

CIRCULATING COPY
Sea Grant Depository

*Computer Simulation in the
Analysis of Surface Current
Drifter Data*

Logic of Computers Group
Technical Report No. 189

PHILIP CHARLES PILGRIM
Department of Computer and Communication Sciences

October 1975

Research Sponsored by:

National Science Foundation
Grant No. DCR71-01997
Washington, D. C.,

City of Chicago
Department of Development and Planning

and

Michigan Sea Grant Program
Grant No. 04-5-158-16
Technical Report No. 46, MICHU-SG-75-204





Computer Simulation in the Analysis of Surface Current Drifter Data

by

Philip Charles Pilgrim

Department of Computer and
Communication Sciences

The University of Michigan
Ann Arbor, Michigan 48104

October, 1975

Logic of Computers Group
Report Number 189

with assistance from:

National Science Foundation
Grant No. DCR71-01997
Washington, D.C.,

City of Chicago
Department of Development and
Planning,

and

Michigan Sea Grant Program
Grant No. 04-5-158-16
TR #46, MICHU-SG-75-204

ABSTRACT

COMPUTER SIMULATION IN THE ANALYSIS OF SURFACE CURRENT DRIFTER DATA

by

Philip Charles Pilgrim

Cochairmen: John H. Holland, Edward C. Monahan

Surface current drifters are nominally floating postcards used to study currents in a body of water. They are released at known points and times and carried by the moving water until they strand ashore, where they may be found by passersby. The finders indicate on each postcard the place and time of recovery and return the cards to the investigator. Drifters recovered in this manner provide a wealth of information about the currents transporting them, if a well-defined means of analysis is available.

Computer simulation provides a methodology for analyzing drifter data on two fronts: hypothesis testing and hypothesis generation. In hypothesis testing, various conjectured current patterns are used in a simulation of the advective and diffusive processes transporting the drifters, and the results of each are compared with the actual data in order to select the hypothesis most compatible with the data. In addition, locations in the investigated water body can be selected which will yield data optimally distinguishing among the available hypotheses, assuming that one is true.

Hypothesis generation starts without any current velocity information and several drifter recovery data and attempts to construct from

the data a current pattern which is compatible with them. A goal-directed simulation, which considers hydrodynamic constraints, lies at the core of this inference process, and current patterns are generated by an iterative scheme. Each pattern generated can then be tested by the hypothesis tester to compare its data compatibility with other hypotheses.

ACKNOWLEDGEMENTS

I would like to thank my cochairmen, Dr. Holland and Dr. Monahan, for the time, suggestions, patience, and encouragement they gave me during this investigation. I am additionally indebted to Dr. Monahan for giving me the idea for the project and the opportunity to participate in the field work. Were it not for the support of these two men, this dissertation would not have come about.

My other committee members, Dr. Burks and Dr. Meyer, have provided a valuable service in their careful evaluation of my work. To Dr. Burks I am particularly grateful for his helpful suggestions along the way and his guidance through the final stages of getting the thesis ready.

Captain Richard Thibault and the crew of the R/V Laurentian have made the field work a rewarding experience. The hundreds of people who returned the drift cards we released have shown the kind of grass-roots response that makes such a study possible. And the efforts of Jim Scherr and Steve Wright have transformed the raw data into machine-intelligible form. My thanks to all.

The Logic of Computers Group has provided me an atmosphere of freedom and stimulation in which to carry out my work. I am very thankful for this and for the financial support which both they and the City of Chicago have generously provided.

Finally, the rare combination of typing skills and technical expertise provided by Vivian Suits has made the preparation of this report a smoother operation than I imagined possible. In addition, the speedy typing of some of the final chapters by Mrs. Alice Gantt enabled me to meet an otherwise impossible deadline. Each was a pleasure to work with, and I am grateful for their help.

PREFACE

The work reported here results from an interest in computerized inference which began several years ago. It was inference, I believed--the ability to discover hidden structure, create usable hypotheses, and make reasonable decisions from incomplete information--which was the hallmark of intelligent behavior. If a computer could only be endowed with inferential capabilities, then the day when machines could presume to act intelligently and get by with it would be near at hand. Needless to say, I had underestimated the complexity of thought processes and was in for some hard, practical lessons.

Before I began my current work, I had made two other attempts at automating inference in widely distinct contexts. Neither was markedly successful. As an undergraduate in physics, I tried designing an algorithm to uncover hidden peaks in gamma ray spectra. From a wavy line, it was supposed to infer the energies of the gamma rays giving rise to the peaks therein and subtract that portion of the curve that each discrete energy was responsible for, revealing further peaks, and so forth. Unfortunately, it ended up creating a few peaks of its own and wasn't very useful.

It wasn't until more than a year later that computer inference caught my fancy again and I began my second attack on the problem. By then I had been exposed to John Holland's ideas about adaptive systems (Holland, 1975), which gave me a framework for investigating the inference process with more confidence. This time I chose a dynamic system as the focal point of my efforts and hoped to hit upon the process of abstraction which begins by recognizing temporal patterns in the

environment and then structures these patterns in a form which facilitates prediction. The problem domain was a maze containing a simulated rat whose task it was to discover the shortest path to the end. The rat was sighted enough to sense a wall ahead and on either side of it and mobile enough to turn or move forward. Its "brain" contained enough complexity to allow it to model its environment, and adaptation occurred in response to punishment every time the rat tried to run into a wall. This it eventually learned not to do; but without having it artificially supplied, it lacked the drive or goal-directedness required to seek the shortest path to the end.

By this time, I had gained a healthy respect for even the simplest sentient functions and began to doubt that any machine could be made to exhibit reasonably intelligent behavior. It seemed that the capabilities of mechanised systems had a distinct outer limit and that the driving force of cognitive process lay somewhere beyond. At any rate my enthusiasm for the rat project waned, and I was disinclined to consider any further attempts at machine intelligence.

In many ways the rat project was an attempt to solve an ill-posed, artificial problem. It lacked the concreteness necessary to determine when the goals of the research would be reached. Reasoning, therefore, that a well-specified problem should be the antecedent of my research efforts, I went looking for a problem.

My search brought me directly to the department of Atmospheric and Oceanic Science and to Edward Monahan, who described to me his experiments with drift cards for determining the currents in a body of water. He pointed out the need for a well-defined method of analyzing drifter data and conjectured that an inference process founded in the maximizing

or minimizing of certain hydrodynamic parameters would add rigor to this elegant, inexpensive method of investigation. The problem immediately fascinated me, but previous experience tempered my enthusiasm for yet another inference project. Its captivating nature got the best of me in the end, though, and I offer the results of my efforts here, but with a cautionary note.

The use of computers in analyzing experimental data has become a tacit expectation in almost all scientific (and some non-scientific) endeavors. Indeed, the age of "hands-free" experimentation has already reached some corners of the scientific community and is anxiously awaited in others. But there is a danger here that could stifle the rapid advances that computers are supposed to induce. By depriving the human senses of raw data, by performing abstractions from them by a pre-defined set of rules, the unwary scientist could be denying himself the opportunity of finding just those quirks--those oddball cases--which don't fit and which could, in that rare case in a thousand, lead to a quantum jump in knowledge.

There is a division of labor between man and machine implied here which I don't believe will ever be erased completely: it involves a well-trained human intuition interposing itself between direct experience and any subsequent abstraction, regardless of the abstractor's sophistication. Unfortunately the pressure of "getting on with it" sometimes affords scant opportunity to look at anything but intermediate and final results, and when these results are from a machine, precision can be easily mistaken for quality. But quality will be elusive unless the inferences drawn from the data are founded in a gut feeling for the data themselves and not just a good mathematical model of the system

that produced them.

It is therefore with more than a little trepidation that I endeavor to do oceanography (or computer science *in* oceanography). For while the research has taught me much, only extensive further experience in the field could grant me the eye to see as an oceanographer sees. So, where I've found it necessary to criticize the published works in oceanography, I hope I've not trespassed recklessly into their authors' intuitive domains. And where I've analyzed new data, I've tried to present the data themselves in a form which permits the practiced, human skills of the oceanographic readership to deal with them. Nevertheless, some of my methods and conclusions involve a liberal dose of my own intuition, both physical and mathematical, and I take full responsibility for whatever errors or misunderstanding might be present as a result.

Some have questioned what all this has to do with computer science in the first place, and I understand their doubts. But computer science, taken alone, is distinguished as an artificial science, and I can't help feeling that much of its *raison d'être* is lost if it remains remote from application in other fields. Moreover, they are often the problems from without a given discipline which pull at its boundaries and tend to expand its frontiers. So it is my hope that by bridging a gap between computer science and oceanography, I've somehow managed to give something to each.

Philip C. Pilgrim
Ann Arbor, Michigan
September 22, 1975

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
PREFACE	iii
LIST OF FIGURES	ix
NOTATION	xiii
CHAPTERS	
1. INTRODUCTION	1
1.1 Overview	10
1.2 Problem Definition	11
1.3 History	25
1.4 Approach	
2. A MODEL OF DRIFTER TRANSPORT	27
2.1 Overview	27
2.2 Prior Assumptions	31
2.3 The Drift Model	36
2.4 Discussion	
3. DRIFTER SIMULATION AND HYPOTHESIS TESTING	42
3.1 Overview	46
3.2 A Monte Carlo Drifter Simulation	55
3.3 Data Interpretation and Goodness-of-Fit	73
3.4 Hypothesis Testing Procedure	77
3.5 Experiments in Lake Michigan	100
3.6 Release Location Bias and Diagnostic Points	105
3.7 Testing Hypotheses about Wind-Driven Currents	117
3.8 Summary	
4. A GOAL DIRECTED DRIFTER SIMULATION AND HYPOTHESIS GENERATION	120
4.1 Overview	123
4.2 Prior Considerations	130
4.3 General Structure of the Inference Process	135
4.4 Algorithmic Structure of the Inference Process	139
4.5 Local Velocity Conditions	148
4.6 Hitting the Target--The Direct Approach	155
4.7 Hitting the Target--The Circumspect Approach	174
4.8 Hitting the Target--The Hybrid Approach	182
4.9 Adaptive Strategy	185
4.10 Velocities of Unvisited Squares	189
4.11 Putting It All Together	199
4.12 Conclusion	

5.	TRIALS OF THE INFERENCE SYSTEM	
5.1	Overview	202
5.2	Testing the Inference Process on Data from Known Fields	203
5.3	Inference Results from Lake Michigan	225
5.4	Summary	247
6.	CONCLUSION AND OUTLOOK	
6.1	Nature of the Results	248
6.2	Methodological Results	248
6.3	Results from Lake Michigan	251
6.4	Work To Be Done on the Analysis Technique	253
6.5	Work To Be Done in Lake Michigan	254
6.6	Future Directions	254
	BIBLIOGRAPHY	255
	SUPPLEMENT	

LIST OF FIGURES

<u>Figure</u>	<u>Page</u>
1.1 A current meter.	5
1.2 A typical drogue assembly.	6
1.3 Typical drifters.	8
1.4 The Eckman current as a function of the wind.	12
1.5 Brucks' (1971) method for determining current speed from drifter returns and dynamic heights measurements.	16
1.6 Current speed based on assumption of straight-line drifter trajectory.	17
1.7 Norcross' and Stanley's (1967) method for estimating current velocities at drifter release points.	18
1.8 Reverse trajectory construction.	20
1.9 An incorrect inference from reverse trajectory construction.	22
2.1 Progressive vector diagram for $\underline{u}'(t')$ and its discretization Δq_1 .	33
2.2 Propagation of drifters in field \underline{s} .	38
2.3 Turbulent evolution of a diffusing cloud as a function of relative diffusion and meandering.	40
3.1 Two time steps in the evolution of $\Psi(\underline{q}, t)$.	43
3.2 Discretization of southern Lake Michigan.	48
3.3 Monte Carlo drifter transitions.	49
3.4 Simulation of 25 drifters.	52
3.5 Simulation of 25 drifters.	53
3.6 Comparison of drifter distributions from two velocity fields, one being shifted slightly along the Y-axis.	57
3.7 Expected time of finding first of several drifters landing on day 1.	61

3.8	Composition of Ψ from local cluster densities.	66
3.9	Superpositions of k local densities centered on points chosen at random from the parent normal distribution.	70
3.10a	Recoveries from a single release.	74
3.10b	Earliest recoveries from each square.	74
3.11	Drifter release locations for 1974.	78
3.12a	Selected recoveries from July, 1974, drifter releases.	80-85
3.12b	Selected recoveries from August, 1974, drifter releases.	86-90
3.12c	Selected recoveries from October, 1974, drifter releases.	91-93
3.13	Hypothesized velocity fields.	94-95
3.14a	Hypothesis testing results for July, 1974.	97
3.14b	Hypothesis testing results for August, 1974.	98
3.14c	Hypothesis testing results for October, 1974.	99
3.15	Contours of equal relative distinguishability for release points to distinguish hypothesis b from e .	106
3.16	Wind observation stations (15 July - 31 October, 1974).	110
3.17	Sample hourly weather observations.	111
3.18	Wind tracks for three observation stations, 15 July - 31 October 1974.	112
3.19	Hypothesis testing results for wind hypotheses.	114
4.1	Trajectories of 10 drifters showing gradual diffusion, then rapid dispersion due to the diverging field.	124
4.2	Scheme of inference system showing separation of goal-directed and adaptive components.	127
4.3	Detailed scheme of inference system.	134
4.4	Local and preferred velocities as considerations in computing the drifter's velocity.	142
4.5	Preferred velocity weighting as a function of remaining time and k , the anti-diffusion coefficient.	146

4.6	Trajectories from the straight-line heuristic.	150
4.7	Final results from straight-line heuristic for various k.	151
4.8	The straight-line heuristic used on circulatory data.	153
4.9	The velocities used for computing the curl of a square.	159
4.10	Parameters used in computing the solenoidal field around a single current dipole.	159
4.11	Velocities in the field of a current dipole.	161
4.12	The contribution of the field velocity at a point to the circulation of a drifter.	164
4.13	Lake Michigan divided into boxes with a few reference points and reference velocities.	166
4.14	Arcs computed by the circulation heuristic in the field of a simple gyre.	169
4.15	Arcs computed by the circulation heuristic in the field of a double gyre.	170
4.16	Trajectories from circulation heuristic.	172-173
4.17	Correction of hybrid velocity for insufficient speed.	176
4.18	Factors pertaining to velocity correction for discrete transitions.	181
4.19	Computation of an undefined field velocity from the field at one release point.	186
5.1	Velocity field and recoveries used for example set.	204
5.2	Inference results for standard parameter values.	205-206
5.3	Inference results for less allowed acceleration.	208-209
5.4	Inference results for less allowed acceleration and higher anti-diffusion values than standard.	210-211
5.5	Inference results for standard parameter values and forward (only) moving drifters.	212-213
5.6	Velocity field and recoveries used for second example.	215

5.7	Inference results for standard parameter values.	216-218
5.8	Inference results for twice-standard allowed acceleration.	219-220
5.9	Velocity field and recoveries used for last example.	222
5.10	Inference results for 4-times standard anti-diffusion.	223-224
5.11	Inferred trajectories for July, 1974, drifter data.	227-234
5.12	Hypothesis generated by inference procedure from July, 1974, drifter data.	235-236
5.13	Hypothesis testing results for July, 1974, comparing generated hypothesis to those of Chapter 3.	238
5.14	Inferred trajectories for August, 1974, drifter data.	239-243
5.15	Hypothesis generated by inference procedure from Aug., 1974, drifter data.	244-245
5.16	Hypothesis testing results for August, 1974, comparing generated hypothesis to those of Chapter 3.	246

NOTATION

The mathematical notation used throughout this report is fairly straightforward, but a couple conventions require explanation. All underlined elements are two-dimensional vectors, vector operators, or functions, or tuples containing at least one two-dimensional vector. A vector is given componentwise as a double of scalars, such as

$$\underline{v} \equiv (v_x, v_y), \quad \text{or} \quad \underline{0} \equiv (0, 0)$$

The subscripts x and y almost always indicate components of such a vector, although exceptions and other subscripts occur, but in a self-explanatory way. Unit vectors are identified by carat marks:

$$\underline{\hat{x}} = (1, 0)$$

Finally, primes are never used to denote time derivatives; rather, the dot notation is sometimes employed:

$$\underline{\dot{q}} = \frac{dq}{dt} = \frac{dq_x}{dt} \underline{\hat{x}} + \frac{dq_y}{dt} \underline{\hat{y}} = \left(\frac{dq_x}{dt}, \frac{dq_y}{dt} \right).$$

CHAPTER 1
INTRODUCTION

1.1 Overview

Notes sealed in bottles and thrown to the waves have served as a means of communication or inquiry for hundreds of years. Since the first desperate attempt of a stranded sailor to be found or remembered or the idle curiosity of someone "just wanting to see where the thing will end up," drifting messages have washed ashore the world over, providing their finders with intriguing tales or earnest requests for aid (Lederer, 1970). As the need for better hydrographic charts showing regional currents became apparent, "bottle papers" found their way into serious oceanographic investigation of these currents. Over the years, the use of such drifters has become more sophisticated, and recently quantitative analysis has been brought to bear on the data they provide. This study considers computer simulation as a means for deriving information about currents from the drifters transported by them.

Currents in a body of water are the motion of the water over time. They are responsible for the distribution of the water's contaminants, be they biological nutrients, heat energy, or pollutants, and play a major role in determining the biological and physical properties of a given region. Knowing the current structure of a water body is fundamental to the prediction of these properties and may be aided by accurate measurement.

Currents are caused by a complex interaction of forces acting on the water. Frictional forces from the wind induce motion on the surface, which is propagated to the depths by internal friction. Inertial considerations lead to the Coriolis effect, the force of which is manifest

in the motion of the water in relation to the rotationally accelerating surface of the earth. Alternate heating and cooling yields density heterogeneity in the water, giving the force of gravity opportunity to set up convection currents. Knowledge of these forces and others, along with a theory of their implication on currents allows the prediction of currents. In developing such a theory, the measurement of currents is essential to the testing of predictions.

Ideally, a comprehensive current measuring experiment will yield, for each point in the water body of interest, for each moment in time, a vector giving the velocity of the water moving through the point. Values for this function or *velocity field* may be obtained by direct or by indirect means. Indirect methods rely on the detection of phenomena which correlate with water movement but are not themselves observations of motion. These techniques frequently involve the measurement of scalar quantities whose gradients align themselves perpendicular to streamlines in the water. A simple example is the aerial observation of suspended silt distributed in a lake or estuary. The lighter colored silt tends to be concentrated along those streamlines intersecting its source, giving a rough visual indication of the direction of flow. Remote sensing from satellites can be applied to a whole spectrum of such quantities, including temperature, to outline the flow patterns where heterogeneities exist. Unfortunately, this general method gives little information on the current speeds involved--only the directions.

Another indirect technique, known as the *dynamic heights* method, relies on the determination of pressures (Neumann and Pierson, 1966). In a body of water at static equilibrium there will be no horizontal pressure gradients. But when Coriolis forces are taken into account,

horizontal pressure gradient forces which counter them will also be present. A system in which these forces exactly balance is said to be in *geostrophic equilibrium*. Since Coriolis forces act at right angles to the current flow, so will the pressure gradients at equilibrium, and their determination will allow the currents to be calculated. Pressure gradients, however, cannot be measured directly but must be deduced from density measurements, which are a function of temperature and salinity. The pressures deduced, though, are relative to the water surface and not absolute, so assumptions must be made about the absolute pressure field at some level. In addition, the measurements should be as synoptic as possible, since stationarity in the density field can seldom be counted on, particularly where internal waves are a problem. When realistic assumptions can be made and measurements can be taken concurrently enough, the method is a valid one, but practical limitations frequently render it awkward.

A third indirect technique takes advantage of electrical potential differences induced by the conducting sea water moving through the earth's magnetic field (von Arx, 1962). The measuring device, a geomagnetic electrokinetograph (GEK), consists of two electrodes placed in the water a known distance apart. The voltage induced between the electrodes is proportional to the velocity of the water passing between them. While elegant in theory, the method suffers from drawbacks in practice. A GEK is difficult to maintain in calibration, and local anomalies in the concentration of dissolved substances yield varying electrical characteristics.

Direct current measuring relies on the observation of motion and is done by two basic methods: "Eulerian" and "Lagrangian." Eulerian

measurements are taken at a fixed reference point and directly yield the velocity of the water flowing past that point. The most frequently used Eulerian current meter consists of a movable vane which aligns itself with the direction of the current and a cupped rotor (Savonius rotor) which turns at a rate proportional to the speed (figure 1.1) (Neumann and Pierson, 1966). The vane and rotor are electrically coupled to either a direct reading meter or a recording device. The latter allows a submerged current meter to be left unattended for several months, and an array of such recording meters will yield good synoptic current data for long time periods. Where cost is no object this kind of Eulerian measurement is the best and most direct method available. Often, though, the expense of sampling a large area with current meters is prohibitive.

Lagrangian measurements are taken by following markers which are fixed in and carried by the moving water. The data they yield are positions as a function of time. As the intervals between observations become small, the instantaneous Eulerian velocities at the observed positions may be determined directly. The commonest example of a Lagrangian measuring device is the drogue (Monahan and Monahan, 1973). A drogue is anything providing a wide cross-section to the current so that it may be carried by the current. It is often tethered to a marker buoy by a fixed length of line (i.e. at a fixed depth) so that its progress may be followed either visually or by more sophisticated radio or radar techniques (see figure 1.2). Drogues can provide excellent current data along their trajectories if observed frequently, but they need to be released in large numbers to provide such information for a wide area. Since they require constant attention, a large scale drogue study can be prohibitively complicated or expensive.




Figure 1.1: A current meter.




Figure 1.2: A typical drogue assembly.

The second major Lagrangian current measuring technique and the one upon which this study focusses, involves the use of drifters. Abstractly, a drifter is a drogue which is released and forgotten. It bears an identifying serial number and a request to whomever may find it to inform the investigators when and where it was found. Each recovered drifter, therefore, provides the final position and elapsed time for a trajectory with known initial position. Physically, a drifter can be a note in a ballasted bottle or some other waterproof container, or simply printed on buoyant, waterproof paper (figure 1.3) (Monahan, *et al.*, 1974). Its cross-section to the current is often augmented by vanes, streamers, or submerged sails. In a typical experiment, hundreds of drifters are released in clusters of 5 to 100, concentrated in a small area or along a line, or scattered in a wide pattern. How many are eventually recovered depends on the proximity of major land masses and how often their beaches are visited, the latter being a function of the season and the weather. On Lake Michigan in the summer, for example, the recovery rate can surpass 30%; in the fall 2% has been observed (Monahan and Pilgrim, 1975). One major drawback of drifter studies is obvious: the uncertainty in the interval between the time a drifter beaches and when it is found. The statistics of this problem are dealt with in Chapter 3. Another problem lies in the necessity of assuming stationarity in the large scale aspects of the average current field over the period of study. The assumption is by and large valid so long as no major change in the current-causing forces (e.g. a shift in the prevailing winds) takes place while the drifters are afloat. Even in stationary situations, though, it might be objected that drifter data provide little information to infer anything about the current velocities yielding them.




Figure 1.3: Typical drifters.

For individual drifters, this is mostly true. But in experiments with many recoveries, certain hydrodynamic constraints concerning what constitutes a reasonable velocity field may be used to coordinate the implications of the data and thus to limit the range of possible interpretations. If these difficulties can be successfully dealt with, the use of drifters for measuring surface currents has one distinct advantage--cost. Due to the inexpense of the materials used in the drifters and the limited amount of ship or plane time needed to release them, a comprehensive drifter study may be done at a small fraction of the expense encountered with other methods.

In any investigation of currents, a mixture of the various measuring techniques will often be employed to take advantage of the positive characteristics of each. In cases where the requisite physical conditions are met, the indirect methods provide an easily obtained overall view of the circulation structure spanning a wide domain. Current meter methods are most appropriate when detailed information is desired regarding instantaneous current velocities as a function of depth or time. Drogue studies are often undertaken when the ultimate goal is to determine the effect of currents on the transport or dispersion of contaminants. Drifter experiments can be done with the same end in mind, as well as to provide an inferred overall circulation structure. It is usually useful to compare the results using one technique with those using another as a check against inaccuracy. Moreover, uncertainties arising in the results of one type of measurement may be most expediently narrowed by using another type (e.g. drifter results which suggest profitable locations for current-metering). It should therefore be apparent that among the various means for determining currents in a

body of water, no single one is best for all occasions, but that each may find profitable employment.


1.2 Problem Definition

The intent of this investigation is to provide an algorithmic approach to the analysis of surface drifter data, which will produce for the experimenter information about the water velocities giving rise to the data. The problem has two aspects. The first is hypothesis testing. Given several hypothesized velocity fields for the same body of water and a set of drifter data from the field, one would like to select from the available hypotheses the one which is most compatible with the data. In other words, a means must be devised for determining which hypothesis, if true, would yield data most like that obtained. Furthermore, it would be desirable to determine in advance which drifter release points could be counted on to yield recoveries most incisively distinguishing among the hypotheses, assuming one to be true.

The second aspect of the problem is hypothesis generation. Here, no prior hypotheses are explicitly given. The goal is to infer from the drifter data alone the velocity field which would most likely yield that data. Naturally, not just any "solution" is acceptable; it must also satisfy certain hydrodynamic constraints, such as minimum speeds, velocity shears, accelerations, and divergence. Given a sufficiently constrained solution space and enough experimental data, it ought to be possible to reconstruct the velocity field yielding that data with a reasonable degree of certainty. If algorithms for doing this and for testing hypotheses can be given, analyses of drifter data from many experiments may be performed uniformly and with a degree of confidence not afforded by less well-defined approaches.

1.3 History

Drifter studies reported in the literature concentrate primarily on hypothesis generation. Hill and Horwood (1974) are an exception, though. They want to find the effect of the wind on the surface velocities of the Irish Sea. According to a theory of Eckman (1905), these velocities will be proportional to the speed of the wind and at a fixed angle θ to the direction of the wind (see figure 1.4). Hill and Horwood attempt to discover the constant of proportionality K and the angle θ using drifters and a computer simulation. While their drifters were afloat, they measured the wind and the subsurface currents at several fixed locations, under the assumption that the drifters would be transported by a net current equal to the sum of the Eckman current and the average, underlying current. Using the measured current values and various assumed values for K and θ as parameters in a drift model, they simulate the advection and dispersion of the drifters and get "recoveries" with which to compare actual drifter returns. By measuring the space-time distances between the points of 50% recovery for each of the simulations and those gotten experimentally, they pick the K and θ values whose average simulation outcomes minimize this distance. As a result, they obtain K values of .01 to .038 and θ values of 5° to 25° to the right of the wind, the smaller K 's and larger θ 's going to drifters extending more deeply into the water. Hill and Horwood's work is an extension of a technique used earlier by Tomczak (1968). Assuming $\theta = 0$ and using horizontal drift envelopes, Tomczak arrives at a K value of .042. His analysis technique neglects any underlying currents except in nearshore areas and any dispersion of the drifters. By plotting trajectories for various K 's and known winds, he can calculate the point where a drifter "should" have



$$\underline{v} = K \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \underline{w}$$

Figure 1.4: The Eckman current as a function of the wind.

ended up after any time T . By counting the number of actual experimental returns after interval T within a fixed radius of the calculated point, he obtains a goodness measure for the assumed K . The best K is the one maximizing this goodness measure over several release points and the appropriate range of T 's. In both articles, the hypotheses being tested are time-varying current patterns given by various values for the Eckman wind parameters K and θ . The chosen hypotheses are the ones which best explain the experimental data *vis à vis* the drift models and goodness criteria used.

The mainstream of published drifter studies focusses on inference or hypothesis generation. The experiments reported are often done to augment or fill in an incomplete body of information about the currents investigated. The methods of analysis used fall into four major categories: (1) qualitative graphic techniques, (2) inference of average speeds given assumed current directions, (3) reverse trajectory construction methods, and (4) an iterative algorithmic technique involving wind correction.

The qualitative graphic techniques involve looking at a chart showing the return data and plotting trajectories from them by hand or mentally, being careful to avoid crossing trajectories or excessively bent or fast ones. Harrington (1895) uses this method in his now classical investigation of the circulation structure of Lake Michigan and the other Great Lakes. From a few drift bottle recoveries he concludes that a counterclockwise gyre appropriately characterizes the current pattern (see figure 3.13a, Chapter 3). Hachey (1935) uses the same approach in Hudson Bay. From a handful of returns he attempts to piece together a general circulation picture of this region. Monahan and Pilgrim (1975)

also take advantage of this approach to draw preliminary conclusions from two drifter studies in southern Lake Michigan. They induce the general sense of the circulation and major changes between the two periods of investigation. The major shortcoming of this method is obvious: it is imprecise, yielding very general, tentative results. Nonetheless it is useful for making quick preliminary conjectures when no other techniques are readily available.

The second category of analysis is useful when the pattern of circulation is known ahead of time, but when the speeds involved are unknown. Stander, *et al.* (1973) released several thousand polyethylene drift cards in the south Atlantic and Indian Oceans. By constructing trajectories determined by the major oceanic currents, they are able to ascribe speeds to these currents based on the drift times for recoveries received over a period of several years from as far away as New Zealand and North Carolina. These speeds are assigned on a regional basis under the assumption that they are constant throughout each known rotational and translational current. Although the trajectory construction technique is rather loosely defined, their results, in one case, correlate well with a theory of larval drift for a population of lobsters on an isolated seamount in the Atlantic. Nevertheless, prior knowledge of current directions is not in general available, particularly in near-shore waters. Moreover the assumption of constant speed over a wide region is a rather strong one, limiting the method to special situations.

Drifters are also used to obtain current speeds when the directions can be measured using some other technique. Brucks (1971) measured density variations in the Caribbean for use in a dynamic heights calculation of the currents. The resulting current directions are used to plot

trajectories from drifter release points to the recovery points. Given the length of the trajectory and the transit time, he is able, in the manner of Stander, *et al.*, to calculate the average speeds involved in the current patterns gotten from the dynamic heights method (see figure 1.5).

Numerous authors simply assume straight-line trajectories for the drifters and calculate the speed as release-to-recovery distance/elapsed time (see figure 1.6). Wyatt, *et al.* (1972) investigated the currents off the Oregon coast using drift bottles to detect seasonal changes in the current flow. They conclude that the current is northward in the winter and southward in the summer and estimate the speeds by the above ratio. Since the actual trajectories are surely longer than a straight line would be, these speeds are admittedly underestimated; and since the initial direction is probably not aimed toward the recoveries, these inferred directions are rather crudely depicted. Norcross and Stanley (1967) are a little more sophisticated in applying this method. In studying the shelf waters off the Chesapeake Bight, they released a large number of drift bottles as well as bottom drifters in a wide-ranging grid pattern. Assuming straight-line trajectories, they plot at each release point vectors representing the average velocity of each recovered drifter released there (see figure 1.7). Next, they pick the 15° interval at each release containing the most vectors. The velocity assigned to a given release point is given by the average direction of the vectors in the interval and their average speed. Since the release points form a grid, they thereby obtain a velocity field for the area spanned by the grid, although the correctness of their inference is subject to the same objections as Wyatt's.

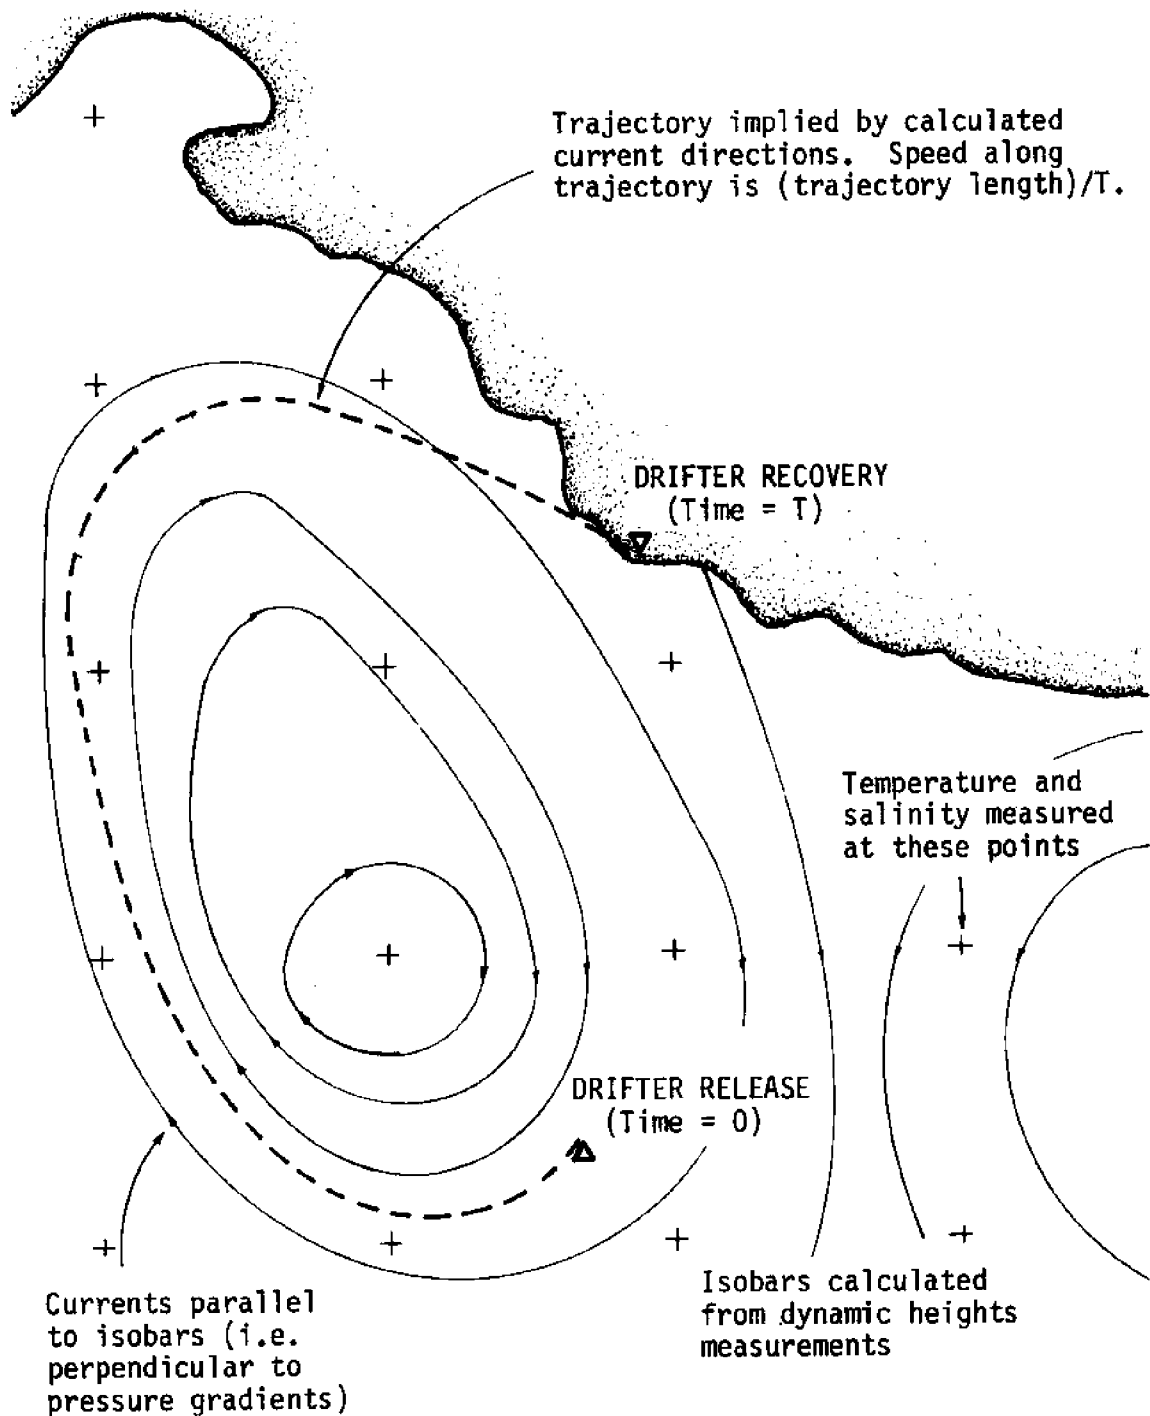


Figure 1.5: Brucks' (1971) method for determining current speed from drifter returns and dynamic heights measurements.




Figure 1.6: Current speed based on assumption of straight-line drifter trajectory.




Figure 1.7: Norcross' and Stanley's (1967) method for estimating current velocities at drifter release points.

As has been shown, current speeds may be estimated from drifter data if their directions are known *a priori* or as the result of another measurement technique, or if these directions are assumed on the basis of straight-line trajectories. The third inference technique seeks to determine the current directions and speeds simultaneously. This method, using reverse trajectory construction, is based on the assumption that the shorter the drift time for a drifter, the more likely its trajectory will have been a straight line. Given several recoveries in the same area from drifters released at different locations, one first picks the one with shortest drift time and traces a straight line to its release point. The drift velocity is the length of the line divided by drift time, as before. Then the drifter with the next shortest drift time is chosen, and its trajectory is traced back along the trajectory of the first drifter, then straight from the release point of the first drifter to its own release point (see figure 1.8). The speed along the first-plotted trajectory section is the same as for the first drifter. The speed along the second section is given by the length of the second section divided by the difference in drift times between the two drifters. The recipe continues for drifters with longer and longer times until all recoveries at the given location are accounted for. Bukin (1974) describes this method, using as an example the current patterns in Lake Issyk-Kul, USSR. Bumpus and Lauzier (1965) use a variant of this algorithm for reconstructing the current velocity field of the continental shelf between Newfoundland and Florida from drift bottle returns gotten between 1948 and 1962. Rather than dealing with only one recovery area at a time, they consider all recoveries in a given season starting with the one having shortest drift time and so forth. They

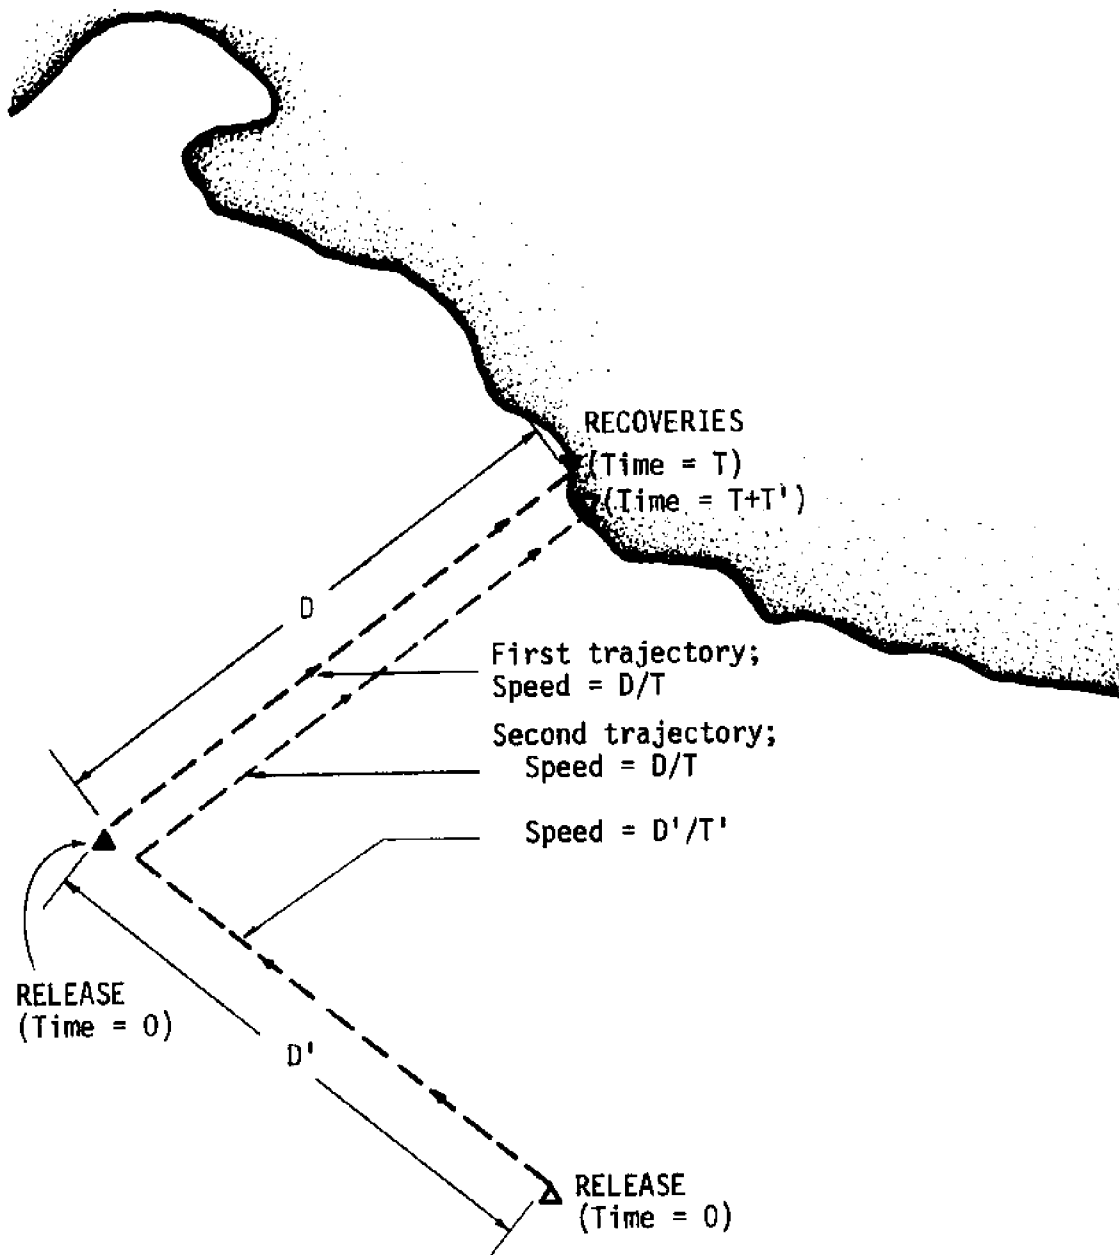


Figure 1.8: Reverse trajectory construction.

also pay attention to inferred drifts in areas neighboring the one being considered to achieve a sort of continuity in the overall velocity field. By this means they are able to outline the seasonal cycle of current patterns along the vast coastal area.

Bukin's method, if followed exactly, can encounter problems. Consider two drifters released from widely separated points and converging onto one recovery area (figure 1.9). Backward trajectory construction will take both drifters back to the release point of the one with shortest drift time and then the remaining one back from that release point to its own. The inferred trajectory section connecting the two release points could entail excessive speeds, depending on the difference in drift times between the two drifters. Bumpus and Lauzier's consideration of neighboring inferred velocities might well counteract this tendency, but the exact effect is not made explicit.

The fourth method found in the literature is an iterative, algorithmic one involving wind correction. It is almost the inverse of Hill and Horwood's hypothesis testing scheme in that it starts with assumptions about the wind effect and attempts to infer the average, underlying currents from drifter data. The method is reported by Pasquay and Bonnot (1971) and elaborated upon in personal correspondence from Mr. Pasquay. Pasquay and Bonnot released large numbers of plastic drift envelopes off the coast of France in and near the English Channel. During the period the cards were adrift, they obtained wind readings over the area, which could then be used in an Eckman wind drift model similar to Tomczak's. They assume *a priori* a K value of .035 and a θ of 0. The area of study is divided into discrete squares. With each square is associated a velocity representing the average, underlying

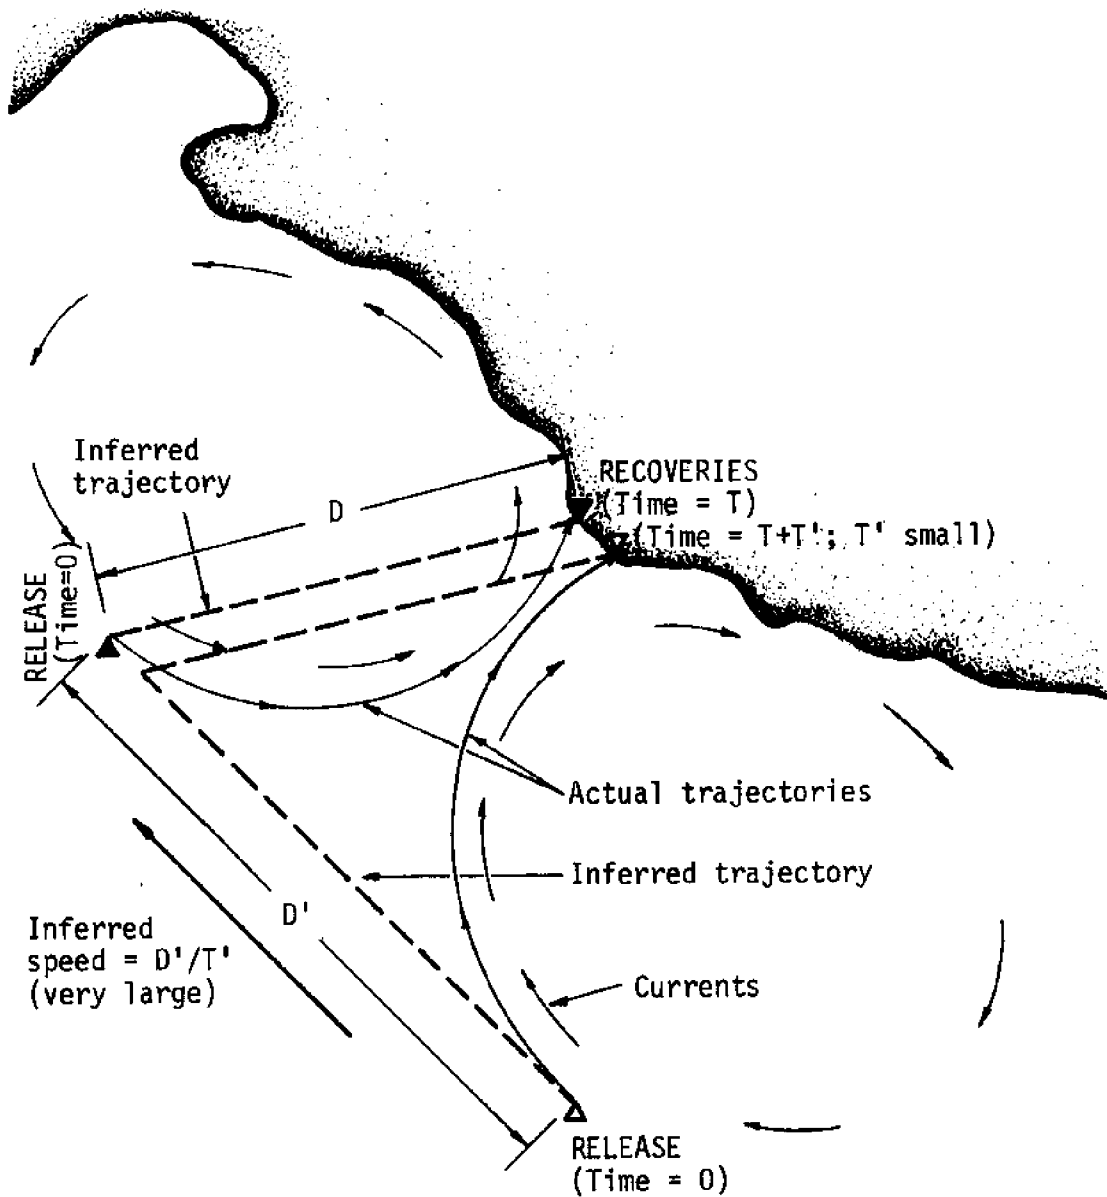


Figure 1.9: An incorrect inference from reverse trajectory construction.

current through the square. As a first approximation, these velocities are assumed to be zero: that is, the drift cards are presumed to be transported by the Eckman (wind-driven) currents only. Using the known wind velocities for each location over time in the Eckman scheme, they simulate the propagation of a single drifter from its known release point and for its known drift time. The end-point of the resulting trajectory is inevitably different from the actual recovery point. To correct this discrepancy, they adjust the velocities in those squares traversed in the simulation to carry the drifter to its proper recovery location. Next, a second drifter is released and propagated by the Eckman currents plus the local currents where they are defined, respecting a "continuity" in these currents and stopping after its appropriate drift time. As for the first drifter, the discrepancy between its stopping point and its known recovery point is corrected by adjusting the local currents in its path. Naturally, this adjustment may readjust those currents introduced in the first iteration. The third and successive drifters are handled the same way in subsequent iterations, and presumably, the process repeats for all the drifters until the local velocities converge. Pasquay and Bonnot derive, in this manner, a chart showing the "permanent" currents in the investigated region. One feature of their results is a rather strong current leading away from the southern coast of Brittany. As there are no large rivers emptying into the area, one is led to examine more closely those aspects of the algorithm leading to this result. The problem seems to lie in the nature of the corrective current applied to wind trajectories ending far inland of the proper recovery location. One must suppose that this current is in the direction of the vector pointing from the calculated endpoint to the

experimental recovery (i.e. offshore). From a physical standpoint a more circuitous corrective velocity field might be more desirable, since coastal currents running parallel to the shore could be derived. Although sources and sinks near the coast (e.g. coastal upwelling) and elsewhere are not impossible, a current inference algorithm may be considered more complete which deals with such divergence and convergence phenomena more explicitly.

In summary, the work reported in the literature includes a few studies using simulation for comparing hypotheses about the effect of the wind on surface current drift. The bulk of the work cited, however, is concerned with inferring the drift field directly from drifter data taken alone or augmented by auxiliary data and assumptions. The assumptions implicit in all the work are (1) that although time may lapse between a drifter's beaching and its being found, the resulting uncertainty is not enough to mislead the data analysis and (2) that any average, underlying currents are stationary during the period of study. In addition, several authors (Hill and Horwood, Tomczak, Pasquay and Bonnot) have presumed a wind effect which gives rise to a surface current with speed directly proportional to the instantaneous windspeed and at a fixed angle to the wind's direction. The constant of proportionality is found or assumed to be a few percent and seems to depend on the type of drifter used (i.e. the depth of the surface involved). Other authors (Stander *et al.*, Brucks) do not depend on wind measurements to determine the currents but assume that the current directions lie along the straight lines connecting release points and their corresponding recovery points and use the time interval data to calculate speeds, although the results are admittedly slow. The reverse trajectory

technique (Bukin, Bumpus and Lauzier) assumes straight lines for the shortest drifts, but by further assuming that all drifters arriving at a given point come from the same direction, it allows the construction of curved trajectories from short, straight segments. A couple papers (Bumpus and Lauzier, Pasquay and Bonnot) take continuity (lack of large velocity gradients) into account, but none explicitly mentions the problem of divergence (sources and sinks) in the current field. Only one paper among those cited (Hill and Horwood) directly considers dispersion of the drifters caused by turbulent diffusion. Even though acceleration, continuity, diffusion, divergence, and dispersion all play roles in determining the results of a drifter study, their relative importance in the total analysis is not often apparent.

1.4 Approach

The key advantages of studying a large-scale system by simulation are the necessity of making one's assumptions concrete enough to parameterize and the resulting ability to adjust the parameters to judge the relative importance of the corresponding assumptions. By applying this strategy to the problem of drifter analysis, the interplay among the prior assumptions made by several authors may be clarified with the result that analysis techniques which coordinate these assumptions may be developed. The plan of attack is twofold. The first phase begins by developing a model of the drifter transport process. In Chapter 2 the assumptions behind such a model are examined in detail and the resulting model is presented. In Chapter 3 a computer simulation based on the model is used for comparing hypotheses about current fields from experiments carried out in Lake Michigan. A method for choosing critical or diagnostic release points for maximally distinguishing between pairs of

hypothesized current fields is put forth. Finally, the simulation is modified to allow the testing of hypotheses involving wind-driven currents.

The second phase involves the inference of a velocity field, given only the data available from recovered drifters. Chapter 4 examines the assumptions necessary for the inference process to work and lays out a goal-directed model of drifter transport based on those assumptions. This model may begin with little or no knowledge about the current field being investigated; but by coordinating the various assumptions or constraints, propagates the drifters from their release points to the points of known recovery in the proper amount of time. By starting a simulation based on this model with no currents at all (*tabula rasa*) and running it until all drifters reach their destinations, first order hypotheses about the currents may be inferred from the resulting trajectories. By iteratively reinserting these hypotheses into the simulation and rerunning it with the same recovery data, a velocity field will begin to emerge which should accurately represent the one giving rise to the original data. That it does is verified in Chapter 5 by using data generated by simulation and comparing the result to the original current field. In addition, the results of using this technique on a drifter study of Lake Michigan are put forth. Chapter 6 outlines the general conclusions and suggests avenues for more research.

CHAPTER 2

A MODEL OF DRIFTER TRANSPORT

2.1 Overview

The key to understanding the implications of any experimental data is a model of the process by which the data were generated. The more explicitly defined the model is, the more clearly these implications are made manifest. Applying this axiom to drifter experiments, one is compelled to understand the forces contributing to the transport of the drifters from release to recovery well enough to allow accurate prediction of the recoveries, given complete information about those forces. The approach taken in this chapter is to model the transport process as a dynamic system in which the state at any time is given by the positions of the drifters, and successive states are dictated by a transition function specified by a parameterization of the forces acting to move the drifters. In order to do this, it is immediately necessary to delineate those physical assumptions which have a bearing on drifter transport, as well as those mathematical assumptions which simplify the model to the point of making predictions calculable. That these two kinds of assumptions conflict is expected, and one is hopefully left with a useful compromise between an accurate model and a manageable one.

2.2 Prior Assumptions

The assumptions made here are most closely aligned with the relevant properties of Lake Michigan, since that is the body of water investigated in subsequent chapters. Application of the results to the open ocean or to smaller lakes should be done with a careful eye on the

implications of violating any of the conditions presumed to prevail for the model presented. The drifters for which the model is designed are surface drifters (i.e. they float) in the "drum" configuration. The ones used in the field were printed on a buoyant plastic-fiber paper and stapled into a loop to form the drum shape. Since various kinds of drifters have different characteristics, those assumptions specific to the type used should be considered carefully when using another type.

The primary assumption is that the drifters are transported solely by water currents. That is, the drifters are embedded in the water and move in union with the water. This rules out any direct effect of the wind. The drum configuration used obeys this assumption nicely since virtually none of it extends above the surface to act as a sail. A further, perhaps obvious consequence of this assumption is that the drifters act independently of each other. In addition, since the drifters float, only currents in the surface layer need be considered. In other words, only those horizontal current velocities defined over the two-dimensional water surface will contribute to drifter motion.

Secondly, when a drifter's path intersects the shore, the drifter remains fixed until found. In some bodies of water, this is a difficult assumption to make, mainly because the definition of the shore is so poor, as in swampy areas. Here it would be possible for a drifter to get snagged by vegetation for a time and then released, thus complicating its trajectory. Luckily, Lake Michigan offers scant opportunity for this problem to arise.

In view of the first assumption, it is natural to establish an equation of motion for the drifters based on the water currents, *viz.*:

$$\frac{dq}{dt} = \underline{v}(q) + \underline{u}(q,t), \text{ where}$$

\underline{q} is the two-dimensional position of a drifter,

\underline{v} is the time-averaged current velocity field defined over the two-dimensional water surface, and

\underline{u} is the instantaneous deviation of the current velocities from the time-averaged field.

If \underline{v} and \underline{u} were known completely, then the equation could be integrated from any given initial position and time to obtain the position at any future time. Because of water's viscosity, \underline{v} is often smooth enough to know to an extent that would allow such predictions. Unfortunately, though, the fine structure of \underline{u} is usually much too complicated to even begin a deterministic solution of the equation. This is because small, random eddies, furnished through friction with energy from atmospheric disturbances and the larger-scale currents, tend to disturb any average or equilibrium current field. The result is that an initially small cluster of drifters placed into such a turbulent field will disperse. Diffusion phenomena like this are most conveniently treated probabilistically.

In the probabilistic formulation, \underline{u} is a stochastic process rather than a deterministic function. In order to simplify things a bit, \underline{u} can be assumed bounded, componentwise independent (separable), and stationary in time and space. The stationarity condition means that no parameters which govern the behavior of \underline{u} can explicitly depend on position or time. Since \underline{v} is completely time-independent, the overall field $\underline{v}(\underline{q}) + \underline{u}(\underline{q}, t)$ will have no explicit, deterministic dependence on time. Consequently, the Eckman wind currents discussed in Chapter 1 cannot be conveniently handled without slight modification to the subsequent model, because they do depend on both time and position in a

deterministic sense. For the most part, it will be assumed that the average current, \underline{v} , predominates over the Eckman currents in the propagation of drifters in Lake Michigan, although several hypotheses based on the Eckman currents will be tested in the next chapter. Fortunately there are no other notable time-varying effects on the surface currents of Lake Michigan, including tidal changes (FWPCA, 1967).

Even after the stationarity assumption, the properties of a realistic turbulence process, \underline{u} , are difficult to come by. Indeed, Csanady (1973, p. 46) credits turbulent motion with being "...one of the most untractable problems of the physical sciences, a full understanding of which is not in sight yet...." Therefore, since any sophisticated treatment of the problem would carry this study far afield of its intended purpose, another simplification is in order: it is assumed that \underline{u} has the memoryless property. This is to say that the turbulent velocity at an instant in space-time is virtually independent of that velocity at a neighboring instant. This assumption implies that turbulent motion is non-inertial. While such an implication is contrary to fact, it allows the propagation of a drifter to be treated much as a random walk--a process for which extensive results have been obtained. That the results under this assumption are, nevertheless, in fair accord with observation will be pointed out later.

Summing up, the propagation of a drifter is determined by the velocity of the two-dimensional surface layer of the water. This velocity field is given by a two-dimensional vector $\underline{v}(\underline{q})$, which is a smooth function of position on the surface, plus a random velocity $\underline{u}(\underline{q}, t)$, which is a bounded stochastic process observing independence between components, stationarity in space-time, and memorylessness. The

shoreline is considered an absorbing barrier: a drifter trajectory determined by the above velocities will terminate upon reaching the shore.

2.3 The Drift Model

Based on the assumptions given, a dynamic state-transition model for drifter motion is derived here. The state of a given drifter is taken to be its coordinates in the two-dimensional space encompassing the water surface and surrounding land masses. State transitions occur over finite intervals of time and are determined by the equation of motion, along with the assumed properties of \underline{v} and \underline{u} , as well as the shoreline conditions.

Since \underline{v} is a smooth function of position, it is reasonable to pick an interval of time Δt for the transitions such that $\underline{v}(\underline{q}) \approx \underline{v}(\underline{q} + \underline{v}(\underline{q})\Delta t)$ for all values of \underline{q} . In other words, Δt is small enough to assume that the contribution of \underline{v} to a one-step position change can be accurately estimated by $\underline{v}(\underline{q})\Delta t$. The interval Δt must be considerably longer, though, than the period of time over which \underline{u} "forgets" its initial value. The memoryless property of \underline{u} guarantees that this period is extremely small.

Given an appropriate Δt , one may begin, similar to the manner of Csanady (1973, p. 31), to find the net displacement of a drifter subject only to the turbulent velocity field \underline{u} . First, Δt is divided into a large number n of smaller increments $\Delta t'$ such that $\Delta t = n\Delta t'$. These smaller increments will still be much larger than the "forgetting period" of \underline{u} . The displacement in the i th increment $\Delta t'$ of Δt is

$$\Delta \underline{q}_i = \int_{(i-1)\Delta t'}^{i\Delta t'} \underline{u}'(t') dt', \text{ where}$$

$\underline{u}'(t')$ is the turbulent velocity of the *drifter* at time t' and is equal to $\underline{u}(\underline{q}, t')$, the turbulent velocity of the *water* at \underline{q} , the location of the drifter, and at the same instant. It is only because the statistical properties of \underline{u} are independent of position that \underline{u}' can be written without explicit dependence on \underline{q} . What has been done is to discretize the stochastic process \underline{u} into finite steps (see figure 2.1) which, on the surface, look like steps in a random walk, although they are not of equal lengths. The net displacement in time Δt is simply the vector sum of the n random steps:

$$\Delta \underline{q} = \sum_{i=1}^n \Delta \underline{q}_i$$

The goal is to determine the statistical properties of $\Delta \underline{q}$.

Since $\Delta \underline{q}$ is a sum of a large number n of random vectors, it would be nice to use the central limit theorem to derive a normal distribution for it. To simplify things a bit, $\Delta \underline{q}$ can be treated component-wise. This is because of the separability condition on \underline{u} (and hence on \underline{u}') which ultimately renders the component processes Δx_i and Δy_i of $\Delta \underline{q}_i$ independent of each other. Because of the memoryless property, the Δx_i 's are mutually independent and, by the additional virtue of stationarity, identically distributed. Since \underline{u} (and thus \underline{u}') is bounded, the Δx_i 's are bounded. Consequently their mean and variance both exist. Moreover, since \underline{u} is the deviation from a time-average field, the mean of the displacements Δx_i must be zero. These implications on Δx_i are sufficient to bring the central limit theorem into play (Feller (1967), p. 244), and one can write immediately for the distribution of Δx (the x -component of $\Delta \underline{q}$),

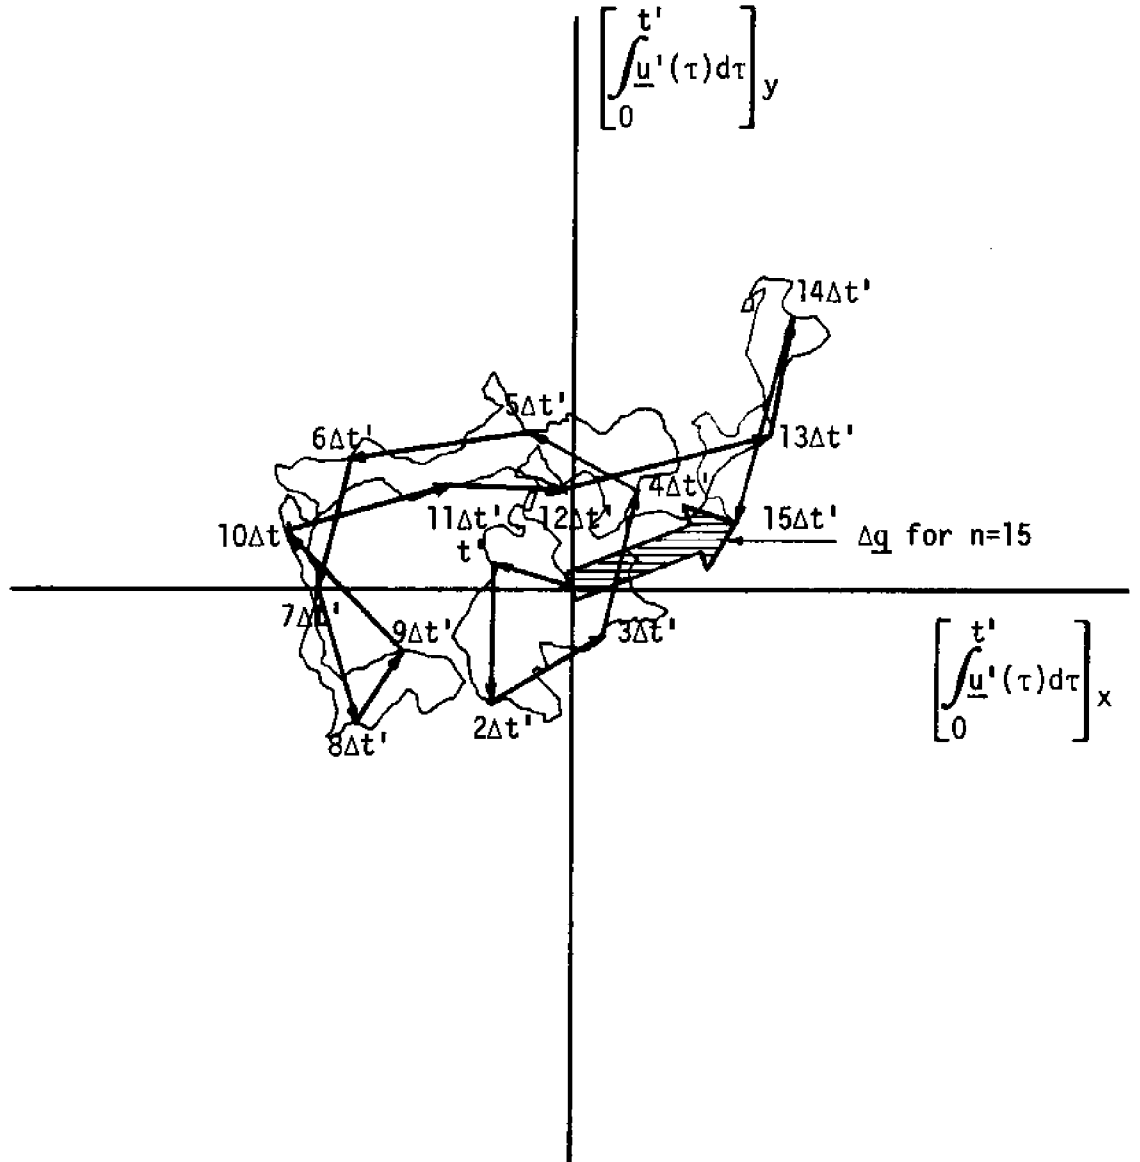


Figure 2.1: Progressive vector diagram for $\underline{u}'(t')$ and its discretization Δq_1 .

$\text{Pr}[\Delta x < \sigma\sqrt{n} r] \rightarrow F_N(r)$, where

σ is the standard deviation of Δx_i , and

$F_N(x)$ is the distribution function of a standard normal random variable.

This means that Δx approximates a normal random variable with mean 0 and variance $n\sigma^2$. By letting $\Delta t'$ be the unit of time (i.e. $\Delta t' = 1$), one gets $n = \Delta t$. Now, σ^2 is a property of Δx_i and ultimately of \underline{u} . It can be thought of as the rate of spreading that an initially small cluster of drifters subjected to \underline{u} can expect to undergo and is often given in terms of a "dispersion coefficient" D as $\sigma^2 = 2D$. Rewriting in new terms, the displacement in one direction Δx of a drifter subjected to the turbulent field \underline{u} for length of time Δt approximates a normal random variable with mean 0 and variance $2D\Delta t$. Therefore, the vector displacement $\Delta \underline{q}$, composed of the independent random variables Δx and Δy , is a random vector having (in the limit) a bivariate normal distribution with mean $\underline{0}$ and covariance matrix

$$\begin{bmatrix} 2D\Delta t & 0 \\ 0 & 2D\Delta t \end{bmatrix} .$$

Coalescing the results obtained so far, the drift model is to be specified as a dynamic system whose state is given by \underline{q} , the position of a drifter on the two-dimensional surface. A discrete state transition occurs in time Δt and represents a displacement $\Delta \underline{q}$ of the drifter from its current position to a new one. This displacement results from the superposition of an advective (mean) displacement given by the velocity field $\underline{v}(\underline{q})$ and a turbulent one completely determined by the dispersion coefficient D as follows:

$$\Delta \underline{q} = \underline{v}(\underline{q})\Delta t + \underline{N}(\underline{0}, 2D\Delta t), \text{ where}$$

$\underline{N}(\underline{\mu}, \sigma^2)$ is a component-wise independent bivariate normal random variable with mean $\underline{\mu}$ and with variances σ^2 in both components.

In order to tell when a shoreline has been crossed, a geographical descriptor function which partitions the state space into land and water areas is necessary. Given this, the model M is defined as the following probabilistic system:

$$M = (\underline{Q}, \underline{v}, D, G, \underline{\delta}), \text{ where}$$

$\underline{Q} = \mathbb{R}^2$ is the state space (location space for the drifter),

$\underline{v}: \underline{Q} \rightarrow \mathbb{R}^2$ is the time-average velocity field over the surface

$D \in [0, \infty)$ is the dispersion coefficient,

$G: \underline{Q} \rightarrow \{\text{land, water}\}$ is the geographical descriptor function, and

$\underline{\delta}: \underline{Q} \rightarrow \underline{Q}$, the state transition function, defines the dynamics of

the system as follows:

If, for any $\underline{q} \in \underline{Q}$, $G(\underline{q}) = \text{land}$, then

$\underline{\delta}(\underline{q}) = \underline{q}$; otherwise

$\underline{\delta}(\underline{q}) = \underline{q} + \underline{v}(\underline{q})\Delta t + \underline{N}(\underline{0}, 2D\Delta t)$, where

Δt is the time interval associated with one discrete transition and

$\underline{N}(\underline{\mu}, \sigma^2)$ is a bivariate normal random variable as defined above.

Since drifters are independent of each other, a system containing multiple drifters may be modelled as the parallel composition (independent, synchronous combination) of several identical single-drifter models.

2.4 Discussion

Certain aspects of this model's behavior can be derived directly from the transition function. To do this, it is simplest to consider a uniform velocity field \underline{v} such that for all \underline{q} , $\underline{v}(\underline{q}) = \underline{s}$. Furthermore, no land areas will be assumed present. The behavior of interest is the position of a drifter after multiple transitions, starting from a point \underline{q}_0 . After the first transition, the drifter will be at $\underline{q}_1 = \underline{q}_0 + \underline{s}\Delta t + \underline{N}(\underline{0}, 2D\Delta t)$. Applying $\underline{\delta}$ to the new position:

$$\underline{q}_2 = \underline{\delta}(\underline{q}_0 + \underline{s}\Delta t + \underline{N}(\underline{0}, 2D\Delta t)) = \underline{q}_0 + 2\underline{s}\Delta t + \underline{N}(\underline{0}, 2D\Delta t) + \underline{N}(\underline{0}, 2D\Delta t).$$

Since the two normal random vectors are independent, they may be combined, taking advantage of a theorem from probability theory stating that the random variable resulting from adding two independent normal random variables is normal, with mean and variance equal to the corresponding sums of the means and variances of the original variables.

Therefore,

$$\underline{q}_2 = \underline{q}_0 + 2\underline{s}\Delta t + \underline{N}(\underline{0}, 4D\Delta t).$$

This can clearly be extended to n transitions, yielding

$$\underline{q}_n = \underline{\delta}^n(\underline{q}_0) = \underline{q}_0 + \underline{sn}\Delta t + \underline{N}(\underline{0}, 2Dn\Delta t),$$

or, letting $t = n\Delta t$ and combining the advective and diffusive terms,

$$\underline{q}(t) = \underline{N}(\underline{q}_0 + \underline{st}, 2Dt).$$

This result exactly parallels the solution to the molecular diffusion (Fokker-Planck) equation:

$$\frac{\partial \Psi(\underline{q}, t)}{\partial t} + \underline{s} \cdot \underline{\nabla} \Psi(\underline{q}, t) = D \underline{\nabla}^2 \Psi(\underline{q}, t),$$

with the solution, starting from an impulse at $(q_0, 0)$, being

$$\Psi(q,t) = \frac{1}{4\pi Dt} \exp\left[-\frac{|q - (q_0 + \underline{s}t)|^2}{4Dt}\right].$$

$\Psi(q,t)$ is the relative "concentration" of diffused substance at position q and time t or, for a single diffusing particle, the probability density of finding it at q at time t . The formula for Ψ is the density function of the bivariate normal random variable $N(q_0 + \underline{s}t, 2Dt)$, the result obtained for $q(t)$ above.

In physical terms, this result implies that an initially small cluster of drifters placed into a constant velocity field \underline{s} will have its center of gravity displaced $\underline{s}t$ in time t . The shape of the cluster as it spreads will be roughly circular with its standard radius (equal to the r.m.s. deviation of the drifters from their center of gravity) expanding proportionally to \sqrt{t} (see figure 2.2). Bearing this in mind, it is possible to compare the model's behavior to the behavior predicted by a more sophisticated treatment of the turbulent processes involved.

Csanady (1973) describes the sequence of events in the turbulent diffusion of an impulse of contaminant as a function of two main factors: $s(t)$, the standard deviation (due to relative diffusion) of a given contaminant distribution about its center of gravity; and $m(t)$, the ensemble standard deviation (due to meandering) of the center of gravity. The overall standard deviation $\sigma(t)$ is related to the other two as $\sigma^2(t) = s^2(t) + m^2(t)$. One may imagine the diffusion process he describes as consisting of three epochs, beginning with the point release of the contaminant. In the first epoch, the entire "cloud" of contaminant is translated together as a parcel by the mean velocity and by the turbulent eddies and grows rather slowly. The "diffusion" is thus




Figure 2.2: Propagation of drifters in field \underline{s} . Circles are of standard radius for the given times.

dominated by the meandering (ensemble) factor $m(t)$ which, Csanady shows, grows linearly with time (figure 2.3). This continues until the slowly enlarging cloud reaches a size comparable in scale to the turbulent eddies, at which point, the second epoch begins. In this "explosive" phase, the eddies have a handle on the cloud itself and cause a rapid expansion of the cloud about its center of gravity. During this period $s(t)$ increases linearly with time and becomes the dominant factor in the overall diffusion. As the cloud grows larger with respect to the turbulent eddies, these eddies have less effect in translating the entire cloud than before, so $m(t)$ begins to taper off. In the last epoch, the cloud has become larger than the typical turbulent eddy, rendering the eddies less effective in dispersing it than before. At this point $s(t)$ takes on the proportionality to \sqrt{t} characteristic of molecular diffusion. $m(t)$ levels out to a constant value, revealing the diminished effect of meandering in furthering diffusion. As a consequence, the overall standard deviation $\sigma(t)$ behaves like $s(t)$ for large t . Csanady points out that experiments in Lake Huron tend to confirm both the growth rates predicted and a Gaussian (normal) shape for the "average" cloud, though the latter conclusion is tentative for certain epochal segments.

In the context of drifters, the "cloud" of contaminant is the cluster of drifters. The "meandering" component of diffusion would correspond to behavioral differences detectable in a large set of drifter releases from the same point at different times. Since it is this ensemble behavior which is to be modeled, direct comparison between the drifter model and Csanady's results is meaningful. As is apparent, the drifter model behaves from time zero as the third epoch of turbulent diffusion outlined above. This agreement comes despite the fact that




Figure 2.3: Turbulent evolution of a diffusing cloud as a function of relative diffusion and meandering. (Adapted from Csanady, 1973, p. 96.)

the drifter model was derived from an assumption of a memoryless field of turbulence, which seems to contradict the observed persistence in typical random eddies. In reality, though, the difference is only one of scale and not of substance, when long drift times are being considered. This is to say that the persistence of an eddy is an insignificant consideration in the long run. What is meant by "the long run" remains to be settled. Experiment has shown (Csanady, 1973) that the third epoch is often underway in a typical ocean or lake by the time the cloud has traveled one kilometer from its source. This is not far enough for the mean field $\underline{v}(q)$ to change much and certainly not far enough from the typical release to reach shore. Therefore, the behavior of an ensemble of drifters is dominated by third epoch statistics; so the drift model, as presented, will be considered sufficient to meet the needs of this investigation.

In conclusion, the drift model presented is a product of compromise. Although the physical assumptions necessary to define the dynamics of the model are simple enough, getting a grip on the mathematics of turbulent currents has required several gross simplifications. The grossest of these is the assumption that turbulent diffusion behaves like molecular diffusion, a process resulting from uncorrelated random motion. Nonetheless, for the time and distance scales involved in a drifter experiment, the resulting model is good enough to make predictions about the ensemble behavior of drifter clusters. What remains is to design a simulation based on the model and to use its predictions for comparing hypotheses about the mean current field \underline{v} .


CHAPTER 3

DRIFTER SIMULATION AND HYPOTHESIS TESTING


3.1 Overview

Two methods are possible for simulating drifter transport based on the model presented. Both have the same end: the determination of $\Psi(\underline{q}, t)$, the drifter "concentration" or probability density function, given an initial distribution (usually an impulse at the release point) subject to the known velocity field $\underline{v}(\underline{q})$. In the first technique, Ψ is calculated directly from the statistics implicit in the model; in the second, Ψ is estimated from the positions of individual drifters propagated Monte Carlo fashion by the model's transition function (see figure 3.1). The advantage of the first technique, assuming it can be carried out, is that Ψ is explicitly expressed from moment to moment: no interpolations need be made. Indeed, this would obviously be the best method if $\underline{v}(\underline{q})$ were constant, as assumed in the discussion of the previous chapter, because Ψ would always have the shape of a bivariate normal density--an easily represented function. Since no interesting \underline{v} has the constant property, the direct computation of Ψ would have to rely on simplifying assumptions. If this proves awkward or unreliable, the Monte Carlo technique, with its advantage of direct correspondence to the model, can be considered a reasonable alternative.


To illustrate the direct calculation of Ψ , the work of Hill and Horwood (1974) is again considered. They assume the same diffusion properties assumed here in the previous chapter and make the same observation that Ψ retains its bivariate normal shape in a constant velocity field.



Equidensity contours for directly calculated (true) ψ .



Monte Carlo drifter positions (samples from ψ).



Equidensity contours for ψ estimated from drifter positions.

Figure 3.1: Two time steps in the evolution of $\Psi(\mathbf{q}, t)$.

They extend this result, though, to cover non-uniform velocity fields. By propagating the center of gravity of the drifter distribution (mean of Ψ) using the advective velocity field (in their case, time-varying), but retaining the characteristic normal shape with its variance increasing linearly with time, they directly obtain a Ψ which is easily calculable. Despite this advantage of quick calculation, though, the implications derived thereby can be greatly misleading. In velocity fields characterized by even moderate inhomogeneities, the normal shape or any other easily represented shape of a drifter cluster can be quickly distorted, as is demonstrated by simulation later in this chapter (figure 3.4 - 3.5).

Because of the problems inherent in a simple, direct parameterization of Ψ , a Monte Carlo approach has been chosen for simulating drifter transport. An advantage of this, of course, is the direct correspondence between the simulation and the model. The model deals with individual, independent drifters, and so does the Monte Carlo simulation. As a consequence, no assumptions need be made about the shape of the overall distribution Ψ : it comes as a result of simulation. Unfortunately it comes at a higher price than Hill and Horwood's estimate. First of all, one must simulate a large number of drifters to get an accurate clue as to what Ψ looks like. The pain of such a price, however, is easily mitigated by the increased reliability. But even then, if one wanted an explicit evaluation of the implied Ψ , he would need a well defined means of translating individual drifters into a more global density map. Although that final step is never literally taken in this investigation, it should be borne in mind that any conclusions drawn from simulated drifters are actually grounded in the density function Ψ

from which the drifters are samples.

Having designed and programmed a drifter simulation, one may then put it to use evaluating hypotheses about $\underline{v}(q)$, the mean velocity field, in the context of experimental return data. The method is simple. Given a \underline{v} to test and an experimental release point, a number of simulated drifters are initialized at the same release point and propagated by the program. A goodness-of-fit criterion which is a function of the simulated drifter positions, the experimental recovery positions, and time is used to obtain a rating of the hypothesis \underline{v} for the given release point. An overall rating for \underline{v} is gotten by combining the individual release point ratings. The number thus obtained may be used as a basis for comparison among all the various \underline{v} 's tested in that way against the data. The one hypothesis yielding the highest rating is then chosen as the one most commensurate with experiment.

In comparing ratings from different release points across several hypotheses, one may notice that some releases indicate larger differences in the goodness-of-fit ratings than others. This observation may be put to good use in designing a drifter experiment when the object is to decide from several given hypotheses which is best. If one can determine ahead of time by simulation which releases will yield these large differences, he may concentrate his experimentation around these "diagnostic points," thus yielding a reduction in the overall experimental effort.

Finally, by making several modifications to the simulation algorithm, one may use the testing scheme to test hypotheses about wind-induced drift in the context of Eckman's (1905) theory (see section 1.3).

3.2 A Monte Carlo Drifter Simulation

To produce a computer simulation from a model, the parameters of the model must be embodied as numbers in the computer, and the transition function must be translated into an algorithm for manipulating these numbers. In terms of the drifter model, this means that the drifter's position q , the velocity field $v(q)$, the dispersion coefficient D , and the geographical descriptor function G have to be given a numerical representation. Finally, the transition function δ must be converted to a program which operates on these quantities to calculate the new drifter position q after any given time step.

First of all, the drifter's domain or state space (Q) must be settled upon. As in the model, this is the set of points in the Euclidean plane and is here represented as the set of pairs (x,y) of real numbers. Since only a finite area is of interest though, everything outside a closed square will be considered "out-of-bounds." For closed bodies of water this presents no problems, since all out-of-bounds regions would correspond to land far inshore and would not be visited by any drifter. With open bodies of water, a drifter could drift out of bounds. In this circumstance, its true position will be maintained, but its behavior will be determined by the characteristics of the nearest in-bounds point. Hopefully, this will be an unusual occurrence. It should be noted that large bodies of water require the use of a non-Euclidean coordinate system due to the earth's spherical shape. The portion of Lake Michigan investigated here is not a large body of water, so plane geometry is adequate.

The drifter domain described is now minced into a finite number of small squares. This is done to accommodate a finite representation for

both \underline{v} and G . It is assumed that within a given square both of these functions are constant, which places a limitation on the fine structures of both the velocity field and the shoreline. In this investigation, the inbounds region consists of 4096 squares arranged 64×64 . Each one corresponds to a square of Lake Michigan 2.75 km on a side. Each of these squares has a value of the binary function $G: \underline{q} \rightarrow \{\text{land, water}\}$ assigned to it, as well as a vector (v_x, v_y) representing the velocity $\underline{v}(\underline{q})$ for all \underline{q} lying within the square. The dispersion coefficient D is the same everywhere and is represented as a positive real number. Figure 3.2 illustrates the portion of Lake Michigan under study and shows the discretization of the shoreline and some representative squares.

The transition function δ is simulated by a procedure which accepts as a parameter the length of time Δt desired for a one-step transition. In the simplest terms, the new position of a drifter is given by its current position \underline{q} plus $\underline{v}(\underline{q})\Delta t + \underline{N}(0, 2D\Delta t)$, where \underline{N} is a pair of normal, pseudo-random numbers with mean 0 and variance $2D\Delta t$ (see figure 3.3a). In many cases, though, Δt could be long enough for a drifter to be displaced several squares (figure 3.3b). Rather than jumping squares and thus possibly creating errors, each drifter iterates through enough subtransitions to yield a high probability of visiting each square in its path. The length of time corresponding to a single subtransition is computed from the local velocity $\underline{v}(\underline{q})$ and D . It is the minimum of the following three times: (1) the time required for the largest component of $\underline{v}(\underline{q})$ to move a drifter $1/2$ square, (2) the mean time required for a drifter to be moved at least $1/2$ square 1% of the time by the random effect \underline{N} , and (3) the amount of time remaining for the entire transition.

LAKE MICHIGAN







Figure 3.2: Discretization of southern Lake Michigan.



a) Drifter transition.



b) Drifter transition jumping a square.



c) Iterated subtransitions.

Figure 3.3: Monte Carlo drifter transitions.

This guarantees that at least 99% of the time, a drifter will land in one of the nine squares adjacent to it (including the square it's in) after one subtransition. The transition of a drifter is iterated in this fashion until all of Δt is used up or until it hits land, whichever comes first. Figure 3.3c illustrates the iterated transition of a single drifter. In practice several drifters are in transit at once. This is easily handled procedurally by iterating each drifter through its transition in its turn.

Having been described verbally, the simulation may now be summarized algorithmically.

Let $(Q(i), i=1, \dots, K)$ be an array containing the position vectors of the K drifters,

$(T(i), i=1, \dots, K)$ be an array containing elapsed drift times for the K drifters (used for bookkeeping purposes),

S be the length along a side of each of the 4096 squares,

$(V(i,j), i=1, \dots, 64; j=1, \dots, 64)$ be an array containing the velocity vectors for each of the 4096 squares,

$(G(i,j), i=1, \dots, 64; j=1, \dots, 64)$ be an array containing either "land" or "water" for each of the 4096 squares,

D be a scalar containing the value of the dispersion coefficient,


ΔT be the time interval associated with one transition, and

$N(x)$ be a pseudo-random number generator generating componentwise independent normal random vectors, each component having mean 0 and variance x .

The transitions take place as follows [comments occur in brackets]:

1. For each $i, i=1, \dots, K$ [For each drifter]
 - 1.1 $ET \leftarrow 0$ [Elapsed time for transition]
 - 1.2 $X \leftarrow \text{truncate}(\min(\max(Q_x(i)/S + 1, 1), 64))$
 - 1.3 $Y \leftarrow \text{truncate}(\min(\max(Q_y(i)/S + 1, 1), 64))$
 [X and Y are the indices of the inbounds square nearest
 drifter i]
 - 1.4 $G(X,Y) = \text{"land"}$?
 Yes: go to 1.13 [Drifter is done]
 No: continue
 - 1.5 $TV \leftarrow 0.5 \cdot S / \max(V_x(X,Y), V_y(X,Y))$
 - 1.6 $TD \leftarrow 0.5 \cdot S / (13 \cdot D)$ [Diffusion time to move 1/2 square maximum]
 - 1.7 $TT \leftarrow \Delta T - ET$ [Time remaining in transition]
 - 1.8 $TG \leftarrow \min(TV, TD, TT)$ [Subtransition time]
 - 1.9 $\underline{Q}(i) \leftarrow \underline{Q}(i) + \underline{V}(X,Y) \cdot TG + \underline{N}(2 \cdot D \cdot TG)$ [Subtransition]
 - 1.10 $T(i) \leftarrow T(i) + TG$ [Cumulative time afloat]
 - 1.11 $ET \leftarrow ET + TG$ [Increment elapsed time]
 - 1.12 $ET = \Delta T$? [Transition complete?]
 Yes: go to 1.13
 No: go to 1.2
 - 1.13 next i
2. END

In a typical simulation run, the drifters are all initialized at the same point. Then the transition algorithm is applied repeatedly until all the drifters wind up ashore, or until a predetermined number of transitions have been made. Examples of typical simulation runs appear in figures 3.4 and 3.5. The velocity fields pictured are sampled at a



Velocity field.





Release: $t=0$. $t=5$. $t=10$.

Figure 3.4: Simulation of 25 drifters.



Velocity field.






Release: $t=0$. $t=5$. $t=10$.

Figure 3.5: Simulation of 25 drifters.


. . . /



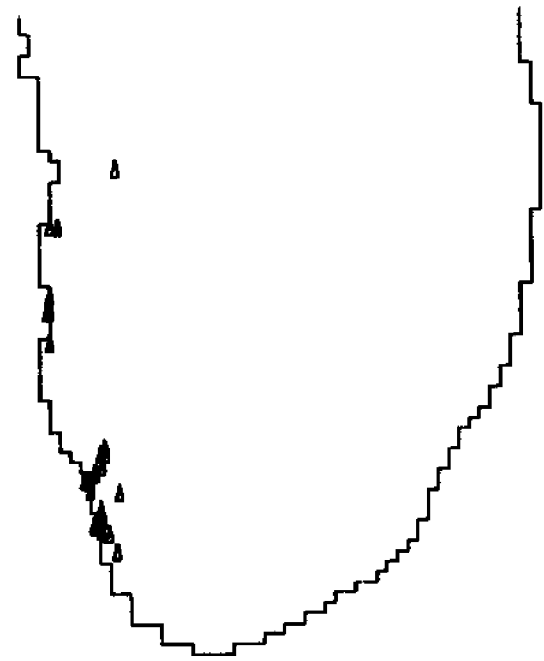
t=15.



t=20.



t=25.



t=30.

/ . . . 3.5 cont'd.

smaller rate than once per cell in order to render the figures less cluttered. As can be seen, the normal shape of the drifter clusters is often not maintained.

3.3 Data Interpretation and Goodness-of-Fit



Before one begins judging hypotheses under the authority of armloads of data, he should decide ahead of time which aspects of the data are important and which are not. This interpretation stands between the raw data and any conclusions drawn from it, so it must be given careful prior consideration. In drifter terms, each datum is the time and location of release and the time and location of recovery. One's task is to decide what relevant information each of these data bears with regard to the selection of an hypothesized current pattern from among the several available.

The most basic assumption concerning recovery data is that every recovered drifter is carried from its release point to its reported recovery point by water currents. Any hypothesis considered to be in concord with a given recovery must have a current pattern which is capable of getting the drifter from its release point to that recovery point within the time interval indicated by the recovery datum. Naturally, "in concord" is a relative term, and the degree of concordance will be some measure of *how closely* an hypothesized field can bring a drifter to its known recovery point in the appropriate amount of time.

The overall degree of compatibility of an hypothesis with all the data can be taken as some overall measure of concord, which, one might suspect, could be some combination of the degrees of agreement of the hypothesis with the individual recovery data. But how to combine them? How do the data interact? First, it is reasonable to assume that if an

hypothesis exhibits some measure of concord with respect to a given recovery, it will automatically have the same degree of agreement with another, identical recovery from the same release. This is to say that two equal data reveal no more about the currents giving rise to them than one does; the important thing is the single observation that whatever currents were responsible were able to transport at least one drifter from point X to point Y in time t . This implies that redundant data may be discarded with no loss of information.




Those with statistical background will surely object that whether 90% of the drifters from a given release end up at point A and 10% at point B or 10% end at A and 90% end at B could be a crucial factor in picking one hypothesis over another. But this kind of behavioral difference can be the result of such a trivial structural difference that any distinctions made on that basis would be overblown indeed. A single example suffices to illustrate the point. Consider the two velocity fields shown in figure 3.6. They are identical except for a small shift along the vertical axis. But drifters released from the same point behave quite differently under the two hypotheses, as far as the relative number going one direction or another is concerned. If one of these current patterns were the real pattern in a system and the other were the hypothesis being tested, one would not want that hypothesis downgraded because of the behavioral differences shown here. By selecting irredundant representatives from the data, one observes a smaller behavioral discrepancy between the real system and the simulated hypothesis. Therefore, in determining the overall degree of compatibility of an hypothesis with the data, multiplicities in the data are best eliminated.



Above: First velocity field.

Above right: Release point.

Right: 14 time steps later.



Above: Second velocity field.

Above right: Same release point.

Right: 14 time steps later.




Figure 3.6: Comparison of drifter distributions from two velocity fields, one being shifted slightly along the Y-axis.

Now consider two recoveries at the same point, from the same release, but at different times. What is implied here? First, of course, they could legitimately have washed ashore at different times just from the diffusive spread of the drifter cluster. Or the latter one might have gotten caught in a persistent eddy or an area of calm water for awhile before resuming its trajectory to the recovery point. Perhaps the latter one may also have come close to the recovery point with the earlier one but, rather than landing, continued in a gyre which brought it back around to that point later. More likely, they both landed at the same time, but the latter one lay on the beach awhile before it was found. While the first three occurrences are natural and may have a legitimate bearing on whatever conclusions are drawn therefrom, the last one can be misleading, and any attempt to minimize its effect will be a boon to the hypothesis testing scheme.

It is tempting just to eliminate from consideration all but the earliest recoveries at a given point from a given release. That this is indeed helpful is shown by a simple probabilistic argument. Let p be the probability that any given drifter lying on a certain section of beach will be picked up on any given day independently of how long it has lain there. The probability of picking it up exactly n days after its arrival is the joint probability of not picking it up on the first $n-1$ days and of picking it up on the n th:

$$\begin{aligned}
 & \text{Pr}[\text{picking up drifter } d \text{ on day } n] \\
 &= \text{Pr}[\text{not picking } d \text{ up on previous } n-1 \text{ days}] \cdot \text{Pr}[\text{picking } d \text{ up} \\
 &\quad \text{on any day}] \\
 &= \text{Pr}[\text{not picking } d \text{ up on a given day}]^{n-1} \cdot p \\
 &= (1-p)^{n-1} p.
 \end{aligned}$$

The expected length of time for lying onshore before being recovered is then:

$$\begin{aligned} E[\text{prerecovery time ashore}] &= \sum_{i=1}^{\infty} i(1-p)^{i-1} p \\ &= p \sum_{i=1}^{\infty} i(1-p)^{i-1} \\ &= \frac{p}{[1 - (1-p)]^2} = \frac{1}{p} \text{ days.} \end{aligned}$$

This length of time is the uncertainty expected in the recovery time due to lying ashore undiscovered when all recoveries from a given area and a given release are considered. This uncertainty can be reduced considerably by considering only the first-found drifter in a given area from a given release. Assuming k drifters arrive on a section of beach from the same release and that each independently has probability p of being picked up on any given day as before, the probability that the first of the k is recovered on day n is:

$$\begin{aligned} &\Pr[\text{first of } k \text{ drifters recovered on day } n] \\ &= \Pr[\text{no recoveries in } k \text{ on previous } n-1 \text{ days}] \cdot \Pr[\text{at least one} \\ &\quad \text{recovery in } k \text{ on a given day}] \\ &= \Pr[\text{no recoveries in } k \text{ on a given day}]^{n-1} \cdot (1 - \Pr[\text{no re-} \\ &\quad \text{coveries in } k \text{ on a given day}]) \\ &= (1-p)^{k(n-1)} [1 - (1-p)^k] \end{aligned}$$

The expected waiting time for the first recovery after arrival on shore is:

$$\begin{aligned} E[\text{prerecovery time for first drifter in } k] \\ &= \sum_{i=1}^{\infty} i(1-p)^{k(i-1)} [1 - (1-p)^k] \\ &= [1 - (1-p)^k] \sum_{i=1}^{\infty} i[(1-p)^k]^{i-1} \end{aligned}$$

$$= \frac{1 - (1-p)^k}{[1 - (1-p)^k]^2} = \frac{1}{1 - (1-p)^k} \text{ days.}$$

This expectation is plotted for various values of k and several p (figure 3.7). For $k=1$, of course, the value is the same as the previously calculated expectation. The rapid decrease in this recovery time uncertainty as k becomes large is readily apparent.

That this probabilistic model is only a crude approximation to reality is unquestionable. Beside the fact that p is not fixed but probably relatively high on weekends and holidays is the tendency for some people, when finding a drifter, to look more intensely for others. This "beachcomber effect" entails an increase in p for drifters found after the first one and a consequent reduction in expected time ashore for any given drifter. Nonetheless, this uncertainty will still be greater than that of the first recovery, so the implication of the crude model is not lost: it makes sense to eliminate all but the first recovery from an isolated group from the same release.

How does this elimination process affect implications drawn from the other causes of sequential recoveries at the same point? In the case of time differences due to diffusion, considering only the first arrivals in a given area would seem to bias the testing process toward hypotheses with higher average current speeds than those actually responsible for the data. But these early arrivals, being from the fringe of a diffuse drifter cluster, will be few and far between and should not interfere with the bulk of arrivals coming from the denser core of the cluster. As for late returns caused by drifters getting caught along the way in a persistent eddy or still water or by drifters stranding on the second or third time around a large gyre, they will be




Figure 3.7: Expected time of finding first of several drifters landing on day 1. p is the probability of finding any drifter on a given day.

eliminated if earlier returns arise from the same location. This implies that there will be no data of this type to support an hypothesis containing the above-mentioned features or to discredit one lacking them. But the implications of data varying only along the time dimension are rather tenuous anyway, taken alone, since it is difficult to tell by which process the time variation arose. To confirm or deny an hypothesis with less equivocation would require first-arrival supporting data varying along spatial dimensions. In the case of a large gyre, this would mean progressively later returns along the circuit of the current. With such data present, the elimination of late arrivals from a given area and the same release is no great loss as far as drift processes are concerned and is a potential benefit in the reduction of recovery time uncertainty due to drifters lying on the beach undiscovered.

Coalescing the decisions made so far reveals two major aspects of the intended data interpretation process. First, the implication of each available datum on an overall compatibility measure for a given hypothesis is a function of how closely that hypothesis can cause drifters to approach the known recovery site in the appropriate time. Second, the data to be made available are only those single recoveries arriving first at their respective recovery sites from each respective release point. The resulting elimination of data redundancy serves to moderate the possibly misleading implications of any distribution of returns from a given release; the discarding of late returns reduces the uncertainty in recovery time caused by drifters lying about undiscovered for awhile. What remains to be settled is the exact form of the goodness-of-fit measure.

Consider a given recovery point (q_R, t_R) and several hypotheses (\underline{v} 's) to be tested against it. By running a simulation under each hypothesis with many drifters starting at the known release point, one gets an evolving distribution of drifters representing the density function $\Psi_{\underline{v}}(q, t)$ for each hypothesized \underline{v} . Assuming a given \underline{v} to be correct, the recovery point (q_R, t_R) must be a sample from the distribution given by the corresponding $\Psi_{\underline{v}}(q, t)$. The measure of concord between a recovery and any given hypothesis, then, should be some indication of the likelihood that the recovery point could have come from the distribution calculated from that hypothesis. In traditional statistical terms (Fisher, 1959), this likelihood may be taken simply as the value of $\Psi_{\underline{v}}(q, t)$ at the point (q_R, t_R) . By assuming this point to be fixed (which it is when the experiment is over) and considering $\Psi_{\underline{v}}(q_R, t_R)$ under every available hypothesis \underline{v} , one can define a function of \underline{v} , $L(\underline{v}; q_R, t_R) = \Psi_{\underline{v}}(q_R, t_R)$, known as the likelihood function. In this formulation, the \underline{v} which maximizes L (the maximum likelihood estimate) will be considered the most compatible with the recovery.

By extending this principle to several recoveries, it is possible to rate the overall compatibility of a given hypothesis with all the data. Consider first the data arising from release $(q_{0,j}, t_{0,j})$, for example. The n_j associated recoveries can be considered independent samples from the same distribution $\Psi_{\underline{v},j}$, yielding a joint density function

$$\begin{aligned} & \Psi_{\underline{v},j,n_j}(q_{1,j}, t_{1,j}, \dots, q_{n_j,j}, t_{n_j,j}) \\ &= \Psi_{\underline{v},j}(q_{1,j}, t_{1,j}) \cdot \Psi_{\underline{v},j}(q_{2,j}, t_{2,j}) \cdots \Psi_{\underline{v},j}(q_{n_j,j}, t_{n_j,j}). \end{aligned}$$

By considering recoveries from the m different releases as independent

too, the overall joint density of all N recoveries can be written

$$\begin{aligned} \psi_{\underline{v}, N}(q_{i,j}, t_{i,j}, i=1, \dots, n_j, j=1, \dots, m) \\ = \prod_{j=1}^m \psi_{\underline{v}, j, n_j}(q_{1,j}, t_{1,j}, \dots, q_{n_j,j}, t_{n_j,j}). \end{aligned}$$

Each of the $(q_{i,j}, t_{i,j})$ can be fixed to its associated recovery $(q_{R,i,j}, t_{R,i,j})$, yielding the likelihood under \underline{v} of the recovery data. This joint likelihood function, then, which is to be maximized over the available hypotheses, is

$$\begin{aligned} L_N(\underline{v}; q_{R,i,j}, t_{R,i,j}, i=1, \dots, n_j, j=1, \dots, m) \\ = \prod_{j=1}^m \prod_{i=1}^{n_j} L_j(\underline{v}; q_{R,i,j}, t_{R,i,j}). \end{aligned}$$


Given a set of data, it is this product which must be estimated for each hypothesis considered in order to pick the best one.

So far, so good. If it were possible to consider every conceivable \underline{v} , then the maximum likelihood technique could be used to choose the one most compatible with the data. But of course it's not possible to do this; instead, only a small set of \underline{v} 's can be considered. So one is constrained to work with the hypotheses available in the hopes that if one of them is close to the overall best \underline{v} , it will be selected as the best from the set considered. But the example of figure 3.6 has already shown that extremely different drifter distributions and, hence, extremely different likelihood functions can be gotten from very similar velocity fields. Therefore, for small sets of hypotheses, the likelihood estimate can be a very fickle indicator of compatibility. What is needed, instead, is an indicator which is relatively invariant over wide fluctuations in the joint likelihood function caused by small


perturbations in the velocity field.

One solution to the problem would be to recognize clusters of drifters. In this formulation, each $\psi_{\underline{v},j}$ could be thought of as a normalized sum of several localized density functions, each corresponding to one of the clusters. (See figure 3.8.) Rather than being calculated from $\psi_{\underline{v},j}$, the likelihood value corresponding to a given recovery would be calculated only from the local density function corresponding to the cluster nearest the recovery point at the time of recovery. This would cancel the effect of any wide differences in drifter density between clusters arising from slight perturbations in \underline{v} . The resulting likelihood measure would therefore reflect the ability of \underline{v} to get *some* drifters to the recovery point from release j in the proper time rather than *most* of them.


The recognition of clusters can be very troublesome, though. Beside the issue of what is a cluster and what is not one but several clusters is the problem of efficiency. This potentially involved cluster analysis would have to be carried out once for every hypothesis tested, for each time a recovery was made, for each release. Rather than pursuing a full blown clustering algorithm to its mega-microsecond-consuming glory, one may consider the extreme of simplifications: one drifter equals one cluster. The primary advantage of this is the unequivocal choice as to which simulated drifters to use in calculating the local density function: the single one closest to the recovery point at the recovery time. The primary disadvantage is the lack of neighboring drifters in the "cluster" to give some idea of the shape and spread of the local density function and, hence, its value at the recovery point. It will be necessary to improvise here.



Drifter clusters.



Equidensity contours for cluster density functions.



Equidensity contours for superposed density function.

Figure 3.8: Composition of Ψ from local cluster densities.

The assumption that each drifter comprises an autonomous cluster implies that the overall density function $\Psi_{\underline{v},j}$ can be represented by an average of many local density functions, each corresponding to one of the drifters. Because each local density function at a given time must be determined only by a drifter's position, there is no basis for assigning other than identical shapes to all of them. Only the means will be allowed to vary. But what shape to choose? The best shape, of course, will be the one which, after the density superposition, yields the composite density $\Psi_{\underline{v},j}$ which best reflects the correct overall density. But in most cases, the correct density is unknown. The trick is to pick a local density that seems to work in cases where the correct $\Psi_{\underline{v},j}$ is known and to hope it works in other cases. The simplest case where $\Psi_{\underline{v},j}$ is known is the one in which $\underline{v}(q) = \underline{s}$, a constant. As was shown in Chapter 2, starting from an impulse at $(\underline{q}_{0,j}, t_{0,j}) = (\underline{0}, 0)$, $\Psi_{\underline{s},j}(q, t) = f_N(q; \underline{s}t, 2Dt)$, the joint density function of two independent normal random variables with variances $2Dt$ and means $s_x t$ and $s_y t$ respectively. Without loss of generality, one can consider the x-dimension only, due to independence, and let $s_x = 0$, leaving $\Psi_{0,j}(x, t) = f_N(x; 0, 2Dt)$. Each simulated drifter represents an independent sample, x_i , from this distribution at any given time step. With n simulated drifters, it is this density function which must be imitated by a superposition of n local density functions.

A likely candidate for the local density function is the normal density. This is neither a random choice nor an inspired one: it simply has nice properties when used in comparing likelihoods, as will be shown later. The mean of the local density for a given drifter will be the drifter's position, x_i . The variance will be σ^2 for each drifter

and will be figured out soon. The composite density function formed by averaging these normal densities is $g(x) = \frac{1}{n} \sum f_N(x; x_i, \sigma^2)$, where this and subsequent unlabeled sums are assumed to run over $i=1, \dots, n$. (See figure 3.9 for examples of this superposition process.)

To test g , one can establish its expected mean and variance to compare with the 0 and $2Dt$ of the drifters' parent distribution. Actually, the following derivations are independent of the normality of the local density or of the parent distribution, but it is helpful to keep these densities in mind as a point of reference. The mean of x , conditional on the x_i is

$$\begin{aligned} E[x | x_1, \dots, x_n] &= \int \frac{x}{n} \sum f_N(x; x_i, \sigma^2) dx \\ &= \frac{1}{n} \sum \int x f_N(x; x_i, \sigma^2) dx \\ &= \frac{1}{n} \sum x_i, \text{ where} \end{aligned}$$

these and subsequent unlabeled integrals are assumed to be over the entire space of the dummy variable. The expectation of this sample average is just the mean of the normal distribution from which the x_i 's came, or 0. So the mean of x under $g(x)$ is the same as its mean under the parent distribution.

The conditional variance of x under g is $E[x^2 | x_1, \dots, x_n] - E[x | x_1, \dots, x_n]^2$. The first term can be calculated

$$\begin{aligned} E[x^2 | x_1, \dots, x_n] &= \int \frac{x^2}{n} \sum f_N(x; x_i, \sigma^2) dx \\ &= \frac{1}{n} \sum \int x^2 f_N(x; x_i, \sigma^2) dx \\ &= \frac{1}{n} \sum (x_i^2 + \sigma^2). \end{aligned}$$

The second term is $E[x | x_1, \dots, x_n]^2 = \frac{1}{n^2} [\sum x_i]^2$, leaving

$$\text{Var}(x | x_1, \dots, x_n) = \frac{1}{n} \sum (x_i^2 + \sigma^2) - \frac{1}{n^2} [\sum x_i]^2.$$

The expectation of this variance over the x_i is

$$\begin{aligned} E[\text{Var}(x)] &= \frac{1}{n} \sum (E[x_i^2] + \sigma^2) - \frac{1}{n^2} E[(\sum x_i)^2] \\ &= \frac{1}{n} \sum ((2Dt) + \sigma^2) - \frac{1}{n^2} (2nDt) * \\ &= 2Dt + \sigma^2 - \frac{2Dt}{n}. \end{aligned}$$

Letting this equal the variance of the parent distribution, get

$$\begin{aligned} \text{Var}(x_i) = 2Dt &= 2Dt + \sigma^2 - \frac{2Dt}{n}, \\ \text{or } \sigma^2 &= \frac{2Dt}{n}. \end{aligned}$$

So a local normal density with mean x_i and variance $\frac{2Dt}{n}$ will yield a composite density $g(x)$ with expected mean and variance equal to those of the density governing the drifter positions.

It might be possible to go on, using variational techniques, to see if a local normal distribution is indeed the optimal one given some expected goodness-of-fit measure. But the fact that g has the proper mean and variance plus a few random examples favorably comparing the composite g with the proper density should suffice to admit the local normal density into the hypothesis testing scheme. Figure 3.9 shows several composite densities formed from x_i 's picked at random from the normal curve shown. As can be seen, the composite density comes about as close as could be hoped for to parent density function--at least when this parent distribution is normal. From here one must cantilever his belief

*The expectations of the cross-product terms in the squared sum are zero because the means of the x_i are zero.




Figure 3.9: Superpositions of k local densities centered on points chosen at random from the parent normal distribution.

in its general applicability.

Returning, finally, to hypothesis testing, one can now calculate the local likelihood that the recovery (q_R, t_R) came from the hypothesized \underline{v} , beginning at release j . This is simply the value of the local density function implied by the simulated drifter nearest q_R at time t_R and can be obtained by extending the normal local density described to two dimensions like so:

$$L(\underline{v}; q_R, t_R) = \frac{n}{4\pi Dt_R} \exp[-|q_R - q_{\underline{v},j}|^2 n/(4Dt_R)], \text{ where}$$

$q_{\underline{v},j}$ is the location of the drifter which was simulated under \underline{v} from release j and nearest q_R at t_R , and n is the number of drifters in the simulation.

The joint likelihood of \underline{v} under all N data points $(q_{R,i,j}, t_{R,i,j})$, $i=1, \dots, n_j$, $j=1, \dots, m$ is the product of the individual likelihoods, or

$$\begin{aligned} L_N(\underline{v}; q_{R,i,j}, i=1, \dots, n_j, j=1, \dots, m) \\ &= \prod_{j=1}^m \prod_{i=1}^{n_j} \frac{n}{4\pi Dt_{R,i,j}} \exp[-|q_{R,i,j} - q_{\underline{v},i,j}|^2 n/(4Dt_{R,i,j})] \\ &= \left[\prod_{i=1}^m \prod_{j=1}^{n_j} \frac{n}{4\pi Dt_{R,i,j}} \right] \\ &\quad \exp\left[\frac{-n}{4D} \sum_{j=1}^m \sum_{i=1}^{n_j} |q_{R,i,j} - q_{\underline{v},i,j}|^2 / t_{R,i,j} \right], \text{ where} \end{aligned}$$

$q_{\underline{v},i,j}$ is the location of the drifter which was simulated under \underline{v} from release j and nearest $q_{R,i,j}$ at $t_{R,i,j}$. In comparing hypotheses, only the relative values of L_N for various \underline{v} (likelihood ratios) are of interest, so the bracketed product above may be dropped without affecting the comparison. Also, the log of L_N may be taken without affecting the relative resultant ranking, yielding a modified "likelihood":

$$L'_N(\underline{v}; \underline{q}_{R,i,j}, i=1,\dots,n_j, j=1,\dots,m) \\ = \frac{-n}{4D} \sum_{j=1}^m \sum_{i=1}^{n_j} |q_{R,i,j} - q_{\underline{v},i,j}|^2 / t_{R,i,j}.$$

From this it is obvious that the best \underline{v} under the original likelihood is the one which minimizes the sum of the $|q_{R,i,j} - q_{\underline{v},i,j}|^2 / t_{R,i,j}$, so it is this sum which will be calculated for each \underline{v} under consideration.

Intuitively, minimizing this sum makes sense. The units of each term are square distance per unit time, the same as the diffusion coefficient D . What these terms are measuring, essentially, is the deviation rate from the hypothesized \underline{v} required to get the drifters to the known recovery points in the requisite time intervals. Therefore, the \underline{v} requiring the least average deviation rate to agree with the data will be the one chosen. But there is a bonus here. As was shown before, uncertainty in recovery time can be reduced by considering only the first of many recoveries at the same point from a given release. This works so long as there is a multitude of recoveries to consider. This multitude is usually manifest in the earlier returns from a release (see figure 3.12); but as the ranks of a propagating drifter cluster thin, strandings tend to be more isolated. Hence late strandings do not have the advantage of multiplicity necessary for reducing the expected time from stranding to first recovery. But this expected time for a single drifter is independent of drift time, and here is where the bonus comes in. The above measure of deviation rate, being inversely proportional to the time of recovery, is less sensitive to constant uncertainties in recovery time for late recoveries than for early ones. Thus, where the data elimination process fails to reduce this uncertainty, the goodness-of-fit measure tends to limit its effects anyway.

With a data selection policy and a compatibility measure, the link between raw drifter returns and hypothesis testing has been forged. The tasks remaining are to implement the selection policy and testing scheme algorithmically, to find or formulate some hypotheses to test, and to go out and get some data to test them against.

3.4 Hypothesis Testing Procedure

Assuming a body of drifter data has been accumulated, one must first comb through it to select those points in agreement with the data acceptance policy already outlined. For any given release, the idea is to pick for each section of shoreline or area in the water, that drifter recovered there first. But what is meant by "section" or "area"? As in the simulation, the domain of the drifters is divided into squares. The size of one square is the same as four squares in the simulation, or 5.5 km on a side in Lake Michigan. Each recovery position can be mapped into one of the squares. When several recoveries fall into one square, all but the earliest one are thrown out. In this manner, each square which contains a recovery will have a number associated with it corresponding to the earliest recovery time for that square (figure 3.10a). But the boundaries between squares are an extremely artificial barrier against comparison of recoveries which are very close together but which happen to fall in adjacent squares. Therefore a second step is required to make comparisons between adjacent squares and eliminate the later recoveries.

Consider a square which has a recovery time assigned to it. If this time is less than or equal to the defined times for all eight adjacent squares, the square will be called a distinguished square. The sieve of acceptance for a recovery is as follows for each square with a




Figure 3.10a: Recoveries from a single release. Numbers indicate time of earliest recovery in each square.




Figure 3.10b: Earliest recoveries from each square. Circles indicate distinguished squares; X's indicate rejected recoveries.

defined time:

1. Is the square a distinguished square?

Yes: Accept the recovery contained in it.

No: Continue.

2. Are any of the eight adjacent squares distinguished squares?

Yes: Reject the recovery.

No: Accept the recovery.

Figure 3.10b illustrates the same squares shown in 3.10a, but containing only the earliest recovery in each. Those containing recoveries accepted on the basis of 1 above are circled, and those with recoveries rejected on the basis of 2 are X-ed. All non-X-ed recoveries are thus accepted. This process is carried out on the body of recoveries from each release. The data used in the hypothesis testing scheme are the union over all the release points in an experiment of the recoveries accepted from each.

Given the prepared data and several hypotheses to test, one may now define the hypothesis testing algorithm.

Let $(QO(i), i=1, \dots, NO)$ be an array of NO release points,

$(TO(i), i=1, \dots, NO)$ be the corresponding array of release times.

$(NR(i), i=1, \dots, NO)$ be an array containing the number of recoveries for each release,

$(QR(i,j), i=1, \dots, NO; j=1, \dots, NR(i))$ be the array of recovery points,

$(TR(i,j), i=1, \dots, NO; j=1, \dots, NR(i))$ be the corresponding recovery times,

$(U(i,j,k), i=1, \dots, 64; j=1, \dots, 64; k=1, \dots, NH)$ be an array of

NH velocity hypotheses to be tested,

(SCORE(i), $i=1, \dots, NH$) be the corresponding scores,

K , Q , T , V , G , D , S , and $N(x)$ be defined as for the simulation algorithm, and

$\Delta T = 1$ day

The algorithm proceeds as follows [comments occur in brackets]:

1. For each kh , $kh = 1, \dots, NH$ [For each hypothesis]
 - 1.1 $V(i, j) \leftarrow U(i, j, kh)$, $i=1, \dots, 64$; $j=1, \dots, 64$ [Velocity field is hypothesis kh]
 - 1.2 SCORE(kh) $\leftarrow 0$
 - 1.3 For each m , $m = 1, \dots, NO$ [For each release]
 - 1.3.1 $Q(i) \leftarrow Q0(m)$, $i=1, \dots, K$ [Initialize drifters]
 - 1.3.2 $T(i) \leftarrow T0(m)$, $i=1, \dots, K$
 - 1.3.3 $TMAX \leftarrow \max(TR(i), i=1, \dots, NR(m))$ [Latest recovery]
 - 1.3.4 For TIME = (T0(m) + 1), ..., TMAX
 - 1.3.4.1 Simulate one transition
 - 1.3.4.2 For each n , $n=1, \dots, NR(m)$ such that $TR(m, n) =$
TIME [For each recovery at time TIME]
 - 1.3.4.2.1 SCORE(kh) \leftarrow SCORE(kh) +
 $\min\{|Q(i) - QR(n)|^2 / (TIME - T0(m))\}$,
 $i=1, \dots, K$; and $T(i) > TIME - 1$
 - 1.3.4.2.2 next n
 - 1.3.4.3 next TIME
 - 1.3.5 next m
 - 1.4 next kh
 2. print SCORE(kh), $kh = 1, \dots, NH$
 3. END

The hypothesis having the lowest score is the one judged most compatible with the data. In many cases, though, two very different hypotheses may tie for the lowest scores. In this situation neither can be accepted over the other just on the basis of the hypothesis testing scheme. Three solutions to such a problem present themselves immediately. The first is to get more data. Here, a method for finding release points optimally distinguishing between the two hypotheses will be very valuable. Such a method is described later. One difficulty with this solution is that by the time the data are analyzed to the point of recognizing conflicting hypotheses, the conditions under which the data were obtained might have changed. The second alternative is to apply external criteria to the selection between the hypotheses. This can involve direct current measurement results or just selection of the simplest hypothesis. Finally, one may try to come up with an additional hypothesis which proves better than either of the ones considered. This may be obtained by some sort of hybridization of the two contenders, or by a direct inference based on the data alone. The latter method is discussed later in this investigation.

3.5 Experiments in Lake Michigan

In the summer and fall of 1974, three research cruises were carried out in lower Lake Michigan. The purpose of these cruises was to determine the coastal current structure near Chicago as well as currents elsewhere in southern Lake Michigan. During each of the July, August, and October cruises, surface current drifters were released. The details of these releases may be read in Monahan and Pilgrim (1975, hereafter referred to as "Coastwise Currents"); a summary of the release points is presented in figure 3.11. The recoveries from





Figure 3.11: Drifter release locations for 1974.


these release points are presented in detail in Coastwise Currents. These data were submitted to the data selection algorithm and the results are summarized in figure 3.12 along with one of the detailed presentations to demonstrate the action of the selection process.

Several hypotheses concerning the currents in this area were found in the literature. They were digitized, and the results are summarized in vector form in figure 3.13. These include Harrington's (1895) inference from drift bottle returns (3.13a), measurements made by the Federal Water Pollution Control Administration (FWPCA, 1967) (3.13b and c), observations reported by Ayers, *et al.* (1958) (3.13e and f), and conjectures made in Coastwise Currents based on a visual inspection of the return data themselves (3.13g and h). Since the absolute speeds implicit in each hypothesis could legitimately vary from those shown without discrediting the hypothesis itself, tests on each hypothesis were made using speed multipliers of 0.1, 0.25, 0.5, 0.75, 1.0, and 1.5. Each multiplier was applied to the entire velocity field, so the speeds shown for any field do not change in ratio to each other.

The hypothesis testing algorithm was run with each of the hypotheses applied to the data from each of the three experiments separately. The reason for this is an apparent shift in the currents (see Coastwise Currents) between the studies. The dispersion coefficient D assumed for the simulation routine was in the neighborhood of $7 \text{ m}^2/\text{sec}$. This value is slightly higher than the $5.5 \text{ m}^2/\text{sec}$ effective dispersion rate observed in a drogue study of Lake Michigan by the FWPCA (1967). Using an inflated value for D , though, has the benefit of moderating the effects of small errors in the digitization of the hypotheses by adding a larger diffusive or "error" term at each time step. Twenty-five



Original recoveries.



Selected recoveries.

- 15 Release
- ▽ Recovery north of chart
- ◇ Recovery, 10 days adrift
- ◇ Recovery, 20 days adrift
- ◇ Recovery, 40 days adrift
- ◆ Recovery used in inference (Chapter 5)












Figure 3.12a: Selected recoveries from July, 1974, drifter releases.

... /







/ . . . 3.12a cont'd .

. . . /







/ . . . 3.12a cont'd.

. . . /







/ . . . 3.12a cont'd.

. . . /



/ . . . 3.12a cont'd.

. . . /



/ . . . 3.12a cont'd .

- 30 Release
 - ▽ Recovery north of chart
 - ◇ Recovery, 10 days adrift
 - ◇ Recovery, 20 days adrift
 - ◇ Recovery, 40 days adrift
 - ◆ Recovery used in inference
- (Chapter 5)














Figure 3.12b: Selected recoveries from August, 1974, drifter releases.

... /







/ . . . 3.12b cont'd.

. . . /







/ . . . 3.12b cont'd.

. . . /



/ . . . 3.12b cont'd.

. . . /



/ . . . 3.12b cont'd.

- 4B Release
- ▽ Recovery north of chart
- ◆ Recovery, 10 days adrift
- ◇ Recovery, 20 days adrift
- ◊ Recovery, 40 days adrift











Figure 3.12c: Selected recoveries from October, 1974, drifter releases.

... /




/ . . . 3.12c cont'd.


. . . /




/ . . . 3.12c cont'd.




a) Harrington (1895), from drift bottle returns.



b) FWPCA (1967) observations (summer, SSW winds).




c) FWPCA (1967) observations (summer, NNE winds).




d) Ayers et al. (1958), summer measurements.

Figure 3.13: Hypothesized velocity fields.


... /




e) Kizlauskas and Katz (1973),
mathematical model (SSW wind).



f) Kizlauskas and Katz (1973),
mathematical model (NNE wind).



g) Monahan and Pilgrim (1975),
from drifters (July, 1974).



h) Monahan and Pilgrim (1975),
from drifters (August, 1974).

/ . . . 3.13 cont'd.

drifters were simulated in each test. Whether or not this is enough can be intuited by looking at the resulting goodness-of-fit values as a function of the velocity multiplier. If an insufficient number of simulated drifters were used, one would expect these curves to be jagged due to poor statistics. That the July and August curves are not (figure 3.14) attests to the adequacy of the number of drifters. That the October curves are rather jagged is probably due mostly to the sparse return data.

The results of the test are plotted in figure 3.14. The compatibility (actually "incompatibility") value is the sum computed by the algorithm, divided by the number of returns. This gives some idea of the average deviation rate required for the hypothesis to agree with the data and forms a basis of comparison with D.

July's results are very much in line with the conjectures made in Coastwise Currents. Of the previously published hypotheses, Kizlauskas and Katz's (figure 3.13e) theoretical results come closest to agreement with the data using a multiplier of 0.25. Very close behind, but without nearly so pronounced a dip is Ayers' hypothesis (figure 3.13d). On the whole, though, the hypothesis put forth in Coastwise Currents based on the data (figure 3.13g) shows to be the most compatible with the data.

The August outcomes paint a different picture, however. Of the previously published hypotheses, the FWPCA measurements (figure 3.13b) have the edge over the Kizlauskas and Katz hypothesis (figure 3.13e) conjectured as best in Coastwise Currents. Moreover, the hypothesis made in that report (figure 3.13h), based on the data, fares less well than either of these.




Figure 3.14a: Hypothesis testing results for July, 1974. Letters refer to figure 3.13.




Figure 3.14b: Hypothesis testing results for August, 1974. Letters refer to figure 3.13.




Figure 3.14c: Hypothesis testing results for October, 1974. Letters refer to figure 3.13.

Conclusions based on the October data are hard to come by legitimately. There were so few returns that selection of any hypothesis would certainly be farfetched. The hypothesis testing graphs illustrate this ambivalence, too, in their lack of sustained trends over the various multiplying factors--particularly under Kizlauskas and Katz's hypotheses. No attempt was made in Coastwise Currents to come up with an hypothesis based on the October return data, and none will be made here.

Under the assumption that the currents in lower Lake Michigan were stationary during each of the July and August experiments, one may conclude from the selected hypotheses that the general trend in July is a clockwise circulation with a north-south split occurring in the eastward current near Grand Haven, Michigan (latitude $43^{\circ}04'N$). There is, moreover, the possibility of a northward coastal countercurrent along the eastern shore. The results from August are less conclusive. While one would not expect a reversal in the overall circulation in one month, the hypothesis favored by the testing procedure (figure 3.13b) is mostly counterclockwise. On the other hand, a primarily clockwise hypothesis (figure 3.13e) is nearly as good a contender. Since both hypotheses have features in common, though, this partial dilemma may be circumvented by restricting one's conclusions to these aspects. Therefore, it can only be said that for August the coastal current along the western shore is northward, as is the coastal current along the northern stretches of the eastern shore. As far as October is concerned, no conclusions can be drawn.

3.6 Release Location Bias and Diagnostic Points

One aspect of drifter experiments not yet discussed is the choice

of release sites in the water body being investigated. This is an important consideration because the release locations as well as the proportionate number of recoveries from each can play critical roles in biasing the hypothesis testing results. The source of this bias lies in treating all the selected data equally. The hypothesized velocities leading from those release areas from which the most recovered drifters originate will naturally have the greatest bearing on the overall success of the hypothesis being tested. To understand why this is so, consider the following experiment: Two bunches of drifters are released, one at point A and one at point B. One hundred drifters originating at A are recovered along with ten beginning at B. In the testing algorithm, the result of each measurement between a recovery and the nearest simulated drifter will hinge on the velocities of every square that drifter visited before the comparison. Ultimately, then, all such comparisons with recovered drifters released at A depend on those velocities leading from A; likewise for B. Since more recovered drifters originated at A than at B, more measurements are likely to depend on the hypothesized velocities leading from there than from B. Therefore these velocities will have the greater effect in making or breaking the hypothesis. How much greater, of course, depends on the interaction of trajectories originating at A with those beginning at B. Though it would be extremely difficult to eradicate this kind of bias, an understanding of the factors contributing to it will help in the interpretation of the test results.

The central issue is the cause of a higher recovery rate from one release area as compared to another. An obvious possibility is that the one had more drifters released in it than the other. On the other hand,

the one point just may be better situated in the current field as far as getting drifters to shore is concerned. Experience in Lake Michigan has shown, for example, that drifters released close to shore have a greater chance of recovery than those loosed far offshore (Coastwise Currents). Although the data selection process will reduce these recovery rate differences when they are manifest as differences in recovery densities along the shore, the experiment designer may wish to exercise some control over the situation in his choice of releases. If the goal is a uniform recovery rate, a uniform distribution of release points would be a good place to start.

Often, though, the distribution of releases must be governed by constraints other than hypothetical ones, such as the amount of latitude allowed in the positions of the release platforms (e.g. ships, planes, drilling rigs). In these cases, one would want to get as much mileage in distinguishing between hypotheses as possible out of the release patterns available. Such was the situation during the July and August cruises on Lake Michigan. There was not enough time to cover the entire southern basin with drifters, so several areas had to be selected which showed the greatest promise of distinguishing among the hypotheses available at the time.

The selection of these release points was accomplished by simulating several releases for each of the available hypotheses (including figures 3.13a, d, e and f). By visual inspection of the strandings resulting from the simulations, it was determined that three areas had a great potential for yielding distinguishing return patterns. These were the nearshore areas off Chicago, Waukegan, and Grand Haven, Michigan. The following table summarizes the general recovery areas indicated for each

of them as a function of each hypothesis:

Hypothesis		Figure 13.3			
Release	a	d	e	f	
Chicago	Eastern shore	—	North on western shore	Southern shore	
Waukegan	Eastern shore	North on western shore	North on western shore	Southern shore	
Grand Haven	North on eastern shore	South on eastern shore	North on eastern shore	South on eastern shore	

Consequently, 50% of the 2390 drifters released were released in these three areas as follows:

	<u>July: 1615 drifters</u>	<u>August: 775 drifters</u>
Chicago	6%	29%
Waukegan	35%	--
Grand Haven	15.5%	6.45%

The percentages are by experiment.

The preselection of such diagnostic release points has largely paid off. For July, the concentration in the Waukegan area has effectively separated hypotheses a and f from d and e. The lesser concentration around Grand Haven has had less effect, though. This is probably due to the diverging current in that area shown by the more universally compatible hypothesis g and not by any of the others. The August results yield the greatest separation among the four original hypotheses with e

the clear favorite. What the August results also display is a close race between hypothesis e and one not considered before the experiment--namely, the FWPCA measurements (figure 3.13b). This one also fares rather well in July, but it is possible that a distinction could have been made with a differently biasing pattern of releases.

Having an hypothesis testing algorithm allows the selection of diagnostic release points with more confidence than that obtained from the visual comparison of simulation results used before the summer, 1974, experiments. This is because the measure of distinguishability between two hypotheses for a given release may be related directly to the expected outcomes of the testing scheme applied to an experiment in which one of the hypotheses is assumed true. The procedure is the following:

1. For each of a uniform distribution of release points and for each of two hypotheses, simulate the transport of some drifters, keeping track of those stranding ashore.
2. Filter the generated "recoveries" through the selection process.
3. For each of the two resulting sets of recoveries, run an hypothesis test on a release-by-release basis using the hypothesis generating the opposite data set. This will result in two compatibility scores for each release.
4. The distinguishability score for each release is taken to be the minimum of the two compatibility scores.

By choosing those release points having the highest distinguishability scores, one is selecting an experimental bias designed to accentuate the incompatibility of an hypothesis with the data provided by its comparand, assuming the comparand is true, *and* vice versa.

This procedure has been applied to the Kizlauskas and Katz

hypothesis (figure 3.13e) and the FWPCA hypothesis (figure 3.13b) with a speed multiplier of 0.5 to determine which release points would have been optimal for separating these hypotheses. The results are summarized in figure 3.15 as a contour plot over the southern Lake Michigan basin. As might have been expected, the southeastern part of the basin provides the richest territory for diagnostic release points. Unfortunately, no releases were made there.

In summary, one must be cautious when interpreting the results of the hypothesis testing scheme, since it is sensitive to locational biasing built into the experiment by the choice of and initial drifter concentration at the release sites. This bias may be used to advantage, though, when the resources for an experiment are limited, by restricting one's efforts to those diagnostic release points promising to distinguish most sharply between pairs of available hypotheses. By using the simulation scheme to generate "data" for each hypothesis over several releases and testing these data against the complementary hypothesis, a measure of distinguishability based on the compatibility indicator can be obtained for each release.

3.7 Testing Hypotheses about Wind-Driven Currents

Chapter 1 outlined a theory by Eckman (1905) relating water velocity to the velocity of the wind above the water. In this theory, the instantaneous water velocity vector $\underline{v}(q,t)$ is a linear function of the wind velocity vector $\underline{w}(q,t)$, viz:

$$\underline{v}(q,t) = \epsilon \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \underline{w}(q,t), \text{ where}$$

LAKE MICHIGAN




Figure 3.15: Contours of equal relative distinguishability for release points to distinguish hypothesis b from e.

e is the ratio of the current speed to the wind speed (usually a few percent), and

θ is the angle of the current to the right of the wind.

Modeling the transport of drifters under the influence of such a current is easily accomplished by expanding the state-space to include t and by redefining \underline{v} as a function of space and time. The new model, expressed in general time-varying terms, is:

$M_w = (\underline{Q}, T, \underline{v}, D, G, \underline{\delta})$, where

$\underline{Q} = \mathbb{R}^2$ is the position space of a drifter,

$T = \mathbb{R}$ is the time domain,

$\underline{Q} \times T$ is the state space,

$\underline{v}: \underline{Q} \times T \rightarrow \mathbb{R}^2$ is the instantaneous velocity field,

$D \in [0, \infty)$ is the dispersion coefficient,

$G: \underline{Q} \rightarrow \{\text{land, water}\}$ is the geographical descriptor function,

and,

$\underline{\delta}: \underline{Q} \times T \rightarrow \underline{Q} \times T$, the state transition function, defines the

dynamics of the system as follows:

If, for any $\underline{q} \in \underline{Q}$, $G(\underline{q}) = \text{land}$, then

$$\underline{\delta}(\underline{q}, t) = (\underline{q}, t), \text{ for all } t; \text{ otherwise}$$

$$\underline{\delta}(\underline{q}, t) = (\underline{q} + \underline{v}(\underline{q}, t)\Delta t + \underline{N}(0, 2D\Delta t), t + \Delta t), \text{ where}$$

Δt is the time interval associated with one discrete transition, and

$\underline{N}(\underline{\mu}, \sigma^2)$ is a pair of independent normal random variables, each with mean $\underline{\mu}$ and variance σ^2 .

The simulation of drifter transport based on this model is the same as that under the stationary model (section 3.2) except that the velocity vector associated with each small square can possibly change at each

step (ΔT) in the simulation. The array T (elapsed times for the drifters) plays more than a bookkeeping role now, as its elements must be combined with the corresponding positions of the drifters to form their states and thus determine the velocities they are subjected to. So, the changes required are an alteration in the description of \underline{V} to read:

"Let $NMAX = TMAX/\Delta T$, where TMAX is the latest time considered, and $(\underline{V}(i,j,k), i=1,\dots,64; j=1,\dots,64; k=0,\dots,NMAX-1)$ be an array containing the velocity vectors for each of the 4096 squares and NMAX time instances."

Furthermore, line 1.9 in the algorithm must be changed to read:

"1.9 $\underline{Q}(i) \leftarrow \underline{Q}(i) + \underline{V}(X,Y,\text{truncate}(T(i)/\Delta T)) \cdot TG + \underline{N}(2D \cdot TG)$
[Subtransition]"

It will be assumed that no simulation will be initialized with any T's less than zero or be allowed to run over more than NMAX iterations.

Hypothesis testing proceeds as before (section 3.4), with the exceptions that the hypotheses are of the form

$(\underline{U}(i,j,m,k), i=1,\dots,64; j=1,\dots,64; m=0,\dots,NMAX-1; k=1,\dots,NH)$,
an array of NH velocity hypotheses, each having the same form as \underline{V} ;

and line 1.1 of the algorithm is changed to read:

"1.1 $\underline{V}(i,j,m) \leftarrow \underline{U}(i,j,m,kh)$, $i=1,\dots,64; j=1,\dots,64; m=0,\dots,NMAX-1$ [Velocity field is hypothesis kh]"

With this new testing procedure, hypotheses about the effect of the wind on currents may be compared with themselves and with those assuming stationary currents. Assuming complete data on the wind field $\underline{w}(q,t)$ over the area studied and the experimental time period, hypotheses \underline{U} may be formed based on an assortment of values for ϵ and θ . In most

cases, though, such synoptic wind data are hard to come by for locations over the water surface, and one must rely on estimates gotten from land-based observations. To do this requires an assumption of a relation between lake winds and land winds. That such an assumption is not unjustified is shown by the FWPCA (1967) in comparing synoptic wind data from Chicago's Midway Airport and a buoy out in the lake. The six-hour-average observations over one month were roughly proportional, with the lake winds tending slightly faster than and to the left of the land winds. Though sufficient correlation for the purposes of this investigation may be acknowledged, it is felt that the data presented are insufficient to justify the utilization of any observed average deviations in the simulation. Therefore, none will be assumed.

In order to estimate the wind field over southern Lake Michigan during the latter half of 1974, hourly weather observations were subscribed to from the National Oceanic and Atmospheric Administration (NOAA) for Milwaukee, Muskegon, and Chicago (figure 3.16). Each hourly observation (figure 3.17) includes a wind measurement accurate to ten degrees and one knot. The hourly wind data for each of the three stations were key-punched to cover the period 15 July through 31 October, 1974--sufficiently long to encompass the returns from the July and August experiments. The hourly observations from each day, beginning and ending at noon, were then added vectorially to give the net wind velocity for that day. Figure 3.18 summarizes these daily wind data using progressive vector diagrams marked to indicate the various drifter releases. The diagrams reveal clear similarities in the winds for the three stations.

To obtain the net wind velocity on a given day for any point q

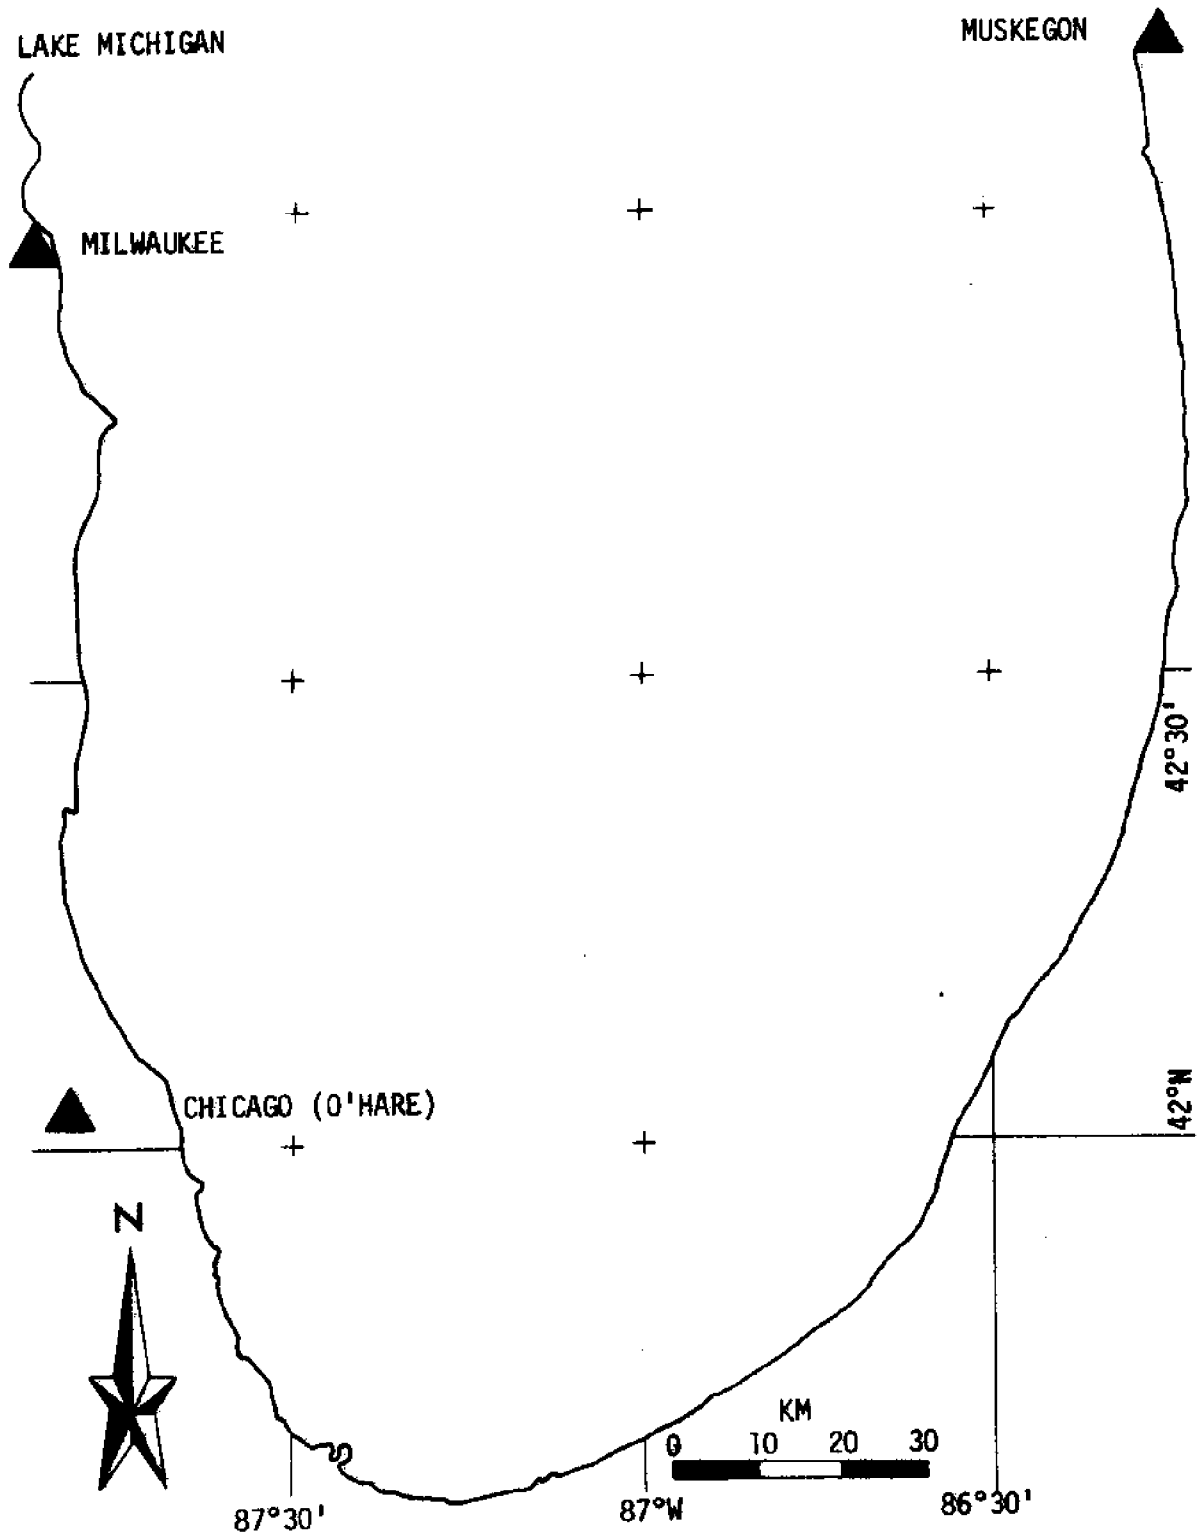


Figure 3.16: Wind observation stations (15 July - 31 October 1974).

STATION		DATE		TIME		REMARKS AND SUPPLEMENTAL CODED DATA		ALTIMETER SETTING		WIND		TEMPERATURE AND DEW POINT		SEA-LEVEL PRESSURE		VISIBILITY		WEATHER AND OBSTRUCTION TO VISION		SEA-LEVEL PRESSURE		TEMPERATURE AND DEW POINT		WIND		ALTIMETER SETTING		RUNWAY VISUAL RANGE		CODED PIREPS								
MTL-104 18-71		WFO		OCT 1 1974		ADD 5 Mz. SUBTRACT																																
U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE		U.S. DEPARTMENT OF COMMERCE NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION NATIONAL WEATHER SERVICE								
SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS		SURFACE WEATHER OBSERVATIONS								
TIME (LST)	TYPE	SKY AND CEILING (Hnds of Feet)	VISIBILITY (Miles)	WEATHER AND OBSTRUCTION TO VISION	SEA-LEVEL PRESSURE (Inches)	TEMPERATURE (F)	DEW POINT (F)	WIND (Mph)	WIND DIRECTION (Mph)	WIND VELOCITY (Mph)	WIND DIRECTION (Mph)	TEMPERATURE (F)	DEW POINT (F)	SEA-LEVEL PRESSURE (Inches)	TEMPERATURE (F)	DEW POINT (F)	WIND (Mph)	WIND DIRECTION (Mph)	WIND VELOCITY (Mph)	WIND DIRECTION (Mph)	TEMPERATURE (F)	DEW POINT (F)	SEA-LEVEL PRESSURE (Inches)	TEMPERATURE (F)	DEW POINT (F)	WIND (Mph)	WIND DIRECTION (Mph)	WIND VELOCITY (Mph)	WIND DIRECTION (Mph)	TEMPERATURE (F)	DEW POINT (F)	SEA-LEVEL PRESSURE (Inches)	TEMPERATURE (F)	DEW POINT (F)	WIND (Mph)	WIND DIRECTION (Mph)	WIND VELOCITY (Mph)	WIND DIRECTION (Mph)
0000	R	M37	27	20507	29.8	45	35	14	010	010	010	45	35	29.8	45	35	14	010	010	010	45	35	29.8	45	35	14	010	010	010	45	35	29.8	45	35	14	010	010	010
0100	R	M32	27	20507	29.8	44	35	13	010	010	010	44	35	29.8	44	35	13	010	010	010	44	35	29.8	44	35	13	010	010	010	44	35	29.8	44	35	13	010	010	010
0200	R	M32	27	20507	29.8	44	33	13	010	010	010	44	33	29.8	44	33	13	010	010	010	44	33	29.8	44	33	13	010	010	010	44	33	29.8	44	33	13	010	010	010
0300	R	M32	27	20507	29.8	44	33	13	010	010	010	44	33	29.8	44	33	13	010	010	010	44	33	29.8	44	33	13	010	010	010	44	33	29.8	44	33	13	010	010	010

KEY TO AVIATION WEATHER REPORTS

LOCATION IDENTIFIER AND TYPE OF REPORT: MKC 150M250

SKY AND CEILING: Sky cover symbols are in ascending order. Figures preceding symbols are heights in hundreds of feet above station. Sky cover symbols are: Clear: Less than 0.1 sky cover. Scattered: 0.1 to 0.5 sky cover. Broken: 0.5 to 0.9 sky cover. Overcast: More than 0.9 sky cover. Thin (When prefixed to the above symbols.) Partial obscuration: 0.1 to less than 1.0 sky hidden by precipitation or obscuration to vision (bases at surface). X Obscuration: 1.0 sky hidden by precipitation or obscuration to vision (bases at surface). Letter preceding height of layer identifies ceiling layer and indicates how ceiling height was obtained. Thus: A Aircraft R Radar B Balloon W Windshears C Ceiling pilot V Immediately following balloon E Estimated M Measured height indicates a variable ceiling.

WEATHER AND OBSTRUCTION TO VISION SYMBOLS: A Blowing dust B Blowing sand C Blowing snow D Dust F Fog G Ground fog H Haze I Ice crystals J Ice fog K Ice pellets L Ice other snows M Smoke N Rain O Rain showers P Rain squalls Q Rain showers with hail R Rain showers with hail and drizzle S Snow T Thunderstorms U Unsettled weather V Variable weather W Windshears X Xthunders

PRECIPITATION INTENSITIES: Light; Moderate; Heavy

WIND: Direction in tens of degrees from true north, speed in knots. 0000 indicates calm. G indicates gusty or peak speed of gusts follows G or Q when gusts or squalls are reported. The contraction WSHFT followed by GMT time group in remarks indicates windshift and its time of occurrence. (Knots X 1.15 = statute mi/hr.)

EXAMPLES: 3627=360 Degrees, 27 knots; 3627044=360 Degrees, 27 knots Peak speed in gusts 40 knots.

ALTIMETER SETTING: The first figure of the actual altimeter setting is always omitted from the report.

Figure 3.17: Sample hourly weather observations.

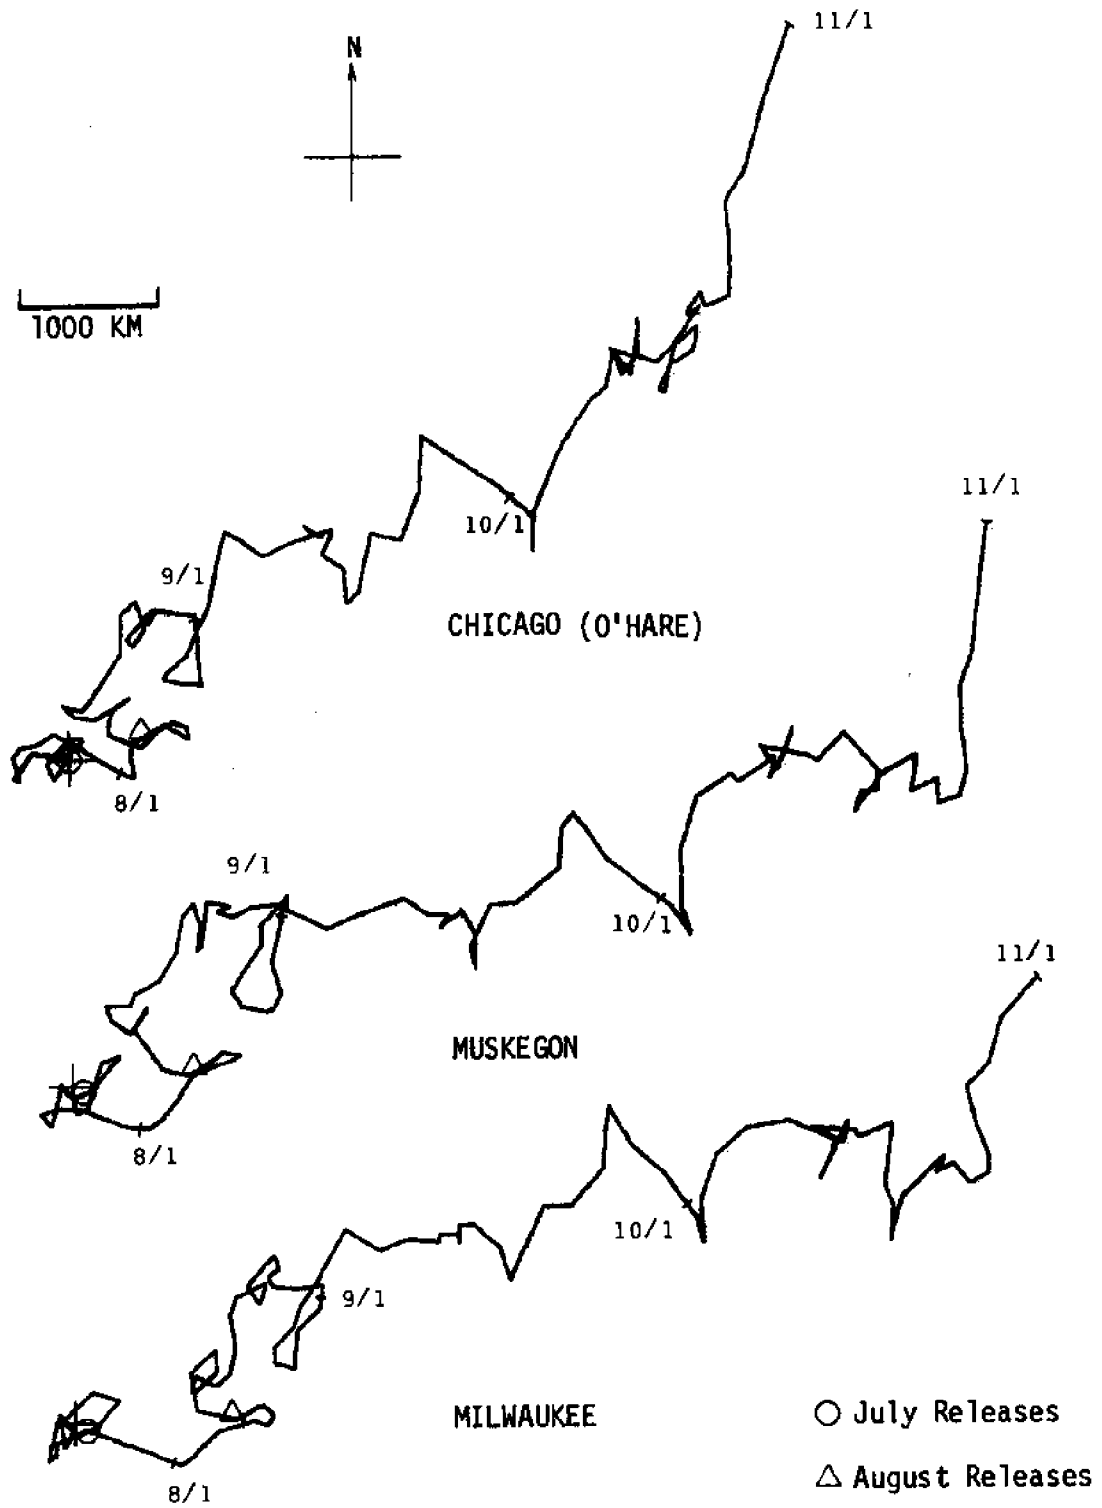


Figure 3.18: Wind tracks for three observation stations, 15 July - 31 October 1974.

over the lake, the observed wind velocities on that day at the three stations are averaged, each in proportion to the inverse of its squared distance to the point in question:

$$\underline{w}(\underline{q}, t) = \frac{\sum_{i=1}^3 \frac{w(\underline{r}_i, t)}{|\underline{q} - \underline{r}_i|^2}}{\sum_{i=1}^3 \frac{1}{|\underline{q} - \underline{r}_i|^2}}, \text{ where}$$

\underline{r}_i is the location of station i .

Nominally, these values are calculated for each square in the drifters' simulation domain, for each day to be simulated. The hypotheses \underline{U} are then determined by applying the Eckman transformation to each of these values for various values of ϵ and θ . Actually, though, for the sake of efficiency, none of these calculations are performed except for velocity values needed during the course of a simulation.

The July recovery data were subjected to hypothesis testing under the wind-driven simulation using values of ϵ ranging from .005 to .075 and values of θ ranging from 0° to 20° to the right of the wind. Again, the compatibility sum was divided by the number of comparisons to get a mean deviation rate. The results are shown in figure 3.19. In general, an ϵ of .03 seems to be preferred, along with the higher values of θ , but not overwhelmingly so.

In comparing these results with the stationary hypotheses, one finds strong competition from the wind hypotheses--at least on the surface. What the graphs of figure 3.19 do not reveal, however, is the tendency of the wind-generated currents to drive all the drifters in the simulation into the shore before the times corresponding to the later recoveries are reached. This means that no position comparisons can be made




Figure 3.19: Hypothesis testing results for wind hypotheses.

for these later recoveries. Consequently only the earlier recoveries contribute to the overall compatibility measure. Since one of the original assumptions of hypothesis testing was that valid hypotheses be able to transport drifters to the known recovery areas, one must temper his enthusiasm for the wind hypotheses with the knowledge that they aren't able to do this.

This tendency of the wind hypotheses to beach all the drifters prematurely raises some questions about the relation of Eckman's theory to drifter experiments. Drifter experiments always involve some near-shore currents, since most drifters are found awash. Though Eckman's theory could be valid on the open water of Lake Michigan, near-shore areas might provide other, more significant influences. Specifically, surface currents running perpendicular towards shore must either roll under and back or divert parallel to the shore. So even if the wind were responsible for some of a drifter's drift, its last days before washing ashore could be controlled by factors other than the Eckman currents simulated.

A possible counter-measure against this effect might be to concentrate the releases as far from shore as possible, rather than in the coastal areas concentrated on in July and August, 1974. This would give the Eckman currents, if any, a chance to rule the drifters' trajectories for most of their time afloat. As a result, the recovery data might better reflect the validity of Eckman's theory for Lake Michigan.

Other investigators who use drifters to study the effect of the wind on currents employ different tactics to account for factors other than the wind. Tomczak (1968), in a study of the North Sea, assumes a coastal current running parallel to the shore in the near-shore areas, which is superimposed on the Eckman currents. Using a drifter model

which neglects dispersion, he propagates a drifter from each release under the influence of various wind factors, stopping along the way at the various recovery times associated with the release or when the drifter reaches shore, whichever comes first. His fitness measure (see section 1.3) is applied to the recoveries at these times or, in the case of a shore intersection, for all subsequent recoveries. So long as his velocity assumptions in the coastal zones are correct, he should have a more accurate picture of the wind influence elsewhere than he would without these assumptions.

Hill and Horwood (1974) go a step further by assuming that the currents everywhere are the sum of an experimentally measured, underlying current and the wind current. By simulating the transport of a drifter cluster's center of gravity under the influence of such a current under various wind factors (see sections 1.3 and 3.1), they calculate the median of the return pattern for comparison to the experimental median. Again, this correction to the wind, if properly applied, should make the testing of Eckman's theory with drifter experiments more plausible.

Neither Tomczak nor Hill and Horwood, though, employ a realistic model of diffusion in their simulations. In cases where *just* the wind effect is considered and where the wind doesn't vary much as a function of position, this is not overly critical, since diffusion will be fairly regular anyway. But, by introducing stationary currents as well, they leave themselves open to the type of dispersion which is augmented by heterogeneities in the stationary velocity field, such as that shown in figures 3.5 and 3.6. Though neither does, it is important to account for this dispersion in any testing scheme which implicitly involves it.

This is particularly vital in the experiments recorded here, where some recovery patterns show a phenomenal spread.

It would be possible, of course, to combine the two testing schemes developed in this chapter to test hypotheses about the wind effect superposed on known, stationary currents. This would involve a prior knowledge of the stationary currents, which was not available for this investigation. Superposing Eckman currents on *hypothetical* stationary currents would be a little farfetched, though, since there would be too many unknowns being tested at once. Nevertheless, the principle of including a diffusion model like the one developed here in such a testing scheme can be considered a necessary condition for the validity of the test results.

3.8 Summary

The model of drifter transport described in Chapter 2 is used as the basis for a Monte Carlo simulation. This simulation forms the foundation of an hypothesis testing scheme which is used to measure the compatibility of each velocity field tested with the experimental recovery data. The measurement of compatibility is a function of how closely the hypothesized velocity field can bring individual drifters from their release site to known areas of recovery in the proper amounts of time. Since this criterion of closeness is not a function of relative numbers of drifters landing one place or another, redundant recovery data are not considered. Moreover, for each recovery, only the simulated drifter coming from the same release as the recovery which is closest to that recovery at the appropriate time is used in calculating compatibility. A further reduction in recovery data is effected by

eliminating late recoveries from a given area and a given release. This is done to moderate the uncertainties produced by drifters lying onshore undiscovered for awhile.

The application of the testing algorithm to recoveries obtained from Lake Michigan in 1974 reveals that, of the hypotheses available, the one (figure 3.13g) involving a clockwise gyre with a north-south splitting in the northeastern lower basin is the most compatible with the July data. Though the August data seem to be explainable by either a clockwise (figure 3.13e) or a counterclockwise (figure 3.13b) gyre, one is led to favor the clockwise hypothesis just on the basis of inertia. What is common to both hypotheses, though, is that northward coastal currents characterize the flow along both shorelines.

It is possible that the choice of release points can bias the results of the hypothesis testing scheme, since no counter-measures are employed in it. This bias can be used to advantage when limitations on the drifter coverage of a body of water must be observed. Release points may be chosen ahead of time to maximize the discrimination between two hypotheses by using the simulation to generate recoveries from various releases and the testing scheme to rate them under the alternate hypothesis. Diagnostic release points are chosen which reveal the greatest co-incompatibility between the two hypotheses.

Finally, the model and simulation can be converted to time-varying systems, allowing the testing of Eckman's theory of wind-driven currents under various parameters. Doing so for the July recovery data reveals that near-shore currents, at least, are not responsive to the wind in the Eckman sense. To test offshore currents for this with drifters would require either far offshore releases and the faith that any data

thus obtained would contain little evidence of other effects felt during their arrival ashore, or else prior knowledge of other aspects of the current field to be combined with the wind hypotheses.

CHAPTER 4

A GOAL DIRECTED DRIFTER SIMULATION AND HYPOTHESIS GENERATION

4.1 Overview

Hypothesis generation is the converse of hypothesis testing. In hypothesis testing, drifter movement is a direct consequence of an hypothesized current field \underline{v} , and the simulation outcome is some measurement of compatibility of that field with a body of known recovery data. In hypothesis generation, a current field is the direct consequence of simulated drifter movement, the movement being guided by a prior assumption of compatibility with the recovery data. Since compatibility is determined by how closely any drifter approaches its recovery point in the proper amount of time, beginning at its release point, its assumption as a guiding factor in drifter transport leads naturally to a goal-directed simulation.

In the goal-directed formulation of drifter transport, each drifter is assigned a target corresponding to a known recovery point and initialized at the corresponding release point. The basic force giving rise to a drifter's motion at any time step is an attraction between the drifter and its target. The intensity of this attraction will be assumed sufficient to guarantee arrival at its target in the proper time. In this way, compatibility with the data in the goal-directed context is assured.

During the course of a goal-directed simulation, each location in the drifters' domain may be visited by several drifters. In the

transport-diffusion model, the average velocity of the drifters at any location \underline{q} is specified by $\underline{v}(\underline{q})$. In hypothesis testing, $\underline{v}(\underline{q})$ is *defined* by the average velocity of the drifters passing through \underline{q} . If a large enough number of points \underline{q} are visited by drifters, then, a velocity field $\underline{v}(\underline{q})$ may be inferred from the drifter velocities. Presumably, then, such a velocity field could be used in the hypothesis testing scheme and show near perfect compatibility with the recovery data.

But such a presumption is premature. The transport-diffusion model on which the testing scheme is based is very specific about not just average drift velocities, but also their variances, which are controlled by the diffusion coefficient D . This implies that not only must $\underline{v}(\underline{q})$ be the average of the drifter velocities at \underline{q} , but that *each* of these velocities be close to $\underline{v}(\underline{q})$. As a result, during the course of the goal-directed inference process, not only must $\underline{v}(\underline{q})$ be a function of drifter velocities, but also these same velocities must depend on \underline{v} . Only if this mutual dependency promotes the satisfaction of the variance constraint can one expect the inferred \underline{v} to show compatibility under the testing scheme.

Even if such compatibility is assured, though, the resulting velocity field may have undesirable characteristics from a hydrodynamic point of view. Indeed, many velocity fields with many different characteristics may be equally compatible with the recovery data. So the inference process must be charged with selecting the field with the most desirable characteristics in addition to compatibility with the data. These characteristics will be determined by the set of trajectories followed by a group of goal-directed drifters; so it is the prior specification of such characteristics which should determine, at the base level,

the nature of the attractive force between a drifter and its target. The success of the inference process, therefore, depends on the ability to design a local transition function (attractive force) whose application in the inference scheme ultimately produces the kind of global velocity field characteristics desired.

The design approach taken here involves the use of heuristics, each one constructed to optimize one or several aspects of the resulting current field. By parameterizing or weighting their application in the transition function, one may control their relative importance in the overall inference process. Finally, the weights may be varied to get some idea of the effectiveness of each heuristic in optimizing its associated velocity field aspects.

Having thus designed an inference procedure, one may test it against "known" velocity fields. Such known fields may be merely hypothetical, of course, so long as a well-defined means for extracting recovery data from them exists. The transport-diffusion model provides just such a means and has been used in the development of the heuristics as a controlled source of data. Here it is used to demonstrate the effectiveness of the inference process evolved and to point out its deficiencies. The method is simple. Starting from a given velocity field, drifter experiments are simulated and their data collected. Using these data in the inference process results in another velocity field. Comparison of this field with the original gives an indication of how well the inference process works.

Finally, the inference scheme may be applied to recovery data from the field. This results in a velocity field which purports to be the actual field giving rise to the recoveries. Even though this actual

field is unknown, such a profession may be scrutinized by testing the inferred field against other hypotheses using the testing scheme. If it performs well in the context of reasonable, alternate hypotheses, the inference procedure may be judged effective.

4.2 Prior Considerations

The ultimate goal of the inference process is to generate hypotheses which will be judged compatible with the known recovery data by the testing process. Since the testing process is based on the transport-diffusion model, those assumptions germane to that model are of equal relevance here. Reiterating, the currents to be inferred are surface currents, which are solely responsible for the transport of the drifters. Until a drifter intersects shore, its velocity is the velocity of the water in which it's embedded. The water velocities are two-dimensional, time-invariant vector functions \underline{v} of the two-dimensional surface position \underline{q} , plus a random component characterized by the diffusion coefficient D . The random component yields dispersion of the drifters, such that the variance of a drifter cluster increases linearly with time, assuming a uniform velocity field. As a consequence, the trajectories of a set of drifters initialized at a point will emanate from that point and slowly spread apart until velocity heterogeneities begin to effect a more rapid relative dispersal. (See figure 4.1.)

Although drifters in the transport-diffusion framework are mutually independent in their motion, closely spaced drifters behave similarly due to their common dependence on the local water velocity. If one were to observe the behavior of such a system without knowing about water velocities, he might well conclude that information was being passed among the drifters in order to regulate their co-evolving trajec-




Figure 4.1: Trajectories of ten drifters showing gradual diffusion, then rapid dispersion due to the diverging field.

tories. This apparent interdependence in the transport-diffusion context becomes real in the hypothesis generation system, because the velocities which control the drifters are, at the same time, determined by the drifters. In this respect, the developing local velocities serve not only as records of the motion of drifters, but also as channels of communication among them. It is assumed, however, that in the hypothesis generation system these velocities will be the *only* media of interaction among drifters.

The state of a drifter in the basic transport-diffusion model is given by its position. Given the velocity field, this is sufficient for the transition function to determine its motion over time. As a result, the drifter goes wherever the velocity field (and chance) leads it. The state of a goal-directed drifter must depend on more than its position, however, because the velocity at that position is uncertain, and because it must reach its recovery point at a certain time. Therefore the state-space of such a drifter is augmented to include the amount of time lapsed since its release. By making the recovery point and time a parameter of the system, the space-time deviation from a drifter to its recovery can always be calculated and a plan for reducing it formulated. Multiple-drifter systems can be modeled, then, as a parallel composition of such single-drifter systems, all identical except for the values of the recovery point parameters.

The velocity field is not in the state space of the goal-directed drifter system. From this viewpoint, the inferential or adaptive parts of the overall system are made responsible for effecting *structural* changes in the drifter system. Such structural changes, of course, may be thought of as state changes in the context of the total hypothesis

generation scheme. (See figure 4.2.)

The transition rule for goal-directed drifters must take into account not only a drifter's state but also those local and global features of the developing velocity field which have a bearing on getting each drifter to its destination at the proper time, with an eye toward maintaining those characteristics of the velocity field considered desirable *a priori*. The driving force of the transition function is an attraction which seeks to draw a drifter ever closer to its goal. Quite often opposed to this force is the local velocity, which the drifter must try to obey, as well as global considerations which a circumspect drifter would do well to heed in order to minimize some measure of "effort" in reaching its target. It is this concept of effort which depends on the characteristics of the field deemed favorable; and it is the minimization of effort which ultimately yields an hypothesis showing parsimony in such a context.

It is now necessary to delineate what constitutes a desirable velocity field. In general, that velocity field is most desirable which contains the least kinetic energy and which requires the least force to maintain it. In the simplest terms, kinetic energy is manifest in the velocity of the water and is proportional to its square. Due to internal friction, a body of water not subjected to external forces will "run down," its energy being dispersed as heat. The more friction there is, the faster it will tend to run down and the greater the driving force which will be required to maintain it. Force is also required to do work on convecting (vertically moving) water parcels as well as to accelerate advecting (horizontally moving) water parcels. Though these latter forces may be present in an equilibrium state (i.e. they cancel), their




Figure 4.2: Scheme of inference system showing separation of goal-directed and adaptive components.

diminution is, nonetheless, desirable.

Friction arises when water elements collide. The frequency and intensity of these collisions increase as the relative velocities of neighboring elements increase. These relative velocities are particularly high in areas characterized by high velocity gradients or shears. In order to reduce friction, then, and consequently the requisite driving force, shears in a velocity field must be minimized. This is to say that a favorable velocity field entertains a high degree of smoothness.

Further characteristics may be derived from water's virtual incompressibility. In three dimensions, this implies that any given volume element of water will neither increase nor decrease in density, meaning that there will be no net inflow or outflow of fluid. As a result, sources and sinks in the three-dimensional velocity field are forbidden in a closed system. This does not rule out sources and sinks on the two-dimensional surface, however; but it does say that sources must be supplied from below, via upwelling, and that sinks must pass the converging surface waters to the depths. Upwelling entails a local increase in potential energy for the water parcels being raised and a net increase in potential energy if an equivalent *mass* of water does not displace them from above. Though an equal *volume* of water will always displace the convected parcels, its density due to temperature differences may be different from the water displaced, resulting in a net mass difference. Specifically, a thermally stratified body of water having cooler, denser water on the bottom will require a driving force to accomplish any vertical mixing because of the concomitant increase in potential energy. Therefore, in such a system, surface sources and sinks correlate directly with a necessary driving force. In the summer, Lake Michigan is such a

system (FWPCA, 1967; Monahan and Pilgrim, 1975), so a desirable surface velocity field in this body of water must display the weakest source/sink structure possible. This principle includes, of course, the shoreline, where a lack of surface sources and sinks implies flow parallel to the shore.

The above hydrodynamic considerations must be taken into account by the goal-directed drifter system in selecting routes from each drifter to its target. The decisions involved, at each time step, though, have to be made on the basis of an incompletely developed velocity field, so the chances of erring are initially rather large. As a result, it is quite likely that a complete run of the drifters from release to recovery will result in an extremely undesirable velocity field. But this resulting field will still contain information of benefit to the goal-directed system if allowed to begin anew at the release points. By an iterative process of rerunning the drifters, beginning with the velocity field evolved from the previous run, each route-establishing decision can be made on better and more complete information than before, yielding a better velocity field. Hopefully, the process will converge to a field having the desired features.

The inference system, then, contains a goal-directed drifter transport system whose structure it modifies. The state of a drifter in the goal-directed context is its position and its time since release. Its recovery position and time are available as parameters. The adaptive element of the inference system monitors the behavior of the goal-directed drifter system and adjusts the evolving velocity field accordingly. The goal-directed drifter system uses both local and global aspects of the velocity field in a heuristic fashion to cause transitions

which will yield a field satisfying certain criteria of desirability. These criteria include minimal velocities, accelerations, and shears, as well as the weakest source-sink structure possible. Finally, the velocity field derived from a given run of the system can be used as an initial condition for a subsequent run, and so forth. The net result of such iterating, if the heuristics are properly designed, should be a velocity field which has favorable hydrodynamic characteristics and is compatible with the data in the hypothesis testing context.

4.3 General Structure of the Inference Process

In line with the considerations discussed in the previous section, a skeletal framework upon which to hang the algorithmic embodiment of the inference process may now be formalized. As shown in figure 4.2, the process consists of a goal-directed simulation system controlled by an adaptive element which modifies its structure via changes in the velocity field. Though the velocity field be modified at every step in the travel of a drifter, it is most convenient to regard it as a parameter of the goal-directed simulation system which can be externally modified, rather than as a state. As a result, the formalism will be a mixture of automata theoretic and algorithmic concepts. In the larger sense, the velocity field is a state of the overall inference system and is even given that title in the definition of the adaptive element.

The goal-directed simulation system for one drifter may be defined by the following structure:

$$M_G = (\underline{Q}, T, \underline{q}_R, t_R, \underline{v}, G, \underline{\gamma}, \underline{\delta}, Y, \underline{\lambda}), \text{ where}$$

$\underline{Q} = \mathbb{R}^2$ is the location space of the drifter,

$T = \mathbb{R}$ is the domain of the time since release,

$\underline{Q} \times T$ is the state space of the drifter,

$q_R \in Q$ is the position of the recovery,

$t_R \in T$ is the time of the recovery,

$\underline{v}: Q \rightarrow \mathbb{R}^2$ is the velocity field,

$G: Q \rightarrow \{\text{land, water}\}$ is the geographical descriptor function,

$\underline{\gamma}: Q \times T \rightarrow \mathbb{R}^2 \times \mathbb{R}$ is a decision function which calculates what velocity the drifter at $q \in Q$ would "prefer" to take and how urgent it is that it take it,

$\underline{\delta}: Q \times T \rightarrow Q \times T$ is the state transition function defining the dynamics of the system,

$Y = \mathbb{R}^2 \times \mathbb{R} \times \mathbb{R}^2 \times \mathbb{R}$ is the output set, and

$\underline{\lambda}: Q \times T \rightarrow Y$, such that $\underline{\lambda}(q,t) = (q,t,\underline{\gamma}(q,t))$, is the output function.

Both $\underline{\gamma}$ and $\underline{\delta}$ are parameterized in \underline{v} , $\underline{\gamma}$ is additionally parameterized in q_R , t_R and G , and $\underline{\delta}$ is additionally parameterized in t_R and the valuation of $\underline{\gamma}$.

The decision function $\underline{\gamma}$ makes up the backbone of the system. In it are contained the heuristics which seek to optimize certain *global* aspects of a drifter's trajectory and hence of the evolving velocity field. In this formalism, its evaluation occurs before that of $\underline{\delta}$; the result of its evaluation remains as a temporary parameter until reevaluation. This permits the adaptive element, when present, to intervene between a decision and the associated transition with a change in \underline{v} . Such a policy allows the effect of a decision on the inference process to be felt without as much delay as would be experienced if it were carried out first.

The transition function $\underline{\delta}$ moves the drifter by a displacement

which is a compromise between the *local* velocity and its preferred velocity. This is in line with the consideration for the variance of the velocities of drifters passing through a certain area. $\underline{\delta}$ also decides how much drift time should elapse during the transition. This decision is founded in keeping the discrete displacements from transition to transition small. Both $\underline{\delta}$ and $\underline{\gamma}$ must be designed to guarantee that a drifter reaches \underline{q}_R at time t_R .

The action sequence of M_G may be defined as follows:

1. Let $\underline{\Gamma} = \underline{\gamma}(q, t)$.

If $t = t_R$, then

$$\underline{\gamma}(q, t) = (0, 0); \text{ otherwise}$$

$$\underline{\gamma}(q, t) = (\underline{u}, h), \text{ where}$$

\underline{u} is the desired velocity, and

h is the urgency measure.

2. Allow the adaptive element to make changes to \underline{v} , using the output $\underline{\lambda}$ as the basis of the changes.
3. Perform the transition $\underline{\delta}$. If $t = t_R$, then

$$\underline{\delta}(q, t) = (q, t); \text{ otherwise}$$

$$\underline{\delta}(q, t) = (\underline{q} + \underline{d}(\underline{v}(q), \underline{\Gamma}, t) \cdot e(\underline{v}(q), \underline{\Gamma}, t), t + e(\underline{v}(q), \underline{\Gamma}, t)),$$

where

\underline{d} is a compromise velocity between $\underline{v}(q)$ and the desired velocity, and

e is the amount of drift time allowed to elapse during the transition.

4. Go to 1.

Strictly speaking, the parameter space of $\underline{\Gamma}$ should be included in the state space of M_G , since it controls and is controlled by M_G . It is

called a temporary parameter instead, because of a desire to have the "state" of M_G correspond to the state of the drifter. Indeed, \underline{y} would not even need an element of memory were it not for the adaptive element, because there would be no changes in \underline{y} that its valuation would have to span.

The adaptive element may be defined as a simple system:

$M_A = (\underline{v}, B, X, \omega)$, where

$\underline{v}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is the same velocity field defined for M_G ,

B is a set of bookkeeping parameters,

$\underline{y} \times B$ is the state of the system,

$X = \mathbb{R}^2 \times \mathbb{R} \times \mathbb{R}^2 \times \mathbb{R}$ is the input to the system (output Y of M_G), and

$\omega: \underline{y} \times B \times X \rightarrow \underline{y} \times B$ is the transition function which accomplishes the modification of \underline{y} .

The interaction of the goal-directed simulation and the adaptive element is detailed in figure 4.3.

In practice, many drifters are simulated at once. A multiple-drifter goal-directed system may be specified by the parallel composition of several single-drifter systems, all sharing the same velocity field. This results in a vector of outputs $\underline{\lambda}$ to the adaptive element. The adaptive element may then consider these outputs as a *sequence* of inputs, updating \underline{y} and B after looking at each one in turn. ω should be designed so that the net change in \underline{y} is independent of the order in which the inputs are considered. When all the inputs have been processed, control returns to the goal-directed system.

The whole of the inference system described must be further embedded in an executive system which initializes the drifters, performs

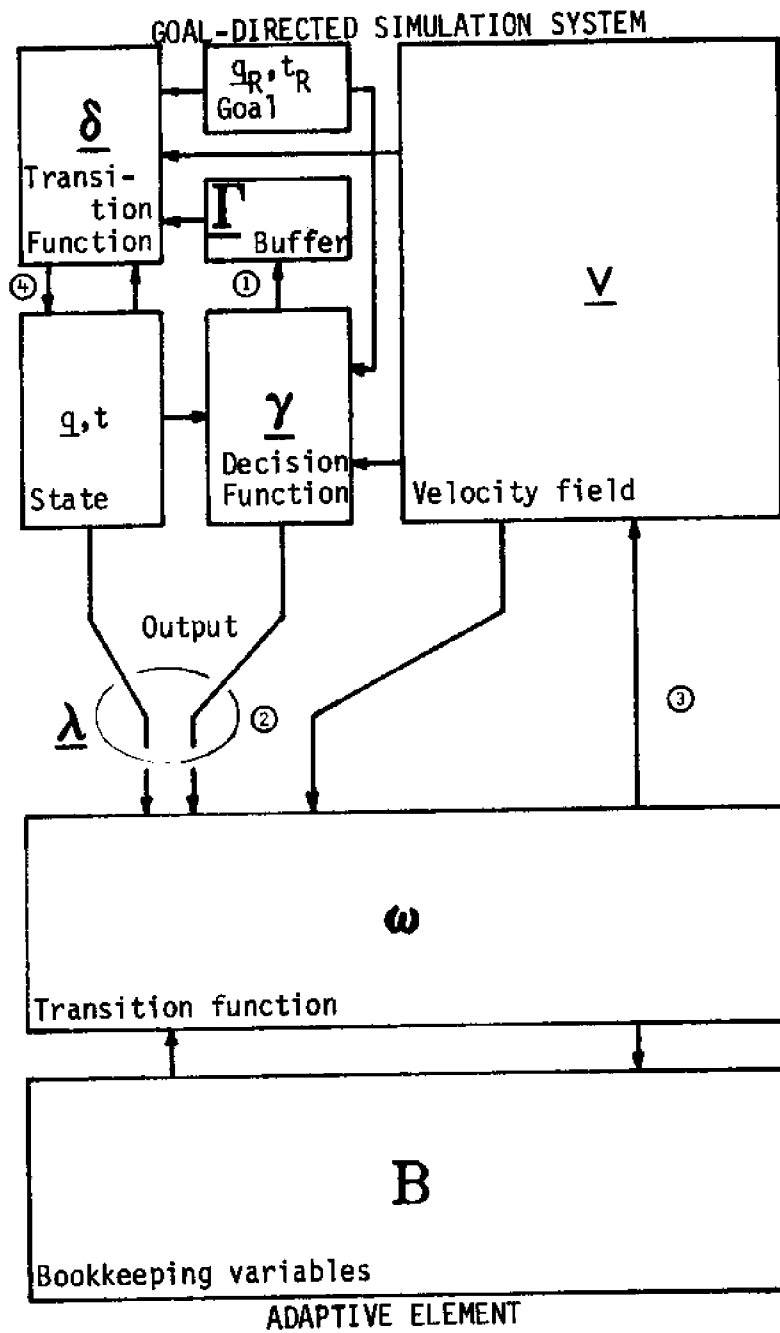


Figure 4.3: Detailed scheme of inference system.

the iterations, controls certain parameters of the inference system, accomplishes input and output, and so forth. Rather than formalizing this system algebraically, it will be developed as necessary in the algorithmic context.

4.4 Algorithmic Structure of the Inference Process

The algorithmic structure of the inference process is related to its general structure as the drifter simulation algorithm is related to the transport-diffusion model. The difference is that here the definitions of the key functions are delayed until they can be expressed in the algorithmic framework, whereas the drifter transport functions could be conveniently expressed algebraically. As before, parameters must be expressed as data structures, and functions must be transformed into procedures.

The inference program consists of four primary procedures corresponding to the four major elements of the general structure: 1) EXEC, the executive, which is responsible for setting up initial conditions, controlling global parameters, and iteratively invoking the other three procedures, which actually perform the drifter transitions and adaptation; 2) DECIDE, the decision procedure, which establishes for each drifter at any given time a preferred velocity based on the drifter's relation to its goal and on circumspection of the current velocity field; 3) ADAPT, the adaptation procedure, which makes local modifications in the velocity field to correspond to the local motion of drifters; and 4) MOVE, the drifter transition procedure, which moves the drifters in accord with their preferred velocities and with the local field velocity.

As in the simulation algorithm, the drifter domain is broken up into a 64×64 array of squares, each one containing two velocity

components valid over the entire square. They are the values of these velocity components which must be determined by the inference system. Most of the variables corresponding to these and other parameters largely parallel those of the simulation system and need only be listed here:

($Q(i)$, $i=1, \dots, K$), the position vectors for K goal-directed drifters,

($T(i)$, $i=1, \dots, K$), the elapsed times for the K drifters,

($U(i)$, $i=1, \dots, K$), an array of velocity vectors, each a preferred velocity for drifter i ,

($H(i)$, $i=1, \dots, K$), an array of urgency coefficients, each corresponding to a drifter's preferred velocity,

($QR(i)$, $i=1, \dots, K$), the position vectors for the K corresponding recovery points (targets),

($TR(i)$, $i=1, \dots, K$), the corresponding recovery times,

($Q0(i)$, $i=1, \dots, K$), the K release point position vectors (not necessarily all different),

($V(i,j)$, $i=1, \dots, 64$; $j=1, \dots, 64$), the array of velocity vectors,

B , a set of bookkeeping variables used to keep track of the state of the adaptation of V ,

($G(i,j)$, $i=1, \dots, 64$; $j=1, \dots, 64$), an array containing "land" or "water" for each square, and

S , the length along a side of each of the squares.

The core of the executive algorithm EXEC begins with a velocity field at some stage of development, and an iteration number ITER and continues as follows [comments occur in brackets]:

1. For each i , $i=1, \dots, K$ [For each drifter]

1.1 $Q(i) \leftarrow Q0(i)$ [Initialize drifter positions]

- 1.2 $T(i) \leftarrow 0$
- 1.3 next i
2. Until each $T(i) = TR(i)$, $i=1, \dots, K$, repeat
 - 2.1 call DECIDE [Get preferred velocities]
 - 2.2 call ADAPT [Alter \underline{V}]
 - 2.3 call MOVE [Move drifters]
3. Has convergence been reached?
 - Yes: clean up and stop
 - No: continue
4. $ITER \leftarrow ITER + 1$ [New iteration]
5. Modify global parameters as needed
6. Go to 1.

The decision procedures DECIDE begins with the drifters wherever they happen to be and calculates for each one its preferred velocity and urgency coefficient:

1. For each i , $i=1, \dots, K$ [For each drifter]
 - 1.1 $T(i) = TR(i)$?
 - Yes: 1.1.Y1 $\underline{U}(i) \leftarrow 0$
 - 1.1.Y2 $H(i) \leftarrow 0$
 - No: 1.1.N1 $\underline{U}(i) \leftarrow \underline{f}(\underline{V}, \underline{Q}(i), T(i), \underline{QR}(i), TR(i), \text{etc.})$
 [Calculate preferred velocity]
 - 1.1.N2 $w(i) \leftarrow \underline{g}(\underline{V}, \underline{Q}(i), T(i), \underline{QR}(i), TR(i), \text{etc.})$
 [Calculate urgency]
 - 1.2 next i
2. return to caller

The adaptation procedure ADAPT uses the preferred velocities and urgency factors defined by DECIDE to alter \underline{v} :

1. For each $i, i=1, \dots, K$ [For each drifter]
 - 1.1 $H(i) = 0?$ [Urgency nil?]
 - Yes: go to 1.7
 - No: continue
 - 1.2 $X \leftarrow \text{truncate}(\min(\max(Q_x(i)/S + 1, 1), 64))$
 - 1.3 $Y \leftarrow \text{truncate}(\min(\max(Q_y(i)/S + 1, 1), 64))$

[X and Y are the indices of the inbounds square nearest drifter i]
 - 1.4 $G(X, Y) = \text{land}?$
 - Yes: go to 1.7 [No velocity appropriate here]
 - No: continue
 - 1.5 $\underline{V}(X, Y) \leftarrow \underline{h}(\underline{V}(X, Y), B, \underline{U}(i), H(i))$ [Modify \underline{V}]
 - 1.6 $B \leftarrow b(B, \underline{V}, X, Y, \underline{U}(i), H(i))$ [Do bookkeeping]
 - 1.7 next i
2. return to caller

The transition algorithm MOVE moves the drifters according to their preferred velocity and the newly altered local velocity. The distance that each drifter moves is at most one square in order to guarantee that each square in its path is visited. The algorithm proceeds as follows:

1. For each $i, i=1, \dots, K$ [For each drifter]
 - 1.1 $T(i) = \text{TR}(i)?$ [Drifter done?]
 - Yes: go to 1.8
 - No: continue
 - 1.2 $X \leftarrow \text{truncate}(\min(\max(Q_x(i)/S + 1, 1), 64))$
 - 1.3 $Y \leftarrow \text{truncate}(\min(\max(Q_y(i)/S + 1, 1), 64))$

[Indices of inbounds square nearest $Q(i)$]

- 1.4 $\underline{D} \leftarrow \underline{d}(\underline{U}(i), \underline{V}(X,Y), H(i))$ [A compromise velocity]
 - 1.5 $\underline{E} \leftarrow \underline{e}(\underline{D}, T(i))$ [Time to elapse for this transition]
 - 1.6 $\underline{Q}(i) \leftarrow \underline{Q}(i) + \underline{D} \cdot \underline{E}$
 - 1.7 $T(i) \leftarrow T(i) + \underline{E}$
 - 1.8 next i
2. return to caller

Thus forms the algorithmic framework of the inference process. What remains is to fill out the missing definitions in the individual procedures.

4.5 Local Velocity Conditions

The average velocity of several drifters moving through a square must be the velocity of the square. The velocity of each should be close to the velocity of the square, in line with the assumed dispersion rate. Thus, trajectories beginning at a point will diverge slowly, and trajectories ending at a point must have converged slowly. This apparent symmetry could be used to advantage if it could be characterized mathematically.

In the transport-diffusion context, the density Ψ of a drifter's position \underline{q} , given its position \underline{q}_0 one step previously, is simply

$$\Psi(\underline{q} | \underline{q}_0; \underline{v}, \Delta t) = \frac{1}{4\pi D \Delta t} \exp[-(\underline{q}_0 + \underline{v} \Delta t) - \underline{q}^2 / (4D \Delta t)].$$

If the question is inverted to seek the density of the drifter's previous position \underline{q}_0 , given its current position \underline{q} , one can rely on Bayes' Theorem (Box and Tiao, 1973), namely:

$$\Psi(\underline{q}_0 | \underline{q}; \underline{v}, \Delta t) = \frac{\Psi(\underline{q} | \underline{q}_0; \underline{v}, \Delta t) p(\underline{q}_0)}{\int \Psi(\underline{q} | \underline{q}_0; \underline{v}, \Delta t) p(\underline{q}_0) d\underline{q}_{x,0} d\underline{q}_{y,0}}$$

$p(\underline{q}_0)$ is some prior distribution on the previous position, and it is assumed that Δt is sufficiently small that $\underline{v}(\underline{q}_0) \approx \underline{v}(\underline{q}) = \underline{v}$.

$p(q_0)$ reflects what is known *a priori* about the previous position of the drifter, with no regard for its current position. Except in near-shore areas, one may express his total ignorance of that location by assigning a density to q_0 which is uniform over a very broad area. Due to this uniformity, p may be taken out of the integral above and cancelled from the equation entirely. The integral then integrates to unity because of the interchangeability of q and q_0 in the definition of $\Psi(q_0 | q; \underline{v}, \Delta t)$ and the fact that it's a density function. This leaves

$$\begin{aligned} \Psi(q_0 | q; \underline{v}, \Delta t) &= \Psi(q | q_0; \underline{v}, \Delta t) \\ &= \frac{1}{4\pi D \Delta t} \exp[-|(q_0 + \underline{v} \Delta t) - q|^2 / (4D \Delta t)] \\ &= \frac{1}{4\pi D \Delta t} \exp[-|(q - \underline{v} \Delta t) - q_0|^2 / (4D \Delta t)]. \end{aligned}$$

This says that the uncertainty in the previous position, given the current position, is the same as that which one would get by starting in the current position, and running the system backwards (i.e. with $-\underline{v}$). Hence the symmetry apparent from considering trajectories is seen probabilistically as a reversibility in the system. Reversibility in this probabilistic sense implies that taking a body of drifter data and calling the releases recoveries and the recoveries releases should have no effect on the conclusions drawn therefrom except in the sign of the velocity field. (As an aside, it may be noted that the data reduction policy has virtually the same effect in either case since it operates only on drifters with approximately the same endpoints.) This principle can be used to advantage in the inference process by starting drifters at *both* ends of their trajectories and running to the opposite ends, making sure that backwards-moving drifters refer to and modify the velocity field with a change in sign. In this way, anomalies at the end of a trajectory

(e.g. large accelerations) caused by early misjudgments of the goal-directed heuristics will be moderated by the effects of the drifter coming the other way. Moreover, the coincidence of complementary trajectories may be taken as one sign of convergence after several iterations.

The preferred velocity of a drifter, calculated by DECIDE, could well be at odds with the field velocity in its square. In such a case as this, ADAPT will have to change the local velocity to agree more with the preferred velocity and MOVE will have to calculate a resultant velocity which is somewhere between the preferred velocity and the local one. All this has to be done while maintaining the principle that the velocity of a square is the average of all the velocities *moved at* by drifters passing through. Nonetheless, the derivation of the requisite functions is quite straightforward.

It is assumed that any preferred velocity \underline{u} calculated by DECIDE must, in some way, minimize the expected acceleration of drifter i and that any deviation from this velocity will result in increased acceleration. By multiplying the desired velocity by Δt , the amount of time remaining for the drifter to reach its recovery, one may obtain a displacement from the drifter's current position which might be thought of as corresponding to a "virtual target" (figure 4.4). The implicit assumption is that after the current time step, since a one-step transition is likely to be small, the location of the virtual target, calculated from the next-determined preferred velocity, will not change much. Hence, one can determine \underline{a} , the minimum acceleration required to correct for making a transition with velocity \underline{q}_i rather than \underline{u}_i as follows:

$$\text{displacement to virtual target} = \underline{u}_i \Delta t_i = \frac{1}{2} \underline{a} (\Delta t_i - e_i)^2 + \underline{q}_i \Delta t_i, \text{ or}$$

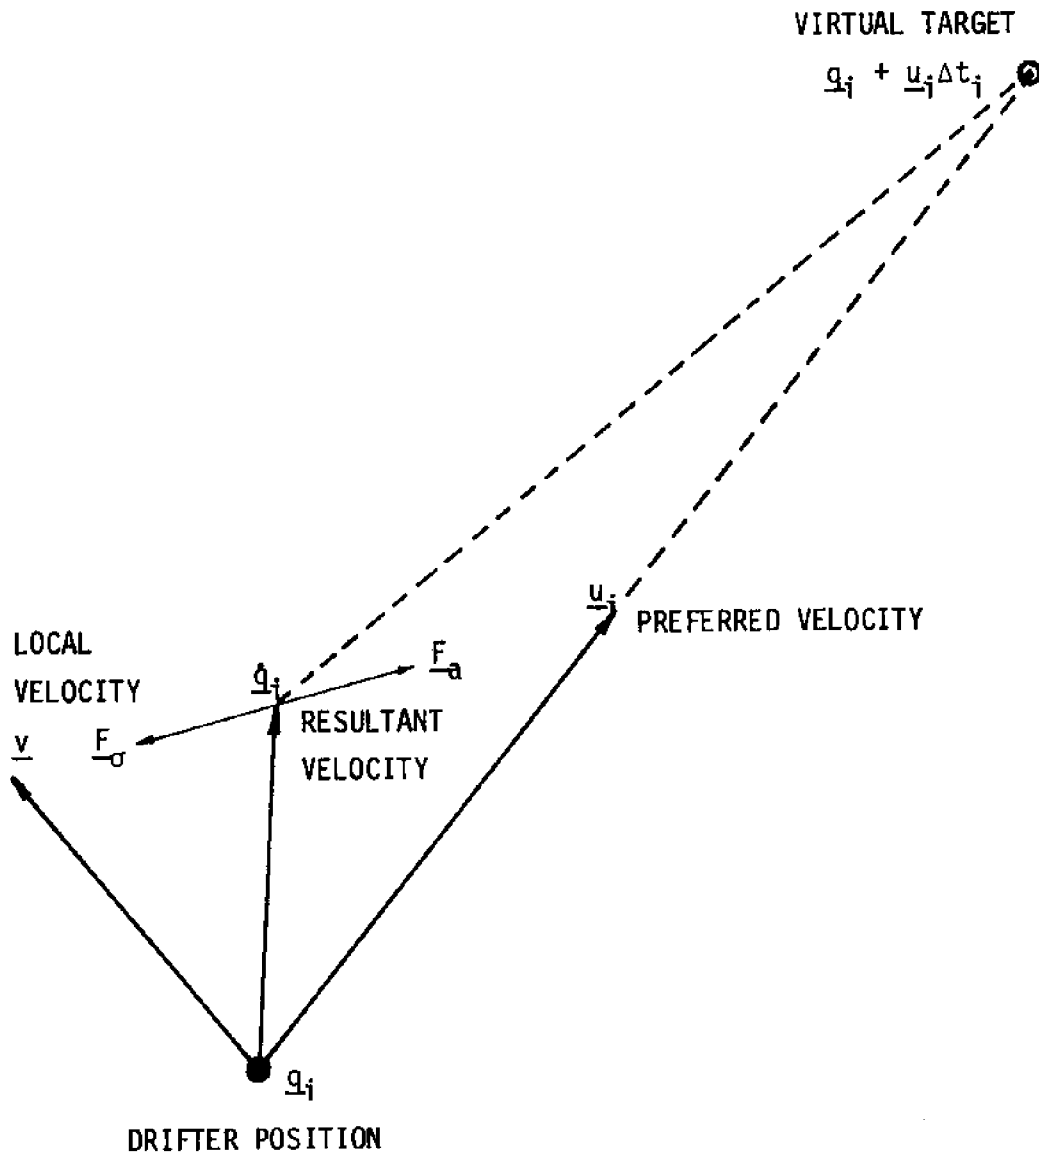


Figure 4.4: Local and preferred velocities as considerations in computing the drifter's velocity.

$$\underline{a} = \frac{2(\underline{u}_i - \dot{\underline{q}}_i) \Delta t_i}{(\Delta t_i - e_i)^2}, \text{ where}$$

e_i is the time elapsed for drifter i in this transition. This acceleration may be thought of as a "force" \underline{F}_a pulling $\dot{\underline{q}}_i$ around to \underline{u}_i , at which coincidence it is zero.

On the other hand, any deviation of the drifter from \underline{v} , the velocity of the square, can expect to meet a counterforce which seeks to keep the drifter in line with the local velocity. Since the aim of such a force is to limit the effective dispersion rate, it can be made an increasing function of that rate, such as

$$\underline{F}_\sigma = k(\underline{v} - \dot{\underline{q}}_i) \sqrt{e_i},$$

which measures the amount of separation of a drifter from its "mean" position $\underline{v}e_i$ after one transition, per radical unit time--the same units as \sqrt{D} , the square root of the diffusion coefficient. k is a parameter of the system and is used to adjust the relative importance of acceleration versus diffusion: a high value for k will tend to limit diffusion; a low value, permit it.

By balancing the two forces \underline{F}_a and \underline{F}_σ , one may obtain the resultant velocity $\dot{\underline{q}}_i$ as follows:

$$\underline{F}_a + \underline{F}_\sigma = \underline{0}, \text{ that is}$$

$$\frac{2(\underline{u}_i - \dot{\underline{q}}_i) \Delta t_i}{(\Delta t_i - e_i)^2} + k(\underline{v} - \dot{\underline{q}}_i) \sqrt{e_i} = \underline{0}, \text{ so}$$

$$k\underline{v} \sqrt{e_i} (\Delta t_i - e_i)^2 + 2\underline{u}_i \Delta t_i = \dot{\underline{q}}_i [k\sqrt{e_i} (\Delta t_i - e_i)^2 + 2\Delta t_i]; \text{ hence}$$

$$\dot{\underline{q}}_i = \frac{k\sqrt{e_i} (\Delta t_i - e_i)^2 \underline{v} + 2\Delta t_i \underline{u}_i}{k\sqrt{e_i} (\Delta t_i - e_i)^2 + 2\Delta t_i} = \frac{k\underline{v} + \underline{u}_i}{k + \underline{\mu}_i},$$

a weighted average of \underline{v} and \underline{u}_i , where

$$\mu_i = \frac{2\Delta t_i}{\sqrt{e_i(\Delta t_i - e_i)^2}}$$

This value for \dot{q}_i is what is calculated by the function \underline{d} in MOVE to determine the "compromise velocity." It uses an assumed value for e_i , the elapsed time of the transition, which, in MOVE, is a function of \underline{d} . It is fairly safe to assume that e_i changes little from transition to transition, so the previous value calculated may be saved to calculate \underline{d} , so long as it is less than Δt_i ; otherwise Δt_i is used. Also entering into the calculation of \underline{d} are the appropriate sign changes for those drifters moving backwards in time.

Since the velocity of the square must be the average velocity of some n drifters moving through it at various times, it can be reckoned as follows:

$$\underline{v} = \frac{1}{n} \sum_{i=1}^n \dot{q}_i = \frac{1}{n} \sum_{i=1}^n \frac{k\underline{v} + \mu_i \underline{u}_i}{k + \mu_i}, \text{ or}$$

$$n\underline{v} = \underline{v} \sum_{i=1}^n \frac{k}{k + \mu_i} + \sum_{i=1}^n \frac{\mu_i \underline{u}_i}{k + \mu_i}, \text{ so}$$

$$\underline{v} = \frac{\sum_{i=1}^n \frac{\mu_i \underline{u}_i}{k + \mu_i}}{1 - \sum_{i=1}^n \frac{k}{k + \mu_i}} = \frac{\sum_{i=1}^n \frac{\mu_i \underline{u}_i}{k + \mu_i}}{\sum_{i=1}^n \frac{\mu_i}{k + \mu_i}},$$

which is a weighted average of the \underline{u}_i .

So, the average of the actual velocities is a weighted average of the desired velocities. The factor μ_i can be considered the urgency coefficient alluded to earlier, and is balanced against k , the anti-dispersion coefficient. The factor $\mu_i/(k + \mu_i)$, which can be thought of

as "relative urgency", is plotted in figure 4.5 for various k , assuming an e_i of one. As can be seen, when time begins to run short, the relative urgency of the situation (i.e. the need to get to the desired location) increases. This increase occurs later and later as k gets larger. The same principle also applies in the formula for \hat{q}_i , where, if time is expected to run out in a current iteration (i.e. if $t_i - e_i = 0$), the weight given to \underline{u}_i is infinite, and no heed is paid to the local velocity.

Since the local field velocity in each square is simply a weighted average of the desired velocities, the adaptive element need only keep an overall weight for each square as its bookkeeping. Therefore, for any given square, each time a new drifter visits it, its desired velocity, weighted by the relative urgency, will be averaged with the local velocity, weighted by the local weight, to form the new local velocity. The new local weight is then the sum of the old one and the relative urgency factor. In this manner, neither the order in which drifters visit a square nor the order in which ADAPT considers them has any bearing on the final result. The order in which drifters visit a square does influence the behavior of intermediately visiting drifters, however, since they do not initially receive the effect of future visitations. But after several iterations, assuming the velocity field and weights carry over between iterations, the effect of these late visitations from previous iterations will be felt by early visitors on the current one; after many iterations, assuming convergence, this effect will be identical to what would be expected from the current iteration.

A final local consideration is the gradient of the velocity field. It might be possible for neighboring squares to contain highly different velocities. To moderate the effect of such a gradient on the motion of




Figure 4.5: Preferred velocity weighting as a function of remaining time and k , the anti-diffusion coefficient.

drifters, the resultant velocity \underline{q}_i of a drifter is made to depend, not only on the velocity of the square it's in, but also on the velocities of the eight adjacent squares. Therefore, \underline{v} in the formula for \underline{q}_i is taken to mean an average over nine squares of their velocities, weighted inversely as the squared distance of the drifter to the center of each square and proportionally to the weight of each square, *viz.*:

$$\underline{v} = \frac{\sum_{j=-1}^1 \sum_{k=-1}^1 \left[\frac{W(X+j, Y+k)}{S^2 + |\underline{q}_i - P_{X+j, Y+k}|^2} \underline{V}(X+j, Y+k) \right]}{\sum_{j=-1}^1 \sum_{k=-1}^1 \frac{W(X+j, Y+k)}{S^2 + |\underline{q}_i - P_{X+j, Y+k}|^2}}, \text{ where}$$

\underline{q}_i is the position of drifter i ,

X, Y are the indices of the square corresponding to \underline{q}_i ,

$\underline{V}(a,b)$ is the velocity of square a,b ,

$W(a,b)$ is the weight value of square a,b ,

S is the length along a side of each square, and

$P_{a,b}$ is the location of the center of square a,b .

This averaging not only reduces the effect of gradients, but also erases the arbitrariness of the boundary locations for the squares. If a drifter lands in a previously unvisited square close to its boundary with a heavily visited square, the information in that neighboring square will play a key role in the movement of the drifter. This resultant movement will be little different from what it might have been had the squares been shifted slightly to include the drifter in the heavily visited square.

From local considerations, three principles have emerged. First, the transport-diffusion system is probabilistically reversible, implying that whatever forces are designed to guide the travel of a goal-directed drifter are equally applicable to a complementary drifter running the

opposite direction in time, so long as appropriate sign changes are accounted for. Secondly, in order to make the velocity of a square the average of the velocities *taken* by visiting drifters, this velocity may be computed as a weighted average of their *preferred* velocities. The weighting scheme further includes a factor for controlling the amount of dispersion (e.g. velocity variance) allowed in a given square. Finally, the effect of velocity gradients and boundary location can be reduced by using an average velocity computed over several squares in calculating the compromise between the local velocity and the preferred one. Despite the fact that the drifters will, by and large, be traveling their whole routes with compromise velocities, DECIDE must come up with preferred velocities which eventually herd them to their respective targets.

4.6 Hitting the Target--The Direct Approach

The primary task of DECIDE is to ensure the arrival of each simulated drifter at its appointed recovery site in the proper time. The simplest way to do this is to redirect the drifter straight towards its target at each transition step and give it a velocity equal to the target displacement divided by the remaining time. Although the early parts of its actual trajectory might be dominated by compromise velocities, the urgency factor will sooner or later become high enough that such an appropriately directed drifter will find its way home.

This straight line heuristic has the effect of minimizing acceleration, so long as the mean field velocity doesn't carry the drifter too far afield, simply because it always advocates the one unaccelerated path. When the overall mean field is given more attention, though, the action of this heuristic begins to look very myopic, since a great deal of overall acceleration might be eliminated by compromising the direct



approach in favor of a long range plan. Nevertheless, this short range planner will always have a role to play in the overall scheme whenever a drifter, nearing the end of its allotted flotation time, needs a sure direction to its goal.

The interaction between this first goal-oriented heuristic and the adaptive element can be demonstrated by considering the hypothetical case of two releases at the same point and associated recoveries at the same distance and time from the release, but separated by some finite distance (figure 4.6a). Drifters running both directions in time are released (one from each recovery point, and two at the release point) in a velocity field having initially zero velocities and weights everywhere. The trajectories of the four drifters are shown for six iterations in figures 4.6b through 4.6g.* As can be seen, accelerations are gradually reduced, yielding fairly smooth trajectories. Figure 4.6h shows the final velocity field.





The effect of k , the anti-diffusion term, can be judged by adjusting it over a wide range. Figures 4.7a - c show the final trajectories and resultant fields of the same drifters run with values of k standing in various ratios with that used for figure 4.6. As the diagrams show, higher values result in a later "breakaway" from the common track. This illustrates an important tradeoff in the inference system: that of diffusion versus divergence. Any pair of recoveries from the same release occurring at the same time but at different points represents a certain amount of dispersion. Such dispersion might be largely due to local dif-

*The action of the adaptive element in this and subsequent examples differs slightly from that so far presented, in line with the adaptive strategy outlined in section 4.9. But the difference is insignificant for the purposes of the examples.

a) Data points
○ ○
RECOVERIES (T=20)



○
RELEASE (T=0)



DIVERGENCE

DIFFUSION




Figure 4.6: Trajectories from the straight-line heuristic. k=1




Figure 4.7: Final results from straight-line heuristic for various k .

fusion processes, or it could be a result of a forking in the velocity field. Where one of these fails to account for it the other must. Therefore when the anti-diffusion term is turned on high, the bulk of the dispersion happens in a late but severe divergence of the velocity field directions.

Now consider the case of two drifters released at the same point and recovered at different points and times (see figure 4.8a). What is implied here? First of all, it might be a result of simple dispersion as before. But dispersion is not a necessary implication since the later-recovered drifter may have stuck with the early one all the way to its recovery point and then proceeded to its own recovery point in the remaining time. Applying the straight-line heuristic to this situation for several iterations yields the trajectories shown in figures 4.8b - e. Notice that the release velocity is always somewhere between the straight lines to the two targets, so dispersion takes place.

This is the most basic manifestation of the straight-line heuristic's myopia: any angular displacement of the recoveries from a common release point virtually requires dispersion of some sort. It is easy to see, though, that this is unnecessary. In the previous example, for instance, trajectories lying on a common circle satisfy the recovery times with no diffusion and minimal acceleration. But the velocity at the release does not lie in the angle subtending the two recoveries. It requires much more global assistance to discover such a path than that provided by the straight-line heuristic.

It is here, however, that most inference efforts reported in the literature stop, and it is this framework in which they can be compared. The primary methods involved are the straight-line schemes (Wyatt, *et al.*,

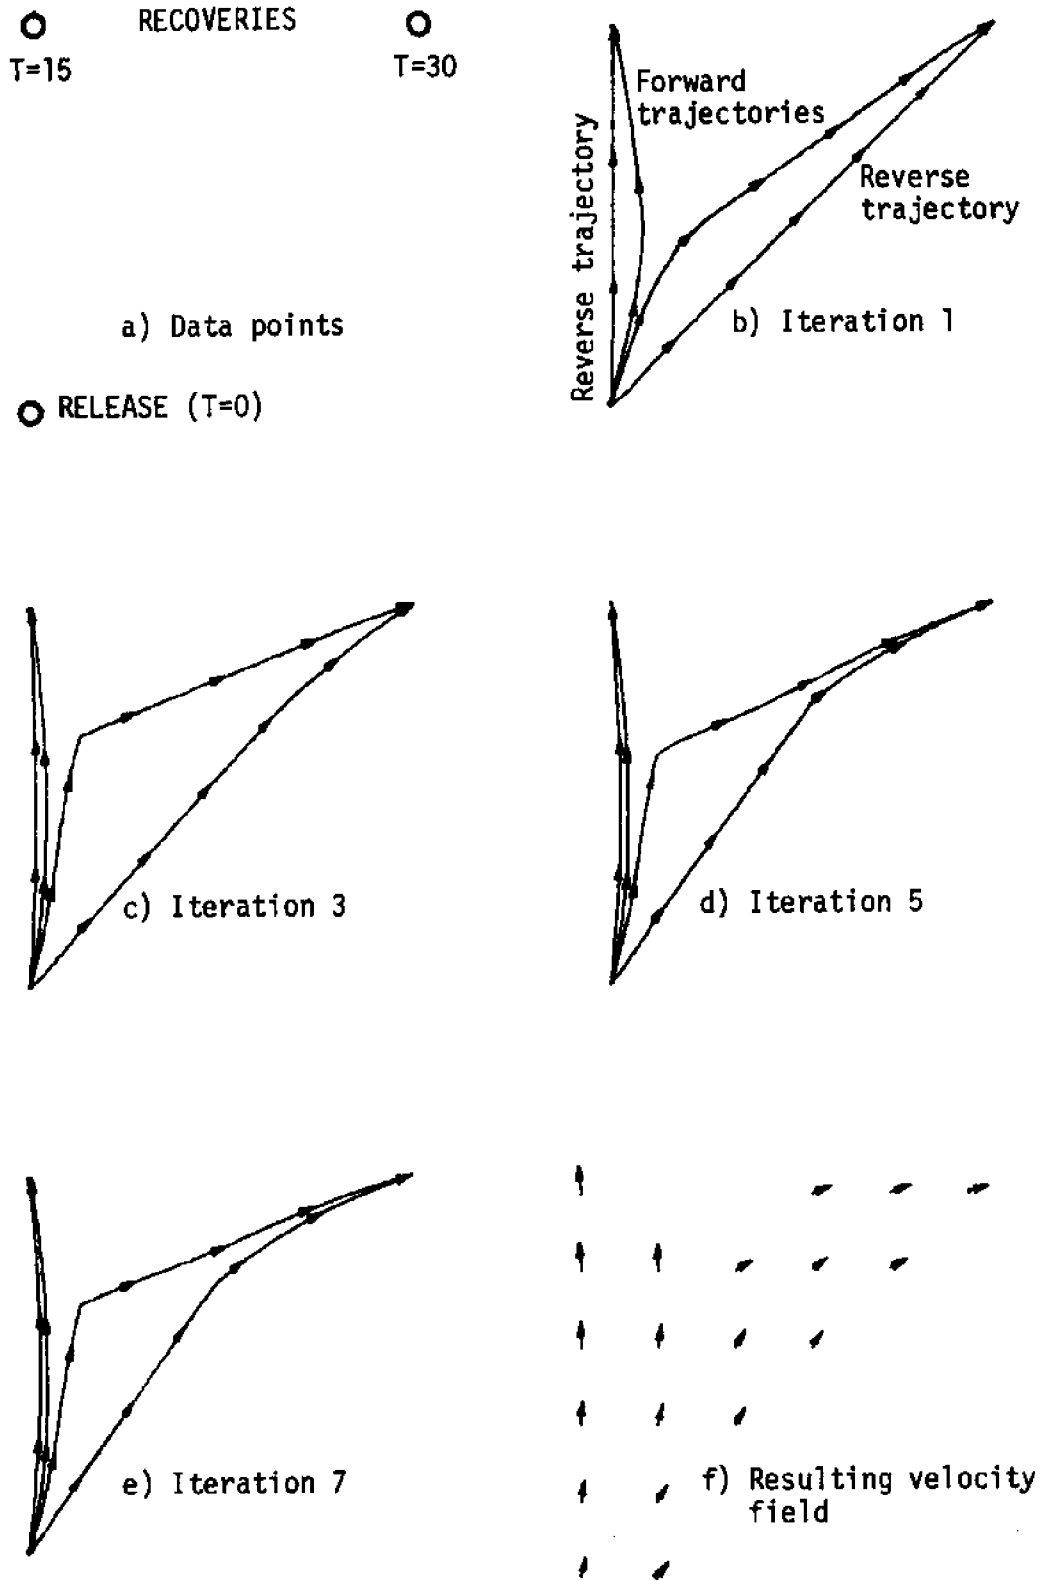


Figure 4.8: The straight-line heuristic used on circulatory data.

1972; Norcross and Stanley, 1967) and the reverse trajectory construction schemes (Bukin, 1974; Bumpus and Lauzier, 1965). These two procedures stand at opposite ends of the diffusion spectrum. The straight line method, in most variants, assumes that each drifter follows the most direct path to its recovery. The field velocity, inferred at each release point is just some average of the implied initial velocities of the drifters released there. In a sense, this corresponds to a k value of zero in the goal-directed system. The correspondence is not strict though since intermediate points of one trajectory may lie near the release point of another and contribute to its velocity there. But the principle is the same.

Reverse trajectory construction, on the other hand, is founded in zero dispersion (actually zero confluence) or an extremely high k value. If one were to consider only the drifters going backward in the goal-directed framework, drifters recovered at a common location would propagate first to the most recent release, then the remainder to the next most recent, and so forth, which is the basis of the reverse trajectory scheme. Again the correspondence is not perfect because of the additional forward-moving drifters used here. Also, neither Bukin nor Bumpus and Lauzier advocate the strict application of their methods to the point of absurd accelerations, as would happen here, but their corrective measures are less well specified than might be desired.

So the straight-line heuristic, though guaranteeing timely target arrivals, falls short of providing the "best" trajectories in every case. Consequently, a more circumspect heuristic must be developed to augment the direct approach.

4.7 Hitting the Target--The Circumspect Approach

There is but one straight line connecting two points; there are myriad arcs joining them. While the direct approach to hitting the target has but one choice for direction, an approach involving circular paths has many, allowing the choice of a path to be governed by external circumstances. So here it is assumed, barring diffusion constraints, that the path from a drifter to its target is the arc of some circle upon which the drifter and target both lie and that its speed is the length of the arc divided by the remaining time. The straight line, of course, is a degenerate case.

Given two distinct points, a vector from one of them determines a unique circle to which it is tangent intersecting the two points. So the problem of selecting an arc to the target is equivalent to the problem of selecting a direction to move from the drifter's current position. What determines this direction should be some function of global velocity field conditions which strives to eliminate from the velocity field certain hydrodynamic features considered undesirable. The circulation heuristic to be described is an attempt to minimize both divergence and velocity gradients in the evolving field. It does this by choosing a path (an arc) which integrates well into the extant field without implying undue acceleration. But before going into this, a lesson in field theory is useful.

The minimization of divergence is the same as the reduction of sources and sinks in a field. Divergence (source density) is defined in two dimensions as

$$\underline{\nabla} \cdot \underline{v} = \frac{\partial v_x}{\partial x} + \frac{\partial v_y}{\partial y} .$$

A field such that $\underline{\nabla} \cdot \underline{v} = 0$ everywhere is said to be *solenoidal*. Any two-

dimensional solenoidal field can be defined in terms of a scalar potential function σ called the *stream function* (Neumann and Pierson, 1966), such that

$$\underline{v} = \frac{\partial \sigma}{\partial y} \hat{x} - \frac{\partial \sigma}{\partial x} \hat{y}.$$

These velocities lie along the contours of constant σ (*streamlines*) and increase as the gradient of σ increases. The key feature of such a field is circulation. Assuming that the velocity vanishes at infinity, every streamline will be a closed loop. Each loop represents a circulation path around either a hill or a valley in σ .

Slight, local deformations in σ will alter the local velocity pattern. Since \underline{v} is linear in σ , one may imagine that superposing a little bump on σ somewhere will result in the original \underline{v} with a small circulation loop superposed at and going around the location of the bump. Indeed one may define σ in terms of "little bumps" (or dips) added to an evenly sloping function, where a "bump density" can be given by the local, two-dimensional deformation rate (the Laplacian):

$$\nabla^2 \sigma = \frac{\partial^2 \sigma}{\partial x^2} + \frac{\partial^2 \sigma}{\partial y^2}.$$

Corresponding to the "bump density" in σ is a circulation density in \underline{v} called the *curl* (or *vorticity*; von Arx, 1962), where

$$\text{curl}(\underline{v}) = \underline{\nabla} \times \underline{v} = -\nabla^2 \sigma, \text{ or}$$

$$\underline{\nabla} \times \underline{v} = \frac{\partial v_y}{\partial x} - \frac{\partial v_x}{\partial y}.$$

Any field in which $\underline{\nabla} \times \underline{v} = 0$ everywhere is said to be *irrotational*. A solenoidal field in which this is true will be bereft of circulation and hence null if it must vanish at infinity.

By a theorem due to Helmholtz (Arfken, 1970), every velocity field

\underline{v} with curl and divergence vanishing at infinity can be represented as the superposition of a solenoidal and an irrotational part. The solenoidal part is given by the stream function σ defined in terms of the curl as

$$\sigma(\underline{q}) = \frac{1}{2\pi} \iint \frac{\underline{v} \times \underline{v}(\underline{r})}{|\underline{q} - \underline{r}|} d\mathbf{r}_x d\mathbf{r}_y$$

Therefore,

$$\underline{v}' = \frac{\partial \sigma}{\partial y} \hat{x} - \frac{\partial \sigma}{\partial x} \hat{y}$$

may be thought of as \underline{v} stripped of its sources and sinks.

In the drifter context, if a transition can be made which results in moving the existing \underline{v} closer to its associated \underline{v}' and still bring the goal within easier reach, then a decision policy which promotes such transitions should diminish the overall divergence. The problem, of course, is that the direction of the associated solenoidal field at a drifter location may result in excessive acceleration if that direction is followed. In other words, one can't neglect getting the drifters to their goals in favor of fixing up the field. A more fundamental problem exists in finding \underline{v}' to begin with, since \underline{v} might not even be defined everywhere (i.e. some local weights may be zero). So a way must be found which takes partially defined circulatory aspects of an evolving field into account when planning a reasonable path from a drifter to its target.

If one were to attempt an evaluation of the integral for $\sigma(\underline{q})$ in the discrete framework used here, all he would need to do is find the circulation density c_{ij} for each square and compute as follows:

$$\sigma_{mn} = \frac{1}{2\pi} \sum_{i=1}^{64} \sum_{j=1}^{64} \frac{c_{ij} S^2}{r_{ij,mn}}, \text{ where}$$

$$r_{ij,mn} = S\sqrt{(i-m)^2 + (j-n)^2}.$$

The circulation density c_{ij} of a square can be thought of as the mean density over the square in the continuous case:

$$\begin{aligned} c_{ij} &= \frac{1}{S^2} \int_{\text{square}} \underline{\nabla} \times \underline{v} \, da \\ &= \frac{1}{S^2} \int_{\text{perimeter}} \underline{v} \cdot d\underline{\ell} \quad , \text{ by Stokes' Theorem.} \end{aligned}$$

Finding \underline{v} along the boundary area between squares can be done by averaging the velocities of the two squares meeting there. Only the component along the boundary need be considered. Therefore c_{ij} in any square is

$$c_{ij} = \frac{1}{S} (a + b + c + d) \quad (\text{see figure 4.9}), \text{ where}$$

$$a = \frac{-1}{2} (v_{x,ij} + v_{x,i,j+1})$$

$$b = \frac{1}{2} (v_{y,ij} + v_{y,i+1,j})$$

$$c = \frac{1}{2} (v_{x,ij} + v_{x,i,j-1})$$

$$d = \frac{-1}{2} (v_{y,ij} + v_{y,i-1,j}).$$

As can be seen, v_{ij} cancels in the calculation of c_{ij} , and only neighboring velocities are used.

It is useful to consider the effect of one component, $v_{y,ij}$ say, in the overall picture. Assume that this is the only non-zero velocity in the whole space (so call it v_y); what is the field \underline{v}' around it?

$v_{y,ij}$ contributes only to $c_{i-1,j}$ and $c_{i+1,j}$. Hence $c_{i-1,j} = \frac{1}{2S} v_y$ and $c_{i+1,j} = \frac{-1}{2S} v_y$. $\sigma(\underline{q})$ for some distant \underline{q} is then

$$\sigma(\underline{q}) = \frac{1}{2\pi} \cdot \frac{1}{2S} \left[\frac{v_y}{r_1} + \frac{v_y}{r_2} \right], \text{ where}$$




Figure 4.9: The velocities used for computing the curl of a square.




Figure 4.10: Parameters used in computing the solenoidal field around a single current dipole.

r_1 is the distance from point $i-1, j$ to \underline{q} , and

r_2 is the distance from point $i+1, j$ to \underline{q} (see figure 4.10).

By the law of cosines,

$$r_1 = [r^2 + S^2 + rS\cos\theta]^{\frac{1}{2}}, \text{ and}$$

$$r_2 = [r^2 + S^2 - rS\cos\theta]^{\frac{1}{2}}.$$

In the far field, where $r \gg S$ (or $\frac{S^2}{r^2} \rightarrow 0$),

$$r_1 \approx r(1 + \frac{S}{2r} \cos\theta), \text{ and}$$

$$r_2 \approx r(1 - \frac{S}{2r} \cos\theta).$$

So

$$\begin{aligned} \sigma(\underline{q}) &\approx \frac{1}{2\pi} \frac{v_y}{2S} \left[\frac{1}{r(1 + S\cos\theta/(2r))} - \frac{1}{r(1 - S\cos\theta/(2r))} \right] \\ &= \frac{1}{2\pi} \cdot \frac{v_y}{2r^2S} \frac{-S\cos\theta}{1 - S^2\cos^2\theta/(4r^2)} \\ &= \frac{1}{2\pi} \cdot \frac{-v_y \cos\theta}{2r^2}. \end{aligned}$$

From this, $v'(\underline{q})$ can be calculated as

$$\begin{aligned} v'(\underline{q}) &= \frac{\partial\sigma(\underline{q})}{\partial y} \hat{x} - \frac{\partial\sigma(\underline{q})}{\partial x} \hat{y} \\ &= \frac{-1}{r} \frac{\partial\sigma(\underline{q})}{\partial\theta} \hat{r} + \frac{\partial\sigma(\underline{q})}{\partial r} \hat{\theta}, \text{ in polar coordinates;} \\ &= \frac{v_y}{4\pi r^3} [\hat{r}\sin\theta - \hat{\theta}\cos\theta]. \end{aligned}$$

\underline{v} ($= v_y$) in the same coordinates is $v_y[\hat{r}\sin\theta + \hat{\theta}\cos\theta]$, so the component of \underline{v}' along the line joining (i, j) and \underline{q} is proportional to that of \underline{v} , and the component of \underline{v}' perpendicular to that line is in the same proportion to \underline{v} , but reflected about the line. Thus both \underline{v} and \underline{v}' are tangent to an arc connecting the two points. The field around such a "current dipole" as v_y is illustrated in figure 4.11.

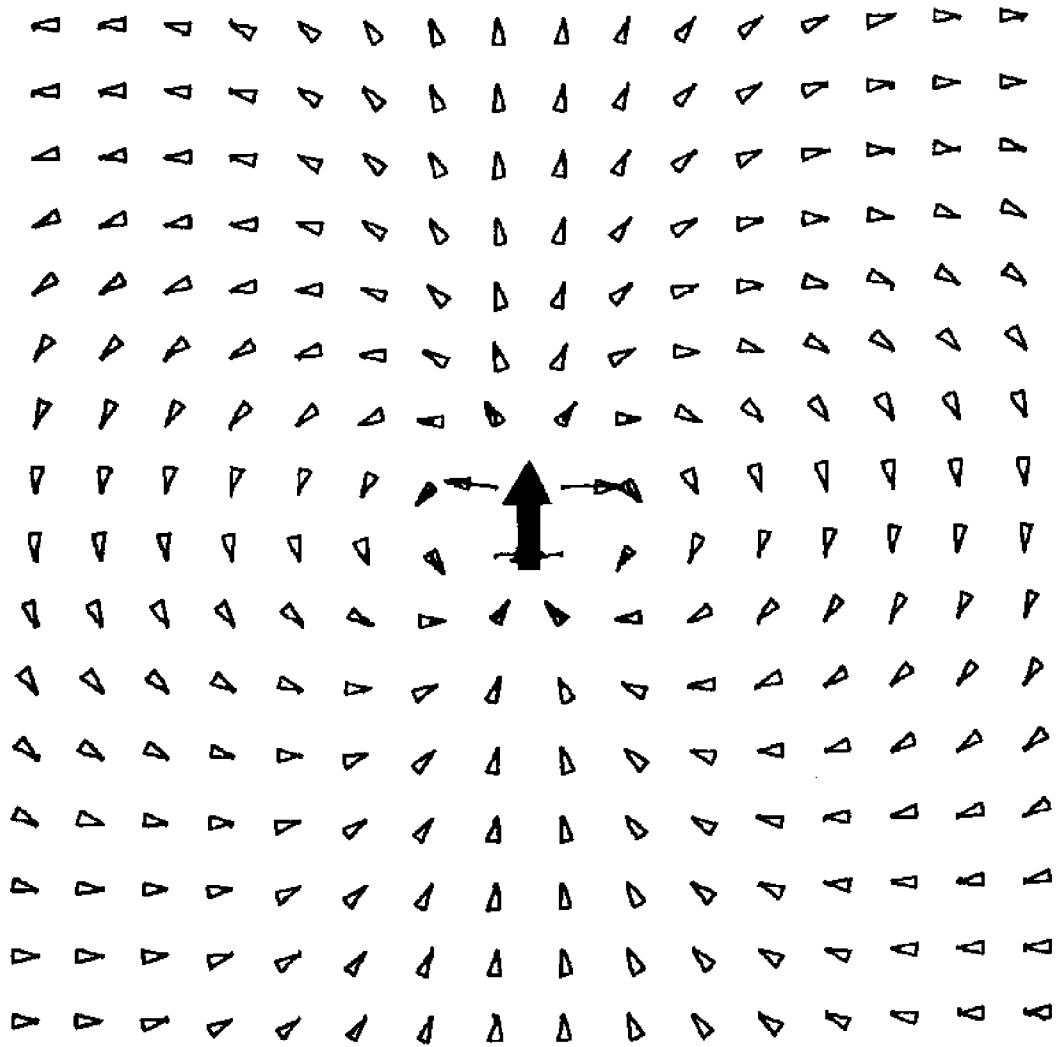


Figure 4.11: Velocities in the field of a current dipole.

Presumably, then, the solenoidal velocity at any point could be computed by the appropriate sum over arcs struck from all the locally defined velocities (except in the near field).

What to do in the near field is problematic. Since the velocity field may be sparsely defined, extremely tight loops may result from such a calculation of \underline{v}' everywhere--a clear contradiction of the minimal gradient criterion. What is needed is some idea of the scale of circulation in the water body--that is, the size of the gyres involved. If the distance from \underline{q} to any square i,j be large with respect to this scale, the arc velocity from this square would best characterize \underline{v}' ; if the distance be small, the velocity of the square itself might better represent $\underline{v}'(\underline{q})$ in line with the smoothness constraint. But how to find the scale of circulation?

The initial assumption in this section was that the path connecting a drifter to its target is an arc. This is all the information needed to arrive at a sense of scale, since both points in this case have to lie in the same gyre. So the appropriate scale of circulation is related to the distance from a drifter to its goal. The effect of each reference velocity on the drifter should be circulatory (along an arc) if distant relative to the goal or viscous (parallel) if close by.

But there is a symmetry here akin to the reversibility discussed earlier. Might not the initial direction of the arc be just as well determined by a calculation of \underline{v}' at the goal, which is then followed backward along an arc to the drifter? Such an approach could be thought of as creating a dipole at the goal which pulls the drifter in along a connecting arc. The answer lies somewhere between the direct dipole influence and that influence directed through the target and includes an ef-


fect which moderates velocity gradients.

Consider the situation pictured in figure 4.12a. Being shown is the influence of the velocity at \underline{r} on the selection of an initial velocity from \underline{q} to \underline{q}_R . The dipole effect of \underline{v} on \underline{q} is \underline{v}'_A ; the dipole effect of \underline{v} on \underline{q}_R is \underline{v}'_B . Finally, the dipole effect of \underline{v}'_R on \underline{q} is \underline{v}'_B . All of these neglect weights due to distance (i.e. speeds do not diminish). The net effect of \underline{v} on \underline{q} is taken to be some average \underline{v}' of \underline{v}'_A and \underline{v}'_B . The weighting of the average is determined by the ratio of the distance between \underline{r} and \underline{q} to that between \underline{r} and \underline{q}_R as follows:

$$\underline{v}' = \frac{|\underline{r} - \underline{q}|^2 \underline{v}'_A + |\underline{r} - \underline{q}_R|^2 \underline{v}'_B}{|\underline{r} - \underline{q}|^2 + |\underline{r} - \underline{q}_R|^2}.$$

That is, the farther \underline{r} is from \underline{q} in relation to its distance from \underline{q}_R , the more \underline{v}' is dominated by the direct dipole influence; the closer it is to \underline{q} , the more \underline{v}' is determined by the dipole influence diverted through \underline{q}_R . As \underline{r} becomes very close to \underline{q} , the dominant, indirect dipole effect virtually parallels \underline{v} itself (see figure 4.12b); hence a reduction of local disparity in current directions, giving hope for a decrease in gradients.

This procedure should be carried out for every \underline{r} in the field, accumulating a weighted sum for \underline{v}' . The formula for the dipole field suggests a weighting inversely proportional to $|\underline{q} - \underline{r}|^3$. This is somewhat inappropriate in the present situation because the averaging process between \underline{v}'_A and \underline{v}'_B will almost always diminish the resultant speed. Indeed, as the distance of \underline{r} from both points \underline{q} and \underline{q}_R increases, the distance ratio approaches unity, and conflicting effects of the direct versus the indirect dipole field exactly cancel. For this reason, an inverse-square weighting is used, *via*:



a) Reference point distant from q and q_R .

b) Reference point close to q .




Figure 4.12: The contribution of the field velocity at a point to the circulation of a drifter.

$$\underline{w}_r = \frac{1}{|\underline{q} - \underline{r}|^2}$$

The net effect of referring to every \underline{r} in the domain is

$$\underline{v}'(\underline{q}) \propto \sum_{i=1}^{64} \sum_{j=1}^{64} \frac{|\underline{r}_{ij} - \underline{q}|^2 \underline{v}'_A(\underline{q}, \underline{r}_{ij}) + |\underline{r}_{ij} - \underline{q}_R|^2 \underline{v}'_B(\underline{q}, \underline{q}_R, \underline{r}_{ij})}{(|\underline{r}_{ij} - \underline{q}|^2 + |\underline{r}_{ij} - \underline{q}_R|^2) \cdot |\underline{r}_{ij} - \underline{q}|^2},$$

where

\underline{r}_{ij} is the position of square ij ,

\underline{v}'_A is the non-distance-biased field value at \underline{q} of the dipole $\underline{v}(\underline{r}_{ij})$, and

\underline{v}'_B is the non-distance-biased field value at \underline{q} of the dipole $\underline{v}'_R(\underline{q}_R)$ obtained from the dipole field of $\underline{v}(\underline{r}_{ij})$.

This is not a cheap function to compute, particularly since it must be figured for each drifter at every transition. Therefore, some simplification is in order.

In the relatively far field, the cumulative effect of the velocities at several \underline{r} in a small area will be close to the effect of their average defined at a single point in the area. Consequently, the velocities in each of several fixed, finite areas may be lumped together for the purposes of this circulation heuristic, reducing the number of terms in the above sum. The method is to average the velocities over each of 64 boxes containing $8 \times 8 = 64$ squares apiece. The defined velocities are weighted in the average by the weights assigned by ADAPT, and a "center of gravity" is established for each box as an average of the square locations, weighted in the same fashion. The velocity and position thus defined are assumed to characterize \underline{v} and \underline{r} for that box of squares (see figure 4.13).

LAKE MICHIGAN




Figure 4.13: Lake Michigan divided into boxes with a few reference points and reference velocities.

This lumping could cause trouble in the near field, however, especially when \underline{q} and \underline{q}_R are relatively close compared to the size of a box. There are two reasons for this. First, the fine structure of the velocities in a nearby box could be far more important in the determination of a trajectory than their average. Second, the location of the box's center of gravity may turn out to be the critical factor in the calculation of $\underline{v}'(\underline{q})$ because of its wide range of possible relative distances to the drifter and its target. There are two solutions to this problem. One is to use the local, square velocities rather than the lumped, box velocities in the near field. This would take the fine structure into account and alleviate the center-of-gravity problem. The other is to rework the distance biasing so that the location of a box's center of gravity doesn't matter so much in the near field. This would not account for the fine structure, but it would take care of the distance ratio problem.

The second approach has been taken here, chiefly because it retains the benefits of a reduction in computation time. Besides, the effect of the fine structure at such a local level should be taken care of by the anti-diffusion computations. So the weighting is adjusted to make the near field look more like the far field by adding to each squared distance a constant equal to the square of half the side length of a box ($= 16S^2$). Therefore, relative distances in the near field will be much less distinguishable in their ratio than before. So, the final form for the calculation of \underline{v}' is

$$\underline{v}'(\underline{q}) = \sum_{i=1}^8 \sum_{j=1}^8 \frac{(|\underline{r}_{ij} - \underline{q}|^2 + 16S^2)\underline{v}'_A + (|\underline{r}_{ij} - \underline{q}_R|^2 + 16S^2)\underline{v}'_B}{(|\underline{r}_{ij} - \underline{q}|^2 + |\underline{r}_{ij} - \underline{q}_R|^2 + 32S^2) \cdot (|\underline{r}_{ij} - \underline{q}|^2 + 16S^2)},$$

where

\underline{r}_{ij} is the center of gravity for box ij ,

$\underline{v}'_A = \underline{cir}(\underline{v}(\underline{r}_{ij}), \underline{r}_{ij}-\underline{q})$, and

$\underline{v}'_B = \underline{cir}(\underline{cir}(\underline{v}(\underline{r}_{ij}), \underline{r}_{ij}-\underline{q}_R), \underline{q}_R - \underline{q})$, where

$\underline{cir}(\underline{v}, \Delta\underline{q})$ is the undiminished field value from a dipole \underline{v} at displacement $\Delta\underline{q}$.

The x-component of $\underline{cir}(\underline{v}, \Delta\underline{q})$ is computed as follows:

$$\text{cir}_x(\underline{v}, \Delta\underline{q}) = \frac{v_x(\Delta q_x^2 - \Delta q_y^2) + 2v_y\Delta q_x\Delta q_y}{\Delta q_x^2 + \Delta q_y^2}$$

The y-component is gotten by switching x and y above. (Note from this formula that $\underline{cir}(\underline{v}, \Delta\underline{q}) = \underline{cir}(\underline{v}, -\Delta\underline{q})$.)

From the direction given by $\underline{v}'(\underline{q})$, one may determine the arc from \underline{q} to \underline{q}_R . The length of this arc is gotten as follows:

Let $\alpha = |\text{angle between } \underline{q}_R-\underline{q} \text{ and } \underline{v}'(\underline{q})|$, such that $\alpha \leq \pi$.

Arc length = $|\text{arc}| = r\theta$, where

$$r = |\underline{q}_R-\underline{q}|/(2\sin\alpha), \text{ and}$$

$$\theta = 2\alpha.$$

The speed in the direction of $\underline{v}'(\underline{q})$ is $|\text{arc}|/\Delta t$, where Δt is the time remaining to get to the target.

So the desired velocity computed by the circulation heuristic to carry a drifter to its target is

$$\underline{v}_c = |\text{arc}|/\Delta t \cdot \frac{\underline{v}'}{|\underline{v}'|}.$$

Figures 4.14 and 4.15 show two partially-defined velocity fields and their effects on the selection of arcs from various hypothetical drifter locations to their respective targets. This does not imply that each




Figure 4.14: Arcs computed by the circulation heuristic in the field of a simple gyre.




Figure 4.15: Arcs computed by the circulation heuristic in the field of a double gyre.

drifter would follow its corresponding arc all the way to its goal, because the circulatory effect at a point further along an initially computed arc may imply a somewhat different arc for the rest of the way, and so on. But the speed assigned to each drifter assumes the arc to be the complete path, and for very circuitous arcs it may be excessively high if time is short. Therefore a way of correcting the velocity in these situations is necessary and is given in the next section in a hybrid heuristic using both the direct and circumspect approaches.

The interaction of just the circulation heuristic with the adaptive element may be demonstrated by considering again the situation of figure 4.8, reiterated in figure 4.16a. Starting with a blank field, several iterations are made with drifters running both directions in time, as before. The trajectories from each are shown in figures 4.16b to k. The resulting velocity field is plotted in 4.16l. In this example the circulation heuristic has almost completely eliminated dispersion. This has been accomplished by setting off from the release point in a direction which is not within the angle subtended by the recoveries. The cost of this policy is a possible increase in overall acceleration, which in the final scheme, will have to be weighed against other considerations.

In summary, the path to a drifter has been assumed to be an arc. This presupposes that the drifter and its target both lie in the same simple gyre. While this is certainly not always the case, it is most likely true in a drifter's middle to final steps when its preferred velocity is most important. Which of the infinity of possible arcs best characterizes the path is derived on the basis of the non-divergent field effect of each defined velocity in the domain at the drifter's location and at its target location. The resultant velocity determines the

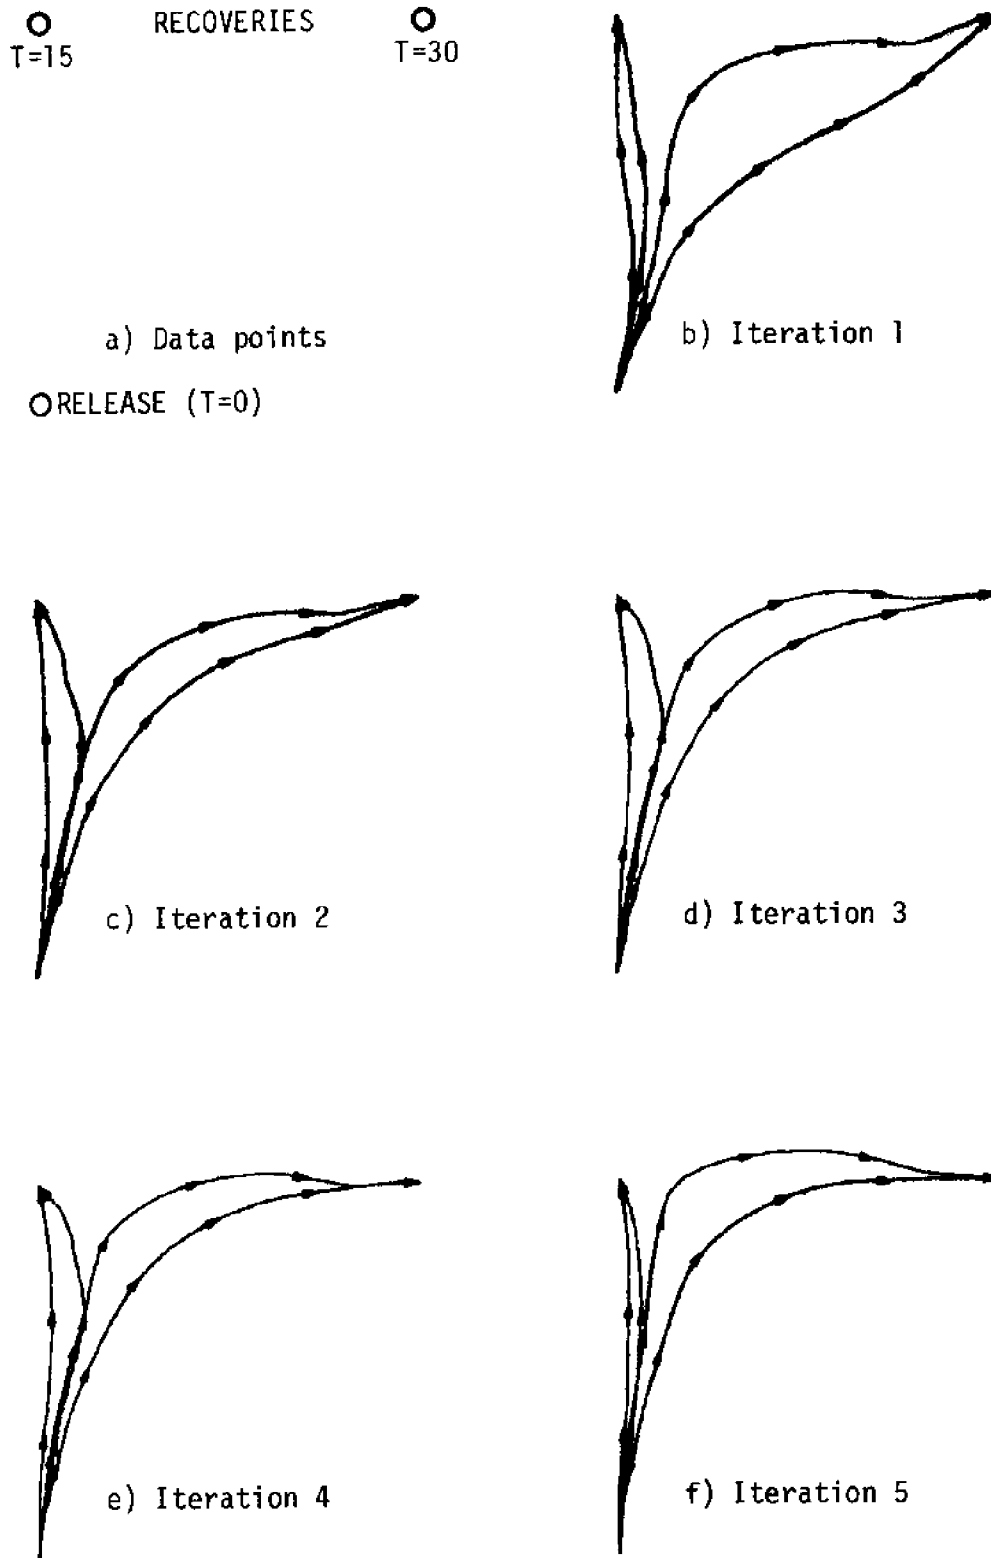


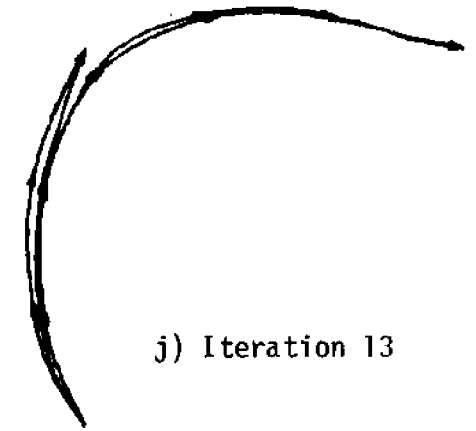
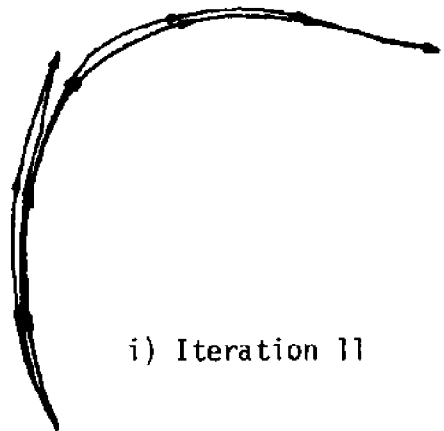
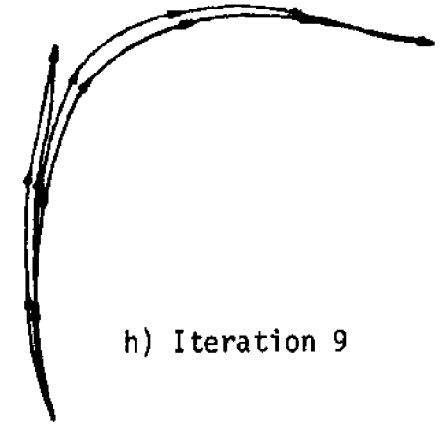



Figure 4.16: Trajectories from circulation heuristic. $k=0.2 \dots /$



/ . . . 4.16, cont'd.

drifter's direction and hence the arc connecting the drifter and its target to which the velocity is assumed tangent. The drifter's speed is gotten from the length of the arc divided by the remaining time to recovery. In some cases, this speed is quite excessive due to the implied arc length, so a way of moderating this effect must be found.

4.8 Hitting the Target--The Hybrid Approach

So far, two goal-directed heuristics have been proposed, each with a specific purpose. The straight-line approach guarantees a drifter's timely arrival at its goal and attempts to minimize its acceleration in getting there, but in a narrow, transition-to-transition sense. The circulatory approach, on the other hand, tries to plot a trajectory to the goal which is cooperative, in some sense, with the smooth and non-divergent aspects of the extant field. These two heuristics must be coordinated, each picking up the reins when the other is incapable of a reasonable, tactical decision as to what direction to go next. Finally, the shoreline must be taken into consideration to avoid running aground prematurely and to redirect currents to run parallel to it when the need and a basis for redirection exists.

When the circulation heuristic demands too much speed, the straight-line heuristic should increase its influence, bringing the preferred velocity back around toward the target. A convenient basis for balancing the two heuristics is acceleration. A high speed determined by the circulation heuristic, which is largely due to more than the straight-line distance to the goal and the amount of time left, will entail a high acceleration as well. This acceleration is given by

$$|a_c| = |v_c|^2/r, \text{ where}$$

r is the implied radius of the arc leading to the target.

The higher this acceleration is, the more weight the straight-line heuristic should demand; so its contribution to the overall preferred velocity can be made linear in $|a_c|$ as follows:

$$\underline{v}_p = \frac{c\underline{v}_c + |a_c|\underline{v}_s}{c + |a_c|}, \text{ where}$$

\underline{v}_s is the straight-line velocity, and

c is some constant used to weight \underline{v}_c against \underline{v}_s .

The constant c implicitly has the dimension of acceleration. Its value will therefore represent some "allowed" acceleration of circulation beyond which the straight-line heuristic must predominate.

The resultant velocity \underline{v}_p should still be meaningful in a goal-directed sense; that is, it should represent some plan for getting to the goal rather than just a compromise between two other conflicting plans. Both heuristics considered chart a course along some arc (if one includes the degenerate case of a line segment), so their combination ought to too. It doesn't. The problem is that the speed resulting from the velocity average is insufficient to cover the arc implied by the direction of \underline{v}_p in the remaining time. Figure 4.17 illustrates this situation. Given two vectors \underline{v}_c and \underline{v}_s emanating from the same point, their average will lie somewhere on the line connecting them. But this average will always lie inside the envelope of vectors corresponding to legitimate arc velocities; hence the insufficient speed.

The answer is to recompute the speed, based on the distance to the goal along the arc to which \underline{v}_p is tangent. Thus, the corrected speed may be computed from the arc-length as before:




Figure 4.17: Correction of hybrid velocity for insufficient speed.

$$\frac{v'}{p} = |\text{arc}| / \Delta t \cdot \frac{v_p}{|v_p|}, \text{ where}$$

$|\text{arc}|$ is calculated from v_p as was done from $v'(q)$.

This new velocity, now, represents a plan for getting to the goal.

By combining the straight-line heuristic with another, the property of the former ensuring arrival at the goal in the proper time is not lost. This is because the acceleration implied by the circulation heuristic is inversely proportional to Δt^2 . As time runs short, any deviation from the straight and narrow path will entail such high acceleration that the straight-line heuristic will virtually control the decision-making.

All of the discussion so far is relevant in open water, but nothing has been said about nearshore areas. In general, a drifter should never run ashore before arriving at its goal. If it does, an adjustment in the factors leading it there is necessary. In the approach used here, the shoreline is ignored until it's almost too late. This is to say that if the velocity computed by the hybrid heuristic is such that the drifter is sure to run aground in the next transition, this velocity is altered to avoid the shore and is given sufficient urgency over and above the usual measure to have an immediate and far-reaching influence. In this sense, running aground is treated as an emergency situation. But one must be careful lest the reaction to such a situation be too extreme and have undesirable side effects.

In a convex body of water, every point is "visible" to every other point without obstruction. This implies that in such a body of water, no drifter following the straight-line path to its goal will run aground; it's only when other factors enter into its trajectory that the possibility

arises. Therefore, when a drifter is about to run aground under the velocity computed by the hybrid heuristic, a deviation from that direction toward the recovery site will steer the drifter away from collision. In a body of water having bays and peninsulae, the solution is more complex. But since the lower basin of Lake Michigan is roughly convex, the intricacies of the non-convex case will not be dealt with here.

The algorithm to correct \underline{v}'_p and avoid a collision is as follows:

1. Collisions \leftarrow 0
2. $(q_R - q) \cdot \underline{v}'_p / (|q_R - q| \cdot |\underline{v}'_p|) > 0.98?$ [\underline{v}'_p already toward q_p ?
(This is done to prevent minor non-convexities from causing an endless loop.)]
 - Yes: go to 7
 - No: continue
3. $G(q + S\underline{v}'_p / |\underline{v}'_p|) = \text{land?}$ [Is adjacent square to which \underline{v}'_p points on land?]
 - Yes: continue
 - No: go to 7
4. Collisions \leftarrow Collisions + 1
5. Rotate \underline{v}'_p .2 radians closer to $q_R - q$
6. go to 2.
7. end

Note that \underline{v}'_p is not corrected for the new implied arc length and will therefore be a bit faster than necessary to get to the goal. Though this is a slight departure from the letter of the hybrid heuristic, no ill-effects have resulted from it. Besides, as alongshore currents become established, the possibility of collision is abated, and the above algo-

rithm should find scarce application.

When this procedure is completed, \underline{v}'_p will be diverted from the shore and $0.2 \cdot \text{Collisions}$ will equal the number of radians through which \underline{v}'_p had to be rotated to avoid running aground. Collisions is then weighted by a constant parameter and added into the urgency factor to be sent to ADAPT. ADAPT is additionally signalled to recognize the emergency, and makes changes not just to the square in which the drifter resides, but also to the eight adjacent squares. What results is a heavy weighting in the velocity field for the corrected velocity which also heavily influences the velocity of the box containing the square containing the drifter. Consequently, the correction will have immediate global influence on the circulation velocities of all the drifters and should eventually tend to redirect those velocities bringing it to shore in the first place.

This collision prevention measure will tend to ensure that velocities close to shore run parallel to the shore and not into it. But this is only possible when the direction alongshore to divert perpendicular velocities is ascertainable, as it plainly is in the case of an imminent collision. When a lone drifter is directed straightway to a nearshore recovery, and when no other clues are available in the velocity field to direct its approach, a current into shore (or away from shore for backwards drifters) will result. The only thing to do in this case is go back to the field if possible and do more experiments to determine the nearshore velocities.

One final adjustment is necessary due to the discrete nature of the transitions. Following a velocity initially tangent to an arc for any finite amount of time will always carry the drifter away from the arc.

Though this effect be small for each one-square transition, an accumulation of such small errors has led to ever-widening arcs in trials without the correction factor. So \underline{v}'_p has to be directed inward slightly so that following it for one transition will land the drifter on the arc. Referring to figure 4.18, let α be the angle between \underline{v}'_p and the line from q to q_R . The angle of the arc will be 2α , which must be covered in time Δt . In elapsed time e , only $\beta = \frac{2\alpha e}{\Delta t}$ of the arc will be covered, so in this time one would like to be on the arc at point P . So \underline{v}'_p has to be rotated from α to α' , assuming the previous elapsed time e will elapse on the current transition too. Now q bears the same relation with P and β as it does with q_R and 2α , so $\alpha - \alpha' = \beta/2 = \alpha e/\Delta t$. Therefore $\alpha' = \alpha(\Delta t - e)/\Delta t$. If the expected elapsed time e is equal to Δt , the time remaining, then $\alpha' = 0$, and the drifter is directed straight to the target. The speed of \underline{v}'_p is not altered to correspond to the length of the chord \overline{qP} , since for small angles β this length is virtually the same as the length of the arc \widehat{qP} . And β will always be a small angle, since e will be small: a large β and small e would imply more acceleration than the balance between the circulation heuristic and the straight-line heuristic would allow. Correcting \underline{v}'_p in this way keeps the drifter on the planned trajectory--barring other forces, of course.

So the hybrid heuristic begins as a compromise between the direct and circumspect approaches. By adjusting the speed of the compromise velocity to correspond to the requisite speed along the implied arc, the hybrid heuristic becomes a plan for getting to the goal in its own right. But this plan is only valid in open water. If a drifter is directed by it to run aground, the velocity computed thereby must be diverted to avoid imminent collision. Finally, the velocity must be changed again




Figure 4.18: Factors pertaining to velocity correction for discrete transitions.

to follow a discretized version of the original arc, since time passes in discrete jumps. The resulting velocity is then the preferred velocity for the drifter computed by DECIDE.

4.9 Adaptive Strategy

Although ADAPT is responsible for adapting the velocity field to reflect the behavior of drifters, it depends on parameters and policies established at a higher level. In addition, changes to the velocity field which reflect considerations other than drifter behavior can be made elsewhere in the system. All these things affect the final result, and their coordination can be thought of as a strategy for adaptation.

The level at which these policies and changes are made is the iteration level, between complete runs of the goal-directed drifter simulation. The main purpose for instituting them is to assist the inference process through the early, transient portion of the velocity field's evolution to keep it from getting stuck on obviously bad inferences. One reason this is necessary is due to the circulation heuristic, which depends from the start on the entire velocity field. When the process begins from a blank field, this heuristic has nothing to go on, so the first, tentative steps of each drifter must rely entirely on the straight-line heuristic. Indeed, until the first iteration is completed, the circulation heuristic may be operating with quite a bit of misinformation. And this misinformation may lead it to worse behavior if allowed to linger through several iterations. Worse yet is the possibility that the heavy weighting commonly associated with the last, desperate accelerations of a drifter rushing headlong toward its target will flavor the initial steps of drifters coming the other way in time from that target for many iterations to come. The solution to these problems lies in massive

limitations and reductions in the weighting of the square velocities during the early iterations.

The weighting limitations are imposed on ADAPT each time it is about to assign a weight to a preferred velocity, based on the calculated urgency. It is allowed to proceed as before but must use the minimum of its calculated weight and the square of the iteration number for the actual weighting of the velocity. In early iterations, therefore, each drifter visiting a square is treated almost equally with the others when it comes to altering the square's velocity. This way, last ditch accelerations and their concomitant high speeds will not be allowed as much influence on the local field as their urgency demands. It is assumed that these cases are illegitimate and will not arise in later iterations. If they do, it might be because they are accurate anyway, and by then the square of the iteration number will be large enough to admit the high weights called for. This does not keep such drifters from reaching their goal, since their weighting in MOVE is unaffected by this policy; it just controls their effect on the velocity field and hence on other drifters. In the long run, as the iteration number becomes very high, the action of ADAPT will be as though no limits on it existed.

Another problem arises from the upper limit imposed on the squares' weights due to their finite representation in the computer; they cannot be allowed to increase unbounded. The solution proposed for this problem violates, to some extent, the condition that late visitors to a square in one iteration have influence on any given drifter equal to that of early visitors during the next; but its application has not caused any difficulties, so it's used anyway. After every iteration, the weight in each square is halved. This not only limits the ultimate size of a weight,

as will be shown, but causes an exponential decay in spurious, transient influences in any given square. At equilibrium (convergence) when the drifter trajectories do not change from iteration to iteration, the net weight added to a given square during one iteration will be some constant w . The result is then halved, and w is added again during the next iteration, and so forth. In the limit, the weight of the square will be

$$\begin{aligned} w_T &= \frac{\frac{w}{2} + w}{2} + w \dots \\ &= w + \frac{1}{2}(w + \frac{1}{2}(w + \frac{1}{2}(w + \dots \\ &= w \sum_{n=0}^{\infty} \left(\frac{1}{2}\right)^n = 2w. \end{aligned}$$

So after a finite number of iterations, w_T will be less than $2w$, assuming equilibrium. This means that the effect on a drifter of visitors to its cell from a previous iteration will be half the effect of those having already been there on the current one. But this does not change the result that the velocity of the square is the average of the velocities of the drifters moving through it; it only (and only slightly) affects these velocities themselves.

The adaptive strategy, then, operates on the iteration level and involves both parameter control and direct intervention. By limiting the effect of high urgencies during the early phases of the field's evolution, the consequences of wild behavioral variations are tethered. What little consequence lingers from such transient effects is diminished over several iterations by halving the square weights after each one. This policy is additionally and primarily carried out to keep a lid on the net weight of each square.

4.10 Velocities of Unvisited Squares

When the inference process has finished, some squares may have undefined velocities (i.e. zero weights) because no drifters visited them during any recent iterations. In order to test the generated hypothesis by use of the scheme involving the simulation algorithm, these blanks have to be filled in. How they are filled in should be compatible with the existing field in terms of smoothness and non-divergence.

As before, the reduction of divergence and gradients is a central aim. This immediately suggests using a variant of the circulation heuristic. Such a procedure would have the benefit of including drifter data in the synthesis of these velocities in addition to using the extant field. In the circulation heuristic, reference points throughout the field are used to compute a velocity at each drifter location. In this situation, the roles are reversed: the velocities at the release locations (which include recovery locations due to reversed drifters) are used to compute the velocity of each undefined reference point. An undefined reference point is one whose box contains all undefined squares, and it lies at the center of the box.

Consider the situation depicted in figure 4.19. \underline{r} is an undefined reference point. $\underline{q}_{0,i}$ is the release; $\underline{q}_{R,i}$, the associated recovery of drifter i . The velocity $\underline{v}_i(\underline{r})$ will be an average of two effects. First is the direct dipole field effect of $\underline{v}(\underline{q}_{0,i})$ at \underline{r} . Second is the same effect felt indirectly through $\underline{q}_{R,i}$. The weighting is the same as before, and $\underline{v}_i(\underline{r})$ as a function of $\underline{v}(\underline{q}_{0,i})$ can be written:

$$\underline{v}_i(\underline{r}) = \frac{(|\underline{r}-\underline{q}_{0,i}|^2 + 16S^2) \underline{v}_A + (|\underline{r} - \underline{q}_{R,i}|^2 + 16S^2) \underline{v}_B}{(|\underline{r}-\underline{q}_{0,i}|^2 + |\underline{r} - \underline{q}_{R,i}|^2 + 32S^2) \cdot (|\underline{r}-\underline{q}_{0,i}|^2 + 16S^2)}, \text{ where}$$




Figure 4.19: Computation of an undefined field velocity from the field at one release point.

i is an index into the vector of release-recovery pairs,

$$\underline{v}_A = \underline{\text{cir}}(\underline{v}(q_{0,i}), \underline{r}-q_{0,i}), \text{ and}$$

$$\underline{v}_B = \underline{\text{cir}}(\underline{\text{cir}}(\underline{v}(q_{0,i}), q_{R,i}-q_{0,i}), \underline{r}-q_{R,i}), \text{ where}$$

$$\underline{\text{cir}}(\underline{v}, \Delta q) \text{ is defined in section 4.7.}$$

The weighting again, has the effect of reducing local changes in direction. If \underline{r} is close to $q_{0,i}$, $\underline{v}_i(\underline{r}) \approx \underline{v}_B$ will be almost proportional to $\underline{v}(q_{0,i})$.

The overall circulation velocity $\underline{v}_c(\underline{r})$ will be proportional to the sum over all releases and recoveries of the $\underline{v}_i(\underline{r})$. What speed to assign to it has to be gotten from something other than an arc length and Δt value, though, because there are none associated with \underline{r} . Since each $\underline{v}_i(\underline{r})$ is diminished by the sum of its square distances to $q_{0,i}$ and $q_{R,i}$, the simplest approach is to divide the result by the sum of these weighting factors $w_i = 1/(|\underline{r}-q_{0,i}|^2 + 16S^2)$ to get:

$$\underline{v}_c = \frac{\sum \underline{v}_i(\underline{r})}{\sum w_i}.$$

This, then, is the velocity of the box containing \underline{r} .

The velocity of each individual square for which none is yet defined will be an average of the local box velocities and the velocities, if any, defined in neighboring squares. It is defined as follows:

$$\underline{v}(q) = 0.2 \cdot \frac{\sum_{i=-1}^1 \sum_{j=-1}^1 \frac{\underline{v}_c(q + \underline{d}_{ij})}{|\underline{r}(q + \underline{d}_{ij}) - q|^2 + 16S^2}}{\sum_{i=-1}^1 \sum_{j=-1}^1 \frac{1}{|\underline{r}(q + \underline{d}_{in}) - q|^2 + 16S^2}} +$$

$$0.8 \cdot \frac{\sum_{m=-2}^2 \sum_{n=-2}^2 \frac{v(\underline{q} + \underline{\delta}_{mn}) - w(\underline{q} + \underline{\delta}_{mn})}{\delta_{ij}^2}}{\sum_{m=-2}^2 \sum_{n=-2}^2 \frac{w(\underline{q} + \underline{\delta}_{mn})}{\delta_{mn}^2}}, \text{ where}$$

$$\underline{\delta}_{ij} = (8S \cdot i, 8S \cdot j),$$

$\underline{v}_c(\underline{q})$ is the box velocity of the box containing \underline{q} ,

$\underline{r}(\underline{q})$ is the box center of gravity (reference location)

for the box containing \underline{q} , if such a box exists;

otherwise any term containing \underline{r} in either sum is

ignored.

$$\underline{\delta}_{mn} = (S \cdot m, S \cdot n),$$

$\underline{v}(\underline{q})$ is the velocity defined for the square containing

\underline{q} , and

$w(\underline{q})$ is the corresponding weight.

The procedure for filling in the blanks is then:

1. Calculate reference velocities for all undefined reference points, using release and recovery points for velocity information.
2. Calculate velocities for all undefined squares as shown in the previous paragraph.
3. Assign a weight of one to each hitherto undefined square.

The last step grants full rights of definition to each newly filled square, so it will be recognized as such by every function referring to it.

This procedure should result in velocities which blend well with the inferred velocities. That it does is demonstrated in the next chapter. What it does not do is take the shoreline directly into account.

Perpendicular nearshore velocities resulting from this algorithm should be tested in the field to find their true, parallel direction, since there is no other basis for assigning one.

4.11 Putting It All Together

In line with the algorithmic structure laid out in section 4.4 and the functions described in the following sections, the inference system can now be defined in its entirety. First the global parameters and some service functions will be defined, and then the expanded EXEC, DECIDE, MOVE, and FILL, the undefined-square defining routine, will be detailed.

The global parameters fall into three groups: those associated with the drifters, those connected with the velocity field, and those corresponding to the adaptive strategy. The drifter parameters are:

($Q(i)$, $i=1, \dots, K$), the drifter positions,
 ($T(i)$, $i=1, \dots, K$), the drifter times (since release),
 ($ET(i)$, $i=1, \dots, K$), the times elapsed on the previous
 transition,
 ($Q0(i)$, $i=1, \dots, K$), the release positions,
 ($QR(i)$, $i=1, \dots, K$), the corresponding recovery positions,
 ($TR(i)$, $i=1, \dots, K$), the recovery times (since release),
 ($U(i)$, $i=1, \dots, K$), preferred velocities for the drifters,
 ($H(i)$, $i=1, \dots, K$), urgency coefficients for the drifters,
 and
 ($C(i)$, $i=1, \dots, K$), shoreline collision coefficients for the
 drifters.

The velocity field parameters are:

($V(i,j)$, $i=1, \dots, 64$; $j=1, \dots, 64$), the velocities of the
 squares,

$(W(i,j), i=1,\dots,64; j=1,\dots,64)$, the weights associated
 with the squares,
 $(G(i,j), i=1,\dots,64; j=1,\dots,64)$, the geographical descrip-
 tor values associated with the squares,
 $(R(i,j), i=1,\dots,8; j=1,\dots,8)$, the reference locations for
 each box of squares,
 $(VR(i,j), i=1,\dots,8; j=1,\dots,8)$, the reference velocities
 for each box of squares,
 $(WR(i,j), i=1,\dots,8; j=1,\dots,8)$, the reference weights for
 each box of squares, and
 S, the width of each square.

The adaptive strategy parameters are:

ITER, the iteration number
 CIR, the "allowed" acceleration of circulation,
 AVG, the anti-diffusion coefficient, and
 SHORE, an anti-collision coefficient.

The service routines are primarily concerned with getting field values from drifter positions. They are:

VEL(Q): [For getting and putting square velocities]
 1. $X \leftarrow \text{truncate}(\min(\max(Q_x/S + 1, 1), 64))$
 2. $Y \leftarrow \text{truncate}(\min(\max(Q_y/S + 1, 1), 64))$
 [Indices of nearest inbounds square]
 3. $\underline{VEL} \leftrightarrow V(X,Y)$ [VEL can be used on both sides of
 an assignment]
 4. return

WT(Q): [For getting and putting weights]
 1. $X \leftarrow \text{truncate}(\min(\max(Q_x/S + 1, 1), 64))$

2. $Y \leftarrow \text{truncate}(\min(\max(Q_y/S + 1, 1), 64))$
3. $\text{WT}(Q) \leftrightarrow W(X,Y)$
4. return

TYP(Q): [For getting and putting "land" and "water" values]

Same as VEL and WT but using G, *mutatis mutandis*.

VES(Q,N) and WTS(Q,N): [For getting "smoothed" velocities and cumulative weights]

1. VES \leftarrow 0
2. WTS \leftarrow 0
3. For each i,j , $i=-N,\dots,N$; $j=-N,\dots,N$
 - 3.1 $P_x \leftarrow Q_x - \text{mod}(Q_x, S) + i \cdot S + S/2$
 - 3.2 $P_y \leftarrow Q_y - \text{mod}(Q_y, S) + j \cdot S + S/2$
 - [P is center location of Q's ij th neighboring square]
 - 3.3 $\text{BIAS} \leftarrow \text{WT}(\underline{Q}) / (|\underline{Q} - \underline{P}|^2 + S^2)$
 - 3.4 $\underline{\text{VES}} \leftarrow \underline{\text{VES}} + \text{BIAS} \cdot \underline{\text{VEL}}(\underline{P})$
 - 3.5 $\text{WTS} \leftarrow \text{WTS} + \text{BIAS}$
 - 3.6 next i,j
4. VES $\leftarrow \underline{\text{VES}}/\text{WTS}$
5. return VES or WTS, whichever was called

REF(Q): [For getting and putting reference positions of boxes]

1. $X \leftarrow Q_x / (S \cdot 8) + 1$
2. $Y \leftarrow Q_y / (S \cdot 8) + 1$
3. Is $1 \leq X \leq 8$ and $1 \leq Y \leq 8$?

Yes: REF \leftrightarrow R(X,Y)

No: REF undefined

4. return

VREF(Q): [For getting and putting reference velocities]

Same as REF but with VR, *mutatis mutandis*.

WREF(Q): [For getting and putting reference weights]

Same as REF but with WR, *mutatis mutandis*.

The executive algorithm is now laid out as follows:

EXEC:

1. read CIR, AVG, SHORE, [Get strategic parameters]
 - ND, (Q0(i), QR(i), TR(i), i=1,...,ND) [Get drifter data]
2. For each i, i = ND+1,...,2·ND [Create reversed drifters in second half of drifter array]
 - 2.1 Q0(i) ← QR(i-ND)
 - 2.2 QR(i) ← Q0(i-ND)
 - 2.3 next i
3. K ND·2 [Total number of drifters]
4. For each i,j, i=1,...,64; j=1,...,64 [Blank out V]
 - 4.1 V(i,j) ← 0
 - 4.2 W(i,j) ← 0
 - 4.3 next i,j
5. For each i,j, i=1,...,8; j=1,...,8 [Blank out references]
 - 5.1 VR(i,j) ← 0
 - 5.2 WR(i,j) ← 0
 - 5.3 next i,j
6. For each ITER, ITER = 1,...
 - 6.1 For each i, i=1,...,K [Restart drifters]
 - 6.1.1 Q(i) = Q0(i)
 - 6.1.2 T(i) = 0


```

6.1.3 ET(i) ← 1 [Fake the first transition time]
6.1.4 next i
6.2 Until each T(i) = TR(i), i=1,...,K, repeat
6.2.1 call DECIDE
6.2.2 call ADAPT
6.2.3 call MOVE
6.2.4 next repetition
6.3 Has convergence been reached? [Determined by visual inspection]
Yes: call FILL and stop
No: continue
6.4 For each i,j, i=1,...,64; j=1,...,64 [Halve weights]
6.4.1 W(i,j) ← W(i,j)/2
6.4.2 next i,j
6.5 For each i,j, i=1,...,8; j=1,...,8
6.5.1 WR(i,j) ← WR(i,j)/2
6.5.2 next i,j
6.6 next ITER [Have another go at it]
7. end

```

The decision routine, which calculates the preferred velocities for each drifter and its urgency, is defined as follows:

DECIDE:

```

1. For each i, i=1,...,K [For each drifter]
1.1 T(i) = TR(i)? [Drifter already finished?]
Yes: 1.1.Y1 H(i) ← 0 [Urgency nil]
      1.1.Y2 C(i) ← 0
      1.1.Y3 go to 1.28

```

No: continue

- 1.2 $DIR \leftarrow \text{sign}(1, K/2-i)$ [1, if drifter is forward-moving;
-1 if reversed]
- 1.3 $VC \leftarrow 0$
- 1.4 $UC \leftarrow 0$
- 1.5 For each $m, n, m=1, \dots, 8; n=1, \dots, 8$ [For each reference point]
- 1.5.1 $DRA \leftarrow R(m, n) - Q(i)$
- 1.5.2 $DRB \leftarrow R(m, n) - QR(i)$
- 1.5.3 $WC \leftarrow (16S^2 + |DRA|^2) / (32S^2 + |DRA|^2 + |DRB|^2)$
- 1.5.4 $DFACT \leftarrow 1 / (16S^2 + |DRA|^2)$
- 1.5.5 $VC \leftarrow VC + WC \cdot \text{cir}(VR(m, n) \cdot DIR, DRA) \cdot DFACT$ [Direct
dipole effect; cir defined in sec-
tion 4.7]
- 1.5.6 $UC \leftarrow UC + (1-WC) \cdot \text{cir}(VR(m, n) \cdot DIR, DRB) \cdot DFACT$
[Dipole effect on QR(i)]
- 1.5.7 next m, n
- 1.6 $VC \leftarrow VC + \text{cir}(UC, Q(i) - QR(i))$ [Add in indirect dipole effect]
- 1.7 $ALPHA \leftarrow |\text{angle between } QR(i) - Q(i) \text{ and } VC|$
- 1.8 $ALPHA > \pi$?
- Yes: $ALPHA \leftarrow 2\pi - ALPHA$
- No: continue
- 1.9 $ARC \leftarrow ALPHA \cdot |QR(i) - Q(i)| / \sin(ALPHA)$ [Arc length]
- 1.10 $VC \leftarrow VC / |VC| \cdot ARC / (TR(i) - T(i))$ [Speed = length of arc/ Δt]
- 1.11 $ACCEL \leftarrow |VC|^2 / (2 \cdot \sin(ALPHA) / |QR(i) - Q(i)|)$
- 1.12 $VP \leftarrow (CIR \cdot VC + ACCEL \cdot (QR(i) - Q(i)) / (TR(i) - T(i))) / (CIR + ACCEL)$
[Average in straight-line velocity]
- 1.13 $ALPHA \leftarrow |\text{angle between } QR(i) - Q(i) \text{ and } VP|$

1.14 $\text{ALPHA} > \pi$?

Yes: $\text{ALPHA} \leftarrow 2\pi - \text{ALPHA}$

No: continue

1.15 $\text{ARC} \leftarrow \text{ALPHA} \cdot |\underline{\text{QR}}(i) - \underline{\text{Q}}(i)| / \sin(\text{ALPHA})$ [Length of new arc]

1.16 $\underline{\text{VP}} \leftarrow \underline{\text{VP}} / |\underline{\text{VP}}| \cdot \text{ARC} / (\text{TR}(i) - \text{T}(i))$ [Fix up speed for hybrid velocity]

1.17 $\text{C}(i) \leftarrow 0$ [No collisions yet]

1.18 $\underline{\text{VP}} \cdot (\underline{\text{QR}}(i) - \underline{\text{Q}}(i)) / (|\underline{\text{QR}}(i) - \underline{\text{Q}}(i)| \cdot |\underline{\text{VP}}|) > 0.98$? [VC directed toward goal?]

Yes: go to 1.23

No: continue

1.19 $\text{TYP}(\underline{\text{Q}}(i) + \underline{\text{VP}} / |\underline{\text{VP}}| \cdot \text{S}) = 1$ and? [Collision imminent?]

Yes: continue

No: go to 1.23

1.20 Rotate $\underline{\text{VP}}$ 0.2 radians closer to $\underline{\text{QR}}(i) - \underline{\text{Q}}(i)$

1.21 $\text{C}(i) \leftarrow \text{C}(i) + 1$ [Increment collision count]

1.22 go to 1.18

1.23 $\text{C}(i) > 0$? [Any collisions?]

Yes: continue [Need to recompute ALPHA]

No: go to 1.26

1.24 $\text{ALPHA} \leftarrow$ |angle between $\underline{\text{QR}}(i) - \underline{\text{Q}}(i)$ and $\underline{\text{VP}}$ |

1.25 $\text{ALPHA} > \pi$?

Yes: $\text{ALPHA} \leftarrow 2\pi - \text{ALPHA}$

No: continue

1.26 $\text{E} \leftarrow \min(\text{ET}(i), \text{TR}(i) - \text{T}(i))$ [Estimated elapsed time for this transition]

1.27 $\text{ALPHAP} \leftarrow \text{ALPHA} \cdot (\text{TR}(i) - \text{T}(i) - \text{E}) / (\text{TR}(i) - \text{T}(i))$ [New angle

- for VP relative to QR(i) - Q(i)]
- 1.28 Rotate VP ALPHAP - ALPHA radians closer to QR(i) - Q(i) [Correct for discrete transitions]
- 1.29 U(i) ← VP [Assign preferred velocity]
- 1.30 $H(i) ← 2 \cdot (TR(i) - T(i)) / (\sqrt{E} \cdot (TR(i) - T(i) - E)^2)$ [Urgency factor; shoreline effect is added in ADAPT using C(i)]
- 1.31 next i
2. return

The adaptation routine which alters the velocity field based on a drifter's desired movement is defined as follows:

ADAPT:

1. For each i, i=1,...,K [For each drifter]
- 1.1 $H(i) + C(i) > 0?$ [Any urgency?]
- Yes: continue
- No: go to 1.6
- 1.2 $DIR ← \text{sign}(1, K/2 - i)$ [Get direction of movement]
- 1.3 $WI ← \min(ITER^2, 100 \cdot H(i) + SHORE \cdot C(i))$ [Weight for velocity change]
- 1.4 $NEIGH ← \min(C(i), 1)$ [Number of adjacent squares changed =1 if C(i)=0; =9 if C(i) > 0]
- 1.5 For each m,n, m = -NEIGH,...,NEIGH; n = -NEIGH,..., NEIGH [For each square in the domain of the change]
- 1.5.1 $\underline{P} ← \underline{Q}(i) + (m,n) \cdot S$ [Location in neighboring square]
- 1.5.2 $\underline{VEL}(\underline{P}) ← (\underline{VEL}(\underline{P}) \cdot WT(\underline{P}) + \underline{U}(i) \cdot WI \cdot DIR) / (WT(\underline{P}) + WI)$ [Alter velocity]

- 1.5.3 $\underline{WT(P)} \leftarrow \underline{WT(P)} + \underline{WI}$ [Alter weight]
- 1.5.4 $\underline{VREF(P)} \leftarrow (\underline{VREF(P)} \cdot \underline{WREF(P)} + \underline{U(i)} \cdot \underline{WI} \cdot \underline{DIR}) / (\underline{WREF(P)} + \underline{WI})$
- 1.5.5 $\underline{REF(P)} \leftarrow (\underline{REF(P)} \cdot \underline{WREF(P)} + \underline{P} \cdot \underline{WI}) / (\underline{WREF(P)} + \underline{WI})$
- 1.5.6 $\underline{WREF(P)} \leftarrow \underline{WREF(P)} + \underline{WI}$ [Alter reference velocity, location, and weight]
- 1.5.7 next m,n
- 1.6 next i
2. return

The procedure for filling in the blanks, once the goal-directed inference process is complete is

FILL:

1. For each m,n, m=1,...,8; n=1,...,8 [For each box]
- 1.1 $\underline{WR(m,n)} > 0?$ [Reference point defined?]
- Yes: go to 1.7
- No: continue
- 1.2 $\underline{VC} \leftarrow 0$
- 1.3 $\underline{WD} \leftarrow 0$
- 1.4 $\underline{R(m,n)} \leftarrow 8 \cdot \underline{S} \cdot (m - 1/2, n - 1/2)$ [Reference point is center of box]
- 1.5 For i=1,...,K [For each drifter]
- 1.5.1 $\underline{DRA} \leftarrow \underline{R(m,n)} - \underline{Q0(i)}$
- 1.5.2 $\underline{DRB} \leftarrow \underline{R(m,n)} - \underline{QR(i)}$
- 1.5.3 $\underline{WC} \leftarrow (16 \cdot \underline{S}^2 + |\underline{DRA}|^2) / (32 \cdot \underline{S}^2 + |\underline{DRA}|^2 + |\underline{DRB}|^2)$
- 1.5.4 $\underline{DFACT} \leftarrow 1 / (16 \cdot \underline{S}^2 + |\underline{DRA}|^2)$
- 1.5.5 $\underline{UC} \leftarrow \underline{cir(VEL(Q0(i)), QR(i) - Q0(i))}$ [Dipole effect of $\underline{v(q_0)}$ on $\underline{q_R}$]
- 1.5.6 $\underline{VC} \leftarrow \underline{VC} + \underline{DFACT} \cdot (\underline{WC} \cdot \underline{cir(VEL(Q0(i)), DRA)} +$

$$(1 - WC) \cdot \underline{\text{cir}}(\underline{UC}, \underline{DRB})$$

1.5.7 $WD \leftarrow WD + DFACT$ [Net distance biasing]

1.5.8 next i

1.6 $\underline{VR}(m,n) \leftarrow \underline{VC}/WD$ [New reference velocity]

1.7 next m,n

2. For each m,n, m=1,...,64; n=1,...,64 [For each square, compute
velocity if in water and undefined]

2.1 $G(m,n) = \text{land?}$

Yes: go to 2.8

No: continue

2.2 $W(m,n) > 0?$ [Velocity defined?]

Yes: go to 2.8

No: continue

2.3 $\underline{VL} \leftarrow 0$

2.4 $WL \leftarrow 0$

2.5 $\underline{QP} \leftarrow S \cdot (m - 1/2, n - 1/2)$ [Location of square's center]

2.6 For each a,b, a=-1,0,1; b=-1,0,1 [For a 9-box neighborhood]

2.6.1 $\underline{P} \leftarrow \underline{QP} + 8 \cdot S \cdot (a,b)$

2.6.2 $\underline{\text{REF}}(\underline{P})$ defined? [\underline{P} inbounds?]

Yes: continue

No: go to 2.6.6

2.6.3 $DFACT \leftarrow 1/(16 \cdot S^2 + |\underline{QP} - \underline{\text{REF}}(\underline{P})|^2)$

2.6.4 $\underline{VL} \leftarrow \underline{VL} + \underline{\text{VREF}}(\underline{P}) \cdot DFACT$

2.6.5 $WL \leftarrow WL + DFACT$

2.6.6 next a,b

2.7 $\underline{V}(m,n) \leftarrow (0.2 \cdot \underline{VL}/WL + 0.8 \cdot \min(1, \text{WTS}(\underline{QP}, 2)) \cdot \underline{\text{VES}}(\underline{QP}, 2))/$

$(0.2 + 0.8 \cdot \min(1, \text{WTS}(\underline{QP}, 2)))$ [Local velocity is average of

box velocity and surrounding local velocities (if any)]

2.8 next m,n

3. For each m,n, m=1,...,64; n=1,...,64 [For each square in the water, assign a weight of one, if zero]

3.1 G(m,n) = land?

Yes: continue

No: $W(m,n) \leftarrow \max(1, W(m,n))$ [Define square if undefined]

3.2 next m,n

4. return

This completes the definition of the inference process.

4.12 Conclusion

A procedure for inferring current velocities from drifter data has been defined. It is composed of three major elements controlled by an executive, plus one routine for filling in the holes when it's done. Two of the three major elements, DECIDE and MOVE, form a goal-directed simulation system. To each drifter in this system corresponds a recovery point or target and a recovery time. It is up to DECIDE to calculate before each transition the velocity which each drifter prefers to have for the transition, based on its space-time displacement from the goal and on local and global hydrodynamic constraints. MOVE performs the transition for each drifter based on its preferred velocity and on the local, field velocity already defined at the drifter's location. DECIDE and MOVE work together via an urgency factor for each drifter to guarantee that the drifter gets to the goal on time.

Interposed between the velocity decision and the actual movement of a drifter is the other major element: ADAPT. ADAPT senses each

drifter's desired velocity and changes the local velocity at its location in a way that each local velocity will equal the average of the actual velocities used by drifters moving through its domain of influence.

Controlling everything is the executive, EXEC, which initializes drifters at their release points and puts them through the simulation-adaptation sequence until they've all exhausted their allotted times and reached their goals. This it does repeatedly, starting with a blank velocity field, until the velocity field converges. At this point, the as-yet-undefined local velocities are filled in by FILL. The hypothesis thus generated may then be tested or used in a simulation, or it may suggest new experiments.

Based on a probabilistic reversibility of the transport-diffusion model, calling recoveries releases and *vice-versa* should lead to no different conclusions than a change in the velocity field's sign. Therefore two drifters are assigned to each release-recovery pair, one running backwards in time with attention to the requisite sign changes.

The criterion for computing a drifter's actual velocity from its preferred velocity and from the local, field velocity is minimum dispersion balanced against the urgency of the drifter's situation. The result is that the velocity a drifter takes will be close to the local, field velocity, just as in the transport-diffusion model.

The criteria for deciding on a preferred velocity are minimum velocity, acceleration, divergence, and gradients. Velocity and acceleration are minimized in a very local way by choosing the straight-line path to the goal each time. But acceleration can occur anyway, along with a high degree of divergence, because straight-line trajectories for all the drifters are usually incompatible with a stationary field having

limited diffusivity. Nonetheless, certain inference procedures reported in the literature (Wyatt, *et al.*, 1972; Norcross and Stanley, 1967; Bukin, 1974; Bumpus and Lauzier, 1975) can be shown more or less equivalent to the approach at this stage of development under various parameter settings. But by extending the set of desired trajectories to the family of arcs connecting a drifter to its goal, non-divergence and gradient considerations may enter into the selection of a preferred velocity. The heuristics represented by the straight-line and arc approaches are combined into one hybrid heuristic in which they are balanced by acceleration constraints.

The adaptation is controlled by a strategy which iterates the goal-directed drifter system through several complete runs at a low level of adaptive vigor in order to get through the early, transient part of the inference without getting hung up on wild velocities. Gradually, the adaptive element is allowed to make more long-lasting modifications to the field until the velocity field at one iteration is enough like the previous one to declare convergence. At this point the field is handed over to FILL to define these velocities at locations not visited by drifters. The result is the inferred hypothesis.

The hypothesis thus generated from the data may then be tested using the testing scheme laid out in Chapter 3. Hopefully it will fare better than any of the alternate hypotheses proposed. This hope is challenged and verified in the next chapter in the context of controlled experiments (from simulation) and the summer, 1974, experiments in Lake Michigan.

CHAPTER 5

TRIALS OF THE INFERENCE SYSTEM

5.1 Overview

The inference system is testable in two ways. First, velocity fields may be set up as known hypotheses and used to simulate recovery data under the transport-diffusion scheme. These data can then be submitted to the inference system along with a blank velocity field. The hypothesis generated by the system from the data may then be visually compared to the one generating the data to judge the effectiveness of the inference process. This method of testing provides a good way of comparing the effects of various parameter settings and of tuning the system.

Secondly, if valid inferences are made in such an artificial environment, application of the inference process to data from the field is justified. But one need not accept its results on blind faith, because any hypothesis generated may be tested with the program described in Chapter 3 against competing hypotheses, whether they be *a priori* or inferred from the data by other techniques. If it survives such a test, and if the competing hypotheses are reasonable alternate explanations, then credence may be given to the inferred velocity field as well as to the inferring agent.

In this chapter, a couple hypothetical fields are used to demonstrate the effectiveness and shortcomings of the inference scheme. Next, the 1974 Lake Michigan data are used to generate hypotheses for July and August. These hypotheses are then tested against those considered in

chapter three.

5.2 Testing the Inference Process on Data from Known Fields

The approach used in this section is to try the inference process on a small set of simulated recovery data from a simple hypothesis for various parameter settings. This will demonstrate the effects of the parameters and justify those values used in later examples and on the Lake Michigan data. Next a more complicated set of data is obtained from the same field and a further parameter modification tried. Finally, simulated recovery data from a fairly complex field are obtained and an inference made. Where applicable, the inferred results are referred back to the field generating the data.

The field used for the first set of examples is Kizlauskas and Katz' hypothesis (figure 3.13e), reiterated in figure 5.1a. Drifters were released, their transport simulated, and offshore "recoveries" recorded. The experiment was focussed in the portion of the lake having the simplest circulation, and the release-recovery pairs obtained there are given in figure 5.1b-d.

These recovery data were fed to the inference system, which was first begun with a blank field and the following parameter settings (hereinafter referred to as the *standard values*): $AVG = 0.2$, $CIR = S/(\text{unit time})^2$, $SHORE = 25$. Convergence had come by the ninth iteration. Representative trajectories are shown in figure 5.2 along with the inferred velocity field before and after the FILL operation. Except for those portions of the original field which the data do not represent, agreement is fairly good. Moreover, the vectors defined by FILL are intuitively sound except for the currents perpendicular to the northwest shore, which FILL was not designed to correct.





Figure 5.1: Velocity field and recoveries used for first example set.

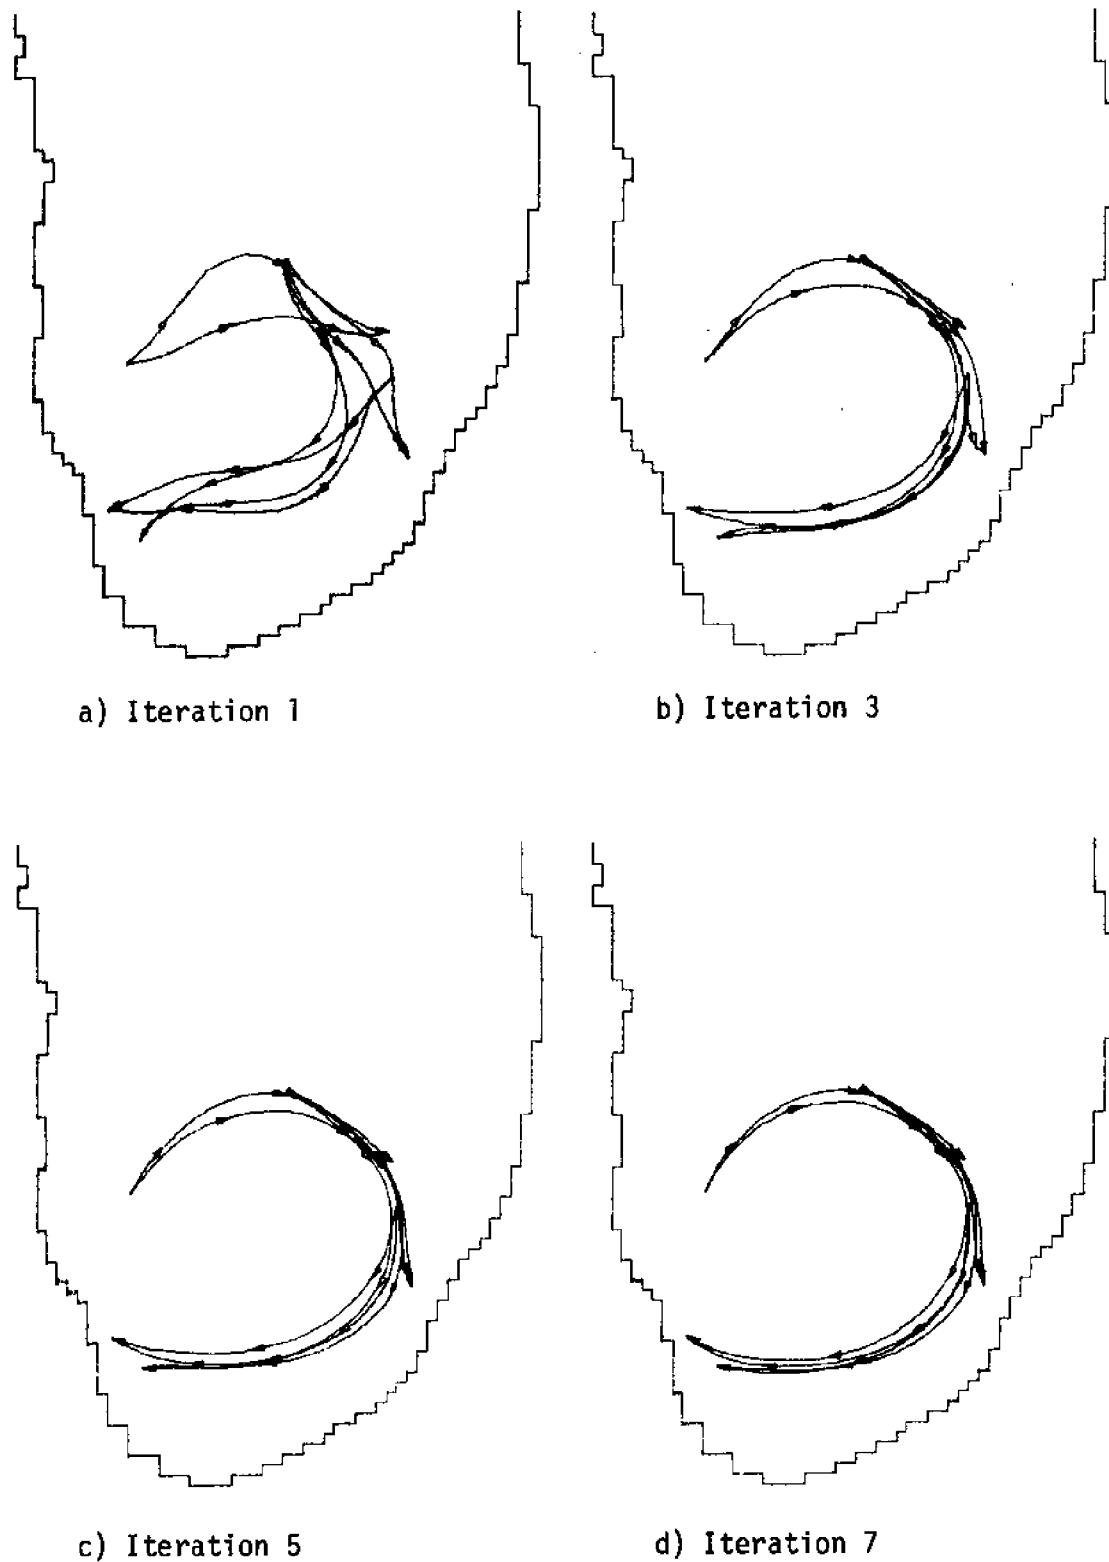





Figure 5.2: Inference results for standard parameter values. . . . /



e) Iteration 9



f) Inferred field (before FILL)



g) Inferred field (after FILL)

/ . . . 5.2, cont'd.

These same data were used with nonstandard parameter values to check the effect of the circulation heuristic and the anti-diffusion coefficient. The sequence shown in figure 5.3 used the same values as figure 5.2 except that $CIR = S/10(\text{time units})^2$. As can be seen, the conflicts among the various trajectories are not resolved nearly as well as with the higher CIR, and the resultant velocity field is much less representative of the original.

In an effort to force conflict resolution by anti-diffusion means rather than through circulation, another trial was made with $CIR = S/10(\text{unit time})^2$ but with $AVG = 2$ rather than 0.2. The results are shown in figure 5.4. The trajectories parallel each other more closely, but at the expense of abrupt directional changes, and the pre-FILL velocity field is an even poorer representation of the original circulation giving rise to the data.

Finally, the effect of running drifters both ways rather than just one was tested by eliminating the reverse drifters. The standard parameter values were used and the process was given nearly twice the number of iterations to come up with the proper velocity field in order to offset its reduced workload per iteration (see figure 5.5). Although convergence ensued much sooner than with drifters going both ways, the result isn't quite as accurate. Of particular note is trajectory A. Without its complementary return trajectory coming from the north as in figure 5.2, it misses its cue to arch a little higher and merge less abruptly with trajectory B. So, while including reversed drifters may slow convergence (at least in this example), the end product can be somewhat better.

The next two examples are from the same velocity field as the

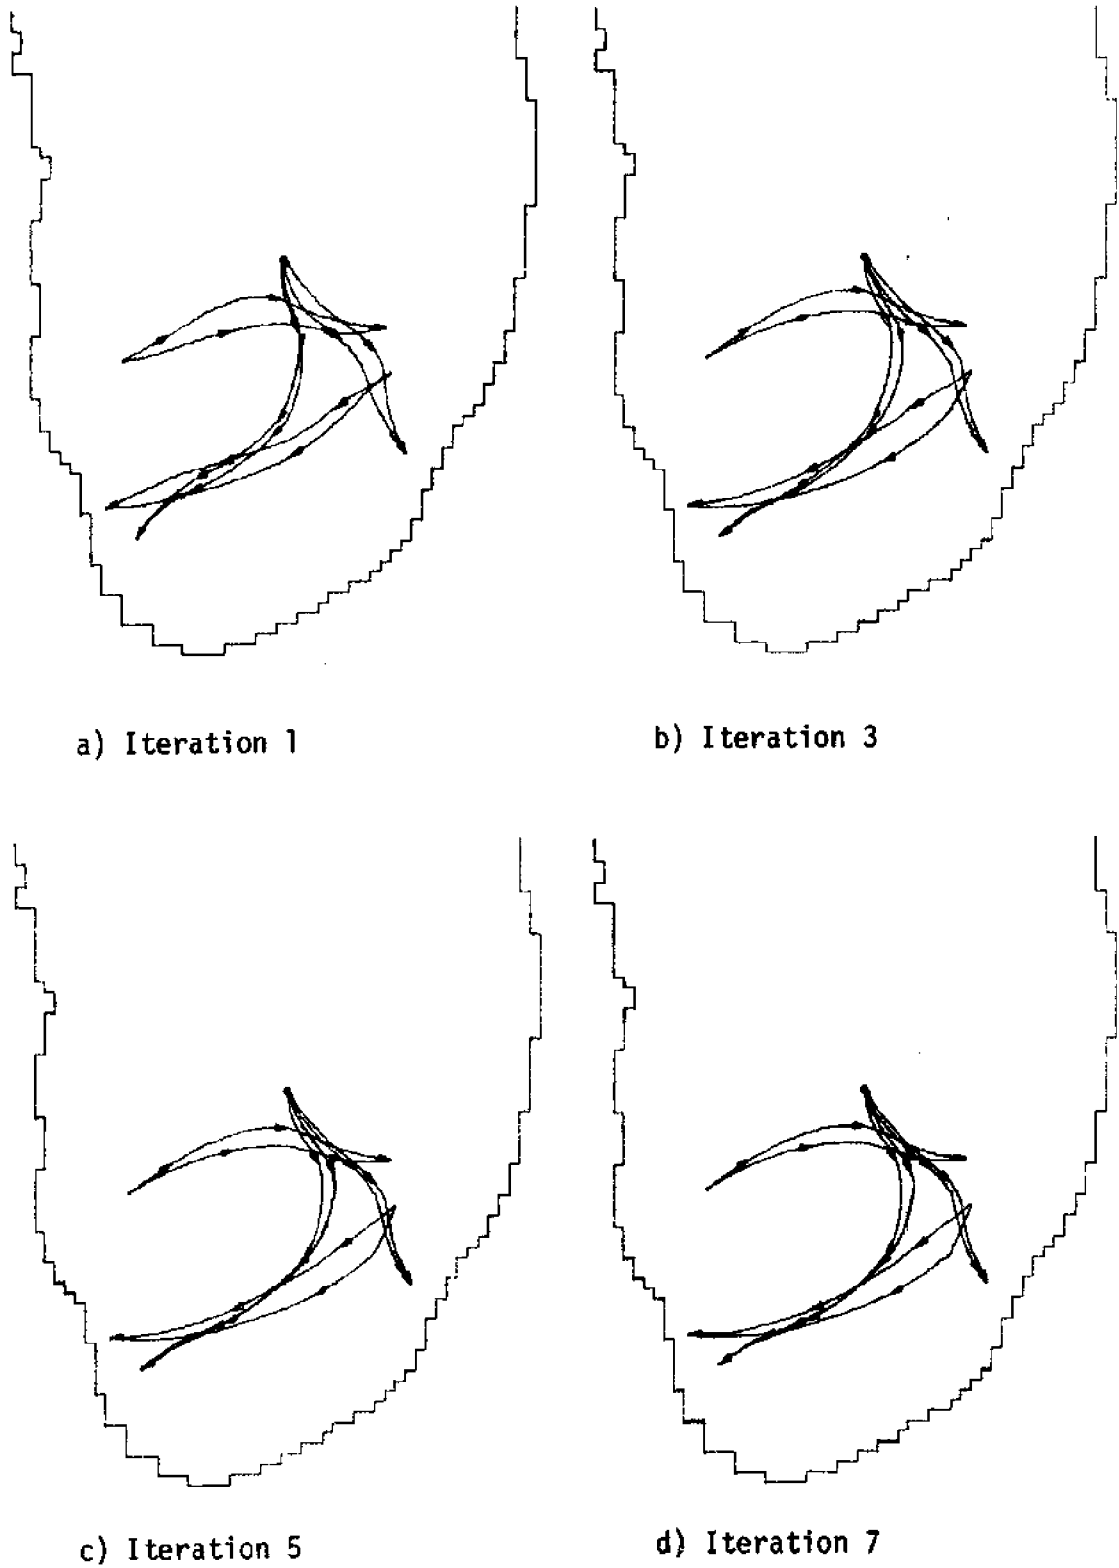





Figure 5.3: Inference results for less allowed acceleration. . . . /



e) Iteration 9




f) Inferred field (before FILL)




g) Inferred field (after FILL)

/ . . . 5.3, cont'd.



a) Iteration 1


b) Iteration 3




c) Iteration 5

d) Iteration 7


Figure 5.4: Inference results for less allowed acceleration and higher anti-diffusion values than standard.



e) Iteration 9




f) Iteration 11




g) Inferred field

/ . . . 5.4, cont'd.



a) Iteration 1


b) Iteration 3




c) Iteration 5

d) Iteration 7


Figure 5.5: Inference results for standard parameter values and forward (only) moving drifters. . . . /




e) Iteration 9



f) Iteration 13



g) Iteration 17



h) Inferred field

/ . . . 5.5, cont'd.

first set, but with more encompassing data (figure 5.6). This should allow the current reversal in the upper right corner to be picked up by the inference procedure. Starting from a blank field, nine iterations were run with the standard parameter values as shown in figure 5.7. At this point, FILL was used to fill in the blanks, and the result can be compared to the original (figure 5.7g and h). Except for the center location of the main gyre, agreement is about as good as could be expected from the available data.

The skewedness of the inferred gyre's center points up a characteristic of the system which will be more apparent in later examples. This is a *trunking* phenomenon in neighboring trajectories which tends to pull them together, thus flaring their ends and causing some unnecessary divergence. It results from the circulation heuristic's tacit treatment of undefined reference velocities as *zero* velocities, which tends to pull each drifter toward the defined velocities if this pull is not countered in undefined areas by an equivalent pull. The centers of the gyres in this example are vast undefined areas, so trajectories will bend away from the centers in response to external attraction, as is present here. A solution to this problem was not attempted in this investigation, although remedies are suggested in the next chapter.

At the ninth iteration, convergence had not truly arrived, and the small dipped trajectory on the extreme right continued to retract. Further iterations were made and the retraction continued, with the partially defined field at the thirteenth iteration shown in figure 5.7k.

This dip represents a rather high angular acceleration, so doubling the standard value of CIR should cause it to be accentuated. This was tried, and the results are shown in figure 5.8. In iteration three, an





Figure 5.6: Velocity field and recoveries used for second example.

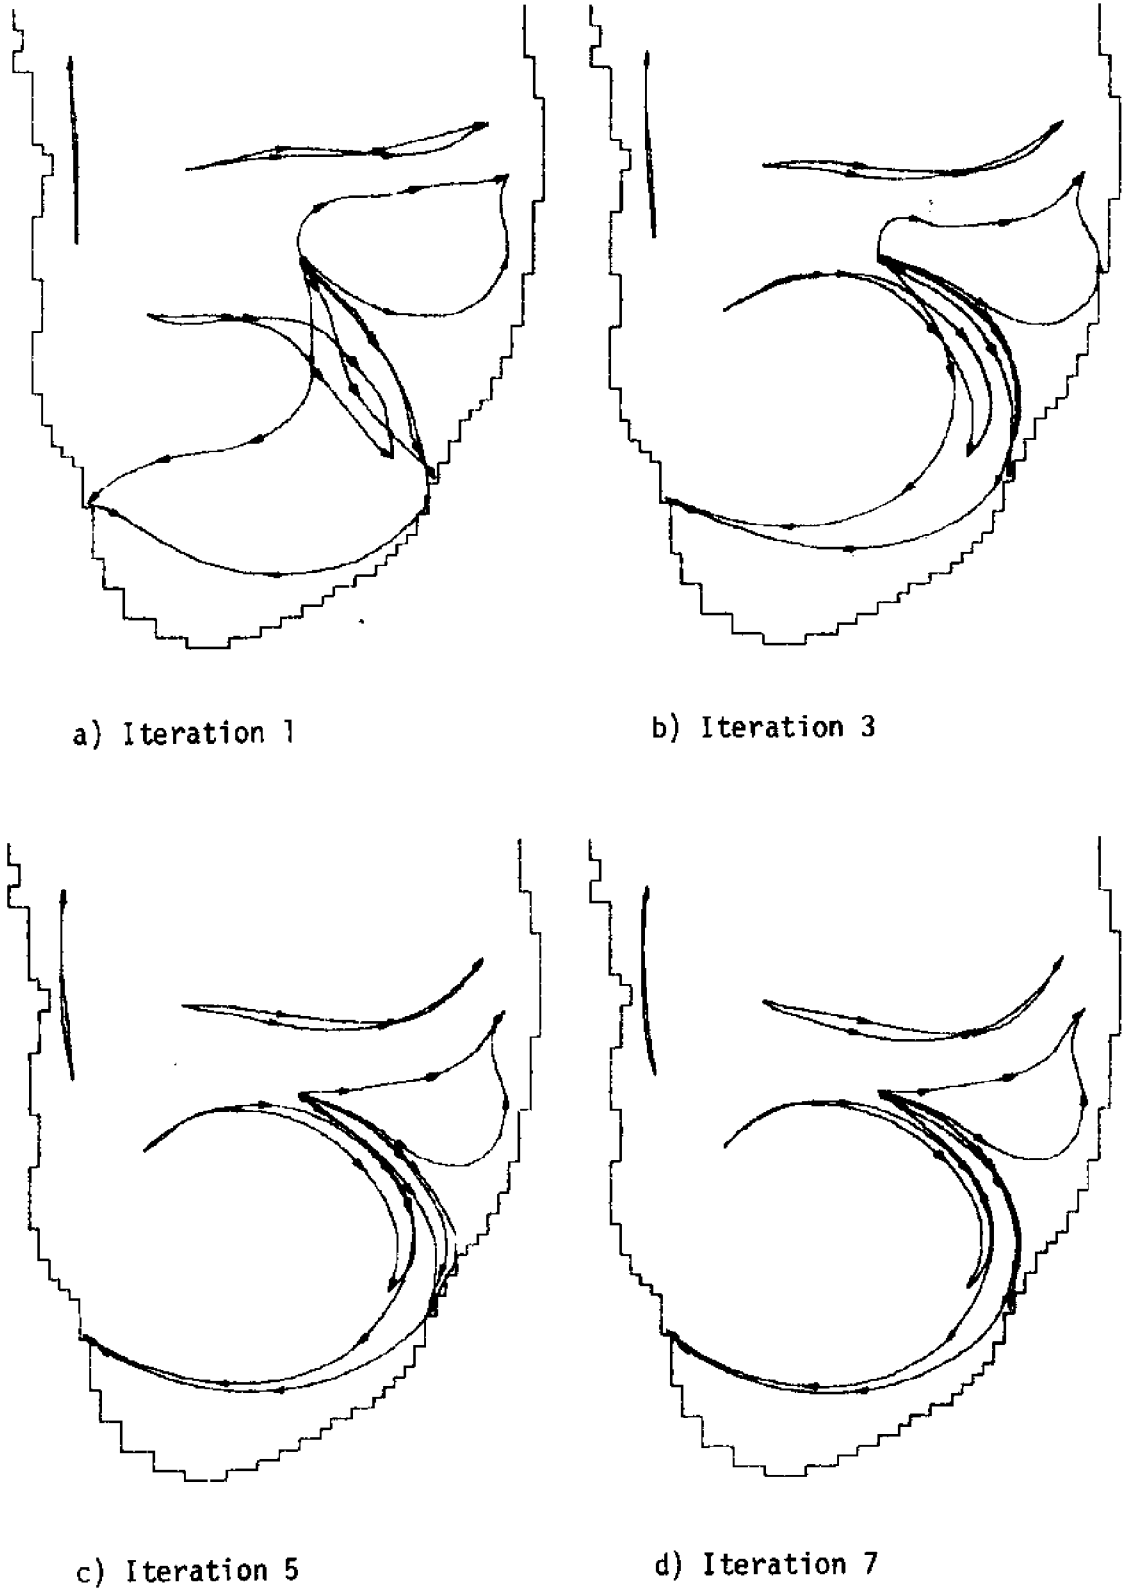





Figure 5.7: Inference results for standard parameter values. . . . /




e) Iteration 9



f) Inferred field




g) Inferred field (FILLED)




h) Original field for comparison

/ . . . 5.7, cont'd.


. . . /



f) Iteration 11




j) Iteration 13




k) Inferred field


/ . . . 5.7, cont'd.




a) Iteration 1



b) Iteration 3




c) Iteration 5




d) Iteration 7


Figure 5.8: Inference results for twice-standard allowed acceleration.
... /




e) Iteration 9



f) Iteration 11



g) Iteration 13



h) Inferred field


/ . . . 5.8, cont'd.

alternate path to the recovery site was tried, but the original one won out in the end, and again began retracting. This behavior may indicate an insufficiency of the data in this area to support the dip initially attempted. At any rate, the increased CIR value does not seem justified by the results.

This example is the first to exercise the shoreline heuristic. In the lower right hand corner, as the large circulating trajectories expand, the shoreline is encountered and the heuristic keeps the trajectories parallel to it, resulting in the parallel coastal current shown.

The last example was contrived to test the system and does not represent a serious hypothesis about Lake Michigan currents (figure 5.9a). It was created to check the system's performance on data obtained from a two-gyre velocity field under difficult acceleration conditions. The recovery data obtained from this field by simulation are shown in figure 5.9b-d. The inference was begun with a blank field and standard parameter values, except that AVG was boosted to 0.8. Representative iterations are shown and the pre- and post-FILL fields laid out in figure 5.10a-g. Figure 5.10h reiterates the original field for comparison purposes.

Qualitatively, agreement is rather good. Both gyres are firmly present. The inferred field has fewer sharp turns in it than the original, though, but that's not to its discredit. On the other hand, trunking is pronounced, particularly with regard to the small, inner trajectory left of center, which flares considerably on its right hand end. On the whole, nonetheless, the inference process seems to be doing the job intended and sufficient validity to use it on data from the field is indicated.



a) Original velocity field



b) Simulated recovery data
(and ff.)


Figure 5.9: Velocity field and recoveries used for last example.



a) Iteration 1



b) Iteration 3



c) Iteration 5




d) Iteration 7


Figure 5.10: Inference results for 4-times standard anti-diffusion.




e) Iteration 10



f) Inferred field



g) Inferred field (FILLED)



h) Original field for comparison

/ . . . 5.10, cont'd.

5.3 Inference Results from Lake Michigan

The drifter recovery data from the July and August, 1974, research cruises in Lake Michigan were used to infer current patterns for the southern basin, one for each set of data. The recoveries used were a subset of those selected for the testing algorithm (figures 3.12a and b, pp. 80-90). There are two reasons for the further culling process. First, for July, too many returns were recorded to be handled by the inference program on the machine available (IBM 1800, 32K core) without a significant reprogramming effort. Second, the algorithm is designed to handle out-of-bounds drifters only as an exceptional case. It was felt that applying it to those recoveries from outlying regions would stretch the allowances made for going out-of-bounds beyond their intended function. Any reasonable subset of the available data, though, should be acceptable so long as the results obtained therefrom compare favorably to other current patterns submitted to the hypothesis test under identical conditions.

The data chosen for the July and August inference runs are denoted by the solid diamonds in figure 3.12. The remaining points were eliminated first on the basis of whether they lay out-of-bounds (triangles in figure 3.12) and then on the basis of other factors (July only). The other factors include the elimination of redundant recoveries from distinct but neighboring releases and the discarding of recoveries having questionable dates.* Both the July and the August hypothesis generations were done with the standard parameter settings, the CIR value


*The scrutiny this entailed turned up one returned card whose date was smudged and was probably recovered ten days later than recorded. Though it was eliminated here, it was retained for the subsequent hypothesis test to permit valid comparison with the results of chapter 3.

corresponding to an acceleration of $3.5 \times 10^{-3} \text{ cm/sec}^2$.

Beginning with a null field and the selected July recoveries, the inference process was run to the sixteenth iteration. Every other iteration along the way is shown in figure 5.11a-h. As can be seen, most of the trajectories are pretty well settled by the last iteration, although a few seem not to have found their slot. For the most part, the tracks are pretty well consistent with each other, but the area around Waukegan ($42^{\circ}20'N$, $87^{\circ}40'W$) is a major exception. Here, the paths seem unwilling to bend to conformity, although the hint of a small clockwise gyre is present. Several factors may contribute to the uncertainty in this area. First, many of its recoveries came very soon after release. In such cases, positional accuracy is most important, since little freedom is possible in the inferred trajectories due to the implied acceleration. Many of these early recoveries, though, were made by boaters several miles offshore. While such data are the best time-wise, much has to be assumed about the boaters' knowledge of their locations, and errors are more likely than for beached drifters. Second, the wind at the time (see figure 3.18, p. 112) was rather capricious and may have had a greater role in upsetting stationarity in the velocity field than would be expected under steadier conditions. Third, the coarseness of the reference grid in the inference algorithm may prevent such small circulations from being recognized without more global supporting data.

The velocity field inferred from the data is shown in figure 5.12a and its completed (by FILL) version in 5.12b. Except for a strong counterclockwise gyre in the northeast corner of the area shown and a southward coastal current in the southwest portion, the circulation is basically clockwise. Only two significant areas of directional divergence


LAKE MICHIGAN



a) Iteration 1

Figure 5.11: Inferred trajectories for July, 1974, drifter data. . . . /

LAKE MICHIGAN



b) Iteration 3

/ . . . 5.11, cont'd.

. . . /

LAKE MICHIGAN



c) Iteration 5

/ . . . 5.11, cont'd.

. . . /

LAKE MICHIGAN




d) Iteration 7

/ . . . 5.11, cont'd.

. . . /

LAKE MICHIGAN




e) Iteration 9

/ . . . 5.11, cont'd.

. . . /

LAKE MICHIGAN



f) Iteration 11

/ . . . 5.11, cont'd.

. . . /

LAKE MICHIGAN



g) Iteration 13

/ . . . 5.11, cont'd.


LAKE MICHIGAN



h) Iteration 15

/ . . . 5.11, cont'd.


. . . /



a) Inferred velocity field

Figure 5.12: Hypothesis generated by inference procedure from July, 1974, drifter data. . . . /

LAKE MICHIGAN



b) Inferred velocity (FILLED)

/ . . . 5.12, cont'd.

are noticeable. One is directly to the east of the Waukegan area, where a stretching occurs between the vague gyre there and the overall trend. The other is in the extreme southern part of the lake where converging flows have nowhere to go. Although an offshore, northward return flow is probable, no data are available to support or deny it.

This velocity field was submitted to the hypothesis testing algorithm along with all the data used in the tests in chapter three. The results are plotted in figure 5.13 along with the curves from the previously considered hypotheses. The inferred field shows to be the most data-compatible of all in the entire speed range between 0.5 and 1.25.

The selected August data were submitted to the same treatment from hypothesis generation to testing. Convergence occurred much more quickly with fewer questionable areas, as can be seen from the trajectories plotted in figure 5.14a-e. The resulting velocity field (figure 5.15a and b) is generally clockwise with a splitting in the northeast corner of the area considered. Of significance in this field is the obvious divergence at the center of the circulation. This is probably due to the early assumption about the field velocity at each point being the average velocity of the drifters moving through it. This principle has actually been applied only to the *recovered* drifters moving through each point. But recovered drifters from a midlake release are special cases indeed if the current pattern is a simple gyre, since they've had to cross "many streamlines" to reach shore. Hence the divergence apparent in the inferred field but surely not present in actuality.

The testing results of the inferred hypothesis from August are plotted in figure 5.16 against those from chapter three. The inferred field is significantly more data-compatible than any of the others and





Figure 5.13: Hypothesis testing results for July, 1974, comparing generated hypothesis to those of chapter 3.

LAKE MICHIGAN




a) Iteration 1

Figure 5.14: Inferred trajectories for August, 1974, drifter data.

... /

LAKE MICHIGAN




b) Iteration 3

/ . . . 5.14, cont'd.

. . . /


LAKE MICHIGAN



c) Iteration 5

/ . . . 5.14, cont'd. . . . /

LAKE MICHIGAN




d) Iteration 7

/ . . . 5.14, cont'd.

. . . . /


LAKE MICHIGAN



e) Iteration 9

/ . . . 5.14, cont'd.


LAKE MICHIGAN



a) Inferred velocity field

Figure 5.15: Hypothesis generated by inference procedure from August, 1974, drifter data. . . . /

LAKE MICHIGAN



b) Inferred velocity field (FILLED)

/ . . . 5.15, cont'd.




Figure 5.16: Hypothesis testing results for August, 1974, comparing generated hypothesis to those of chapter 3.

tends to tip the scale in the clockwise versus counterclockwise question heavily on the clockwise side.

5.4 Summary

Three examples of hypothesis generation applied to simulated recovery data have shown its effectiveness in dealing with data of varying complexity. The necessity for the circulation heuristic is apparent from trials using it at a low level, as is the desirability of having drifters running both ways. In addition, the proper functioning of the shoreline heuristic has been demonstrated. Although a possible problem in the circulation heuristic is manifest by a trunking phenomenon in the trajectories, application of the inference process to data from the field seems justified.

Inferences drawn from both the July and August data fare rather well in the context of previously considered hypotheses, adding further credibility to the hypothesis generation scheme. Both inferred hypotheses indicate a clockwise circulation, although anomalies are present in each one. Differences between the two are a shift in the Chicago area coastal current from south to north as well as the elimination of the small reverse gyre in the northeast corner of the lower basin.

CHAPTER 6

CONCLUSION AND OUTLOOK

6.1 Nature of the Results

A model of surface current drifter transport and dispersion which is equivalent to the molecular diffusion or random walk formulation has been established. Two simulations are based on this model. The first is a Monte-Carlo simulation which requires a known or hypothesized current field and which is used for making predictions about drifter data. It is employed in an hypothesis-testing scheme for evaluating the compatibility of proposed velocity fields with a set of recovery data from the field. The second simulation is goal-directed and requires a set of recovery data to operate. It is used with an adaptive scheme for inferring a velocity field from the data. The results arising from the construction and application of each simulation are of both a methodological and an oceanographical (pertaining to Lake Michigan) nature. These two aspects are set forth separately.

6.2 Methodological Results

The primary methodological result from the hypothesis testing scheme is a measure of data compatibility which is directly relatable to diffusion. It measures, for each recovered drifter, the rate of deviation from an hypothesized velocity field required to get the drifter from its release to its recovery. The measure is most effective if recoveries which are redundant or returned later from a given recovery

site than others of the same release are eliminated. The measure is still biased, however, by the selection of release locations. This bias can be used to advantage for preselecting release points designed to distinguish optimally among a set of hypotheses. A method has been established using simulated recovery data for finding such diagnostic points for pairs of hypotheses. Finally, by extending the simulation to cover non-stationary velocity fields, hypotheses about the wind influence on currents can be tested. Experiments with the latter have shown the need to consider not just the compatibility metric in the evaluation of an hypothesis, but also whether it runs all the drifters aground prematurely.

The use of a diffusion model as the basis of a testing scheme solves some of the problems inherent in other testing schemes reported in the literature (Tomczak, 1968; Hill and Horwood, 1974). In cases where currents other than Eckman (1905) wind currents are present, dispersion can be augmented by heterogeneities in the velocity field. This sometimes results in a splitting of trajectories which neither a single-drifter simulation (Tomczak) nor a normal cloud simulation (Hill and Horwood) can properly take into account. A multi-drifter Monte-Carlo simulation does account for such behavior though, and when combined with a compatibility measure which further accounts for it by considering only the simulated drifter closest to a given recovery at the recovery time, provides an effective means for dealing with complex hypotheses.

The important technical aspects of the inference process are its goal-directed structure and the use of iteratively applied adaptation on the velocity field to yield the inferred hypothesis. The goal-directed approach is a natural one for solving the type of boundary value problem that drifter data provide, since trajectory end points

are all it nominally needs to run. In actuality, a purely goal-directed simulation which works well would be tough to design because of the complex interaction of the drifters implicit in this framework. The significance of adding adaptation lies in the decoupling of individual drifters from each other, but allowing them to communicate via the changing velocity field. The connectedness thus lost in the spatial (drifter-to-drifter) domain is made up for as necessary in the temporal (iterative) domain. The complexity of the inference problem thereby becomes linear in the number of recoveries considered rather than in its square, so long as the requisite number of iterations is independent of the number of recoveries. This it seems to be, as convergence usually approached near the tenth iteration in examples having from four to fifty recoveries.

The use of an heuristically guided decision algorithm in the goal-directed simulation seems to have paid off too. The isolation of the decision process from the rest of the system facilitates modification as new ways of satisfying the hydrodynamic constraints are discovered. The heuristics used (straight-line and circulation) have shown to be adequate for unraveling conflicting trajectories in several difficult cases, including data from the field.

More specifically, several results have come out of the individual techniques used. First, the advantage of running drifters both directions in time has been argued and demonstrated. Second, the use of the straight-line heuristic alone with the adaptive element corresponds roughly to inferential techniques appearing in the literature, depending on the diffusion assumptions made. By setting the anti-diffusion coefficient AVG very low, the velocity at the release points will be the aver-

age of the straight-line velocities of the drifters released there. This compares closely with the approaches taken by Wyatt, *et al.* (1972) and Norcross and Stanley (1967). By setting it high, the reverse trajectory construction techniques are approximated (Bukin, 1974; Bumpus and Lauzier, 1965). Finally, the consideration of global, circulatory aspects in an evolving velocity field has provided the computational boost necessary to counter the problems inherent in the pure straight-line approach.

This inference process is not unique in its use of adaptation. The algorithm described by Pasquay and Bonnot (1971) is also adaptive, but is not coupled to a goal-directed simulation. Modifications to the velocity field are only made after an entire run of each drifter, which moves as a function of the wind and an underlying field velocity. The corrections made to the field velocities as a result of a simulated drifter's missing its recovery point do not consider circulation.

6.3 Results from Lake Michigan

The hypothesis-testing results for Lake Michigan during July, 1974, heavily favor the velocity field generated from the return data by the inference process (figure 5.11b). This field has a predominantly clockwise trend with a prominent counterclockwise gyre appended in the vicinity of 43°N , $86^{\circ}30'\text{W}$. This gyre, rather than the coastal counter-current conjectured in *Coastwise Currents* (Monahan and Pilgrim, 1975), seems to account for the wide, time-uncorrelated spread in the recoveries along the eastern shore. Near the western shore at $42^{\circ}30'\text{N}$, there is the possibility of a clockwise gyre, but such an anomaly would go against the grain of the main circulation, causing the divergence apparent in the generated hypothesis. It is also probable that the coastal

current along the southwest shore (near Chicago) was southward during the prime recovery weeks, though no data are available to indicate a return path from the southernmost extremity of the lake. Of the published hypotheses made prior to the data, Kizlauskas and Katz' (1973) result from a finite difference model (figure 3.13e) is the most compatible with the data obtained.

For August, the hypothesis generated from the data (figure 5.13b) is, again, the most data-compatible. Its trend is predominantly clockwise, the only exception being a north-south splitting on the northern reaches of the eastern shore. Several significant changes are apparent between July and August, if these generated hypotheses be true. The prominent counterclockwise gyre in the northeastern part of the southern basin has either been replaced by a splitting near the shore, or else the splitting is simply its lower extremity, the gyre having moved north. In addition, the surface coastal current near Chicago underwent a reversal during this period from south to north. Finally, the current anomaly near the western shore seems to have disappeared between July and August, although it's hard to tell since no releases were made in that area during August.

As for the Eckman wind effect, the results are rather inconclusive. The testing results alone give some credence to a wind hypothesis, though not as much as to the inferred, stationary hypotheses. With the additional observation that drifters simulated with the Eckman currents calculated for this period run into shore *en masse*, one is led to suspect that Eckman's theory is inapplicable in nearshore areas and possibly in offshore regions as well.

6.4 Work To Be Done on the Analysis Technique

Several improvements in the testing and inference techniques suggest themselves immediately from using the system. First, the problem of all the simulated drifters running aground prematurely, though insignificant for the stationary hypotheses considered, needs to be eliminated. A possible method for doing this would be to replace each newly beached drifter at its previous position in the water. This would keep all the drifters active for positional comparison at all times. Secondly, some technique for positionally unbiaseding the compatibility measure when desired to do so could be useful. Unfortunately this could involve some rather extensive record-keeping of the positions of each drifter simulated in the testing process in order to come up with an estimate for how many other drifters used in comparisons have been where it has been. A simpler technique based on release points alone might be more feasible.

In the inference process, the circulation heuristic could be improved to reduce the trunking effect. It would have to estimate or have estimated for it the reference velocities for undefined reference points in the field. FILL might be used between iterations to accomplish this, but it takes rather long to execute to use that often. It might also be desirable to separate shear considerations from divergence considerations in the decision algorithm. This task would be less easy, since they are both closely related to the scale of circulation. Finally, the position calculations of the reference points should be dispensed with. They are unnecessary, considering the effort gone to to disguise the actual positions afterwards. The reference points could merely be assumed to occupy the centers of their respective boxes whenever a distinct position is necessary. This calls for more careful consideration in the near field,

though, of a reference velocity defined at a distinct point and probably implies an estimate for the integrated effect over the entire box.

6.5 Work To Be Done in Lake Michigan

Two drifter studies are inadequate to cover the complex current behavior of Lake Michigan. It is almost trite to suggest "more experiments", but a couple specific ones readily suggest themselves. First, the northeast corner of the southern basin deserves intensive investigation. A program of concentrated drifter releases there over several months, coupled with other measuring techniques, might determine more closely the waxing and waning character of the gyre that seems to occupy that area. Second, the effect of the wind on the surface currents in the lake deserves more scrutiny too. In such a confined area, drifters are probably not the best means of investigating this effect, but a concentration of far-offshore releases could prove beneficial in moderating the effects of non-Eckman, nearshore currents on the recovery data.

6.6 Future Directions

The most fertile area for development of the methods presented here lies in furthering the man-machine interaction between an experienced scientist using the inference system and the system itself. Although the system, as programmed, allows some intervention to change parameter settings, much could be gained by letting a human operator correct mistakes in the trajectories inferred or suggest new lines of attack to the system during the course of its running. This would allow the powerful, intuitive capabilities of a human to be amplified by the computer's rapid computational ability, resulting in a significantly more useful investigative tool.

BIBLIOGRAPHY

- Arfken, George (1970). *Mathematical Methods for Physicists*. Academic Press (New York).
- Ayers, J.C., D.C. Chandler, G.H. Lauff, and C.F. Powers (1958). *Currents and Water Masses of Lake Michigan*. University of Michigan Great Lakes Research Division Publication No. 3 (Ann Arbor)
- Box, George E.P. and George C. Tiao (1973). *Bayesian Inference in Statistical Analysis*. Addison-Wesley (Boston).
- Brucks, J.T. (1971). "Currents of the Caribbean and Adjacent Regions as Deduced from Drift-Bottle Studies," *Bulletin of Marine Science--Miami* 21 (2), pp. 455-465.
- Bukin, V.M. (1974). "Special Postcards for Gathering Information about Surface Currents in Water Bodies," *Meteorologiya i Gidrologiya* 2. Soviet Hydrometeorological Service (Moscow). Trans. U.S. Joint Publications Service (Arlington, Va.), pp. 94-98.
- Bumpus, D.F. and L.M. Lauzier (1965). "Surface Circulation on the Continental Shelf off Eastern North American between Newfoundland and Florida," *Folio 7, Serial Atlas of the Marine Environment*. American Geographic Society (New York).
- Csanady, G.T. (1973). *Turbulent Diffusion in the Environment*. D. Reidel (Dordrecht, Holland).
- Eckman, V.W. (1905). "On the Influence of the Earth's Rotation on Ocean Currents," *Arkiv för Matematik, Astronomi och Fysik* 2 (11), pp. 1-52.
- Federal Water Pollution Control Administration (1967). *Lake Currents. Water Quality Investigations: Lake Michigan Basin*. FWPCA Great Lakes Region (Chicago).
- Feller, William (1967). *An Introduction to Probability Theory and Its Applications*. Wiley (New York).
- Fisher, R.A. (1959). *Statistical Methods and Scientific Inference*. Oliver and Boyd (London).
- Hachey, H.B. (1935). "The Circulation of Hudson Bay Water as Indicated by Drift Bottles," *Science* 82 (2125), pp. 275-276.
- Harrington, M.W. (1895). "Surface Currents of the Great Lakes as Deduced from the Movements of Bottle Papers during the Seasons of 1892, 1893, and 1894," *U.S. Department of Agriculture, Weather Bureau, Bulletin B*.

- Hill, H.W., and J.W. Horwood (1974). "A Computer Simulation of Surface Drifter Returns," *Journal du Conseil* 35 (2). Conseil International pour l'Exploration de la Mer. pp. 158-164.
- Holland, John H. (1975). *Adaptation in Natural and Artificial Systems*. The University of Michigan Press (Ann Arbor).
- Kizlauskas, A.G. and P.L. Katz (1973). "A Two-Layer Finite Difference Model for Flows in Thermally Stratified Lake Michigan," *Proceedings of the 18th Conference on Great Lakes Research*. International Association of Great Lakes Research, pp. 743-753.
- Lederer, Muriel (1970). "Letters from the Sea," *Oceans* 3 (5), pp. 31-33
- Monahan, E.C., P.C. Hawkins, and E.A. Monahan (1974). "Surface Current Drifters: Their Evolution and Application," Michigan Sea Grant Program Technical Report MICHU-SG-74-603 (Ann Arbor).
- Monahan, Edward C., and Elizabeth A. Monahan (1973). "Trends in Drogue Design," *Limnology and Oceanography* 18 (6), pp. 981-985.
- Monahan, Edward C., and Philip C. Pilgrim (1975). "Coastwise Currents in the Vicinity of Chicago, and Currents Elsewhere in Southern Lake Michigan," University of Michigan Technical Report (Ann Arbor).
- National Oceanographic and Atmospheric Administration (U.S. Department of Commerce). Surface Weather Observations, July 15 - October 31, 1974; National Weather Service: Muskegon, Michigan; Milwaukee, Wisconsin; Chicago, Illinois.
- Neumann, Gerhard, and Willard J. Pierson, Jr. (1966). *Principles of Physical Oceanography*. Prentice-Hall (Englewood Cliffs, NJ).
- Norcross, J.J. and E.M. Stanley (1967). "Inferred Surface and Bottom Drift," in Harrison, W., J.J. Norcross, N.A. Pore, and E.M. Stanley, "Circulation of Shelf-waters Off the Chesapeake Bight," *ESSA Professional Paper* 3, (Washington, D.C.), pp. 11-42.
- Pasquay J.-N. and J. Bonnot (1971). "Utilization de Cartes-Flotteurs pour l'Etude des Derives de Surface et Application à la Prevision des Pollutions Côtières," *La Houille Blanche* 26 (8), pp. 769-778.
- Shannon, L.V., G.H. Stander and J.A. Campbell (1973). "Oceanic Circulation Deduced from Plastic Drift Cards," *Sea Fisheries Branch Investigational Report No. 108* (South Africa), pp. 1-31.
- Standar, G.H., L.V. Shannon, and J.A. Campbell (1969). "Average Velocities of Some Ocean Currents as Deduced from the Recovery of Plastic Drift Cards," *Journal of Marine Research* 27 (3), pp. 293-300.

- Tomczak, G. (1968). "Investigations with Drift Cards to Determine the Influence of the Wind on Surface Currents," in *Studies in Oceanography* (Tokyo), pp. 130-139.
- von Arx, W.S. (1962). *An Introduction to Physical Oceanography*. Addison-Wesley (Boston).
- Wyatt, B., W.V. Burt and J.G. Paltallo (1972). "Surface Currents Off Oregon as Determined from Drift-bottle Returns," *Journal of Physical Oceanography* 2 (3), pp. 286-293.

SUPPLEMENT

S.1 Prefacatory Remarks

This supplement outlines the implementation of the system described in the main body of the report with emphasis on the graphic input/output utilized. Also, it contains a case study of a further drifter experiment carried out in Lake Michigan during July, 1975.

In addition to those individuals thanked in the Acknowledgements, I am grateful to Rick Boyce and Jane Matthews for handling the drift cards as they came in from the 1975 releases, establishing the recovery coordinates, and sending acknowledgements to the finders. I would also like to thank Paul Lambarth of Ann Arbor Aero Service for his help in planning and executing the flight over Lake Michigan for the drifter releases and Mark Wagner for his skillful piloting. Finally, I am indebted to Mike Geary and the pilots of Wolverine Aviation who, although they were not able to fly the mission over Lake Michigan, provided me with several flights for test-dropping the drifter release package.

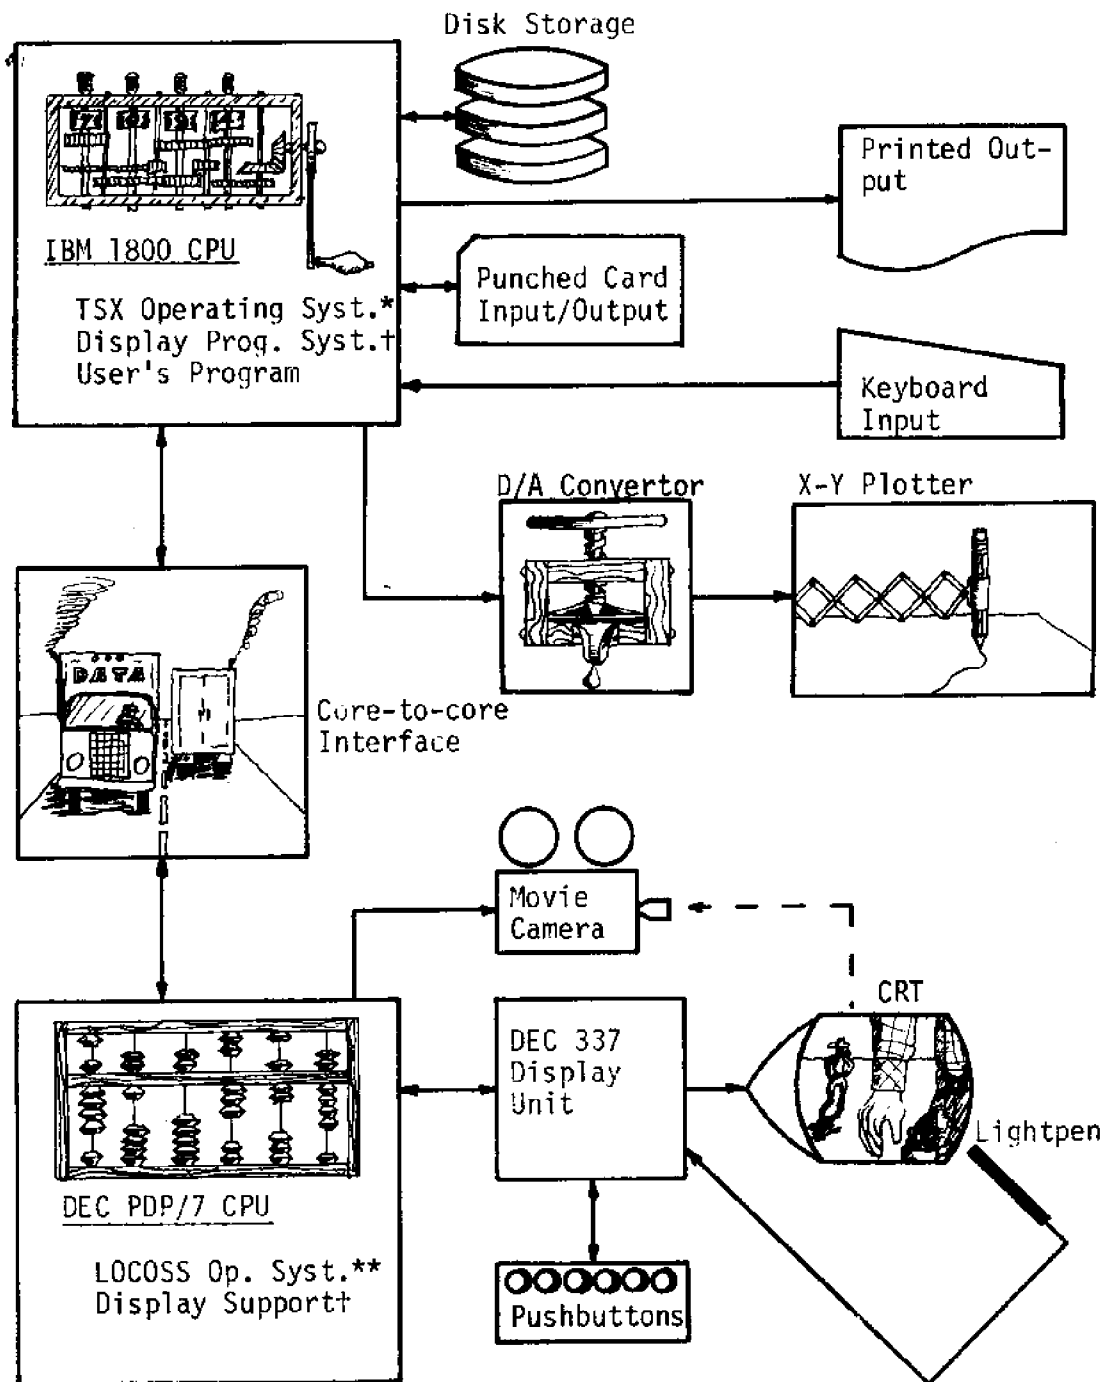
S.2 Use of the Computer in Drifter Data Analysis

The computer work described in this report was carried out in the computation lab of the Logic of Computers Group of the Department of Computer and Communication Sciences. This lab houses an IBM 1800 processor with a 32K core, card reader/punch, line printer, disk drives, and a D/A convertor which can drive an X-Y plotter. Connected to the processor through a high-speed core-to-core interface is a PDP7 with an attached DEC 337 display, push-button box, and lightpen. To the PDP7

may be connected a movie camera for recording output from the display. Under the drifter system, the PDP7/337 was used only as an "intelligent" terminal for graphic I/O, the calculation being carried out in the 1800. Figure S.1 illustrates the relationships among the various components and the software minimally resident in each computer for the drifter system.

The major steps in doing a drifter study, from a computational point of view, are 1) defining the geographical area studied--that is, establishing the land and water areas in the grid and outlining the shore for display purposes, 2) defining hypothetical velocity fields in the computer for use in the hypothesis-testing scheme, 3) generating test data, 4) culling the experimental data, 5) performing hypothesis tests, 6) performing inferences from the data, and 7) monitoring and outputting the results. Steps 4, 5, and 6 are pretty well covered in the main body of the report; the rest will be dealt with in sequence here.

In defining the area studied, an array of 4096 squares (64×64) must be partitioned into land and water areas and the dividing line (shoreline) located. To do this, the interactive capabilities of the 337 display are used to full advantage. Rather than defining the shoreline and computing the land/water function, land masses are entered (beginning with water everywhere), and the shoreline is determined from them. These land masses are entered on the CRT screen with the lightpen as blocks of 16×16 , 8×8 , 4×4 , 2×2 , or 1×1 squares using a moveable and alterable, prototype block. The size of this block is selected or altered by touching the lightpen to one choice in a "menu" of items displayed on the screen, and its location




*International Business Machines (1970)
 **Frantz, *et al.* (1968)
 †Pilgrim (1974)

Figure S.1: The Logic of Computers Group's computing system.


determined by touching the pen to one point in a grid of illuminated points or to one of four points orthogonally neighboring the displayed block. By the latter method, a block can be translated across the screen to its desired position. Once positioned, the implied land mass may be fixed by selecting the "define" option from the menu, at which instant a second, stationary copy of the block is displayed on the screen and the data base updated to reflect the new land mass. From there the defining block may be translated or resized and used to define further land masses. If a mistake is made, all the land mass in any of sixteen square regions (4×4) may be erased and redefined. The shoreline is determined by locating a 1×1 block over a square of land neighboring the water and selecting the "shore" option from the menu. At this point a line is drawn on the screen which is gotten algorithmically by following the land/water interface with the land on the right, much as a blindfolded person would follow a wall with an outstretched hand, until the starting point or an out-of-bound area is reached. This must be done for each isolated land mass. Finally, the land/water function and the shoreline information may be stored on a disk file for subsequent reference by selecting the "write" option. The sequence of events for defining a simple lake is sketched in figure S.2.

Velocity fields can also be defined by entering them with the display unit and the lightpen. The process begins with an initially blank (zero velocities, zero weights) field. Velocity vectors are entered by "drawing" them on the screen with the aid of a vector-cursor displayed on the screen (figure S.3). The cursor has two modes, translate and alter, which can be switched by a menu selection. In the translate mode, the cursor may be moved by "towing" it with the lightpen.



- 1 - 16 Set land mass size
- D Define land mass
- X Destroy land mass in area
- S Follow and outline shore
- W Write result to disk
- ① Lightpen, step 9

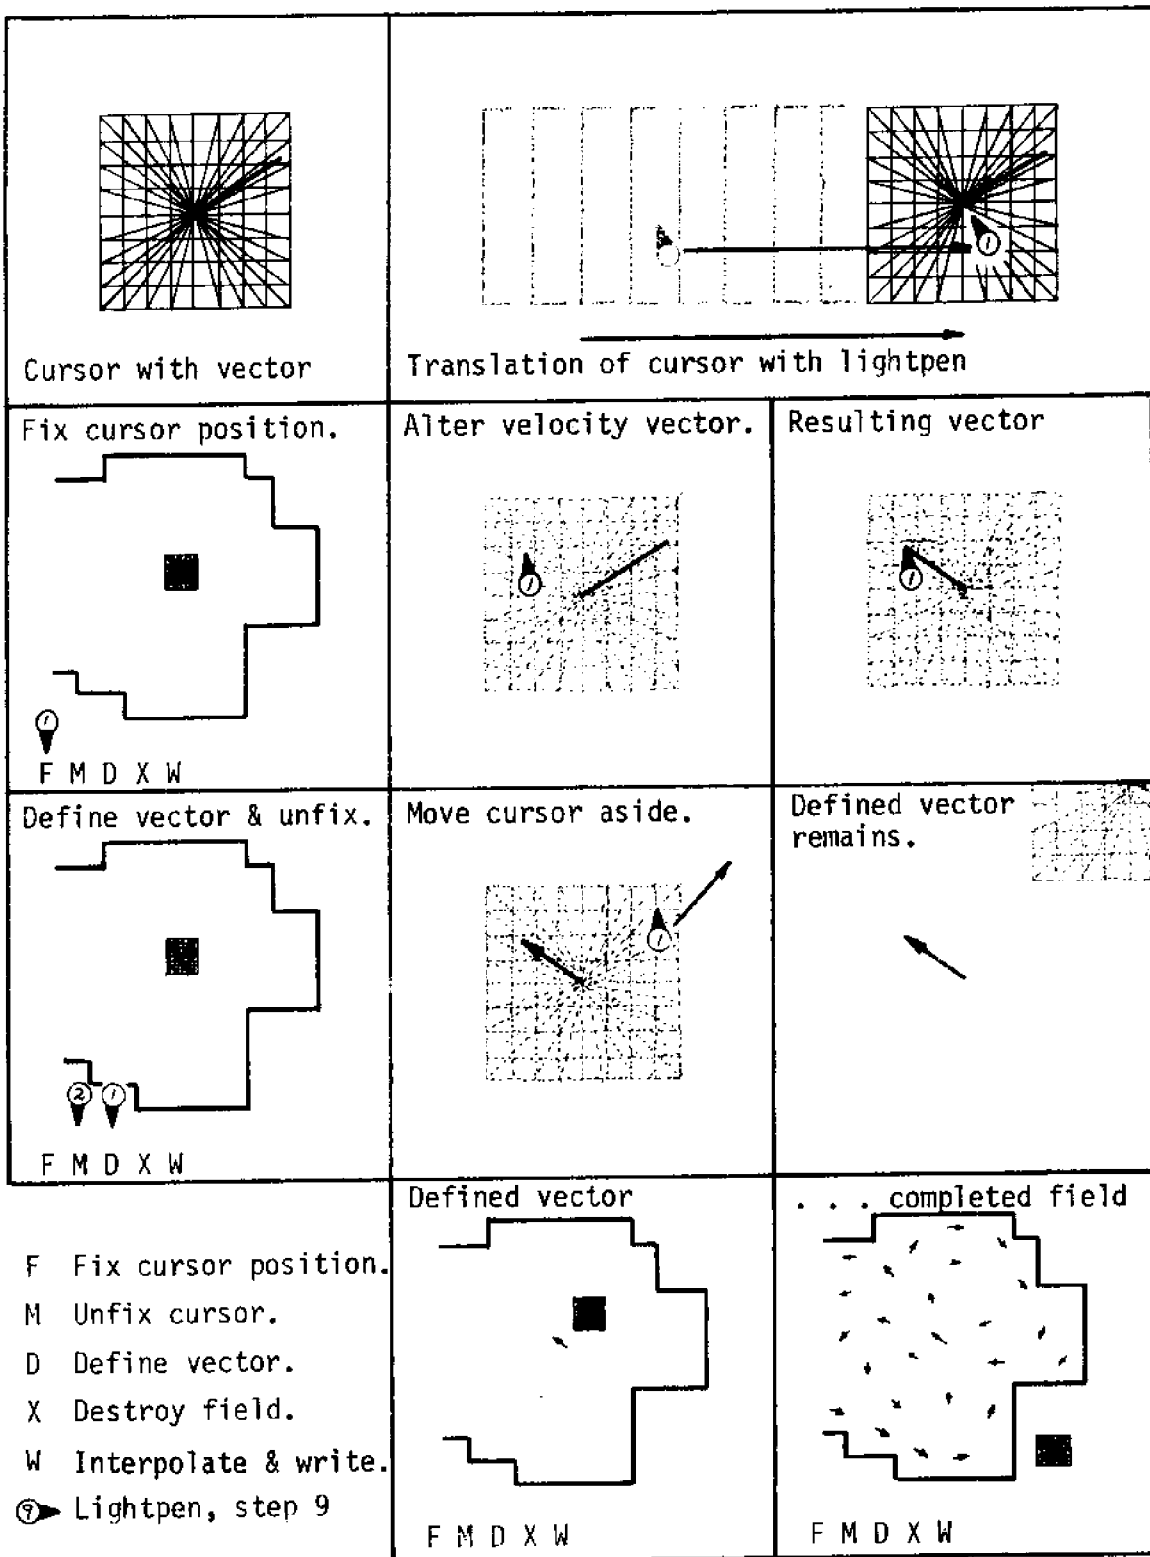
Figure S.2: Definition of a simple lake on the display using the lightpen. . . . /



/ . . . S.2, cont'd.

This is accomplished algorithmically by recentering the cursor under the lightpen whenever the lightpen is triggered by any part of it. Once the cursor is centered over a point at which a velocity is to be defined, the alter mode can be selected. The vector in the cursor can then be modified to agree with the desired velocity by touching the lightpen to any part of the cursor, at which time the head of the vector moves to the point under the pen. When the desired vector is thus established, the "define" menu option is selected and a fixed copy of the vector is displayed on the screen. At the same time, the velocity field at the cursor position is altered to reflect the newly-defined velocity value. At this point the cursor may be moved to a new location, and so on. If an error is made, all the vectors in any of sixteen (4×4) square regions may be erased and re-defined. Finally, when the velocity field is adequately represented by drawn-in vectors, the "interpolate" option is selected and the as-yet-blank squares are filled in. Then the resulting, totally defined field is written to a disk file. Figure S.3 illustrates several steps in the definition of a velocity field.

The generation of test data for the development and demonstration of the inference system was also done interactively with a program using a simple command language geared to the push-buttons of the 337 display. With this program a release point may be established at any point in the displayed body of water and one to one hundred drifters released there. The drifters appear on the screen as points and move as the simulation algorithm is activated from transition-to-transition. The simulation can be stopped at a pre-established point or interrupted and continued later. During any such interim the lightpen may be used to select drifters to act as recoveries. For each selection, a card is punched with



- F Fix cursor position.
- M Unfix cursor.
- D Define vector.
- X Destroy field.
- W Interpolate & write.
- ☉ Lightpen, step 9

Figure S.3: Several steps in the definition of a velocity field.

the release and recovery locations and intervening time, which may later be used by the inference program. Also, during any interim, the current velocity field may be displayed or a different release point defined. By means of this interactive program, the user has considerable control over the generation of test data.

All drifter motion, simulated or inferred, can be monitored on the 337 at all times. Such a feature enables the user to know when errors have arisen and when intervention is needed. This monitoring is accomplished by a single subroutine which is called every time a transition occurs. It displays drifters either as distinct points, or as line segments connecting previous positions to current ones, thus forming trajectories. In addition, it automatically handles the drawing of the shoreline and can display a set of vectors representing any velocity field extant in the system. Everything displayed by this monitoring routine can be double-buffered, so that a displayed picture can remain on the screen while another is being prepared. Finally, the routine can be made to trigger the movie camera for a cinemagraphic record of anything displayed.

Hardcopy graphic output can be tedious to produce if it has to be done as the results to be plotted become available, since the X-Y plotter requires constant attention with respect to changing the paper, starting the servos, etc. In order to alleviate the waiting between plots, complete drifter histories and velocity fields may be saved on a disk file as they evolve. They may then be retrieved later and plotted at the user's convenience in rapid sequential fashion. As in the

monitoring routine, the plotting routine plots shorelines*, velocity fields, drifters, or trajectories to any scale.

A summary of the drifter system implementation is sketched in block form in figure S.4.

S.3 Summer, 1975, Drifter Experiment in Lake Michigan--A Case Study

On July 8, 1975, 970 surface drifters were released in Lake Michigan from a plane making several passes over the southern basin. The intent was to get as broad a coverage of this area as possible in order to obtain a global picture of the current patterns therein. In this section the considerations and steps leading to the release are discussed, as well as the results obtained by an analysis of the consequent recovery data.

In order to execute a broad, uniform pattern of drifter releases, a raster involving considerable lineal distance must be followed. Doing this by boat in the lower basin of Lake Michigan would take several days. By using a plane, however, the entire area may be covered in a few hours and at a fraction of the cost of a fully-equipped lake-going research vessel. For these reasons, it was decided to release the drifters involved in the 1975 Lake Michigan study from the air.

The central problems associated with an airborne release are:

- 1) accurate positioning of the release points, and 2) precise deployment of the drifters at the known release locations. The first is a problem of navigation. It is critical over a body of water when flying at low altitudes because the usual VHF navigational aids soon disappear

*The plotted shorelines in this report are the ones with the squared corners. The others were hand-drawn and offset printed as blank forms for use with the plotter.

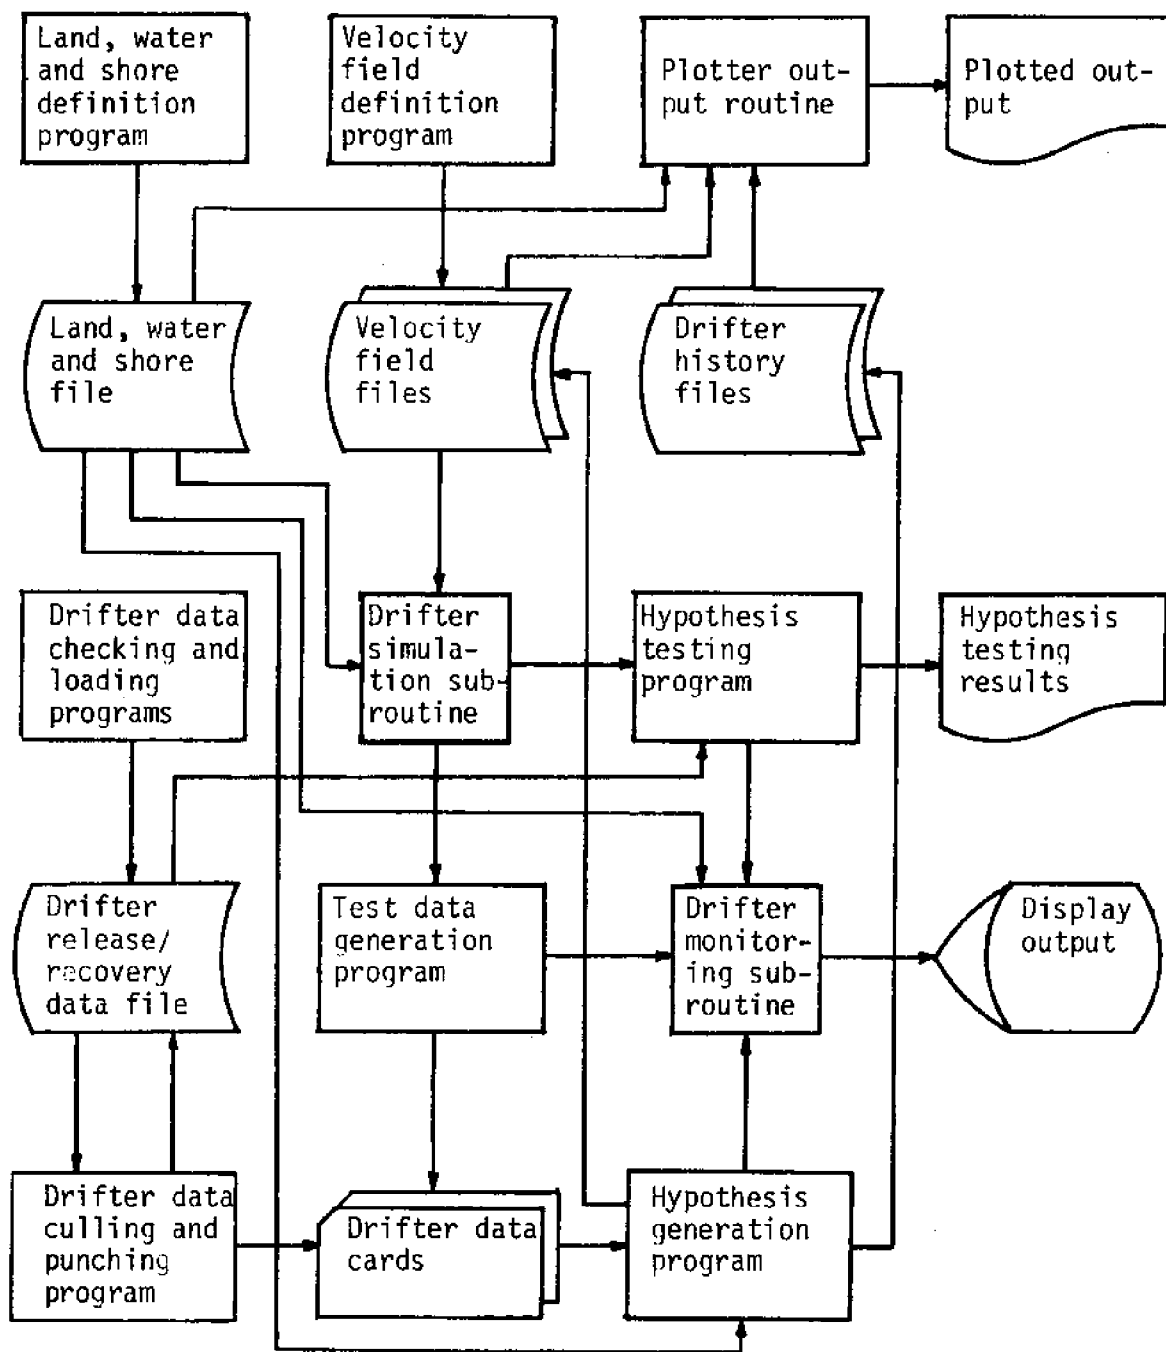


Figure S.4: Organization of drifter system.

over the horizon. The second problem involves getting the drifters from the plane to the water without their scattering in the air before hitting the water's surface. This problem will be addressed first.

The drifters used in the Lake Michigan experiments are made of a plastic paper (see Monahan, *et al.*, 1974), each one printed on both sides of a $3\frac{1}{4} \times 11$ " sheet (figure S.5) and stapled into a drum. The problem of their scattering before hitting the water is significant because they are so light. This problem can be alleviated, however, by either of two release methods. First, the plane can fly very low and very slowly to simulate, as closely as possible, release of the individual drifters from a ship. Second, the drifters can be bundled in packages designed not to come apart until they hit the water, the bundles being released from a higher altitude and at a higher speed. The first method is not only slightly dangerous (at least from a charter company's point of view), but could interfere with accurate navigating; hence, the second technique was chosen.

Any drifter bundle released from a plane has to satisfy certain criteria, namely:

1. The bundle must be small enough to fit through the plane's release port,
2. the bundle must be large enough not to permanently deform the drifters contained therein,
3. the bundle must not come apart in the plane's slipstream or anywhere else in the air, and
4. the bundle must come apart after entering the water to deploy the drifters reliably.

Return Postage Will Be Paid By—
CIRCULATION PROJECT
DEPT. of ATMOS. & OCEANIC SCIENCE
 Space Research Building
 The University of Michigan
 Ann Arbor, Michigan 48105

BUSINESS REPLY CARD
 No postage necessary if mailed in United States

FIRST CLASS
 Permit No. 1949
 Ann Arbor,
 Michigan



Figure S.5: Sample drifter.

OBVERSE

REVERSE

PLEASE REMOVE CARD, LET DRY IF WET,
 FILL IN INFORMATION AS REQUESTED,
 AND SEND BY MAIL.

U.M.S.G. No 10966

Where found (name of beach or place on shore, near what landmark in what town or city,
 on what lake).....

When found (date).....

Your Name (print).....

Your Home Address (print).....

Your return of this card will assist the addressee in a study of Lake Currents. The location
 and date of release will be sent to finder upon return of the card. **THANK YOU!**

CUT ALONG DOTTED LINES BEFORE MAILING

This drifter which you have found was launched by scientists from the University of Michigan as part of a study of currents in southern Lake Michigan. This study is a joint undertaking of the University of Michigan Sea Grant Program and the City of Chicago.

Drifters such as this act as tracers of the motion of lake water. Each card returned provides additional information that helps in the development of a picture of the current patterns that appear in Lake Michigan under certain wind and temperature conditions. The assistance you provide by filling in and mailing the attached postcard is gratefully acknowledged.

After several test flights involving the release of drifters bundled in various ways, the package shown in figure S.6 was settled upon. In practice it operates as follows: The box is introduced into the plane's slipstream through a small vent window next to the pilot's seat. It is kept together in the slipstream by the two rubberbands, joined in the middle by a Lifesaver.* When the box hits the water, the sand ballasting ensures that the hole sinks below the surface, allowing both the Lifesaver and Froot-loop[†] to dissolve. This takes place in about fifteen minutes, at which point the two halves of the box fall away and the drifters disperse. Although the drifters must be deformed somewhat to get twenty-five or so into a single box, they do tend to their original drum shape after being in the water a day. Even if they become creased badly, they will still present a broad cross-section to the current because of the convoluted pattern of folding. The boxes themselves were dropped from a plane over land to test their strength both aloft and on impact and put in the water to test their deployment reliability. In each respect their performance was judged more than adequate.

The problem of accurate positioning has to be tackled by a mixed bag of techniques. Nearshore, the VOR/DME stations can be used for an accurate fix, but further offshore below 1500 feet, say, they cannot be picked up, and dead-reckoning must be relied upon. Because of the wind, however, one cannot be certain of his position during the first couple

*It was found that fruit-flavored Lifesavers were much stronger than the mint-flavored ones, especially in humid weather, although they became quite sticky.

[†]Froot-loops were used in lieu of Cheerios because of their better resistance to the humidity in the air.




Figure S.6: Blowup of drifter deployment package.

transects of the lake except by retrospective correction. After that, though, the wind can be accurately compensated for and positioning should be possible with a half-mile square.

The deployment and navigational methods having been established, thirty-nine boxes were made, stuffed with twenty-five drifters each (with one exception) and marked on the outside with a code which could be quickly transcribed into a release log during the flight. When the first conjunction of an available plane and good weather came around, the flight was made, taking less than seven hours total. The log for this flight is shown in figure S.7; the release points, in figure S.8.

In all, 93 of the 970 released drifters were returned. A graph showing the cumulative recovery number versus time is given in figure S.9. The recovery data were submitted to the data selection program outlined in Chapter 3 of this report, and the data selected are shown in figure S.10. These data, minus the recoveries north of the chart extreme, were submitted to the inference program using standard parameter settings (Chapter 5).

The inference procedure had pretty well converged by iteration sixteen, although certain aspects of the inferred trajectories are rather confused. Selected iterations are shown in figure S.11, and the resulting field is laid out in figure S.12. The major feature of this field is its counterclockwise sense, in contrast to the clockwise nature of the previous summer. Also of note is the southward coastal countercurrent along the eastern shore. Beyond that, the other local anomalies are probably more artifactitious than real, considering the complicated nature of the returns (especially those from releases 90, 91, 104, and 107). Particularly suspicious is the small gyre near 42°N , 87°W . A

FLIGHT LOG for July 8, 1975

Drifter Releases over Lake Michigan

Prepared by Philip C. Pilgrim

0900 EST: Four cartons of prepackaged surface current drifter buoys were hauled to Ann Arbor Municipal Airport and loaded aboard a Piper Aztec twin-engine airplane for release over southern Lake Michigan. At 0925 the pilot (Mark Wagner) and scientific crew (Phil Pilgrim) took off, reaching the lakeshore south of Grand Haven at 1015. Setting course due west at 43°00'N and cruising at 120 knots and 700 feet above the lake, bundles of 25 drifters apiece were released from the air vent at the following times and release points:

Time	Release	Time	Release
1017	76	1037	80
1021	77	1042	81
1027	78	1047	82
1032	79	1048	83

Time	Release	Time	Release
1058	84	1120	88
1103	85	1125	89
1108	86	1130	90
1114	87		

At 1048, the western shore was reached south of the planned Milwaukee arrival point (due to a north wind) and the release point positions were altered linearly to compensate.

Heading back east at 42°45'N, seven more bundles were dropped as follows:

Land was reached north of the planned arrival point, and again the release points were altered to compensate. Setting course due west at 42°30'N, the following drops were made:

Time	Release	Time	Release
1147	91	1207	95
1151	92	1213	96
1157	93	1219	97
1202	94		

At 1225 the plane landed at Waukegan for lunch and refueling.

Taking off from Waukegan at 1310, three more transects of the lake were made with drops at the following times:

LATITUDE HEADING	42°15'N East		42°00'N West		41°45'N East	
	Time	Release	Time	Release	Time	Release
	1320	98	1404	105	1433	111
	1326	99	1406	106	1439	112
	1331	100	1412	107	1445	113
	1337	101	1418	108	1447	114
	1343	102	1423	109	1448	Arr. shore
	1349	103	1425	110		
	1351	104	1426	Arr. shore		
	1353	Arr. shore				

At 1540 the plane landed at Ann Arbor Municipal Airport.

Figure S.7: Flight log for 1975 drifter releases.

LAKE MICHIGAN




Figure S.8: Drifter release points for July 8, 1975.




Figure S.9: Number of recoveries versus time since release for July, 1975, drifter experiment.

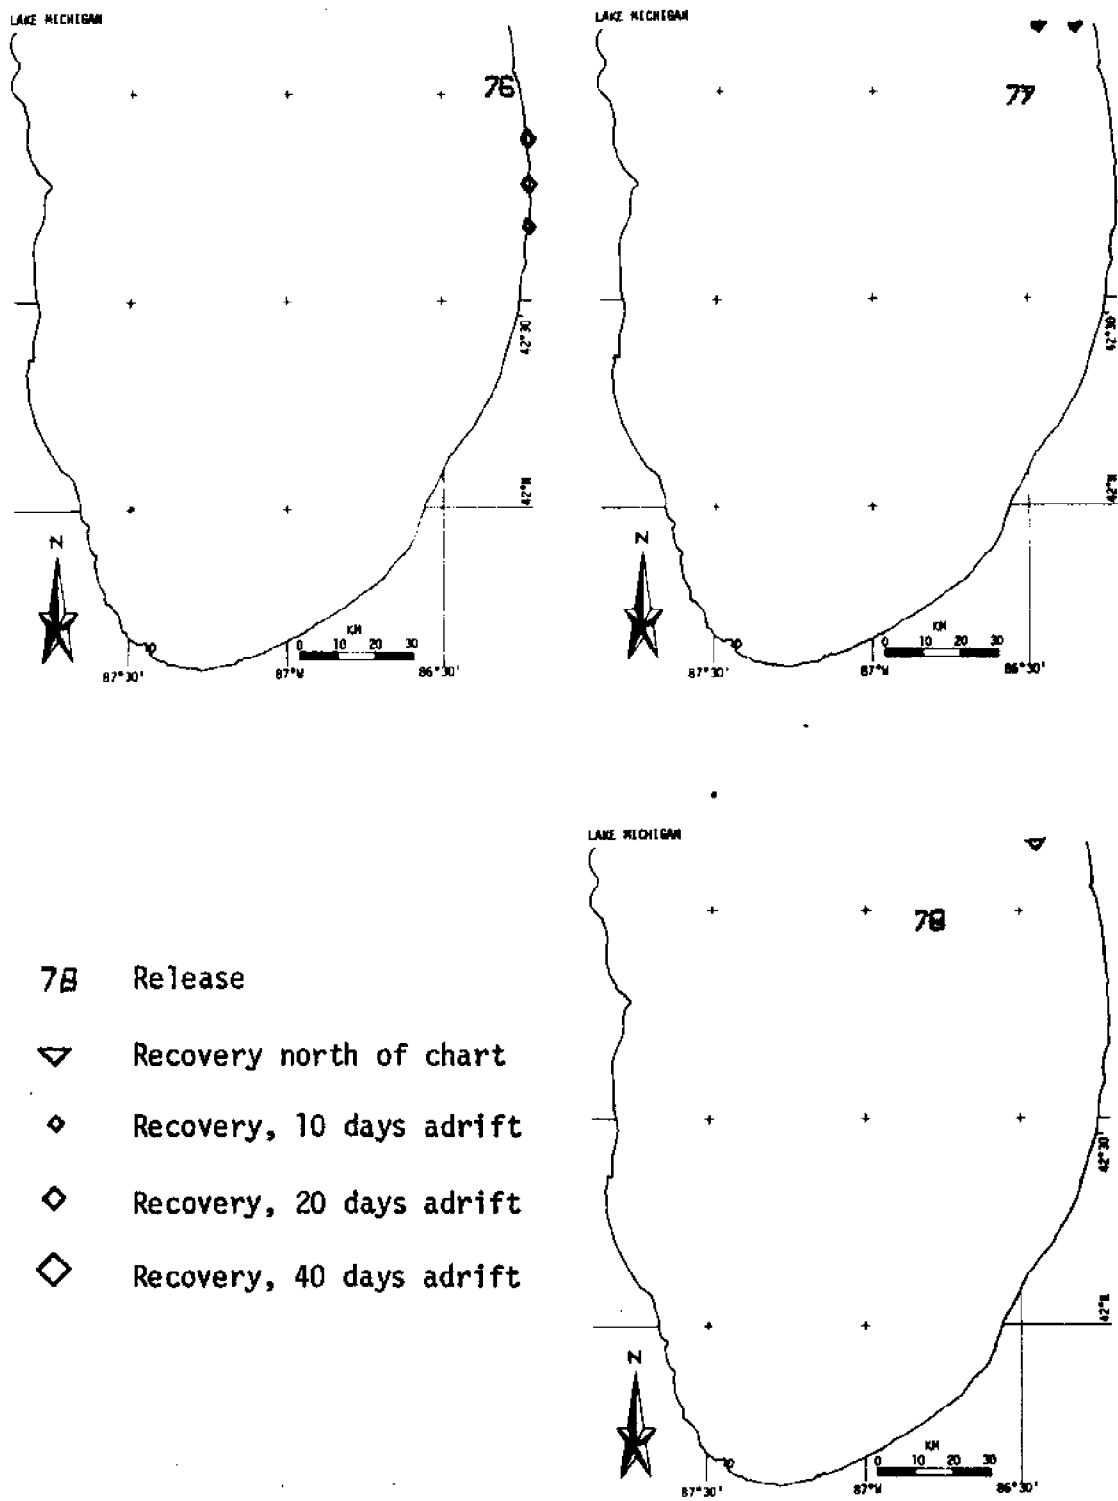










Figure S.10: Charts showing selected release/recovery data for July, 1975, drifter experiment. . . . /







/ . . . S.10, cont'd.

. . . /







/ . . . S.10, cont'd.

. . . /







/ . . . S.10, cont'd.

. . . /




/ . . . S.10, cont'd.


. . . /




/ . . . S.10, cont'd.




a) Iteration 1



b) Iteration 3




c) Iteration 5



d) Iteration 7


Figure S.11: Iterations of inference program on July, 1975, drifter data. . . . /




e) Iteration 9



f) Iteration 11



g) Iteration 13




h) Iteration 15

/ . . . S.11, cont'd.


. . . /

LAKE MICHIGAN



i) Iteration 16

/ . . . S.11, cont'd.




a) Inferred velocity field

Figure S.12: Inference results for July, 1975, drifter experiment.

. . . . /

LAKE MICHIGAN



b) Inferred velocity field (FILLED)

/ . . . S.12, cont'd.

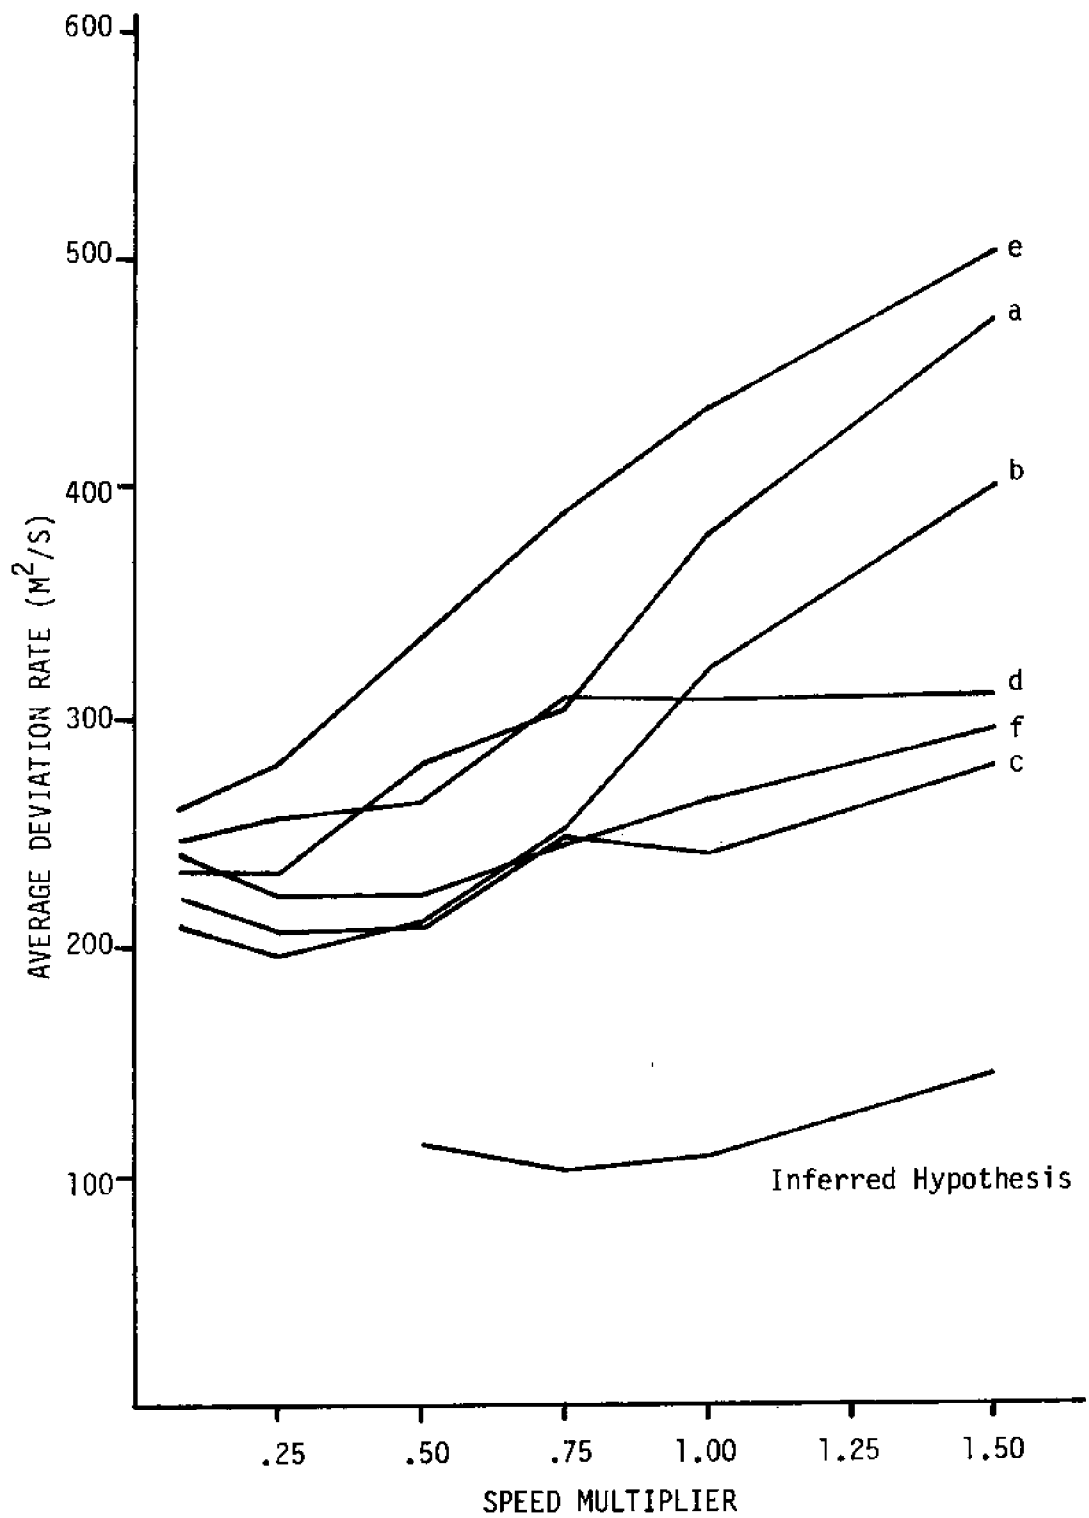


Figure S.13: Hypothesis testing results for July, 1975, drifter data.

Letters refer to figure 3.13.

possible explanation for such irregularities might be a slight violation of the stationarity assumption for the velocity field during the time of the experiment.

The field generated by the inference program as well as those hypothesized fields shown in figures 3.13a-f were submitted to the hypothesis testing scheme, using all the selected recovery data shown in figure S.10. Compatibility curves for the hypotheses are given in figure S.13. As can be seen, the generated hypothesis is the preferred one.

S.4 Supplementary Bibliography

Frantz, D.R., R.F. Brender, and J.L. Foy, Jr. (1968). "LOCOS: A Multiprogramming Monitor for the DEC PDP-7." CONCOMP, Technical Report 10, The University of Michigan (Ann Arbor).

International Business Machines (1970). *IBM 1800 Time-Sharing Executive System Operating Procedures*, Fifth Edition, no. 1800-36.

Pilgrim, Philip C. (1974). "Display Program System for 1800/PDP7," unpublished internal memo, Logic of Computers Group, Department of Computer and Communication Sciences, The University of Michigan (Ann Arbor).

