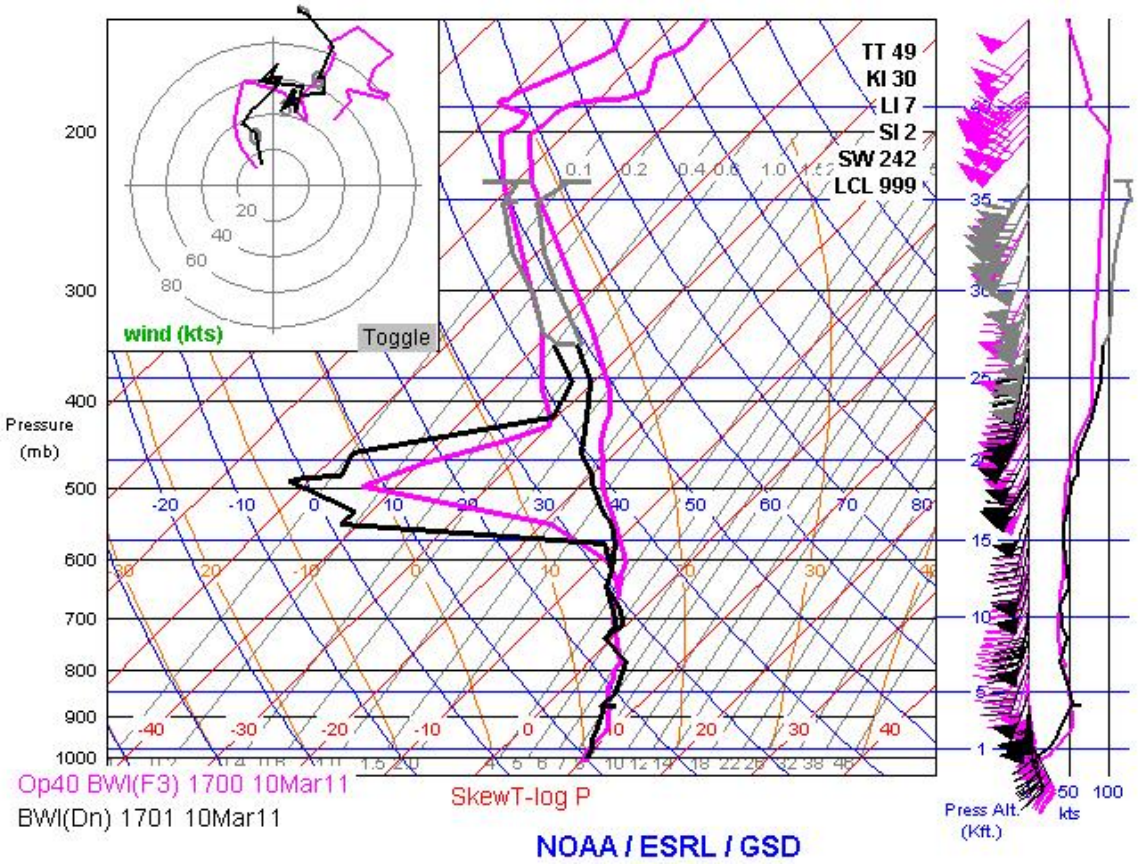

Fig. 2.21. The two soundings from Figs. 2.19 and 2.20 superposed for comparison. The aircraft data are plotted in black and gray, the RUC model data in pink.

2.2.7. Current Icing Product

Like the GTG2 turbulence product, the Current Icing Product (CIP) is a fully automated, unrestricted, supplementary product based upon multiple sources of information relevant to the hazard. CIP was thoroughly described in Bernstein et al. (2005), and the icing potential calculation has changed little since that paper was published. The abbreviated description given here focuses on the flowchart that appeared in Fig. 2 of Bernstein et al. (2005).

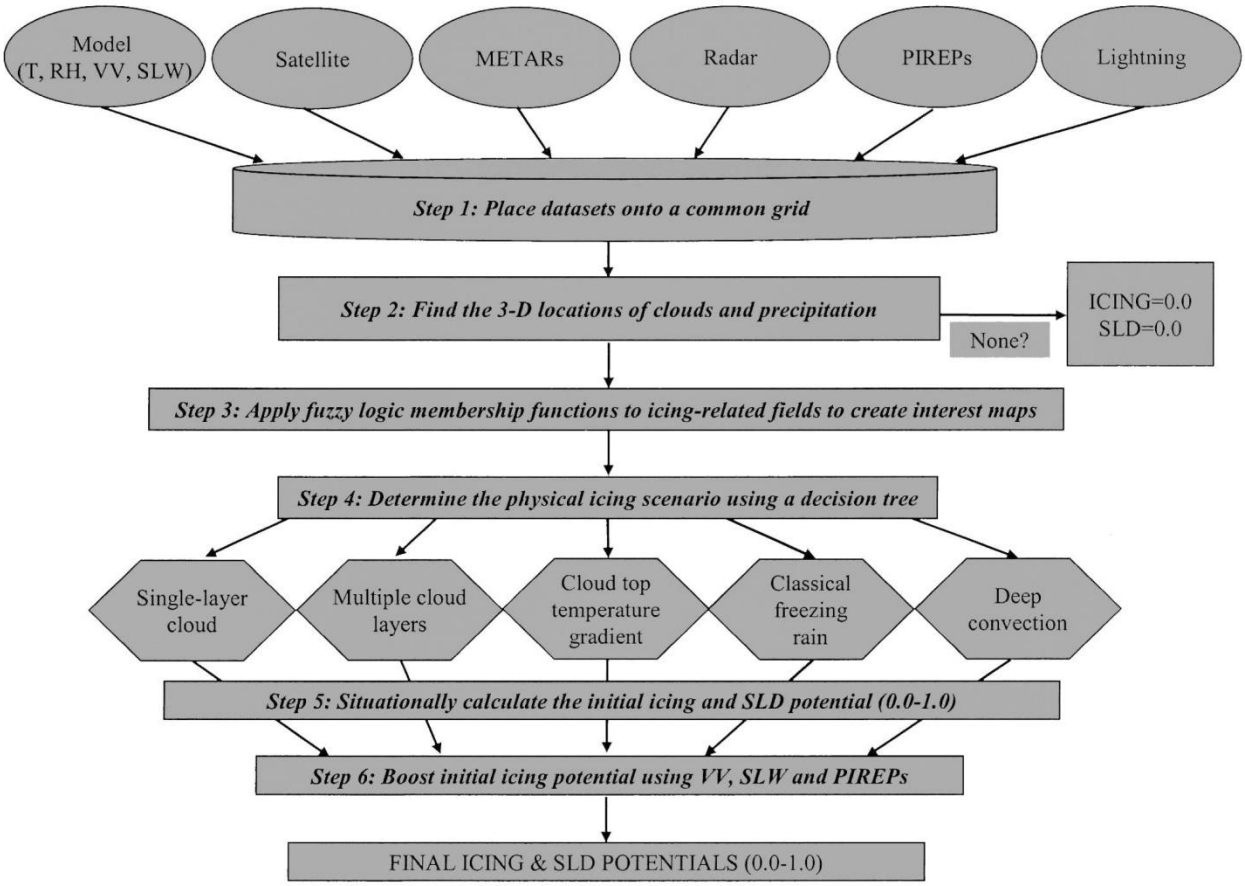

Fig. 2.22. Flowchart for the Current Icing Potential (CIP) and the Supercooled Large Droplet (SLD) Potential, including the information sources and the data processing steps that lead to a scaled estimate of icing and the presence or absence of supercooled large drops. Details are in the text. Reprinted with permission of the American Meteorological Society.

Six information sources feed the CIP algorithm.

- Output from the RUC model, specifically, temperature T, relative humidity RH, vertical velocity VV, and supercooled liquid water (SLW). From a modeling standpoint, icing is expected where T < 0°C but not much less than -20°C; RH is high, indicating clouds; VV indicates upward motion (guarantees continued saturation where clouds are); and SLW is present. Liquid water content is one of the microphysical variables carried in the RUC model.
- GOES satellites view clouds both in the visible and infrared. Knowing where clouds are is immediately useful because icing does not occur outside of clouds. Infrared images show cloud-top temperature (CTT). At CTT = -40°C, ice is certain at cloud top; with increasing CTT, liquid becomes more likely, until, for CTT > 0°C, liquid is certain.

- METARs (a type of surface observation approved by the WMO. Surface observations are useful in several ways. Precipitation type is useful for inferring conditions aloft. Snow indicates that the entire vertical column is below freezing except perhaps near the surface. Graupel is very heavily rimed snow; it indicates supercooled water in the clouds above. Sleet or freezing rain indicates a layer of above-freezing air above the surface cold air. A gentle rain neither confirms nor precludes icing aloft. A thunderstorm always indicates icing aloft. Cloud coverage and height of the ceiling layer are also part of a METAR report. Again, icing does not occur with clear skies.
- In the U.S. 10-cm, WSR-88D Doppler radars indicate in three dimensions where precipitation is forming in clouds and falling to the ground. Sometimes it detects non-precipitating clouds if the hydrometeors are large enough and in sufficient concentrations. The icing algorithm calculates the 25^{th} and 75^{th} percentiles of reflectivity in each model grid box. It currently uses only the scan at the lowest beam elevation.
- PIREPs give direct evidence of icing and its intensity. They are fundamental both for generating the CIP and for verifying icing forecasts.
- Lightning data. The government purchases lightning data from the National Lightning Detection Network (NLDN), operated by Vaisala, a private vendor. These data indicate the time and location of cloud-to-ground lightning strikes and define where thunderstorms are.

Step 1 of the data processing is to reference the above data to the model grid. In the case of 88D or GOES data, multiple values are placed in each grid box, the number depending upon the spatial resolution of the data source. METARS and PIREPs are less dense than the points on a model grid. This information is spread to nearby grid points but given less weight the greater the distance between the report and the grid point. The radius of influence of METARS and PIREPs is thus limited. Lightning strikes are recorded at the grid point if they occur no more than 25 km away.

Step 2 finds the 3-D locations of clouds and precipitation. GOES satellite images at three wavelengths, 0.65 μm (Channel 1, visible), 3.9 μm (Channel 2, shortwave IR), and 10.7 μm (Channel 4, longwave IR) are used extensively to decide whether clouds are present. 3.9 and 10.7 μm are wavelengths at which absorption by atmospheric water vapor is minimal but cloud detection is effective.

In daytime (sun more than 20^{o} above the horizon), the algorithm is looking for a bright surface, such as snow or clouds. The so-called albedo in the GOES visible channel is really the brightness of the scene on a gray scale ranging from 0 (black) to 255 (white). Because the brightness is greatest at solar noon and decreases as solar zenith angle increases, the measured brightness is

divided by the cosine of the solar zenith angle. This "normalized" brightness count must be at least 51 (20% of 255) to diagnose clouds. In addition, at least one other criterion must be met. The other criteria help to distinguish clouds from the surface, and clouds from snow-covered ground. One criterion looks for a METAR observation of broken or overcast skies, another looks for a model surface temperature considerably higher than the 10.7 μm brightness temperature (highlights contrast between cloud-top and ground temperature), and a third looks for a reflectance ≥ 20% at 3.9 μm.

If the sun is below the horizon (night), the visible channel is no longer useful, and the brightness temperatures in Channels 2 and 4 (BT_2 and BT_4) help determine whether clouds are present. Channels 2 and 4 respond differently enough to IR radiation from the surface and clouds that the difference $BT_2 - BT_4$ figures prominently in various tests. Clouds are assumed to be present if any of the following are true:

- $BT_4 < 5^oC$ *and* either $(BT_2 - BT_4) \leq -2.1^oC$ or $(BT_2 - BT_4) \geq 4.1^oC$.
- $BT_4 < -35^oC$ (high, cold clouds inferred except in polar regions)
- $5^oC \geq BT_4 \geq -35^oC$ *and* METAR cloud cover is at least broken
- $20^oC > BT_4 > 5^oC$ *and* $(BT_2 - BT_4) \leq -2.1^oC$

If the sun is near the horizon (early and late in the day when solar zenith angle is greater than 70^o), cloud detection is less certain than at other hours of the day. If any of the following criteria are met, clouds are assumed to be present.

- $BT_4 < -35^o$
- $20^oC \geq BT_4 \geq -35^oC$ *and* METAR cloud cover is at least broken
- BT_4 is more than 15^oC lower than the model temperature at the ground.

Once the areal extent of clouds is determined, it remains to estimate the height of cloud top and cloud base. BT_4 provides a good estimate of cloud-top temperature because at 10.7 μm there is very little absorption of the IR radiation between cloud top and the satellite. To get cloud-top height, the algorithm examines model temperatures at each level starting from the top and working down. As soon as it finds a model temperature higher than BT_4, the search stops, and the cloud height is assigned as the height of the next model level above. Icing is very unlikely above this height.

The algorithm searches METAR reports for cloud base height. It interpolates cloud base heights to points on the model grid, giving more influence to reports that are nearby, and less influence to more distant reports. In no event can a reported cloud base influence the grid point estimate if it is more than 125 km away.

It is important to know the type of precipitation and whether it is falling, and the type of precipitation. Six types are of interest: drizzle, rain, freezing drizzle, freezing rain, ice pellets (sleet), and snow. Any of the first five precipitation types do not preclude icing below cloud base. If no precipitation is falling, the air space below cloud base is considered free of icing. For grid points that have no METAR reports nearby, the 88-D radar is consulted for the presence of echoes. Because the radar does not clearly distinguish among precipitation types and it is difficult to tell whether precipitation forming in clouds is reaching the ground, icing cannot be ruled out in the sub-cloud layer when echoes are present.

Step 3 is to develop *interest maps* for icing-related parameters. As used here, an interest map indicates the relative likelihood, on a scale of minus one to one, of icing for given values of a single parameter. A value of zero indicates that the parameter has no interest for icing. A value of one does *not* indicate that icing is certain, but rather, at the parameter value for which the interest map has a value of one, icing occurs the most often. Similarly, a value of minus one does not totally negate the presence of icing; rather, it indicates that the parameter is usually associated with no-icing cases. Figure 2.23, from Bernstein et al. (2005), illustrates interest maps for temperature (T_{map}), temperature in convective situations ($T_{map\text{-}convective}$), cloud-top temperature (CTT_{map}), relative humidity (RH_{map}), and vertical velocity (VV_{map}).

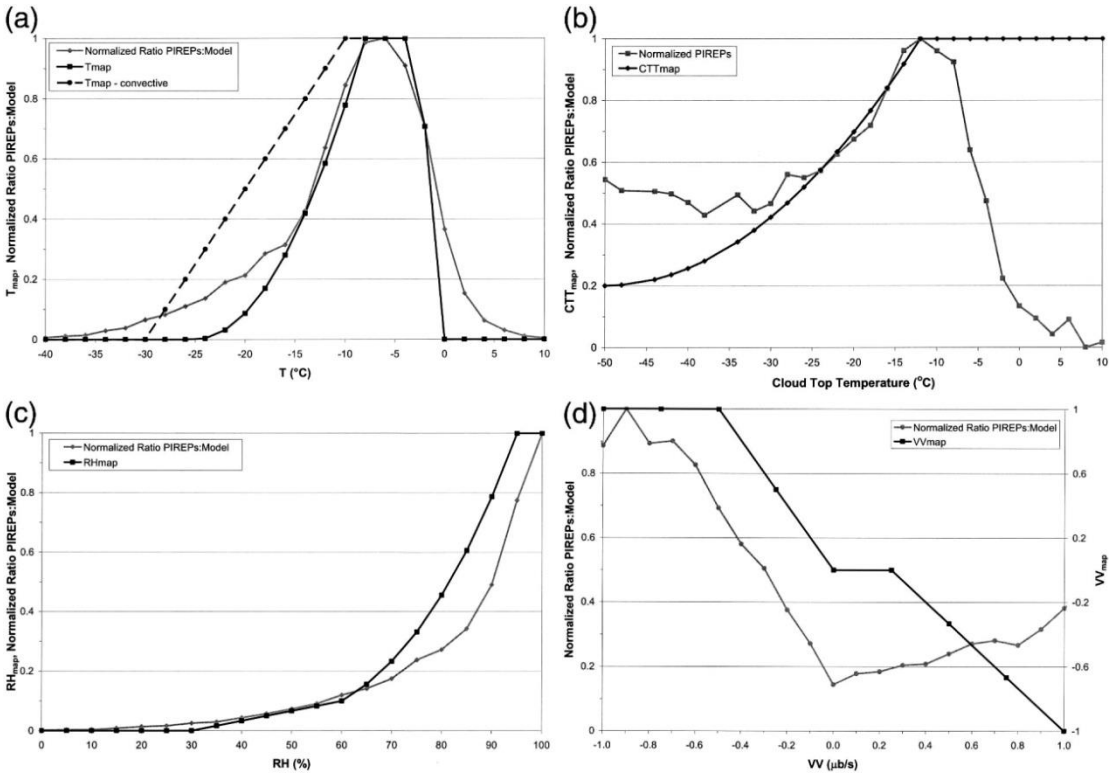

Fig. 2.23. Interest maps for various parameters used in diagnosing icing (bold curves) plus normalized icing PIREP ratio (see text for definition—thin gray curves). (a) T_{map} for non-convective situations and $T_{map\text{-}convective}$. (b) RH_{map}. (c) CTT_{map}. (d) VV_{map}. Used with permission of the American Meteorological Society.

The normalized icing PIREP ratio appears in all four panels; it indicates the relative frequency of icing PIREPs as a function of the value inferred by the model for T, RH and VV, and observed by satellite for CTT. The normalized PIREPs ratio in Fig. 2.23a peaks between -6° and -8°C, which says that, given an icing PIREP, the most likely temperature lies in this two-degree range. At successively lower temperatures, an icing report becomes less likely. Homogeneous nucleation (spontaneous freezing of cloud droplets in the absence of ice nuclei) occurs at -40°C, so it is no surprise that icing reports at still lower temperatures are not received. At temperatures above -6°C, the frequency of icing falls off rapidly, but not to zero at 0°C, as one might expect. In such cases, either the temperature analyzed by the model is too high, or the flight level reported by the pilot is incorrect. (The model temperature is computed from the flight level.)

The interest maps for temperature, with and without convection, are also shown in Fig. 2.23a. These are the curves employed by the icing algorithm. The T_{map} curve lies below the normalized PIREP ratio at low temperatures because experience shows that diagnosing icing more frequently at temperatures less than -20°C results in too many false alarms. The

$T_{map\text{-}convective}$ curve lies well to the left of the T_{map} curve, reflecting the fact that strong updrafts in convective clouds raise SLW to higher altitudes than would be observed in stratiform clouds.

Figure 2.23b shows the normalized PIREP curve for cloud-top temperature and the corresponding interest map CTT_{map}. The resolution and accuracy of the RUC model are insufficient for determining cloud-top temperature, so the GOES satellite is used instead to match icing PIREPs with Channel 4 brightness temperature for the highest clouds in the near vicinity. Thus in Fig. 2.23b, the thin gray curve is for a normalized PIREP ratio using CTT instead of model temperatures. Clouds with very cold tops can still contain SLW. Ice crystals falling through SLW droplets lower in the cloud will glaciate (freeze) them upon contact, resulting in a mixed-phase or completely glaciated cloud above the freezing level. As the cloud-top temperature increases, liquid water at the top becomes increasingly likely. Few PIREPs are received when cloud-top temperatures nearby lie above freezing, but the number is non-negligible. It is unusual but possible that an inversion and sub-freezing temperatures exist within a cloud even when the top is relatively warm.

CTT_{map} is intended to give the relative likelihood of liquid water within the cloud, given the temperature at cloud top. CTT_{map} increases steadily with cloud-top temperature, reaching 1.0 at -12^{o}C, indicating that liquid in the cloud is highly likely when CTT exceeds this value. CTT_{map} can stay at 1.0 for CTT > 0^{o} because it is always multiplied by T_{map} when the final icing index is calculated, and T_{map} = 0 for temperatures above freezing.

Fig. 2.23c shows the normalized PIREP curve for model relative humidity. Few pilot reports are received when the model RH is less than 70% (and this is due more to model errors in RH than pilot error), but the number increases rapidly between 70 and 100%, as expected. RH varies on spatial scales of tens of kilometers. Clouds form only when the RH is close to 100% (saturation with respect to liquid water). Their spatial variability is even greater than that for water vapor and is not well resolved by the RUC model. For this reason, the RH_{MAP} curve is more liberal in diagnosing icing than the normalized PIREP ratio. RH_{MAP} is 1.0 for very high values of RH.

Fig. 2.23d shows the normalized PIREP curve for model vertical velocity as expressed in microbars per second (μb/s). Rising air parcels experience decreasing pressure whereas sinking parcels experience increasing pressure. Thus the left side of the diagram represents rising air and the right side sinking air. Note that icing PIREPs become more likely as rising motion increases, consistent with the formation of clouds, some containing SLW, when rising motion persists. A significant number of PIREPs are associated with downward motion. Though the concentration of SLW must decrease with downward motion, it must still be sufficient for icing,

but perhaps not for long. Because VV forecasts are a challenge for any model, and there is little chance for direct verification of the forecast, the VV_{maps} curve allows for a wide range of vertical velocity, with good chances of icing (VV_{maps} = 1.0) from -1.0 to -0.5 μb/s, then decreasing chances in steps, starting from 1.0 at -0.5 μb/s and ending at 0.0 at +1.0 μb/s.

The CIP algorithm uses many more interest maps; two of special significance will be discussed further. The interest map for pilot reports is shown in Fig. 2.24. A pilot report of icing, regardless of severity, should clearly be included in any product that portrays icing. Because icing reports are not particularly dense, the influence of a single report is spread horizontally and vertically as in Fig. 2.24. The vertical influence of a report is limited to 300 m above and below on the assumption that a pilot will not deliberately fly into a thunderstorm but will normally encounter icing in a stratiform cloud. In such a cloud, small changes in flight altitude can result in large changes in icing conditions.

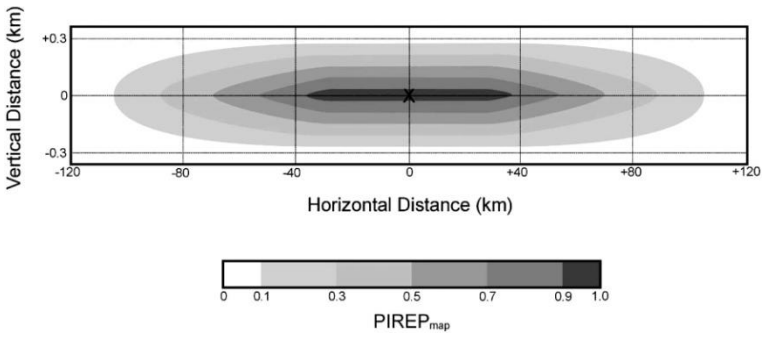

Fig. 2.24. The two-dimensional interest map for PIREPs. The scale (0.0 to 1.0) is at the bottom.

The final interest map pertains to a model prediction of SLW. Wherever the RUC model predicts SLW, it captures roughly 40% of all icing events. On the other hand, many icing reports come from atmospheric volumes where the RUC has *not* predicted SLW. The interest map for SLW is intended to indicate either the presence or absence of SLW, and so SLW_{map} is set to one when the RUC predicts SLW and to zero when it does not. Used with permission of the American Meteorological Society.

Step 4 uses the interest maps to calculate the initial icing potential and the initial SLD potential. The algorithm considers five different cases: 1) a single cloud layer, 2) multiple cloud layers, 3) a gradient in cloud-top temperature, 4) classical freezing rain structure with above-freezing air overlying sub-freezing air, 5) and deep convection. It is noteworthy that the product T_{map}*CTT_{map}*RH_{map} appears in the calculation of all initial icing potentials for cases 1, 2 and 3.

A single cloud layer is inferred whenever high RH in the model is found at all levels between the cloud top obtained from an IR image and the cloud base obtained from a METAR. When

neither radar nor surface reports indicate precipitation, the SLD potential is set to "unknown." Liquid precipitation and relatively warm cloud tops suggest that the collision-coalescence process is active in forming larger drops. The potential for icing is high and the possibility of SLD must be considered. When only snow is observed at the surface, the crystals scavenge supercooled cloud droplets, and, the heavier the snowfall as indicated by radar, the less likely icing will occur.

Multiple cloud layers are inferred whenever the model indicates one or more layers of RH < 50% (each at least 75 mb thick) in between cloud top as defined by GOES and cloud base defined by METARs. Precipitation at the surface is attributed only to the lowest cloud layer. Each layer has it own top and base; the CIP algorithm calculates icing and SLD potentials for each one.

When a gradient in CTT is present within a grid box, variations in icing conditions may be occurring. Clouds with CTT near -12^{o} may contain considerable SLW whereas those with CTT around -25^{o} may be mostly glaciated so that an aircraft traversing the grid box at fixed altitude may encounter rapidly changing icing conditions. For each flight altitude within cloud, the CIP algorithm infers more or less icing potential depending upon the CTT above the aircraft.

The classical freezing rain structure consists of precipitation falling through a layer of above-freezing air sandwiched between two layers of below-freezing air. The model temperature profile defines the depths and vertical extent of the three layers. Because no melting occurs in the upper cold layer, it is treated as a single-layer cloud. Since this layer is usually cold and full of snow crystals or flakes, which scavenge cloud droplets, the icing potential is often small. When liquid precipitation (freezing or not) or ice pellets are observed at the surface, SLD is likely in the lower sub-freezing layer. Icing potential in the warm layer is zero.

When lightning is observed within 25 km of a grid point, deep convection is assumed. The initial icing potential and the final SLD potential are set to $T_{map\text{-}convective}$.

The SLD potential is set to "unknown" whenever no precipitation is observed at the ground or only snow is observed because not enough information is present to infer the presence or absence of SLD. In many cases, the SLD calculation is affected by the radar reflectivity..

The explicit calculations for all cases noted above and a few others appear in Table 2 of Bernstein et al. (2005).

Step 5 alters the initial estimates of icing potential (but not SLD potential) on the basis of the three remaining interest maps not yet used: SLW_{map}, $PIREP_{map}$, and VV_{map}. Recent PIREPs, the presence of SLW, and upward motion in the model can all increase the initial icing potential. Only downward motion can decrease it. Any increase cannot make the final icing potential greater than 1.0, nor can any decrease make it less than 0.0. SLW_{map}, $PIREP_{map}$, and VV_{map} can boost the initial icing potential by as much as 40%, 35%, and 25%, respectively, of the difference between 1.0 and the initial value. VV_{map} can decrease the initial icing potential by as much as 25% of the difference between the initial value and zero.

2.2.8. Forecast Icing Product

Construction of the Forecast Icing Product (FIP) is very similar to that for the CIP, the major difference being that model output from the RUC is used exclusively as a substitute for the observations used in the CIP. The best description of FIP methodology is in McDonough et al. (2004). Without the benefit of cloud observations from GOES and METARs and precipitation observations from radar and METARs, the presence of clouds and precipitation must be inferred from RUC model output.

Working from the top model level down, cloud tops are assumed whenever RH rises through 70%. Starting from 1,000 ft above ground and working up, cloud bases are assumed whenever the RH rises through 80%. Dry layers are defined by an RH 50% or less at three consecutive model levels sandwiched between layers where RH is 70% or more. Cloud layers separated by dry layers are treated independently because precipitation falling through a dry layer is likely to evaporate before reaching a cloud layer below.

If the model forecasts more than 0.1 mm of precipitation in 3 h, it is expected to reach the surface. By checking the temperature and humidity profiles between where precipitation forms in the cloud and on the surface, it is possible to infer precipitation type at the ground.

The interest maps used in CIP (T_{map}, RH_{map}, CTT_{map}, VV_{map}, and SLW_{map}) are also used in FIP, and they are virtually identical. FIP uses the quantitative precipitation forecast (QPF) as the basis for a new interest map QPF_{map}, shown in Fig. 2.25. If only snow is falling at the ground, QPF_{map} acts to reduce the likelihood of icing, more so for greater snowfall rates. If sleet or freezing precipitation is reaching the ground, QPF_{map} acts in the opposite direction, to increase the likelihood of icing as the precipitation rate increases like radar reflectivity in CIP.

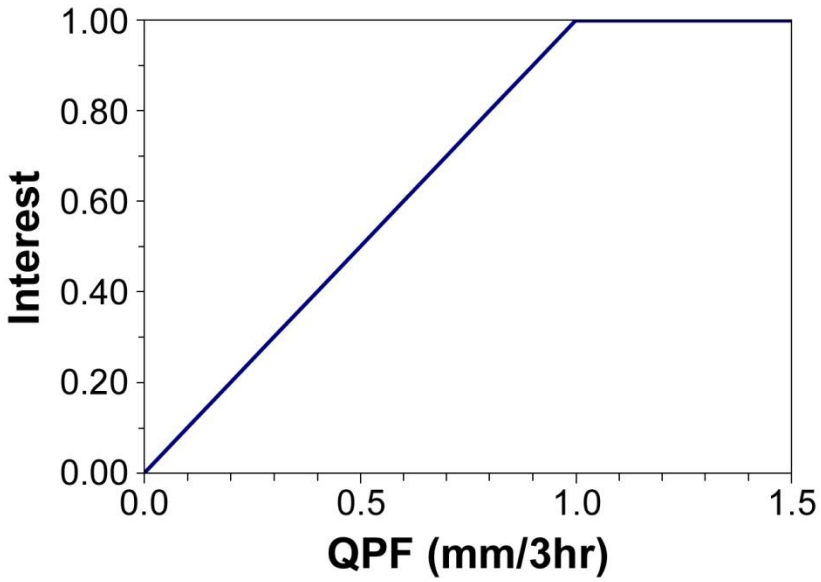

Fig. 2.25. Interest map for a quantitative precipitation forecast.

The decision tree for arriving at final values for forecasting icing potential and SLD is similar to that described in Step 4 above for CIP or, with greater detail, in Bernstein et al. (2005) except for the use of QPF_{map}.

The SLD potential is disseminated along with the icing severity product to be described in Section 2.2.10. This index is *not* calibrated. It is used only to suggest where supercooled large drops may occur. Over 18,300 icing reports from 1 Jan thru 31 Mar 2005 were used to assess the performance of the SLD algorithm. The results are summarized in Table 2.4.

Table 2.4

The computed SLD index at the initial time with corresponding pilot reports and surface observations. Extracted from Fowler et al. (2006).

SLD (0.0 – 1.0) Computed Value	PIREP observations of icing				Surface Observations	
	Severe	FZRA/FZDZ	Moderate or greater	No Icing	FZRA/FZDZ	No FZRA/FZDZ
Yes (SLD > 0)	3.3% (3)	18.8% (9)	12.2% (593)	5.5% (286)	26.6% (81)	1.6% (200)
No (SLD = 0)	0% 0	43.7% (21)	15.8% (773)	52.9% (2,767)	2.6% (8)	0.6% (78)
Unknown	96.7% (88)	37.5% (18)	72.0% (3,508)	41.6% (2,175)	70.8% (216)	97.8% (12,399)
Totals	100% (91)	100% (48)	100% (4,874)	100% (5,228)	100% (305)	100% (12,677)

The left-hand column of Table 2.4 gives the SLD potential, a number from 0.0 to 1.0. As noted earlier, in many areas where icing is possible but precipitation is either absent or entirely snow, the SLD potential is listed as unknown. During winter, the atmospheric volume where SLD is positive comprises only about .005 of the total airspace volume. Because surface observations of freezing rain (FZRA) or freezing drizzle (FZDZ) are common when SLD > 0, these observations are listed at right. The middle of the table lists pilot reports of icing or lack of it, and of FZRA or FZDZ. The latter reports are unusual but particularly valuable for assessing the utility of SLD diagnoses.

The areas where supercooled large drops were diagnosed contained nearly 19% of all moderate or greater icing reports and almost 27% of all surface reports of FZRA or FZDZ. Less than 6% of null reports fell within SLD areas and less than 2% of the time did FZRA or FZDZ *not* occur within SLD areas. Wherever SLD is 0.0, so is the icing potential. About 53% of null reports came from areas where SLD was zero. Yet pilots reported FZRA or FZDZ 21 times in these same areas. It is problematic that so many moderate or greater icing reports fell within "unknown" areas, but it is impossible to say how many of these were associated with SLD. At least icing is considered possible in the "unknown" areas.

The icing potentials calculated in CIP and FIP are not disseminated; however, they are the basis for computing the probability of icing, described in the next section.

2.2.9 Probability of Icing

By design, the icing potential in CIP and FIP ranges from 0.0 to 1.0. From the time of invention, CIP and FIP have been treated as uncalibrated measures of the probability of icing, that is, higher values suggested a greater likelihood of icing. In 2006 and 2007, NCAR conducted exercises in which icing PIREPs were matched with CIP and FIP values. Investigators wanted to estimate the probability Z_i that icing conditions would be observed, given a particular range of CIP or FIP values t, say, between t_i and t_{i+1}, where i specifies a particular interval. Using Bayes' Formula, one can derive the following expression:

$$Z_i = \frac{p_x(PODy_{i-1} - PODy_i)}{(p_{f,i-1} - p_{f,i})},$$

where p_x is the overall probability of an icing observation; $PODy_i$ is the conditional probability that icing conditions were diagnosed ($CIP, FIP \geq t_i$), given that icing conditions were observed; and $p_{f,i}$ is the probability of a "yes" forecast when the CIP/FIP threshold is t_i. $p_{f,i}$ is

the fraction of the grid volume for which there is a "yes" diagnosis or forecast of icing when the threshold is t_i.

The problematic variable is p_x because the atmospheric volume is not uniformly sampled for icing. Pilots are motivated to report icing conditions but disinclined to report lack of icing. Fortunately, TAMDAR, first mentioned in Section 2.2.2, includes an icing sensor that reports at regular intervals during flight. AIRDAT allowed NCAR to analyze sensor data during the winter of 2005–2006, from aircraft flying in the Great Lakes and Upper Midwest that were made available for research. As expected, null reports far outnumbered positive reports of icing, but a uniform sampling of atmospheric volume is precisely what is needed to form an icing climatology. Using the TAMDAR icing data, NCAR estimated that the climatological probability of icing p_x is 0.046.

The other numbers in the equation for Z_i come from matched pairs of icing PIREPs and CIP/FIP values. The major finding of this study was that CIP slightly exaggerates the icing threat at the initial time, and FIP overforecasts icing at all lead times but more so as lead times increase. To obtain the probability of icing from icing potential, the remedy is to multiply CIP by 0.85, and FIP by 0.73, 0.62, 0.56, and 0.44 for lead times of 3, 6, 9, and 12 h, respectively. The icing probability can never exceed 0.85; it is available as a product on the ADDS website at http://aviationweather.gov/adds/icing/.

Figure 2.26 shows a sample map of the maximum probability of icing at *any* level at the initial time. Available PIREPs are superimposed. One can also display the probability of icing at individual flight levels and for 3-, 6-, 9-, and 12-h lead times.

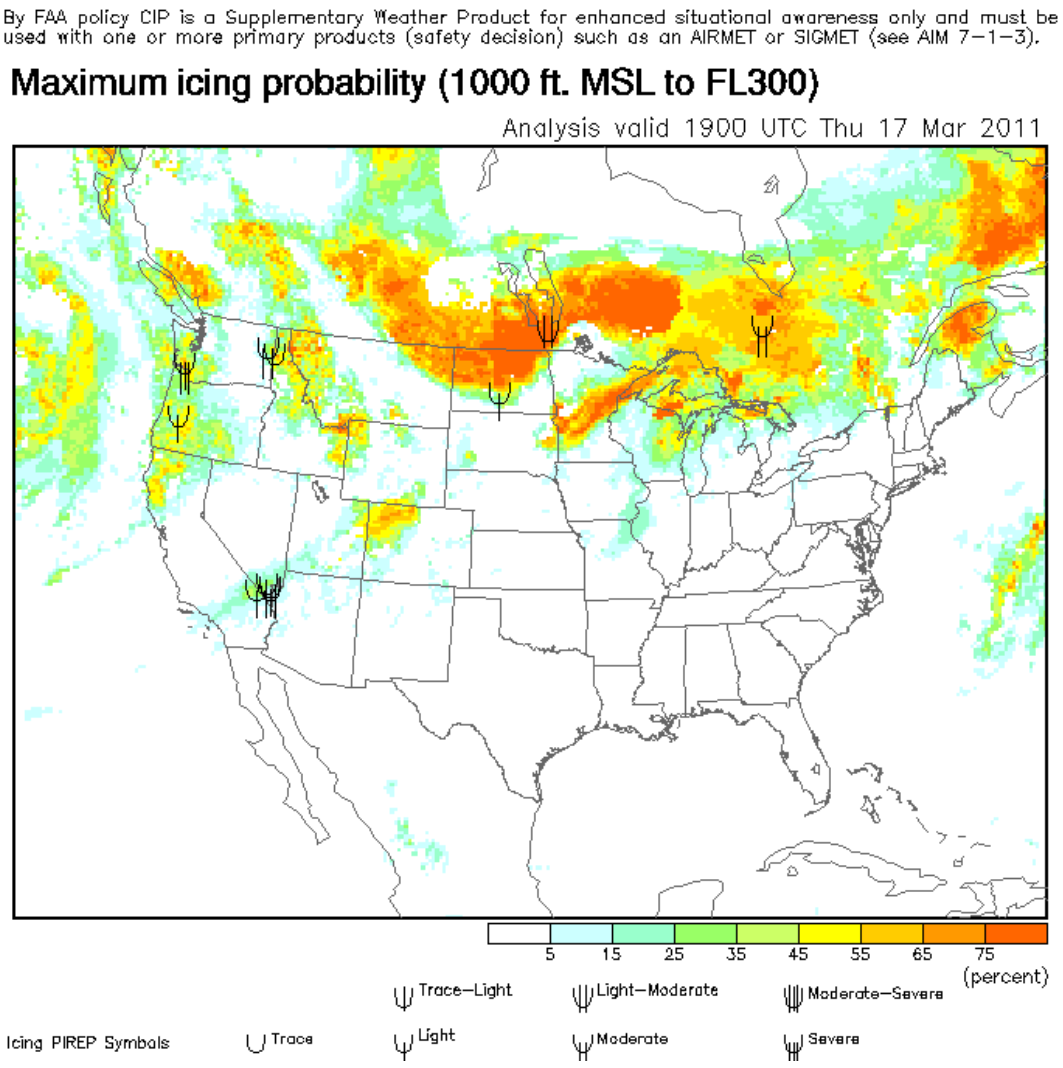

Fig. 2.26. The maximum probability of icing at any level for the initial time of 1900 GMT, 17 March 2011. Icing reports at all levels are plotted on the map.

Figure 2.27 shows the ROC curve for the winter icing season from 1 Nov 2009 through 31 Mar 2010: a measure of performance of the probability product. Table 2.5 gives the data used to generate Fig. 2.27.

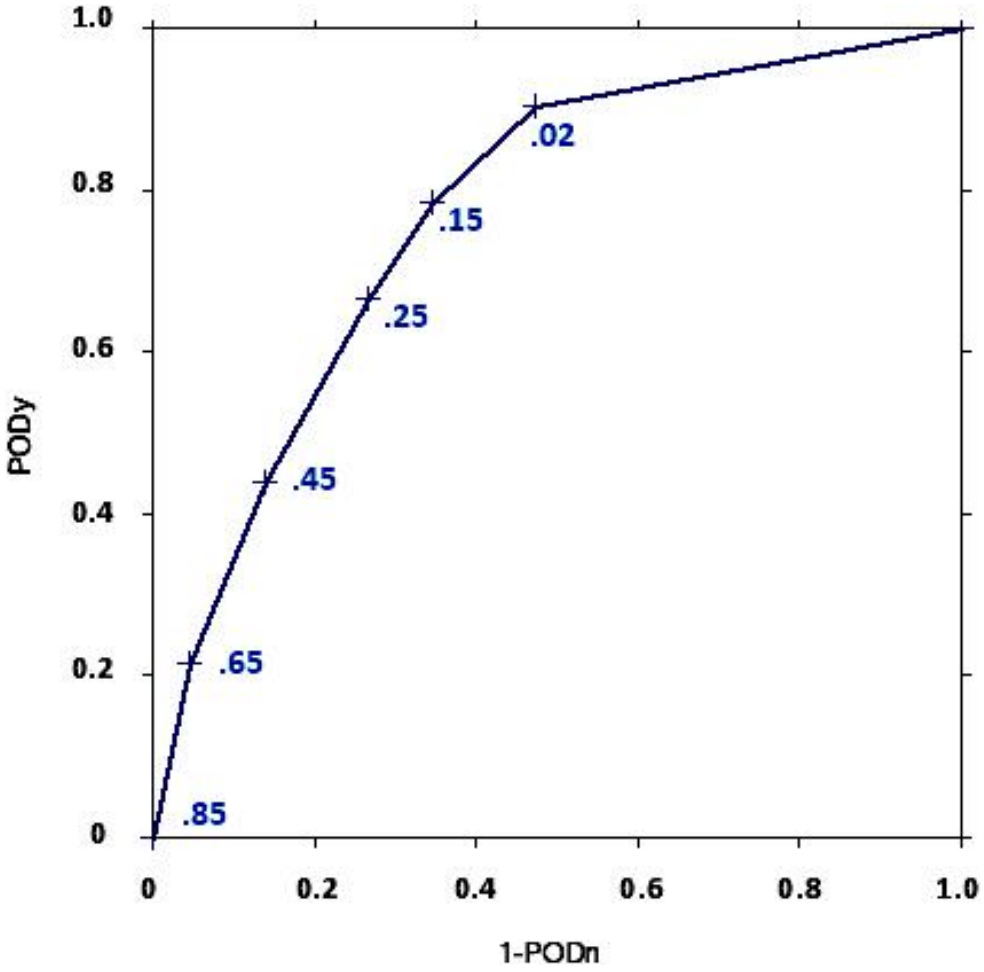

Fig. 2.27. The ROC diagram for icing probability forecasts at the initial time made at any level from 01 Nov 2009 through 31 Mar 2010. Probability values from Table 2.5 are plotted alongside the ROC curve.

Table 2.5

Data used in constructing the ROC diagram in Fig. 2.27 for the probability of icing at the initial time.

Data Source	Variable	Obs Threshold	Start Date	End Date	Product Time	Lead Time	Region	Flight Level	ROC Area
CIP Operational (20 km)	Probability	N = No Y ≥ Light-Moderate	2009 01 Nov	2010 31 Mar	All hours	00 h (initial Time)	National	Ground to 30,000 ft	0.773

Fcst Thesh	YY	YN	NY	NN	1-PODn	PODy
0.85	0	0	29080	16195	0.000	0.000
0.65	6239	743	22841	15452	0.046	0.215
0.45	12747	2228	16333	13967	0.138	0.438
0.25	19403	4310	9677	11885	0.266	0.667
0.15	22821	5597	6259	10598	0.346	0.785
0.02	26308	7638	2772	8557	0.472	0.905

In Table 2.5, the probability thresholds are listed in the left-hand column. These probabilities are the same as those given in the color bar of Fig. 2.26, just with different break points. In the first row, because there were no icing forecasts at the 0.85 probability level, one quickly notices that there were 29,080 reports of icing and 16,195 null reports in the data sample. Starting with the bottom row, the area where the probability of icing is at least 0.02 is fairly large, and so most of the icing reports fall in this area, giving a PODy of 0.905. As the threshold probability increases to .15, the area enclosed shrinks, but 0.785 of icing reports still fall inside. In areas where the probability of icing was 0.65 or more, there were 6239 reports of icing and only 743 null reports, but 22,841 icing reports lay *outside* these areas. Such behavior is typical in ROC diagrams. The area under the ROC curve is often considered an overall measure of skill. A perfect score is 1.0. For this product, the score is a respectable 0.773.

The icing probability product runs operationally at AWC.

2.2.10 Icing Severity: Current and Forecast

Having produced an algorithm for icing potential, NCAR capitalized on that experience to write an algorithm for icing severity. The latter algorithm uses many more icing-related parameters

than the former one, and it must run later because icing and SLD potentials are two of the parameters used. Figure 2.28 shows the flow chart for computing icing severity. It may be compared with Fig. 2.22.

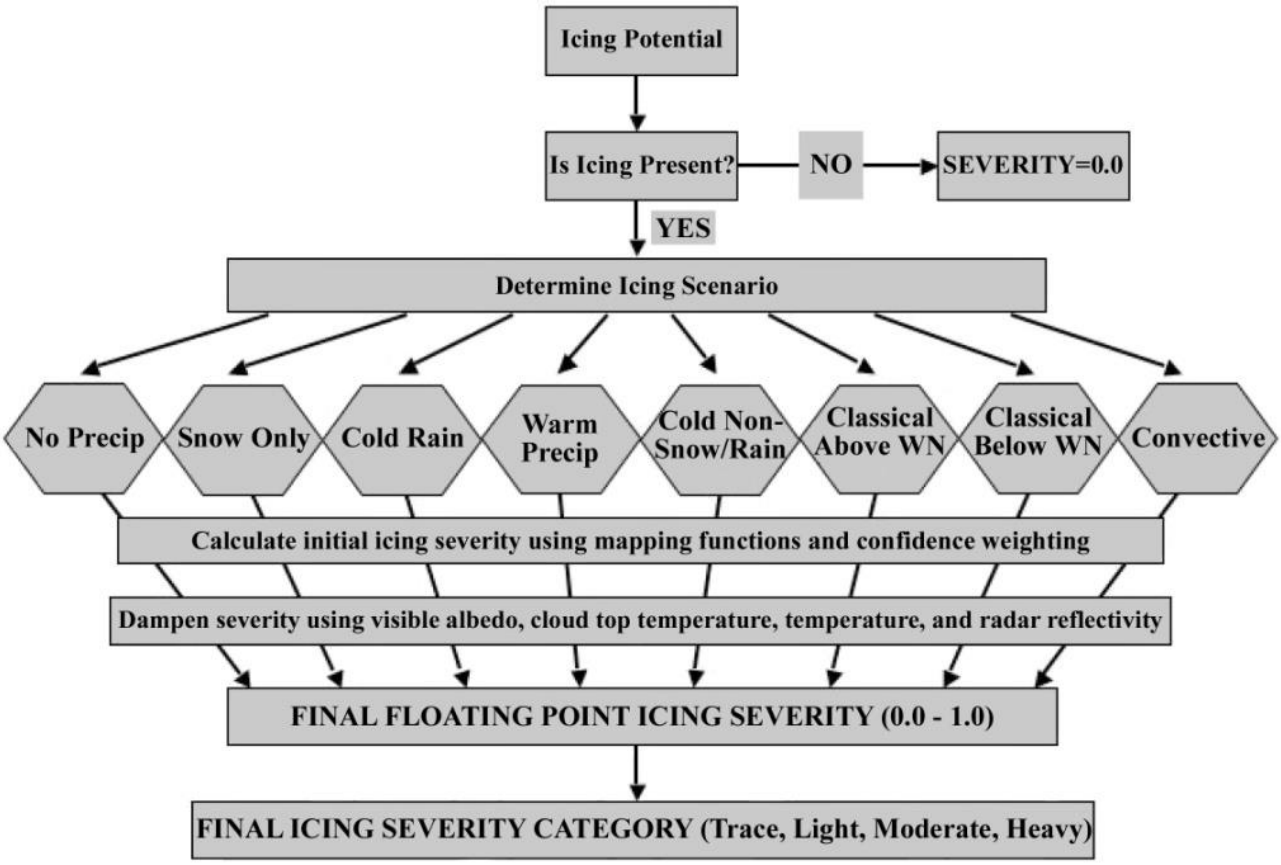

Fig. 2.28. Flow chart for generation of the Icing Severity Index. From Politovich et al. (2006)

The computation of "final icing severity" in Fig. 2.28 is a long and involved process, described in Appendix B, Section B2.2.10.

Icing severity products are available on the ADDS website at http://aviationweather.gov/adds/icing/icing_nav.php. The user can choose to view the current best estimate of icing severity, or forecasts for 1, 2, 3, 6, 9, and 12 h in the future. At any of these times, the user can view 1) the worst icing conditions expected at any level, or the icing severity at specific levels; 2) icing severity only in areas where the probability of any icing at all equals or exceeds 25%, or 50%; 3) icing severity with or without a hatched overlay showing where supercooled large drops are possible. The hatching appears wherever the SLD index

exceeds 0.05 (on a scale of 0.0 to 1.0). The SLD index is still an uncalibrated product, and so gradations in the concentration of large supercooled droplets are not specified.

Here are three sample products. Figure 2.29, a 6-h forecast, shows the maximum icing severity at any level with a hatched red overlay showing the potential for SLD: the darker the shade of blue, the more severe the icing. The potential for SLD stretches in a broad band from southern Alberta to northern Virginia. Most rawinsonde soundings taken along this band featured a frontal inversion; from Green Bay southeastward, several soundings showed a distinct warm nose. A well-defined cyclonic storm approaching the West Coast is spreading icing conditions inland.

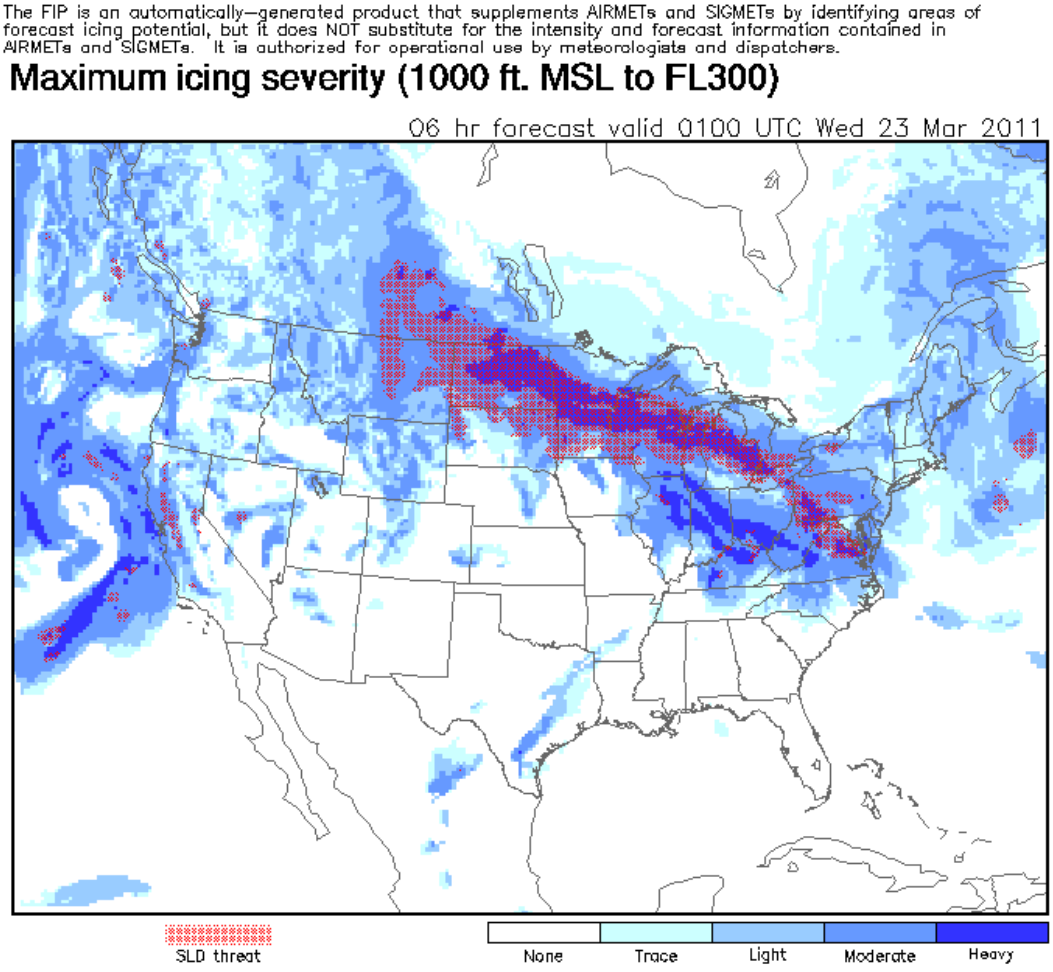

Fig. 2.29. Six-hour forecast of maximum icing severity at any level, valid at 0100 GMT 23 March 2011. The hatched red overlay indicates where supercooled large drops are possible. A white

background indicates that no icing is expected. The shades of blue indicate the severity of icing, as indicated by the color bar.

Figure 2.30 indicates icing severity at 9,000 ft MSL. Icing on this map covers much less area than on the composite map (Fig. 2.29) because much of the icing in Canada lies below 9,000 ft whereas icing in the southern U.S. lies mostly above 9,000 ft. Moderate or greater icing and the potential for SLD are predicted at most levels up to 15,000 ft (not shown) in the band from Montana to Lake Erie because the freezing level is low and the cloud system is deep. The tan coloring indicates where terrain rises above 9,000 ft.

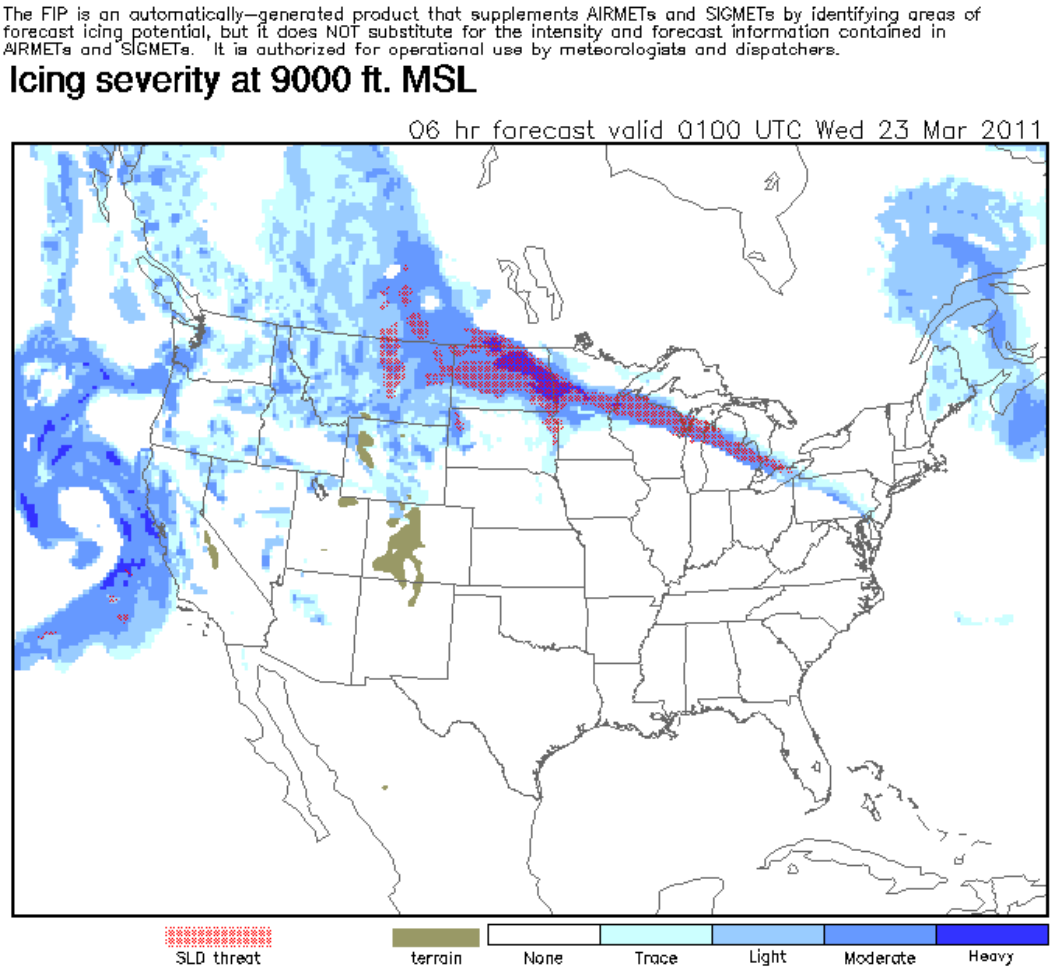

Fig. 2.30. Six-hour forecast of icing severity at 9,000 ft MSL, valid at 0100 GMT 23 March 2011. The red overprint indicates where large supercooled drops are possible. The tan color indicates where terrain exceeds the 9,000-ft elevation.

Like Fig. 2.30, Fig. 2.31 indicates icing severity at 9,000 ft MSL, but only where the probability of icing is 50% or more. A gray mask covers areas where the probability of icing is less. Note the pronounced shrinkage of blue areas when this mask is applied. Pilots should be wary of flying into areas where moderate or heavy icing is predicted with a probability of 50% or more.

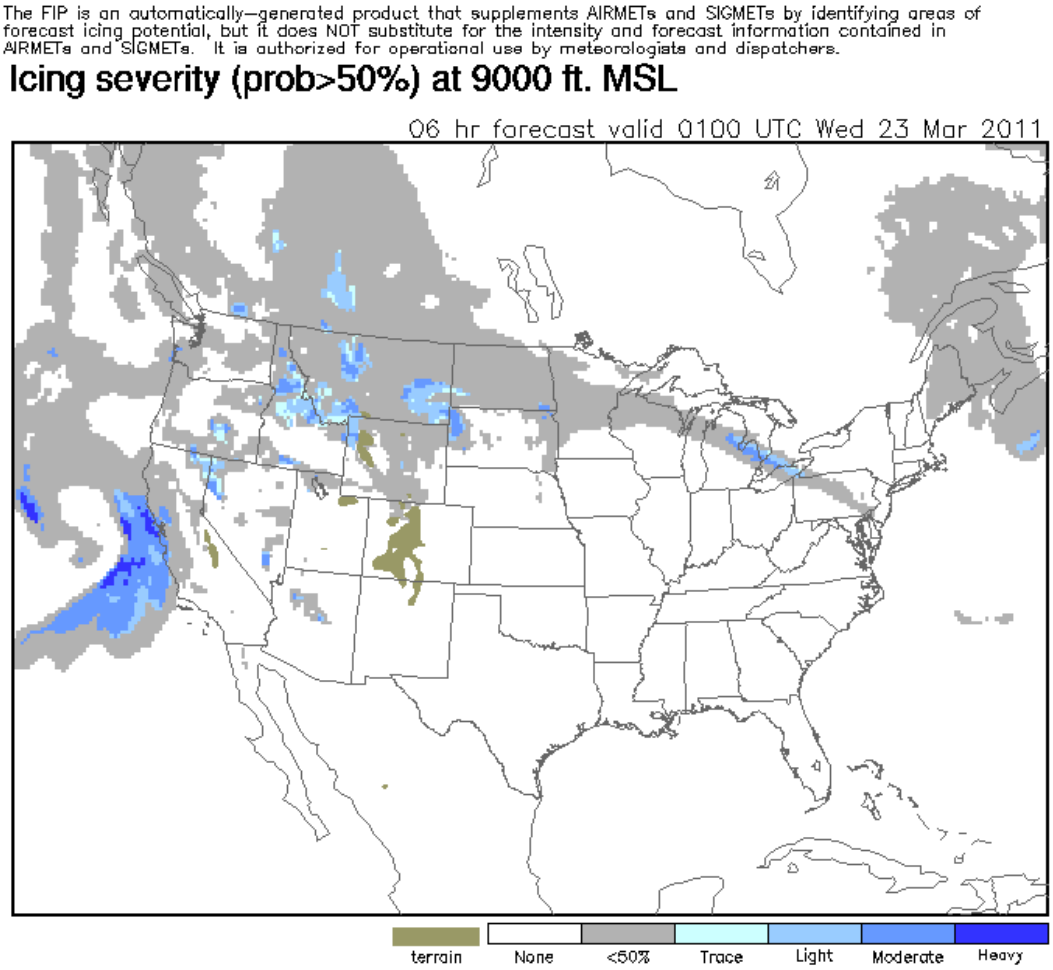

Fig. 2.31. Six-hour forecast of icing severity at 9,000 ft MSL, valid at 0100 GMT, 2300 March 2011, but only where the probability of any icing is greater than 50%. Areas where the probability of icing is less are masked in gray.

How good is the icing severity product? Figure 2.32 is the result of matching over 12,000 PIREPs, collected over the U.S. from January through March 2005, with current and predicted icing severity index. The curves show the distribution of icing severity values SEV_{final} in bins

0.08 wide, for icing PIREPs in five different categories: none (black), trace (dashed red), light (dotted green), moderate (dash-dotted purple), and heavy (cyan). The vertical axis, labeled *density*, is equivalent to counts of SEV_{final} values in each bin, as in a histogram. For example, when pilots reported no icing, a large majority of severity values was close to 0.0. Very few were higher than 0.10. When pilots reported a trace of icing, the distribution of SEV_{final} values peaked at 0.30. When they reported light icing, it peaked at 0.35. The distributions of SEV_{final} for reports of moderate and severe icing both peaked near 0.40. For reference, the vertical dash-dotted lines in Fig. 2.32 mark the breakpoints between levels of icing severity. The value 0.175 separates a trace of icing from light icing; 0.375 separates light and moderate icing; and 0.70 separates moderate and severe icing. Ideally, the distributions would not overlap so much and their peaks would be more distinct. One concludes that the icing severity index shows skill, but the discrimination between various categories of pilot-reported icing is marginal.

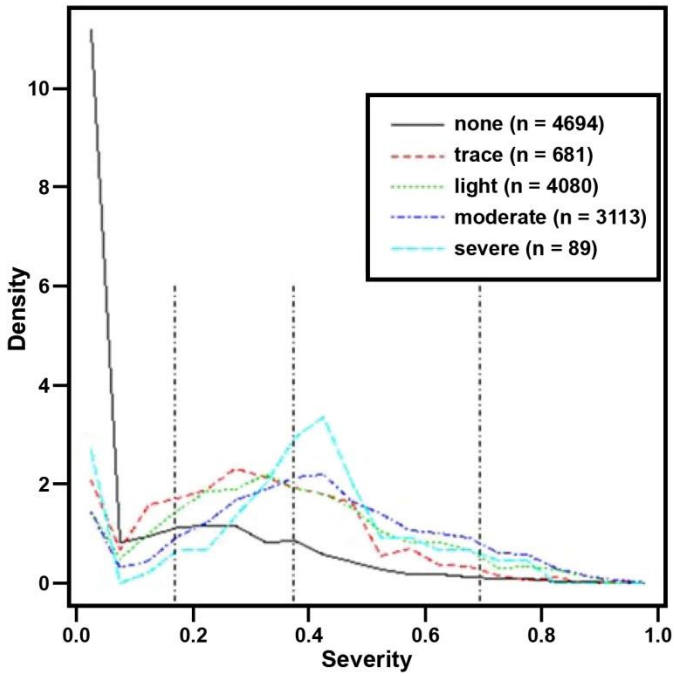

Fig. 2.32. A discrimination plot showing the distribution of icing severity values for all valid times (current and predicted) at locations with PIREPs. Each curve shows the distribution of SEV_{final} values for PIREPs of a given category of icing: none, trace, light, moderate, or heavy. PIREPS were collected over the CONUS from January through March 2005. Vertical dash-dotted lines show threshold values used to separate light, moderate, and heavy icing on the ADDS icing severity product. From Fowler et al. (2006).

Another way of looking at the icing severity product is based upon a different collection of PIREPs made in January through March of 2006 and 2007. The PIREPs were matched to icing severity forecasts valid at 1, 2, 3, and 6 h. The union of areas where *any* icing was predicted at these lead times was overlaid on the area covered by an AIRMET, which is valid for six hours at a time. The intersection of the AIRMET area with the aggregated icing severity areas was the area used for calculating statistics. This is described in more detail in an internal report by Madine et al. (2008).

Figure 2.33 shows the PODy and PODn scores in 5,000-ft altitude ranges. To repeat earlier definitions, PODy gives the fraction of moderate or greater icing reports that were forecast correctly. PODn gives the fraction of all other reports (no icing, trace or light icing) that were forecast correctly. The total number of reports (icing and no icing) in each altitude range is given at right. The figure at left is without regard to the probability forecast; the figure at right includes only those areas where the probability of icing was expected to be 25% or more.

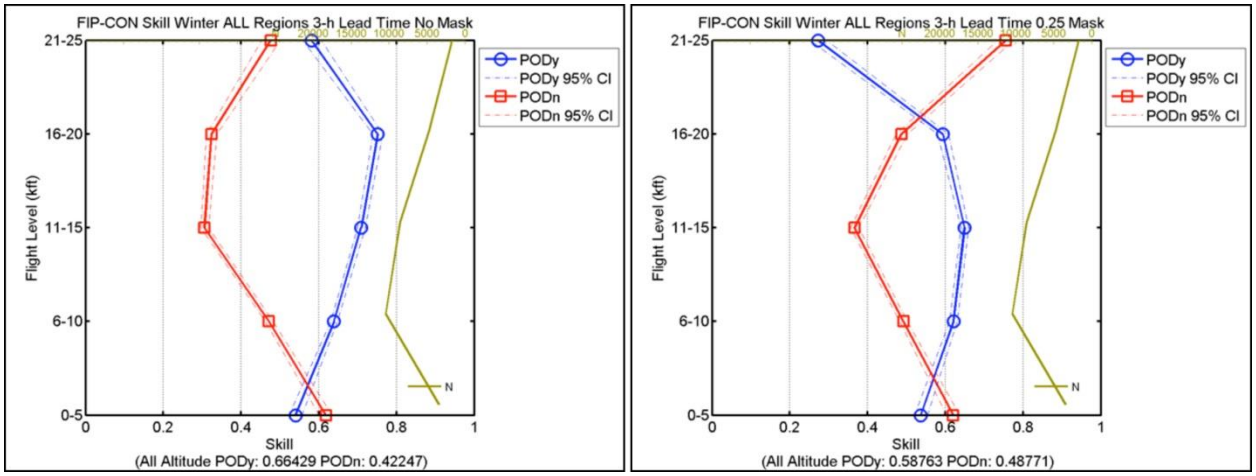

Fig. 2.33. Skill of CONUS icing severity forecasts, when "yes" means moderate or greater icing, and "no" covers everything from no icing through light icing. Values of PODy (blue circles) and PODn (red squares) are plotted in 5,000-ft altitude ranges, bracketed by 95% confidence intervals in the dashed lines. The total number of PIREPs (icing and no icing) considered in each altitude range is the gold-colored line at right. PIREP counts are given at top right of each figure.

The icing severity forecast adds value to the AIRMET by effectively reducing the area in which moderate or greater (MOG) icing is predicted. The PODn averaged over all altitudes in Fig. 2.33 (left) is 0.42, whereas the corresponding PODy average is 0.66. In Fig. 2.33 (right) the verification area is reduced by masking the area where the probability of icing is less than 25%.

This has the effect of increasing the PODn average to 0.49 but decreasing the PODy average to 0.59, a trade-off not necessarily worthwhile.

The icing severity product runs operationally at AWC.

Upgrades to the icing probability, icing severity, and SLD algorithms fall into three categories:

- Changes to the numerical prediction models that supply much of the input. This includes changes in spatial resolution, geographical coverage, and cloud and precipitation physics.
- Data from new sources, some available but not yet exploited (satellite instruments that measure cloud-top properties), some soon to be available (dual-polarization capability for the WSR-88D radars).
- Changes to the algorithm to improve performance, in particular, new or modified interest maps.

Within a year, the Rapid Update Cycle model will be upgraded to the Rapid Refresh model. The latter has finer spatial resolution, greater areal coverage, and upgraded physics. The icing algorithms will be modified to accept output from this new model. Within two years, the algorithms will incorporate products generated by the Advanced Satellite Aviation weather Program (ASAP) at the NASA Langley Research Center (LaRC), in particular, for identifying the phase (liquid or ice) at cloud top, measuring the cloud-top temperature, and estimating the liquid water path. ASAP is testing these products routinely. They are described next.

2.2.11 Cloud Properties from GOES Satellites

As of this writing, GOES-11 provides data for the western U.S. and GOES-13 for the eastern U.S. GOES-13 replaced GOES-12 on 14 April 2010. The algorithm descriptions below apply equally to GOES-12 and GOES-13. The imagers aboard these satellites provide high-resolution views of the Earth and its atmosphere and hence are useful for locating clouds and inferring their properties. Table 2.6 lists the imaging channels for GOES-11 and 12, their wavelength ranges, and the size of the spot viewed on the Earth's surface when the scanning instrument points straight down (technical name: instantaneous geometric field of view (IGFOV) at nadir).

Table 2.6
Imaging channels for GOES-11, GOES-12, and GOES-13

Channel	Central Wavelength (μm)	Size of IGFOV at nadir
1 (visible), GOES 11, 12, 13	0.67	1 km
2 (near IR), GOES 11, 12, 13	3.9	4 km
3 (moisture), GOES 11	6.7	8 km
3 (moisture), GOES 12, 13	6.7	4 km
4 (longwave 1), GOES 11, 12, 13	10.8	4 km
5 (longwave 2), GOES 11	12.0	4 km
6 (longwave 2), GOES 12, 13	13.3	4 km

Remote sensing of cloud properties by satellite has a fairly long history of development at NASA LaRC dating from before the mid 1980s. This section describes icing-related products derived from GOES-11 and GOES-13 radiances.

A few definitions from the *Glossary of Meteorology* (Glickman, 2000) are in order before proceeding. *Radiance* is the rate at which radiant energy in a set of directions confined to a unit solid angle around a particular direction is transferred across a unit area. The unit area is oriented normal to the central direction. The radiance may pertain to radiation at a given wavelength, in which case it is called *spectral radiance*, or it may refer to the energy integrated over a spectrum of wavelengths. The units of spectral radiance are watts per square meter per steradian, per wavelength interval: $W\ m^{-2}\ sr^{-1}\ \mu m^{-1}$. *Irradiance* has a different meaning; it is the rate at which radiant energy in a radiation field is transferred across a unit area of surface in a hemisphere of directions. The radiation may be at a single wavelength or integrated over many wavelengths. For *spectral irradiance*, the units are $W\ m^{-2}\ \mu m^{-1}$. *Radiance* is an intrinsic property of the radiation field. *Irradiance* depends upon the orientation of the surface. For example, installers of solar panels will be more interested in the irradiance than the radiance. In model calculations of surface heating caused by the sun, irradiance is the relevant quantity.

Finally, *reflectance* is the radiation reflected from, or scattered back through, a given surface in response to radiation incident on that surface at the same wavelength or within the same wavelength range. It is a dimensionless ratio. Reflectance in terms of the ratio of reflected to incident irradiance is generally termed *albedo*. To continue the above example, the fraction of solar radiation reflected at the ground, the albedo, is important for computing surface heating. The term *albedo* is generally used when considering the hemispheric reflectance at solar wavelengths either in narrow or broad bands. *Reflectance* is more often used when referring to narrowband visible light reflected in specific directions. For example, the reflectance of radiance from a specified incident direction into a specified reflected direction is called *bidirectional reflectance.* These terms arise again in the following discussion.

The flow diagram in Fig. 2.34 shows how various cloud properties are calculated by ASAP. The items in green may be considered constant, at least for a month at a time. The clear-sky albedo refers to the albedo that would be observed with the sun overhead and the direction of view straight down. It depends upon the brightness of the Earth's surface, which in turn depends upon the composition of the surface (sand, dirt, rock, water), vegetation (growing or dormant), and the presence or absence of snow and ice cover. The surface type map classifies the surface into 19 types according to a scheme proposed by the International Geosphere Biosphere Programme (IGBP), for example, evergreen forest, desert, grassland, open water, snow-ice, cropland, etc. The surface emissivity map (wavelength dependent) gives the ratio of power emitted by the land surface compared with the power that would be emitted if it were a black body (which follows Planck's radiation law). All three of these parameters are derived from satellite data and are updated as necessary. They are primarily used to 1) estimate the clear sky or "background" radiances needed to differentiate clouds from cloud-free areas and 2) derive accurate cloud properties.

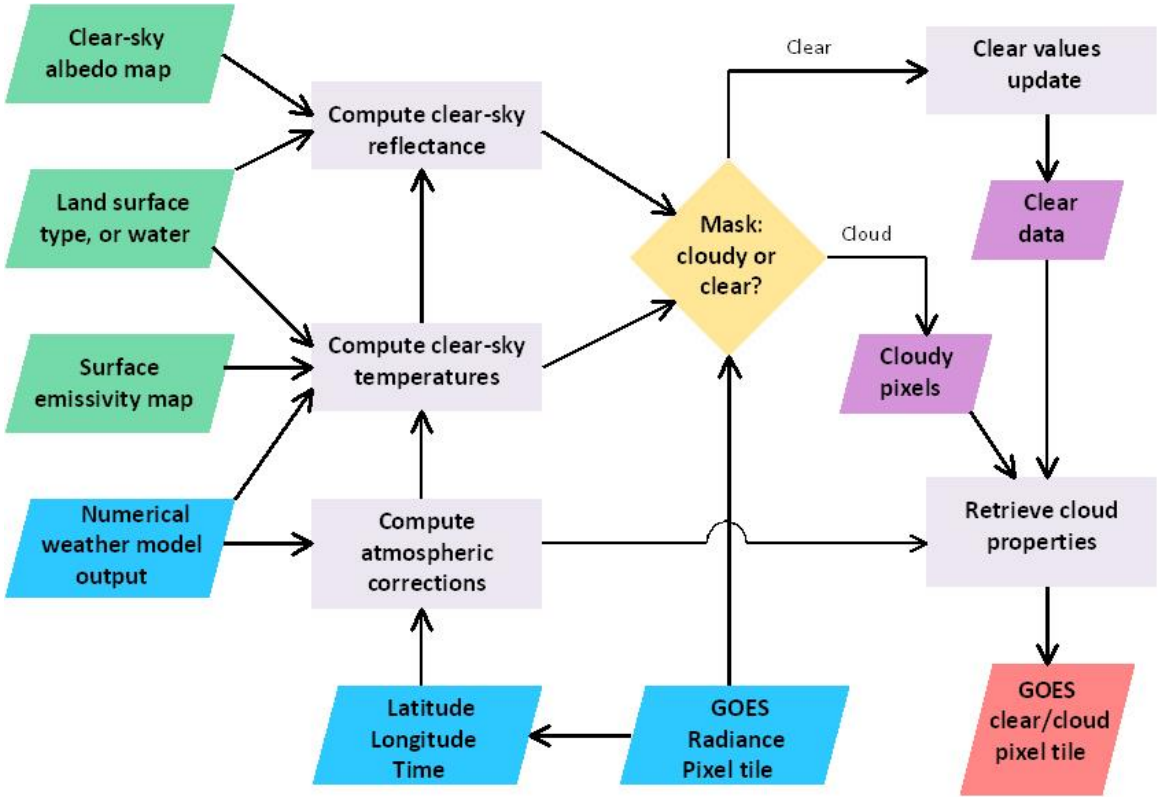

Fig. 2.34. Flow diagram illustrating the computation of cloud properties from GOES image data at the NASA Langley Research Center. Figure from Minnis et al. (2004).

Other time-varying input data (blue boxes in Fig. 2.34) come from the GOES imagers, which measure radiances as the instrument scans the Earth's surface. The radiances are processed into arrays covering an area roughly one-degree square in latitude and longitude and into products at the pixel level. Finally, information on current surface temperature and atmospheric conditions comes from a numerical prediction model.

Sun-Mack et al. (1999) describe the calculation of clear-sky reflectance. The observed clear-sky reflectance ρ is

$$\rho_\lambda = \rho_\lambda(K; LAT, LON; \mu_0, \mu, \psi),$$

where λ is the wavelength, K is the IGBP surface type, LAT and LON are the latitude and longitude, respectively, μ_0 and μ are the cosines of the solar and viewing zenith angle, respectively, and ψ is the azimuth angle measured between the two viewing zenith angles. For both the visible and near-IR channels on GOES, an expected reflectance corresponding to the observed one is computed from

$$\rho_\lambda = \alpha_{s\lambda}\delta_\lambda(K, \mu_0)\chi_\lambda(K, \mu_0, \mu, \psi),$$

where $\alpha_{s\lambda}$ is the surface albedo with both the sun and viewer overhead, δ_λ is the normalized directional reflectance model, which accounts for the land surface and variation of clear-sky albedo with the solar zenith angle, and χ_λ is the bidirectional reflectance model that accounts for, in addition, the viewing zenith and relative azimuth angles. The calculated reflectances do not account for atmospheric effects, either on the downward-directed solar radiation or the upward-reflected radiation. That is taken care of in the box labeled "Compute atmospheric corrections. As an example, for the near-infrared (NIR), the effects of water vapor absorption are accounted for by a multiplicative factor t_{NIR} applied to ρ_λ (with $\lambda = NIR$):

$$t_{NIR} = \exp\left[-\tau_{NIR}\left(\frac{1}{\mu_0} + 2.04\right)\right],$$

where τ_{NIR} is the water vapor optical depth, parameterized as a function of column-integrated precipitable water, surface pressure, μ_0, and LAT using detailed radiative transfer calculations and satellite data. For the visible channel, atmospheric corrections are made for Rayleigh scattering and ozone. The numerical model assists in the computation of atmospheric corrections because it contains vertical profiles of temperature and moisture.

The model surface temperature is useful for calculating the outgoing longwave radiances for clear skies such as that measured by Channels 4 or 5. This is not explicit in Fig. 2.34.

A set of decision trees determines whether a pixel is clear or cloudy. The decision trees were first discussed in Trepte et al. (1999) and a revised version more recently in Minnis et al. (2008). They are too involved for inclusion in this survey. The decisions involve clear-sky reflectances calculated for the visible and near-IR channels, the corresponding observed reflectances, and the brightness temperature for at least one longwave channel. At night, the visible channel

cannot be used. The final decision, clear or cloudy pixel, is rendered in the yellow box of Fig. 2.34.

The next major step in Fig. 2.34 is to retrieve cloud properties from computed clear-sky temperatures and reflectances and a comparison of observed vs. computed reflectances for cloudy conditions. The computations are made for the wavelengths observed in the different satellite channels. They are also made for a variety of solar zenith angles, viewing zenith angles, and relative azimuth angles: μ_0, μ, and ψ. Finally, they are made for visible optical depths ranging from 0.25 to 125, droplet effective radii from 2 to 32 µm, and for ice crystal effective diameters from 6 to 135 µm. Radiative transfer equations are used to compute the reflectances. The computations are so extensive that most of the results are stored in look-up tables, thus avoiding the need to recompute them every time a new set of reflectances is measured. Day and night calculations differ because only IR channels are useful at night. By matching observed with computed reflectances, using differences in observed brightness temperatures between channels, and the clear-sky values, it is possible to infer cloud-top temperature, cloud-top phase (liquid or ice), cloud optical depth τ_{VIS}, and effective drop radius r_e. The logic for doing so is complicated. The earliest comprehensive description is in Minnis et al. (1995). A more recent treatment is in Minnis et al. (2011a).

As noted earlier, knowing the cloud-top temperature and phase is important for estimating the likelihood of icing. Mean cloud droplet size relates directly to SLD. The liquid water path (LWP) may be computed simply from

$$LWP = \frac{4\delta_{liq} r_e \tau_{VIS}}{3Q_{VIS}},$$

where δ_{liq} is the density of liquid water (1 g cm^{-3}). Q_{VIS} is called the extinction efficiency. It ranges from 2.03 to 2.19 as r_e ranges from 32 to 4 µm.

Here are a few samples of cloud products produced routinely by NASA LaRC and available on their website http://www-angler.larc.nasa.gov. One can click on “East CONUS,” “West CONUS,” “Merged CONUS,” or over “North America (RR)” for a variety of products covering the U.S. or most of North America.

Figure 2.35 shows the 0.65 µm reflectance and a pseudo-color image for 1815 GMT, 05 February 2010 from GOES-E (GOES-12). Much of the area appears to be covered by bright clouds in the visible channel image (Fig. 2.35a). However, the pseudo-color image (Fig. 2.35b) reveals a complex weather situation, one that poses challenges for any satellite cloud analysis algorithm. This day featured a trough near the Mississippi River Valley, a surface frontal zone in the southeast U.S., a huge cloud shield from the Carolinas to the Dakotas, a strong upper tropospheric jet along the Eastern Seaboard, low-level cold air streaming offshore in New England, and, where the sky is clear, both bare and snow-covered ground.

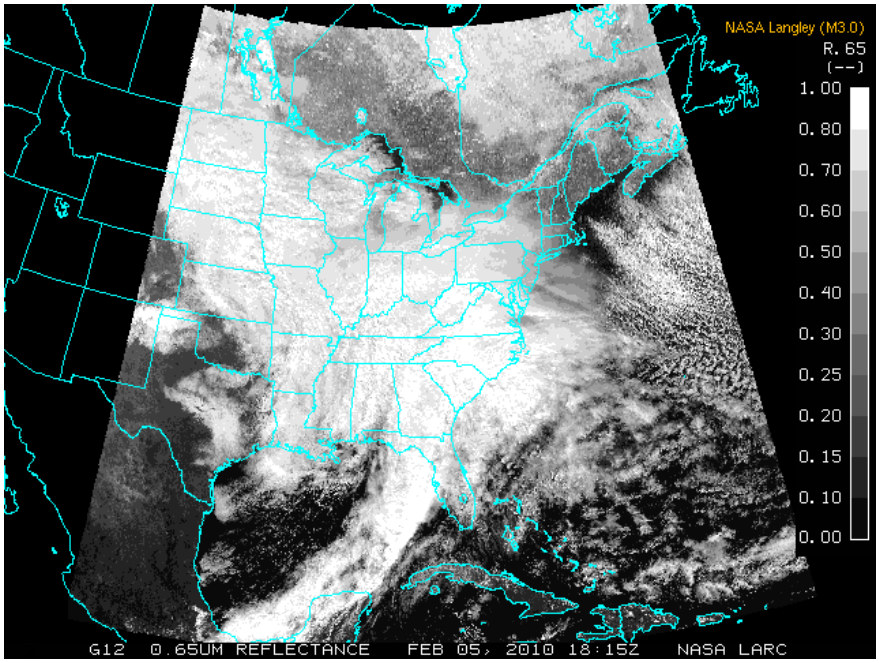

(a)

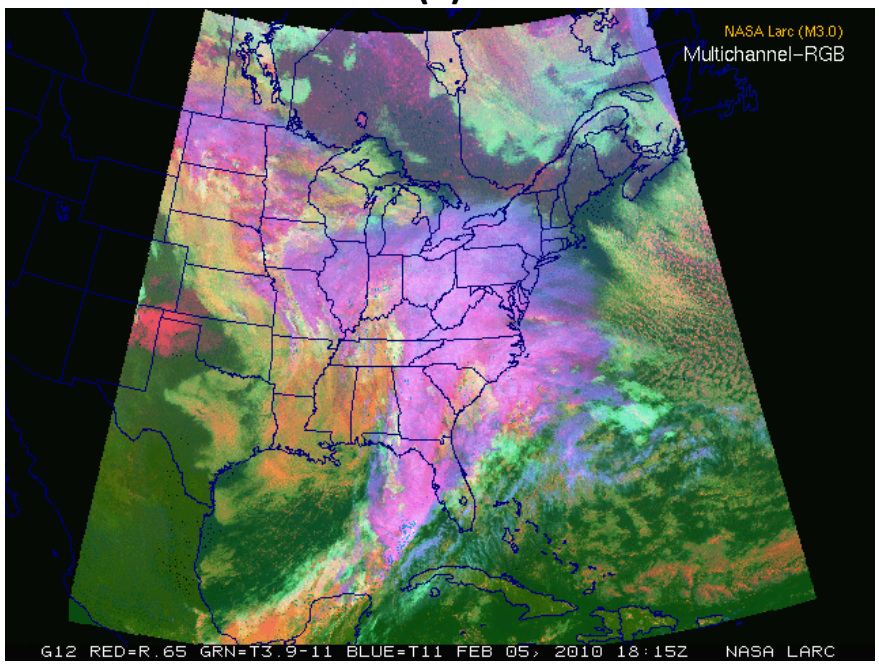

(b)

Fig. 2.35. (a) 0.65 µm reflectance from GOES-12, 1815 GMT, 05 February 2010. The gray scale for reflectance is at right. Values range from 0.0 to 1.0. (b) a corresponding multichannel image mixing colors from red, green, and blue guns. The intensity of red indicates the 0.65 µm reflectance, the intensity of green indicates the difference between the 3.9 µm and 10.8 µm

brightness temperatures, and the intensity of blue varies inversely with the 1.8 μm brightness temperature (bluer is colder).

As noted earlier, the phase of hydrometeors near cloud top is important for inferring icing. If the diagnosis is liquid phase in a sub-freezing cloud, the largest droplets are often found near cloud top. Figure 2.36 is a map of cloud phase for the same date and time as Fig. 2.35. The GOES imagers cannot see deeply into clouds at the longer absorbing wavelengths, and so the properties of cloud targets relying on these channels, such as cloud phase and effective drop radius, apply near the cloud top. However, since clouds absorb very little visible light, a wide range of cloud optical depths can be derived from the visible reflectance during the daytime, providing important information on the cloud geometric thickness and the vertically integrated density of cloud water.

If there are no clouds, the imagers see the ground. The algorithm detects snow-covered ground in and near the Texas Panhandle and in southeast Canada. Ice clouds cover many of the eastern states, but lower clouds with liquid near their tops cover a broad area surrounding the ice clouds. Their tops are colored dark blue if they are above freezing, and light blue if they are sub-freezing. Determining icing conditions at lower altitudes within those clouds whose tops contain ice crystals is more difficult but an active area of research. Clear areas with no snow cover are colored green.

The cloud phase product has limitations. Only the cloud phase of the highest cloud layer can be determined. Optically thin clouds containing supercooled water or both ice and liquid phases may be misclassified as warm liquid water clouds because of IR radiation penetrating these clouds from below.

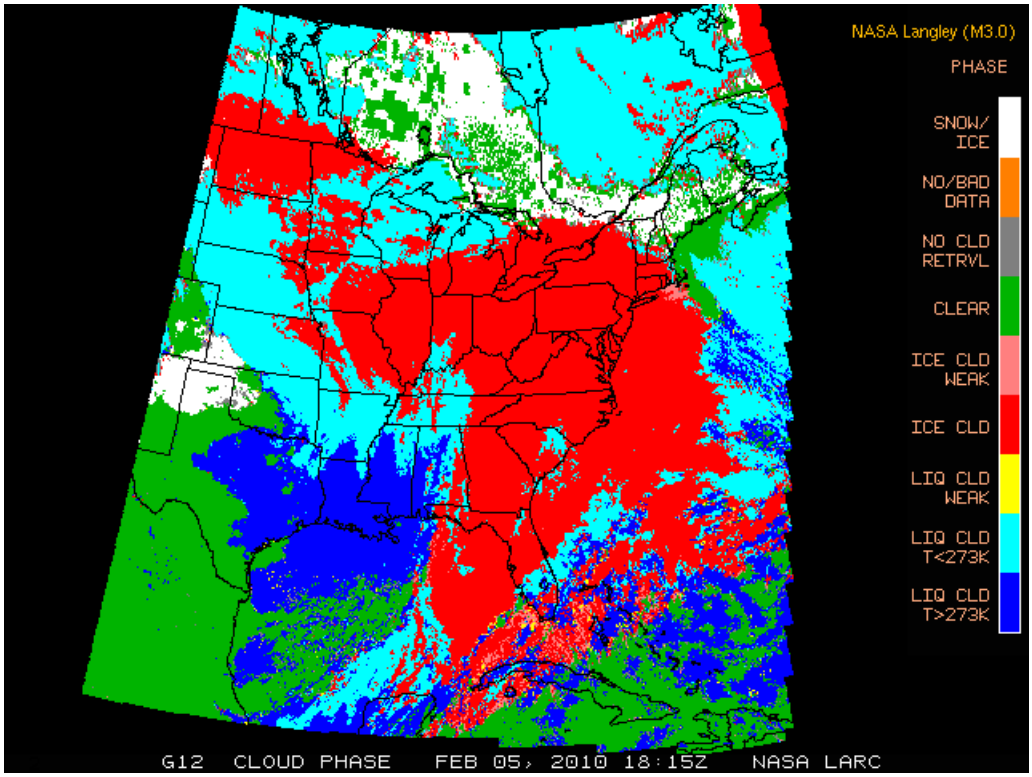

Fig. 2.36 The phase (liquid or ice) of hydrometeors near cloud top as inferred from GOES-12 imager data, 1815 GMT, 05 February 2010. The color bar at right allows interpretation of cloud properties. The word "weak" implies that there is considerable uncertainty in the diagnosis.

Figure 2.37 is a map of cloud optical depth τ_{VIS}. From the *Glossary of* Meteorology (Glickman, 2000), the cloud optical depth is the optical thickness between the top and bottom of a cloud. Optical depth is relatively independent of wavelength in the visible spectrum, but rises rapidly in the infrared due to absorption by water. In the visible spectrum, optical depth is controlled by scattering of light by cloud droplets and ice crystals. The scattering properties of ice crystals is uncertain because crystals come in many different shapes and sizes. Optical depth ranges from 0.1 in a thin cirrus cloud to over 1000 for cumulonimbus clouds. State-of-the-art retrievals of cloud optical depth at visible wavelengths is limited to values less than about 150 (128 in the current LaRC retrieval system), providing good sensitivity for all but the deepest cloud systems, which occur relatively infrequently over the Earth.

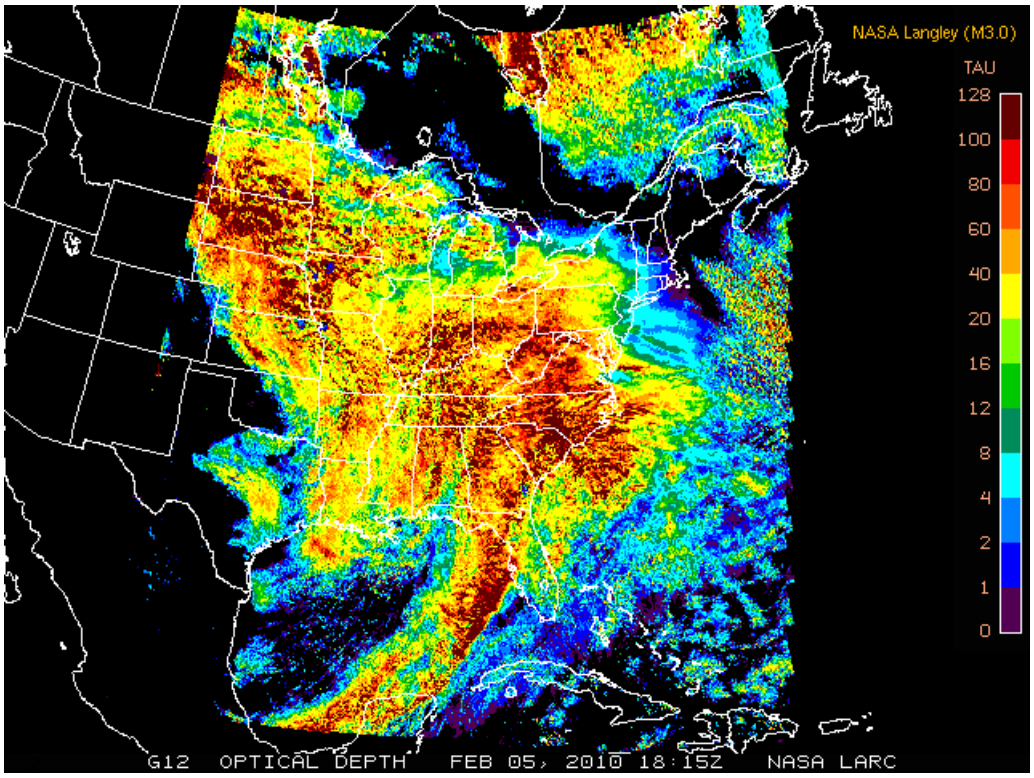

Fig. 2.37. Cloud optical depth τ_{VIS} (a dimensionless quantity) as inferred from GOES-12 imager data, 1815 GMT, 05 February 2010. The color bar at right gives the scale.

Figure 2.38 is a map of effective cloud droplet radius. This quantity is not relevant in the case of high and cold ice clouds, which are colored gray in the figure. As often noted above, larger droplets usually pose a greater risk of icing.

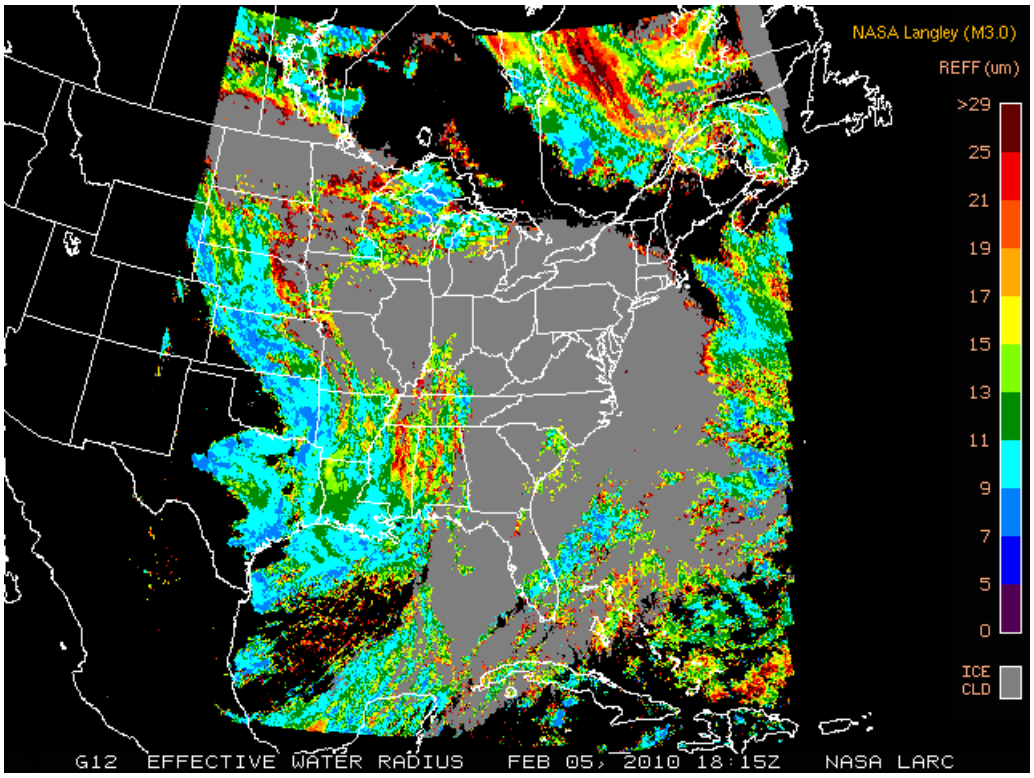

Fig. 2.38. Effective droplet radius as inferred from GOES-12 imager data at 1815 GMT, 05 February 2010. The color bar at right gives the scale in micrometers.

The liquid water path is very useful for inferring icing in supercooled clouds. As explained earlier, it can be approximated from the effective droplet radius and the optical depth. Figure 2.39 shows the liquid water path in g m^{-2}. In this image, many of the highest values occur near the edges of the ice clouds, a small percentage overall, but an indication that some of the pixels retrieved as liquid may in fact be contaminated by higher level ice clouds. A liquid water path of 100 g m^{-2} is sufficient to make a cloud a black body in the infrared (having an emissivity of 1.0), which is another way of saying that the IR channels on GOES cannot see very deeply into a water droplet cloud. This, in turn, explains the retrieval limitations at night.

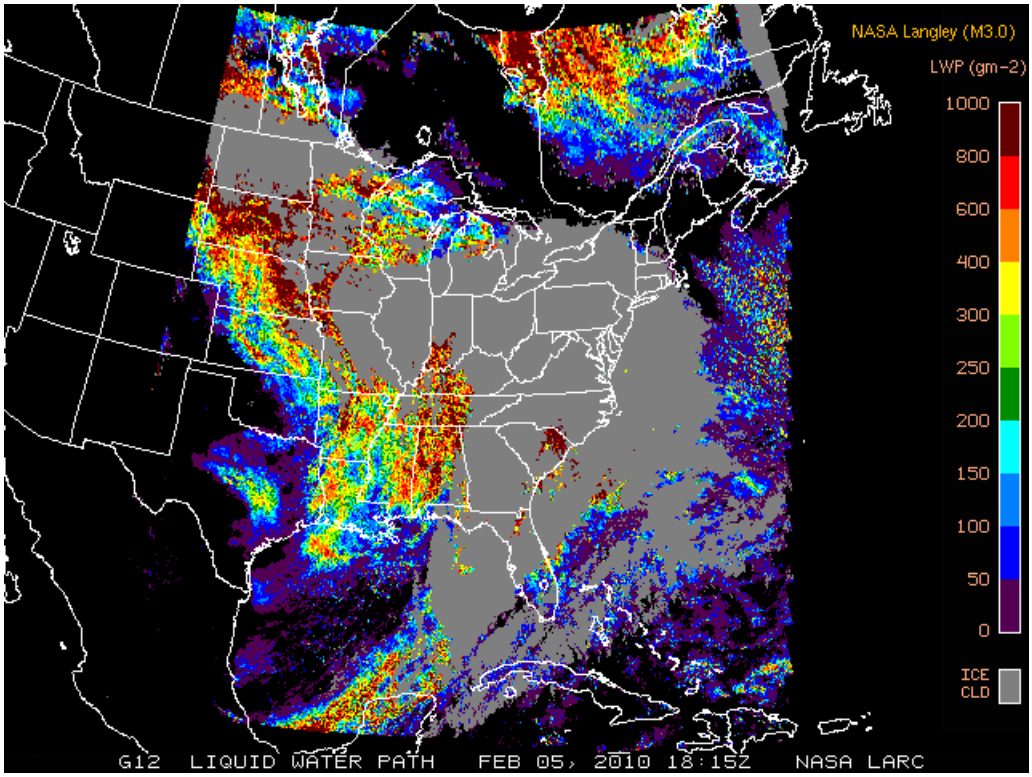

Fig. 2.39. The liquid water path as inferred from GOES-12 imager data at 1815, 05 February 2010. The color bar at right gives values in g m^{-2}.

Dong et al. (2002) compared values of r_e, LWP, and τ_{VIS} (last three figures) inferred from GOES images to corresponding values measured by aircraft and surface-based instruments, including a microwave radiometer that measures LWP. Figure 2.40 gives the results of four separate measurement episodes totaling ten hours, all of them in stratus clouds, when stratus were the only clouds present. The measurements were part of a field project in Oklahoma during March 2000.

There is moderate scatter in the points plotted in Fig. 2.40. The three mean values of LWP are quite close. Those for τ_{VIS} differ by as much as 13% and those for r_e differ by as much as 25%. GOES estimates of r_e are consistently higher than those from aircraft or surface instruments, but then GOES sees only the cloud tops, where droplet size is often greater. Both GOES and aircraft measurements of τ_{VIS} seem to be biased low with respect to surface measurements. GOES estimates of LWP compare favorably with both surface and aircraft measurements.

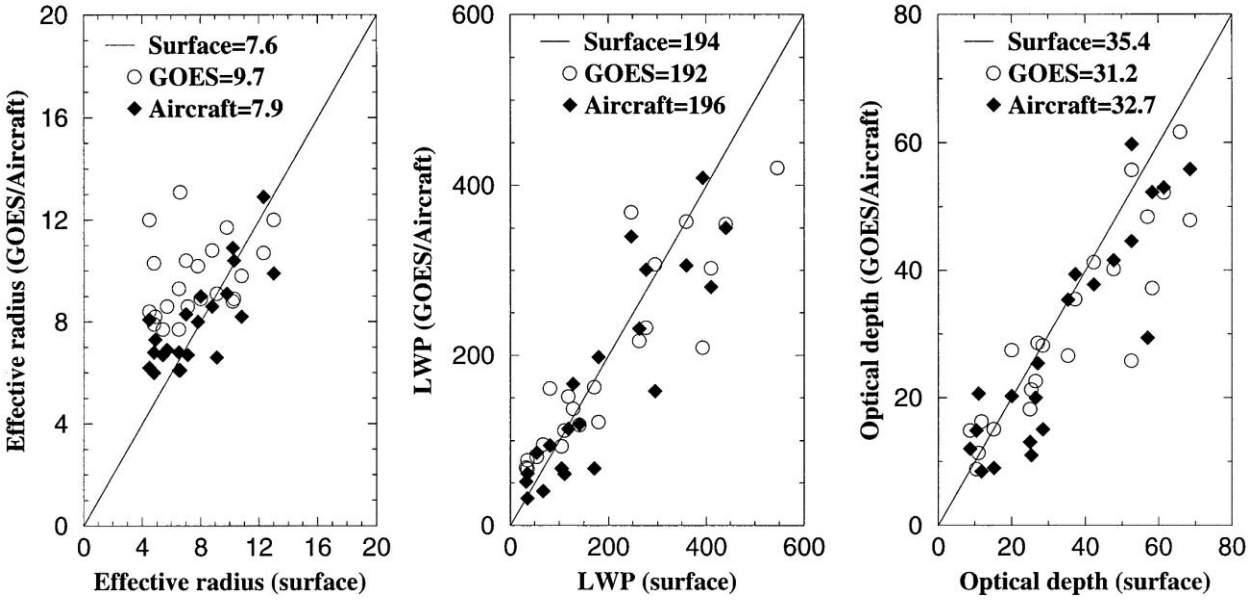

Fig. 2.40. Comparison of r_e, LWP, and τ_{VIS} from GOES, aircraft, and surface measurements taken in stratus clouds during four intensive observation periods in March 2002 over Oklahoma. Note that surface measurement values lie along the horizontal axis. The vertical axis is used for two data sources: GOES (open circles) and aircraft (solid diamonds). r_e is in μm, LWP is in g m^{-2}, and optical depth is dimensionless. The overall average values measured by each observing system—surface, GOES, and aircraft—are listed at the top of each panel. Used with permission of the American Meteorological Society.

In addition to providing input to the ASAP icing algorithms, the products can be used to diagnose icing directly, a valuable nowcasting tool, especially in remote areas where high-resolution models are currently unavailable. Minnis et al. (2004) developed an icing index that uses the satellite-retrieved cloud liquid water path, temperature, and effective radius to assign a probability of icing to a given area viewed by the satellite. This algorithm has been improved (Smith et al. 2011) and serves as the baseline algorithm for the suite of products planned for the first GOES-R satellite to be launched in 2015.

Figure 2.41a shows the icing probabilities determined from the results in Figs. 2.36-2.39. Most of the supercooled clouds yield a high probability of icing, either light (orange) or heavy (red). Where large optical depths and ice clouds occur, the icing probability cannot be determined with that algorithm and is deemed indeterminate (white). The indeterminate areas are extensive in this case, and limit the utility of using the data for diagnosing an icing threat, probably in some of the most dangerous conditions.

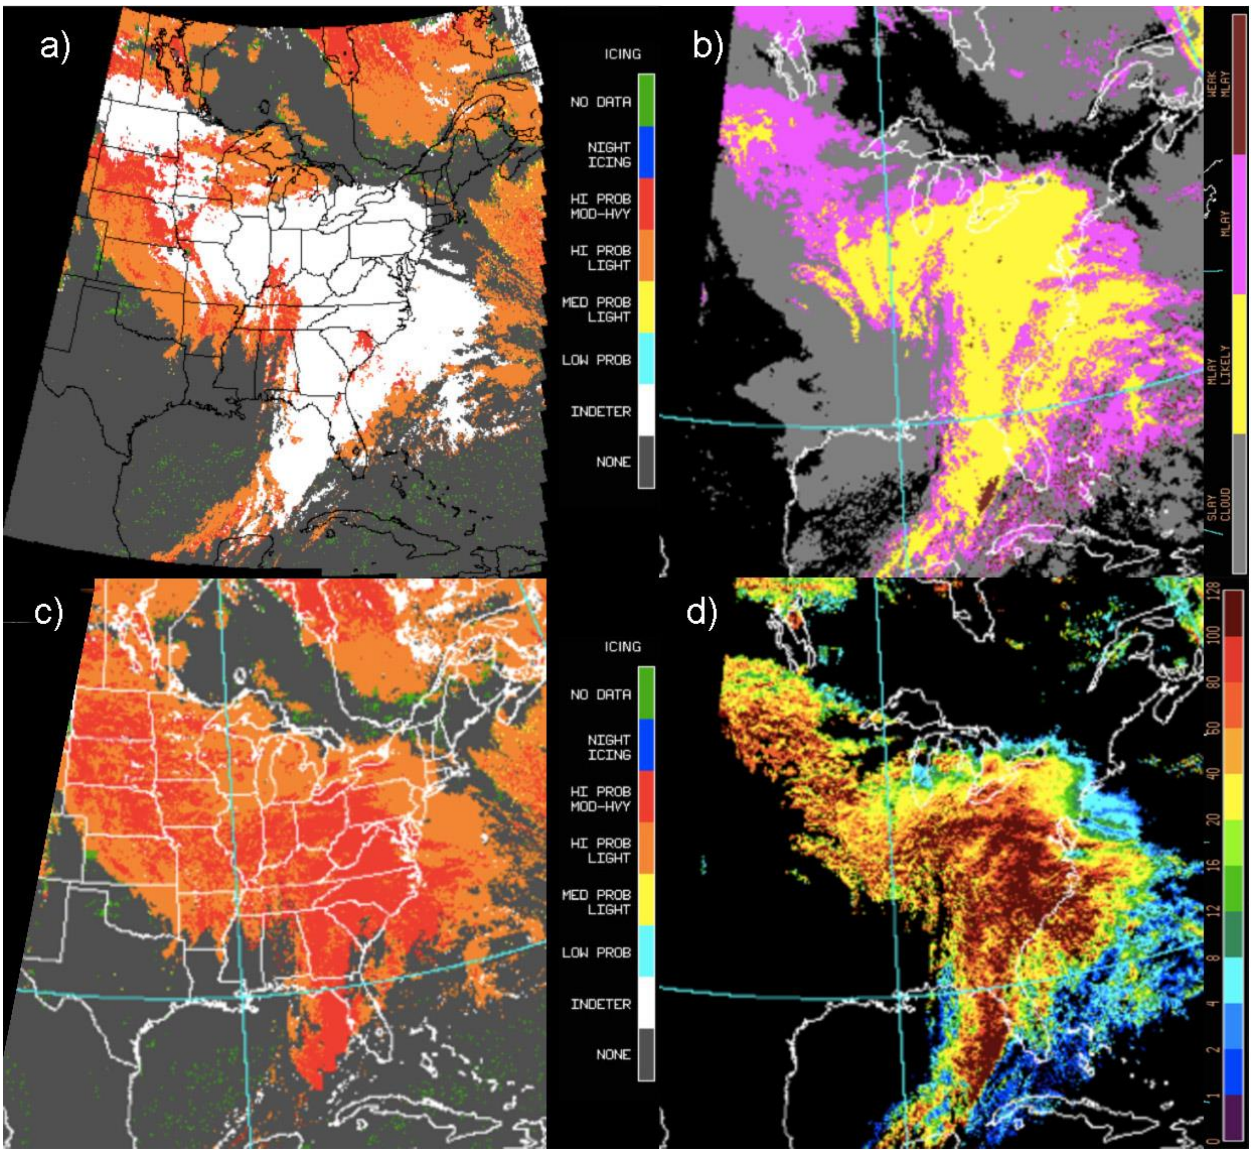

Fig. 2.41. Probability of icing and multilayer cloud properties for scene in Fig. 2.35, valid at 1815 GMT, 5 February 2010. (a) Icing probabilities for single-layer retrieval, determined from information in Figs. 2.36-2.39. Areas in white indicate that icing conditions cannot be determined because of the presence of an upper level ice cloud. (b) Likelihood of multiple cloud layers: pink – very likely, yellow – likely, brown – unlikely. Gray indicates single-layered clouds; (c) Icing probability combining single- and multi-layered cloud retrievals; (d) Lower (liquid) cloud layer optical depth under the upper ice cloud layer. Courtesy of Patrick Minnis, NASA LaRC.

The images in Fig. 2.41 are available on the web. Go to http://www-angler.larc.nasa.gov/ . Under "Cloud Products," click on "East CONUS." Specify date/time: Year 2010, Month 02, Day 05, Time 1815 UTC.
Panel (a): Under "Single Image," pick "Icing Potential."
Panel (b): Under "Multi-Layer," pick "Multi Layer ID."
Panel (c): Under "Multi-Layer," pick "Multi Layer Icing Potential"
Panel (d): Under "Multi-Layer," pick "Lower Layer Optical Depth."

To reduce the indeterminate area and enhance the utility of the satellite data, a method for identifying multilayered clouds and separately retrieving the properties of the lower and upper cloud layers has been developed and is part of the LaRC product suite. The method is currently applicable only to imagers having a CO_2 absorption channel (wavelength ~ 13.4 μm) along with the other channels typically used for the cloud property retrievals. Chang et al. (2010a) developed a two-channel technique that uses the GOES 10.8 and 13.3-μm channels to determine the pressure level and optical depth of the highest cloud in the scene, as well as the radiating temperature of the background under that cloud. With that information and the standard retrievals (e.g., Figs. 2.36-2.39), the algorithm determines for every pixel having a cloud pressure level lower than 400 hPa, whether the cloud is likely to be thin cirrus (optical depth < 5) over a lower level cloud. It is categorized as being multilayered (high probability), likely multilayered (medium probability), weak multilayered (probably contiguous ice and liquid water clouds, e.g., cumulonimbus), or not multilayered (Chang et al. 2010b).

Figure 2.41b shows the multilayered cloud probability for the case from Fig. 2.35. In this instance, it determined that most of the ice clouds overlay lower level clouds. For those clouds, it attempts a retrieval of the properties of both cloud layers. The optical depth retrieved for the lower cloud layer is shown in Fig. 2.41d. It indicates that most of the ice clouds were optically thin and most of the cloud optical depth can be attributed to the low clouds. The temperatures retrieved for the lower cloud layers were mostly less than 0°C, and the icing index algorithm concludes that icing threat is highly probable for those supercooled clouds underneath the cirrus deck (Fig. 2.41c). In this case, the indeterminate area has nearly been eliminated with the use of the multi-layered cloud retrieval algorithm.

Initial validations of both the standard (Smith et al. 2011) and the multilayered (Minnis et al. 2010) icing probability algorithms show that the positive icing diagnoses are extremely accurate (~98%) compared to pilot reports, but the skill for negative icing diagnoses is low. The icing probabilities for the multi-layered clouds are as accurate as those for single-layered clouds. The low skill for negative icing is attributable to many factors, for example, low pilot reporting of null icing, variations in aircraft sensitivity to icing, and aircraft location uncertainty, among others. More research is needed to refine the negative icing assessments. The single and multilayered icing probability products are both available, and ASAP is working to include the multilayered properties in their assimilation of the cloud products.

GOES products relating to cloud properties and icing similar to the above, except for the multilayered cloud properties, are also generated in real time at CIMSS. They are available at http://cimss.sssec.wisc.edu/geocat.

2.2.12 Algorithms under Development That Use Data from More Advanced Satellites

As part of the GOES-R Proving Ground program mentioned in Section 2.1.6.1, CIMMS, NASA LaRC and NASA Goddard are collaborating on algorithms that should improve detection of icing

conditions by satellite. On board GOES-R will be an advanced baseline imager (ABI), with 16 channels as indicated in Table 2.7. A quick comparison with Table 2.6 indicates many more imaging channels on GOES-R than on the current GOES and considerably higher spatial resolution, enabling detection of smaller clouds within the field of few. Researchers are simulating GOES-R measurements with data from other satellites, all in polar orbits but providing much the same spectral information as GOES-R will have on the ABI. For example, wavelength intervals for the channels on the MODIS instrument, also mentioned in Section 2.1.6.1, overlap the wavelength intervals for 14 of the 16 channels on GOES-R.

Table 2.7

Information about the Advanced Baseline Imager (ABI) to be flown on the GOES-R satellite with projected launch date in 2015. The channel number is at left. The wavelength interval for each channel is in the second column. This wavelength interval defines the limits within which the received signal is at least half maximum. The IGFOV is the instantaneous geometric field of view, the minimum width of objects that may be discerned when the view is straight down. From Schmit et al. (2005). Reprinted with permission of the American Meteorological Society.

Future GOES imager (ABI) band	Wavelength range (μm)	Central wavelength (μm)	Nominal subsatellite IGFOV (km)	Sample use
1	0.45–0.49	0.47	1	Daytime aerosol over land, coastal water mapping
2	0.59–0.69	0.64	0.5	Daytime clouds fog, insolation, winds
3	0.846–0.885	0.865	1	Daytime vegetation/burn scar and aerosol over water, winds
4	1.371–1.386	1.378	2	Daytime cirrus cloud
5	1.58–1.64	1.61	1	Daytime cloud-top phase and particle size, snow
6	2.225–2.275	2.25	2	Daytime land/cloud properties, particle size, vegetation, snow
7	3.80–4.00	3.90	2	Surface and cloud, fog at night, fire, winds
8	5.77–6.6	6.19	2	High-level atmospheric water vapor, winds, rainfall
9	6.75–7.15	6.95	2	Midlevel atmospheric water vapor, winds, rainfall
10	7.24–7.44	7.34	2	Lower-level water vapor, winds, and SO_2
11	8.3–8.7	8.5	2	Total water for stability, cloud phase, dust, SO_2 rainfall
12	9.42–9.8	9.61	2	Total ozone, turbulence, and winds
13	10.1–10.6	10.35	2	Surface and cloud
14	10.8–11.6	11.2	2	Imagery, SST, clouds, rainfall
15	11.8–12.8	12.3	2	Total water, ash, and SST
16	13.0–13.6	13.3	2	Air temperature, cloud heights and amounts

Refer to the right-hand column of Table 2.7 for sample uses of each channel. No less than eight channels (2, 5, 6, 7, 11, 13, 14, and 16) can be used for cloud detection and determination of properties relevant to icing. Here is a list of properties being studied and related experimental products generated by CIMMS, NASA LaRC, and NASA Goddard within the GOES-R Proving Ground program:

- Cloud Detection
- Temperature and pressure at cloud top
- Phase (liquid or ice) at cloud top (Inference of supercooled water requires a liquid signature and an IR temperature measurement.)

- Infrared emissivity
- Infrared microphysics (how much ice is in a cloud)
- Day optical depth and water path
- Day hydrometeor size
- Night optical depth and water path
- Night hydrometeor size

A CIMMS website http://cimss.ssec.wisc.edu/goes_r/proving-ground/SPC/SPC.html provides an early sampling of real-time products from the GOES-R Proving Ground program, but, as of this writing, not much is there pertaining to icing. NASA LaRC has two recently accepted publications (Minnis et al. 2011a, b) on the determination of cloud properties from TRMM VIRS and Terra and Aqua MODIS data; these papers give a preview of the capabilities to be anticipated from the ABI imager on GOES-R.

Two NASA satellites launched on 28 April 2006, CALIPSO and CloudSat, provide measurements for independent verification of some of the experimental products listed immediately above. CALIPSO (Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations) is a joint project of NASA and CNES (French space agency). This satellite measures aerosols and clouds 24 hours a day with a three-channel backscatter lidar. It flies in formation in the NASA "A-Train" with several other satellites, one of which is CloudSat. CloudSat carries a cloud profiling radar operating at 3.2 mm wavelength. It is always looking straight down and providing a vertical profile of reflectivity within cloud along the satellite's path. Zhang et al. (2010) used CALIPSO and CloudSat data to examine midlevel, liquid-layer topped stratiform clouds.

Andrew Heidinger, NOAA/NESDIS, University of Wisconsin at Madison (Email: Andrew.heidinger@noaa.gov), is a team lead within the GOES-R Proving Ground Project, specializing in the use of advanced satellites for detecting clouds, estimating their height, and determining their optical properties. He is a good contact for recent developments in this area.

2.3 Convection

Quoting from the *Glossary of Meteorology* (Glickman, 2000), convection involves "motions that are predominantly vertical and driven by buoyancy forces arising from static instability." The buoyancy forces give rise to accelerations in the vertical, which represent a departure from hydrostatic equilibrium. Thunderstorms spawn practically every hazard discussed in this survey: turbulence, icing, poor visibility, lightning, low-level wind shear, hail and tornadoes. Though thunderstorms get most of the publicity, moderate convection in thick stratocumulus clouds trailing a cold front and growing cumulus clouds in a warm, humid air mass also present

hazards, most notably, turbulence and icing. This section, however, emphasizes the detection and prediction of deep convection.

2.3.1 Radar and Lightning Observations

The U.S. network of approximately 130 WSR-88D Doppler radars gives fairly complete coverage of convective storms because they grow to such great heights. The radar reflectivity and radial velocity displays for each site are available on the ADDS website at: http://aviationweather.gov/adds/radar/. Click on the site desired, or view the radar data at national or regional scales. Figure 2.42 is an 88D reflectivity image from Vance Air Force Base, OK (VNX), showing several severe storms and the four warning areas associated with them.

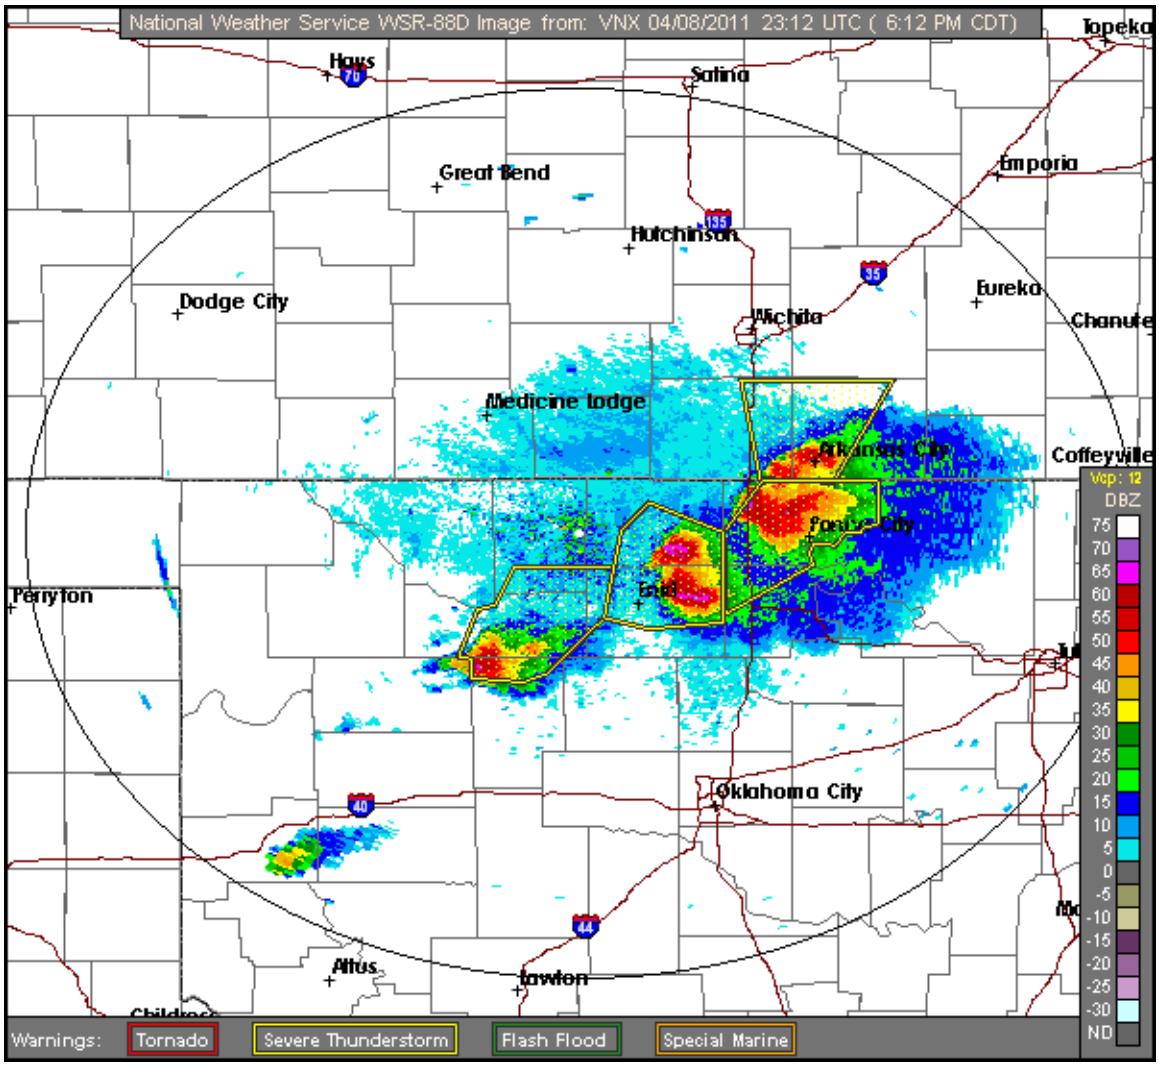

Fig. 2.42. Composite reflectivity image for WSR-88D radar at Vance Air Force Base, Oklahoma, 2312 GMT, 08 April 2011. In a composite image, the greatest reflectivity at any altitude is

displayed. Yellow polygons outline severe weather warning areas. Color bar gives the reflectivity scale.

Vaisala, a company that manufactures and sells meteorological instruments, operates the National Lightning Detection Network (NLDN), which includes more than 100 sensors that detect the time, location, polarity (positive or negative stroke), and current in cloud-to-ground lightning flashes. Sensors on the ground detect an electromagnetic signal unique to cloud-to-ground flashes. Data processing software then determines the location of the flash to within a median accuracy of 500 m. The flash detection rate is close to 95%. Thunderstorm detection rate is 99%. Lightning data are available only by subscription. NOAA purchases the data, which are displayed on Advanced Weather Interactive Processing System (AWIPS) workstations at NWS offices around the country. AWIPS displays a wide assortment of 88D radar data, not only reflectivity and radial velocity data but also the output of a number of algorithms for detecting severe storm characteristics.

Figure 2.43 is a sample image from the Memphis, TN, AWIPS workstation, displaying the local radar reflectivity with NLDN lightning data overlaid. Bright pink dashes show negative strokes and white plus signs show positive strokes occurring in the 15 min prior to 1630 GMT. As is often the case, the ratio of positive strokes to the total number is small.

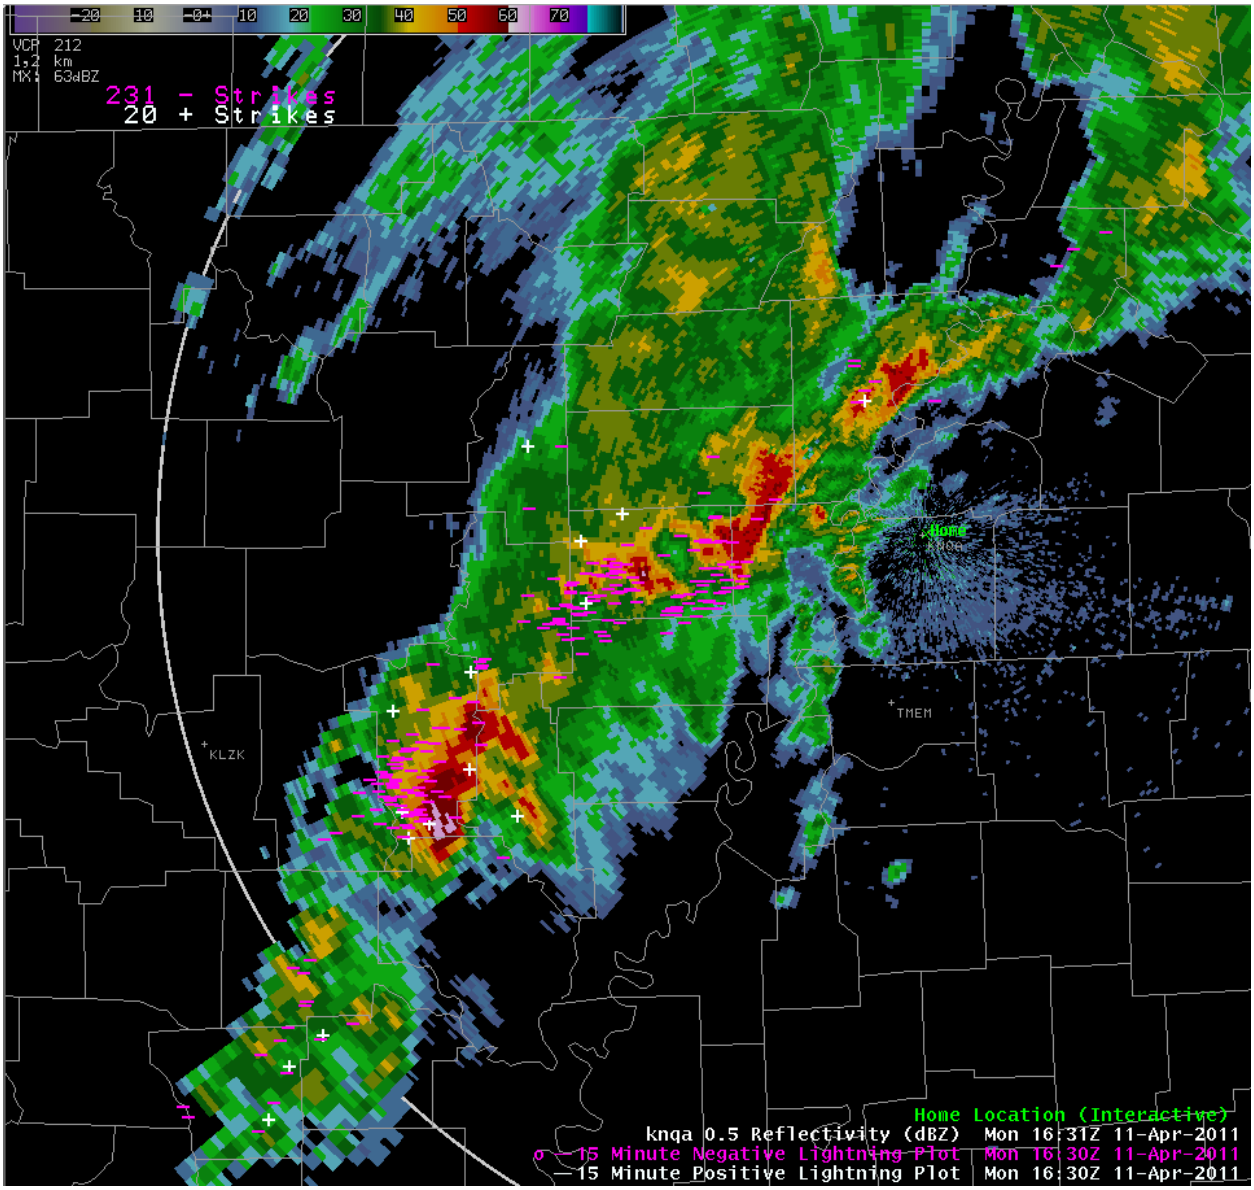

Fig. 2.43. Reflectivity data on the AWIPS workstation from the Memphis, TN, WSR-88D radar (NQA) with cloud-to-ground lightning strike data overlaid. The radar data are for 1631 GMT 11 April 2011. The bright pink dashes indicate negative strokes and the white plus signs indicate positive strokes.

2.3.2 Convective SIGMETs

The Aviation Weather Center issues convective SIGMETs whenever a NWS forecast office issues a severe weather warning for thunderstorm surface winds ≥ 50 kt, hail at the surface ≥ 1 inch in diameter, or tornadoes. Apart from this, they issue convective SIGMETs whenever one or more

of the following events is expected to exist for more than 30 min: a severe thunderstorm; thunderstorms embedded in stratiform precipitation; lines of thunderstorms at least 60 miles long, with thunderstorms affecting at least 40% of its length; and thunderstorms with VIP Level ≥ 4, affecting 40% or more of an area of at least 3,000 square miles. (VIP refers to the Video Integrator and Processor of the 88D radar). VIP Level 4 corresponds to reflectivity values between 44 and 50 dBZ.

One AWC forecaster per shift is responsible for issuing SIGMETs over the entire CONUS, which is divided into three regions: west (WSUS 33), central (WSUS 32), and east (WSUS 31). The WSUS designation appears at the top of each SIGMET message. SIGMETs are issued and/or updated hourly for each region and are valid for two hours. An outlook for 2-6 h accompanies each SIGMET. A SIGMET may be issued at any time based upon a severe weather report or under rapidly changing conditions. AWC issued nearly 30,000 SIGMETs during 2008. Fifteen SIGMETs an hour is not uncommon during busy periods.

The convective SIGMET forecaster relies heavily on lightning and radar data for information on current thunderstorm activity. For short-term changes in convection, forecasters go to the Storm Prediction Center website http://www.spc.noaa.gov/exper/. This site displays numerous diagnostic maps and selected predictions. In broad categories, they are:

- Upper-air analyses based upon rawinsonde data acquired at 0000 and 1200 GMT
- Analyses of sounding data acquired by all U.S. rawinsonde stations, including a host of parameters used in convective forecasting
- Mesoanalyses for nine different regions of the country at the surface and aloft
- Forecasts out to 87 h at 3-h intervals from the Short-Range Ensemble Forecast (SREF) produced every 6 h by EMC. The ensemble contains 22 members. SPC post processing of model output focuses on thunderstorms and other mission-critical events.
- Composite maps for overlaying selected fields from the RUC analyses and forecasts

Convective SIGMETs are available in both text and graphical formats at: http://aviationweather.gov/adds/airmets/. To see text messages, click on “Convective” at the top right of the page, choose the U.S. region of interest, then click “Retrieve.” Here is a text sample from 05 April 2011.

```
WSUS32 KKCI 041355
SIGC
CONVECTIVE SIGMET 35C
VALID UNTIL 1555Z
TN KY IN IL MS MO AR
FROM 40E STL-40E PXV-50SE MEM-60ESE LIT-40ENE FSM-40E STL
AREA SEV TS MOV FROM 24050KT. TOPS TO FL450.
TORNADOES...HAIL TO 2 IN...WIND GUSTS TO 60KT POSS.

OUTLOOK VALID 041555-041955
FROM BVT-40W MGM-140SSW LCH-30S SAT-50N SAT-BVT
REF WW 87 89.
```

```
WST ISSUANCES EXPD. REFER TO MOST RECENT ACUS01 KWNS FROM STORM
PREDICTION CENTER FOR SYNOPSIS AND METEOROLOGICAL DETAILS.
```

The graphical equivalent of the above text (but showing the whole country) is in Fig. 2.44.

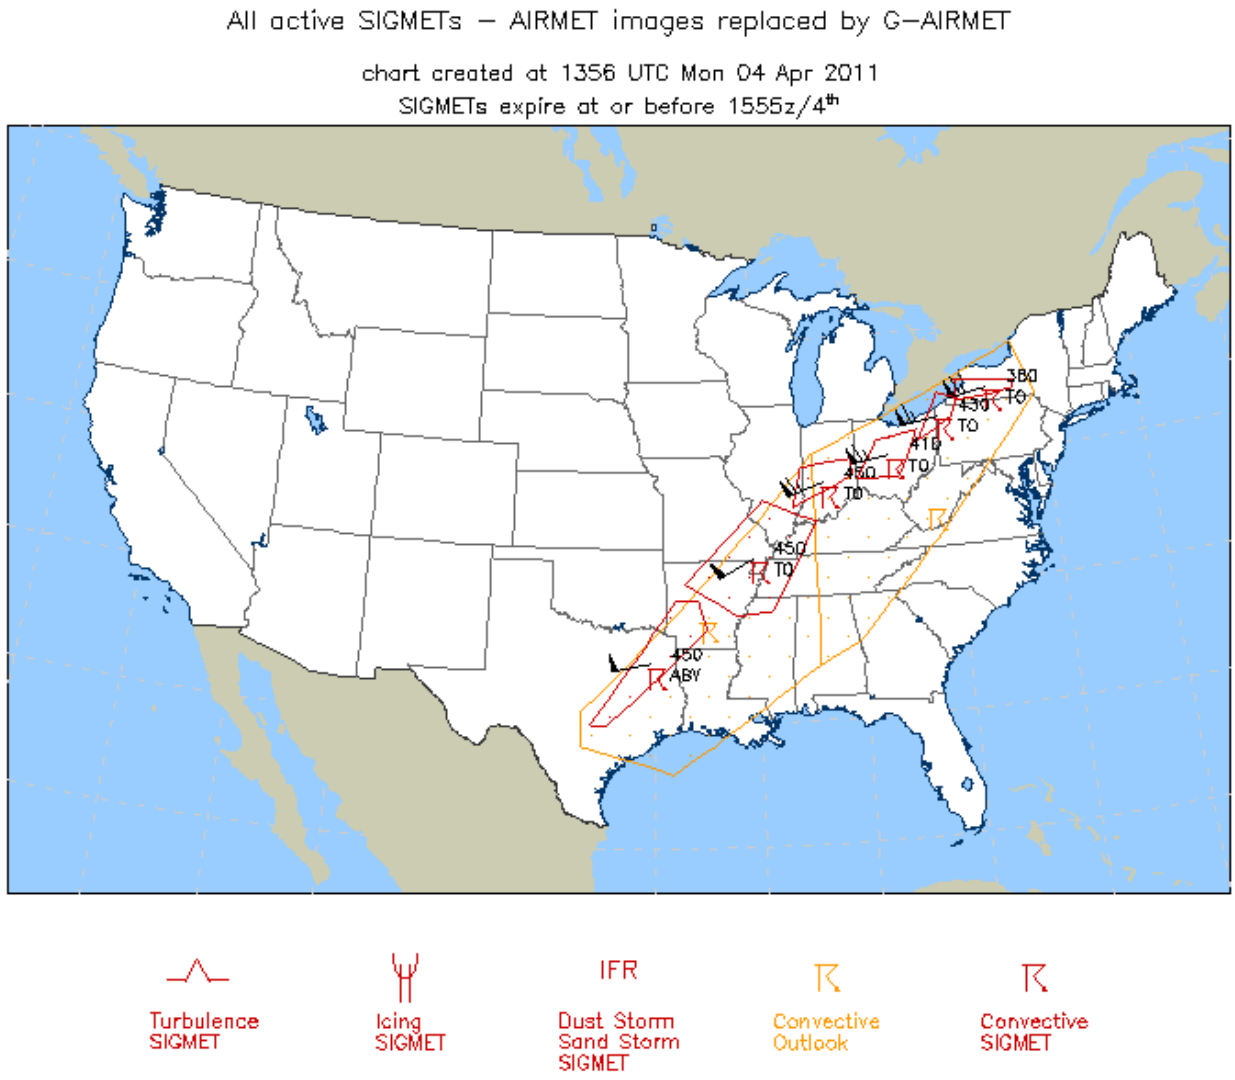

Fig. 2.44. Convective SIGMET issued at 1356 GMT on 04 April 2011. Six SIGMETS, outlined in red, are in effect at this time. The text message corresponding to the SIGMET covering part of Missouri has just been examined. So was the text for the western outlook area, one of two, outlined in orange.

The outlook message refers to severe weather watches Nos. 87 and 89, issued by the Storm Prediction Center in Norman Oklahoma. To view watches, go to http://www.spc.noaa.gov/products/watch/. Figure 2.45 gives details for Watch #89.

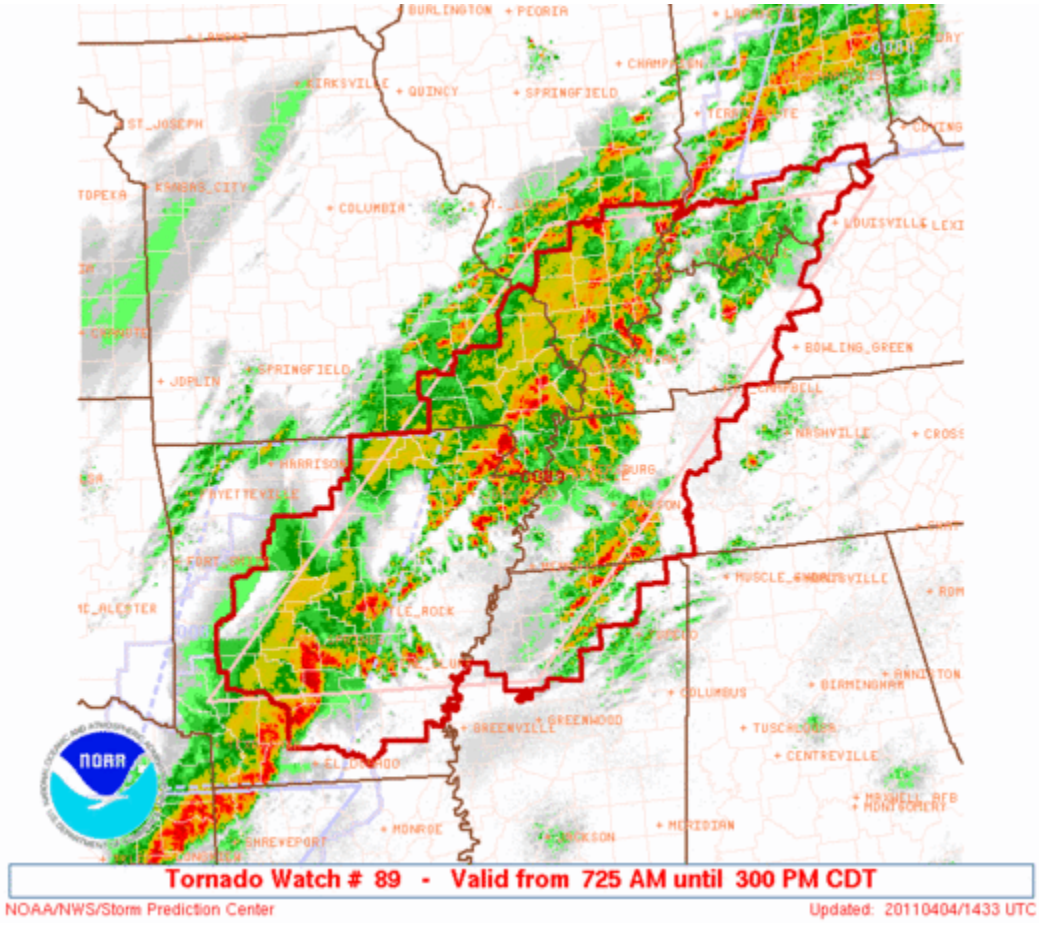

Fig. 2.45. Watch #89 for tornadoes, issued by the Storm Prediction Center on 04 April 2011. The watch box is a parallelogram in pink. The scalloped red boundary enclosing the watch box shows the borders of counties near the edge of the watch box. Radar echoes from 1433 GMT (0933 CDT) are shown. The progression of colors from green through gold to red indicates increasing echo intensity.

Here is the text that accompanies the image in Fig. 2.45.

```
URGENT - IMMEDIATE BROADCAST REQUESTED
   TORNADO WATCH NUMBER 89
   NWS STORM PREDICTION CENTER NORMAN OK
   725 AM CDT MON APR 4 2011

   THE NWS STORM PREDICTION CENTER HAS ISSUED A TORNADO WATCH FOR PORTIONS OF
          CENTRAL AND EASTERN ARKANSAS
          SOUTHERN ILLINOIS
```

SOUTHERN INDIANA
WESTERN KENTUCKY
SOUTHEAST MISSOURI
NORTHERN MISSISSIPPI
WESTERN TENNESSEE
EFFECTIVE THIS MONDAY MORNING AND AFTERNOON FROM 725 AM UNTIL 300 PM CDT.

TORNADOES...HAIL TO 1.5 INCHES IN DIAMETER...THUNDERSTORM WIND GUSTS TO 70 MPH...AND DANGEROUS LIGHTNING ARE POSSIBLE IN THESE AREAS.

THE TORNADO WATCH AREA IS APPROXIMATELY ALONG AND 115 STATUTE MILES EAST AND WEST OF A LINE FROM 20 MILES SOUTH SOUTHWEST OF PINE BLUFF ARKANSAS TO 35 MILES NORTH NORTHEAST OF EVANSVILLE INDIANA. FOR A COMPLETE DEPICTION OF THE WATCH SEE THE ASSOCIATED WATCH OUTLINE UPDATE (WOUS64 KWNS WOU9).

DISCUSSION...TSTMS ARE GRADUALLY INTENSIFYING THIS MORNING ALONG AND AHEAD OF COLD FRONT FROM AR INTO WRN KY. 12Z LIT SOUNDING SHOWED THAT STRONG CAP WHICH WAS OBSERVED BY 04/00Z SOUNDINGS HAD BEEN REMOVED...YIELDING AND MOIST AND STRONGLY SHEARED AIR MASS. EXPECT A COMBINATION OF LINE SEGMENTS AND SUPERCELLS TO EVOLVE WITH A THREAT FOR DAMAGING WINDS...HAIL AND TORNADOES. UPSCALE GROWTH INTO A QLCS IS ANTICIPATED LATER TODAY WITH THE PRIMARY HAZARD TRANSITIONING TO PRIMARILY DAMAGING WINDS.

AVIATION...TORNADOES AND A FEW SEVERE THUNDERSTORMS WITH HAIL SURFACE AND ALOFT TO 1.5 INCHES. EXTREME TURBULENCE AND SURFACE WIND GUSTS TO 60 KNOTS. A FEW CUMULONIMBI WITH MAXIMUM TOPS TO 500. MEAN STORM MOTION VECTOR 25040.

...MEAD

Despite the fact that most SIGMETs are issued on the basis of existing convection, the skill scores at two hours are not particularly impressive. Table 2.8 gives verification scores during the most active convection months for 2010. Definitions were given in section 1.4 on verification. SIGMETs are verified against the National Convective Weather Detection (NCWD) product that relies on a calculation of Vertically Integrated Liquid (VIL) from the WSR-88D reflectivity data at each tilt angle. The acronym for all of this is NCWD-VIL. NCWD and the corresponding forecast product (NCWF) will be described in more detail in section 2.3.6. For now, it is sufficient to know that NCWD-VIL depicts VIL values of 3.5 kg m^{-2} or greater, but only for radar cloud tops lying above 17,000 ft. This product is one of the standard NEXRAD algorithms; it is generated on a 4-km grid.

A “yes” forecast occurs inside any SIGMET box. A “yes” observation occurs within any NCWD grid box showing VIL ≥ 3.0 kg m^{-2}. A “YY” is presumed anytime one 4-km x 4-km “yes” box lies within, or partly within, a SIGMET. Thus, many YY events (480 or more) can theoretically occur within a given SIGMET since its area is at least 3000 mi^2. There are very roughly a million 4-km x 4-km VIL boxes in the CONUS (roughly 3000 km by 5000 km in area), the great majority without convection in them. Thus, it is not surprising that the “NN” counts are in the hundreds of millions for each month.

Table 2.8

Verification of SIGMETS during the most active convection months in 2010. Verification is against the National Convective Weather Forecast detection field.

Month	YY (1000s)	YN (1000s)	NY (1000s)	NN (1000s)	PODy	FAR	CSI
Apr 2010	266	3,955	262	467,044	0.504	0.937	0.059
May 2010	665	9,118	737	492,491	0.474	0.932	0.063
Jun 2010	816	14,222	1,043	471,904	0.439	0.946	0.051
Jul 2010	797	19,222	1,164	486,836	0.407	0.960	0.038
Aug 2010	695	16,441	971	492,059	0.417	0.959	0.038
Sep 2010	449	9,235	548	484,193	0.450	0.954	0.044

SIGMETs covered less than 4% of the CONUS area, even for July 2010, the most active month. Except in April 2010, note that more convective events were observed outside of SIGMET areas than inside; the PODy scores reflect this directly. The PODy scores are higher in the spring and fall than in summer probably because convection is more strongly forced by atmospheric dynamics in the transition seasons than in summer. "Popcorn" convection is the rule in summer. The false alarm rate is quite high, greater than 93% in all months. The CSI (the number of correct "yes" forecasts relative to the number of "yes" forecasts and observations) is disappointingly low at two hours, illustrating the difficulty in making accurate convective forecasts. A perfect score is 1.0.

2.3.3 Center Weather Advisories and Meteorological Impact Statements for Convection

Large-scale convective events can usually be anticipated, and their effects on air travel mitigated ahead of time. Unanticipated events, however, can cause havoc, especially in the early evening, the busiest travel time, and in the Northeast Urban Corridor, which has the most congested airspace. Worst of all, are thunderstorms lying across arrival gates, of which there are usually just four at the biggest airports. CWSUs maintain a continuous watch, especially around busy airports, for thunderstorms affecting approach and departure paths. They issue a CWA whenever criteria for a convective SIGMET are met but the thunderstorm coverage is too small (under 3000 mi^2) to merit a SIGMET from AWC, or if more widespread activity develops that is not covered by a current SIGMET. CWSU meteorologists are supposed to coordinate with AWC when they issue a convective CWA. Air traffic controllers must transmit SIGMETs and CWAs to pilots at least once upon issuance. Here is an example, issued 5 April 2011.

Miami ARTCC (KZMA)

Center Weather Advisory (CWA) ZMA CWA 102 VALID UNTIL 051933
FROM 80E PBI TO 105SE MIA
LINE...30NM WIDE...OF SCT TS WITH HVY-EXTRM PCPN MOV 24035KT. MAX
TOPS TO FL450. EXP LTL CHG THRU PD.

In plain language, this advisory is issued by the Miami Air Route Traffic Control Center, valid until 1933 GMT, 5 April. From 80 nautical miles east of West Palm Beach, FL, to 105 nautical miles southeast of Miami, FL, there is a line of scattered thunderstorms, 30 nautical miles wide, with heavy to extreme precipitation, moving from 240° at 35 knots. Thunderstorm tops extend to flight level 45,000 ft. Expect little change through the period.

CWSU meteorologists can also issue a Meteorological Impact Statement (MIS) for convection. An MIS is a discussion product that summarizes weather conditions that may affect air traffic routing or flight operations in the CWSU area of responsibility. The MIS is not intended for pilots. Here is a sample MIS for convection, issued by the CWSU at the Jacksonville, FL, ARTCC.

Jacksonville ARTCC (KZJX)
Meteorological Impact Statement (MIS) ZJX MIS 02 VALID 051330-052130
...FOR ATC PLANNING PURPOSES ONLY...
OVR SC/GA ATLC WATERS/FL PEN SCT TO NMRS TS MOV
FM 30050KT. ISOLD SEV TS/TORNADOES POSS. MAX TOPS
FL500. TS CONTG OVR ATLC WATERS THRU 2130Z ENDG
OVR FL ZJX BY 17Z. SC/GA/FL MOD TURB BLW FL400. ...PR...

Translation: The valid times for this MIS are from 1330 to 2130 GMT, 05 April 2011. Over South Carolina, Georgia Atlantic waters, and the Florida Peninsula, scattered to numerous thunderstorms moving from 300° at 50 knots. Isolated severe thunderstorms with possible tornado and highest cloud tops to flight level 50,000 ft. Thunderstorms continuing over Atlantic waters through 2130 GMT, but ending over Florida Jacksonville ARTCC area by 1700 GMT. Moderate turbulence below flight level 40,000 ft over South Carolina, Georgia, and Florida.

CWAs and MISs are distributed to ARTCC and collocated Traffic Management Unit (TMU) personnel. They appear on the website for CWSU products at http://aviationweather.gov/products/cwsu/.

2.3.4 Collaborative Convective Forecast Product

The Collaborative Convective Forecast Product (CCFP) is generated by a meteorologist at AWC for the CONUS and parts of Quebec and Ontario near the Great Lakes. The CCFP is generated from March through October for the CONUS and from April through September for southeast Canada. The issue times are between 0300 and 2300 Eastern Local Time (Eastern Standard Time or Eastern Local Time, depending upon the season). Once every two hours, the meteorologist prepares three forecast maps with lead times of two, four, and six hours. A forecast is issued if all the following conditions are met:

- The polygon depicting the hazard must cover at least 3,000 square miles.

- At least 25% of the polygon must be covered with echoes of at least 40 dBZ composite reflectivity. (This corresponds to VIP level 3.)
- At least 25% of the polygon must be covered with echo tops of at least 25,000 ft.
- In the forecaster's judgment, the likelihood of these conditions occurring must be at least 25%.

Note that the above conditions do not necessarily imply a thunderstorm. The CCFP is basically an echo-tops product because, whenever tops exceed 25,000 ft, the reflectivity almost always exceeds 40 dBZ. CCFP echo tops are not the maximum tops within the polygon but rather the maximum altitude at which at least 25% of the area is cloud covered.

In addition to current radar data and satellite imagery, the forecaster examines output from the RUC and NAM models. The RUC model produces a new analysis of atmospheric conditions every hour, which drives a forecast out to 18 h. The NAM model runs every six hours and produces forecasts out to 84 h. Of particular interest are the convective precipitation forecasts and the high-resolution point soundings, processed from model output, which aid in the estimation of convective cloud tops. Other useful products available on AWIPS are stability indices and maps of Convective Available Potential Energy (CAPE) and Convective Inhibition (CIN). The former is a vertically integrated measure of the buoyant energy that can be released once volumes of air are lifted to their level of free convection; the latter is a measure of the strength of the "cap," a stable layer usually topping the boundary layer that must be either removed by mid-tropospheric cooling or breached before deep convection can occur.

The Storm Prediction Center (SPC) issues Mesoscale Discussions helpful to the CCFP forecaster. These are available at http://www.spc.noaa.gov/products/md/ and pertain to the next few hours. Figure 2.46 is an example. The text accompanying the figure follows.

```
MESOSCALE DISCUSSION 0491
NWS STORM PREDICTION CENTER NORMAN OK
0116 PM CDT WED APR 20 2011

AREAS AFFECTED...FAR ERN MS...NCNTRL TO SCNTRL AL...SW GA

CONCERNING...SEVERE POTENTIAL...WATCH POSSIBLE

VALID 201816Z - 201945Z

A SEVERE THREAT IS EXPECTED TO DEVELOP THIS AFTERNOON ACROSS CNTRL TO ERN
MS...NCNTRL TO SCNTRL AL AND SW GA. LARGE HAIL AND WIND DAMAGE WILL LIKELY
ACCOMPANY THE STRONGER CELLS. A WW MAY BECOME NECESSARY AS THE THREAT
INCREASES THIS AFTERNOON.

A CLUSTER OF STRONG TO SEVERE THUNDERSTORMS IS LOCATED IN NCNTRL MS ALONG A
GRADIENT OF MODERATE INSTABILITY. A SECOND CLUSTER OF DEVELOPING STORMS IS
LOCATED IN SCNTRL AL. AS SFC TEMPS CONTINUE TO WARM THIS AFTERNOON.
CONVECTIVE COVERAGE SHOULD INCREASE. THE MODERATE DEEP LAYER SHEAR EVIDENT ON
REGIONAL WSR-88D VWPS ALONG WITH STEEPENING LOW-LEVEL LAPSE RATES SHOULD
RESULT IN A WIND DAMAGE THREAT WITH THE MORE INTENSE CELLS. AN ENHANCED WIND
```

DAMAGE THREAT COULD DEVELOP IF A COLD POOL CAN DEVELOP OR A THE STORM CLUSTERS CAN CONGEAL INTO A FASTER-MOVING LINE-SEGMENT. IN ADDITION...A LARGE HAIL THREAT MAY ALSO DEVELOP WITH ANY STORMS THAT BECOME DOMINANT OR OBTAIN SUPERCELL STRUCTURE.

..BROYLES.. 04/20/2011

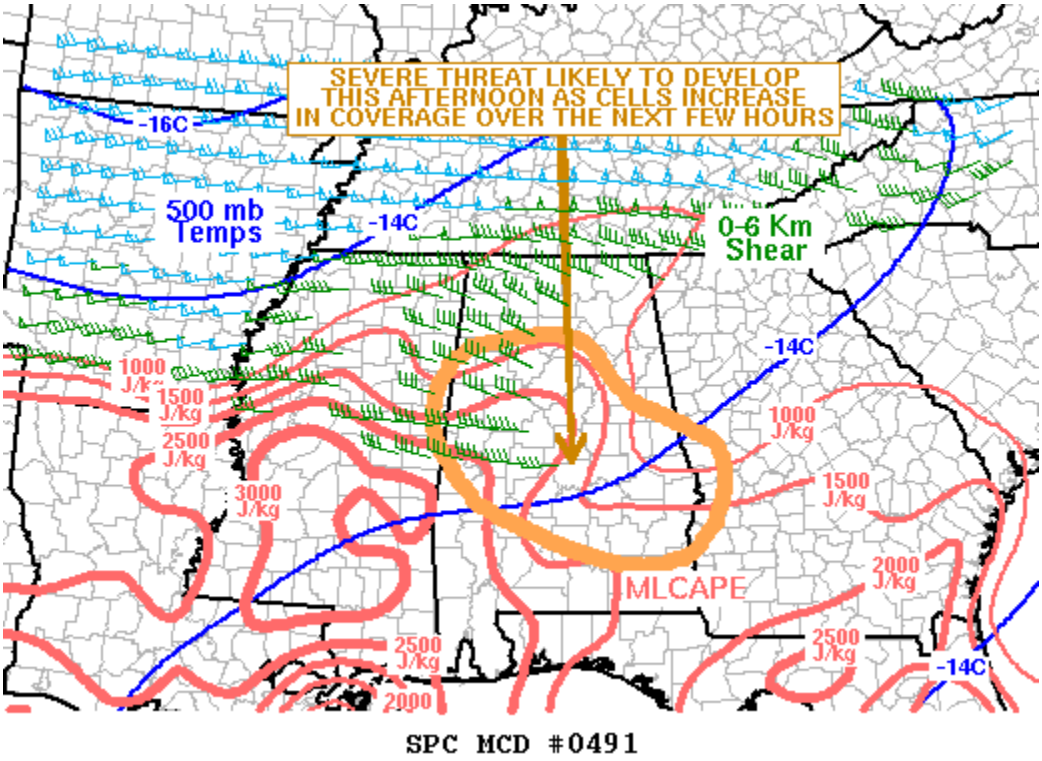

Fig. 2.46 Weather map germane to Mesoscale Discussion #491 issued by the Storm Prediction Center and valid from 1816-1945 GMT, 20 April 2011. The dark blue contours are 500-mb isotherms. The wind barbs in green and blue give the vertical shear vector from 0-6 km. The pink contours are for CAPE in J kg-1. The wide orange contour outlines the area under threat for severe weather.

A successor to the RUC model, the Rapid Refresh (RAP), should become operational by late 2011. The computational domain will cover most of North America at 13-km resolution. The Global Systems Division (GSD) of NOAA's Earth System Research Laboratory (ESRL) regularly runs an experimental, high-resolution version of the RAP (HRRR) over the CONUS at 3-km resolution. Output from this model is available at http://rapidrefresh.noaa.gov/hrrrconus/ most of the time and has proven quite useful for CCFP meteorologists.

The CCFP is a *strategic* forecast in the sense that it affects the future routing and scheduling of aircraft. It is primarily intended for air traffic management, which includes both the FAA and the airline industry. The product is unique in that those who have a stake in aircraft routing (which affects arrival and departure times) can comment on a draft CCFP before it becomes final.

Chat room sessions open no later than 45 min before the final CCFP is issued and close automatically 15 min before issue time. The Air Traffic Control System Command Center (ATCSCC) in Warrenton, Virginia, balances air traffic demand with system capacity in the National Airspace System (NAS). ATCSCC participates in a session early each morning while defining the problems of the day. Other participants at various times are CWSUs, ARTCCs, SPC, airline meteorologists, and selected NWS forecast offices. Points raised by participants may lead to alterations of the CCFP, but the AWC forecaster makes the final decisions.

Figure 2.47 shows a sample CCFP. The legend deserves a brief explanation. *Tops* refers to radar echo tops. The maximum altitude at which at least 25% of the polygon area is covered by at least 18.5 dBZ reflectivity is relevant. (Echo tops is a NEXRAD product.) Say that this altitude is 32,000 ft MSL. This falls within the range 30,000 to 34,000 ft, and so the tops listed for the polygon would be 340, representing the top of this altitude range in hundreds of feet. If tops within this altitude range cover 25-39% of the polygon area, the coverage is called sparse, and the polygon is hatched with dashed lines. Medium coverage (hatching with solid lines) means that tops cover 40-74% of the area. Solid coverage (solid shading) means that tops cover 75-100% of the area. If forecaster confidence is between 25 and 49% (low) that criteria for

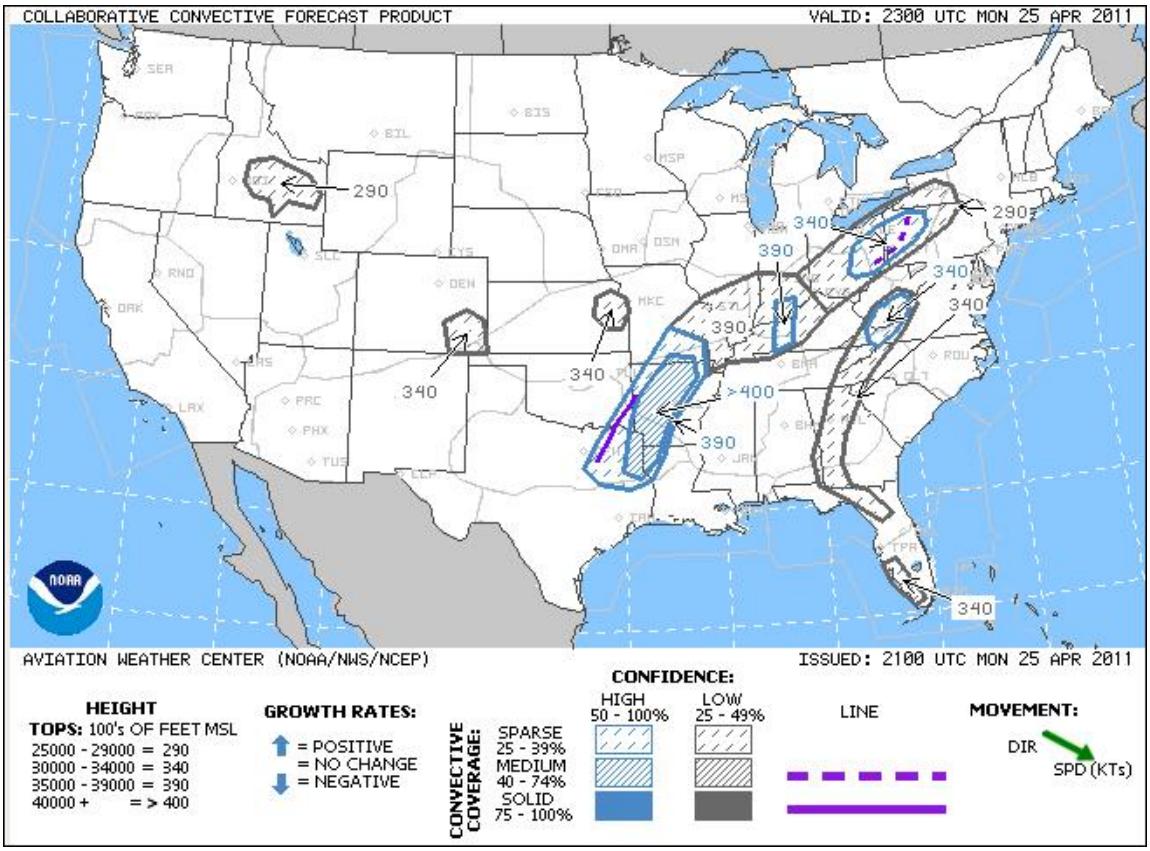

Fig. 2.47. A CCFP issued at 2100 GMT and verifying at 2300 GMT 25 Apr 2011. The legend at bottom explains the information conveyed.

drawing a polygon are met, the polygon and the hatching or shading inside are colored gray. If forecaster confidence is 50% or more (high), blue is used. For medium or high coverage, confidence is assumed to be high. Finally, convective lines may be included in the CCFP. Solid purple indicates that CCFP criteria are met along 75-100% of the line. Dashed purple indicates criteria are met along 40-74% of the line. The lines may be inside or outside of polygons.

CCFPs are routinely verified. Go to http://aviationweather.gov/products/ccfp/ and click on "2-hr verification." Figure 2.48 is an example.

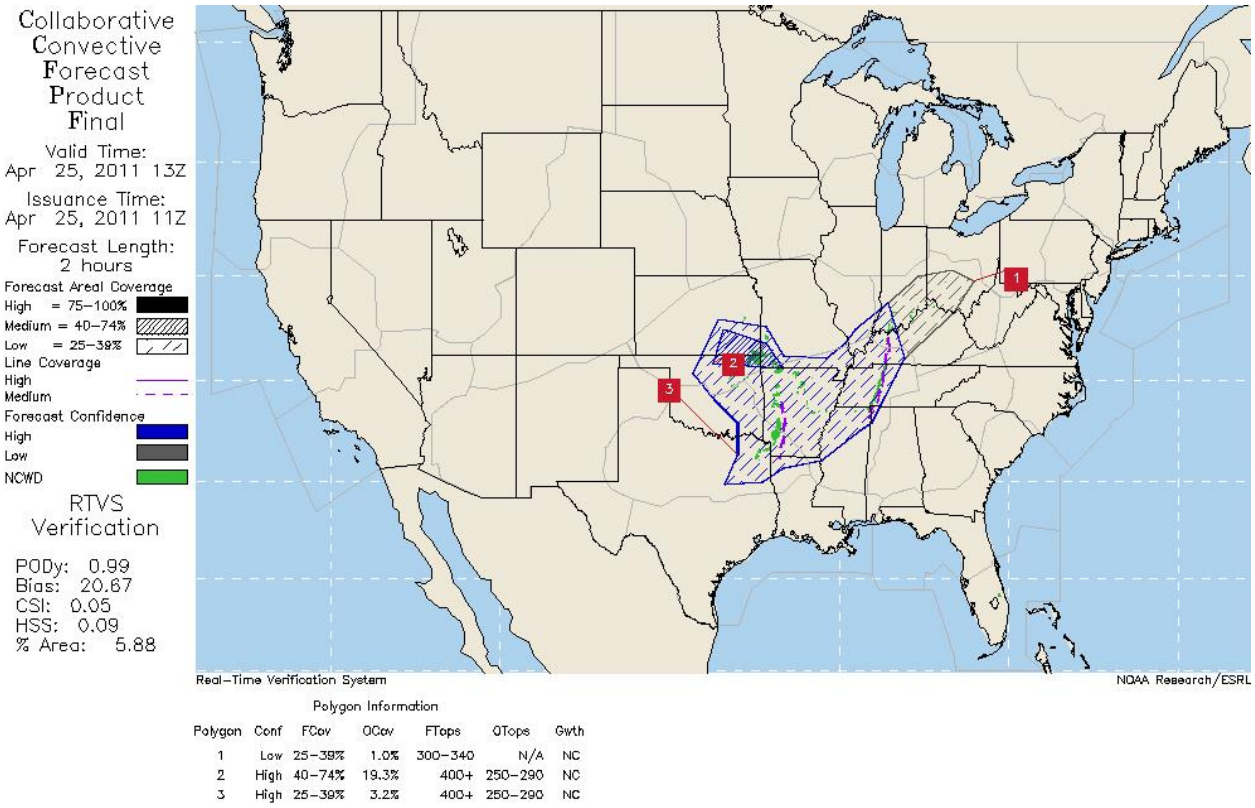

Fig. 2.48. Verification of a 2-h CCFP forecast valid at 1300 GMT 25 April 2011. The forecast identified three convective regions and two convective lines, using the coloration and shading indicated to the left of the figure. Observed convection is shaded green.

The three CCFP regions are numbered. In regions 1 and 2, forecaster confidence was high (50% or more); in region 3, confidence was low (25-49%). Coverage (definition: VIP levels 3 or greater and radar tops 25,000 ft or greater) was expected to be low (25-39%) in regions 1 and 3 but medium (40-74%) in region 2. The 2-h forecast of the location of the two convective lines

was accurate. 99% of the observed convection occurred within the CCFP regions (PODy). The areal coverage within the regions was grossly overforecast (bias 20.67). Note the polygon information at the bottom of Fig. 2.48. Area 1 had only 1% coverage, area 2, where medium coverage (40-74%) was expected, actually had 19.3% coverage—not a bad convective forecast. Area 3 had 3.2% coverage. Most CCFPs overforecast convection, but the long-term bias is known and can be applied as a correction. Finally, note that the predicted cloud tops in areas 2 and 3 were greater than observed. In area 1, not enough convection was observed to assess the cloud-top prediction.

2.3.5 Next-Day Convective Guidance

NCEP's Short-Range Ensemble Forecast (SREF) was first mentioned in section 3.2.2. The SREF initialized at 09Z is the basis for an experimental Extended Convective Forecast Product (ECFP), valid between 1800 and 2400 GMT the next day, a period of maximum air traffic and, coincidentally, also maximum convective activity. Customers include traffic managers at ARTCCs, the ATCSCC, airline and aviation industry dispatch and flight planners, and private weather vendors supporting the airlines and the FAA, who need a day-ahead look at potential air traffic problems. The 22 members comprising the SREF predict convective precipitation and thus make possible the calculation of thunderstorm probabilities. If most members of the ensemble agree that convection will occur in a specific area, the probability of thunderstorms will be rated high. The format for the ECFP is similar to that of the CCFP. Unlike the CCFP, however, the ECFP is a fully automated product. A sample is given in Fig. 2.49.

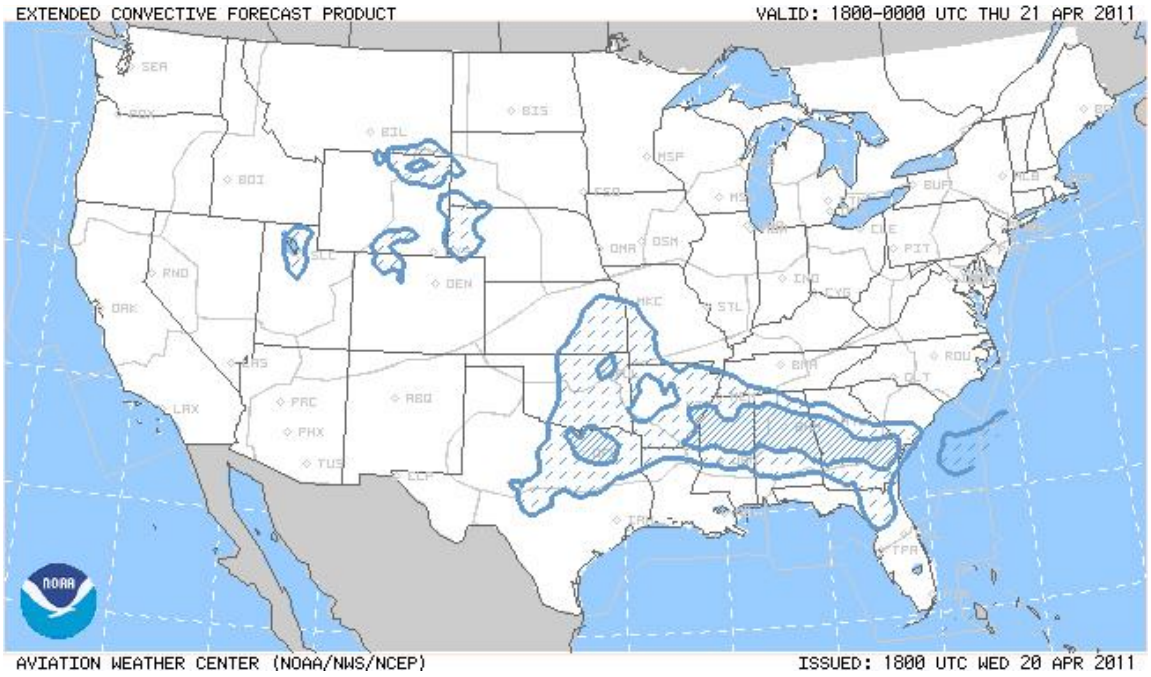

Fig. 2.49. The Extended Convective Forecast Product issued at 1800 GMT 20 April 2011 but valid a day later, from 1800-2400 GMT. Hashed areas represent 40-59% probability, solid lined areas represent 60-79% probability, and solid blue fill (not seen in this example) represents ≥ 80% probability. This product is available at http://aviationweather.gov/testbed/ccfpoutlook/ .

The Storm Prediction Center issues storm guidance for days 1, 2, and 3. Figure 2.50 shows guidance for day 2. The product is available at http://aviationweather.gov/adds/convection/ (click on the image at lower right beneath "Severe Weather Products") or go directly to http://www.spc.noaa.gov/products/outlook/day2otlk.html . This product is useful for longer range aviation planning. A technical discussion accompanies Day-1, Day-2, and Day-3 products.

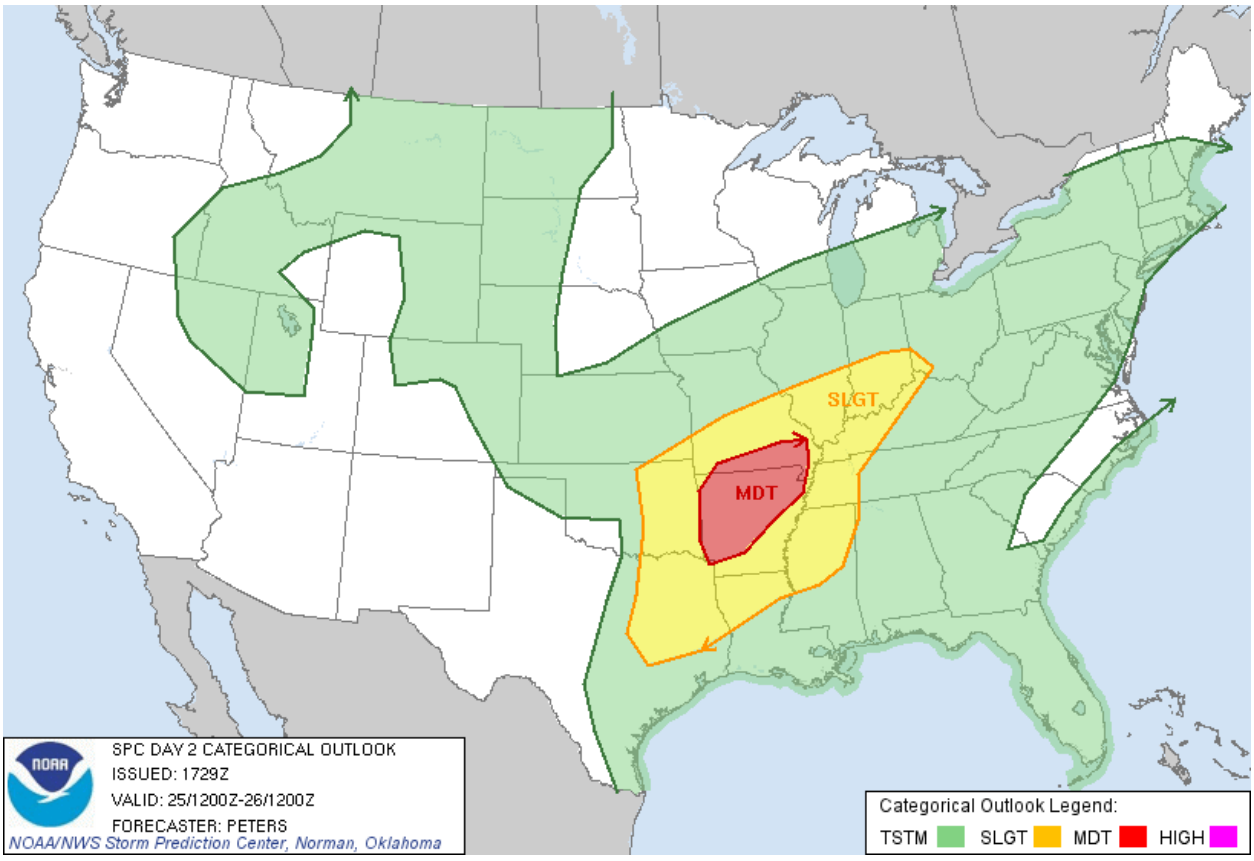

Fig. 2.50. Convective outlook for Day 2, issued by the Storm Prediction Center at 1729 GMT 24 April but valid from 1200 GMT 25 April through 1200 GMT 26 April 2011. Thunderstorms are possible in the area shaded in green. A slight risk of severe weather (yellow shading) means there is a 15-30% chance of severe weather of any type (tornado, hail ≥ 1" diameter, or wind gusts ≥ 50 kt) with 25 miles of any point within the contour. A moderate risk (red shading) implies a 45% probability.

The last two sections described SIGMETs and CCFPs, primary aviation weather products generated by senior meteorologists at the Aviation Weather Center. The next three products are supplemental and generated by computer.

2.3.6 National Convective Weather Detection and Forecast Products

This section covers the National Convective Weather Detection (NCWD) product and the corresponding forecast (NCWF) product. These products were developed within the FAA's Aviation Weather Research Program, in particular, by the Convective Weather Product Development Team, a collaboration among, MIT Lincoln Laboratories (MITLL), NCAR, the National Severe Storms Laboratory (NSSL), and AWC. NCWD and NCWF products are intended for airline dispatchers, general aviation, and FAA Traffic Management Units. As of this writing,

two versions of these products exist: on the ADDS website at: http://aviationweather.gov/adds/convection/ and on the experimental ADDS site at: http://weather.aero/tools/weatherproducts/convection . Support for the latter version will end after the summer of 2011 so it will not be discussed here. Documentation for the former version will be summarized here. Details are at: http://aviationweather.gov/products/ncwf/webtrng/ .

NCWD also goes by the name Convective Weather Hazard. This product depicts current conditions based upon data from WSR-88D radars operated by NWS and the National Lightning Detection Network (NLDN), operated by Vaisala, Inc. In determining the locations of hazardous convective weather, Table 2.9 is used; it gives equivalences between radar reflectivity, VIP levels, VIL, and cloud-to-ground lightning strike frequency.

Table 2.9

Approximate equivalences between WSR-88D reflectivity, VIP level, VIL, and cloud-to-ground lightning (LTGCG) frequency. A VIL of 1 kg m^{-2} is equivalent to 1 mm of suspended water.

Reflectivity (dBZ)	VIP Level/ Color Coding	VIL kg m^{-2}	LTGCG Strikes/10 min
> 50	5-6	> 12	> 15
45-49	4	6.9 – 12.0	6 - 14
40-44	3	3.5 – 6.9	3 – 5
0-39	1-2	0.0 – 3.5	NA

The NCFD product begins when the NEXRAD Information and Data System (NIDS) calculates VIL from reflectivity data acquired by each WSR-88D radar. UNYSYS, a commercial vendor, then processes the VIL data into a national mosaic on a 4-km grid, with colors assigned as in Table 2.9. Step 2 involves removing radar data where the echo tops (another NIDS product) are less than 16,000 ft. This has the desirable effect of removing ground clutter and artifacts due to anomalous propagation. It also removes considerable light and stratiform precipitation.

In Step 3, cloud-to-ground lightning strikes in the past ten minutes are plotted on the national map. For each point on the 4-km radar grid, strikes within an 8-km radius are counted, and the corresponding VIP number assigned if it exceeds the number already in that position. For example, if nine strikes in ten minutes occurred within 8 km of a grid point, the corresponding VIP Level is 4. If the reflectivity at that point was 42 dBZ, it would have been assigned a VIP Level of 3. The value 4 would then replace 3. This step is taken because, in regions where strong storms are numerous, the radar beam may be attenuated while passing through a nearby shaft of heavy precipitation and fail to detect a more distant strong storm along the same radial. The lightning data help solve this problem.
Both the NCWD and NCWF products are updated every 5 min. To see looping of either one at either 5-min or 1-h intervals for 24 frames, go to http://aviationweather.gov/products/ncwf/. Fixed views are available on the national, regional, and ARTCC scales. This option is best for

capturing the evolution and movement of storm systems. Two other viewing options are available, both capable of zooming in on windows of arbitrary size and location, but neither one allows looping: http://aviationweather.gov/adds/convection/java/ and http://aviationweather.gov/adds/convection/java/?appletsize=large . The latter web site gives higher resolution images; it supplied Figs. 2.51, 2.52, and 2.53.

NCWD or NCWF at 5-min frequency for 24 frames

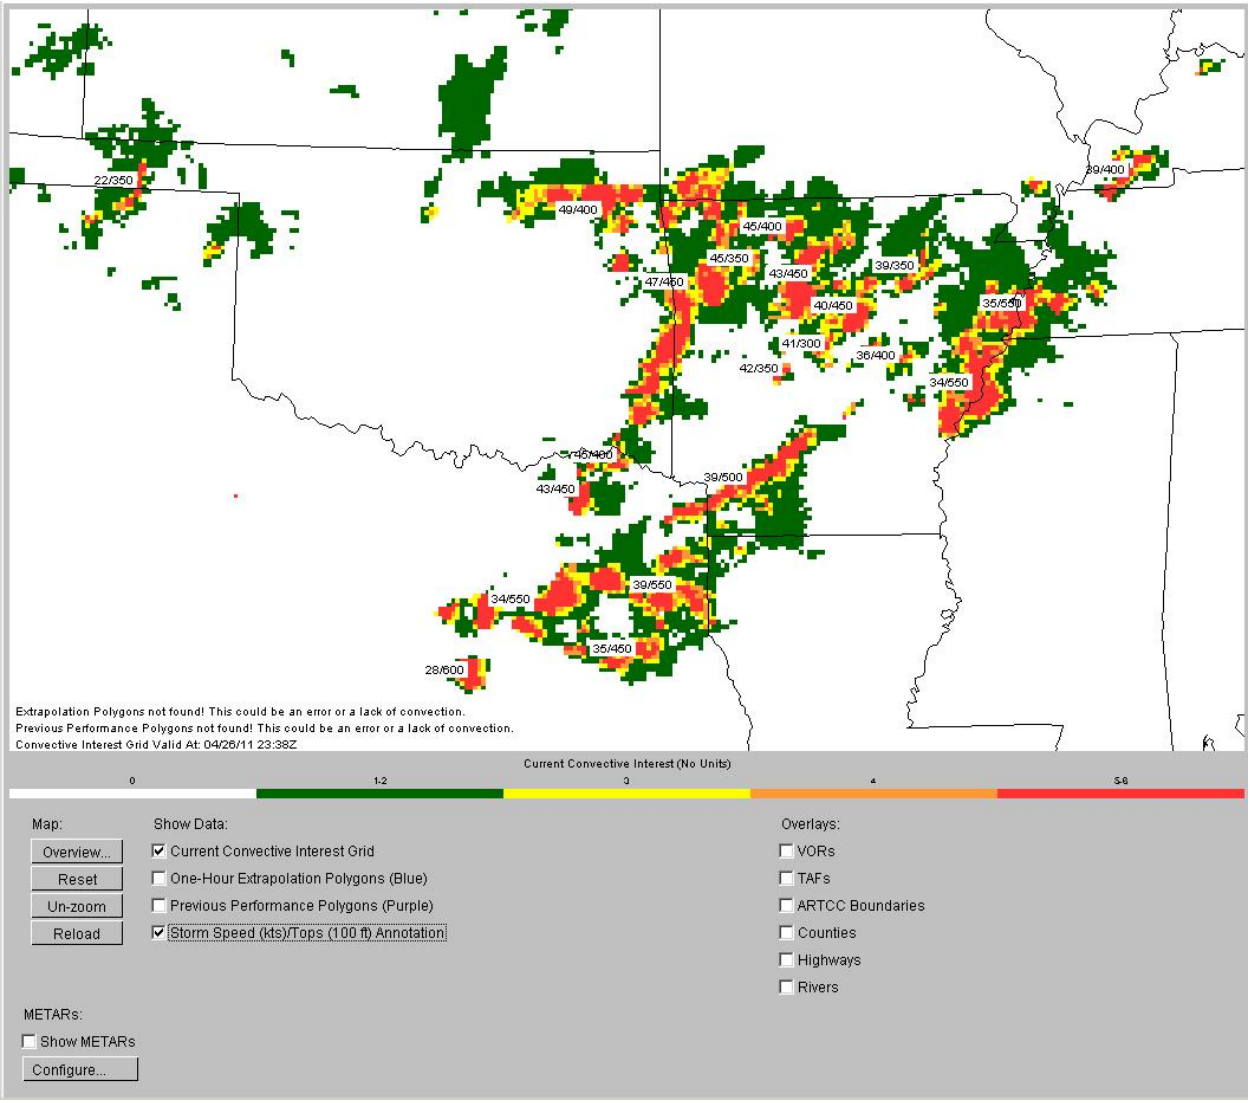

Fig. 2.51. The NCWD product for 2338 GMT on 26 April 2011, showing the current convective hazard. The color bar corresponds to information in Table 2.9. Individual storm cells are annotated with speed of movement in knots and echo tops in hundreds of feet.
Though no overlays are activated in Fig. 2.51, note that many are available, including a plot of METAR observations at various densities.

The VIL field (without the lightning correction) is the starting point for generation of a 1-h extrapolated forecast. Step 1 is the partitioning of the 4-km VIL field into stratiform and convective components by a procedure due to Steiner et al. (1995). If the gridded VIL is at least 3.5 kg m^{-2}, it is automatically designated convective. If VIL < 3.5 kg m^{-2}, but the value exceeds the average of surrounding grid values (within an 11-km radius) by a certain minimum amount, then it, too, is designated convective. Otherwise, grid points with non-zero VIL are designated stratiform. A narrow fringe area surrounding each designated convective grid point may also be designated convective. This area may extend outward up to 5 km if the VIL value exceeds 12 kg m^{-2}, but the extension decreases to 1 km as VIL decreases to 0.15 kg m^{-2}. Step 1 concludes with the removal of all echoes designated stratiform. Ideally, only active convection should remain, and this is what is to be tracked.

Step 2 uses cloud-to-ground lightning data to account for possible attenuation of the radar beam by nearby strong storms. As for the NCWD product, convective areas are augmented if a 10-min count of lightning strikes within an 8-km radius suggests a VIL value exceeding 3.5 kg m^{-2} where no convection was previously indicated.

Following step 2, small-scale features not likely to persist are eliminated by means of an elliptical filter (step 3) suggested by Wolfson et al. (1999). To each convective VIL value on the 4-km grid, a narrow 5 x 21 filter (see Fig. 2.52) is applied, covering an area of 20 km x 84 km.

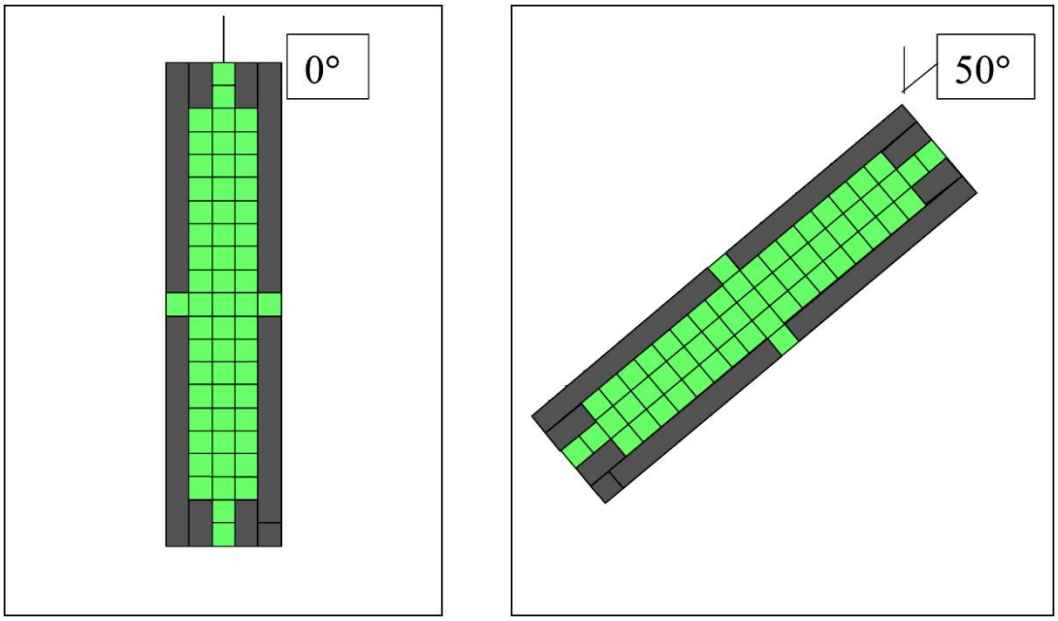

Fig. 2.52. The elliptical filter of Wolfson et al. (1999) applied to convective VIL values. See text for details.

The center of the filter lies over the VIL value. The initial orientation of the filter is vertical as in Fig. 2.52. The average of all VIL values under the green pixels is computed for each orientation, from 0^o through 170^o (18 values). VIL values under the gray pixels are ignored. The largest of the 18 average values replaces the original VIL value under the center of the filter. This filter eliminates small-scale (and, most likely, transient) features in the VIL image and retains the larger scales, facilitating tracking of convective elements.

Step 4 employs a tracking algorithm called TITAN: Thunderstorm Identification, Tracking, Analysis, and Nowcasting (Dixon and Wiener, 1993). Only the most basic aspects of TITAN can be described, and that happens in Appendix B, Section 2.3.6, in order to spare the more casual reader some mathematics.

Once TITAN identifies storms and gives them a speed and direction, Step 5 assigns an echo top to each storm. As noted earlier, the NIDS provides echo top values. The average of the top 5% of values contained within the storm polygon becomes the assigned echo top.

Step 6 smooths the storm polygon defined by the end of Step 3, and moves the smoothed outline of the storm in fixed orientation to the location specified by the storm motion vector. See Fig. 2.53.

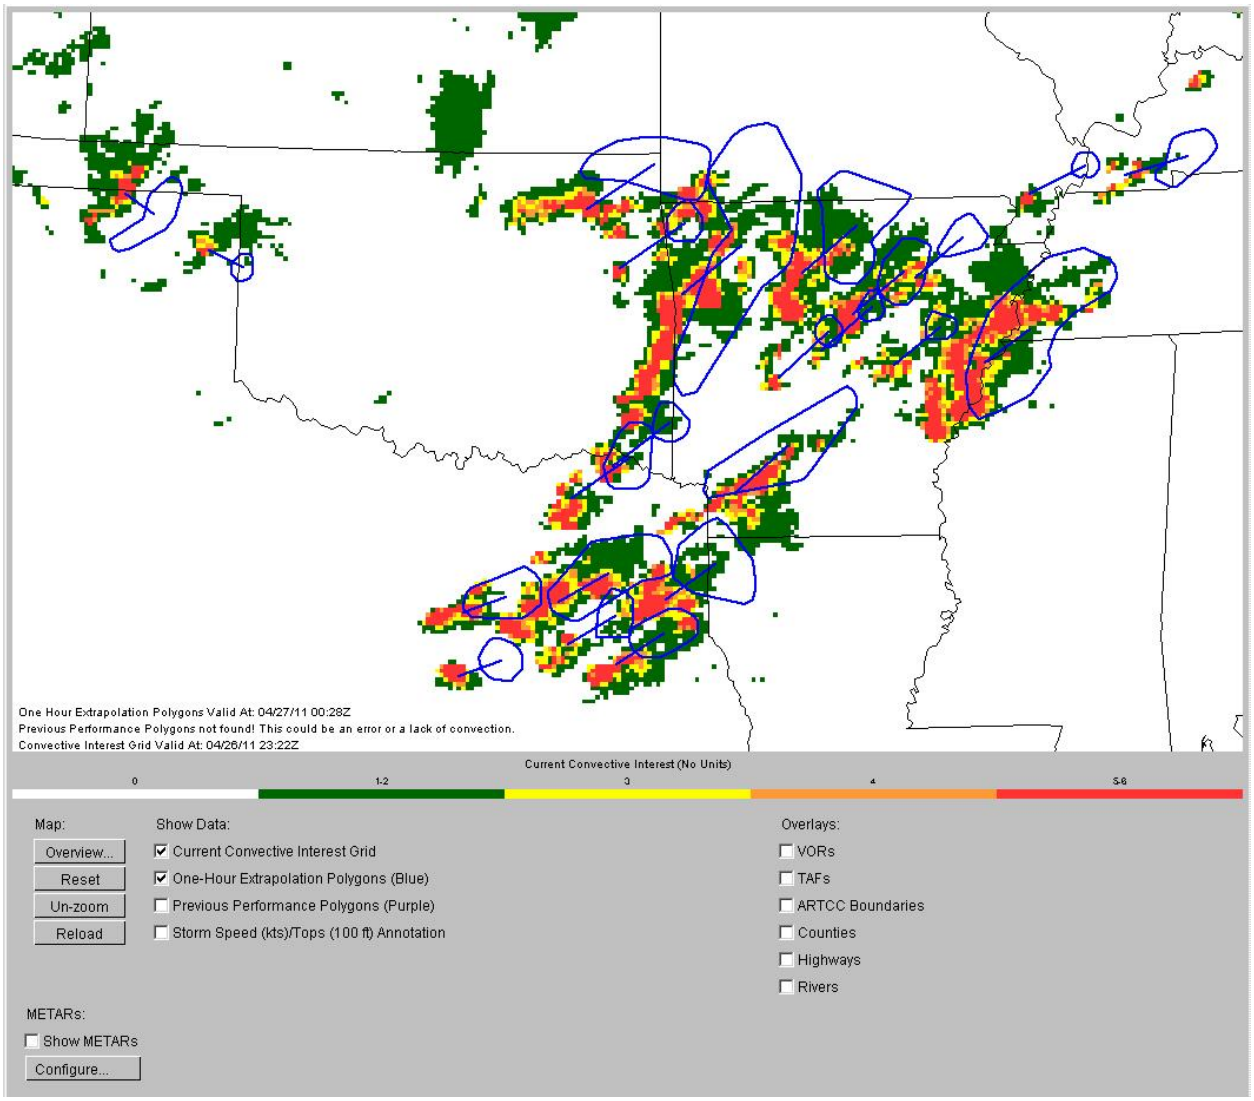

Fig. 2.53. The NCWF product, showing the storms identified at 2322 GMT, 4 April 2011, and a one-hour extrapolation of their positions outlined in blue. The straight blue lines are one-hour motion vectors for each storm. The storm tops are not annotated so as to show all the details of the VIL image. Table 2.9 gives the correspondence between the VIL level and the map colors.

TITAN cannot forecast storm generation or dissipation; it is strictly an extrapolation tool, which performs best with long-lived, dynamically forced convection but not so well with poorly organized, weakly forced, and short-lived convection.

Verification of the NCWF product is part of the product display. One can click on “Previous Performance Polygons” to see the results of the most recent extrapolations as in Fig. 2.54.

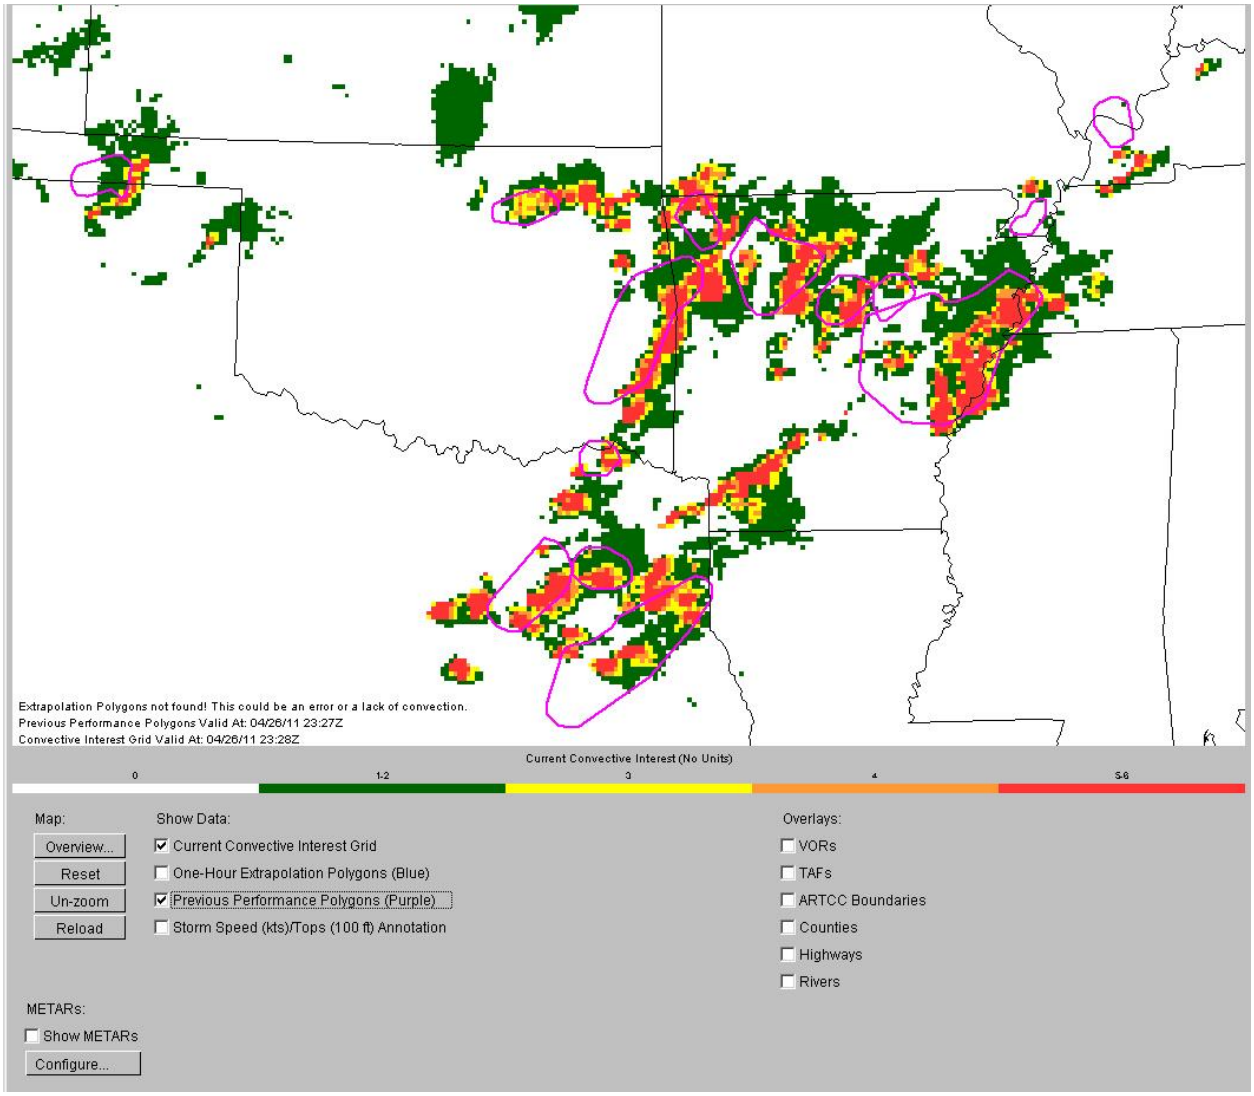

Fig. 2.54. Verification of the 1-h extrapolation of storms (pink contours) valid at 2327 GMT 26 April 2011. The color-coded echoes marking the real storms were observed one minute later.

Note that the extrapolation of some storms is quite successful, whereas others, especially in northern Arkansas, move more quickly than anticipated and are already starting to pass outside their eastern extrapolated boundaries. The storm in southwest Arkansas, also visible in Fig. 2.53 at 2322 GMT, is apparently less than one hour old because its forecast track is not yet available.

Table 2.10 examines the performance of the NCWF for April 2011, a record month for severe thunderstorm and tornado activity in the U.S. If one compares the area of the country covered by convective storms with the area free of storms, it is easy to see why the NN count is in the hundreds of millions and the YY count in the tens of thousands. A storm at a given location is the exception rather than the rule, and thus a good score on PODn (very close to one) is not

hard to achieve. PODy scores in April (0.41 – 0.47) were higher than normal because much of the convection was strongly forced by an active jet stream with embedded short waves. Many storms were long-lived with rotating updrafts. In contrast, the weekly PODy scores for August 2010, a late summer month when convection is forced more by diurnal heating than by jet stream dynamics, ranged from 0.26 – 0.33. The false alarm rate for April, in the 0.64 - 0.68 range is fairly high, almost surely because the speed and shape of storms is hard to estimate by extrapolation one hour in advance, but it was even higher in the preceding August (0.73 - 0.75). The CSI scores for April 2011 (0.22 – 0.25) are notably higher than those for August 2010 (0.14 - 0.17).

Table 2.10

Skill scores by week for the 1-h NCWF based on pure extrapolation for April 2011. Note that numbers in the contingency table (YY, NN, NY, NN) are in thousands, based on a national 4-km by 4-km grid.

Week Beginning	**YY** (1000s)	**YN** (1000s)	**NY** (1000s)	**NN** (1000s)	**PODy**	**PODn**	**FAR**	**CSI**	**% Area**
01 Apr 2011	51	93	58	118,574	0.468	0.999	0.646	0.253	0.12
08 Apr 2011	51	108	71	114,969	0.414	0.999	0.680	0.220	0.13
15 Apr 2011	93	194	105	118,384	0.471	0.998	0.675	0.238	0.24
22 Apr 2011	187	361	218	119,441	0.462	0.997	0.658	0.245	0.45

The percent of the predicted storm polygons actually filled by storm echoes is given in the last column of Table 2.10; it varies widely by week. The final week of April featured numerous, devastating severe storms in the southeast U.S., possibly accounting for the large percent of area covered. Convective storms typically move more slowly in August than in April and hence tend to stay inside the 1-h predicted polygons more reliably in August. The percent area average for the first four weeks of August 2010 was 0.42

2.3.7 Corridor Integrated Weather System

The Corridor Integrated Weather System (CIWS) draws upon research conducted over the past two decades, mostly at MITLL. The development of CIWS has been sponsored by the FAA's Aviation Weather Research Program and carried out by the Convective Weather Product Development Team, a collaboration among MITLL, NCAR, NSSL, and AWC.

CIWS is a fully automated analysis and forecast system. It provides high-resolution images of precipitation (in the form of VIL) with emphasis on convection, echo tops, lightning, and satellite cloud fields within the CONUS and southeast Canada around the Great Lakes. The product updates every five minutes. It also produces storm motion vectors, estimates storm growth and decay, and extrapolates echo motion, storm intensity, and storm top data every five minutes out to two hours. One can loop over past and forecast images from -120 min to

+120 min at five-minute intervals. Products are archived, and the website supports playback of past cases.

CIWS is designed to provide rapidly updated weather information to air traffic flow managers. It is used by the ATCSCC, ARTCCs, TMUs, CWSUs and airline dispatchers for routing and metering air traffic. Though the system is available to air traffic managers, it is not available to the public as of this writing. Evans and Ducot (2006) give a high-level description of CIWS.

Figure 2.55 is a sample CIWS "echo tops" image on the national scale. The image may be looped through past data at five-minute increments. Other products for display are listed across the bottom of the image; they will be discussed later. The cloud-top forecast (valid 120 min later than the image in Fig. 2.55) is given in Fig. 2.56. Note the use of different colors to portray predicted echo-top heights.

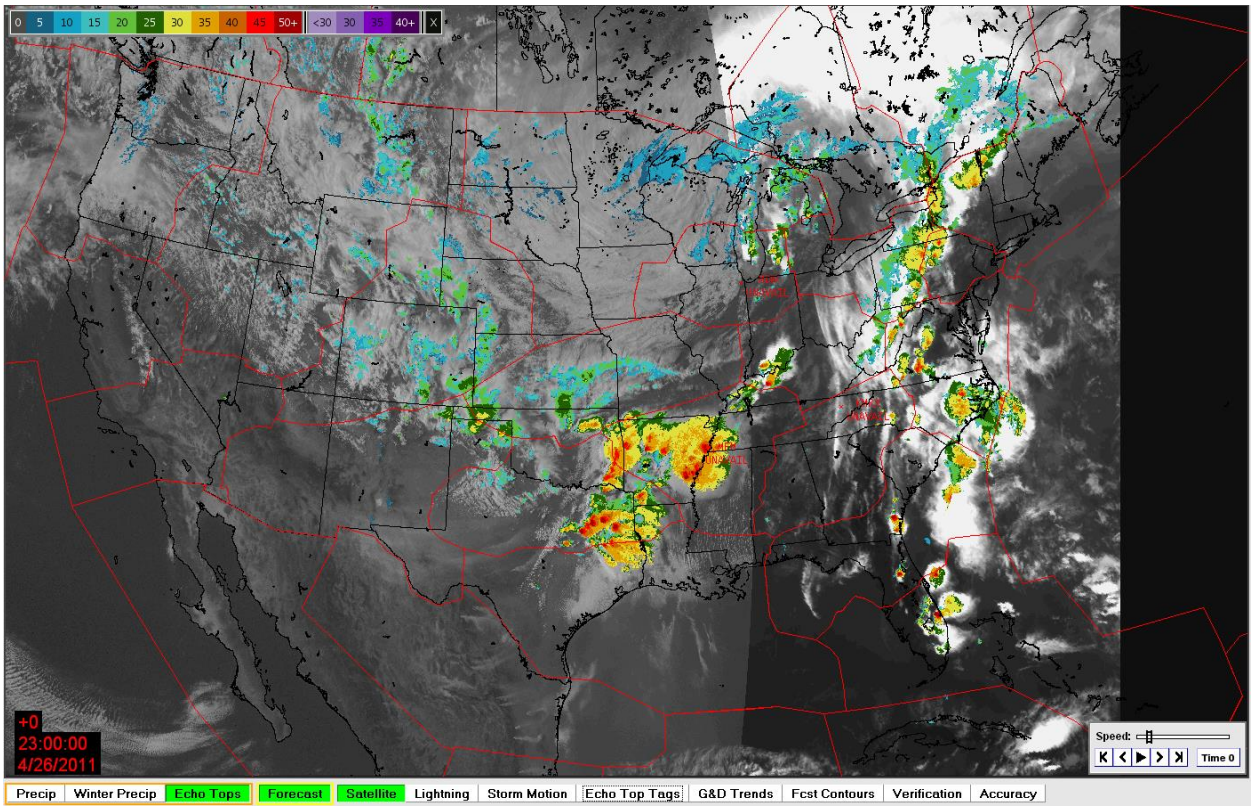

Fig. 2.55. CIWS echo top image, current at 2300 GMT, 26 April 2011. Echo tops are color-coded in thousands of feet as in the color bar at top left. As an option, echo tops are superimposed on a GOES satellite image for the same time. Sunset is occurring along the eastern borders of Mississippi, Illinois, and Wisconsin. Note the line of transition between visible and infrared imagery along the sunset boundary. ARTCC boundaries are in red.

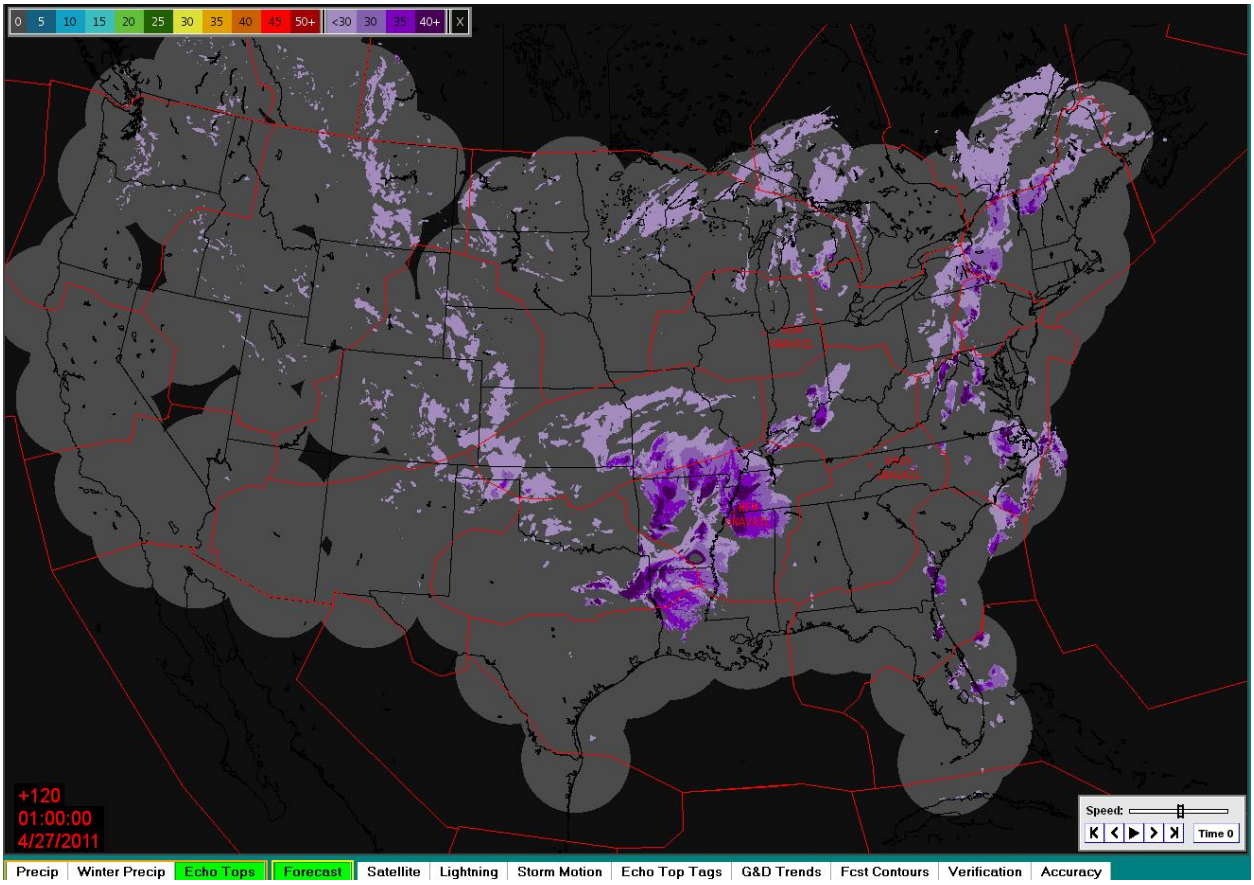

Fig. 2.56. 120-min cloud-top forecast proceeding from the conditions shown in Fig. 2.55. Shades of purple, given on the right side of the color bar at upper left, show predicted echo-top heights.

This brief introduction to CIWS leads to a discussion of the scientific underpinnings of product generation. Details follow.

2.3.7.1 Classification of image features

VIL and echo top data from each volume scan of a single radar provide the input data at 1-km spatial resolution for classification of image features. Dupree et al. (2002) performed the initial classification work, and Wolfson et al. (2004) broadened it. The latter reference and Dupree et al. (2005) provided much of the information for this subsection.

Figure 2.57 diagrams the classification procedure, which includes four major steps: 1) distinguishing convective from nonconvective precipitation, 2) sorting precipitation regions and individual cells by size, 3) assigning sub-types by growth and decay trends (this requires data from successive volume scans), and 4) final classification.

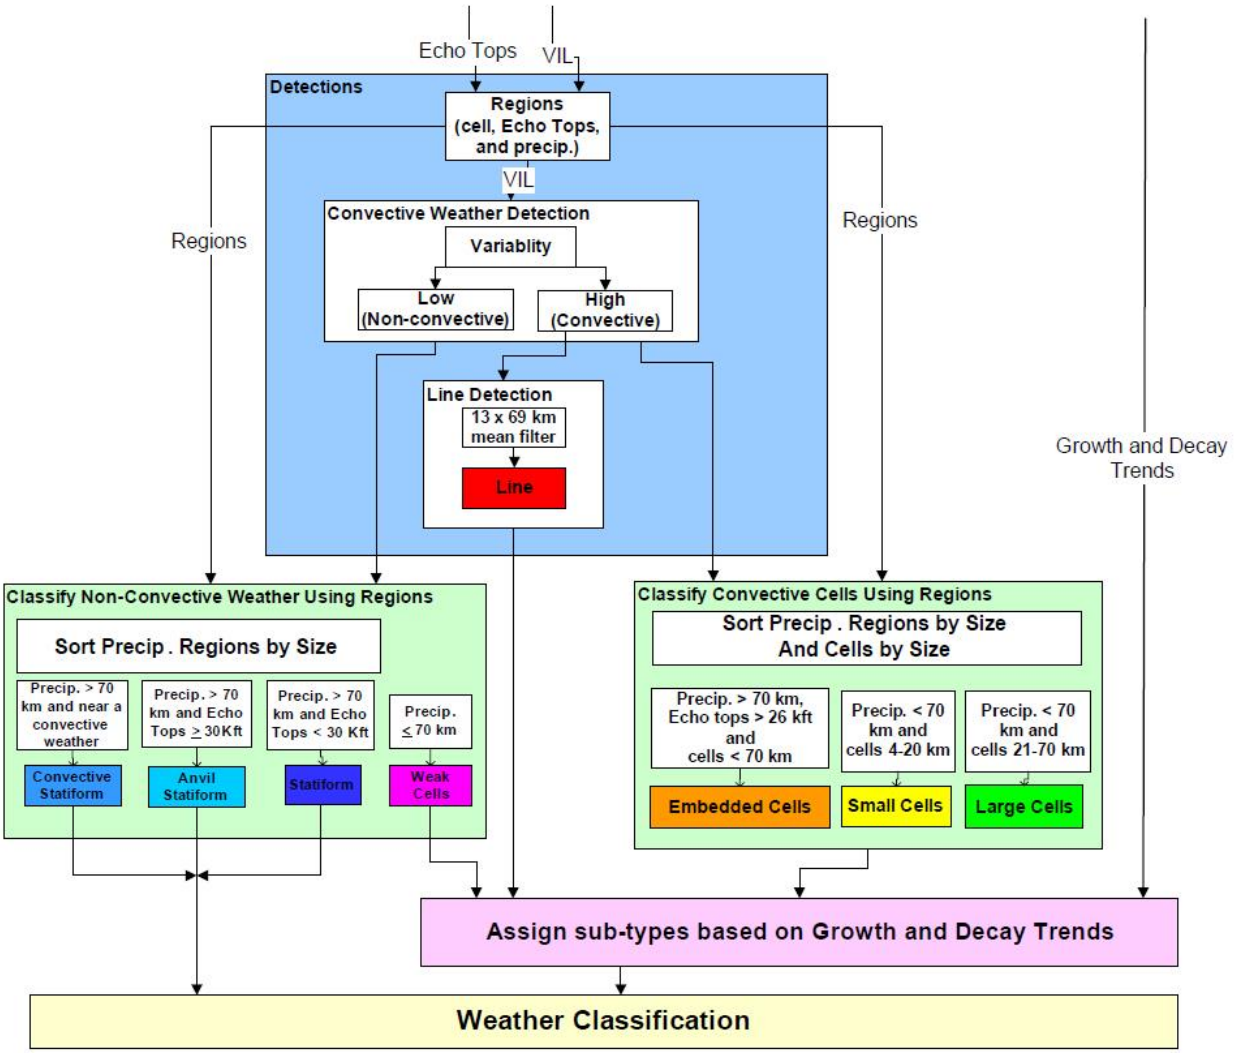

Fig. 2.57. Procedure for classifying features in VIL and echo top images. See details in the text. From Wolfson et al. (2004).

Classifying features involves thresholding the input images at various levels and then applying image processing techniques that characterize features by shape, size, and statistical variability. Two masks are applied to the VIL image. Masking VIL values less than Level 2 highlights areas with heavier precipitation; masking VIL values greater than a certain threshold highlights

sometimes large areas of lighter precipitation. An echo-tops mask helps to identify areas of deep convection.

A spatial variability map is created by calculating the local standard deviation of VIL values at each point. Large variability implies convective precipitation; low variability implies "non-convective" precipitation. An elongated, 69-km by 13-km, elliptical filter, rotated every 5°, is applied to the original, thresholded VIL image. The smoothing accomplished by this filter retains large-scale features, and the elliptical shape highlights linear features, namely convective lines. Wherever linear features coincide with high variability in VIL, convective lines are indicated.

On the non-convective side of Fig. 2.57, large precipitation echoes (> 70 km across) are classified as "convective-stratiform" if they are near convective weather (e.g., stratiform rain occurring behind a squall line), "anvil stratiform" if echo tops exceed 30,000 ft (e.g., a dissipating air mass thunderstorm), and merely "stratiform" if echo tops lie below 30,000 ft. If the precipitating region is less than 70 km across, the echo is designated "weak cell."

On the convective side of Fig. 2.57, precipitating regions (except for linearly organized storms) are sorted by size. If the precipitating region is more than 70 km across, any smaller cells within this region that have tops over 26,000 ft are designated "embedded tops." For precipitating regions less than 70 km across, cells 4-20 km in size are called "small" and those 21-70 km in size are called "large."

The final classification step is to decide whether the convective entities are growing, decaying, or staying the same. "Weak cells" on the non-convective side of Fig. 2.57 are included in this step to allow for the possibility that they are young but growing convective cells. To assign "growth" and "decay" designators, it is necessary to consult a brief history of VIL and echo tops and identify individual features in successive images. For that, feature tracking is necessary.

2.3.7.2 Feature tracking

Two different filters are applied to the original VIL image at 1-km resolution: the 69-km x 13-km rotated elliptical filter described in the preceding section, and a 13-km circular filter. Both filters have a smoothing effect, but the first of these preserves larger-scale features while the second filters only the smallest scales. The reason for separating scales in this way is that lines or clusters of convective cells often move in a direction different from the individual cells

contained within them. The goal is to obtain track vectors not only for the individual cells (preserved in the image resulting from the 13-km circular filter) but also for envelopes of cells (highlighted by the much larger elliptical filter).

For each filter type, two successive filtered images are subjected to a cross-correlation tracker. The two images are divided into sub-areas. The corresponding sub-areas in each image are then superposed. The correlation between pixel values of the sub-areas are computed as the later image is shifted up and down, right and left, over the earlier image. The displacement of the later image with respect to the earlier image that gives the highest correlation yields the track vector. There is more to this procedure than described here, including *hierarchical analysis* based upon decreasing/increasing image resolution. See Chornoboy et al. (1994) for details. *Cell* vectors are used for the short-term trends (growth or decay) in VIL and echo tops. *Envelope* vectors capture the larger-scale motion of convective lines and storm clusters.

The earlier image is shifted according to the track vector and then superposed upon the later image. The difference between the two images, so positioned, gives the trend.

The short-term trend image is obtained from two consecutive images (about 6 min apart), differenced utilizing the cell vectors. Because difference images are quite noisy, several consecutive difference images are averaged together. This minimizes the noise but preserves the persistent trends.

The longer-term trend image is generated by differencing two images that are several radar scans apart, say, 18 min. Envelope vectors determined after each radar scan are used to advect the earlier image in steps to its final position atop the later image, whereupon the difference (trend) image is generated.

Features classified as small cells (not more than 20 km across in Fig. 2.57) are moved with cell vectors. All other features, convective or not, are moved with envelope vectors.

Two other feature detectors (filter types) are related to growth and decay. One is a Boundary Growth Feature Detector, which identifies thin bands of moderate to strong growth in an otherwise echo-free area and gives line orientation. The other identifies small, growing storm cells well removed from other convection.

The foregoing paragraphs covered determination of growth and decay for VIL features. The growth and decay of echo tops are determined similarly, but the procedure requires echo-top

images and cell vectors (from images derived with the 13-km circular filter). Only short-term trends are computed for echo tops.

Table 2.11 shows the feature types classified in CIWS. The "major types" in the first column correspond to those in the two light green boxes (non-convective and convective) in Fig. 2.57. The sub-types merely indicate whether the feature is stable, growing, or decaying with time. Abbreviations for each sub-type appear in the third column.

In most cases, pixels receive only one classification, but it may happen that some pixels receive two. For example, lines may contain embedded cells, so that a pixel may receive both LINE and EB classifications. The fourth column of Table 2.11 shows which classification takes precedence. In the example, the LINE classification would prevail.

Table 2.11

Weather types resulting from the classification of image features. From Wolfson et al. (2004)

Major Type	Sub-Type	Abbreviation	Precedence Order
Line	Line Boundary	LINB	1
	Line Growing	GLIN	2
	Line Decaying	DLIN	3
	Line	LINE	4
Large Cell	Large Cell Boundary		5
	Large Cell Growing	GLC	6
	Large Cell Decaying	DLC	7
	Large Cell	LC	8
Small Cell	Small Cell Boundary		9
	Small Cell Growing	GSC	10
	Small Cell Decaying	DSC	11
	Small Cell	SC	12
Embedded Cell	Embedded Cell Growing	GEB	13
	Embedded Cell Decaying	DEB	14
	Embedded Cell	EB	15
Weak Cell	Weak Cell Growing	GWC	16
	Weak Cell Decaying	DWC	17
	Weak Cell	WC	18
Stratiform	Stratiform Anvil	SA	19
	Stratiform Convective	SC	20
	Stratiform	S	21
No type	No type		22

Figure 2.58 shows a sample weather classification image for Florida. Many of the classification types in Table 2.11 are represented.

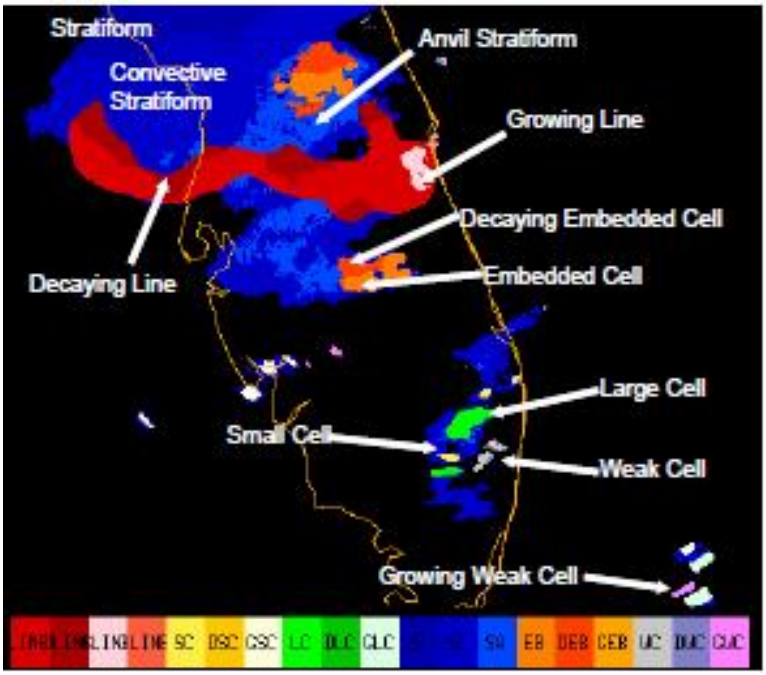

Fig. 2.58. A sample weather classification image over Florida. From Wolfson et al. (2004).

2.3.7.3 Multiscale tracking for a two-hour forecast

The feature tracking discussed in the previous section was adequate for deciding whether convective features were growing, decaying, or maintaining status quo, but more sophisticated tracking procedures are needed to make the 2-h forecast of VIL and echo tops, which is a major component of CIWS. The problem to be solved anew every 5 min is to use the envelope and cell vectors and the trend information to produce new VIL and echo-tops fields every 15 min out to 2 h.

Before track vectors from individual radars are combined into a national mosaic, they undergo quality control and temporal smoothing. Local mean vectors are calculated within sub-regions. Any vector whose direction differs from that of the local mean vector by more than 45° is rejected. A weighted running mean is calculated for both cell and envelope vectors. For cell vectors, the most recent set receives 90% weight and the previous average receives 10% weight. This allows for quick changes in track vectors at the smallest scales. For envelope

vectors, the most recent set receives 30%, and the previous average receives 70% weight. Thus envelope vectors are constrained to change more slowly.

The calculations discussed so far were for individual radars out to a range of 230 km. At this point, the results for all radars are merged into a single national mosaic and synchronized in time. This involves a transformation from the local radar coordinates (range and azimuth), centered at the latitude and longitude of each radar, to the horizontal grid used for the national display. The 230-km range rings for WSR 88-D radars overlap in many parts of the U.S., especially in the east, where radar coverage is fairly dense. Three to six values, each from a different radar, may be present at a given pixel location. The maximum plausible VIL value at a given pixel location is chosen for the mosaic, and the VIL trend value from the radar that had the max VIL is accepted as well. The same procedure is used for echo tops. The maximum echo top reported by any radar at a given location, and the echo-top trend corresponding to this maximum value are inserted into the mosaic.

Two national mosaics of track vectors, one of envelope vectors and the other of cell vectors, undergo further processing. The statistical variance of vectors is calculated under a 57-km diameter kernel. The vectors included in this calculation are ordered from top to bottom according to the size of their contribution to the vector variance, and the top 15% are rejected.

As noted earlier, only “small cells” (a weather classification type) advect with cell vectors. All other types advect with envelope vectors. At this point, all weather classification types have exactly one associated track vector, except the 15% just rejected. Every pixel is assigned a track vector. Within each specific weather classification area, all pixels share the same track vector. But there are still large voids on the map with no track vectors. A pixel with no track vector receives a track vector that is a weighted average of all track vectors at pixels not more than 50 km distant. The weight is 1/r, where r is the distance between the central pixel and any pixel contributing to the average. This fills voids where track vectors have been rejected and also within 50 km of precipitation. Remaining voids are filled by means of bilinear interpolation from the nearest points having track vectors.

2.3.7.4 Two-hour VIL and echo-top forecasts

The full field of track vectors tells the CIWS prediction algorithm how to move information from each pixel downstream over the next two hours, in 15-min increments. The information to be moved includes the weather classification, the initial VIL and cloud-top values, the two VIL

"interest images" regarding linear convection and isolated growing cells, and the growth and decay trends for VIL and cloud tops. Sometimes more than one track vector will point to the same pixel. If that happens, the last vector that pointed to the pixel is the one used. The information from the origin of each track vector is transferred to the destination pixel. The next step is to determine new VIL values at each forecast time.

The new VIL value at each pixel depends upon the values imported to that location: VIL, VIL trend, "interest" from linear convection and "interest" from isolated growing cells. At the majority of pixel locations, all these values will be zero. They are non-zero only where precipitation is predicted. The interest maps for linear convection and isolated growing cells contain mostly zeroes. The new VIL value is a linear combination of these values for which the weights

- are normalized to unity, that is, they add up to one;
- depend upon the forecast time (15, 30, 45, 60, 75, 91, 105, or 120 min into the future);
- depend upon the weather classification type; and
- are determined empirically from historical observations of the evolution of different weather types and the roles interest images played in that evolution.

Finally, a median filter is applied to the VIL forecast at each time. In a median filter, the values within the filter template are ordered from top to bottom, and the middle value in the ordered list replaces the value at the center of the template.

An accurate echo top forecast is important so that controllers can decide whether to route flights over the top or divert them around the side of a storm. In CIWS, the RUC model supplies an estimate of the maximum possible storm top. From the predicted temperature and moisture profiles, the model computes the convective available potential energy (CAPE), which, in turn, allows computation of an updraft profile within potential storms, assuming no lateral entrainment. The altitude at which the updraft velocity falls to zero is the maximum possible cloud top. CIWS does not allow any cloud top to exceed this limit.

The echo tops move according to the track vectors during each 15-min forecast interval. Once their new positions are determined, the appropriate amount of growth or decay from the trend data is applied, but only to convective elements, i.e., weather types involving lines, and large, small, and embedded cells, and with the constraint that echo tops cannot exceed the altitude limit just mentioned. In a process called dilatation, the size of growing echoes may be augmented, especially if they have reached maximum altitude. Similarly, convective lines may be extended. To suppress noise in the final echo-top forecast, a 5-km by 5-km median filter is applied to the image.

2.3.7.5 VIL forecast and verification

Figures 2.55 and 2.56 showed a current echo-tops image and 2-h forecast on the national scale. This section illustrates the high-resolution capability of CIWS when the emphasis is on a much smaller region. The focus is on the historic outbreak of tornadic thunderstorms in the southeast U.S. on the afternoon and evening of 27 April 2011.

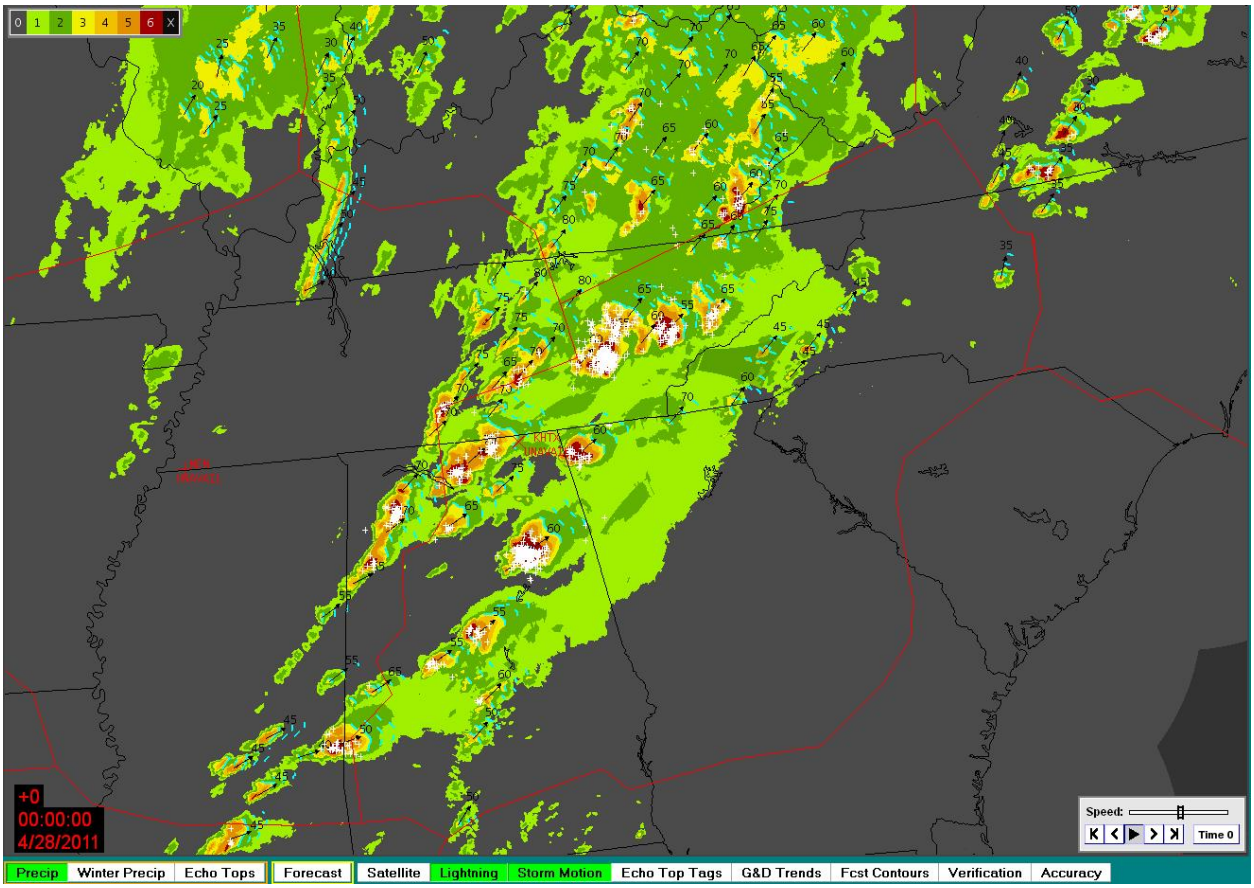

Fig. 2.59. A color-coded VIL image for 0000 GMT 28 April 2011 during the historic tornado outbreak in Alabama and Mississippi. The correspondence between VIL and reflectivity (dBZ) is given in Table 2.9, but levels 5 and 6 (color-coded orange and red) represent "extreme" precipitation and reflectivity values greater than 50 dBZ. The white "plus" symbols locate cloud-to-ground lightning strikes occurring in the past 6 min. Thin black vectors show the direction of storm movement. The black numbers near vector tips indicate storm speed to the nearest 5 knots. Solid cyan lines indicate the leading edge of Level 3 cells. Dashed cyan lines indicate leading edge positions 10 and 20 min in the future. Zoom on image for more detail.

Figure 2.59 is a current VIL image with superposed lightning strike locations (white “plus” symbols), storm motion vectors (black arrows and numbers), and projected positions of the leading edge of Level 3 echoes 10 and 20 min hence. Note lightning so intense that the “plus” symbols almost completely mask the Level 6 storms. Note also the very rapid movement of severe storms toward the northeast, many at speeds from 50-75 knots. One can see large and small cells, some isolated, embedded cells, well-defined convective lines, and stratiform precipitation in this image.

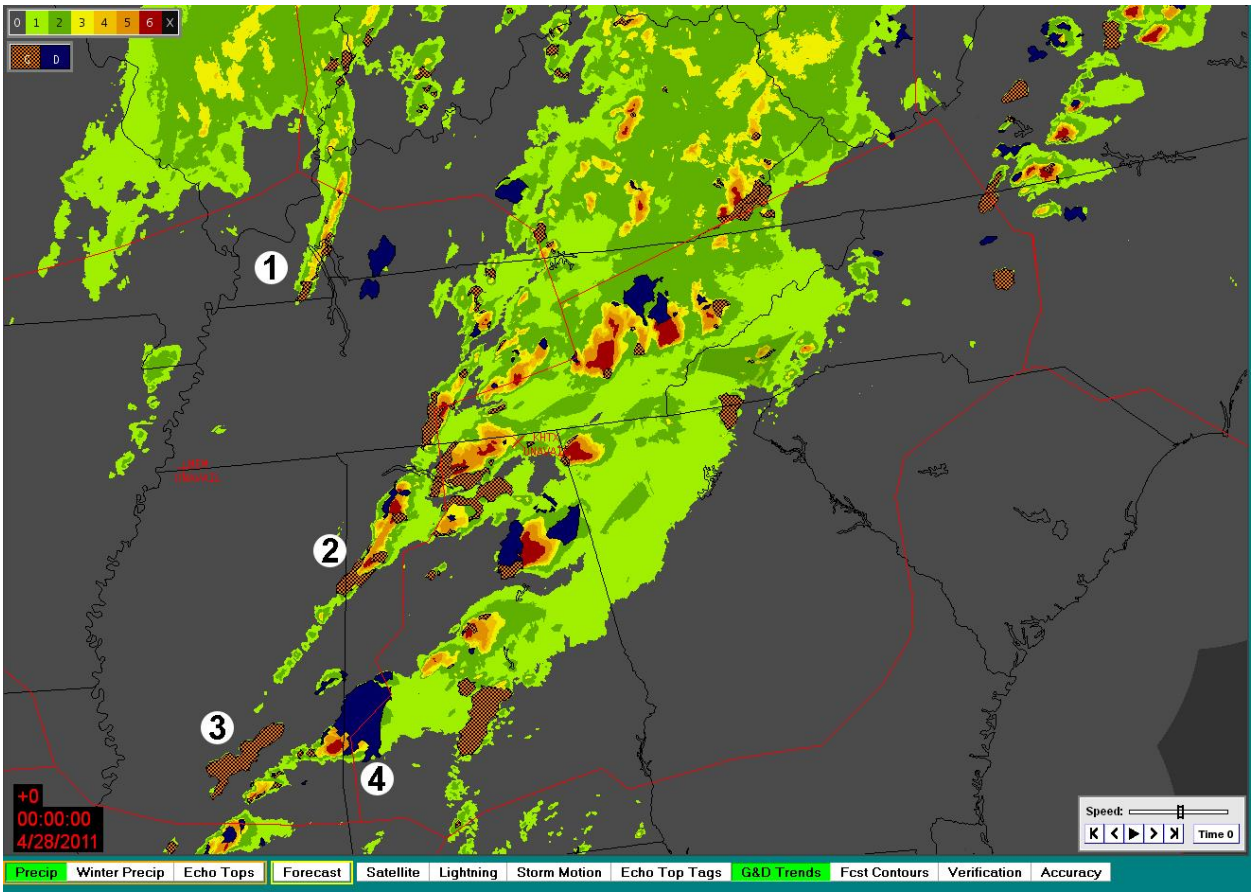

Fig. 2.60. The same VIL image as in Fig. 2.59 but with growth and decay information superposed. Growth regions, where VIL is increasing, are hatched with an orange and black pattern. Decay regions, where VIL is decreasing, are dark blue.

Figure 2.60 shows the same VIL image as Fig. 2.59 but with regions of growth and decay superposed. The numbers inside of circles highlight several features of the growth and decay algorithm. The convective line labeled 1 shows growth near its middle and southern end. Convective line 2 shows growth across the Mississippi-Alabama state line. The Level 4 and 5

cells at 3 (see Fig. 2.59) show pure growth. Lightning activity in the cell at 4 is modest compared with other cells (Fig. 2.59), and its eastern half is decaying. What does the CIWS prediction scheme do with this information? See Fig. 2.61, a 2-h VIL forecast.

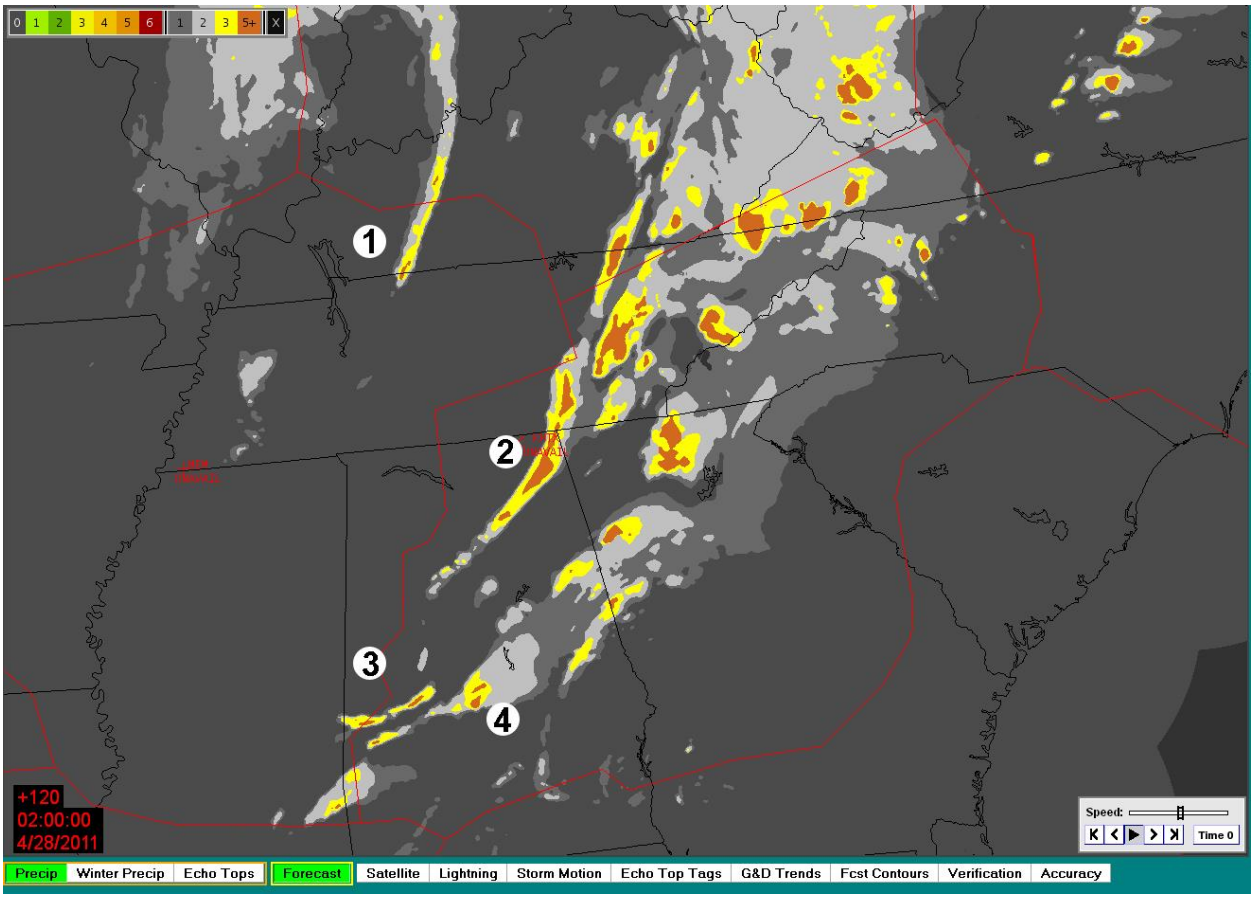

Fig. 2.61. A 2-h VIL forecast valid at 0200 GMT, 28 April 2011, initialized with information contained in Fig. 2.59. Note the altered color scale, upper left, for predicted values of VIL.

Though the line motion at 1 was 40-50 kt to the northeast (Fig. 2.59), the prediction extended the line southward (dilatation) because its southern end was growing at the initial time. The squall line at 2 raced quickly northeastward across northern Alabama, partly entering Tennessee, its southern half largely unchanged (it initially showed neither growth nor decay), but its northern portion retained vigorous convection, reflecting the growth diagnosed at the initial time. The growing cells at 3, initially in south central Mississippi in Fig. 2.60 moved to west central Alabama, containing Level 5+ echoes. The eastern portion of the cell at 4 weakened as it moved from the Mississippi-Alabama state line to central Alabama, consistent

with the decaying trend at the initial time. The forecast maintained the vigor of the western portion.

How good was this forecast? Figure 2.62 gives at least a qualitative answer. Contours in three different colors are drawn where Level 3 and higher echoes are predicted to be at 30 min (dark blue), 60 min (hot pink), and 120 min (white). White contours correspond to the VIL echoes forecast in Fig. 2.61.

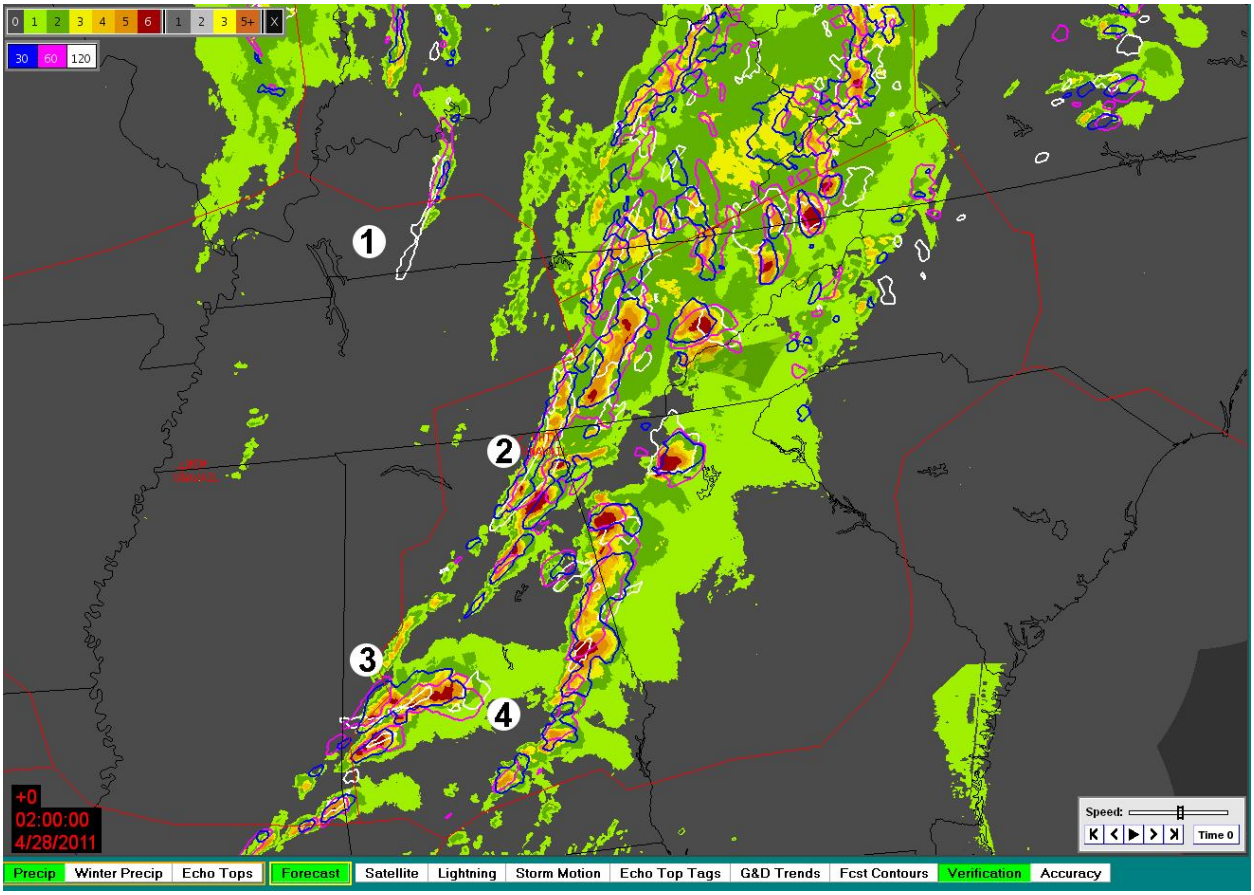

Fig. 2.62. Verification of the 2-h forecast appearing in Fig. 2.61. The observed VIL field at 0200 GMT 28 April is shown. Superposed are color contours enclosing Level 3 or higher echoes predicted 30 (dark blue), 60 (hot pink), and 120 (white) min before this map time. All three forecasts verify at 0200 GMT, but only the white contours correspond to the forecast image in Fig. 2.61.

One notices quickly that the 30-min forecast of cell positions (blue contours) is fairly good for both shapes and positions, as one might expect for a pure extrapolation with trending. The 2-h forecast is not so good. The southern end of the line at 1 was expected to develop, but it

dissipated. Without magnification, it is hard to see the white contour associated with the squall line at 2, but it overlays the squall line well, just slightly east of where it should be. Likewise the growing cells located south of 3 were correctly predicted, both the intensity and location. The weakening of the eastern portion of the cell north of 4 was well anticipated, but not the weakening of the western portion. Only Level 1 and 2 echoes lie within the white contour immediately northwest of 4. Criticizing shortcomings in a 2-h forecast of convection may seem like nitpicking, but, from the perspective of an air traffic controller or a pilot, errors of 15 minutes or 15 miles count. These examples illustrate well both substantial progress in forecasting convection but also the weaknesses in an approach that relies primarily on extrapolation.

Another way of looking at forecast skill in CIWS is to click the "Accuracy" button at the far right of the "weather" control bar across the bottom of the display, as in Fig. 2.63. In addition, one must specify a location for "Home," using the control bar across the top of the image. The "Home" button is left of center. Tampa (TPA) is chosen as home for this example.

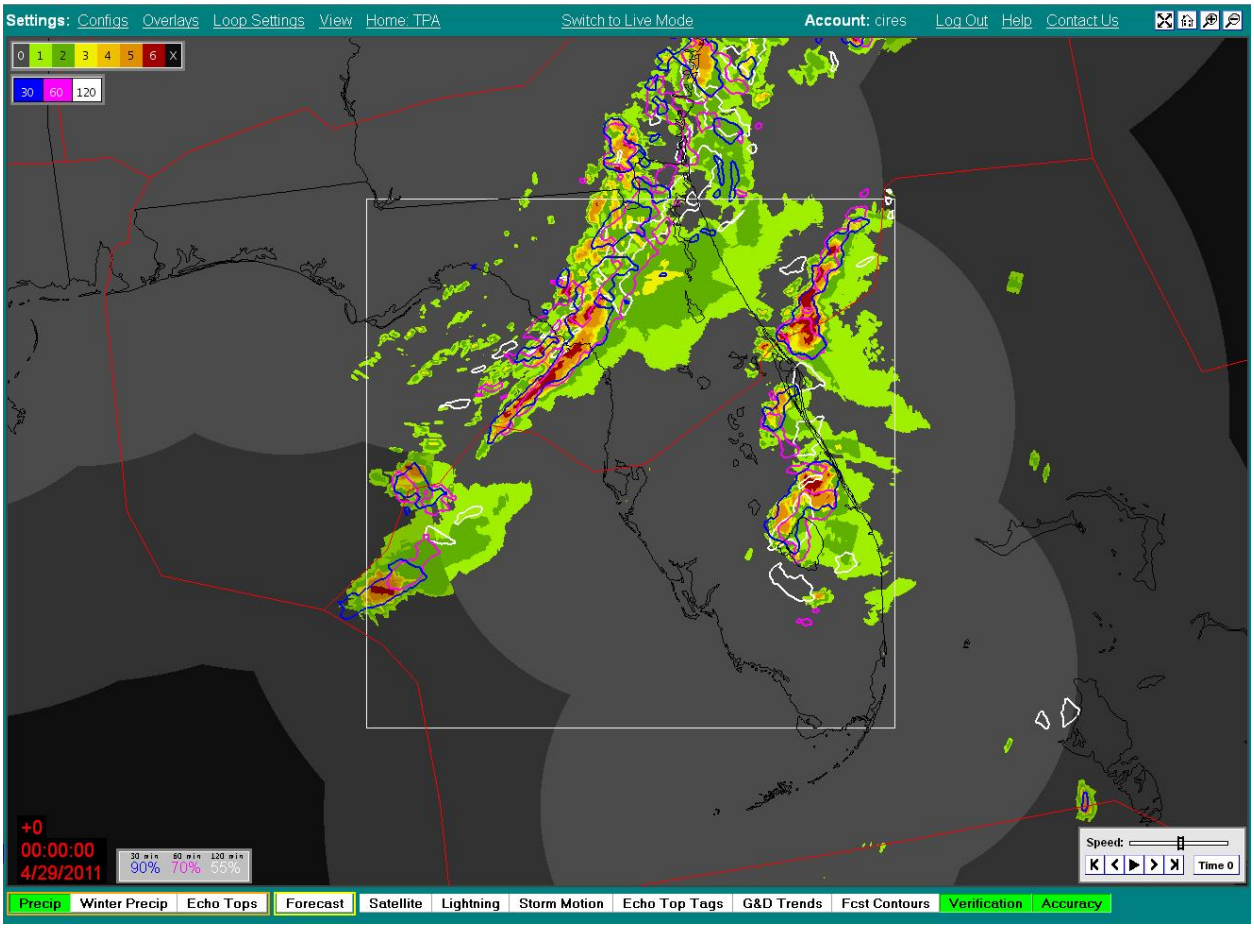

Fig. 2.63. The square, centered on Tampa, Florida, is 300 nautical miles on a side. Verification of the 2-h VIL forecast valid at 0000 GMT 29 April 2011 takes place inside the square. The verification method is described in the text. The scores for 30-, 60-, and 120-min forecasts appear at lower left in blue, pink, and white: 90%, 70%, and 55%, respectively.

The three scores at the lower left of Fig. 2.63 are based upon CSI scores, first defined in Table 1.2.

CSI = Hits / (Hits + False Alarms + Misses)

The score is computed from the observed and forecast values of VIL at each pixel within the square box centered on Tampa in Fig. 2.63. In the simplest case, a hit at a pixel would be defined as an observed value of VIL ≥ 3 coinciding with a predicted value of VIL ≥ 3. A miss is scored if VIL ≥ 3 is observed but not predicted, and a false alarm if VIL ≥ 3 is predicted but not observed.

The CSI could be computed in this straightforward way, but two concessions are made: 1) If VIL ≥ 3 is observed at a pixel and not predicted there, but is still predicted within 10 km of that pixel, a hit is scored. This concession recognizes a desire of users that the forecast be accurate to within 10 km and 10 min. 2) If no predicted VIL ≥ 3 pixels are found within 10 km of an observed VIL ≥ 3 pixel, then the algorithm searches for predicted VIL = 2 pixels in the same area. If at least one is found, a partial hit (0.25) and a partial miss (0.75) are scored. If a VIP ≥ 3 forecast coincides with an observed VIP = 2, a partial hit (0.25) and a partial false alarm are scored. Both these concessions have the effect of increasing the CSI score. The scoring procedure is described in more detail in Theriault et al. (2001), pp. 55-57.

2.3.7.6 Summary and commentary

The CIWS product is valuable for its five-minute updating; its high resolution; the sophistication of the image processing that allows identification of weather types and determination of trends in VIL and cloud-top heights; and its forecast ability, which capitalizes on track vectors at multiple scales and trend information. Yet it is still mostly an extrapolation technique. It can identify growing cells that may later reach thunderstorm status, but it cannot anticipate convective cloud growth where there are no clouds. The technique is skillful in moving severe thunderstorms that are strongly forced by atmospheric instability and the dynamics of the jet stream. Such storms, featured in examples above, are well organized and persist for many hours because of rotation.

Extrapolation performs rather poorly in cases of weakly forced, semi-random, short-lived, and slow-moving convection. This is not to say that such storms are unpredictable but only that the inhomogeneities in the boundary layer that prompt development occur on scales of a kilometer or so that are seldom fully observed. Weckwerth (2000) and Weckwerth et al. (2008) discuss the challenges in forecasting convection in weakly forced situations.

Numerical prediction models at kilometer-scale resolution may eventually show skill in predicting weakly forced convection, but the skill is marginal now. The skill is already substantial in predicting strongly forced convection. Another challenge in using model convective forecasts in the first two hours is a practical matter. It is difficult to collect and process mesoscale observations from all relevant sources, assimilate them effectively into a prediction model at kilometer resolution, and then generate a short-term forecast in time to be useful before two hours have elapsed. These logistical difficulties will probably be overcome before the more difficult problem of predicting weakly forced convection is solved.

2.3.8 CoSPA

2.3.8.1 Introduction

The FAA is supporting development of NextGen, which is shorthand for the Next-Generation National Airspace System. Because convective weather is a major hazard for those who fly and the primary cause of weather-related delays in the form of en route detours or ground stops, a major component of NextGen is the development and refinement of a relatively new product, CoSPA, formerly known as Consolidated Storm Prediction for Aviation, but now referred to only by the acronym. CoSPA marshals scientific and computational resources at MITLL, the NCAR Research Applications Laboratory, and the Global Systems Division of NOAA's ESRL. The FAA expects that CoSPA incorporates the best features of several existing convective products and, in some cases, improves upon them.

CoSPA has been generating products since 2008. They are available on a password-protected website to developers and selected users, primarily those who participate in formal evaluations. A user evaluation ran from 1 June through 31 October 2010. The subsequent evaluation occurred for the same period in 2011, but results are not yet available. The FAA stresses that CoSPA is a research prototype, not an operational system. Despite the limited access and the constantly evolving nature of CoSPA, it deserves space here because it will probably become the primary tool for aviation convective weather forecasts.

CoSPA is a major extension of CIWS in the following ways:

- Forecasts extend beyond two hours to eight hours.
- Beyond two hours, extrapolation forecasts are blended with output from a mesoscale numerical weather prediction (NWP) model.

A flow chart for CoSPA appears in Fig. 2.64. The next four subsections will address changes to CIWS data processing that improve the echo-tops growth and decay algorithm and allow the extrapolation forecasts to extend beyond two hours; introduce the High-Resolution Rapid Refresh (HRRR) mesoscale model, which is increasingly necessary beyond two hours as the skill of computer-generated forecasts begins to exceed that of extrapolation forecasts; and explain how information is blended from both types of forecasts during the "in-between" time (2-6 h) when both offer useful information.

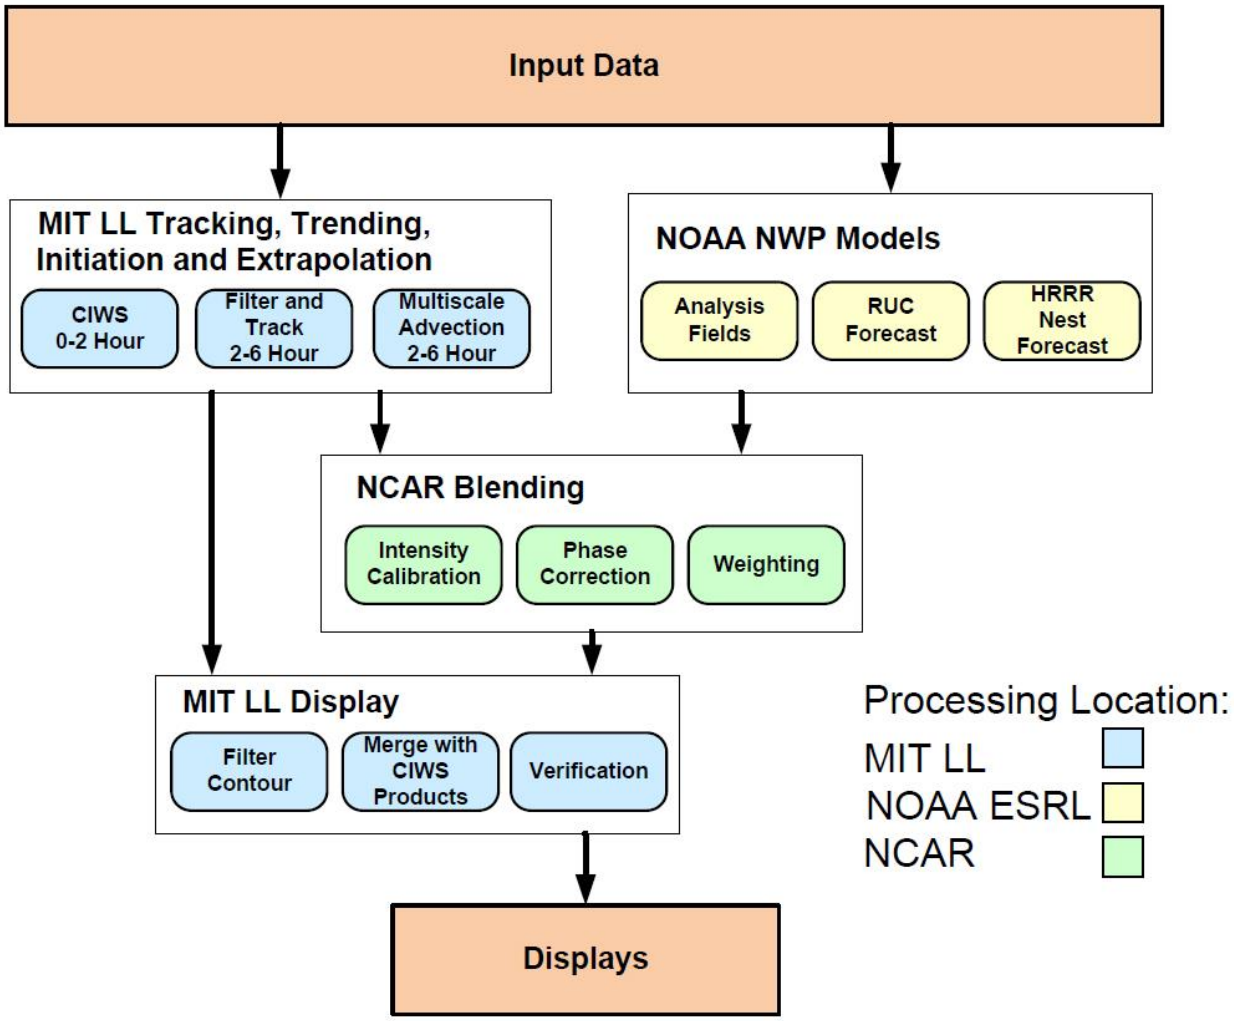

Fig. 2.64. A flow chart for CoSPA. The color coding not only designates the agency responsible for major modules but also the physical location of the computer processing associated with

the modules. From Pinto et al. (2010). Reprinted with permission of the American Meteorological Society.

2.3.8.2 Enhancements to CIWS data processing

The version of CIWS used in CoSPA has two fairly recent enhancements not described in Section 2.3.7. The first is the incorporation of the SATCAST (SATallite Convection AnalySis and Tracking) algorithm on 01 June 2011, which uses data from GOES-13 and GOES-15 to identify convective clouds in early stages of development. SATCAST is described later in Section 2.3.12.1.

The second enhancement blends cloud-to-ground lightning data with echo-tops data to improve the realisim of growth and decay in the echo-top forecast algorithm. The process uses filtered lightning density images, which are created by collecting lightning flashes over the prior 15 minutes and applying a 51-km by 51-km Gaussian filter to the flashes. After three such images have been created, the stack of images is sampled at each pixel along the time axis, and a trend line is determined using linear regression. The trend line for each pixel shows how quickly the values in the lightning density images are increasing or decreasing with time. The echo tops trend is calculated in a similar way, by collecting several Gaussian-filtered echo tops images to estimate how quickly the echo tops are increasing or decreasing in time.

Once both the lightning trend images and echo top images have been created, they are combined using a fuzzy weighted average to create a combined trend image. This image is filtered using a 201-km by 51-km rotated Gaussian filter to create a revised filtered trend image. The new filtered trend image is averaged with the previous filtered trend image from five minutes earlier to create a *time-averaged trend image*.

The process then applies a region labeling function to the echo tops image to identify contiguous regions of echo tops with heights greater than a 20,000-ft threshold. This generates an *echo tops regions image*. A flood-fill technique is then used to expand the regions to and slightly beyond the boundaries of the storms to create a *flood-filled regions image*. The latter is combined with the time-averaged trend image to create a growth and decay predictor, by assigning to each region the average of the time-averaged trend image within the area of that region. The process then applies this growth and decay predictor image to echo-top forecasts by linearly growing or decaying echo tops according to the predicted growth and decay rates. Additionally, regions which are predicted to decay are eroded from the outside in, decreasing their size as well as their height over time.

2.3.8.3 Extending the extrapolation forecast out to six hours.

CoSPA incorporates the CIWS techniques during the first two hours of the forecast, which are partly extrapolative and partly heuristic, in that there are expert-derived rule sets employed to model growth and decay. The CIWS forecast must be extended beyond two hours and blended with a high-resolution model forecast. Because the latter is computationally very demanding, it is not ready until at least two hours have passed. To facilitate the extension of the CIWS forecast well beyond two hours, a third scale, the *synoptic* scale is added to the cell and envelope scales already described for CIWS. As noted in section 2.3.7.2, a 13-km circular mean filter, applied to the 1-km VIL data, defines the cell scale. A 69-km by 13-km rotated elliptical filter, discussed in section 2.3.7.1, defines the envelope scale. The new synoptic scale is created through application of a 201-km by 101-km rotated elliptical filter to the 1-km VIL data. Track vectors are derived through cross-correlation of earlier and later images, separated in time by 6 min for cells, 18 min for envelopes, and 45 min for synoptic features. Figure 2.65 shows images after the three filters have been applied to the full-resolution VIL field.

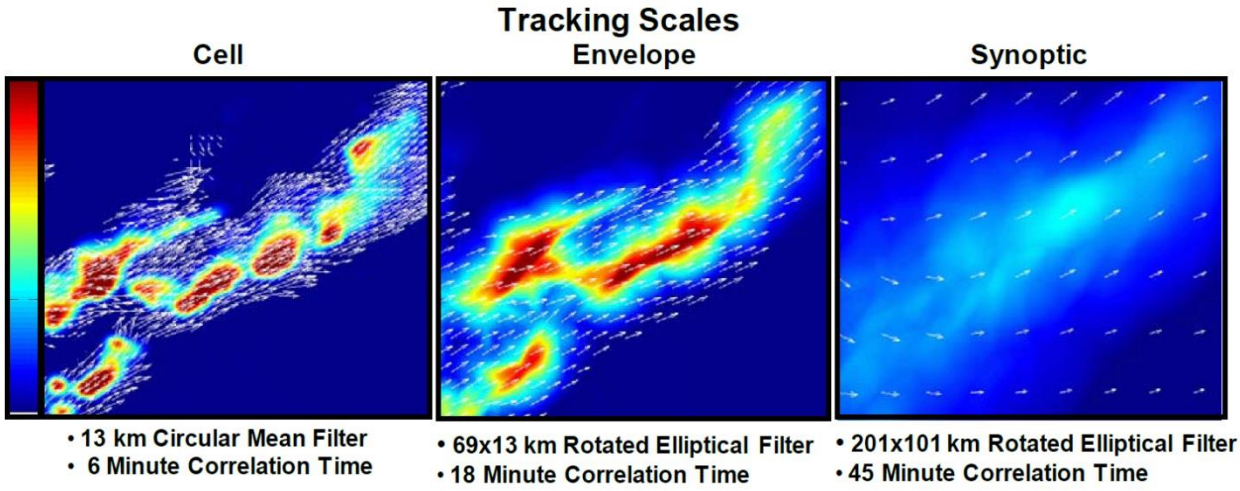

Fig. 2.65. The three scales used to derive track vectors. Track vectors are in white. The background colors correspond to VIL, as in the color bar at left, blue corresponding to the lowest value, dark red to the highest. At each scale, the current image is cross-correlated with previous images to obtain the track vectors.

Following filtering and cross-correlation at each scale, any grid points lacking a track vector receive one through lateral interpolation. Then the synoptic-scale vector field is subtracted

from both the cell and envelope vector fields, thereby removing larger-scale motions from the arrays of cell and envelope vectors. Using a very small time step of a minute or so, each pixel value is advected according to its track vector. Then, using the track vector at the new location, the pixel is again advected to a third position, and so on, out to 10 min. Then envelope vectors advect the pixel information in small steps from 10 to 90 min. Using a different track vector at each time step is equivalent to advecting information along a streamline in the small-scale motion field. Beyond 90 min, information moves according to a blend of the envelope and synoptic-scale vectors, with the latter receiving increasing weight at later times.

2.3.8.4 The High-Resolution Rapid Refresh Model

With the extension of thunderstorm forecasts out to eight hours, CoSPA needed a high-resolution prediction model to supplement its extrapolation forecasts, whose skill diminishes rapidly after two hours. Among the desired characteristics of this model were 1) resolution high enough that convection could be handled explicitly, without parameterization, 2) ability to assimilate data from both Doppler radars and satellites, the observing systems providing the most frequent and highest-resolution information about incipient and mature thunderstorms, and 3) update and output frequency compatible with the need to revise convective products at sub-hourly intervals.

The High-Resolution Rapid Refresh Model (HRRR) meets these specifications. It operates on a 3-km horizontal grid covering the CONUS, with 50 levels in the vertical. It generates convection with explicit calculations. It assimilates reflectivity data from the WSR-88D Doppler radars and cloud data from the GOES satellites. And finally, its update frequency is hourly, with model output produced every 15 min. A description of the HRRR and access to model output are available at http://ruc.noaa.gov/hrrr/.

As of 14 April 2011, the HRRR model is nested within, and derives its initial and lateral boundary conditions from, the Rapid Refresh (RAP) model, which will succeed the Rapid Update Cycle (RUC) at EMC in the spring of 2012. The RAP model has 50 levels, covers most of North America at 13-km resolution, and updates hourly. It relies on a dynamical core from the Advanced Research version of WRF, the Weather Research and Forecasting model, supported and developed by NCAR. The dynamical core refers to advection, pressure gradients, the Coriolis force, buoyancy, filters, diffusion, and time-stepping. In other words, the dynamical core handles the principal terms in the equations of motion and finite differencing techniques. The data assimilation component of the RAP model relies on an adaptation of the Grid point Statistical Interpolation (GSI) analysis employed in the operational Global Forecast System at

EMC. A major advantage of GSI is its capability to assimilate a great variety of satellite data, including radiances, something the operational RUC model cannot do. As of this writing, the HRRR does not have its own data assimilation package; its prognostic variables are initialized by interpolating analyzed fields from the 13-km RAP grid to the 3-km HRRR grid.

Two fairly complete references on the RUC, RAP, and HRRR models and how they are related are Benjamin et al. (2009) and Weygandt et al. (2009). Both are Powerpoint presentations from the same meeting, the 2009 WMO Symposium on Nowcasting held in Whistler, British Columbia, Canada. The papers (3.1 and 3.17, respectively) may be found at http://www.nowcasting2009.ca/images/stories/Presentations/WSN09_Presentations.htm.

The HRRR model is far from operational. EMC has barely enough computing resources to implement the RAP in spring of 2012. The HRRR, which requires far more horsepower, runs in real time on a massively parallel machine at NOAA/ESRL/GSD, requiring roughly 1000 processors. Even so, the output is not available until just over two hours after the initial time. Because the HRRR operates on a 1-h cycle, this implies that three runs must be in various stages of completion at any given time. The oldest run will be in the post-processing stage while the middle run computes a 15-h forecast and the most recent run will have begun processesing observations and assimilating them into an analysis depicting the initial state of the atmosphere.

2.3.8.5 Merging extrapolation forecasts with output from the HRRR model

The CIWS forecasts, combining extrapolation and heuristics, can anticipate 40-dBZ and stronger echos, those of greatest interest to aviation. They also incorporate comprehensive data, many of them at 1-km resolution. Numerical predictions do not yet run routinely at such high resolution, but they do assimilate a very wide variety of observational data and incorporate comprehensive physical processes, thus giving them the chance to successfully capture the complete life cycles of convection, starting with the earliest growth of cumulus clouds and continuing through storm evolution and dissipation. The reason for merging the best aspects of extrapolation, heuristics, and numerical prediction, specifically the HRRR model prediction, is made clear by the diagram at the center of Fig. 2.66.

The blue curve in Fig. 2.66 represents the accuracy of the extrapolative forecast on an arbitrary scale. The accuracy decreases rapidly at first because of errors in motion vectors, assumptions of a constant trend in terms of both motion and intensity, and the somewhat chaotic nature of storms. Model skill (green curve) starts lower than extrapolative skill because of 1) the

disparity between radar pulse volumes and model grid box volumes (some smoothing of details in the reflectivity field is inevitable) and 2) computational noise early in the forecast due to inconsistencies between assimilated fields and model physics. However, because the model can capture nonlinear development and decay of weather features, its skill spins up after two hours and is maintained beyond the first few hours, typically exceeding that of extrapolation by about four hours. Model data assimilation (dashed brown), improves the forecast, causing the skill cross-over point (lead time at which model skill exceeds extrapolative skill) to shift to the left. A blend of the two forecasts is expected to be better than either forecast separately, as shown by the red curve.

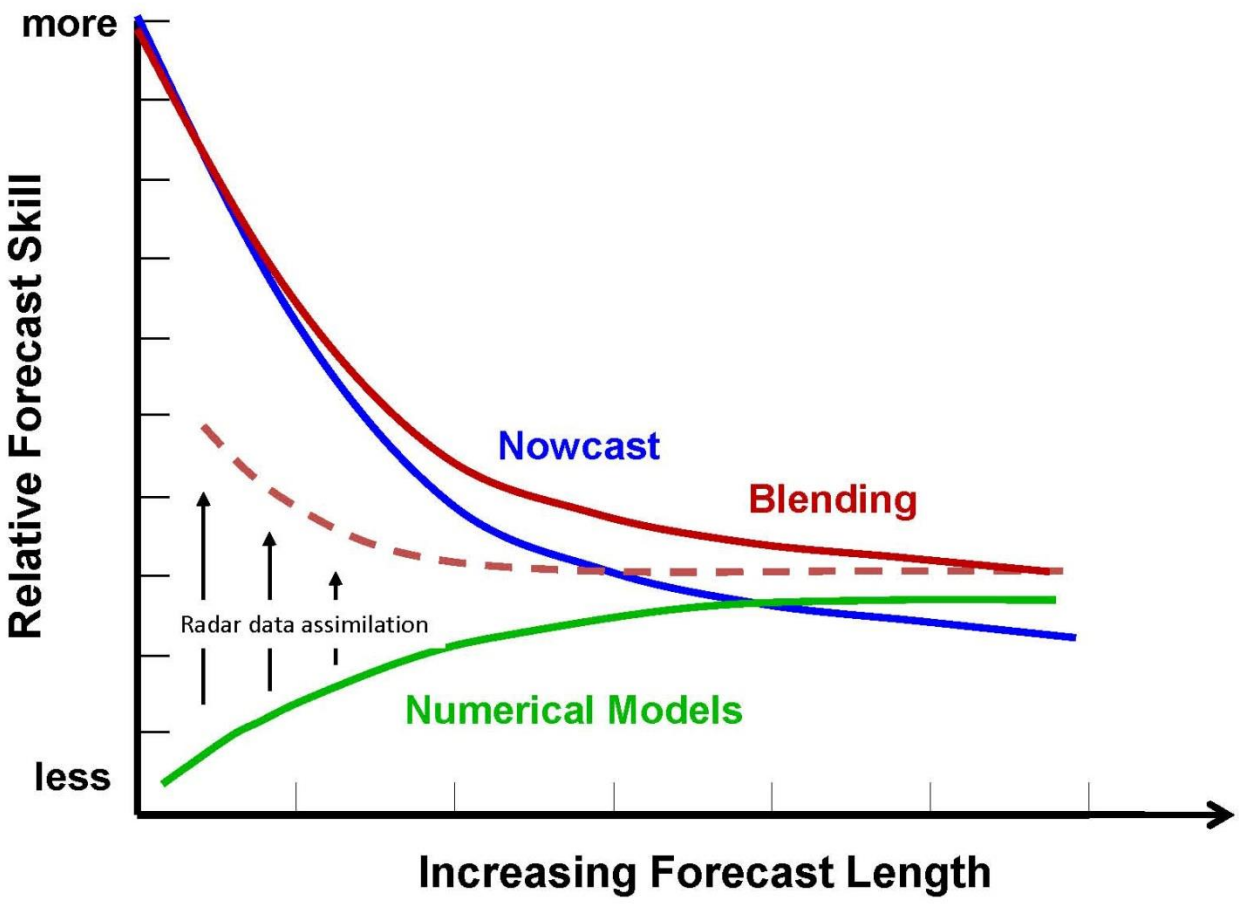

Fig. 2.66. Illustrating relative skill of extrapolative nowcasting techniques (blue curve) versus numerical prediction models (green curve) in the first few hours of the forecast. Frequent assimilation of radar and other high-frequency observations can significantly improve model performance (dash brown curve with black arrows). A statistical blend of extrapolative and numerical forecasts gives the best result (red curve). Figure courtesy of James Pinto, NCAR.

Figure 2.67 illustrates the merging process graphically. At left are the radar echoes as portrayed by the numerical model (top) and the extrapolation forecast (bottom). At bottom middle is a map of the relative weight applied to the forecast at five hours. The color-coded weights range from about 0.4 to about 0.7. The extrapolation forecast would get one minus these weights. The top middle graph shows the overall relative weights received by the two forecasts as a function of lead time. All the weight goes to the extrapolation forecast for hours zero to two, and all the weight goes to the model forecast for hours six to eight. In between, the relative weight given to the extrapolation decreases with time while that given to the model increases. From 1500 to 0000 GMT, the model gets relatively more weight because it has some skill with convective initiation during that part of the day.

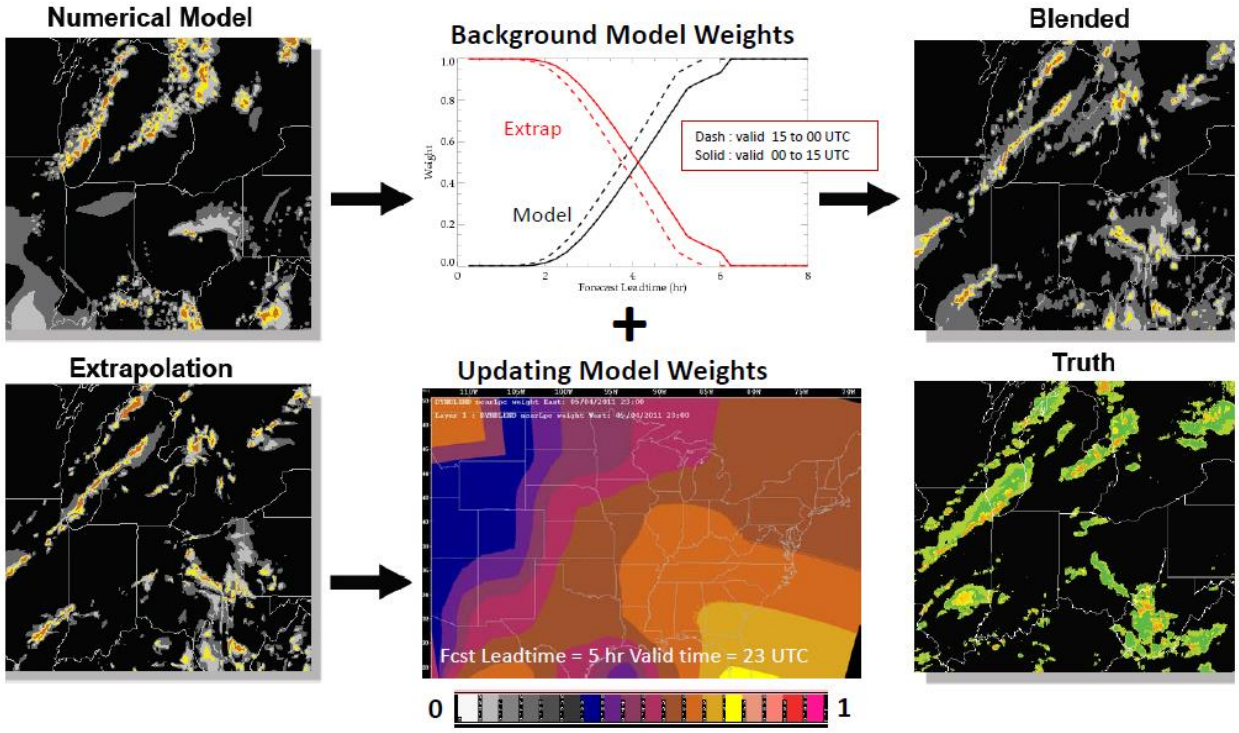

Fig. 2.67. Graphical illustration of the blending of extrapolation and numerical forecasts of convection. Blending extrapolation and numerical forecasts exploits the initial details and trends in the observational data at earlier times and the interacting physics and dynamics of a model at later times. See text for details. Courtesy of James Pinto, NCAR.

The merging process is subject to several constraints: 1) the extrapolation forecast out to two hours from CIWS, available at 5-min intervals, should be untouched. This constraint is

understandable because the HRRR forecast skill is much less than that for extrapolation between 0 and 2 hours. 2) The HRRR forecast, with output available at 15-min intervals, should dominate from 6-8 hours, by which time the skill of extrapolation is much less than that of the HRRR. 3) Because the HRRR forecast is available 2-3 hours after the initial time, it can be phase-corrected (the model radar echoes can be shifted) to match the latest observed radar echoes. This correction is applied at all lead times in the 8-h forecast still ahead of real clock time.

Phillips et al. (2010) and Pinto et al. (2010) provide a high-level description of the merge. The goal is to compare a national radar mosaic of VIL with the VIL field computed from hydrometeor mixing ratios carried by the HRRR model. The comparison occurs every 15 min as the radar mosaic of VIL is updated.

The HRRR model computes mixing ratios for five classes of hydrometeors: rainwater, snow, graupel (snow pellets or heavily rimed snow that can serve as embryos for hail growth), cloud water, and cloud ice. Because of their small diameters, the latter two hydrometeors contribute very little to radar reflectivity and so are not considered further. Hail is not considered because the HRRR model does not predict it. The mixing ratios for rainwater, snow, and graupel, q_r, q_{sn}, q_g, respectively, are used to estimate the equivalent reflectivity factor Z_e (in dBZ). Details are given in Appendix B, Section 2.3.8.4.

Once Z_e has been calculated, the echo top is assumed to be the highest level at which Z_e exceeds 18 dBZ. Next, Z_e is converted to liquid water mass for each model level, then integrated vertically to give a VIL value. The model-derived VIL field is calibrated against WSR-88D VIL values to correct for any bias and converted to a digital scale (Pinto et al. 2010). Finally, the calibrated VIL field from the most recently available model forecast and the corresponding observed VIL field from the NEXRAD network are mapped to a common coordinate system in order to correct position errors in the model echoes. Brewster (2003) described the procedure for this "phase correction."

Figure 2.68, taken from Phillips et al. (2010) shows how the grid with the superposed VIL images from the HRRR model and the latest WSR-88D mosaic is subdivided into overlapping tiles. The model VIL field is shifted up and down, right and left, with respect to the observed WSR-88 VIL field, to get the best match. The best match is obtained by minimizing a "penalty function" as described by Brewster (2003) and explained in Appendix B, Section 2.3.8.4. Minimizing the penalty function yields just a single shift vector for each tile, but this procedure is repeated in several iterations, each time with considerably smaller tile sizes, so that the final

shift vector field has mesoscale detail. The final shift vectors are rather heavily smoothed. If there is a non-zero VIL or echo top value at a grid point, they are moved according to the local

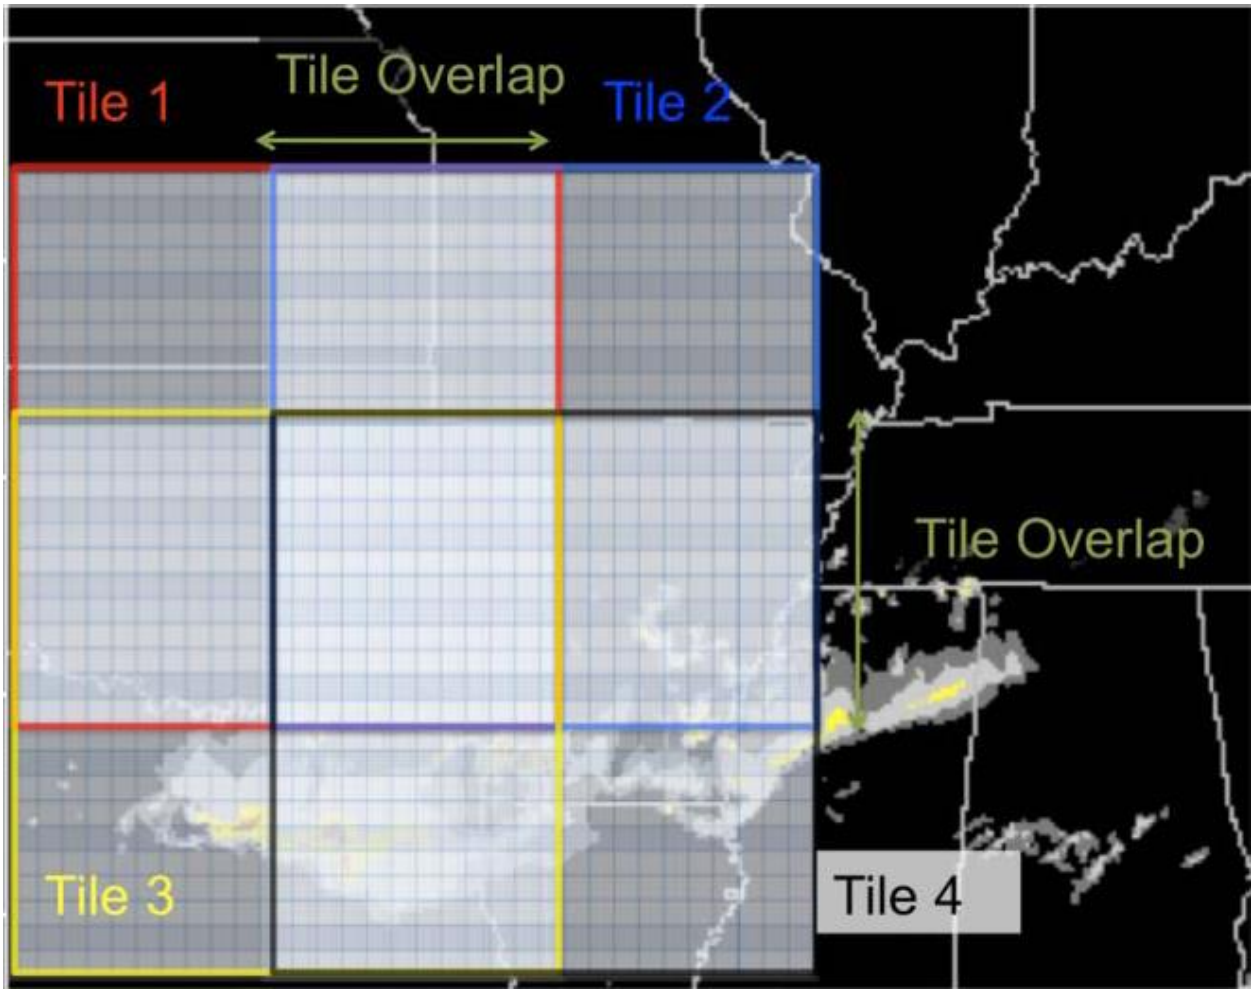

Fig. 2.68. Schematic representation of the configuration of tiles used to compute the phase shifts. A phase correction vector is independently calculated for each tile. Tile 1 in the upper left is outlined (mostly) in red; Tile 2, upper right, blue outline; Tile 3, lower left, yellow outline; Tile 4, lower right, black outline. Each tile overlaps adjacent tiles by half a tile. All four tiles overlap in the middle. Used with permission of the American Meteorological Society.

shift vector. The same shift vectors are applied at all subsequent model forecast times, in practice from three to eight hours, but they are advected with the model-derived motion field. In other words, the same shift vector applied at the current time would also be applied at a later time but at a different location, downwind. Since assumptions of steady state break down

after a few hours, the amount of phase correction applied is also a function of forecast lead time – tapering to near zero by the 8-h forecast lead time.

With phase-corrected VIL and echo-top forecasts in hand from the HRRR model, the final step is to merge these forecasts with the corresponding extrapolative forecasts. The blending occurs between forecast hours two and six. Before forecast hour two, the CoSPA products are constrained to match the currently operational CIWS forecast. Beyond forecast hour six, extrapolation of convection has very little skill, and so the HRRR phase-corrected forecast is allowed to prevail. In between, the two forecasts are blended with weights chosen to minimize bias and maximize skill. This strategy results in weights that depend upon the forecast time, location, and time of day. Generally, up to about four hours into the forecast, the extrapolative forecast receives more weight because it usually has greater skill (except during the time of day characterized by rapid storm growth between 1500 and 0000 GMT, when the model is more skillful). Beyond four hours, the phase-corrected HRRR model is usually more skillful, and so it receives more weight. In addition, the model receives relatively more weight during the afternoon and evening hours, the period favored for storm initiation and growth, because extrapolative forecasts do not handle storm initiation and evolution well. Finally, by maintaining a running mean record of past performance, which is then converted to weights, CoSPA blending can respond to relative changes in forecast skill that might be due to seasonal variability in performance or improvements in the forecast system inputs used in the blending. This method for updating the weights allows for the treatment of regional variations in forecast performance. In the example of Fig. 2.67, the phase corrected model five-hour forecast is given substantially more weight in the southeastern US than across the Great Plains due to variations in the relative performance of the extrapolative and model forecasts for this time of day.

2.3.8.6 Convective Weather Avoidance Model and the CoSPA Probability Product

Delaura and Evans (2006) describe development of the Convective Weather Avoidance Model (CWAM). This begins with an examination of cases in which convection was deemed to be the cause of deviations from planned flight tracks. A deviation occurs if the pilot departs by more than a specified distance from the clear-air flight path. Figure 2.69 illustrates the result of classifying 490 flights as deviating or not from the planned route following a *weather encounter*, which is defined as the planned flight track passing through *either* a VIL Level 2 echo

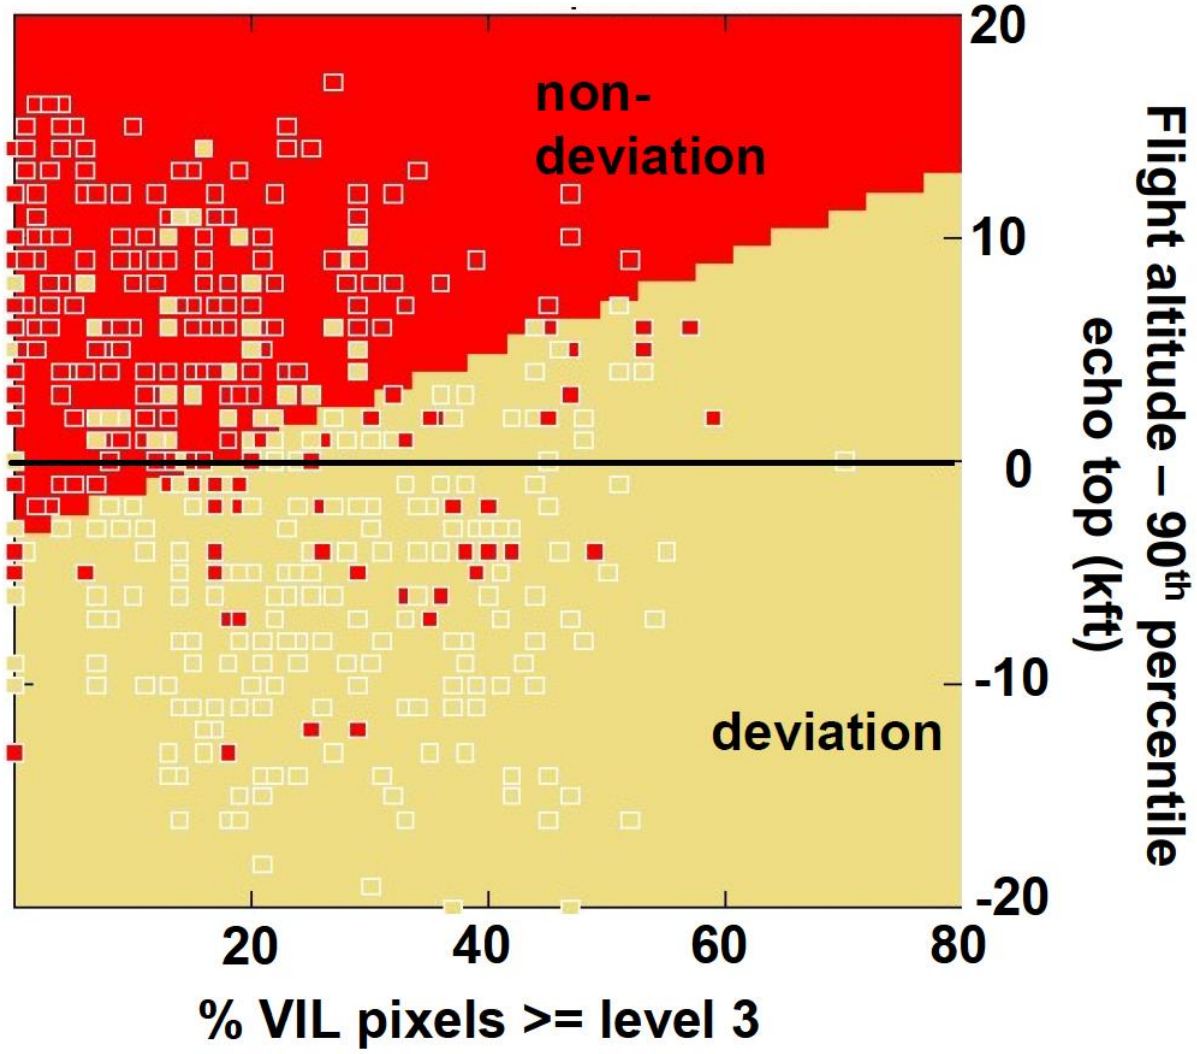

Fig. 2.69. Classification of 490 flights experiencing a *weather encounter* (see text), as a function of percent echo coverage (horizontal axis) and the difference between flight altitude and radar echo top (thousands of feet) along the planned flight track (vertical axis). The boundary between a deviation (gold background) and no deviation (red background) is determined by an empirical model. Incorrect decisions by the model are indicated by gold squares within the red background and red squares within the gold background. From DeLaura and Evans (2006), used with permission of the American Meteorological Society.

or greater *or* echo tops of 25,000 ft or greater for at least two minutes. In Fig. 2.69, the horizontal axis gives the percentage of area covered by VIL 3 or greater echoes. The area is a square 60 km on a side centered on the flight track. The vertical axis is the difference between the flight level and the 90th percentile echo top within an area 16 km on a side, centered on the flight track. The latter parameter is the best predictor of a flight deviation. If echo tops are

much higher than the flight altitude, a deviation is highly probable. The percent coverage of Level ≥ 3 VIL echoes is the second most useful predictor of a flight deviation.

Figure 2.70 shows histograms that indicate how effective the two predictors are in discriminating between flight deviation and non-deviation. The top panel shows the distribution of all flights, and deviations and non-deviations separately, as a function of flight altitude minus the 90th percentile echo top. The bottom panel shows the same, but as a function of areal coverage (%) of Level ≥ 3 VIL echoes.

Note that the two distributions (red and green) in the top panel of Fig. 2.70 are rather well separated, indicating that the altitude difference is a good predictor. The distributions in the bottom panel are less well separated, but the VIL ≥ 3 echo coverage nonetheless adds value as a predictor.

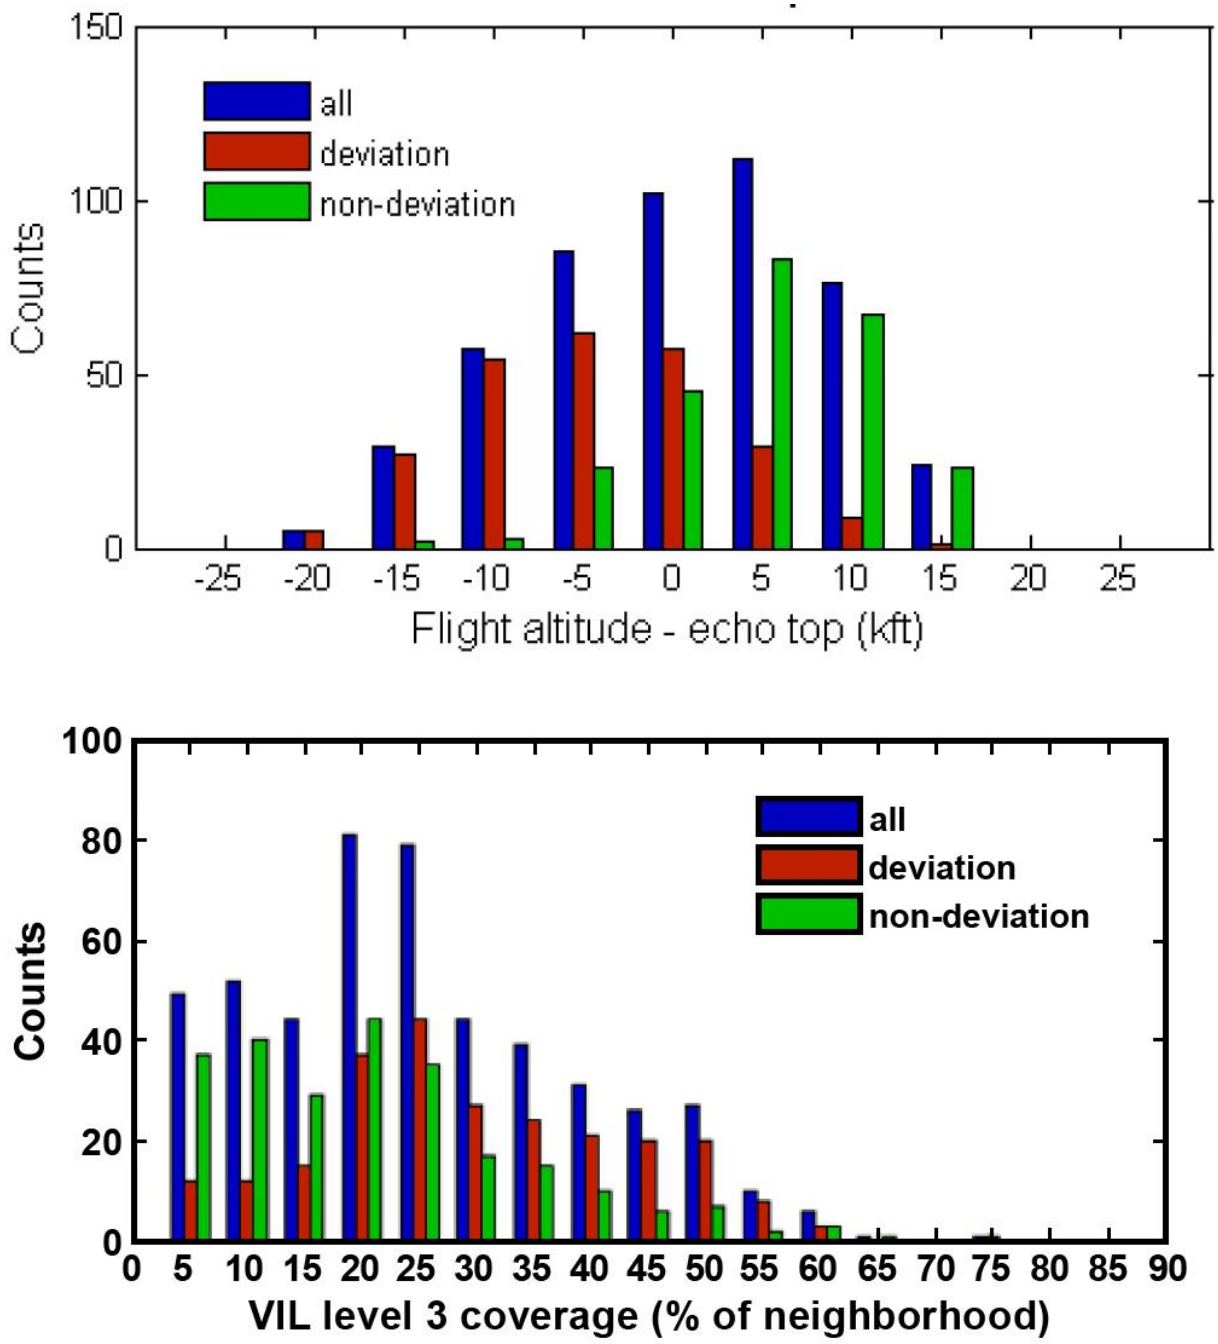

Fig. 2.70. Histogram illustrating the number of flights with a *weather encounter* (blue), and, of those, the number that deviated from the planned route (red) and did not deviate (green). Top: as a function of the difference between flight altitude and the 90th percentile echo top in thousands of feet; bottom: as a function of VIL Level ≥3 areal coverage in percent within a square 16 km on a side centered on the planned flight track. From DeLaura and Evans (2006), used with permission of the American Meteorological Society.

DeLaura et al. (2008) investigated several different predictors of flight deviation beyond those examined by DeLaura and Evans (2006). In both papers, the difference in altitude between flight level and the highest echo tops along the planned route was by far the best predictor. In DeLaura et al. (2008), the second best predictor was percent coverage of echo tops above 30,000 ft. in a 16-km square centered along the planned route. Raw data such as those in Fig. 2.69 were used to construct tables of the probability of flight deviation, given a specific category of altitude difference (flight level minus echo top) and percent coverage of echo tops ≥ 30,000 ft. Figure 2.71 presents the probabilities in color-coded format. The top panel depicts the unedited raw probability data from a scatter diagram resembling that in Fig. 2.69. The bottom panel depicts a Convective Weather Avoidance Model (CWAM), based on the raw probability data. The CWAM includes both smoothing and augmentation (blank squares filled in). Note that the major variation in probability lies in the direction of the vertical axis, not the horizontal axis.

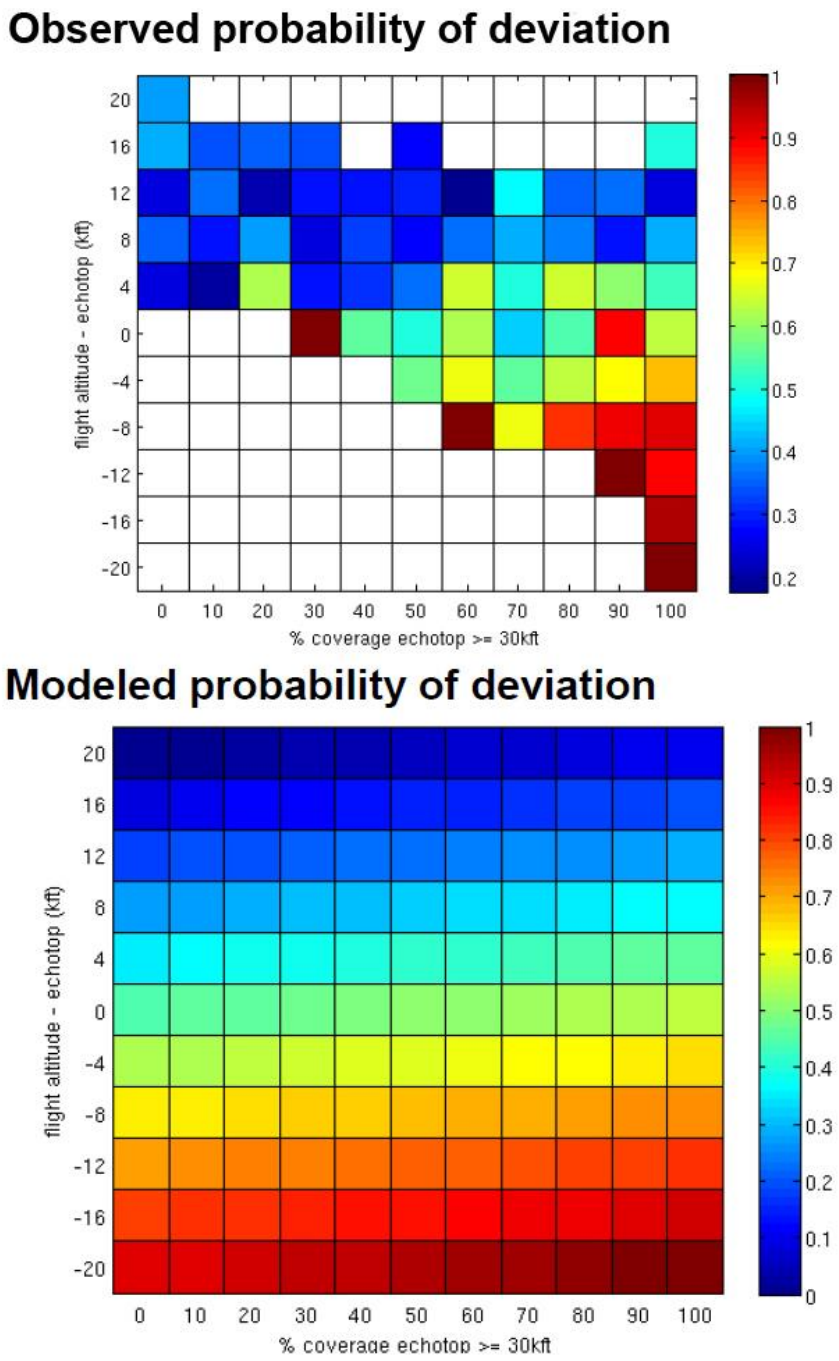

Fig. 2.71. Color-coded probability of a flight deviation, given categories of altitude difference (flight level minus highest echo tops along planned route) and percent coverage of echo tops ≥ 30,000 ft. Top: from observations of hundreds of flight tracks in the vicinity of convection. Bottom: a smoothed version of the observed probabilities with blank boxes filled in. From DeLaura et al. (2008). Used with permission of the American Meteorological Society.

DeLaura et al. (2009) took the idea of probability of flight deviation one step further to create Weather Avoidance Fields (WAFs). This is based upon a two-predictor model, again with the difference in altitude between flight level and the highest tops along the planned route being the principal predictor. The probability of deviation may be based upon either current radar observations or a prediction of VIL and echo tops. Figure 2.72 gives an example of the former. As discussed, the deviation decision model is based upon a historical database of flight paths in the vicinity of convection.

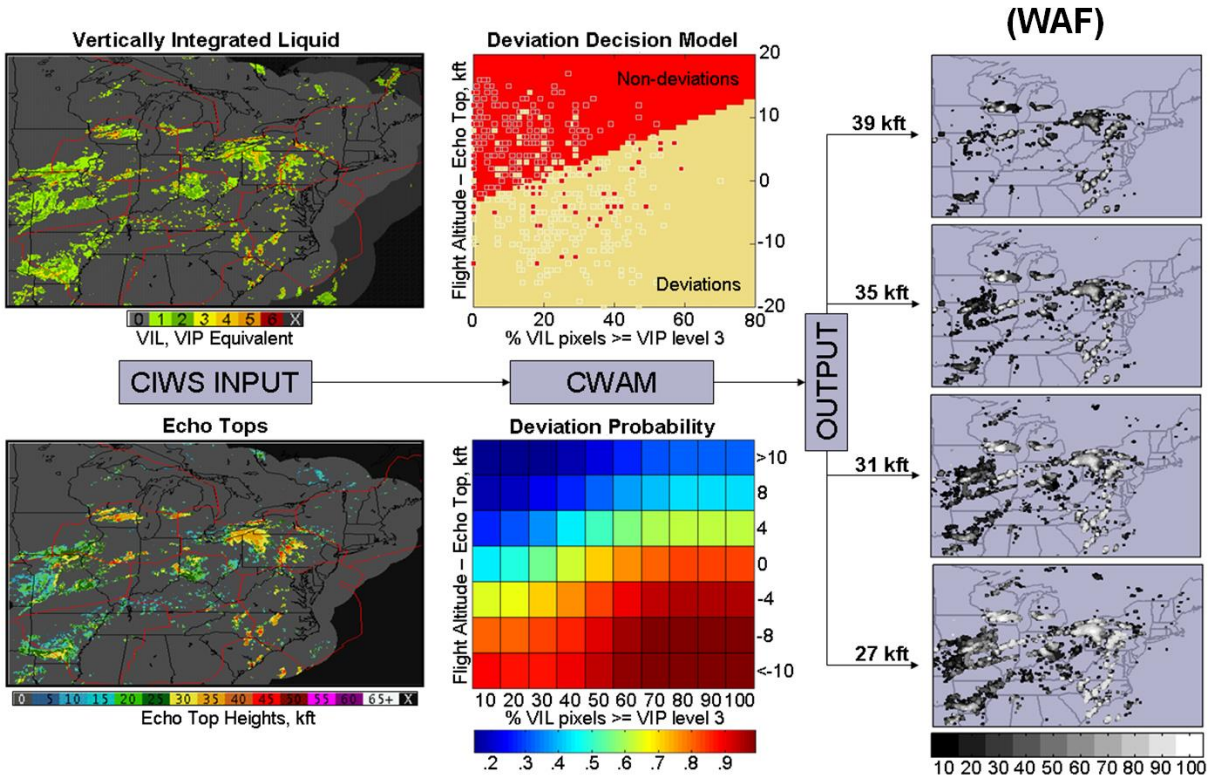

Fig. 2.72. A diagram for the Weather Avoidance Fields (WAFs). The current VIL and echo top products from the Corridor Integrated Weather System (CIWS) at left are the input for a Convective Weather Avoidance Model (CWAM), bottom middle, which gives the probability of a deviation from planned flight path as a function of flight level and the echo tops and VIL values along that path. The probability of a deviation at various flight altitudes is given at right. In this example, the higher the flight level, the less likely it is that a deviation will occur. Figure courtesy of Dave Pace (FAA), his Slide 14 from a Powerpoint presentation on 27 May 2009 regarding the Air Traffic Management Weather Integration Plan.

It is worth asking why pilots deviate from a planned route. If convection is the general cause, visual cues from the cockpit may prompt the deviation, for example, convective towers blocking

the flight path ahead. If the aircraft is flying VFR, on-board radar may indicate a cell ahead with high reflectivity. The pilot normally wants to avoid a bumpy ride and so will generally detour around active convection whenever possible. Lightning, graupel, and hail are additional, coincident hazards. Though a CWAM cannot always anticipate what a pilot will do in convective situations, the WAF product is at least a qualitative guide for deciding how likely a flight deviation will be, along with the corresponding delay in arrival time.

WAFs have been used in operational environments in the New York and Chicago areas within the context of the FAA's Route Availability Planning Tool (RAPT). As discussed by Davison Reynolds et al. (2010), RAPT is an automated decision support tool designed to help air traffic managers determine if and when specific departure times and routes will be affected by convective weather. When convective weather closes departure routes at large hub airports, schedule delays ripple through the entire NAS.

Robinson et al. (2009) discuss the use of WAFs in the RAPT during a multi-year evaluation in the New York City area. Every five minutes, WAFs generates the probability of a flight deviation at each pixel of the RAPT domain. These probabilities apply specifically along departure routes, and they help controllers decide whether these route should be wide open, patially blocked, or closed. One of the most useful findings of that study is that departure routes often remained closed longer than necessary after RAPT indicated (reliably) that weather impacts had cleared.

Though WAFs themselves are not CoSPA products, they could easily be used in a CoSPA probability product. At the time of this writing (January 2012), neither the NWS nor the FAA has sanctioned a probability product for development and use in CoSPA, though candidate products have been proposed. A convective probability product is certainly desired by the avation community and would prove useful for decision-making, especially in route planning.

A probability product based upon an ensemble of forecasts has been proposed and tested by Steiner et al. (2009). It is a possible candidate for a CoSPA probability product. Before discussing this product, however, it is necessary to introduce the concept of maximum flow/minimum cut or, more simply maxflow/mincut. The concept is used in many applications, including traffic flow on streets and highways, but here the interest is in obtaining the maximum throughput of air traffic in a given airspace sector, given that convective storms prevent flying direct paths through the sector and block or at least limit certain pathways. A simple, intuitive description of maxflow/mincut ideas is available at: http://www.sce.carleton.ca/faculty/chinneck/po/Chapter9.pdf.

In simplest form, maxflow/mincut problem is illustrated in Fig. 2.73. The flow is from left to right in a corridor ten squares wide. To the left of the figure, the flow is unrestricted, and so the full width of ten squares is available for air traffic. Two barriers to the flow, considered impenetrable, occupy the middle of the figure. The greatest bottleneck to the flow is shown by the short arrows, indicating the mincut. Evidently, only four squares are available for passage through these idealized linear storms. The obstructed mincut is four. The available flow capacity ratio (AFCR) is the mincut when convection is present divided by the mincut when there is no blockage by the weather, in this case, 4/10.

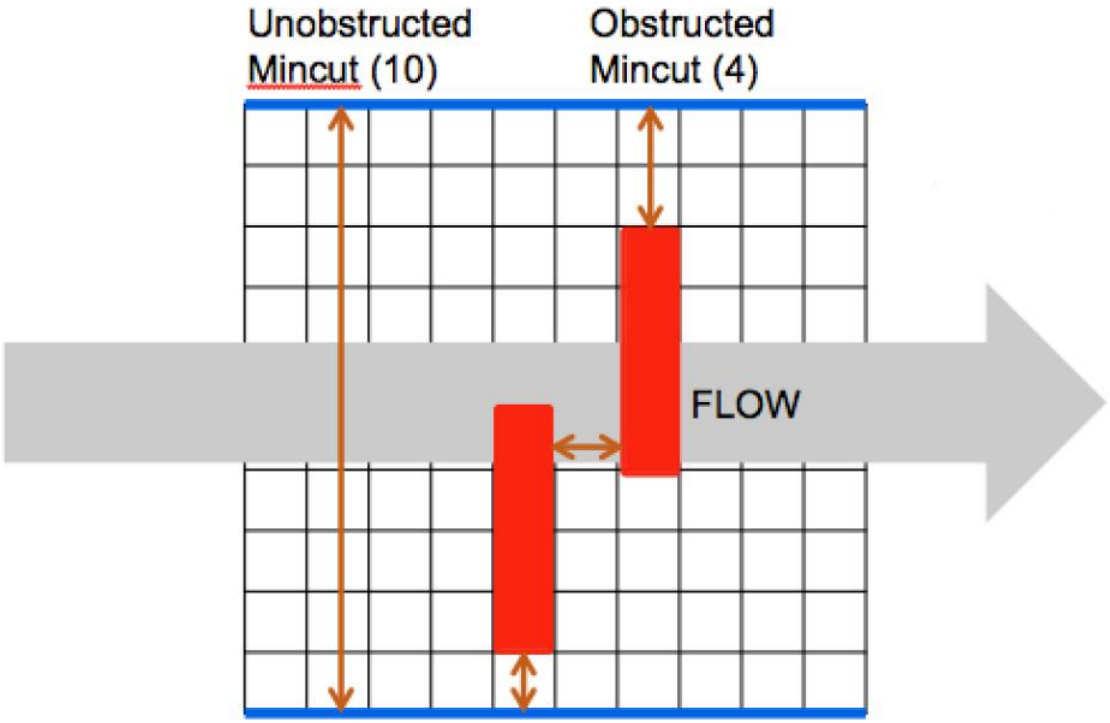

Fig. 2.73. A simple example of flow blockage by two impenetrable lines of idealized thunderstorms shown in red. The greatest constriction of the flow occurs along the path shown by the short arrows. This path defines the minimum cut (mincut), linking the top and bottom of the flight corridor through the obstacles. From Lack et al. (2011).

Figure 2.73 illustrates flow blockage in a single direction. The report by Lack et al. (2011) uses hexagons to compute the flow blockage in three directions, as shown in Fig. 2.74. The left side of the figure illustrates significant blockage (0.75) when the flow is either toward 240° or 60°. In this case, the mincut is the shortest distance between the lower edge of the echo and the edge of the flight corridor (black arrow). But traffic flow could also be in the N-S direction, in which case the mincut would extend from the left edge of the northern boundary of the

hexagon to the left edge of the echo, and then across the shortest path from the echo to the NE edge of the hexagon. This would result in a blockage of only 0.40. Finally, the blockage to traffic travelling in the 300° or 120° directions would be minimal (0.05) because the echo is approximately aligned along the flight path and off to the side. On the right side of Fig. 2.74, the lengths of the lines drawn through the center of the hexagon show the blockages for the three directions of flight traversing the hexagon.

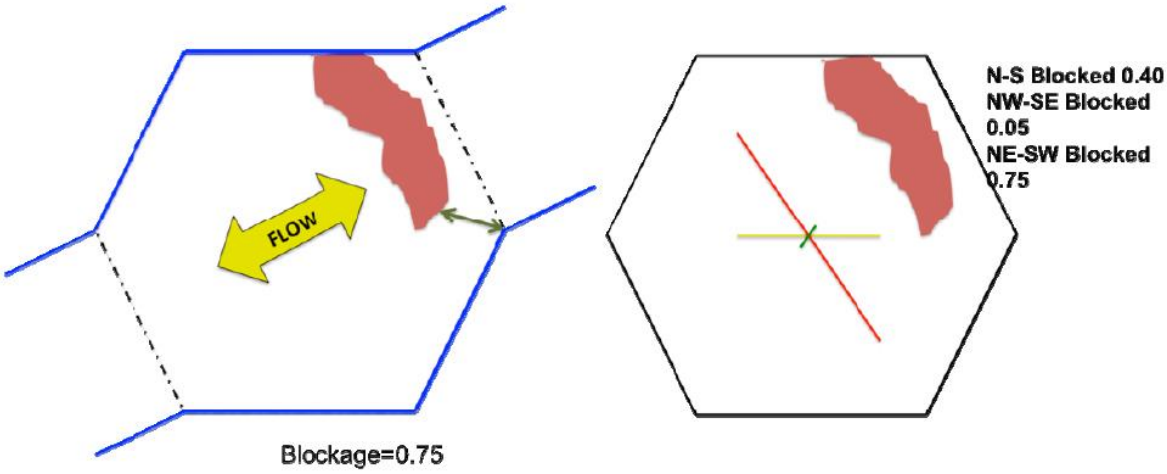

Fig. 2.74. Using hexagons in a honeycomb pattern to fill air traffic control space, allows estimation of route blockage in three different directions. As shown here, the amount of blockage depends upon the size and orientation of the echo within the hexagon. From Lack et al. (2011).

Finally, blockage can be computed for Air Traffic Control sectors, which are contiguous polygons. One such polygon, Sector B, is shown in Fig. 2.75 from Steiner et al. (2009). In this example, traffic enters the polygon from Sector A and exits into Sector C. Three obstacles constrict the passageways through the polygon. The mincut without weather hazards is the width of the entry edge of the polygon because nowhere within the polygon or on its exit edge is the unobstructed flight corridor narrower. The mincut when three convective cells block part of the airspace is indicated by the thick red lines.

Finding the mincut in an irregularly shaped polygon when many convective echoes are present is a complicated problem. Mitchell et al. (2006) give an entry point to the mathematics of

maxflow/mincut theory, which was developed mostly before 1990. Algorithms are available to find the mincut when convective storms are impenetrable but also when they are permeable.

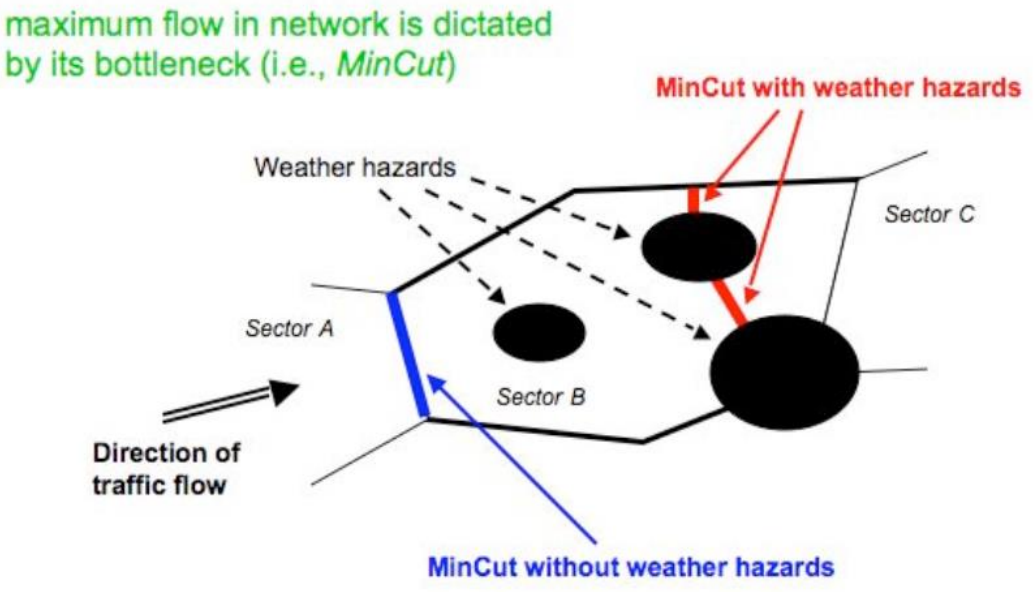

$$\text{Available Flow Capacity Ratio} = \frac{MinCut(weather)}{MinCut(clear)}$$

Fig. 2.75. Illustrating how the maximum flow rate through an airspace sector is determined by the mincut for clear weather and the mincut that accounts for convective obstacles. From Steiner et al. (2009)

Consider Fig. 2.76, presenting a situation more complicated than that in Fig. 2.73. The two convective lines previously considered impenetrable are now considered permeable. There is a 50% probability of blockage in the northern convective line (rose color) and a 75% probability of blockage in the southern line (brownish rose). Moreover, a large area in light pink has a 1/3 chance of blockage. The mincut through the airspace now involves all squares along the mincut path because none of the squares is completely blocked. The calculation of the (partially) obstructed mincut length is given at top right of the figure. Note that the calculation uses the probability of penetration rather than the probability of blockage. The AFCR is 6/10, better than the 4/10 calculated in Fig. 2.73.

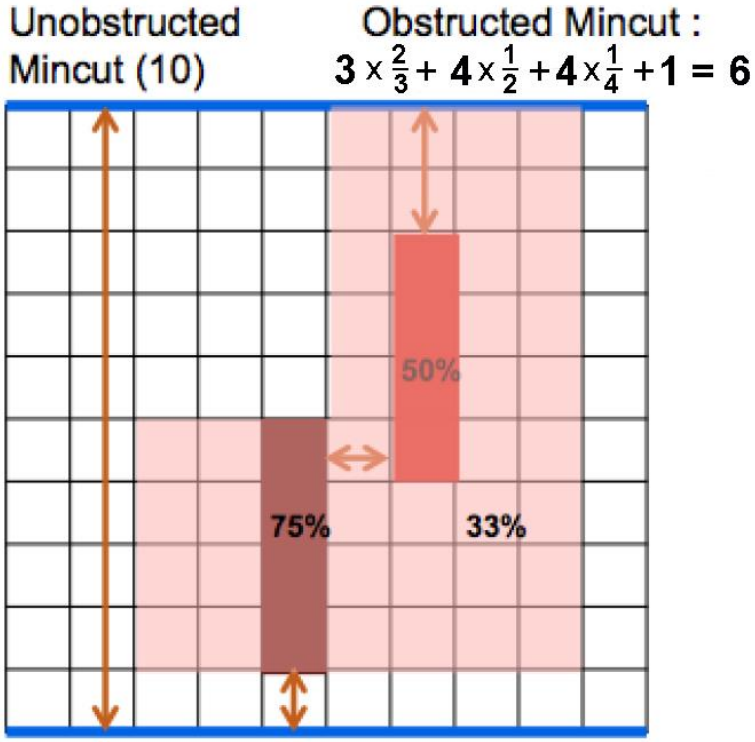

Fig. 2.76. An example of the mincut calculation for a situation similar to that in Fig. 2.73, except that squares are colored according to the probability of blockage, whereas in Fig. 2.73, the blockage was either total or nonexistent. The calculation of the obstructed mincut is at top right; it uses the probability of penetration (one minus the probability of blockage). From Lack et al. (2011).

The foregoing introduction to maxflow/mincut concepts prepares the way for a discussion of the paper by Steiner et al. (2009), which shows how an ensemble of 28 mesoscale (10-km resolution) forecasts for the northeast U.S. can be used to estimate the probability of flow reductions. First, criteria for a blockage are specified. For example, locations where model VIL values or cloud-top heights exceed a certain threshold define a blockage. The remaining steps are illustrated in Fig. 2.77.

In the upper left panel, a grid of squares, say, 50 km on a side, overlies a map of whatever model parameter defines a blockage, either analyzed or predicted. For *each* box of the grid, the

AFCR is calculated based upon the strength and coverage of convective storms within that box. This is accomplished for *each* member of the ensemble, resulting in a histogram (upper right) of AFCR values (28 in all, one for each forecast). The histogram values are then binned by AFCR in increments of 0.1, from 0.0 to 1.0. The number of counts in each bin is shown at lower right. Note that many forecasts for this particular box produced an AFCR between 0.7 and 0.8,

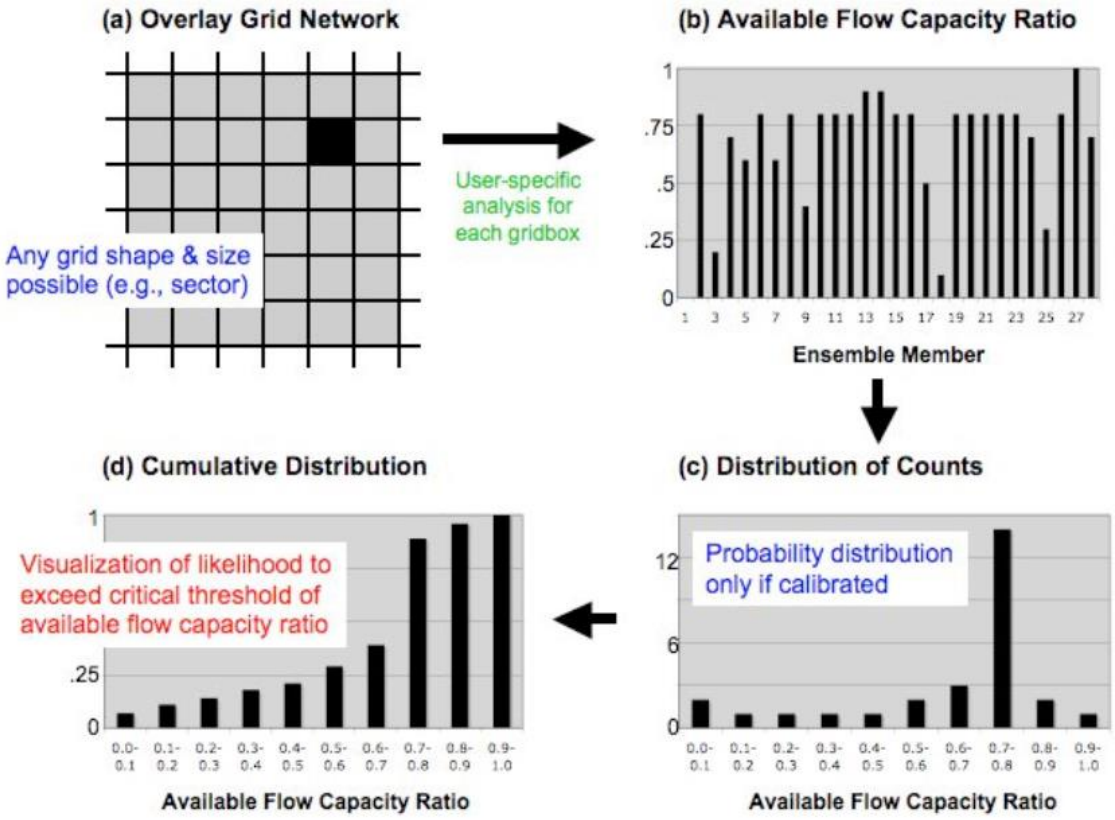

Fig. 2.77. The procedure for determining the probability of an Available Flow Capacity Ratio (AFCR) for any given threshold. (a) For a fixed area within a grid (the black square) compute the AFCR resulting from each member of an ensemble of forecasts. (b) The AFCR computed within the black square for each of 28 forecasts. (c) Counting the number of forecasts for which the AFCR was in specified ranges from 0.0 to 1.0 in increments of 0.1. (d) The cumulative probability distribution of AFCR values derived from (c). From Steiner et al. (2009).

equivalent to a blockage of 20-30%. If the distribution of counts is properly calibrated, it may be viewed as a probability distribution. This, in turn is easily converted to a cumulative probability distribution (lower left). The interpretation is straightforward. The probability of an AFCR < 0.8 is about .91; the probability of an AFCR < 0.7 is about .38, and so on.

If the information in Fig. 2.77d is compiled for every square of the grid in Fig. 2.77a, then a basis exists for generating a probability map such as that on the left side of Fig. 2.78. For this map an

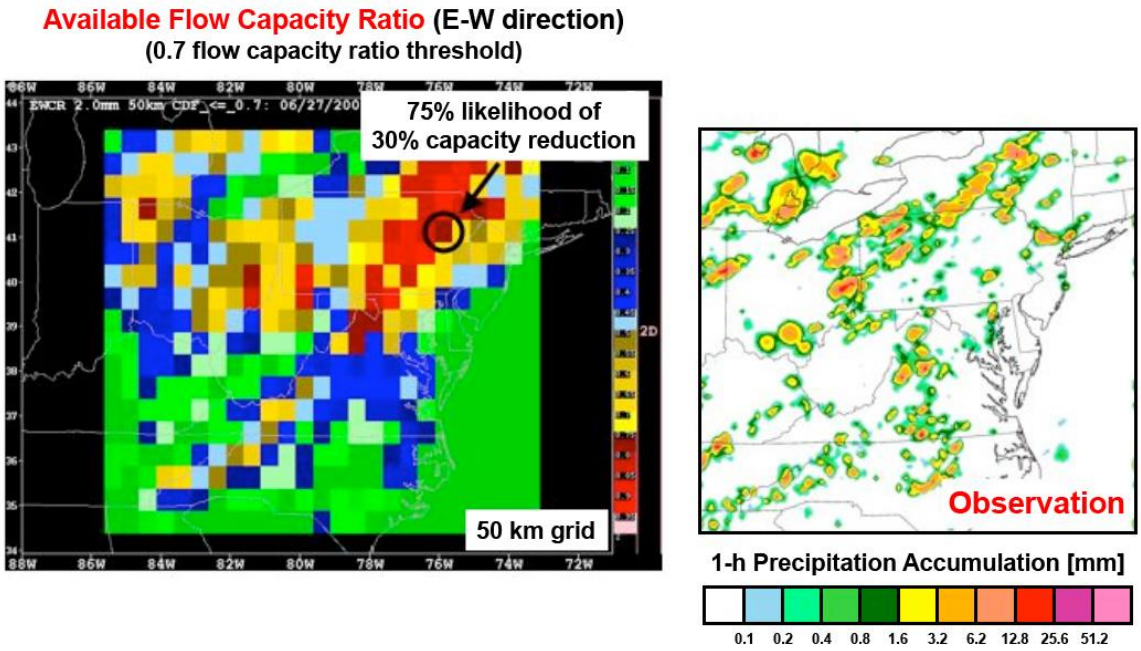

Fig. 2.78. At left, the probability of a 30% traffic reduction within grid boxes 50 km on a side generated from information such as that shown in Fig. 2.77d. The probabilities are color-coded as indicated in the color bar at right. At right, the observations of hourly precipitation rate corresponding to the flow reduction map at left. VIL and hourly precipitation rate are well correlated. From Steiner et al. (2009).

AFCR threshold of 0.7 was specified, which is equivalent to a 30% capacity reduction (blockage). In the example just discussed, the probability of a 30% capacity reduction was about .38. In Fig. 2.78, the probabilities are color coded, dark green indicating no chance of flow reduction, and yellows, oranges, and reds indicating increasing chances of flow reduction. information is directly relevant to Air Traffic Control in that it translates the weather-related products of CoSPA into a decision-making tool.

Keep in mind that convection is not the *only* cause of flight deviations and delays. Equipment malfunctions, controller workload, and excessive air traffic are other causes.

2.3.8.7 Verification of CoSPA products

CoSPA has built-in verification for both current and archived forecasts. This is a valuable asset because convective forecast accuracy is highly situation-dependent: a 4-h track forecast of a rotating updraft thunderstorm may be surprisingly accurate, whereas a 1-h forecast of a single-cell thunderstorm developing in an unsheared environment may be poor indeed. The next two figures are from the CoSPA website. Figure 2.79 shows an echo-top analysis for 1525 GMT, 01 July 2011. The colored contours are for cloud tops ≥ 30,000 as predicted 2, 4, 6, and 8 h earlier. White, red, purple, and black contours correspond to 2-, 4-, 6-, and 8-h forecasts, respectively.

The convection in South Dakota and Nebraska had persisted all of the previous night. It was forced by a compact short-wave with high values of cyclonic curvature vorticity. The mesoscale convective system (MCS) centered over Indiana had developed late on the night of 30 June, forced by a strong low-level jet mostly below 850 mb centered on a moist tongue stretching almost all the way back to the Gulf of Mexico. Though both these features persisted for many hours, the forecasts favored dissipation by mid-morning on 01 July. There were no purple or black contours for the storms over South Dakota and Nebraska, and no black contours near the MCS in northern Indiana. The purple contours associated with this MCS were too far north and south. Still, one could say that the 4-h forecasts for both systems were fairly accurate.

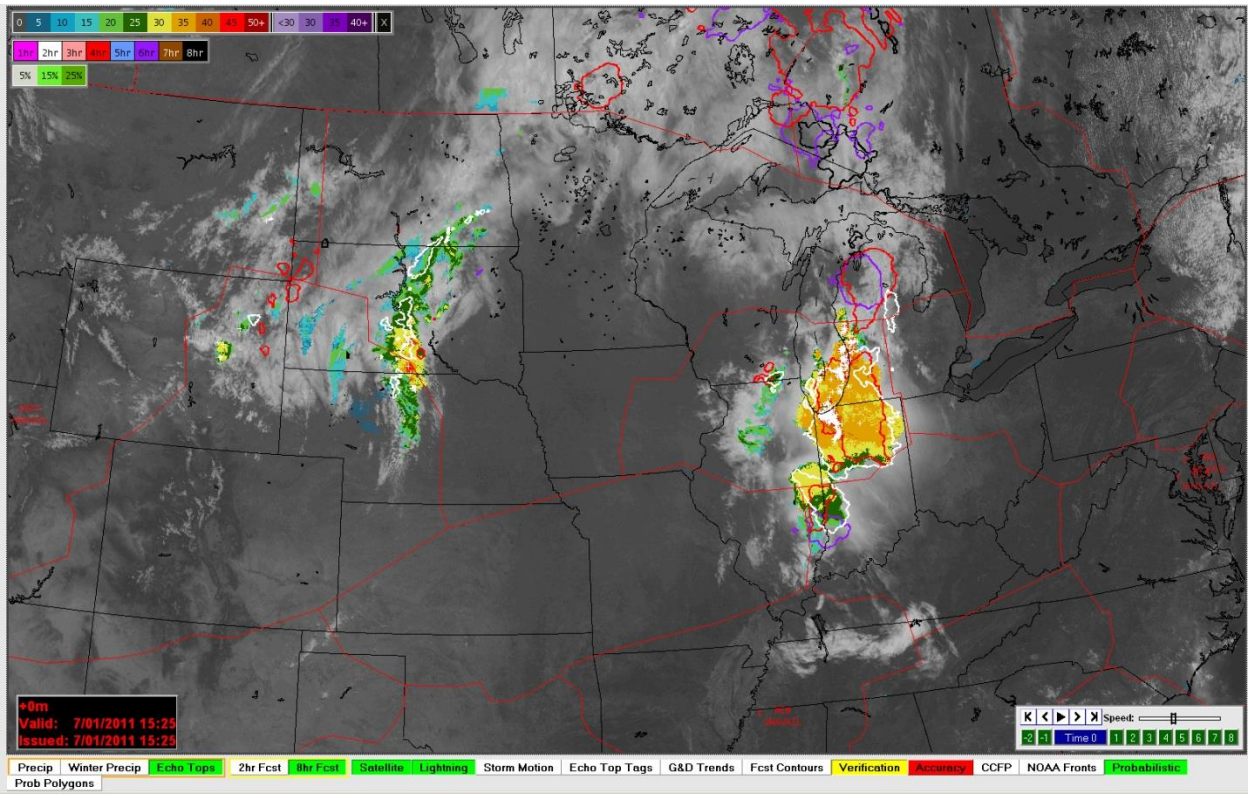

Fig. 2.79. The echo top analysis for 1525 GMT, 01 July 2011 superposed on a GOES visible satellite image. The color bar for echo-top heights is at upper left. Cloud-to-ground lightning strikes appear as tiny white crosses (+), sometimes densely clustered. Areas where tops were predicted to be 30,000 ft and higher are contoured. White contours are for 2-h forecasts verifying at 1525 GMT. Red contours are for 4-h hour forecasts verifying at the same time. Purple contours are for 6-h forecasts, and black contours (none visible in this image) are for 8-h forecasts.

By contrast, examine verification for the echo-top forecasts along the Gulf Coast in Fig. 2.80, valid at the same time as Fig. 2.79.

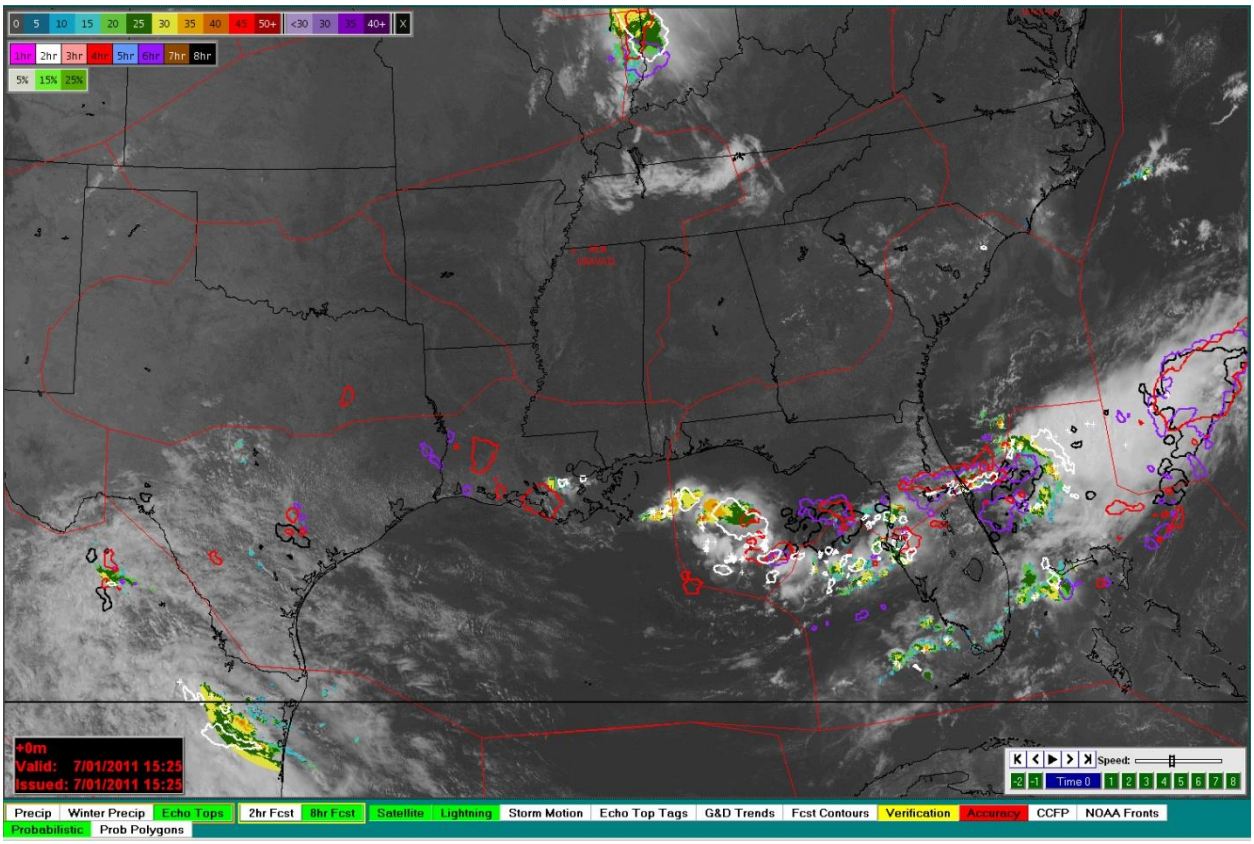

Fig. 2.80. Same as Fig. 2.79 except that verification is shown for the Gulf Coast states instead of the Great Plains and mid-Mississippi Valley.

This convection was mostly offshore, as is typical in the morning, before the sea breeze circulation becomes well established. The convection lies on the southeast periphery of a closed anticyclone aloft, is mostly weak, short-lived, not dynamically forced, and the individual cells are small. This kind of convection is very difficult to forecast, and a comparison between the observed and predicted echo tops confirms this. Except at two hours, there is little correspondence between the observed and predicted convection. The models consistently developed convection offshore within the dense cloud shield east of Florida, but there are no radars close enough to confirm whether this occurred.

The colored contours of Figs. 2.79 and 2.80 portray how the forecast is performing on individual cases. Section 2.3.10 assesses longer-term performance of CoSPA, but in comparison with several other convective products, one of which is LAMP. The latter has yet to be described in Section 2.3.9.

2.3.8.8 Commentary on CoSPA

In that CIWS provides the first two hours of convective weather guidance for aviation, comments made regarding CIWS in Section 2.3.7.6 apply equally to CoSPA. CIWS and its temporal extension provide the best convective forecast currently available for about three hours. After that, mesoscale prediction models begin to show greater skill. In extending the convective forecast out to eight hours, CoSPA capitalizes on the strengths of CIWS and the HRRR model by blending output from the two during the period when the relative skill of CIWS is waning and that of the HRRR is growing. The combination of the two results in a forecast that is usually superior to either one of them alone over the blending period.

Another strength of CoSPA is the high-resolution displays, which present both analyses and forecast fields with equal clarity and resolution. To the credit of developers and users who provided feedback, the color codes for the predicted fields are different from those used for the analyzed fields, so that the two are always easily distinguished.

The computation of VIL from hydrometeor mixing ratios in the HRRR model seems convoluted. The model mixing ratios are first converted to reflectivities at each level, these are converted to water content, then vertically integrated to give VIL. It would be much simpler and more direct to go straight from the mixing ratios of precipitation-size particles, namely, rain, snow, and graupel, to VIL.

CoSPA needs a probability product, but none is supported at this time. There are at least three possible candidates: 1) use of successive WAF forecasts, 2) the product developed by Steiner et al. (2009), involving the Available Flow Capacity Ratio and illustrated in Figs. 2.77 and 2.78, and 3) mesoscale ensemble forecasts, though the skill of the ensemble average forecast might not match that of CIWS in the first hour or two.

The length of this section on CoSPA is testament to decades of development in image processing, filtering, numerical modeling, and computer graphics. Some of the forecast products resemble Doppler radar images, a feature many aviation decision managers say they like. They appreciate that the accuracy of convective forecasts decreases with each additional hour of lead time, but they must be vigilant that the realism of the forecast product and its resemblance to standard radar displays is not interpreted as equivalence in accuracy. In a sense, display technology has outstripped advances in convective forecasting. The theoretical predictability of convective storms, especially in weakly forced situations, may be only an hour or two, which means that high-resolution models of the future, benefitting from

comprehensive prescriptions for cloud physics, dense observational coverage at convective scales, and enough computational horsepower to generate forecasts at sub-kilometer resolution in less than an hour, may still perform poorly at two hours. Even so, there is still much room for progress, and CoSPA is the designated vehicle for convective aviation forecasts for NextGen.

2.3.9 The LAMP convective forecast

LAMP stands for Localized Aviation Model Output Statistics (MOS) Program (LAMP). LAMP was developed by the Meteorological Development Laboratory of the NWS. Since November 2008, LAMP has provided statistical guidance for aviation forecasting out to 24 hours, updated hourly. LAMP produces meteograms of surface temperature, dewpoint, and wind direction and speed: graphical depictions of the expected variations in these parameters over the next 24 h. In addition, LAMP gives categorical and probabilistic forecasts of ceiling height, visibility, and obstructions to vision. These will be discussed at length in Section 3.1. Finally, LAMP produces categorical and probabilistic forecasts of measurable precipitation, precipitation type, and thunderstorms. Only the thunderstorm forecast is relevant here.

Ghiradelli and Glahn (2010) describe LAMP in detail. LAMP relies on five primary sources of information: 1) METAR surface observations, 2) gridded (10-km) radar and lightning observations, the latter from the National Lightning Detection Network, 3) MOS guidance from the Global Forecast System (GFS), the operational global model run by the EMC, 4) output from advective and other simple models, and 5) and several miscellaneous predictors, some relevant to thunderstorms.

MOS uses observations, persistence, and conditional climatologies in a multiple linear regression framework in part to remove biases in numerical forecasts and in part to estimate weather conditions at the surface and aloft at locations that may not be explicitly forecast by the model, for example, the cloud ceiling and surface visibility seven hours from now at Minneapolis. MOS processes output from the GFS every three hours to make these forecasts and provides them at over 1500 locations throughout the U.S.

The LAMP convective probability product is available at: http://www.nws.noaa.gov/mdl/lamp/convection.php. This product gives the probability of convection in 20-km grid boxes for 2-h periods in the 3- to 25-h range over the CONUS. In addition to the information sources noted above, this product exploits information from the NCEP North American Mesoscale (NAM) model, which has the advantages over the GFS of

being nonhydrostatic and running at 13-km resolution. Convection in a grid box is defined as the occurrence of either radar reflectivity ≥ 40 dBZ or one or more cloud-to-ground lightning strikes (or both) during the 2-h valid period. Figure 2.81 gives the color-coded probability of convection between six and eight hours.

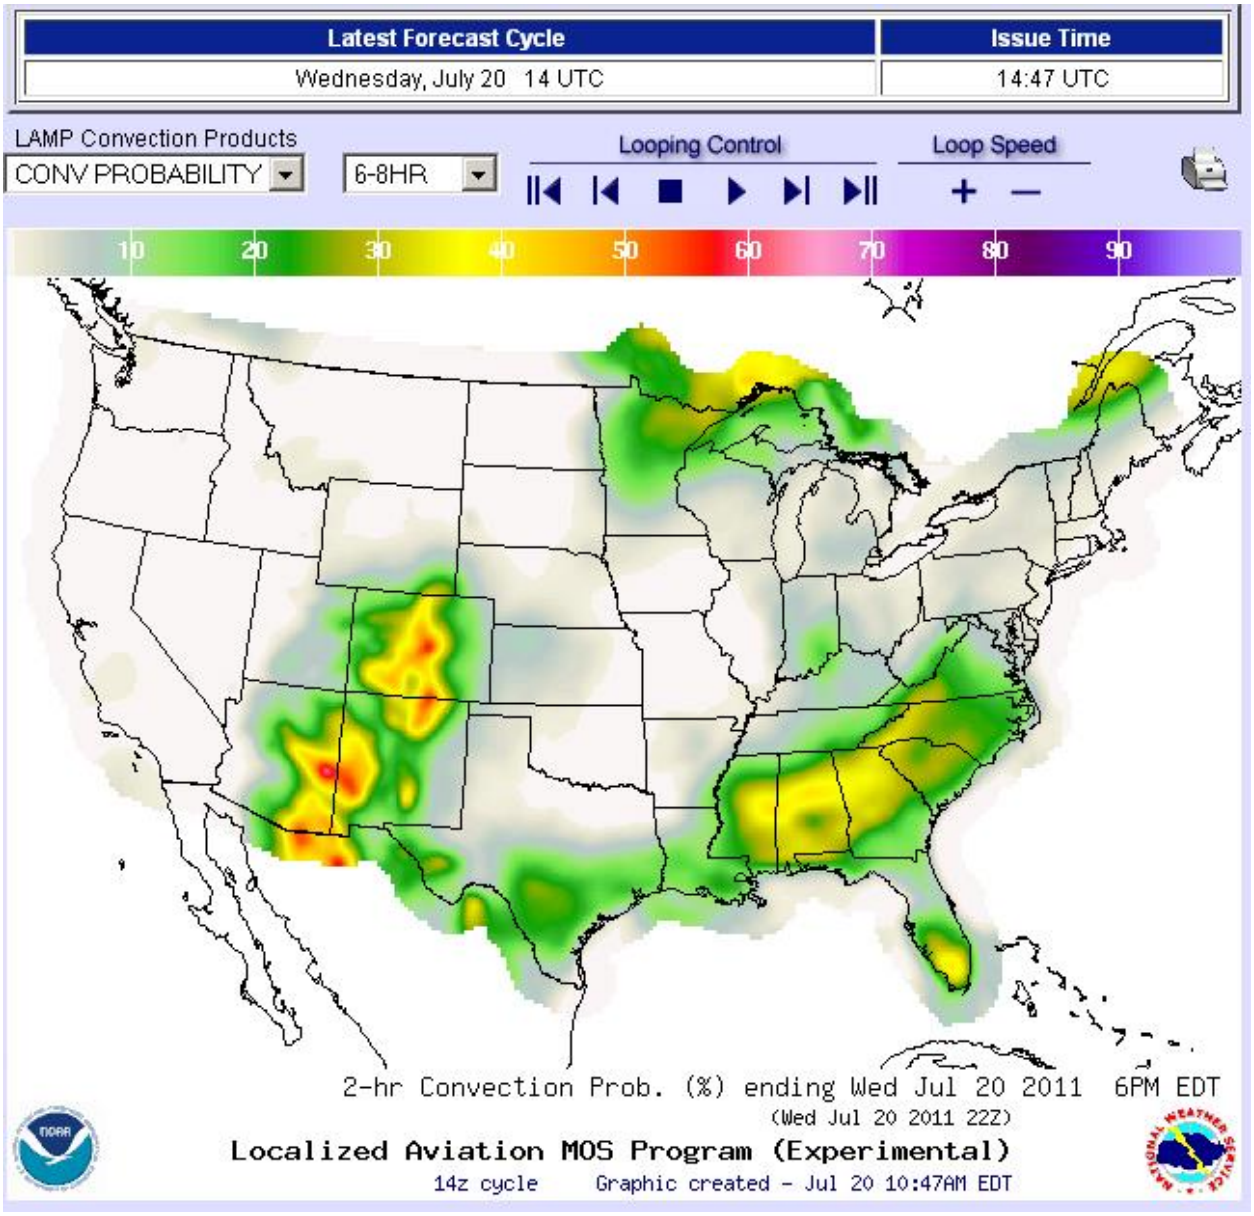

Fig. 2.81. A sample LAMP convective probability product, showing the probability of convection in 20-km squares for 6-8 hours in the future.

LAMP also generates categorical forecasts of convection potential: zero, low (L), medium (M), and high (H). The probability threshold for each category is derived in such a way that the CSI score is maximized within a prescribed bias range. The average bias is specified as 2.7 for L, M, and H potentials combined. This means that, for every 27 forecasts of L, M, or H potential, 10 events are observed. The average bias for M and H forecasts is specified at 1.1, and for H potential only, the average bias is 0.4 (H events are deliberately underforecast). Figure 2.82 shows a categorical forecast for 6-8 hours.

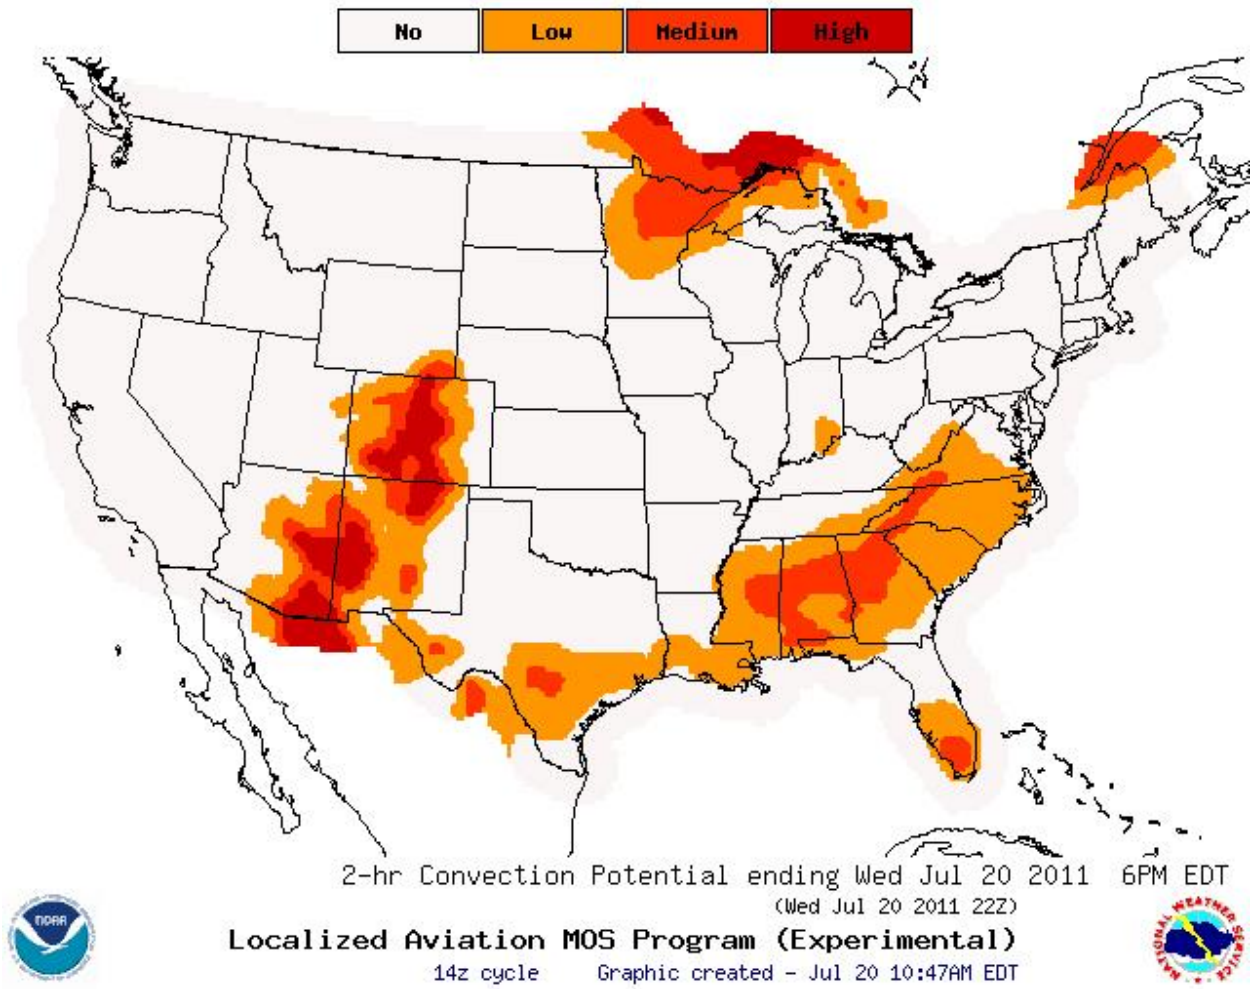

Fig. 2.82. LAMP categorical forecasts of convection potential for 6-8 h in the future. Shades of orange and red give the potential for convection: low, medium, or high.

Charba and Samplatsky (2009) verified forecasts illustrated in Figs. 2.81 and 2.82 during warm and cool seasons, 01 Jul – 15 Oct 2007 and 16 Oct 2007 – 15 Mar 2008, respectively. They used the Brier Skill Score to verify probability forecasts and the CSI to verify the categorical forecasts. CSI was defined in Table 1.2; the Brier Skill Score is explained here.

If p_i is the probability that an event i will occur, a probabilistic forecast is verified by an observation that the event either occurred ($o_i = 1$) or did not ($o_i = 0$). For a sample of N events, the *Brier Score* is defined by

$$BS = \frac{1}{N}\sum_{i=1}^{N}(p_i - o_i)^2 .$$

The range of BS is zero (perfect score) to one (no skill whatever). The Brier Score has the disadvantage that, the rarer an event, the easier it is to get a good score without having any real skill. Thus, if one predicts zero probability all the time for a three-inch deluge, the score will be very nearly zero, but not much has been accomplished by such a forecast. The *Brier Skill Score* BSS overcomes this disadvantage by introducing a reference forecast, usually climatology, and evaluating the forecast skill relative to this reference.

$$BSS = 1 - \frac{BS}{BS_{ref}}.$$

A perfect BSS score is one. A score of zero indicates no skill beyond that of the reference forecast. In the example, p_i for the reference forecast would be the climatological probability of a three-inch rainfall (a very small number). A score near unity would be very difficult to attain because it is hard to correctly forecast a rare event while significantly beating climatology.

Figure 2.83 from Charba and Samplatsky (2009) shows skill scores for the LAMP products illustrated in Figs. 2.81 and 2.82. In Fig. 2.83a, convective probability forecasts are verified with the Brier Skill Score. In Fig. 2.83b the occurrence / non-occurrence forecasts for each category, low, medium and high potential, are verified with the threat score (same as CSI). All forecasts are issued at 0900 GMT. The greater skill in the early hours reflects the positive influence of radar, lightning, and surface weather observations, but this influence is short-lived. Cool season forecasts (blue) are almost always more accurate than warm season forecasts because cool season convection is dynamically forced far more often than warm season convection. This reflects the positive influence of models, but even that has disappeared after about seven hours. The secondary maximum in forecast accuracy at 13 hours is probably related to the common afternoon peak in convective activity (2200 GMT).

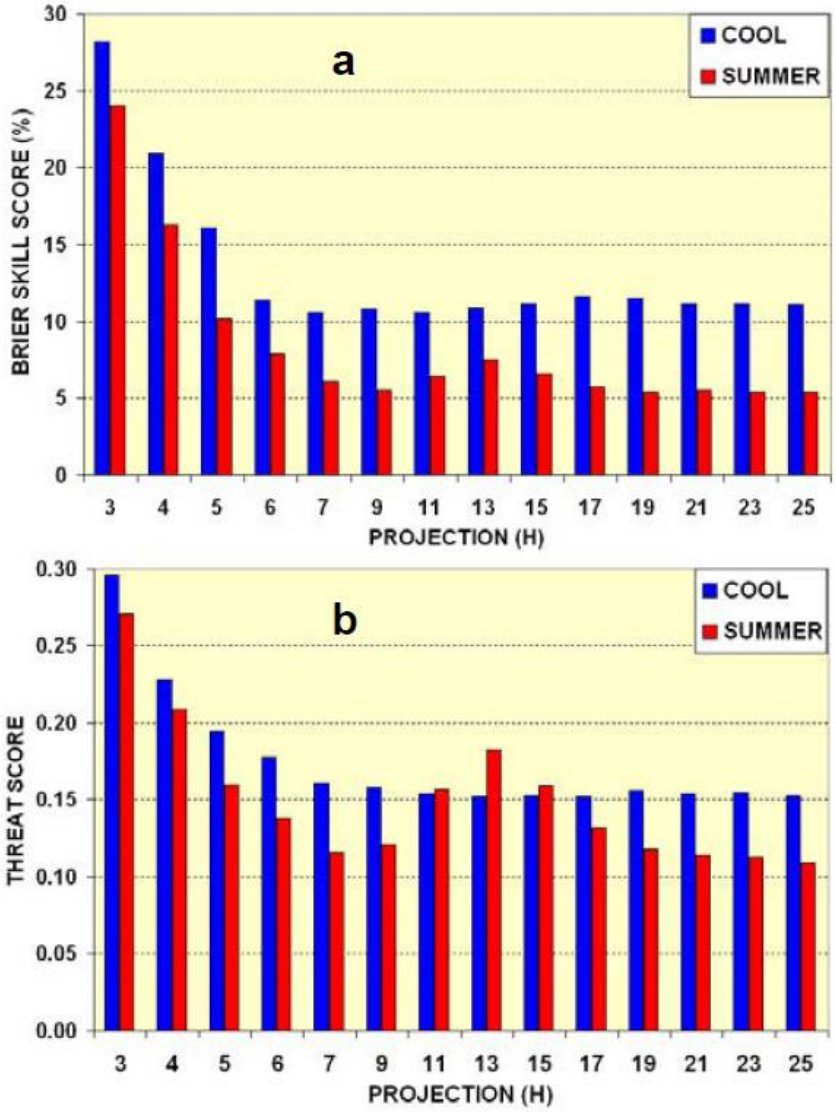

Fig. 2.83. (a) Brier Skill Score for 2-h LAMP convective probability forecasts for the cool season (blue) and warm season (red). (b) Threat Score (same as CSI) for 2-h categorical forecasts of convection. A hit is scored if the observed category matches the forecast category (low, medium or high potential for convection). All forecasts are issued at 0900 GMT. The horizontal axis gives the end time for each 2-h forecast interval. Note that forecasts are issued every hour out to seven hours, and thus the early forecasts overlap. After seven hours, forecasts are issued at 2-h intervals. Reprinted with permission of the American Meteorological Society.

2.3.10 Verification of Multiple Convective Forecast Products

Lack et al. (2011) verified several aviation weather products related to convection using the same observational database, namely the CIWS analysis fields for VIL. These products are:

- The Collaborative Convective Forecast Product (CCFP) described in Section 2.3.4
- A calibrated CCFP. When AWC forecasters issue a CCFP, they estimate areal coverage as sparse, medium, or high. They also express their confidence in the forecast as low or high, high confidence being assumed for medium or high coverage. Table 2.12 illustrates the coverage actually observed for all categories of CCFP forecasts.

Table 2.12

Observed climatological percent coverage in CCFP by forecast type and lead time. CIWS analyses were used to obtain the climatological coverage.

	2-h Lead	4-h Lead	6-h Lead
Sparse Coverage / Low Confidence	4.64	3.93	3.62
Sparse Coverage / High Confidence	9.77	7.75	6.76
Medium or High Coverage	20.38	14.87	13.15

As an example of using this table, if a forecaster predicts sparse coverage ($\leq$ 39% areal coverage) with high confidence (probability $\geq$ 50%) at 4-h lead time, then the calibrated areal coverage of the CCFP polygon is set to 7.75%.

- The LAMP product for probability of convection, as in Fig. 2.81.
- A straight (unaltered) forecast by the High-Resolution Rapid Refresh (HRRR) model
- CoSPA
- Climatology. Specifically, the probability that a thunderstorm will be in progress within a grid box at a given hour of the day during a given season.
- Uniform. The same as climatology except a single number, an average over all grid boxes in the CONUS, is used for each hour of the day and season.

Figure 2.84 from Lack et al. (2011) examines the bias in areal coverage of thunderstorms (VIP $\geq$ 3) as predicted six hours in advance for all the above products except the last two. Note that the vertical scale is logarithmic. A bias of one is perfect. The top panel is valid for the northeast quadrant of the U.S.; the bottom panel applies to the southeast quadrant.

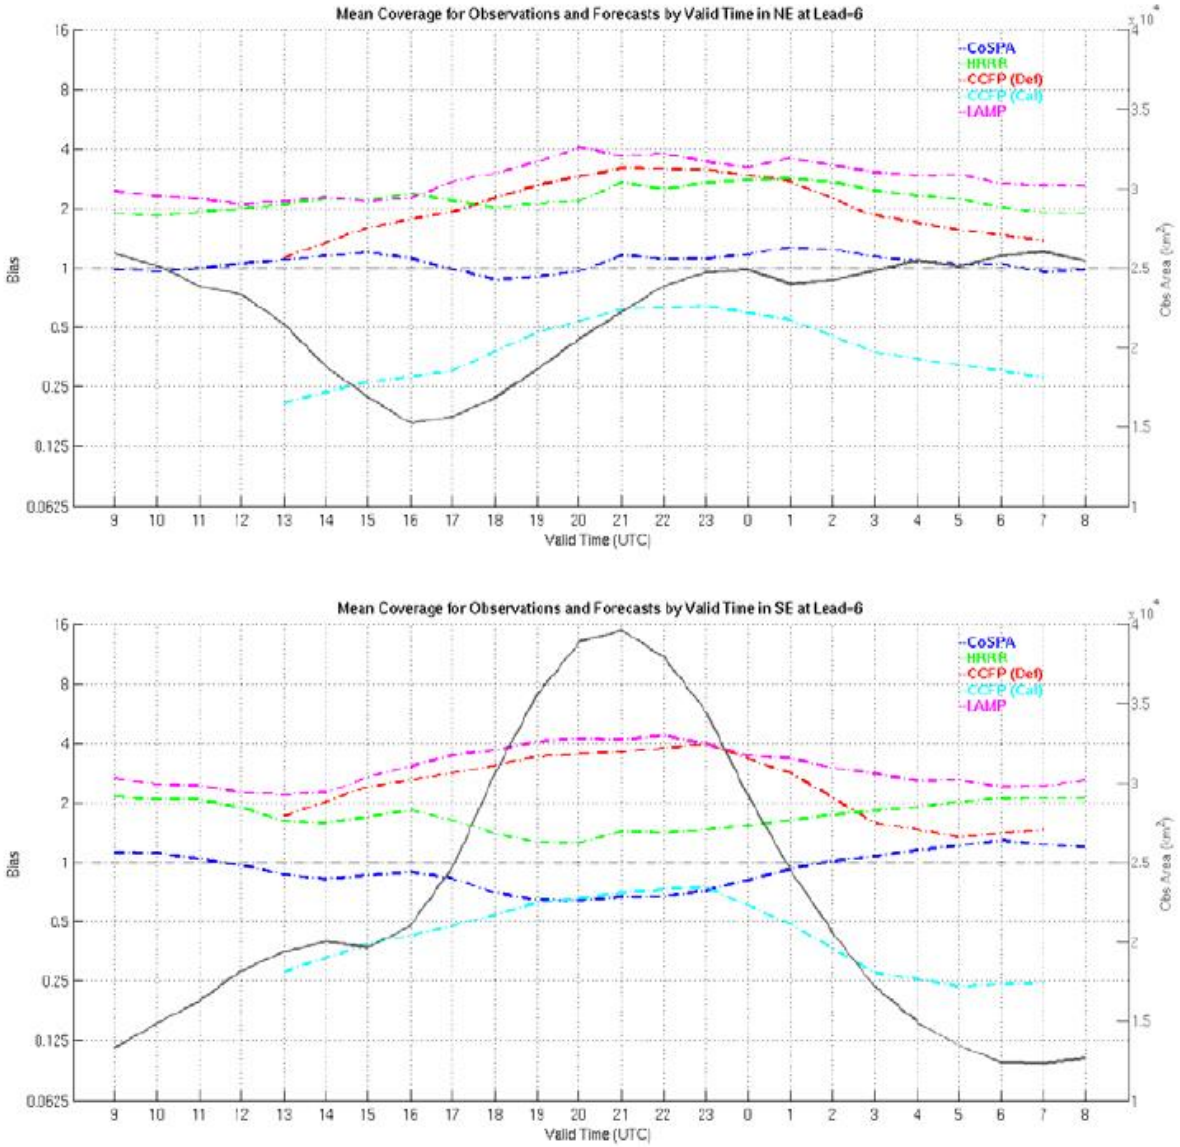

Fig. 2.84. The bias in predicted areal coverage of thunderstorms (VIP ≥ 3) for each hour of the day (GMT) for the northeast (top) and southeast U.S. (bottom) from June through September 2010. The forecast lead time is six hours. Note that the vertical scale for bias is logarithmic. A perfect bias is 1.0. The hours of the day (GMT) are labeled along the horizontal axis. The charts are centered at 2100 GMT, close to the time of maximum convective activity. The colored curves correspond to different forecasts as indicated upper right. The black curve is for observed mean areal coverage in units of 10,000 km^2, given along the vertical axis at right. CCFP (Def) is the unaltered CCFP product (default); CCFP (Cal) is the calibrated version. From Lack et al. (2011).

The observed areal coverage of convection during the warm season (black curve) differs substantially between the northeast and southeast U.S. In the northeast, convection is dynamically forced more often; the morning minimum is the most obvious feature. In the southeast, convection is diurnally forced more often and covers larger areas; the afternoon maximum is pronounced. The CoSPA product has the lowest bias, followed by the HRRR model. The bias for both varies little throughout the day. The unaltered CCFP tends to overforecast convection whereas the calibrated CCFP tends to underforecast it.

Lack et al. (2011) verified forecasts of areal coverage of thunderstorms at many different scales. They covered the CONUS with a mosaic of hexagons (honeycomb pattern) of various sizes. The size of a hexagon is defined as the distance between two parallel edges. Lack et al. (2011) used CIWS analyses (Section 2.3.7) with a threshold of VIL ≥ 3.5 kg m^{-2} to verify all convective products from June through September 2010. One measure of skill is related to the Mincuts calculated for each hexagon. Figure 2.85 is a specific CIWS analysis with hexagons 300 nautical miles (nm) on a side. VIL values ≥ 3.5 are shown. Convective cells with these values are considered impenetrable.

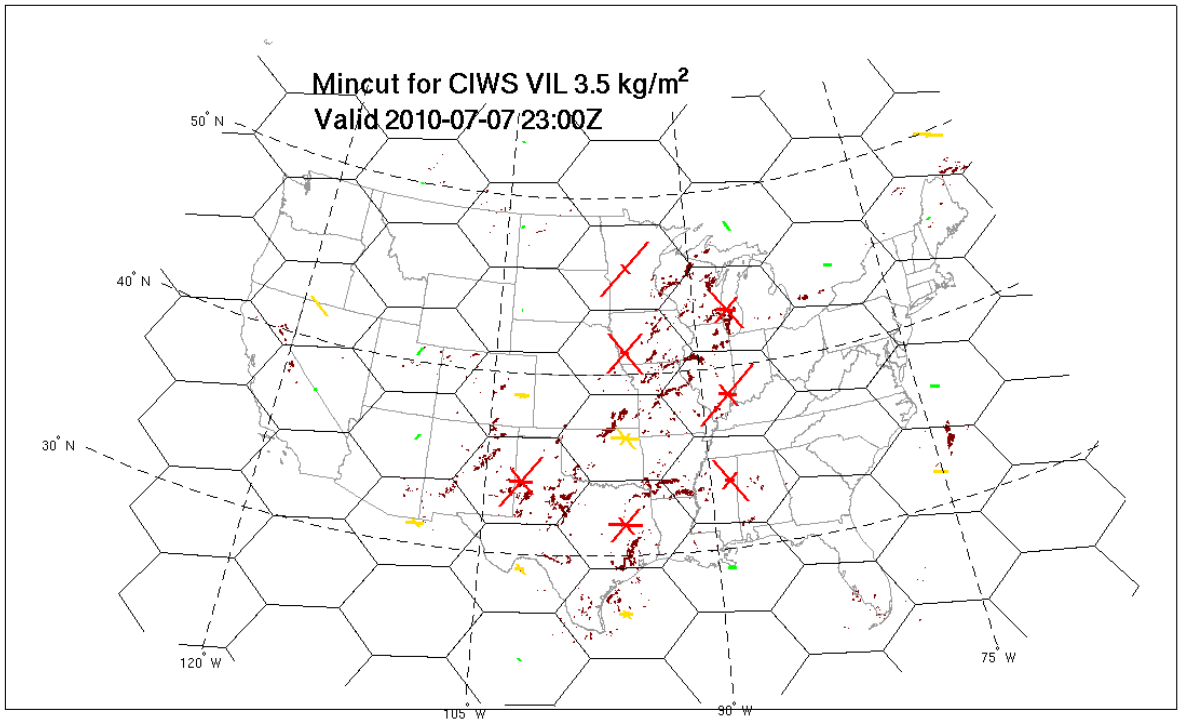

Fig. 2.85. Observed Mincuts for a CIWS analysis, calculated as in Fig. 2.73 (cells impenetrable). The three lines at the center of each hexagon show the relative flow blockages in three directions. Courtesy of Steve Lack, NOAA/ESRL.

Figure 2.86 shows results of Mincut calculations for a 2-h LAMP forecast in hexagons 300 nm wide. This forecast verifies at the same time as the CIWS analysis in Fig. 2.85. The probability of convection from LAMP was used directly in the calculation of the Mincut as illustrated earlier in Fig. 2.76. Of the three Mincut values available in each hexagon (see Fig. 2.74), the largest is used for verification.

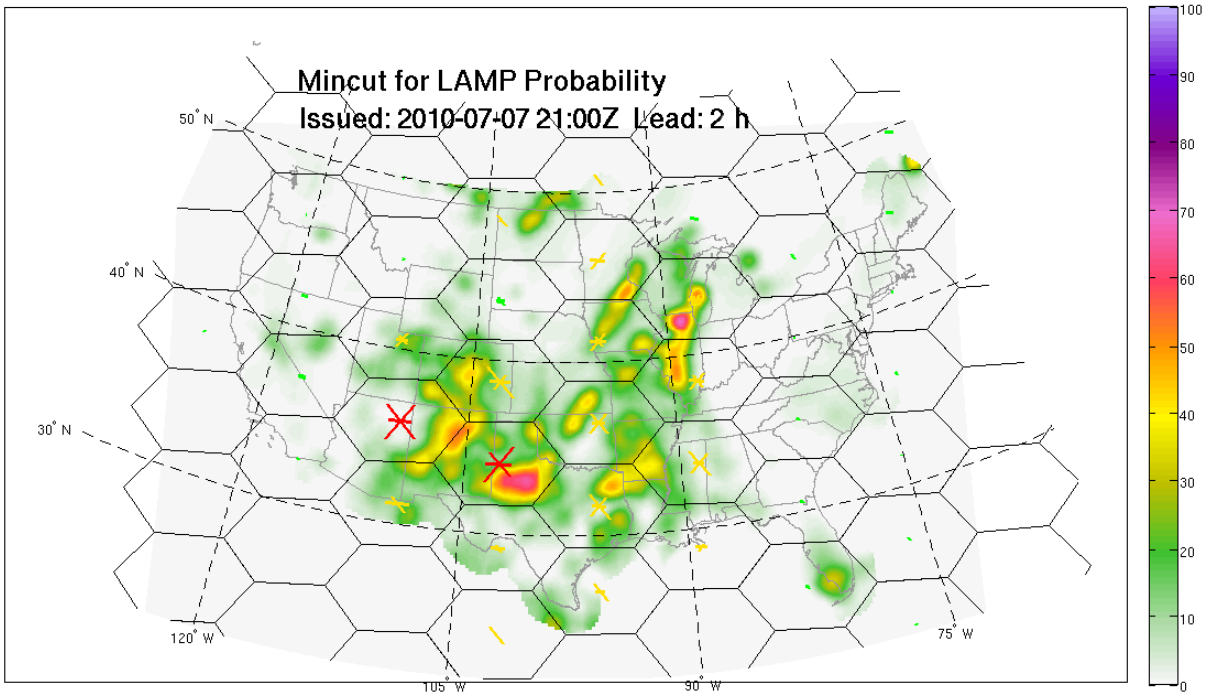

Fig. 2.86. Mincuts illustrated by the three lines at the center of each 300-nm hexagon for a 2-h LAMP probability forecast of convection. See Fig. 2.74 for more details about the Mincut calculation. Courtesy of Steve Lack, NOAA/ESRL.

If one uses the Mincut analysis from CIWS in Fig. 2.85 to verify the LAMP forecast in Fig. 2.86 (red threshold for flow reduction), one finds one hit, six misses, and one false alarm, giving a CSI of 0.125 and a bias of 2/7.

Lack et al. (2011) verified deterministic CoSPA forecasts. They assumed that any cells with predicted VIL values ≥ 3.5 km m^{-2} are impenetrable for the purpose of Mincut calculations. This is consistent with the assumption made for verifying CIWS analyses.

In Fig. 2.87, only two hexagon sizes are considered: 300 nm (5° of latitude or 555 km) (left) and 75 nm (1.25° of latitude or 139 km) (right). This figure gives the CSI scores for five products (color-coded in the figure) from which Mincut thresholds for the three directions of travel were computed within each hexagon (see again Fig. 2.74) in the northeast U.S. The threshold refers to the percent reduction in air traffic ($1 - AFCR$) resulting from convection that constricts or blocks the traffic lanes. Six-hour forecasts issued at 1500 GMT present a severe challenge because mid-morning usually marks a diurnal mimimum in convection, whereas the verification time, 2100 GMT, is in mid-afternoon, when initiation is most likely.

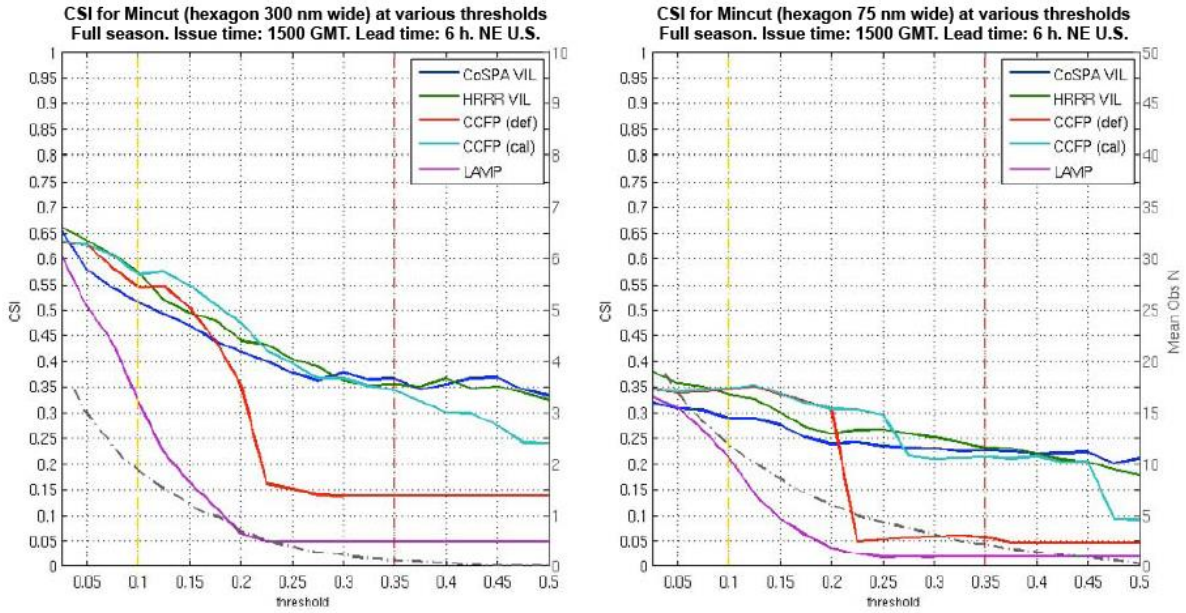

Fig. 2.87. CSI score (vertical axis at left) for different Mincut thresholds (horizontal axis). The threshold refers to the percent decrease in air traffic flow because of convection. A "hit" occurs when the product correctly predicts a blockage at or above the threshold. Five products (see color codes at upper right) are scored here. Both graphs apply for the northeast quadrant of the U.S. The left panel is for verification hexagons 300 nautical miles (nm) wide. The right panel is for verification hexagons 75 nm wide. All forecasts issued at 1500 GMT and valid 6 h later for each day from 1 June through 30 September 2011 are included. The dash-dotted black curve refers to the vertical axis at right. It indicates the mean number of hexagons affected at each Mincut threshold. The vertical yellow line is the threshold for medium-impact events (10% or more reduction in capacity); the vertical maroon line is the threshold for high-impact events (35% or more reduction in capacity). From Lack et al. (2011).

For CSI scores computed at the larger scale (300-nm hexagons), the HRRR model outperforms CoSPA at six hours except at high thresholds. The calibrated CCFP performs surprisingly well

except at the highest thresholds. The unaltered CCFP falters for thresholds ≥ 0.2, but its bias at 6 h (Fig. 2.84) is high, and it is only allowed three categories of coverage at 6 h: 25, 40, and 75%, and so its forecasts are not very sharp. *Sharpness* refers to the tendency to forecast probabilities near zero or one, as opposed to values clustered around the mean. LAMP performs poorly at six hours because its forecasts lack sharpness. It performs better at shorter lead times.

In Fig. 2.87, note the general trend of decreasing accuracy as the threshold increases. The likelihood of a threshold of 0.1 (a reduction in capacity of 10% or more indicated by the vertical yellow line) is much greater than that for threshold of 0.35 (vertical maroon line). Clearly, the less common events are more difficult to forecast correctly. Also note that it is harder to forecast reductions in capacity when the size of the area shrinks (left graph vs. right).

Another measure of skill for evaluating convective products, not yet introduced, is the Fractions Skill Score (FSS); it evaluates the accuracy of forecasts of the percent of area covered by convection, PC_{fcst} . If the percent of area covered by the observations (radar-based estimate) is PC_{obs} , then

$$FSS = 1 - \frac{\frac{1}{N}\sum_{i=1}^{N}(PC_{fcst} - PC_{obs})^2}{\frac{1}{N}\sum_{i=1}^{N}PC_{fcst}^2 + \frac{1}{N}\sum_{i=1}^{N}PC_{obs}^2}$$

A perfect forecast scores FSS = 1; the worst possible forecast scores FSS = 0.

Figure 2.88 shows FSS as a function of the resolution of the honeycomb grid. To repeat, the resolution is defined by the distance h between parallel edges of the hexagon. Each side of the hexagon is $s = h/\sqrt{3}$, and the area of the hexagon is $A = (\sqrt{3}\,h^2)/2$. CCFPs are not issued unless the area of convection exceeds 3000 mi^2 or 2265 nm^2. This area corresponds to a resolution h of 51 nm. Thus, CCFPs for $h < 50$ nm have little meaning. In the southeast U.S., where convection is often fairly random, the calibrated 6-h CCFP wins the day at all resolutions coarser than 45 nm. In the northeast U.S. (not shown), CoSPA beats the calibrated CCFP for the same range of resolutions, but at 6 h CoSPA is heavily influenced by the HRRR forecast, which exhibits skill at 6 h for dynamically forced convection. Returning to Fig. 2.88, the HRRR model has a slight edge over CoSPA at all resolutions. Climatology fares well at coarse resolution, which might be expected if the convection is fairly random. The *uniform* curve is for reference

only. It represents a single area-averaged probability, appropriate for the time of day, applied across the entire domain. Scores improve for all products as the size of the hexagon increases, i.e., as the resolution coarsens.

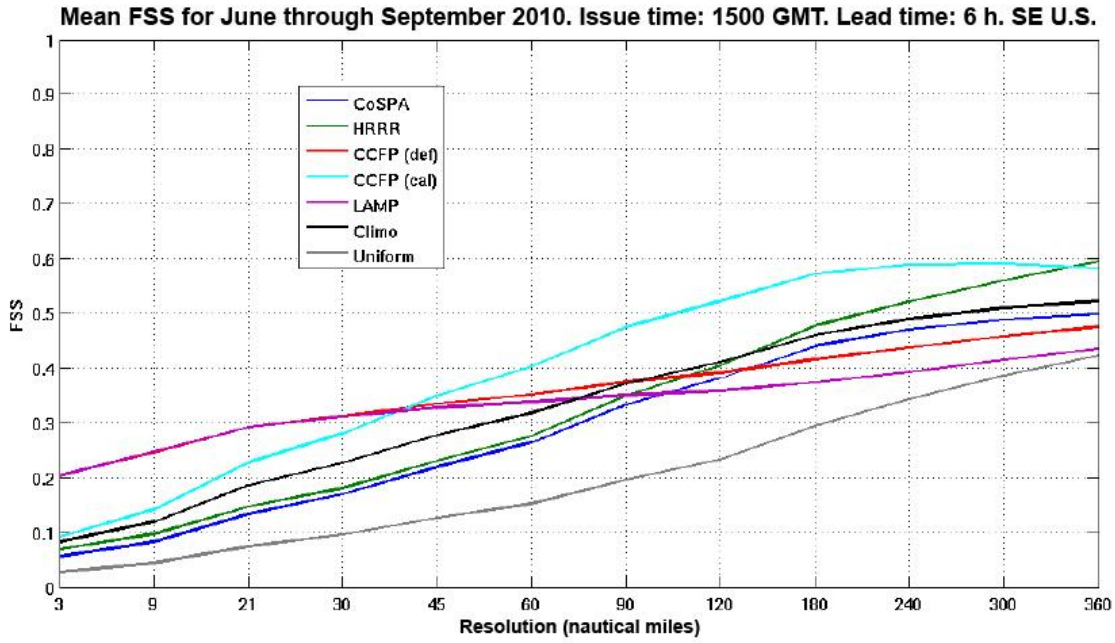

Fig. 2.88. Fractions Skill Score (FSS) for June through September 2010 for five products (color-coded) and two forms of climatology appropriate for the hour of the day: spatially varying (dark black curve) and area-averaged (*Uniform*–gray curve). For six-hour forecasts issued at 1500 GMT in the southeast quadrant of the U.S. The size of the hexagons in which the FSS was calculated is indicated along the abscissa. Note that the scale in resolution is not linear. From Lack et al. (2011).

FSS scores for forecasts with 2-h lead times (Fig. 2.89) are quite different from those shown in Fig. 2.88. CoSPA is the clear winner for all resolutions coarser than 21 km. The calibrated CCFP is considerably better than the uncalibrated version at all meaningful resolutions (≥ 50 nm) and a close second to CoSPA, indicating the value of experienced forecasters in the loop at short lead times. LAMP performs well at the finest resolution but quickly loses its advantage at coarse resolution. FSS scores for climatology and the HRRR model rise in tandem as resolution coarsens. Recall that the HRRR model does not influence the CoSPA product at all during the first two hours.

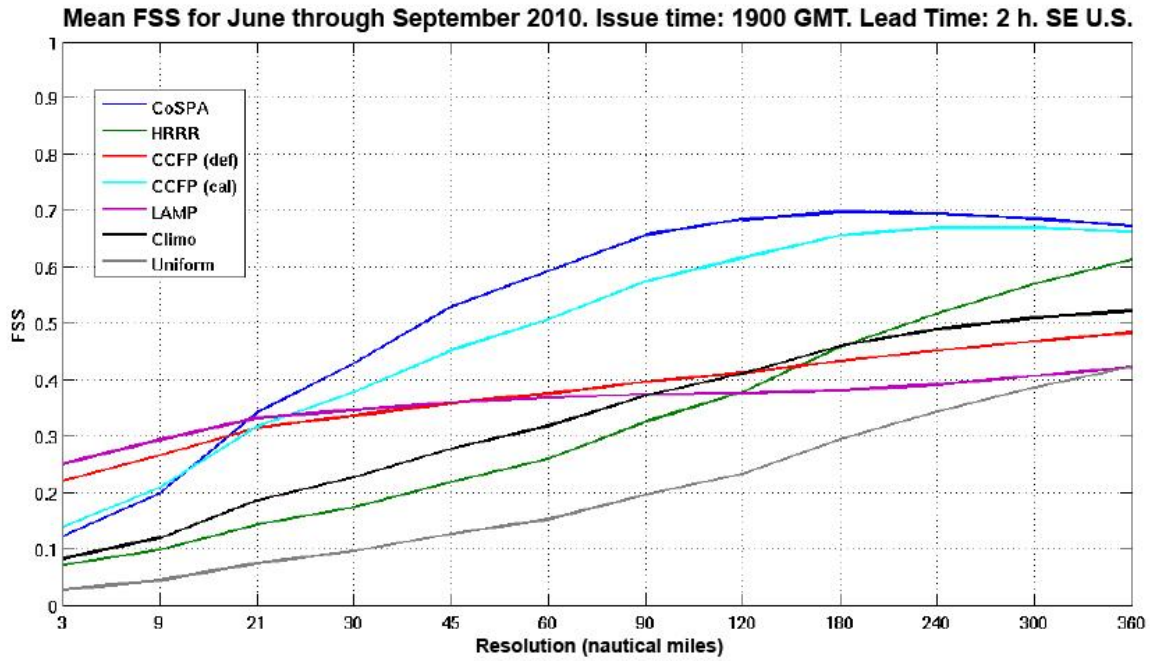

Fig. 2.89. Same as Fig. 2.88 except mean FSS scores for 2-h forecasts issued at 1900 GMT.

Verification scores relevant to flight planning are valuable. For example, if flight controllers have high confidence in 6-h forecasts of Available Flow Capacity Ratio (AFCR) early in the day, when they try to anticipate where problems will arise, they will be better equipped to plan strategically, thus improving efficiency, rather than merely react later after problems develop. Forecasts of AFCR, or any other parameter that assists in decision making, are only as good as the convection forecast on which they depend. To emphasize an oft-made point, the prediction of short-lived convection, from initiation to dissipation, is probably the most challenging weather problem facing aviation today.

2.3.11 Autonowcaster

2.3.11.1 Introduction

The Autonowcaster (ANC) is an expert system that relies on observations, computer models, human input, and *fuzzy logic* (see second paragraph after this one) to produce 0-1 h projections of storm initiation, growth, and decay. According to the World Meteorological Organization (WMO), the word *nowcast* means a forecast including local detail, prepared by any method,

over a period from the present to a few hours ahead. A nowcast includes a detailed description of present weather as a starting point.

The primary emphasis of the ANC has always been on initiation, particularly under conditions of weak synoptic forcing, diurnal heating, and convergence lines that develop in the boundary layer. The ANC had its origins in a paper by Wilson and Mueller (1993), which described results of field experiments involving experienced forecasters. The forecasters attempted to identify features in Doppler radar images, surface weather reports, and cumulus cloud observations that are precursors to thunderstorm development. Forecasters used these features to predict the initiation of convection but were only moderately successful. They were even less successful in predicting the subsequent evolution of convection. The limited success was attributed to 1) limited understanding of the processes that lead to initiation, 2) lack of detailed observations in the boundary layer, and 3) manually intensive procedures imposed on forecasters already pressed for time. Development of the ANC was an attempt to automate many of the labor-intensive tasks while still allowing forecaster intervention.

Whereas CoSPA is being developed strictly for aviation uses with FAA support, the ANC is intended for convective weather forecasting in general and is supported primarily by NWS. The first real-time tests of the ANC forecaster-computer mix occurred during the Sydney, Australia 2000 Olympics Forecast Demonstration Program (Keenan et al. 2003). ANC has since been tested extensively at NWS offices in Melbourne, Florida, Sterling, Virginia, and Dallas-Fort Worth, Texas. It is now undergoing a regional test in the so-called Golden Triangle, formed by Chicago, Atlanta, and New York City, where air traffic is heaviest. It is also used routinely at White Sands Missile Range, New Mexico.

Though the technique was not named, *fuzzy logic* (McNeill and Freiberger, 1993) was used earlier in this survey to describe generation of the Current Icing Potential (CIP) and Forecast Icing Potential (FIP) products in Sections 2.2.7 and 2.2.8. Fuzzy logic is a convenient method for combining various data sources and incorporating conceptual models to produce a likelihood estimate for specific events. The method is called “fuzzy” because it does not produce a “yes-no” answer, but it does provide a logical framework that can incorporate almost any information that decision-makers would normally consider. *Interest maps*, a tool of fuzzy logic, used extensively in Sections 2.2.7 and 2.2.8, are also used in the Autonowcaster.

The most complete description of the Autonowcaster in the refereed literature (Mueller et al. 2003) is summarized here. The information provided may not be completely up to date, but it illustrates the concepts employed and the salient features of the system.

2.3.11.2 Conceptual models and rules of thumb incorporated in the Autonowcaster

Field experience with weakly forced convection has led to a number of predictors of thunderstorm initiation, growth, and decay (Mueller et al. 2003). They are divided into three classes: recent behavior of existing storms, features in the boundary layer, and cloud characteristics. The following bulleted list gives a name for each predictor, a brief description, a weight assigned to this predictor, and the name of the algorithm that develops predictor information from surface, Doppler radar, and satellite observations. The weights for each predictor are not intended to add to one; they merely indicate the relative importance of each predictor when the Autonowcaster considers all of them together. The algorithms themselves will be discussed in the next section.

Recent behavior of existing storms

- Extrapolated reflectivity (dBZ). Indicates the extrapolated position of radar echoes.
  Weight: 0.20 Algorithm: TREC
- Extrapolated reflectivity (dBZ) with stratiform regions removed. Highlights convective regions. This predictor receives the largest weight because behavior of an existing storm in the next hour is much easier to predict than the development of a new one.
  Weight: 0.40 Algorithm: Stratiform Filter; TITAN (vectors)
- Storm area (km^2) . Provides the area enclosed by the 35-dBZ contour; storm size is used as an indicator of storm longevity.
  Weight: 0.20 Algorithm: TITAN
- Negative and positive growth rates ($km^2\ h^{-1}$). Used alone, recent changes in the size of a convective cell or cluster is not a very good predictor of future growth or decay. However, when used in combination with boundary layer properties in the near vicinity, this information can be valuable. For example, a decaying thunderstorm crossing a dry line may suddenly experience inflow with higher dewpoints, thus promoting reintensification.
  Weight: 0.15 Algorithm: TITAN
- Precipitation accumulation (mm). Estimated from WSR-88D radar data based upon standard relationships between radar reflectivity and rainfall. Convective storms cool the boundary layer through the evaporation of rain. Areas covered by rain-cooled air are more stable and thus unlikely to experience further convection in synoptically quiet situations.
- Weight: 0.15 Algorithm: NEXRAD-based, reflectivity-rainfall relationship (not discussed further)

Boundary layer characteristics

- Boundary location and movement (m s^{-1}). Indicates where convective initiation is more likely because of a convergence line.
  Weight: 0.20 Algorithms: COLIDE for detection and extrapolation; VDRAS for dectection
- Boundary collisions and storm-boundary collisions. If two moving boundaries (e.g. gust fronts from separate thunderstorms) collide, vertical velocity is enhanced. There is a better chance that air parcels will be lifted to their level of free convection. If a storm moves over a boundary, the influx of moist air into that storm strengthens, at least temporarily.
  Weight 0.25 Algorithms: COLIDE for detection and extrapolation; TITAN for storm motion; rawinsonde for steering level winds
- Boundary-relative steering flow (m s^{-1}). A measure of the differential motion between the surface-based convergence line and cumuliform clouds above, which are steered by winds in the cloud layer, often the mid-tropospheric winds. Example: if a gust front outruns its parent thunderstorm, the latter becomes separated from its warm, moist, low-level inflow and will weaken. In this case, the boundary-relative steering flow is large. If the boundary-relative steering flow is small (< 6 m s^{-1}), a convective cloud has a better chance of remaining close to the low-level boundary and will benefit from a stronger and more sustained updraft.
  Weight: 0.20 Algorithm: COLIDE in conjunction with raob or profiler wind data
- Boundary-relative low-level shear [m s^{-1} (2 km)$^{-1}$]. The boundary in this case refers to a pool of cool air, for example, a sea-breeze front or a thunderstorm outflow. The boundary may be moving or stationary, but the shear is relative to any motion of the boundary. It is measured in the first 2 km above ground on the warm side of the boundary. Figure 2.90 from Rotunno et al. (1988) illustrates why this predictor is important.

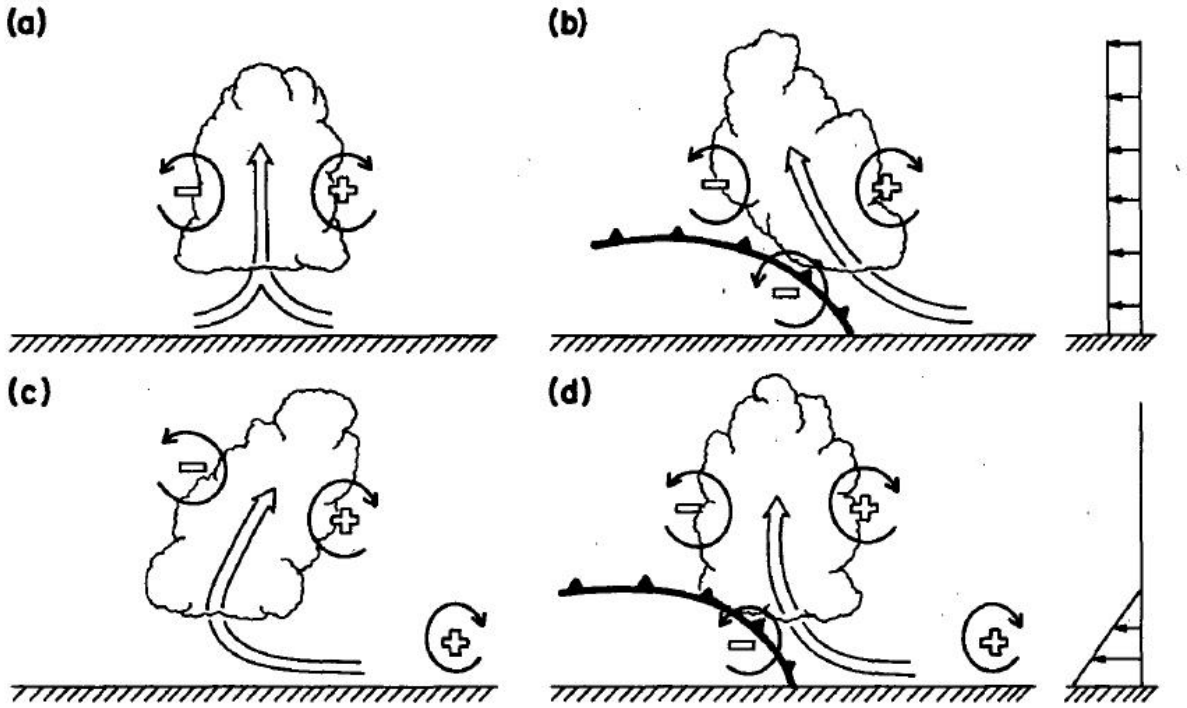

Fig. 2.90. Updrafts (wide arrows) with and without environmental shear (indicated by the stack of vectors, far right) and with and without a low-level boundary (cold front symbols). Plus and minus signs give the sense of rotation about horizontal vortex tubes. The text explains these configurations. From Rotunno et al. (1988). Reprinted with permission of the AMS.

A cumulus cloud grows in an unsheared environment in Fig. 2.90a. There is no precipitation yet. The updraft stands erect. The horizontal shear of the vertical wind across the cloud generates horizontal vortex tubes on either side of the cloud (the tubes are oriented into and out of the page), with the sense of rotation indicated by the plus and minus signs. If a cold pool noses under the left side of the cloud as in Fig. 2.90b, with the strongest breezes at the surface, decreasing upward, this generates an additional vortex with a negative sense of rotation. The two negatives on the left side of the cloud cause the updraft to lean to the left, over the cold pool. The bottom panels indicate changes that occur if there is low-level shear. In Figs. 2.90c and d, low-level environmental shear (independent of the cold pool) induces horizontal vorticity with a positive sense of rotation. In (c), the two positives on the right cause the updraft to lean to the right. In (d), with the addition of a cold pool, there is a fourth source of horizontal vorticity (minus sign). If the opposing effects of (b) and (c) balance each other, there is a chance for an erect updraft. For this to happen, the boundary-*relative* low-level shear should be small.

Weight: 0.20 Algorithms: VDRAS provides the boundary layer wind field in the lowest 2 km; COLIDE provides additional information on locations of convergence zones.

- Maximum updraft (m s^{-1}). The maximum vertical velocity associated with the boundary.
  Weight: 0.20 Algorithm: VDRAS

Cloud characteristics

- Radar-observed towering cumulus. Experience indicates that reflectivity of cells can increase from 10 to 40 dBZ in 15 min or less. Growth detected at early stages has short-term predictive value. First, the maximum reflectivity in a column between 3 and 6 km AGL is determined. Next, the field is filtered using the stratiform-convective partitioner of Steiner et al. (1995), described briefly in Section 2.3.6. After the filtering, only convective cells remain. Finally, the field is advected with steering-level winds.
  Weight: 0.25 Algorithm: just described (components: stratiform filter, TITAN and TITAN vectors)
- Satellite cloud type. Identifies cloud type on the basis of GOES visible and infrared images. High albedo (visible wavelengths) and very sharp gradients in infrared brightness temperatures indicate cumuliform clouds. Areas where cumulus clouds mark organized low-level convergence (boundaries) are favored for development. Organized low-level convergence has a variety of causes: density currents (thunderstorm outflow), cold fronts (a large-scale density current), topographically induced circulations (mountain-valley breeze), and circulations induced by differential heating (boundary between wet and dry ground, land and ocean, or clear and cloudy sky).
  Weight: 0.15 Algorithm: Naval Research Laboratory's cloud classification algorithm using GOES image processing
- Cloud-top cooling (^{o}C / 15 min). Gives cooling rate of cloud tops. Assumes ability to track same cloud on successive images, 15 min apart. Cloud-top temperature is closely related to IR brightness temperature. If tops of cumulus clouds are cooling rapidly with a boundary in the vicinity nearby, growth is indicated.
  Weight: 0.25 Algorithm: Time rate of change of infrared temperature from GOES

- The GOES infrared imager passively senses the temperature of the target, whether the top of a cloud or the ground. When the brightness temperature in the infrared image is close to the observed surface temperature, this indicates clear skies and too little moisture in the boundary layer to produce any clouds—a negative indicator of thunderstorm development.
  Weight: 0.15 Algorithm: GOES image processing

- Terrain elevation (m). Can be used to increase the likelihood of thunderstorms over ridges and mountains and decrease it in basins.
  Weight: 0.15 Algorithm: None needed

The predictors listed above were those used in the early years of Autonowcaster testing. Since then, another category has been added (Roberts et al. 2011, Table 1), with the predictors derived from output of the RUC 13-km NWP model. Roberts and coauthors do not give weights for these predictors:

Characterization of the environment

- Frontal forcing. A combination of equivalent potential temperature (θ_e), vorticity and convergence
- Maximum Convective Available Potential Energy (CAPE) in the 900–700 mb layer
- Average Convective INhibition (CIN) in the 900–975 mb layer. CIN is a measure of the amount of energy required to lift a parcel to its level of free convection.
- Average relative humidity (RH) in the 875–725 mb layer
- 700-mb vertical velocity
- Layered instability (CAPE, CIN, RH, and vertical shear combined)

Experience has indicated that forecaster intervention can improve the final one-hour projection. Using an AWIPS workstation, installed at NWS offices throughout the country, forecasters can enter boundaries undetected by the Autonowcaster, but sometimes obvious to the human eye. Forecasters can change the interest maps belonging to individual predictors (samples to be given in section 2.3.11.4), and the relative weights given to these predictors. Finally, forecasters can reconfigure the list of predictors so that it is more appropriate for the problem of the day (Nelson et al. 2008). For example, separate sets of predictors have been developed for seven different scenarios:

- Mixed. Intended for events where the synoptic situation is ill-defined and a variety of initiation mechanisms are considered possible. Human-specified boundaries are heavily weighted. Most predictors listed earlier are included.
- Cold Front. In addition to the predictors above, this scenario uses precipitation accumulated during the past three hours to suppress initiation forecasts. This counteracts the use of surface-based stability indices from models, which may not be accounting for cold pools generated very recently. Boundaries entered by forecasters are heavily weighted.

- Dryline. A new predictor in this scenario is the gradient of equivalent potential temperature, intended to boost interest along the dryline. Boundaries entered by forecasters are heavily weighted.
- Stationary Front. Similar to the mixed regime, but with modifications to account for the possibility that elevated convection may occur on the cold side of the front. Layer-dependent stability predictors are included in this scenario. Forecaster-entered boundaries receive less weight because convection is not always anchored to the position of the surface front but may occur some distance away where the frontal surface is most strongly sloped.
- Pulse Storms. Often called "popcorn convection," these storms occur when vertical shear is weak. This scenario uses the fewest predictors and increases weights for satellite-based predictor fields and the objectively analyzed convergence field. The weight given to forecaster-entered boundaries is reduced because of the likelihood that interacting outflow boundaries may generate numerous, small, short-lived thunderstorms. This regime is not used if high clouds prevent the satellite from viewing cumulus clouds below them.
- No Storms. Used when convective initiation is not expected. No forecaster input is expected.
- Advecting Mesoscale Convective System (MCS). Relies heavily on predictors associated with extrapolation and intensity trends. Initiation forecasts primarily depend upon forecaster input of gusts fronts or other boundaries.

The extent to which forecaster intervention improves forecasts, and the nature of the intervention are still being examined by the NWS's Meteorological Development Laboratory (MDL). The most firm conclusion so far is that subjective insertion of boundaries is valuable.

2.3.11.3 Major algorithms employed by the Autonowcaster

This section describes the four major algorithms developed to support the Autonowcaster:

- TITAN: Thunderstorm Identification, Tracking, Analysis, and Nowcasting (Dixon and Wiener 1993).
- COLIDE: surface COnvergence LIne Detection and Extrapolation algorithm (no published reference)
- TREC: Tracking clear-air Radar Echoes by Correlation (Tuttle and Foote 1990)
- VDRAS: Variational Doppler Radar Analysis System (Sun and Crook 2001)

- NCAR has written algorithms that process GOES satellite data to infer cloud characteristics and retrieve specific predictor features. The RUC mesoscale model also supplies predictor values,

as indicated in the previous section. In the interest of brevity, these two contributions to the Autonowcaster will not be discussed here.

TITAN

This algorithm is used extensively in generation of the National Convective Weather Detection and Forecast (NCWD and NCWF) products. It is described in Appendix B (B2.3.6). In the Autonowcaster, it is used primarily to detect and track convective cells. It tracks the three-dimensional centroid, recording size, position, and intensity of cells at frequent intervals. It is thus a good vehicle for development of convective climatologies stratified by month, time of day, region, and size. Such climatologies may themselves be used as interest fields in the Autonowcaster. If TITAN identifies a cell near a boundary, there is increased likelihood of cell intensification or development of a new cell nearby.

COLIDE

This algorithm processes reflectivity and radial velocity data from a single WSR-88D radar. COLIDE looks for convergence features in the boundary layer. It focuses on scans at the lowest elevation angles and, if using data collected at higher tilts, it imposes a range limitation beyond which radar data are not used because the radar beam extends too high above the boundary layer. COLIDE applies configurable templates (filters) to the reflectivity data to find features of interest in the boundary layer, for example, long, narrow lines of enhanced (but still weak) reflectivity that mark convergence lines. Insects captured in the convergence zone are responsible for the higher reflectivity. The templates are similar to those used in CoSPA and belong to a large class of filters used in image processing.

COLIDE also examines radial velocity data for convergence/divergence signatures, i.e. a change in the sign of radial velocity along a given radial. If adjacent radials show the same convergence signature, this may confirm a boundary detected in the reflectivity data. Boundaries oriented along a radial cannot be detected. COLIDE tracks any features it finds, provided the same feature shows up on successive volume scans.

TREC

TREC tracks storms and envelopes (clusters) of storms and derives their vector motion. It was originally developed to track clear-air echoes in the boundary layer, but that function is now provided by VDRAS, to be discussed next. TREC employs cross correlation (a form of pattern matching) to assign motion vectors to storms. In this, it differs from TITAN, which tracks storm centroids and assumes that echo area changes little between radar scans and motion is regular.

Figure 2.91 illustrates the cross correlation procedure. At Time 1 (left) a box is drawn, large enough to include one or more echoes or a cluster of echoes. The object is to find an echo pattern at Time 2 (right) in an identically sized box that most closely matches that in the box at Time 1. The search for the best match proceeds within a radius large enough to account for echo motion between radar scans but small enough to avoid searching in areas where a match would never be found. To conduct the search, the pattern in the box at left is positioned at the center of the circle at right, and the correlation between it and the underlying pattern for Time 2 is computed. The correlation is recorded at the position of the center of the box. The box with the pattern from Time 1 is then moved up and down, right and left, over the radar image for Time 2 in small increments, a kilometer or less, each time recording the correlation between the two patterns. Eventually an entire field of correlations is generated within the search area. The vector connecting the center of the circle with the center of the box with the highest correlation is the estimate of echo motion.

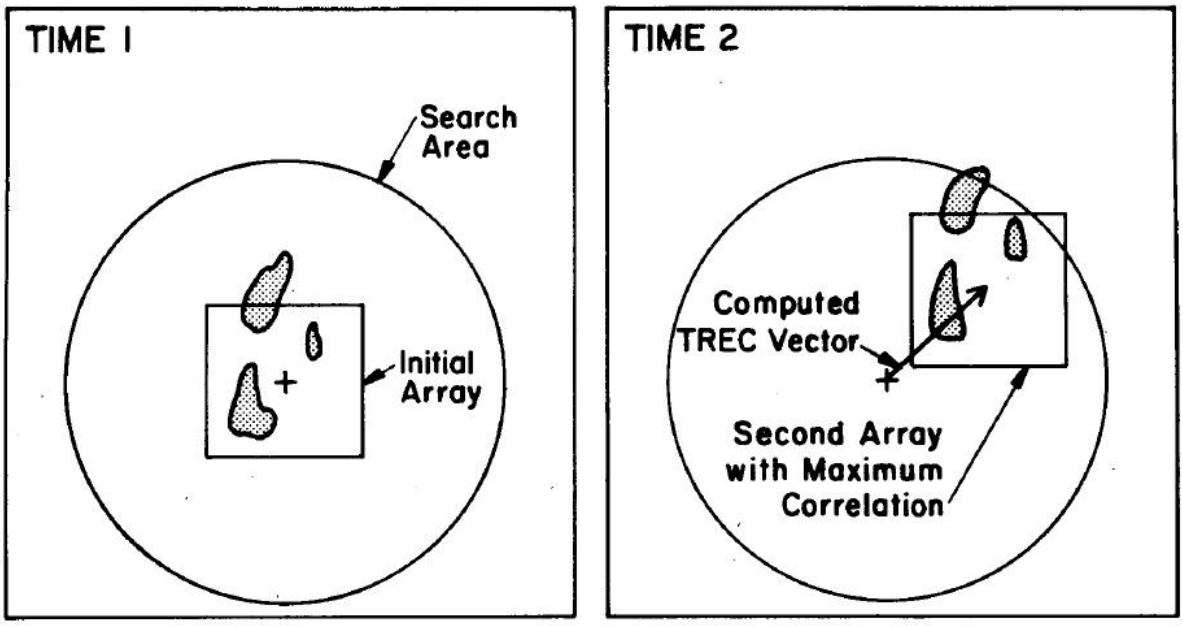

Fig. 2.91. Illustrating how cross-correlation can be used to compute cell motion vectors. At left, a box containing two echoes and part of a third is defined at Time 1. The object is to find the most similar appearing box at Time 2 lying within a specified search radius (the circle) and see how far its center is from the center of the original box at Time 1. From Tuttle and Foote (1990). Reprinted with permission of the American Meteorological Society.

With this procedure, TREC generates cell or envelope motion vectors over the area covered by the radar scan. The vectors are edited to eliminate artifacts caused by ground clutter and second-trip echoes.

VDRAS

Of the four algorithms contributing to the Autonowcaster, *VDRAS* is, by far, the most complex, computationally demanding, and difficult to describe. The idea for VDRAS originated with a paper by Sun et al. (1991). She and her coauthors presented a method whereby reflectivity and radial velocity data from a single Doppler radar could be assimilated into a simplified numerical model in order to infer three-dimensional wind and temperature fields in the boundary layer. They used simulated data in the early tests. Sun and Crook (1994) successfully incorporated reflectivity and radial velocity observations from multiple radar scans into the numerical model. Their procedure became known as the Variational Doppler Radar Analysis System (VDRAS). Through the late 1990s, VDRAS was used primarily as a research tool, but Sun and Crook (2001) converted it to a real-time data assimilation system using WSR-88D radar data.

VDRAS is a form of four-dimensional variational data assimilation (4DVAR). The latter attempts to find a particular evolution of model states that best fits a time sequence of observations. When this goal is achieved, an atmospheric state is defined on the model grid that is not only consistent with observations but also with the model equations. This is a highly desirable state because it is dynamically balanced and will not generate any spurious waves in the early hours of the forecast. A highly developed form of 4DVAR is used in global predictions of the European Centre for Medium Range Weather Forecasts, but the computational demands are so huge that a simplified version is essential for VDRAS.

Two numerical models are available for VDRAS, depending upon the application and the available computing power. The first of these is appropriate for a dry boundary layer with a specified lapse rate. The earth's rotation and mean horizontal temperature gradients are neglected because the time interval for assimilation is short (nominally 12 min), the area covered by a single Doppler radar is small, and emphasis is on weakly forced situations. The prognostic variables are u, v, w (the eastward, northward, and vertical components of the wind, respectively) and potential temperature. The three-dimensional divergence is zero, which implies that horizontal divergence results in vertical motion. Because VDRAS assimilates reflectivity data, the model needs a governing equation for reflectivity. Since the model is dry, reflectivity from hydrometeors is not considered, only clear-air reflectivity. Reflectivity is considered a passive tracer, carried by the wind. It has no sources or sinks. Readers interested

in a more detailed description of the dry model equations applied to the boundary layer should consult Sun et al. (1991).

The second numerical model available for VDRAS was developed by Sun and Crook (1997). It differs in several ways from the boundary-layer model just described. It includes the entire troposphere. It is not dry but allows several moist processes: condensation and evaporation of cloud water, evaporation of raindrops in subsaturated air, conversion of cloud water to rainwater, raindrop growth by coalescence, and sedimentation of rain, which is defined by a Marshall-Palmer drop-size distribution and a fall velocity. This "wet" model has two additional prognostic variables besides (u,v,w) and potential temperature: rainwater and total water (the sum of cloud water and rainwater). The model produces only "warm" rain, that is, there are no ice processes, the justification being that the emphasis is on accurate boundary layer wind predictions, and the boundary layer is above freezing during the summer convective season.

Because the second model generates precipitation, reflectivity is no longer conserved. In fact, the moist model assimilates rainwater "observations" calculated from the reflectivity using either

$$Z = 2.04 \times 10^4 (\rho q_r)^{1.75}$$

if Z is expressed in units of mm^6 m^{-3} or

$$Z = 43.1 + 17.5 \log(\rho q_r)$$

if Z is expressed in units of dBZ. ρ is the air density, and q_r is the rainwater mixing ratio. The appropriate expression is back-solved for q_r, and the observed reflectivity is plugged in. Whereas clear-air reflectivity was directly assimilated in the boundary-layer model and treated as a passive tracer, here reflectivity is a proxy observation for rainwater mixing ratio in the model. Clear-air reflectivity plays no role in the moist version of VDRAS.

Sun and Crook (2001) describe the assimilation of radar data. First, quality control procedures remove second-trip echoes, ground clutter, and artifacts of anomalous propagation. Next, the plan position indicator (PPI) radar data are interpolated from a polar grid (their native coordinates) to the 3-km model Cartesian grid, while remaining on the original constant elevation (PPI) surface. Justification for this action is that the horizontal resolution of the radar data is several times finer than that of the model grid. Henceforth, the interpolated radar data on constant elevation surfaces are considered the "observations."

Any 4DVAR method like VDRAS starts with a model forecast from a previous time, called the *first guess* or *background*, and a set of observations. It attempts to find a corrected model state that is close to both the observations and the background. How close to each is determined by

the accuracy of each information source. For example, if the observations are trusted more than the model, the observations receive greater weight.

A *cost function* is a mathematical expression, usually the sum of three terms. The first term measures the misfit between the actual observation and an estimate of the observed value obtained from the corrected model state (the one to be determined). This term is calculated at the observation locations. The second term measures the misfit between the first guess and the corrected model state. This term is calculated on the model grid. The first guess can be the previous analysis or a short forecast (typically 5 min) from the analysis. The third term usually involves constraints on the solution (the corrected model state). In the case of VDRAS, the constraint is for a smooth solution. The smoother the solution, the smaller this term will be. The desired model state minimizes the cost function, that is, it minimizes the misfits and yields smooth model fields. A number of numerical techniques have been invented over the years to find a solution. The solution is not trivial because the cost function involves all the model variables stored on the grid and all the interpolated radar observations on constant elevation surfaces. The solution is iterative, in that each new trial solution, involving all the corrected model variables, must be tested to see if it reduces the cost function from the value calculated with the previous solution. In VDRAS, the cost function usually levels out near a minimum value after 30 to 50 iterations.

One other aspect of the cost function should be mentioned. The first term involves differences between each observed value and a corresponding value estimated from the corrected model. An observed reflectivity can be matched with a model reflectivity following interpolation of reflectivity values from the model grid to the observation location. (For the wet model in VDRAS, the "observed" rainwater mixing ratio is matched with the corresponding model value.) But the observed radial velocity does not match the model prognostic variables u, v, and w. Thus, the component of the three-dimensional velocity along the appropriate radial in the PPI must be computed prior to interpolation from the model grid to the location of the observation. Transformation of variables and spatial interpolation are standard procedures in all but the simplest variational analysis schemes. For more details and the equations behind this generic description, see Sun and Crook (1997 and 2001).

4DVAR (and, by implication VDRAS) involves a continuing cycle of 1) observation ingest and quality control, 2) assimilation of the observations into a model, which results in an analysis, the initial state for the next forecast, then 3) a forecast, which evolves the atmospheric state until it is time to ingest more observations. The VDRAS cycle is illustrated in Fig. 2.92. Within

each 12-min assimilation window, VDRAS uses reflectivity and radial velocity data from three volume scans to find the optimal model trajectory by minimizing the cost function. The end of this trajectory becomes the analysis, from which many of the predictors discussed above are calculated. It also becomes the starting point for the next forecast. Sometimes the cycling must start cold, and often the radar data do not fully cover the model domain because the signal-to-noise ratio is too low. In these cases (usually), an analysis of surface mesonet observations (Barnes, 1964) fills gaps at the surface, and a Velocity-Azimuth Display (VAD) analysis (Lhermitte and Atlas, 1961), using different elevation angles at constant range, provides an area-wide average of the wind aloft.

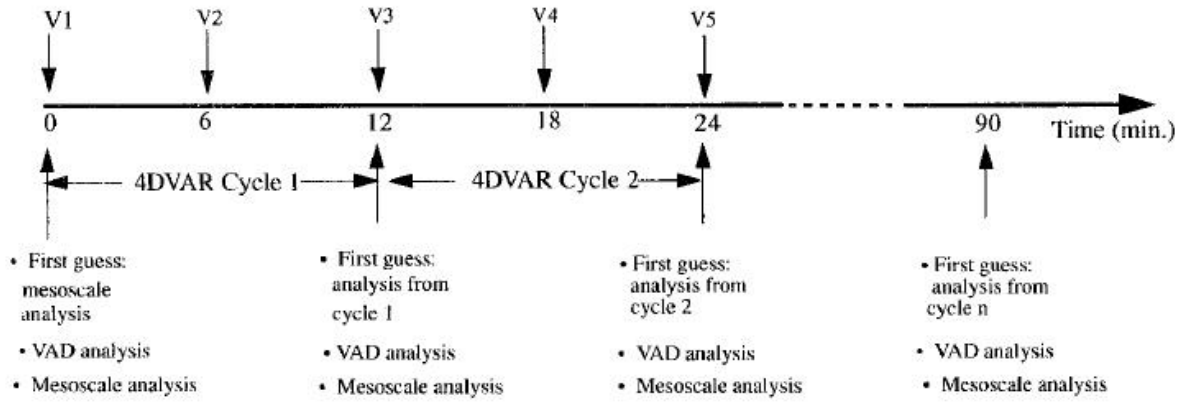

Fig. 2.92. Illustrating the VDRAS 4DVAR assimilation cycle. The time line, labeled in minutes, runs from right to left. V1, V2,…. illustrate the ingest of data from successive volume scans of the WSR-88D radar. The first guess for the first cycle is an analysis based on a Velocity-Azimuth Display (VAD) and surface mesonet reports (see text for details). Subsequent first guesses come from the previous analysis in this particular example. Data from three successive volume scans are used to compute each analysis, which is valid at the end of each 12-min assimilation period. From Sun and Crook (2001). Reprinted with permission of the American Meteorological Society.

2.3.11.4 The fuzzy logic algorithm and nowcast out to one hour

Fuzzy logic uses *membership functions* to translate data fields into interest fields. The user supplies the range of data values for the membership functions and specifies the weight (relative importance) of each predictor field in the Autonowcaster algorithm. Figure 2.93 gives a sample of membership functions. Each function is associated with a predictor. The interest value ranges from -1 to +1. A negative interest indicates that the predictor inhibits convective initiation; a positive interest indicates that the predictor promotes initiation. For example, in

Fig. 2.93g, if an echo is growing at a rate between 0 and 100 $km^2\ h^{-1}$ the interest value increases sharply with growth rate. All growth rates greater than 100 $km^2\ h^{-1}$ receive an interest value of one. In Fig. 2.93h, as the amount of precipitation in the past hour increases from zero to one millimeter, the interest rate falls from zero to minus one, indicating increasing suppression of convective initiation. The interest values are subjective, but they reflect forecaster experience.

Readers may be puzzled by the two binary predictors in Fig. 2.93, (d) and (k). The steep ramp for values less than one just ensures that when the predictor value is zero, the interest is zero, and when the predictor value is one, the only other choice, the interest is one. The details of the ramp are irrelevant, as long as the value lies between zero and one.

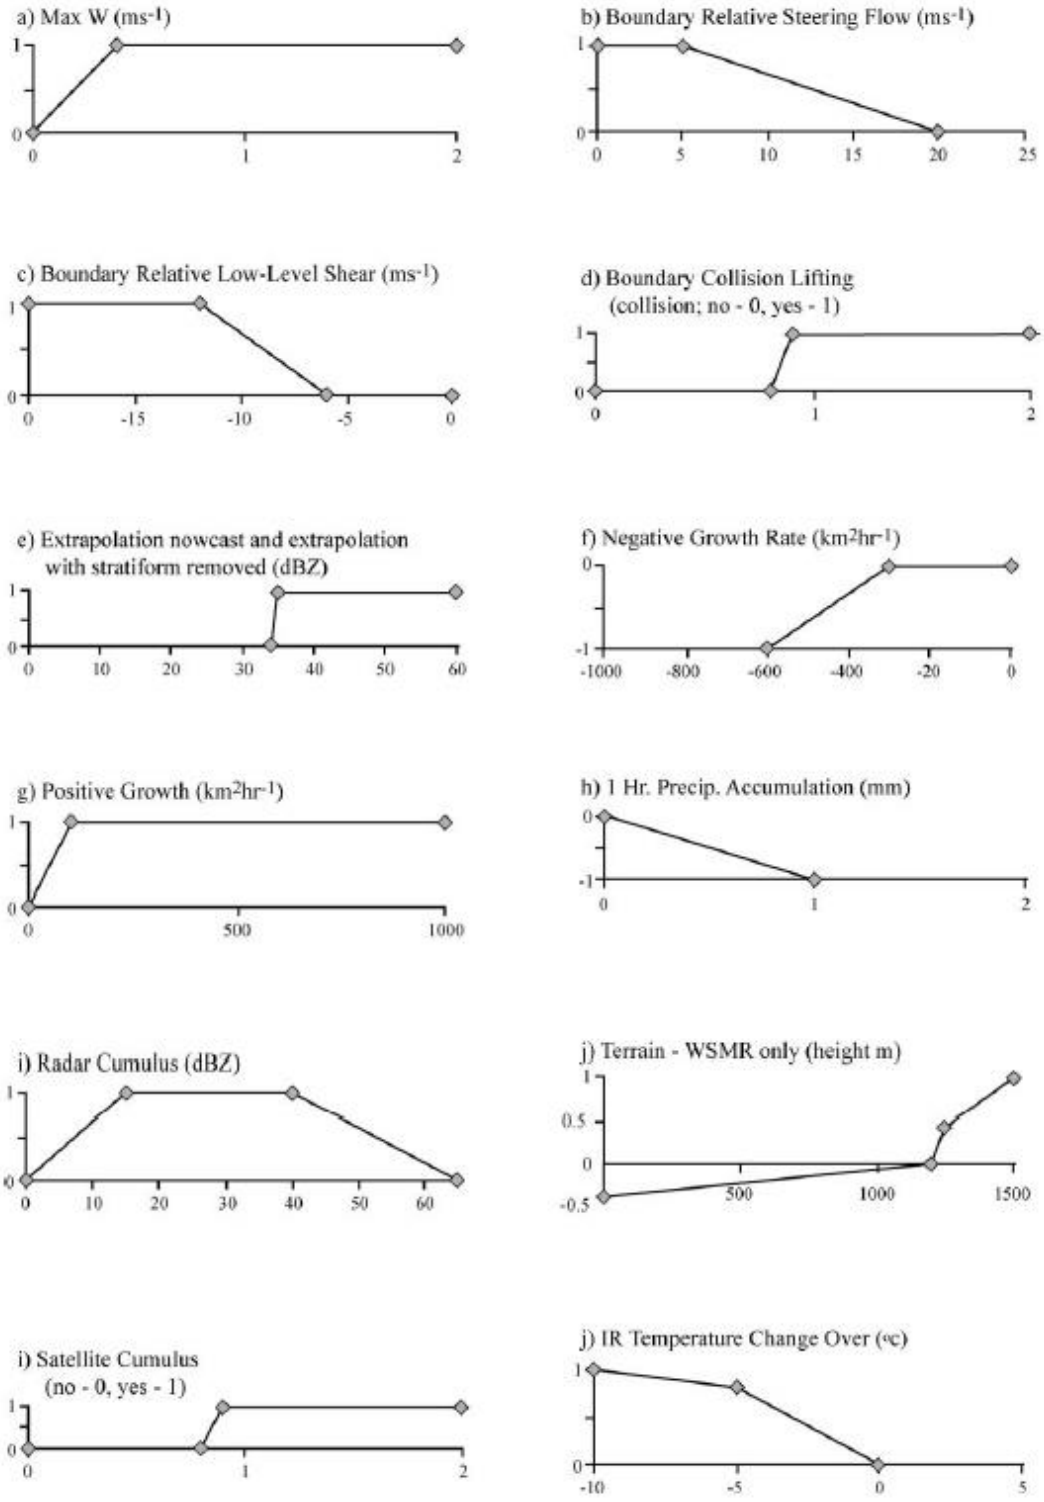

Fig. 2.93. A sample of interest maps used in the Autonowcaster. The vertical axis gives the interest value. The horizontal axis gives the value of the predictor. From Mueller et al. (2003). Reprinted with permission of the AMS.

Once all the predictors are computed and the interest values (all dimensionless) are obtained from the functions shown in Fig. 2.93, the combined likelihood is computed by applying the predictor weights (discussed earlier) to the interest values and then summing. The combined likelihood field is filtered, smoothed, and thresholded to produce the final nowcast, for either 30 or 60 min. The likelihood field is just that: higher numbers imply a greater likelihood of initiation. The combined likelihood can be greater than one and should not be mistaken for a probability.

2.3.11.5 Sample Products

The NWS office at Melbourne, Florida, has been testing the Autonowcaster for some time. Products are generated by the NWS Meteorological Development Laboratory (MDL) and transmitted to Melbourne. They are available for viewing at http://www.mdl.nws.noaa.gov/mlb/ancView/, but the product mix may change from time to time. Sample products, relying heavily on data from six WSR-88D radars in Florida, are described in this section.

The first three figures, all depicting conditions a few minutes after 2000 GMT (1600 EDT) on 04 August 2011, provide information for some of the membership functions that, together, lead to the 60-min forecast of initiation and decay.

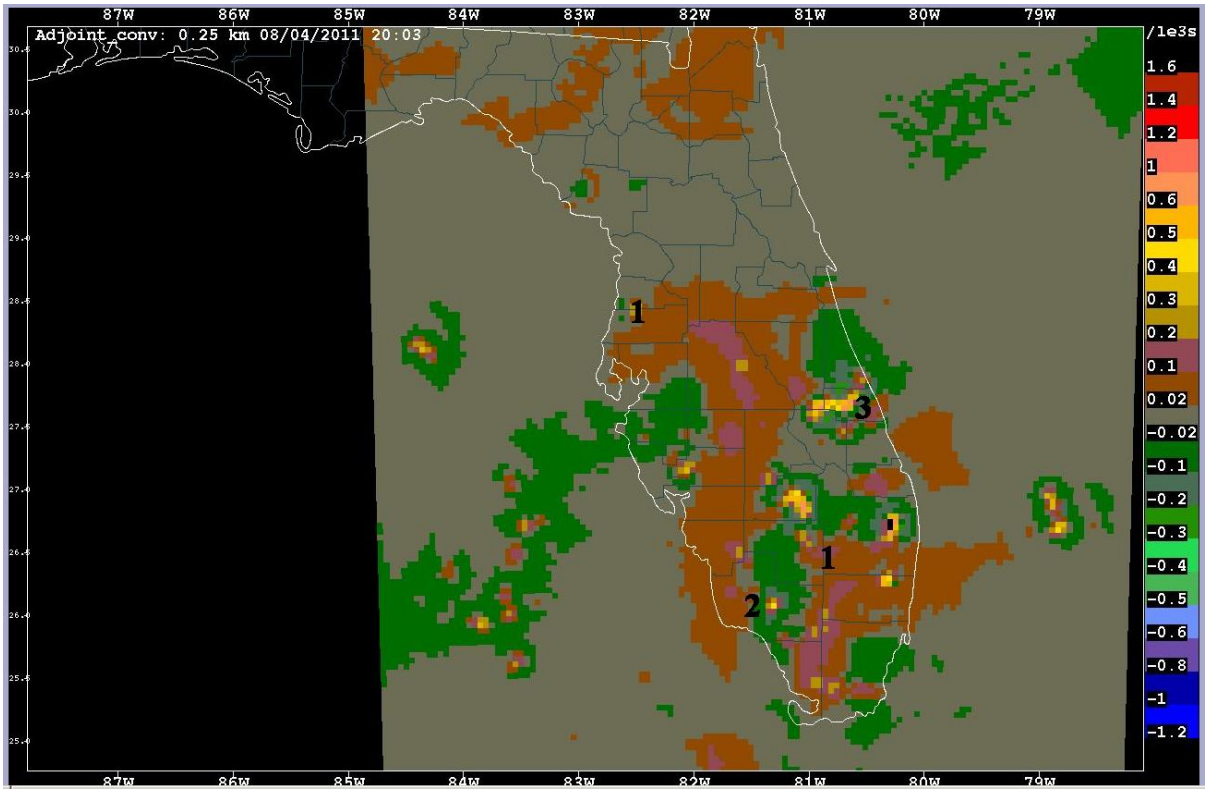

Fig. 2.94. Convergence at 250 m MSL, computed by VDRAS. The color code for convergence is at right in units of 10^{-3} s^{-1}. Negative values indicate divergence. The numbers 1, 2, and 3 are discussed in the text.

Figure 2.94 shows convergence at 250 m MSL, calculated by VDRAS and based primarily on radial velocity data. (The greatest land elevation in Florida is 105 m, so 250 m MSL is low in the boundary layer.) Convergence is color-coded at right in units of 10^{-3} s^{-1}. Warm colors (yellow, orange, red) denote convergence; cool colors (green, blue, purple) denote divergence. Note the numbers 1, 2, and 3 plotted in Fig. 2.93. The number one appears twice. These numbers will appear in the same locations in subsequent figures. The two number ones mark locations where the ANC considers initiation most likely in the coming 60 min (look ahead to Fig. 2.97). Low-level convergence is substantial only near the northern number one. Convergence is also substantial at 2 and 3. It is strongest of all ENE of the southern number one, near the coastline. We shall see what happens 60 min hence at all these locations, but bear in mind that low-level convergence is only one of many interest fields.

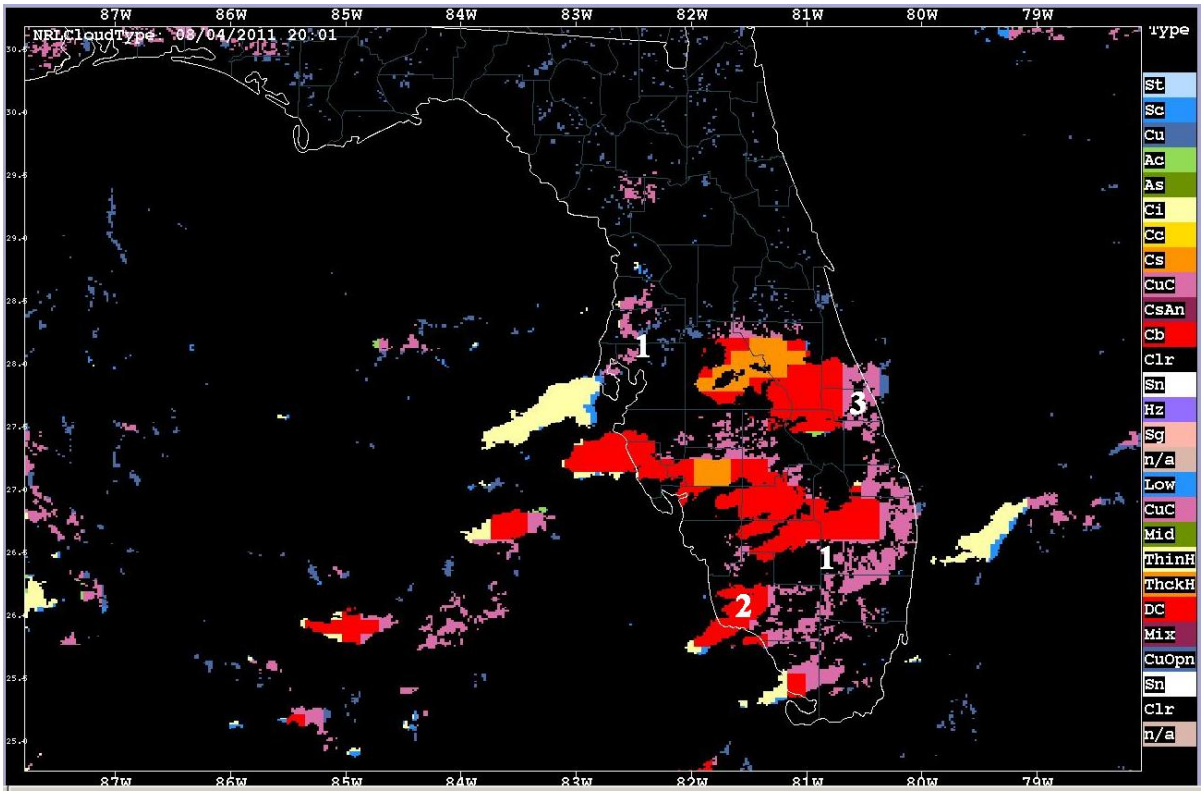

Fig. 2.95. Cloud analysis derived from visible and infrared images (GOES East) and analyzed by the Naval Research Laboratory. The color code at right identifies cloud type. Only the top half of the code is used during daylight. In order from the top, the cloud types are: St – stratus, Sc – stratocumulus, Cu – cumulus, Ac – altocumulus, As – altostratus, Ci – cirrus, Cc – cirrocumulus, Cs – cirrostratus, CuC – cumulus congestus, CsAn – cirrostratus anvil, Cb – cumulonimbus, Clr – clear. The numbers 1, 2, and 3 are discussed in the text.

Figure 2.95 is a cloud analysis produced by the Naval Research Lab (NRL) from GOES East visible and infrared images. The visible image gives cloud shape and texture, whereas the infrared image gives cloud-top temperature, which is easily converted to cloud-top height with the aid of sounding information from numerical prediction models. The red thunderstorm anvils are most prominent in this analysis. The eye discerns a westward drift of the anvils, caused by 20-knot ENE winds in the high troposphere. Note the cumulus congestus (towering cumulus) clouds (magenta) at 1 (both) and 3. Number 2 is already under an anvil, but cumulus congestus lie just to the east. At this time of day (1600 EDT, close to the time of maximum surface

temperature), cumulus congestus, large and growing cumuli with well-developed updrafts, are far more prominent than the smaller cumulus clouds (dark blue) speckled across the map.

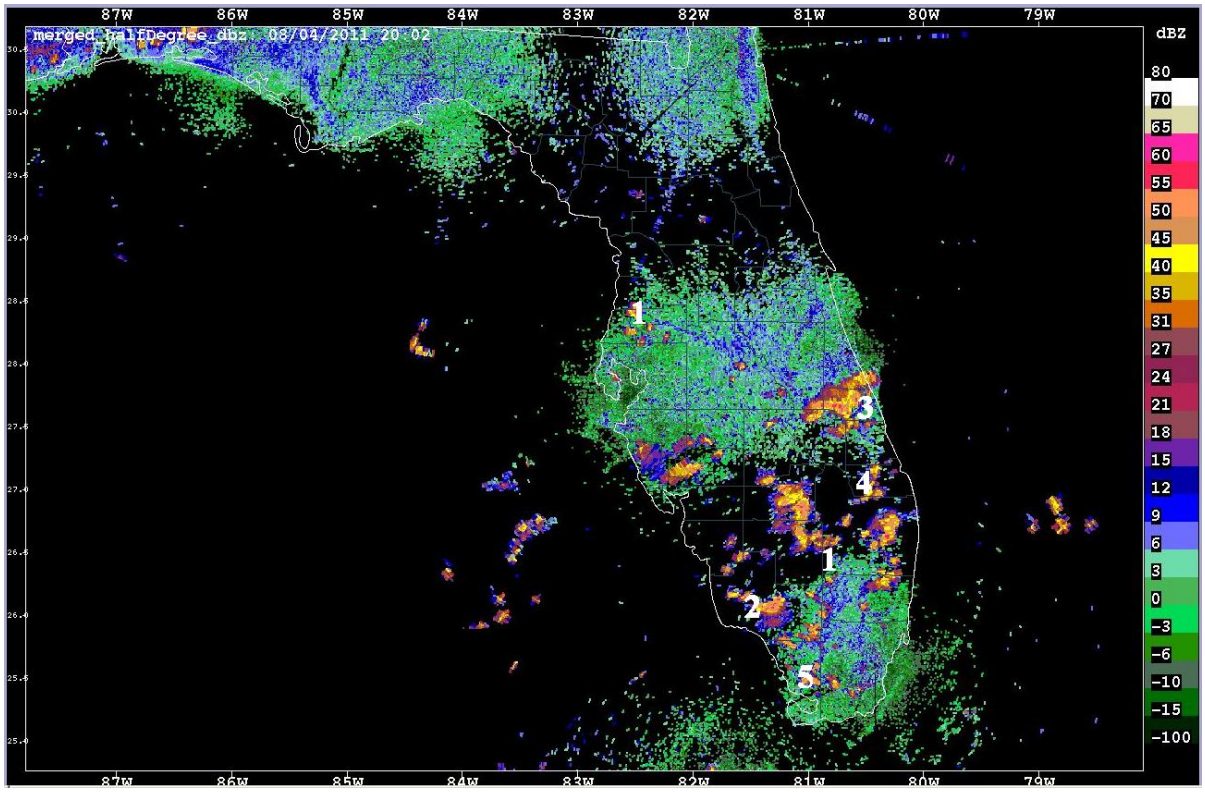

Fig. 2.96. Merged PPI scans from six WSR-88D radars, each at 0.5° elevation for 2002 GMT, 04 August 2011. The color-coded reflectivity scale in dBZ is at right. The numbers 1 – 5 are discussed in the text.

Figure 2.96 shows the PPI scans from six WSR-88D radars around Florida, each at 0.5° elevation. Clear air returns (blues and greens) dominate the image. Several weak boundaries are evident (blue) in central and southern Florida. In areas devoid of echoes, the distance to the nearest radar is great enough that the half-degree beams no longer lie in the boundary layer. Precipitation echoes generally have dBZ values above 20. Note the small showers near 1 (both locations). Larger thunderstorms are near 2 and 3. The numbers 4 and 5 appear in this figure and the next three. Significant convective development occurs near 4 and 5 in the next hour (compare Figs. 2.96 and 2.99).

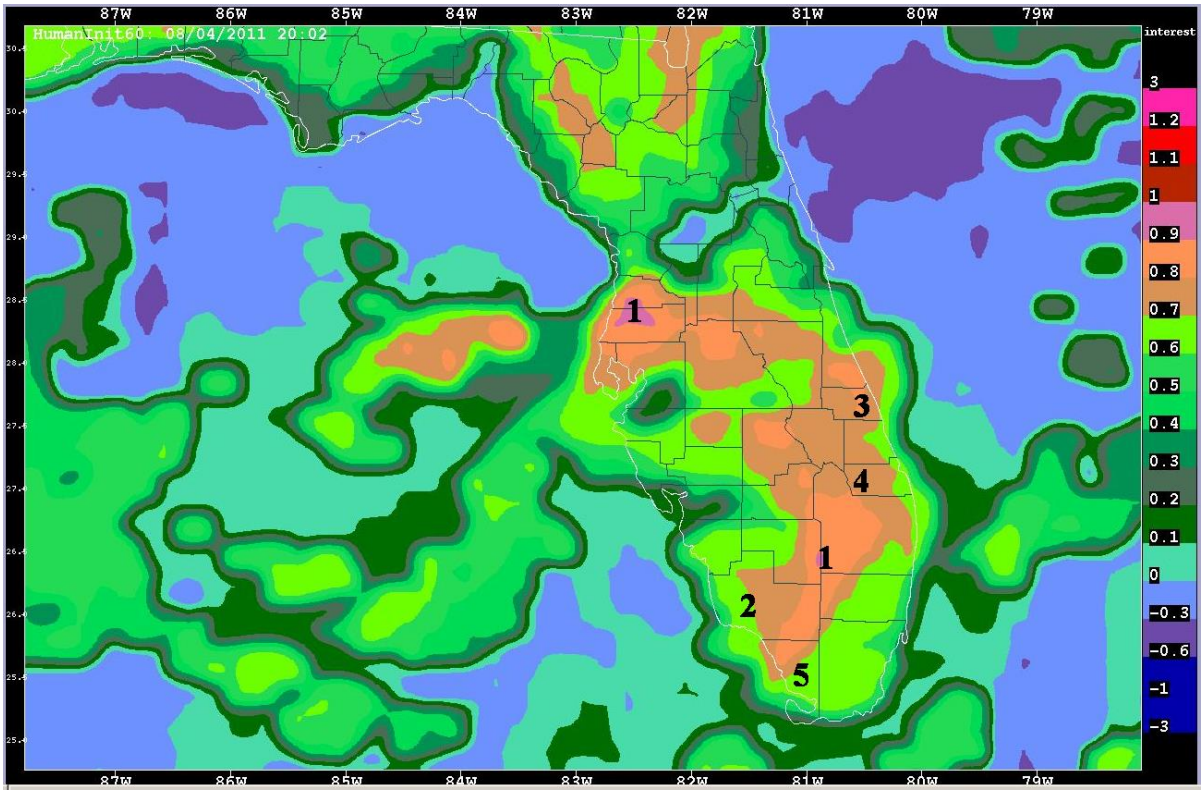

Fig. 2.97. The combined interest field for convective initiation in the 60 minutes following 2002 GMT, 04 August 2011. The color-coded interest values are at right. Where interest exceeds 0.7 (coral colors), initiation is favored. Where interest is negative (purple or dark blue), initiation is unlikely. The numbers 1 – 5 are discussed in the text.

Figure 2.97 shows the combined interest field for convective initiation in the 60 min following 2002 GMT. The numbers 1 mark the two locations with highest interest (0.9). Experience indicates that where interest is ≥ 0.7, forecasters pay attention. Only one area has negative interest (blue-purple) — near 82°W longitude, about one-third of the way down from the top of the figure. In this case, no convection developed there. Melbourne forecasters did not intervene in Fig. 2.97 by inserting their own boundaries, but ANC supports this capability. Such boundaries receive more weight than any other predictor. The combined interest field in Fig. 2.97 is the basis for the 60-min nowcast in Fig. 2.98. We shall follow the numbers 1–5 in Figs. 2.98 and 2.99. The latter figure verifies the forecast with radar data.

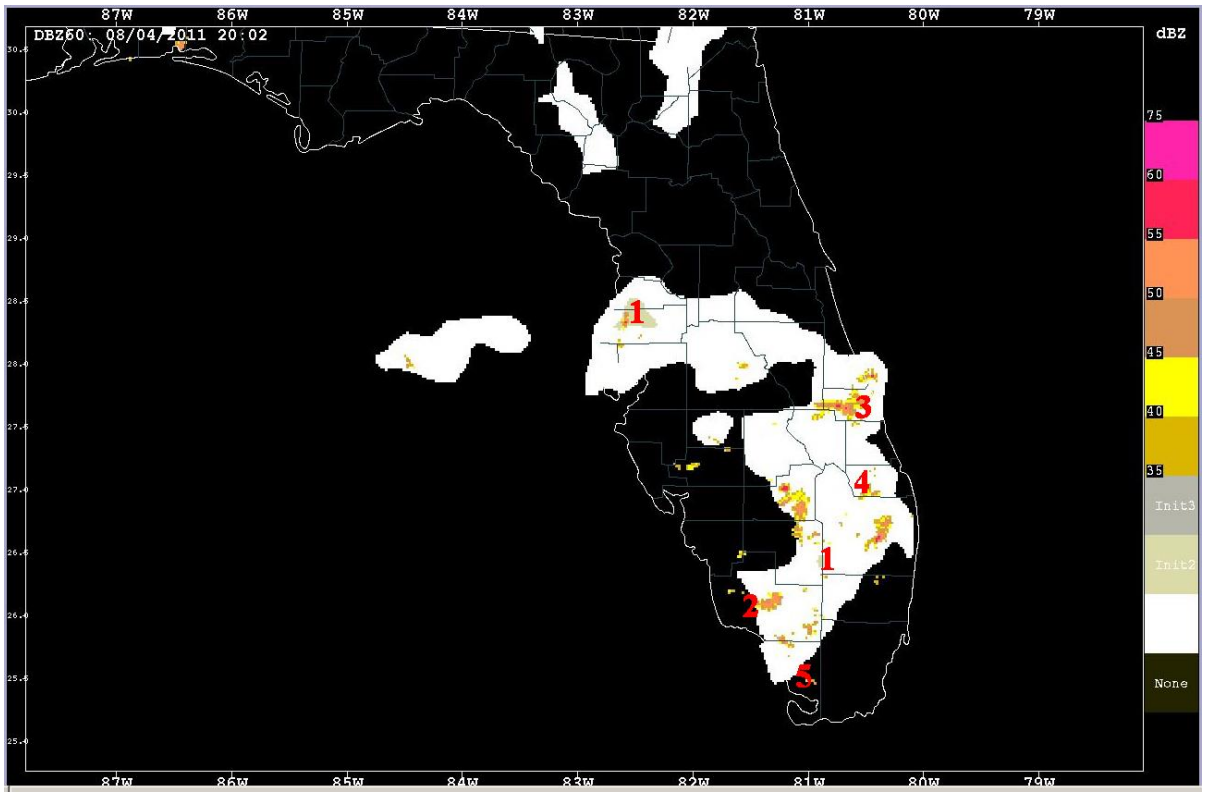

Fig. 2.98. The 60-min nowcast for convection made at 2002 GMT, 04 August 2011. White (Init 1) shows where the interest values in Fig. 2.97 are 0.7 or more. Light brown (Init 2), indicates interest values of 0.9 or more in Fig. 2.97. The gray color (Init 3) would be used for still higher threshold values, but none appears in this figure. The brighter colors indicate echo locations 60 min hence along with their dBZ values, thresholded at 35. The numbers are discussed in the text.

Figure 2.98 is the nowcast verifying at 2102 GMT. The information in Fig. 2.97 is again conveyed here, but interest values less than 0.7 in Fig. 2.97 are not shown here. Initiation is considered possible within the white regions. The echoes shown here are the predicted positions of echoes already present at the initial time (Fig. 2.97). The reflectivity values of the predicted echoes are given in the color bar at right. Note that a few small echoes are expected to persist at locations 1, 4, and 5. Larger echoes are expected to persist near 2 and 3.

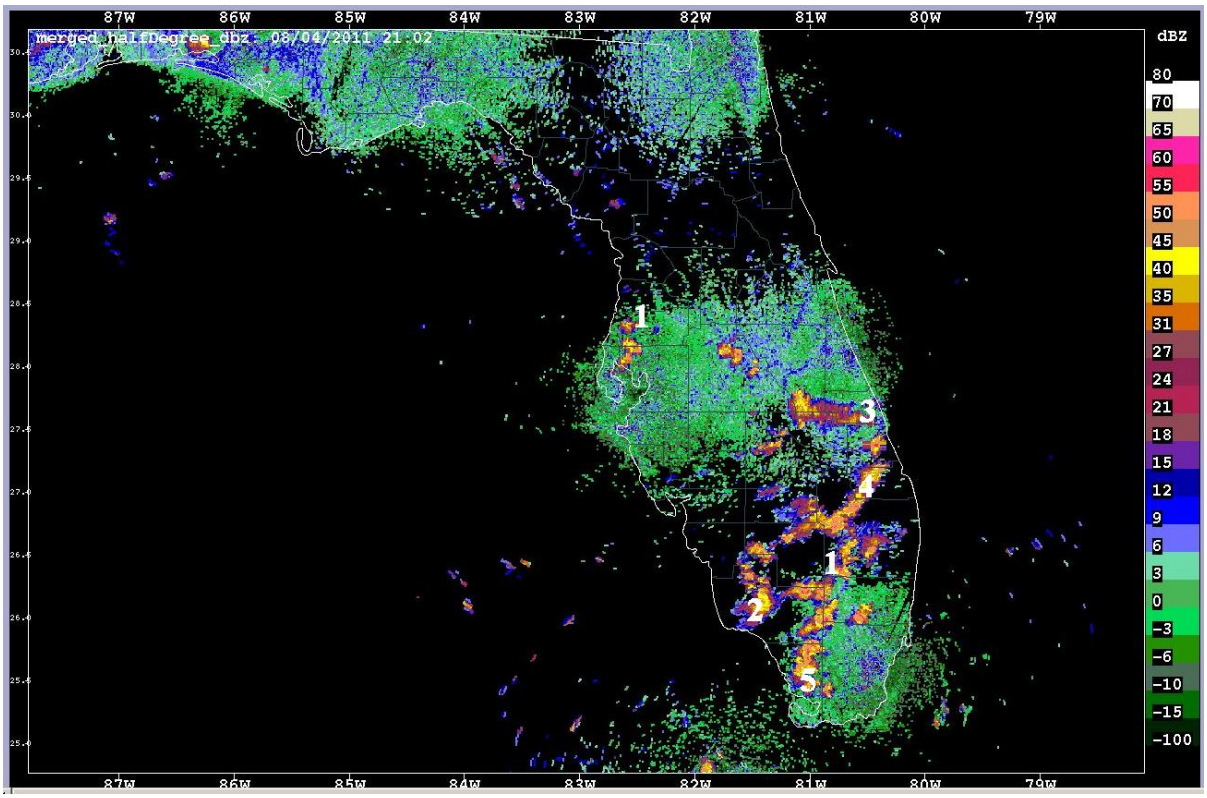

Fig. 2.99. Verifying reflectivity map for the 1-h nowcast in Fig. 2.98. Valid time is 2102 GMT, 04 August 2011. Data are from PPI scans at 0.5° elevation from six WSR-88D radars. The numbers 1 – 5 are discussed in the text.

Figure 2.99 is the reflectivity map that verifies the 1-h nowcast made in Fig. 2.98. Most echoes expected to persist in Fig. 2.98 appear in approximately correct locations in Fig. 2.99. The 1s mark where initiation was most likely. Convection developed within 10 km of each 1. Convective activity decreased west and north of 3, contrary to expectation (compare Figs. 2.96 and 2.99). This is considered a missed forecast. At 4, a line of convection oriented NE-SW blossomed during the hour. This lies within the large thresholded region where the interest value is ≥ 0.7 (Fig. 2.98) and so is considered a hit. Convection already in progress at 2 increased in area. Convection at 5 was minimal at the start of the hour but substantial at the end. Both 2 and 5 lie along the edge of the thresholded area but may still be considered hits. The false alarm rate is large; it is related to that portion of the area in white (Fig. 2.98) where

initiation did *not* occur. Still, given the difficulty of forecasting initiation in the first place, this ANC forecast was moderately successful.

The foregoing section introduced a sample of ANC products and gave a rough qualitative evaluation of performance during one afternoon of convection in Florida. The next section treats quantitative verification.

2.3.11.6 Verification

Roberts et al. (2011) have recently reported verification results from a test of the ANC at the Dallas – Fort Worth, Texas NWS Office (DFW) from 2005 to 2010. During this demonstration, forecasters could exert control over operation of the ANC by choosing the scenario of the day, which, in turn, determined the mix of predictors used. They could also enter surface convergence boundaries, specify their orientation and speed of motion, control the weight given to them as predictors in the ANC, and even delete them if they ceased to be important. The main goal of the DFW demonstration was to improve the consistency, reliability, and accuracy of 1-h convective storms nowcasts for inclusion in automated aviation weather digital products planned for NextGen. The main outcome of the verification was to show that forecaster intervention consistently improves the automated ANC nowcast. Only a brief sample of results can be included in this survey. See Roberts et al. (2011) for many more examples.

Verification scores for the 60-min nowcasts are computed every six minutes. Roberts et al. (2011) present scores computed over the full verification area, which includes the coverage of eight WSR-88D radars in northeast Texas and extreme southern Oklahoma. A hit is defined as an echo (≥ 35 dBZ) occurring within a 2-km grid box at the verifying time of the 60-min nowcast, when the ANC predicted such an echo. This is a harsh measure, in that the timing *and* location must be nearly perfect. Echoes extrapolated from one location to another during the hour are mixed together with echoes that formed during the hour, and the former greatly outnumber the latter. This is reflected in the forecast scores of Fig. 2.100.

Figure 2.100 shows five curves. Init 1 (dotted), Init 2 (dashed), and Init 3 (solid) were introduced in Fig. 2.98. They represent successively smaller areas, each one inside the preceding one, in which convective initiation is increasingly likely. The persistence forecast (no change during the hour – dot-dashed gray curve) and the extrapolation forecast (existing echoes move during the hour – solid gray curve) are also shown. Initiation occurs between the two bold vertical lines, and that is where differences among forecasts are more noticeable. As

expected, the larger the area where initiation is predicted, the greater the POD and also the larger the bias (many more storms predicted than observed). The FAR is uniformly high for all forecasts, near 0.8, with little distinction among them. The CSI scores for the Init 1, 2, and 3 forecasts are only slightly higher than those for extrapolation and persistence during initiation time.

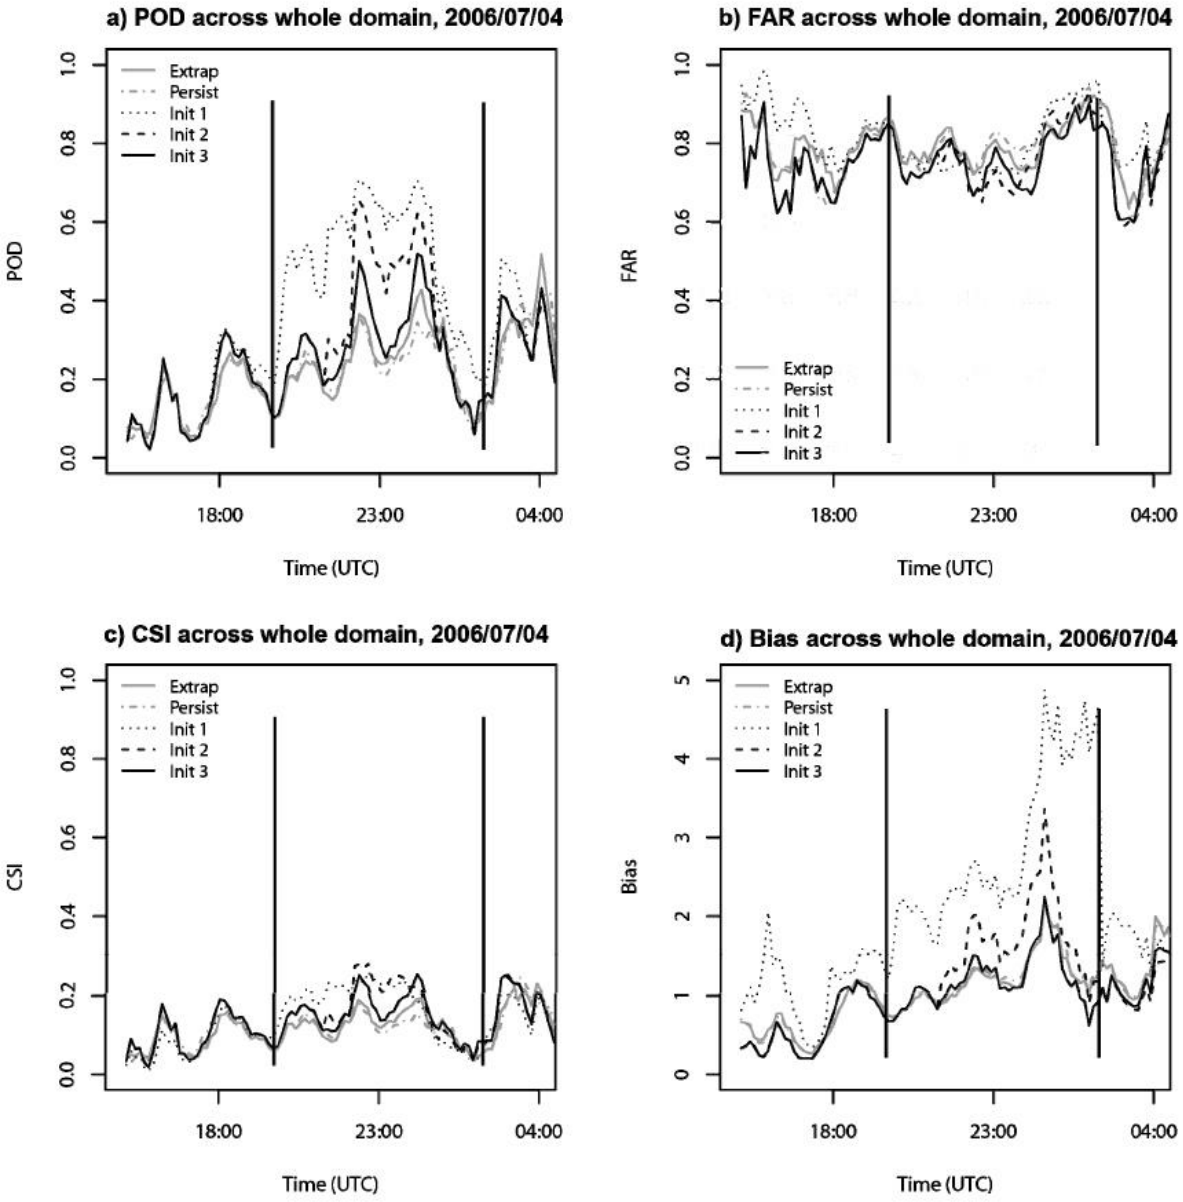

Fig. 2.100. Standard skill scores (POD, FAR, CSI, and bias) for five different forecasts made on 04 July 2006: extrapolation (Extrap), persistence (Persist), and Init 1, 2, and 3, successively smaller areas, one inside the other, where convective initiation becomes increasingly likely. The two bold vertical bars bracket the times when convective initiation occurred. ANC issues a

new 60-min nowcast every 12 min. Reprinted with permission of the American Meteorological Society.

The scoring depicted in Fig. 2.100 is not particularly useful for illustrating the value added to the ANC product when forecasters manually insert boundaries. Roberts et al. (2011) suggested that verifying in smaller areas (1° latitude by 1° longitude) would highlight this benefit more clearly. Figure 2.101 shows what happens to CSI scores when verification is in 1° lat-lon squares and boundaries lie within these squares.

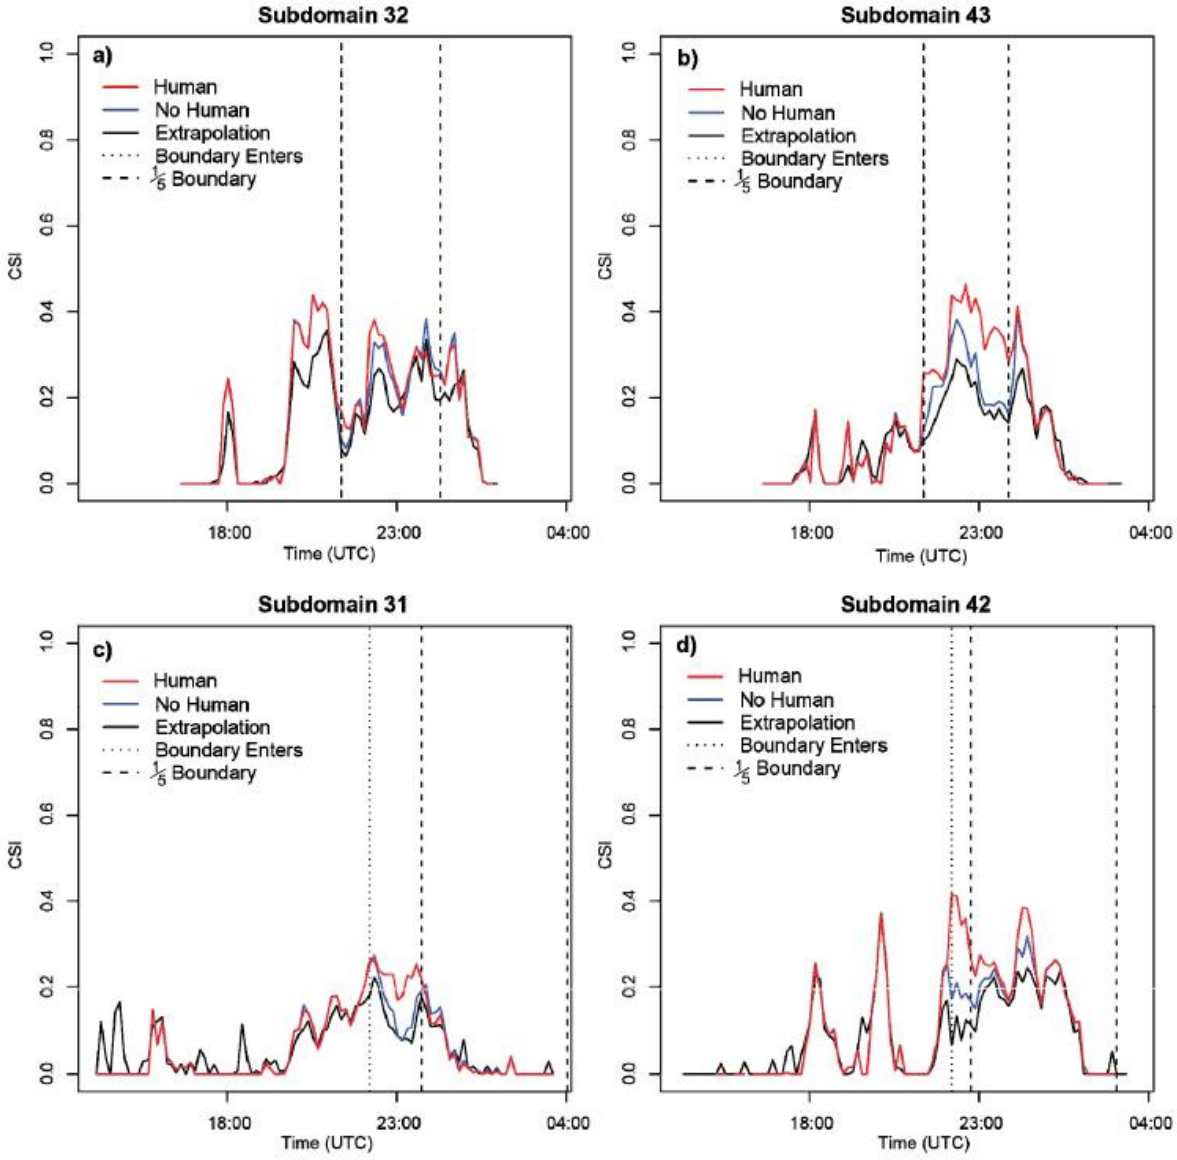

Fig. 2.101. CSI scores for 04 July 2006 within one-degree latitude-longitude squares, each identified by a subdomain number. The subdomains all lie within the full domain for which

scores were presented in Fig. 2.100. These subdomains were chosen because they had boundaries within them. The vertical dotted line indicates the time when the boundary first appeared within the subdomain, either because the ANC algorithm identified it or a forecaster inserted it. The two dashed vertical lines bracket times when a boundary occupied at least one-fifth of the subdomain area. In subdomains 32 and 43, the dotted vertical line is coincident with the leftmost dashed vertical line. CSI curves are shown for a human in the loop (red), no human (blue), and extrapolation forecasts only (black). From Roberts et al. (2011). Reprinted with permission of the American Meteorological Society.

For times between the vertical dashed lines in Fig. 2.101, note that CSI scores are significantly higher when forecasters enter boundaries manually into the ANC.

Another form of ANC verification, a *performance diagram*, was introduced by Roebber (2009). This diagram takes advantage of rather simple mathematical relationships among several skill scores for "yes-no" forecasts: POD, FAR, CSI, and bias, all introduced in Section 1.4, Table 1.2. The only case not included in these scores is the "no-no" forecast, when the event is not forecast and does not occur. If the event is relatively uncommon, however, success in predicting its non-occurrence is nothing to brag about.

Figure 2.102 shows a performance diagram. The horizontal axis is the success ratio, defined as one minus the false alarm ratio (1 – FAR). The vertical axis is the probability of detection (POD). The dashed lines radiating out from the origin represent the bias. These lines are labeled from 0.3 to 10.0. The solid curves within the diagram are contours of constant Critical Success Ratio (CSI). The area above and to the right of the curve CSI = 0.6 is shaded in light blue. Ideally, verification scores should fall within the shaded area because the POD and CSI are both high, the FAR is low, and the bias is close to one.

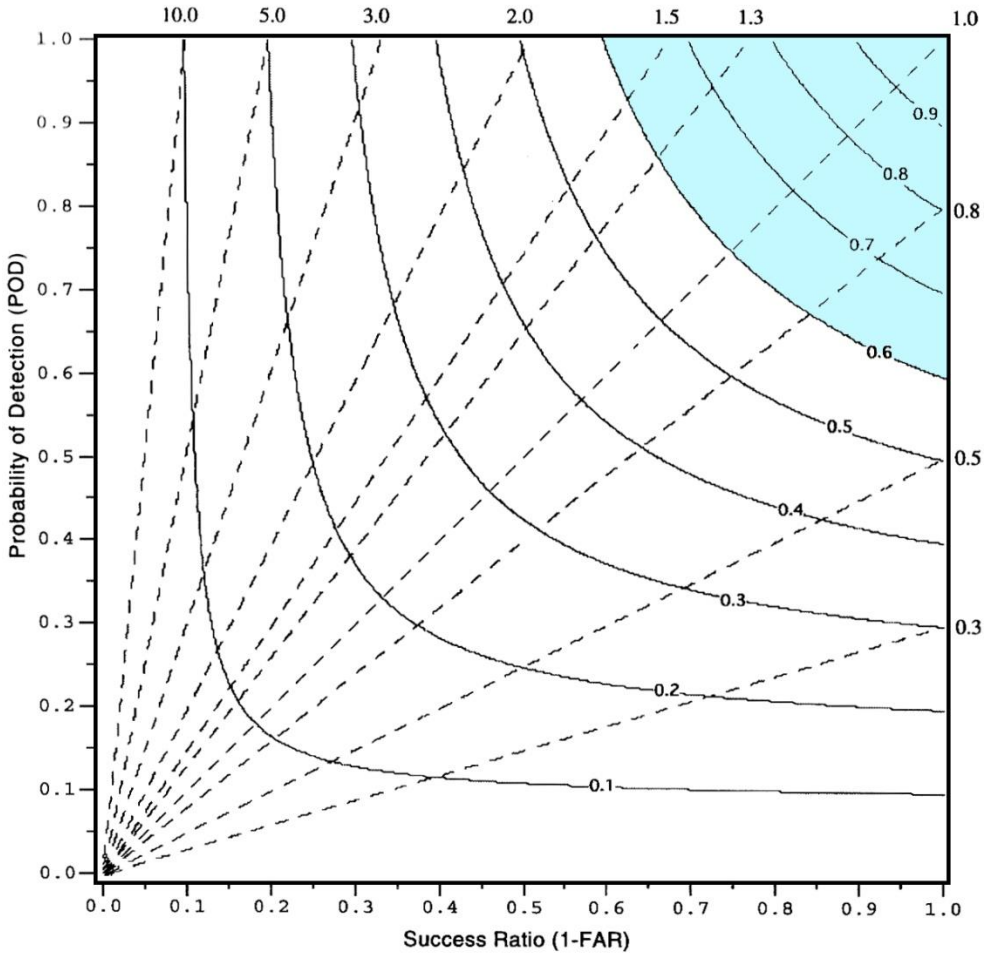

Fig. 2.102. A performance diagram (Roebber 2009). See text for details. Reprinted with permission of the American Meteorological Society.

Figure 2.103 illustrates the use of a performance diagram to verify a one-hour nowcast. The left side of the figure shows the combined interest field for convective initiation, similar to Fig. 2.97, a one-hour nowcast valid at 1933 GMT, 27 August 2006. A forecaster identified a boundary (a cold front in this case) with the yellow line. The predicted position of this line is indicated. The squiggly black outlines show actual positions of the echoes (≥ 35 dBZ) that verify the forecast. Interest values are high along the cold front, as indicated by the color bar across the bottom. The verification area is divided into 30 subdomains (one-degree latitude-longitude rectangles overlaid on the combined interest field). The 30 colored boxes at far right correspond to these subdomains. On this day, the forecaster chose to operate the ANC under the mixed scenario, discussed in Section 2.3.11.2.

In Fig. 2.103, right, verification data are plotted on a performance diagram. The colors of the symbols correspond to boxes occupied by convergence boundaries. These boxes extend diagonally from upper right to lower left in the combined interest image. Triangles represent

nowcasts with forecaster input. Circles represent nowcasts with no forecaster input. Thin gray lines connect triangles and circles of like color so that CSI scores with and without forecaster involvement may be compared within individual boxes. The size of the symbols scales with the ratio of predicted events to all events. In the notation of the contingency table (Table 1.2, Y=yes, N=no), this ratio is (YY + YN) / (YY + YN + NY + NN). A practical outcome is that boxes with more convective events to forecast get larger symbols. All triangle-circle pairs show improved performance when forecasters insert boundaries. The lone exception is the purple pair with CSI values slightly above 0.2. In this case, the CSI resulting from forecaster input was slightly lower than the CSI for the fully automated nowcast. Forecaster input resulted in a greater bias and a decrease in the success ratio.

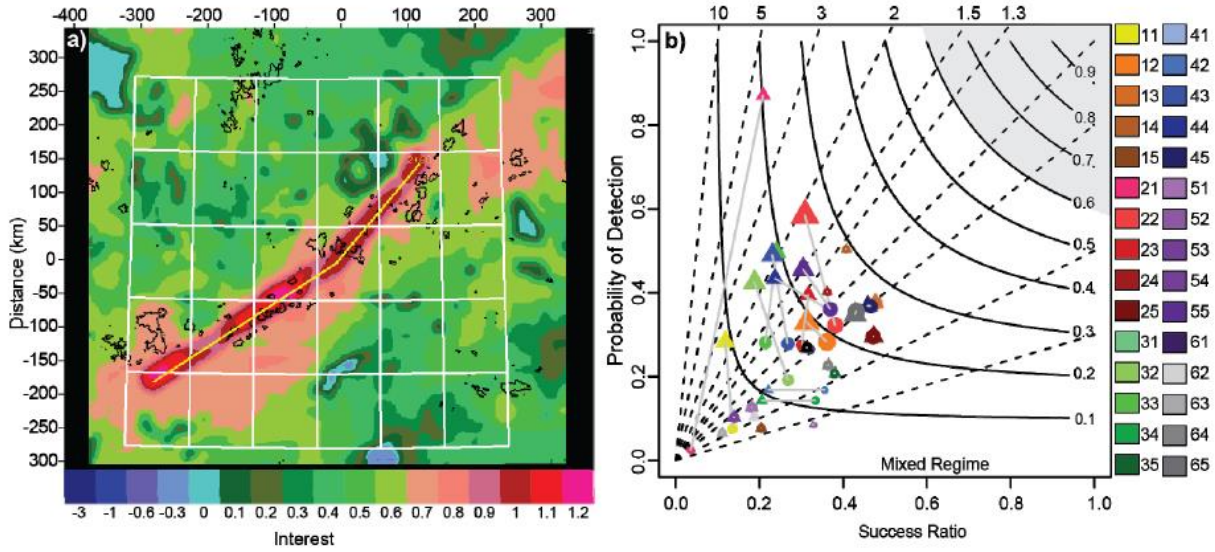

Fig. 2.103. Left: combined interest field for a one-hour nowcast verifying at 1933 GMT 27 April 2006. The color code for interest is at bottom. Initiation is not anticipated unless the interest value is 0.7 or more. Thirty 1° latitude-longitude subdomains are overlaid for verification purposes. These 30 boxes are color-coded at the far right of the figure. Right: performance diagram but only for those boxes that contained boundaries during the one-hour nowcast. Triangles represent nowcasts with forecaster input. Circles represent nowcasts with no forecaster input. Thin gray lines connect like colors to facilitate comparison of nowcasts within a given box, with and without forecaster input. The size of the symbols is roughly proportional to the convective activity within the box. Larger symbols mean more activity. From Roberts et al. (2011). Reprinted with permission of the American Meteorological Society.

Field trials and verification scores indicate that the ANC has difficulty in the following situations:

- Rapid changes in the motion of severe storms when they become right-movers or develop a bow. In the ANC and for most other storm tracking algorithms, giving too much weight to recent changes in the track of the storm centroid runs the risk of creating erratic storm tracks for normal, non-severe storms.
- Boundaries stalling or accelerating. This is a problem for forecaster-entered boundaries but not for the COLIDE algorithm.
- Motion of new storms incorrectly tied to steering level winds rather than remaining attached to the boundary feature that prompted their development, or anchored to a topographical feature. This is a problem for all storm-tracking algorithms.
- Elevated convection. According to Wilson and Roberts (2006), nearly 50% of all convective initiation over the southern Great Plains is elevated convection. With its focus on near-surface convergence boundaries, the ANC cannot anticipate elevated convection; it can only extrapolate the position of existing storms that have their roots above the boundary layer.

2.3.11.7 Commentary

The ANC focuses on true convective initiation, even before there are clouds or radar echoes. It relies heavily on processing reflectivity and radial velocity data from single Doppler radars. VDRAS assimilates these data into a cloud model for warm rain and predicts the wind and potential temperature field anew every 12 min. For awhile at least, MDL will continue to run ANC in small regions and disseminate the products to NWS forecast offices that are testing the ANC in real time and using it in operational forecasts. Eventually, the NWS will consider whether data from individual 88D radars or regional clusters should be stitched together to produce a national ANC product.

The HRRR model used in CoSPA is a full-physics model. Though it runs at a horizontal resolution comparable to that of VDRAS, it does not yet assimilate radial velocity data as VDRAS does, and the HRRR is more computationally demanding per unit area of coverage. Computational demand is the principal reason why HRRR output is not available until two hours past the initial time. It is therefore not useful in a system that produces one-hour nowcasts. Thus VDRAS is the only vehicle at present to assimilate the full-resolution output of 88D radars and use it in real-time predictions.

Under a program called the Collaborative Adaptive Sensing of the Atmosphere (CASA), supported by the National Science Foundation (NSF), a consortium of universities is investigating how a large number of low-power Doppler radars might effectively fill the gaps in

boundary-layer coverage (called *cones of silence*) in the WSR-88D network. Many of the CASA radars could be mounted on cell phone towers. CASA radars are optimized for tracking storms. If they were also made sensitive to clear-air returns in the boundary layer, then systems like the ANC would be more effective. More information about CASA is in McLaughlin et al. (2009) or at http://www.casa.umass.edu/main/research/highlights/casa_in_awips/.

2.3.12 Satellite Products Related to Convective Initiation

2.3.12.1 SATellite Convection AnalySis and Tracking (SATCAST)

SATCAST stands for the SATellite Convection AnalySis and Tracking system. It was developed under the NASA Advanced Satellite Aviation weather Products (ASAP) program by researchers at the University of Alabama–Huntsville (UAH) and the University of Wisconsin Cooperative Institute for Meteorological Satellite Studies (CIMSS) to identify convective clouds in early stages of development. It uses data from both GOES-13 and GOES-15. A recent description of SATCAST is in Iskenderian et al. (2010). The SATCAST algorithm is included in this survey because it was implemented as part of CIWS (see Section 2.3.7) on 01 June 2011. SATCAST defines convective initiation as the first appearance of radar echoes with a reflectivity of at least 35 dBZ.

The SATCAST algorithm has three components that lead to creation of eight interest fields related to convective initiation: 1) a cloud mask component that isolates cumulus clouds (Berendes et al. 2008), 2) a cloud tracking component (Bedka and Mecikalski 2005), and 3) a component to combine cloud type, cloud motion, and various infrared brightness temperatures to produce the eight interest fields.

The eight interest fields are described in Table 2.13, which combines information from tables in Mecikalski et al. (2008) and Mecikalski and Bedka (2006).

Table 2.13

Interest fields for convective cloud growth and glaciation as calculated from infrared channels on GOES-12. (The central wavelengths of some GOES-11 channels differ from those shown here.) IF – Interest Field; T_B – brightness temperature. The resolution is the size of the spot (km) when viewed at nadir. The greater the number of criteria satisfied, the greater the likelihood of convective initiation in the next hour.

Convective Initiation Interest Field	**Resolution and Purpose**	**Critical Values**
6.5-10.7-µm difference – IF1	4-km; cloud-top height relative to the tropopause	–35° to –10°C
13.3-10.7-µm difference –IF2	8-km; cloud-top height changes	–25° to –5°C
10.7-µm T_B – IF3	4-km; cloud-top glaciation	$-20° < T_B < 0°C$
10.7-µm T_B drop below 0°C – IF4	4-km; start of cloud-top glaciation	Within prior 30 min
10.7-µm T_B time trend (IF5 – 15 min; IF6 – 30 min)	4-km; cloud-top cooling rate	$< -4°C\ (15)\ min)^{-1}$ $\Delta T_B\ (30\ min)^{-1} < \Delta T_B\ (15\ min)^{-1}$
6.5-10.7-µm time trend –IF7	4-km; Rate of change in cloud height relative to the tropopause	$> 3°C\ (15\ min)^{-1}$
13.3-10.7-µm time trend –IF8	8-km; rate of change in cloud-top height	$> 3°C\ (15\ min)^{-1}$

Figure 2.104, adapted from Mecikalski and Bedka (2006), is a sketch that neatly explains how these interest fields are used. The vertical axis is temperature (°C). The horizontal axis indicates cloud height. Four of the measures in Table 2.13 are represented. As IF1 increases (dotted line), the temperature difference between the tropopause and cloud top decreases in magnitude as the cloud grows toward the tropopause. Large negative differences indicate a recently formed cumulus cloud; small negative differences indicate a cumulus cloud large enough to generate precipitation. Similarly, IF2 (dashed line) increases as the cumulus cloud grows. IF3 varies with cloud-top temperature; colder cloud tops are associated with taller clouds. If IF5 indicates 3°C or more cloud-top cooling in 15 min, the cloud is clearly growing. The dash-dotted horizontal line at -3°C marks this threshold.

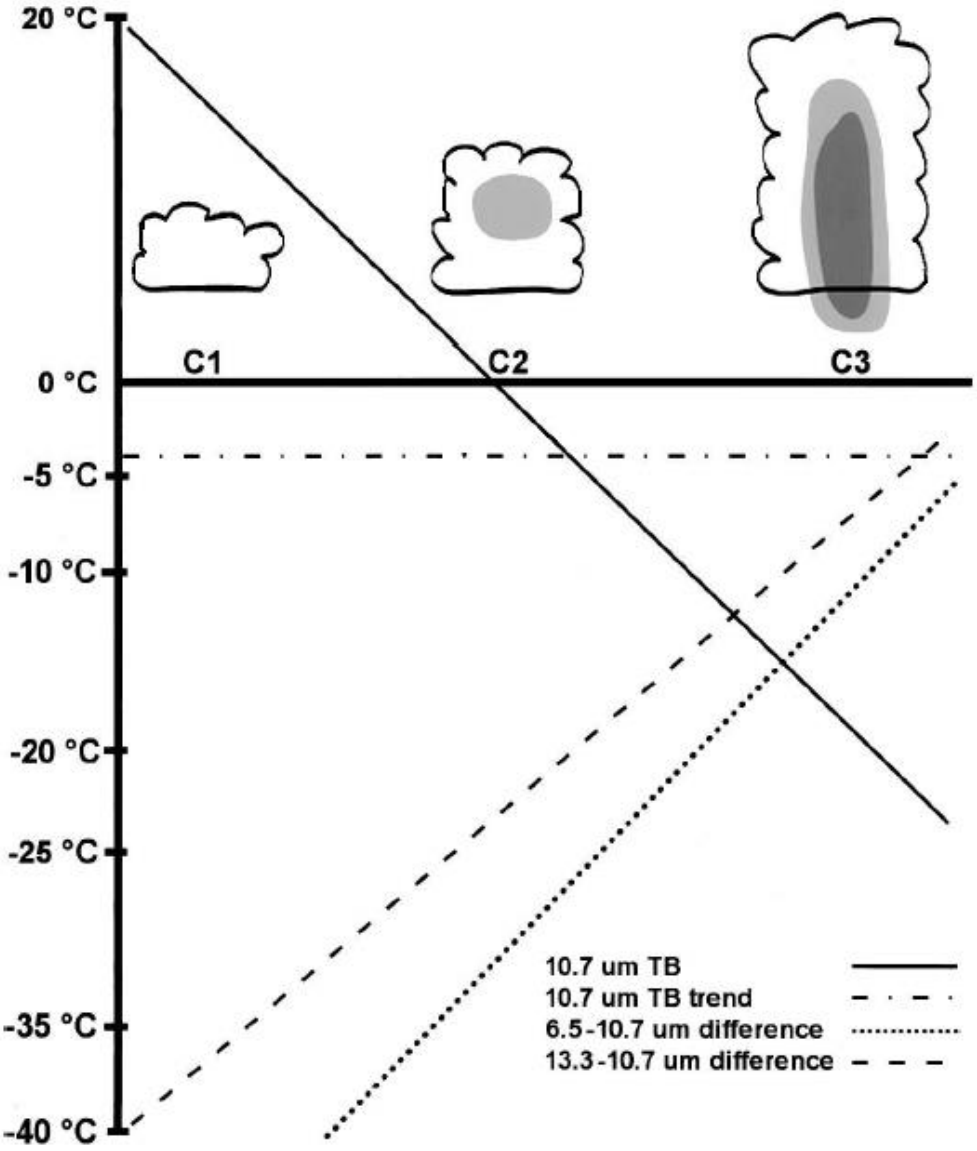

Fig. 2.104. Sketch illustrating the connection between brightness temperatures (TB) in different IR channels, channel differences in TB, and time trends in TB as cumulus clouds grow. Channels are indicated by center wavelength, e.g., 10.7 μm. C1 is a small cumulus cloud with no radar echo, sometimes called cumulus humilis or fair-weather cumulus. C2 is a towering cumulus (cumulus congestus) producing an early radar echo (≥ 10 dbZ, light gray shading). C3 is a young cumulonimbus producing precipitation. Darker shading indicates a radar echo > 35 dBZ. The temperature scale at left relates to critical values of interest fields in Table 2.13. The sloping lines indicate values typically found during various stages of growth; values are consistent with those in Table 2.13. The trend in this figure is for 15 min. The horizontal dash-dot line indicates a threshold. As indicated in Table 2.13, values below the threshold indicate growing cumulus and, above the threshold, no growth. Reprinted with permission of the American Meteorological Society.

The eight interest fields are combined into a single convective initiation nowcast field. At each pixel location, the number of interest fields for which the critical values are satisfied is tallied. Wherever the count is ≥ 7, there is high potential for convective initiation in the next hour.

The connection between SATCAST and CoSPA is illustrated in Fig. 2.105. The eight SATCAST indicators are used together with three additional fields.

The first of the three fields is a mask (Fig. 2.105a), which is a combination of the instability near the surface and the instability aloft. The instability near the surface is determined by the Convective Aviailable Potential Energy (CAPE) in the lowest 50 mb of the troposphere and nearness of the surface temperature to the convective temperature. Areas of higher CAPE and a surface temperature near or at the convective temperature will have higher values of the mask (light shades of gray ranging to white). The instability aloft is determined by the CAPE from 50 mb above the surface to 600 mb. Again, higher values of CAPE result in higher mask values. Data from the RUC model (soon, the RAP model) (hourly surface temperature and tropospheric analyses) and the NOAA Space-Time Mesoscale Analysis System (STMAS — 15-min surface analyses; Xie et al. 2010) are used to calculate the latest values of CAPE and surface temperature.

The second field used with the eight SATCAST indicators is the radar growth (Fig. 2.105b), dilated with a 20-km circular kernel.

The third field is the solar elevation angle to incorporate the time of day in the convective initiation algorithm (Fig. 2.105c)

These three fields are combined in a fuzzy–logic, weighted average and applied to the eight SATCAST indicators in Fig. 2.105d to create a downselected (reduced) number of SATCAST indicators (Fig. 2.105e) that are used to initiate convection. Areas of high instability near radar growth at times of large solar elevation angle require fewer SATCAST indicators to initiate convection. The downselected SATCAST indicators are then used to create a convective initiation interest field (Fig. 2.105f), which is passed to the CIWS forecast engine to include convective initiation in the 0-2 hour forecast. This 0-2 hour forecast is then supplied to the CoSPA 0-8 hour forecast.

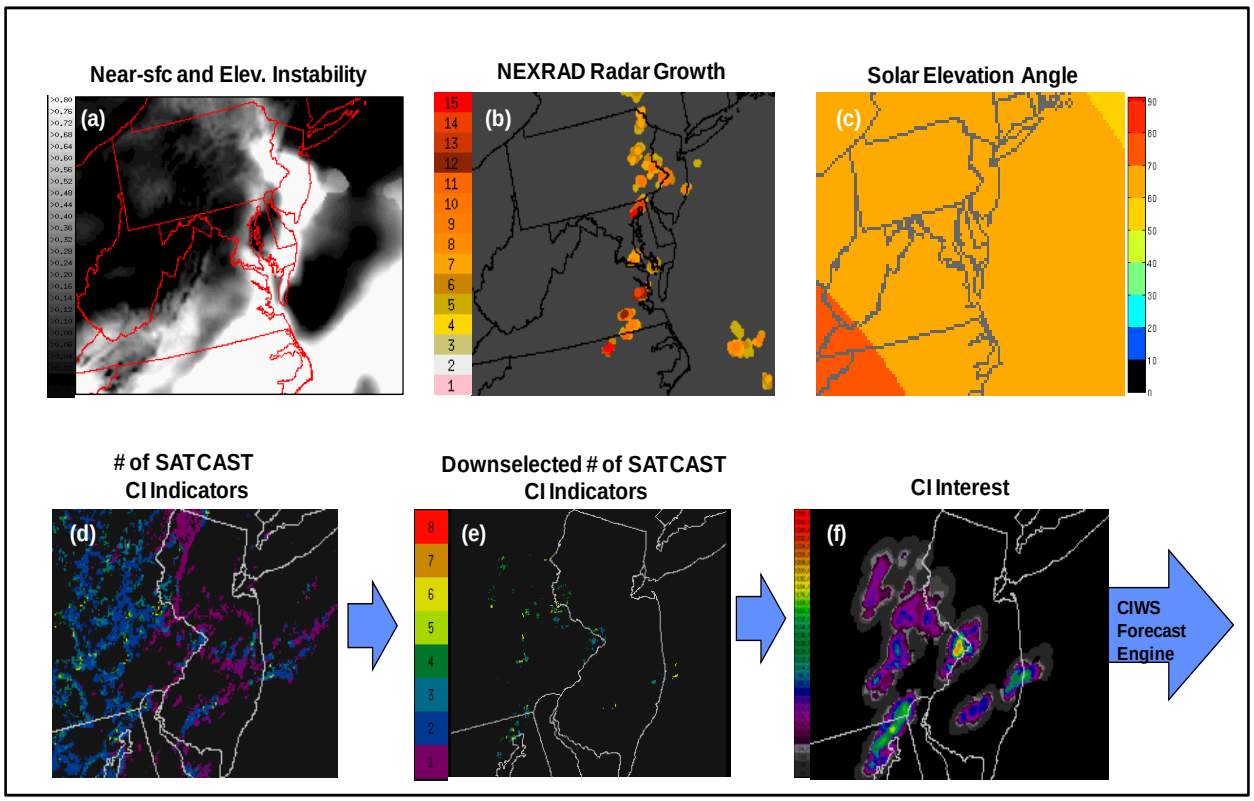

Fig. 2.105. How SATCAST generates a convective initiation (CI) interest field. (a) Composite measure of near-surface and elevated instability. Very stable areas are black; unstable areas are white. (b) NEXRAD (WSR-88D) radar echo growth. Higher numbers indicate more rapid growth. (c) Solar elevation angle in degrees (see color bar at right). (d) The number of SATCAST indicators in Table 2.13 that favor convection. The color bar for (d) is at left in (e); it gives the number of favorable indicators. The information in (a) through (d) is combined to eliminate many of the indicators from consideration in (e). Finally, the remaining indicators in (e) are once again weighed with the instability, growth, and solar elevation fields to generate a final CI interest field. More detail is in the text. Courtesy of Haig Iskendarian of MIT Lincoln Labs.

2.3.12.2 University of Wisconsin Convective Initiation (UWCI) algorithm

The University of Wisconsin Convective Initiation (UWCI) algorithm is thoroughly described in Sieglaff et al. (2011). The UWCI and SATCAST approaches to the problem of convective initiation are different. For the UWCI algorithm, the first cloud-to-ground lightning strike marks convective initiation, not the first appearance of a 35-dBZ echo. The reason for this choice is that there are gaps in U.S. coverage for 88D radar but not for the National Lightning Detection

Network (NLDN), which is used for verification. SATCAST tracks clouds; the UWCI algorithm does not. One objective of the UWCI algorithm is to minimize the false alarm rate; it does so at the expense of a reduced probability of detection compared with SATCAST. The UWCI algorithm uses data from GOES-12 but not GOES-11. Thus its coverage is limited from the Great Plains eastward. SATCAST focuses on cloud-top temperature; the UWCI algorithm tries to determine the phase of cloud-top hydrometeors. Finally, the UWCI algorithm does not explicitly track clouds; SATCAST does.

The UWCI algorithm relies only upon infrared radiances and so is equally applicable day and night. The first step is to mask pixels that do not contain cloud. The next step is to classify the phase of hydrometeors appearing in the pixel: liquid water, supercooled liquid water, mixed phase (liquid mixed with ice), opaque ice (also called "thick ice"), non-opaque ice (also called "thin cirrus"), and multilayered ice cloud. Technically, "opaque" refers to infrared optical depth, but in practice, if the ice cloud is dense and thick, it will be opaque.

The cloud classification scheme is due to Pavolonis (2010). Brightness temperature differences (BTD) between two infrared channels have long been used to infer cloud composition. But Pavolonis showed that, for clouds with an 11-µm emissivity of less than 0.85, commonly used BTDs constructed from 8.5-, 11-, and 12-µm brightness temperatures are more sensitive to non-cloud variables than to cloud microphysics. Non-cloud variables include surface temperature, surface emissivity, and clear-sky atmospheric transmittance. A different parameter, called the effective absorption optical depth ratio (β), shows much greater sensitivity to cloud microphysics. Pavolonis computes β from the same infrared radiances that yield BTDs. He first computes effective cloud emissivities ε from radiances at two wavelengths, λ_1 and λ_2, in the range 8–13 µm. These emissivities are, in most cases, proportional to the fraction of radiation incident at cloud base that is absorbed by the cloud. He then computes β from the two emissivities:

$$\beta = \frac{ln[1 - \varepsilon(\lambda_1)]}{ln[1 - \varepsilon(\lambda_2)]}$$

Different values of β arise from different choices of the two wavelengths. Figure 2.106, resulting from theoretical calculations reported by Pavolonis (2010), illustrates how two commonly used pairings of wavelengths provide information about cloud composition.

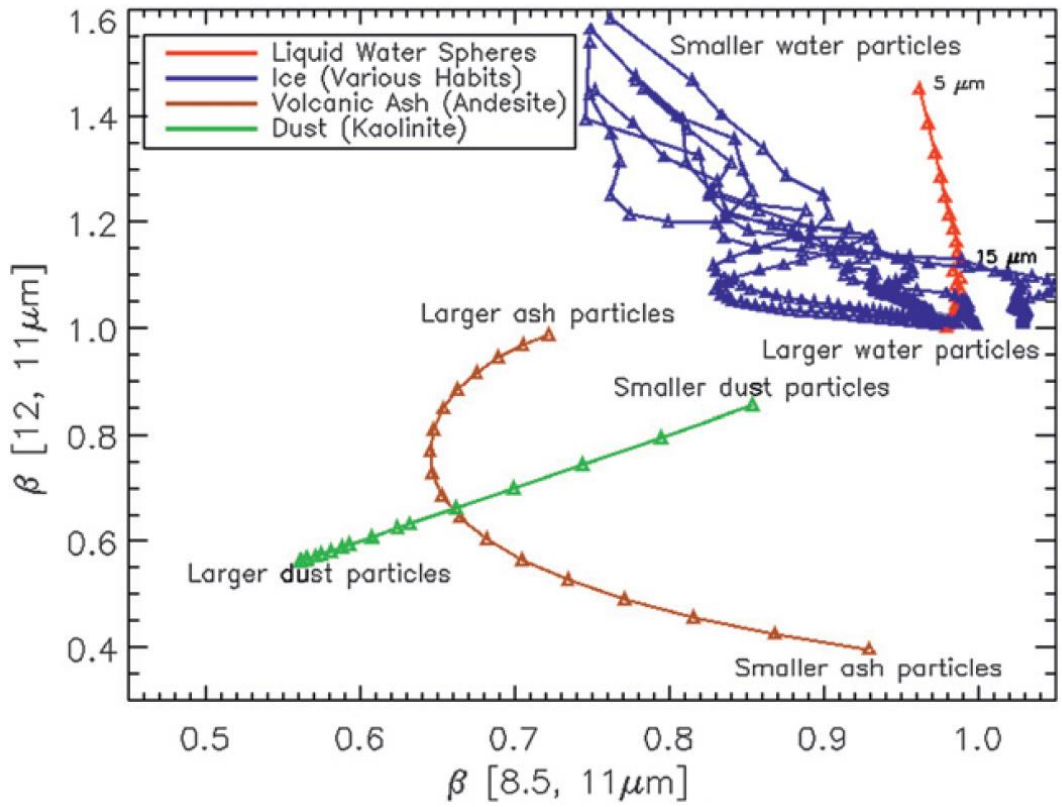

Fig. 2.106. Results of theoretical calculations illustrating how cloud composition varies with two pairings of wavelengths for β, the effective absorption optical depth ratio. Effective emissivity at 11 μm is used in the denominator in the expression for β above. Red indicates liquid water spheres. Size increases from 5 μm at top to 30 μm at bottom. The 15 μm size is marked. Blue indicates various types of ice crystals (plates, aggregates, columns, bullet rosettes, etc.). Brown refers to spherical volcanic ash particles (andesite is a mineral found in volcanic ash) arranged by size from larger at top to smaller at bottom. Green represents spherical dust particles, also arranged by size along the straight line. Kaolinite is a mineral dust picked up by the wind. Reprinted with permission of the American Meteorological Society. Compare with Fig. 4.13.

Figure 2.106 indicates that solid and liquid hydrometeors are well separated in beta space. Only for large cloud droplets is there ambiguity in distinguishing between the two. Mineral particles are well separated from hydrometeors, but, because the curves for mineral dust and volcanic ash cross near their midpoints, it is more difficult to distinguish between these two. Water and supercooled water are indistinguishable in beta space, but cloud-top temperature is readily estimated from infrared brightness temperatures. If a liquid cloud top is sub-freezing, it is supercooled.

Once cloud-top properties are identified, pixel by pixel, the main part of the UWCI algorithm executes. It is based upon box averaging. A square box, seven pixels on a side (28 km on a side when the satellite views straight down), is centered on each cloudy pixel. The size of this box is chosen so that a convective cloud at the center of the box could move much faster than normal in 15 min (the nominal time interval between successive GOES images) and still remain in the box. The IR-window brightness temperature (BT) is averaged within each 7 x 7 pixel box, but only those pixels identified as cloudy are included. Clear-sky pixels and pixels with fog are excluded. At least two cloudy pixels must be present in each box. Otherwise, no average is computed. If a valid average is present for the same pixel in two successive images, the change in average BT is computed over time. This is used as a proxy for cloud-top cooling or, equivalently, vertical cloud growth.

Figure 2.107, from Sieglaff et al. (2011) shows the resulting unfiltered cloud-top cooling (CTC) rate, with the lower limit fixed at -4^{o}C in 15 min. A large complex of thunderstorms (not shown) with very cold anvil tops extended from north central Texas, across the eastern two-thirds of Oklahoma, and into eastern Kansas and southwest Missouri. Significant cloud-top cooling occurred over only a small part of this area, most of it from overshooting tops. Such cloud top cooling is a sign of vigorous convection but has nothing to do with convective initiation. It is also possible that some of the indicated cooling might be caused by the advection of higher clouds into the 7 x 7 box rather than by the vertical growth of convective clouds. Clearly, filtering is necessary in order to isolate growing cumulus clouds of small or intermediate size.

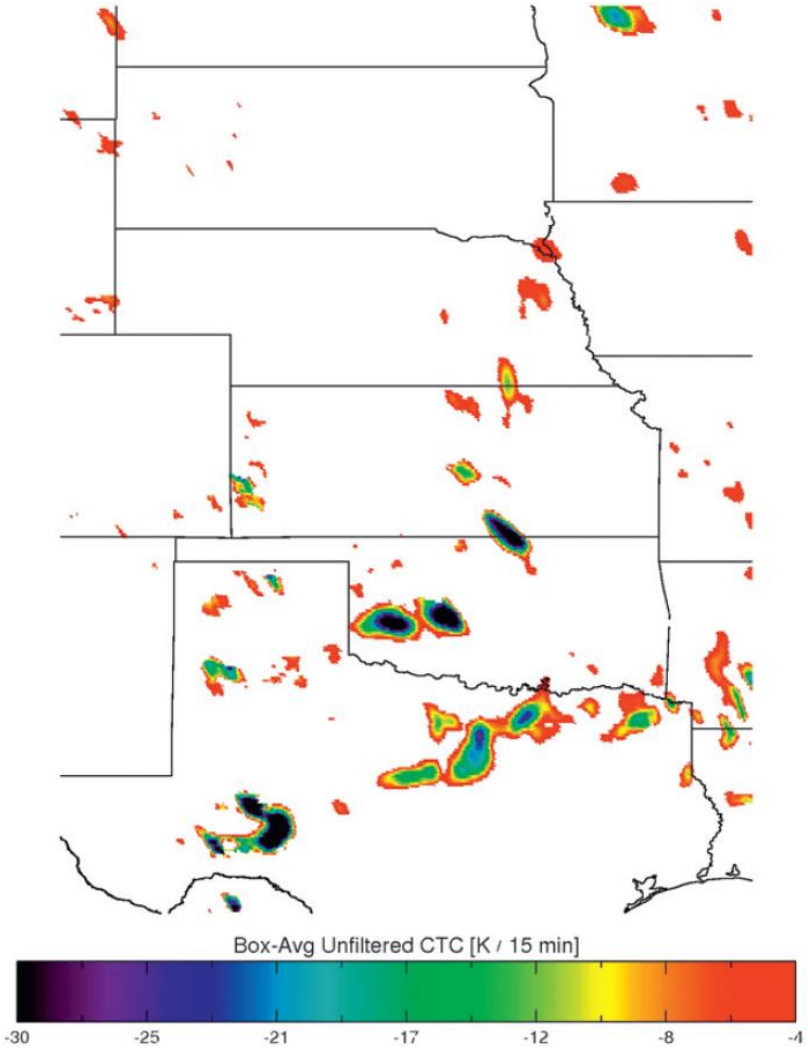

Fig. 2.107. Unfiltered cloud-top cooling (CTC) rate in ^{o}C (15 min)$^{-1}$, valid at 2015 GMT, 29 April 2009. Where there is color, the cooling rate exceeds -4^{o}C in the previous 15 min. The color bar indicates the amount of cooling. The calculation of cooling rates at each pixel is based upon box averages, differenced over time, as explained in the text. Where the background is white, either the sky was clear, there were not enough cloudy pixels in the 7 x 7 pixel box to constitute an average, or there was little or no cooling of cloud tops. Reprinted with permission of the American Meteorological Society.

At this stage, the UWCI algorithm has a 7-by-7 pixel average (BT_{avg}) stored at each pixel location where cloud-top cooling apparently occurred, and it also has the cloud-top classification for each pixel. The next step is to center a 13 by 13 pixel box on each pixel. The raw pixel values (BT – not averaged) *and* the 7-by-7 pixel average values within this larger box are used to

remove colored pixels in Fig. 2.107 that are not relevant to convective initiation. One more source of information is required for the removal of colored pixels. The percentage of pixels in each cloud-type category is computed for all cloudy pixels within the small and large boxes. To repeat, the categories are liquid water, supercooled liquid water, mixed phase, opaque (thick) ice, non-opaque ice (thin cirrus) , and multilayered ice cloud. A colored CTC pixel must survive seven tests for it to be declared relevant to convective initiation. Otherwise, it is removed. These tests are summarized here but described in detail in Sieglaff et al. (2011).

Refer to Fig. 2.108, which shows the center pixel (dark cyan), the periphery of the inner box (7-by-7 pixels – cyan) and the periphery of the outer box (13-by-13 pixels – light cyan). Let P_1 refer to the central pixel, P_7 refer to all pixels on the periphery of the 7-by-7 array, and P_{13} refer to all pixels on the periphery of the 13-by-13 array.

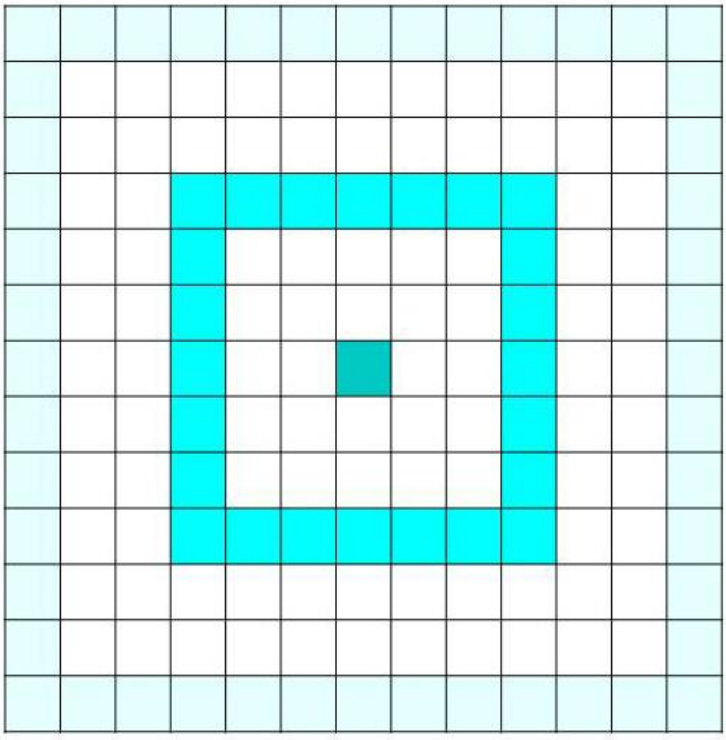

Fig. 2.108. An array of 13 by 13 infrared pixels used to determine whether cooling apparent in the central pixel is explained by vertical growth of cumulus clouds or some other cause.

Test 1. Is cooling at the central pixel due to advection of higher clouds into the 7-by-7 array? If so, discard CTC pixel value. To pass this test, the BT_{avg} at P_1 at the current time must be less than the lowest BT_{avg} in P_{13} at the previous satellite image time.

Test 2. Is the cooling occurring outside the main updraft near the central pixel location? If so, discard CTC pixel value. Three subtests must be passed; all three involve values from the current scan. 1) The BT_{avg} in P_1 must be less than the minimum BT_{avg} in both P_7 and P_{13}. The minimum BT_{avg} in P_7 must be less than the minimum BT_{avg} in P_{13}. 3) For BT *not* averaged, the BT in P_1 must be less than the minimum BT in both P_7 and P_{13}. All three subtests ensure that the greatest cooling is near the center of the 13-by-13 array.

Test 3. Are there too many pixels with ice (thick ice, thin cirrus, or multi-layered ice cloud)? If so, discard CTC pixel value. If less than 50% of the cloudy pixels in the 13-by-13 array are ice pixels, the CTC value survives. If ice clouds predominate in the large array, convective initiation has probably already occurred.

Test 4. When both ice and water clouds are present, do cirrus and multi-layered ice clouds dominate the scene? If so, discard the CTC pixel value. In the 13-by-13 array, when only water clouds, thin cirrus, and multi-layered ice clouds are present, if pixels occupied by the latter two types outnumber the pixels occupied by water clouds, then the CTC pixel value is discarded.

Test 5. Do observed changes in cloud-top microphysics proceed in an unexpected order? If so, discard the CTC pixel value. The tops of convective clouds typically contain liquid water droplets early in their development, then mixed phase (supercooled droplets and ice crystals), and then all ice crystals at maturity. In the absence of thick-ice pixels, if the expected progression does not occur, for example, if cloud tops in the box are first dominated by ice and later with mixed-phase hydrometeors, the CTC pixel value is discarded.

Test 6. Are there too many multi-layered ice clouds obscuring the view of lower clouds? If so, discard the CTC pixel value. In the 7-by-7 box, if multi-layered ice clouds are the dominant cloud type (more pixels with that type than any other), discard the CTC pixel value.

Test 7. This test is too complicated to describe here, but it tends to screen out clouds whose updrafts are inferred to be weak.

Not many colored pixels in Fig. 2.107 survive all seven tests; the survivors are shown in Fig. 2.109, from Sieglaff et al. (2011). These pixels mark where convective initiation is likely: in southwest Kansas, the Texas Panhandle, and south of the Rio Grande River in Mexico.

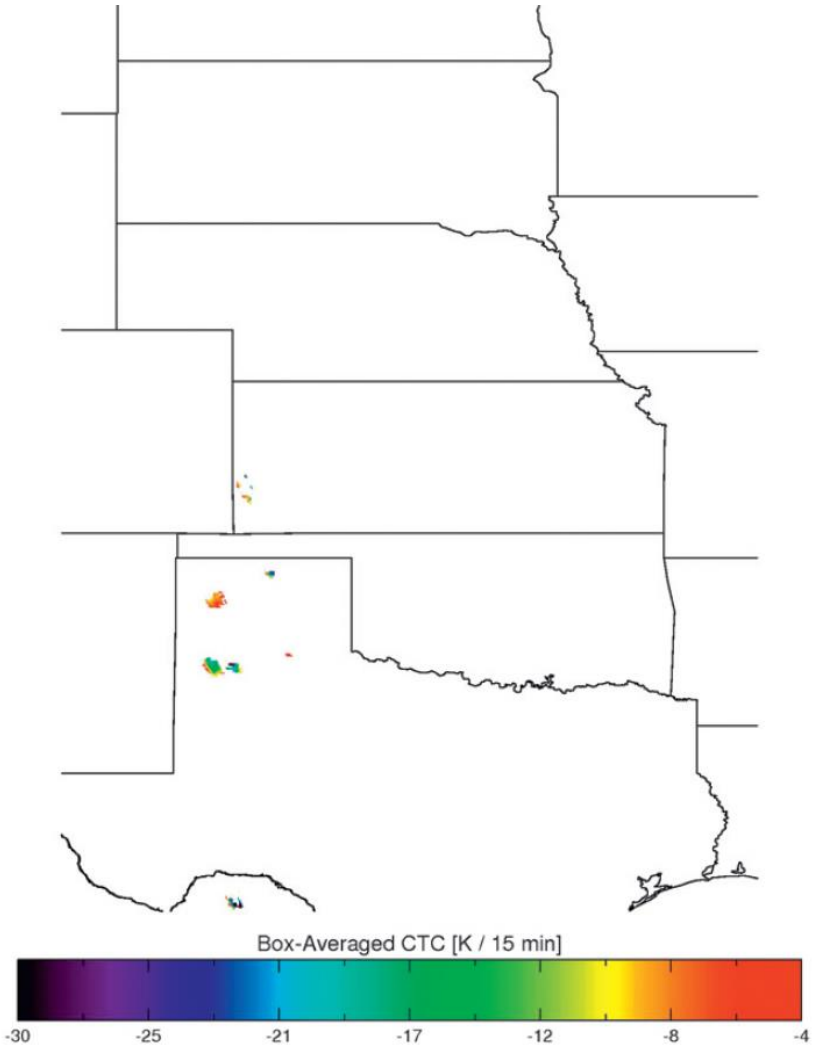

Fig. 2.109. The filtered CTC rate. The colored pixels are a small subset of those shown in Fig. 2.107. They mark the most likely locations for convective initiation, where the cloud-top cooling rate is most likely due to vertical growth of cumulus clouds and not some other cause. Reprinted with permission of the American Meteorological Society.

Figure 2.109 is the basis for the UWCI nowcast. The latter is not shown because it is very similar to Fig. 2.109. The nowcast categories are as follows: *Category 1*. Early growth well before the first cloud-to-ground lightning strike. At least 10% water cloud in the 7-by-7 box and less than 5% for each of supercooled-water, mixed-phase, and thick-ice clouds at the current GOES image time. *Category 2*. Convective initiation likely. At least 5% supercooled-water or mixed-phase clouds in the 7-by-7 box and less than 5% thick-ice cloud at the current GOES image time.

Category 3. Convective initiation occurring. At least 5% thick-ice clouds in the 7-by-7 array at the current GOES image time, indicating glaciation.

Sieglaff et al. (2011) verified the UWCI algorithm during 23 convective afternoons and 11 convective nights in 2008 and 2009. A total of 509 convective initiation nowcasts were issued. As previously noted, all were verified against the time of the first cloud-to-ground lightning strike. The POD for all cases was 47.0%; the FAR was 34.8%. If verification is limited to regions where the Storm Prediction Center (SPC) indicated a slight or greater risk of severe weather (380 nowcasts), the scores improve to 56.3% for POD and 25.5% for FAR. Another way of phrasing the latter result is that the UWCI algorithm produced a successful nowcast for more than half of all storms generating cloud-to-ground lightning within an SPC risk area, and roughly three of every four clouds identified in UWCI nowcasts produced lightning at a later time.

To conclude Section 2.3.12, there are two major satellite products for convective initiation, SATCAST and the UWCI algorithm. Though these products had many developers in common, their methods differ by design. The UWCI algorithm is intended to minimize the FAR. The penalty is missed storms and a somewhat lower POD. The SATCAST emphasis is on detection; it has higher PODs but at the expense of higher FARs. These statements were confirmed by NWS forecasters, who evaluated both products at the SPC Hazardous Weather Testbed in spring 2011.

2.3.13 Multiple Radar / Multiple Sensor (MRMS) System

The Multiple Radar / Multiple Sensor (MRMS) System was developed at the National Severe Storms Laboratory (NSSL) in Norman, Oklahoma, primarily to improve decision-making in hazardous weather situations and sharpen estimates of precipitation at the ground. Parts of MRMS have direct aviation applications, and so NSSL transferred it to FAA's William J. Hughes Technical Center, in Atlantic City, New Jersey, in November of 2010. MRMS will be tested in the context of NextGen.

Probably the most useful element of MRMS for aviation is the national 3-D radar mosaic. Zhang et al. (2005) describe its construction. The raw data for the mosaic comes from full-volume scans at the radar sites indicated in Fig. 2.110. Note that radars in southern Canada are included. The first step is to remove non-meteorological echoes using multiple sensor quality control. As a simple example, if a radar echo appears at a location where the GOES satellite indicates clear skies, the echo is removed. Five-minute reflectivity data in spherical coordinates

(range, azimuth, elevation angle) at each radar site are mapped to 3-D Cartesian coordinates (x–east, y–north, z–up) at roughly 1-km resolution in the horizontal and 31 levels in the vertical, ranging from 500 to 18,000 m MSL.

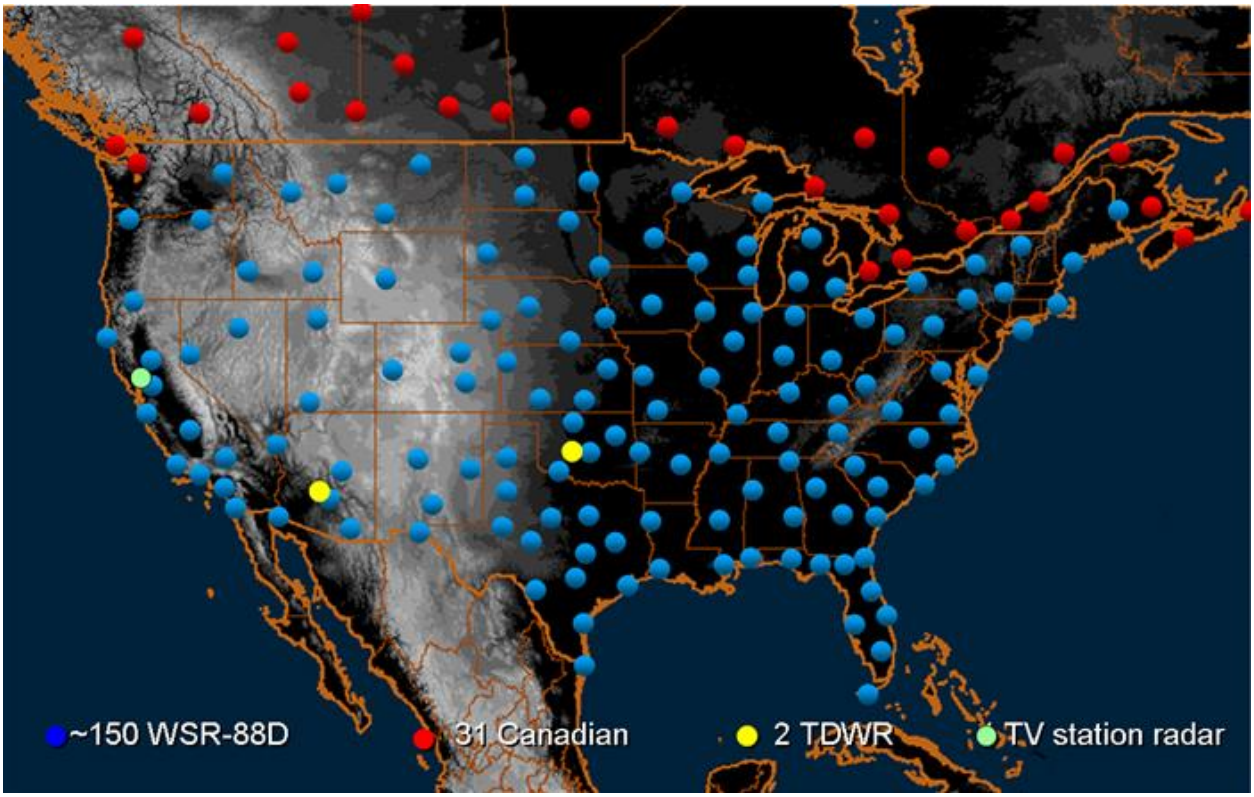

Fig. 2.110. Location of radars included in the MRMS 3-D radar mosaic. TDWR means Terminal Doppler Weather Radar.

The interpolation of reflectivity data from spherical to Cartesian coordinates for each full-volume scan proceeds as follows. To obtain the value analyzed at a given grid cell i in Cartesian coordinates f_i^a,

1) Compute the range, azimuth, and elevation at the center of the grid cell i.
2) Find two observed values, f_1^o and f_2^o, on the two adjacent tilts below and above the grid cell, respectively, and at the radar range and azimuth nearest those for the grid cell.
3) Compute the analyzed value from

$$f_i^a = (w_1 f_1^o + w_2 f_2^o)/(w_1 + w_2)$$

w_1 and w_2 are the interpolation weights assigned to the reflectivity observations below and above the grid cell, respectively. They are computed from

$$w_2 = \frac{\theta_i - \theta_1^o}{\theta_2^o - \theta_1^o} \quad \text{and} \quad w_1 = \frac{\theta_2^o - \theta_i}{\theta_2^o - \theta_1^o} \ ,$$

where θ_i, θ_1^o, and θ_2^o are the elevation angles of the grid cell and the radar bins below and above, respectively. The weights sum to one. This procedure amounts to linear interpolation in the vertical and a nearest neighbor approach in range and azimuth. Voids will exist wherever reflectivity values are lacking in the immediate vicinity of the grid cell.

The radars in Fig. 2.110 have overlapping coverage, and so the Cartesian grids, whose edges are overlapping range circles, must be merged. If reflectivity values from several different radars happen to fall at the same Cartesian grid location, then the value belonging to the closest radar receives the highest weight, values from more distant radars receive proportionally smaller weights, and the weighted average applies at the grid location.

Figure 2.111 shows a sample radar reflectivity mosaic at 3,000 m MSL. Other products are available at http://nmq.ou.edu (click on "Mosaic3D Levels") and http://wdssii.nssl.noaa.gov/web/wdss2/products/radar/index.shtml (choose products under "Multi-Sensor").

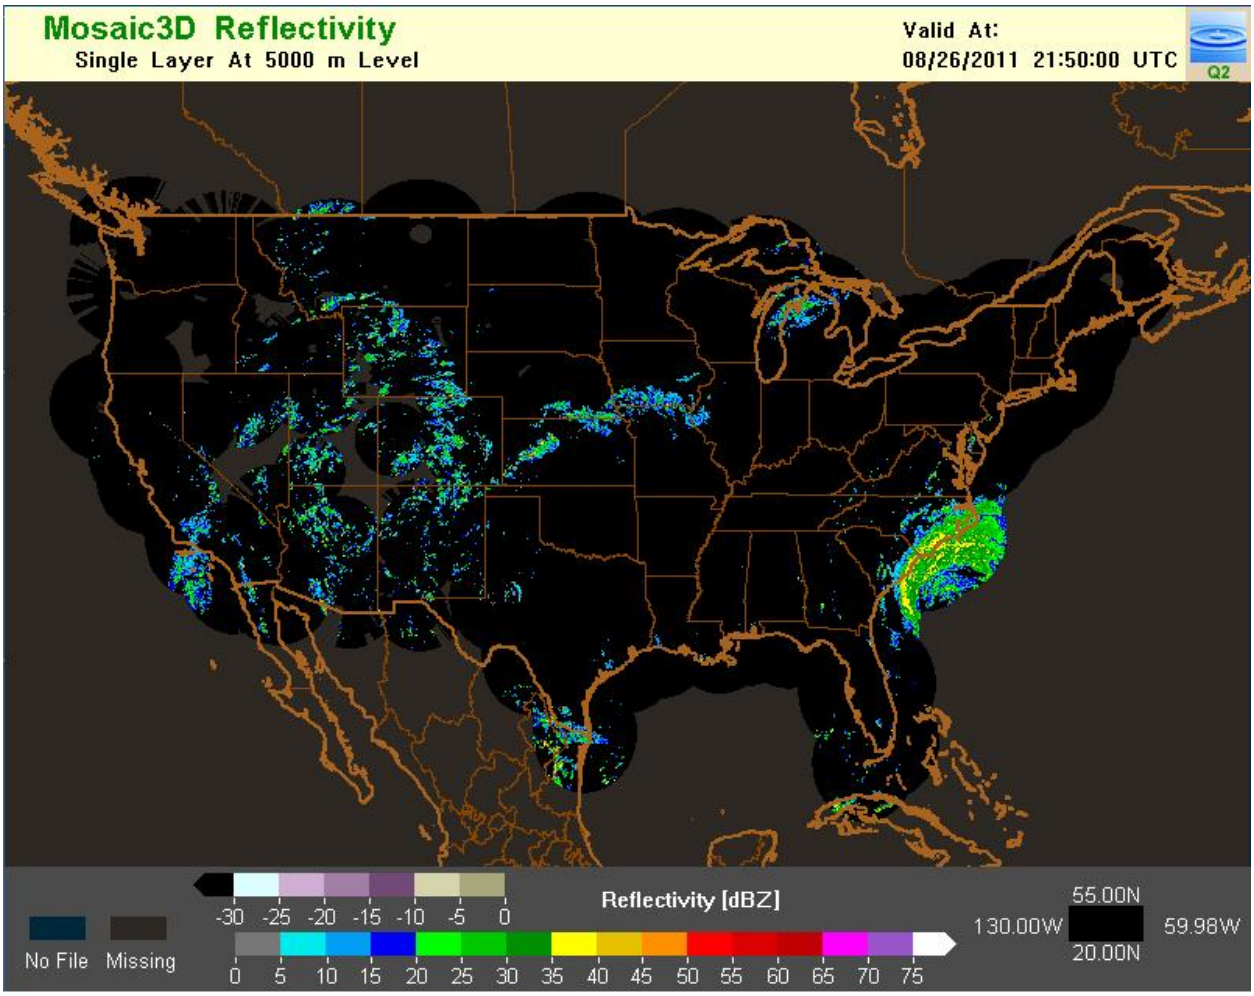

Fig. 2.111. National mosaic of radar reflectivity at 5,000 m MSL for 2150 GMT 26 Aug 2011. Rain bands from Hurricane Irene are visible along the coasts of North and South Carolina. Scattered afternoon convection dots the Great Plains and the Intermountain West. The color bar at bottom gives reflectivity values. Where the background is dark gray instead of black (difficult to see), radar coverage is lacking: over parts of the West, southern Canada, Mexico, and the oceans.

Having a national radar mosaic at 31 levels constitutes a 3-D data cube, an entity touted by the NextGen program. This particular data cube is already the means whereby reflectivity data are made available to the RUC and HRRR models, which supply analyses and predictions for the majority of applications discussed in this survey. The 3-D reflectivity grid is used to derive a 3-D latent heating field that helps ensure not only that precipitating clouds are in the right places at the initial time of the forecast, but that they endure as the forecast proceeds. RUC and HRRR precipitation forecasts have improved because of the assimilation of reflectivity data.

Many aviation weather products already require processing of Doppler radar data. The processing is accomplished in a variety of ways by several organizations. The MRMS algorithm that generates the reflectivity cube was developed by an organization that excels in building and operating radars, and in interpreting radar data. Perhaps *the* question to ask is to what extent the MRMS algorithm should supplant other algorithms that process radar data.

3. Hazards Near the Ground

This major section describes aviation weather products pertaining to hazards encountered on takeoff, final approach, or during ground operations at airports. The subtopics covered below are:

- Low ceiling and visibility
- Aircraft de-icing prior to departure
- Lightning and airport operations
- Low-level wind shear (microbursts)
- CWSU continuous weather watch and weather-related changes in acceptance rates

3.1 Low Ceiling and Visibility

The ceiling is the height above the surface of the base of the lowest layer of clouds that obscures more than half the sky. If the sky is obscured (when features of the cloud base are not discernable), the ceiling is defined as the vertical visibility into the obscuring phenomena aloft.

According to the *Glossary of Meteorology* (Glickman, 2000) visibility is the greatest distance in a given direction at which it is just possible to see and identify with the unaided eye 1) in the daytime, a prominent dark object against the sky at the horizon and 2) at night, a known, preferably unfocused, moderately intense light source.

A study of National Transportation Safety Board statistics indicated that low ceilings and poor visibility were contributing factors in 24% of all general aviation accidents between 1989 and and early 1997 (Kulesa, 2002). They were also cited as contributing factors in 37% of commuter/air taxi accidents during the same period. Such accidents occur when pilots rated for VFR flight encounter IFR conditions or when they are flying aircraft not equipped with instrumentation for “flying blind.” Commercial aircraft accidents due to low ceilings and poor

visibility are relatively rare because pilots are experienced, and instrument-assisted landings are commonplace under adverse conditions.

Low ceilings and poor visibility are not just a safety issue. The acceptance rates for aircraft landing at major airports are reduced under these conditions. This can lead to airborne or ground delays that result in diversions, cancellations, missed connections, and extra operational costs (Kulesa, 2002). About one-third of all weather-related delays are caused by low ceilings or poor visibility in the terminal area.

Table 3.1 defines categories of flight rules which apply to ceiling and visibility conditions. These categories will be mentioned frequently in this section.

Table 3.1
Definition of Flight Rules Categories. AGL means above ground level.
LIFR is a subcategory of IFR; MVFR is a subcategory of VFR.

Category	**Ceiling (Cig)**		**Visibility (Vsby)**
Low Instrument Flight Rules (LIFR)	Cig $<$ 500 ft AGL	and/or	Vsby $<$ 1 mile
Instrument Flight Rules (IFR)	500 $\leq$ Cig $<$ 1000 ft AGL	and/or	1 $\leq$ Vsby $<$ 3 miles
Marginal Visual Flight Rules (MVFR)	1000 $\leq$ Cig $\leq$ 3000 ft AGL	and/or	3 $\leq$ Vsby $\leq$ 5 miles
Visual Flight Rules (VFR)	Cig $>$ 3000 ft AGL	and	Vsby $>$ 5 miles

3.1.1 The Area Forecast (FA)

An Area Forecast (FA) is a forecast of VFR cloud and weather conditions for specific regions, normally covering several states. See Fig.3.1. FAs for the contiguous states, the Gulf of Mexico, and the Caribbean are issued three times daily by forecasters at the Aviation Weather Center in Kansas City. FAs for Alaska and Hawaii are issued by National Weather Service offices in those states.

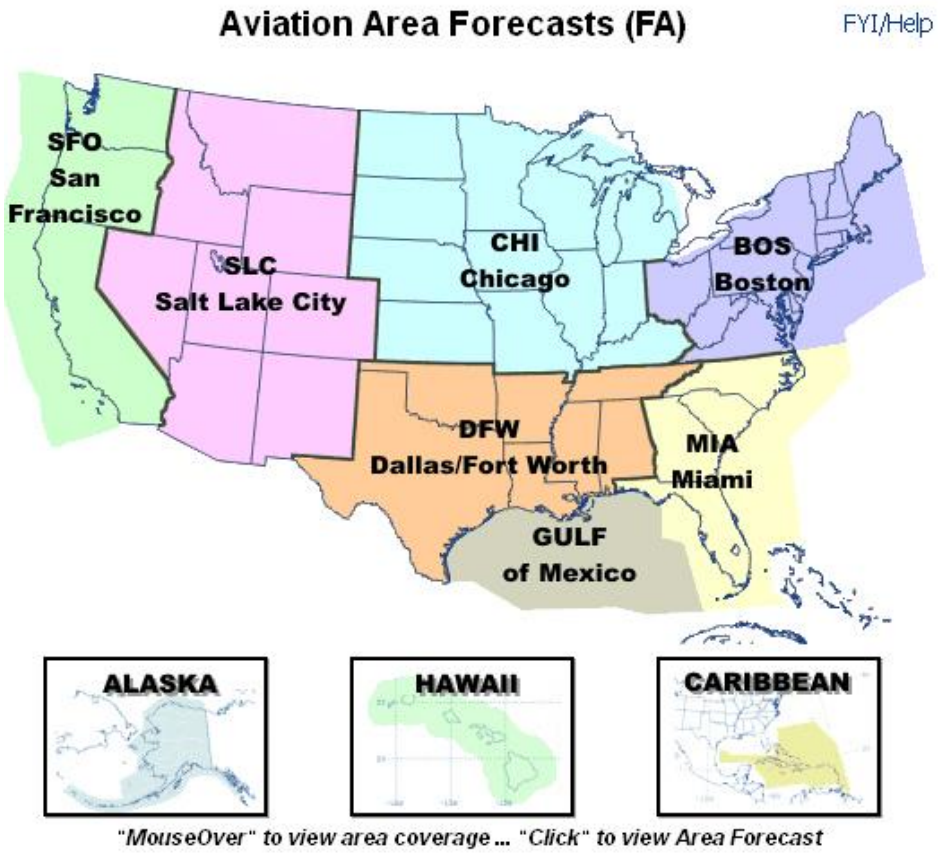

Fig.3.1. Regions for which Area Forecasts are issued.

It is vital that FAs be used *in conjunction* with AIRMET Sierra (IFR) bulletins (discussed in the next section) because IFR conditions are quite frequently embedded within regions of VFR conditions. The FAs offer guidance on flight conditions en route and interpolate conditions at airports for which no terminal forecasts (TAFs) are issued. TAFs will also be discussed later.

To see an FA, go to http://www.aviationweather.gov/adds/tafs/, click on "Area Forecast (FA) Product,' then click on region of interest.
Here is a sample Chicago Area Forecast (FA), issued at 0953 GMT, 28 September 2011. Many contractions and abbreviations are used, so a rough "translation" follows each subsection.

```
FAUS43 KKCI 280953 CCA
FA3W
CHIC FA 280953
SYNOPSIS AND VFR CLDS/WX
SYNOPSIS VALID UNTIL 290400
CLDS/WX VALID UNTIL 282200...OTLK VALID 282200-290400
ND SD NE KS MN IA MO WI LM LS MI LH IL IN KY
```

Chicago FA
Weather synopsis and information about VFR conditions, clouds and weather.
Synopsis is valid until 29 September [2011] at 0400 GMT.
Information about clouds and weather is valid until 28 September at 2200 GMT.
The outlook is valid from 28 September at 2200 GMT until 29 September at 0400 GMT.
The states covered by this FA are listed with the usual abbreviations and are also shown in Fig. 3.1 in light blue. In addition, LM is Lake Michigan, LS is Lake Superior, and LH is Lake Huron.

```
SEE AIRMET SIERRA FOR IFR CONDS AND MTN OBSCN.
TS IMPLY SEV OR GTR TURB SEV ICE LLWS AND IFR CONDS.
NON MSL HGTS DENOTED BY AGL OR CIG.
```

See AIRMET Sierra for IFR conditions and mountain obscuration.
Thunderstorms imply severe or greater turbulence, severe icing, low-level wind shear, and IFR conditions. Cloud heights that are not about mean sea level (MSL) and denoted by AGL (above ground level) or CIG (ceiling).

```
SYNOPSIS...CDFNT NERN ND-XTRM NWRN SD FCST 04Z NERN MN-NWRN IA-
SWRN NEB-NERN CO.
```

Synopsis: A cold front stretches from northeastern North Dakota to extreme northwestern South Dakota. By 0400 GMT it should stretch from northeastern Minnesota to northwestern Iowa, to southwestern Nebraska to northeastern Colorado.

The full FA contains VFR forecasts for every state in the Chicago region. For brevity, only three states are represented in this sample.

```
MN
W HLF...SKC. VIS 3-5SM BR. 15Z SCT CI. OTLK...VFR.
NERN...SCT010 SCT CI. VIS 3SM BR. 15Z SCT080-100. OTLK...VFR.
SERN...OVC020 LYRD FL250. VIS 3SM BR. 15Z BKN050. OTLK...VFR.
```

Minnesota: Western half, sky clear, visibility 3-5 statute miles (SM) in mist (BR). After 0400 GMT, scattered cirrus clouds. Outlook (after 2200 GMT, 28 September) is for VFR.
Northeastern: scattered clouds at 1,000 ft AGL, scattered cirrus clouds. Visibility 3 miles in haze. After 1500 GMT, scattered clouds 8,000–10,000 ft AGL. Outlook after 2200 GMT: VFR.
Southeastern: Overcast clouds with bases at 2,000 ft AGL, then layered clouds up to flight level 25,000 ft. Visibility 3 miles in haze. At 1500 GMT, broken clouds at 5,000 ft. After 2200 GMT, VFR conditions.

```
WI
SERN PTNS...BKN025 LYRD 130 BKN CI. OCNL VIS 3-5SM BR. ISOL -
```

```
SHRA. AFT 13Z SCT -SHRA. OTLK...VFR.
RMNDR...OVC020 TOP 080-100. OCNL VIS 3-5SM BR. 15Z SCT035 BKN100
LYRD FL200. OTLK...VFR.
```

Wisconsin: Southeastern portions, broken clouds with bases at 2,500 ft AGL, then clouds layered to 13,000 ft with broken cirrus clouds above that. Occasional visibility of 3-5 miles in haze. Isolated light rainshowers. After 1300 GMT, scattered light rainshowers. After 2200 GMT, VFR conditions. Remainder of state, overcast clouds at 2,000 ft AGL, cloud tops 8,000–10,000 ft. Visibility occasionally 3–5 miles in haze. At 1500 GMT, scattered clouds at 3,500 ft, broken clouds at 10,000 ft, layered up to flight level 20,000 ft. VFR conditions after 2200 GMT.

```
.
LS UPR MI
OVC010-020 LYRD FL200. VIS 3SM -DZ/-RA BR. 16Z BKN015. 19Z SCT035
SCT CI. OTLK...VFR.
```

Lake Superior and upper Michigan: Overcast at 1,000-2,000 ft AGL, the layered clouds to flight level 20,000 ft. Visibility 3 miles in light drizzle, light rain, and haze. By 1600 GMT, broken clouds at 1,500 ft AGL. By 1900 GMT, scattered clouds at 3,500 ft AGL and scattered cirrus clouds. Outlook (after 2200 GMT, VFR conditions.

```
.
LM LWR MI LH
N HLF...OVC020 LYRD FL200. VIS 3SM BR. 15Z BKN050-060.
     OTLK...VFR.
S HLF...OVC020-030 LYRD FL240. VIS 3-5SM -RA BR. 18Z OVC050.
  OTLK...VFR.
```

Lake Michigan, lower Michigan, and Lake Huron: northern half, overcast clouds at 2,000 ft AGL, then layered clouds up to flight level 20,000 ft. Visibility 3 miles in haze. By 1500 GMT, broken clouds with bases from 5,000 to 6,000 ft AGL. VFR conditions after 2200 GMT. Southern half, overcast clouds at 2,000-3,000 ft AGL, then layered clouds up to 24,000 ft AGL. Visibility 3-5 miles in light rain and haze. By 1800 GMT, overcast clouds at 5,000 ft AGL. VFR conditions after 2200 GMT.

3.1.2 AIRMETS and Graphical AIRMETS

As noted above, FAs must be used *along with* AIRMETs because the latter pertain to IFR conditions, for which the need for pilot training and specialized instrumentation is greater. AIRMETS come in the form of either a text message or a graphical product on the web (G-AIRMETs). To see a text AIRMET, go to http://www.aviationweather.gov/adds/airmets/. Under "Text AIRMETs/SIGMETs," click on the boxes labeled "IFR" and the choose a region, say "CHI." The example below is for the Chicago Region (CHI). The "SIERRA" designation on an AIRMET indicates that it pertains to IFR conditions. All three AIRMETs below were issued at 1445 GMT, 28 September 2011 (the month and year are not specified in the message) and are

valid until 2100 GMT on 28 September. Plain language translations follow each AIRMET.

```
WAUS43 KKCI 281445
CHIS WA 281445
AIRMET SIERRA UPDT 3 FOR IFR VALID UNTIL 282100
AIRMET IFR...LM MI IN
FROM 20WNW MKG TO 40SSW FNT TO FWA TO 20WNW IND TO 40SSE ORD TO
20WNW MKG
CIG BLW 010/VIS BLW 3SM PCPN/BR. CONDS ENDG 15-18Z.
```

This AIRMET is for parts of Lake Michigan, Michigan, and Indiana and pertains to an area bounded by a polygon described by these locations: 20 miles WNW of Muskegon, Michigan, to 40 miles SSW of Flint, Michigan, to Fort Wayne, Indiana, to 20 miles WNW of Indianapolis, Indiana, to 40 miles SSE of Chicago O'Hare Airport, to 20 miles WNW of Muskegon, the starting point. The expected conditions are a ceiling below 1,000 ft AGL, visibility below 3 statute miles in precipitation and haze. Conditions are expected to end between 1500 and 1800 GMT.

```
WAUS43 KKCI 281445
CHIS WA 281445
AIRMET SIERRA UPDT 3 FOR IFR VALID UNTIL 282100
AIRMET IFR...MN IA WI LS MI IL
FROM 50ESE YQT TO 60NW SSM TO SAW TO 60SSW RHI TO 20ENE UIN TO
40WSW IOW TO 50SSE MSP TO 80E DLH TO 50ESE YQT
CIG BLW 010/VIS BLW 3SM PCPN/BR. CONDS ENDG 15-18Z.
```

This AIRMET is for parts of Minnesota, Iowa, Wisconsin, Lake Superior, Michigan and Illinois. IFR conditions are expected within a polygon defined by these locations: 50 miles ESE of Thunder Bay, Ontario, to 60 miles NW of Sault Ste Marie, Michigan, to Marquette, Michigan, to 60 miles SSW of Rhinelander, Wisconsin, to 20 miles ENE of Quincy, Illinois, to 40 miles WSW of Iowa City, Iowa, to 50 miles SSE of Minneapolis–St. Paul, Minnesota, to 80 miles E of Duluth, Minnesota to 50 miles ESE of Thunder Bay, the starting point. The expected conditions are a ceiling below 1,000 ft AGL and visibility below 3 statute miles in precipitation and mist. These conditions will end between 1500 and 1800 GMT.

```
WAUS43 KKCI 281445
CHIS WA 281445
AIRMET SIERRA UPDT 3 FOR IFR VALID UNTIL 282100
AIRMET IFR...WI LM LS MI LH
FROM 40ENE SAW TO SSM TO 50ESE SSM TO 30WNW ASP TO 60WSW SSM TO
40ENE GRB TO 40ENE SAW
CIG BLW 010/VIS BLW 3SM BR. CONDS ENDG 15-18Z.
```

This AIRMET is for parts of Wisconsin, Lake Michigan, Lake Superior, Michigan, and Lake Huron. IFR conditions are expected within a polygon defined by these locations: 40 ENE of Marquette, Michigan, to Sault Ste. Marie, Michigan, to 50 miles ESE of Sault Ste. Marie, to 30 miles WNW of Oscoda, Michigan, to 60 miles WSW of Sault Ste. Marie to 40 miles ENE of Green Bay, Wisconsin, to 40 miles ENE of Marquette, the starting point. The expected conditions are

ceiling below 1,000 ft AGL and visibility below 3 statute miles in mist. These conditions will end between 1500 and 1800 GMT.

Comment: The text AIRMET is an anachronism. Most probably, it continues to be generated because some pilots still obtain voice briefings, but the nearly ubiquitous use of web displays on mobile phones and portable computers is rapidly making this product obsolete. How many pilots or meteorologists are familiar with all the call signs and, of these, how many can visualize the polygons described in these messages? Very few, if any. Of far greater utility is the Graphical Airmet (G-AIRMET), which plots this information on a background map. Figure 3.2 is an example. The three polygons plotted in the Upper Midwest and Great Lakes are the same as those described in the text message above. Clearly, the graphical information is far easier to assimilate than the audio information.

To see a current product, to
http://www.aviationweather.gov/products/gairmet/display_gairmet.php.
Find time series of "Sierra Graphics," AIRMETs for IFR conditions at various forecast times, and click on desired time for a static image.

Fig. 3.2. A graphical AIRMET for Instrument Flight Rule (IFR) conditions, valid at 1500 GMT, 28 Sep 2011. The three polygons in the Upper Midwest and covering parts of the Great Lakes were also discussed as a text file above. The abbreviations BR, FG, PCPN, and CLDS mean mist, fog, precipitation, and clouds, respectively. The dashed pink outline over parts of the

Appalachians and the symbols in pink resembling two side-by-side "A"s with the upper parts of the letters filled in designate mountain obscuration.

Verification data for Sierra-AIRMETs (those for IFR conditions) are available at http://rtvs.noaa.gov/cv/stats/. To obtain the statistics immediately below in Table 3.2, choose the product "AIRMETs," a beginning date of 01 March 2011, and end date of 30 April 2011, and output in the form of "Summary Table." A springtime period was chosen because of the prevalence of low ceilings are poor visibility. The contingency table is at left (refer back to Section 1.4 on verification). The corresponding event counts are at right. "YY" means IFR conditions were predicted, and IFR conditions occurred. "YN" means IFR conditions were predicted but did not occur, and so on.

Table 3.2

A contingency table for verification of IFR AIRMETs issued for 6 h in the future by AWC from 01 March through 30 April 2011 anywhere in the contiguous 48 states. Surface METAR observations verify the AIRMETs.

	IFR Conditions Observed		
IFR Conditions Predicted		Yes	No
	Yes	220,721	323,080
	No	57,433	1,960,056

From this table, one finds that the fraction of IFR events forecast correctly (PODy) is 0.79, the fraction of non-IFR events forecast correctly (PODn) is 0.86, the false alarm rate (FAR) is 0.59, and the critical success index (CSI) is 0.37. To gauge the frequency of IFR conditions, it is useful to know that 17.5% of the CONUS was covered by IFR conditions during March–April of 2011.

3.1.3 SIGMETs

With regard to ceiling and visibility, the only SIGMET issued is for a dust storm or sand storm that lowers visibility below three miles, either at the ground or at flight level. Because these conditions must apply within an area of 3,000 mi^2 or larger, such SIGMETs are rare. The

haboobs that afflicted Phoenix, Arizona, multiple times during the summer of 2011, raising dust and sand many of thousands of feet into the air probably qualified for a SIGMET. Advisories for haboobs affecting areas smaller than 3,000 mi^2 would be issued by a CWSU as a Center Weather Advisory.

3.1.4 TAFs

The official government document describing TAFs (Terminal Aerodrome Forecasts) is available at http://www.nws.noaa.gov/directives/. Click on "National Weather Service Instruction 10.813." The latest version is dated 14 December 2010. TAFs specify expected meteorological conditions significant to aviation at airports for a specified period, in the U.S. for 24 hours (30 hours at most airports handling international traffic). A TAF is a *primary* weather product; it includes a forecast of surface wind, visibility, weather (rain, snow, etc.), obstructions to vision, clouds, low-level wind shear, and any significant changes in one or more of these elements during the forecast period.

TAFs are important because they affect the movement of air traffic. In particular, if conditions at the time of landing are predicted to be below minimums for a particular airport or aircraft, the pilot may neither take off for that airport nor use that airport as an alternative landing site when filing a flight plan. Not only that, but as conditions deteriorate, acceptance rates for incoming traffic decrease, some departing flights may be temporarily held, and delays accumulate. Many decisions relating to air traffic control are based upon TAFs.

The TAF applies to the terminal area, defined as a circle of radius 5 miles, centered on the airport's runway complex. A TAF must also note fog, showers, or thunderstorms expected to occur between 5 and 10 miles of the center point.

NWS meteorologists at WFOs issue TAFs four times a day based upon surface, rawinsonde, profiler, satellite, and radar observations and short-term predictions by numerical models. In addition they rely on probability information provided by LAMP, to be discussed later. Each WFO issues TAFs for up to ten airports within its area of responsibility.

The flight categories listed in Table 3.1 are relevant for amending TAFs, when they no longer represent actual conditions. To these categories, the NWS has added one more, the Very Low IFR (VLIFR) category: ceiling < 200 ft and visibility < ½ mile. Whenever the observed condition passes between categories and the change was not forecast, the TAF is amended.

To get a TAF, go to http://www.aviationweather.gov/adds/tafs/index.php and type in the four-character ICAO identifier for the desired airport. In the contiguous U.S., a "K" is always the first character. The next three characters are the familiar call sign. For example, typing "KSTL" will call up the TAF for St. Louis, Missouri (Lambert Field). (ICAO stands for the International Civil Aviation Organization, which sets rules for, and keeps track of, station identifiers.)

Here is a sample TAF for Bangor, Maine, covering from 1800 GMT, 04 October 2011 to 1800 GMT 05 October 2011. The official TAF is in bold; the plain-language translation is immediately below. Note that the month and year are not given in the first line of the TAF since this information will be obvious to the user.

Forecast for: KBGR (BANGOR , ME, US)
Text: **KBGR 041730Z 0418/0518 01012KT 2SM -RA BR OVC008**
Forecast period: 1800 UTC 04 October 2011 to 0300 UTC 05 October 2011
Forecast: Winds: from the N (10 degrees) at 14 MPH (12 knots; 6.2 m/s) **Visibility:** 2.00 miles (3.22 km) **Ceiling:** 800 feet AGL **Clouds:** overcast cloud deck at 800 feet AGL **Weather:** -RA BR (light rain, mist)

Text: **FM050300 35011KT 1SM BR OVC003**
Forecast period: 0300 to 1200 UTC 05 October 2011
Forecast: Winds: from the N (350 degrees) at 13 MPH (11 knots; 5.7 m/s) **Visibility:** 1.00 miles (1.61 km) **Ceiling:** 300 feet AGL **Clouds:** overcast cloud deck at 300 feet AGL **Weather:** BR (mist)

Text: **FM051200 34009KT P6SM BKN010**
Forecast period: 1200 to 1400 UTC 05 October 2011
Forecast: Winds: from the NNW (340 degrees) at 10 MPH (9 knots; 4.7 m/s) **Visibility:** 6 or more miles (10+ km) **Ceiling:** 1000 feet AGL **Clouds:** broken clouds at 1000 feet AGL **Weather:** no significant weather forecast for this period

Text: **FM051400 34013KT P6SM BKN025**
Forecast period: 1400 to 1700 UTC 05 October 2011
Forecast: Winds: from the NNW (340 degrees) at 15 MPH (13 knots; 6.8 m/s) **Visibility:** 6 or more miles (10+ km) **Ceiling:** 2500 feet AGL **Clouds:** broken clouds at 2500 feet AGL **Weather:** no significant weather forecast for this period

Text: **FM051700 34017KT P6SM BKN030**

Forecast period: 1700 to 1800 UTC 05 October 2011
Forecast: Winds: from the NNW (340 degrees) at 20 MPH (17 knots; 8.8 m/s) **Visibility:** 6 or more miles (10+ km) **Ceiling:** 3000 feet AGL **Clouds:** broken clouds at 3000 feet AGL **Weather:** no significant weather forecast for this period

The preceding example is for a single station. To view TAFs for a given area, go to http://www.aviationweather.gov/adds/tafs/. Under "Plotted TAFs," click on "Prevailing" or "Temporary" conditions, then choose the area to view by clicking a city on the map below. A map of TAFs for the coming hour appears, but one can choose any other lead time from the pull-down menu. Figure 3.3 is a sample map of TAFs for the northeast U.S., valid at 2000 GMT, 04 October 2011. Conditions plotted for Bangor (BGR), in central Maine, match those indicated in the first line of the TAF printed for Bangor above. The symbols used in Fig. 3.3 are the same as those used in METAR plots, to be discussed below.

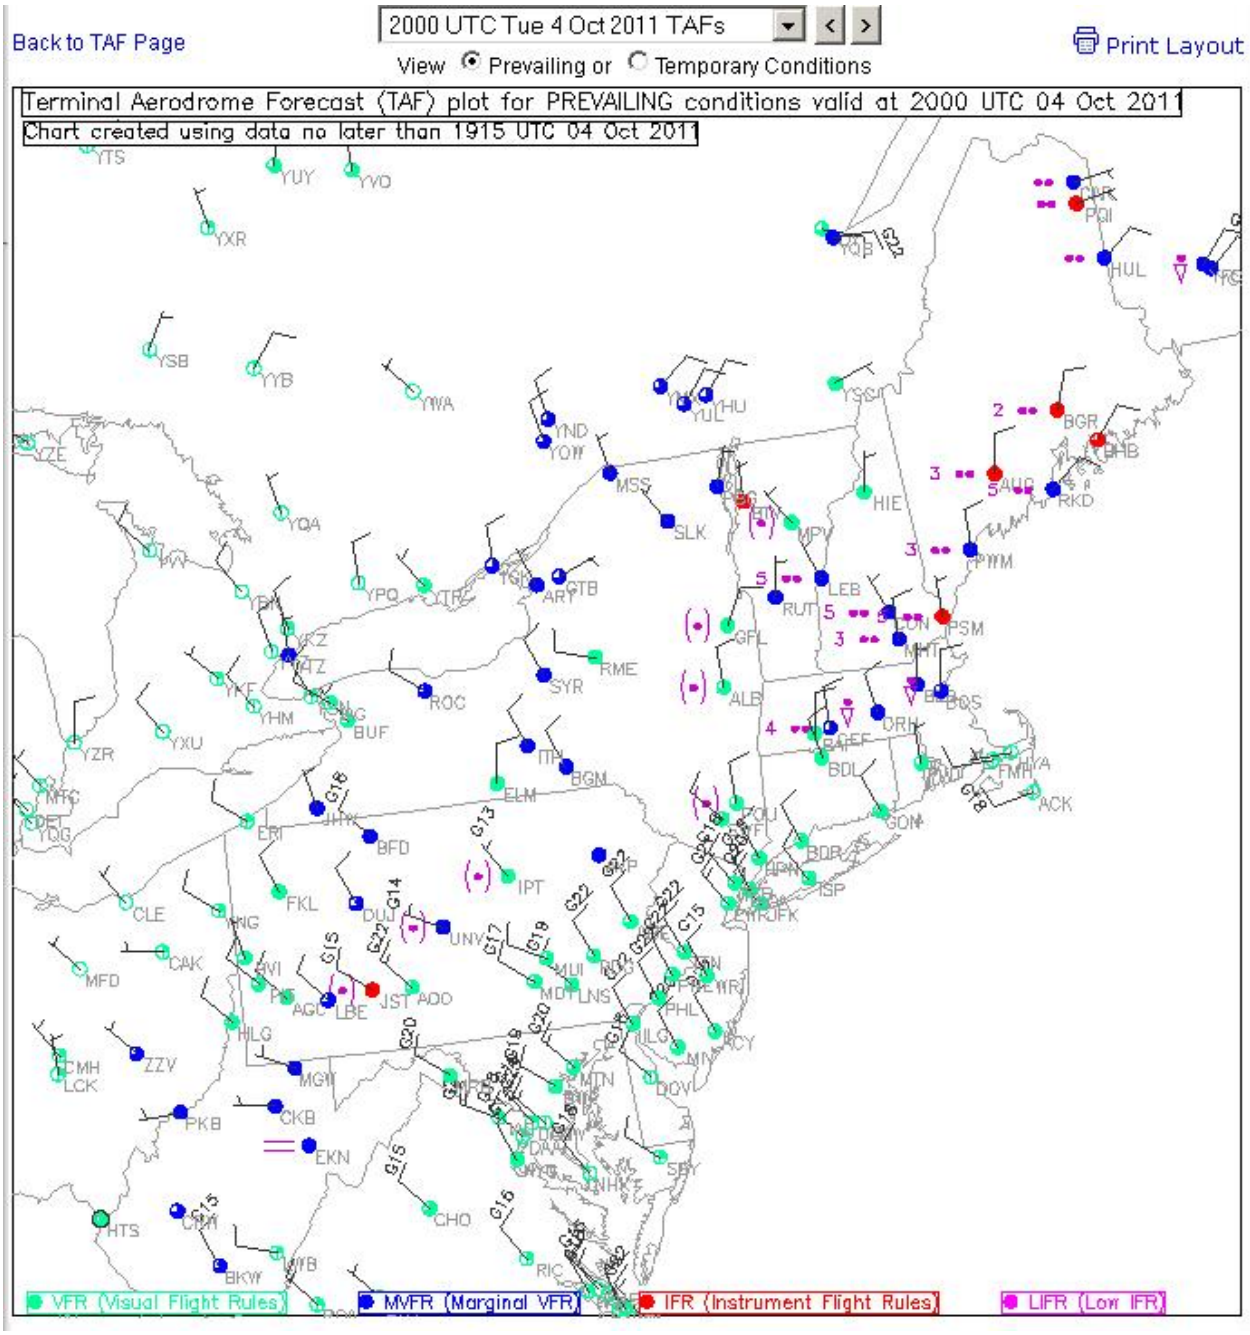

Fig. 3.3. A map of TAFs for the northeast U.S. valid at 2000 GMT, 04 October 2011. The station circles are colored to indicate the predicted flight category, as indicated at the bottom of the figure.

The NWS verifies TAFs routinely to provide feedback to the forecasters who generate them. Table 3.3 provides sample statistics for ceiling forecasts from an internal NWS website.

Table 3.3

A sample report from a NWS internal web site for verification of TAF ceiling forecasts. All TAFs issued within the listed states from 01 November 2010 through 30 April 2011 are included. The number of reports in each category, forecast and observed, is listed in the contingency table. The magnitude of the forecast error is color-coded as indicated. Data courtesy of Charles Kluepfel, NWS.

TAF Report Page

Element	Ceiling
TAF Type	Prevailing (Scheduled Only)
Date Range	11/01/2010 TO 04/30/2011
Months Selected	JAN, FEB, MAR, APR, NOV, DEC
State	IL, IN, MI, MN, OH, WI, PA
Ceiling Threshold	1000 FEET
Cycle Times	0000Z, 0600Z, 1200Z, 1800Z
Projections	>0 - 3, >3 - 6
Report Format	Hours | Minutes | Percent | Frequency | Toggle Legend

Hit	1-Category Error	2-Category Error	3-Category Error	4 or More Category Error

Show All Hide All

MULTICATEGORY CONTINGENCY TABLE WITH ASSOCIATED SCORES

OBS\FORECASTS	1	2	3	4	5	6	TOTAL
<200 (1)	6,709	9,874	8,078	3,435	1,490	6,267	35,853
200-400 (2)	2,056	32,587	43,409	17,478	5,395	9,147	110,072
500-900 (3)	1,276	20,636	115,883	80,589	24,948	23,899	267,231
1000-1900 (4)	912	8,508	73,091	275,744	120,764	81,325	560,344
2000-3000 (5)	346	2,283	18,606	108,112	232,608	154,370	516,325
>3000 (6)	1,377	5,762	26,326	108,612	235,671	2,658,463	3,036,211
TOTAL	12,676	79,650	285,393	593,970	620,876	2,933,471	4,526,036
BIAS	0.354	0.724	1.068	1.060	1.202	0.966	----

The parameters listed at the top of Table 3.3 were chosen to include the cold half of the year, November through April, when ceilings are generally lower, and normally cloudy locations—seven states bordering the Great Lakes. Only TAFs verifying within the first six hours are included. Verification is against METAR observations within the seven-state area.

The contingency table includes six ceiling categories. Hits are color-coded in bright green. Misses are also color-coded, with light green, yellow, orange, and red indicating increasingly serious errors of one, two, three, and four categories, respectively. The number of observations and forecasts in each category are summed at the end of each row and at the bottom of each column, respectively. The total number of TAFs exceeds 4.5 million. Note the bias scores in the last row of Table 3.3, which indicate that the lowest two ceiling categories are substantially underpredicted. For example, the number of predictions of Category 1 ceiling (< 200 ft) was 12,676; the observed number was 35,853. The ratio of these two numbers is the bias, 0.354. The ceiling category from 2,000 to 3,000 ft was somewhat overforecast (bias 1.202). For all other categories, the bias was close to one.

Table 3.4 shows a 2 x 2 contingency table for a ceiling threshold of 1,000 ft. This table is derived for the same dataset as Table 3.3.

Table 3.4

Verification data for TAF forecasts issued from 01 November 2010 through 30 April 2011 within the states of IL, IN, MI, MN, OH, PA, and WI with lead times up to six hours. The ceiling threshold for this table is 1,000 ft.

		Observed Ceiling		
Forecast Ceiling		< 1,000 ft	≥ 1,000 ft	Total
	< 1,000 ft	240,508	172,648	413,156
	≥ 1,000 ft	137,211	3,975,669	4,112,880
	Total	377,319	4,148,317	4,526,036

From the values in this table, one can derive a POD of 0.582, a FAR of 0.363, and a CSI of 0.437 for ceilings less than 1,000 ft.

The Real-Time Verification System (RTVS) of ESRL provides additional ways of examining similar data. At this web address, http://rtvs.noaa.gov/cv/taf_leadtime/, one finds a map of the U.S. with the locations of the nation's largest airports prominently displayed. At lower left, there is

a button for tutorial materials that include both text and video describing the products available at this site. The following figures illustrate a few of them.

Figure 3.4 is a zoomed version of the U.S. map, on which 47 major and mid-sized airports in seven states bordering the Great Lakes have been selected for TAF analysis. As in Table 3.3, the period of analysis is from 01 November 2010 through 30 April 2011. Whereas Table 3.3 considered *all* airports in these seven states, only the larger airports are considered here, so the total number of TAFs considered is much smaller. On the other hand, the entire TAF, out to 24 h, is included, not just the first 6 h.

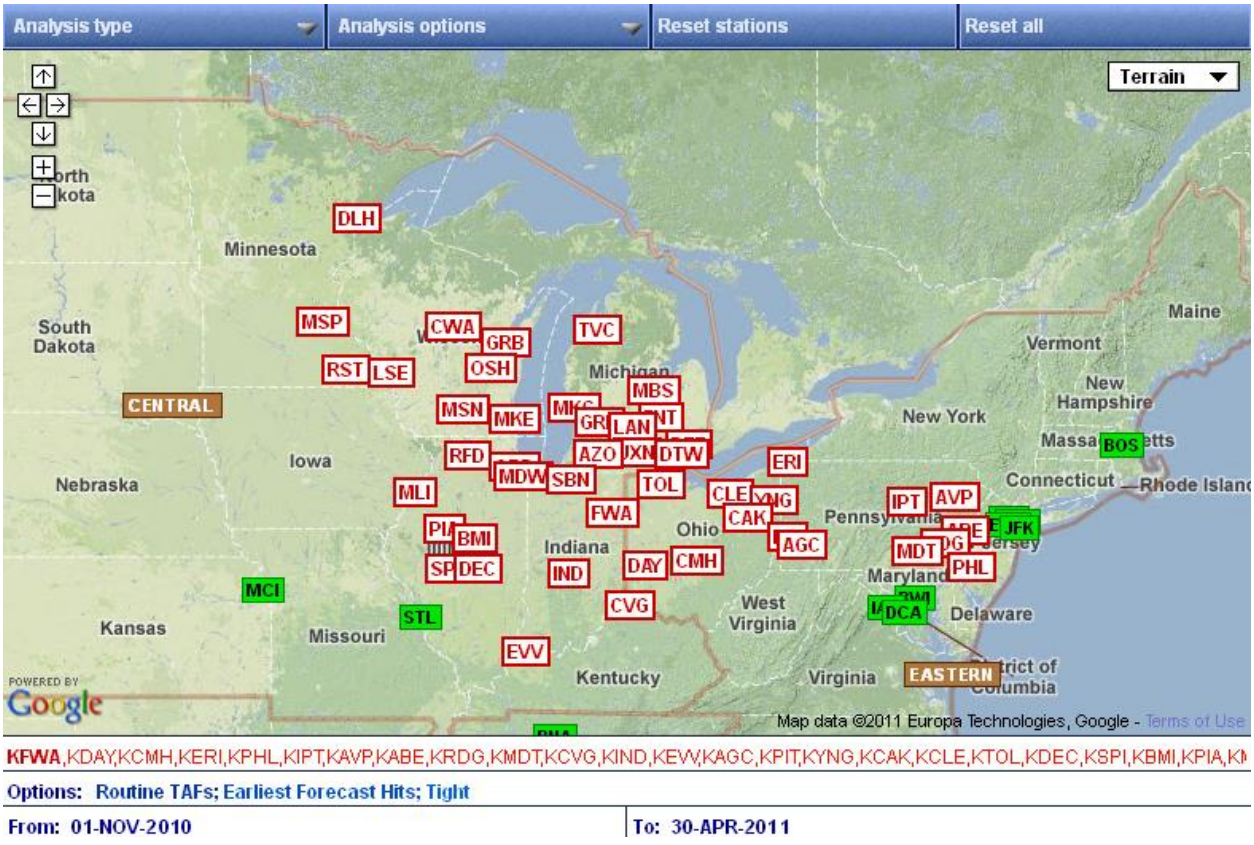

Fig. 3.4. Locations of 47 airports (in red rectangles) selected for TAF analysis. The analysis includes TAFs issued from 01 November 2010 through 30 April 2011.

Figures 3.5 compares the onset times of IFR conditions as observed at METAR sites, and as predicted in TAF forecasts. All onset times occurring within each 24-h TAF are included.

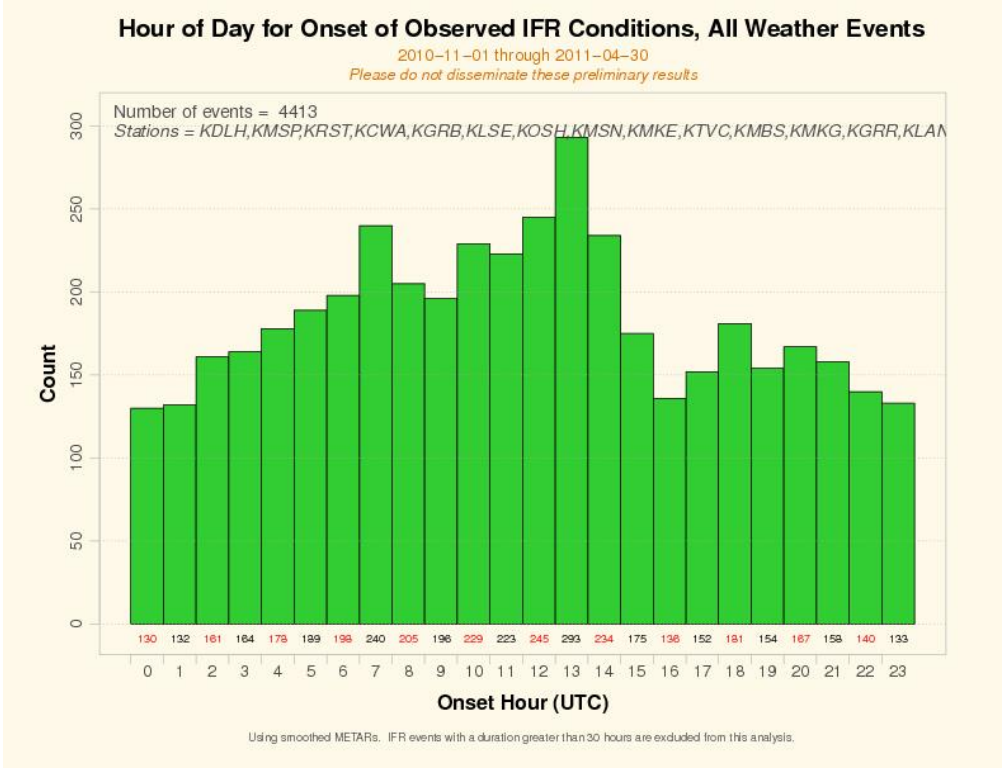
Hour of Day for Onset of Observed IFR Conditions, All Weather Events
2010-11-01 through 2011-04-30
Please do not disseminate these preliminary results
Number of events = 4413
Stations = KDLH,KMSP,KRST,KCWA,KGRB,KLSE,KOSH,KMSN,KMKE,KTVC,KMBS,KMKG,KGRR,KLA
Count
0
50
100
150
200
250
300
130 132 161 164 179 189 198 240 205 196 229 223 245 293 234 175 135 152 181 154 167 158 140 133
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Onset Hour (UTC)
Using smoothed METARs. IFR events with a duration greater than 30 hours are excluded from this analysis.

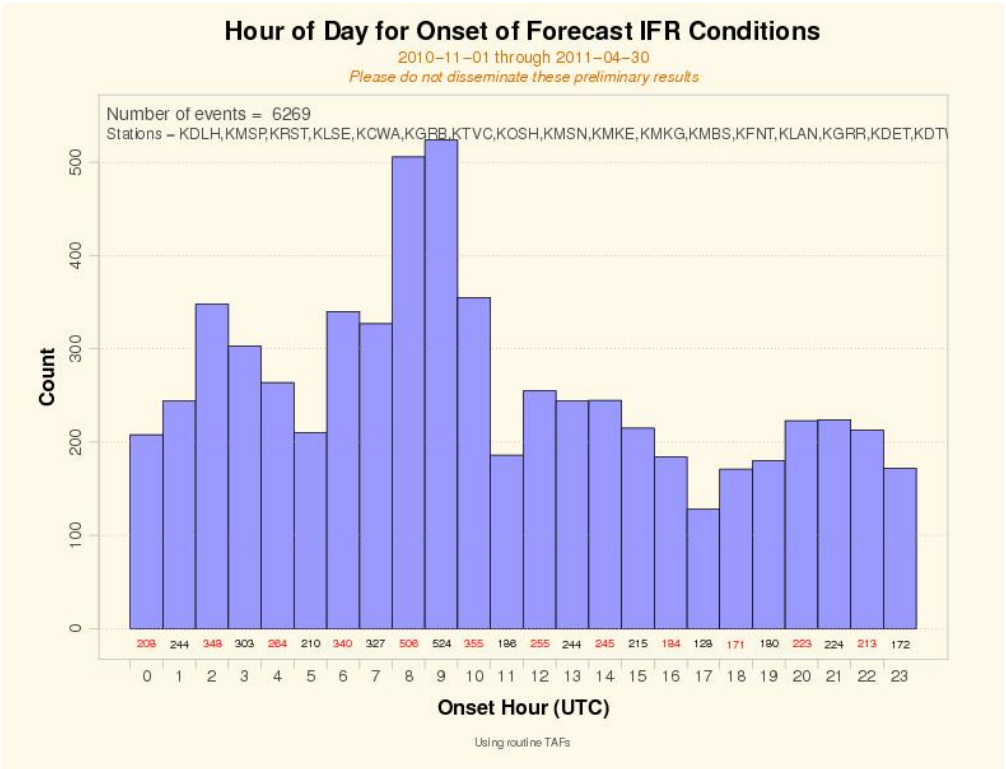
Hour of Day for Onset of Forecast IFR Conditions
2010-11-01 through 2011-04-30
Please do not disseminate these preliminary results
Number of events = 6269
Stations = KDLH,KMSP,KRST,KLSE,KCWA,KGRB,KTVC,KOSH,KMSN,KMKE,KMKG,KMBS,KFNT,KLAN,KGRR,KDET,KDT
Count
0
100
200
300
400
500
208 244 348 303 264 210 340 327 506 524 355 186 255 244 245 215 184 129 171 180 223 224 213 172
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23
Onset Hour (UTC)
Using routine TAFs

Fig. 3.5. Hour of onset of IFR conditions, top, as recorded in METAR observations and, bottom, as predicted in TAFs, November 2010 through April 2011 for 47 selected airports around the Great Lakes.

Aviation meteorologists predicted the onset of IFR conditions about 1.5 times more often than they occurred. Climatologically, IFR conditions are most likely to begin around 1200 and 1300 GMT, which is early in the day across the U.S. That is a favored time for ground fog and low stratus clouds, and poor surface visibility occasioned by fog and pollutants trapped below the nocturnal inversion. The distribution of onset times in the TAFs peaks several hours earlier, at 0800 and 0900 GMT. In addition, the distribution of actual onset times is considerably smoother than that of the predicted onset times. Unfortunately, it is anything but straightforward to bring the predicted distribution of onset times more in line with the observed distribution.

It is also instructive to compare observed durations of IFR events with predicted durations. Two distributions of IFR event durations appear in Fig. 3.6, one obtained from METARs (top) and the other from TAFs (bottom). Figure 3.6 uses the same database as Fig. 3.5. A sharp maximum in observed IFR events occurs at 1-2 h, and the observed distribution decreases rapidly with increasing duration. The distribution of predicted durations also has a sharp peak, but at 3-4 h. Moreover, for longer durations, it is much broader than the observed distribution. Evidently, the forecasts are biased toward longer durations than that actually occur.

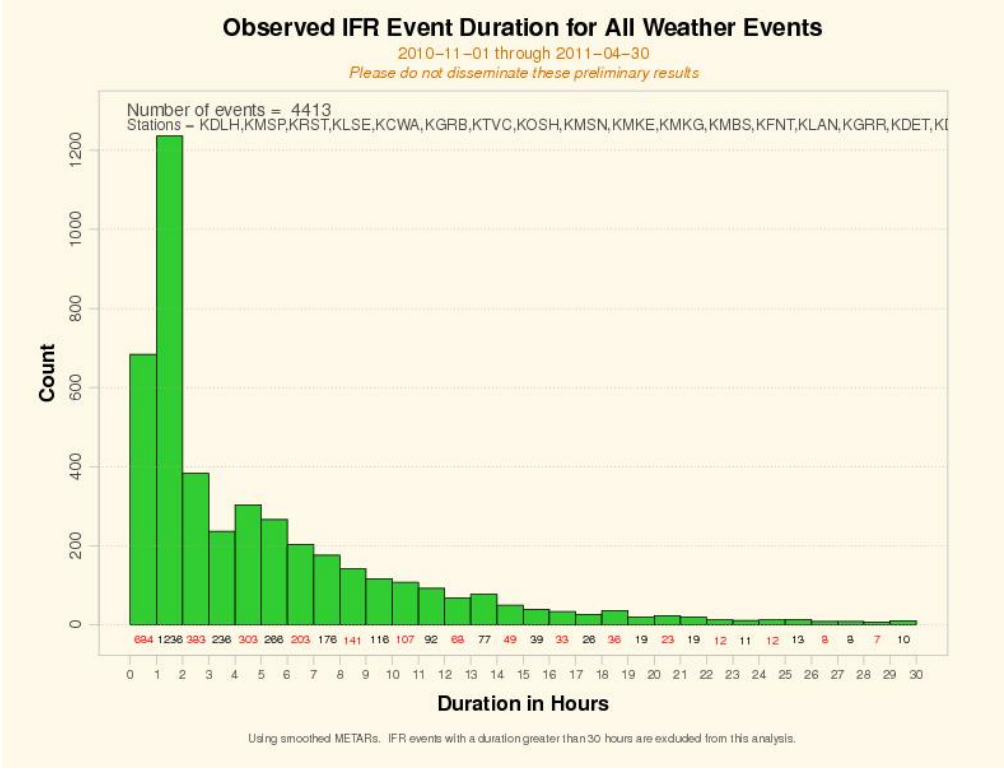
Observed IFR Event Duration for All Weather Events
2010–11–01 through 2011–04–30
Please do not disseminate these preliminary results
Number of events = 4413
Stations – KDLH,KMSP,KRST,KLSE,KCWA,KGRB,KTVC,KOSH,KMSN,KMKE,KMKG,KMBS,KFNT,KLAN,KGRR,KDET,KI
Count
0 200 400 600 800 1000 1200
684 1236 383 236 303 268 203 176 141 116 107 92 69 77 49 39 33 26 36 19 23 19 12 11 12 13 9 8 7 10
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30
Duration in Hours
Using smoothed METARs. IFR events with a duration greater than 30 hours are excluded from this analysis.

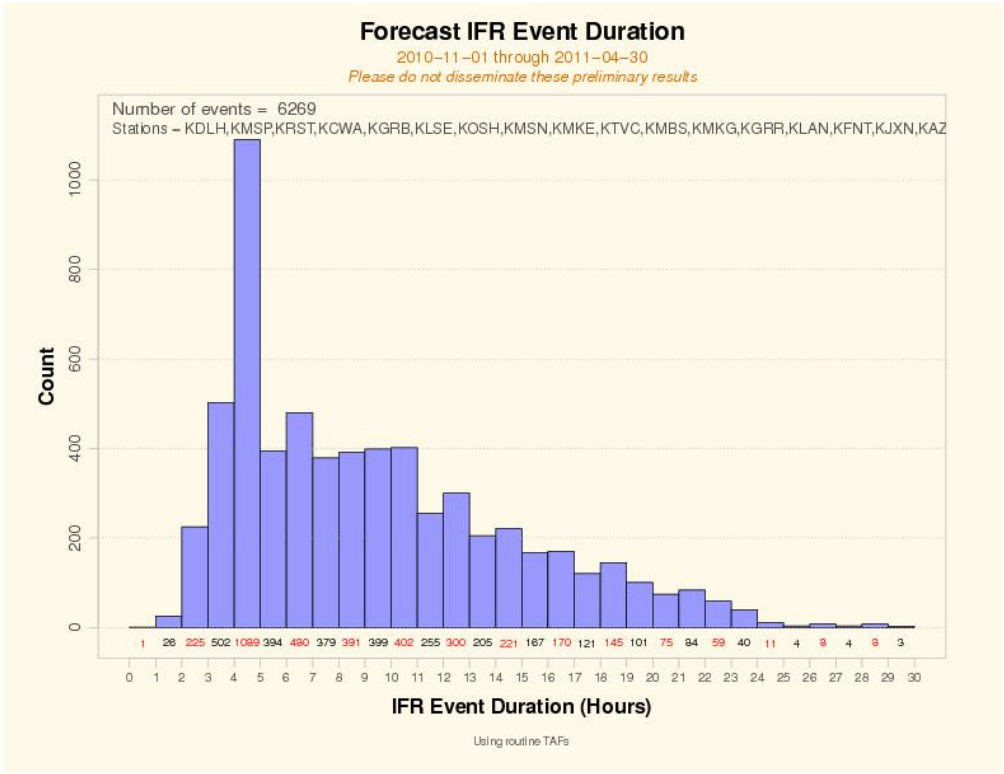
Forecast IFR Event Duration
2010–11–01 through 2011–04–30
Please do not disseminate these preliminary results
Number of events = 6269
Stations – KDLH,KMSP,KRST,KCWA,KGRB,KLSE,KOSH,KMSN,KMKE,KTVC,KMBS,KMKG,KGRR,KLAN,KFNT,KJXN,KAZ
Count
0 200 400 600 800 1000
1 26 225 502 1089 394 480 379 391 399 402 255 300 205 221 167 170 121 145 101 75 84 59 40 11 4 8 4 8 3
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30
IFR Event Duration (Hours)
Using routine TAFs

Fig. 3.6. The distribution of IFR event duration, in hours, for 47 airports around the Great Lakes. Top: obtained from METAR observations. Bottom: obtained from TAF forecasts. Event counts are near the top of each panel. IFR data for these histograms were collected from November 2010 through April 2011.

The RTVS website http://rtvs.noaa.gov/cv/taf_leadtime/ also has an "Event Viewer," illustrated here. Click on "Analysis Type" at the upper left and then on "Event Viewer" on the pull-down menu. This option allows examination of TAFs issued for a single airport for any period of up to one month long. Figure 3.7 shows two examples for Chicago's O'Hare Airport (KORD) from March 2011.

Two thin gray lines run horizontally across each panel. The bottom line is for IFR conditions, and the top line is for non-IFR conditions (which include MVFR and VFR). The bold green line indicates the observed flight conditions (IFR or non-IFR) and transitions between them, obtained from METARs. The bold red line indicates the same but for the TAF forecast. The little squares along the green line, indicate the times of METAR observations. The somewhat larger squares along the red line indicate divisions between the time blocks of the TAF. The time line, labeled by hour (GMT) and date, extends for 30 hours, the standard length of TAFs issued for O'Hare Airport. The TAF itself appears at lower left. Displays like this enable meteorologists to see precisely how well they timed transitions between IFR and non-IFR conditions.

The top panel of Fig. 3.7 illustrates a simple case. This routine [ROU] TAF runs from 0000 GMT on 09 March to 0600 GMT on 10 March 2011. It calls for rain most of the forecast period, becoming mixed with snow after 0000 GMT 10 March. IFR conditions occur when either the visibility is less than three miles or the ceiling is less than 1000 ft. The TAF indicates a transition to IFR conditions (a ceiling of broken clouds at 700 ft AGL) at 0600 GMT, three hours before it actually occurred. Similarly, the TAF indicated improvement to MVFR conditions at 2100 GMT (an overcast at 1,200 ft), six hours before this actually occurred. Given the state of the science, this is not a bad forecast. The predicted duration of IFR conditions was about right; the only fault was in the timing: predicted conditions deteriorated too soon and recovered too soon.

The bottom panel of Fig. 3.7 illustrates a more complicated situation. This amended [AMD] TAF is issued four hours after the original one and runs for 26 hours starting at 1000 GMT, 22 March 2011. Rain showers and thunderstorms are predicted during parts of the 26 hours, but IFR conditions are not expected until 2200 GMT, when the TEMPO group forecasts a visibility of 1½ miles in thunderstorms and a ceiling of 700 ft. By definition, the TEMPO conditions cannot last more than one hour and must occur for less than half the total time for which the TEMPO group is valid, in this case from 2200 GMT on 22 March to 0200 GMT on 23 March. The FROM (FM) group, bracketing the TEMPO group, forecasts MVFR conditions (a 1,500-ft overcast) from 2200 GMT on 22 March to 0300 GMT on 23 March. Since the worst conditions (those in the TEMPO group) determine whether a pilot can land or not, the TEMPO group explains why the bold red line is on the IFR line from 2200 to 0200 GMT but on the non-IFR line from 0200 to 0300 GMT.

After 0300 GMT, 23 March, the predicted ceiling at 400 ft means that conditions are in the IFR category through the end of the forecast. This forecast was obviously challenging. Considering the two categories—IFR and non-IFR—the TAF was in the wrong category most of the time after 1900 GMT.

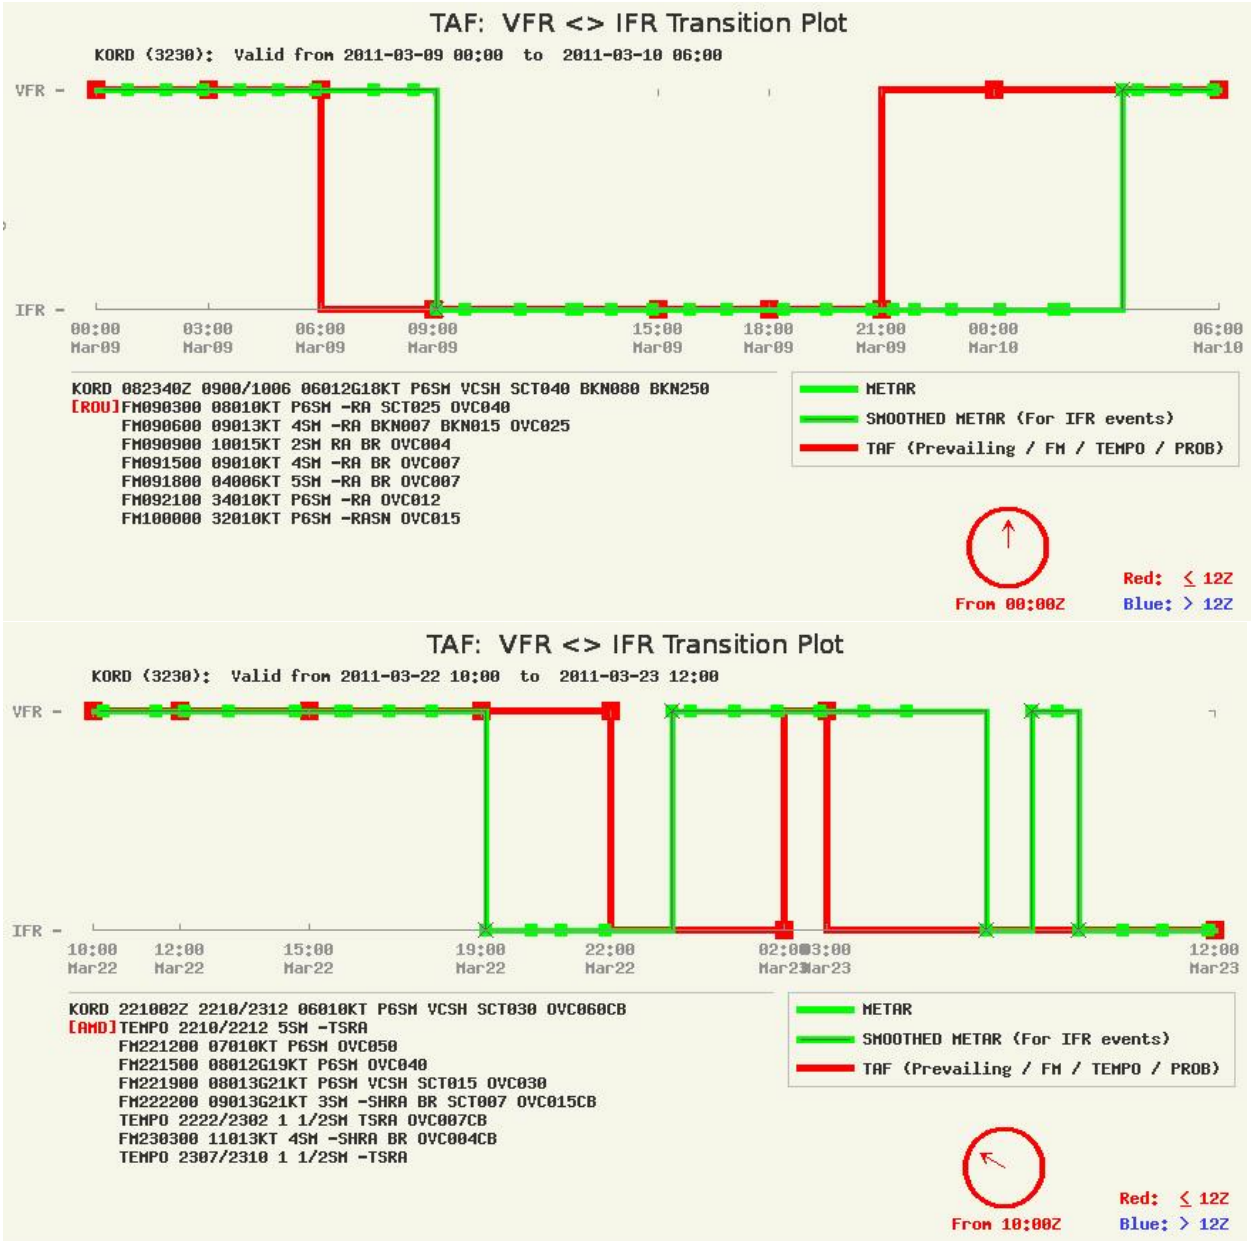

Fig. 3.7. Verification of two TAFs, issued for Chicago O'Hare (KORD). Top: Routine TAF issued at 0000 GMT, 09 March 2011. Bottom: Amended TAF issued at 1000 GMT, 22 March, 2011. See text for complete explanation.

3.1.5 Plots of Ceiling, Cloud, and Visibility Observations

3.1.5.1 METAR plots

METARs are reports of surface weather conditions that adhere to fairly high standards for reliability, accuracy, and comprehensiveness. Surface measurements important for aviation are invariably included in these reports. Figure 3.8 illustrates the plotting convention for METARs. A typical report will include temperature, dewpoint, present weather (if any), visibility, altimeter setting wind direction and speed, fractional cloud cover, and flight conditions (LIFR, IFR, MVFR, and VFR). For a complete listing of present weather symbols, go to http://www.aviationweather.gov/adds/metars/description3.php. For a description of plotting conventions for cloud cover and ceiling height (equivalently, flight conditions), go to http://www.aviationweather.gov/adds/metars/description_ifr.php.

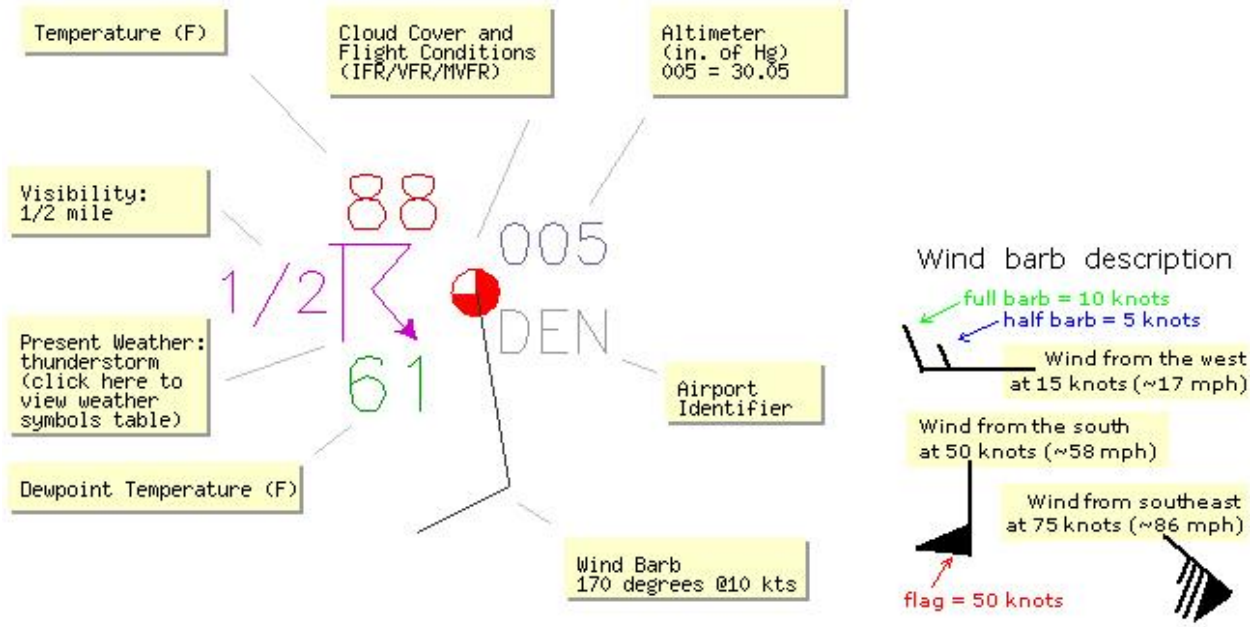

Fig. 3.8. The plotting convention for METAR observations.

For regional plots of METARs, go to http://www.aviationweather.gov/adds/metars/, and click on any of the listed major airports to display METAR data in that area. Figure 3.9 is a plot of METARS near Denver, Colorado, and surrounding states at 1800 GMT, 08 October 2011. It

depicts an early-season mix of rain, snow, and low clouds along the Front Range of the Rocky Mountains and adjacent plains.

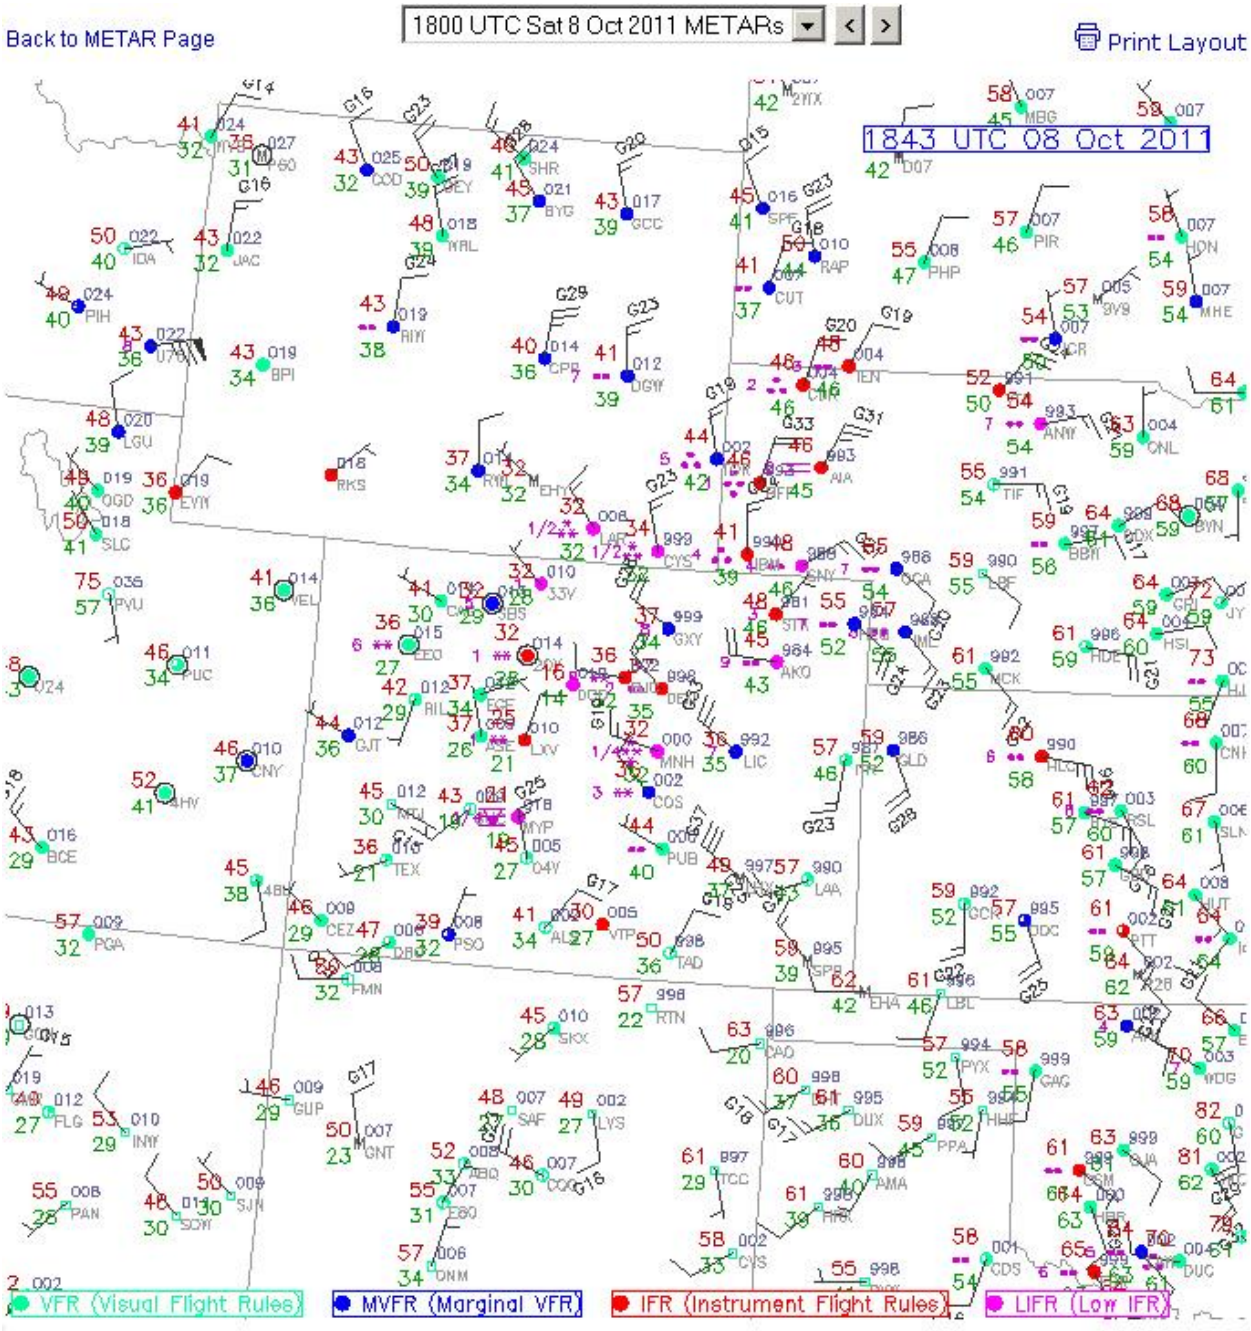

Fig. 3.9. A regional plot of METAR observations for the central Rocky Mountains and adjacent Great Plains centered on Denver International Airport, Colorado (DEN). Websites listed in the text explain the plotting convention and symbols used.

3.1.5.2 Satellite images with symbols for cloud cover and ceiling overlaid

A very useful display of METAR cloud and ceiling observations (from METARs) superimposed on a GOES satellite image is available at http://www.aviationweather.gov/adds/satellite/. Click on the type of satellite image desired (visible, infrared, etc.), choose the latest image or a loop of images, then click on "Western U.S." or "Eastern U.S." Figure 3.10 is a GOES visible cloud image for the western U.S. with an overlay of LIFR, IFR, MVFR, and VFR icons and station circles filled, partially filled, or open depending upon the total cloud cover.

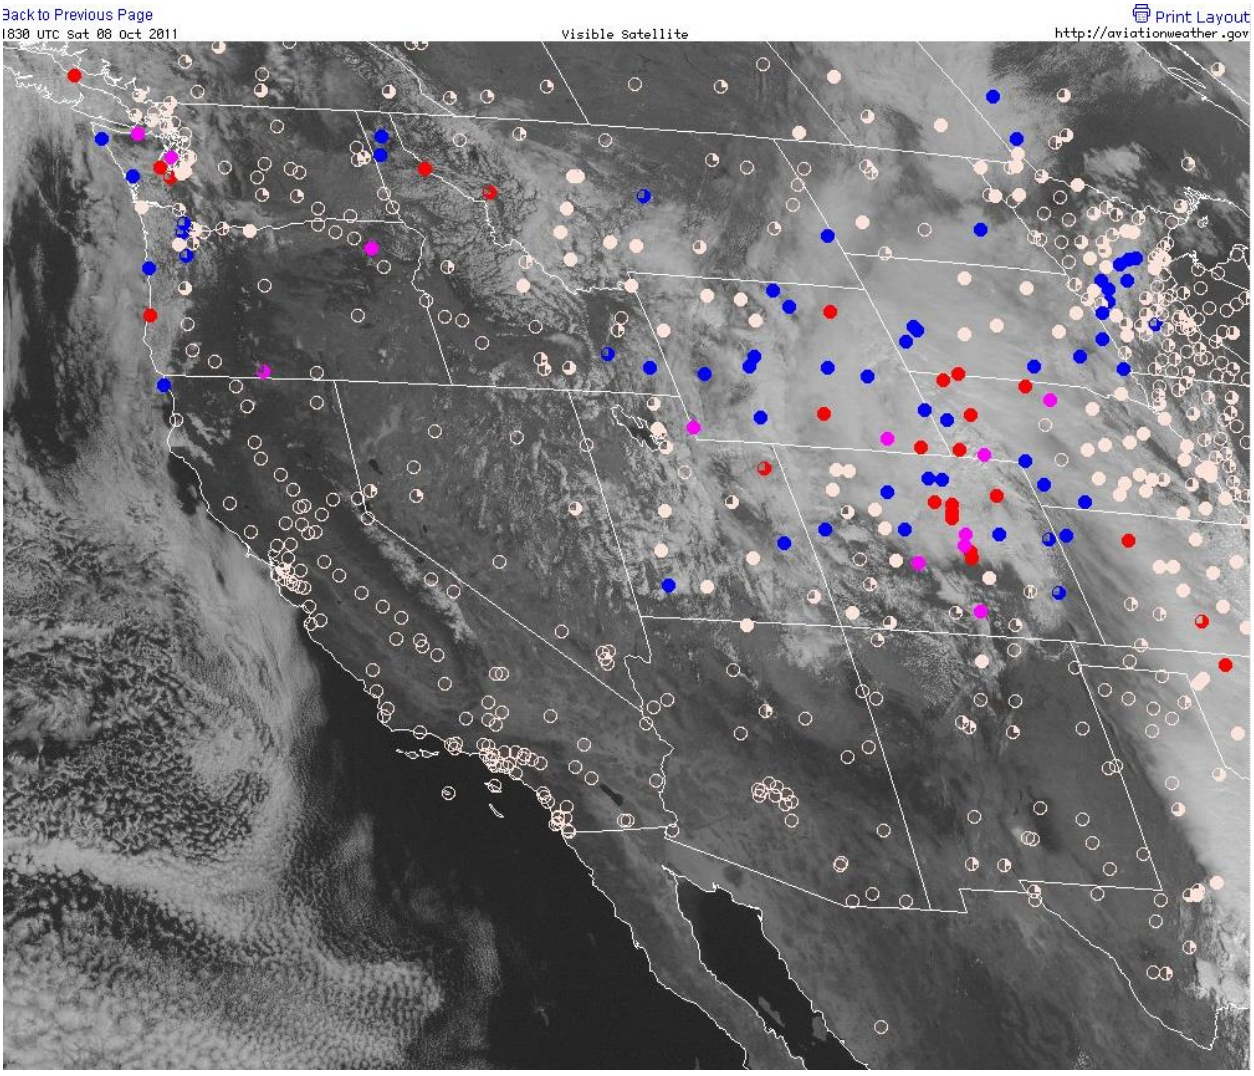

Fig. 3.10. GOES visible cloud image for the western U.S., 1830 GMT, 08 October 2011. Flight categories are color coded: LIFR, IFR, MVFR, and VFR. Station circles are open if the sky is clear, filled if the sky is overcast. Partially filled circles indicate cloud cover in octas (eighths).

3.1.5.3 Pilot reports of cloud bases and tops

General aviation pilots sometimes report cloud bases and tops, as in Fig. 3.11. The key to interpretation is at the bottom of the figure. During the valid period for this plot, IFR conditions existed in western North Carolina, and almost all of South Carolina and Georgia. One would not expect many GA pilots to be flying in such weather, and such was the case. Most reports indicate MVFR conditions or better. Plots like this one are available at http://www.aviationweather.gov/adds/pireps/. Click on "Weather/Sky" and then on any of the six regions on the U.S. map below.

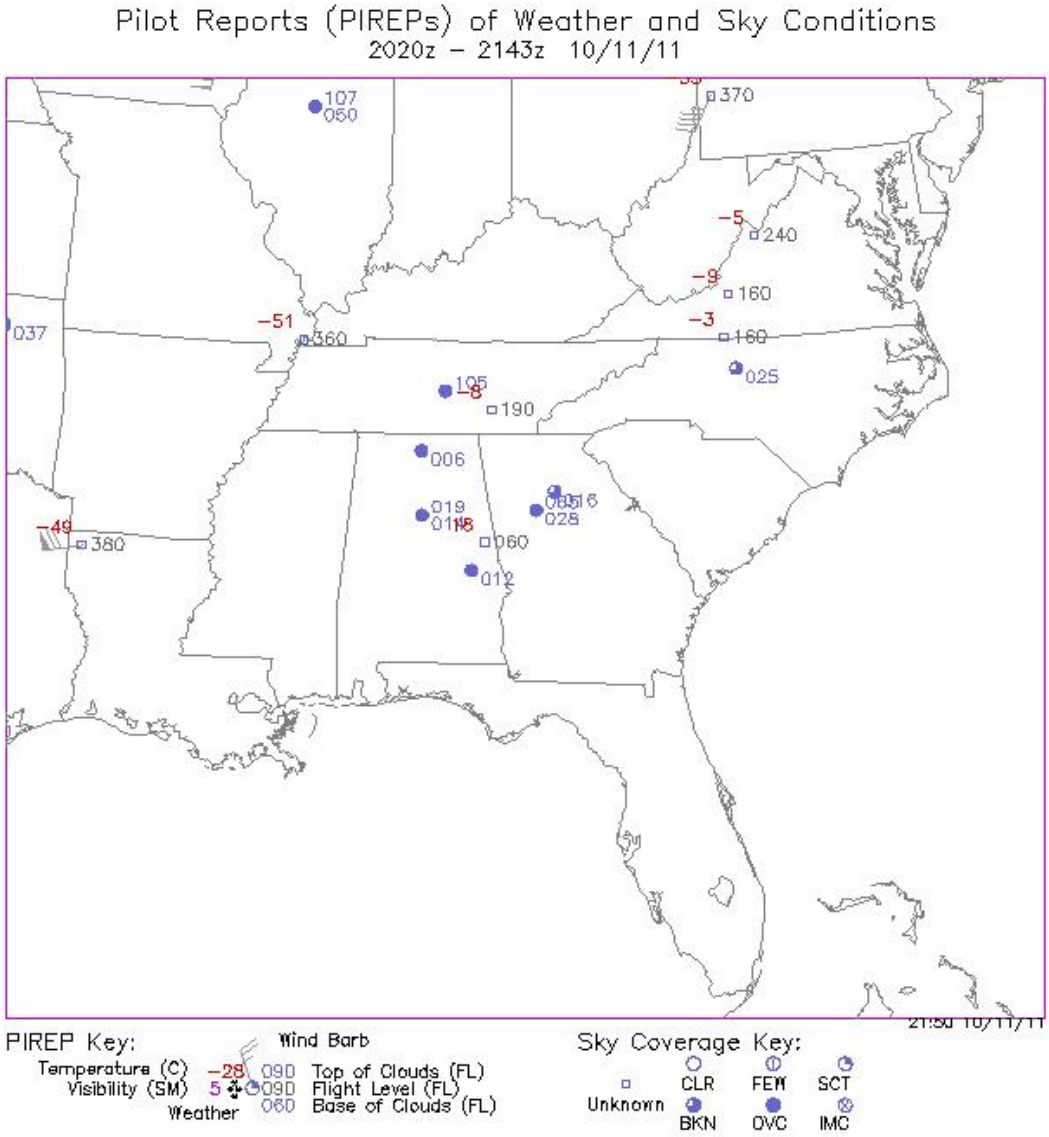

Fig. 3.11. Pilot reports of cloud bases, tops, outside air temperature, and wind for the SE U.S.

3.1.6 Localized Aviation MOS (Model Output Statistics) Product (LAMP)

LAMP was introduced in Section 2.3.9 in connection with products relating to convection. Because LAMP is more fundamental to ceiling and visibility products than it is to convection products, it receives more space here.

The Meteorological Development Laboratory (MDL) of NWS developed LAMP, motivated by a desire to provide useful guidance to meteorologists who generate TAFs. It was clear from the start that such guidance would have to be hourly. Figure 3.12, from Ghirardelli and Glahn (2010), illustrates the LAMP paradigm. LAMP melds information from the most recent surface observations (METARs) and MOS guidance from GFS model output. The persistence of current conditions as defined by METARs is a good forecast, especially in the first hour or two. A GFS

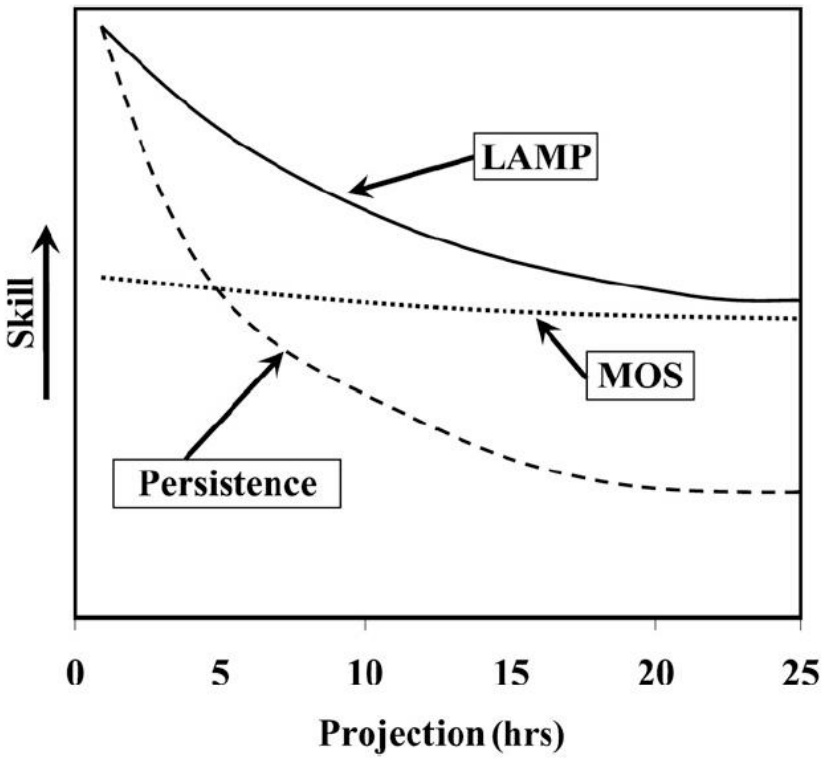

Fig. 3.12. Illustrating the LAMP paradigm. Persistence of current conditions is a good forecast for the first few hours. MOS starts with a handicap because, at issue time, the GFS forecast, on which it is based, is already four hours old, but it soon beats persistence. LAMP is a statistical amalgam of persistence, as defined by observations, simple model output, and GFS MOS. Used with permission of the American Meteorological Society.

MOS forecast beats the GFS model forecast routinely because it uses statistics of past model performance to remove biases. GFS MOS is not available until four hours after the initial time of the GFS forecasts at 0000, 0600, 1200, and 1800 GMT and so it has little chance of beating persistence in the early hours. On the other hand, a statistical combination of both information sources is expected to be better than either one alone (the expectation is borne out in practice), and that is the basis of LAMP. A minor logistical detail is that the GFS MOS output is available only in 3-h increments, and LAMP requires hourly increments. Linear interpolation in time solves this problem.

Figure 3.13, also from Ghirardelli and Glahn (2010), gives a time line for generation of LAMP. At 0400, 1000, 1600, and 2200 GMT each day, the most recent GFS model run is complete, and the MOS guidance (light blue squares) has been generated from that model run. Note that the 0300 GMT METARs are incorporated in the MOS guidance issued at 0400 GMT, the 0900 GMT METARs are incorporated in the MOS guidance issued at 1000 GMT, and so on. In each color-coded 6-h time block, LAMP (red diamonds) updates the GFS MOS every hour on the half hour with the most recent METAR observations (green circles). Not by coincidence, routine TAFs (yellow stars) are due between 80 and 100 min after GFS MOS becomes available, thus giving meteorologists ample time to consult the LAMP guidance while the GFS MOS product is still "fresh."

For ceiling and visibility forecasts, LAMP relies primarily on three sources of predictors (Weiss and Ghirardelli, 2005): GFS MOS, METARs, and a simple model that advects clouds, surface temperature, and saturation deficit. (The saturation deficit measures the degree of saturation in the 1000–500 hPa layer. If the amount of moisture in this layer is held constant while the layer cools, its thickness will decrease, and the relative humidity in the layer will rise until saturation occurs. The difference between the thickness reached at saturation and the actual thickness is the saturation deficit.)

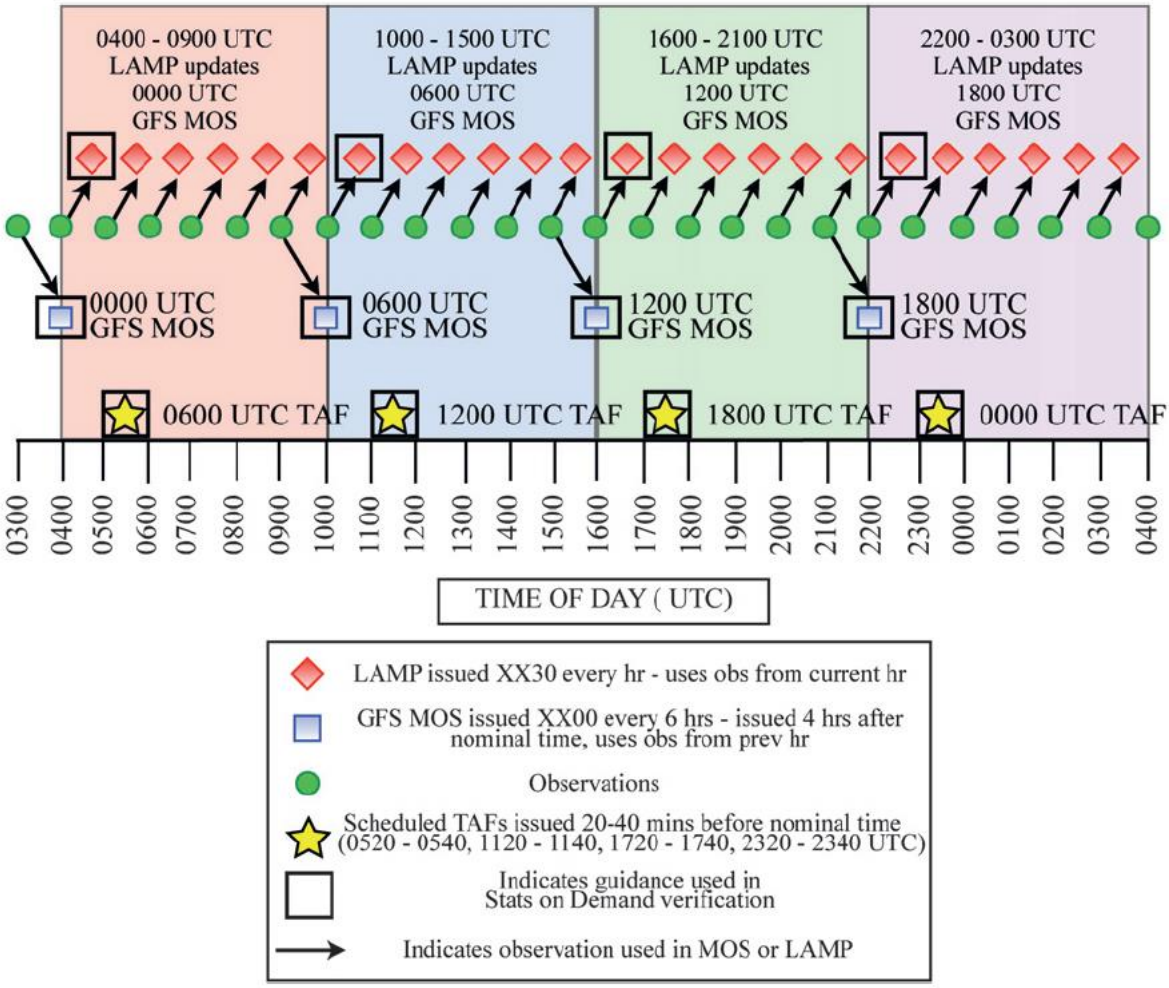

Fig. 3.13. The time line for generating LAMP, issued on the half hour, every hour. Used with permission of the American Meteorological Society.

The LAMP equations rely on multiple–linear regression, and the predictors are chosen with a forward-screening process (Glahn and Lowry, 1972). A separate set of regression equations is used for each hourly cycle. For ceiling and visibility forecasts, some events are relatively rare, e.g. ceiling below 500 ft or visibility below ½ mile, so, instead of developing regression coefficients that apply to an individual station, LAMPS groups stations regionally to increase the sample size and make the statistics more robust. For elements that are clearly related, such as ceiling height and total cloud cover, the regression equations are developed simultaneously to

minimize the chance of obvious inconsistencies in the forecast.[1] Finally, the regression coefficients change with the season: the cold season runs from October through March and the warm season from April through September.

LAMP generates hourly projections for more than 1,500 sites in the U.S., of which more than 1,400 are in the CONUS. METAR observations at these same locations contribute to the forecasts and also provide the verification data after the fact. LAMP has been updating forecasts hourly since November 2008.

3.1.6.1 LAMP text product

To see a text forecast for an individual site, go to http://www.nws.noaa.gov/mdl/gfslamp/statebull.shtml, then click a state either on the U.S. map or in the list of state abbreviations below. A LAMP text product for each site in the chosen state will appear. A sample appears below in Table 3.5.

Table 3.5

The LAMP text product for Concord, New Hampshire, running for 25 hours, beginning at 1500 GMT, 13 October 2011. Information for decoding the numbers in the table is immediately below the table.

KCON	CONCORD						GFS LAMP GUIDANCE			10/13/2011		1400 UTC												
UTC	15	16	17	18	19	20	21	22	23	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15
TMP	56	56	57	57	58	58	58	57	57	56	56	56	56	55	55	55	55	56	56	55	55	57	59	60	62
DPT	53	54	54	55	56	56	57	57	57	56	56	56	56	55	55	55	55	56	56	55	55	57	59	60	61
WDR	03	03	03	02	03	04	04	04	04	04	05	00	00	00	00	00	00	00	00	00	00	00	10	10	10
WSP	06	06	05	05	04	03	02	02	02	02	01	00	00	00	00	00	00	00	00	00	00	00	01	02	02
WGS	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG	NG
PPO	90	67	71	57	47	29	23	22	21	22	24	23	23	26	28	31	38	43	48	55	61	65	68	71	74
PCO	Y	Y	Y	Y	Y	N	N	N	N	N	N	N	N	N	N	N	N	Y	Y	Y	Y	Y	Y	Y	Y
P06										37						38						83			
TP2			0	0	0	0	0	0		0		0		0		1		2		3		3		3	
TC2			N	N	N	N	N	N		N		N		N		N		N		N		N		N	
POZ	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
POS	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
TYP	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R	R
CLD	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV	OV
CIG	4	3	3	3	3	3	3	3	3	3	2	2	2	1	1	1	1	1	1	1	1	1	1	2	3
CCG	3	3	3	3	3	3	3	3	3	3	3	2	2	2	2	2	2	2	2	1	1	2	2	2	3
VIS	6	5	5	5	5	5	5	5	5	5	5	4	2	2	2	2	2	2	1	1	1	1	1	1	4
CVS	5	5	5	5	5	5	5	5	6	7	6	5	5	5	5	5	5	5	5	5	5	5	5	5	5
OBV	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	BR	FG	BR	FG	FG	FG	FG	FG	FG	FG	FG	BR

[1] According to Weiss and Ghirardelli (2009), the next version of LAMP ceiling and sky cover forecasts, not yet implemented, will *not* group stations regionally, and the statistics for ceiling and sky cover will be developed separately.

The decoding information below is available at
http://www.nws.noa.gov/mdl/gfslamp/docs/LAMP_description.html.

- **UTC** = Hour of the day in UTC time. This is the hour at which the forecast is valid, or if the forecast is valid for a period, the end of the forecast period.
- **TMP** = surface temperature valid at that hour.
- **DPT** = surface dewpoint valid at that hour.
- **WDR** = forecasts of the 10-meter wind direction at the hour, given in tens of degrees.
- **WSP** = forecasts of the 10-meter wind speed at the hour, given in knots.
- **WGS** = forecasts of the 10-meter wind gust at the hour, given in knots. "NG" means that no gust is forecast.
- **PPO** = probability of precipitation occurring on the hour. The precipitation need not be measurable.
- **P06** = probability of measurable precipitation (PoP) during a 6-h period ending at that time.
- **PCO** = categorical forecast of yes (Y) or no (N) indicating if precipitation, not necessarily measurable, will occur on the hour.
- **TP2** = probability of thunderstorms during the 2-hr period ending at the indicated time.
- **TC2** = categorical forecast of yes (Y) or no (N) indicating if thunderstorms will occur during the 2-hr period ending at the indicated time.
- **POZ*** = conditional probability of freezing pcp occurring at the hour. This probability is conditional on precipitation occurring.
- **POS*** = conditional probability of snow occurring at the hour. This probability is conditional on precipitation occurring.
- **TYP*** = conditional precipitation type at the hour. This category forecast is conditional on precipitation occurring.
- **CLD** = forecast categories of total sky cover valid at that hour.
- **CIG** = ceiling height categorical forecasts at the hour.
- **CCG** = conditional ceiling height categorical forecasts at the hour. This category forecast is conditional on precipitation occurring.
- **VIS** = visibility categorical forecasts at the hour.
- **CVS** = conditional visibility categorical forecasts at the hour. This category forecast is conditional on precipitation occurring.
- **OBV** = obstruction to vision categorical forecasts at the hour.

***In the contiguous United States, these products are available September 1 – May 31.**

The parameters of particular interest in this section are highlighted in yellow in Table 3.5. The hourly predictions of ceiling (CIG) and visibility (VIS) are represented by code numbers, with lower numbers corresponding to lower ceilings and visibilities (see Table 3.6). The probability of precipitation on the hour (PPO) is high at the beginning (90%), decreases to less than 30% by 2000 GMT, then increases steadily after 0300 GMT. If precipitation occurs, the precipitation type (TYP) is expected to be rain (R) for every hour. The cloud cover (CLD) is overcast (OV) every hour. The obstruction to visibility (OBV) is due to mist (BR) or fog (FG), depending upon the hour. Two conditional probabilities are listed, the condition being that precipitation is falling. CCG gives the ceiling category and CVS gives the visibility category, provided that precipitation is falling. These values are sometimes higher than the corresponding values for

CIG and VIS, particularly in the latter half of the forecast for Concord. It might be a surprise that the statistics deliver this result, but the physical cause is that rain scavenges some of the very small droplets associated with fog and actually raises the ceiling and improves the surface visibility. This is not always the case, however. For example, the onset of snowfall often lowers both ceiling and visibility. Rudack (2009) discusses how the onset of precipitation affects ceiling and visibility.

Table 3.6

Definitions of category numbers appearing in the LAMP of Table 3.5. Because ceiling values are given only to the nearest 100 ft, the ranges of ceiling height are not continuous across categories.

LAV Ceiling Height (CIG) and Conditional Ceiling Height (CCG) Categories	
1	< 200 feet
2	200 - 400 feet
3	500 - 900 feet
4	1000 - 1900 feet
5	2000 - 3000 feet
6	3100 - 6500 feet
7	6600 - 12,000 feet
8	> 12,000 feet or unlimited ceiling

LAV Visibility (VIS) and Conditional Visibility (CVIS) Categories	
1	< 1/2 miles
2	1/2 - < 1 miles
3	1 - < 2 miles
4	2 - < 3 miles
5	3 - 5 miles
6	6 miles
7	> 6 miles

For numerical forecast elements, missing values are indicated by a ***99*** or ***999***

LAV Cloud (CLD) Categories	
CL	clear
FW	few > 0 to 2 octas
SC	scattered > 2 to 4 octas
BK	broken > 4 to < 8 octas
OV	overcast

LAV Obstruction to Vision (OBV) Categories	
N	none of the following
HZ	haze, smoke, dust
BR	mist (fog with visibility >= 5/8 mile)
FG	fog or ground fog (visibility < 5/8 mile)
BL	blowing dust, sand, snow

LAV Precipitation Type (TYP) Categories	
S	pure snow or snow grains
Z	freezing rain/drizzle, ice pellets, or
	anything mixed with freezing precip
R	pure rain/drizzle or rain mixed with snow
X	forecast is missing for this hour

3.1.6.2 LAMP categorical forecasts of ceiling and visibility

Though Table 3.5 (a sample LAMP) lists the probability of precipitation occurring on the hour, it did *not* list the probability of a ceiling or visibility occurring in a given category. Nonetheless, such probabilities lie behind the categorical forecasts. Figure 3.14 illustrates the derivation of categorical forecasts of ceiling height from the probability data for the same location (Concord, NH) and the same period as considered in Table 3.5. To obtain histograms like those in Fig. 3.14, go to http://www.nws.noaa.gov/mdl/gfslamp/uncertplots.php, choose "Threshold Plot," type in the four-character station name, for example, KSTL for St. Louis, Missouri, and then choose a threshold in the LIFR, IFR, or MVFR columns. (One can also look at "Uncertainty Plots" on the same web page. Though these provide equivalent information, they are not so easy to interpret.)

Fig. 3.14a shows the probability of a ceiling height less than 500 ft for each hour of the forecast. A horizontal black line for each hour indicates a threshold probability for this ceiling category. The threshold is determined iteratively from a dependent data sample in such a way that the CSI score is maximized within a targeted bias range. In other words, adjust the threshold (which helps one decide whether to forecast a ceiling of < 500 ft or not) repeatedly and look at the bias and CSI scores. For those cases where the bias is acceptably close to 1.0, find the threshold that delivers the highest CSI score, and that becomes the threshold shown in Fig. 3.14a.

What is the meaning of the colors in the histograms of Fig. 3.14? Let ΔP represent the threshold probability minus the predicted probability. Table 3.7 interprets the colors.

Table 3.7
How to read the colors in Fig. 3.14

ΔP	**Color**	**Interpretation**
> 10%	Light blue	Event will not occur.
0 to 10%	Yellow	Chance that the event will occur.
-10 to 0%	Orange	Event is likely.
< -10%	Red	Event will occur.

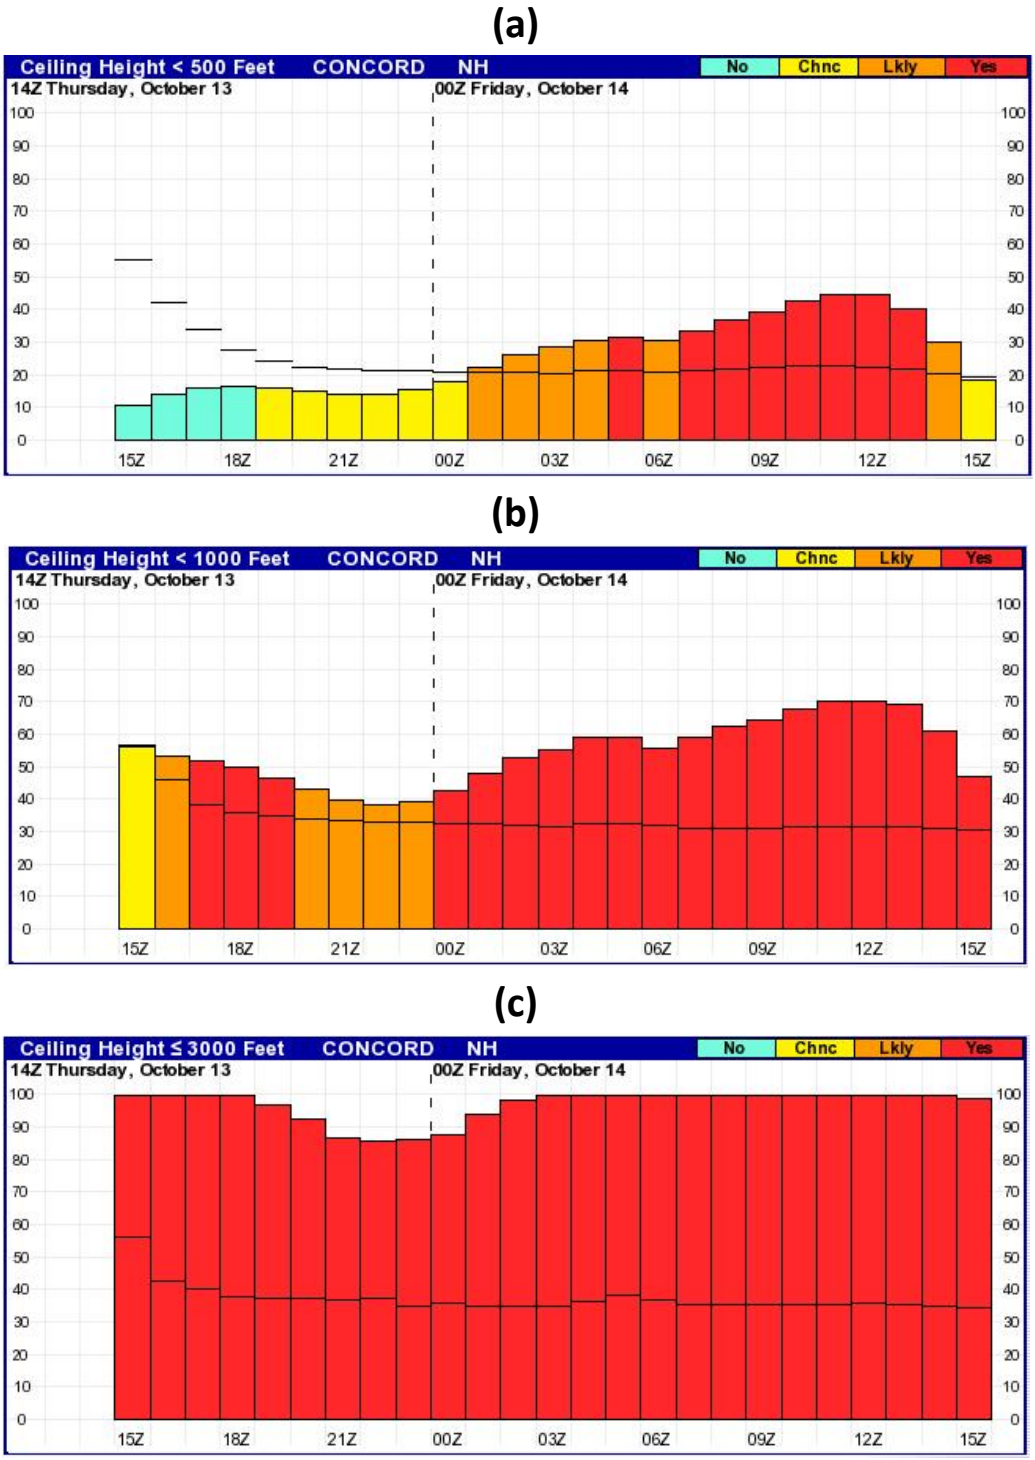

Fig. 3.14. The derivation of categorical ceiling forecasts from probability forecasts of ceilings at three thresholds: (a) < 500 ft, (b) < 1000 ft, and (c) ≤ 3000 ft. Horizontal black lines in these histograms represent category thresholds. See text for further explanation.

3.1.6.3 Meteograms from LAMP forecasts

MDL generates meteograms for many of the parameters listed in the LAMP forecast, that is, graphs of parameter values for each hour of the forecast. To obtain meteograms, go to http://www.nws.noaa.gov/mdl/gfslamp/meteoform.php, then click on the state and airport of interest. One can view up to 12 different parameters in meteogram form. Figure 3.15 shows meteograms for parameters related to ceiling and visibility.

The meteograms in Fig. 3.15 portray graphically the LAMP forecast in Table 3.5. The ceiling height forecasts are categorical, but they depend on probability histograms such as those shown in Fig. 3.15. To make categorical forecasts for all eight categories defined in Table 3.6 one must examine the corresponding eight histograms, starting with the one for the lowest ceiling height and working up. For a given hour, if the predicted probability is less than the threshold probability, look at the next higher category. If the predicted probability is still less than the threshold for that category, continue the process. The category for which the predicted probability first exceeds the threshold becomes the "best category" forecast. If this never happens, then the default category (No. 8: ceiling above 12,000 ft) is chosen.

A similar procedure is followed to derive the best category forecast for visibility.

Fig. 3.15. Meteograms related to ceiling and visibility for Concord, New Hampshire. These are consistent with the LAMP forecast in Table 3.5 and the histograms of probability in Fig. 125.

3.1.6.4 Gridded LAMP forecasts of ceiling and visibility

The LAMP hourly forecasts apply at over 1400 stations in the CONUS, for which METAR observations are available. The small pink dots on Fig. 3.16 locate these stations. Recognizing that the "4-D data cube" envisioned for NextGen will create a need for high-resolution gridded forecasts in the years to come, MDL began developing a gridded version of LAMP several years ago. This version became operational on 27 September 2011.

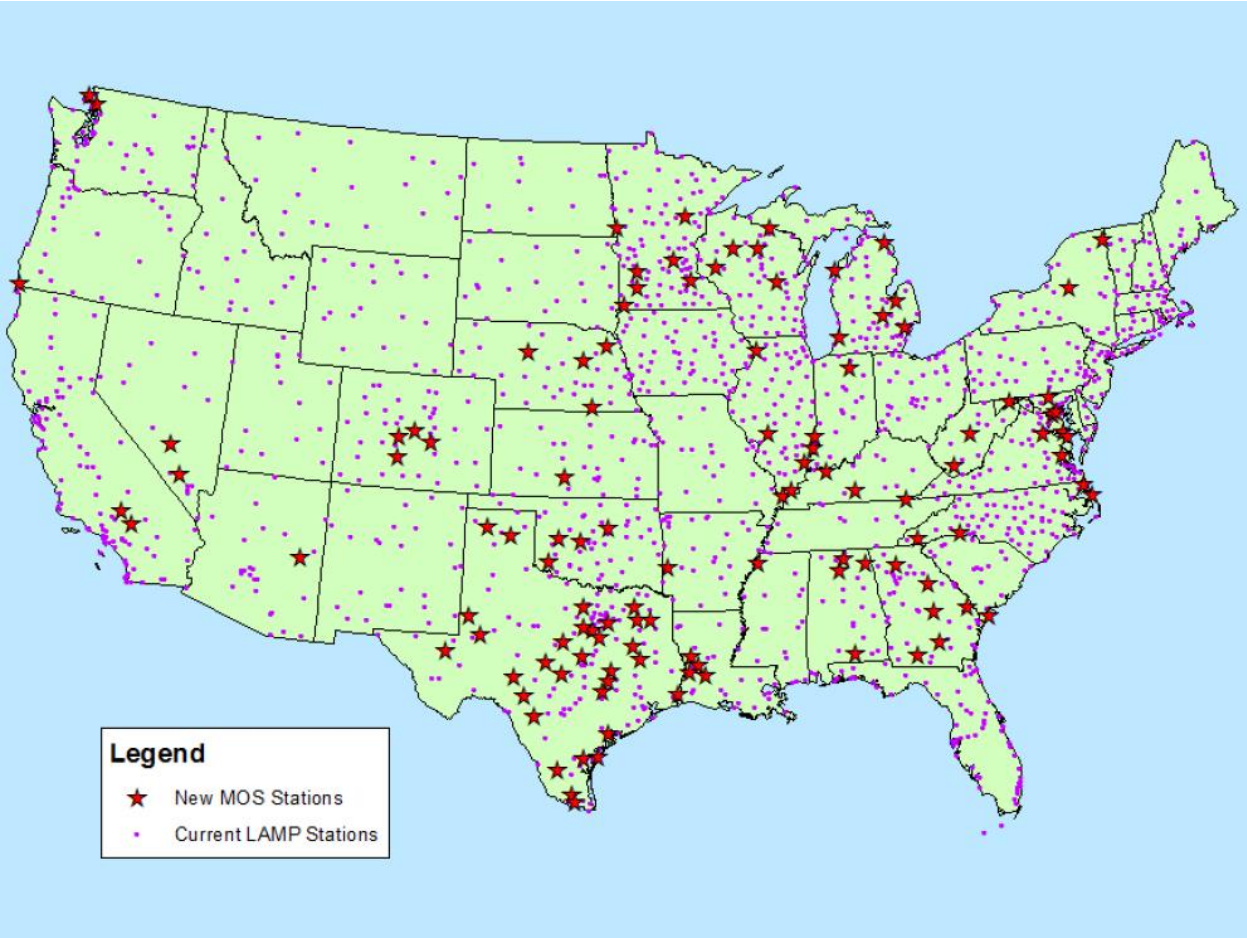

Fig. 3.16. Pink dots identify the 1400 plus METAR sites for which LAMP forecasts are issued. The stars identify new MOS (METAR) stations, for which LAMP does *not* make a forecast. The purpose of the starred stations is to verify LAMP gridded forecasts at points where spatial interpolation of the station forecasts is required. Used with permission of the American Meteorological Society.

Ghirardelli and Glahn (2011) describe the generation of gridded LAMP guidance. One challenge is to interpolate the ceiling and visibility forecasts, which are in the form of discrete categories, onto a 2.5-degree grid. The first step in this process is to convert the integer values for ceiling category (1, 2, 3, 4, 5, 6, 7, 8) to a continuous range of values from 1.0 to 9.0, and the integer values for visibility category (1, 2, 3, 4, 5, 6, 7) into a continuous range from 1.0 to 8.0.

An example illustrates this procedure. Suppose the forecast for a particular hour and METAR location is for a Category 2 ceiling. Find the lowest and highest probabilities associated with a Category 2 ceiling anywhere in the CONUS for that hour. If the probability of Category 2 at the METAR location is close to the highest value found, the METAR location would receive a continuous ceiling value slightly above 2.0. If the probability is close to the lowest value found, the METAR location would receive a continuous ceiling value slightly below 3.0. A higher probability means a lower continuous ceiling value within that category. This logic holds for all categories except Category 8, which applies to ceilings above 12,000 ft. The higher the probability of a Category 8 ceiling, the closer the continuous value is to 9.0.

Once continuous category values have been derived for the LAMP forecast, they are interpolated onto a grid via a method described by Glahn and Im (2011). The method is essentially the successive correction method used extensively in the early days of numerical prediction, but with a number of modifications. The method begins with a first guess, a specification of the initial values on the grid, all subject to a series of corrections. Over land, the first guess can be as simple as a constant field, for example, a nationwide average of all LAMP station values. Over water, where surface observations are sparse, the first guess can be a short-term model forecast of cloud ceiling.

The first round of corrections begins with specification of a *radius of influence*, which determines the size of a circle, centered at the grid point where an estimate of ceiling height is to be derived. The radius of this circle will vary with the density of LAMP stations surrounding the grid point, but the circle will be large enough to include perhaps tens of stations. The first guess value is interpolated from the grid to each LAMP station within the circle and then the difference between the LAMP value and the first guess is computed. The correction at the grid point is a linear combination of these differences, with individual weights depending upon the distance between the grid point and the LAMP station. The weights are normalized so that their sum is one. This weighting scheme ensures that LAMP stations closest to the grid point will have the greatest influence on the correction to be applied at the grid point.

On the first pass, the collections of LAMP sites that determine the corrections at adjacent grid points are likely to be very similar. This results in a fairly smooth field of gridded LAMP values. But the true field of ceiling height or visibility values is not likely to be smooth, so additional corrections are required. The corrected values at each grid point become the first guess for a second round of corrections, this time with a smaller radius of influence. Fewer LAMP stations lie within the smaller circle, and so the second round of corrections will result in a field that is less smooth than the previous one, and more heavily influenced by LAMP values near the grid point. Four of five corrections, each one employing a smaller radius of influence, may be computed.

Glahn and Im (2011) employ special procedures near land-sea boundaries and in complex terrain to minimize unreasonable artifacts in the final analysis.

Observations of ceiling (in hundreds of feet) and visibility (in miles and fractions thereof) at LAMP sites are also converted to a continuous scale of categories. For example, a ceiling height of 1800 ft falls in Category 4 (Table 3.6). Category 4 includes ceilings from 1,000 ft up to, but not including, 2,000 ft. The continuous category number would thus be 4.8. The continuous category numbers at LAMP stations are interpolated to the 2.5-km grid by the successive correction method just described.

The categorical values analyzed on the 2.5-km grid are converted back to ceiling height (hundreds of feet) or visibility (miles) in preparation for display. For gridded LAMP products, go to http://www.mdl.nws.noaa.gov/~glmp/glmp_expr.php. Choose the desired product and the forecast time. Figure 3.17 shows a color-coded analysis of observed ceiling height. The transitions between colors becomes pronounced only for low ceilings, which have the greatest operational impact. All ceilings between 3,000 and 12,000 ft are colored in various shades of green. Wherever the ceiling is unlimited (above 12,000 ft), the color is beige.

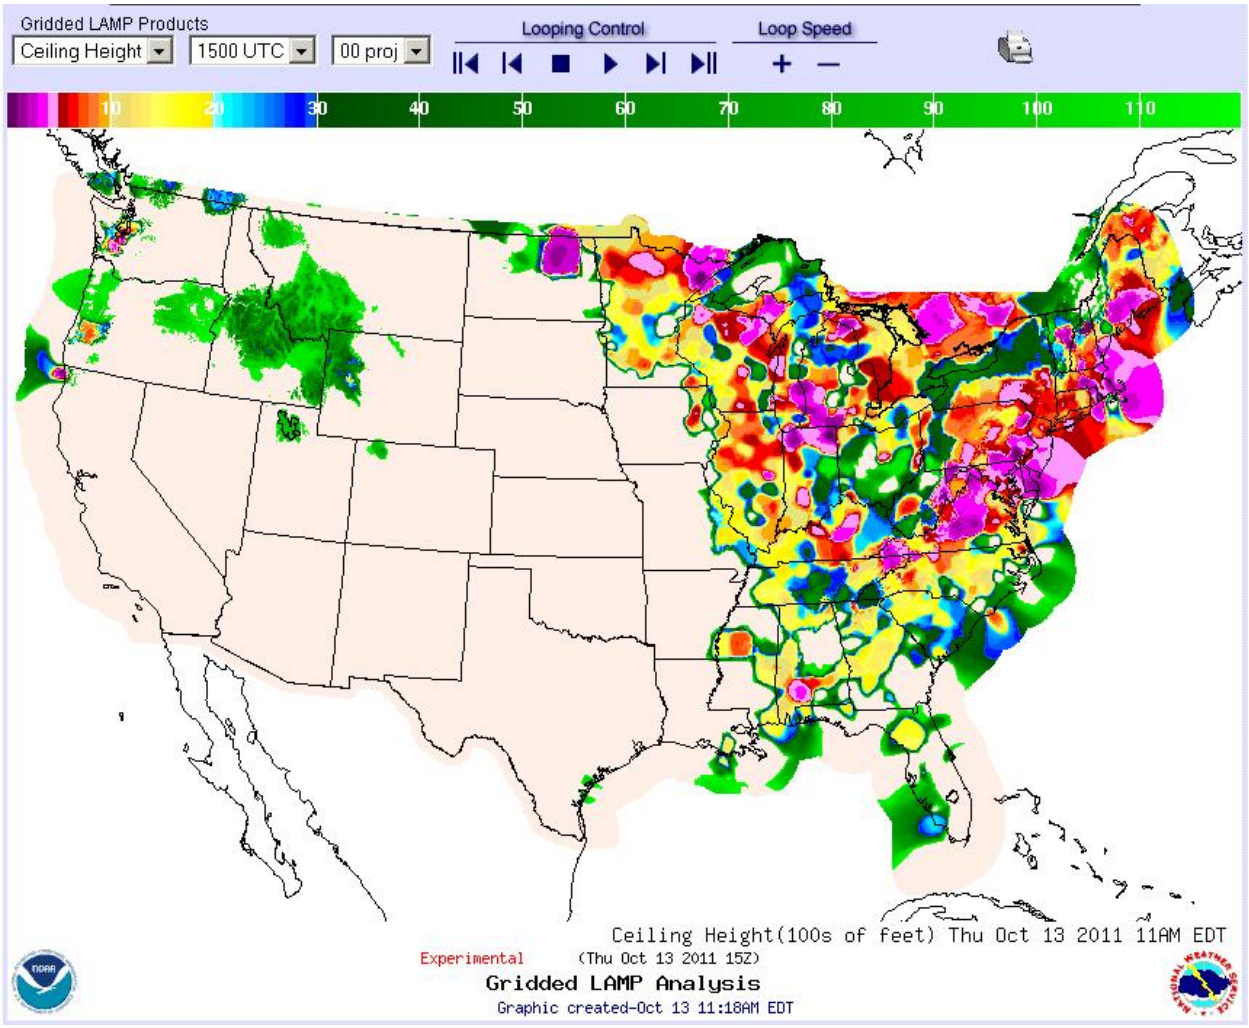

Fig. 3.17. Current ceiling height, color-coded in hundreds of feet AGL, for 1500 GMT, 13 October 2011, analyzed to a 2.5-km grid. The color bar used to depict ceiling height lies just above the image.

Figure 3.18 depicts a gridded, 12-h forecast of ceiling heights derived from LAMP forecasts at 1400+ LAMP stations. The LAMP forecasts are analyzed to a 2.5-km grid. The color bar at the top of the image corresponds to the ceiling height in hundreds of feet, as in Fig. 3.18. The forecast is valid at 0300 GMT, 14 October 2011. Note the deterioration of conditions along the Mid-Atlantic and New England coasts expected to occur between the initial time (Fig. 3.18) and the valid time.

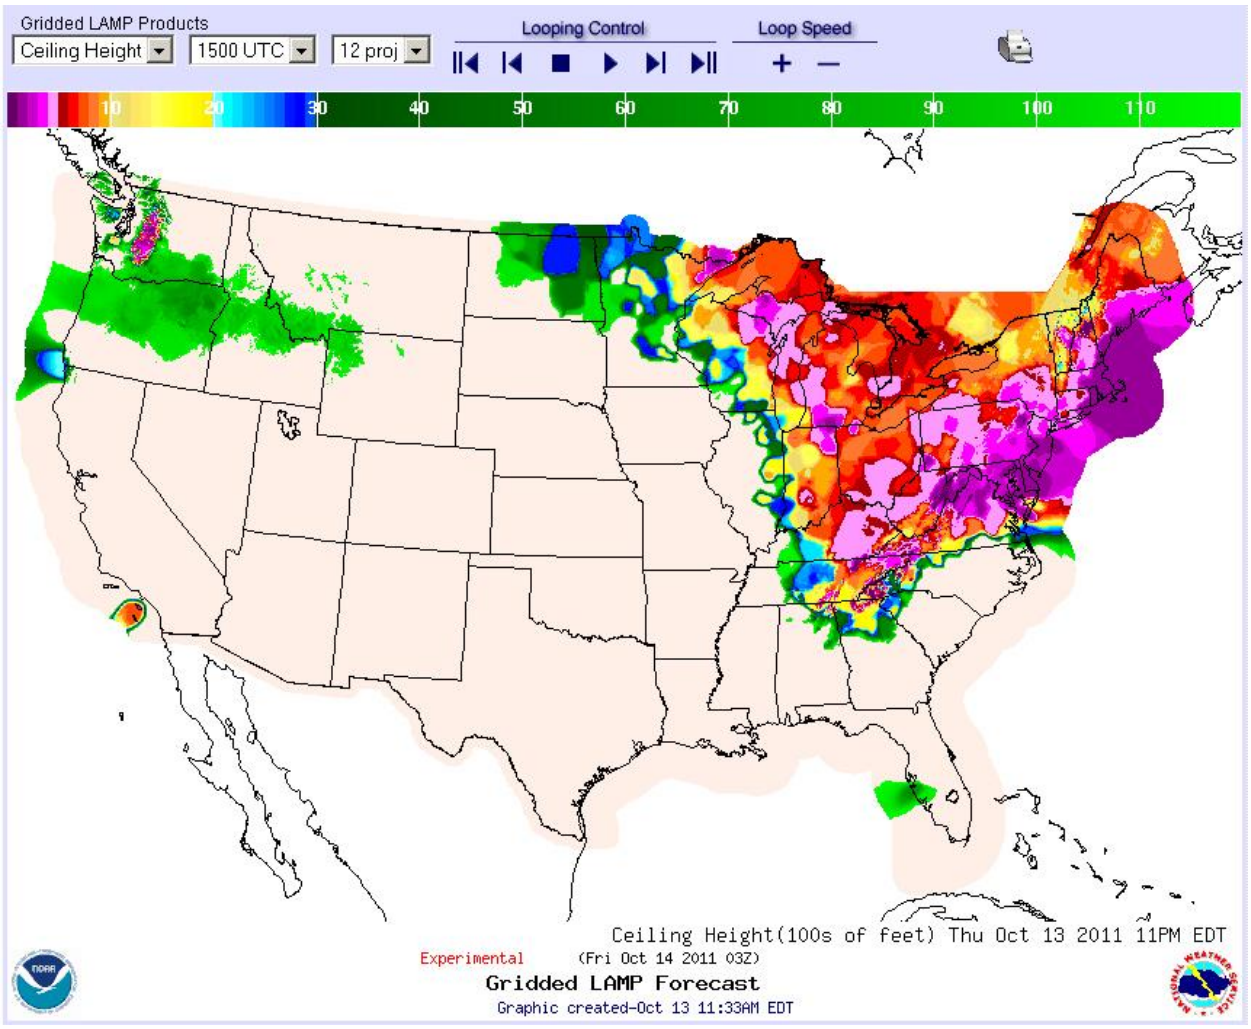

Fig. 3.18. A gridded 12-h LAMP forecast of ceiling height, valid at 0300 GMT, 14 October 2011. The color bar across the top of the image indicates ceiling height in hundreds of feet.

As noted earlier, the purpose of gridded LAMP is to provide statistical guidance about aviation weather conditions at arbitrary locations, not just where the LAMP METAR stations are located. According to Ghirardelli and Glahn (2011), interpolation of ceiling and visibility forecasts from the 2.5-km grid to the 1400+ LAMP stations locations does not degrade the accuracy of the original forecasts valid at LAMP stations, as expected. See Fig. 3.19. These forecasts are as accurate as persistence for the first three hours, after which they beat persistence. But when the gridded forecast is compared with observed values of ceiling and visibility at 115 *non*-LAMP locations which were *not* included in the analysis (identified by the red stars in Fig. 3.16), the

accuracy is somewhat degraded, especially for IFR conditions, because such conditions tend to be patchy and discontinuous. Forecast accuracy at the non-LAMP stations is inferior to that of persistence until about three hours after the initial time.

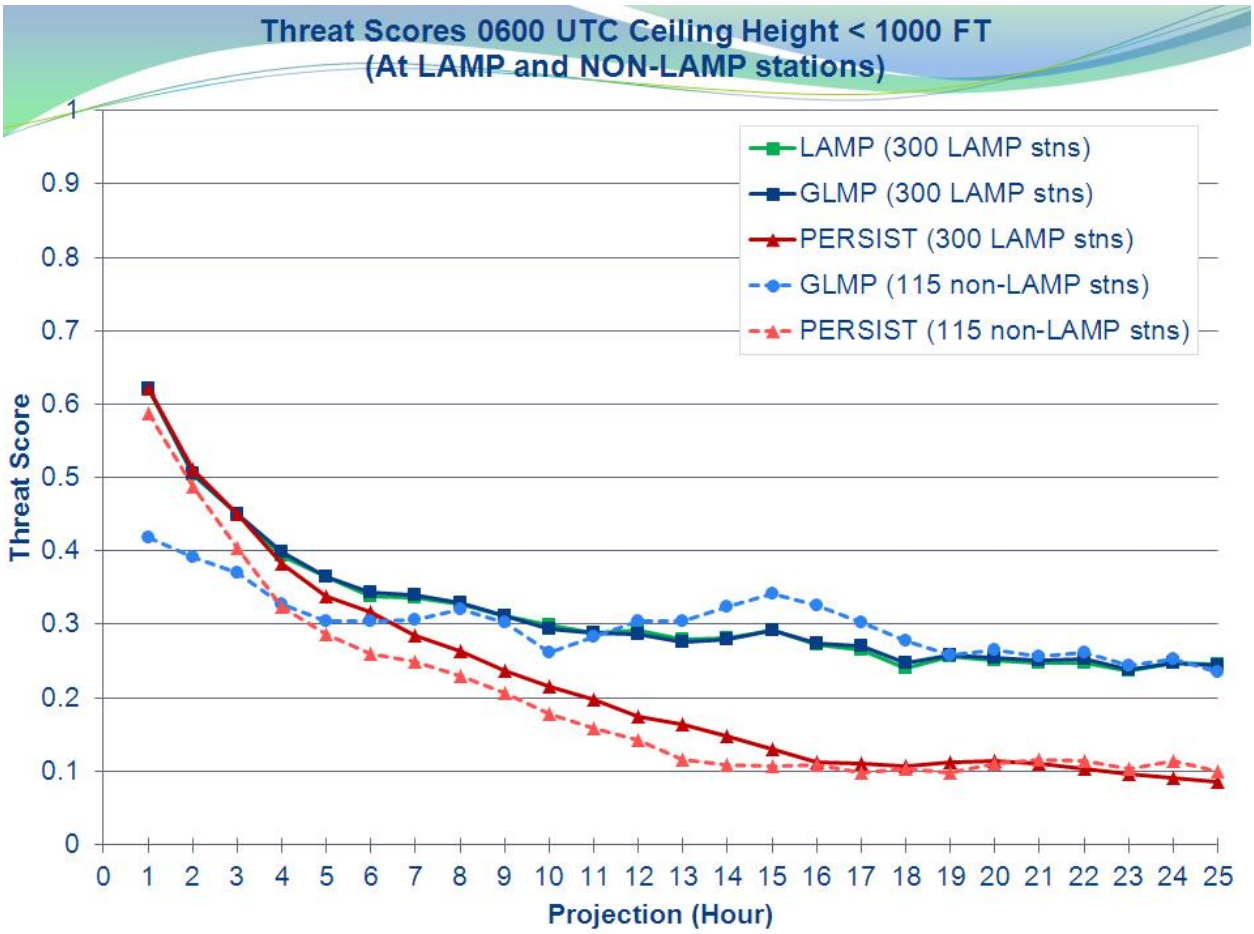

Fig. 3.19. Threat scores (same as CSI) for a categorical forecast of ceiling height less than 1000 ft at 300 LAMP and 115 non-LAMP locations. This chart demonstrates degradation in forecast accuracy at locations in between the LAMP stations, particularly during the first few hours, the result of interpolating a noisy field. Figure courtesy of Judy Ghirardelli, NWS/MDL.

3.1.6.5 Verification of LAMP ceiling and visibility products

Ghirardelli and Glahn (2010) present some of the most recent verification statistics in the refereed literature. Figure 3.20 compares threat scores (same as CSI) for yes-no forecasts of

ceilings below 1000 ft for April through September 2007 (top) and October 2007 through March 2008 (bottom). Three forecasts are verified: LAMP, GFS-MOS, and persistence. The verification occurred at 1,522 METAR stations in the CONUS, Alaska, Hawaii, and Puerto Rico.

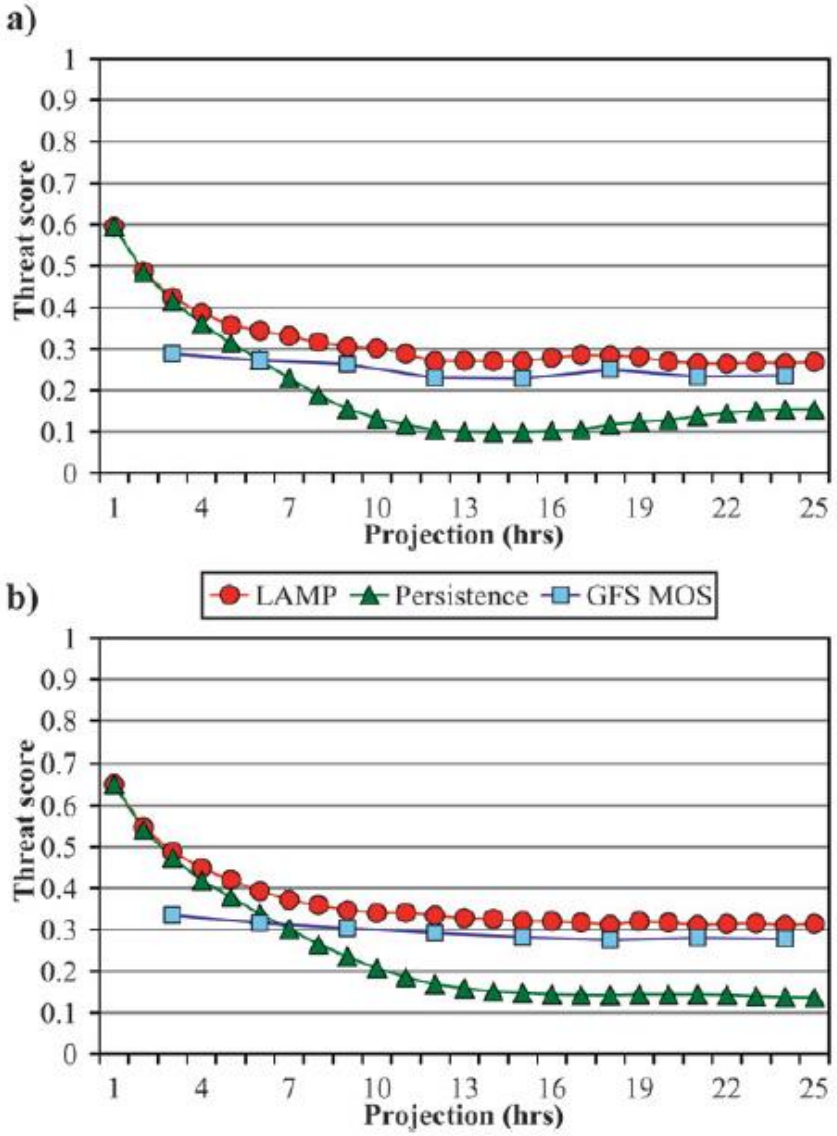

Fig. 3.20. Verification of LAMP (red circles), GFS-MOS (blue squares), and persistence (green triangles) forecasts of ceiling below 1000 ft. The vertical axis is the threat (CSI) score. The horizontal axis is the length of the projection from 1 to 25 hours. (a) Warm–season verification (April through September 2007). (b) Cold–season verification (October 2007 through March 2008).

Because LAMP is a statistical blend of persistence and GFS-MOS, one expects its performance to be equal to or better than either one alone, and that is the case at all hours. As noted earlier, GFS-MOS is not available until about four hours after the GFS initial time, and the official GFS-MOS output occurs at 3-h intervals, hence the spacing of the blue squares. The cold-season forecasts are slightly better than the warm-season forecasts, probably because the cloud systems of winter are larger, more cohesive, and more strongly forced by dynamics than those of summer.

The information in Fig. 3.21 is similar to that in Fig. 3.20, the only difference being that the forecast is for visibility less than three miles (IFR conditions). Forecasts of low visibility are less challenging in winter (b) than in summer (a) because, when it snows, visibility frequently goes below three miles, but snow only falls during the cold half of the year. LAMP visibility has no advantage over GFS-MOS forecasts of visibility in the last 12 h. IFR visibility forecasts receive lower scores than IFR ceiling height forecasts.

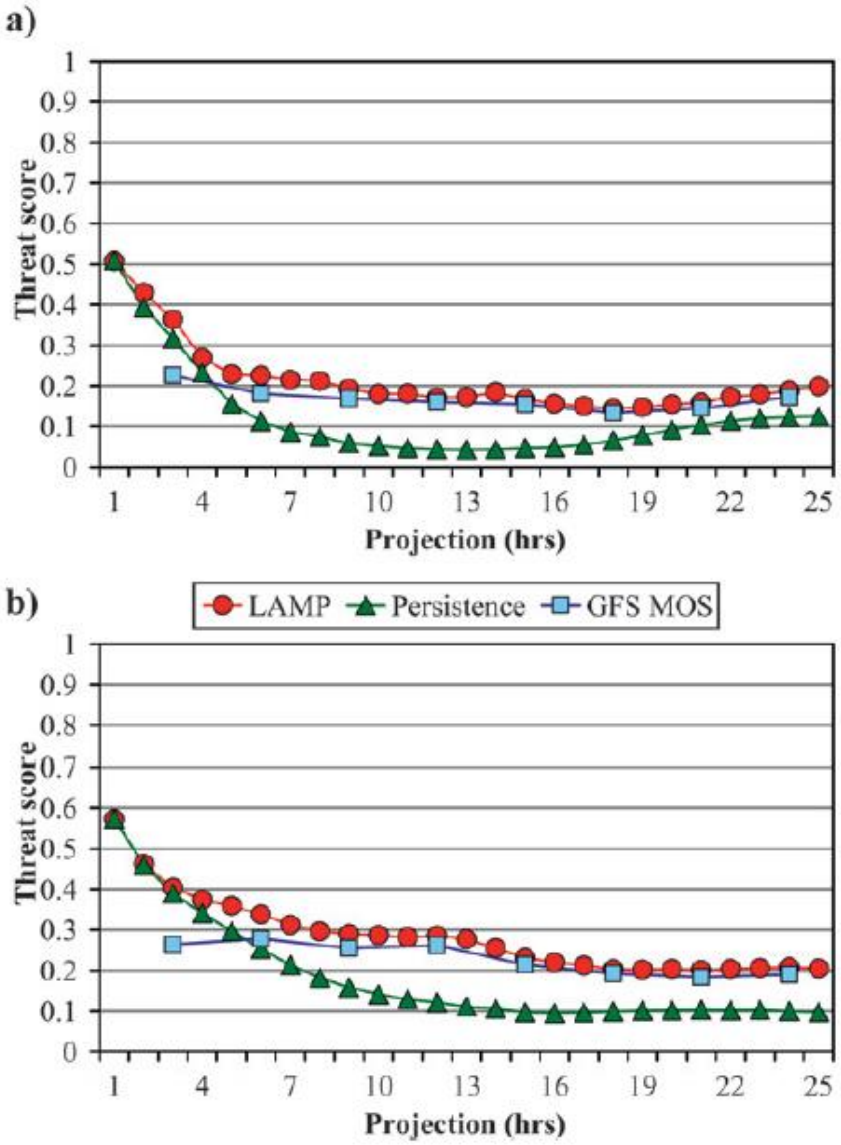

Fig. 3.21. Same as Fig. 3.20, except for forecasts of visibility less than three miles.

Rudack and Ghirardelli (2010) compared the skill of ceiling and visibility forecasts by LAMP and several numerical weather prediction models. Included in the comparison were the 20-km Rapid Update Cycle[2] (RUC20), the Weather Research and Forecasting Nonhydrostatic Mesoscale Model (WRF-NMM), and the Short-Range Ensemble Forecast (SREF) system, all operational at NCEP/EMC. The SREF system generates an ensemble (a collection) of 21

[2] By 2006, the operational version of the RUC ran at 13-km resolution, but model output was not archived. The only archive available for this study was from a 20-km version of the RUC.

forecasts, including 10 perturbations from NCEP's Eta model, five from NCEP's Regional Spectral Model, three from the WRF-NMM, and three from NCAR's version of the WRF model, which is distinct from the WRF-NMM. The database consisted of forecasts generated from October 2006 through March 2007. Only two examples of verification from among many can be shown here.

Figure 3.22 shows the CSI score (left) and bias (right) for categorical forecasts of ceiling height: (a) $\leq$ 3,000 ft, (b) < 1,000 ft, and (c) < 500 ft. Four different forecasts are verified: LAMP, WRF-NMM, RUC20, and persistence. Because persistence is a very good forecast and its influence on LAMP is strong during the first few hours, the CSI scores for LAMP and persistence forecasts track together through the third hour. LAMP beats the two model forecasts at all hours, handily at the beginning, and in all three ceiling categories. The RUC forecast ends at 12 h, and its skill as measured by CSI is relatively steady during this period. The skill of the WRF-NMM model matches that of the RUC20 for ceilings $\leq$ 3,000 ft but is somewhat less for lower ceiling categories. The bias of LAMP forecasts is close to one for all ceiling categories, but the two numerical models overforecast ceilings under 1,000 ft and especially ceilings under 500 ft.

The verification data in Rudack and Ghirardelli (2010) are already five years old. LAMP has been revised since then, as have the numerical models. If the model comparisons were repeated today, the results might be different.

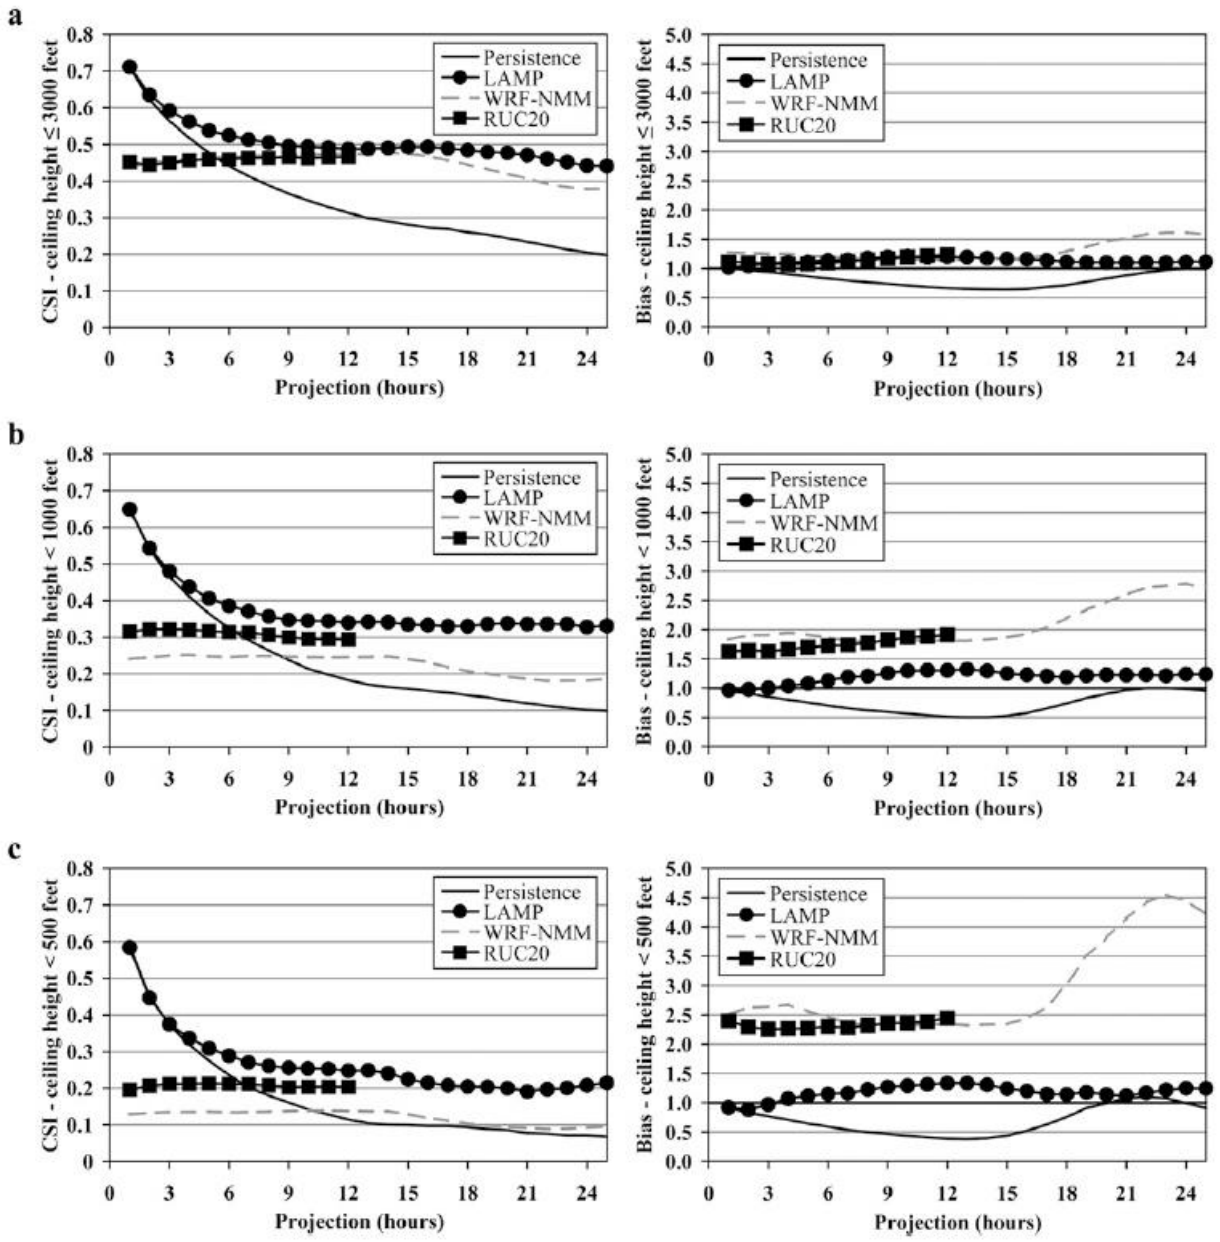

Fig. 3.22. CSI scores (left) and bias (right) for forecasts of ceiling height: (a) ≤ 3,000 ft, (b) < 1,000 ft, and (c) < 500 ft. Forecasts are from four sources: persistence, LAMP, WRF-NMM, and RUC20. Forecasts generated from October 2006 through March 2007 are included in the data sample. From Rudack and Ghirardelli (2010). Used with permission of the American Meteorological Society.

Figure 3.23 compares LAMP and SREF forecasts of the probability that visibility will be less than three miles. Panels (a), (b), and (c) identify 6-h, 12-h, and 18-h forecasts, respectively. At left

are reliability diagrams; at right are histograms counting the number of forecasts in various probability intervals for each model, issued from October 2006 through March 2007. The hollow symbols indicate small samples sizes (less than 1% of total cases). The diagonal straight

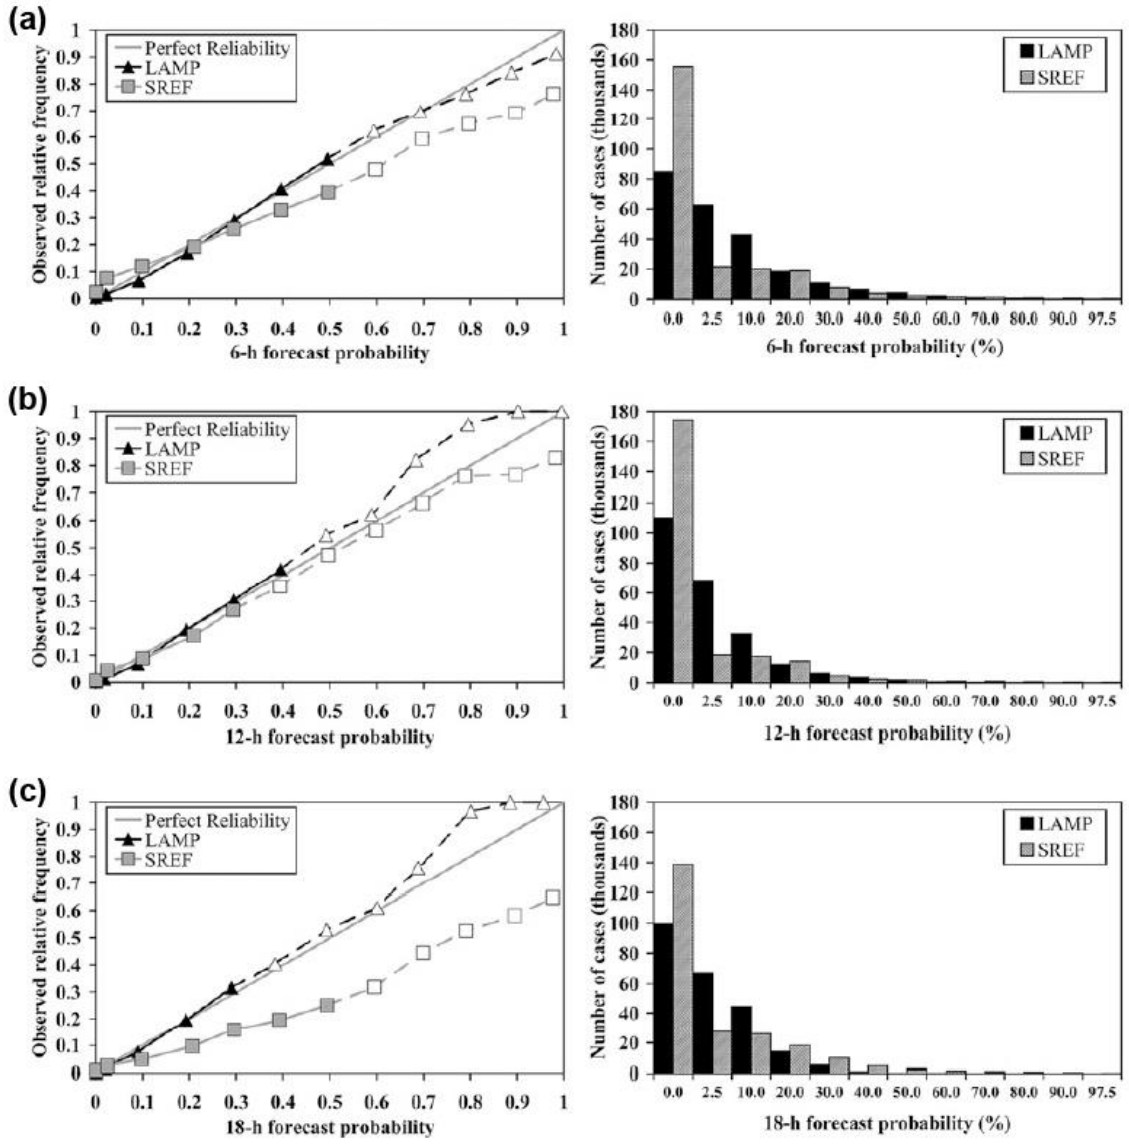

Fig. 3.23. Reliability diagrams (left) and histograms for probabilities of visibility less than three miles for (a) 6-h, (b) 12-h, and (c) 18-h projections of LAMP and SREF. Only forecasts issued at 0900 GMT are verified here. Reliability squares and triangles representing less than 1% of the total number of cases are hollow, as a reminder to be wary of small sample sizes. From Rudack and Ghirardelli (2010). Used with permission of the American Meteorological Society.

line in the reliability plot represents perfect calibration, namely, one forecasts the event with the same frequency as it is observed. If 10,000 forecasts are issued, indicating a 30% chance of visibility less than three miles, perfect reliability is when 3,000 of the verifying observations indicate visibility below three miles and 7,000 indicate greater visibility. The LAMP forecast is very well calibrated for probabilities between 0 and 60%. The SREF is less well calibrated; it overforecasts the probability of visibility under three miles.

The histograms show that most forecasts of visibility of less than three miles are associated with low probability, less than 20%. Since visibility less than three miles is a rare event, this is appropriate. The SREF model generates more forecasts with very low probabilities (< 2.5%) than the LAMP.

3.1.6.6 Commentary on LAMP ceiling and visibility forecasts

Advantages: Primarily through its use of METAR observations, LAMP improves upon GFS MOS, particularly during the first quarter of its 24-h forecast period. LAMP generates probabilistic guidance by means of the regression equations it employs. These are based upon long-term performance statistics for the GFS model. LAMP corrects for systematic biases in GFS forecasts. Because LAMP uses output from a global model, the procedures employed to generate ceiling and visibility forecasts are applicable anywhere in the world. LAMP generates probabilistic forecasts from a single model run; ensemble forecasts are not required. Paul Herzegh (NCAR) has performed recent verification of ceiling and visibility forecasts by LAMP and the RUC and RAP models. His results are similar to those resulting from the LAMP vs. model comparison of nearly five years ago and shown in Fig. 3.23: the LAMP forecasts of ceiling and visibility almost always beat the model forecasts in the first six hours.

Limitations: Output from the GFS model is not processed by LAMP until more than three hours after the initial time, the model runs once every six hours, and the output is available only at three-hour intervals. If a mesoscale model such as the RUC or RAP were used (these run every hour), output would be available within one hour of the initial time and at 1-h intervals. Another advantage of using a mesoscale model is the higher spatial resolution, which incorporates more precise terrain elevation data and allows more sophisticated prescriptions of moist physical processes important for predicting ceiling and visibility. The GFS model (spectral resolution T574, triangular truncation) has roughly 35-km resolution. The RUC and Rapid Refresh models have 13-km resolution. The main obstacle to using mesoscale models is that

they evolve with time, and it is difficult to develop long-term, stable statistics for use in LAMP because the training period is long.

3.1.7 Ceiling and Visibility Analysis

The Ceiling and Visibility Analysis (CVA), also sometimes called the National Ceiling and Visibility Analysis (NCVA) product, is used very heavily by general aviation and helicopter pilots because they are particularly vulnerable to the impacts of low ceiling and visibility conditions, both en route and at airports. The CVA is produced on the 5-km National Digital Forecast Database (NDFD) grid. Its main purpose is to provide ceiling and visibility information at arbitrary locations within the CONUS *between* METAR locations. Information for this description was extracted from informal notes by Herzegh and Cowie (2011) of NCAR.

The CVA updates every five minutes. It relies on METAR observations, transmitted hourly by default but more often whenever ceiling and visibility pass through thresholds that affect flying conditions. Some of the METARs lie just across the border in Canada and Mexico, but product coverage does not extend far beyond the border, nor does it extend much beyond the U.S. seacoasts. The CVA also relies on GOES satellite cloud images, primarily to confirm where there is cloud cover.

3.1.7.1 Use of METARs in CVA

Visibility at METAR sites is measured by a forward-scatter visibility sensor. The instrument generates pulses of light, some of which is scattered toward a detector at an angle to the forward direction (usually between 20^{o} and 50^{o}). The amount of light scattered depends upon the number density and sizes of particles within the beam—aerosols, fog droplets, snowflakes, etc. The instrument is calibrated to deliver a visibility measurement as a function of the amount of scattered light detected. The METAR visibility report is a 10-min average of 1-min measurements; it is assumed to be representative of conditions within two to three miles of the sensor. The instrument cannot distinguish among visibilities greater than 10 miles.

METAR ceiling measurements are made with a laser beam ceilometer, which emits very short pulses of light and measures energy backscattered from various altitudes. Cloud base usually supplies a strong signal unless the cloud is tenuous or very thin. The ceiling height is a 30-min average computed from 30-s samples. Samples from the last ten minutes are double weighted.

The number of samples that are “cloudy” helps determine cloud coverage: clear, partly cloudy, broken, or overcast. Cloud observations are assumed to be representative of conditions within three to five miles of the site. The ceilometer does not measure cloud heights above 12,000 ft AGL.

Ceiling and visibility are notoriously patchy fields, that is, large changes in both are possible within distances of less than a kilometer. For this reason, analysis of these parameters on a 5-km grid using linear interpolation of values from nearby METAR sites is inadvisable because linear interpolation results in smoothing. Instead, a *nearest-neighbor* interpolation scheme is used. At each grid point, provisional values of ceiling and visibility are taken from the nearest METAR station report, following a simple quality check of the sensor data. While there are still errors associated with nearest-neighbor interpolation, verification studies show that this approach yields better skill than linear interpolation.

If the elevation of the grid point and the METAR site differ, for example, if the grid point is 1,000 ft higher than the METAR station, simple adoption of the METAR ceiling value could be hazardous because ceiling height tends to follow isentropic surfaces more closely than terrain. This implies that the ceiling will probably be lower at the grid point than at the METAR site. To avoid this problem, the METAR ceiling value is first converted from AGL to MSL. The elevation of the grid point is then subtracted from the METAR MSL ceiling value to obtain the grid point AGL ceiling value. If this process results in a ceiling value at the grid point of zero or less, the diagnosis is “terrain obscured by cloud.”

If the ceiling or visibility measurement is missing from the regularly reporting METAR site nearest the grid point, the measurement(s) from the next nearest METAR site is used.

3.1.7.2 Use of GOES satellite data in CVA

Research is under way to deduce regions of low clouds and poor visibility from data provided by GOES and polar orbiting satellites. This research will be described at the end of Section 3.1. For the present, however, the greatest utility of GOES images and sounder data is in distinguishing clear areas from cloudy areas. GOES views atmospheric cloudiness from above; METAR sensors “look” at clouds from below. The object is to corroborate the two views at METAR sites and exploit the high resolution and high frequency of GOES images to infer cloud conditions in between METAR sites.

Jedlovec et al. (2003) wrote an algorithm for detecting clouds in GOES images and sounder data. Clouds have different emissivities when viewed at wavelengths of 10.7 and 3.9 μm. The emissivity of clouds is substantially less at 3.9 μm than at 10.7 μm. During the day, however, the reflected solar radiation from clouds combined with cloud emissivity at 3.9 μm, makes the brightness temperature BT(3.9 μm) much greater than BT(10.7 μm). At night, when reflected solar radiation no longer contributes to BT(3.9 μm), BT (3.9 μm) < BT (10.7 μm). During the day in clear areas, the two brightness temperatures are roughly comparable. This behavior is the basis for the Jedlovec et al. (2003) algorithm, which is known as the Bi-spectral ThresHold and Height (BTH) approach. Estimating cloud-top height is not a concern here, so only the first part of the algorithm will be described. One complication is that cloud emissivity varies with cloud type and solar zenith angle.

The algorithm, which runs every 15 min, performs four tests, each one a refinement of the one before, in order to decide whether a pixel is cloudy or clear. The tests involve several quantities defined here:
DI(i) = BT(10.7 μm) – BT(3.9 μm)
DI is the difference in brightness temperature between the longwave channel (10.7 μm) and the shortwave channel (3.9 μm), measured at pixel i, located along a scan line in a GOES image. A difference image can be created from the DI values at each pixel location.

For each hour of the day, difference images are examined over the preceding 20 days. Three new images are created:

- CNDI refers to an image in which the smallest *negative* difference (closest to zero) observed over 20 days at a given hour appears at each pixel location.
- CPDI refers to an image in which the smallest *positive* difference observed over 20 days at a given hour appears at each pixel location.
- CWLW refers to an image in which the greatest BT(10.7 μm) observed over 20 days at a given hour appears at each pixel location.

Figure 3.24 illustrates the logic of the BTH algorithm. In Step 1, the brightness temperature differences at two adjacent pixels are compared. If the difference is large and negative, then it is almost certainly daytime, and a cloud at pixel i is assumed to be responsible.

Step 2 deals with smaller differences observed during the daytime. First, determine whether pixel i-1 is cloudy or potentially clear by forming the difference DI(i-1) – DI(i-2) (a repeat of Step 1 but with the two grid points shifted one grid interval). At this point, pixel i is potentially clear. If pixel i-1 is cloudy, check for condition DI(i) – DI(i-1) < 0.0 K. If satisfied, pixel i is probably

cloudy; the 3.9-μm reflectance at pixel i is probably greater than that at pixel i-1. If pixel i-1 is potentially clear, check for condition DI(i) – DI(i-1) < –3.1K. If satisfied, pixel i is probably cloudy; the 3.9 μm reflectance from clouds at pixel i is greater than that from a potentially clear sky at pixel i-1.

Step 3 begins with pixel i still potentially clear. If DI(i) is negative (likely for high sun angles), and if it is at least 5.1K more negative than the smallest negative value, CNDI(i), observed at the same pixel in the last 20 days at the same hour of day, then pixel i is probably cloudy. This test is for daytime. If DI(i) is positive, and it is least 2.0K more positive than the smallest positive value, CPDI(i), observed at the same pixel in the last 20 days at the same hour of day, then pixel i is probably cloudy. This test is for nighttime.

Step 4 is the final test. Again it begins with pixel i potentially clear. Let LW(i) represent the current value of BT(10.7 μm). Barring continuous cloudiness for 20 days at pixel i, CWLW(i) approximates the highest ground temperature observed there during that period. If the condition CWLW(i) – LW(i) > 18.5K is satisfied, then pixel i is assumed to be cloudy. Otherwise, it is assumed to be clear. Figure 3.24 shows the four-step testing for clear vs. cloudy pixels.

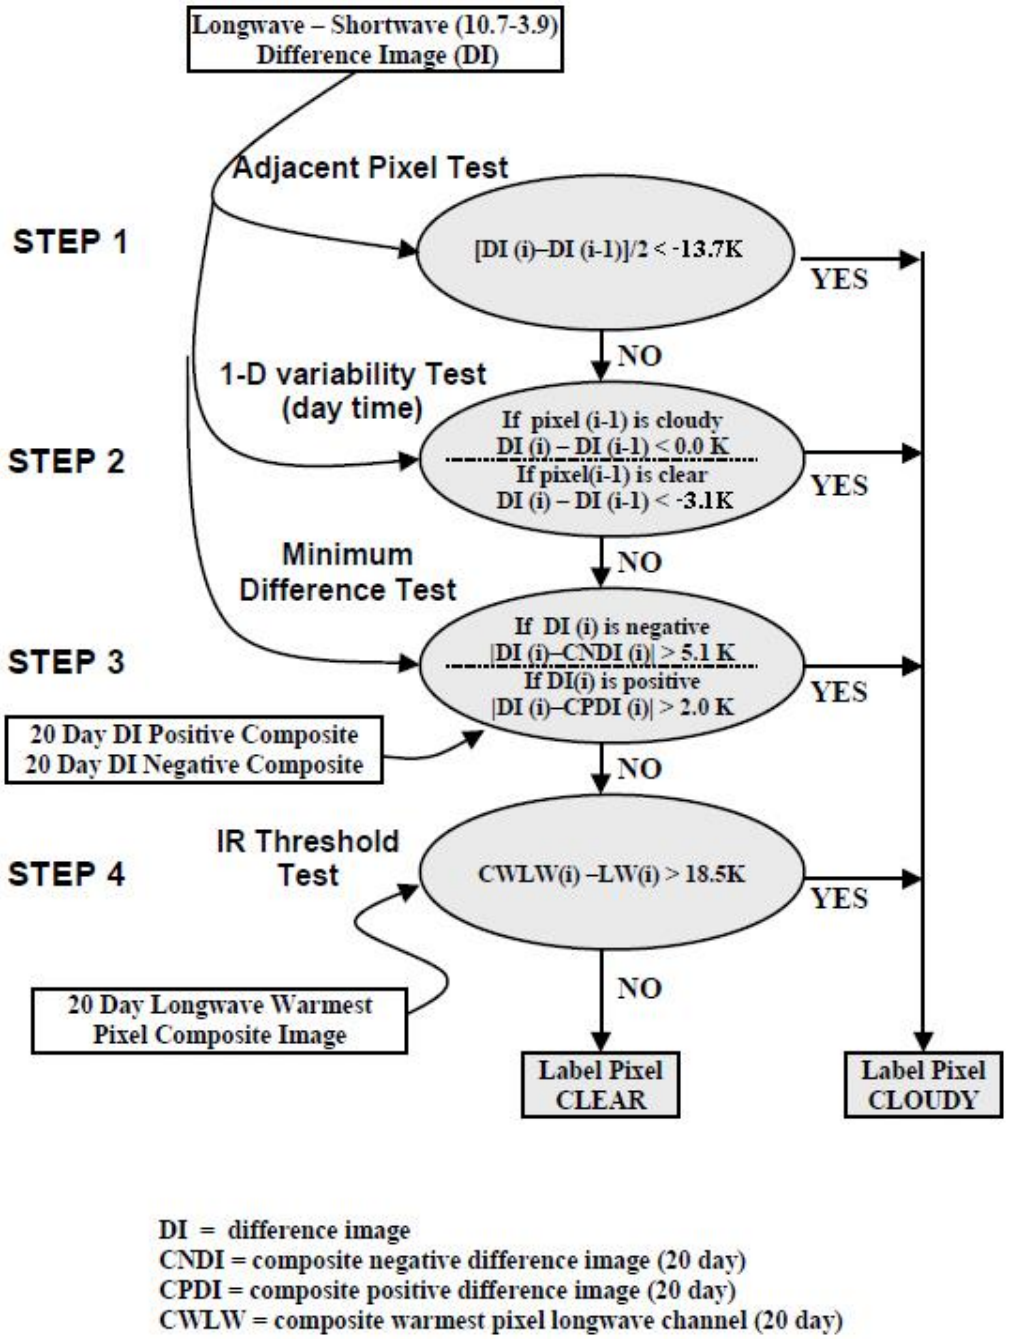

DI = difference image
CNDI = composite negative difference image (20 day)
CPDI = composite positive difference image (20 day)
CWLW = composite warmest pixel longwave channel (20 day)

Fig. 3.24. Logic for the BTH algorithm by Jedlovec et al. (2003) for determining whether a pixel is clear or cloudy. Details are in the text. The scheme is used with slight modification of several of the numeric thresholds shown. Figure used with permission of the American Meteorological Society.

The four-step procedure above applies to both GOES East and GOES West images. Both satellites observe the zone between 91^{o} and 113^{o}W longitude. If tests on data from either satellite declare a pixel "cloudy," the pixel is designated "cloudy." Outside of this zone, data from the satellite that has the best viewing angle determine whether a pixel is clear or cloudy.

All the "cloudy" pixels in a CONUS image define a *cloud mask*, which is compared with cloud observations from METARs. The comparison is necessary because satellite cloud products sometimes fail to detect very low clouds in the day-night terminator region and at night. This happens when the infrared BT for the cloud is close to that for the ground, for example, in a nocturnal inversion. Every five minutes, the cloud mask is compared with the METAR cloud observation. If the METAR reports a ceiling, then every pixel with a 10-km radius must be "cloudy." If not, the satellite information is discarded within the nearest neighbor domain of the METAR. Details of the comparison are given in Table 3.8.

Table 3.8

Actions resulting from comparison between the GOES-derived cloud mask and METARS. Courtesy of Paul Herzegh.

Region of Comparison	METAR Indication	GOES Indication	Comment	CVA Action in METAR Nearest Neighbor Domain
Inside 10-km radius around METAR site. (Note: This region is typically a fraction of the METAR's nearest neighbor domain.)	Clear to 12K ft AGL (no ceiling exists).	Some or all GOES grid points show clear conditions.	Assume no ceiling exists below 12K ft. GOES detection assumed to be related to higher level cloud.	Ceiling height 'clear' assigned at grid points shown clear by GOES. Ceiling height '12K or higher' assigned at grid points shown cloudy by GOES.
		GOES data unavailable.	Applies if GOES data is missing or older than 30 min.	Ceiling height '12K or higher' assigned to all grid points.
	A ceiling is found at or below 12K ft AGL.	All GOES grid points show cloud.	No disagreement between METAR & GOES. GOES data is assumed to be representative of cloud cover.	Ceiling height 'clear' assigned for grid points outside 10 km radius shown clear by GOES, if any. METAR ceiling height (with routine terrain adjustment) is assigned for all remaining grid points.
		Some or all GOES grid points show clear conditions	GOES cloud detection may have failed, or cloud cover is scattered.	METAR ceiling height (with routine terrain adjustment) is assigned throughout METAR domain.
		GOES data unavailable.	Applies if GOES data is missing or older than 30 min.	
	METAR ceiling report missing.	Any	No comparison with GOES cloud mask is possible.	Ceiling height from the next-nearest METAR is assigned. New comparison test is run for the next-nearest METAR.

3.1.7.3 CVA sample products

Samples of regional CVA products are shown in Figs. 3.25 (ceiling) and 3.26 (visibility). A national product is not shown because it lacks the METAR details and only shows the shading for regions with terrain obscuration, IFR, and non-IFR. To see similar products, go to http://weather.aero/tools/weatherproducts/cva. One can choose to view flight category, determined by both ceiling and visibility, a ceiling product alone, or a visibility product alone. One can choose to view the entire contiguous U.S. or specific regions. Finally, one can view current or previous products (up to 24 h ago).

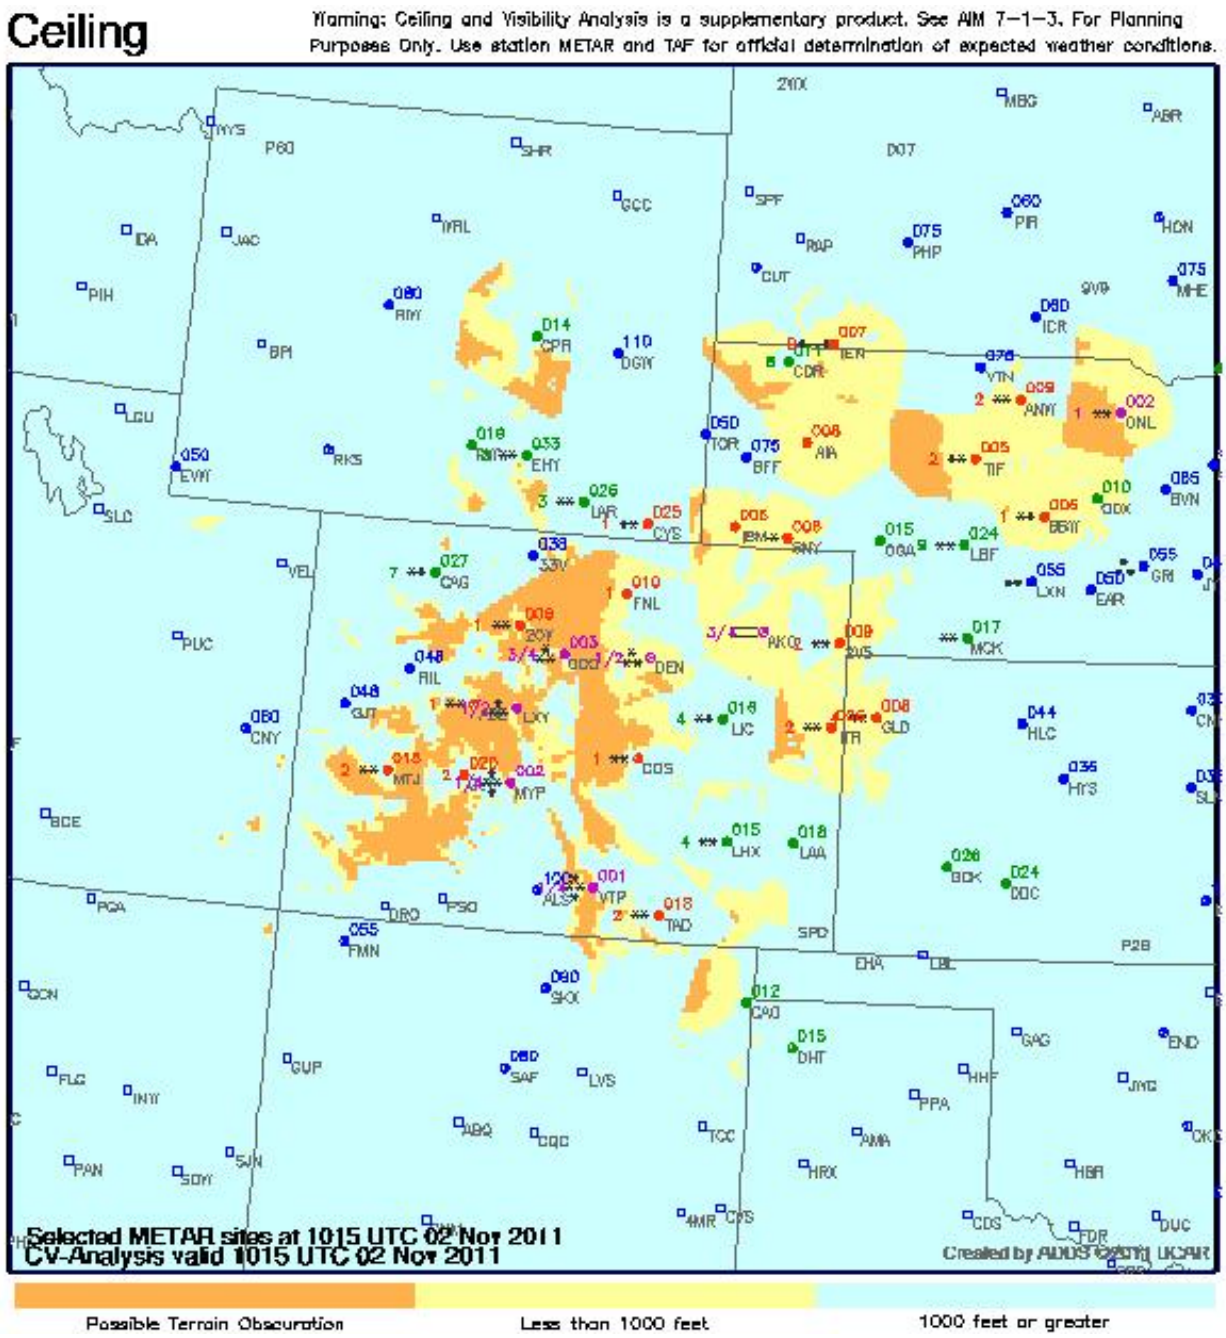

Fig. 3.25. The CVA product for ceiling for a region centered about Denver, Colorado. An early November snowstorm was affecting the region. Orange shading indicates possible terrain obscuration by clouds or fog, yellow shading indicates IFR conditions, and light blue indicates better than IFR conditions. Station identifiers, ceiling height, if any, in hundreds of feet (015 means 1,500 ft AGL), visibility (miles) and obstructions to vision are plotted at each METAR site. The colors of the station location circles indicate LIFR (purple), IFR (red), MVFR (green), and VFR (blue) conditions. Blue squares indicate ceiling above 12,000 ft or no ceiling.

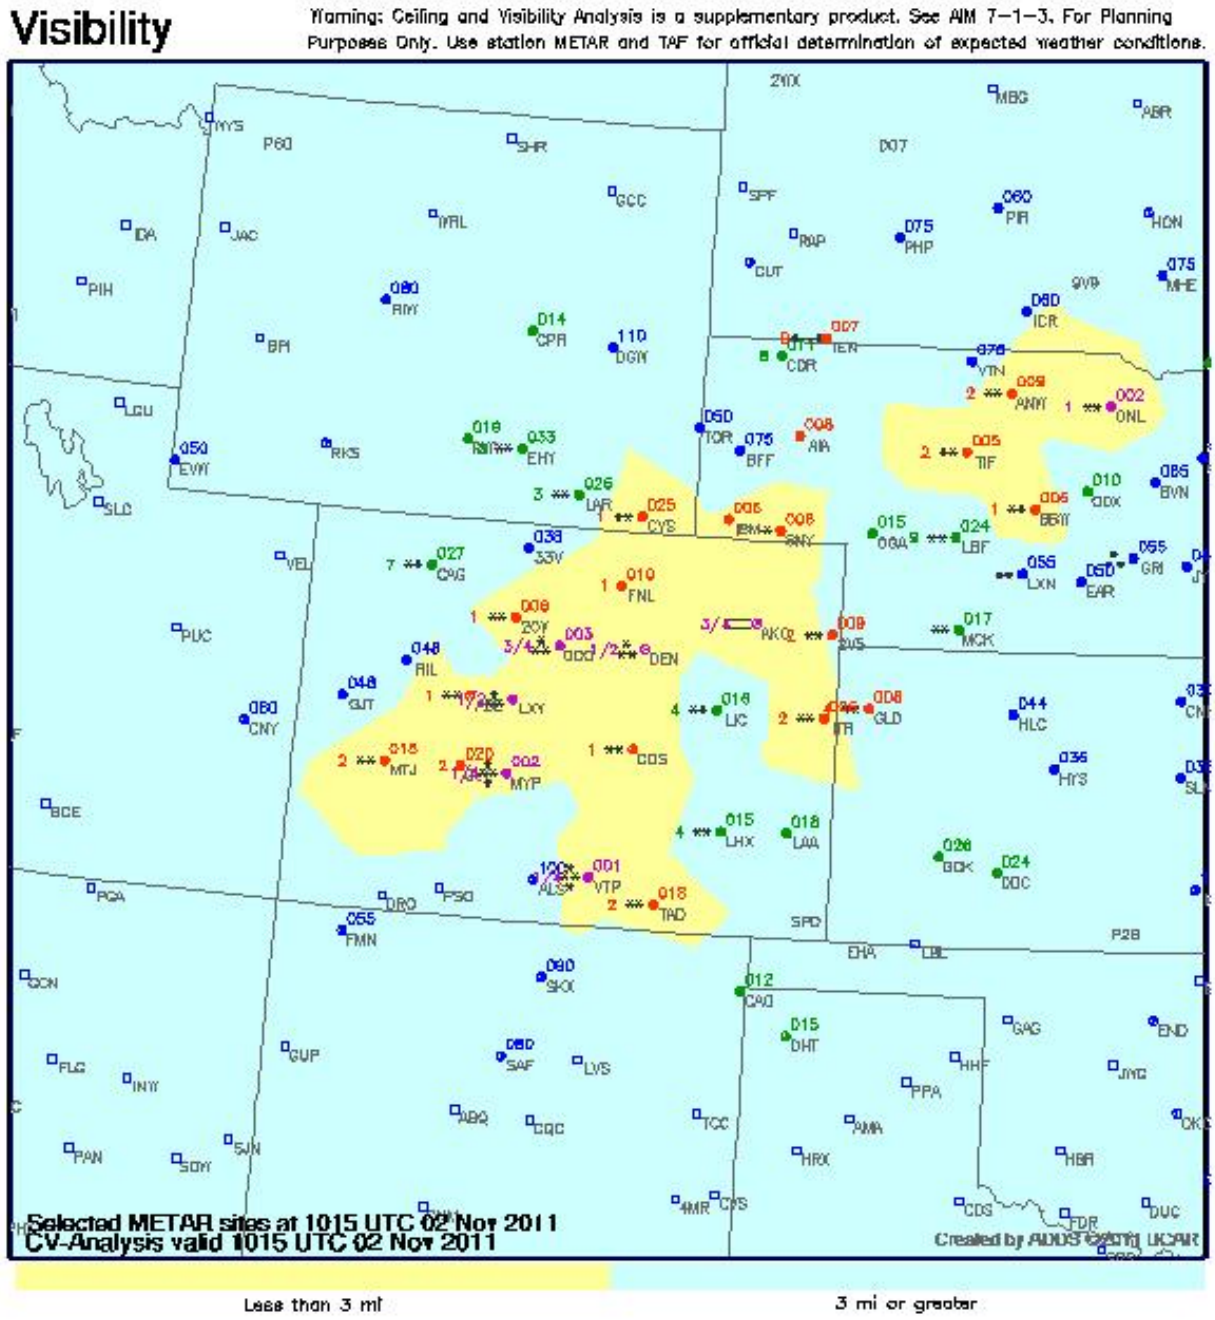

Fig. 3.26. The CVA product for visibility for a region centered around Denver, Colorado. Yellow shading indicates IFR conditions. Light blue indicates better than IFR conditions. The color coding of the plotted data is the same as for Fig. 3.25.

3.1.7.4 Verification of CVA

This section excerpts information from a quality assessment report by Loughe et al. (2010). One expects the CVA to match METAR ceiling and visibility observations at METAR locations because the CVA is based on METAR observations. To assess the performance of the CVA in between METAR locations, Loughe et al. use a cross validation technique, whereby 20% of randomly chosen METAR sites are withheld from the CVA. They then verify the CVA against ceiling and visibility observations at the withheld sites. Nine more times they randomly choose 20% of the METAR sites and withhold them from the analysis. Figure 3.27 shows the distribution of included and withheld METAR sites for just one of the ten random selections. The blue triangles show the approximately 1360 stations included in the analysis; the red squares show the approximately 340 sites excluded. METAR sites must report at least 90% of the time to be included in the CVA and the cross validation.

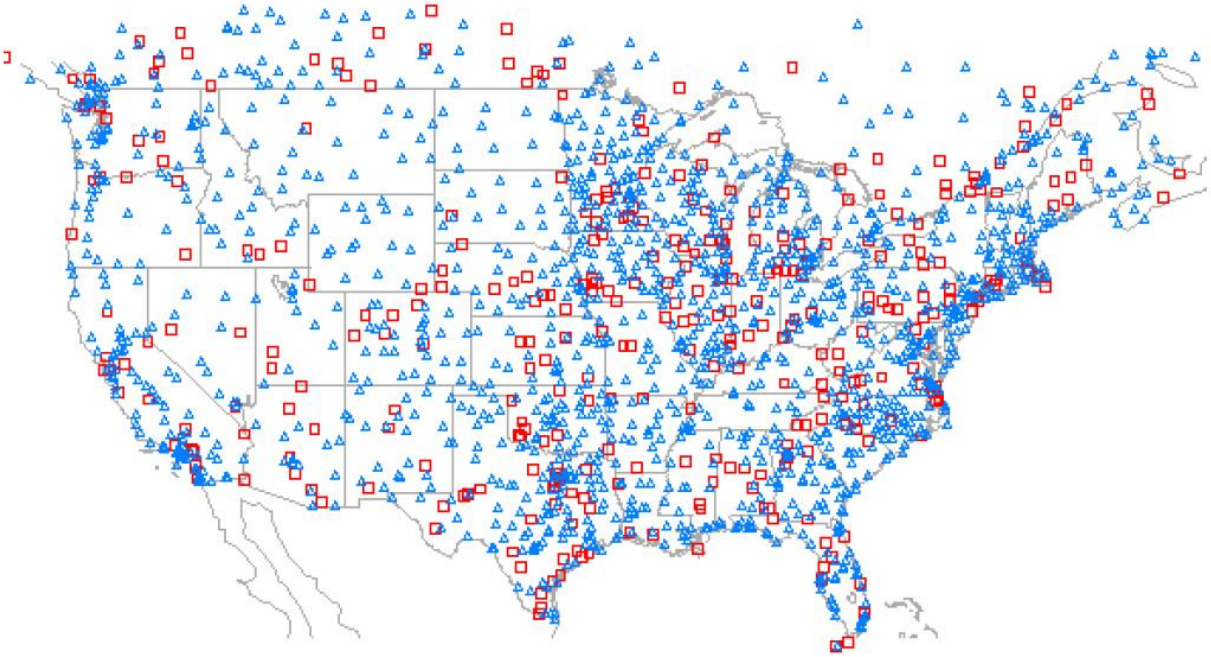

Fig. 3.27. One of ten randomly chosen distributions of METAR sites whose data are withheld from the CVA (red squares), 20% of the total number. The CVA is computed from METARs at the blue triangles and verified against METARs at the red squares—a cross validation technique.

The verification took place from June through August 2008 and December 2008 through February 2009. Scores were based on a dichotomous forecast of IFR conditions vs. non-IFR conditions (MVFR or VFR) and computed from values in a contingency table like the one below.

		Withheld METARs	
		IFR	**non-IFR**
CVA	**IFR**	YY	YN
	non-IFR	NY	NN

Figure 3.28 shows the probability of detection (PODy) of IFR conditions for all seasons and for summer and winter separately for 16 different climatological regions of the U.S. The colors used in each region indicate whether the regional score is lower or higher than the national average and by how much. In every region except WCS (West Coast South), the winter score is higher than the summer score. Along the California coast, fog is reliably present in summertime, whereas it is usually clear and dry inland. The METAR network is dense along the California coast and in the interior valleys, and so one would expect the CVA to perform well. Convection is common from the Intermountain West to the Atlantic Coast in summer, and consequently IFR conditions tend to be more patchy in summer than in winter, when large-scale storms often cause widespread low ceilings and visibility. Note that PODy scores over the mountainous West, where METAR coverage is limited, are significantly lower than east of the Rockies in all seasons. Large variations in terrain elevation limit the skill in extrapolating ceilings and visibility values more than a few kilometers from a METAR site.

Figure 3.29 is similar to Fig. 3.28 except that it shows false alarm rates (FAR) for IFR in the CVA. False alarm rates are higher in summer than in winter in all sectors except the western half of California because of the regularity of fog along the coast and the high density of METAR stations. Everywhere else in the country, the spotty nature of convection means that IFR conditions are highly localized. The situation is worst in the Intermountain West, primarily because of the rugged terrain and frequent obscuration of peaks by clouds and precipitation.

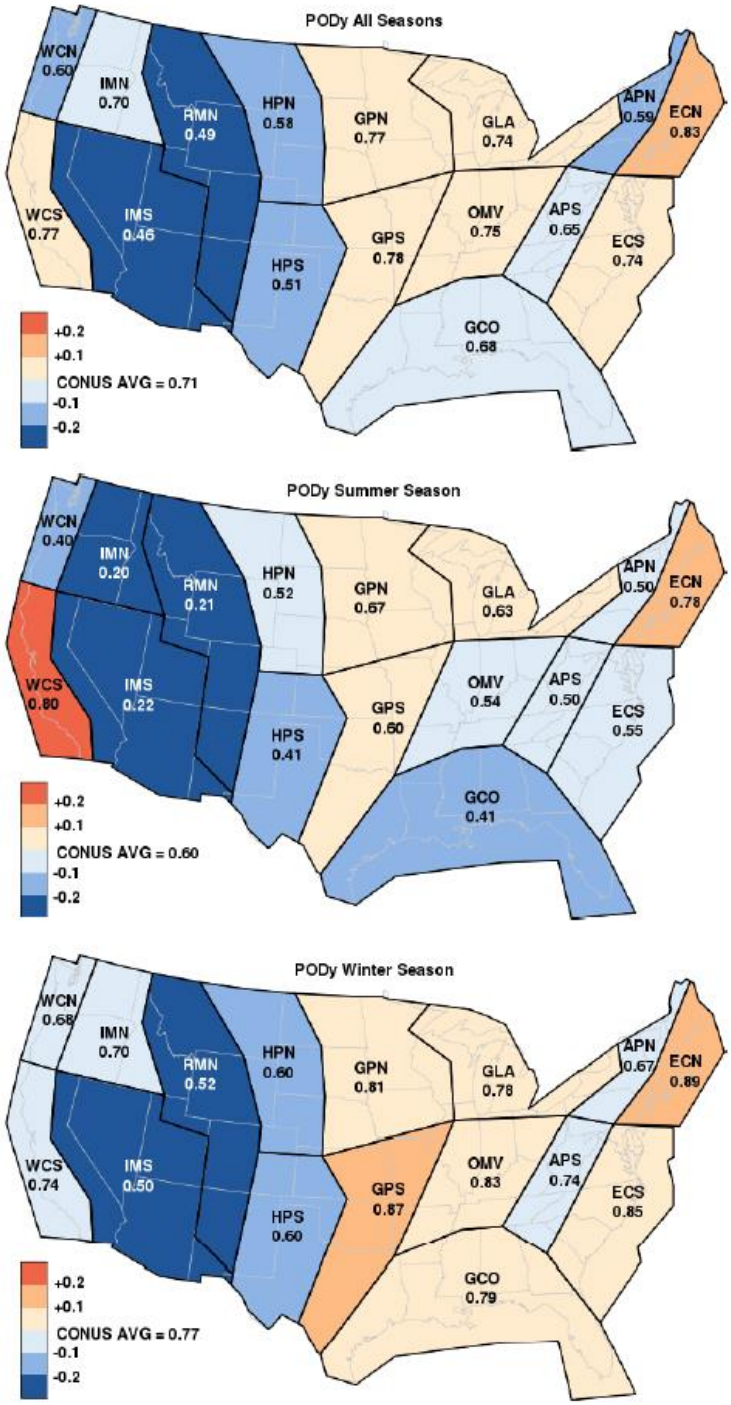

Fig. 3.28. Probability of detection (PODy) of IFR ceilings in the CVA for summer and winter combined (top), summer only (middle–June through August 2008), and winter only (bottom–December 2008 through February 2009). PODy scores are indicated for each of 16 regions. Colors indicate whether the regional score is higher or lower than the national average and by how much. From Loughe et al. (2010).

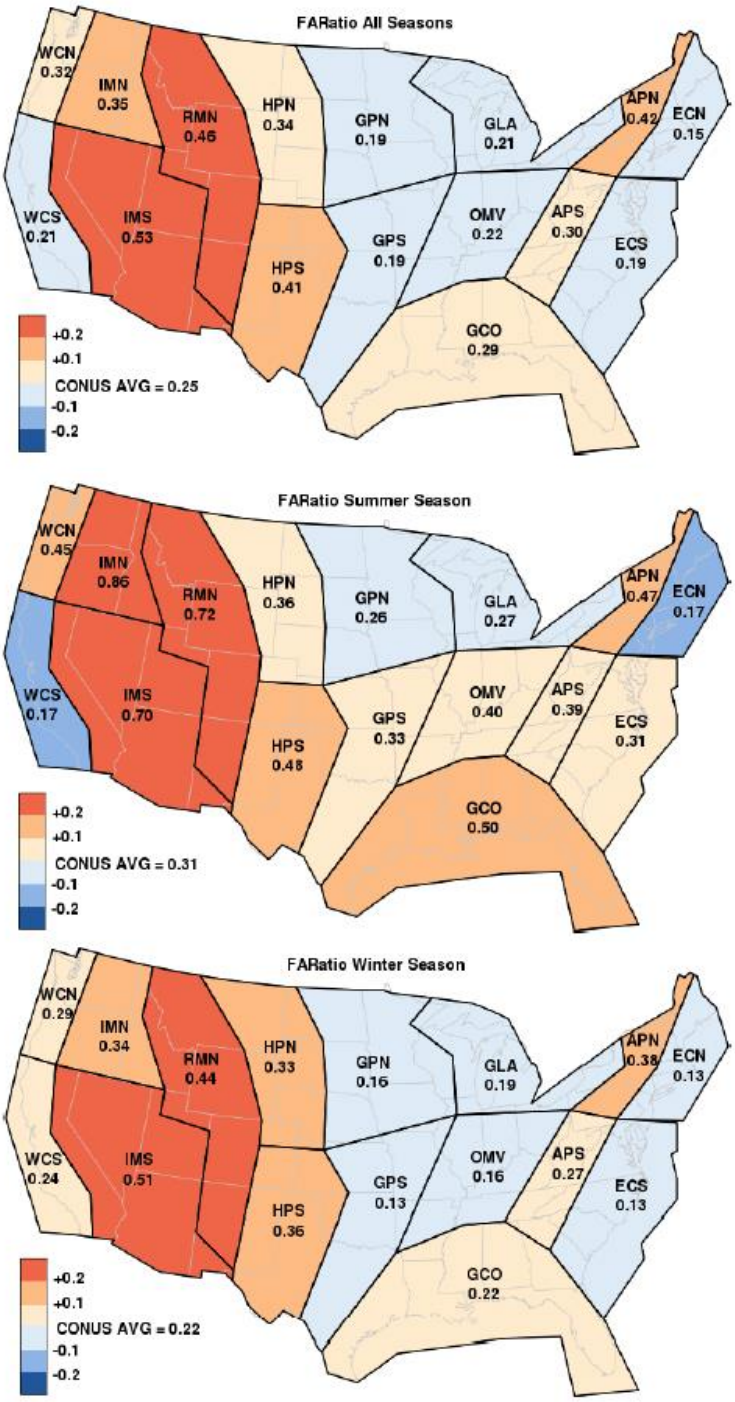

Fig. 3.29. Same as Fig. 3.28, except for false–alarm ratio. See text for interpretation. From Loughe et al. (2010).

Loughe et al. (2010) also investigated the effect of the cloud mask on the CVA. By design, the cloud mask can alter the flight category classification in only one direction: if the CVA indicates IFR conditions at a grid point solely on the basis of a METAR some distance away, and the satellite indicates that the sky is clear over the grid point, the analyzed condition at the grid point is revised to "non-IFR." If the original CVA decision (without the cloud mask) was correct (YY), then the changed decision results in a miss, an undesirable outcome. If the original decision was incorrect (YN–a false alarm), then the changed decision results in a correct negative (NN), a desirable outcome. The situation is illustrated in Fig. 3.30.

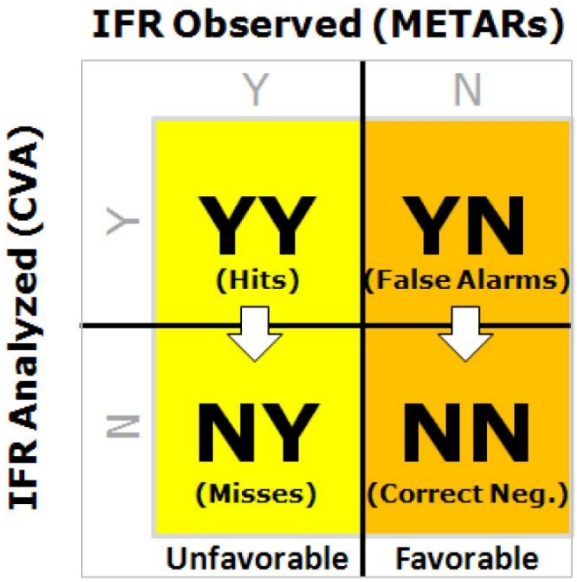

Fig. 3.30. Illustrating how the cloud mask can change the CVA ceiling category in only one direction: from IFR to non-IFR. When this happens, hits are converted to misses (unfavorable), or false alarms are converted to correct negatives (favorable).

Table 3.5 illustrates how the cloud mask alters the scores for the CVA. It increases the number of misses slightly, overall, by about half a percent, but it decreases the number of false alarms by more than 5%. The positive effect of the cloud mask is slightly greater in summer than winter and slightly greater in daytime than at night, primarily because the algorithm by Jedlovec et al. (2003) works better during the daytime at discriminating between low clouds and the ground. Infrared images alone are less capable of doing so at night.

Table 3.5

Change in misses and false alarms for the contingency table shown in Fig. 3.30 following the application of the cloud mask in the CVA. The top figure is the number of category changes (IFR to non-IFR). The number in parentheses is the percentage of values changed.

Time Period of Study	*Increase* in Misses with Application of Cloud Mask	*Decrease* in False Alarms with Application of Cloud Mask
Summer and Winter	137 (0.52%)	1,356 (5.87%)
Summer (July 2008) (06, 09, 12, 15, 18 GMT)	63 (0.53%)	635 (7.35%)
Winter (January 2009) (06, 09, 12, 15, 18 GMT)	74 (0.51%)	721 6.52%)
Daytime (15 and 18 GMT)	28 (0.33%)	503 (6.52%)
Nighttime (06 and 09 GMT)	81 (0.72%)	524 (5.52%)

3.1.7.5 Commentary on CVA

The CVA improves upon Fig. 3.10 (a display of flight conditions at METAR sites overlaid on a visible satellite image) in that it gives the same information on a 5-km grid, thereby providing flight conditions in between METAR sites. The cloud mask adds slightly to the number of misses (failure to detect IFR) but reduces the number of false alarms significantly. Cross validation experiments indicate high probability of detection of IFR conditions (~0.80) and low false–alarm rates (~0.20) during winter in non-mountainous areas. Performance during the summer convective season and in mountainous areas all year round is less impressive.

The most fundamental barrier to improved performance is the fact that large variations in ceiling and visibility often occur at spatial scales of a few kilometers or less. Average spacing between METARs is many tens of kilometers. Current data from GOES satellites provide high-resolution information about cloud cover, but methods for detecting patchy low clouds and low visibility are limited by the available instrumentation. Instruments with greater spectral, temporal, and spatial resolution aboard GOES-R, now scheduled for launch in 2015, should improve this situation. For example, the detection of radiation fog, almost exclusively a nighttime phenomenon, should improve with GOES-R.

The shading on the regional images (Figs. 3.25 and 3.26) discriminates only between IFR and better-than-IFR conditions. The national scale product (not shown) is the same; it shows only IFR, non-IFR, and terrain obscuration. Both regional- and national-scale products could be improved by adding enough colors to portray all four categories: LIFR, IFR, MVFR, and VFR. The discerning reader will have noticed that the plotted data in Figs. 3.25 and 3.26 *are* color-coded according to LIFR, IFR, MVFR, and VFR, but these are the only indicators of the finer distinction and are easily missed.

Until recently, the CVA product displayed confidence fields. The shading of the product indicated either normal or low confidence in the analyzed values of ceiling and visibility. The confidence level depended upon several criteria:

- If a grid point lies within R kilometers of a METAR site, confidence is normal. The minimum R, 10 km, is assigned when the ceiling at the METAR site is less than 8,000 ft or visibility is less than seven miles –OR– for any value of ceiling or visibility when the relative humidity is high. With decreasing values of humidity and a ceiling ≥ 8,000 ft and visibility is greater than or equal to seven miles, R increases up to a maximum of 40 km.
- If the flight category between neighboring METAR sites is different, there is no way to know where the transition occurs, so low confidence exists at any point more than 10 km distant from both sites.
- If the elevation difference between the grid point and the nearest METAR exceeds 300 ft, the confidence at the grid point is rated low.
- Outside of radius R, the GOES cloud mask determines whether the confidence is normal or low. During the day, the decision is somewhat complicated. At sunrise, sunset, and during the night, confidence is low because GOES detection of fog, low cloud, or thin cloud may be unreliable.

The confidence values have recently been removed, presumably because the METAR sites appeared to be islands of confidence in a sea of reduced confidence. Without the confidence fields, GA pilots especially must be vigilant for abrupt changes in ceiling and visibility along low–altitude routes that may not be captured by the CVA.

Implementation of the CVA as an operational product (with *supplementary* status) at the AWC is scheduled to occur in the latter half of 2012.

3.1.8 Ceiling and Visibility Forecast

So far, there is no Ceiling and Visibility Forecast (CVF) to complement the Ceiling and Visibility Analysis (CVA), but one is under development at the Research Applications Laboratory (RAL) of NCAR. The aim of the CVF is to generate an optimal blend of three different information sources on ceiling and visibility: LAMP, described earlier, a *time-lagged ensemble* of forecasts from the Rapid Refresh (RAP) model, and an "obcast."

Time-lagged ensemble. Because the RAP runs hourly out to 18 h, it generates multiple forecasts of varying lengths that verify at the same time. A subset of these forecasts constitutes a time-lagged ensemble, whose average usually results in a more accurate forecast than the forecasts provided by individual members. This is a poor-man's ensemble in that its members are not statistically independent: they are temporally correlated. A more proper ensemble would contain forecasts all of the same length, each starting from different initial conditions and perhaps incorporating different dynamics and physics. Such an ensemble is prohibitively expensive because it would require far more computing resources than are currently available.

Obcast. The obcast is a statistical projection of current conditions at METAR sites from one to five hours ahead. The obcast is based upon logistic regression, which is briefly described in Appendix B, Section 3.1.8. The probability that ceiling or visibility will be in a specific category one to five hours hence is estimated from current observations at each METAR site and its two nearest neighbors.

Several steps lead to the CVF. Step 1 determines what blend of LAMP and the RAP time-lagged ensemble forecast provides the best forecast. Since the objective is to produce a skillful probabilistic forecast, several metrics specific to probabilistic quantities have been examined. Based on requirements for good performance and computational simplicity, the Brier Skill Score (BSS, discussed in Section 2.3.9), which yields a generalized measure of forecast improvement over climatology, is used to guide the CVF optimization process. By evaluating the BSS over a look-back period for a series of candidate weights in real time, an optimum blend yielding the highest value of BSS is found. An optimum blend is determined for each variable (ceiling or visibility), each projection (one to ten hours), each METAR site, and each valid time. The blending weights are updated through a fresh look-back every hour. The *length* of the look-back period is dynamically determined each hour, and typically varies from one to as many as twenty days, depending upon how frequently the condition being forecast has recently occurred.

The probability forecast that results from merging LAMP and a RAP time-lagged ensemble forecast can be labeled the LAMP-Model Blend (LMB). Step 2 in generating the CVF is to merge the LMB with the obcast. The merge occurs only for forecast projections of one to five hours because the obcast has little residual skill at later times. The weights for merging the LMB with the obcast are determined over a full season. Unlike the look-back weights, which are updated hourly, these weights are fixed for at least a season. The main advantage of doing so is that the METAR observations exert a relatively fixed influence on the first five hours of the forecast, a desirable trait, because persistence of ceiling and visibility conditions is usually a very good forecast during the first few hours.

Figure 3.31 hints at the superiority of the experimental CVF over the operational LAMP, which occurs most of the time. This example is for forecasts of ceiling below 500 ft at Boston, made from 1 November 2010 through 22 June 2011. Note that model forecast skill (small blue squares) rises with lead time of the forecast. The model forecast is competitive with CVF beyond six hours, at least for this one case.

Notice the short "wings" on the model squares at two and six hours. These are "mini-curves" and they represent the spread in the time-lagged ensemble RAP model forecast. The wings are very short (compared to the much broader LAMP curve for example) because there is very little spread or dynamic range in the ensemble forecast itself. The lags are too similar to produce much spread. There are no "wings" at 10 h because the most recent deterministic 10-h forecast is usually better than the time-lagged ensemble of 10-h forecasts, and so the former forecast is used.

As explained earlier, the CVF is computed only at METAR sites. The extension of CVF to a fine-scale national grid can be accomplished in the same way as the original LAMP was extended from METAR sites to a 2.5-degree grid. In fact, the techniques employed in generating the CVF may be subsumed into LAMP as an enhancement. NCAR/RAL and NWS/MDL are discussing this option.

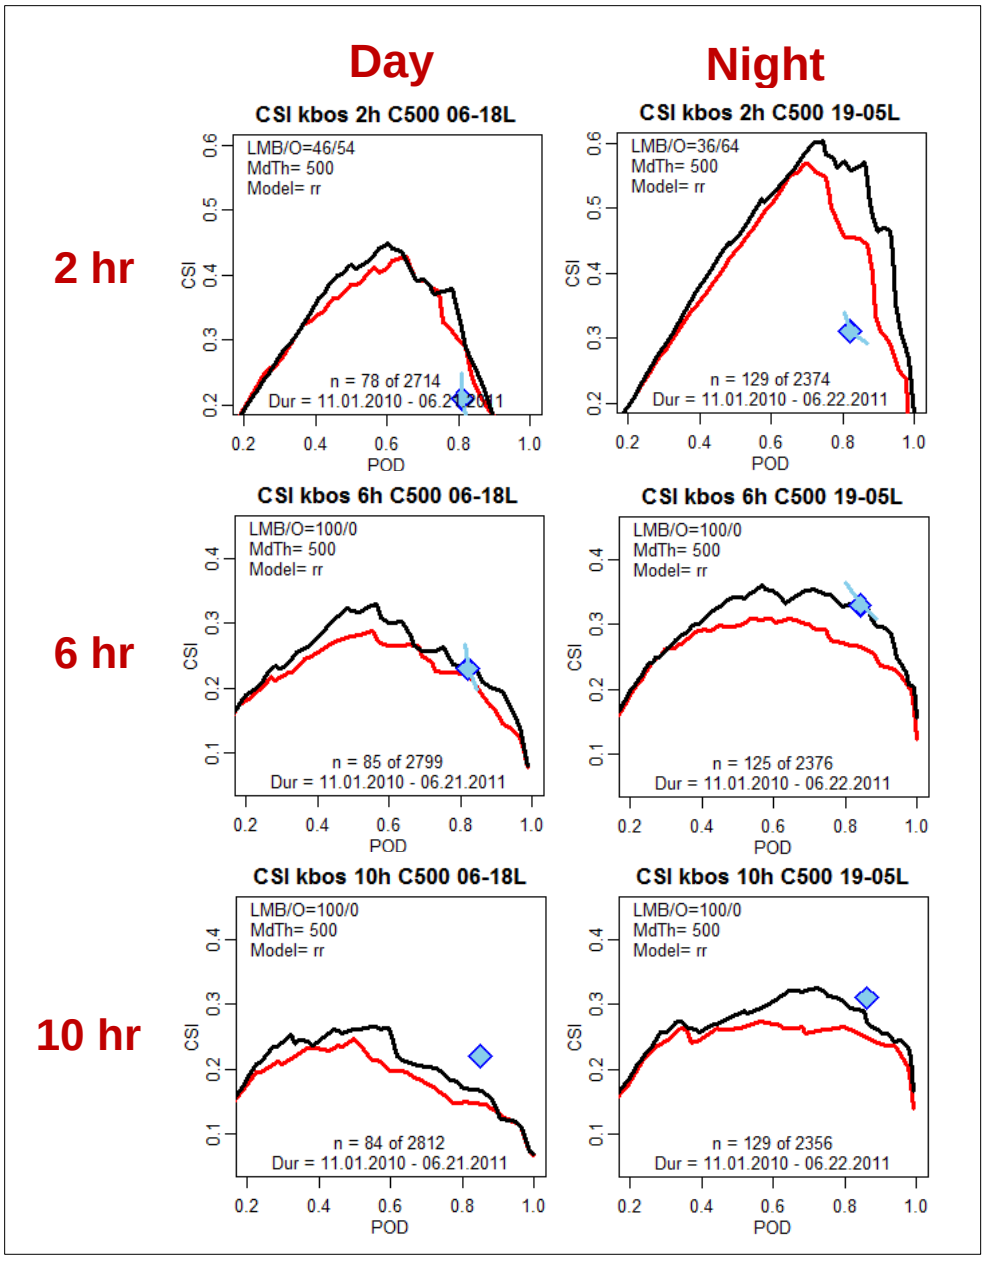

Fig. 3.31. Verification of probability forecasts for a ceiling of less than 500 ft at Boston, Massachusetts. Verification period: 1 November 2010 through 22 June 2011. Daytime forecasts at left verify between 0600 and 1800 local time. Nighttime forecasts at right verify between 1900 and 0500 local time. The length of the forecasts increases from top to bottom: 2 h, 6 h, and 10 h. Each graph plots CSI vs POD. The red curve is for LAMP; the black curve is for CVF. The light blue squares represent the RAP time-lagged ensemble ceiling forecasts after correction for forecast bias.

3.1.9 Helicopter Emergency Medical Services (HEMS) Java Tool

The Helicopter Emergency Medical Services (HEMS) Java Tool provides guidance to helicopter pilots involved in rescue operations or ferrying accident victims to hospitals. Helicopters can operate with low ceilings, but they cannot fly safely in clouds. Of 27 fatal accidents reported by commercial operators from 1998 through 2004, 21 occurred at night. Of these, 16 accidents ensued after helicopters flying in VFR conditions suddenly encountered instrument meteorological conditions and collided with the terrain. Other accidents were caused by loss of control or poor visibility.

Though they don't normally deal with emergencies, helicopter pilots in the Gulf of Mexico comprise another large class of HEMS users. There are approximately 4000 helicopter operations per day (average duration of 20 min) by 650 helicopters in the northern Gulf of Mexico. 350 of these helicopters belong to Petroleum Helicopters, Inc. with headquarters in Lafayette, Louisiana, but there are many smaller operators, too. The helicopters ferry personnel and supplies back and forth between the coast and oil rigs. (Heavy supplies go by ship.) The FAA has installed Automated Weather Observing Systems (AWOS) at some of these sites. Fifteen AWOS were already deployed by fall 2010, and 20 more were planned. These provide valuable offshore flight information for pilots.

Unlike many aviation weather products, whose development is supported by the FAA's Aviation Weather Research Program, the HEMS product is supported by FAA Flight Standards. HEMS presents a variety of weather graphics available on ADDS and Experimental ADDS through an enhanced format that includes Geographic Information System (GIS) background information and high–resolution displays. Together, these features enable helicopter pilots to better examine available data as they plan and conduct operations. Real-time ceiling and visibility information in HEMS is the diagnosis provided by CVA, already described in Section 3.1.7, so the story of its development needn't be repeated here. Future CVF ceiling and visibility forecast data will also appear in HEMS when this product reaches operational status. A few of the unique display capabilities of HEMS are described here because HEMS is the only aviation weather product specifically designed to support helicopter operations.

To access HEMS products, go to http://weather.aero, click on "Desktop Apps" at top center, then on "HEMS Tool" and "Launch Application." By default, the HEMS tool displays the most recent analysis of flight category. The flight category product is discussed here because it is germane to the subject matter of this section, but there are many other types of information for display. One can display ceiling, visibility, and radar reflectivity as separate fields;

temperature, relative humidity, wind speed, and icing data at 500-ft increments above ground up to 5,000 ft AGL; and overlays of METARs, TAFs, PIREPs, AIRMETs, SIGMETs, and VORTAC locations. (VORTACs send out signals that help pilots navigate along the flight path.) All these display options are available under “Weather” and “Overlays” at the top of the HEMS display.

Figure 3.32 illustrates the national-scale flight category product. If you are viewing this document on line, you can zoom in to see details. The background shading indicates the analyzed flight condition: possible obscuration of terrain, LIFR, IFR, MVFR, and VFR. The color of the station circle corresponds to the observed flight condition. As with the CVA, the gridded flight condition is inferred from nearby METAR data with the aid of GOES satellite images, but may not match the actual flight condition between METARs. Open squares indicate clear sky. Partially filled circles indicate the amount of cloudiness; completely filled circles indicate an overcast. At 2245 GMT, 14 November 2011, conditions across the U.S. were mostly benign: VFR. Note the AWOS reports plotted in the northern Gulf of Mexico.

The time line across the bottom of Fig. 3.32 enables the user to view past, present, or near-term conditions. The vertical red line indicates the current time. The orange bar indicates product valid time. The user can loop through the product at 15-min increments or pause at any chosen time. The inset at bottom right indicates the area displayed. One can roam and zoom on the product.

In Fig. 3.33, the display default has been changed to ceiling rather than flight category, and the image has been zoomed several times to focus on the lower ceilings over Alabama. Note the expanded scale for ceilings in the color bar, with more emphasis given to the lower ceilings. This zoomed view shows the detail available from GIS: cities, counties, highways, rivers, reservoirs, and smaller airports, the latter highlighted in orange. The text box with the METAR report resulted from clicking on the METAR station circle just to the left of the box.

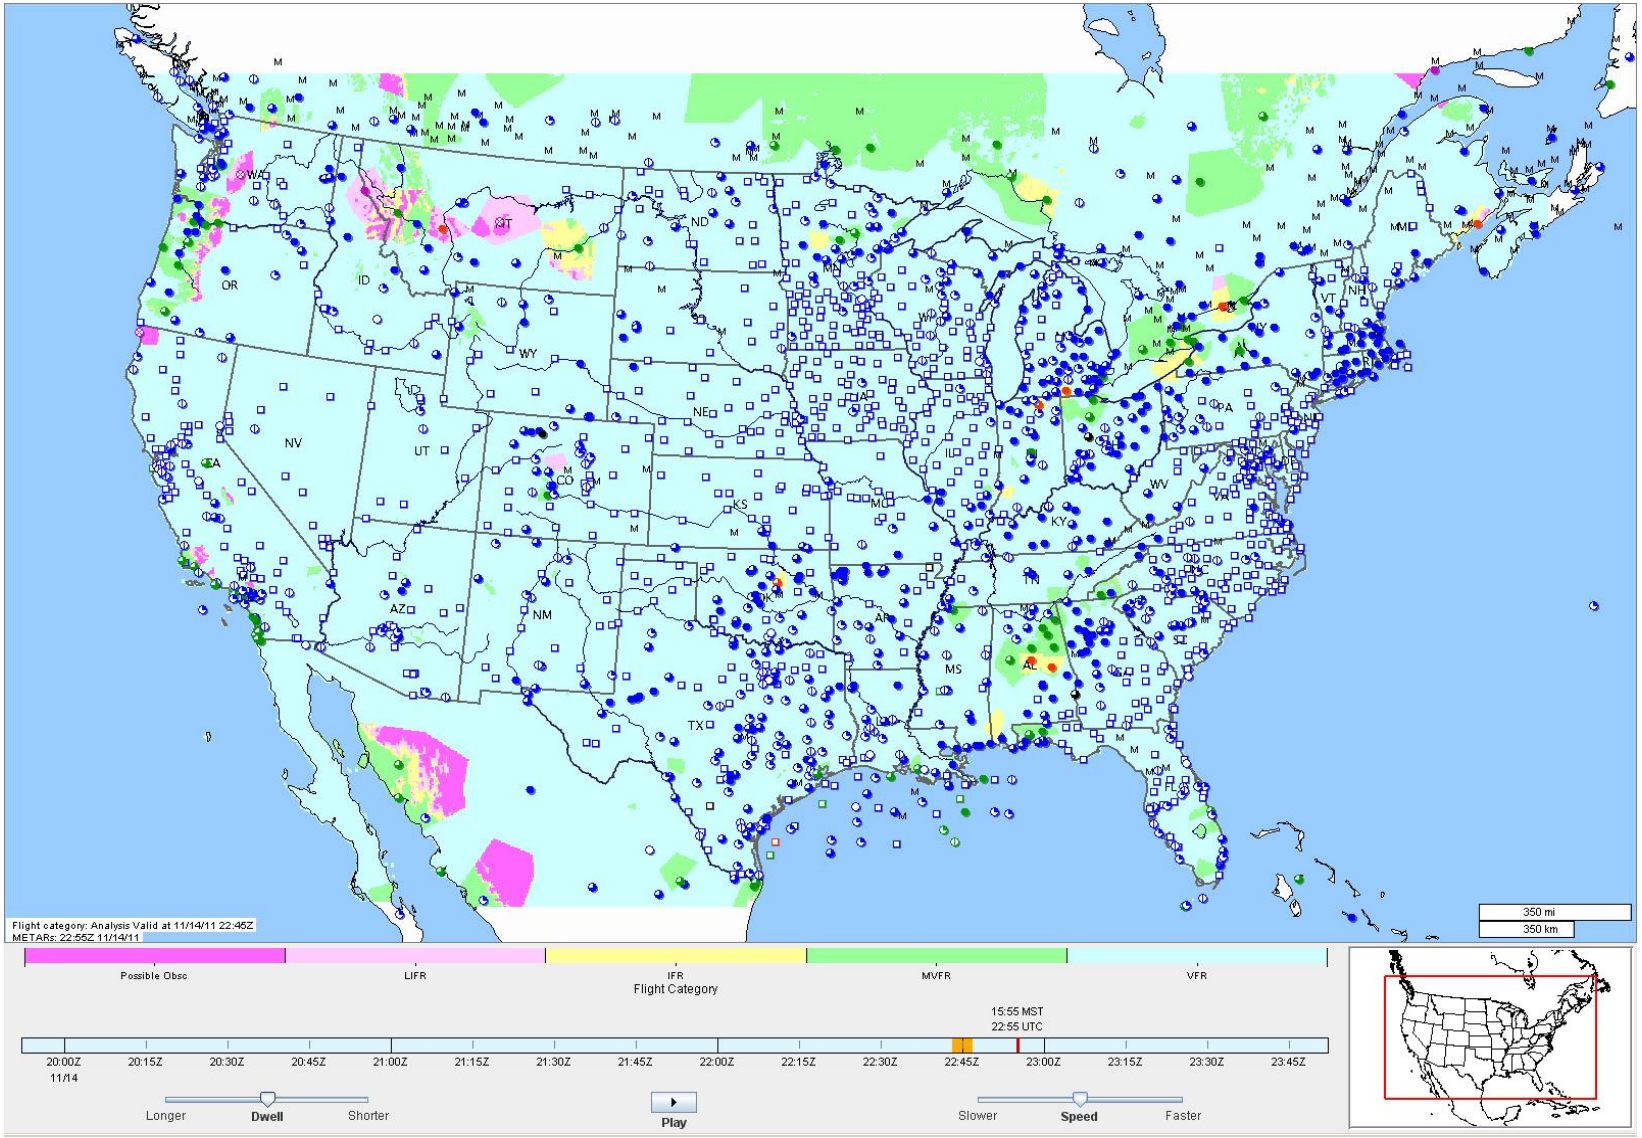

Fig. 3.32. The HEMS analysis for 2245 GMT 14 Nov 2011 showing flight category, color-coded as shown. Observed flight categories from METARs are overlaid. See text for additional details.

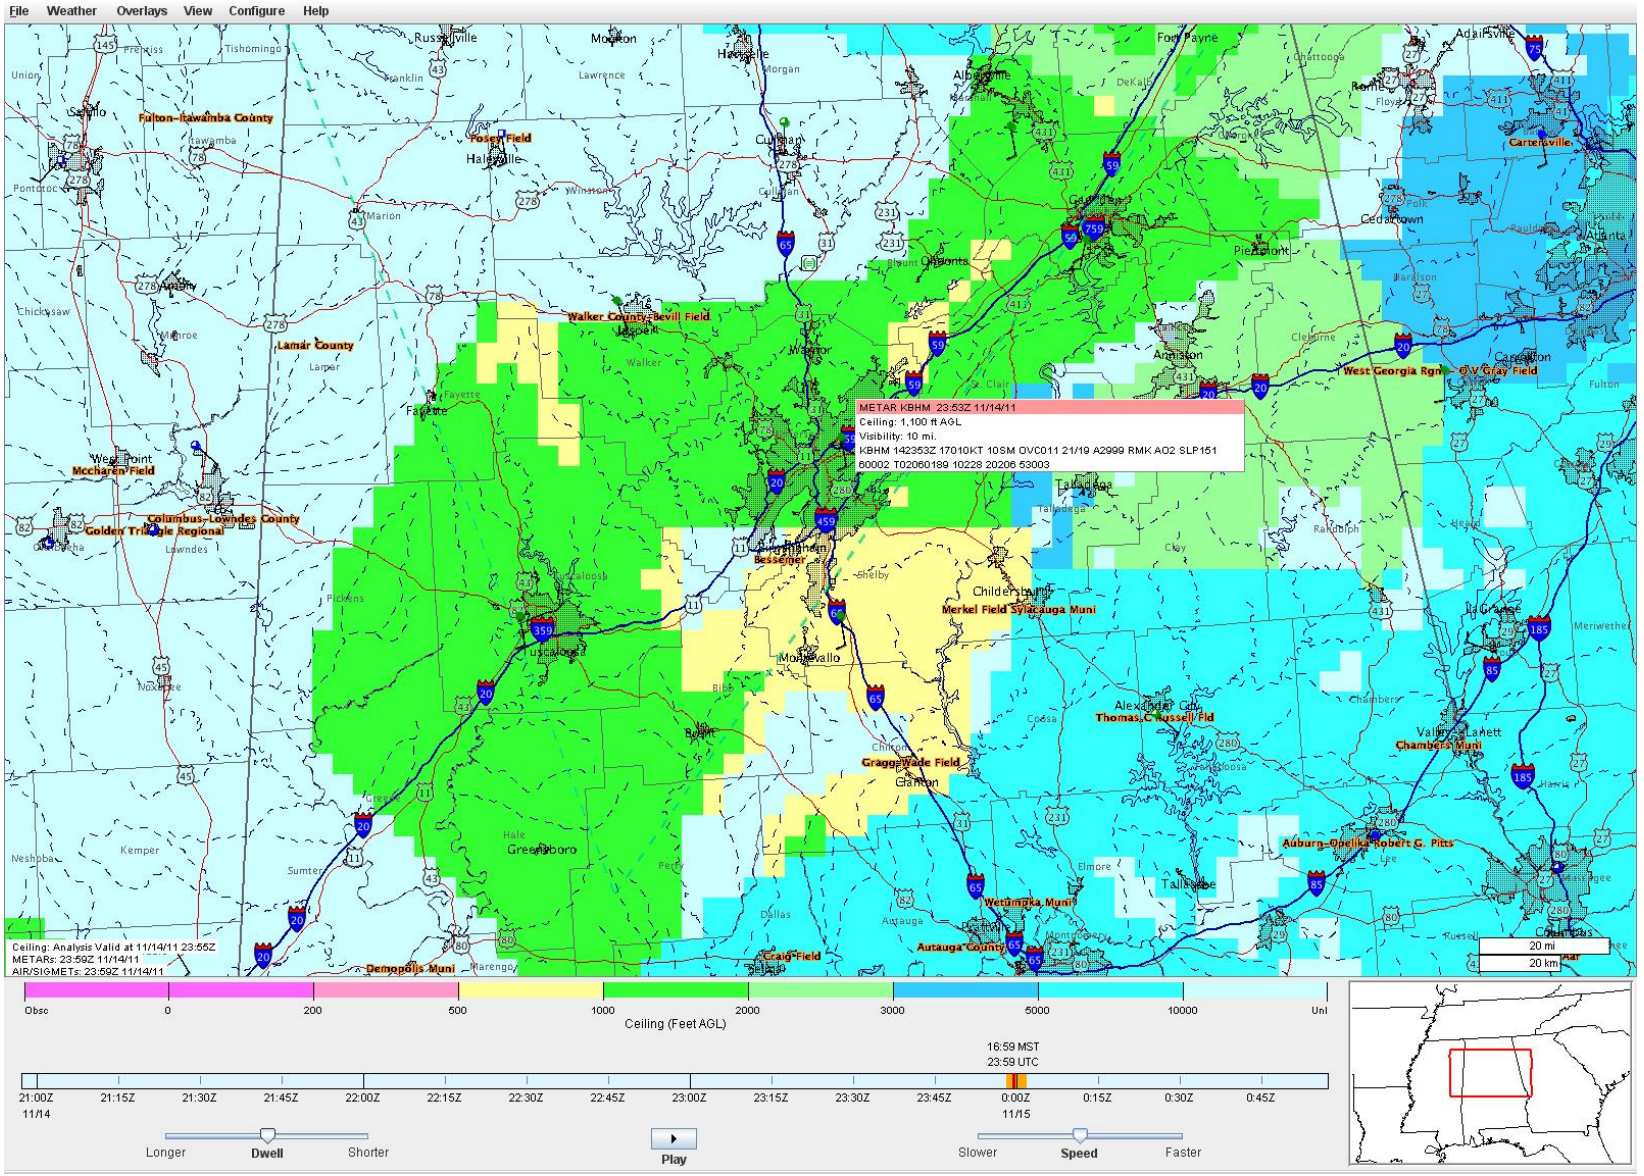

Fig. 3.33. Ceiling analysis on a 5-km grid generated by HEMS. Image is highly zoomed. The background colors indicate ceiling values corresponding to the color bar. Note background map detail from GIS available at this zoom level. See text for additional details.

3.1.10 GOES-R Algorithm for inferring probability of IFR conditions

The GOES-R Proving Ground project, first described in Section 2.1.6.1, engages many collaborators to develop algorithms that will exploit the capabilities of GOES-R, to be launched in 2015. Among these is the Advanced Baseline Imager (ABI), having 16 channels with spatial resolution of 2 km or better. A product developed in anticipation of GOES-R that estimates the probability of IFR conditions is described in this section. Data from the MODIS instrument aboard the NASA Terra and Aqua Satellites approximate the data to be available from GOES-R.

This section describes a Fog/Low Stratus (FLS) product documented by Calvert and Pavolonis (2011). The product is based upon anticipated GOES-R capabilities but applies to current GOES satellite data as well. It is the probability of ceiling less than 1,000 ft (IFR) without regard to surface visibility. Because the product relies on infrared sensing, the probability can be estimated only if no cloud layers lie above the layer responsible for the low ceiling or fog. Probabilities are not calculated when the sun is near the horizon (85^o < solar zenith angle < 90^o).

The primary sensor data used by the fog/low-cloud algorithm are:

- Calibrated reflectances for ABI channels 2 (0.65 μm) and 7 (3.9 μm). Refer to Table 2.7 for channel descriptions.
- Calibrated radiances for ABI channel 7 (3.9 μm) and 14 (11 μm)
- Calibrated brightness temperature for ABI channel 14 (11 μm)
- Solar zenith angle (depends upon the location of the volume sampled by the ABI, time of day, and day of the year)

The FLS algorithm also uses three derived products: 1) an ABI cloud mask, which helps eliminate cloud-free areas from consideration; 2) ABI cloud phase, which indicates whether hydrometeors near cloud top are liquid, solid, or mixed (the cloud algorithm does not run for ice fog); and 3) the liquid water path, a measure of how much water is contained within a vertical column of the cloud. These parameters were discussed in more detail in Section 2.2.11.

Finally, the FLS algorithm needs monthly mean surface emissivity for ABI channels 7 and 14, and an estimate of current temperature and relative humidity profiles from a numerical prediction model. The surface emissivity at 3.9 μm is interpolated from monthly mean values from MODIS measurements at 3.6 and 4.3 μm). The surface emissivity at 11 μm is estimated from monthly mean MODIS values at 10.8 and 12.1 μm.

The detection of fog and low stratus on satellite images relies on the following physical properties of fog:

- Composed mainly of liquid water (detection of ice fog is not attempted)
- Low cloud base or sky obscured
- Top of fog layer is highly uniform in both brightness temperature and reflectance because vertical velocities are typically very small.
- The difference in temperature between the top of the fog layer and the ground is relatively small.
- Because the boundary layer has a high concentration of cloud condensation nuclei, fog droplets are usually small and have a high number density.
- The water content of fog (liquid water path) is relatively low.

A technical description of the FLS algorithm is in Appendix B, Section B3.1.10.
The algorithm is routinely applied to observations from GOES-West and GOES-East. Sample products are available at http://cimss/ssec/wisc/edu/geocat/. Click on a geographic region, and see a selection of nearly thirty products. The one illustrated in Fig. 3.34 is for GOES-West CONUS: GOES-R IFR Fog Probability (Day/Night). In this example, black shading indicates where the daytime GOES-R cloud mask determined clear sky; gray shading indicates where the GOES-R cloud phase algorithm determined ice clouds. The greatest likelihood of IFR conditions is along the northern California coast and offshore, where marine stratus is present.

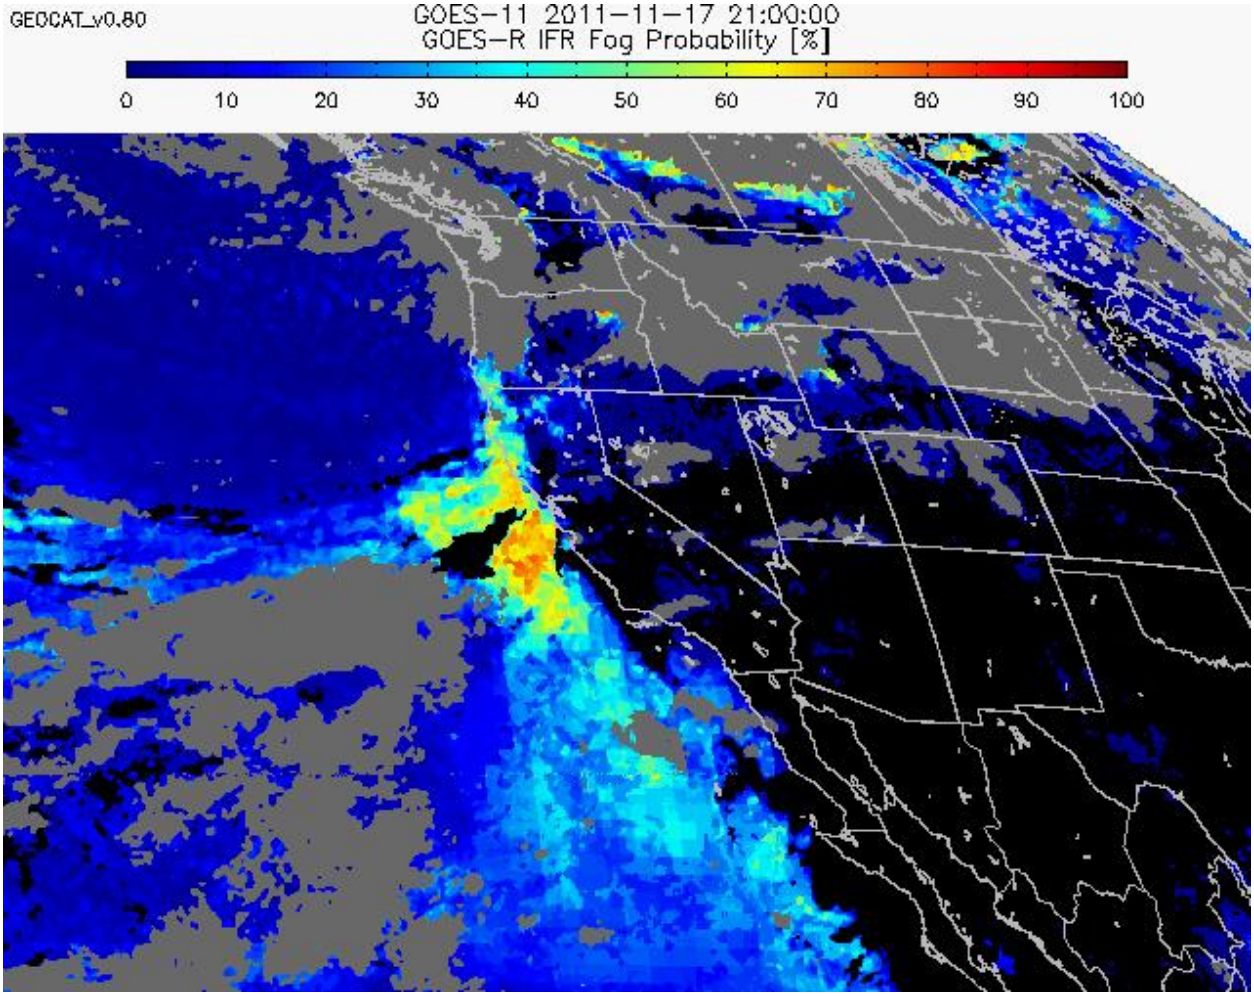

Fig. 3.34. The probability of IFR conditions estimated with the GOES-R algorithm for fog/ low stratus, as applied to GOES-11 data. Valid time is 17 Nov 2011, 2100 GMT. The color bar at the top indicates probability. Black shading indicates clear skies; gray shading indicates ice clouds.

Limited verification of the GOES-R fog/ low-cloud product is summarized in Table 3.10, courtesy of Mike Pavalonis (Center for Satellite Applications and Research, NESDIS/University of Wisconsin). The threshold probability for this table is 40% during the day and 26% at night, that is, a categorical forecast of IFR conditions occurs if the probability equals or exceeds the threshold; a non-IFR forecast occurs otherwise. The "skill" at right is just the difference between the POD and FAR. The target skill for this product is 0.70. An "event" is a forecast of IFR conditions (yes) or non-IFR conditions (no) at a pixel. Surface observations of cloud ceiling determine hit and missed forecasts. Care is taken to verify forecasts only when higher clouds do not obscure clouds near the ground. Daytime PODs are higher, and FARs are lower, than the corresponding nighttime values.

Table 3.10
Validation of the GOES-R Fog/Low-Cloud product, courtesy of Mike Pavolonis

Date		# events	POD	FAR	Skill
13 Dec 2008	Day	7,139	0.803	0.070	0.733
	Night	14,926	0.764	0.125	0.639
	Both	22,065	0.779	0.110	0.669
16 Jan 2010	Day	5,423	0.823	0.064	0.759
	Night	19,952	0.637	0.178	0.458
	Both	25,375	0.664	0.151	0.513

3.2 Aircraft De-icing Prior to Takeoff

Aircraft are subject to two separate hazards while waiting to push back from the gate, during taxi, or while awaiting clearance for takeoff. Snow or ice accumulating on the surfaces of the aircraft, especially the wings, can greatly diminish lift and affect aircraft control. Also problematic is the accretion of ice on the engine fan blades during freezing drizzle. This ice can be shed during engine run-up, thus accelerating shards of ice to high velocity inside the engine and causing damage to the fan blades themselves (Rasmussen et al. 2006).

If snow, freezing rain, freezing drizzle, freezing fog, or ice pellets occur while the aircraft is on the ground, the pilot proceeds to a de-icing area where the fuselage is sprayed with a de/anti-icing agent before takeoff. The most common agent, Type 1, is a mixture of propylene glycol or ethylene glycol and water. It is sprayed at high temperature and pressure to melt and remove snow, ice, and frost (which may accumulate if the aircraft is parked at the gate overnight). If snow or freezing precipitation is still falling after de-icing, the aircraft must take off within 5-15 minutes (depending upon the precipitation rate) because this low-viscosity fluid quickly runs off the fuselage. If an aircraft cannot become airborne in that time, Type 2 (in Europe) or Type 4 anti-icing fluid is applied to the aircraft. Type 2 and 4 fluids contain a thickening agent that prevents the fluid from running off until high air speeds are attained (at takeoff), thus affording longer protection.

The major hazard associated with de/anti-icing operations is that additional ice or snow will accumulate on the aircraft between the time of de-icing and takeoff. Furthermore, during poor weather conditions, departure rates are often reduced, causing longer wait times near the head of the runway. Several commercial aircraft accidents, some with fatalities, have occurred when aircraft took off with ice on the wings (Rasmussen et al. 2000).

The amount of time a de-iced aircraft can be considered safe for takeoff depends upon the type of de-icing fluid used, the rate of precipitation, and the temperature. In years past, the rate of precipitation was usually estimated on the basis of visibility. Researchers at NCAR have

demonstrated that visibility is a poor estimator of precipitation rate and have advocated ways of measuring the liquid equivalent rate of precipitation directly. This is important because the amount of water in the precipitation dilutes the de/anti-icing fluid and reduces the protection it affords.

Rasmussen et al. (1999) examined the correspondence between visibility and liquid-equivalent snowfall rate during the winters of 1994–‘5 and 1995–‘6. They used two state-of-the-art precipitation gauges, an ETI NOAH II precipitation gauge inside a Wyoming wind shield, and a GEONOR T-200b gauge inside an Alter wind shield. An HSS VPF-730 sensor was also used for automated visibility measurements. Figure 3.35 reproduces two of their figures. Note the wide scatter in both data samples, bearing in mind the log-log scaling.

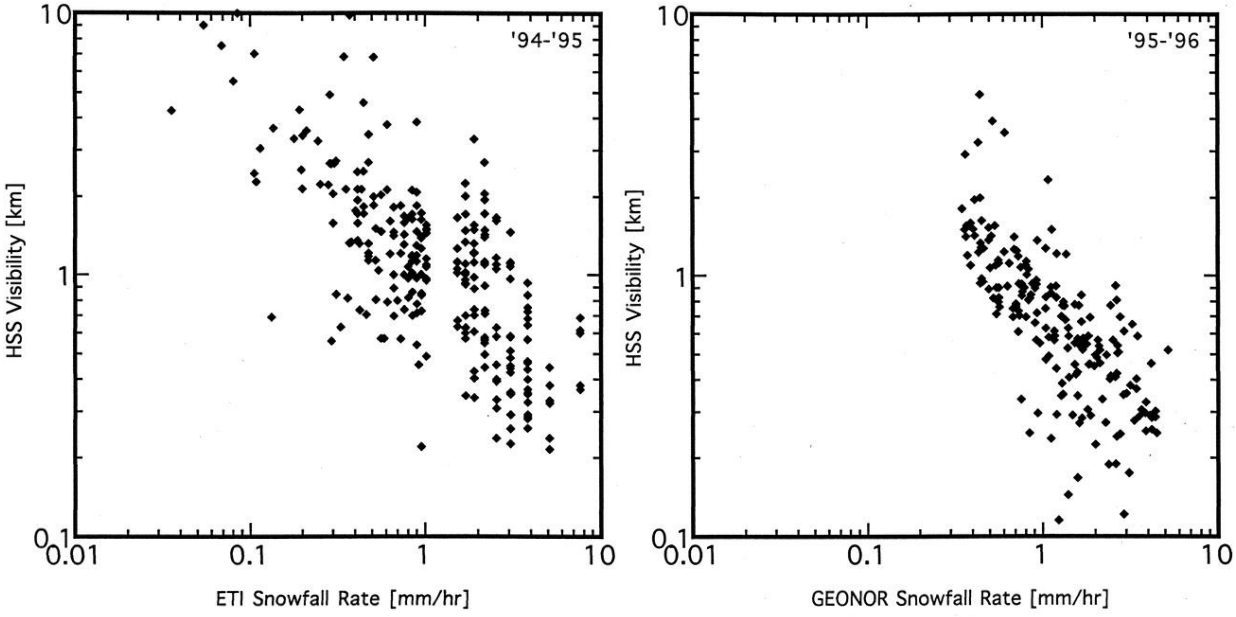

Fig. 3.35. Observed snowfall rates as a function of visibility, as measured by an ETI precipitation gauge during the winter of 1994–’5 and a GEONOR precipitation gauge during the winter of 1995–’6. Given the log-log scale, the scatter of points is very wide. From Rasmussen et al. (1999). Reprinted with permission of the American Meteorological Society.

Theoretical calculations of visibility during snowfall indicate that it can vary by up to two orders of magnitude for the same liquid-equivalent rate of snowfall, depending upon crystal size and habit (plates, needles, stellars, columns, etc.), and whether the snowfall consists of single, unrimed crystals (low fall speeds); crystal aggregates (higher fall speeds); or partially melted or heavily rimed crystals (high fall speeds).

Clearly, visibility is not helpful in estimating the liquid-equivalent rate of snowfall. Neither is it useful in estimating the intensity of drizzle because fog often influences visibility when drizzle is present.

Unheated precipitation gauges are not especially accurate at measuring the liquid-equivalent snowfall rate. When the temperature is near freezing, snow can accumulate on the lip and interior walls of the gauge, eventually closing off the orifice and decreasing the catch. Airflow around the gauge can also reduce the catch, particularly if the wind is strong. A strong wind can carry snow crystals over the top of the gauge and prevent them from falling in. The Automated Surface Observing System (ASOS), the basis for many METAR reports, originally employed a heated tipping bucket (HTB) gauge to measure precipitation. This device systematically underestimated the liquid equivalent of snow, partly because of wind effects but especially because of evaporative losses between the time the snow melts upon contact with the heated collector and the time the melt water reaches the tipping bucket. The problem is most acute during light snowfalls and also occurs with freezing drizzle. In recent years, the FAA and NWS have addressed this problem by replacing the HTB gauge with an OTT Pluvio gauge at all ASOS sites that experience winter weather precipitation. The OTT Pluvio gauge is similar in design and concept to the GEONOR gauge. It is a weighing gauge with a heated orifice that can measure liquid equivalent precipitation amounts down to 0.01 inches.

The hotplate precipitation gauge overcomes many of the difficulties of measuring liquid-equivalent precipitation rates when snow is falling. Described by Rasmussen et al. (2011), and developed by NCAR and the Desert Research Institute, this gauge has no moving parts, presents only a thin, disk-shaped obstacle to the wind, and can measure precipitation rate each minute, an especially valuable feature for pilots who need to know how fast ice is accumulating on airfoils between de-icing and takeoff. Figure 3.36 is a photo of a commercial version of the hotplate, manufactured by Yankee Environmental Systems, Inc.

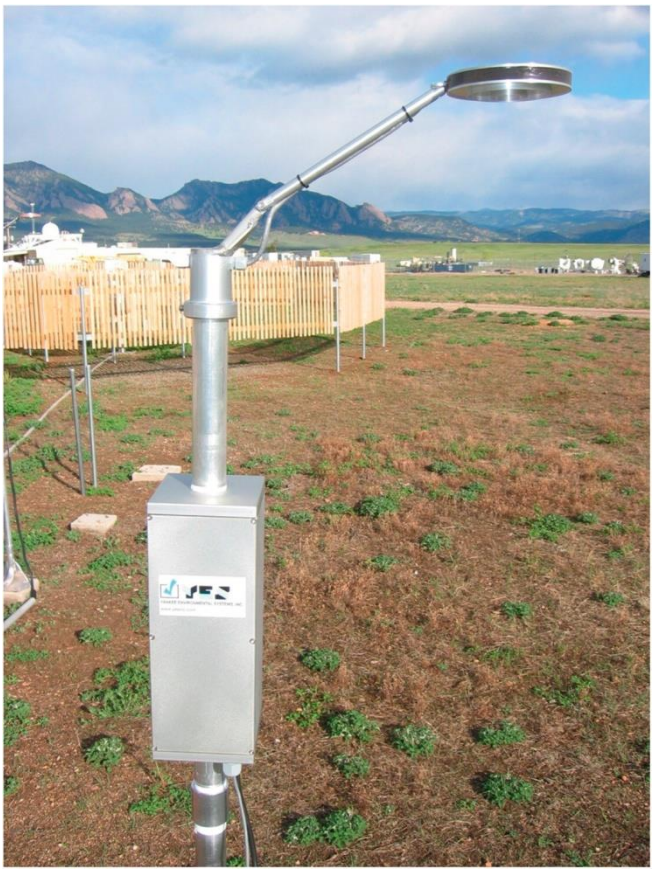

Fig. 3.36. A commercial version of the hotplate precipitation gauge, manufactured by Yankee Environmental Systems. From Rasmussen et al. (2011). Reprinted with permission of the American Meteorological Society.

The hotplate gauge consists of two metal plates, one on top of the other and thermally insulated from each other. The raw output of the sensor is the difference in power required to maintain the top and bottom plates at a constant temperature of about 75°C. The power supplied to the bottom plate is determined only by the air temperature and the wind speed. The top plate is affected by these and, in addition, captures falling precipitation. The power required to maintain it at constant temperature is the same as that required by the bottom plate plus the power needed to evaporate rainwater or melt and evaporate snow. A temperature of 75°C is enough to melt even large snowflakes within a few seconds. The hotplate is calibrated so that the difference in power supplied to the top and bottom plates is directly related to the precipitation rate.

To properly calibrate the hotplate, it is necessary to know whether the precipitation is liquid or solid. The decision on precipitation type is currently based solely on temperature. If the temperature is less than 4°C, the precipitation is assumed to be snow; if it is above 4°C, the precipitation is assumed to be rain. If the temperature lies between 0° and 4°C, a linear

interpolation between these calibration points determines what fraction of the precipitation is liquid. A better parameter for making this decision would be wet-bulb temperature (0°C). A humidity measurement, together with temperature, is needed to calculate wet-bulb temperature. This idea is being tested on the latest version of the hotplate.

All ASOS sites include a Light-Emitting Diode Weather Identifier (LEDWI), capable of distinguishing rain from snow, as well as an icing sensor that detects the accretion of freezing precipitation. The LEDWI examines how precipitation falling through a pulsed infrared beam, about 5 cm in diameter, alters the frequency of the transmitted beam over a distance of 0.8 m. The LEDWI has three channels: a "high" channel sensitive to faster falling precipitation, a "low" channel sensitive to slower falling precipitation, and a "particle" channel that records the relative size of the largest particle observed during the 1-min sampling interval. Using data from all three channels, an algorithm determines both precipitation type and intensity.

The icing sensor is a thick, rapidly vibrating rod. When the rod is clear of precipitation, it vibrates at roughly 40,000 Hz. As ice builds up on the rod, the frequency decreases. The *rate* of decrease in frequency is related to the precipitation intensity. Unfortunately, the LEDWI was not designed to detect drizzle, but the current ASOS algorithm can still infer freezing drizzle (FZDZ) from multiple parameters as indicated in Table 3.11.

Table 3.11

The ASOS algorithm for inference of rain, freezing rain, snow, and freezing drizzle (RA, FZRA, SN, FZDZ, respectively). UP – unknown precipitation; OVC – overcast; VV – vertical visibility. From Ramsay (2002).

Ice Detector	**LEDWI Present Weather Type**	**Temperature**	**Sky Cover**	**Present Weather Reported**
Accretion ≥ 0.13 mm AND 15-min accretion rate > 0.2 mm h^{-1}	RA, UP	< 2.8°C	Any	FZRA
	SN	< 3.3°C	Any	SN
Accretion ≥ 0.13 mm AND 15-min accretion rate > 0.1 mm h^{-1}	No precipitation	≤ 0°C	OVC or VV	FZDZ
			Not OVC	None

Ramsay (2000) determined the following correspondence between the icing sensor's drop in frequency over 15 minutes and the precipitation rate as shown in Table 3.12. The more the frequency drops in 15 minutes, the greater the precipitation rate. Values in this table are now being tested in ASOS with a view toward reporting the intensity of FZDZ.

Table 3.12

Intensity of freezing drizzle based upon the drop in vibration frequency of the ASOS icing sensor in 15 min (Hz/15 min).

Intensity	Weather Type		Precipitation Rate	Frequency Drop
	Freezing fog	FZFG	.001–.004 in/h	1–5 Hz/15 min
Light	Light freezing drizzle	–FZDZ	.004–.010 in/h	6–16 Hz/15 min
Moderate	Moderate freezing drizzle	FZDZ	.010–.020 in/h	17–33 Hz/15 min
Heavy	Heavy freezing drizzle	+FZDZ	>.020 in/h	>33 Hz/15 min

The foregoing description of sensors for measuring liquid equivalent snowfall rates and ice accretion rates is appropriate because, without minute-by-minute estimates of these parameters, it would be difficult for aircraft de-icing operators to decide which de-icing fluid to use and for pilots to know how long they can afford to wait before taking off after being de-iced.

During the 1990s, scientists and engineers at the Research Applications Laboratory of NCAR developed the Weather Support to De-icing Decision Making (WSDDM) program, an automated system to provide up-to-the-minute guidance to airports, airlines, pilots, and tower controllers about ice accretion and liquid-equivalent snowfall rates during winter storms. WSDDM also includes a 60-minute nowcast based upon extrapolated radar echoes and continuing calibration of the relationship between radar reflectivity and measured liquid-equivalent snowfall rate.

NCAR developed WSDDM in collaboration with several airlines, several airports, and FAA air traffic managers. In the mid and late 1990s, prototype versions were tested at Denver International Airport, Chicago O'Hare Airport, and airports operated by the Port Authority of New York and New Jersey. The technology and software for WSDDM is government owned but currently licensed to Vaisala. Denver International Airport and the Minneapolis-St. Paul Airport have contracts with Vaisala for WSDDM support. Many other airports have expressed interest in purchasing WSDDM for their winter operations.

At the heart of the WSDDM algorithm is a real-time calibration of the relationship between radar reflectivity and liquid-equivalent snowfall rate. The calibration is described in Rasmussen et al. (2003). Many studies have tried to relate radar reflectivity factor Z [mm^6 m^{-3}] to liquid-equivalent snowfall rate S [mm h^{-1}] in terms of a power law:

$$Z = aS^b ,$$

where a and b are empirical constants. For pure rain events (T > 2°C), $b = 1.6$. For snow events (T < –3°), $b = 1.75$. For mixed precipitation, b decreases linearly with increasing T for –3 < T < 2°C. Though b is thus determined a priori, a is determined in real time.

The next step is to estimate the radar echo motion $\boldsymbol{V}_r$ directly over the precipitation gauge using TREC. TREC matches radar features on successive radar scans by shifting the two images with respect to each other until their cross correlation is maximized, as described earlier in Section 2.3.11.3. Experiments have determined that the best motion vector to use for calibration is a combination of $\boldsymbol{V}_r$ and the surface wind vector at the gauge $\boldsymbol{V}_s$:

$$\boldsymbol{V}_g = c_1\boldsymbol{V}_r + c_2\boldsymbol{V}_s \, .$$

c_1 and c_2 currently have the value 0.5. The fall speed of snow s_f is assumed to be 0.9 m s^{-1}, that of rain, 10.0 m s^{-1}. For mixed precipitation, the fall speed varies linearly between these two values. The fall time of snow t_f is the vertical distance from the altitude of the radar volume to ground divided by the fall speed:

$$t_f = \frac{D \tan\theta}{s_f} \, .$$

D is the distance from the radar to the gauge location, and θ is the elevation of the radar beam.

The next step is to use the vector $-\boldsymbol{V}_g(t_f)$ to look upwind and find the average reflectivity S of the echo at an earlier time ($t = -t_f$) that delivered the liquid equivalent precipitation measured currently at the gauge ($t = 0$). Denote the horizontal location of this echo by L_0. If the power law cited earlier is integrated over time and then solved for a, the result is

$$a = \left[\frac{\int Z^{(1/b)}}{\int S dt}\right]^b \, .$$

The integration over time is from $-t_f$ to 0. Multiple measures of radar reflectivity factor and liquid-equivalent precipitation will have been taken over this interval. Determination of a completes the real-time calibration.

The final step is to look upwind from location L_0 to find the echoes that will deliver precipitation at the gauge at future times t_i. First, determine the radar-derived echo motion $\boldsymbol{V}_0$ above L_0 using TREC vectors. Then, choose a particular projection time t_i, and look upstream with the vector $-\boldsymbol{V}_0(t_i)$. Apply the power law at time ($t = -t_f$) to the echo, if any, at this new location, to estimate the liquid-equivalent precipitation rate expected at the gauge location at time $t = t_f + t_i$. The accumulated precipitation may be calculated from the liquid-equivalent rates found at multiple projection times t_i.

Rasmussen et al. (2003) performed verification on 60-min forecasts of liquid equivalent precipitation made during four storms occurring at New York Airports during the winter of 1997–'8. A forecast was issued following every new radar scan, at 6–10-min intervals, depending upon the volume coverage pattern currently in use. Each predicted liquid-equivalent precipitation rate was assigned a rate of light (< 1.0 mm h^{-1}), moderate (between 1.0 and 2.5 mm h^{-1}) or heavy (> 2.5 mm h^{-1}). A hit (YY) occurs when the forecast rate verifies, a

miss (NY) is an underforecast of the actual rate, and a false alarm (YN) is an overforecast of the actual rate. The WSDDM forecasts were compared with persistence forecasts, defined as the continuation of the current observed liquid-equivalent precipitation rate for the next 60 min. Figure 3.37 gives standard verification scores—POD, FAR, and CSI—for four winter storms in the New York area. Note that WSDDM forecasts have better PODs, and much better FARs than the persistence forecasts. Also, recall that persistence forecasts can be difficult to beat in the first 60 min.

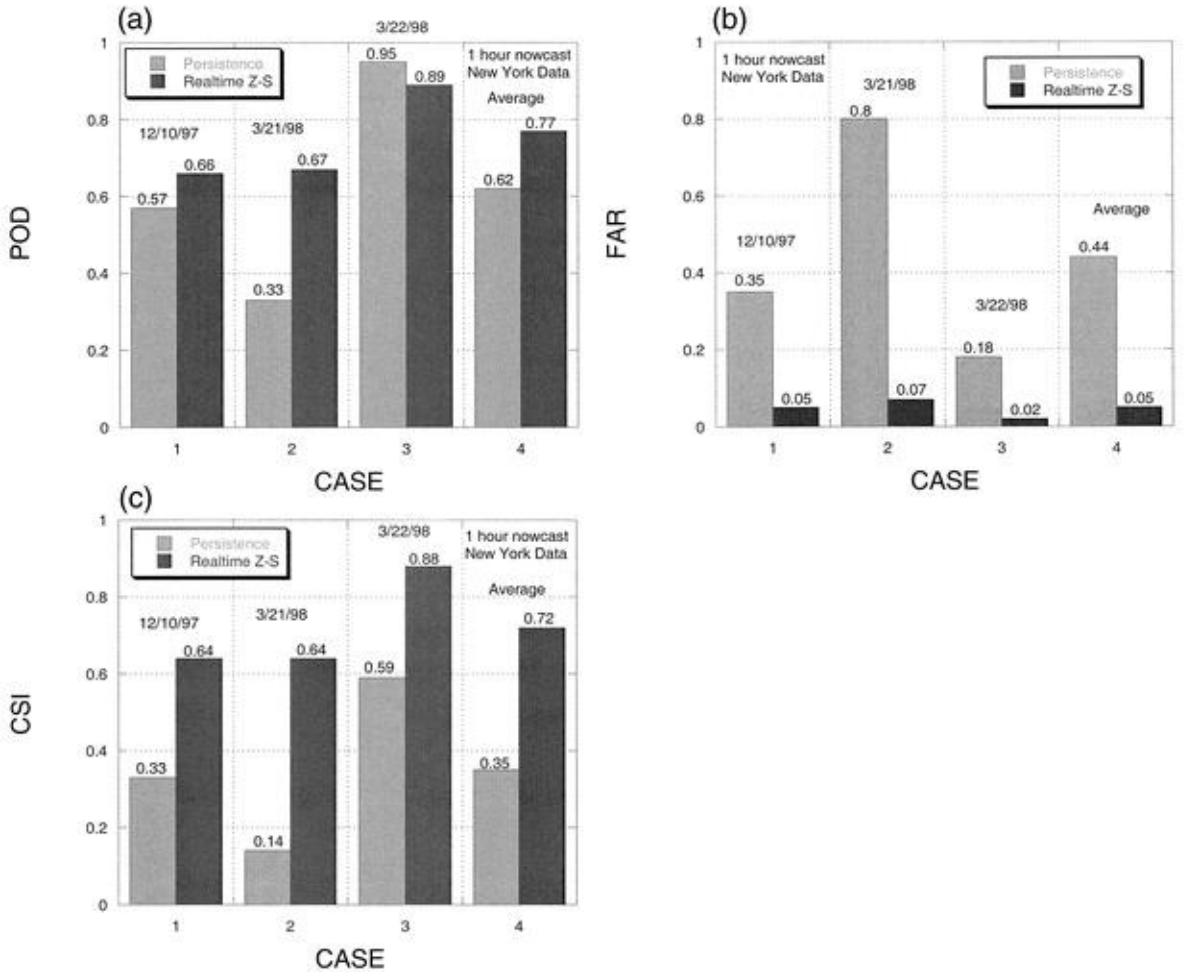

Fig. 3.37. Probability of Detection (POD), False–Alarm Ratio (FAR), and Critical Success Index (CSI) for 60-min nowcasts of the intensity of liquid-equivalent precipitation: light, moderate, or heavy, as defined in the text. Four cases are shown in each panel, with dates marked at the top. Gray bars identify persistence forecasts; black bars identify WSDDM forecasts. From Rasmussen et al. (2003), used with permission of the American Meteorological Society.

Ruzanski and Chandrasekar (2011) have proposed an alternative to TREC vectors for the determination of echo motion. Called Dynamic Adaptive Radar Tracking of Storms (DARTS), this

alternative leads to a 14% average improvement in the accuracy of 60-min nowcasts of liquid-equivalent precipitation and reduces the runtime of WSDDM software markedly.

The Society of Automotive Engineering Ground De-icing Committee publishes Holdover Tables that indicate how long a de-icing fluid can prevent icing formation on an aircraft for a given precipitation type, rate, and temperature. An add-on module for WSDDM can automatically provide estimated holdover times using this information. This module, called "Checktime," incorporates the current liquid equivalent precipitation rate, air temperature, and Holdover Table values which result in a "check time" or estimated time that a de/anti-icing fluid will no longer provide protection. If an aircraft receives clearance to take off before the check time, the de-icing fluid will still be effective. If the aircraft is still waiting for takeoff past the check time, pilots must visually inspect the fluid and, if necessary, return for another de-icing.

3.3 Lightning and Airport Operations

Cloud-to-ground (CG) lightning in the vicinity of air terminals is a serious hazard to outdoor ramp operations such as positioning jet bridges (jetways), refueling the aircraft, replacing food service racks, handling baggage, operating tugs, or simply standing on the tarmac directing aircraft to their assigned gates. Decisions to suspend ramp operations are made by airports or airlines rather than the FAA, and procedures and standards vary greatly. A nominal guideline is the "30/30 rule," which says that ramp operations are suspended following the first occurrence of a CG strike within a six-mile radius. About 30 s pass between perception of a flash six miles away and the sound of thunder. Operations do not resume until 30 min after the most recent CG strike within a six-mile radius. The first CG strike is always the most dangerous, especially if it comes from a thunderstorm developing overhead. An electrically active thunderstorm usually makes its presence known well before its arrival overhead.

The decision to suspend ramp operations is not taken lightly. Safety of personnel is paramount, but unnecessary suspensions are costly. Delays at Chicago's O'Hare Airport (ORD) and Orlando International Airport (MCO) both total over 3,000 min yr^{-1} spread over more than 50 ramp closure events at each airport. If the ramp closure interval could be reduced by an average of ten minutes per event, the annual savings are estimated to be over $6M at ORD and $2.8M at MCO.

The best reference for lightning detection and safety procedures at airports is Heitkemper et al. (2008), one in a series of reports by the Airport Cooperative Research Program (ACRP) under the Transportation Research Board of the National Academies. The research was sponsored by the FAA and carried out jointly by MDA Federal, Inc., and NCAR. Information in that report is summarized here.

There are several forms of lightning detection. The most elementary is seeing the lightning flash and hearing the thunder. This method might work on golf courses or ball fields but has

little utility at airports. Lightning emits electromagnetic energy pulses, primarily at radio frequencies. Low-cost detectors that rely on the amplitude of such emissions can deliver a crude estimate of the range to a CG stroke but give no information on the direction. These detectors, too, are unsatisfactory for airport applications.

Electric field mills measure the strength of the electrical field. The electrical field is measured in the vertical direction. In fair weather, the field strength is very roughly 100 volts per meter and points downward (positive charge in the atmosphere and negative charge on the earth's surface. Beneath a thunderstorm the field is usually reversed in sign, and its strength can increase by two orders of magnitude. When lightning causes an electrical discharge, the field strength decreases immediately, though often only temporarily. Though field mills are useful for alerting outdoor workers of the buildup of the electrical field over a period of ten minutes or so, they are less helpful in providing information about the timing or location of strikes.

The most accurate information about CG strikes comes from networks of electromagnetic sensors. Two different technologies are employed. One uses magnetic direction finders to locate strikes. A magnetic direction finder is a crossed-loop antenna, consisting of two vertical loop antennas mounted perpendicular to each other. A microprocessor analyzes signals and accepts only those from return strokes (unique to CG lightning), keying on the peak in the waveform radiated by the return stroke, approximately 100 m above ground. A flat-plate electrical field antenna is also part of the system; it determines the polarity of the electrical field associated with the CG stroke and thus whether it transferred positive or negative charge to the ground. Data from two or more direction-finding antennas are needed to compute by triangulation the location of ground strikes.

The other technology for CG lightning detection is based on time of arrival (TOA), that is, the precise time at which the electromagnetic signal from a CG strike arrives at multiple stations. The clocks at each station are accurately synchronized by reference to Global Positioning System satellites. Differences in time of arrival of the signal at multiple stations are the basis for calculating the location of the strike using sophisticated processing and spherical trigonometry. Because the signal travels at near the speed of light (3×10^8 m s^{-1}), time differences of small fractions of a microsecond must be measurable.

Two independent networks provide lightning detection services in the U.S. Vaisala owns and operates the National Lightning Detection Network (NLDN), originally based on magnetic direction-finding technology but recently adding TOA technology as well. More than 114 sites in the CONUS are part of NLDN. Vaisala is under contract with the National Weather Service to provide lightning detection data to government agencies. NLDN data are on Automation of Field Operations and Services (AFOS) workstations at NWS offices around the country.

The other lightning detection network is the United States Precision Lightning Network (USPLN), operated by WSI (Weather Services International), one of four companies owned by

The Weather Channel, in partnership with TOA Systems. This network is based upon TOA technology and serves a number of major airports around the country and several airlines.

Both networks advertise better than 95% detection accuracy for CG flashes and position accuracy of about 250 m in the CONUS.

When lightning location data are plotted on a map of radar echo locations (both available every few minutes), a loop of these images is the basis for a nowcast. Both Vaisala and WSI provide variants of this product to their customers.

A Geostationary Lightning Mapper (GLM) is being developed for GOES-R, scheduled for launch in 2015. An optical filter, designed to detect short-burst lightning emissions in the near infrared (777.4 nm), should be able to detect 70-90% of all flashes, day and night, with 8-km resolution at nadir and 14-km resolution at the edge of the scan. Data from GLM should be a valuable adjunct to ground-based lightning networks because in-cloud discharges almost invariably precede CG flashes by at least several minutes.

Several *lightning mapping arrays* are now operating in the U.S. A dense network of ground-based sensors is able to detect not only the flashes but also the actual three-dimensional paths of lightning channels, whether in cloud or cloud-to-ground. These arrays are used primarily for research, but SPoRT, the Short-term Prediction Research and Transition Center in Huntsville, Alabama, a NASA project, is providing data from lightning mapping arrays on an experimental basis to a number of NWS forecast offices.

3.4 Low-Level Wind Shear Detection and Alerting

A sudden wind shift at an airport prompts controllers to change the traffic patterns for landing and departing aircraft because, whenever possible, pilots avoid taking off or landing in tailwinds or substantial crosswinds. Operationally significant wind shifts are caused by frontal passages (usually easy to predict well in advance), moving sea breeze boundaries, gust fronts marking the arrival of rain-cooled air from distant storms, and microbursts. Of these, the most dangerous are microbursts, a downrush of air from showers or thunderstorms over the airport or off the ends of runways.

Wet microbursts are usually easy to see, at least from the outside: a nearly opaque shaft of heavy precipitation, often with curls around its lower edges where strong, outflowing winds at the ground push raindrops laterally. A *dry microburst* may leave few visual cues. It usually originates from a high-based shower cloud, sometimes not producing lightning, often producing only light showers of small drops from an innocuous-looking cloud. As raindrops fall from cloud base into dry, boundary-layer air below, evaporational cooling gives the downdraft negative buoyancy, and it accelerates toward the ground. Upon reaching the ground, the downdraft

rushes out in all directions, sometimes creating a ring of dust around the periphery of the outflow. Downbursts are hazardous to aircraft on final approach, as illustrated in Fig. 3.38.

If the downdraft lies off the end of the runway, the aircraft first experiences an increase in air speed. The pilot may be tempted to point the noise down to reenter the glide path, but he will soon experience first the downdraft, and then a rapid decrease in air speed as he passes to its far side of the microburst. If he has not been warned of this hazard and is not trained in the appropriate procedures for this encounter, he is in danger of touching down short of the runway. A similar hazard exists for pilots accelerating down the runway toward takeoff. A microburst near the end of the runway will cause a sudden increase in air speed, perhaps prompting the pilot to lift off, but the following sudden decrease in air speed may bring the plane perilously close to objects off the end of the runway.

Controllers alert pilots whenever sudden increases or decreases in air speed exceeding 30 kt are expected.

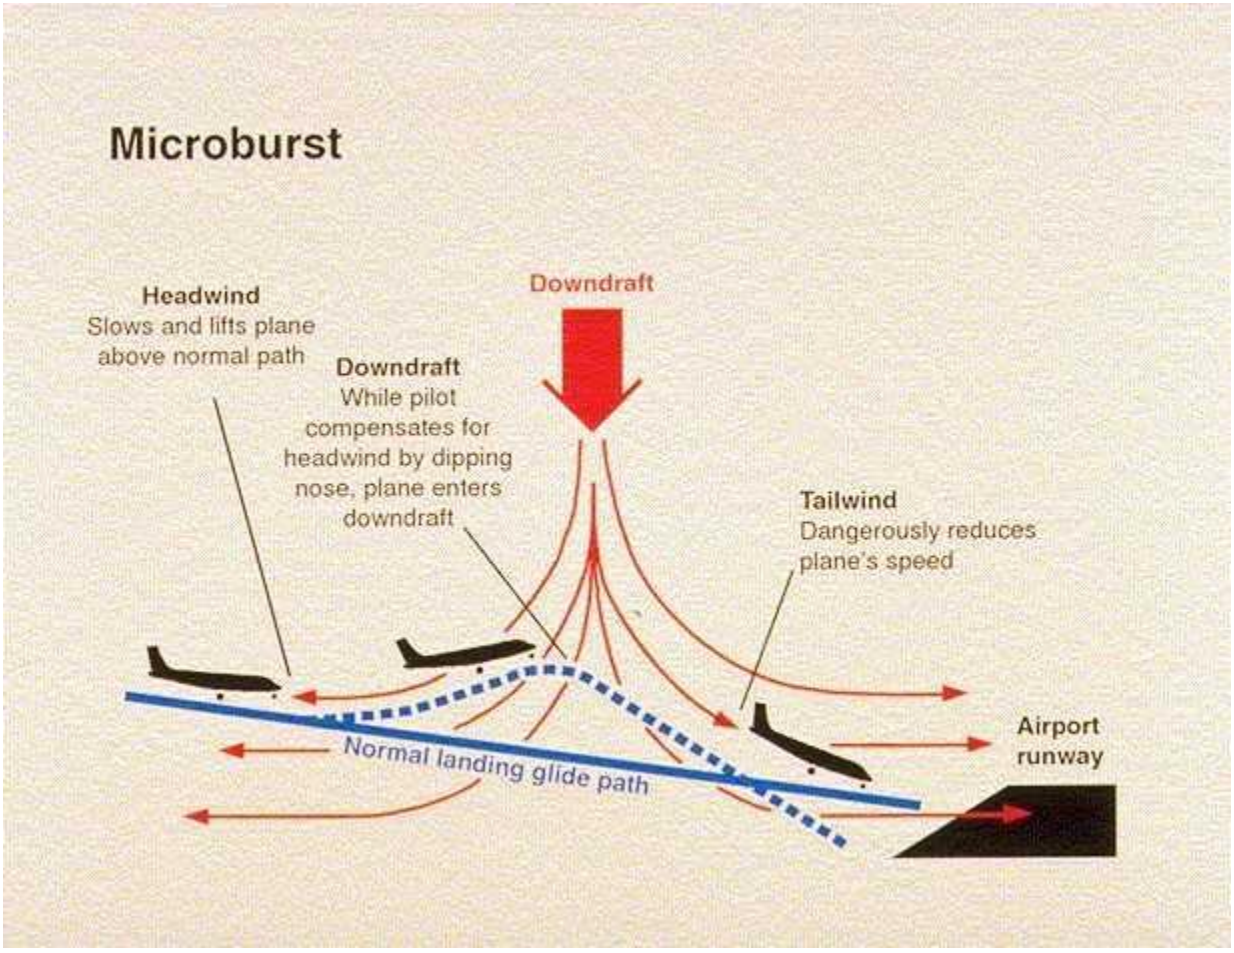

Fig. 3.38. Illustration of a microburst, courtesy of NCAR, and the hazard it poses to aircraft on final approach.

Airports use one of four methods to warn pilots of low-level wind shear: Terminal Doppler Weather Radar (TDWR), the Low-Level Wind shear Alert System (LLWAS), the Weather Systems Processor (WSP), or the Integrated Terminal Weather System (ITWS). Very large, high-traffic airports may use a combination of TDWR and LLWAS. Keohan (2007) reviews the capabilities of all four systems.

3.4.1 Terminal Doppler Weather Radar (TDWR)

The Terminal Doppler Weather Radar (TDWR) was developed in the early 1990s by MIT Lincoln Lab and is manufactured by Raytheon. The principal application of this radar, then and now, is the detection hazardous low-level wind shear. As of 2011, 47 TDWRS are deployed in the CONUS near the busiest airports subject to wind shear (Fig. 3.39). Note the absence of these radars along the West Coast, where low-level wind shear, particularly from convection, is absent.

Fig. 3.39. Locations of the 47 Terminal Doppler Weather Radars in the CONUS as of 2011. A 48th radar in Puerto Rico is not shown. From Wikipedia.

Table 3.13 compares the properties of the TWDR with those of the WSR-88D. The 5-cm wavelength of the TDWR makes it more sensitive to precipitation than the WSR-88D. Concentrating on low-level scans in order to detect wind shear near runways, the TDWR can complete a scan at 0.2° elevation in just one minute. The TDWR beam width is narrower than

Table 3.13

A comparison of properties of two radars, the TWDR and the WSR-88D.

	TDWR	**WSR-88D**
Wavelength	5 cm	10 cm
Volume Scan Time	1 min (0.2° elevation angle)	6 min (full volume scan)
Beam Width	0.55 degrees	1.25 degrees
Max Unambiguous Velocity	30 kt	62 kt
Range resolution	150 m	250 m
Max Doppler Range	90 km	230 km

that of the WSR-88D, and the range resolution is finer, both properties apt for the detection of small–scale microbursts on or near the airport property. Velocity folding is more of a problem for the TWDR than for WSR-88D, but software corrects for this most of the time. With emphasis on wind shear detection near airports, the 90-km maximum Doppler range for the TWDR is not a problem.

Because microbursts generally produce strong gusts of wind that fan out in all directions from a point directly beneath the downrushing air, radial velocity displays have little trouble detecting strong wind shear as long as the reflectivity is adequate.

The probability of detection of hazardous wind shear is better than 90%, when precipitation backscatters adequate returned power to the radar. In parts of the arid Intermountain West and the desert Southwest, however, the detection rate is lower because some downbursts in these areas are dry close to the ground. A few sprinkles of very small drops may be all that reaches the ground.

NWS offices near airports that have TWDR radars now receive the TDWR displays as a supplement to the WSR-88D displays.

3.4.2 Low-Level Wind shear Alert System (LLWAS)

An early version of today's Low-Level Wind shear Alert System was developed in the 1970s in response to aircraft accidents resulting from sudden and unexpected wind shifts. The early system compared center-field wind measurements with five other measurements around the airport. The network of sensors was sub-optimal, the false–alarm rate was high, and microbursts were just beginning to be understood.

In 1983, the FAA asked NCAR to develop a version of LLWAS appropriate for microburst detection. NCAR did so, and tested the system from 1983 through 1988 by alerting pilots along each runway whenever significant changes in head- or tailwinds were expected. Today, roughly 50 U.S. airports have what are called LLWAS-Phase 3 systems, of which about ten of the highest-traffic airports have TWDRs as well. Most Phase 3 systems have between 12 and 16 wind sensors. Denver International Airport, a very busy airport with many microburst days each summer, has the most: 32. The sensors are spaced 2 km apart and cover out to 4 km from the end of each major runway.

The FAA owns and operates LLWAS systems in the U.S. The University Corporation for Atmospheric Research (parent of NCAR), owns the intellectual property, namely the wind shear detection algorithm and licenses it to several private companies: MTECH Systems, Climatronics, Telvent, Vaisala, and Vitrociset. The output of the wind shear detection algorithm is sent to the Air Traffic Control Tower, Terminal Radar Approach CONtrol (TRACON), and the Center Weather Service Units (CWSUs) in Air Route Traffic Control Centers (ARTCCs).

Figure 3.40 is a sample LLWAS display, courtesy of Vaisala. Based upon wind sensor measurements along the north-south runway, the wind shear detection algorithm has determined that a microburst alert (MBA) should be issued along runway 18-36 (this gives the runway orientation (180^o or 360^o, depending upon the direction of travel). The alert is for a 45-kt loss in air speed along the runway or just off the north end.

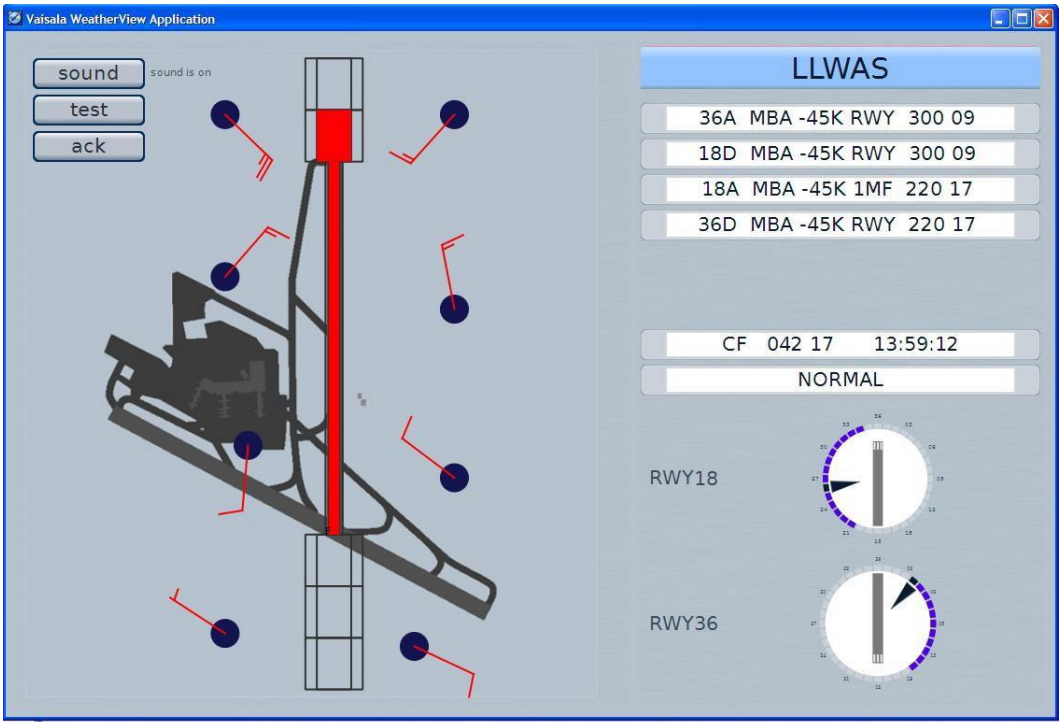

Fig. 3.40. A sample Microburst Alert (MBA) graphic from the Vaisala LLWAS product. The alert indicates a 45-kt loss in airspeed along the runway during landing or takeoff, the only exception being a 45-kt loss one mile from touchdown when landing toward the south. The wind vector at the threshold (point of landing or liftoff) is given at the right of each alert message: 220° at 17 kt at the north end of the runway and 300° at 9 kt at the south end.

One recommendation for NextGen is that wind observations from LLWAS not just be used for wind shear alerts but instead be made more widely available for other applications.

3.4.3 Weather Systems Processor (WSP)

The Weather Systems Processor (WSP) is a hardware and software modification to the existing Airport Surveillance Radar-9 (ASR-9) system. There are 135 ASR-9 systems in the CONUS, mostly at medium capacity airports. Of these, 34 are equipped with WSP. The main function of the ASR-9 is to detect aircraft in the vicinity, and so the antenna spins at 12.5 rpm, much faster than either the TWDR or WSR-88D radars. Though the update rate is very fast, the volume surveyed by the antenna is large: the azimuthal beamwidth is 1.4°, and the elevation beamwidth is 5.0°. Consequently, the ASR-9 is handicapped in detecting microbursts except at close range. False alarms are sometimes triggered by birds or ordinary vertical wind shear

detected because of the wide vertical beam width. Still the WSP provides some protection from strong low-level wind shear at airports that would not otherwise have it.

3.4.4 Integrated Terminal Weather System (ITWS)

The Integrated Terminal Weather System (ITWS) was developed in the early 1990s by MIT Lincoln Laboratory (Evans and Ducot, 1994). The purpose of ITWS is to bring together various sources of data relevant to air terminal operations and generate a suite of products, requiring little meteorological interpretation that could serve tower controllers, ramp controllers, CWSU meteorologists, the TRACON, the TMU, pilots, and airline dispatchers. The inputs include radar data from WSR-88D, TDWR, and ASR-9 radars; LLWAS wind measurements and any other surface observations near the airport; cloud-to-ground lightning data; and RUC analyses and forecasts. Among the products generated are wind shear and microburst detection and prediction (for a few minutes ahead), storm cell intensity and direction of motion, lightning information, detailed winds in the terminal area, and a one-hour storm prediction.

Raytheon Company built production versions of ITWS and began installing them at major airports in 2001. As of July 2010, 32 production systems were commissioned for live operations at major terminal areas, and two more sites are expected to be commissioned shortly.

When ITWS has both TWDR and LLWAS wind sensor data at its disposal, the probability of detection of a microburst causing significant, sudden loss in airspeed is at least 95%, and the false-alarm rate is under 5% (Keohan, 2007). This success is attributed to the TWDR data that provide estimates of the liquid water content in shower clouds above the surface wind sensors, thus giving some indication of the amount of rain that could be evaporated into the downdrafts. The detection algorithm is less effective in dry climates, when radar reflectivity of the shower producing a microburst is minimal.

To conclude this section on low-level wind shear detection and alerting, it is noteworthy that no fatal aircraft accidents due to low-level wind shear have occurred at any airport where the four systems described above have been deployed.

3.5 CWSU Continuous Local and Regional Weather Watch

As noted in earlier sections, Center Weather Service Units (CWSUs) are collocated with Air Route Traffic Control Centers and staffed by NWS meteorologists. This survey has described the two major products generated by CWSUs, namely, Center Weather Advisories (CWAs) and Meteorological Impact Statements (MISs), in the sections on turbulence, icing, and convection. Another important aspect of CWSU operation is discussed here: monitoring continuously the local and regional weather for anything that might affect traffic flow, either en route across the ARTCC region or into or out of major airports. When airspace becomes congested, air traffic control becomes more complicated, and the potential for error increases.

The main customers of a CWSU are the TMU, TRACON, and the control tower. The CWSU is collocated with the TMU at the ARTCC. Communication between CWSU and TMU is mainly verbal, and the two share common displays. TMU personnel have been controllers earlier in their careers and must spend at least eight hours per month *as* a controller. They communicate directly with ARTCC area supervisors about weather events that affect air traffic flow. The air traffic controllers themselves are responsible for aircraft separation, but they need advanced warning to implement various strategies when acceptance and departure rates must be reduced. Events that reduce the acceptance rates at airports or the number of aircraft that can be in transit across an ARTCC region are the focus of this section.

En Route Spacing. Clear-air turbulence or thunderstorms within an ARTCC region may constrict or block standard flight corridors. Controllers may modify flight routes, increase the miles-in-trail, or instruct pilots to change altitudes in order to maintain traffic flow. Widespread turbulence or thunderstorms in adjacent (Tier 1) ARTCC regions may force aircraft to divert from normal routes into another region, thereby increasing traffic flow there. The CSU meteorologist, acting through TMU personnel, must make controllers aware of these situations in advance.

Metering. CWSU meteorologists must inform their customers of any phenomena that might change the flight pattern around the airport or decrease the Airport Arrival Rate (AAR). Specific examples for Denver International Airport (DIA) will clarify how this works. Figure 3.41 shows the layout of runways at DIA.

Fig. 3.41. An aerial view of Denver International Airport. The direction north is up.

Two parallel runways oriented along 160°–340° lie to the northwest of the main terminal and concourses. Two more parallel runways oriented along 170°–350° lie to the southeast. Two single runways, both oriented along 81°–261° lie to the northeast and southwest. Under Visual Meteorological Conditions (VMC), DIA can land 114 flights per hour with three arrival runways in operation, and 152 when four arrival runways are available. AARs during Instrument Landing System (ILS) approaches drop to 96 per hour with three runways open and 64 with two runways open. Rarely, strong west winds will preclude landings and takeoffs on the north-south runways, leaving only two active runways. In this case AARs drop to 48-72 per hour.

During de-icing operations, when the wind is usually from the north, Runways 35R/35L are the primary arrival runways, and 34L/34R are the primary departure runways. De-icing equipment is west of the terminal, and using 34L/34R for departures minimizes travel time from the de-icing area to the head of the runway.

DIA has frequent convective activity during the warm season. When cloud base is high and the boundary layer is dry, thunderstorms or even light showers are often accompanied by

microbursts, which can shut down all runways for many minutes at a time. Microbursts and thunderstorms in the wrong places typically disrupt landings and takeoffs for 30-60 min.

CWSU meteorologists at the Denver ARTCC try to anticipate the following conditions:

- Wind shifts that result in tailwinds of 8 kt or more on landing or takeoff can cause reconfiguration of traffic patterns, which takes an average of 30 min. Frontal passages are usually easy to anticipate because fronts stretch for hundreds of miles and usually move at uniform speed. Gust fronts from distant thunderstorms are inherently mesoscale phenomena. They are not hard to track with dense networks of surface observations, but thin lines on radar can sometimes be the only clue that they are on the way. Thunderstorms approaching DIA may be tracked with radar, but some thunderstorms form *over* DIA and are therefore less predictable. In particular, when a convergence zone embedded within a mesoscale cyclonic eddy forms in the vicinity of DIA, severe thunderstorms are possible (Szoke et al. 1984). Microbursts are hardest of all to anticipate because they descend from nearly overhead, often with little warning. All phenomena associated with wind shifts require quick action by controllers, who do not like to be surprised.
- Thunderstorms or microbursts not over the airport but near arrival or departure gates (sometimes miles off the end of the runway) can prompt gate closure. This forces rerouting of incoming traffic and may force arriving and departing traffic to share the same runway.
- Cloud-to-ground lightning within 5 nm affects ground operations.
- As noted earlier, low ceiling and poor visibility causes reduced AARs. Dense fog decreases AARs even more when airfield conditions drop below minimums for certain types of aircraft, which are then not permitted to land. There are three categories of ILS approaches:
  Category I – A precision instrument approach and landing with a decision height (to land or abort) of at least 200 ft above the touchdown zone elevation, and with a Runway Visual Range (RVR) of at least 1,800 ft.
  Category II – A precision instrument approach and landing with a decision height between 100 and 200 ft and an RVR of at least 980 ft (for most aircraft).
  Category III – This category is further subdivided, but the decision height is always less than 100 ft and the RVR is 660 ft or less.

Major airports have Tactical Decision Aids, based upon TAFs issued for that airport. Figure 3.42 is an example for Denver.

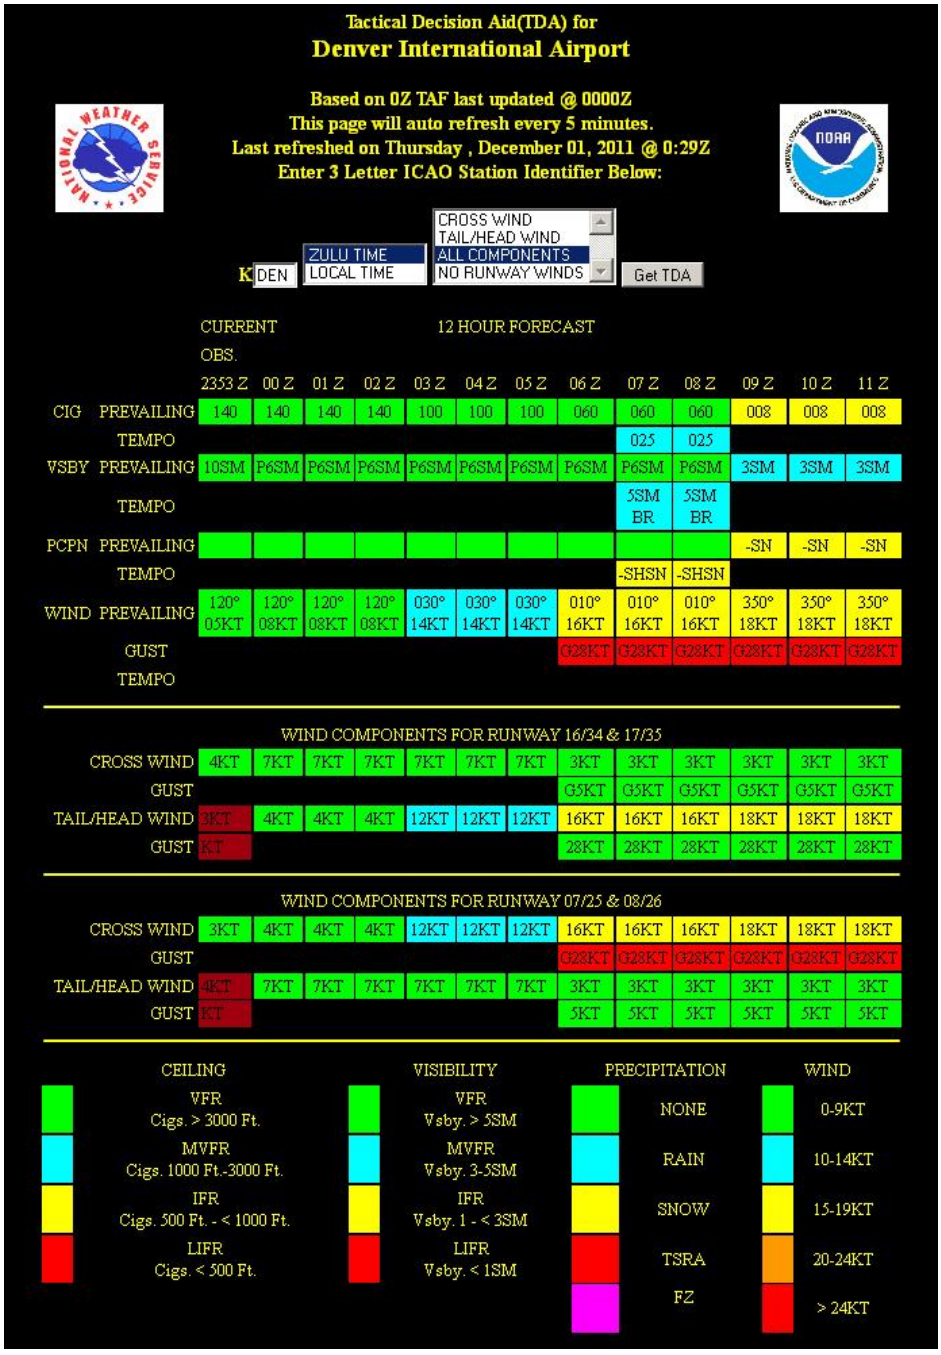

Fig. 3.42. Tactical Decision Aid for Denver. The hour-by-hour TAF is portrayed in the top panel. Wind components along the N-S runways and the E-W runways are shown in the middle panels. The color codes corresponding to various categories of information appear at bottom.

The top panel in Fig. 3.42 gives 12 h of TAFs in color-coded form. A wind shift and frontal passage is expected between 02 and 03 GMT, followed by increasing northerly winds by 06 GMT and light snow by 09 GMT. The middle two panels indicate light crosswinds on the N-S

runways but stiff crosswinds on the E-W runways. This display makes it easy for users to interpret TAFs without reading text.

Ten CWSUs serving major airports have created experimental displays showing arrival and departure gates and their availability. These are available at http://aviationweather.gov/cwsu/. Figure 3.43 is an example from the Dallas-Fort Worth Airport.

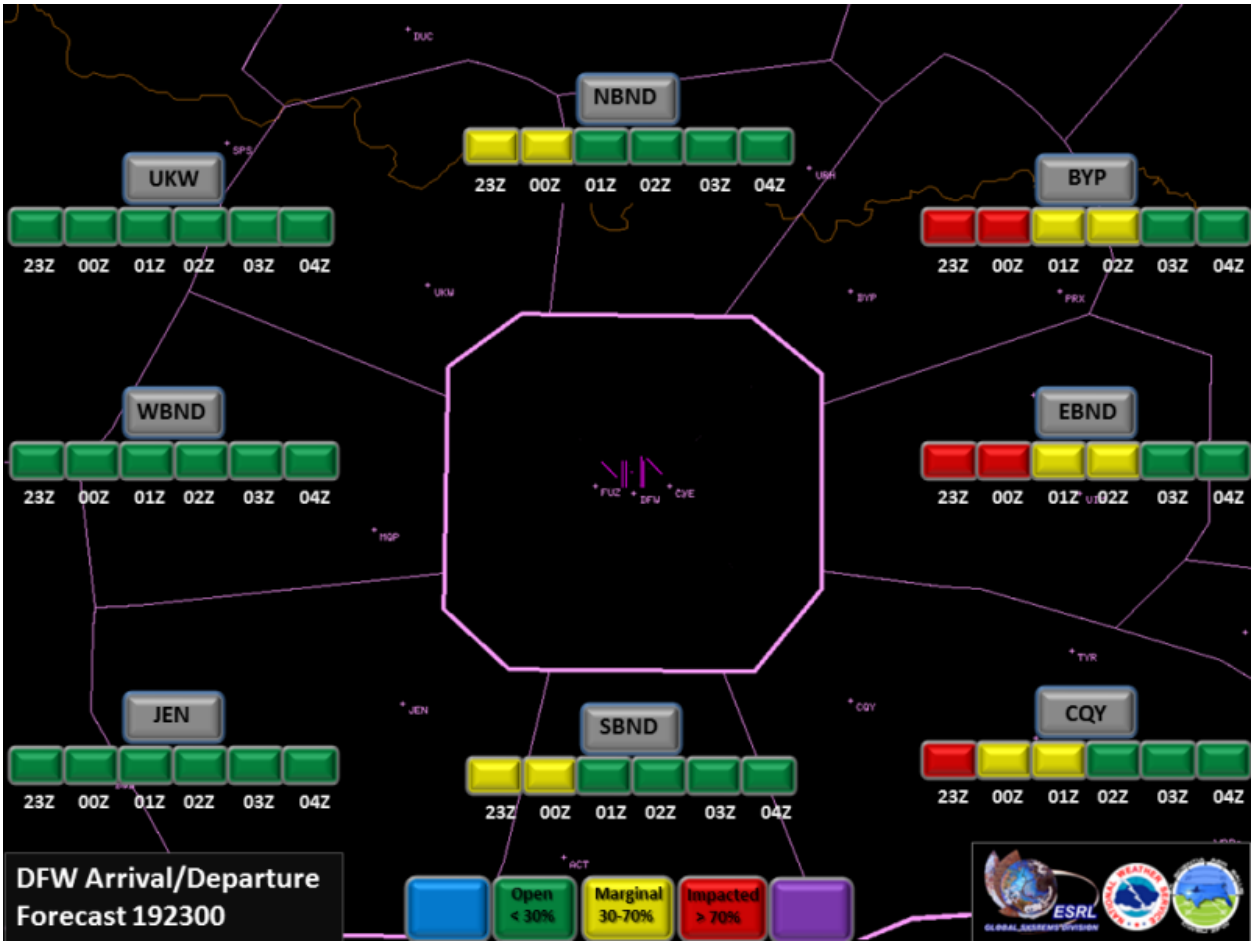

Fig. 3.43. Gate forecast for Dallas-Fort Worth (DFW) Airport. The gates are labeled. The color code indicates the likelihood that convection will block a gate at the indicated time.

Gates to the east are expected to be at least partially blocked for the first few hours. The reason for this is apparent from Fig. 3.44, a weather depiction product generated by the Fort Worth CWSU, current at 2200 GMT, 19 December 2011. The radar reflectivity image shows a north-south line of thunderstorms passing over the DFW area (the purple octagon at the center of Fig. 3.43 and (much smaller) at the center of Fig. 3.44, which covers a larger area). While passing over DFW, these storms caused a number of ground holds. As Fig. 3.43 indicates, these

storms are expected to be east of the airport by 2300 GMT, at least partially blocking some of the approach and departure routes.

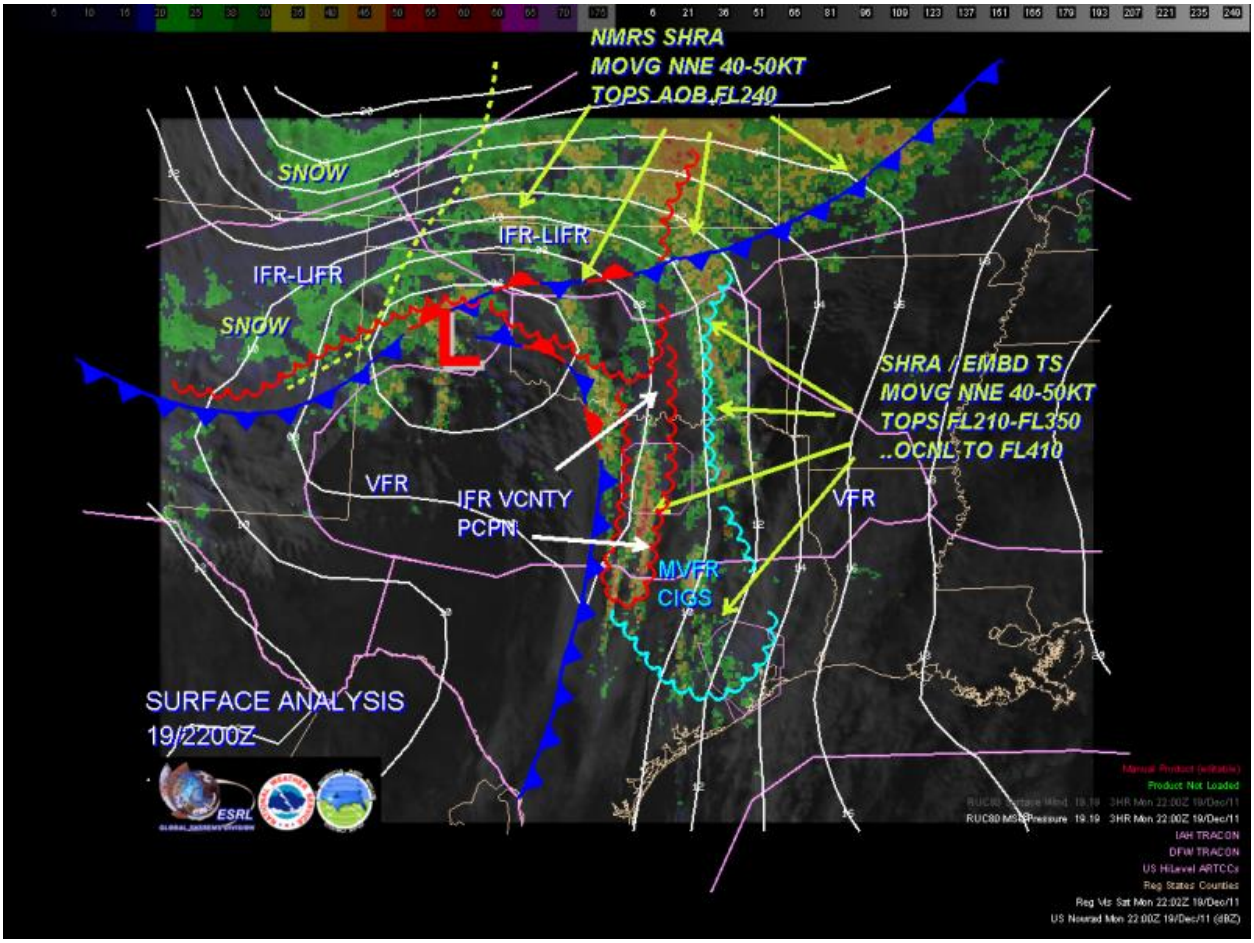

Fig. 3.44. Weather depiction product issued by the Fort Worth CWSU, valid at 2200 GMT, 19 December 2011. The radar image indicates a north-south line of thunderstorms progressing slowly eastward across the TRACON area for DFW (the thin purple octagon at the center of the image). The area enclosed by red scallops has IFR or LIFR conditions. Cyan scallops enclose areas with MVFR conditions. Cold, stationary, and occluded fronts are identified by the bold blue and red symbols. White contours are sea-level pressure contours. ARTCC area boundaries are in purple.

Compression is the word used to describe a decrease in aircraft separation, especially when there is a long line of aircraft approaching an airport. En route, a common in-trail separation is 12 nm. As aircraft descend and slow for final approach, their in-trail spacing decreases. Within 10 nm of the runway it may decrease to 2.5 nm, provided that landing aircraft can clear the runway within 50 s. Wind shear encountered en route or on descent can cause a decrease in separation beyond that effected by merely slowing down. For example, consider two aircraft in

a straight-line flight, one trailing the other by 10 nm and both having an air speed of 300 kt. As they pass Point A, they are enjoying a 125-kt tailwind. During the next 200 nm of travel, suppose the tail wind decreases steadily to 75 kt. At the end of 200 nm, their spacing has dropped to 8.8 nm, even though the airspeed of both is still 300 kt. A similar but more pronounced compression occurs if aircraft with a strong tailwind at altitude descend toward an airport and experience a substantial headwind on final approach. TRACON can prevent unsafe bunching of aircraft by alerting controllers of the likelihood of "meteorological compression." Controllers, in turn, can begin slowing aircraft at greater distances from the airport or lengthening the legs flown on final approach.

CWSUs can alert TRACON personnel of the potential for compression with a Meteorological Impact Statement (MIS) like the following, transmitted by the Denver Center CWSU on 01 December 2011 and valid for five hours from 2000 GMT.

ZDV MIS 02 VALID 012000-020100
...FOR ATC PLANNING PURPOSES ONLY...
...DEN TRACON BRIEFING...
...REFER TO CURRENT SIGMET/AIRMET/CWA ADVISORIES...
MAX UPPER LEVEL WINDS: SW 90-110 KT.
***** ARRIVAL GATE (FL180) WINDS: SW 25-35 KT *******
2000FT AGL WINDS: N 10-15 KT.
TRACON AREA SFC WINDS: N 10-15 KT.
TURBULENCE AOB FL200: OCNL MOD.
ALTIMETERS: ABV 30.50.
TSTMS WRN HALF: NONE
TSTMS ERN HALF: NONE
MICROBURST/LLWS POTENTIAL: NONE.
FREEZING LEVEL: SFC.
ICING AOB FL200: OCNL LGT-MOD.
MOUNTAIN OBSCURATION: WDSPD.
VISUALS/ILS OPERATIONS: ILS APCHS.
WIDESPREAD LOW IFR (CIG LT 500 FT/VIS LT 1 SM): NONE
REMARKS: EARLY EVNG OUTLOOK: ILS APCHS. CHC -SN THRU 02Z. NW-N 06-12KT.

This message implies substantial shear between flight level (max winds SW 90-100 KT), the arrival gate winds at FL 180 (SW 25-35 KT), and the winds 2,000 ft above ground (N 10-15 KT). Flights arriving from the southwest would experience substantial meteorological compression on descent into Denver. (Note also mention of light to moderate icing at or below FL 200, clouds obscuring mountains, and the need for an ILS approach.)

4. The Special Case of Alaska

4.1 Alaska's Uniqueness

Why should the State of Alaska be considered a special case with regard to aviation weather hazards? There are many reasons:

- *Size and spread*. With a land mass of 570,000 square miles Alaska spans nearly 3,000 miles in the east-west direction. The Alaska Region of the NWS includes three forecast offices. When overlaid on the CONUS, their contiguous areas of responsibility cover the majority of the lower 48 states (Fig. 4.1). The Alaska Aviation Weather Unit (AAWU) in Anchorage is Alaska's equivalent of the Aviation Weather Center in Kansas City. The AAWU and collocated CWSU are responsible for an area even larger than that shown in Fig. 4.1. It extends northward from Alaska to the North Pole in a pie-shaped wedge bounded roughly by 140° and 170° W longitude (Fig. 4.2).

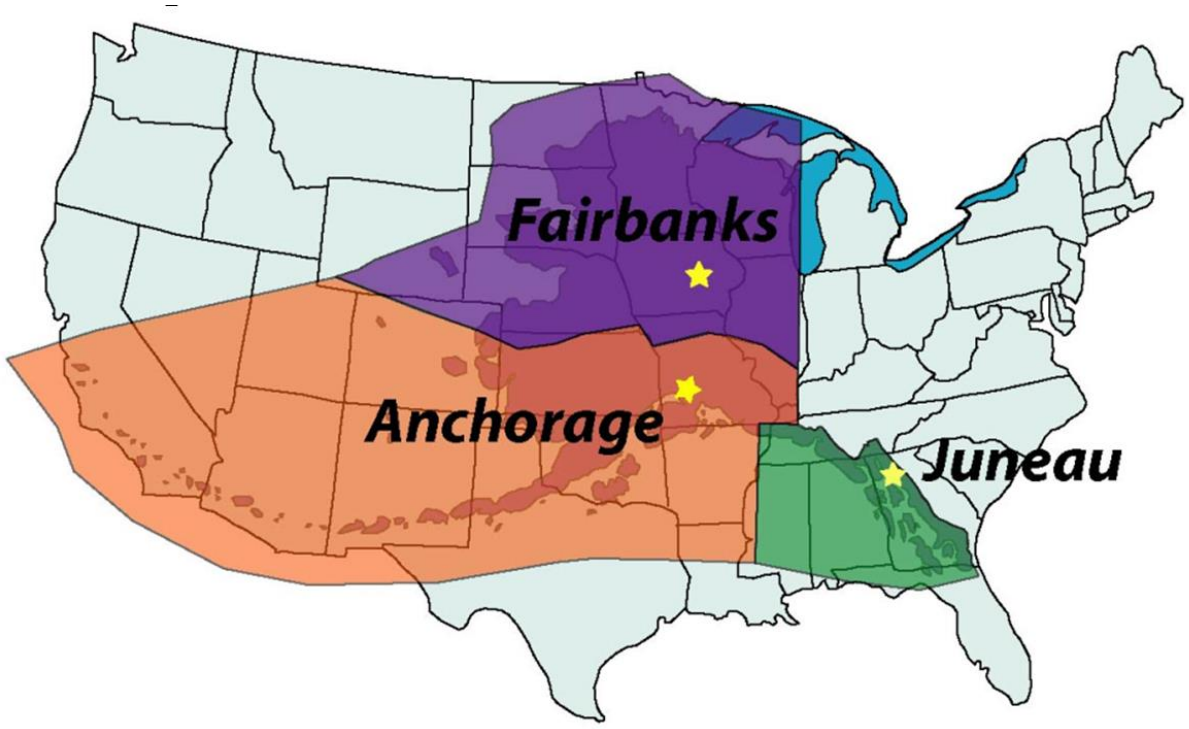

Fig. 4.1. The Alaska Region of the National Weather Service has a huge area of responsibility that would span most of the CONUS when overlaid on top. The Anchorage forecast office covers southern Alaska and the Aleutian Islands. The Juneau office covers southeast Alaska. The Fairbanks office covers northern Alaska.

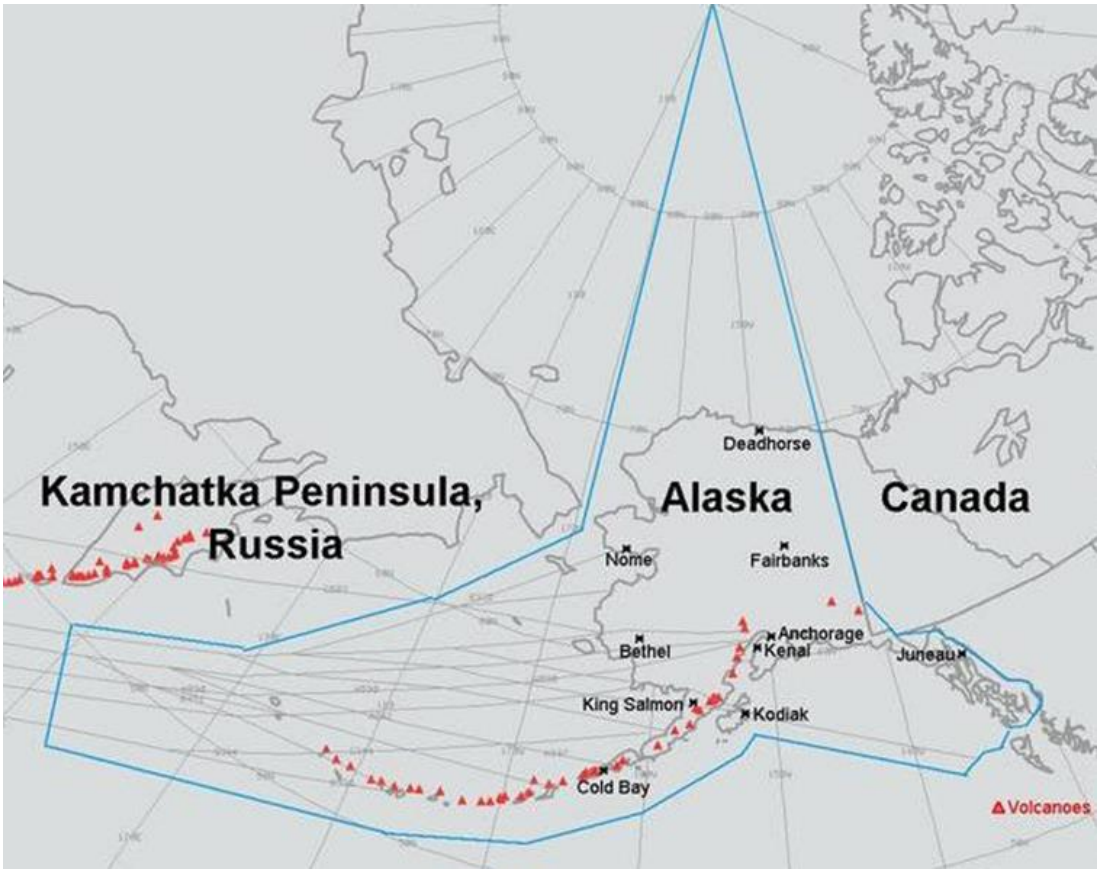

Fig. 4.2. The area of responsibility for the AAWU/CWSU is outlined in blue. Small red triangles locate historically active volcanoes along the Aleutian Island chain at bottom and the Kamchatka Peninsula, left. Major jet routes from the Far East are indicated.

- *Volcanoes*. More than 50 volcanoes in Alaska have been active in historical times, as shown in Fig. 4.2. As the recent eruption of the Eyjafjallajökull Volcano in Iceland amply demonstrated, the injection of volcanic ash into the upper troposphere and lower stratosphere can cause major disruptions in air travel, even after the ash has traveled thousands of miles. Volcanic ash can damage jet engines and even cause in-flight failure. The number concentration of ash particles that would allow safe travel through a volcanic plume is not known. Thus, the currently accepted practices are detection and avoidance.
- *Complex terrain*. Alaska has by far the most glaciers of any state, some reaching to sea level; the highest mountain in North America (Denali, 20,320 ft); the 17 tallest peaks in the U.S.; rugged and mountainous coastlines; and numerous islands and inlets. There are relatively few roads in the state, making air travel the principal means of transportation in many areas.

- *High latitude*. Alaska's northerly location gives it both a cold climate and large seasonal variations in the length of day. Alaska's interior can be exceptionally cold (-60^{o}) in winter.
- *Maritime influence*. Surrounded on three sides by water—the Gulf of Alaska on the south, the Bering Sea on the West, the Chukchi Sea on the northwest, and the Beaufort Sea on the north, Alaska has two-thirds of all the nation's coastline. Much of this coastline is exposed to fog, low clouds, strong wind, and precipitation for much of the year.
- *Reliance on air travel*. According to the Alaska Department of Transportation (AK DOT), approximately 82% of Alaskan communities are not served by roads. The practical alternatives are air travel, dogsled, snowmobiles, a few ice roads in winter, and sea travel, which is viable only after sea ice has melted. Alaska's bush air services are highly developed. Alaska has six times as many pilots (over 8500 out of an estimated state population of 664,000 in 2005) and sixteen times as many aircraft per capita, compared with the rest of the U.S. The AK DOT owns 255 rural airports. Municipalities own several dozen more, and there are hundreds of privately owned strips. Bush pilots land on glaciers, lakes, and grassy or gravel strips far more than on paved runways.

Controlling airspace around Anchorage is unusually difficult because of the presence of four airports within a circle of 6-mile radius. Foremost among these is Ted Stevens Anchorage International Airport (ANC). During 2006, the airport handled an average of just under 800 aircraft operations per day. In 2010, it served over 5 million passengers. ANC is perhaps more noted for its cargo traffic than its passenger traffic. According to AK DOT, it ranks fifth in the world for cargo throughput after Hong Kong, Memphis (site of the FedEx "super-hub"), Shanghai, and Incheon (South Korea). In the U.S., it ranks second (after Memphis) for landed weight of cargo aircraft. ANC boasts of being 9.5 h or less flight time from 90% of the industrial world. Both east- and westbound cargo carriers, which benefit from short route segments, use ANC frequently.

The Lake Hood Airport is billed as the world's largest and busiest for floatplanes. It is the only facility of its kind requiring a control tower to deal with the traffic, mostly bound for remote parts of Alaska. Merrill Field is a general aviation (GA) airport. Aircraft using the field cannot weigh more than 12,500 lb. Of over 10,000 aircraft registered in Alaska, 877 are based at Merrill Field. Merrill Field was the 72nd busiest *in the nation* in 2010, with 145,000 flight operations. Finally, Anchorage is home to a large military base, the Joint Base Elmendorf Richardson. Many high-performance jet aircraft fly from there.

A special Federal Air Regulation establishes how air traffic is segregated among these four airports. Control of the airspace is both complex and nonstandard.

4.2 Alaska Aviation Weather Unit (AAWU)

The AAWU is located just south of ANC airport. It serves the same functions as the Aviation Weather Center in Kansas City and almost as large an area. Its website is at http://aawu.arh.noaa.gov. Figure 4.3 shows the home page, with "AIRMETs" selected. By dragging the cursor slowly over the map, one can see text versions of AIRMETs for each subregion. The reader can explore a great variety of products listed at the top of the page.

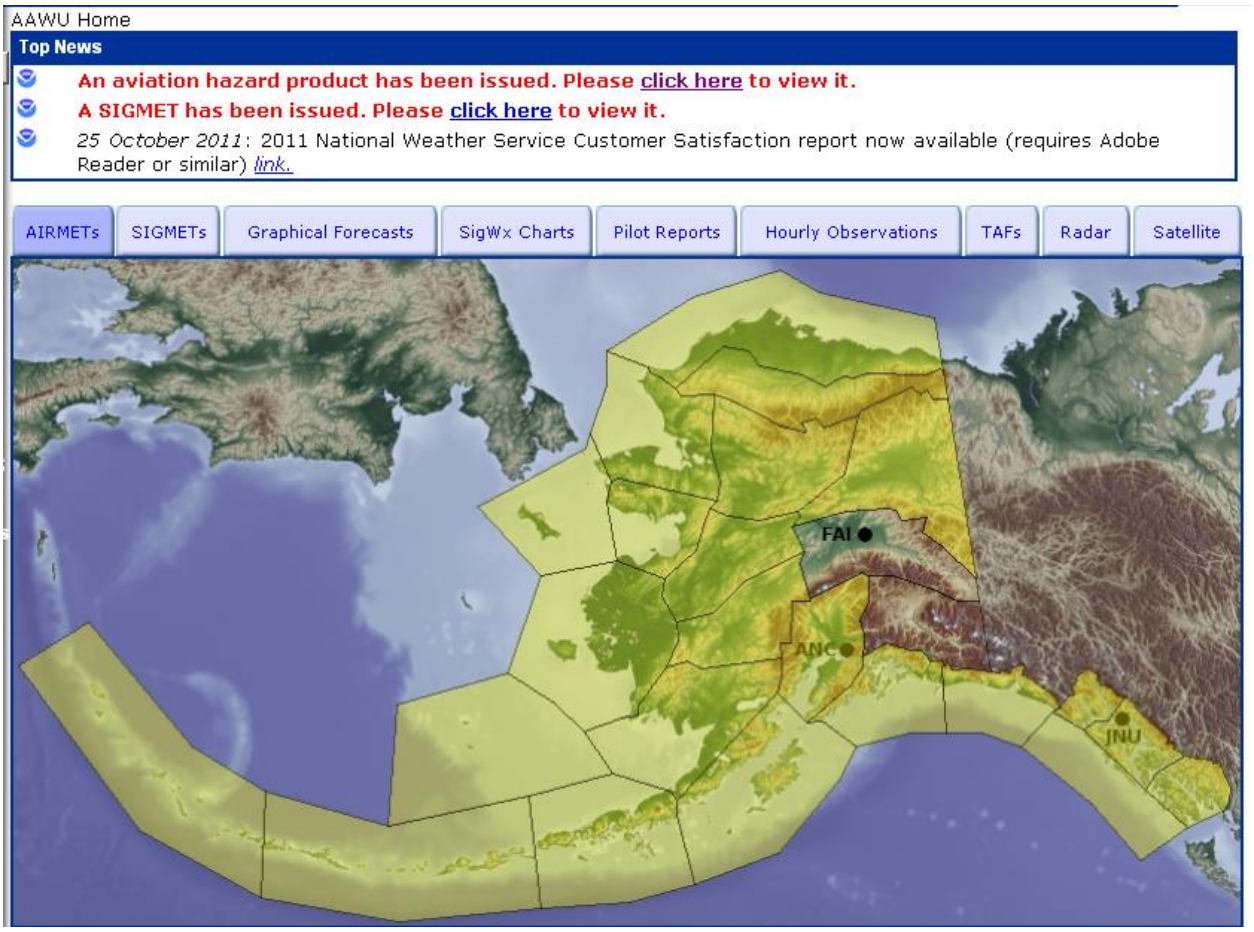

Fig. 4.3. The home page of the Alaska Aviation Weather Unit. The map shows the subregions that the AAWU is responsible for.

At the very top of the product in Fig. 4.3 is the message **An Aviation Hazard Message has been issued. Please click here to view it.** Clicking brings up the graphic in Fig. 4.4. The hazard covers much of south central Alaska, including the Anchorage area.

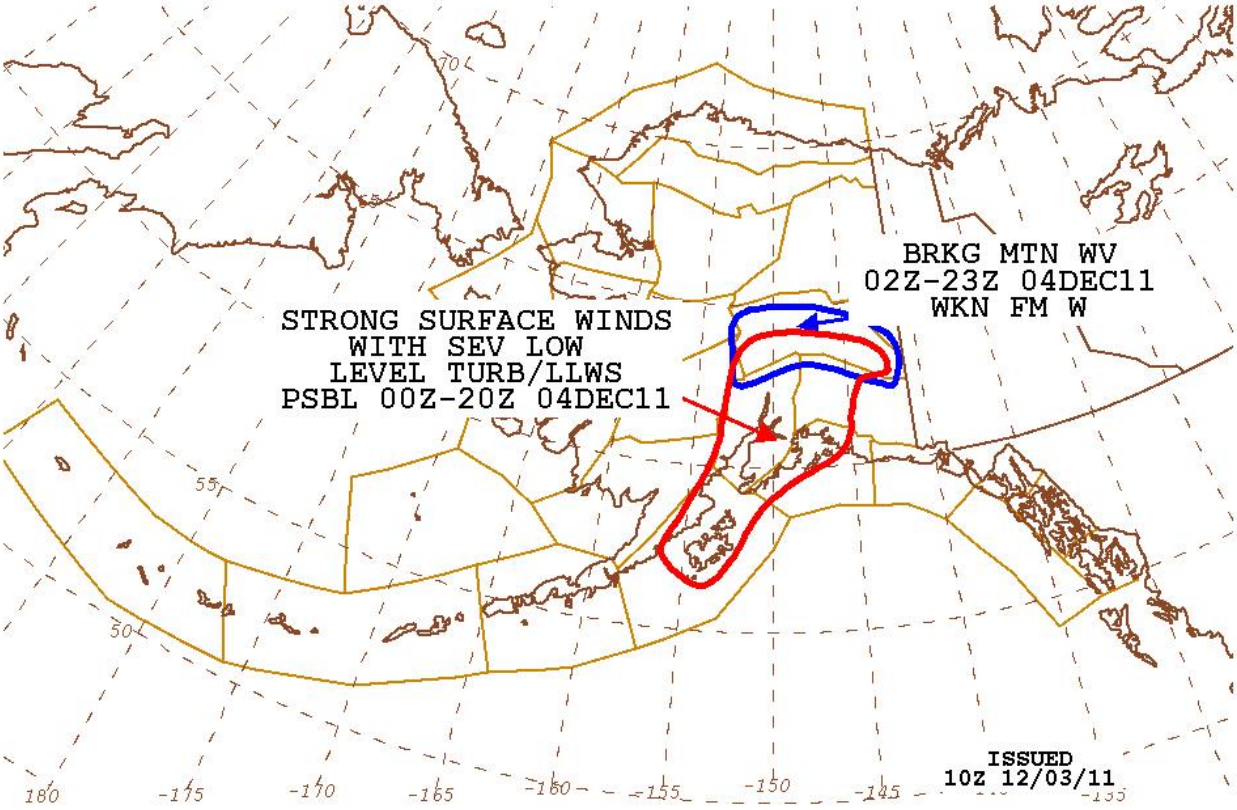

Fig. 4.4. A hazard map issued by AAWU, calling for severe low-level turbulence and the possibility of low-level wind shear early on 04 December 2011 (GMT).

The SIGMET associated with Fig. 4.3 reads as follows:

```
WSAK02 PAWU 031754
SIGAK2
NCJ WS 031750
PAZA SIGMET JULIET 1 VALID 031750/032150 PANC-ANCHORAGE FIR.
OCNL SEV TURB BTN SFC-FL030 AREA WI 50 NM N DLG - 80 NM NE AKN -
30 NM SE PDN - 70 NM W DLG - 50 NM N DLG. MOV NE 10KT. NC.
LLWS 30-40 KTS.
NCE DEC 2011 AAWU
```

The additional details provided by the SIGMET are that the severe turbulence can occur between the surface and 3,000 ft AGL, and the low-level wind shear can result in a sudden loss of airspeed, 30-40 kt.

The AAWU is handicapped by a relative lack of atmospheric data as compared with the lower 48 states. TAFs are prepared by the Anchorage, Juneau, and Fairbanks WFOs, as shown in Fig. 4.5, but compare the number of TAFS (37) with the many hundreds of airports and landing strips in Alaska, and the difficulty of tracking surface weather conditions at all of them becomes immediately apparent.

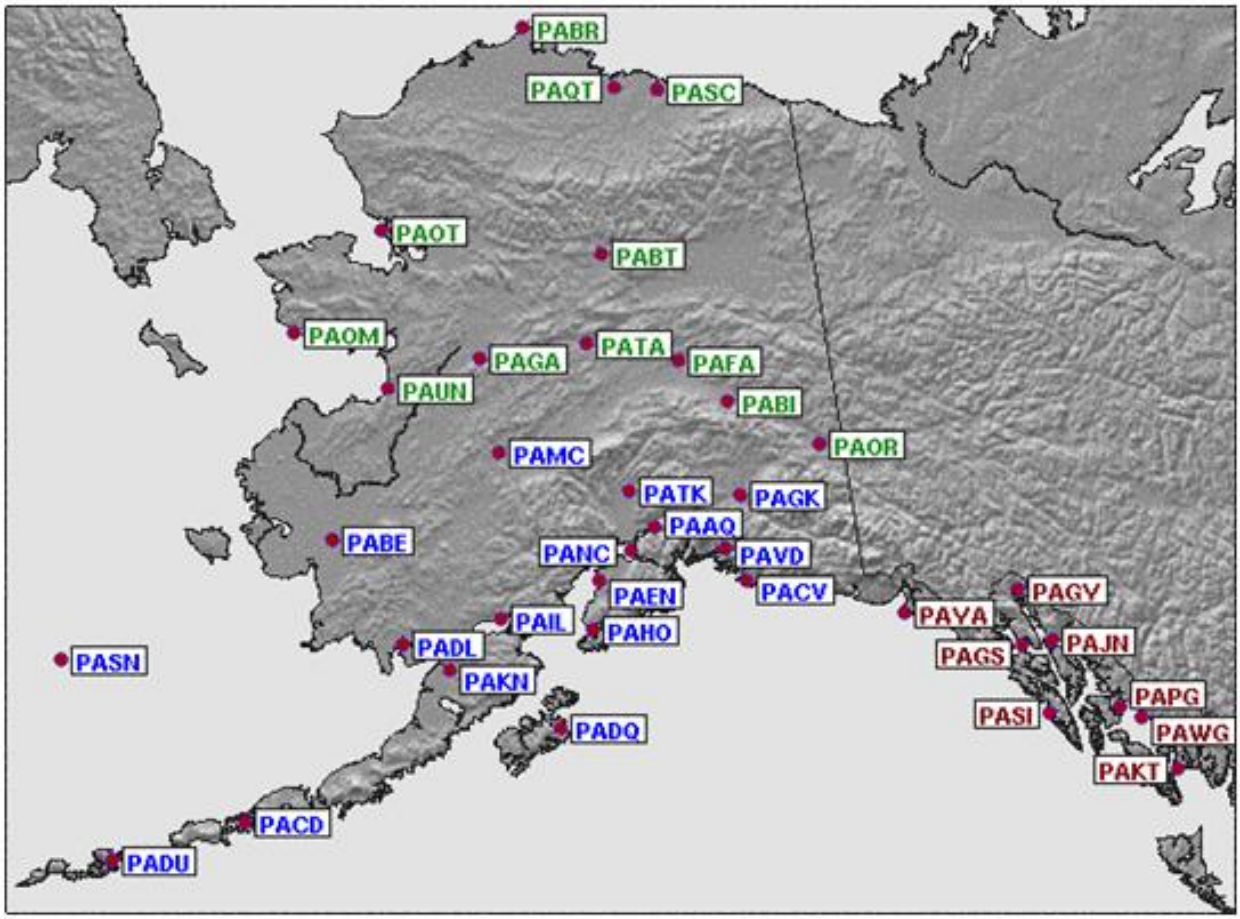

Fig. 4.5. The 37 Alaska locations for which TAFs are prepared. The Fairbanks WFO issues TAFs for northern Alaska (green), the Anchorage WFO issues TAFS for south central and southwest Alaska (blue), and the Juneau WFO issues TAFs for southeast Alaska (red).

Alaska has 13 rawinsonde stations compared with 72 in the CONUS, three wind profilers compared with about 30 in the CONUS. Automated aircraft reports from commercial airliners, both en route and on ascent or descent, are far more numerous in the CONUS than in Alaska because most flights terminate or originate in Anchorage and arrive from southwest or southeast directions, giving sparse coverage inland. There is significant traffic only between four airports: Anchorage, Fairbanks, Bethel, and Juneau, and hence not many ascent or descent soundings. For these reasons, satellite data are far more critical in Alaska than in the CONUS for aviation weather applications. Here again, Alaska has a handicap in that GOES, routinely providing 15-min images of the same area in the CONUS, has only a very oblique view of Alaska, given its high latitude. Thus, polar orbiting satellites supply the bulk of the imagery used. In winter, daylight is short in southern Alaska, and non-existent for a time north of the Arctic

Circle, so visible images are of little or no use. Many of the sophisticated algorithms used to detect and predict turbulence, icing, and convection are not available for Alaska. Fortunately, convection is not the problem in Alaska that it is in the CONUS.

Mesoscale forecasts from the operational North American Model are heavily used in aviation weather applications in Alaska. The Rapid Update Cycle (RUC) extends little beyond the CONUS, and so has not been helpful. Its successor, the Rapid Refresh (RAP), to become operational in January 2012, will cover Alaska and provide hourly analyses and forecasts for the first time in that area.

In summary, though the aviation weather products in Alaska address the same hazards as in the CONUS, they lack the sophistication of many of the CONUS products, principally because the observational assets in Alaska do not match those of the lower 48 states, and the model guidance lacks the refresh rate available in the RUC.

Alaska contends with the same aviation weather hazards as the CONUS, with one very notable exception: volcanoes. Mt. St. Helens is the only volcanic eruption in the lower 48 states in most people's lifetime. Even the drift of a much-diluted ash cloud across the CONUS is a rare event. Thus, Alaska, with at least one volcano threatening air routes on most days of the year, devotes considerable attention to this hazard, the subject the next section.

4.3 Volcanic Ash

Explosive eruptions release volcanic ash into the atmosphere. Composed of small, jagged pieces of rock, minerals, and volcanic glass, volcanic ash particles range in size from 1 μm to 1 mm. They are very hard and abrasive, do not dissolve in water, are mildly corrosive, and conduct electricity when wet. The largest particles have substantial fall speeds and settle out of the plume within a short time, close to the eruption. The smaller particles fall out more slowly and drift with the wind. The tiniest particles can remain suspended for many months and encircle the globe, especially if a very energetic eruption propels them into the stratosphere, where there is no precipitation to wash them out. Engines of jet aircraft have failed after flying through clouds of even thinly dispersed ash, and several disastrous crashes have been narrowly averted. The first few hours after an eruption are the most critical because that is when ash concentrations are high, and any aircraft unfortunate enough to fly into an ash cloud during this time, especially at night when visual cues are lacking, could quickly be in trouble. This section

describes strategies for quick detection of eruptions and predicting where the ash plume will drift.

4.3.1 Agency Coordination

The AAWU manages the Anchorage Volcanic Ash Advisory Center (VAAC) and is collocated with it. (Two VAACs are responsible for volcano alerts in the U.S., the one in Anchorage and another in Washington, D.C. at the NESDIS Satellite Analysis Branch.) Each VAAC

- Develops and executes volcanic ash dispersion models in real time;
- Continuously uses satellite information to identify volcanic ash clouds and to discriminate them from weather clouds;
- Issues Volcanic Ash Advisory Statements, and provides guidance to AAWU or AWC regarding volcanic ash SIGMETs;
- Coordinates with the aviation community, the public, and neighboring VAACs about volcanic activity.

The Anchorage VAAC in particular has a close working relationship with the Alaska Volcano Observatory (AVO), with offices in Anchorage and Fairbanks. The AVO is operated jointly by the U.S. Geological Survey (USGS), the Geophysical Institute of the University of Alaska Fairbanks, and the State of Alaska Division of Geological and Geophysical Surveys.

In the event of an eruption, here is the sequence of events:

- AAWU/VAAC meteorologists use input from the AVO, satellite images, radar images (if any), and pilot reports to determine that an eruption has occurred and to estimate its intensity.
- AAWU issues an eruption SIGMET to warn pilots of the danger.
- One or more computer models are run to predict the movement of the ash cloud, using the latest and most detailed information about the plume.
- A Volcanic Ash Advisory Statement is issued describing the three-dimensional location of the ash plume.
- SIGMETs and advisories are updated to keep everyone informed of the situation.

Figure 4.6 is a sample Volcanic Ash Advisory issued by the Anchorage VAAC for Sheveluch Volcano on the Kamchatka Peninsula in October 2010. Sample products from the archive are available at http://vaac.arh.noaa.gov/list_vaas.php

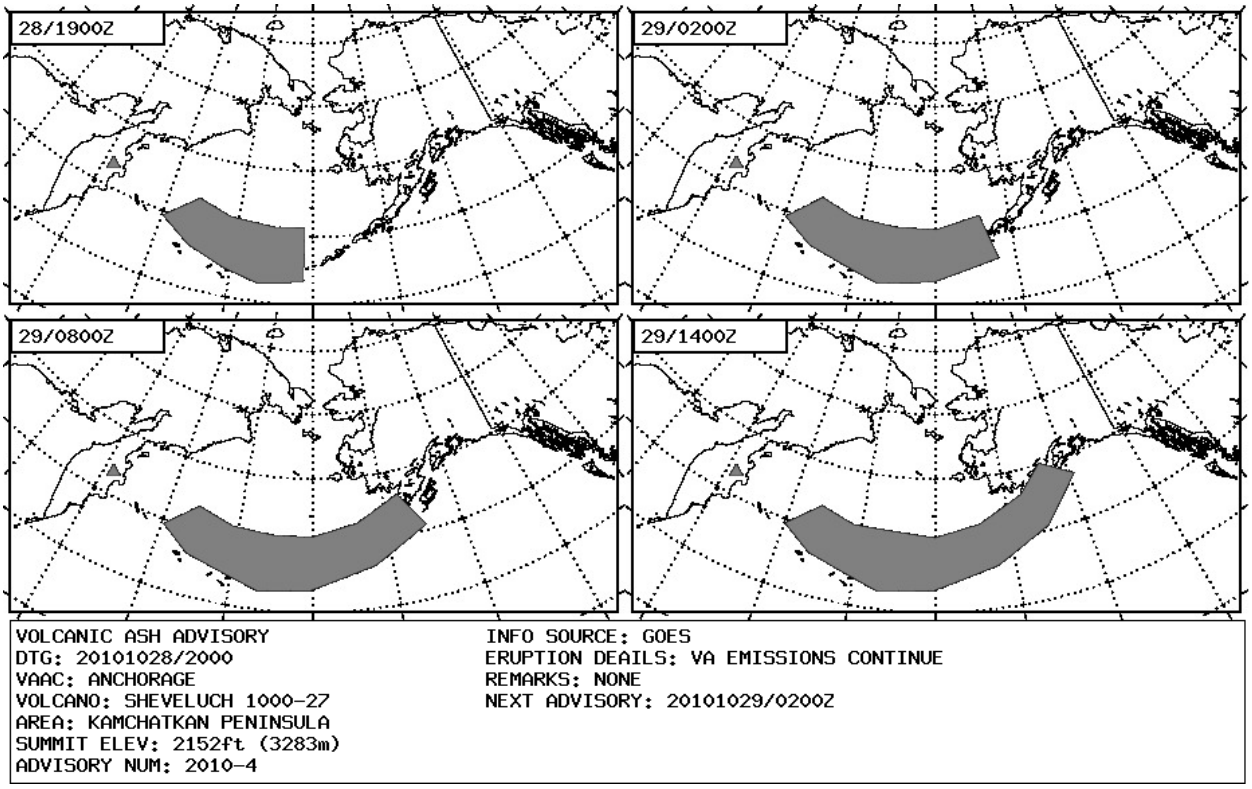

Fig. 4.6. A Volcanic Ash Advisory issued by the Anchorage VAAC at 2000 GMT, 28 October 2010. The volcano (triangle, middle left) is no longer erupting. The ash cloud is drifting over the Aleutian Islands. The listed summit elevation is correct in meters, but not in feet.

AVO monitors volcanoes by analyzing satellite images twice daily for thermal anomalies and ash plumes. (Satellite downlinks are in Fairbanks.) Thermal anomalies can sometimes be detected several weeks before large eruptions. AVO also operates networks of recording seismometers, deployed around more than 20 volcanoes on the Alaska Peninsula and the Aleutian Islands. Seismic activity often increases months to weeks prior to volcanic eruptions. Part of AVO's mission is to provide accurate and timely information on volcanic hazards and warnings of dangerous activity to local, state, and federal agencies and the public. AVO provides advisory assistance to the Kamchatka Volcanic Eruption Response Team when volcanoes erupt over Kamchatka. Detailed information on all of Alaska's volcanoes is available at the AVO website, http://www.avo.alaska.edu.

4.3.2 Selected Volcanic Ash Products

Probably of greatest interest is the map of current volcanic activity, shown in Fig. 4.7, issued by the AVO. (The ANC VAAC issues a similar map at http://vaac.arh.noaa.gov/.) Cleveland Volcano (yellow triangle, bottom middle), was the only active Alaskan volcano as of this writing. Details on the Cleveland Volcano (below) are available on the AVO website at http://www.avo.alaska.edu/activity/Cleveland.php. See below.

December 5, 2011 12:18 PM (Monday, December 5, 2011 21:18 UTC)
CLEVELAND VOLCANO (CAVW #1101-24-)
52°49'20" N 169°56'42" W, Summit Elevation 5676 ft (1730 m)
Current Volcano Alert Level: **ADVISORY**
Current Aviation Color Code: **YELLOW**

A clear satellite image in the past 24 hours detected elevated surface temperatures at the summit. AVO has received no reports of activity.

The new lava dome began growing at Cleveland in late July 2011 and eventually reached about 10 m (33 feet) below the crater rim. Recent views of the summit show gradual subsidence and disintegration of the summit dome. Throughout this period of unrest and eruption, AVO has not observed deposits or features indicative of explosions or ash emission. Cleveland's activity from late July to present can be characterized as an effusive lava eruption. The volume of the lava dome is about one million cubic meters (1.3 million cubic yards).

We are uncertain if the eruption is over or has merely paused. It is possible for effusion to resume at any time, which could send lava over the crater rim and down the steep flanks of the volcano. Such lava flows could collapse and produce avalanches of hot debris that reach the sea and may be accompanied by small ash clouds.

It also remains possible for sudden explosions and ash emission to occur, and ash clouds exceeding 20,000 feet above sea level may develop. Such explosions and their associated ash clouds may go undetected in satellite imagery for hours. However, in cooperation with the University of Washington, AVO has implemented a lightning alarm system that may detect significant ash-producing events within minutes of onset. If a large explosive event occurs, seismic signals may be recorded on AVO seismic networks at nearby volcanoes. There is no real-time seismic monitoring network on Mt. Cleveland.

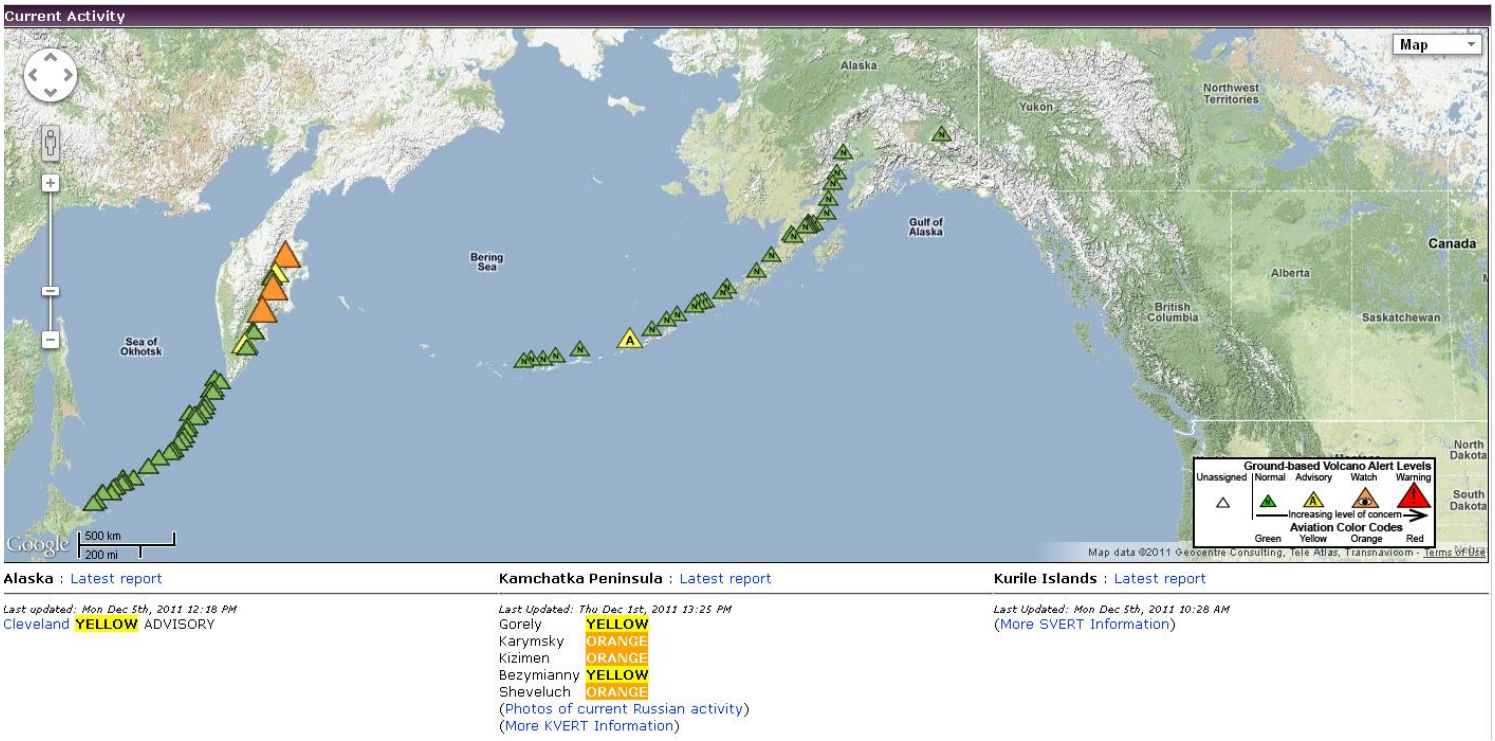

Fig. 4.7. The status of volcanoes along the Aleutian Island Chain (middle), the Kamchatka Peninsula (middle left), and the Kurile Islands (lower left) as of 01 December 2011. The colors of the triangles indicate the level of alert, as indicated at the bottom of the figure and in the description immediately below. Cleveland (the lone yellow triangle) was the lone active volcano in the Aleutian Islands as of 01 December 2011. The aviation color code shows at lower right and is described in more detail immediately below.

Aviation Color Codes for Volcanic Activity

GREEN Volcano is in typical background, noneruptive state or, *after a change from a higher level*, volcanic activity has ceased and volcano has returned to noneruptive background state.
YELLOW Volcano is exhibiting signs of elevated unrest above known background level or, *after a change from a higher level*, volcanic activity has decreased significantly but continues to be closely monitored for possible renewed increase.
ORANGE Volcano is exhibiting heightened or escalating unrest with increased potential of eruption, timeframe uncertain, **OR** eruption is underway with no or minor volcanic-ash emissions [ash-plume height specified, if possible].
RED Eruption is imminent with significant emission of volcanic ash into the atmosphere likely **OR** eruption is underway or suspected with significant emission of volcanic ash into the atmosphere [ash-plume height specified, if possible].

NOAA's Air Resources Laboratory (ARL) computes hypothetical ash trajectories every six hours for several volcanoes, two in Alaska and one on the Kamchatka Peninsula, so that they are immediately available if an eruption is detected. The trajectories are computed with the HYbrid Single-Particle Lagrangian Integrated Trajectory (HYSPLIT) model, developed by ARL (Draxler and Hess 1998a,b), using output from both the GFS and NAM models. Products are available at http://ready.arl.noaa.gov/READY_traj_alaska.php. Sample output from HYSPLIT and the GFS model is shown in Fig. 4.8.

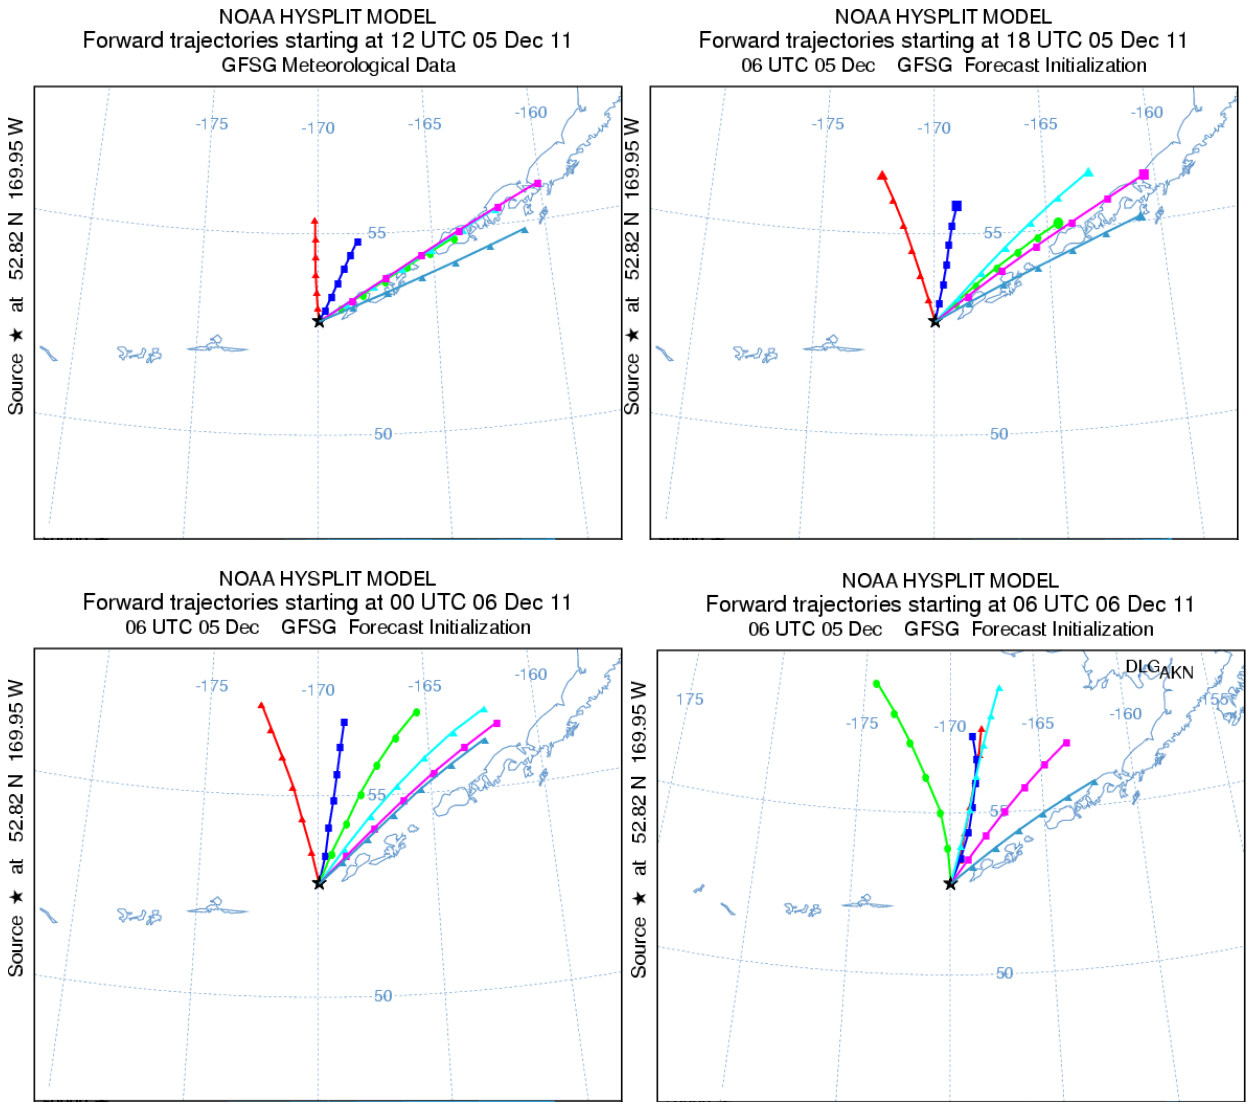

Fig. 4.8. Six-hour hypothetical trajectories computed from GFS model output with the HYSPLIT trajectory model. Trajectories begin at 1200 GMT (upper left) and 1800 GMT (upper right) on 05 December, and at 0000 GMT (lower left) and 0600 GMT (lower right) on 06 December 2011. Red, dark blue, green, cyan, pink, and light blue indicate trajectories of particles released at 5,000, 10,000, 20,000, 30,000, 40,000, and 50,000 ft, respectively. The marks along the trajectories represent one-hour positions.

Each panel of Fig. 4.8 should be accompanied by an indicator of the vertical displacement of each trajectory by the hour. To save space, only the vertical displacement for parcel trajectories starting at 0600 GMT on 06 December 2011 (bottom right of Fig. 4.8) are shown in Fig. 4.9. The vertical displacements for the bottom and top trajectories are not listed and are difficult to

estimate from the figure, but ash injected at 10,000 ft rises to 14,000 ft down wind six hours later, ash injected at 20,000 ft rises to 23,000 ft, etc.

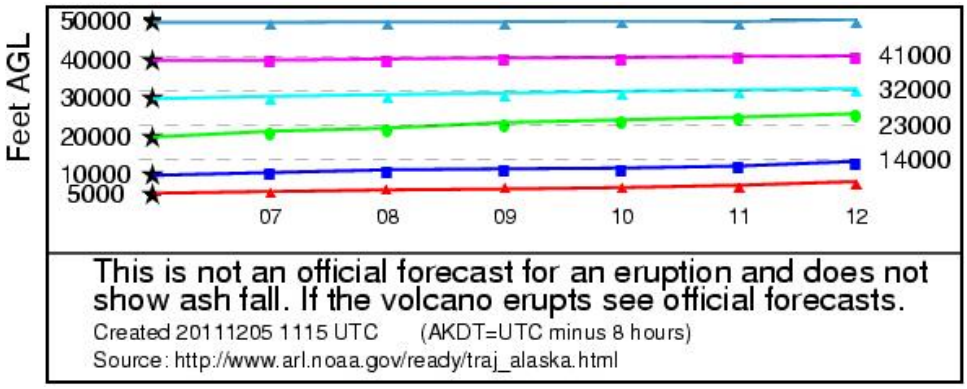

Fig. 4.9. The vertical displacement of volcanic ash during six-hour trajectories, whose horizontal paths are indicated in the lower right panel of Fig. 4.8. The color code used here is the same as in Fig. 4.8. The release altitudes are at left; the final altitudes are at right.

Another way of looking at trajectories is illustrated in Fig. 4.10. These trajectories come from the PUFF Model, developed at the University of Alaska Fairbanks (Searcy et al. 1998). At the time of this writing, the PUFF model was generating new trajectory forecasts every three hours for 14 volcanoes ringing the North Pacific Ocean, including Cleveland Volcano. Access to these trajectories, computed with the help of the GFS model, is available at http://puff.images.alaska.edu/monitoring.shtml. Click on "Cleveland," then on the desired altitude of ash injection, "10 km" in the case of Fig. 4.10. The panels show the plume 1 h, 3 h, 6 h, and 10 h after injection. In this hypothetical example, injection continues for 3 h, then shuts off. Subsequently, the plume drifts away from the eruption site. The color code at the bottom of Fig. 4.10 shows the altitude of particles initially injected at 10 km. Most particles remain near 10 km, but some drift or settle to lower altitudes.

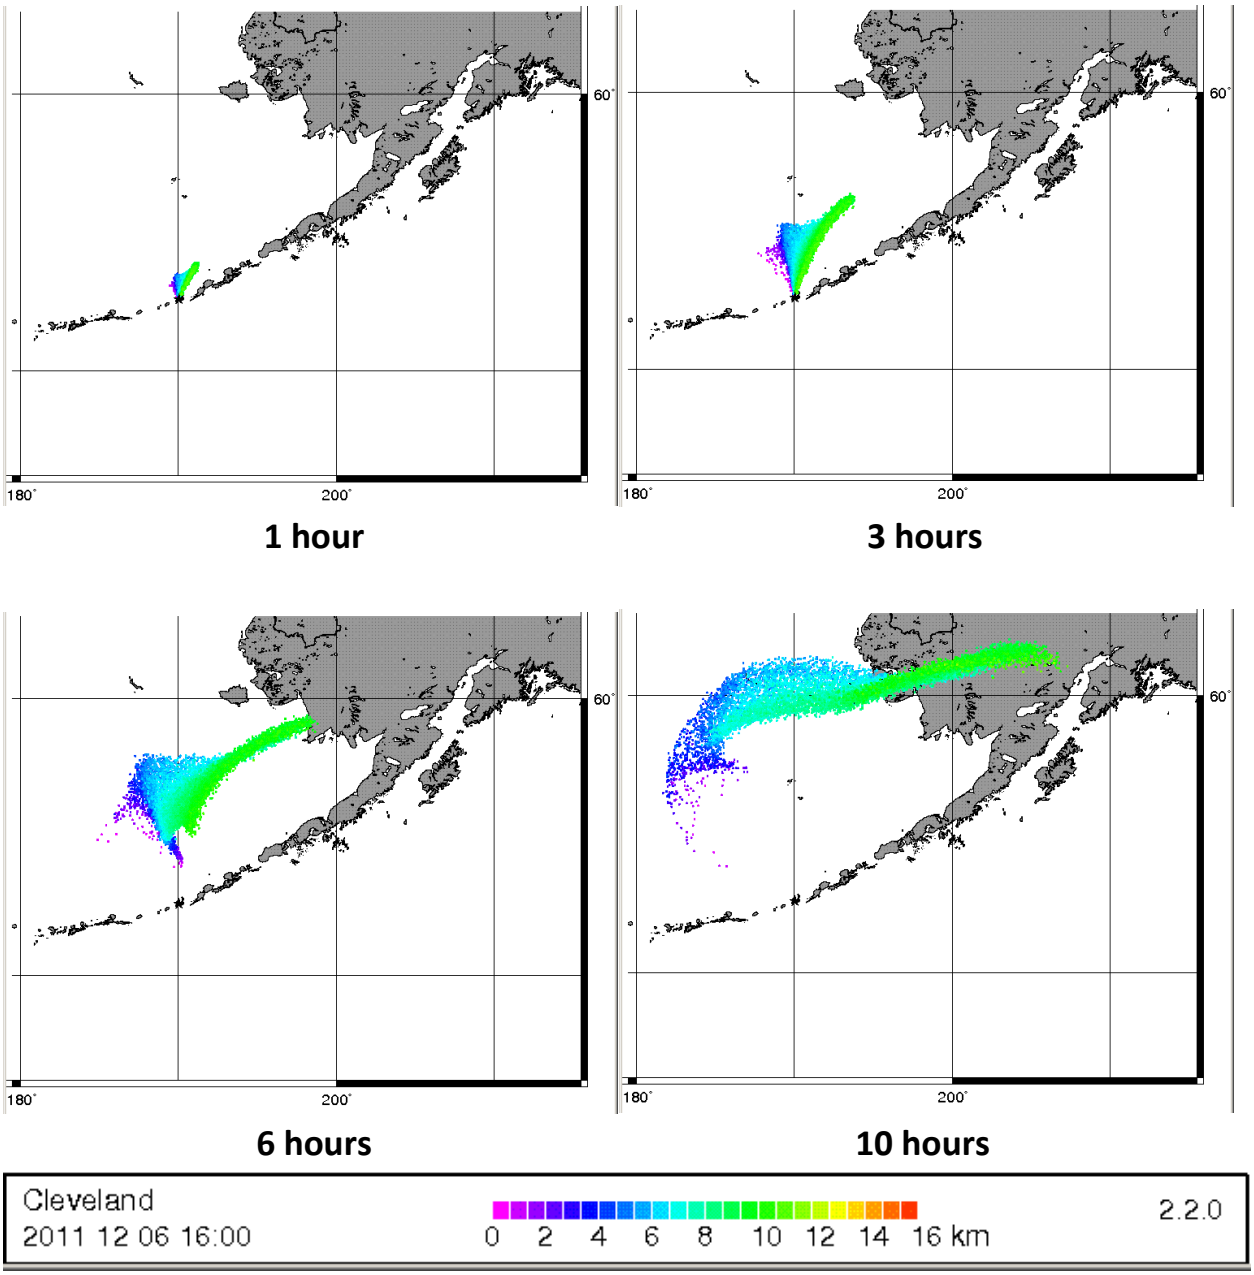

Fig. 4.10. PUFF model trajectories 1, 3, 6, and 10 h after injection by Cleveland Volcano at 10-km altitude. Release is at 1500 GMT 06 December 2011. The altitude of individual particles is color coded as indicated at bottom. These are hypothetical trajectories; Cleveland Volcano was not erupting at this time.

The two trajectory models represented in Figs. 4.8 and 4.10 have a lengthy heritage and are discussed at greater length in the following section.

4.3.3 Trajectory models

A comprehensive review of trajectory calculations was published by Stohl (1998). At that time, errors of 20% of the distance travelled were typical for trajectories computed from a succession of analyzed wind fields. Larger errors, up to 30% of the distance travelled, were typical for trajectories calculated from model forecasts. A more recent paper by Witham et al. (2007) tested four trajectory models used around the world and concluded that all produced similar dispersion patterns of *visual* ash plumes (nominally a concentration of 10^{-4} grams of ash per cubic meter). The models differed in predicted ash concentrations at specific flight levels.

The two trajectory models used by the AAWU are described here.

4.3.3.1 The PUFF model

As the word implies, a "PUFF" model deals with an eruption of short duration, minutes to hours, as opposed to a continuous eruption that may last for many days. The PUFF model originated in 1991 with H. L. Tanaka at the University of Tsukuba in Japan. It was further developed and expanded by Craig Searcy at the University of Alaska Fairbanks (UAF) (Searcy et al. 1998), and still later with the support of the Air Weather Service Agency and NWS.

The PUFF model is designed to track particles from young ash clouds. It typically uses real-time wind forecasts from the GFS model, though it can use NAM output as well. It initializes an ash cloud characterized by a large number of particles and then calculates transport, turbulent dispersion (spreading of particles in the plume), and fallout. At the heart of this model is a Lagrangian calculation of the position vector $\boldsymbol{R}$ of particle i at time $t + \Delta t$, where t may be considered the current time and Δt a time increment, say, a few minutes. The calculation is written as

$$\boldsymbol{R}_i(t + \Delta t) = \boldsymbol{R}_i(t) + \boldsymbol{W}(t)\Delta t + \boldsymbol{Z}(t)\Delta t + \boldsymbol{S}_i(t)\Delta t\,.$$

$\boldsymbol{R}_i(t)$ is the position of the ith ash particle at time t, $\boldsymbol{W}(t)$ is the three-dimensional local wind velocity obtained from a model forecast, $\boldsymbol{Z}(t)$ is a vector representing turbulent dispersion, and $\boldsymbol{S}_i(t)$ is the terminal gravitational fallout vector. A new position is computed for each particle following each time increment.

The turbulent dispersion of the puff is simulated by a random walk in three dimensions. All random displacements follow a Gaussian distribution with zero mean. The standard deviation

of the horizontal displacements is considerably greater than the standard deviation of the vertical displacements. Properly specifying the standard deviations is important, because they determine how rapidly the plume will spread.

The default specification of particle fall speed s is given by Stoke's Law. (s is the vertical component of $\boldsymbol{S}_i$ and the only nonzero component.)

$$s = \frac{2(\rho_{ash} - \rho_a)gr_1^2}{9\eta}$$

ρ_{ash} is the density of the ash particle, ρ_a is the density of air, g is the acceleration due to gravity (9.8 m s^{-2}) , r_i is the diameter of the ith particle, and η is the dynamic viscosity of air (1.79 x 10^{-7} kg m^{-1} s^{-2}). The PUFF model lumps $\rho g/\eta$ into a single constant: 1.08 x 10^9 m^{-1} s^{-1}, which corresponds to an ash particle density of roughly 2000 kg m^{-3}. Air density is negligible compared to ash particle density.

The distribution of ash particle sizes must be specified at the outset. One available distribution is the Gaussian distribution of the logarithm of particle size, i.e. x = $\log_{10}(r_i)$. For example, to specify the mean particle size as 10^{-5} m, one would set the mean of x to -5. Similarly one could set the standard deviation of x to 2, implying that 68.3% of all particle radii would lie between 10^{-3} and 10^{-7} m. (The exponents arise from $10^{-5\pm2}$.)

Several other parameters must be specified before the PUFF model can run:

- the latitude and longitude of the volcano
- the total number of particles to be released (usually many thousands)
- the maximum plume height
- the number distribution of particles (how many per cubic meter) as a function of altitude
- the shape of the initial plume (e.g. spreading with height or not)

4.3.3.2 The HYbrid Single-Particle Lagrangian Integrated Trajectory (HYSPLIT) model

Roland R. Draxler and Albion D. Taylor of NOAA's Air Resources Lab published a paper (Draxler and Taylor, 1982) that would become the basis for the HYSPLIT model. Since then, with the help of the Australian Bureau of Meteorology Research Center, it has grown through several versions and become widely available, nationally and internationally. The best references for the current version are Draxler and Hess (1998a, b). As noted earlier, ARL runs a version of HYSPLIT routinely for several volcanoes ringing the North Pacific Ocean. If one of them erupts,

simple trajectory predictions are immediately available, but then NCEP/EMC would begin running more detailed calculations that include plume dispersion.
HYSPLIT incorporates more sophisticated physics than the PUFF model, appropriate for many applications besides volcanic eruptions, for example, the release of toxic substances in the boundary layer. The capabilities used to calculate volcanic plume dispersion are described here.

To start a dedicated run of HYSPLIT, the user specifies the latitude and longitude of the volcano, the start time and duration of the eruption, the summit elevation of the volcano (the bottom of the ash cloud), and the height of the cloud. HYSPLIT uniformly distributes the ash between these two heights, in effect, creating a vertical line source of ash. The user has the option to specify the mass of pollutants released in each of four particle-size bins covering the range from 0.3 to 30.0 µm, for example:

Bin Number	Middle Value Particle Diameter	Percent by Mass
1	0.6 µm	1%
2	2.0 µm	7%
3	6.0 µm	25%
4	20.0 µm	67%

The density of all particles is assumed to be 2,500 kg m^{-3}, typical for a mix of pumice, basalt, and glass. If the rate of mass injection into the atmosphere can be estimated, then values in the table permit calculation of the number of particles in each bin. If no estimate is possible, the model will assume a unit mass injection rate, and all particle concentrations will be calculated relative to that rate. 99% of the mass normally resides in particles greater than 3 µm in diameter. The largest fall out first, while the smallest may be carried considerable distances.

HYSPLIT can be run in any of three modes: a particle mode, a puff mode, and a hybrid mode. The NCEP/EMC version of HYSPLIT uses the particle mode. A large number of particles is released at altitudes reached by the ash cloud above the volcano. A three-dimensional wind field obtained from a numerical prediction model (GFS or NAM) advects (carries) each particle downstream. HYSPLIT uses horizontal and vertical mixing coefficients obtained from model data to calculate the dispersion of ash particles. In the vertical direction, the dispersion depends heavily upon the stability, also available from the model. After the model wind carries a particle for one time step, a random displacement, whose magnitude depends upon the mixing coefficients, is added to each particle. In this way, the particles spread apart with time,

as occurs in nature. The number density of particles at a given time is equivalent to the ash concentration, which is of great concern to pilots and airlines.

HYSPLIT also runs in *puff* mode. As used here, "puff" refers to a discrete cloud of particles whose centroid is tracked rather than the individual particles within it. This gives the puff mode a computational advantage. This puff should not be confused with the PUFF model described earlier, designed to track individual particles of ash released over a short time. The spatial distribution of particles within the puff is specified, and each puff contains a specified fraction of the total mass in the eruption. HYSPLIT tracks the centroid of the puff and computes its dispersion (expansion) based on turbulence information available from the numerical prediction model. The volume of the puff is related to the standard deviation of the difference between particle positions and the centroid position, in both horizontal and vertical directions. When puffs become larger than the volume of the grid boxes, they split, sharing the mass equally between them, and then two centroids are tracked. If splitting produces so many puffs that the computation of their motion and dispersion becomes too demanding computationally, the puffs can be merged into larger entities.

Finally, HYSPLIT can run in a hybrid mode, in which the horizontal dispersion is handled like a puff, and the vertical dispersion is computed as it would be for particles.

The NCEP/EMC version of HYSPLIT includes gravitational settling, which acts to deplete larger particles from the ash cloud faster than smaller particles. It does not yet include another HYSPLIT feature—wet deposition—whereby precipitation washes ash out of the plume. The model clouds and the model precipitation rate determine how much ash is removed. Only ash within or below the cloud is subject to removal in this way.

A large selection of experimental HYSPLIT products is available from the Air Resources Laboratory at http://ready.arl.noaa.gov/ready2-bin/ashhypo.pl. Two examples are shown here for a hypothetical eruption of Mt. Cleveland. Figure 4.11 shows the cloud of ash particles on a map, color-coded by altitude, 6 h (upper left), 12 h (upper right), and 18 h (lower left) after release at 1800 GMT 13 December 2011. The text box at lower right gives details on the hypothetical release. The release occurs from the volcano's summit up to 12,000-m altitude. In all, 20,000 identical particles are released, and they are uniformly distributed in the vertical. The lower portion of each panel shows the projection of all points onto a vertical plane defined by the dashed red line. The GFS model, initialized at 0600 GMT 13 December 2011, provided the data for computing the particle trajectories.

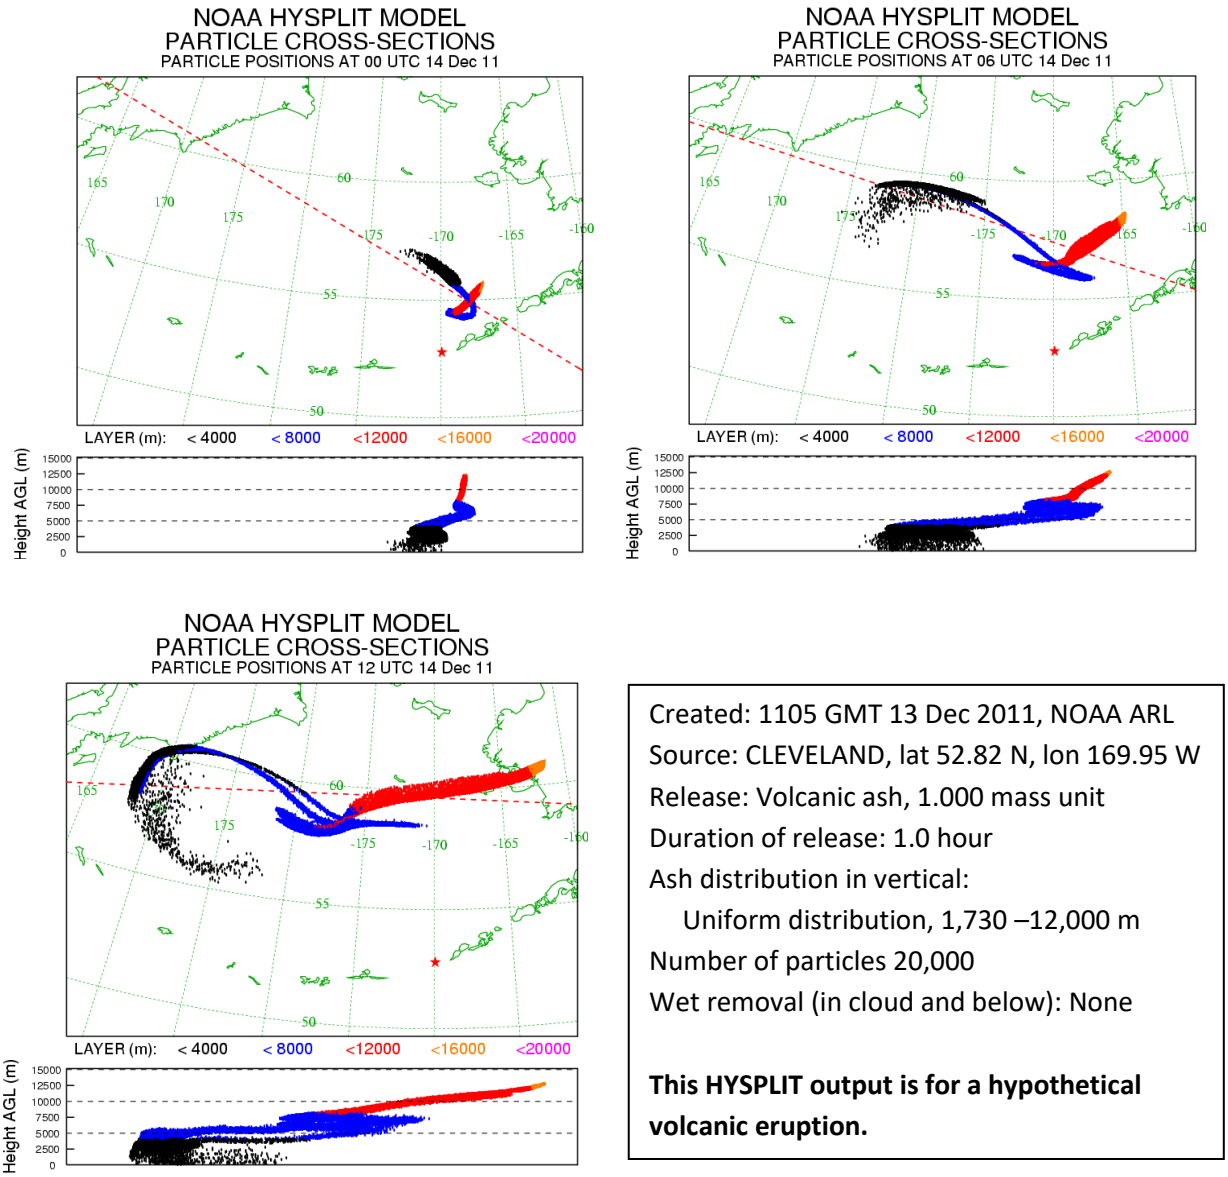

Fig. 4.11. A cloud of 20,000 particles is released for one hour, starting at 1800 GMT 13 Dec 2011 from the location of the red star. The GFS model initialized at 0600 GMT 13 Dec 2011 provided the data for computing the HYSPLIT trajectories. The panels show the cloud of particles 6 h (upper left), 12 h (upper right), and 18 h (lower left) after release. Other details of the release are in the box at lower right. The bottom of each panel shows the projection of all points onto a vertical cross section defined by the red dashed line. Source: NOAA Air Resources Laboratory.

The second example, Fig. 4.12, shows the concentration of ash particles 6,12, and 18 h after the release of a unit mass of ash, evenly distributed across 60,000 particles. The concentration is

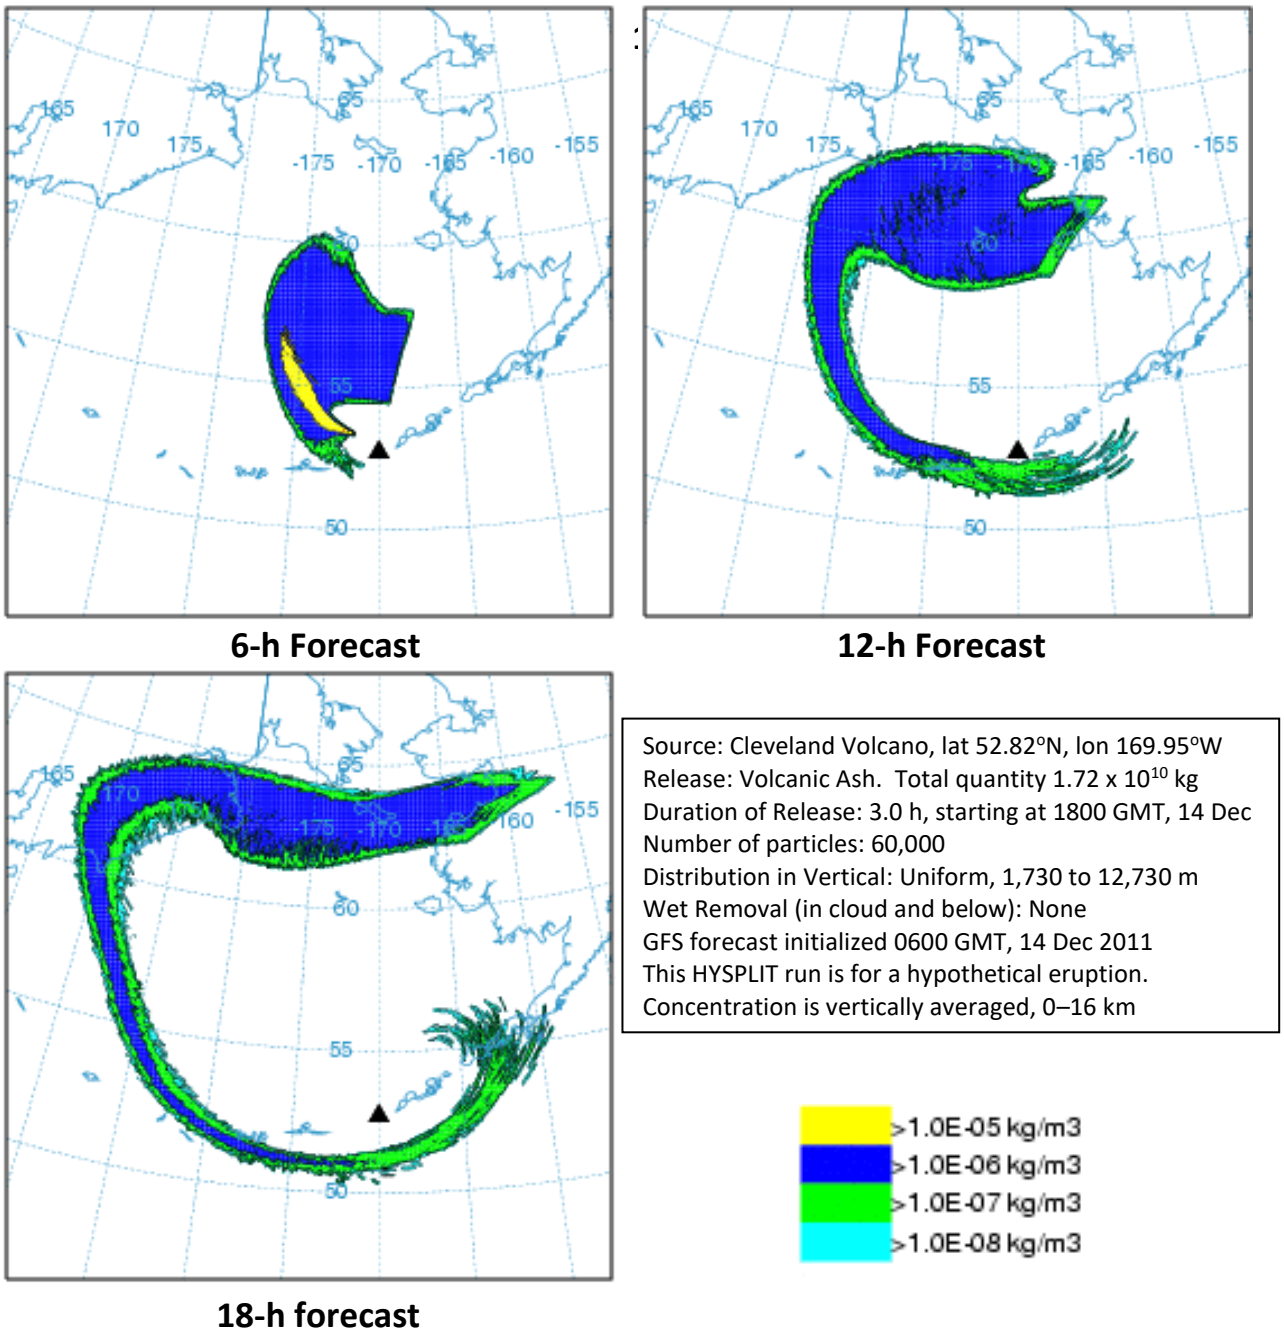

Fig. 4.12. HYSPLIT simulation of ash density 6, 12, and 18 h after a hypothetical release of a unit mass of ash at 1800 GMT 14 Dec 2011 from the Cleveland Volcano (black triangle). Details of the release are in the text box, lower right. This release is one day later than the one in Fig. 4.11. The color-coded ash density (lower right) is averaged through a vertical column of

atmosphere, 16 km high. Ash concentrations at each time step in the hour preceding the valid time of each display are averaged to produce a smoother image. Source: NOAA ARL.

expressed in units of kilograms per cubic meter, vertically averaged between sea level and 16 km, and is color-coded as at lower right. The ash density is so small numerically because the total mass released (1.72×10^{10} kg) is shared among 60,000 particles, and, after many hours, those particles are widely dispersed, horizontally over hundreds of thousands of square kilometers, and vertically over more than 12 km.

A major challenge for initializing trajectory and dispersion models is the lack of detailed information on the eruption itself, namely, the precise time of eruption, its duration, the total mass of ejecta, the maximum height reached by the eruption plume, and the size distribution of the ash particles. Another challenge arises when one tries to verify a prediction of the plume location and the particle concentrations within it, especially days after the eruption, when the plume has become very diffuse and tenuous. Given the remote locations of most of the world's volcanoes, remote sensing techniques by satellite hold the greatest promise for addressing these challenges.

4.3.4 Volcano Surveillance by Satellite

4.3.4.1 Detecting and characterizing volcanic ash

The focus of this section is on volcanic ash products developed at the University of Wisconsin by NOAA/NESDIS/Center for SaTellite Applications and Research (STAR) and CIMSS over a period of more than 15 years. The products will eventually depend upon observations from the GOES-R satellite, yet to be launched. They are generated now using similar data from the MODIS instrument, flying on NASA satellites in low earth orbit, as described earlier in Section 2.1.6.1. The products are being used experimentally by the Alaska Region of the NWS.

The scientific basis of the GOES-R volcanic ash products is described in Pavolonis and Sieglaff (2010), a lengthy document replete with technical and mathematical details. Only a brief sketch of this work can be presented here. The algorithm has two main sections: ash detection and the derivation of ash plume properties: height of the ash cloud, ash loading, and effective radius of ash particles.

The input data for the GOES-R volcanic ash product will ultimately come from the Advanced Baseline Imager. For now similar data are available from the NASA MODIS instrument:

- Calibrated radiances from MODIS channels having central wavelengths of 7.32, 8.55, 11.03, 12.02, and 13.3 μm,
- Calibrated brightness temperatures from MODIS for the channels with central wavelengths at 11.03, 12.02, and 13.3 μm, and
- Local zenith angle of the satellite.

Supporting data of the following types are also needed:

- Land surface classification in 14 classes (from http://glcf.umiacs.umd.edu/data/landcover/),
- Global surface emissivity at 11 and 12 μm. Monthly mean values are available from MODIS.
- Current profiles of temperature and moisture as a function of pressure from global prediction models (GFS). The tropopause level is deduced from the temperature profile.

Finally, the algorithm requires the use of two radiative transfer models:

- Model of "black"-cloud radiance at 7.32, 8.55, 11.03, 12.02, and 13.3 μm. More specifically, the model computes the radiance emitted upward by a black body surface, transmitted through a non-cloudy atmosphere with gaseous absorption, to the top of the atmosphere as a function of the pressure level of the black-body surface.
- Model to compute the clear-sky radiance at each of the above wavelengths from vertical profiles of temperature and moisture obtained from an NWP model.

At the heart of ash detection is a diagram similar to the one in Fig. 2.106, except that the scale is expanded along the horizontal axis. Effective absorption optical depth ratios (the β-ratios described in Section 2.3.12.2) are plotted along both axes in Fig. 4.13. The plotted information is the same as in Fig. 2.106 except that regions of the figure are shaded according to the likelihood of ash detection. The curve for dust is not shown. Dust cannot be distinguished from ash. Fortunately, dust raised from the ground by the wind and transported long distances is not usually mixed with an ash cloud.

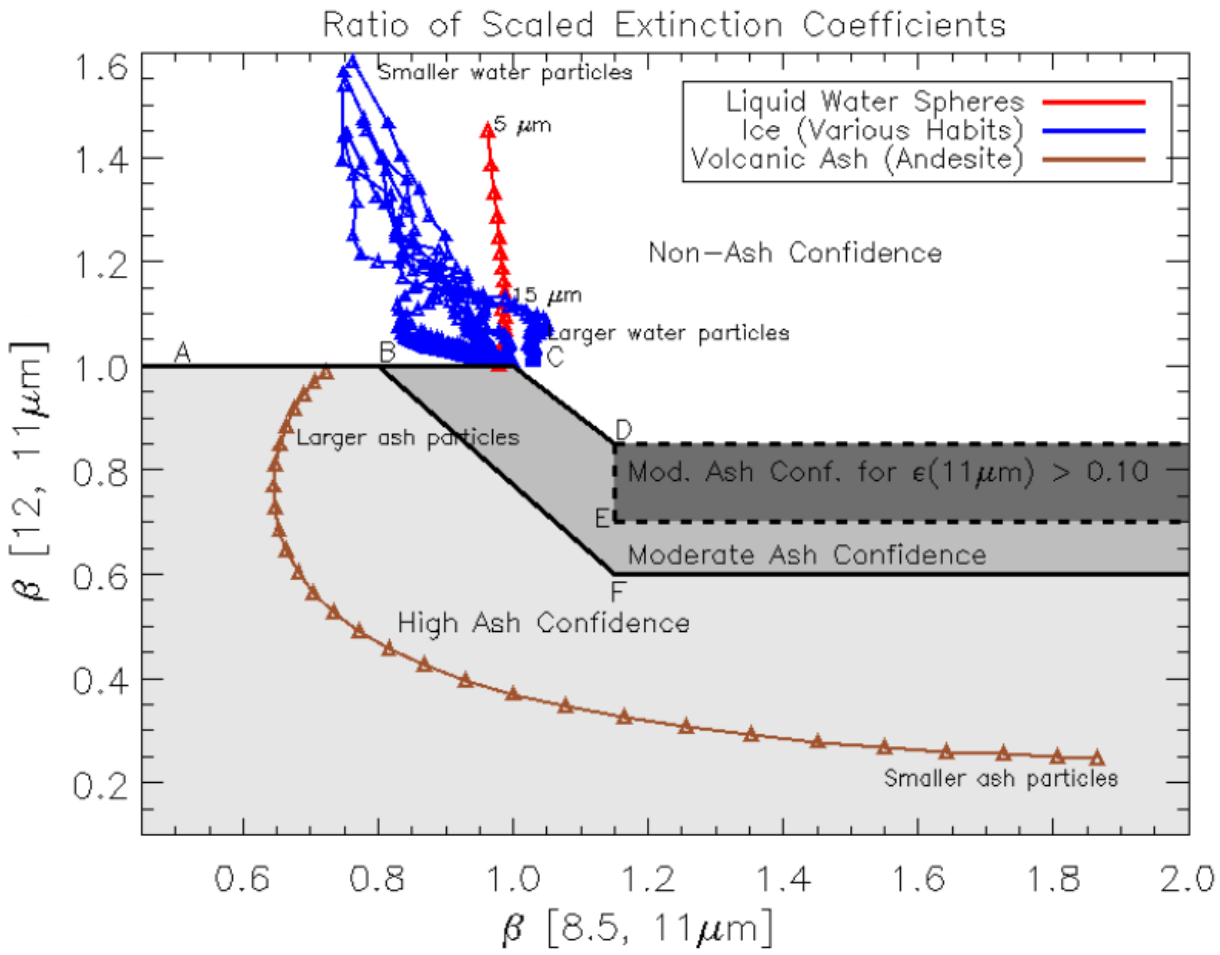

Fig. 4.13. Diagram for determining confidence in ash detection. Values along both axes are effective absorption optical depth ratios, involving the emissivity of the first-listed wavelength in the numerator, and the emissivity at 11 μm in the denominator. Refer back to Section 2.3.12.2 for details. In this beta space, the confidence (or lack of it) in ash cloud detection is shown in shades of gray. The brown curve is for volcanic ash, the blue curves are for various ice crystal habits, and the red curve is for liquid droplets. Note the gradation in size from one end to the other of the red and brown curves.

Once an ash cloud has been identified, the algorithm retrieves the ash cloud effective temperature T_{eff}, its effective emissivity ε at 11 μm , and a new estimate of $\beta(12, 11\ \mu\mathrm{m})$ specific to the ash cloud. The input data are the brightness temperature BT(11 μm), the difference in brightness temperatures BT(11 μm) – BT(12 μm), and another difference in brightness temperatures BT(11 μm) – BT(13.3 μm). The solution for all three unknowns is

obtained simultaneously by means of statistical optimal estimation. The solution involves the two radiative transfer models mentioned at the beginning of this section.

The three parameters displayed by the GOES-R volcanic ash algorithm are computed from the retrieved values of ash cloud effective temperature, its emissivity at 11 μm, and $\beta(12, 11\ \mu\mathrm{m})$.

The height of the ash cloud is estimated from its effective temperature by searching a temperature profile at the volcano location provided by an NWP model. Starting from the top of the profile, one looks for the first layer that contains the effective temperature. The height of the top and bottom of this layer are known by the model. Linear interpolation between these two heights produces the height of the ash cloud.

The computation of ash mass loading (metric tons per square kilometer) and effective particle radius is more complicated. It proceeds in a number of steps as described in Pavolonis and Sieglaff (2010, pp. 52-53).

1) The effective optical depth $\tau(11\ \mu\mathrm{m})$ of the ash cloud is computed from $\varepsilon(11\ \mu\mathrm{m})$ and the local zenith angle of the satellite.
2) The extinction cross section $\sigma_{ext}(11\ \mu\mathrm{m})$ and the effective particle radius r_{eff} are computed from regression equations involving $\beta(12, 11\ \mu\mathrm{m})$. The latter quantity contains most of the microphysical information about the ash cloud.
3) The distribution of ash particle size is assumed to be log-normal. With increasing values of particle size, this distribution rises sharply to a peak value, then decreases more slowly. The value of r_{eff} and a width parameter determine the exact shape of this distribution. The number of particles in each size bin, which this distribution supplies, is needed to calculate the mass of ash per unit area.
4) The total number of particles per unit area N_0 is determined from $\tau(11\ \mu\mathrm{m})$ and $\sigma_{ext}(11\ \mu\mathrm{m})$.
5) The mass loading per unit area is computed essentially by integrating the log-normal distribution. N_0 (preceding item) and the density of the ash are needed for this calculation. Pavolonis and Sieglaff use 2,600 kg m^{-3} for the density.

The GOES-R ash cloud algorithm, just described, generates four basic products. The ash detection product is illustrated in Fig. 4.14 This false-color image gives the location of an ash cloud spewed out by the Kasatochi Volcano on the Aleutian Island chain, 7 August 2008. The yellow color indicates the presence of ash with certainty. The pink and magenta colors indicate high probability of an ash cloud, but less certainty. The enclosed, irregularly shaped brown area

northwest of picture center indicates ice cloud either mixed with or obscuring the ash cloud. Determination of ash cloud properties is not possible within this area.

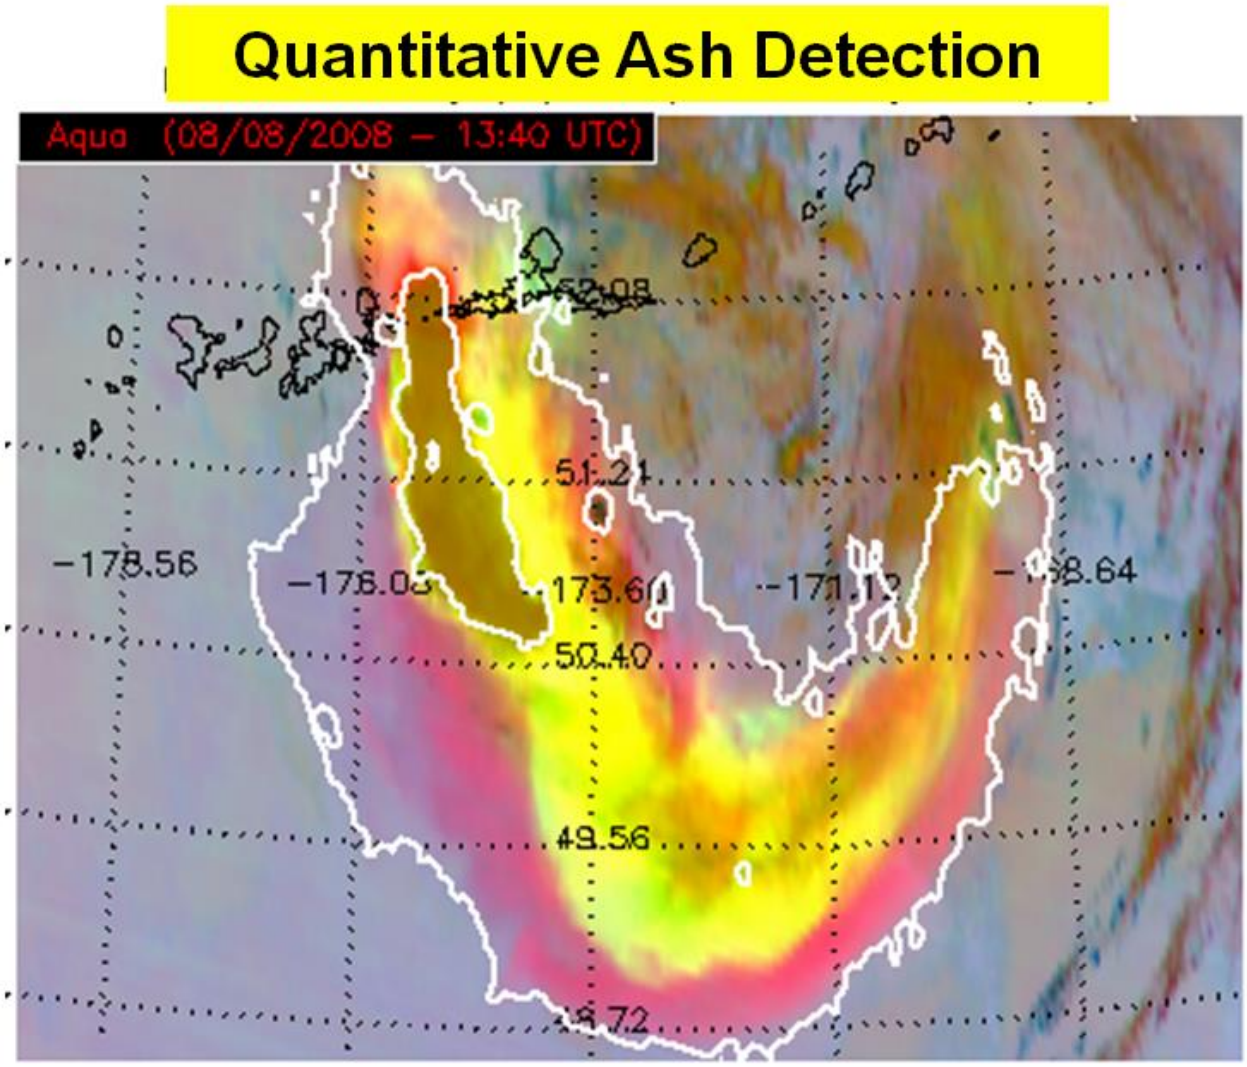

Fig. 4.14. A false-color image of an ash cloud released by the Kasatochi Volcano on one of the Aleutian Islands. The image is valid at 1340 GMT, 8 August 2008. Yellow, pink, and magenta colors indicate high probability of ash. The enclosed brown area, near the Aleutian Islands and surrounded by ash, indicates an ice cloud, presumably mixed with the ash. Retrieval of ash cloud properties is not possible within the brown area. Courtesy of Mike Pavolonis and Justin Sieglaff.

The Anchorage VAAC receives the next three products on AWIPS whenever there is an eruption in their area of responsibility. Figure 4.15 illustrates the amount of ash (metric tons per square kilometer) that would be deposited on the ground if all of it fell out instantaneously. This

product gives pilots a good idea of where the ash cloud is dense and where it is more tenuous. Current practice is to avoid even very low concentrations of ash.

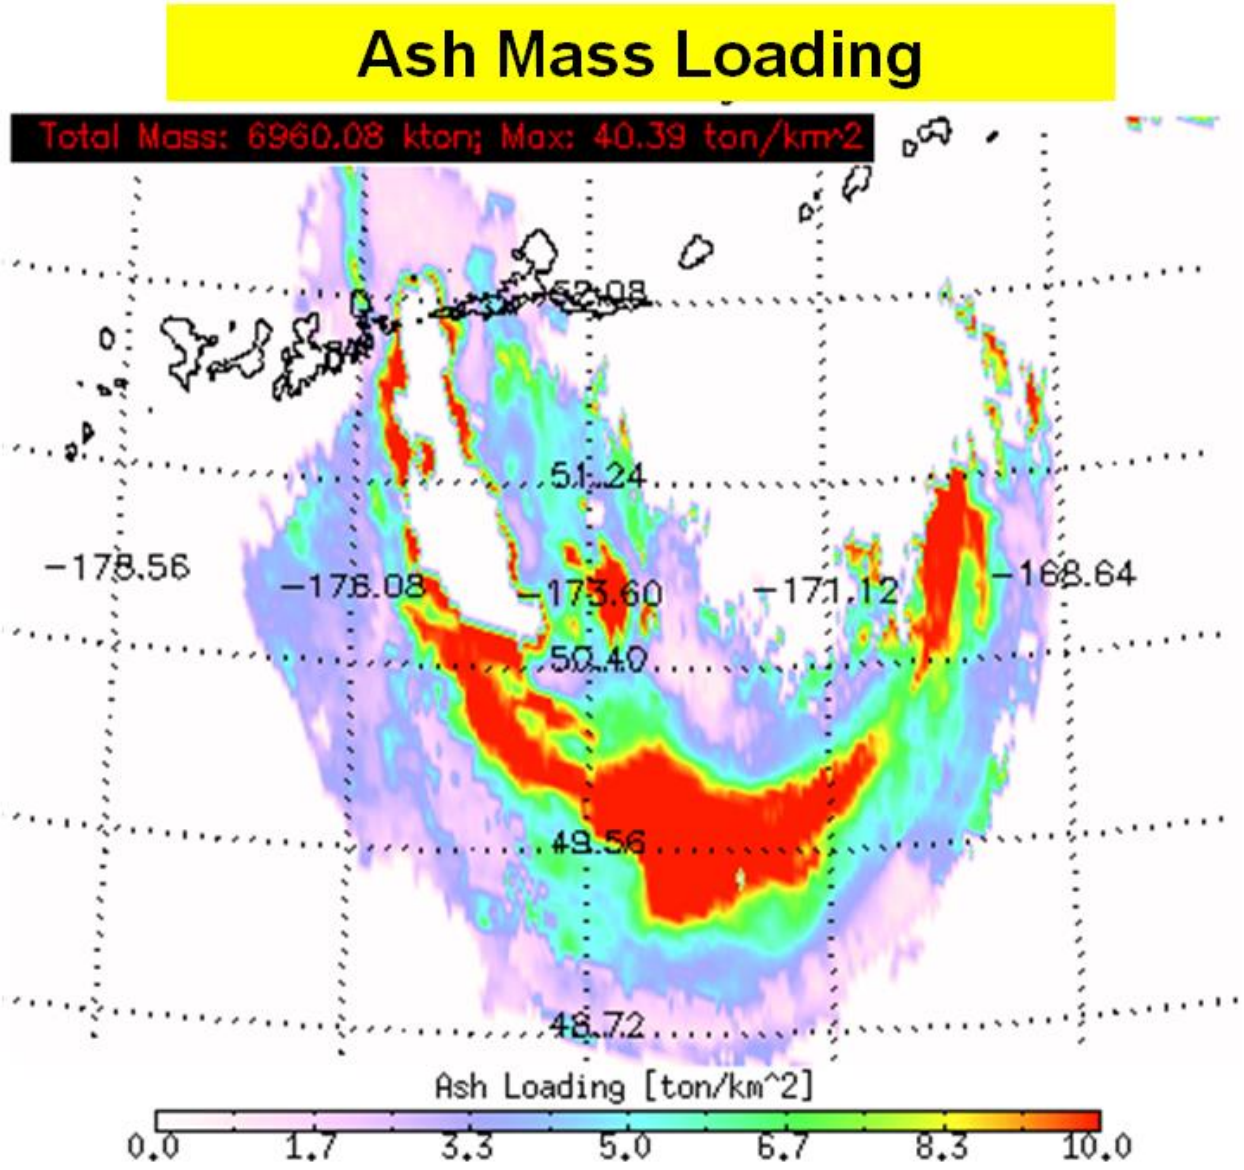

Fig. 4.15. The mass of ash that would accumulate on the ground if all of it fell out of the cloud instantaneously. The color bar gives the scale in metric tons per square kilometer. The estimated total mass of ash in the cloud is 6.96 x 10^6 kg. For Kasatochi Volcano, 1340 GMT, 8 August 2008. Courtesy of Mike Pavolonis and Justin Sieglaff.

Figure 4.16 shows the estimated height of the ash cloud in kilometers (see color bar at bottom). En route flights would avoid any area shaded in blue.

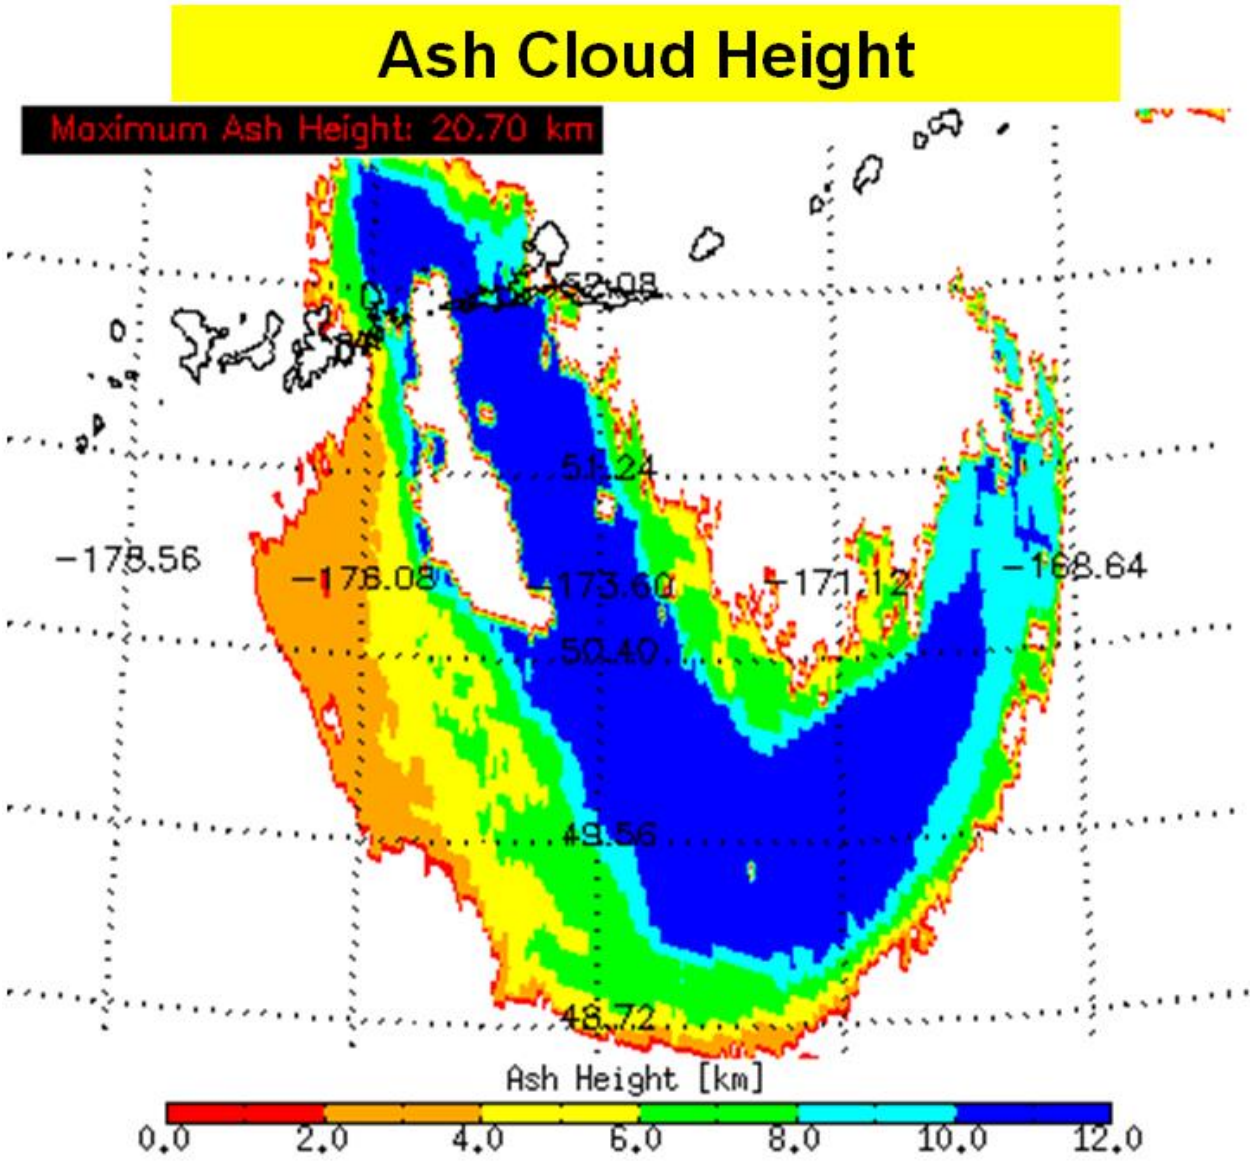

Fig. 4.16. Estimated height of the ash cloud from Kasatochi Volcano, 1340 GMT, 8 August 2008. The color bar at bottom gives cloud height in kilometers. Courtesy of Mike Pavolonis and Justin Sieglaff.

Figure 4.17 gives the ash effective radius of particles in the ash cloud from the Kasatochi eruption. The eruption is less than 24 h old and so many ash particles larger than 6 µm remain suspended. Because larger particles settle out of the cloud first, one would expect to find more yellows and oranges in an image produced one day later.

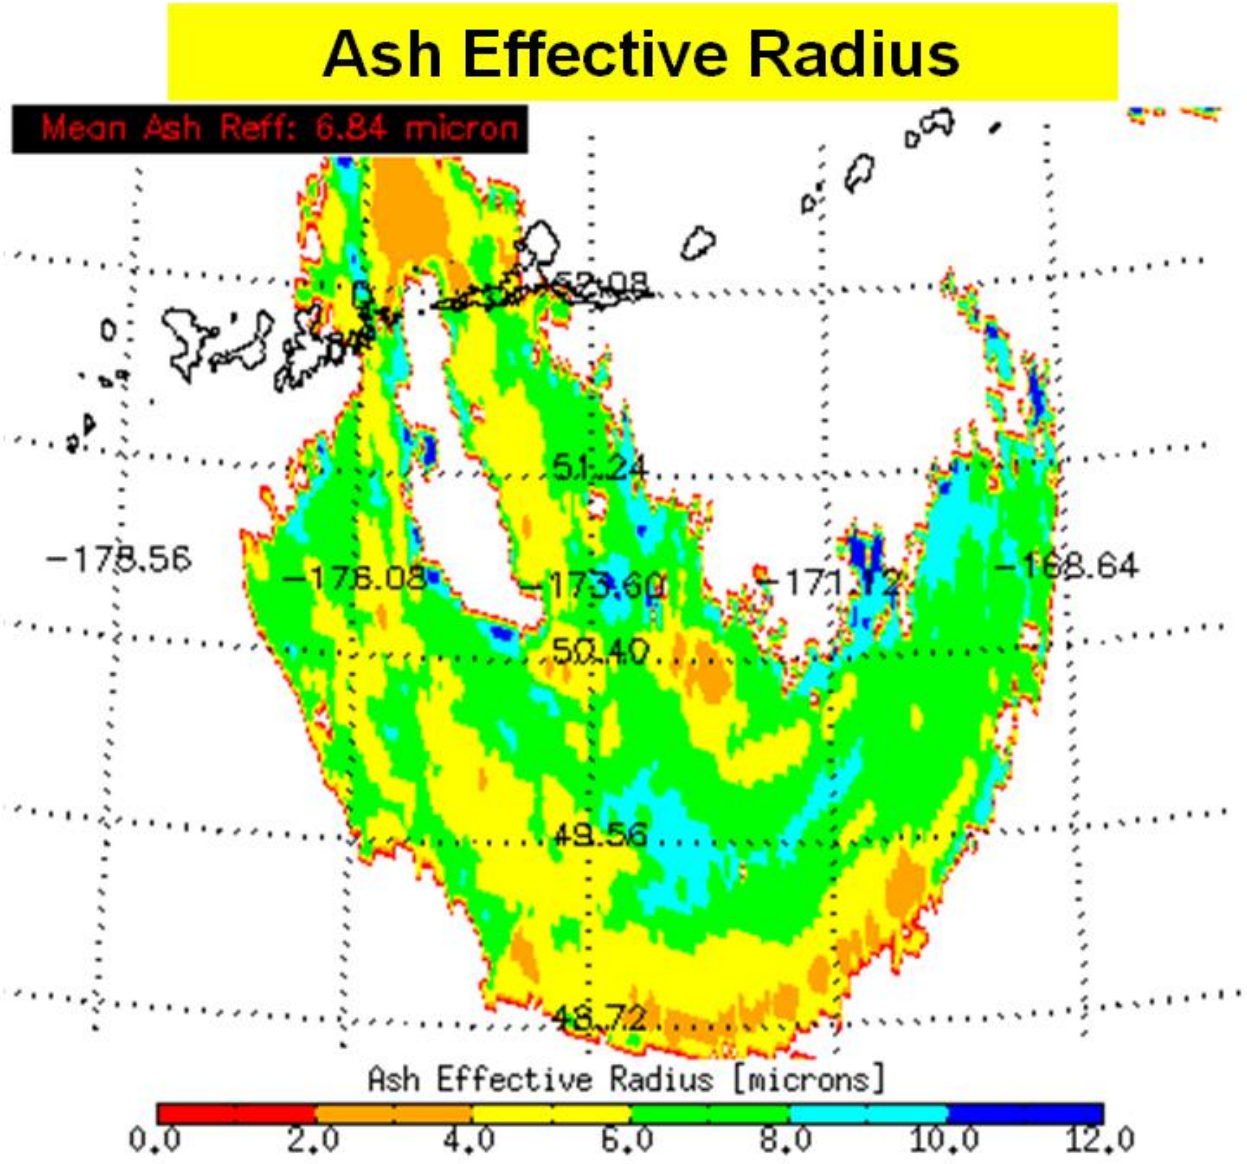

Fig. 4.17. The effective radius of ash particles in a plume produced by Kasatochi Volcano. Image valid at 1340 GMT, 8 August 2008. The color bar gives the effective radius in µm. Courtesy of Mike Pavolonis and Justin Sieglaff.

The previous four products are not yet available to the public, but an informative training package for their interpretation may be found at http://cimss/ssec/wisc/edu/goes_r/proving-ground/SPC/SPC.html. In the first table, find the entry under "Volcanic Ash" and "Training," and click on it.

Pavolonis and Sieglaff (2010) have assessed the accuracy of the GOES-R retrievals of ash cloud height, ash mass loading, and effective particle radius. They have obtained data from the Cloud-Aerosol Lidar with Orthogonal Polarization (CALIOP) instrument aboard the Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observation, one of five satellites included in NASA's A-Train, launched in April 2006. The lidar transmits pulses at two wavelengths, 532 and 1064 nanometers, in a pencil beam. It can make very detailed measurements of clouds, ash plumes, and other aerosol particles and is a good source for verification data.

Data from the CALIOP lidar were collocated with 3866 ash pixels, positively identified by the GOES-R algorithm. Histograms comparing the two estimates of ash cloud height, mass loading and effective particle radius are shown in Fig. 4.18. Note that the GOES-R algorithm estimates

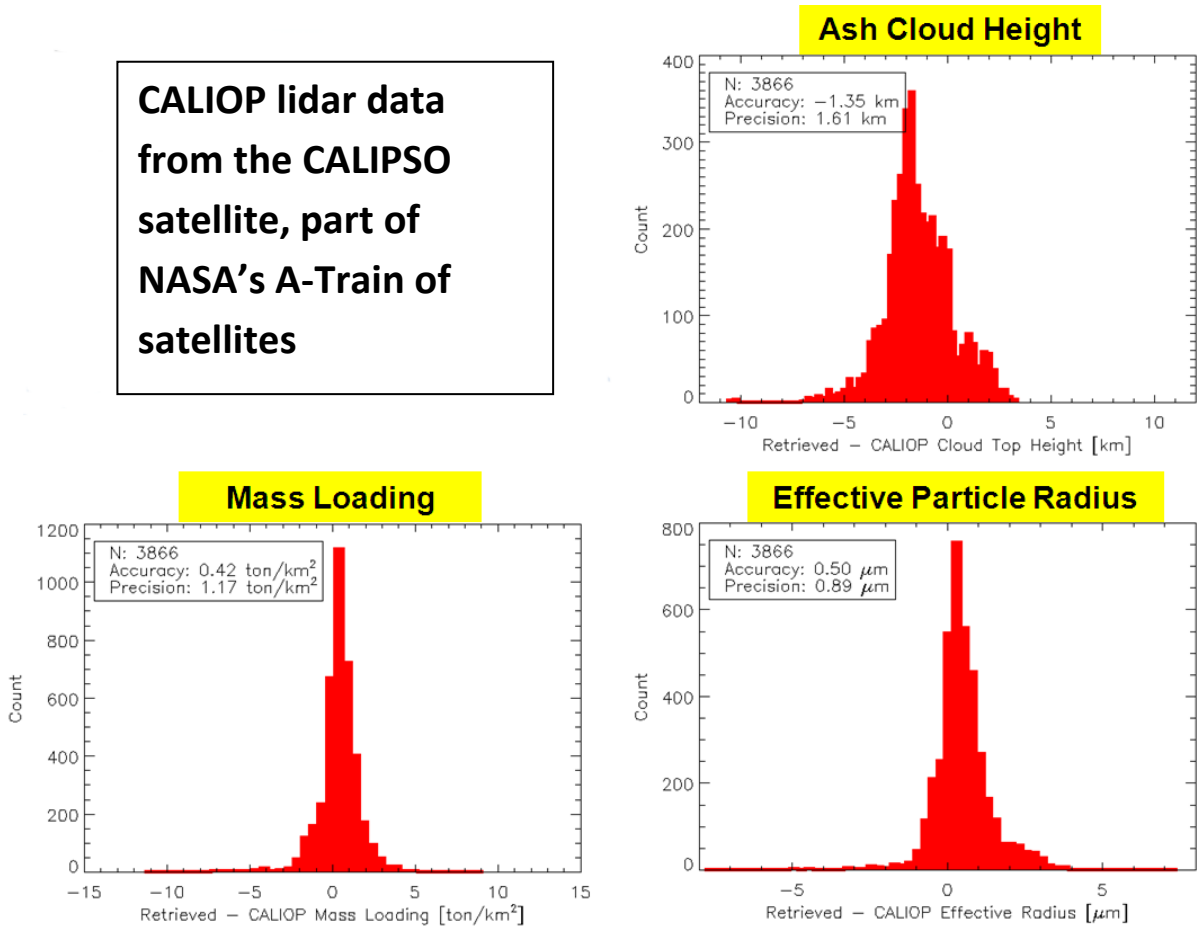

Fig. 4.18. A comparison of lidar and GOES-R algorithm estimates of ash cloud height, mass loading, and effective particle radius. The CALIOP lidar is the standard for accuracy. The quoted accuracy is the mean bias between the two estimates. The precision is the standard deviation of the difference. From Pavolonis and Sieglaff (2010).

of ash cloud height are, on average, 1.35 km less than the CALIOP measurement. The operational implication of this finding is that pilots should stay well above GOES-R estimates of ash cloud height, at least 1.35 km (4,500 ft). The bias in GOES-R estimates of mass loading is +0.42 metric tons per square kilometer (too high). If the ash cloud is assumed to be one kilometer thick, the mass loading in metric tons per square kilometer is numerically equivalent to the ash concentration in milligrams per cubic meter. It is important to make accurate estimates of ash loading because the threshold ash concentration that causes damage to jet engines is still unknown. Even very diffuse ash clouds are considered dangerous. The bias in estimates of effective particle radius is +0.50 μm (too high).

4.3.4.2 Detecting sulfur dioxide plumes

Sulfur dioxide, SO_2 is a trace atmospheric gas that enters the atmosphere because of human activity (fossil fuel combustion and biomass burning) and natural phenomena (oxidation of organic materials in soils and especially volcanic eruptions). SO_2 readily combines with water molecules to form sulfuric acid aerosols, which are corrosive. High concentrations of SO_2 are almost invariably associated with volcanic eruptions. SO_2 remains in the atmosphere longer than most ash particles; therefore its presence may indicate a remnant volcanic plume after the ash cloud itself becomes sub-visible.

Satellite detection of SO_2 occurs in the ultraviolet-visible (UV-VIS) portion of the spectrum and the infrared (IR) portion.

Instruments aboard three satellites are used to detect SO_2 at UV-VIS wavelengths. The SCanning Imaging Absorption spectroMeter for Atmospheric CartograpHY (SCIAMACHY) is an imaging spectrometer aboard the Envisat satellite, launched by the European Space Agency (ESA) in March 2002. It covers the spectrum from 240 to 1,700 nanometers (nm) at high resolution, ranging from 0.2 to 0.5 nm. The ground pixel size is 40 km by 80 km. The Global Ozone Monitoring Experiment, version 2 (GOME-2) sensor is on the first Metop Satellite, launched in October 2006. The Metop series of polar orbiting satellites is operated by the European Organization for the Exploitation of Meteorological Satellites (EUMETSAT). The GOME-2 instrument takes measurements from 240-790 nm with a spectral resolution of 0.24 to 0.53 nm. The ground pixel size is 40 km by 40 km. Finally, the Ozone Monitoring Instrument (OMI) was launched on NASA's Aura satellite, part of the A-Train, in July 2004. The OMI is a contribution to NASA's Aura satellite by the Netherlands Agency for Aerospace Programs and

the Finnish Meteorological Institute. OMI takes measurements from 270 to 380 nm with a spectral resolution of 0.45 to 1.00 nm. The ground pixel size is 13 km by 24 km.

SO_2 and ozone O_3 both absorb in the UV-VIS part of the spectrum. Sensors operating in this portion of the spectrum measure the solar energy scattered by the atmosphere and reflected at the ground, but absorption by SO_2 and O_3 alters the spectrum from what would be measured at the top of the atmosphere. In simplest terms, comparison of the measured spectra with the "pure" spectrum of sunlight allows estimation of the amount of SO_2 and O_3 along the slant paths taken by sunlight to the earth's surface and then to the satellite detector. The effects of SO_2 and O_3 can be separated, thus permitting an estimate of slant-path SO_2. The conversion from slant-path SO_2 to vertical-path SO_2 is accomplished by means of an *air mass factor*, a number that depends upon the length of the light path through the atmosphere, the satellite look angle, the albedo of the earth's surface, and the vertical distribution of the absorbing gases. Computation of the air mass factor is not trivial.

The slant-path SO_2 is called *slant column density* (SCD); the vertical-path SO_2 is called *vertical column density* (VCD). Both are expressed in terms of Dobson units (DU). A Dobson unit was originally used to express the amount of O_3 in a vertical column, but it can be applied to any atmospheric gas. One Dobson unit refers to the amount of pure gas that would be 10 µm thick under conditions of standard temperature (0^oC) and pressure (1013 mb).

The Support to Aviation Control Service (SACS), headquartered in Belgium, supplies VAACs around the world with SO_2 alert products. The SACS website is at http://sacs.aeronomie.be/. Click on "products" to find extensive documentation on the alert products, including the science behind them. Another good reference is van Geffen et al. 2009. An alert is issued if a ground pixel has a VCD value of more than 2 DU *and* at least five of the eight surrounding pixels also have values this large. Click on "Volcano Websites" for comprehensive information about all of the satellite sensors mentioned in this section.

Figure 4.19 is a sample SO_2 alert product issued by SACS. This alert involves SO_2 erupted by the Grimsvoetn Volcano in Iceland, starting on 21 May 2011. The image is valid at 1343 GMT on 23 May. 252 flights were cancelled when this plume covered Scotland and Northern Ireland. The text box at the top gives the average latitude and longitude of the SO_2 plume, the solar zenith angle (the smaller the angle, the more accurate the estimate), and both the VCD and SCD estimates of SO_2 in Dobson units. The VCD estimate was corrected for the presence of clouds.

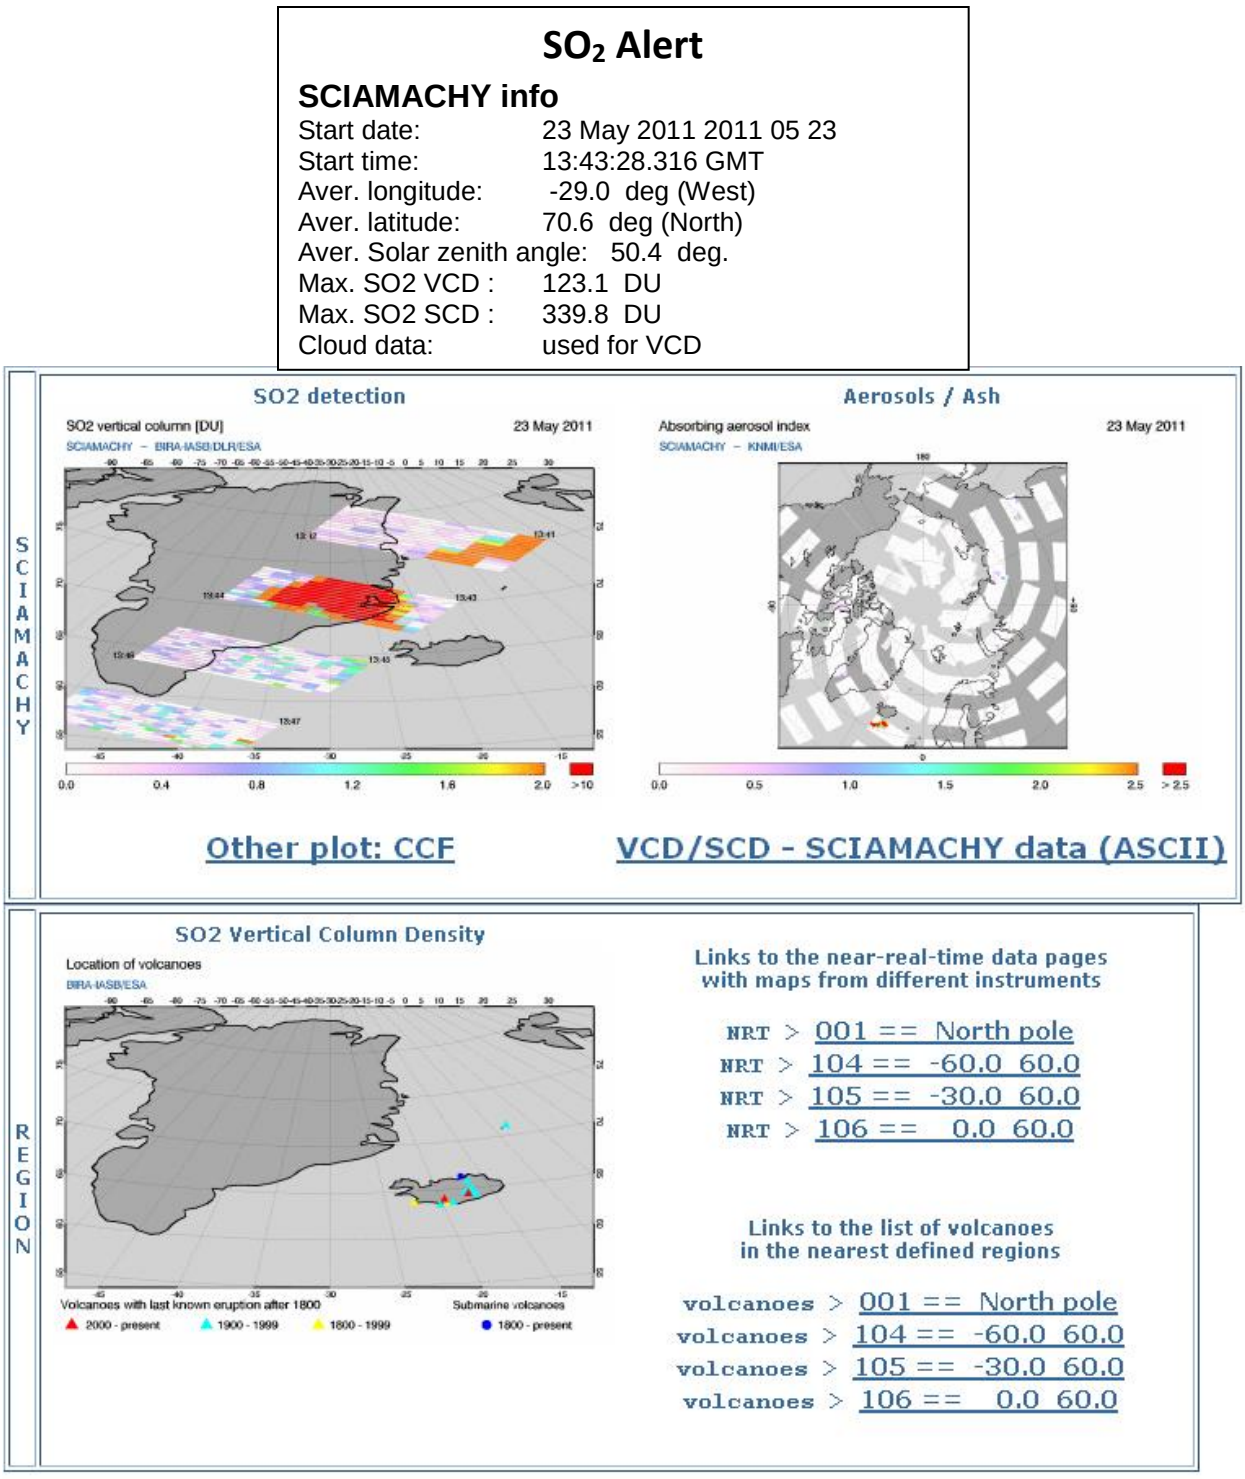

Fig. 4.19. An alert for exceptional SO_2 concentration, issued by SACS related to the 2011 eruption of the Grimsvoetn Volcano in Iceland. Box at top gives product details. Color bars in top two panels give Dobson units (DU). The red dots in the bottom left panel locate the

Grimsvoetn and Eyjafjallajökull Volcanoes. Other colors identify other volcanoes not recently active. Left panels give regional view; right panel gives high-latitude Northern Hemisphere view.

It is important to remember that estimates of SO_2 concentration utilizing UV-VIS spectral measurements are available *only during daytime*. For this reason researchers are investigating ways of detecting SO_2 at night with infrared sensors, two in particular: 1) The Infrared Atmospheric Sounding Interferometer (IASI) that flies on a Metop satellite, launched in October 2006 and 2) the Atmospheric InfraRed Sounder (AIRS), one of six instruments on board the Aqua satellite, launched in May, 2002. Like IASI, AIRS has the equivalent of thousands of channels; this endows both instruments with very high spectral resolution. The sensitivity of IASI and AIRS to SO_2 is fairly low, but, by differencing brightness temperatures at judiciously chosen wavelengths, it is possible to track SO_2 plumes for many days after an eruption and over long distances. Clarisse et al. (2008) describe an alert product based on IASI and implemeted at Université Libre de Bruxelles (ULB). Figure 4.20 is a sample of this product, issued for the eruption of the Grimsvoetn Volcano in Iceland, mentioned earlier. The figure and others like it from an archive are available at http://cpm-ws4.ulb.ac.be/Alerts/index.php.

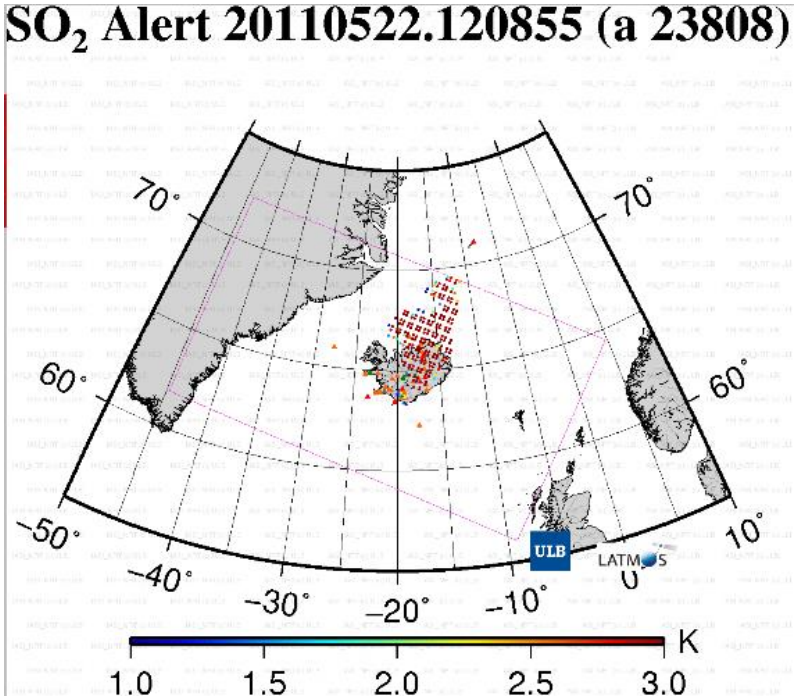

Fig. 4.20. An SO_2 alert, 1208 GMT 22 May 2011,based upon IASI data from the Metop satellite. Alerts like this are issued by the Université Libre de Bruxelles (ULB) in Belgium. The color bar at bottom indicates a brightness temperature difference. Red colors are a clear indication of SO_2.

4.3.5 Commentary on Volcanic Ash Products

The threshold of ash concentration below which air travel may be considered safe is unknown. A concentration of 2 x 10^{-6} kg m^{-3} has been suggested as this threshold. For a 1-km thick cloud, this is equivalent to 2.0 on the color-bar scale of Fig. 4.15. Many questions surround this proposed threshold. Will *any* commercial jet engine operate normally and indefinitely at this threshold? Presumably, the operating temperature of the jet engine is a factor because the hotter the engine, the more magmatic silicate glass melts and resolidifies internally. If the concentration threshold is exceeded, how long does the pilot have before an emergency develops? Do the outside air temperature and pressure affect the level of hazard? If this threshold is near the minimal level for satellite detection, the contour of threshold ash concentration may be ill defined. The meteorological satellite community may have to exploit further the spectral information already available. Failing that, it may need to improve the sensitivity of instruments that detect ash. In addition, the aircraft manufacturing industry may have to perform many engine inspections and tests before these questions can be answered satisfactorily.

A study of aircraft encounters with aircraft (Guffanti et al. 2010) resulted in several recommendations appropriate for inclusion here.

- About half the world's 550 historically active volcanoes now have real-time, ground-based geophysical monitoring. More extensive instrumentation would give volcano observatories such as AVO better awareness of precursory unrest and more timely eruption reports.
- Early information on the volume of erupted material and the height and density of the ash cloud is seldom available. Ground-based radars are seldom close to the eruptions. Airborne measurements performed remotely from outside the cloud and ground reports of ash fall would be valuable input to dispersion models. Scollo et al. (2010) have used data from the Multi-angle Imaging Spectroradiometer (MISR) on NASA's Terra Satellite to detect the height and shape of the eruption column with stereoscopic retrievals.
- Refining and testing communication protocols would make information flow more effective among the diverse parties who deal with ash hazards to aviation. Several problems prompt this recommendation. 1) VAACs sometimes conveyed different messages during the eruption of Eyjafjallajökull. This was a matter of international coordination. 2) Regional coordination can also be problematic. According to Christina Neal of the AVO, a frequent critique from users following an eruption is that they must go to USGS for some products and NWS for others. Centralization of all volcano-warning products for Alaska would help solve this problem. 3) USGS and NWS share responsibility for warning the public of volcano

hazards. The two agencies work hard to coordinate messages during events in order to avoid potentially dangerous confusion, but time pressures during an eruption, a barrage of requests for information, offices in different locations, and differing workloads can make this difficult. Also, the channels for disseminating information are different: NWS products are formalized (SIGMETs, VAAs, etc.): they go out under a WMO heading through the AAWU and Anchorage VAAC, and over NOAA Weather Radio and Emergency Medical Services Broadcasts. AVO information goes out mainly via email, fax, and web posting, which is not as efficient in reaching all the affected aviation users.

Finally, a word about numerical prediction models. Today's dispersion models take much of their input from EMC's operational prediction models, for example, the wind fields, vertical profiles of temperature and humidity, and turbulence parameters. On-line chemistry models are being tested in the U.S. and Europe, in which trace gases and aerosols, their chemical transformations wrought by exposure to sunlight or water, and their trajectories as affected by wind and turbulence are all explicitly predicted. The computational horsepower to do this is at last an order of magnitude greater than that required by today's operational prediction models, but the advantages of online versus offline chemistry are already apparent. In the future, volcanic ash dispersion is likely to be just another process incorporated in operational NWP models.

5. The Essential and Pervasive Role of Computer Prediction Models in the Generation of Aviation Weather Products

Figure 5.1 illustrates the role of numerical prediction models in the generation of aviation weather products. The information is embedded in a matrix. The columns of the matrix describe the progression of aviation weather information from researchers (who discover it) through operational meteorologists (who use it in numerical prediction and for advisories, watches, and warnings) to the end users (who require it tailored in the form of specific products). The rows of the matrix represent the chronological use of weather information, first in the form of raw observations or analyses (diagnosis) and then as input for numerical forecasts (prediction). The arrows represent the direction of information flow. Except for a few products that can be generated from observations alone, the entire body of aviation weather information depends upon gridded analyses and forecasts from computer models. The role of models is mentioned in every previous section of this report. Models and their associated data assimilation systems are an essential component of the aviation weather enterprise.

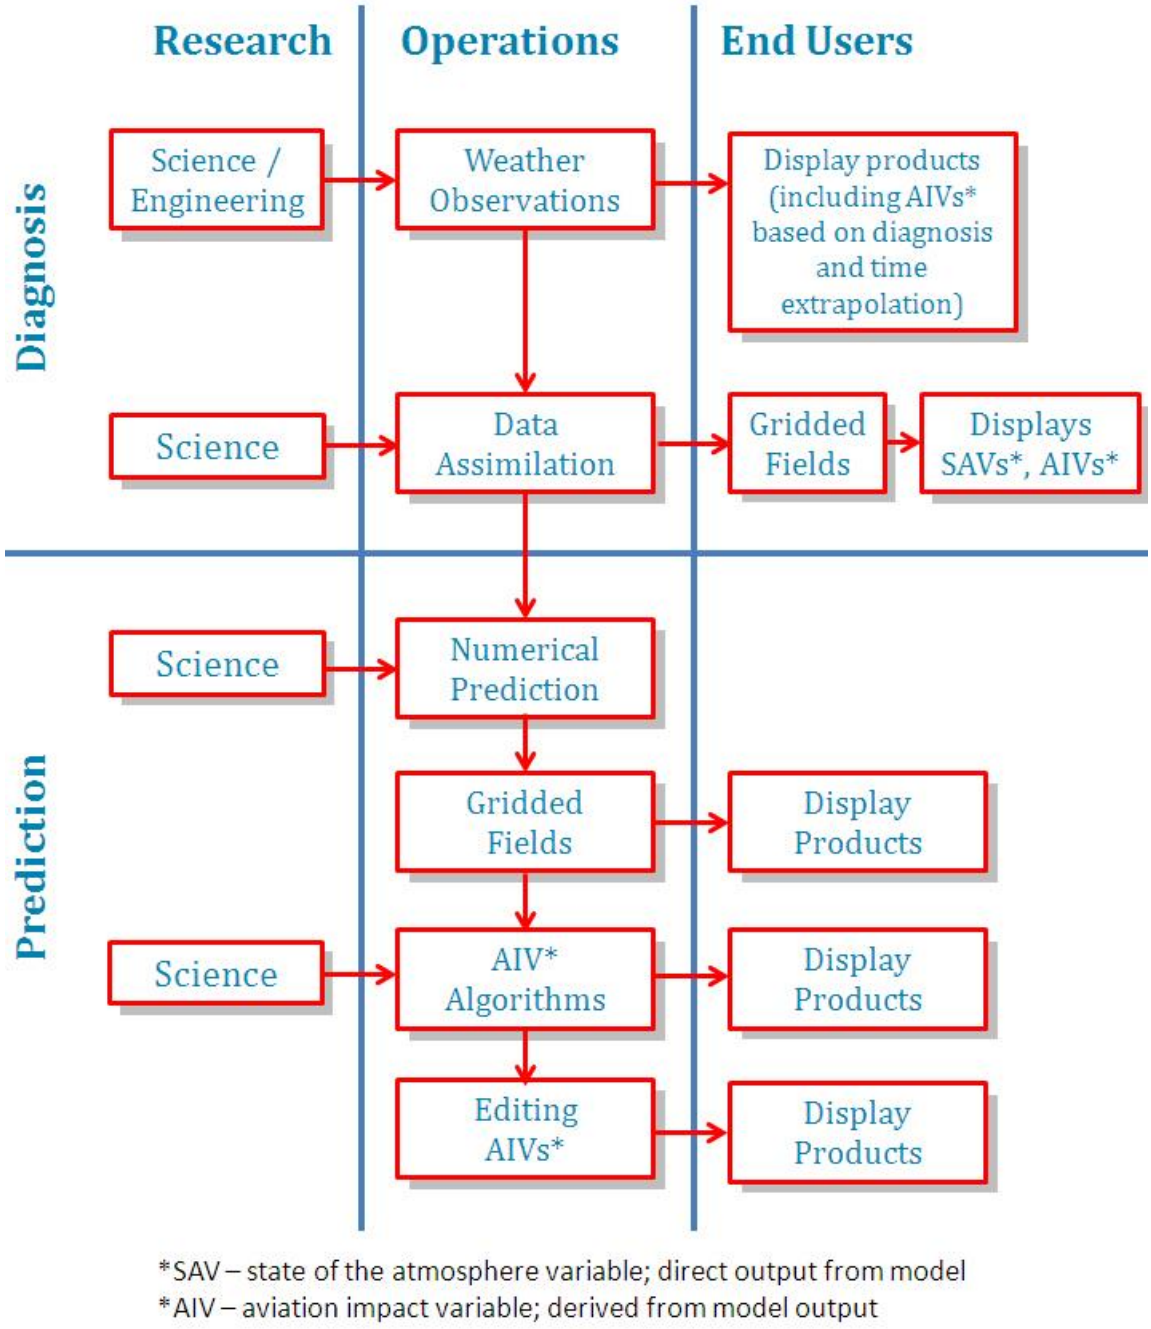

Fig. 5.1. The process of generating aviation weather products starts with meteorological observations. A few products relating to aviation weather hazards are based upon raw observations and occasionally simple time extrapolation. However, the great majority of products arise from data assimilation, which produces analyses of current conditions based upon all observational data sources; numerical prediction, which calculates future atmospheric states in the form of gridded fields; and algorithms for generating aviation-specific information for decision making (AIVs).

The remainder of this section will treat three major forecast systems: the Global Forecast System (GFS), the North American Mesoscale (NAM) Forecast System, and the Rapid Update Cycle (RUC) and its two progenies, the Rapid Refresh (RAP) and the High-Resolution Rapid Refresh (HRRR). All have been mentioned previously in this report. Atmospheric prediction models such as these are written in some of the largest and most computationally demanding codes in any scientific field. More staff years have been devoted to their development than for any algorithm earlier described in this report. The descriptions here can only scratch the surface. A word of caution: all these models are constantly evolving. Every three to five years, they change substantively, with incremental changes occurring more frequently. Therefore, what follows should be considered a snapshot in time as of late 2011.

5.1. Global Forecast System (GFS)

The GFS is the global spectral model running operationally at EMC. It supplies lateral boundary conditions to the NAM and RAP models, and its data assimilation system, Gridpoint Statistical Interpolation (GSI), has been adapted for regional application in those two models.

In a spectral model, the values of the prognostic variables are described globally in terms of expansions in spherical harmonics. The coefficients in these expansions carry the information about the variables while the spherical harmonic functions (the mathematical basis functions for the expansion) determine the spatial scales that can be resolved. The horizontal resolution of the model is determined by the triangular truncation of the basis functions, represented as T574, roughly equivalent to 27-km horizontal resolution at 35°N. There are 64 unequally spaced levels in the vertical, with closer spacing near the earth's surface. About 15 levels are below 800 hPa, 24 levels are above 100 hPa, with the top level at 0.3 hPa (about 57-km altitude in a standard atmosphere). Certain operations, such as the analysis of the initial atmospheric state, calculation of nonlinear quantities, and parameterization of physical processes are performed in grid space on a so-called Gaussian grid instead of in spectral space.

The prognostic variables are vorticity, divergence, the logarithm of surface pressure, specific humidity, virtual temperature, and cloud condensate. The model treats the following physical processes: horizontal and vertical diffusion; gravity-wave drag, longwave and shortwave radiation; deep and shallow convection; condensation, precipitation, and clouds; and heat, moisture, and momentum transfer in the planetary boundary layer and at the ground. The model accounts for land and sea surface properties, including surface characteristics (land, ocean, vegetation, snow or ice cover). More information and many references are available at http://www.emc.ncep.noaa.gov/GFS/doc.php .

The GSI provides an analysis of the initial atmospheric state, required to start the GFS model running. The GSI became operational in the GFS in May 2007, replacing an earlier spectral

analysis. The GSI is a three-dimensional variational (3DVAR) data assimilation system, suited to both global and regional applications. In its global application, it blends a great number and variety of atmospheric observations with a six-hour GFS forecast verifying at the same time as the analysis is desired. The "prior" information in the model forecast supplements the information in the observations, and the blending accounts for the expected errors in each source. The blending is subject to two constraints: 1) The water vapor mixing ratio must be greater than zero but less than the saturation value and 2) the dry mass of the global atmosphere must be conserved.

The GSI must take into consideration that the set of observed parameters does not, in general, match the set of prognostic variables in the model. The GSI employs *forward models* to transform the model variables into an observed variable and to interpolate from the model grid to the location of the observation, thus enabling a direct comparison between an observation and the model's best guess as to what the observed value should be.

John Derber of NCEP/EMC presented a very accessible guide to the GSI in June 2011, at the Developmental Test Center Community GSI Tutorial at NCAR in Boulder, Colorado. Go to http://www.dtcernter.org/com-GSI/users/docs/tutorial_presentation_2011.php and click on "Overview of GSI" by John Derber.

This section concludes with a list of the observations processed by the GSI, courtesy of John Derber. First are global satellite data, Table 5.1, thinned to roughly 145-km spacing to keep the volume of data manageable.

Table 5.1

Global satellite data processed by the Grid-Point Statistical Interpolation (GSI) analysis. See notes at the end that explain satellite instruments not previously defined.

GOES-13 and GOES-15 Sounders

Channels 1-15

Individual fields of view

Four detectors treated separately

Used over the ocean only, in part, because the underlying surface there is radiatively simple

AMSU-A

NOAA-15	Channels 1-10, 12-13, 15
NOAA-18	Channels 1-8, 10-13, 15
NOAA-19	Channels 1-7, 9-13, 15
EUMETSAT Metop	Channels 1-6, 8-13, 15
NASA AQUA	Channels 6, 8-13

AMSU-B/MHS	
NOAA 15	Channels 1-3, 5
NOAA-18	Channels 1-5
EUMETSAT Metop	Channels 1-5
HIRS	
NOAA-17	Channels 2-15
NOAA-19	Channels 2-15
EUMETSAT Metop	Channels 2-15
AIRS	
NASA AQUA	148 channels
IASI	
EUMETSAT Metop	165 channels

Notes for Table 5.1: AMSU stands for Advanced Microwave Sounding Unit. They have flown on the polar orbiting satellites listed above, beginning in 1998. Almost from the beginning, observations from AMSU led to improved global predictions. From the radiance data AMSU provides, temperature and water vapor profiles, snow and ice coverage, cloud liquid water, and rain rates may be inferred. AMSU-A has 15 channels between 23.8 and 89.0 GHz, used primarily for temperature sounding. It has a ground resolution of 45 km at nadir. AMSU-B has five channels between 89.0 and 183.3 GHz, used primarily for moisture sounding. It has a ground resolution of 15 km at nadir. MHS stands for the Microwave Humidity Sounder, an instrument similar to AMSU-B, which flew on NOAA-18 and Metop.

The second class of observations, often called *conventional*, includes just about everything that is *not* a radiance observation. New acronyms are defined within the list.

- Radiosondes
- Winds from pilot balloons
- Conventional aircraft reports
- Automated aircraft reports
  AMDAR – Aircraft Meteorological DAta Relay: international collection of automated reports from commercial aircraft
  MDCRS – Meteorological Data Collection and Reporting System, the U.S. contribution to AMDAR. MDCRS is operated by Aeronautical Radio, Incorporated (ARINC)
- MODIS infrared and water vapor winds (tracking the motion of high-latitude features in infrared cloud images and water vapor images to infer wind speed and direction)

- GMS, METEOSAT, and GOES cloud drift infrared and visible winds (feature tracking to infer wind vectors). GMS – Geosynchronous Meteorological Satellite operated by the Japanese Meteorological Agency (JMA). METEOSAT – the European geosynchronous satellite, operated by EUMETSAT.
- GOES water vapor cloud drift winds
- Land observations of surface pressure
- Surface ship and buoy observations of temperature, dewpoint, and surface pressure
- SSM/I wind speeds. The Special Sensor Microwave/Imager is a passive microwave radiometer, measuring at four frequencies and flown on the U.S. Air Force's Defense Meteorological Satellite Program (DMSP) satellites. From measurements of the ocean surface, it can infer wind speed.
- ASCAT wind speed and direction. ASCAT – Advanced SCATterometer that flies on the European Metop satellite. Measurements of the ocean surface permit inference of wind speed and direction.
- SSM/I and TRMM TMI. TRMM refers to the Tropical Rainfall Measuring Mission, a satellite launched jointly by NASA and the Japan Aerospace Exploration Agency in November 1997. One of the instruments aboard this satellite is the TMI, the TRMM Microwave Image. From TMI measurements at five different frequencies, one can estimate precipitation.
- Doppler radial velocities from WSR-88D radars (used in the form of VAD profiles)
- Global Positioning System (GPS): bending angle profiles or refractivity profiles observed during radio occultation
- SBUV ozone profiles and OMI total ozone. SBUV – Solar Backscatter Ultraviolet Radiometer, an instrument aboard NOAA satellites that measures radiation in channels from 252 to 340 nm that are sensitive to the presence of ozone.

A main advantage of GSI, and the reason why a regional adaptation of the GSI is used to initialize the NAM and RAP regional models, is that it includes virtually all observation sources that demonstrably improve global forecasts.

5.2 North American Mesoscale (NAM) Forecast System

The operational NAM Forecast system underwent a major upgrade in October 2011. The upgraded version goes by the name Nonhydrostatic Multiscale Model, B-grid (NMM-B) and is described here. Detailed documentation should be available by spring 2012 (Janjic and Gall 2012).

The model has 60 vertical levels, defined by hybrid sigma-pressure coordinates (Eckermann 2009). Such coordinates are dimensionless and run from one at the Earth's surface to zero at the top of the model. They are the same vertical coordinates η as used in the GFS, defined by

$$\eta = \frac{A(\tilde{\eta}) - p_{top}}{p_s - p_{top}} + B(\tilde{\eta}) \, .$$

p_s is the surface pressure; p_{top} is the pressure at the top of the model. $\tilde{\eta}$ is defined by

$$\tilde{\eta} = \frac{p - p_{top}}{p_0 - p_{top}} ,$$

where p_0 is some pressure near sea level, nominally, 1000 mb. $A(\tilde{\eta})$ and $B(\tilde{\eta})$ must satisfy the following properties:

$$p_{top} \leq A(\tilde{\eta}) \leq p$$
$$0 \leq B(\tilde{\eta}) \leq 1$$
$$\frac{dB}{d\tilde{\eta}} \geq 0$$

Finally, $A(\tilde{\eta})$ and $B(\tilde{\eta})$ are related by

$$B(\tilde{\eta}) = \frac{p - A(\tilde{\eta})}{p_s - p_{top}} .$$

With these prescriptions, the vertical coordinate η approximates a terrain-following coordinate close to the ground and becomes a pressure coordinate high in the atmosphere (in the new NAM, at pressures below about 300 mb; in the GFS at pressures below about 70 mb).

The vertical distribution of the 60 η layers is shown in Fig. 5.2. The old distribution (before the model upgrade in October 2011) is shown in blue. The new distribution is shown in pink. Along the vertical axis, the layers are numbered from top to bottom. The top layer (60) is at 2.0 mb. Layers above about 43 are in the stratosphere. Layer 20 is near 850 mb. The horizontal axis gives the thickness of each layer in terms of η. If one added up the $\Delta\eta$ (same as *DSIGMA*) values for the 60 layers, the sum would be 1.0. The graph clearly shows that the layers near the surface and the top of the model are more bunched in the new version than in the old version of the model. The increased vertical resolution at low altitudes should resolve boundary layer processes more accurately and give better forecasts of mountain waves and downslope windstorms.

The horizontal grid for the NAM NMM-B is called an Arakawa B-grid, after its inventor Akio Arakawa. The B-grid is staggered, as shown in Fig. 5.3. That means that different variables are stored at different locations. In particular, the (u,v) components of the horizontal wind are stored at the yellow dots, and the temperature and moisture (T,q) are stored at the black dots. In Fig. 5.3, Δx refers to the grid-point spacing, approximately 12 km for the NAM NMM-B.

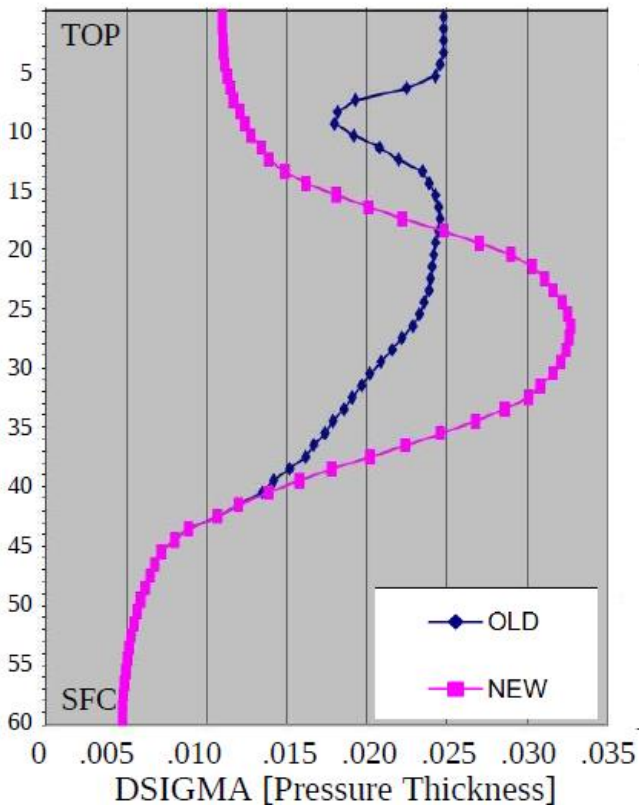

Fig. 5.2. The distribution of the 60 levels in the hybrid sigma-pressure coordinates. The graphs give the relative thickness of each layer for the old NAM (blue) and the new NAM (pink).

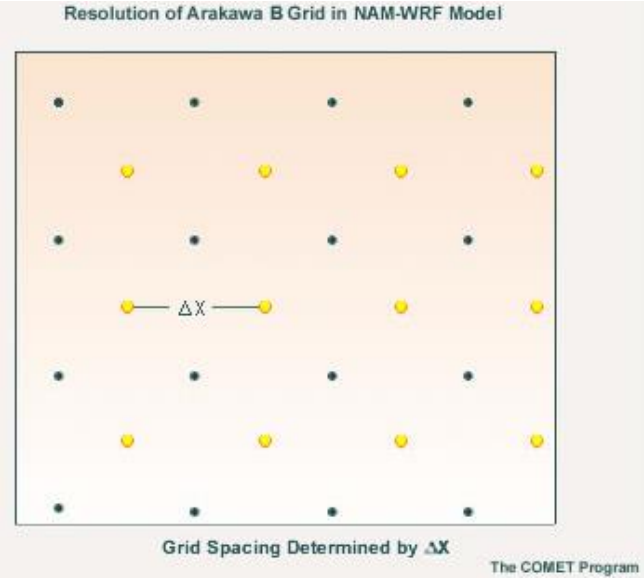

Fig. 5.3. The staggered Arakawa B-grid. (u,v) stored at yellow dots; (T,q) stored at black dots. Courtesy of the COMET Program.

The geographical coverage of the NAM NMM-B is shown in Fig. 5.4. The Mercator projection, a shape-preserving projection, is often used to portray computational grids for global models. In this projection, longitude lines are vertical and evenly spaced; latitude lines are horizontal, but spacing between latitude lines increases from equator to pole. The spacing beyond 80^{o} of latitude becomes so large that very high latitudes are usually not shown. Now imagine a Mercator projection rotated so that the equator passes through the center of the map shown in Fig. 5.4, that is, near the northern border of the U.S. The staggered B-grid is laid out along the latitude-longitude lines of this rotated grid, which appears as the rectangular array of points in Fig. 5.4. T and q are stored at the green dots or crosses; u and v are stored at the red dots. The *real* latitude and longitude lines are mapped to the rotated grid. The NAM grid covers all of North America, the North Pole, and much of the North Pacific and North Atlantic Oceans. Rotation of the Mercator projection allows the grid-point spacing to be much more uniform over the surface of the earth than it would otherwise be. The dots in Fig. 5.4 show the grid layout; the actual spacing of points on a 12-km grid would be much closer.

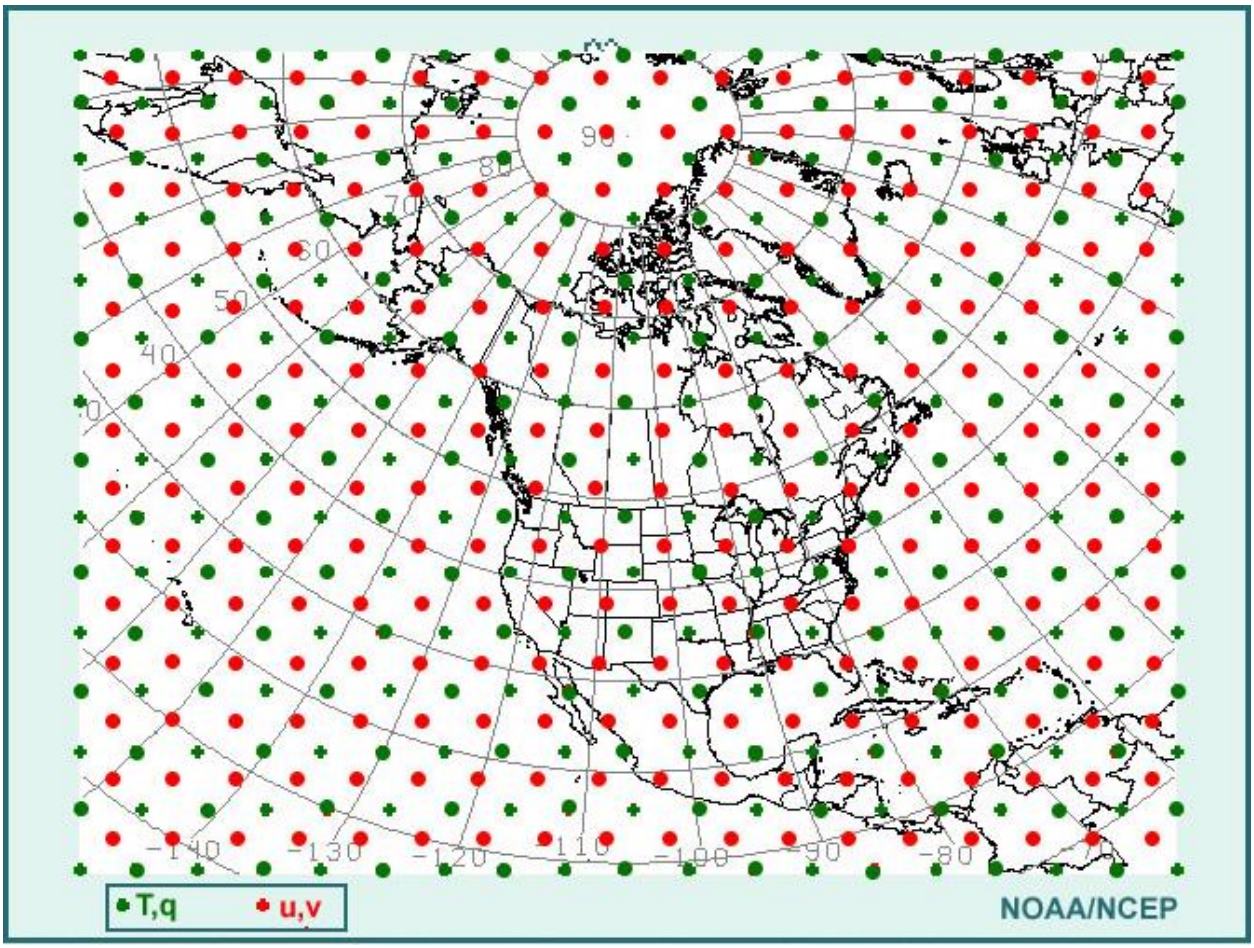

Fig. 5.4. The layout of the Arakawa B-grid in a transformed Mercator projection, in which the equator runs through the center of the image. A very coarse grid is shown. The actual 12-km NAM NMM-B grid is much denser.

A major improvement for aviation forecasting in Alaska and the CONUS in October 2011 was the inauguration of nested grid forecasts within the parent 12-km NAM grid. As Fig. 5.5 hints, the computational grids within the nests are all compatible with the parent grid.

The parent NAM model runs four times per day at 0000, 0600, 1200 and 1800 GMT at 12-km resolution out to 84 h. The nested forecasts run on the same schedule.

- The fixed CONUS nest (largest rectangle) runs at 4-km resolution out to 60 h.
- The fixed Alaska nest (large square) runs at 6-km resolution out to 60 h.
- The fixed nests for Hawaii and Puerto Rico (small rectangles) run at 3-km resolution out to 60 h.
- The nests for fire weather or other emergencies are movable. The CONUS nest runs at 1.33-km resolution; the Alaska nest runs at 1.5-km resolution. Both forecasts run out to 36 h.

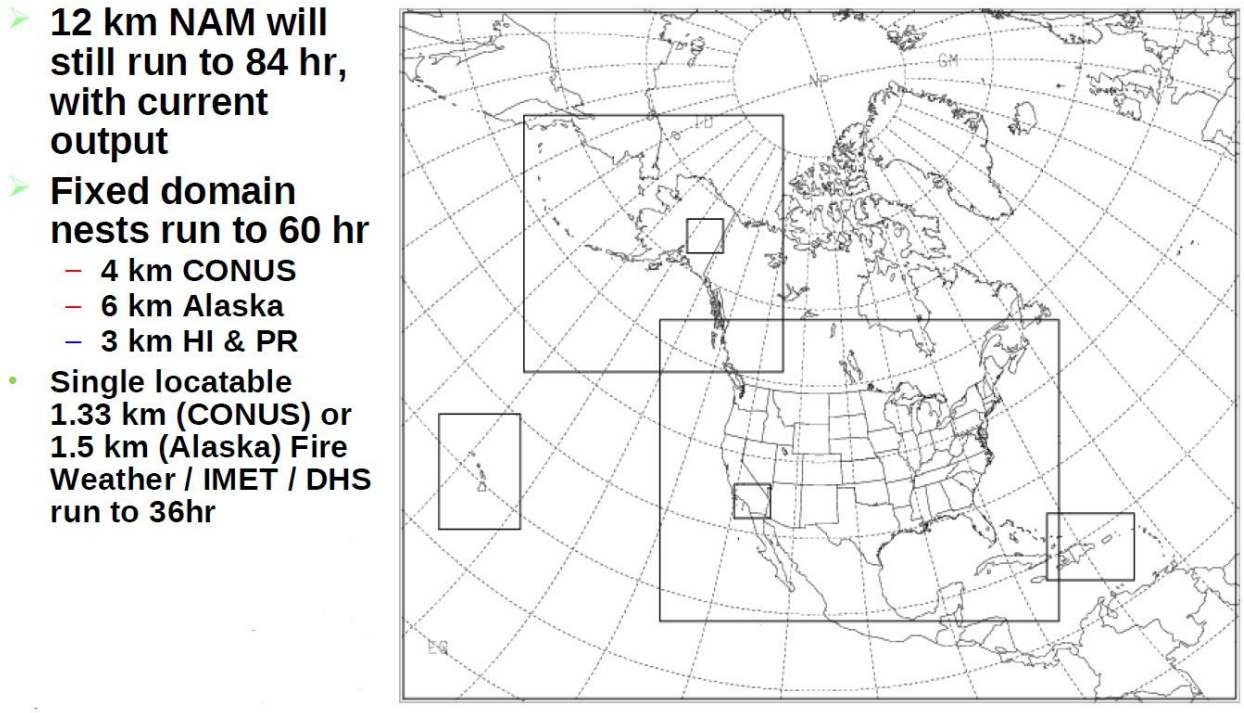

Fig. 5.5. Nests within the parent NAM NMM-B 12-km model. See text for details. IMET – Incident METeorologist trained to provide weather information on site when wildfires or other disasters occur. DHS – Department of Homeland Security.

Comprehensive documentation of the NAM NMM-B model is available in a training module generated by COMET, the Cooperative Program for Operational Meteorology, Education and Training, run by the University Corporation for Atmospheric Research, the corporate parent of NCAR. Go to http://www.meted.ucar.edu/nwp/pcu2/ and view items in the leftmost column of the matrix for information on the NAM NMM-B model.

Here are the major physical processes included in the NAM NMM-B model.

At the surface of the Earth:

- Absorption of solar radiation
- Emission and absorption of longwave radiation
- Sea-surface temperature prescribed by satellite observations
- Snow cover and the liquid water equivalent of snow updated from NESDIS and U.S. Air Force sources
- Sea ice updated with NESDIS satellite data
- Vegetation type and vegetation fraction from MODIS and the International Geosphere-Biosphere Programme
- Soil type from U.S. Department of Agriculture database
- Soil moisture from a four-layer model

In the atmosphere

- Atmospheric radiation: gaseous absorption in clear sky conditions of solar and IR radiation; effects of clouds on the radiation field
- Turbulence (Mellor-Yamada-Janjic 2.5-Level Closure – see web site above for details)
- Convection (Betts-Miller-Janjic parameterization – see web site above for details)
- Cloud microphysics and precipitation (see immediately below)

Because clouds and precipitation are so important for aviation, extra detail on the NAM NMM-B parameterization is provided here. The parameterization is due to Brad Ferrier at NCEP/EMC. See http://www.meted.ucar.edu/nwp/pcu2/nambpcp1.htm for a very accessible description of this clever and computationally economical algorithm. The scheme allows four classes of hydrometeors: suspended cloud liquid water droplets, rain, *large ice* (snow, graupel, sleet, etc.) and *small Ice* (suspended cloud ice). Together, these four classes comprise the total condensate. Though all four mixing ratios are tracked within a grid box, only the total condensate is advected. The partitioning of total condensate into the four classes is accomplished empirically. It depends upon the air temperature, and the vapor pressure e. Let e_{si} and e_{sl} be the saturation vapor pressures with respect to ice and liquid water, respectively. Different partitioning of total condensate occurs for the temperature classes $T > 0^{\circ}\mathrm{C}$, $-40 < T \leq 0^{\circ}\mathrm{C}$, and $T < -40^{\circ}\mathrm{C}$ in combination with the water vapor classes $e < e_{si}$, $e_{si} \leq e < e_{sl}$, and $e \geq e_{sl}$. The partitioning is summarized in Fig. 5.6.

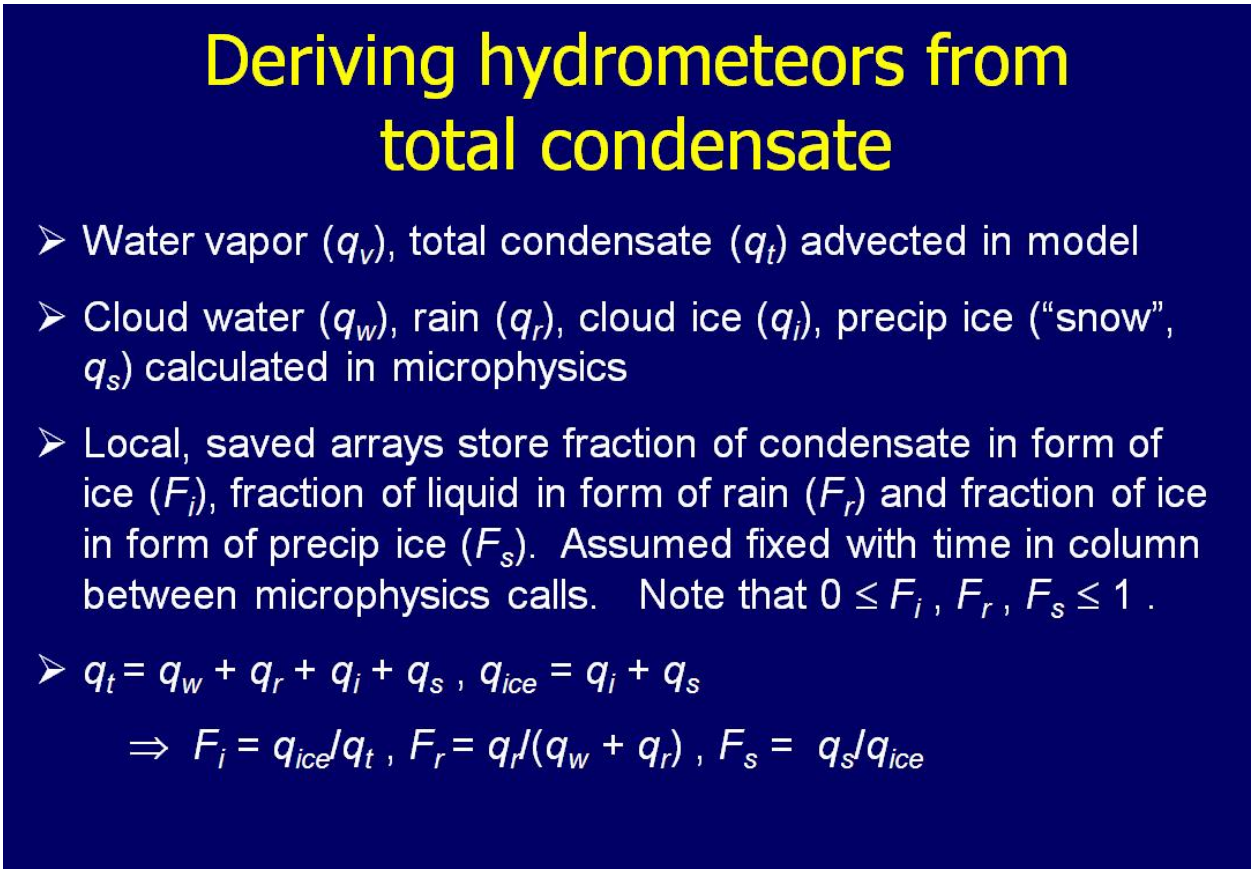

Fig. 5.6. The partitioning of total condensate into four hydrometeor classes in the Ferrier microphysics scheme.

The Ferrier scheme allows for many microphysical processes. Hydrometeors can grow and shrink in size. Precipitation can fall to the ground, but it is subject to evaporation in sub-saturated air below cloud base. In a sub-freezing cloud containing both liquid and ice particles, snow crystals can grow at the expense of liquid droplets, a process called glaciation. Raindrops can grow in size through collision and coalescence. Fall speeds of hydrometeors depend upon hydrometeor type and size. For example, larger raindrops fall faster than smaller raindrops, and raindrops fall faster than snow crystals. Supercooled drops can freeze on snow crystals in a process called riming. Melting ice can shed small raindrops.

Like the GFS, the NAM draws upon a great variety of observations, which are assimilated with a regional adaptation of the GSI. Figure 5.7 lists the conventional (non-radiance) observations as of September 2011. Sources of mass observations (related to pressure, moisture, and temperature) are listed in the left column. Sources of wind observations are listed in the right column. Note the many sources of wind information derived from tracking features in sequences of visible, infrared, and water vapor satellite images.

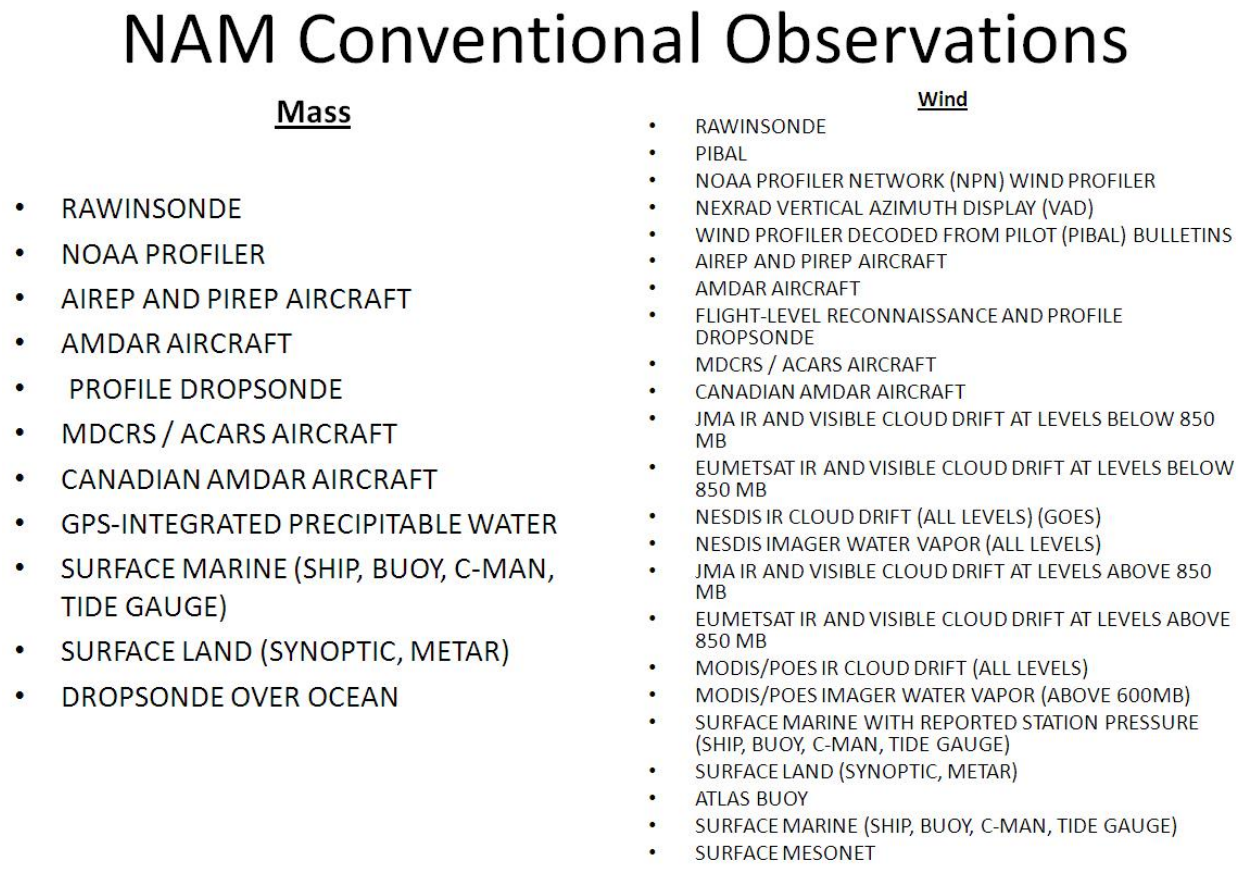

Fig. 5.7. Sources of meteorological observations used in NAM NMM-B forecasts (radiance data excepted) as of September 2011. Sources in the left column give information about the distribution of atmospheric mass in the vertical. Sources in the right column give wind information.

Several new sources of observations were added in the October 2011 upgrade of the NAM:

- Water Vapor Sensing System, Version 2 (WVSS-2) moisture observations included in automated aircraft reports, currently (early 2012) from aircraft flown by United Parcel Service and Southwest
- Observations of temperature, surface pressure, and humidity from mesonet stations around the country, using a dynamic reject list (data quality is variable)
- Vertical wind profilers obtained from multiple agencies
- Profiles of virtual temperature (confined to boundary layer) from Radio Acoustic Sounding Systems (RASS)
- Ocean surface winds obtained from a polarimetric microwave radiometer (WindSat) aboard the U.S. Air Force's Coriolis satellite and the Advanced Scatterometer aboard EUMETSAT's Metop satellite.

5.3 The RUC, RAP, and HRRR Models

The Rapid Update Cycle (RUC) is a data-assimilation / short-range prediction system operational at NCEP/EMC (Benjamin et al. 2004a, Benjamin et al. 2010). It provides a new analysis of atmospheric conditions and an 18-h forecast every hour. The geographic coverage of the RUC is shown in red in Fig. 5.8. The Rapid Refresh (RAP), to become operational early in 2012, is the successor to the RUC. It will run in the area enclosed in blue. The High-Resolution Rapid Refresh (HRRR) is nested within the RAP; it runs over the CONUS (green rectangle) at 3-km resolution. The HRRR currently runs in real time on NOAA research computers but is not currently planned to be operational on NCEP computers for several more years. The RAP and HRRR are already being used to generate a number of aviation weather products. This section will describe all three models and the differences among them.

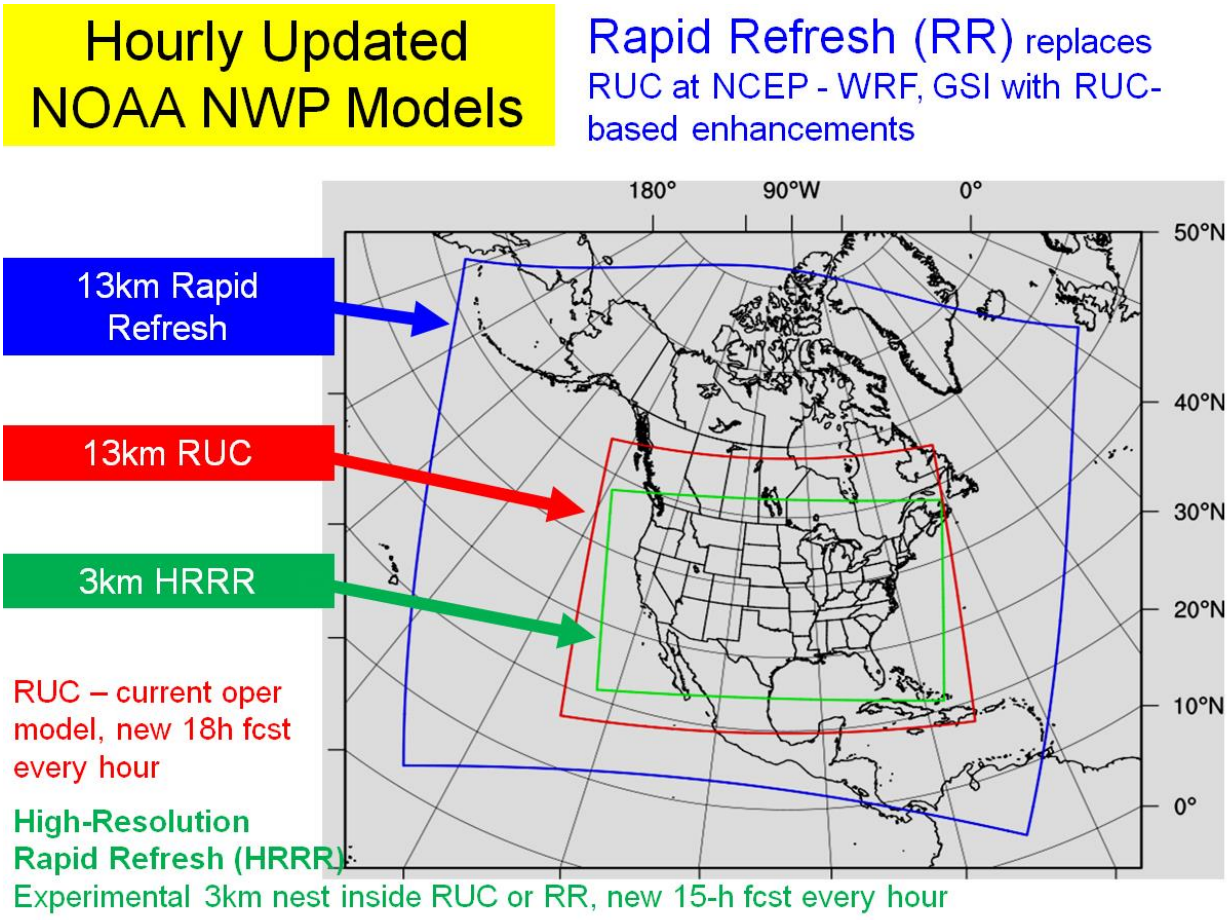

Fig. 5.8. Geographic coverage of the RUC, RAP, and HRRR models. Courtesy of Stan Benjamin.

Table 5.2 facilitates a comparison among the RUC, RAP, and HRRR models. Explanations follow the table.

Table 5.2

A comparison of properties of the RUC, RAP, and HRRR models.

Model Characteristics	**Rapid Update Cycle (RUC)**	**Rapid Refresh (RAP)**	**High-Resolution Rapid Refresh (HRRR)**
Domain	CONUS +	North America	CONUS
# Horizontal Grid Points	451 x 337	758 x 567	1799 x 1059
Map Projection	Lambert Conformal	Rotated lat/lon (like NAM)	Lambert Conformal
Grid Spacing	13 km	13 km	3 km
Horizontal Differencing	Arakawa C-grid	Arakawa C-grid	Arakawa C-grid
Vertical Levels	50	50	50
Vertical Coordinate	Hybrid Sigma-isentropic	Sigma	Sigma
Height of Lowest Level	5 m AGL	~8 m AGL	~8 m AGL
Pressure Top	500 K (40-85 mb)	10 mb	20 mb
Where Model Runs	NCEP	GSD/EMC; becomes operational early 2012	ESRL/GSD and NESCC (starting spring 2012)
Time Step	18 s	60 s	22 s
Length of Forecast	18 h	18 h	15 h
Analysis	RUC 3DVAR	GSI 3DVAR	None: Initial conditions from RAP model
Frequency of Runs	Hourly (cycled)	Hourly (cycled)	Hourly (not cycled)
Lateral Boundaries	NAM	GFS	RAP
Run time	~25 min	~25 min	~ 50 min
Latency	T + 50 min	T + 50 min	T + 2 h

Several items in Table 5.2 deserve explanation. All three map projections are conformal, that is, actual shapes on the Earth's surface are preserved in the projection. The projections are chosen so that variation in the map scale is minimized. All three models use the Arakawa C-grid for finite differencing, as in Fig. 5.9. On the C-grid, the u- and v-components of the wind are stored at different locations, and the mass variables are stored at points separate from u and v.

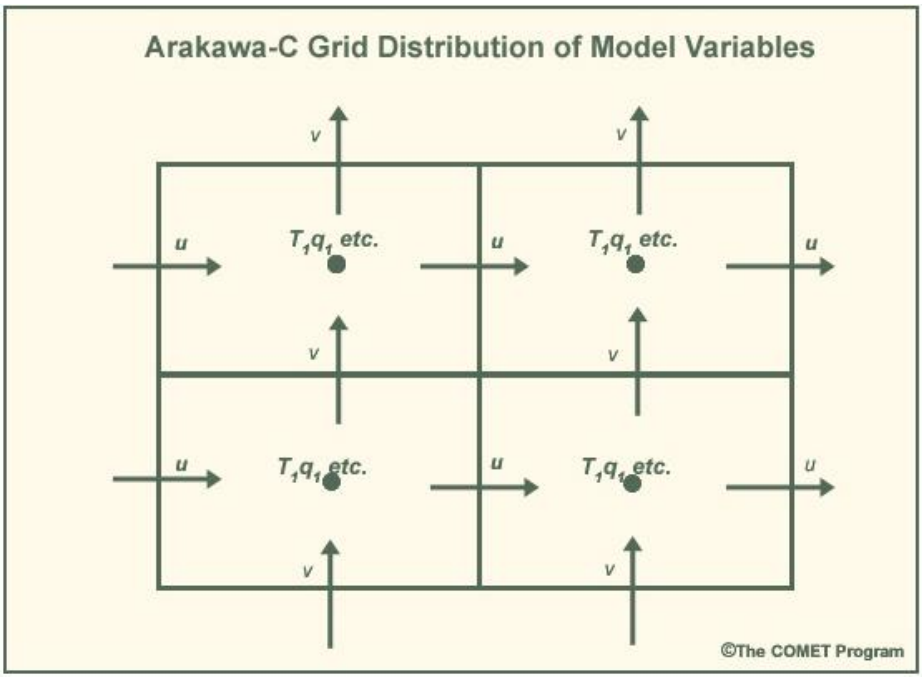

Fig. 5.9. The Arakawa C-grid used by the RUC, RAP, and HRRR models. Courtesy of the COMET Program.

Vertical Coordinate. The RUC model uses a hybrid sigma-isentropic vertical coordinate. This coordinate has the advantage of following the terrain close to the ground but gradually transitions to an isentropic coordinate in the free atmosphere, that is, following surfaces of constant potential temperature. In the absence of diabatic effects, air flow remains on isentropic surfaces, which motivates the choice of this coordinate system. The predictive component of the RUC is *not* based upon the Weather Research and Forecasting (WRF) model. The decision to change prediction models in the transition from the RUC to the RAP was prompted partly by the fact that the WRF model has the advantages of mature infrastructure, shared NOAA and NCAR development, and broad community support. The predictive component of the RAP thus uses the Advanced Research version of the WRF model (ARW), which has a sigma (terrain-following) vertical coordinate.

Pressure Top. Because the RUC model top is defined in terms of potential temperature (500 K), the pressure of this surface (in the stratosphere) varies from 40–85 mb, depending upon the season and atmospheric wave patterns.

Where the Model Runs. At the end of 2011, the HRRR model ran regularly but experimentally at ESRL/GSD. To improve reliability of the runs (the HRRR is used in the generation of real-time aviation weather products), the HRRR will begin running also at the National Environmental Security Computing Center (NESCC), Fairmont, West Virginia, in spring or summer 2012.

Analysis. The 3DVAR analysis used in the RUC was adequate for a system that covered the CONUS, southern Canada, and northern Mexico, but did not include large ocean areas. The expanded coverage of the RAP (see again Fig. 5.8) demanded a data assimilation package that could accommodate a wide range of national and international satellite data because satellites provide the vast majority of oceanic observations. Accordingly, the RAP employs a regional adaptation of the GSI algorithm used in the GFS model. The HRRR does not have a cycled data assimilation system. Instead, it depends upon the hourly RAP analysis, which is interpolated to the finer HRRR grid.

Run time. The run time is the wall clock time required to compute the forecast.

Latency. Latency refers to the time elapsed since the initial time T before the product is available on the web.

Table 5.3 references papers that describe how various physical processes are treated in the NAM, RUC, RAP, and HRRR models. Following the table, extra information is provided regarding cloud microphysics and hydrometeors.

Table 5.3

How several different models treat various physical processes. The table entries lead to references where parameterizations are discussed in detail.

Model	Radiation	Cloud Physics / Hydrometeors	Deep Convection	Boundary Layer	Shallow Cumulus	Land-Surface Model
NAM	Solar: Lacis and Hanson (1974) Infrared: Schwarzkopf and Fels (1991)	Ferrier scheme described in Rogers et al. (2001)	Betts-Miller-Janjić scheme (BMJ) Janjić (1994)	Mellor-Yamada-Janjic scheme (MYJ) Janjić (2001)	BMJ Janjić (1994)	Noah[2] (an acronym) Pan and Mahrt (1987)
RUC	Solar: Dudhia (1989). Infrared: Mlawer et al. (1997)	Thompson (2004)	Grell and Devenyi (2002)	Burk and Thompson (1989)	None	Smirnova (1997, 2000) Benjamin et al. (2004b)
RAP	Mlawer et al. (1997)	Thompson (2008)	Grell G3[1]	MYJ Janjić (2001)	Grell G3[1]	Smirnova (1997, 2000) Benjamin et al. (2004b)
HRRR	Mlawer et al. (1997)	Thompson (2008)	Calculated explicitly	MYJ Janjić (2001)	None	Smirnova (2000) Benjamin et al. (2004b)

[1]Grell G3: http://www.mmm.ucar.edu/wrf/users/workshops/WS2008/presentations/10-2.pdf
[2]Noah: http://www.atmos.illinois.edu/~snesbitt/ATMS597R/notes/noahLSM-tutorial.pdf

Just as Ferrier's treatment of cloud microphysics and precipitation received extra coverage earlier, so does the Thompson (2008) microphysics scheme in this case. This is the scheme used in the RAP and HRRR models. Unlike the Ferrier scheme, which treats only water vapor and total condensate explicitly, the Thompson scheme predicts mixing ratios of water vapor, cloud water (tiny, suspended droplets), rainwater (drops large enough to precipitate), cloud ice (tiny individual crystals, which also precipitate but more slowly), snow (larger crystals and aggregates), and graupel (heavily rimed, soft ice pellets). In addition, the scheme predicts the size distribution of raindrops. The explicit treatment of these hydrometeors is an advantage for

predicting icing within supercooled clouds and the type of precipitation that reaches the ground.

The interactions among water vapor and the five hydrometeor types include:

- condensation and evaporation
- deposition and sublimation
- freezing and melting
- accretion (through collision and coalescence: liquid-liquid or liquid-solid)
- Bergeron-Findeisen process, whereby ice crystals grow at the expense of liquid droplets because of the difference in saturation vapor pressures of these two species
- aggregation of crystals into flakes
- mixing ratios in a grid box depleted by precipitation at a rate commensurate with hydrometeor type and fall speed.

The motivation for the original development of the RUC still exists: to capitalize on the great variety of meteorological observations available at frequent intervals over the United States. Aviation often requires sub-hourly guidance on changing weather conditions, and rapid updating of analyses and forecasts directly responds to this need. The hourly cycle of data assimilation and prediction in the RAP is illustrated in Fig. 5.10. A one-hour forecast by the WRF-ARW model provides a starting point, a *background*, for the analysis of current observations. The GSI analysis "corrects" the one-hour forecast through the assimilation of observations. The forecast and the observations are weighted according to the errors known to be associated with each in the generation of *analysis fields*, which are the initial conditions for a subsequent one-hour forecast. In this way, the dynamical and physical information embodied in the model and frequent atmospheric measurements are exploited to move the state of the atmosphere forward incrementally in time. Efforts continue to incorporate new sources of high-frequency, high-density observations.

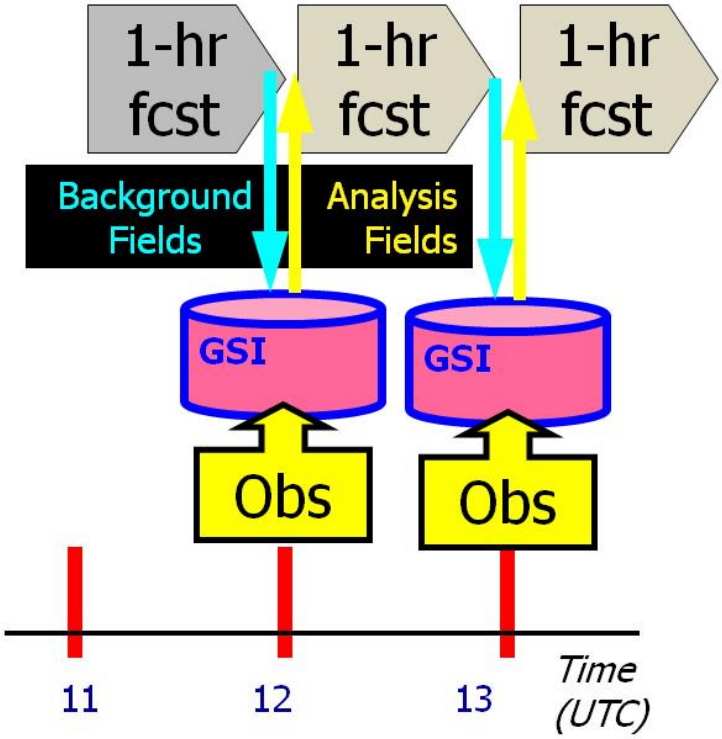

Fig. 5.10. Illustrating the hourly assimilation and prediction cycle for the Rapid Refresh model and the GSI assimilation algorithm that incorporate both observations and forecast information to produce a new analysis every hour. Courtesy of Stan Benjamin.

Table 5.4 lists the hourly observations (rawinsondes excepted) assimilated in the RAP.

Table 5.4
Sources of hourly observations used in the Rapid Refresh.

Data Type	**Approximate number per hour**
Rawinsonde (00 and 12 GMT)	120 (once every 12 h)
NOAA profilers	21
VAD winds	~125
PBL – profiler/RASS	~25
Aircraft (wind, temperature)	2,000-15,000 (average 7,000)
WVSS (RH)	0-800 (average 520)
Surface/METAR	~2,500
Buoy/ship	200-400
GOES cloud winds	4,000-8,000
GPS integrated precipitable water	~260
GOES cloud-top pressure	10-km resolution
METARs (clouds, visibility, present weather)	2,000

WSR-88D radar reflectivity 1-km resolution
AMSU A/B, HIRS MHS

(Below – for ESRL RAP and HRRR only, as of early 2012)
Lightning (proxy for radar reflectivity)
Radar radial wind observations
Sodar and wind tower wind observations

Assimilation of the GOES cloud-top pressure, METAR, and 88D data in Table 5.4 has led to improved forecasts of clouds and precipitation. Additional data sources will be added to the NCEP RAP in an upgrade later in 2012:

- Surface mesonet data (temperature, dewpoint, and wind). These data vary widely in quality and must be carefully monitored.
- Lightning detection data as a proxy for reflectivity in areas with poor radar coverage
- Radial winds from WSR-88D radars.

5.4 Ensemble Prediction: Dealing with Uncertainty in the Forecast

The accuracy of a deterministic forecast, that is, one with a single outcome based upon prescribed initial conditions, is hard to judge a priori. Meteorologists expect less accuracy in rapidly changing conditions than in quiescent conditions, but forecast uncertainty is hard to quantify ahead of time. When decisions regarding flight safety or efficiency must be made, there is great value in knowing the odds that weather conditions will fall within a certain range, when a wrong decision could involve personal injury or economic loss.

Ensemble prediction is a method for estimating the probability of a whole range of conditions. It involves running a collection of forecasts under different but still plausible conditions to generate a range of outcomes. If all forecasts in the collection are very similar, one has confidence that the outcome is correctly forecast. On the other hand, if the member forecasts in the ensemble have a large spread, one places little confidence in the forecast.

Over the long term, experience shows that the mean of an ensemble of forecasts will be closer to the truth than the individual members. The standard deviation about the mean is a measure of the spread in the ensemble members.

There are many ways to generate an ensemble of forecasts:

- Perturb the initial conditions or use analyses from different data assimilation schemes.

- Include models with different dynamical cores in the ensemble.
- Use different parameterizations of physical processes.
- Change the horizontal or vertical resolution.
- Use different surface boundary conditions.
- For a regional model, use different lateral boundary conditions.

The most basic requirement is that each member of the ensemble produce a plausible forecast.

NCEP/EMC generates several kinds of ensemble forecasts at global and regional scales. The focus here is on two ensembles that have the most direct application to aviation weather. They generate short-range forecasts at reasonably high resolution, and they cover the CONUS and Alaska. They are the Short-Range Ensemble Forecast (SREF) system and the North America Rapid Refresh Ensemble – Time Lagged (NARRE-TL).

As of the most recent upgrade in October 2009, the SREF system included 21 member forecasts running at 32-km grid resolution (or its equivalent) out to 87 hours. Another upgrade in resolution to 16 km is expected late in 2012. Six members are variations of the NAM/Eta model, the predecessor of the NAM/NMM model, and five members each are variations of the WRF-ARW model, the NAM-NMM model, and NCEP's Regional Spectral Model (RSM) (Huang and Kanamitsu, 1997). The principal differences between model runs in the same class are in the initial conditions and the model physics. To examine SREF model output, go to http://mag.ncep.noaa.gov/NCOMAGWEB/appcontroller, click on "Model Guidance," then on "SREF."

The NARRE-TL was formerly called the Very Short-Range Ensemble Forecast. It differs substantively from the SREF system in that it consists not of a collection of forecasts starting and stopping at the same times but rather a collection starting at staggered times yet all verifying at the same time. These are called *time-lagged* forecasts and, together, they comprise an ensemble. The major computational advantage of time-lagged ensemble forecasts is that they require no additional computation. All the forecasts are produced in the normal operational sequence, but the overlaps in their valid times means that computation of the mean and spread of the ensemble requires only post-processing of the model output. The disadvantage of such ensembles is that members have less statistical independence than SREF members.

Figure 5.11 shows a NARRE-TL whose members participate in the computation of ensemble mean and spread for each of the hours from 0700 through 1800 GMT. (The RAP forecasts go out to 18 h.) As noted in the figure, forecasts get varying weights depending upon their age.

The NARRE-TL will become operational some months after the RAP becomes operational.

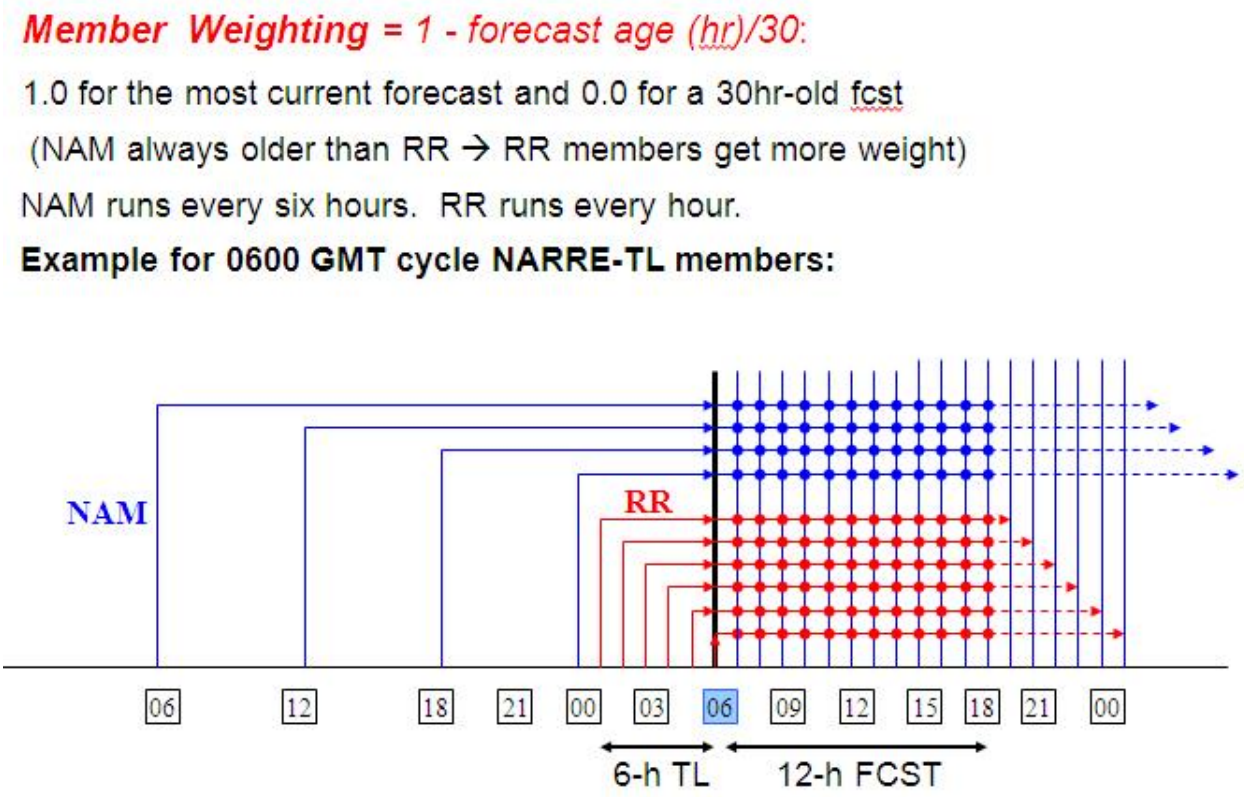

Fig.5.11. Example of a NARRE-TL with initial time of 0600 GMT. Four NAM forecasts and six RAP forecasts participate in the ensemble, for which statistics are computed from 0700 through 1800 GMT. Figure courtesy of Jun Du (NCEP/EMC).

Table 5.5 lists some of the aviation weather products that will be available from the NARRE-TL.

Table 5.5
Aviation weather information to be extracted from NARRE-TL (courtesy of Jun Du, NCEP/EMC).

Field	Ensemble Product
Icing	Probability of occurrence on eight flight levels
Turbulence (CAT)	Three levels of severity, probability of occurrence on nine flight levels
Ceiling (ht of cloud base)	Mean and spread; probability of occurrence in four ranges
Surface visibility	Mean and spread; probability of occurrence in four ranges
Low-level wind shear	Mean and spread; probability of occurrence
Jet stream	Probability on three levels
Fog (light or dense)	Mean, spread, and probability of occurrence
Convection	Probability of occurrence
Reflectivity	Probability of occurrence, four thresholds
Freezing level	Mean and spread
Precipitation type	Probability of rain and snow types
Accumulated precipitation	Probability of precipitation thresholds at three and six hours
Lightning	Probability of occurrence
Severe thunderstorm	Probability of occurrence

ESRL also produces a 3-km, time-lagged ensemble product called the HRRR Convective Probabilistic Forecast (HCPF) that incorporates spatial filter averaging and use of time-integrated vertical motion (maximum values from each grid point over the most recent hour). The HCPF presages a future HRRR Ensemble, ultimately planned to be run at NCEP with 6 members initialized by a 6-member ~12-km North American Rapid Refresh Ensemble.

6. Final Remarks

Remarks here give opinions of the author and do not necessarily reflect the policy of the sponsoring organization (NWS), the FAA, or any other government agency or private entity that generates aviation weather products.

Given the state of the science, aviation weather researchers have done well in generating products. However:

1) The ability to create slick graphics and to make forecasts look like observations (radar and satellite), i.e., the technology, is well ahead of the ability to make accurate forecasts of many aviation hazards more than a couple of hours in advance.

2) There is still a great amount of empiricism behind most of the products, which is probably appropriate, given the state of the science, but building more physics into product generation while reducing empiricism should be the goal. The use of probabilities in predictions is valuable for end users and should be continued. Product developers should try to reduce the use of statistical correlations as understanding of cause and effect improves. Some phenomena important for aviation occur at such small scales, notably turbulence, that they will be impossible to predict explicitly for a long time to come. For such phenomena, statistical inference will remain a necessary tool.

3) Mesoscale models at high resolution (currently 3 km for the High-Resolution Rapid Refresh model) and with sophisticated representations of physical processes should play an increasing role in product generation as forecast accuracy improves. It should not be forgotten that most aviation hazards occur at very small scales, still far below the resolution of operational models and even most research models. Improvements in cloud physics and the parameterization of turbulence, and more dense and frequent observations in the boundary layer, especially of wind, temperature, and moisture, will go a long way toward addressing the vexing problems of clear-air turbulence, icing (by implication, cloud forecasts in general), and the initiation of convection.

4) Computational horsepower, especially at the Environmental Modeling Center, is insufficient to support models running at the high resolutions demanded by many aviation applications. This is slowing progress in making many worthwhile products more widely available.

5) The distinction between *primary* and *supplemental* products is legalistic rather than practical. Primary products have very high reliability, but costs to bring a product from 99.0% to 99.9% reliability are high, and the FAA and NWS do not want to be held legally responsible for decisions made on the basis of supplemental products. Still, there is probably more useful information available in supplemental products than in primary products, and conscientious pilots and the dispatchers who support them want to base their decisions on the most comprehensive information available.

6) The information in this report is *dated.* To the author's knowledge, this is the only comprehensive survey of web-based products relating to aviation weather hazards and the only

attempt to summarize the science behind them. The information in this report is believed to be current as of late 2011, but, within a year or two, many of the details will be out of date because the development of new products and the numerical prediction models upon which they rely continues at a rapid pace.

Acknowledgments

I especially thank John C. Osborn of ESRL/GSD for comprehensive graphics support during the entire year this report was in preparation. Special thanks also to Ann M. Reiser (ESRL), who edited the entire manuscript and linked the Table of Contents with the text. Barbara Schlatter helped prepare the list of acronyms. Jeannine Schlatter checked the reference list.

Visit to the Aviation Weather Center in Kansas City, October, 2010
Thanks to Bruce Entwistle, who hosted the visit, planned the itinerary, and spent hours explaining AWC operations. I spent time on shift with Nolan Duke at the tropical desk, Jim Roets (FA East), Gregory Harris (convective SIGMETs), and Don Baker (CCFP). All these meteorologists were generous with their time and eager to share their experiences and past presentations.

Visit to the University of Wisconsin, Madison, January, 2011
Thanks to Ralph Petersen and Wayne Feltz, who coordinated the visit and planned the itinerary. I conducted interviews with the following on the use of satellite information for diagnosis of aviation weather hazards:
Wayne Feltz introduced the GOES-R Proving Ground concept.
Justin Sieglaff, Lee Cronce, and Kris Bedka (NASA, Hampton, VA, participated remotely) discussed convective initiation.
Jason Brunner, Scott Lindstrom, and Kris Bedka (again, participating remotely from NASA) explained detection of overshooting tops.
Mike Pavolonis and Justin Sieglaff explained an algorithm for the detection of volcanic ash and several of its properties.
Mike Pavolonis and Corey Calvert presented the GOES-R algorithm for detecting fog and low clouds.
Brad Pierce discussed his research on inferring surface visibility from GOES-R ABI data.
Tim Schmidt summarized current and expected GOES satellite capabilities.
Andy Heidinger explained how MODIS data are being used to simulate GOES-R data for the detection of cloud properties especially as they relate to icing.

Tony Wimmers explained algorithms for detecting tropopause folds, gravity wave, and mountain waves—all related to turbulence.

Chris Velden gave a status report on feature tracking by satellite for the inference of atmospheric winds.

Ralph Petersen gave an update on WVSS-II (water vapor sensing system) and explained his research on the inference of convective initiation and dissipation from satellite data.

Introduction

Jennifer Mahoney (ESRL/GSD provided verification statistics and verification reports for many products. Her group's RTVS was the source of several verification figures and tables.

Stephen Walden, Weather Systems Engineering, introduced me to a large body of regulations about flying and how they relate to the use of aviation weather information.

Turbulence

Bob Sharman (NCAR) gave the initial interview on GTG2. He helped with the GTG2 turbulence section and reviewed a draft.

Rod Frehlich (NCAR) also participated in discussions about GTG2 and provided references.

John Williams (NCAR) and Michael Istok (NWS Office of Science and Technology) provided information on the NEXRAD Turbulence Detection Algorithm.

Tom Fahey (Delta Airlines) sent information on the Delta Turbulence Plot System.

Icing

Marcia Politovich (NCAR) gave the initial interview on the icing products.

Marcia, Cory Wolff, and Tressa Fowler, also of NCAR, provided references and reviewed the draft segment on CIP/FIP.

Patrick Minnis steered me to a host of references and web sites relating to NASA LaRCs work on cloud properties and their connection with icing.

William L. Smith, Jr. (NASA LaRC), and Patrick Minnis reviewed the LaRC segment on cloud Properties.

William Moninger (ESRL/GSD) provided plots of atmospheric soundings used to infer icing.

Convection

Marilyn Wolfson (MIT Lincoln Laboratory) gave several interviews on CIWS and CoSPA, was quick to exchange emails, and provided numerous formal and informal references.

Haig Iskenderian, (MIT Lincoln Laboratory), with help from Marilyn Wolfson, reviewed many pages of manuscript pertaining to CIWS and CoSPA and offered extensive, constructive comments.

Colleen Reiche (MIT Lincoln Laboratory) reviewed the segment on CoSPA.

Greg Thompson (NCAR) and John Brown (ESRL/GSD) explained the calculation of radar reflectivity from model mixing ratios.

Keith Brewster (CAPS/University of Oklahoma) gave a phone interview and provided references on merging the CoSPA extrapolation forecast with the HRRR model forecast.

James Pinto (NCAR) helped me understand how merging VIL fields from the HRRR model and the 88D radar occurs; he reviewed this portion of the CoSPA writeup.

Steven A. Lack (ESRL/GSD and CIRES) explained verification of the CoSPA product and provided several figures.

Judy Ghirardelli (NWS/MDL) helped with a phone interview about LAMP and pointed me to many references. She and her colleague, David Rudack, read the draft description of LAMP and sent comments.

Rita Roberts, Jim Wilson, and James Pinto, NCAR, gave the initial interview on the Autonowcaster and provided many references.

Rita Roberts and Juanzhen (Jenny) Sun (NCAR), reviewed the draft description of Autonowcaster. Jenny Sun also provided information about VDRAS.

David Sharp and Matthew Volkmer, NWS/Melbourne, FL, described their experience with the Autonowcaster.

Haig Iskenderian, MIT Lincoln Labs, provided extensive documentation on SATCAST.

Justin Sieglaff, SSEC, University of Wisconsin, pointed me to information sources about UWCI.

Fog and Low Ceiling

Paul Herzegh gave two interviews on CVA and CVF and provided a number of references and figures used in the report. He also reviewed the draft description of CVA and CVF and sent comments and corrections.

Andrew Loughe (ESRL/GSD) provided a report on the critical evaluation of CVA.

Charles Kluepfel, Performance Branch, NWS, spoke with me about verification, especially of TAFS.

Mike Pavolonis (STAR/NESDIS/U. Wisc.) and Corey Calvert (CIMMS, U. Wisc.) helped me understand the GOES-R product for fog/low clouds. They pointed me to extensive formal and informal documentation. Corey and Mike reviewed my description of the GOES-R product.

Paul Herzegh (NCAR) informed me about HEMS.

Hazards at and near Airports

Roy Rasmussen (NCAR) interviewed on ground de-icing of aircraft and the WSDDM project.

Scott Landolt (NCAR) gave a very quick review and helpful comments on the segment on

ground de-icing.

David Johnson (NCAR) enlightened me on lightning protection at airports.

William Mahoney and David Johnson (NCAR) sent information on LLWAS.

William Bobb (MIC, Denver CWSU) arranged for a half-day visit to his facility, including interviews and a tour of the collocated ARTCC. He also provided considerable documentation about the CWSU.

Larry Larkin (TMU, Denver ARTCC) interviewed and explained operations at the Traffic Management Unit.

Alaska and Volcanoes

Kristine A. Nelson (MIC, Anchorage CWSU) gave a personal interview and provided documentation.

Tony Hall (MIC, AAWU, Anchorage), gave briefings and arranged for access to his staff during the day shift.

Jeffrey M. Osiensky (SSD, Alaska Region, NWS) arranged my itinerary for a visit to Anchorage in January 2011, gave personal briefings, and offered extensive documentation regarding special problems faced by aviation in Alaska.

Christina A. Neal, Volcanologist (USGS, AVO, Anchorage), gave an interview and facility tour at AVO. She also provided extensive documentation on Alaska volcanoes.

I thank Gary Hufford (SSD) and Phil Dutton (Alaska Field Support Unit), both at the Alaska Region headquarters, NWS, in Anchorage for briefings and discussions.

Nathan Eckstein and Dave Hadley, forecasters at AAWU, Anchorage, took time out from their shift to discuss their jobs with me.

Barbara Stunder (ARL) helped with documentation on HYSPLIT.

Mike Pavolonis and Justin Sieglaff (SSEC, U.Wisc) provided documentation and consultation on satellite products relating to volcanic ash.

Atmospheric Modeling

Stan Benjamin and John M. Brown (ESRL/GSD) provided interviews, Powerpoint presentations, and documentation on RUC, RAP, and HRRR models. They also reviewed my draft description of these models.

Geoff DiMego (NCEP/EMC) interviewed on GFS and NAM models and provided extensive documentation. He reviewed what I wrote about these two models.

Appendix A. List of Acronyms

3DVAR	Three-Dimensional VARiational (a form of data assimilation)
4DVAR	Four-Dimensional VARiational (a form of data assimilation)
AAR	Airport Arrival Rate
AAWU	Alaska Aviation Weather Unit, Anchorage, Alaska
ABI	Advanced Baseline Imager
ACRP	Airport Cooperative Research Program

ADDS	Aviation Digital Data Service
AFCR	Available Flow-Capacity Ratio
AFOS	Automation of Field Operations and Services
AGL	Above Ground Level
AIM	Aeronautical Information Manual
AIRMET	AIRmen's METeorological information
AIRS	Atmospheric InfraRed Sounder
AMDAR	Aircraft Meteorological DAta Relay
AMS	American Meteorological Society
AMSU	Advanced Microwave Sounding Unit
ANC	Anchorage International Airport
ANC	AutoNowCaster
ARINC	Aeronautical Radio INCorporated
ARL	Air Resources Laboratory (NOAA)
ARTCC	Air Route Traffic Control Center
ARW	Advanced Research WRF (name given to one version of the WRF model)
ASAP	Advanced Satellite Aviation weather Program
ASCAT	Advanced SCATterometer
ASOS	Automated Surface Observing System
ASR	Airport Surveillance Radar
ATC	Air Traffic Control
ATCSCC	Air Traffic Control System Command Center
ATLAS	Autonomous Temperature Line Acquisition System (a moored buoy used in the tropics)
AVO	Alaska Volcano Observatory
AWC	Aviation Weather Center, Kansas City, Missouri
AWIPS	Advanced Weather Information Processing System
AWOS	Automated Weather Observing System
BMJ	Betts-Miller-Janjić (authors of a convective parameterization scheme)
BSS	Brier Skill Score
BT	Brightness Temperature
BTD	Brightness Temperature Difference
BTH	Bi-spectral Threshold and Height
CALIOP	Cloud-Aerosol Lidar With Orthogonal Polarization
CALIPSO	Cloud-Aerosol Lidar and Infrared Pathfinder Satellite Observations
CAPE	Convective Available Potential Energy
CAPS	Center for the Analysis and Prediction of Storms (Norman, OK)
CASA	Collaborative Adaptive Sensing of the Atmosphere
CAT	Clear Air Turbulence
CCFP	Collaborative Convective Forecast Product
CDT	Central Daylight Time
CFR	Code of Federal Regulations
CG	Cloud-to-Ground (lightning)

CI	Convective Initiation
CIMSS	Cooperative Institute for Meteorological Satellite Studies
CIN	Convective INhibition
CIP	Current Icing Potential
CIRA	Cooperative Institute for Research in the Atmosphere
CIRES	Cooperative Institute for Research in Environmental Sciences
CIWS	Corridor Integrated Weather System
C-MAN	Coastal Marine Automated Network, operated by the National Data Buoy Center
CNES	French Space Agency
CO_2	Carbon Dioxide
COLIDE	surface COnvergence LIne Detection and Extrapolation
COMET	Cooperative Program for Operational Meteorology, Education and Training
CONUS	CONtinental U.S.
CoSPA	Consolidated Storm Prediction for Aviation
CSI	Critical Success Index
CTC	Cloud-Top Cooling
CTT	Cloud-Top Temperature
CVA	Ceiling and Visibility Analysis
CVF	Ceiling and Visibility Forecast
CWA	Center Weather Advisory
CWAM	Convective Weather Avoidance Model
CWSU	Center Weather Service Unit
DARTS	Dynamic Adaptive Radar Tracking of Storms
dBZ	decibels of reflectivity (unit used for equivalent reflectivity factor)
DFW	Dallas-Fort Worth (airport)
DHS	Department of Homeland Security
DIA	Denver International Airport
DMSP	Defense Meteorological Satellite Program
DOT	Department of Transportation
DTC	Data Transformation Corporation
DU	Dobson Unit
DUAT	Direct User Access Terminal
ECFP	Extended Convective Forecast Product
EDR	Eddy Dissipation Rate
EMC	Environmental Modeling Center
ESA	European Space Agency
ESRL	Earth System Research Laboratory (NOAA)
EUMETSAT	EUropean organization for the exploitation of METeorological SATellites
EWINS	Enhanced Weather INformation Systems
FA	Area Forecast
FAA	Federal Aviation Administration
FAR	False-Alarm Ratio
FAR	Federal Aviation Regulation

FIMS	Flight Standards Information Management System
FIP	Forecast Icing Potential
FL	Flight Level
FLS	Fog / Low Stratus
FMF	Flight Movement Forecast
FSS	Fractions Skill Score
FZDZ	FreeZing DriZzle
FZRA	FreeZing RAin
GA	General Aviation
G-AIRMET	Graphical AIRMET
GFDL	Geophysical Fluid Dynamics Laboratory (NOAA)
GFS	Global Forecast System
GIS	Geographic Information System
GLM	Geostationary Lightning Mapper
GMS	Geosynchronous Meteorological Satellite
GMT	Greenwich Mean Time
GOES	Geostationary Operational Environmental Satellite
GOME	Global Ozone Monitoring Experiment
GPS	Global Positioning System
GSD	Global Systems Division (division of ESRL)
GSI	Grid-point Statistical Interpolation
GTG	Graphical Turbulence Guidance
HCPF	HRRR Convective Probabilistic Forecast
HEMS	Helicopter Emergency Medical Services
HIRS	High-resolution Infrared Radiation Sounder
HRRR	High-Resolution Rapid Refresh
HTB	Heated Tipping Bucket
HYSPLIT	HYbrid Single-Particle Lagrangian Integrated Trajectory
IASI	Infrared Atmospheric Sounding Interferometer
ICAO	International Civil Aviation Organization
IFR	Instrument Flight Rules
IGBP	International Geosphere Biosphere Programme
IGFOV	Instantaneous Geometric Field of View
ILS	Instrument Landing System
IMET	Incident METeorologist (provides weather support at accident/disaster sites)
IPS	Information Processing Systems
IR	InfraRed
ITWS	Integrated Terminal Weather System
JMA	Japanese Meteorological Agency
LAMP	Localized Aviation Model Output Statistics (MOS) Program
LaRC	Langley Research Center
LAT	LATitude
LEDWI	Light-Emitting Diode Weather Identifier

LIFR	Low Instrument Flight Rules
LLWAS	Low-Level Wind shear Alert System
LMB	LAMP-Model Blend
LSM	Land Surface Model
LUT	Look-Up Table
LWC	Liquid Water Content
LWP	Liquid Water Path
MCS	Mesoscale Convective System
MDCRS	Meteorological Data Collection and Reporting System
MDL	Meteorological Development Laboratory (NWS)
METAR	A type of surface observation approved by WMO
MHS	Microwave Humidity Sounder
MIC	Meteorologist in Charge
MIT	Massachusetts Institute of Technology
MITLL	Massachusetts Institute of Technology Lincoln Laboratory
MODIS	MODerate resolution Imaging Spectroradiometer
MOG	Moderate or Greater
MOS	Model Output Statistics
MRMS	Multiple Radar / Multiple Sensor
MSL	Mean Sea Level
MVFR	Marginal Visual Flight Rules
MYJ	Mellor-Yamada-Janjić (authors of a turbulence closure scheme)
NAM	North American Mesoscale (a computer model)
NARRE-TL	North America Rapid Refresh Ensemble – Time Lagged
NAS	National Airspace System
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
NCEP	National Centers for Environmental Prediction
NCWD	National Convective Weather Detection
NCWF	National Convective Weather Forecast
NDFD	National Digital Forecast Database
NESCC	National Environmental Security Computing Center (Fairmont, West Virginia)
NESDIS	National Environmental Satellite Data and Information Service
NEXRAD	NEXt-generation RADar
NextGen	Next-Generation (comprehensive overhaul) of the National Airspace System. FAA is the responsible agency.
NIDS	NEXRAD Information and Data System
NIR	Near InfraRed
NLDN	National Lightning Detection Network
NMM	Nonhydrostatic Mesoscale Model
NN	No-No (number of times an event was not predicted and did not occur)
Noah	Nested acronym referring to the four organizations that developed a community

	land-surface model: N=NCEP, O=Oregon State University (Dept. of Atmospheric Sciences), A=Air Force (AF Weather Agency and AF Research Lab), H=Hydrologic Research Lab, NWS. Current usage favors lower case “oah.”
NOTAM	NOtice To AirMen
NRL	Naval Research Laboratory
NSF	National Science Foundation
NSSL	National Severe Storms Laboratory
NTDA	National Turbulence Detection Algorithm
NTDA	NEXRAD Turbulence Detection Algorithm
NWP	Numerical Weather Prediction
NWS	National Weather Service
NY	No-Yes (number of events that occurred but were not predicted)
O_3	Ozone
OCP	Overshooting Center Pixel
OMI	Ozone Monitoring Instrument
OT	Overshooting Top
PIREP	PIlot REPort
POD	Probability Of Detection
POES	Polar Operational Environmental Satellite (NOAA)
PPI	Plan Position Indicator (most common weather radar display)
QICP	Qualified Internet Communication Provider
QPF	Quantitative Precipitation Forecast
RAL	Research Applications Laboratory (NCAR)
RAP	New acronym for Rapid Refresh (a prediction model); old acronym: RR
RAPT	Route Availability Planning Tool
RASS	Radio Acoustic Sounding System
RH	Relative Humidity
ROC	Receiver Operating Characteristic
RR	Old acronym for Rapid Refresh (a prediction model); new acronym: RAP
RRTM	Rapid Radiative Transfer Model (GFDL)
RSM	Regional Spectral Model (NCEP)
RTVS	Real-Time Verification System
RUC	Rapid Update Cycle
RVR	Runway Visual Range
SACS	Support to Aviation Control Service
SATCAST	SATallite Convection AnalySis and Tracking
SBUV	Solar Backscatter Ultraviolet Radiometer
SCD	Slant Column Density
SCIAMACHY	SCanning Imaging Absorption spectroMeter for Atmospheric CartograpHY
SIGMET	SIGnificant METeorological information
SLD	Supercooled Large Drops
SLW	Supercooled Liquid Water
SNR	Signal-to-Noise Ratio

SO_2	Sulfur Dioxide
SPC	Storm Prediction Center
SPoRT	Short-term Prediction Research and Transition center (NASA)
SREF	Short-Range Ensemble Forecast
SSD	Scientific Services Division (one in each NWS region)
SSEC	Space Science and Engineering Center (University of Wisconsin, Madison)
SSM/I	Special Sensor Microwave / Imager
STAR	SaTellite Applications and Research
STMAS	Space-Time Mesoscale Analysis System
SWA	SouthWest Airlines
T	temperature
TAF	Terminal Aerodrome Forecast
TAMDAR	Tropospheric Airborne Meteorological DAta Reporting (source of automated aircraft reports)
TDWR	Terminal Doppler Weather Radar
TITAN	Thunderstorm Identification, Tracking, Analysis, and Nowcasting
TKE	Turbulence Kinetic Energy
TMI	TRMM Microwave Image
TMU	Traffic Management Unit
TOA	Time Of Arrival
TP	Turbulence Plot
TRACON	Terminal Radar Approach CONtrol
TREC	Tracking clear-air Radar Echoes by Correlation
TRMM	Tropical Rainfall Measuring Mission
TSS	True Skill Statistic
U.S.	United States
UAF	University of Alaska Fairbanks
UAH	University of Alabama—Huntsville
ULB	Université Libre de Bruxelles
USGS	United States Geological Survey
USPLN	United States Precision Lightning Network
UTC	Coordinated Universal Time (acronym is a compromise between English and French translations)
UV	UltraViolet
UWCI	University of Wisconsin Convective Initiation (name of an algorithm)
VAAC	Volcanic Ash Advisory Center
VAD	Velocity-Azimuth Display
VCD	Vertical Column Density
VDRAS	Variational Doppler Radar Analysis System
VFR	Visual Flight Rules
VIL	Vertically Integrated Liquid
VIP	Video Integrator and Processor
VIIRS	Visible Infrared Imager Radiometer Suite

VIS	VISible
VLIFR	Very Low Instrument Flight Rules
VMC	Visual Meteorological Conditions
VORTAC	VHF (Very High Frequency) Omnidirectional Range / Tactical Aircraft Control
VV	Vertical Velocity
WAF	Weather Avoidance Field
WFO	Weather Forecast Office
WMO	World Meteorological Organization
WRF	Weather Research and Forecasting (name of a community mesoscale prediction model)
WSDDM	Weather Support to De-icing Decision Making
WSI	Weather Services International
WSP	Weather Systems Processor
WSR-88D	Weather Service Radar (1988 Doppler)
WVSS	Water Vapor Sensing System
YN	Yes-No (number of predicted events that did not occur)
YY	Yes-Yes (number of predicted events that did occur)

Appendix B. Technical Details

This appendix includes technical and mathematical details essential for developers of aviation weather products but not for the general reader. The numbering system in this appendix corresponds to that used in the main body of the text.

B2.1.5 Algorithms for Calculating Turbulence Indices Used in the GTG2 Product

Each algorithm computes a turbulence index from Rapid Refresh (RAP) model output. Most indices are divided by the Richardson number Ri, where $Ri = N^2/S_V^2$, N^2 is the Brunt-Väisälä frequency, a measure of stability in statically stable air, and S_V is the magnitude of the vertical shear of the horizontal wind vector. To be more specific, $N^2 = \frac{g}{\theta}\frac{\partial \theta}{\partial z}$ and $S_V^2 = \left(\frac{\partial u}{\partial z}\right)^2 + \left(\frac{\partial v}{\partial z}\right)^2$, where *g* is the acceleration due to gravity, θ is potential temperature, *z* is the vertical coordinate, and *u* and *v* are the eastward and northward components, respectively, of the horizontal wind. Here is a description of indices calculated from output of the Rapid Refresh (RAP) model. The RAP model is the successor to the RUC model; it is undergoing final checkout at EMC and should become operational in the spring of 2012.
At upper levels, FL 200 to FL 450:

1. Frontogenesis: F_θ. This index looks at the time rate of change $\frac{d}{dt}$ in the vertical shear of the horizontal vector wind as computed in isentropic coordinates (in which the vertical coordinate is *Θ*, the potential temperature). $F_\theta = \frac{d}{dt}\left|\frac{\partial \boldsymbol{v}}{\partial \theta}\right|^2$, where |…|^2 means vector magnitude, ***v*** is the horizontal wind vector with components *u* toward the east and *v* toward the north. The vertical shear of the horizontal wind is related to the horizontal temperature gradient, so its time rate of change is an appropriate measure of the strengthening or weakening of a frontal zone. The total derivative $\frac{d}{dt}$ in isentropic coordinates is $\frac{d}{dt} = \frac{\partial}{\partial t} + u\left(\frac{\partial}{\partial x}\right)_\theta + v\left(\frac{\partial}{\partial y}\right)_\theta + \frac{d\theta}{dt}\frac{\partial}{\partial \theta}$, where x is the eastward pointing coordinate, and y is the northward pointing coordinate. The expression for F_θ is expanded with the help of the equations of motion in isentropic coordinates. Neglecting the friction force, $\frac{du}{dt} = fv - \frac{\partial M}{\partial x}$ and $\frac{dv}{dt} = -fu - \frac{\partial M}{\partial y}$, where $M = c_pT + gz$ is the Montgomery stream function, c_p is the specific heat of air at constant pressure, f is the Coriolis parameter, T is temperature, g is the acceleration due to gravity, and z is the height of an

isentropic surface. The most important term in this expansion (Robert Sharman, personal communication) is rendered as the turbulence index, namely,

$$\frac{\partial \boldsymbol{v}}{\partial \theta} \cdot \left(\frac{\partial \boldsymbol{v}}{\partial \theta} \cdot \nabla\right) \boldsymbol{v} = \frac{\partial u}{\partial \theta}\left(\frac{\partial u}{\partial \theta}\frac{\partial u}{\partial x} + \frac{\partial v}{\partial \theta}\frac{\partial u}{\partial y}\right) + \frac{\partial v}{\partial \theta}\left(\frac{\partial u}{\partial \theta}\frac{\partial v}{\partial x} + \frac{\partial v}{\partial \theta}\frac{\partial v}{\partial y}\right).$$

All partial derivatives with respect to x and y are calculated on isentropic surfaces interpolated from the native grid coordinates.

2. Deformation: DEF^2/Ri. $DEF^2 = D_{SH}^2 + D_{ST}^2$, where $D_{SH} = \frac{\partial v}{\partial x} + \frac{\partial u}{\partial y}$ is the shearing deformation, and $D_{ST} = \frac{\partial u}{\partial x} - \frac{\partial v}{\partial y}$ is the stretching deformation.
3. Inertial Advective Wind (looks for turbulence in regions of unbalanced flow): v_i/Ri. $v_i = |\boldsymbol{v} \cdot \nabla \boldsymbol{v}|/f$, where $\boldsymbol{v}$ is the horizontal wind vector with components (*u,v*) and $\nabla = \boldsymbol{i}\frac{\partial}{\partial x} + \boldsymbol{j}\frac{\partial}{\partial y}$ is the horizontal gradient operator. Reference: McCann (2001).
4. Unbalanced flow: UBF/Ri. The starting point for this index is a wave equation (Eq. 1 in Knox et al. 2008). The right-hand side of this equation pertains to gravity wave generation and, by implication, turbulence. In continuously stratified flows, the right-hand side may be expressed as

$$UBF = \frac{\partial}{\partial t}(\nabla \cdot \boldsymbol{G}) + f\boldsymbol{k} \cdot \nabla \times \boldsymbol{G}$$

where $\boldsymbol{G} = \boldsymbol{v}\nabla \cdot \boldsymbol{v} + (\boldsymbol{v} \cdot \nabla)\boldsymbol{v}$, $\boldsymbol{v} = \boldsymbol{i}u + \boldsymbol{j}v$ is the horizontal wind vector, and $\boldsymbol{i}, \boldsymbol{j}$, and $\boldsymbol{k}$ are unit vectors in the eastward, northward, and vertical directions, respectively. The first term may be expanded as

$$\frac{\partial}{\partial t}(\nabla \cdot \boldsymbol{G}) = 2\left[\frac{\partial}{\partial t}(DIV)^2 + \frac{\partial}{\partial t}(\boldsymbol{v} \cdot \nabla(DIV)) - \frac{\partial}{\partial t}J(u,v)\right]$$

where $DIV = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}$ is the divergence, and $J(u,v) = \frac{\partial u}{\partial x}\frac{\partial v}{\partial y} - \frac{\partial u}{\partial y}\frac{\partial v}{\partial x}$ is the Jacobian determinant. The second term may be expressed as

$$f\boldsymbol{k} \cdot \nabla \times \boldsymbol{G} = 2fD\zeta + f\boldsymbol{v} \cdot \nabla\zeta + f\left(v\frac{\partial}{\partial x}DIV - u\frac{\partial}{\partial y}DIV\right)$$

where $\zeta = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}$ is the vertical component of relative vorticity. Viewed in this way, the two terms of UBF are plausibly related to turbulence because one associates large values of divergence and vorticity and large spatial gradients of these two quantities with stormy weather.

The algorithm computes UBF in several steps. First, it computes the local time derivatives of the u and v components of the wind from the horizontal equations of motion:

$\frac{\partial u}{\partial t} = -\frac{1}{\rho}\frac{\partial p}{\partial x} + fv - u\frac{\partial u}{\partial x} - v\frac{\partial u}{\partial y}$ and $\frac{\partial v}{\partial t} = -\frac{1}{\rho}\frac{\partial p}{\partial y} - fu - u\frac{\partial v}{\partial x} - v\frac{\partial v}{\partial y}$.

Second, it computes the divergence tendency from $\frac{\partial}{\partial t}DIV = \frac{\partial}{\partial x}\left(\frac{\partial u}{\partial t}\right) + \frac{\partial}{\partial y}\left(\frac{\partial v}{\partial t}\right)$.

Third, it computes $G_x = uDIV + u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y}$ and $G_y = vDIV + u\frac{\partial v}{\partial x} + v\frac{\partial v}{\partial y}$.

Fourth, it computes the second term of UBF: $f\boldsymbol{k}\cdot\nabla\times\boldsymbol{G} = f\left(\frac{\partial G_y}{\partial x} - \frac{\partial G_x}{\partial y}\right)$.

Fifth, it computes

$$\frac{\partial G_x}{\partial t} = DIV\frac{\partial u}{\partial t} + u\frac{\partial}{\partial t}DIV + \frac{\partial u}{\partial t}\frac{\partial u}{\partial x} + u\frac{\partial^2 u}{\partial x\partial t} + \frac{\partial v}{\partial t}\frac{\partial u}{\partial y} + v\frac{\partial^2 u}{\partial y\partial t}$$

$$\frac{\partial G_y}{\partial t} = DIV\frac{\partial v}{\partial t} + v\frac{\partial}{\partial t}DIV + \frac{\partial u}{\partial t}\frac{\partial v}{\partial x} + u\frac{\partial^2 v}{\partial x\partial t} + \frac{\partial v}{\partial t}\frac{\partial v}{\partial y} + v\frac{\partial^2 v}{\partial y\partial t}.$$

Finally, it computes the first term of UBF: $\frac{\partial}{\partial t}(\nabla\cdot\boldsymbol{G}) = \frac{\partial}{\partial x}\left(\frac{\partial G_x}{\partial t}\right) + \frac{\partial}{\partial y}\left(\frac{\partial G_y}{\partial t}\right)$.

5. Horizontal Divergence: $|DIV|/Ri$. $DIV = \nabla\cdot\boldsymbol{v} = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}$
6. Negative Vorticity Advection: NVA/Ri. $NVA = \max\{[-u\frac{\partial}{\partial x}(\zeta+f) - v\frac{\partial}{\partial y}(\zeta+f)],0\}$
7. Richardson Number Tendency, Φ. This index was developed by Brown (1973) starting from an expression for the Richardson number tendency $\frac{d}{dt}Ri$. The index is $\Phi = (0.3\,\zeta_a{}^2 + D_{SH}^2 + D_{ST}^2)^{1/2}$, where $\zeta_a = \zeta + f$ is the absolute vorticity.
8. Structure constant for temperature: C_T^2/Ri. Structure functions are the basis for several turbulence indices. In one dimension, a structure function D for a variable q is defined by $D_q(s) = \langle[q(x) - q(x+s)]^2\rangle$, where $\langle\ldots\rangle$ means an ensemble average. For homogeneous turbulence, which is assumed here, the statistics do not depend upon the location x, rather only upon the distance s separating the two points where q is measured. In turbulent air, one would expect large differences in q over short distances, especially if q happens to be the u or v component of the wind or temperature T. If the differences are generally large, then the structure function will be large and presumably correlated with the occurrence of turbulence. A structure function is computed at each grid point on a given model level as follows.

   For the arbitrary variable q, the two-dimensional grid indices (i,j) give the location. Thus $q(i,j) = q(i\Delta x, j\Delta y) = q(x,y)$. In practice, the structure function is computed along each coordinate axis of the model grid. In the x direction, for each (i,j),

$$D_q(l,0) = \frac{1}{(2N_i + 1 - l)}\sum_{p=i-N_i}^{i+N_i-l}[q(p,j) - q(p+l,j)]^2$$

   The separation distance (also called the lag) in the x direction is l grid points, a distance of $l\Delta x$. In the y-direction, the lag is zero. A similar equation exists for the structure function along the y-axis, where l would represent a separation distance of $l\Delta y$. In practice, $N_i =$

$N_j = 2$, and l ranges from 1 to $2N_i = 4$. This ensures proper behavior of the summation index p. With these conventions, one can rewrite the above equation to get the structure function in each direction. For each (i, j),

$$D_q(l,0) = \frac{1}{5-l}\sum_{p=i-2}^{i+2-l}[q(p,j) - q(p+l,j)]^2$$

$$D_q(0,l) = \frac{1}{5-l}\sum_{p=j-2}^{j+2-l}[q(i,p) - q(i,p+l)]^2$$

These two calculations are combined to form the model structure function for the variable q at each grid point (i, j). $D_{q-model}(l) = 0.5[D_q(l,0) + D_q(0,l)]$

Frehlich and Sharman (2004) suggested an empirical equivalent to the model-derived structure function immediately above: $D_{q-model}(s_k) = K_q D_{cor}(s_k) D_{q-ref}(s_k)$, where l and s_k are equivalent, l representing the separation in terms of grid points and s_k in terms of actual distance. K_q is a constant that depends upon the variable q. If q is temperature T, then $K = C_T$, the turbulence parameter to be calculated. D_{ref} is a reference structure function suggested by Lindborg (1999). It has the form

$$D_{T-ref}(s) = a_4 s^{2/3} + b_4 s^2 - c_4 s^2 \ln s,$$

where a_4, b_4, and c_4 are constants given in Frehlich and Sharman (2004). D_{cor} is a correction function that accounts for smoothing of model fields at the smallest resolvable scales. It does not depend upon the variable q.

$$D_{cor}(s) = \frac{\left(\frac{s}{p_1}\right)^{4/3}}{1 + \left(\frac{s}{p_1}\right)^{4/3} + p_2\left(\frac{s}{p_1}\right)^{2/3}}$$

p_1 depends upon the length scale of the model-specific filter; p_2 is a fitting parameter. Both constants are fixed for a given model.

It was mentioned above that $K = C_T$ is a constant. As explained in Frehlich et al. (2010), the best value of K, for the values of s_k considered in practice, is $K = C_T = NUM/DEN$, where

$$NUM = \sum_{l=1}^{4}\frac{D_{T-model}^2(l)}{l D_{cor}^2(s_l) D_{T-ref}^2(s_l)}$$

$$DEN = \sum_{l=1}^{4} \frac{D_{T-model}(l)}{lD_{cor}(s_l)D_{T-ref}(s_l)}$$

Again, l is the lag in terms of grid distances, and s_l is the lag in terms of actual distance. $D_{T-model}(l) = 0.5[D_T(l,0) + D_T(0,l)]$, the same equation as earlier, except T replaces q.

9. Eddy Dissipation Rate: $\varepsilon^{1/3}/Ri$. As discussed earlier, the eddy dissipation rate is ε, but the automated aircraft reports of turbulence give $\varepsilon^{1/3}$ in units of $m^{2/3}s^{-1}$. As for C_T, the $\varepsilon^{1/3}$ measure of turbulence comes from the calculation of structure functions on a model grid, but in this case for the horizontal components of the wind, u or v. For example,

$$D_{u-model}(l) = 0.5[D_u(l,0) + D_u(0,l)]$$

is computed for lags $l = 1,2,3,4$. These computations are compared with the an empirical model: $D_{u-model}(s_k) = K_u D_{cor}(s_k) D_{u-ref}(s_k)$, where $K \sim 2\varepsilon^{2/3}$ (which provide a means for calculating $\varepsilon^{1/3}$ once K is determined), $D_{cor}(s_k)$ is the same as before, and

$$D_{u-ref}(s_k) = s^{2/3} + \frac{b_1}{a_1}s^2 - \frac{c_1}{a_1}s^2 \ln s$$

This has the same form as $D_{T-ref}(s)$, but the constants are different. See Frehlich and Sharman (2004) for the actual values. K_u is determined as before but with $D_{T-model}$ replaced by $D_{u-model}$ and D_{T-ref} by D_{u-ref}. From K it is easy to obtain $\varepsilon^{1/3}$.

An alternative way to calculate $\varepsilon^{1/3}$ is to use structure functions of the v-component of the wind. $K_v \sim \frac{10}{3}\varepsilon^{2/3}$ may be calculated from $D_{v-model}(s_k) = K_v D_{cor}(s_k) D_{v-ref}(s_k)$, where

$$D_{v-model}(l) = 0.5[D_v(l,0) + D_v(0,l)]$$

$$D_{v-ref}(s_k) = a_3 s^{2/3} + b_3 s^2 - c_3 s^2 \ln s$$

The constants a_3, b_3, and c_3 are related to a_1, b_1, and c_1. Their values are given in Frehlich and Sharman (2004).

The average of the two $\varepsilon^{1/3}$ values (one computed for the u-component of the wind, one for v) becomes the final estimate.

At mid levels, 10,000 ft MSL to FL 200

The same calculations as described above for upper levels are also used at mid levels with the exceptions of Index 1 and Index 4. In addition, the following three algorithms are used.

10. Frontogenesis in constant height coordinates: $F2D/Ri$. This index is very similar to Index 1, but the computation proceeds in constant height (z) coordinates instead of isentropic coordinates. The relevant expression is

$$F2D = \frac{\partial \boldsymbol{v}}{\partial z} \cdot \left(\frac{\partial \boldsymbol{v}}{\partial z} \cdot \nabla \right) \boldsymbol{v} = \frac{\partial u}{\partial z} \left(\frac{\partial u}{\partial z} \frac{\partial u}{\partial x} + \frac{\partial v}{\partial z} \frac{\partial u}{\partial y} \right) + \frac{\partial v}{\partial z} \left(\frac{\partial u}{\partial z} \frac{\partial v}{\partial x} + \frac{\partial v}{\partial z} \frac{\partial v}{\partial y} \right).$$

11. Predictor of CAT from Model Output Statistics: NGM/Ri. Reap (1996) found that this predictor, computed from the output of NCEP's Nested-Grid Model, had skill in diagnosing clear-air turbulence. This index is simply the product of the magnitude of the horizontal wind times the deformation DEF as defined for Index 2.

$$NGM = (u^2 + v^2)^{1/2} DEF$$

12. Horizontal temperature gradient: $|\nabla_H T|/Ri$. The horizontal temperature gradient is related to the vertical shear of the horizontal wind through the thermal wind approximation. When the shear is strong, turbulence is likely. Airline forecasters use this index frequently.

$$|\nabla_H T| = \left[\left(\frac{\partial T}{\partial x} \right)^2 + \left(\frac{\partial T}{\partial y} \right)^2 \right]^{1/2}$$

B2.1.7. Calculations Relating to the NEXRAD Turbulence Detection Algorithm (NTDA)

This section summarizes the NTDA algorithm described in more detail in Williams et al. (2006).

The first step in processing 88-D (also called NEXRAD) radar data is to remove contaminated data. When the radar is pointed directly at the sun, increased reflectivity, noisy radial velocities, and large spectrum widths occur. These are removed whenever the radar points within a degree or two of the sun. Another artifact consists of four or more repeated values of radial velocity and spectrum width at adjacent range bins. These may occur in several adjacent azimuths but are fairly easy to remove.

The next step is to determine a confidence value for each spectrum width measurement. The overall confidence value C_{SW} is the product of five separate confidence values, each ranging from zero to one, zero implying no confidence and one implying high confidence.

- The signal-to-noise ratio (SNR) is calculated solely as function of range.

$$SNR(r) \approx dBZ(r) + 20\log_{10}\left(\frac{230}{r}\right)$$

dBZ is the reflectivity in decibels, and the range is expressed in kilometers. The confidence value C_{SNR} rises linearly from 0.0 at 10 dBZ to 1.0 at 20 dBZ. Between 20 and 70 dBZ, C_{SNR} is 1.0. Between 70 and 80 dBZ, it drops linearly back to zero. There is little confidence in radial velocity or spectrum width values when reflectivity is low because small hydrometeors in clouds backscatter little more energy than clear air. On the other hand, very high values of SNR also decrease confidence because they suggest contamination by non-atmospheric sources, including ground clutter.

- Second-trip echoes: The pulse repetition frequency of the radar determines the maximum distance a transmitted pulse can travel, have some of its energy scattered back toward the radar, and be received at the antenna before the next pulse is emitted. This is called the unambiguous range R_{max} of the radar. If strong echoes lie beyond R_{max}, energy can be scattered back toward the radar and be received at the same time as backscattered energy from much nearer targets. These are called second- or multiple-trip echoes. If the returned power at range $r < R_{max}$ is not significantly larger than that from multiple-trip echoes beyond R_{max}, the confidence factor C_{PR} quickly falls from 1.0 to 0.0.

- Spectrum width variability: This confidence measure C_{SWV} depends upon the variability of the spectrum width in neighboring pulse volumes. If the spectrum width shows high variance in the neighborhood, the quality of the measurement becomes suspect, and C_{SWV} is reduced linearly from 1.0 at a threshold of 4 m^2s^{-2} to 0.0 at the value 16 m^2s^{-1}.

- Clear air returns close to the ground: In summertime, especially, high values of spectrum width are often observed close to the ground in "clear-air" conditions, when reflectivity values are low. These high values are most probably the result of flying insects rather than turbulent air, and so a confidence factor C_{ZH} is designed to filter out these events. The sum of the returned power in decibels and 3.5 times the height in km of the radar beam above ground is computed. For example, a reflectivity of 20 dBZ at a height of 1 km above ground would give a sum of 23.5. The height of the radar beam is relevant because the concentration of insects drops as height above ground increases. If the sum is less than 15, the confidence factor C_{ZH} is zero. Between values of 15 and 25, C_{ZH} increases linearly from 0.0 to 1.0; it remains at 1.0 for values above 25.

- Effect of radar range: Half the radiated power of the 88-D radar lies in a narrow cone subtending an angle of about 1°. Because the beam broadens with distance from the radar,

the effect of vertical shear on the spectrum width becomes increasingly important with increasing range, so that there is less confidence in spectrum width as a measure of turbulence. The confidence factor C_{RNG} accounts for this. C_{RNG} rises from 0.0 to 1.0 at very short ranges (0-5 km), stays at 1.0 from 5 to 140 km, then decreases back to 0.0 between 140 and 275 km.

The single confidence value for spectrum width C_W ranges from 0.0 to 1.0. It is the product of the five confidence values just discussed: $C_W = C_{SNR}^{2/3} C_{PR} C_{SWV} C_{ZH} C_{RNG}$

For any spectrum width measurement $SW(r)$ with a nonzero confidence value, a raw eddy dissipation rate is calculated from

$$EDR_{raw}(r) = \varepsilon^{1/3} = SW(r) f(r, L_0)$$

where $f(r, L_0)$ is a scaling function depending upon radar range and certain assumptions about the kind of turbulence contained within a radar pulse volume. L_0 is called an outer length scale for turbulence. Values of f for L_0 = 500 m and r ranging from 0 to 250 km are shown in Fig. 2 of Williams et al. (2006). The value of f is 0.23 at r=0; at first, it decreases rapidly with range, then more slowly, reaching a value of 0.115 at r=200 km.

As a means of reducing noise in the EDR_{raw} field, a local average is taken, and each value weighted by its confidence factor C_W.

B2.2.10 Computing Icing Severity

The flow chart below, Fig. B1, repeats Fig. 2.28 in the main text. Each step in the flow chart on the way to computing final icing potential is described in this section.

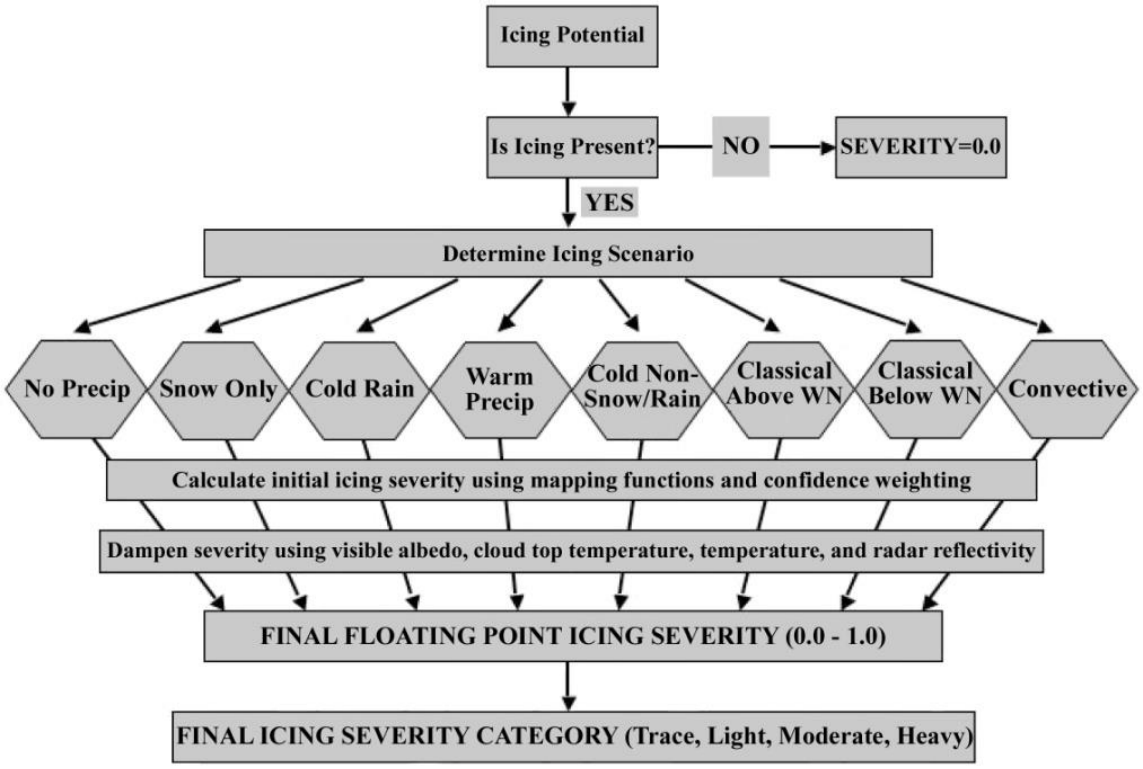

Fig. B1. Flow chart for generation of the Icing Severity Index. From Politovich et al. (2006)

Two definitions will be helpful prior to the description of the flow chart. First, reflectivity values (dBZ) at 4-km resolution from the WSR-88D radar are inserted within the model grid boxes. When ordered by value from greatest to least, the value one quarter of the way down from the top is called the 75th percentile value; similarly, the value three quarters of the way down is the 25th percentile value. Second, a *warm nose* refers to a layer of above-freezing air sandwiched between two sub-freezing layers.

Step 1. If the icing potential is less than 0.01, then no icing severity is calculated.
Step 2. If the icing potential $\geq$ 0.01, then precipitation at the surface and vertical profiles of temperature, humidity, and SLW are examined to see which of the scenarios below is most likely. If severity is being estimated for the current time, observations are also consulted.

Icing Scenarios

A Non-precipitating cloud: All evidence points to no precipitation.

B Snow only: The only precipitation type identified is snow.

C Cold rain: CTT < -12^{o}C *and*
either rain is a precipitation type

or the 75th percentile of reflectivity from 4-km 88D radar observations ≥ 5 dBZ.

D "Warm" precipitation: CTT ≥ -12°C *and*
either the precipitation type is not snow
or the 75th percentile reflectivity is ≥ 5 dBZ.

E **P**recipitation type is identified as drizzle, freezing drizzle, freezing rain, or ice pellets and the CTT < -12°C.

F Precipitation above the warm nose: CTT < -12°C, a warm nose is present, the precipitation type is not all snow, and the level in question is above the warm nose.

G Precipitation below the warm nose: CTT < -12°, a warm nose is present, the precipitation type is not all snow, and the level in question is below the warm nose.

H Deep convection: The distance to the lightning strike is ≤ 25 km.

Associated with these eight scenarios are 22 parameters, each with its own interest map, similar to those shown in Fig. 2.23.

Step 3. Calculate the initial value of icing severity. The output from a given interest map, say map k, is a number i_k between 0.0 and 1.0; it is assigned a weight w_k and a confidence value c_k, the latter indicating how much the information source can be trusted. All confidence values are currently set to 1.0 except for satellite and PIREP data. The general form of the initial severity equation is

$$SEV_{init} = \frac{\sum_{k=1}^{n} c_k w_k i_k}{\sum_{k=1}^{n} c_k w_k},$$

where n is the number of parameters. If a particular interest map is not applicable to a given scenario, then its term is simply dropped from the sum in both numerator and denominator. The denominator ensures that the weights are normalized (their sum is one). Here is a list of the 22 parameters currently considered.

- "Satellite Combo." GOES Channels 1, 2, and 4 are used, the same channels used to compute icing potential. Visible albedo is useful only for solar zenith angles less than 70° and virtually useless otherwise. The confidence factor goes from 1.0 to 0.0 as the solar zenith angle goes from 70° to 75° (transition to nighttime). The difference between brightness temperatures in Channel 2 and Channel 4 gives information on cloud-top phase, with larger values supporting the existence of small droplets and smaller values supporting the existence of large ice crystals. An interest function called "satellite combo" is formed from the product of two interest functions, one for the visible albedo and the other for the difference in brightness temperatures.
- Model Vertical Velocity

- Model RH + Condensate. This parameter combines the RH interest map RH_{map} with an interest map for liquid and/or ice condensate (g cm^{-3}), CL_{map} or CI_{map}, to form a hybrid moisture interest map m_{map}.

$$m_{map} = \frac{RH_{map}}{2} + \left(1 - \frac{RH_{map}}{2}\right)\left(0.6CL_{map} + 0.4CI_{map}\right)$$

  The interest maps for liquid and ice condensate are the same.
- Maximum Potential Liquid Water Content (LWC). This parameter is calculated on the assumption that moist adiabatic ascent is occurring within the cloud and that no condensate is lost to precipitation. At cloud base, there is no condensate, and the saturation mixing ratio w_s measures the vapor content of air parcels entering the cloud. The higher the parcel ascends, the cooler it becomes and the more vapor condenses. At any level within the cloud z, the condensate is just the difference between the saturation mixing ratio at cloud base $w_s(z_{bot})$ and that at level z, namely, $w_s(z)$.

$$w_s(z) = \frac{0.622e_s}{p(z) - e_s},$$

  where p is total pressure and e_s is the saturation vapor pressure, which depends only upon temperature T. There are many accurate formulas for e_s. Among the simpler ones is this equation from Bolton (1980):

$$e_s = 6.1121\ exp\left(\frac{17.67T}{T + 243.5}\right),$$

  where T is expressed in Kelvin degrees, and e_s is in millibars. Because temperature decreases upward when the lapse rate is moist adiabatic, this guarantees that $w(z_{bot}) - w(z)$ increases upward. If the condensate is supercooled water, the prospect of icing increases upward as well. The in-cloud lapse rate will not be moist adiabatic if precipitation is falling, if dry air is mixing into the updraft from the sides of the cloud, or if a stable layer (perhaps a frontal inversion) lies within the cloud.
- Icing Potential (discussed at length in Section 2.2.7)
- Supercooled Large Droplet (SLD) Potential (discussed at length in Section 2.2.7). Used only for Scenario D.
- PIREP Severity. For each grid point, all PIREPs within 200 km horizontally and 4,000 ft vertically are considered. The value v of a PIREP ranges from 0.0 to 1.0, decreasing from 1.0 once it is more than 15 min old, 30 km distant from the grid point horizontally, and 1000 ft vertically. The severity s of reported icing also ranges from 0.0 (no icing) to 1.0 (severe icing). The severity assigned to each grid point is a weighted sum

$$rep\_sev = \frac{v_1 s_1 + v_2 s_2 + \cdots}{v_1 + v_2 + \cdots}$$

The confidence value for PIREP severity is related to the sum of PIREP values in the denominator. A substantial number of nearby PIREPs gives a confidence of 1.0. A few distant PIREPs would get a low confidence.

- Total Water Path. This parameter is calculated from model values of condensate, both solid and liquid, for each cloud layer. Here, two or more consecutive model layers with non-zero condensate define a cloud layer. The model computes condensate in g m^{-3}, and this must be multiplied by the depth (m) of the model layer to get the water path for that layer (g m^{-2}). The path lengths in each consecutive model layer with condensate must be added to get the path length in each cloud layer.
- Total Water Path for Scenario D (warm precipitation). Because LWC is likely to form a larger fraction of the total condensate in warmer clouds than in colder clouds, the severity of icing will be greater in Scenario D than in other scenarios for the same total water path.
- Cloud Thickness. From model data, the icing algorithm computes the height above cloud base. As noted in the discussion of maximum potential LWC, the higher an aircraft is above cloud base, the greater the potential for icing, unless the cloud top is very cold, in which case, the hydrometeors are ice crystals. Precipitation forming within the cloud will either deplete or augment SLW depending upon the type. Snow collects tiny supercooled cloud droplets during collisions with them. The droplets freeze to the snow crystals or flakes in a process called riming, thereby reducing the icing hazard. On the other hand, freezing rain falling below a warm nose can collect cloud droplets, thereby increasing the threat of icing by SLD. Five interest maps relate to cloud thickness, depending upon whether there is no precipitation (Scenario A), warm precipitation (Scenario D), snow only (Scenario B), cold rain (Scenario C), or freezing rain at the ground (Scenario G).
- Temperature Below Warm Nose. Used only with Scenario G. If the temperature below the warm nose is above freezing, there is no icing. Between 0^{o} and -4^{o}C, the interest map value rises steeply from 0.0 to 1.0 and remains at 1.0 for temperatures lower than -4^{o}C.
- Temperature with Convection. Used only with Scenario H. In convective clouds, icing is likely at considerably lower temperatures than in stratiform clouds because strong updrafts raise supercooled droplets to great heights. The interest map reflects this.
- 75th Percentile, Reflectivity Spectrum. Used only with Scenario G. As noted earlier, 4-km WSR-88D radar reflectivity data (in decibels, dBZ) are mapped to the model grid. Since the grid is much coarser than the radar data, a multitude of dBZ values lies in each grid box. The 75th percentile dBZ value represents some of the higher reflectivity values when precipitation, not all snow, is falling below the warm nose. The interest map value rises linearly from 0.0 to 1.0 as reflectivity rises from 15 to 35 dBZ.
- Effect of CTT Below Warm Nose. Used only with Scenario G. The reasoning goes like this. The colder the cloud top above the warm nose, the deeper the cloud generating the

precipitation that falls through the warm nose, and the more likely that this precipitation will be substantial, leading to freezing rain in the cold air below the warm nose.

- Vertical Velocity (VV) Near the -15^{o}C isotherm. Used only with Scenario G. At in-cloud temperatures near -15^{o}C, the growth of large, dendritic (delicate and star-like) crystals is favored. If VV is significant, these crystals grow quickly. When they fall through the warm nose, they melt into sizeable drops, which contribute to freezing rain below the warm nose.
- Distance from Lightning. Used only with Scenario H. Lightning less than 15 km away receives an interest map value of 1.0. For lightning between 15 and 25 km away, the value drops from 1.0 to 0.4, then abruptly to 0.0 beyond 25 km.
- Difference between 75th and 25th Percentile Values of Reflectivity. Used only with Scenario H. If the difference is large, new and vigorous convection is suggested within the grid box and interest map values are raised.
- Number of Lightning Strokes in the Past 15 min. Used only with Scenario G. Strikes must be within 25 km of the grid point. As the number of strikes within 15 min rises from one through ten, the interest values rise from 0.0 to 1.0

Space does not permit displaying the many interest maps described above, but all of them are contained within an internal NCAR Report by Politovich et al. (2006).

Many of the 22 parameters mentioned above are used in only one scenario, but they are all assigned a weight, as described at the beginning of Step 3. Table B1 shows which parameters are used for each scenario and gives the weights assigned to each.

Table B1

Weights assigned to each icing parameter for Scenarios A through H described in Step 2.

Scenario	Satellite combo	Vertical velocity	Model RH + condensate / m^3	Max potential LWC	Icing Potential	SLD Potential	PIREP severity	Total water path (TWP)	TWP, warm precip	Cloud thickness (**Δz**), no precip	**Δz**, precip all snow	**Δz**, precip cold rain	**Δz**, warm precip	**Δz**, precip, freezing rain	T below warm nose	T with convection	75^{th} percentile, dBZ spectrum	CTT below warm nose	Vertical velocity near −15°	Distance from lightning	Difference, $75^{th} - 25^{th}$ percentile, dBZ	# of lightning strikes past 15 min
A	5	4	4	3	3		5			3.5												
B	4	3.5	4		3		5	3			3.5											
C	4	3.5	4		2		5	3				3.5										
D	4	3	3.5	2		4	5		3				3									
E	4	3.5	4		2		5	3				3.5										
F	4	3.5	4		2		5	3														
G			3.5		2		5							3	3		4	3	2			
H	3															3				3.5	2	5

The weights in Table B1, the confidence values (most of them 1.0), and the interest map values described above for each parameter are combined in the formula given earlier to obtain the initial icing severity, a number that ranges from 0.0 to 1.0.

Step 4. Refer again to the flow chart in Fig. 2.28. Step 4 acts to reduce the initial icing severity, just calculated. As Table B1 indicates, there is very little direct effect of temperature on the icing severity calculation except through "Icing Potential" (which receives a fairly low weight, mostly 2), "T below warm nose," and "T with convection." To compensate for this, wherever the cloud is cold (T < -12°C), the icing potential is adjusted downward. Let the adjustment be T_{adj}. The maximum adjustment is for T < -25°C, but cannot exceed 0.5 of the initial severity.

The downward adjustment in icing severity due to CTT is allowed only near cloud top. It is phased out with increasing distance downward from cloud top because the likelihood of ice decreases as one descends into the cloud. The phase out CT_{adj} is complete 10,000 ft down from cloud top. The maximum adjustment cannot exceed 0.2 of the initial severity.

Similarly, the downward adjustment in icing severity due to reflectivity is phased out with increasing distance upward from cloud base. The phase out CB_{adj} occurs beginning at 1,000 ft above cloud base and is complete at 5,000 ft above cloud base. The reduction in the initial icing severity near cloud base cannot exceed 0.15 of the initial icing severity, and that only occurs if

$CB_{adj} = 1.0$. CB_{adj} depends not only on the height above cloud base but also on the radar reflectivity at the 25th percentile. It is maximized for dBZ > 30 dBZ; it is 0.0 for dBZ < 5. The rationale is that higher precipitation rates act to decrease the probability of icing in most scenarios.

The downward adjustment to the initial value of icing severity is computed from

$$SEV_{init2} = SEV_{init} - T_{adj} - CT_{adj} - CB_{adj}$$

A final adjustment follows only in daytime, when the solar zenith angle is less than 70°. Clouds with low visible albedo are likely to be thinner than clouds with high visible albedo, and the number density of hydrometeors, liquid and solid, is probably less near cloud top. For this reason, the estimate SEV_{init2} is reduced once more unless the visible cloud albedo is greater than 0.65. The reduction VIS_{adj} is greatest for an albedo of .25 or less but in no case may it exceed 0.75 of $SEV_{init2.}$

$$SEV_{final} = SEV_{init2} - VIS_{adj}$$

Step 5. Convert continuous values of Icing Severity (the scale runs from 0.0 to 1.0) to categories of icing severity. See Table B2.

Table B2
Correspondence between icing categories and icing severity.

Category	Range, Icing Severity
Trace	$0.01 \leq SEV_{final} \leq 0.175$
Light	$0.175 < SEV_{final} \leq 0.375$
Moderate	$0.375 < SEV_{final} \leq 0.7$
Heavy	$0.7 < SEV_{final} \leq 1.0$

B2.3.6. Computing Motion Vectors with the TITAN Algorithm

Following application of the elliptical filter shown in Fig. 2.52, the VIL grid, defined in Cartesian coordinates, is examined for contiguous non-zero VIL values. Each blob of contiguous non-zero VIL values is identified as a storm. The area of the storm A is calculated from the number of grid boxes within the blob times the grid box area. If this area is less than 520 km^2, TITAN eliminates it. Otherwise, it calculates a VIL-weighted centroid (x_c, y_c) from

$$x_c = \frac{1}{\Sigma VIL}\sum_{i=1}^{n} VIL_i x_i, \qquad y_c = \frac{1}{\Sigma VIL}\sum_{i=1}^{n} VIL_i y_i, \qquad \Sigma VIL = \sum_{i=1}^{n} VIL_i$$

where VIL_i is the VIL value at (x_i, y_i) and the summation includes all n points in the area.

Each time radars produce a new full-volume scan, there will be a new set of storms to contend with. Since the TITAN algorithm tracks storms, it must decide which of the storms at a later time t_2 correspond to the storms identified from the previous scan at time t_1. Consider the case where there are n_1 storms at t_1 and n_2 storms at t_2, where n_1 may or may not equal n_2.

TITAN must consider all possible paths connecting the storms at t_1 with those at t_2 and choose the most reasonable ones. In solving this problem, TITAN makes the following intuitive assumptions. 1) Because only a few minutes elapse between t_1 and t_2, the correct set of paths connecting storms will be shorter rather than longer. 2) The correct set of paths will join storms occupying similar areas. 3) There is an upper limit to the distance a storm can move between t_1 and t_2. Paths greater than this distance should not be allowed.

TITAN solves an optimization problem that minimizes a cost function incorporating the three assumptions above. Suppose the state $\boldsymbol{S}$ of storm j at time t_1 is defined by its VIL-weighted centroid and its area: $\boldsymbol{S}_j(x_{cj}, y_{cj}, A_j)$. The state of storm k at time t_2 is similarly defined: $\boldsymbol{S}_k(x_{ck}, y_{ck}, A_k)$. The "cost," in units of distance, of connecting storm j to storm k is defined by

$$C_{jk} = d_p + d_v,$$

where d_p is the distance between the centroids of the two storms

$$d_p = \left[(x_{cj} - x_{ck})^2 + (y_{cj} - y_{ck})^2\right]^{1/2}$$

(This incorporates assumption 1), and d_v is the absolute value of the difference between the square roots of their areas.

$$d_v = \left|A_j^{1/2} - A_k^{1/2}\right|$$

(This incorporates assumption 2). Note that $A^{1/2}$ has units of distance. To incorporate assumption 3, C_{jk} is set to a large number if the speed of movement between storms j and k is greater than s_{max}, say 60 km h^{-1}, that is, if

$$\frac{d_p}{t_2 - t_1} > s_{max}.$$

TITAN seeks a match between storms at t_1 and t_2 that minimizes the cost function

$$Q = \sum_{j=1}^{n_1}\sum_{k=1}^{n_2} C_{jk}.$$

The solution to this problem is a number of paths connecting storms at t_1 with those at t_2, with the number not exceeding the minimum of n_1 and n_2. The solution is far from trivial; it involves combinatorial optimization. It is complicated by new storms forming and old storms dissipating between t_1 and t_2. In other words, some storms at t_1 or t_2 may not have a match. TITAN deals with further complications caused by merging and splitting storms, but the extra

logic required for addressing these problems is too much to discuss here. Together, the length and orientation of the optimum connecting paths give storm motion vectors. In practice, the motion vectors are computed over *two* volume scans rather than one in order to minimize sometimes erratic, very short-term motions.

B2.3.8.4 Computations Related to the Merging of Extrapolation and Model Forecasts of Convection

Computing Equivalent Radar Reflectivity from Model Mixing Ratios

A simple form of the gamma distribution for hydrometeor size is usually assumed for all three hydrometeor types (rain, snow, and hail):

$$N(D) = N_0 \exp(-\lambda D)$$

$N(D)$ is the number of hydrometeors of a given type per unit volume having a diameter within an infinitesimal size interval centered on D; N_0 is called an intercept parameter, a constant in the simplest case; and λ is called the slope factor, to be determined from the model mixing ratios. When N_0 is considered constant, it takes typical values of 8 x 10^6, 2 x 10^7, and 4 x 10^6 m^{-4} for rain, snow, and graupel, respectively.

If the hydrometeors are spherical raindrops, the water content of these raindrops per cubic meter is

$$\rho_d q_r = \int_0^\infty N(D)\rho_l \left(\frac{4}{3}\pi r^3\right) dD \, .$$

ρ_d is the density of dry air in kg m^{-3}. $q_r = \rho_l/\rho_d$ is the rainwater mixing ratio in kilograms of liquid water per kilogram of dry air, and ρ_l = 10^3 kg m^{-3} is the density of liquid water. Substituting for the mixing ratio, the drop radius $r = D/2$, and the drop size distribution in the above equation, one obtains

$$\rho_d q_r = \frac{\pi N_0 \rho_l}{6} \int_0^\infty D^3 e^{-\lambda D} dD \, .$$

The left-hand side is known from model output, and everything on the right-hand side is known except λ. The integral may be evaluated and the resulting equation solved for λ. The result is

$$\lambda = \left(\frac{\pi N_0 \rho_l}{\rho_d q_r}\right)^{\frac{1}{4}}$$

Now that λ is known, the reflectivity factor Z_r may be calculated from the sixth moment of the raindrop size distribution (Doviak and Zrnic, 1984, pp. 184-191).

$$Z_r = \int_0^\infty D^6 N(D)\, dD = N_0 \int_0^\infty D^6 e^{-\lambda D} dD$$

This formula is valid for Rayleigh scattering, when the droplet diameter is much smaller than the radar wavelength (10 cm in the case of WSR-88D radars). A good rule of thumb is that the droplet diameter should be less than 1/16 of the radar wavelength. The integral on the right has the value of $\Gamma(7)\lambda^{-7}$, where Γ is the gamma function. For integers m, $\Gamma(m) = (m-1)!$ so that $\Gamma(7) = 6 \cdot 5 \cdot 4 \cdot 3 \cdot 2 \cdot 1 = 720$. Finally, $Z_r = \Gamma(7) N_0 \lambda^{-7}$, with λ defined as above. Because N_0 has units of [m^{-4}] and λ has units of [m^{-1}], Z has units of [m^3]. Z is often expressed in the equivalent units [$(10^3 \text{ mm})^6 \text{ m}^{-3}$] to bring out the physical meaning of a concentration of millimeter-sized drops per unit volume contributing to the reflectivity.

The radar equation (Doviak and Zrnić, 1984, p. 58) relates the returned power (energy per second backscattered to the radar by targets in the pulse volume) to the distance of the target, properties of Rayleigh scattering, engineering aspects of the radar, and the reflectivity factor. Whenever Rayleigh scattering is not a good approximation, however, an equivalent reflectivity factor Z_e is substituted for the reflectivity factor in the radar equation. Z_e, expressed in decibels (dBZ), is the quantity shown on conventional radar displays, and thus it is the preferred quantity to calculate from model output to portray convective activity. For raindrops, $Z_r = Z_{er}$.

Stoelinga (2005, personal communication) gives formulas for calculating equivalent reflectivity factors for snow Z_s and graupel Z_g. For Rayleigh scattering, the radar cross section of a snow crystal or an aggregate of crystals is the same as that of a solid ice sphere of the same mass. As noted earlier, for the purpose of calculating Z_e, raindrops, snow crystals or flakes, and graupel, are all assumed to be spherical, but their densities differ significantly. The density of liquid water was already noted: $\rho_l = 1000$ kg m^{-3}. The densities for snow and graupel are, approximately, $\rho_s = 100$ kg m^{-3} and $\rho_g = 400$ kg m^{-3}. The density of solid ice is $\rho_i = 917$ kg m^{-3}. If the diameter of a snow hydrometeor is D_s and the diameter of a sphere of solid ice with equivalent mass is D_i , then the two diameters are related by

$$D_i = D_s(\rho_s/\rho_i)^{1/3}$$

The sixth power of this relationship appears in the calculation of the equivalent reflectivity factor for snow Z_{es}, and so the constant factor $(\rho_s/\rho_i)^2$ is part of the answer. In addition, the reflective capacity for ice is less than that for liquid water by a factor that is the ratio of two dielectric constants, that for ice ($|K_i|^2 = 0.176$) to that for liquid ($|K_l|^2 = 0.930$). With these facts in mind, the equivalent reflectivity factor for snow is

$$Z_{es} = \Gamma(7) N_0 \lambda^{-7} \left(\frac{\rho_s}{\rho_i}\right)^2 \frac{|K_i|^2}{|K_l|^2} \quad , \qquad \text{where } \lambda = \left(\frac{\pi N_0 \rho_s}{\rho_d q_s}\right)^{\frac{1}{4}}$$

now includes constants appropriate for snow, including N_0 (2 x 10^7 m^{-4}). q_s is the mixing ratio for snow, obtained from the model.

Similar arguments lead to the equivalent reflectivity factor for graupel.

$$Z_{eg} = \Gamma(7) N_0 \lambda^{-7} \left(\frac{\rho_g}{\rho_i}\right)^2 \frac{|K_i|^2}{|K_l|^2} \ , \qquad \text{where } \lambda = \left(\frac{\pi N_0 \rho_g}{\rho_d q_g}\right)^{\frac{1}{4}} .$$

Here, $N_0 = 4 \times 10^6$ m^{-4}, the value appropriate for graupel, and q_g is the mixing ratio for graupel from the model.

Should more than one type of hydrometeor be found in the same volume, the reflectivity factors are additive, that is, if the model computes non-zero mixing ratios for rain, snow, and graupel at the same location, then $Z_e = Z_{er} + Z_{es} + Z_{eg}$. As noted earlier, the formulas above give Z_e in units of [m^3], which is equivalent to [10^{18} mm^6 m^{-3}], and that is why the equivalent reflectivity factors for each hydrometeor are additive. To express the cumulative Z_e in units normally used in radar displays (dBZ), one has to make a final conversion:

Z_e in dBZ = 10 $\log_{10}$ [Z_e in mm^6 m^{-3}]

Phase-Correcting a Predicted VIL Field Based upon Current WSR-88D Radar Observations

This procedure follows Brewster (2003). Consider two VIL fields mapped to a common grid, one derived from a model forecast, F, and the other from a nationwide mosaic of WSR-88D radar observations, O. To keep the problem manageable, consider only those N observations within a given tile (subdivision) of the entire grid. (See Fig. 2.68 in the main body of text.) Within each tile, minimize the penalty function J defined by

$$J[\delta \boldsymbol{x}] = \frac{s(|\delta \boldsymbol{x}| l^{-1})}{N} \sum_{i=1}^{N} [\bar{F}(\boldsymbol{x}_i + \delta \boldsymbol{x}) - \bar{O}(\boldsymbol{x}_i)]^2$$

$\boldsymbol{x}_i$ is the horizontal position vector to an observation within the tile. $\bar{F}$ and $\bar{O}$ are the forecasted and observed VIL fields after smoothing by a nine-point filter in two dimensions. $\delta \boldsymbol{x}$ is the horizontal shift vector, valid for the whole tile, which is to be determined by minimizing J. $|\delta \boldsymbol{x}|$ is the magnitude of the shift vector. N is the number of VIL observations in the tile. s is a distance-dependent function that imposes a penalty as the magnitude of the potential shift increases.

$$s(|\delta \boldsymbol{x}| l^{-1}) = \frac{\exp(|\delta \boldsymbol{x}| l^{-1})}{1 + |\delta \boldsymbol{x}| l^{-1}}$$

l is a length scale parameter defined by

$$l = 0.5 \sqrt{L_x^2 + L_y^2},$$

where L_x and L_y are the lengths of the sides of the tile under consideration.

Instead of minimizing the penalty function J mathematically, which could be problematic because of sharp gradients in the VIL field, many different shifts between the predicted and observed VIL fields are tried in both horizontal dimensions, first shifting three grid lengths at a time to find the coarse shift that produces a minimum J, then shifting only one grid length at a time to refine the shift vector and get closer to the true minimum.

The shift vector $\delta\boldsymbol{x}$ that minimizes the penalty function is applied to every grid point in the tile. Shift vectors are determined for all tiles contained within the VIL map. Because the tiles overlap considerably, most grid points will receive multiple shift vectors. After minimization has concluded for all tiles, the shift vectors accumulated at each grid point are averaged, which produces a smoother field of shift vectors.

B3.1.8 Logistic Regression and Its Use in the Ceiling and Visibility Forecast (CVF)

The logistic function

$$f(x;\mu,s) = \frac{1}{1+e^{-(x-\mu)/s}}$$

is a cumulative distribution function. The free parameters are the mean μ of the distribution and *s*, which is proportional to the standard deviation of the distribution. This function is derived from a probability density function (pdf) not unlike that for a Gaussian (normal) distribution. The pdf for the logistic function is symmetric about the mean, rises steeply toward the peak value, and has values that decrease toward zero in the tails. The logistic function, graphed below in Fig. B2 with $\mu = 0$ and $s = 1$, rises monotonically from near zero at large negative values to near one at large positive values of x.

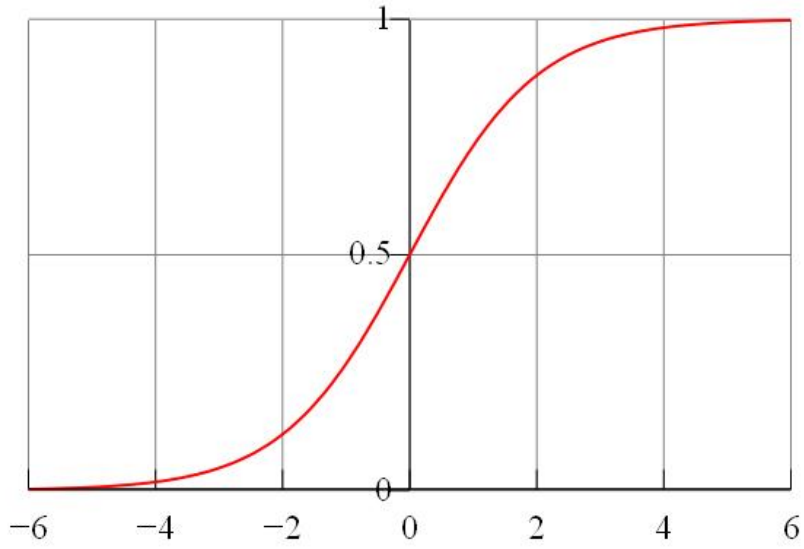

Fig. B2. A graph of the logistic function with $\mu = 0$ and $s = 1$.

The variable x usually takes the form $x = \alpha_1 x_1 + \alpha_2 x_2 + \cdots + \alpha_k x_k$, where the x_k are predictors of x and the α_k are regression coefficients. For the CVF, x is computed separately for each variable (ceiling or visibility), each category (LIFR, IFR, MVFR, VFR or others), each METAR site (over 1400 of them), each valid time, and each forecast lead time (hours 1, 2, 3, 4, and 5). The predictor set is drawn from observations at the target METAR site and its two nearest neighbors. The regression coefficients are determined from a multi-year archive of METAR data. They are updated only infrequently.

B3.1.10 GOES-R Algorithm for inferring probability of IFR conditions

The reference for this algorithm is Calvert and Pavolonis (2011). The preliminary estimate of the probability of fog/ low stratus depends upon the following metrics: at night, upon the 3.9-µm pseudo-emissivity, the radiometric surface temperature bias and the low-level profile of relative humidity. During the day, it depends upon the 3.9-µm reflectance, the spatial uniformity of the 0.65-µm reflectance measured with reference to the cloud's local radiative center, the radiometric surface temperature bias and low-level profile of relative humidity. These metrics are discussed in more detail immediately below. The probability of fog or low stratus is not computed if the cloud mask indicates clear sky (only during the day) or if the cloud top consists of ice crystals.

The 3.9 µm pseudo-emissivity is a ratio. The numerator is the observed 3.9-µm radiance,

$R_{obs}(3.9\ \mu\text{m})$. The denominator is the 3.9-µm blackbody radiance calculated from the Planck Function B using the 11-µm brightness temperature, $BT(11\ \mu\text{m})$, that is,

$$ems(3.9\ \mu\text{m}) = \frac{R_{obs}(3.9\ \mu\text{m})}{B\big(3.9\ \mu\text{m}\,, BT(11\ \mu\text{m})\big)}.$$

The Planck Function for blackbody radiation at wavelength λ and absolute temperature T is

$$B(\lambda, T) = \frac{2hc^2}{\lambda^5}\,\frac{1}{\exp\left(\frac{hc}{\lambda k_B T}\right) - 1}.$$

h is the Planck constant, 6.626 x 10^{-34} J s.
c is the speed of light in a vacuum, 2.998 x 10^{8} m s^{-1}.
k_B is the Boltzmann constant, 1.381 x 10^{-23} J K^{-1}.
When using the Planck Function, T must be expressed in degrees K, and λ in meters. To calculate $B\big(3.9\ \mu\text{m}\,, BT(11\ \mu\text{m})\big)$, use λ = 3.9 x 10^{-6} m and $T = BT(11\mu m)$ in [K] in the equation for B.

According to Calvert and Pavolonis (2011), for the detection of fog and low stratus, it is better to use the 3.9 µm pseudo-emissivity than the brightness temperature difference, $BT(3.9\ \mu\text{m}) - \text{BT}(11\ \mu m)$, because the former is less sensitive to the scene temperature.

The radiometric surface temperature bias is calculated in three steps. The first step estimates the 11-µm radiance emitted by the target, R_{tar}, which may be the land surface or a cloud top.

$$R_{tar}(11\ \mu\text{m}) = \frac{R_{obs}(11\ \mu\text{m}) - R_{atm}(11\ \mu\text{m})}{t_{atm}(11\ \mu\text{m})}$$

R_{obs} is the observed radiance; R_{atm} is the top-of-the-atmosphere, upwelling clear-sky radiance; and t_{atm} is the total gaseous atmospheric transmittance (no clouds), all at a wavelength of 11 µm. The latter two quantities are estimated using a clear-sky radiative transfer model with input from an NWP model that can provide surface temperature and vertical profiles of temperature, water vapor, and ozone (Pavolonis, 2010). Surface emissivity and viewing angle are also needed. The former is available from MODIS measurements (monthly means); the latter is known from satellite viewing geometry.

The second step is to estimate the radiative temperature of the target T_{tar} by backsolving the Planck Function, using $R_{tar}(11\ \mu\text{m})$ for $B(\lambda, T)$ and λ = 1.1 x 10^{-5} m in the following equation:

$$T_{tar}(11\ \mu\text{m}) = \frac{1}{\varepsilon_{sfc}(11\ \mu\text{m})} \frac{hc}{\lambda k_T \ln\left[1 + \frac{2hc}{\lambda^5 B(\lambda, T)}\right]}$$

The surface emissivity at 11 μm, $\varepsilon_{sfc}(11\ \mu\text{m})$, is a monthly mean, obtained from MODIS observations. The third and final step is to calculate the radiometric surface temperature bias. It is the difference between the estimated target temperature and the surface temperature obtained from a numerical weather prediction (NWP) model:

$$T_{bias} = T_{tar}(11\ \mu\text{m}) - T_{sfc}(NWP)\ .$$

If the sky is clear and the calculated atmospheric transmissivity, the upwelling radiance at the top of the atmosphere, the surface emissivity (all at 11 μm), and the model surface temperature are all accurate, the radiometric surface temperature bias should be close to zero. Because the calculations are not perfect, the bias is seldom zero, even with clear skies. Even so, when $T_{bias} \geq -10$K, the probability of fog/ low stratus is substantially greater than when T_{bias} is strongly negative (< -10K). Whenever $T_{tar} \ll T_{sfc}(NWP)$, the implied lapse rate suggests that the observed 11-μm radiance is probably coming from cloud tops well above the surface and thus considerably colder than the surface. Due to the lack of vertical motion within low stratus clouds, they generally occur in thin layers. Thus, if the cloud tops are well above the surface, the cloud bases usually are, too, thus reducing the probability of fog or low stratus.

The 3.9-μm reflectance helps to identify droplet sizes in clouds. Considering only liquid-droplet clouds, fog or low stratus clouds usually contain smaller water droplets than elevated cloud layers. Smaller water droplets have a higher reflectivity at 3.9 μm than larger droplets, and that is why the 3.9-μm reflectance helps to identify clouds with small droplet sizes.

Stratiform clouds are spatially uniform in both temperature and reflectance. The spatial uniformity of 0.65-μm (visible) reflectance can be calculated as the standard deviation of this parameter, σ_{Refl}(0.65 μm) within a 3x3 pixel array but this calculation can be troublesome near cloud edges where both cloud and non-cloud pixels make up the 3x3 pixel array. In order to avoid this issue, the spatial uniformity for each pixel is calculated with reference to the local radiative center of the cloud that contains the pixel. The 0.65-μm reflectance decreases toward cloud edges, so the local radiative center is found by following the gradient of the 0.65-μm

reflectance to a nearby local maximum in brightness. Calculating the spatial uniformity for each pixel with reference to its local radiative center ensures that pixels filled with cloud but near the cloud edge will not be included with clear pixels or pixels partially filled with cloud when computing spatial uniformity. The spatial uniformity is stored at each pixel location. When σ_{Refl}(0.65 μm) is small, the field is considered to be spatially uniform, increasing the probability that the cloud layer is stratiform.

Fog and low stratus clouds form when the air becomes saturated and water molecules migrate to condensation nuclei and form water droplets. For the GOES-R fog/low-status algorithm, relative humidity (RH) data are obtained from a NWP model such as the Global Forecast System (GFS). Fog is defined as a cloud with a base that reaches the surface. Low stratus clouds meeting IFR criteria have bases above the surface, in which case relative humidity increases upward to cloud base. For this reason the modeled surface RH and RH profiles are used to find the highest RH within 1000 ft of the surface. The maximum RH and its altitude are stored to help identify areas that meet IFR criteria.

The preliminary determination of fog/low-stratus probability depends upon an extensive statistical database, including the five parameters just discussed, and surface observations of ceiling. The fraction of events for which specific combinations of parameter values result in IFR conditions (as determined from surface observations) is the basis for estimating probabilities.

Nighttime probability of IFR conditions. The parameters used are $ems(3.9\ \mu\mathrm{m})$, T_{bias} , and RH. $ems(3.9\ \mu\mathrm{m})$ is used only at night because, during the day, BT(3.9 μm) is contaminated by reflected solar radiation. Calvert and Pavolonis (2010) generated look-up tables (LUT) that indicate the probability of IFR conditions, given values of the first two parameters. Figure B3 illustrates the LUT graphically. The vertical axis is $ems(3.9\ \mu\mathrm{m})$, and the horizontal axis is T_{bias} . In the LUT, the probability of IFR conditions rises with decreasing $ems(3.9\mu m)$ particularly as T_{bias} rises from -10K toward zero.

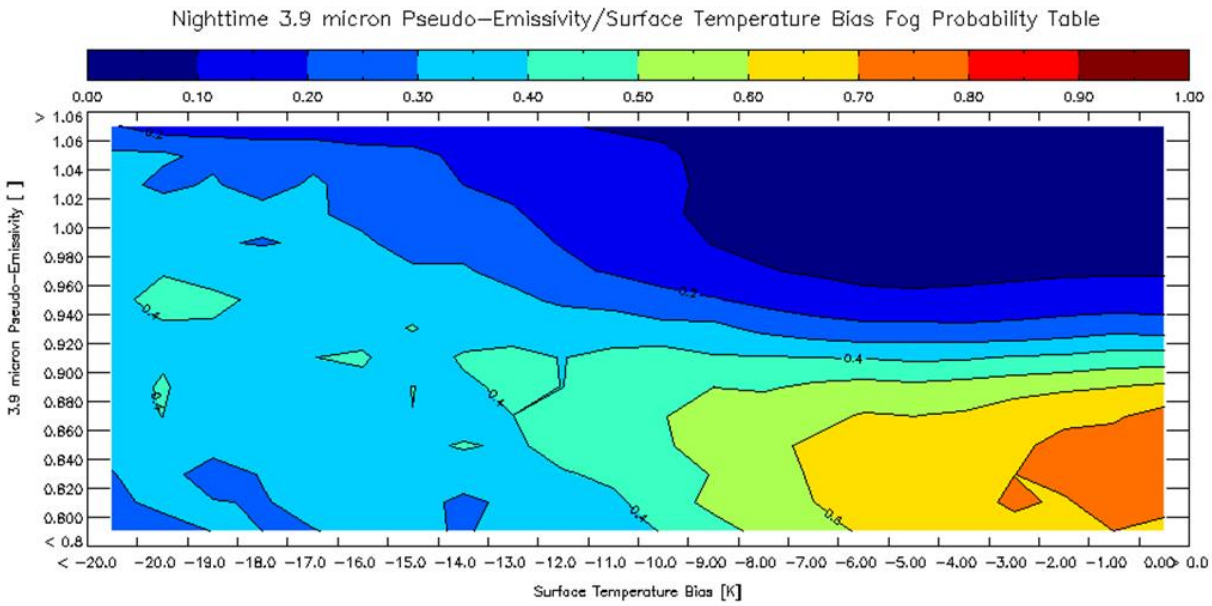

Fig. B3. Nighttime look-up table in graphical form, giving the probability of IFR conditions as a function of the 3.9-µm pseudo-emissivity (vertical axis) and the radiometric surface temperature bias (horizontal axis). The color bar at top gives the probabilities. From Calvert and Pavolonis (2011).

A second LUT is based on the graph in Fig. B4, a plot of relative humidity near the surface against the probability of fog. The figure shows two curves, one for night (black) and one for daytime (red).

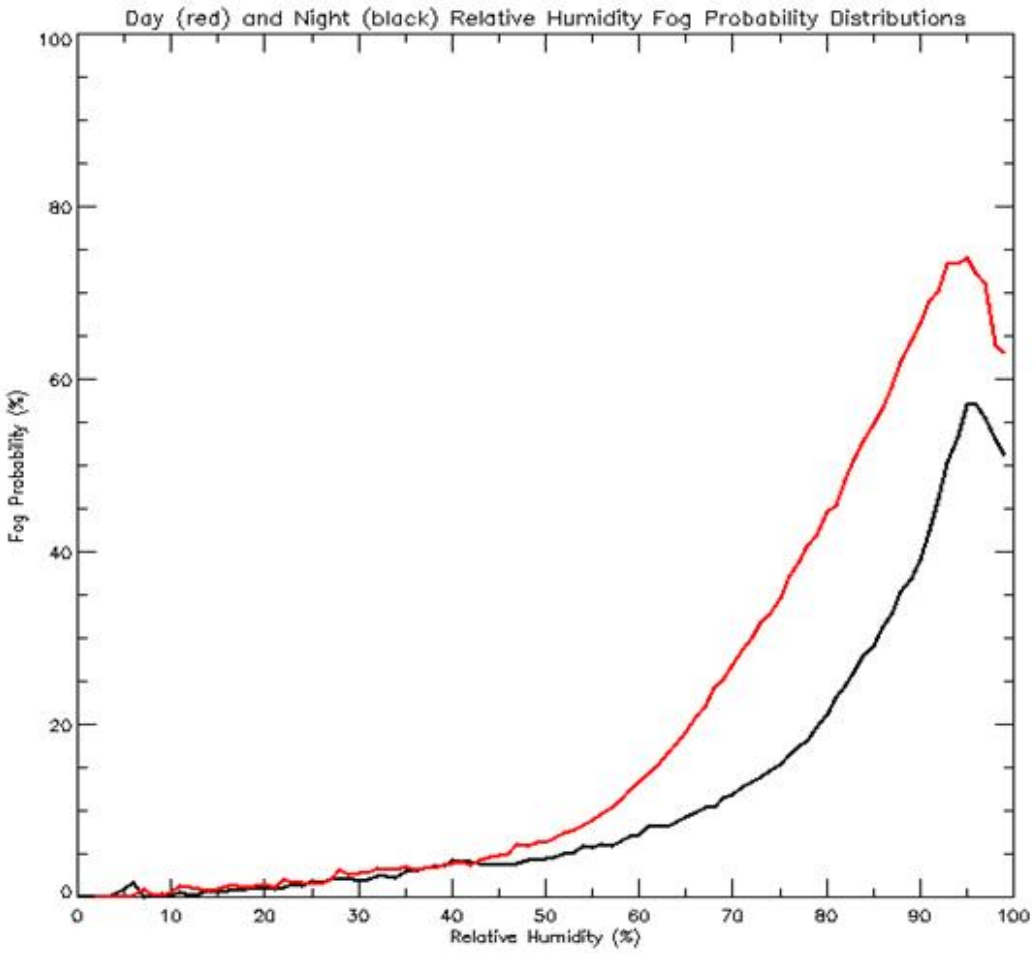

Fig. B4 Graph showing the probability of fog, given values of relative humidity close to the ground as obtained from a numerical prediction model. The red curve is for daytime; the black curve is for night. From Calvert and Pavolonis (2011).

Daytime probability of IFR conditions. The parameters used are the 3.9-µm reflectance, σ_{Refl}(0.65 µm), T_{bias} , and RH. Figure B5 shows a sample daytime LUT in color-coded form. For daytime IFR conditions, there are separate LUTs for each of four bins of 3.9-µm reflectance: (0.0 to 0.10, 0.10 to 0.15, 0.15 to 0.20, and greater than 0.20). The figure shows the LUT for one of these bins: 0.15 to 0.20. The probabilities are highest for smooth fields of 0.65-µm reflectance. The probability of IFR conditions depends rather strongly on spatial uniformity of the cloud top as seen at 0.65 µm but only weakly on the surface temperature bias. The look-up table for daytime RH is part of the previous figure.

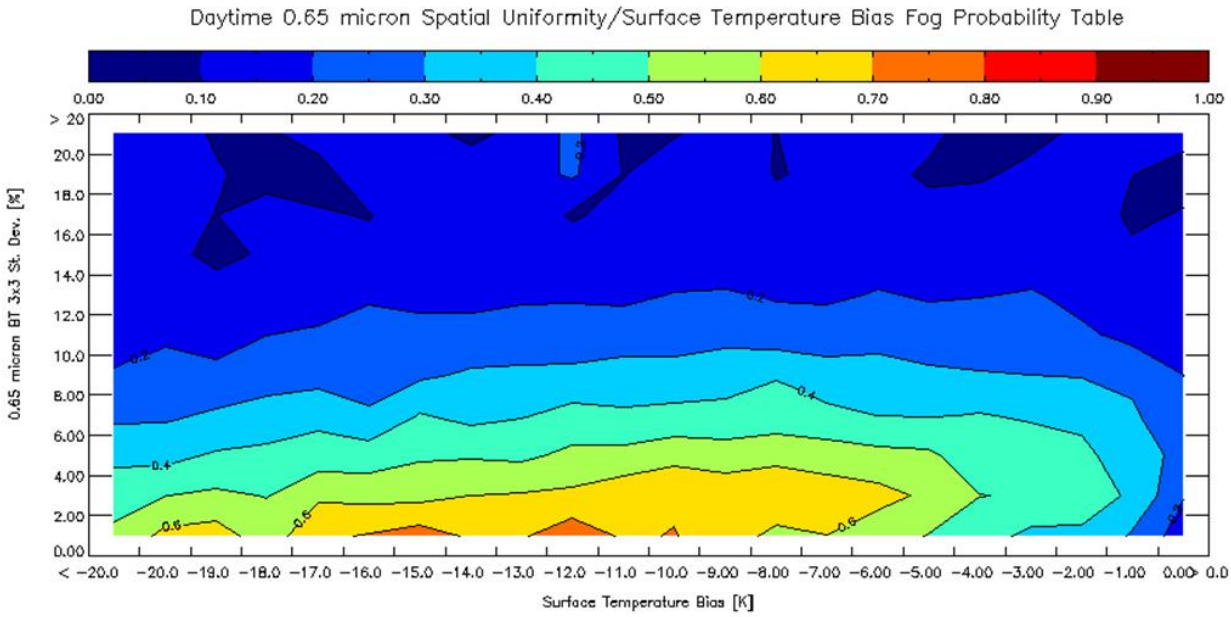

Fig. B5. Daytime look-up table in graphical form, giving the probability of IFR conditions as a function of the standard deviation of 0.65-µm reflectance (a measure of spatial uniformity of the cloud top – vertical axis) and the radiometric surface temperature bias (horizontal axis). This LUT is for a range of 3.9-µm reflectance between 0.15 and 0.20. From Calvert and Pavolonis (2011).

The lookup tables just presented illustrate how various parameters, taken one or two at a time, influence the likelihood of IFR conditions. The overall probabilities, based upon all measured parameters, are computed using a Bayesian probability model as summarized in Calvert and Pavolonis (2011). This model returns a conditional probability that an event will occur, given the set of satellite measurements discussed above.

Once probabilities of fog/low stratus have been assigned at pixels with liquid-water cloud, further screening occurs. During the daytime, pixels where the daytime probability is less than 40% are rejected. At night, the threshold is 26%. Other tests are performed to increase the probability of detection of fog/low clouds and decrease the false alarm rate. These tests involve 0.65-µm and 3.9-µm reflectances, T_{bias} , BT(11 µm); and various ratios, smoothness tests, and cumulative distribution functions involving these parameters, too involved to describe here.

References

Bader, J.J., G.S. Forbes, J.R. Grant, R.B.E. Lilley, and A.J. Waters, 1995: *Images in Weather Forecasting. A Practical Guide for Interpreting Satellite and Radar Imagery*. Cambridge University Press, 499 pp.

Barnes, S., 1964: A technique for maximizing details in numerical map analysis. *Journal of Applied Meteorology*, **3**, 395-409.

Bedka, K., J. Brunner, R. Dworak, W. Feltz, J. Otkin, and T. Greenwald, 2010: Objective satellite-based detection of overshooting tops using infrared window channel brightness temperature gradients. *Journal of Applied Meteorology and Climatology*, **49**, 181-202.

Bedka, K.M., and J.R. Mecikalski, 2005: Application of satellite-derived atmospheric motion vectors for estimating mesoscale flows. *Journal of Applied Meteorology*, **44**, 1761-1772.

Benjamin, S.G., B.D. Jamison, W.R. Moninger, S.R. Sahm, B.E. Schwartz, and T.W. Schlatter, 2010: Relative short-range forecast impact from aircraft, profiler, radiosonde, VAD, GPS-PW, METAR, and mesonet observations via the RUC hourly assimilation cycle. *Monthly Weather Review*, **138**, 1319-1343.

Benjamin S., M. Hu, S. Weygandt, and D. Devenyi, 2009: Integrated assimilation of radar, satellite, and METAR cloud data for initial hydrometeor/divergence fields to improve hourly updated short-range forecasts from the RUC, Rapid Refresh, and HRRR. *WMO Symposium on Nowcasting*, 31 August – 4 September, Whistler, B.C., Canada, Paper 3.1, Powerpoint presentation.

Benjamin, S.G., D. Dévényi, S.S. Weygandt, K.J. Brundage, J.M. Brown, G.A. Grell, D. Kim, B.E. Schwartz, T.G. Smirnova, T.L. Smith, and G.F. Manikin, 2004a: An hourly assimilation–forecast cycle: The RUC. *Monthly Weather Review*, **132**, 495-518.

Benjamin, S.G., G.A. Grell, J.M. Brown, T.G. Smirnova, and R.Bleck, 2004b: Mesoscale weather prediction with the RUC hybrid isentropic–terrain-following coordinate model. *Monthly Weather Review*, **132**, 473-494.

Berendes, T.A., J.R. Mecikalski, W.M. MacKenzie, Jr., K.M. Bedka, and U.S. Nair, 2008: Convective cloud identification and classification in daytime satellite imagery using Standard Deviation Limited Adaptive Clustering. *Journal of Geophysical Research, Atmospheres*, **113**, D20207, doi:10.1029/2008/JD010287, 25 October, 9 pp.

Bernstein, B.C., F. McDonough, M.K. Politovich, B.G. Brown, T P. Ratvasky, D.R. Miller, C.A. Wolff, and G. Cunning, 2005: Current Icing Potential: Algorithm description and comparison with aircraft observations. *Journal of Applied Meteorology*, **44**, 969-986.

Bolton, D., 1980: The computation of equivalent potential temperature. *Monthly Weather Review*, **108**, 1046-1053.

Brewster, K.A., 2003: Phase-correcting data assimilation and application to storm-scale numerical weather prediction. Part I: Method description and simulation testing. *Monthly Weather Review*, **131**, 480-492.

Brown, R., 1973: New indices to locate clear-air turbulence. *Meteorological Magazine*, **102**, 347-361.

Burk, S.D., and W.T. Thompson, 1989: A vertically nested regional numerical prediction model with second-order closure physics. *Monthly Weather Review*, **117**, 2305-2324.

Calvert, C., and M. Pavolonis, 2011: GOES-R Advanced Baseline Imager (ABI) algorithm theoretical basis document for low cloud and fog, NOAA NESDIS Center for Satellite Applications and Research (Version 2.0, 01 July), 72 pp., available at http://cimss.ssec.wisc.edu/~mpav/goesr/GOESR_ABI_ATBD_Aviation_Fog_v2.0_01July2011_MS2004.doc.

Chang, F.-L., P. Minnis, B. Lin, M. Khaiyer, R. Palikonda, and D. Spangenberg, 2010a: A modified method for inferring cloud top height using GOES-12 imager 10.7- and 13.3-μm data. *Journal of Geophysical Research*, **115**, D06208, doi:10.1029/2009JD012304.

Chang, F.-L., P. Minnis, S. Sun-Mack, L. Nyugen, and Yan Chen, 2010b: On the satellite determination of multi-layered, multi-phase cloud properties. Proceedings, AMS 13th Conference on Atmospheric Radiation and Cloud Physics, Portland, OR, 27 June – 2 July, Paper JP1.10.

Charba, J.P., and F.G. Samplatsky, 2009: Operational 2-h thunderstorm guidance forecasts to 24 hours on a 20-km grid. 23rd Conference on Weather Analysis and Forecasting/ 19th Conference on Numerical Weather Prediction, 1-5 June 2007, Omaha, Nebraska, American Meteorological Society, 8 pp.

Chornoboy, E.S., A.M. Matlin, and J.P. Morgan, 1994: Automated storm tracking for terminal air traffic control. *The Lincoln Laboratory Journal*, **7** (2), 427-447.

Clarisse, L., P.F.Coheur, A.J. Prata, D. Hurtmans, A. Razavi, T. Phulpin, J. Hadji-Lazaro, and C. Clerbaux, 2008: Tracking and quantifying volcanic SO_2 with IASI, the September 2007 eruption at Jebel at Tair. *Atmospheric Chemistry and Physics*, **8**, 7723-7734.

Cornman, L.B., G. Meymaris, and M. Limber, 2004: An update on the FAA Aviation Weather Research Program's in situ turbulence measurement and report system. Preprints. 11th Conference on Aviation, Range, and Aerospace Meteorology, Hyannis, MA, American Meteorological Society, CD-ROM, P4.3, 14 pp.

Davison Reynolds, H.J., R.A. DeLaura, and M. Robinson, 2010 : Field & (Data) Stream: A method for functional evolution of the Air Traffic Management Route Availability Planning Tool (RAPT). *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, **54**, No. 1 (September), 104-108. doi:10.1177/154193121005400123.

DeLaura, R, B. Crowe, R. Ferris, J.F. Love, and W.N. Chan, 2009: Comparing convective weather avoidance models and aircraft-based data. Aviation, Range and Aerospace Meteorology, Special Symposium on Weather – Air Traffic Management Integration, 11-15 January, Phoenix, Arizona, 89th Annual Meeting, American Meteorological Society, 17 pp.

DeLaura, R., M. Robinson, M. Pawlak, and J. Evans. 2008: Modeling convective weather avoidance in en route airspace. 13th Conference on Aviation, Range and Aerospace Meteorology, 21-24 January, New Orleans, Louisiana, 88th Annual Meeting, American Meteorological Society, 12 pp.

DeLaura, R., and J. Evans, 2006: An exploratory study of modeling en route pilot convective storm flight deviation behavior. 12th Conference on Aviation, Range and Aerospace Meteorology, 30 January – 2 February, Atlanta, Georgia, 86th Annual Meeting, American Meteorological Society, 15 pp.

Dixon, M., and G. Wiener, 1993: TITAN: Thunderstorm identification, tracking, analysis, and nowcasting—A radar-based methodology. *Journal of Atmospheric and Oceanic Technology*, **10**, 785–797.

Dong, X., P. Minnis, G.G. Mace, W.L. Smith, Jr., M. Poellot, R. T. Marchand, and A.D. Rapp, 2002: Comparison of stratus cloud properties deduced from surface, GOES, and aircraft data during the March 2000 ARM Cloud IOP. *Journal of the Atmospheric Sciences*, **59**, 3265-3284.

Doviak, R.J., and D.S. Zrnić, 1984: *Doppler Radar and Weather Observations*. Academic Press, Inc., New York, NY, 458 pp.

Draxler, R.R., and G.D. Hess, 1998a: An overview of the HYSPLIT_4 modelling system for trajectories, dispersion, and deposition. *Australian Meteorological Magazine*, **47**, 295-308.

Draxler, R.R., and G.D. Hess, 1998b: Description of the HYSPLIT_4 modeling system. NOAA Technical Memorandum ERL ARL-224 (August), Air Resources Laboratory, 24 pp.

Draxler, R.R., and A.D. Taylor, 1982: Horizontal dispersion parameters for long-range transport modeling. *Journal of Applied Meteorology*, **21**, 367-372.

Dudhia, J., 1989: Numerical study of convection observed during the winter monsoon experiment using a mesoscale two-dimensional model. Journal of the Atmospheric Sciences, 46, 3077–3107. (See Appendix C for description of short-wave radiation scheme.)

Dupree, W.J., M.M. Wolfson, R.J. Johnson, Jr., R.A. Boldi, E.B. Mann, K.T. Calden, C.A. Wilson, P.E. Bieringer, B.D. Martin, and H. Iskenderian, 2005: FAA tactical weather forecasting in the United States national Airspace. World Weather Research Program Symposium on Nowcasting and Very Short Term Forecasts, 5-9 September 2005, Toulouse, France, 29 pp.

Dupree, W.J., R.J. Johnson, M.M. Wolfson, K.E. Theriault, B.E. Forman, R.A. Boldi, and C.A. Wilson, 2002: Forecasting convective weather using multiscale detectors and weather classification – Enhancements to the MIT Lincoln Laboratory Terminal Convective Weather Forecast. *10th Conference on Aviation, Range, and Aerospace Meteorology*, Portland, Oregon, American Meteorological Society, 132-145.

Eckermann, S., 2009: Hybrid $\sigma - p$ coordinate choices for a global model. *Monthly Weather Review*, **137**, 224-245.

Evans, J.E., and E.R. Ducot, 2006: Corridor Integrated Weather System. *Lincoln Laboratory Journal*, **16** (1), 59-80.

Evans, J.E., and E.R. Ducot, 1994: The Integrated Terminal Weather System (ITWS). *The Lincoln Laboratory Journal*. **7**, 449-474.

Fowler, T.L., M. Chapman, A. Holmes, B.G. Brown, J.L. Mahoney, J.T. Braid, and P. Boylan, 2006: Quality assessment report. Current Icing Potential (CIP): Probability, severity, and SLD. Internal report issued 31 May 2006 jointly by NCAR and NOAA's Earth System Research Laboratory, Boulder, Colorado, 47 pp.

Frehlich, R., R. Sharman, F. Vandenberghe, W. Yu, Y. Liu, and J. Knievel, 2010: Estimates of C_n^2 from numerical weather prediction model output and comparision with thermosonde data. *Journal of Applied Meteorology and Climatology*, **49**, 1742-1755.

Frehlich, R., and R. Sharman, 2004: Estimates of turbulence from numerical weather prediction model output with applications to turbulence diagnosis and data assimilation. *Monthly Weather Review*, **132,** 2308-2324.

Ghirardelli, J.E., and B. Glahn, 2011: Gridded Localized Aviation MOS Program (LAMP) guidance for aviation forecasting. *Preprints, 15th Conference on Aviation, Range, and Aerospace Meteorology*, 1-4 August, Los Angeles, California, American Meteorological Society, Paper 4.4, 22 pp.

Ghirardelli, J.E., and B. Glahn, 2010: The Meteorological Development Laboratory's aviation weather prediction system. *Weather and Forecasting*, **25**, 1027-1051.

Glahn, H.R., and D.A. Lowry, 1972: The use of model output statistics (MOS) in objective weather forecasting. *Journal of Applied Meteorology*, **11**, 1203-1211.

Glahn, B., and J.-S. Im, 2011: Algorithms for effective objective analysis of surface weather variables. *Preprints, 24th Conference on Weather Analysis and Forecasting*, 24-27 January, Seattle, Washington, American Meteorological Society, Paper J19.4, 11 pp.

Glickman, T.S., 2000 (Ed.): *Glossary of Meteorology*. American Meteorological Society, Boston, MA, 855 pp.

Grell, G.A. and D. Devenyi, 2002: A generalized approach to parameterizing convection combining ensemble and data assimilation techniques, *Geophysical Research Letters*, **29**, No. 14, doi:10.1029/2002GL015311.

Guffanti, M., T.J. Casadevall, and K. Budding, 2010: Encounters of aircraft with volcanic ash clouds: A compilation of known incidents, 1953-2009. U.S. Geological Survey Data Series 545, Version 1.0, 12 pp. + 4 appendices including the compilation database. Available at http://pubs.usgs.gov/ds/545.

Herzegh, P. and J. Cowie, 2011: CVA – The NCV CONUS Analysis Product. CVA technical description. Informal notes, 24 May 2011, National Center for Atmospheric Research, 9 pp.

Heitkemper, L., R.F. Price, and D.B. Johnson, 2008: Lightning-warning systems for use by airports. ACRP Report 8, Airport Cooperative Research Program, Transportation Research Board, National Academies, Washington, D.C., 24 pp.

Juang, H.-M. H., S.-Y. Hong and M. Kanamitsu, 1997: The NCEP regional spectral model: an update. *Bulletin of the American Meteorological Society,* **78**, 2125-2143.

Iskenderian, H., C.F. Ivaldi, M. Wolfson, J. Mecikalski, W. MacKenzie, Jr., W. Feltz, J. Sieglaff, R. Dworak, and K. Bedka, 2010: Satellite Data Applications for Nowcasting of Convective Initiation. 14th Conference on Aviation, Range, and Aerospace Meteorology, 18-21 January, Atlanta, Georgia, American Meteorological Society, Paper 5.2, 11 pp. (This paper did not get into the conference proceedings but is available from the first author: haig@ll.mit.edu.)

Janjić, Z., and R. Gall, 2012: Scientific documentation of the NCEP Nonhydrostatic Multiscale Model on the B grid (NMMB). Part 1 Dynamics. In review for publication as an NCAR Technical Note, possibly also as an NCEP/EMC Technical Note. To be available at the Developmental Testbed Center at http://www.dtcenter.org/ by spring 2012.

Janjić, Z.I., 2001: Nonsingular implementation of the Mellor-Yamada level 2.5 scheme in the NCEP meso model. Office Note 437, National Centers for Environmental Prediction, 61 pp.

Janjić, Z.I., 1994: The step-mountain Eta coordinate model: Further developments of the convection, viscous sublayer, and turbulence closure schemes. *Monthly Weather Review*, **122**, 927-945.

Jedlovec, G.J., and K. Laws, 2003: GOES cloud detection at the Global Hydrology and Climate Center. *12th Conference on Satellite Meteorology and Oceanography*, 10-13 February, Long Beach, California, American Meteorological Society, Paper P1.21, 8 pp.

Keenan, T., P. Joe, J. Wilson, C. Collier, B. Golding, D. Burgess, P. May, C. Pierce, J. Bally, A. Crook, A. Seed, D. Sills, L. Berry, R. Potts, I. Bell, N. Fox, E. Ebert, M. Eilts, K. O'Loughlin, R. Webb, R. Carbone, K. Browning, R. Roberts, And C. Mueller, 2003: The Sydney 2000 World Weather Research Programme forecast demonstration project: Overview and current status. *Bulletin of the American Meteorological Society,* **84**, 1041-1054.

Keohan, C., 2007: Ground-based wind shear detection systems have become vital to safe operations. *ICAO Journal* (the magazine of the International Civil Aviation Organization), **62**, 16-19, 33-34.

Knox, J.A., D. W. McCann, and P.D. Williams, 2008: Application of the Lighthill-Ford theory of spontaneous imbalance to clear-air turbulence forecasting. *Journal of the Atmospheric Sciences*, **65**, 3292-3304.

Kulesa, G., 2002: Weather and aviation: How does weather affect the safety and operations of airports and aviation, and how does FAA work to manage weather-related effects? Discussion paper presented at *The Potential Impacts of Climate Change on Transportation*, a Federal Research Partnership Workshop, 1-2 October, the Brookings Institution, Washington, D.C., hosted by the U.S. Department of Transportation Center for Climate Change and Environmental Forecasting, 10 pp.

Lacis, A.A. and J.E. Hansen, 1974: A parameterization for the absorption of solar radiation in the earth's atmosphere. *Journal of the Atmospheric Sciences*, **31**, 118-133.

Lack, S.A., G.J. Layne, S. Madine, and J.L. Mahoney, 2011: Quality Assessment of CoSPA. Unpublished report by the Quality Assessment Product Development Team, NOAA/ESRL/GSD, 31 March, Boulder, Colorado, 34 pp.

Lhermitte, R.M., and D. Atlas, 1961: Precipitation motion by pulse Doppler radar. Proceedings of the Ninth Weather Radar Conference, Kansas City, Missouri, American Meteorological Society, 218-223.

Lindborg, E., 1999: Can the atmospheric kinetic energy spectrum be explained by two-dimensional turbulence? *Journal of Fluid Mechanics*, **388,** 259-288.

Loughe, A.F., B.P. Pettegrew, J.K. Henderson, J.E. Hart, S. Madine, and J.L. Mahoney, 2010: Quality assessment report: National Ceiling and Visibility Analysis Product. Prepared by the Quality Assessment Product Development Team as part of the FAA's Aviation Weather Research Program, NOAA/Earth Systems Research Laboratory/Global Systems Division, 30 November, 38 pp. + 3 appendices.

Madine, S., S.A Lack, S.A. Early, M. Chapman, J.K. Henderson, J.E. Hart, and J.L. Mahoney, 2008: Quality Assessment Report: Forecast Icing Product (FIP). Internal report, Global Systems Division, NOAA Earth System Research Laboratory, Boulder, Colorado, 40 pp.

McCann, D. W., 2001: Gravity waves, unbalanced flow, and aircraft clear air turbulence. *National Weather Digest*, **25**, 3-14.

McDonough, F., B.C. Bernstein, M.K. Politovich, and C.A. Wolff, 2004: The Forecast Icing Potential (FIP) algorithm. Paper 3.2, *20th International Conference on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology*, Seattle, WA, 11-15 January 2004, American Meteorological Society, 6 pp.

McLaughlin, D. and 28 coauthors, 2009: Short-wavelength technology and the potential for distributed networks of small radar systems. *Bulletin of the American Meteorological Society*, **90**, 1797-1817.

McNeill, D., and P. Freiberger, 1993: *Fuzzy Logic: The Revolutionary Computer Technology That is Changing our World*. Simon and Schuster, 319 pp.

Mecikalski, J.R., K.M. Bedka, S.J. Paech, and L.A. Litten, 2008: A statistical evaluation of GOES cloud-top properties for nowcasting convective initiation. *Monthly Weather Review*, **136**, 4899-4914.

Mecikalski, J.R., and K.M. Bedka, 2006: Forecasting convective initiation by monitoring the evolution of moving cumulus in daytime GOES imagery. *Monthly Weather Review*, **134**, 49-78.

Minnis, P., D.P. Kratz, J.A Coakley, Jr., M.D. King, D. Garber, P. Heck, S. Mayor, D.F. Young, and R. Arduini, 1995: Cloud Optical Property Retrieval (Subsystem 4.3) in *Clouds and the Earth's*

Radiant Energy System (CERES), Algorithm Theoretical Basis Document, Volume III: Cloud Analysis and Radiance Inversions (Subsystem 4). NASA RP 1376, Vol. 3, 135-176.

Minnis, P., L. Nguyen, W.L. Smith, Jr., M.M. Khaiyer, R. Palikonda, D.A. Spangenberg. D.R. Doelling, D. PHan, G.D. Nowicki, P.W. Heck, and C. Wolff, 2004: Real-time cloud, radiation, and aircraft icing parameters from GOES over the USA. *Proceedings, 13th Conference on Satellite Oceanography and Meteorology*, Norfolk, VA, 20-24 September 2004, CD-ROM, Paper 7.1, American Meteorological Society, 6 pp.

Minnis, P., Q.Z. Trepte, S. Sun-Mack, Y. Chen, D.R. Doelling, D.F. Young, D.A. Spangenberg, W.F. Miller, B.A. Wielicki, R.R. Brown, S.C. Gibson, and E.B. Geier, 2008: Cloud detection in nonpolar regions for CERES using TRMM VIRS and Terra and Aqua MODIS data. *IEEE Transactions on Geoscience and Remote Sensing*, **46***,* No. 11 (November), 3857-3884.

Minnis, P. W. L. Smith, Jr., D. A. Spangenberg, C. Fleeger, F.-L. Chang, R. Palikonda, S. Bedka, G. Hong, P. W. Heck, and T. Hauf, 2010: Nowcasting aircraft icing conditions at night and in the presence of multilayered clouds using meteorological satellite data. 2010 EUMETSAT Satellite Conference, Cordoba, Spain, September 20-24.

Minnis, P., S. Sun-Mack, D.F. Young., P.W. Heck, D.P. Garber, Y. Chen, D.A. Spangenberg, R.F. Arduini, Q.Z. Trepte, W.L. Smith, Jr., J.K. Ayers, S.C. Gibson, W.F. Miller, G. Hong, V. Chakrapani, Y. Takano, K.-N. Liou, Y. Xie, and P. Yang, 2011a: CERES Edition-2 cloud property retrievals using TRMM VIIRS and Terra and Aqua MODIS Data. Part I: Algorithms. *IEEE Transactions on Geoscience and Remote Sensing*, **49**, in press, 27 pp.

Minnis, P., S. Sun-Mack, Y. Chen, M.M Khaiyer, Y. Yi, J.K. Ayers, R.R. Brown, X. Dong, S.C. Gibson, P.W. Heck, B. Lin, M.L. Nordeen, L. Nguyen, R. Palikonda, W.L. Smith, Jr., D.A. Spangenberg, Q.Z. Trepte, and B. Xi, 2011b: CERES Edition-2 cloud property retrievals using TRMM VIRS and Terra and Aqua MODIS data, Part II: Examples of average results and comparisons with other data. *IEEE Transactions on Geoscience and Remote Sensing,* **49**, in press, 30 pp.

Mitchell, J.S.B., V. Polishchuk, and J. Krozel, 2006: Airspace throughput analysis considering stochastic weather. AIAA Guidance, Navigation, and Control Conference, Keystone, Colorado, August, 19 pp.

Mlawer, E.J., S.J. Taubman, P.D. Brown, M.J. Iacono and S.A. Clough, 1997: RRTM, a validated correlated-k model for the longwave. *Journal of Geophysical Research*, **102**, 16,663-16,682.

Mueller, C., T. Saxen, R. Roberts, J. Wilson, T. Betancourt, S. Dettling, N. Oien, and J. Yee, 2003: NCAR Auto-Nowcast System. *Weather and Forecasting*, **18**, 545-561.

Nelson, E, R. Roberts, T. Saxen, H. Cai, J. Pinto, and M. Steiner, 2008: Evaluating the performance of different synoptic regime settings on the NCAR Auto-Nowcaster. 13th Conference on Aviation, Range and Aerospace Meteorology, 21-24 January, New Orleans, Louisiana, American Meteorological Society, Paper 3.2, 6 pp.

Pan, H-L. and L. Mahrt, 1987: Interaction between soil hydrology and boundary layer developments. *Boundary-Layer Meteorology*, **38**, 185-202.

Pavolonis, M.J., 2010: Advances in extracting cloud composition information from spaceborne infrared radiances: A robust alternative to brightness temperatures. Part I: Theory. *Journal of Applied Meteorology and Climatology*, **49**, 1992-2012.

Pavolonis, M., and J. Sieglaff, 2010: GOES-R Advanced Baseline Imager (ABI) algorithm theoretical basis. Document for volcanic ash (detection and height). Version 2.0 (15 September), NOAA NESDIS Center for Satellite Applications and Research, 84 pp. Available at http://goesrtest2.woc.noaa.gov/products/ATBDs/baseline/Aviation_VolAsh_v2.0.pdf.

Phillips, C. , J. Pinto, M. Steiner, and D. Albo, 2010: Diagnosis of phase errors in high-resolution NWP model forecasts of precipitation and application to improved aviation weather forecasts. Paper 3.4, *14th Conference on Aviation, Range, and Aerospace Meteorology*, 18-21 June, Atlanta, Georgia, American Meteorological Society, 5 pp.

Pinto, J., W. Dupree, S. Weygandt, M. Wolfson, S. Benjamin, and M. Steiner, 2010: Advances in the Consolidated Storm Prediction for Aviation. Paper J11.2, *14th Conference on Aviation, Range, and Aerospace Meteorology*, Atlanta, Georgia, 18-21 January, American Meteorological Society, 10 pp.

Politovich, M.K., C. Wolff, B. Bernstein, and F. McDonough, 2006: CIP severity scientific and technical document. Unpublished document prepared for the Aviation Weather Technology Transfer Technical Review Panel, 1 June 2006, 43 pp.

Ramsay, A.C. 2002: Freezing drizzle (FZDZ) identification from the Automated Surface Observing System (ASOS): Status of the ASOS multi-sensor FZDZ algorithm. *6th Conference on*

Integrated Observing Systems, 13-17 January, Orlando, Florida, American Meteorological Society, 241-247.

Ramsay, A.C., 2000: Surface ice accretion rates from the Automated Surface Observing System (ASOS): An issue for deicing holdover times. *Ninth Conference on Aviation, Range, and Aerospace Meteorology*, 11-15 September, Orlando, Florida, American Meteorological Society, 312-316.

Rasmussen, R.M., J. Hallett, R. Purcell, Landolt, S.D., and J,. Cole, 2011: The Hotplate Precipitation Gauge. *Journal of Atmospheric and Oceanic Technology*, **28**, 148-164.

Rasmussen, R., C. Wade, F. Hage, S. Landolt, M. Tryhane, J. Cole, A. Ramsay, D. Fleming, R.K. Moore, A. Davis, B. Reis, T. Lisi, M. Kjolleberg, and K.M. Rosenlund, 2006: New ground deicing hazard associated with freezing drizzle ingestion by jet engines: *Journal of Aircraft*, **43**, 1448-1457.

Rasmussen, R., M. Dixon, S. Vasiloff, F. Hage, S. Knight, J. Vivekanandan, and M. Xu, 2003: Snow nowcasting using a real-time correlation of reflectivity with snow gauge accumulation. *Journal of Applied Meteorology*, **42**, 20-36.

Rasmussen, R., J. Cole, R.K. Moore, and M. Kuperman, 2000: Common snowfall conditions associated with aircraft takeoff accidents. *Journal of Aircraft*, **37**, 110-116.

Rasmussen, R.M., J. Vivekanandan, F. Cole, B Myers, and C. Masters, 1999: The estimation of snowfall rate using visibility. *Journal of Applied Meteorology*, **38**, 1542-1563.

Reap, R.M., 1996: Probability forecasts of clear-air turbulence for the contiguous U.S. National Weather Service Office of Meteorology, Technical Procedures Bulletin No. 430, 15 pp.

Roberts, R.D., A.R.S. Anderson, E. Nelson, B. Brown, M. Pocernich, T. Saxen, and J.W. Wilson, 2011: Operational convective storm initiation nowcasting: Forecaster involvement and evaluation. *Weather and Forecasting*, **26**, submitted.

Robinson, M., R. DeLaura, N. and Underhill, 2009: The Route Availability Planning Tool (RAPT): Evaluation of departure management decision support in New Yorik during the 2009 convective weather Season. 8th USA/Europe Air Traffic Management Research and Development Seminar (ATM2009), Napa, California, 29 June – 02 July, 10 pp. Paper available at

http://www.ll.mit.edu/mission/aviation/publications/publication-files/ms-papers/Robinson_2009_ATM_WW-16318.pdf

Roebber, P.J., 2009: Visualizing multiple measures of forecast quality. *Weather and Forecasting*, **24**, 601-608.

Rogers, E., T. Black, B. Ferrier, Y. Lin, D. Parrish, and G. DiMego, 2001: Changes to the NCEP Meso Eta Analysis and Forecast System: Increase in resolution, new cloud microphysics, modified precipitation assimilation, modified 3DVAR analysis. Eta-12 Technical Procedures Bulletin, NCEP/EMC. (Ferrier microphysics described on pp. 2-7).

Rotunno, R., J.B. Klemp, and M.L. Weisman, 1988: A theory for strong, long-lived squall lines. *Journal of the Atmospheric Sciences*, **45**, 463-485.

Rudack, D.E., 2009: Aviation weather observations vs. LAMP forecasts with the onset of precipitation. *Bulletin of the American Meteorological Society*, **90**, 929-938.

Rudack, D.E., and J.E. Ghirardelli, 2010: A comparative verification of Localized Aviation Model Output Statistics Program (LAMP) and numerical weather prediction (NWP) model forecasts of ceiling height and visibility. *Weather and Forecasting*, **25**, 1161-1178.

Ruzanski, E., and V. Chandrasekar, 2011: Improved liquid water equivalent nowcasting using the Weather Support to Deicing Decision Making system. *Journal of Atmospheric and Oceanic Technology*, **50**. In press. May be published in 2012, Vol. 51.

Schmit, T.J., M.M. Gunshor, W.P. Menzel, J.J. Gurka, J. Li, and S. Bachmeier, 2005: Introducing the Next-Generation Advanced Baseline Imager on GOES-R. *Bulletin of the American Meteorological Society*, **86**, 1079-1096.

Schwarzkopf, M. D., and S. B. Fels (1991): The Simplified Exchange Method revisited: An accurate, rapid method for computation of infrared cooling rates and fluxes. Journal of Geophysical Research, 96 (D5), 9075–9096, doi:10.1029/89JD01598.

Scollo, S., A. Folch, M. Coltelli, and V. J. Realmuto 2010: Three-dimensional volcanic aerosol dispersal: A comparison between Multiangle Imaging Spectroradiometer (MISR) data and numerical simulations, *Journal of Geophysical Research, Atmospheres*, **115**, D24210, doi:10.1029/2009JD013162, 14 pp.

Searcy, C., K. Dean, and W. Stringer, 1998: PUFF: A high-resolution volcanic ash tracking model. *Journal of Volcanology and Geothermal Research*, **80**, 1-16.

Sharman, R., C. Tebaldi, G. Wiener, and J. Wolff, 2006: An integrated approach to mid- and upper-level turbulence forecasting. *Weather and Forecasting*, **21,** 268-287.

Sieglaff, J.M. L.M. Cronce, W.F. Feltz, K.M. Bedka, M.J. Pavolonis, and A.K. Heidinger, 2011: Nowcasting convective storm initiation using satellite-based, box-averaged cloud-top cooling and cloud-type trends. *Journal of Applied Meteorology and Climatology*, **50**, 110-126.

Smirnova, T.G., J.M. Brown, S.G. Benjamin, and D. Kim, 2000: Parameterization of cold-season processes in the MAPS land-surface scheme. *Journal of Geophysical Research*, **105**, D3, 4077-4086.

Smirnova, T.G., J.M. Brown, and S.G. Benjamin, 1997: Performance of different soild model configurations in simulating ground surface temperature and surface fluxes. *Monthly Weather Review*, **125**, 1870-1884.

Smith, W. L., Jr., P. Minnis, and C. Fleeger, 2011: Flight icing threat, algorithm theoretical basis document. GOES-R Program Office.

Steiner, M., R.E. Bateman, D.L. Megenhardt, and J.O. Pinto, 2009: Evaluation of ensemble-based probabilistic weather information for air traffic management. Aviation, Range and Aerospace Meteorology (ARAM) Special Symposium on Weather – Air Traffic Management Integration, 11-15 January, Phoenix, Arizona, American Meteorological Society, Paper 4.3, 12 pp.

Steiner, M., R.A. Houze, Jr., and S.E Yuter, 1995: Climatological characterization of three-dimensional storm structure from operational radar and rain guage data. *Journal of Applied Meteorology*. **34**, 1978-2007.

Stoelinga, M.T., 2005: Simulated equivalent reflectivity factor as currently formulated in RIP: Description and possible improvements. Informal notes, University of Washington, Seattle, 5 pp. (RIP is a software package that processes output from mesoscale models and generates a wide variety of diagnostic quantities.)

Stohl, A., 1998: Computation, accuracy and applications of trajectories—a review and bibliography. *Atmospheric Environment*, **32**, 947-966.

Sun, J., D.W. Flicker, and D.K. Lilly, 1991: Recovery of three-dimensional wind and temperature fields from simulated single-Doppler radar data. *Journal of the Atmospheric Sciences*, **48**, 876-890.

Sun, J., and A. Crook, 1994: Wind and thermodynamic retrieval from single-Doppler measurements of a gust front observed during Phoenix II. *Monthly Weather Review*, **122**, 1075-1091.

Sun, J., and N.A. Crook, 1997: Dynamical and microphysical retrieval from Doppler radar observations using a cloud model and its adjoint. Part I: Model development and simulated data experiments. *Journal of the Atmospheric Sciences*, **54**, 1642-1661.

Sun, J., and N.A. Crook, 2001: Real-time low-level wind and temperature analysis using single WSR-88D data. *Weather and Forecasting*, **16**, 117-132.

Sun-Mack, S., Y. Chen, T.D. Murray, P. Minnis, and D.F. Young, 1999: Visible clear-sky and near-infrared surface albedos derived from VIRS Data for CERES. *Proceedings, 10th Conference on Atmospheric Radiation*, 28 June – 02 July, Madison, WI, American Meteorological Society, 422-425.

Szoke, E.J., M.L. Weisman, J.M. Brown, F. Caracena, and T.W. Schlatter, 1984: A subsynoptic analysis of the Denver tornadoes of 3 June 1981. *Monthly Weather Review*, **112**, 790-808.

Theriault, K.E., M.M. Wolfson, W.J. Dupree, B.E. Forman, R.G. Hallowell, R.J. Johnson, Jr., and M.P. Moore, 2001: TCWF Algorithm Assessment – Memphis 2000, Project Report ATC-297 (July), Appendix B, Lincoln Laboratory, Massachusetts Institute of Technology, 55-57.

Thompson, G., P.R. Field, R.M. Rasmussen, and W.D. Hall, 2008: Explicit forecasts of winter precipitation using an improved bulk microphysics scheme. Part II: Implementation of a new snow parameterization. *Monthly Weather Review*, **136**, 5095-5115.

Thompson, G., R. M. Rasmussen, and K. Manning, 2004: Explicit forecasts of winter precipitation using an improved bulk microphysics scheme. Part I: Description and sensitivity analysis. *Monthly Weather Review*, **132**, 519–542.

Trepte, Q., Y. Chen, S. Sun-Mack, P. Minnis, D.F. Young, B.A. Baum, and P.W. Heck, 1999: Scene identification for the CERES cloud analysis subsystem. *Proceedings, 10th Conference on Atmospheric Radiation*, 28 June – 02 July, Madison, WI, American Meteorological Society, 169-172.

Tuttle, J.D., and G.B. Foote, 1990: Determination of the boundary layer airflow from a single Doppler radar. *Journal of Atmospheric and Oceanic Technology*, **7**, 218-232.

Van Geffen, J., M. van Roozendael, J. van Gent, P. Valks, M. Rix, R van der A, P. Coheur, L. Clarisse, and C. Clerbaux, 2009: An alert system for volcanic SO_2 emissions using satellite measurements. *Proceedings of the 2009 EUMETSAT Meteorological Satellite Conference*, 21-25 September, Bath, England, 8 pp.

Weckwerth, T.M., H.V. Murphey, C Flamant, J. Goldstein, and C.R. Pettet, 2008: An observational study of convection initiation on 12 June 2002 during IHOP 2002. *Monthly Weather Review*, **136** (7), 2283-2304.

Weckwerth, T.M., 2000: The effect of small-scale moisture variability on thunderstorm initiation. *Monthly Weather Review*, **128** (12), 4017-4030.

Weiss, M., and J.E. Ghirardelli, 2009: Improvements to the Localized Aviation MOS Program (LAMP). Statistical guidance for ceiling a height and sky cover. Preprints, *23rd Conference on Weather Analysis and Forecasting / 19th Conference on Numerical Weather Prediction*, 1-5 June, Omaha, Nebraska, American Meteorological Society, Paper 6A.4, 9 pp.

Weiss, M., and J. Ghirardelli, 2005: A summary of ceiling height and total sky cover short-term statistical forecasts in the Localized Aviation MOS Program (LAMP). *Preprints, 21st Conference on Weather Analysis and Forecasting / 17th Conference on Numerical Weather Prediction*, 1-5 August 2005, Washington, D.C., American Meteorological Society, Paper 13B.6, 4 pp.

Weygandt, S., S. Benjamin, T. Smirnova, K. Brundage, C. Alexander, M. Hu, B. Jamison, and S Sahm, 2009: Evaluation of the High Resolution Rapid Refresh (HRRR): an hourly updated convecteion resoling model utilizing radar reflectivity assimilation from the RUC/RAP. *WMO Symposium on Nowcasting*, 31 August – 4 September, Whistler, British Columbia, Canada, Paper 3.17, Powerpoint presentation.

Williams, J.K., L.B. Cornman, J. Yee, S.G Carson, G. Blackburn, and J. Craig, 2006: NEXRAD detection of hazardous turbulence. Paper presented at the 44th Aerospace Sciences Meeting and Exhibit of the American Institute of Aeronautics and Astronautics, 9-12 January, 2006, Reno, Nevada, 16 pp.

Wilson, J.W., and C.K. Mueller, 1993: Nowcasts of thunderstorm initiation and evolution. *Weather and Forecasting*, **8**, 113-131.

Wilson, J.W., and R.D. Roberts, 2006: Summary of convective storm initiation and evolution during IHOP: Observational and modeling perspective. *Monthly Weather Review,* **134**, 23-47.

Witham, C.S., M.C. Hort, R. Potts, R. Servranckx, P. Husson, and F. Bonnardot, 2007: Comparison of VAAC atmospheric dispersion models using the 1 November 2004 Grimsvötn eruption. *Meteorological Applications*, **14**, Royal Meteorological Society, 27-38.

Wolfson, M.M., B.E. Forman, K.T. Calden, W.J. Dupree, R.J. Johnson, Fr., R.A. Boldi, C.A. Wilson, P.E. Bieringer, E.B. Mann, and J.P. Morgan, 2004: Tactical 0-2 hour convective weather forecasts for FAA. *Proceedings, 11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, Massachusetts, American Meteorological Society, 35 pp.

Wolfson, M.M., B.E.Forman, R.G. Hallowell, and M.P. Moore, 1999: The growth and decay storm tracker. 8th Conference on Aviation, Range, and Aerospace Meteorology (ARAM), Dallas, Texas, 10-15 January 1999, American Meteorological Society, Paper 4.3, 5 pp.

Xie, Y., S. Koch, J. McGinley, S. Albers, P.E. Bieringer, M. Wolfson, and M. Chan, 2011: A Space–Time Multiscale Analysis System: A sequential variational analysis approach. *Monthly Weather Review*, **139**, 1224-1240.

Zhang, D., Z. Wang, and D. Liu, 2010: A global view of midlevel liquid-layer topped stratiform cloud distribution and phase partition from CALIPSO and CloudSat measurements, *Journal of Geophysical Research, Atmospheres*, **115**, D00H13, doi:10.1029/2009JD012143, 10 pp.

Zhang, J., K. Howard, and J.J. Gourley, 2005: Constructing three-dimensional multiple-radar reflectivity mosaics: examples of convective storms and stratiform rain echoes. *Journal of Atmospheric and Oceanic Technology*, **22**, 30-42.