

U. S. DEPARTMENT OF COMMERCE
NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION
NATIONAL-WEATHER SERVICE
NATIONAL METEOROLOGICAL CENTER

OFFICE NOTE 179

Minutes of a Workshop on Optimum Interpolation
Held 19-20 September 1977 at the
National Meteorological Center, Camp Springs, Maryland

Development Division

MAY 1978

This is an unreviewed manuscript, primarily
intended for informal exchange of information
among NMC staff members.

MINUTES OF A WORKSHOP ON OPTIMUM INTERPOLATION HELD 19-20 SEPTEMBER 1977
AT THE NATIONAL METEOROLOGICAL CENTER, CAMP SPRINGS, MARYLAND

1. Foreword

Two large international meetings concerned with the four-dimensional assimilation of meteorological data have been held in this decade. The first was in Princeton, New Jersey, in April 1971, and the second in Paris in November 1975. At the latter meeting it seemed that many of the major modelling groups expected to participate in the First GARP Global Experiment (FGGE) were moving in the direction of statistical, or "optimum," methods of interpolation to assimilate the strongly heterogeneous FGGE data base.

By early 1977 all of the major modelling groups in the United States were at least thinking in terms of optimum interpolation, and several possessed functioning systems. Recognizing this, a proposal was made at a meeting on 3 March 1977 at the Goddard Space Flight Center that a meeting of the people involved in building optimum interpolation assimilation systems should be held.

A poll of the major groups was conducted to ascertain first if such a meeting would be useful and, if so, what its format should be. The consensus was that an informal workshop with a small attendance would indeed be useful, and that the format should focus on the engineering details of constructing a working assimilation system. The list of participants and the agenda reflect this emphasis on a small meeting concerned with application rather than theory.

As the organization of this Workshop proceeded, it became evident that interest in it was wider than had first been supposed. The number of prospective participants appeared to threaten the informal structure which was deemed essential to the success of the Workshop. Accordingly, attendance was limited by requiring that all participants be actively involved in the application of optimum interpolation methods to global data assimilation. Ruthless adherence to this requirement contributed to the success of the Workshop.

The minutes of the Workshop have been compiled by the NMC participants primarily from notes and tape recordings, occasionally supplemented by published material. We have generally adopted the attitude that the third and fourth sections represent the substance of the Workshop. Consequently, the minutes of those discussions reflect very little editing--they are nearly verbatim transcripts of the recordings. On the other hand, the second and fifth sections generally cover material which has been published, or will be published in the near future. We have not thought it necessary to include details of these topics, but instead have sought to summarize them rather severely.

The result is a document of rather uneven character. Rather than risk emasculating the content of the main discussions, we have elected to accept such unevenness as a necessary if unfortunate condition. Nevertheless, we hope the document will prove useful to those who participated in the Workshop as well as to others who have an interest in this area.

List of Participants

Ian Rutherford	Recherche en Prévision Numérique, Canadian Meteorological Centre (CMC)
Gorm Larsen Andrew Lorenc	European Centre for Medium-Range Weather Forecasting (ECMWF)
Kikuro Miyakoda	Geophysical Fluid Dynamics Laboratory (GFDL)
Michael Ghil	Goddard Laboratory for Atmospheric Sciences (GLAS)
Thomas Schlatter	National Center for Atmospheric Research (NCAR)
Edward Barker	Naval Environmental Prediction Research Facility (NEPRF)
Kenneth Bergman Doris Gordon Robert Kistler Ronald McPherson Glenn Rasch	National Meteorological Center (NMC)

Agenda

Monday, 19 September:

- 8:15 1. Opening

- 8:30 2. General Structure of Each System
 - a. CMC (Rutherford)
 - b. ECMWF (Larsen)
 - c. GFDL (Miyakoda)
 - d. GLAS (Ghil)
 - e. NCAR (Schlatter)
 - f. NEPRF (Barker)
 - g. NMC (Kistler)

- 11:30 Lunch

- 1:00 3. Input Problems
 - a. Observation Error Statistics
 - b. Data Search and Selection
 - c. Data Checking

- 3:30 4. Internal Problems
 - a. Prediction Error Statistics
 - b. Matrix Inversion Algorithms

- 4:30 Close of Session

Agenda

Tuesday, 20 September:

- 8:30 4. Internal Problems (continued)
 - c. Calculation of Residuals: horizontal and vertical interpolation
 - d. Horizontal Update Mesh in High Latitudes
 - e. Filtering
 - f. Insertion Procedures
 - g. Initialization Procedures
 - h. Treatment of Moisture

- 11:30 Lunch

- 1:30 5. Characteristics of the Results
 - a. NMC
 - b. GLAS
 - c. GFDL
 - d. ECMWF
 - e. CMC

- 4:00 6. Summary of Meeting

2. General Structure of the Several Assimilation Systems
(First Session, Monday morning, September 19, 1977)

In the first session, each group described the main characteristics of their assimilation system in order that detailed discussion might begin from a common point. The following paragraphs summarize the presentations:

a. CMC (Ian Rutherford)¹

General: This system is the only one of the seven discussed at the Workshop which is in daily operational use.

Prediction model: The updating cycle uses a primitive-equation model solved on a polar stereographic grid of 51x55 points covering most of the Northern Hemisphere. Horizontal resolution is 381 km at 60N. The vertical coordinate is normalized pressure (σ) and resolves the atmosphere into seven layers. Time differencing is accomplished by the semi-implicit method using a time step of 30 minutes. The history variables of the model are geopotential (ϕ), wind components (u and v), and dewpoint depression ($T-T_d$). Ambient temperature is a diagnostic variable. This model is expected to be replaced in the near future by a spectral prediction model with ten to twelve layers.

Updating: The updating is intermittent with a period of 6 hours. At the end of each interval, the prediction is interpolated from the σ vertical coordinate to mandatory isobaric levels prior to the analysis. The procedure is performed twice each 6 hours; once with a 1.5-hour data cutoff for early operational use, and once with a 6-hour cutoff to use late-arriving data.

Analysis: The analysis is multivariate in that ϕ , T , u , v , are treated simultaneously. In the vertical, ϕ and T influence each other. In the horizontal, ϕ , u , v influence each other, and u and v influence T . Dewpoint depression is treated separately. It is a "split" three-dimensional system: vertical interpolation of single-level data or other incomplete reports is done first, followed by horizontal interpolation. The vertical coordinate of the analysis is pressure; the analysis is done at ten isobaric levels on the same horizontal grid mesh as is used by the prediction model. Quantities interpolated are actually differences, or residuals, between the observations and a "background" field interpolated to the locations of the observations. A 6-hour forecast is used as the background field; the residuals may then be considered as forecast errors.

¹For a detailed account of the CMC system, see: Rutherford, I. D., 1976: An operational three-dimensional multivariate statistical objective analysis scheme. Proc. JOC Conference on Four Dimensional Data Assimilation, Nov. 17-21, 1975, Paris, 98-121.

The statistical interpolation requires spatial correlations of the forecast errors. A set of correlations of forecast geopotential errors, modelled on data, is assumed; wind correlations and wind-geopotential cross-correlations are then derived assuming geostrophy. Corrections to the geopotential and wind fields thus tend to be related through the geostrophic wind law. The forecast error correlations are also split into vertical and horizontal components.

Observational errors are explicitly specified for different types of data. All observations are subjected to both vertical and horizontal consistency checks. The latter is performed by a pre-analysis: a preliminary scan through the observations with the horizontal part of the analysis procedure.

b. ECMWF (Gorm Larsen)¹

General: The ECMWF assimilation system is designed to provide, on an operational basis, initial states for a prediction model, as well as Level III analyses during the FGGE. Its design emphasizes the efficient handling of large quantities of data through constructing large correlation matrices but solving each one only once, an operation which is well-suited for large vector-processor computers. Development of the ECMWF system is proceeding along the lines of constructing independent modules for prediction, analysis, and peripheral codes. The intent is that these modules will be largely interchangeable, thus maintaining a high degree of flexibility. The several components have not yet been fully linked together, so that specific characteristics of the combined system are not yet available.

Prediction model: The assimilation system will use a simplified version of the ECMWF medium-range prediction model. It is a global, primitive equation model in σ -coordinates and uses semi-implicit differencing. Both spectral and finite difference versions are being developed. The version to be used, and its resolution, have yet to be decided.

Updating: The updating will be intermittent, probably at intervals of 6 hours.

Analysis: The analysis is statistical and multivariate, treating height and wind simultaneously. It is three-dimensional, with the error correlations separated into horizontal and vertical components. Deviations of the height and wind field from a background field are the quantities analyzed. It is intended that a short-range (6 hour) prediction will serve as the background field once the complete system is linked together.

¹The ECMWF system is described in Lorenc, A., I. Rutherford, and G. Larsen, 1976: The ECMWF analysis and data-assimilation scheme: analysis of mass and wind fields, ECMWF Technical Report No. 6, Sept. 1977.

The vertical coordinate of the analysis is pressure. A coordinate transformation is therefore necessary between the prediction and analysis. The possibility of transforming only the analyzed residuals and the forecast changes is being explored, in order to minimize the error of the transformation.

A central feature of the analysis is that the correlation matrix is formed for $6^{\circ} \times 6^{\circ}$ latitude boxes rather than for individual grid points. Each matrix then need be inverted only once and the analysis points within each box may have any desired configuration or density. Furthermore, the inverted matrices may be used for horizontal consistency checks as well as for the grid point analysis.

Within each $6^{\circ} \times 6^{\circ}$ box, observations are classified as "primary"--data used both in the correlation matrix for the box and in those for neighboring boxes--and "secondary," used only locally. Each correlation matrix uses all observations within its own box, plus the primary observations from immediately adjacent boxes. A maximum of 151 observations are used; if there are too many observations in a box, "super" observations are formed by combining closely-spaced similar observations; if too few observations are in a box and its adjacent neighbors, then the neighbors' neighbors are included.

Forecast error statistics are modelled; presently, a Gaussian form is assumed. Observational errors of different types of data are explicitly specified; horizontal and vertical correlations of some errors are included.

c. GFDL (Kikuro Miyakoda)¹

General: The GFDL system features continuous assimilation of asynoptic data combined with standard synoptic analysis at 12-hour intervals. It has the greatest horizontal resolution of the global systems discussed at the Workshop. It has been used to analyze data from the GATE experiment as well as from the Data Systems Test (DST).

Prediction model: The system uses as an assimilator a global primitive equation model solved on a modified Kurihara grid mesh. Horizontal resolution is approximately 2° latitude. Nine layers are used to resolve the vertical structure of the atmosphere. Normalized pressure (σ) is the vertical coordinate. The Euler-backward time integration scheme is used continuously during the assimilation in order to reduce the gravitational noise resulting from insertion of data. History variables to be updated are temperature, surface pressure, and the horizontal wind components; moisture is not updated. A semi-implicit spectral model with 18 layers and 30 modes (rhomboidal truncation) is expected to replace the grid point model in the near future.

Updating: Asynoptic data are stratified into time blocks of 2 hours each. These data include remote sounding data, cloud-tracked winds, and aircraft reports, in addition to radiosonde observations interpolated

¹A description of the GFDL system may be found in Miyakoda, K., L. Umsheid, D. H. Lee, J. Sirutis, R. Lusen, and F. Pratte, 1976: The near-real-time, global, four-dimensional analysis experiment during the GATE period, Part I, J. Atmos. Sci., vol. 33, pp. 561-591.

linearly with time. The local analysis of these data is inserted at grid points within a 250-km radius of the observation locations at each time step during a 2-hour interval. Each 12 hours of assimilation is blended with conventional data in a global analysis.

Analysis: Both local and global analyses are statistical and univariate. The analyses are performed two-dimensionally on isobaric surfaces for temperature and for the u- and v-components. Surface pressure is also analyzed. The analysis grid is the same as the prediction grid, a modified Kurihara grid with 48 points between equator and pole. Both analyses use climatology as a background field. Up to eight observations are permitted to affect a single grid point. The local analysis has a 250-km search radius and does not use the previous forecast in any way. The global analysis has a 500-km search radius and uses the result of the previous assimilation as pseudo-data. Observational errors, including those of the pseudo-data, are explicitly assigned. The local analysis therefore represents a blend of the asynoptic data and climatology, while the global analysis blends together synoptic observations, the results of the assimilation, and climatology according to weights based on pre-assigned error levels.

d. GLAS (Michael Ghil)¹

General: The GLAS system has been specially designed to match the characteristics of remote sounding data. Continuous assimilation of satellite-derived temperatures is its central feature. In this respect, it is similar to the GFDL system except that the GLAS assimilation follows the orbiting-spacecraft rather than sorting the data into 2-hour blocks, and only remote sounding data are continuously assimilated rather than all asynoptic data.

Prediction model: The prediction model is a global, finite-difference, primitive-equation model on a latitude-longitude grid. Horizontal resolution is 4° in latitude and 5° in longitude. Normalized pressure is the vertical coordinate, with resolution of nine layers. An explicit, time-alternating, space-uncentered finite-differencing method is used, with a 10-minute time step. History variables are temperature, wind, surface pressure, and moisture.

¹Further details of the GLAS system are available in the following sources:

Halem, M., M. Ghil, R. Atlas, J. Susskind, and W. Quirk, 1978: The GISS sounding temperature impact test. NASA Technical Memorandum 78063, NASA Goddard Spaceflight Center, Greenbelt, Maryland 20771.

Ghil, M., R. Dilling, and H. Carus, 1977: A statistical method for the time-continuous assimilation of satellite-derived temperatures. Preprint Vol., Fifth Conf. on Probability and Statistics in Atmos. Sci., Nov. 15-18, 1977, Las Vegas, Nevada (AMS), 320-324.

Updating: Each 12 hours, a conventional global synoptic analysis is done. Between these updates, remote sounding temperature data are assimilated each time step following the path of the spacecraft; other asynoptic data are used in 12-hour blocks by the conventional analysis.

Analysis: The conventional synoptic analysis at 12-hour intervals is done by a successive-corrections scheme resembling that designed by Cressman. It is two-dimensional and is applied on isobaric surfaces with no vertical coupling. Temperature, wind, and surface pressure are analyzed univariately on the prediction model grid. In this type of analysis, all observations within a prescribed influence radius are used; the weight of each is proportional to the distance of the observation from the grid point. No allowance for observational errors is made. The background field, or "first guess," is the result of the assimilation of remote sounding data during the previous 12 hours.

Between the conventional analyses at 12-hour intervals, a time-continuous statistical assimilation procedure is applied. In this procedure, radiometrically-observed temperatures are used to update the prediction model at each 10-minute time step. Only the sounding data available within the previous 10-minute interval are included; spatially, this results in the updates being performed over limited areas along the path of the satellite.

Within each analysis area, the observations are first converted to deviations from the current prediction. These deviations are then interpolated statistically to the model grid points using modelled spatial correlations of the deviations. Like the ECMWF system, the correlation matrix reflects all of the sounding data within each analysis area corresponding to a 10-minute interval. Because the elements of the correlation matrix depend only on the horizontal locations of the soundings, only one large matrix inversion is required for each update, regardless of the number of isobaric levels updated or the number of grid points within the analysis area. At present, observational errors are not accounted for in the formation of the correlation matrix.¹

The statistical interpolation procedure is applied only to remote temperature soundings. Subsequent to each update, however, the winds in the analysis area are adjusted by the method of geostrophic increments² to agree with the updated temperatures.

¹In the interim between the Workshop and the publication of this document, a study of the effect of observation errors on the results of the assimilation was performed, and is currently being evaluated.

²Stone, P., L. Tsang, and D. Schneider, 1973: Balanced winds for assimilation of temperature and pressure data. GISS Research Review 1973, Part 2 (Applications), NASA, New York, N.Y. 10025, 160-163.

e. NCAR (Thomas Schlatter)¹

General: The NCAR system has been used in both global and limited-area analysis experiments, the latter having the greatest horizontal resolution of any of the seven systems discussed at the Workshop. The system has also been applied in reference level experiments with remote sounding data.

Prediction model: The NCAR General Circulation Model (GCM) is used in global assimilation experiments. It is a primitive equation model on a 2.5° latitude-longitude grid. The vertical coordinate is geometric height, with six layers each 3-km thick. History variables are the pressures of the coordinate surfaces and the wind components.

Updating: Intermittent updating is used, with an interval of 6 hours.

Analysis: The scheme is statistical in nature and is applied two-dimensionally on isobaric surfaces without vertical coupling. It is multivariate in that the updated variables--geopotential and the two wind components--are treated simultaneously, subject to a weak geostrophic constraint outside the Tropics. A feature unique to the NCAR system is that geopotentials are not analyzed in the Tropics (22.5N-22.5S); rather, the tropical winds are analyzed subject to a weak constraint of nondivergence, and the geopotentials obtained from a solution of the balance equation. Surface pressure is analyzed separately; the isobaric analyses are done for the ten standard levels from 850 mb to 70 mb.

For global analyses, the NCAR GCM provides the background field at 6-hour intervals. Since the vertical coordinate of the prediction model is height while that of the analysis is pressure, vertical interpolation is necessary. Linear interpolation of winds and logarithmic interpolation of heights is used. In some limited-area experiments over data-rich areas, persistence is used as a background field. Higher spatial resolution (1.25°) has been used in these experiments. Resolution in the global analyses is 2.5°.

¹Pertinent references for the NCAR system are:

Schlatter, T. W., 1975: Some experiments with a multivariate statistical objective analysis scheme. Mon. Wea. Rev., vol. 103, 246-257.

Schlatter, T. W., G. W. Branstator, and L. G. Thiel, 1976: Testing a global multivariate statistical objective analysis scheme with observed data. Mon. Wea. Rev., vol. 104, 765-783.

Schlatter, T. W., G. W. Branstator, and L. G. Thiel, 1977: Reply. Mon. Wea. Rev., vol. 105, 1465-1468.

The correlation functions are based on spatial correlations of deviations of observed quantities from climatology; the height-height correlation is modelled by a damped cosine curve. This function represents the observed negative correlations at large distances better than the commonly-used Gaussian curve. Other needed correlation functions are derived from the height-height correlation through the assumption of geostrophy or nondivergence.

Random observational errors are approximately specified. Data selection for the correlation matrix depends only on distance; the nearest five observations are used. The search procedure is in spherical geometry rather than "boxes," as in the ECMWF system.

f. NEPRF (Edward Barker)¹

General: The NEPRF system is in the initial stages of development. It is expected that a version of the data assimilation system will be under-going evaluation by the summer of 1978 and will be operational by the fall of 1979, when a new computer will be functioning. Several features of the planned system are unique, including the use of Sasaki's "noise-freezing" technique to assimilate remote sounding data, and the use of variational initialization.

Prediction model: The assimilation system will use a spectral form of the UCLA general circulation model. Details of the configuration, including horizontal and vertical resolution, have not yet been decided.

Updating: Intermittent updating at 6-hour intervals will be used for conventional data. Remote sounding data will be continuously assimilated using Sasaki's method.

¹Documentation of the NEPRF approach to data assimilation is contained in the following references:

Haltiner, G., Y. Sasaki, and E. Barker, 1976: An initialization technique for primitive equation models using a balance-equation constraint. Proc. JOC Conference on Four Dimensional Data Assimilation, Nov. 17-21, 1975, Paris, 198-223.

Barker, E., G. Haltiner, and Y. Sasaki, 1977: Three-dimensional initialization using variational analysis. Vol. Conf. Papers, Third Conf. on Numerical Weather Prediction, April 26-28, 1977, Omaha, Nebraska, (AMS), 169-181.

Sasaki, Y., and T. Baxter, 1977: Assimilation by the noise-freezing technique. Vol. Conf. Papers, Third Conf. on Numerical Weather Prediction, April 26-28, 1977, Omaha, Nebraska, (AMS), 460-471.

Analysis: The analysis will consist of three parts: the conventional analysis of synoptic data to provide independent estimates of temperature and wind, the continuous assimilation of remote sounding data to produce an estimate of 1000-300 mb thickness, and a blending of temperature, wind, and thickness through variational initialization.

The analysis will be carried out on a 2.5° global grid with pressure as the vertical coordinate; eleven isobaric surfaces will be used. Presently, the conventional analysis uses the successive-corrections method. It produces univariate analyses of geopotential and the two wind components using the prediction as a background field. Statistical interpolation is expected to replace the current method for operational use. Remote sounding data will be treated as observations of 1000-300 mb thickness primarily for use in the Southern Hemisphere. They will be assimilated through Sasaki's method, which uses a global shallow-water-equation model and a technique to isolate gravitational modes and "freeze" them while carrying the meteorological model forward in time. The result of this step is a blending of the univariate analyses of temperature, thickness, and winds through a variational technique. The technique is three-dimensional and incorporates dynamic constraints which improve the initial balance of the mass and motion fields.

g. NMC (Robert Kistler)¹

General: The NMC system is intended for operational use with a long data hold (~ 10 hours) during the First GARP Global Experiment. Its purpose is twofold: to provide FGGE Level IIIa analyses, and to supply a background field for the NMC operational (3.5-hour data cutoff) cycle. The design of the system emphasizes updating in the prediction model's vertical coordinate, thereby reducing the amount of required vertical interpolation.

Prediction model: A global, primitive-equation model on a latitude-longitude grid is used. The vertical coordinate is a modified σ (normalized pressure) system which divides the atmosphere into two

¹Preliminary descriptions of the NMC system may be found in:

Bergman, K. H., 1976: Multivariate objective analysis of temperature and wind fields using the thermal wind relationship. Preprints, Sixth Conf. on Weather Forecasting and Analysis, May 10-13, 1976, Albany, N.Y. (AMS) 187-190.

McPherson, R., K. Bergman, R. Kistler, G. Rasch, and D. Gordon, 1977: Global data assimilation by local optimum interpolation. Vol. Conf. Papers, Third Conf. on Numerical Weather Prediction, April 26-28, 1977, Omaha, Nebraska, (AMS), 444-459.

domains corresponding to the troposphere and the stratosphere. A material-surface "tropopause" separates the two domains. Nine layers model the vertical structure of the atmosphere; six are in the troposphere and three in the stratosphere. The upper boundary of the prediction model is at 50 mb. Horizontal resolution is 2.5° . History variables are layer mean temperature, the two wind components, specific humidity, and the pressure thicknesses of the two σ domains. Time differencing is by the leapfrog method with pressure gradient averaging to permit a 10-minute time step. A strong time filter is included to damp the update shock.

Updating: The system features intermittent updating, in the σ -coordinate, at 6-hour intervals.

Analysis: The history variables (u, v, T) are updated using three-dimensional statistical interpolation of deviations from a background field provided by the prediction model. First in the sequence, however, is the treatment of the pressure-thickness of the σ -domains. This is accomplished by separate updating of the surface pressure and the "tropopause."

The surface pressure is updated by statistical interpolation of differences between the predicted pressure at the model terrain height and the observed station pressure. Both observation and background field have had the standard atmosphere pressure for their respective altitudes subtracted. In addition, a hydrostatic correction is applied to the observations in areas of high terrain in order to account for the difference between the model terrain height and the height of the observing station. The interpolation procedure is two-dimensional on a 2.5° latitude-longitude grid and uses only pressure data over land, but includes wind data also over marine areas. Observational errors are explicitly specified.

For the "tropopause," the background field is the model's prediction. Updating is done on a 5° mesh by univariate statistical interpolation of reported tropopause pressure. Radiosonde observations are used directly; tropopause pressures are calculated for remote sounding profiles. Stringent controls are enforced so that the analyzed surface does not stray outside climatological limits, and does not change drastically from one update to the next.

Updating the surface pressure and "tropopause" also redefines the vertical structure of the prediction model. Temperature and the winds are updated at the midpoints of the redefined layers of the model, at intersections of a 5° latitude-longitude horizontal mesh. At present the residuals (observation minus forecast) to be interpolated are formed on isobaric levels, since most observations are available there. The treatment is, in principle, multivariate; residuals of all three variables are interpolated simultaneously. The required horizontal correlations are obtained by modelling the temperature autocorrelations by a Gaussian curve and deriving the remainder through the thermal wind equation. The

data base consists of all reports of geopotential and wind within ± 3 hours of the update time. All geopotentials are converted to thickness temperatures prior to the interpolation. Random observational errors are explicitly specified. Spatially correlated errors of remote sounding data are also accounted for. The error of the prediction model is obtained by augmenting the estimated error of the previous analysis by an amount approximating the error growth rate of the model. The prediction error is therefore a function of space and time. At a particular point, it depends exclusively on how recently the point has been updated and on the quality, quantity, and distribution of data that affected the update.

The analysis proceeds on a gridpoint-by-gridpoint basis; if there are no observations within a 15° (latitude) radius of a point, the background field at that point is unaltered. A maximum of ten observations are permitted to affect each update; data selection is done on the basis of maximizing the correlation between the observations and the gridpoint to be updated. All data are subjected to a hierarchy of checks, varying from internal consistency tests in the operational preprocessing step to a spatial consistency check among neighboring observations in the analysis itself.

Specific humidity is updated in the lowest six layers of the prediction model, on a 5° mesh. Univariate statistical interpolation is used.

The interpolated residuals of thickness temperature, wind, and specific humidity on the 5° update mesh are interpolated to the 2.5° prediction mesh through a spherical harmonic interpolation function with 36-mode triangular truncation. The residuals are then recombined with the background field, thus completing the update.

h. General discussion

With respect to the NMC system, Miyakoda expressed the opinion that updating in the prediction model's coordinate makes the analyses too model-dependent. Isobaric analyses are preferred because of their greater generality.

McPherson replied that the philosophy underlying the NMC choice of updating in the σ -coordinate was to avoid incidental or unintentional changes to the background fields in the absence of data. In isobaric analyses, the background field is interpolated from σ to pressure before the analysis; in the NMC system, the analogous interpolation is done after the analysis and is done for display only. Interpolation errors are thus not permitted to feed back into the cycle.

Ghil voiced a concern that, while it is necessary to go forward with the implementation of existing assimilation systems because of the impending start of the FGGE, it is well to not lose sight of the fact that the present systems are far from optimum and require much further research. In particular, there is a need for integration of the efforts of those involved in temperature retrievals from satellite radiance data and those concerned with the assimilation of those data. Retrievals up to the present have been done one column at a time. Suggestions have recently been made that incorporation of radiance data from surrounding areas would improve the quality of the retrieved temperature profiles. This amounts to analysis of radiance fields, which is the province of the assimilators.

Rutherford expressed strong support for the view that assimilating retrievals is the wrong approach; the retrieved temperature is an unnecessary intermediate variable. Rather, radiance data should be assimilated directly.

Lorenc, however, pointed out that a fixed interface is necessary between data producers and data users. If radiances from a set of channels are used as the interface, they would be subject to changes in instrumentation. Retrieved temperatures now serve the interface function quite effectively.

Larsen, in commenting on Ghil's call for further research in assimilation methods, stated his view that the assimilation method is a problem mainly in data-sparse areas; as the data base improves, the influence and importance of the assimilation method diminishes. As long as data-sparse areas exist, it may be more profitable to improve the prediction model as a means of transferring information into those regions, rather than developing more sophisticated analysis schemes.

Miyakoda expressed the philosophy at GFDL as moving in the direction of less model dependence in the analysis procedure. To the extent that model errors contaminate the analysis, future research which may depend on those analyses is hindered.

3. Input Problems (Second Session, Monday afternoon, September 19, 1977)

a. Observation error statistics

Rasch presented values of observational errors used in the NMC global optimum interpolation analysis. The satellite thickness temperature errors as a function of height were from a study by Hayden, and their horizontal correlation from studies by Schlatter and Polger. The remaining errors were based on estimates available in the literature or were merely intelligent guesses.

Bergman pointed out that it is the ratio of observational error to forecast error which actually appears in the analysis equations. Kistler explained how the analysis error is updated to obtain the forecast error for the next cycle in NMC global data assimilation. At each gridpoint, the analysis error is incremented by an assumed forecast growth of error for a 6-hour period. This incremented value is the forecast error for the next cycle's analysis. Lack of advection and diffusion of forecast errors so obtained leads to unrealistically large forecast errors in some areas after several cycles. Kistler plans to incorporate advection and/or diffusion of forecast errors plus a ceiling on the maximum allowable error.

Schlatter showed the results of his study of the horizontal correlation of Nimbus satellite observational errors. He interpolated an analysis based on rawinsonde data only to the locations of satellite soundings and compared values. His results compared closely with those of Polger, although the methods used were different. (Polger compared rawinsonde-satellite sounding pairs located with a 2° latitude radius, ± 3 -hour space-time "window.") The horizontal correlation at three levels fell to zero at about 900-km separation.

Schlatter also showed satellite error (as compared to the analysis) as a function of height for 1589 Nimbus soundings. RMS error is greatest at surface and near tropopause, least in midtroposphere. Mean error is always less than 1°C . These results are similar to those of W. Smith and C. Hayden (NESS). Although these errors were not corrected for error in the analysis (presumably smaller), only a small reduction in the magnitude of satellite errors would be expected to result.

Finally, Schlatter showed the standard deviation of satellite error as a function of altitude, noting that it is about two-thirds the natural variation of temperature with altitude. McPherson commented that the satellite errors are too large for the data to be of value in the Northern Hemisphere.

Rutherford noted that he found VTPR satellite-minus-forecast temperature differences to be consistently greater than rawinsonde-minus-forecast differences, but that this may merely indicate that the forecast error is larger in areas observed by satellites. A discussion followed on the difficulties of comparing observational errors of different types.

Larsen noted that satellite observations should provide good gradient information because of the spatially correlated errors.

Ghil mentioned that error characteristics may be quite different for VTPR and for Nimbus data. He also indicated the need to allow for the presence of mean error (bias) in the satellite observations.

Bergman showed the effect of correlated observational errors on the analysis at a grid point by preparing an example analysis with correlated errors and comparing it to the analysis obtained when the same observations have uncorrelated errors. The tendency of observations with correlated errors to extrapolate the gradient between them was noted.

Rutherford noted that the effect of correlated errors on analysis error is such as to increase it for absolute values but to decrease it for gradient measurements. Lorenc mentioned that he has run experiments which show that horizontally correlated temperature errors lead to an improved wind analysis, but worsened temperature and height analyses. The inclusion of positive vertical error correlation makes everything worse.

Rutherford asked if anyone has looked at the vertical correlation of satellite errors. Lorenc stated that he suspects the satellite vertical errors for thickness temperature may be negatively correlated because of compensating thickness errors. McPherson said that NESS studies indicate such for VTPR data. Bergman indicated that Systems Evaluation Branch of NMC plans to compute vertical, as well as horizontal, error correlations for Nimbus data. Lorenc speculated that satellite error correlation characteristics may be a function of retrieval method, and stated that information needs to be provided at the interface between NESS and users of satellite data.

b. Data search

A review of the number of predictors used in updating grid point values was made. GFDL uses eight predictors, NMC ten, Canada eight, NCAR uses five data points (up to fifteen numbers if all observations are complete). GLAS and ECMWF have a variable number of predictors. ECMWF uses

"super-observations," up to a maximum of 151 numbers; GLAS has no a priori limit, but in the present implementation the number rarely exceeds 70.

Gordon described the NMC data search. In it, every piece of data is treated as a separate observation. The data are sorted by latitude belts and by longitude within each belt. The data are not sorted by level. A stepwise search of the data about each grid point to be updated is made. Successive boxes of approximately 5° , 10° , 20° , and 30° on a side, centered on the grid point, are scanned until the following criterion is met. A combination of complete radiosonde and satellite soundings is required which adds up to 6, with each RAOB = 2 and each SATOB = 1. Once this condition is satisfied, the search stops and all the correlations between each level of the grid point and the observations within the largest box scanned are approximately determined from a table look-up. To allow for varying quality of observations, these correlations are divided by $(1 + \sigma_{\epsilon}^2)$ where σ_{ϵ} is the assumed error standard deviation of the instrument which took the observation. Those ten observations with the highest resulting values are the ones actually used to update the grid point at the particular level.

Lorenc reviewed the ECMWF method which uses up to 151 observations at a time. Some of these observations are formed by combining several individual observations located close together. The same 151×151 matrix of correlations is used for all variables (only the right-side vector changes) and for several neighboring grid points. An assumed observational error is added to the diagonal terms of the matrix. Among other things, this prevents the matrix from becoming very ill-conditioned.

Lorenc commented that it is preferable to use both the u and v component of a wind observation in the updating of a grid point, rather than just one or the other (as is possible with the NMC selection method). Observations which have low correlation with the grid point value may nevertheless turn out to be very useful observations when the inter-observational correlations are taken into account.

Bergman indicated that it is difficult and expensive to allow for inter-observational correlations in the selection process. Selecting solely on the basis of the observation-grid point correlation (adjusted for observational error) is quick, although it admittedly doesn't always select the ten predictors which would receive the most weight if all available predictors were used to update the grid point. In the majority of cases, only two or three of the ten observations receive most of the assigned weight.

Lorenc mentioned that in the ECMWF scheme, all of the data (some of it lumped into super-observations) in the analysis area is used, thus the problem of getting the best predictors is avoided. The same is true

of the GLAS scheme, where the lumping is done by averaging satellite observations to model grid points.

Rutherford stated that in all cases where the data selected for performing the update represents a severe restriction on the total number of observations possessing correlations above the noise level, experience indicates that a reasonable result is obtained in most cases in that the analysis is not very sensitive to the selection procedure. Experiments indicate that use of as few as four or as many as one hundred data points usually makes little difference because most of the information is in a few (usually close) points and the others contribute little. One can devise individual analysis situations where including or rejecting an observation may make a lot of difference, but such situations are uncommon in practice.

One can get around the problem of the inter-observational correlation effect by doing stepwise regression on a larger set of observations than those actually used in the analysis, but this procedure can be expensive.

Ghil mentioned the possibility of using a matrix approximation method, common in finite element theory, where one matrix is replaced by another which is diagonal or nearly so. Preliminary studies in this direction were undertaken by GLAS, but are not completed. Rutherford and Larsen indicated that ECMWF does something like this in forming "super-obs." This procedure deletes many rows and columns of the original matrix, getting it down to a reasonable size.

Rutherford indicated that pre-combining observations into "super-obs" should save time overall, since analyzing grid point by grid point uses the same observations over and over again in data-dense areas.

Larsen mentioned that it might be worthwhile to reorder the matrix in order to make it more diagonal. Ghil added that there are ordering algorithms available in the form of software packages. One of these is due to J. Alan George at Waterloo, Ontario, Canada. He suggested making contact with George.

Lorenc then gave some details on the "superobing" process used at ECMWF. It is convenient to combine nearby compatible observations at the pre-analysis stage to form a "super-observation." This super-observation is then used in the analysis as an ordinary observation of increased accuracy.

If the standard optimum interpolation technique were used to create the super-observation, then the interpolated value would contain information from, and hence be correlated with, the predicted value. This is inconvenient since for ordinary observations it is assumed that there is no correlation with prediction. So the interpolation equations

used for super-observation formation are modified by imposing a constraint that no information is taken from the local predicted value.

Most super-observations are formed between two or more similar very close observations on a univariate basis. In this case, the simplifying assumption can be made that the observations are co-located for the purpose of forming the super-observation.

c. Data checking

Rutherford outlined the data checking procedure at CMC. A preliminary interpolation is done to each data point with its datum not included. A datum is rejected if the difference between it and the interpolated value exceeds some specified multiple of the computed interpolation error. This generally works well, but on occasion one bad report may cause rejection of nearby good ones. One way to handle this problem is to have consecutive interpolation passes excluding one-at-a-time each datum used in the interpolation. This has been tried successfully but is fairly expensive.

The rejection criteria are set such that a bad or marginal datum is accepted occasionally, with the hope that it will be controlled by nearby good data, rather than risk throwing out good data with the bad. Too tight a rejection criterion tends to throw out data which is good but "unusual." This is precisely the data which is most useful in the analysis.

A manual monitor allows data which has been rejected by the objective scheme to be reinserted. The total error-checking routine takes about 1 minute of the 6 required for the entire analysis package. This is a function of the ratio of observations to the number of grid points.

At the ECMWF, error checking takes very little additional time since the inverted matrices give the analysis at any point, including both data and grid points, within a specified region, and a datum is rejected by suppressing the weight that the datum would otherwise receive.

The biggest problem in error checking occurs when two observations are close together. An error checking scheme cannot distinguish if one is bad and the other is good, but will either accept or reject them both. At the ECMWF, a two-part scheme is being considered where first all the data are checked, and second all the rejected data are rechecked against the accepted data, with the option of reinserting rejected data. This procedure might permit the distinction between good and bad observations which are close together to be made.

Schlatter described the error checking routine used at NCAR. The nearest five neighboring observations to a given observation are found. These observations are averaged four at a time and compared with the

observation being checked. This is done for the five possible combinations. Rejection criteria depend on the difference between the five averages and the checked observation. All five averages have to disagree with the observation in order to toss it. This procedure prevents bad observations from tossing good ones and works well where data density is relatively high. If data are sparse, problems may occur, most notably the rejection of an isolated report of a strong wind. The choice of five observations averaged four at a time is arbitrary.

Bergman next described the NMC error checking scheme, currently in the process of being implemented in the global analysis. The scheme saves some time by using the matrix of correlations between observations that has already been computed for the analysis. The differences between residuals of observations are computed for each pair of the same type. One or both of the observations are flagged if the difference does not satisfy the following inequality:

$$\text{DIFF}_{ij} \leq (a - b\rho_{ij})\sigma_{\text{guess}}$$

where ρ_{ij} is the correlation between the two observations, σ_{guess} is the standard deviation of the forecast error, and a and b are arbitrary constants.* If the observations are of equal quality, both are flagged; otherwise, only the observation judged (by preliminary vertical consistency checks) to be of lower quality is flagged.

A matrix of flags between pairs of observations is generated in this manner. The total number of flags which each observation receives is determined. That observation with the largest number (≥ 2) of flags is tossed, and its row and column deleted from both the correlation matrix and the matrix of flags. Then, for the remaining observations, the same procedure is repeated until all remaining observations have no more than one flag.

The performance of this scheme has been satisfactory in the regional analysis version and in the global surface analysis and is very economical. One drawback is that the observations must be tossed anew at each grid point. This means that it is possible for an observation to be tossed at one grid point but retained at a neighboring grid point. This is most likely to happen when the suspect observation is peripheral to the other nine observations used in the analysis of a grid point. The observation is likely to be retained though probably bad. Since it usually receives a very low weight in this event, it does minimum harm to the analysis.

*Current values are:

	a	b
Temperatures (°C)	8.4	5.4
Winds (m/sec)	11.2	7.2
Surface pressure (mb)	12.0	10.0

The NMC regional analysis package doesn't toss much data all told, and the gross error check catches most of these. In the global surface analysis, Kistler has noted that observations with large negative pressure residuals occur in rapidly deepening storm situations. These observations are flagged by the gross check and are used if they are not tossed by other observations in the buddy check, but they are not allowed to toss other observations.

There is no manual intervention in the data check for the experimental analysis program; however, manual intervention to restore tossed data may be necessary in operational programs.

4. Internal Problems

a. Prediction error statistics

Rutherford described the system used in Montreal. Prediction error statistics are used; these are model dependent to some degree. The correlation statistics are modelled with analytic functions which incorporate certain constraints between them. Several sets of statistics have been computed for various models, but the set still used operationally is the one derived from the earliest model used at Montreal. Initially, "guessed" values of the correlations were used to get the cycling started, and the prediction error statistics were updated at every cycle. There's a possibility of incestuous feedback in such a system, so that the ultimate model statistics may depend markedly on the original guessed values. Correlations of T, z, u, and v were computed at ten levels and checked to see how well geostrophic and hydrostatic constraints were actually obeyed by statistics. It turned out that they were obeyed fairly well.

A scheme was evolved for separating the vertical observational error correlations from the prediction error correlations. The z prediction error statistics from that study are tabulated in a 10 x 10 matrix, and all the other vertical correlations are computed from that one assuming that geostrophy and hydrostaticity apply. The horizontal z and T correlations are modelled with a Gaussian function, which may not be the best form to use, but it is convenient and reasonably satisfactory.

In the Canadian system the correlations are independent of horizontal location, but the variances are assigned values which vary with latitude and level. Actually, the wind variance is assumed to be independent of latitude, and the height variance is computed as a function of latitude from that assumption. This causes the height prediction error variance to vanish at the equator.

As for initialization, one could analyze heights and winds independently and rely on some other initialization procedure to couple them so that integration proceeds with a minimum of gravity wave noise.

Lorenc noted that the ECMWF scheme uses all observations over an area in the analysis of several grid points within the area, thus the constraints are applied precisely between the height and wind corrections. Thus, there is the possibility of a jump in the analyzed fields when moving across a box boundary to an adjacent box, where a partially different set of observations also has the geostrophic constraint applied precisely.

Rutherford mentioned that there is an escape from such a precise enforcement of the geostrophic constraint. If one knows that the constraint is 70% correct, the cross-correlations can be multiplied by 0.7, hence relaxing the enforcement of the geostrophic constraint.

Larsen stated that they found that multiplying the cross-correlations by .99 eliminated some of the ill-conditioning problems that they were having. Rutherford suggested that the ill-conditioning arises in such cases because the combined effect of wind observations and of height gradients implied by height observations may result in the height gradient in an area being specified with a very low effective error level. Reducing the correlation slightly seems to prevent this.

In response to a question about how much the cross-correlation between height and winds should be reduced, Bergman noted that cross-correlations are being computed at NMC for both forecast and climatology error statistics. The cross-correlation should rise to a maximum of 0.605 according to geostrophic theory if the height autocorrelation is modelled with a Gaussian function. The maximum cross-correlation actually observed will be something less than that. The percent difference should be the desired reduction. These computations will soon be available, but the result obtained from the forecast error statistics may be model dependent.

It was noted by Bergman and Lorenc that the shapes of the cross-correlation functions are fixed once the basic autocorrelation function is specified. This led to problems when the cross-correlation function was replaced with an approximate table of correlations in the NMC scheme. Lorenc noted problems when the correlation was allowed to vary as a function of latitude within one of their analysis boxes. Thus, the correlations must be "preset" so they do not vary within a box in the ECMWF system.

b. Matrix inversion

Lorenc and Ghil each presented their respective methods of matrix inversion. (Since these presentations involved putting equations on the blackboard, it is not possible to give details here.) From the presentations, it appears that the ECMWF method is direct matrix inversion, giving the estimated analysis error at each grid point as a by-product. Ghil's method does not invert the matrix directly and is possibly faster, but does not give the analysis errors.

Experiments were done with the Canadian barotropic model which showed that when geostrophic coupling is used, the gravity noise is not too large. Actually, mass-motion balance was introduced in the analysis scheme not for the purpose of removing gravity waves from the analysis and resulting forecast, but rather to retain the analyzed change in the balanced mode. One could also use variational methods to bring this about.

The analysis scheme used at CMC tends to be univariate in data-rich areas because of the data selection procedure. This is an advantage rather than a difficulty. It doesn't make sense to insist on strict satisfaction of constraints which are only approximately satisfied by the atmosphere. It's not the way to reduce gravity wave activity to zero. Provided sufficient data are available, the data themselves should define the correctional relationship between two fields reasonably well--probably better than the constraint--provided the data are relatively free of large errors. On the other hand, it's important to relate the correlations with a constraint in areas of limited information, such as information on the mass field alone. Studies by Hayden, and by McPherson and Kistler, confirm that, if only mass information is provided, it helps to make an adjustment to the motion field at the same time. Presumably, the reverse is also true. The multivariate scheme incorporates this kind of adjustment.

The actual constraints on the atmosphere are much more complicated than the geostrophic constraint and involve triple correlations. One can derive these constraints; they involve several more terms. But the effect that these terms would have on analysis results is probably too small to be worth the effort of including them.

The geostrophic constraint is actually satisfied fairly well by the atmosphere in a gross statistical sense. This doesn't say that it's a good approximation in particular instances. But of course it is used in particular analysis situations, so sometimes it's a good approximation and sometimes a bad one; hopefully, more often good than bad.

A short discussion followed on the magnitude of the ageostrophic wind component. Ghil cited Neiburger, et al., (1948), as giving the mean ageostrophic wind component as 30% of the geostrophic wind. Schlatter mentioned that the Cressman analysis scheme assumes the geostrophic wind is 108% of the actual wind. Miyakoda noted that the geostrophicity of the wind must be a function of latitude and Bergman noted that a statistical study of the geostrophicity of the wind in the tropics shows that this is indeed the case.

Rutherford noted that although it's not particularly logical to let the degree of geostrophic relation or unrelation depend on the arbitrariness of the data selection procedure, in fact that is what happens with most of the multivariate schemes.

Rutherford described the iterative scheme used at Montreal. This scheme, the method of steepest descent, is one of several similar schemes which usually all give similar results. But the steepest descent method has the advantage that it approximates the calculation of what is called the deflated inverse. To elaborate, one could compute the matrix inverse by first calculating the matrix eigenvalues and then calculating the inverse matrix as an expansion over the eigenvectors of the matrix. One way of removing ill-conditioning from a matrix is to delete the eigenvectors with very small eigenvalues. But the eigenvector expansion is expensive. Steepest descent has the property that it starts with a first guess solution, and then removes the error from that guess in such a way that it attacks first that part of the error associated with the largest eigenvalues. So, if the iteration is truncated well before complete convergence, the solution will have converged to the eigenvectors with large values but not to those with small eigenvalues. This solution approximates that obtained with the deflated matrix, and it is relatively insensitive to ill-conditioning. Besides, it is fast, requiring only a small number of scans rather than convergence, and it seems to work quite well in practice.

Bergman noted that the iterative method of conjugate gradients is used in the NMC system. Both schemes were tried and conjugate gradients chosen over steepest descent simply because it converges more rapidly. This may not always be desirable in light of the above discussion.

Rutherford stated that he has done experiments with standard situations which give trouble (ill-conditioning) and has found one or the other method to be better, depending on the particular situation. But generally, although the method of conjugate gradients does converge more rapidly in well-conditioned cases, it tends to give a slightly worse result in poorly-conditioned cases.

Bergman noted that, perhaps because the NMC scheme pushes convergence too far in ill-conditioned cases, the resulting estimated analysis error for the solution occasionally is imaginary. Those cases are handled by scanning the matrix and deleting one of the pair of observations which are most highly correlated and re-solving the system, which usually gives a satisfactory result. This is because two observations located close together are a common cause of ill-conditioning. Currently, the deleted observation is not used at all, but it would be preferable to assign one-half the weight obtained for the remaining observation of the pair to each of the observations. But the number of analysis situations where the problem occurs is a very small percentage of the total.

Lorenc commented that when ECMWF investigated the source of ill-conditioning in very large matrices, it was found that many observations taken together cause the ill-conditioning. A typical example would be several height observations in the eastern part of the box combining with several heights in the western part to give essentially the same

gradient information as meridional wind components near the middle. Two observations close together do not appear to cause ill-conditioning in the ECMWF scheme.

McPherson noted that the ill-conditioned problems don't occur often enough in practice to be of great concern. More important in choosing a scheme is the amount of time it uses. In an operational framework, one has to be very careful of time spent. A study comparing the economy of various schemes indicated that the conjugate gradient iterative scheme is the quickest overall.

Ghil and Lorenc remarked that what method is faster may depend on the computer system used and on how the method is coded for that computer.

Schlatter mentioned that the NCAR system uses Cholesky decomposition followed by Gaussian elimination. (This is apparently similar to, but not precisely the same as the ECMWF method.) The program runs on a CDC 7600, and the matrix inversions take about 20 to 25% of the total time. The very few cases where trouble occurred were not associated with two observations close together but rather with a peculiar overall configuration of the observations. When it occurred close to the boundaries (of the regional version), it appeared to be an extrapolation problem between the edge of the grid and the nearest observations inside.

4. Internal problems (continued) (Third Session, Tuesday morning, 9/20)

c. Calculation of residuals: horizontal and vertical interpolation

Since analyses are done on pressure surfaces, and only climatology is used as a background field so far, and since observations are moved to standard levels, ECMWF has not yet had to perform any vertical interpolations. GFDL uses cubic splines for vertical interpolation (in log p). Cubic splines were found to be superior in the vicinity of the tropopause. Horizontal interpolation is bilinear. Vertical interpolation in the NCAR model is quadratic in log p, but is nonreversible, and bilinear interpolation is used in the horizontal. CMC uses cubic in the vertical (in log p) and bi-cubic in the horizontal. NMC and GLAS both use linear in log p in the vertical and bilinear in the horizontal.

Barker asked if anyone had tried p^{2051} for vertical interpolation, which is linear for a standard atmosphere. The UCLA model uses it because it results in reduced noise around high terrain. It works well for the troposphere. No one else in the group had tried such a scheme.

d. Horizontal update mesh in high latitudes

This item was placed on the agenda by NMC because of the problem NMC has with converging meridians at high latitudes. CMC has no similar problem since analysis is done on a polar stereographic grid. ECMWF is

also somewhat divorced from the problem because only one matrix need be inverted for each box. GFDL uses a modified Kurihara grid to get around the problem NMC has. NMC has considered doing analyses on an equal area grid and forecasts on another grid. This would present no particular disadvantage since NMC already uses a spectral filter and could use the same filter to interpolate the analysis from the equal area grid to the forecast grid.

A question was brought up by Larsen as to what resolution ought to be used generally for global assimilation. Rutherford stated that it depends upon data density. Resolution should be no finer than the finest station spacing.

McPherson asked if anyone had given consideration to doing analyses on a coarser resolution grid than the prediction model. Miyakoda responded by saying that there is no reason for the resolutions to be the same. In fact, they should not be the same at high latitudes. GFDL updates on a Gaussian grid when doing dynamic assimilations. The global analysis, however, employs a modified Kurihara grid which also has somewhat denser grid points near the poles than the Kurihara grid. The modified Kurihara grid has about 2° resolution.

Rutherford pointed out that the count of radiosondes over the United States is about the same as the number of grid points over the United States on a 381-km grid. Ghil stated that if satellite data are included, they can be produced at whatever density the user desires. Also, data density should be considered in four-dimensional space.

Larsen commented that maybe it should be required that we be able to analyze waves down to a length of 1000 km. At least four grid points are needed to represent such a wave, which gives a rough estimate of what grid length would be required. Schlatter pointed out though that what you actually can obtain is still a function of station spacing. But if we assume that data producers ask us what we want, Lorenc stated, then maybe we can use Larsen's argument.

Bergman commented that the density of grid points could be determined locally by the density of data in that area. This approach, however, requires considerable programming effort.

Lorenc pointed out that analyses must be sent to users, and this may require sending too many tapes, since we are talking about a full year. It may be desirable to archive analyses in spectral form.

Miyakoda said there is evidence to suggest that spectral representation with rhomboidal wave number 25 is comparable in resolution to a 381-km NMC grid. GFDL is trying to produce Level IIIb data using a spectral model and is thinking of using rhomboidal wave number 30, which is considered barely adequate for Level IIIb purposes.

Discussion arose concerning use of rhomboidal rather than triangular truncation. Triangular theoretically should be better, but rhomboidal is easier to program and is higher in resolution at high latitudes.

Lorenc pointed out that resolution cannot be decided solely on the basis of initializing numerical models. Users of Level IIIb data will be researchers who may want to locate fronts or have other uses for the data than initializing forecast models.

e. Filtering

NMC filters its forecasts with a spectral filter prior to updating, and also filters the residual fields with the same filter using 36 modes. Filtering is done in part because of NMC's decision to limit the number of observations affecting a grid point to 10, which results in some noisiness.

Lorenc stated that ECMWF has a problem which has not yet been solved and would like some comments on it. ECMWF has a choice of analyzing increments in wind and height or streamfunction and height, and from these increments the nondivergent winds can be derived. Within boxes, the two procedures should be identical. Across boundaries, if discontinuities occur, should the jumps be nondivergent and geostrophically balanced, or should noise be in gravity waves? It would be best to select enough data from a neighboring box so jumps do not occur between boxes.

Rutherford suggested that the streamfield only could be selected and heights corrected to fit the streamfield. However, discontinuities could also occur in the wind field. The severity of discontinuities should become less when corrections are smaller, i.e., when ECMWF progresses to the point where it uses forecasts rather than climatology as a first guess.

Schlatter said he encountered noisy analyses in a cycling mode doing limited area analyses over the United States. Contours became jumpier with time using persistence as a first guess. The problem was cured by using a 9-point smoother.

Lorenc pointed out that such instability in a repeated application of a statistical analysis method can be due to using inappropriate statistics, similar to having too large a diffusion term in a forecast model.

McPherson stated that one motivation for NMC's use of a filter on the residual field is to reduce the effect of an occasional bad piece of data.

Miyakoda asked whether such filtering damages the representation of fronts. McPherson responded that although 36 modes are used, it is suspected that there is some loss. However, since the analysis is on a 5° grid, the loss is probably not terribly great.

Rutherford found that when 29 waves are used in CMC's spectral assimilation model (6-hour cycle), noise accumulates, but that if the spectrum is truncated at 26 waves, the problem is cured. Barker stated that using a variational method and the UCLA model, no smoothing is required. The variational method has the advantage of adding only meteorological modes to the model.

f. Insertion procedures

GFDL's local optimum interpolation analysis is done at 2-hour intervals. The same updated value is inserted at every time step of the model over the 2-hour period for those grid points updated by new data. The local optimum interpolation analysis is done on constant pressure surfaces and interpolated in the vertical using cubic splines. The local optimum interpolation analysis is univariate. Winds are not adjusted by temperature data. In the future, GFDL will try a balance relationship--but only in the Southern Hemisphere. The only balancing at present is that which is performed by the model during the forecast. The global optimum interpolation analysis is also univariate. In the local analysis each piece of grid point data is inserted with full weight for those grid points which are updated. The first local update occurs 2 hours after synoptic time, so it is possible that an orbit of satellite data can override the analysis in an area that has just been updated with many radiosonde reports.

In the GLAS method, the analysis of the difference between forecast and satellite data is simply added to the forecast value at those grid points affected by data. Future plans call for other data types to be weighted along with satellite observations in the time-continuous assimilation. In addition, radiosonde data will have an influence which decreases in time from full-weight to no-weight over a 6-hour time period.¹ A geostrophic correction is currently made to the wind whenever temperature data are inserted and whenever the surface pressure field is updated. Satellite observations and radiosondes are assigned relative weights based upon observational errors similar to those quoted by NMC during the session on observation error statistics.

Larsen asked if anyone had considered the feasibility of doing optimum interpolation analyses in time when using asynoptic data. For the FGGE period data will be available forward in time, so it would be possible to avoid shocking the model by pre-introducing data in areas which have been void for some time. Rutherford feels it may be dangerous to look at the time behavior of a model over ± 3 hours because the model may not be too well damped in such a short time period and the tendency may not be correct. The error made by assuming all data in a ± 3 -hour time window are synoptic is probably smaller than the tendency error in a 6-hour forecast.

¹This experiment has been performed in the interim between the Workshop and the publication of this document, and will be reported shortly.

g. Initialization

This topic was introduced by reading an excerpt by Temperton summarized as follows: If analyzed fields are used directly as initial conditions for a forecast model, the integration may be contaminated by spurious, high frequency gravity oscillations caused by imbalances between the variables. Imbalances may be removed by a variety of initialization procedures. There is disagreement among forecasters on the impact of initialization on the accuracy of subsequent forecasts. Constraints imposed during analysis and the link between analysis and forecast might ideally render initialization unnecessary.

Rutherford opened the discussion by stating that two types of initialization must be considered separately: that needed for an assimilation cycle, and that needed for an extended forecast run. In an assimilation cycle if there is no initialization, there will be gravity waves in the assimilation model. There are two ways of coping with such noise: (1) damp the gravity waves by dynamic initialization (backward and forward integration), or (2) apply some kind of diagnostic initialization procedure. No such diagnostic procedure available up to now has proven entirely successful operationally. However, it now looks like the nonlinear normal mode approach will provide a satisfactory procedure for removing all gravity wave oscillations. The technique allows selective removal. High frequency noise is easy to remove by dynamic initialization, but lower frequency modes (e.g., large-scale tropical modes) cannot be removed by such a technique. The main problem with the nonlinear normal mode technique is identifying the modes to be removed or retained. The feasibility of the normal mode technique has not been proven, but hopefully it will be soon.

McPherson asked if gravity wave noise is really a problem in assimilation models. Rutherford responded that it is in the CMC model, at least for low-frequency tropical modes. A large part of this noise problem is due to the boundary in the tropics.

McPherson then asked how expensive the nonlinear normal mode technique is. Rutherford replied that according to Daley's results, it is no more expensive than about three time steps of the model, but somewhat more expensive for multilevel models.

Ghil commented that it is not clear that one wants to remove all gravity wave modes since some are there in the atmosphere. Rutherford reiterated that the nonlinear normal mode technique allows one to be selective. Ghil remarked (without going into any detail) that there is another approach along the lines of generalizing the balance equation.¹

¹Ghil, et al. (1977). Mon. Wea. Rev., 105, 1223-1238.

Much discussion took place about what should be damped and what the best techniques are. Rutherford feels that, though all this discussion is interesting, in the final analysis initialization is not much of a problem in the assimilation cycle. Its importance depends to some extent on what kinds of motion are to be depicted in the assimilation, especially in the tropics. Miyakoda expressed the view that analysis and initialization should be two separate steps. NMC analyses are used by the meteorological community for research purposes and should be as high quality as possible. A separate initialization step should be used to adjust this high quality analysis to suit a particular forecast model, which may not be good enough to accept a high quality analysis as an initial state.

h. Treatment of moisture

GFDL performs a global analysis of dewpoint depression optimally. Using a model forecast as a background field is a problem since the forecast is very model-dependent (depends upon parameterization scheme, saturation criteria, etc.). Using climatology as a background field is impossible since data are too sparse. "Bulls-eyes" result. Therefore, GFDL decided to use the result of its 4-D analysis as data during the DST-5 experiment. The moisture analysis is done at all levels up to 300 mb. Water vapor data are not inserted during local updates.

The present NMC system (based on Hough functions) uses a forecast as first guess. The lowest two levels are especially model dependent. NMC is therefore moving away from giving the forecast much weight in the analysis. Some sort of climatology will have to be used in the low levels. Lorenc remarked that a better approach would be to try to correct the model. McPherson agreed, but feels such an approach is a long-range solution.

Rutherford commented that two problems arise in moisture analysis: (1) difficulty in representing small-scale details at the resolution used, and (2) the difficulty in drawing in the saturation or near-saturation areas. The variable CMC uses tends to magnify the errors in areas at or near saturation. Dewpoint depression is the variable analyzed. The available upper-air net is not sufficient to define the moisture field correctly.

Lorenc asked if anyone knows what sort of moisture data can be expected from satellites in the future. McPherson responded that VTPR contains one integrated channel of moisture, which is not very helpful. Tiros N will have more channels, but this data likewise will probably not be very useful. Lorenc remarked that it seems that satellites are our only hope. They might someday provide an integrated area value for moisture which is really what models require.

Bergman pointed out that there is another source of data, which is surface observed clouds from surface observations. NMC uses these in its regional model; also NESS transforms observations from satellite pictures into bogus moisture soundings.

Barker stated that vertical motion is of overriding importance. Even a very dry air mass, when lifted over a small distance, quickly becomes saturated. Incorrect initial imbalances can set off oscillations which can quickly ruin a good initial moisture field.

Rutherford has found that the importance of the moisture field depends a great deal on resolution of the model. With full physics, a 20-wave spectral model is insensitive to the initial moisture field, whereas a 29-wave model is quite sensitive to the moisture analysis.

Rutherford and Bergman mentioned that both CMC and NMC analyze moisture using optimum interpolation simply because it is the most convenient way, not because statistical methods have any particular advantage over other methods. One problem with using optimum interpolation is knowing what correlation function to use. For lack of anything else, NMC uses the same correlation function as for analyzing temperature. In the regional analysis, NMC is experimenting with horizontal elliptic-shaped functions oriented along the axis of the wind direction and vertical functions which drop off as a function of static stability. NMC has not evaluated this approach sufficiently to know if it is useful.

Ghil pointed out that moisture is much more model dependent than any other analyzed variable since treatment of moisture varies greatly from model to model, whereas other variables are treated pretty much the same since all models solve the same dynamical equations. Kistler remarked that the best moisture variable to analyze might also vary depending on the model. NMC has chosen to analyze specific humidity because that is the moisture variable forecast by the model.

McPherson stated that because of phase errors, moist areas in the NMC model tend to lag behind where they occur in the real atmosphere. Satellite moisture soundings prove useful in correcting such phase errors by indicating that the areas behind migratory cyclones are drier than the model shows. However, satellite moisture soundings appear to be almost useless in specifying the moisture field in advance of such systems.

Lorenc asked whether anyone feels it is even necessary to perform a moisture analysis when making 10-day forecasts. Rutherford expressed the view that moisture analysis is important in a 10-day forecast if the model has high resolution. With the 29-wave spectral model, CMC has found large forecast differences to result from a small difference in the initial moisture analysis by 48 hours. In a 10-day forecast, slight differences in the moisture field may result in major differences in the forecast long-wave pattern.

5. Characteristics of the Results
(Fourth Session, Tuesday afternoon, September 20, 1977)

a. NMC

The results from NMC were presented by McPherson. Comparisons were made for the current operational system (Hough analysis) and the experimental system based upon optimum interpolation (OI). Verification scores of root-mean-square (RMS) error were tabulated for 80 Northern Hemisphere (NH) rawinsonde stations for the DST-6 period 02-09 February 1976 for two assimilation cycles: one including Nimbus-6 and NH VTPR soundings (SAT), and one without (NO SAT). Both analyses and predictions valid at 00 GMT and 12 GMT were verified. In reality the analyses were zero hour forecasts, i.e., Hough analyses were interpolated from pressure into sigma and back to pressure coordinates. The predictions were 6-hour forecasts from initial conditions at 06 GMT or 12 GMT. The results presented were from the NO SAT mode.

On the average at 850 mb the Hough analyses fit the heights more closely, whereas there was little difference between the two systems for the 6-hour predictions. One interpretation of this may be to view the Hough system as having a more rapid error growth rate. The fits at 850 mb to temperatures and winds had insignificant differences.

The 300-mb results paralleled those at 850 mb. Again the fits to temperature, winds, and predicted heights were nearly equal, while the Hough system analysis fit the heights more closely.

A point was made regarding the relative stability of each cycle. Considerable difficulty had been encountered and overcome in implementing a 6-hour Hough system due to the accumulation of noise in the cycling process, which occurred in spite of the Hough system's nondivergent characteristic and global mass-momentum balance. The experimental system, on the other hand, retains a smoothed predicted divergence field and has a far weaker mass-momentum balance. The impact of these differences was examined by comparing the evolution of the global RMS divergence. The graph of the experimental system showed a sustained growth, then a leveling-off after 2 to 3 days. While this indicates that the system is stable, it does not answer the question of whether the level of divergence is harmful or not. In addition the graph had a sawtooth appearance, the consequence of filtering noise components at each update.

A second measure of the level of noise was presented, the global RMS pressure tendency. It was quite evident from this graph that the experimental system excited more noise following each update than the Hough system. In addition the lower values at 06 and 18 GMT compared to 00 and 12 GMT indicated the shock was partially dependent upon the number of upper air data analyzed.

The energetics of the two NO SAT systems were presented next. A diagram of eddy kinetic energy was shown, the value having been integrated from 20°N to 90°N and between 1000 and 200 mb. Especially noticeable was that the Hough system lowered the value at 06 and 18 GMT from the levels found at 00 and 12 GMT, giving a sawtooth appearance to the evolution. However, the experimental system showed none of this behavior; rather, it consistently had a higher level of eddy kinetic energy. The difference between the systems is attributed to what is viewed as excessive vertical interpolation in the Hough system, a feature which was deliberately avoided in the experimental system by performing a three-dimensional analysis directly in the model coordinates.

Next the 300-mb height and isotach analyses were displayed for each system. It was shown that the experimental system had higher wind speeds, particularly in jet streaks over the oceans. Over data-rich areas the two were essentially identical. A difference chart of the isotach analyses strikingly depicted the higher wind speeds of the OI system over the data-sparse regions. A diagram of the 300-mb height differences pointed out two 180m centers. The one located off the coast of Greenland indicated OI to be the deeper of the two systems. This is viewed as a reflection of a 944-mb cyclone at the surface in the experimental system which corresponded fairly well to ship observations. The second 180m center was found over the Himalayas, and is a consequence of the difference between the two systems in handling model terrain pressure. As was the case with the isotach differences, there was little apparent disagreement in the height fields over data-rich areas.

A point was made about the relative geostrophic balance in each system. It was noted that the Hough system has faster wind speeds in the troughs and slower wind speeds in ridges, indicating a more nearly geostrophic mass-motion balance.

Next came the verification of 6L PE predictions at 48 and 72 hours started from 00 GMT analyses on 3, 5, 7, and 9 February 1976. It was evident that the experimental system scored better for the first two forecasts, then worse and much worse in the latter two, respectively. It was speculated that this apparent deterioration of the quality of the predictions produced by the experimental system during the assimilation period may be attributed to: (1) an excessive accumulation of noise; (2) inadequate quality control of data, especially in data-sparse areas (i.e., rejection of "good" reports due to excessive prediction error); (3) poor tuning of the prediction error-observation error ratio in the experimental analysis.

An example of a particular result due to poor error checking was shown. Three observations were rejected at 00 GMT 07 February 1976 that were indicating that the experimental system had not deepened a cyclone in the North Atlantic sufficiently. The data were rejected because they failed a single check, based on the size of the difference between data and the prediction. On subsequent analyses, more and more data were

rejected in the vicinity of the deepening system. When the case was repeated with a horizontal consistency check instituted, the evolution of the deepening system proceeded in concert with observations in its vicinity.

As a final remark, it was noted that this was the first extended test of NMC's experimental system. The system is due for many refinements before it is implemented into the NMC operations.

b. GLAS (M. Ghil)

Results of analyses and predictions conducted at GLAS for a DST-6 period 1-21 February 1976 were presented. Individual experiments covering this period consisted of a time-continuous data assimilation cycle and 72-hour predictions made every other day from 00 GMT initial conditions involving both a NO SAT mode and a number of SAT modes.

A table of "percent impact" was shown summarizing the various experimental configurations. One of the quantities summarized was the SAT to NO SAT ratio of S1 scores tabulated for both 48 hours and 72 hours for all 11 forecasts and for two verification areas. The NMC analysis served as truth for the S1 scores. Another quantity was a ratio measuring the statistical significance level of the experimental result; in viewing the table, a ratio greater than 2 was considered significant.

The results from cycles based upon a successive correction method (SCM) analysis were shown to have little or no impact. However, when a statistical analysis method (SAM) was employed, the criterion for significance was exceeded; the corresponding S1 score improvements were of the order of 5%. Several SAM experiments were run, differing only in their handling of observational bias; none of the methods proved outstanding with respect to the others in this regard. It was noted that the correlation of satellite observations was accounted for only indirectly--through the use of calculated correlation functions.

Attention was focused on 2 days when large positive impacts occurred. As would be expected, large differences between the SAT and NO SAT initial conditions existed. When these large-scale numerical predictions were verified by an automated local precipitation forecasting method developed by R. Atlas at 128 United States cities, it became obvious that the positive impact was not evident at 24 hours. However, by 48 and 72 hours the influence of satellite data was advected over the United States and improved the SAT scores.

c. GFDL (K. Miyakoda)

Results from the DST-5 (August 1975) period were presented. First, a comparison between SAT and NO SAT modes was made for 850-mb temperature fields averaged over a 17-day period. The RMS difference map showed, as would be expected, the largest values occurring over the entire Southern Hemisphere and areas void of conventional data in the Northern Hemisphere. The maps of the individual means indicated that the SAT fields were more zonal in character. A similar comparison for 200 mb yielded the same results.

In order to explain the large differences over normally data-sparse areas, the NMC analyses for the same period were examined. However, NMC's NO SAT mode included VTPR data; GFDL's NO SAT did not. In spite of this inconsistency, both GFDL and NMC produced positive biases over the oceans west of North and South America. Overall, it was noted that NMC's difference maps tend to be negative while GFDL's were positive.

d. ECMWF (A. Lorenc)

The results of a test of the mass-motion coupling in the ECMWF analysis was presented. Four analyses were produced. The first included the proposed mass-motion balance. It displayed a reasonable representation of contours and wind vectors at 500 mb. In the second analysis, the mass and motion fields were decoupled. The resulting 500-mb chart depicted cross-contour flow in several areas, while intensifying a cut-off low over the Balkans and a trough over Greenland.

The third analysis was produced using only height and thickness data. The most striking feature of this analysis is the strong easterly wind and the phase error between the contours and wind vectors in the vicinity of the Balkan cutoff. The fourth analysis was produced with only wind data. It produced a gross loss of amplitude in the contour field.

(Rutherford then presented results of a study performed during his tenure at ECMWF)

Global fields of the analysis error were calculated by the method of optimum interpolation for various configurations of the observing system for FGGE. It was assumed that the analysis was done by updating a global model using the three-dimensional multivariate scheme of ECMWF, as described earlier. Prediction error coveriances were modelled by suitable functions of latitude and level. They were assumed to be independent of the amount of data available on previous updates, which is certainly incorrect, but experiments were made with two different prediction error models in order to assess the sensitivity of the results to this factor.

Full allowance was made for the different values of observation error of the various observing system components and in particular for

their spatial correlations. Several different forms of the correlation for satellite sounding errors were tested.

The main results of the calculations, which were presented in the form of zonal average profiles of the analysis error of wind and geopotential at several levels, were:

1. With the expected FGGE observing system, errors in the Southern Hemisphere will still be considerably larger than at corresponding latitudes in the Northern Hemisphere (more than twice as large, at about latitude 60 in the upper troposphere). During the special observing periods the system of drifting buoys was fairly effective in controlling the lower tropospheric error level but had little effect on the errors above 500 mb. A fleet of constant level balloons at 300 mb throughout the Southern Hemisphere (not to be available during FGGE) provided a significant reduction of the upper tropospheric error.

2. In the tropics, the expected special observing system additions should provide wind analyses with an error level not much higher than that for northern mid-latitudes.

3. The impact of satellite data in the Southern Hemisphere was substantial, about as great as that of all other sources of information. The wind data from satellite-tracked clouds (at two levels) provided about as much information as the satellite temperature soundings (with errors of 2-3°). Halving the temperature errors doubled the impact of the soundings. Broadening the horizontal correlation lowered the wind analysis error and increased the geopotential error but the effects were very small. However, there was a greater sensitivity to the vertical correlation of the sounding errors; it will be important to determine these accurately.

4. The results were quite sensitive to the assumed prediction error profile, except north of about 45°, where there are apparently sufficient observations to determine the analysis nearly independently of the quality of the forecast.

Acknowledgments. The NMC participants would like to express their appreciation to Ian Rutherford, Gorm Larsen, Andrew Lorenc, Kikuro Miyakoda, Michael Chil, Thomas Schlatter, and Edward Barker for attending the Workshop and contributing generously of their expertise to enliven the discussions and ensure its success. We acknowledge with gratitude the efforts of Mary Daigle and Al Snow in making the arrangements for the Workshop and for the typing and publication of these Minutes.