

Optimal Sites for Supplementary Weather Observations: Simulation with a Small Model

EDWARD N. LORENZ AND KERRY A. EMANUEL

Department of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts

(Manuscript received 12 December 1996, in final form 22 June 1997)

ABSTRACT

Anticipating the opportunity to make supplementary observations at locations that can depend upon the current weather situation, the question is posed as to what strategy should be adopted to select the locations, if the greatest improvement in analyses and forecasts is to be realized. To seek a preliminary answer, the authors introduce a model consisting of 40 ordinary differential equations, with the dependent variables representing values of some atmospheric quantity at 40 sites spaced equally about a latitude circle. The equations contain quadratic, linear, and constant terms representing advection, dissipation, and external forcing. Numerical integration indicates that small errors (differences between solutions) tend to double in about 2 days. Localized errors tend to spread eastward as they grow, encircling the globe after about 14 days.

In the experiments presented, 20 consecutive sites lie over the ocean and 20 over land. A particular solution is chosen as the true weather. Every 6 h observations are made, consisting of the true weather plus small random errors, at every land site, and at one ocean site to be selected by the strategy being considered. An analysis is then made, consisting of observations where observations are made and previously made 6-h forecasts elsewhere. Forecasts are made for each site at ranges from 6 h to 10 days. In all forecasts, a slightly weakened external forcing is used to simulate the model error. This process continues for 5 years, and mean-square forecast errors at each site at each range are accumulated.

Strategies that attempt to locate the site where the current analysis, as made without a supplementary observation, is most greatly in error are found to perform better than those that seek the oceanic site to which a chosen land site is most sensitive at a chosen range. Among the former are strategies based on the "breeding" method, a variant of singular vectors, and ensembles of "replicated" observations; the last of these outperforms the others. The authors speculate as to the applicability of these findings to models with more realistic dynamics or without extensive regions devoid of routine observations, and to the real world.

1. Introduction

Early in the present century synoptic weather analyses were based mainly on surface observations, taken at fixed stations over land and moving ships at sea. Before midcentury these had been augmented by upper-level reports, from balloons released at weather stations and aircraft flying diverse routes. More recently remote sensing became a reality, with measurements of certain quantities from satellites extending over much of the globe, and observations by such devices as radar covering more localized regions. Yet despite the present wealth of data—more, in fact, than we know how to use to full advantage—large gaps remain in our picture of the global weather pattern, particularly over the less frequently visited areas of the oceans.

Sporadic observations, whose locations have been dictated by current weather conditions, have over the

years been taken in time of need; for example, instrumented aircraft have flown through the eyewalls of tropical hurricanes, when better fixes on storm positions have been desired. Very recently, however, it has appeared that a limited number of special platforms, such as drone aircraft carrying dropsondes, may become available on a regular basis for supplementary weather observations, sometimes referred to as "adaptive" or "targeted" observations (e.g., Snyder 1996; Palmer et al. 1998), whose locations can depend upon the synoptic situation. The question then arises as to where these platforms should be deployed at any given time, if the data that they gather are to be most effective in improving the analyses and forecasts. It is this question to which the present study is addressed.

The proper answer is not immediately obvious. Presumably it will depend upon the nature of the data to be gathered—whether it is conventional rawinsonde data or something more exotic. It will depend upon how the data, once obtained, are to be assimilated into the analyses. Finally, it will depend upon what is considered to be maximum effectiveness—greatest improvement in forecasts at selected local sites or over extensive regions,

Corresponding author address: Dr. Edward N. Lorenz, Dept. of Earth, Atmospheric, and Planetary Sciences, Massachusetts Institute of Technology, Room 54-1622, Cambridge, MA 02139.

and at short or extended range. A search for an answer has constituted a part of the recent Fronts and Atlantic Storm-Tracks Experiment (FASTEX) (e.g., Snyder 1996; Joly et al. 1997).

Various strategies for locating the new observations suggest themselves. Some seek the regions where the analyses, as performed without the new observations, will be most greatly in error. Others try to target the locations where the present weather conditions will most strongly influence the subsequent weather. Each strategy possesses many possible variants. To test a large number of them adequately in the field within a reasonable period, once the platforms are available, seems to be out of the question. In these days when mathematical models of the weather are rife, it is virtually an axiom that one or more of them should be used for our first tests. This can be done even before any platforms are ready; nevertheless, we should anticipate that, whatever strategy we may decide upon, the need for modifications, at least of the details, will become apparent as soon as the new observations become a reality.

Because of the large number of specific procedures that might be tested, and the considerable number of simulated weather situations to which each must be applied before a definitive choice among them can be made, a full-scale experiment using a reasonably sophisticated model, such as the operational model of a major forecasting center, would be a vast undertaking, even if less formidable than a real-world test. We have therefore deemed it advisable, before embarking upon any such experiment, to perform the same sort of study with a very small model, where the time wasted by trying out untenable procedures or temporarily failing to discover program bugs will not constitute a serious loss. The model need not be meteorologically very realistic, but it should have enough in common with the atmosphere and its underlying ocean and land surfaces to allow some chance that the *relative* merits of various strategies will be properly indicated. In particular, it should behave chaotically; otherwise the forecasting problem will virtually disappear, and the experiments will have little meaning. We shall presently see that the use of simple models can at the very least cast doubt upon certain general approaches whose study with a large model would be rather time-consuming. At the same time, we cannot be certain of the extent to which our findings will hold up in the real world.

2. The model

It will be convenient to let the separate variables of our model represent conditions at separate geographical sites, and to let some of the variables be "observed" regularly and fairly accurately, while others are observed poorly or not at all. If the model is to serve its purpose, it must possess enough of the former to allow us to make "analyses" and enough of the latter to allow a targeting strategy to select sites in a rational manner.

We have chosen a model with J variables, denoted by X_1, \dots, X_J ; in most of our experiments we have let $J = 40$. The governing equations are

$$dX_j/dt = (X_{j+1} - X_{j-2})X_{j-1} - X_j + F, \quad (1)$$

for $j = 1, \dots, J$. To make Eq. (1) meaningful for all values of j we define $X_{-1} = X_{J-1}$, $X_0 = X_J$, and $X_{J+1} = X_1$, so that the variables form a cyclic chain, and may be looked at as values of some unspecified scalar meteorological quantity, perhaps vorticity or temperature, at J equally spaced sites extending around a latitude circle. Nothing will simulate the atmosphere's latitudinal or vertical extent.

We know of no way that the model can be produced by truncating a more comprehensive set of meteorological equations. We have merely formulated it as one of the simplest possible systems that treats all variables alike and shares certain properties with many atmospheric models, namely,

- 1) the nonlinear terms, intended to simulate advection, are quadratic and together conserve the total energy, defined as $(X_1^2 + \dots + X_J^2)/2$;
- 2) the linear terms, representing mechanical or thermal dissipation, decrease the total energy;
- 3) the constant terms, representing external forcing, prevent the total energy from decaying to zero.

One of us has, in fact, used the model previously in another context (Lorenz 1996). Note that the variables are scaled so that the coefficients of the quadratic and linear terms are unity; the time unit is thus the dissipative decay time, which we assume to equal 5 days.

As with typical nonlinear systems, any solutions found analytically are likely to be rather specialized, but certain properties may be deduced without solving the equations at all. First, if a bar ($\bar{\quad}$) over a quantity denotes an average over all values of j and over a long enough time to make average time derivatives negligibly small, it follows from multiplying (1) by X_j and averaging that

$$\overline{X^2} = F\bar{X}, \quad (2)$$

whence

$$\overline{X^2} - \bar{X}^2 = \bar{X}(F - \bar{X}). \quad (3)$$

Since the variance $\sigma^2 = \overline{X^2} - \bar{X}^2$ of X is nonnegative, it follows from (3) that the mean \bar{X} of X lies in the interval $[0, F]$, whence the standard deviation σ lies in the interval $[0, F/2]$. In the obvious steady solution where $X_j = F$ for each j , $\bar{X} = F$ and $\sigma = 0$.

Small perturbations x_j about the steady solution obey the equation

$$dx_j/dt = F(x_{j+1} - x_{j-2}) - x_j. \quad (4)$$

If we let

$$x_j = \sum_k p_k \exp(ikj), \quad (5)$$

we find that




FIG. 1. Longitudinal profiles of X_j at 6-h intervals, as determined by Eq. (1) with $N = 40$ and $F = 8.0$, when initially $X_{20} = F + 0.008$ and $X_j = F$ when $j \neq 20$. On horizontal portion of each curve, $X_j = F$. Interval between successive short marks at left and right is 0.01 units.

$$dp_k/dt = [(e^{ik} - e^{-2ik})F - 1]p_k. \quad (6)$$

The steady state is therefore unstable if $(\cos k - \cos 2k)F > 1$ for some k . The factor $\cos k - \cos 2k$ assumes its maximum positive value $9/8$ when $\cos k = 1/4$; thus the steady solution becomes unstable with respect to waves of length $2\pi/\cos^{-1}(1/4) = 4.77$ zones when F exceeds $8/9$. Since the actual number of zones in a wave must be a divisor of J , we find, for $J = 40$, that waves of five zone lengths, or wavenumber 8, will begin to grow when F exceeds $(\cos 2\pi/5 - \cos 4\pi/5)^{-1} = (4/5)^{1/2} = 0.894$, a value only slightly exceeding $8/9$.

The incipient waves will move with velocity $c = -(\sin k + \sin 2k)(F/k)$, which, when $k = 2\pi/5$ and $F = 0.894$, equals -1.09 ; thus they will drift westward. The group velocity $c_g = -(\cos k + 2 \cos 2k)F$, on the other hand, equals $+1.17$, implying eastward propagation of regions of enhanced activity; this feature has been an important consideration in our decision to adopt the model.

3. Numerical documentation of the model

Regardless of how well or how poorly the *equations* of the model resemble those of the atmosphere, it is essential to know, before proceeding with our experi-

ments, how closely the model resembles the atmosphere in its *behavior*. This we can discover by solving the equations numerically. We shall use a fourth-order Runge-Kutta scheme. By trial and error we have found that, with the numerical values that we shall be using, the system is computationally stable with a time step of 0.05 units, or 6 h—a convenient value, since in our subsequent experiments we shall assume that observations are taken at 6-h intervals.

We begin with small values of F and find, in almost perfect agreement with our theoretical deductions, that when $J = 40$ the steady solution becomes unstable when F exceeds 0.895. In Fig. 1, produced with the strongly supercritical value $F = 8.0$, each curve is a longitudinal “profile” of X , constructed by connecting simultaneous values of the separate variables with straight-line segments. In the leading curve, the steady state $X_j = F$ has been perturbed by multiplying the single variable X_{20} by 1.001. The remaining curves, obtained by numerical integration, follow their predecessors at 6-h intervals. After 2 days the detectable perturbations extend halfway around the latitude circle. The westward progression of the individual maxima and minima and the eastward progression of the center of activity are plainly revealed.

Similar initial behavior, with less rapid amplification




FIG. 2. Longitudinal profiles of X_j as in Fig. 1 but at 12-h intervals, and with the profile that would follow the initial profile in Fig. 1 by 3 years used as the new initial profile. Horizontal lines are zero lines. Interval between zero lines and short marks at left and right is 10.0 units.

and progression, is found when F is less strongly supercritical. When $F < 4.0$, the perturbations ultimately develop into a perfect wavenumber 8, progressing westward. This finding might not have been anticipated for the larger of these values of F since, even when F just exceeds 2.0, the steady solution is unstable with respect to wavenumbers 4 through 12 simultaneously. Evidently the growing waves gain their energy by extracting it from the energy source represented by the mean \bar{X} of X , and in so doing reduce \bar{X} to the point where the state is unstable only with respect to wavenumber 8.

When F exceeds 4.0, the waves no longer extract energy fast enough to offset the effect of the external forcing, and a spatially irregular pattern with chaotic time variations appears. In Fig. 2, the curves are constructed as in Fig. 1, again with $F = 8.0$ but at half-day intervals, while the leading profile is the one that would result from extending the integration of Fig. 1 for 3 years; presumably the transient effects have had ample time to die out. Wavenumber 8 appears prominently in the spectrum, but the individual “highs” and “lows” alter their shapes and intensities rather irregularly as they progress slowly and not invariably westward. Figure 3 shows the extension of Fig. 2 to 14 days, at 2-day intervals; a second solution, differing from the

first by $F/2$ at just one site in the leading profile, has been superposed. Areas where the new profile lies above the old one have been shaded; where it lies below, they are unshaded. Individual shaded areas seem to drift westward, but the entire shaded region spreads rapidly eastward, with very little accompanying westward penetration, until, by day 14, it extends almost completely around the latitude circle. The qualitative resemblance between the alternating positive and negative areas and the peaks and troughs in the incipient waves in Fig. 1 is apparent.

Figure 4 shows an 8-month time series of X_{10} , originating from the state that forms the leading profile in Fig. 3, displayed as four consecutive 60-day segments. Because of the symmetric form of Eq. (1), any other variable would vary in a qualitatively similar way. The lack of observable periodicity, despite some near repetitions, is typical of chaotically varying systems. To demonstrate the presence of chaos more conclusively, we have superposed on the segments extending from day 60 to 240 a second time series for X_{10} , produced by adding 0.0001 units to X_{10} at the start of day 60, while leaving the remaining variables unchanged. For about a month the difference between the solutions remains too small to be visually detectable, but then it




FIG. 3. Longitudinal profiles of X_j as in Fig. 2, but at 2-day intervals, with the initial profile of Fig. 2, and with a second set of profiles superposed. The superposed initial profile is formed by adding 4.0 units to X_{10} . Where the second profile lies above the original one, the area between the profiles is shaded.

quickly appears, and for the last four months each solution seems to go its own way.

Extended integrations indicate that $\bar{X} = 2.3$ and $\sigma = 3.6$, values well within the theoretical bounds. The leading Lyapunov exponent corresponds to a doubling time of 0.42 units, or 2.1 days, a value close to one that seems to prevail in some large atmospheric models (Lorenz 1982). Growth rates during limited intervals can be considerably larger or smaller. Systematic growth ceases, and “saturation” occurs, when the error in each variable reaches $\sigma\sqrt{2}$, or 5.1.

Altogether there are 13 positive Lyapunov exponents, while the fractional dimension of the attractor, as estimated from the formula of Kaplan and Yorke (1979), is about 27.1. Clearly the simultaneous values of all the variables are not implied by the instantaneous values of just a few.

Increasing F to 10.0 has little effect on the qualitative appearance of curves like those in Figs. 1–4. Quantitatively, the doubling time is decreased to 1.5 days, a value also compatible with recent large models (Simmons et al. 1995), while there are 14 positive exponents and a fractional dimension of about 29.4. Increasing J to 80 does not increase the spatial resolution of the

individual highs and lows; it simply doubles the number of them, leaving the dominant wavelength unchanged. Effectively it doubles the circumference of the earth.

4. The experimental setup

There are not only numerous targeting strategies that one might consider testing, each with numerous conceivable variants; there are also many ways in which one might design an experiment to test a given variant of a given strategy. To keep our total effort within bounds, we shall introduce a single standard format.

We begin by letting the consecutive sites numbered 1–20 lie over the ocean, while sites 21–40 lie over land. Whether a site is over ocean or land will not affect the governing dynamics; however, “observations” will be taken every 6 h at each land site, while there will be no routine observations at all over the ocean. A supplementary observation will be made at a single oceanic site every 6 h; our task is to evaluate the various possible strategies for choosing this site.

Real oceans are, of course, far from completely devoid of observations. In any event, the closest real-world analog to our ocean would be any extensive region,




FIG. 4. A time series of X_{10} , as determined by Eq. (1), originating from the initial state of Fig. 2 and extending for 240 days, displayed as four consecutive 60-day segments. A second time series, produced by adding 0.0001 units to the single variable X_{10} at the beginning of the second segment, is superposed on the last three segments. Horizontal lines are zero lines. Interval between short marks at left and right is 5.0 units.

oceanic or otherwise, where the observations are very sparse. Our choice of a “worst possible situation” may well affect our qualitative results.

The sequences of “meteorological” quantities that will appear in each experiment will extend over N 6-h time steps; we have worked with numerous values of N , but, in our “production runs,” $N = 7200$ (i.e., 5 years). The quantities consist of the true value X_{jn} of the variable X_j at the end of the n th time step, the observed value Y_{jn} of X_j at land and targeted ocean sites, the analyzed value Z_{jn} of X_j at all sites, and the value Z_{jnm} of X_j forecast from the analysis m time steps before the end of the n th step, also at all sites. The forecast range will extend up to M time steps; in production runs $M = 40$ (i.e., 10 days). Since an analysis is equivalent to a zero-day forecast, $Z_{jn0} = Z_{jn}$. We shall write the symbols with commas separating the subscripts when any subscript consists of more than a single letter or digit.

The “true” sequence of weather will be a particular time-dependent solution of Eq. (1). To obtain the true initial values X_{j0} , we first choose the J values $X_{j,-L}$ randomly from an approximately Gaussian distribution with mean $F/4$ and standard deviation $F/2$, and then

integrate forward for L steps to remove transient behavior; we have let $L = 360$ (i.e., 90 days). We shall denote these transient values by X_{jn} with negative values of n .

For the observations Y_{jn} we add an “observational error” a_{jn} , chosen randomly from an approximately Gaussian distribution with mean zero and standard deviation ϵ , to the true value X_{jn} ; in production runs ϵ equals $F/40$.

Prior to each analysis we make a “first guess”—the value Z_{jn1} forecast from the analysis 6 h earlier. For the analyzed value Z_{jn} , we use the first guess, sometimes called the “background” value, at sites where there are no observations. At sites with observations, including the targeted site, we simply replace the first guess by Y_{jn} . One might suppose that when an observation replaces a first guess at an isolated oceanic site, or a coastal land site, some adjustment ought to be made at neighboring sites to preserve the spatial continuity. Our rationale for not doing so is that solutions of Eq. (1) exhibit very little site-to-site continuity; observe Fig. 2. The time correlation between X_j and X_{j+1} —values at adjacent sites j and $j + 1$ —is close to zero. We have not investigated the correlation between errors in X_j and




FIG. 5. Five-year-average forecast errors at each site, at ranges up to 10 days, when there are no supplementary observations. The shaded and unshaded areas are produced by rounding each error down to a whole number and plotting the numbers, which range from 0 to 6, except that 0's and 3's are replaced by blank spaces. The site number increases toward the right, and the forecast range increases downward. A 0-day forecast is an analysis.

X_{j+1} , which in any event would depend upon our method of assimilating the observations.

Having completed an analysis, we determine the true state 6 hours later by integrating Eq. (1). At the same time, we make a 6-h forecast by applying Eq. (1) to the analysis, a 12-h forecast by applying the equation to a previously made 6-h forecast, and so on, to a range of $6M$ h. To simulate the “model error” in the forecasts, we replace F in Eq. (1) by the slightly lower value F' ; in our production runs $F' = 0.95F$. We retain F when applying Eq. (1) to the true weather.

Since a forecast depends on an analysis and an analysis depends on a previous forecast, we need some way to start things up. We do this by adding errors $a_{j,-L}$ to the true values $X_{j,-L}$ at all sites at the preinitial time step $-L$, to obtain the first guesses $Z_{j,-L,1}$. We carry out the observation–analysis–forecast routine throughout the transient period, when $n < 0$, but accumulate verification statistics only after n passes 0.

We shall represent the effectiveness of a particular strategy in terms of the root-mean-square difference D_{jm} between the analysis or forecast $Z_{j,m}$ and the true state $X_{j,m}$ for each site j and each forecast range m . Note that in practice a forecaster, not knowing the true state, would have to use some other measure. Our format for

presenting these $M(J + 1)$ values, 1640 of them in production runs, is modeled after a practice that was popular in the early days of numerical weather prediction, when computer graphics was in its infancy. We round each root-mean-square error D_{jm} down to a whole number, which will prove to range from 0 to 6, and then print the rounded numbers in a block, with j increasing horizontally and m increasing downward, except that we replace a number by a blank space if the number is 0 or 3.

Figure 5 shows such a block of values, determined with $F = 8.0$, $\epsilon = 0.2$, $F' = 7.6$, and $N = 7200$, for the special case when there are no supplementary observations at all. We see two shaded areas. One, composed of printed 1's and 2's, and representing fairly good forecasts with errors lying between 1.0 and 3.0 units, appears here as two areas, one on the extreme left and one on the right; it would appear as a single area had we chosen to construct the diagram with land on the left and ocean on the right. The other, composed of printed 4's and 5's and just a few 6's, represents poor or worthless forecasts with errors exceeding 4.0 and covers much of the central part of the diagram. Separating these shaded areas is an unshaded area representing mediocre forecasts, with errors between 3.0 and

4.0; it also appears here as two areas. Finally there is an unshaded area representing rather good forecasts, with errors below 1.0, confined here to shorter ranges at land sites.

We note that a “climatology” forecast, where every variable is predicted to equal \bar{X} , would have a root-mean-square error of σ , or 3.6 units, and would therefore fall in the “mediocre” unshaded area. We do not feel that this means that such a forecast is superior to those represented by 4’s, which, although poor, may still give some indication of the locations of highs and lows. Instead we feel that the better score for climatology represents an inadequacy of the root-mean-square as a measure of forecast skill. We do not expect this shortcoming to matter when we compare several schemes, each of which produces prognostic charts that “look like weather maps”—that possess realistic spatial variability—which climatology certainly does not. We consequently feel that the root-mean-square is acceptable in our experiments, since each strategy is expected to produce prognostic profiles that look somewhat like those in Fig. 2.

5. Strategies seeking maximum analysis errors

Let us examine Fig. 5 in greater detail. In the upper-right portion, the area of “good” forecasts, where $D_{jm} < 1.0$, extends at the eastern interior land sites to a 4-day range. Since the prescribed analysis error over land is 0.2 units, this range simply indicates an error growth with a doubling time slightly below 2 days. The larger forecast errors farther eastward and especially westward are due not to a faster growth in situ but instead to the propagation of already large errors from the oceanic sites.

Turning to the upper left, we encounter progressively larger analysis errors as we then move eastward across the ocean, until near the continent the analyses are no better than random guesses; indeed, since saturation occurs when $D_{jm} = 5.1$, the 6’s indicate that they are significantly worse. Obviously the analyses would be equal in quality to random guesses if there were no data assimilation at all, since they would effectively be infinite-range forecasts. With the assumption that assimilation of data over land ought to produce improved analyses everywhere, we can conclude that our assimilation procedure—simple substitution of observed values—is not optimal. More refined assimilation techniques, including the variational schemes known as 3D-VAR and 4D-VAR, have been developed for use with large operational models as well as smaller ones (e.g., Lorenc 1981; Thépaut and Courtier 1991; Parrish and Derber 1992). We shall nevertheless retain our simple scheme in comparing different targeting strategies, considering afterward how our results might be changed by refining the assimilation process.

Since the greatest errors at oceanic sites are just off the coast, it might seem that a good strategy would be to target site 19 or 20 at every observation time. A test

has shown, however, that the resulting overall improvement in analysis and prediction is disappointingly small, and this finding is not surprising when we note that by always targeting site 20 we effectively do no more than replace a model with 20 land and 20 ocean sites by one with 21 land and 19 ocean sites. Choosing some other fixed oceanic site, say site 10, as an “island,” that is, targeting it every time, does only slightly better.

When we target a particular oceanic site at a particular time, the analysis immediately becomes good at this site, and it usually remains fairly good for 6 or 12 h since neither the analyzed nor the true value changes too rapidly. If we immediately target the same site again, we shall be replacing a fairly good rather than a poor value by a good value and we shall not greatly reduce the total analysis error. In other words, the new observation will be largely redundant. It thus seems that some strategy that does not pick any one site or small subset of sites too frequently is needed. We might add that there is much redundancy among the *land* observations.

One possible strategy is simply to select the targeted ocean sites at random. When this is done, Fig. 5 is replaced by Fig. 6, which is constructed with the same format. We see that the areas representing any given degree of skill—good, fairly good, mediocre, or poor—are displaced downward, that is, toward longer forecast ranges. The analyses at all ocean sites are much better than guesswork. The greatest improvements over land are near the coasts, especially the west coast, but even in the interior the gain is detectable. Since the targeted sites are chosen without concern for the formulation or behavior of the model—without any “meteorological” knowledge, any strategy that purports to be skillful should outperform this one.

The random targeting strategy tends to avoid choosing a site where, simply because the same site has just been chosen, the analysis error is small, but it still may choose a site where the error simply happens to be small. A scheme that would avoid the latter possibility would be to choose at every time the site where the analyzed and true values differ most greatly. Figure 7 shows the results that would be obtained. The analysis errors at all ocean sites would be reduced essentially to those over land, while the forecast errors would not appreciably exceed those demanded by the inevitable two-day doubling. It is necessary to observe, then, that this strategy, while easy to apply when working with a simulation, cannot be applied under operational conditions where the forecaster, not knowing the true state, does not know where the analysis is most in error. At best it represents a goal that the forecaster can strive to approach.

It is hardly realistic to believe that one additional observation in the real atmosphere every 6 hours would cure the bulk of the forecaster’s troubles, and we might question the adequacy of the model on this account. However, in the model, one oceanic observation every 6 h amounts to 5% of the number made over land. It is




FIG. 6. As in Fig. 5 except that each set of routine synoptic observations is augmented by a single supplementary observation at a randomly selected oceanic site.

perhaps not so far fetched to believe that additional real atmospheric observations, totaling 5% of the number already assimilated into the analyses, and located strategically, might make spectacular improvements.

In any case, in the model the greatest achievable reduction of the analysis and forecast errors, with any specified assimilation scheme, is highly dependent upon the parameters of the model, including the number of land and ocean sites and the error-doubling time. With a larger value of F , and consequently more rapid error growth, the improvements produced by any targeting strategy would be less pronounced.

Of the attainable strategies that might be expected to be superior to random sampling, we first consider one patterned after the “breeding” procedure of Toth and Kalnay (1993), although not identical to it. The underlying reasoning is as follows.

Among the infinitesimal perturbations x_{jn} that may be superposed on X_{jn} , there are some special ones that possess their own characteristic long-term average rates of amplification or decay; these rates correspond to the Lyapunov exponents of the system. For convenience we shall refer to these perturbations as “modes,” even though it is sometimes objected that the term is inappropriate since the structure of such a mode, that is, its profile of relative values at the separate sites, is time variable, as is its instantaneous rate of amplification or decay. A randomly chosen infinitesimal perturbation in-

roduced at a specified time presumably consists of a linear combination of the modes and, after the growing modes have grown and the decaying modes have decayed for a while, it should lose its randomness and have approximately the structure of a combination of the more rapidly growing modes, rather than a combination of all modes.

To the extent that current forecast errors have resulted from the amplification of errors present at a somewhat earlier time, we may expect the forecast-error field to have one of these preferred structures. To the extent that current analysis errors have resulted from first guesses based on erroneous forecasts, we may expect the analysis-error field to possess a similar structure. The largest analysis error should then have a better-than-average chance of appearing at the site where some previously introduced perturbation now assumes its maximum absolute value.

Although this reasoning is patently quite speculative, we are still at liberty to see whether targeting such a site actually can lead to improved analyses and forecasts. Some modifications from the most direct procedure are needed. First, if we are to simulate a strategy that a forecaster can pursue in practice, we must, a number of hours or days before we are to choose our target, perturb the analysis Z_{jn} , rather than the true state X_{jn} , which the forecaster will not know. Next, if such a perturbation z_{jn} is not to become large instantaneously, we




FIG. 7. As in Fig. 6 except that the supplementary observation occurs at the oceanic site where the error in the first guess is greatest.

must, as long as we assimilate observations into the analysis Z_{jn} every 6 hours, assimilate these same or slightly modified observations, including the one oceanic observation, into the perturbed analysis $Z_{jn} + z_{jn}$. Effectively, instead of perturbing a solution of Eq. (1), we shall be perturbing a realization of the complete observation–analysis–forecasting routine. Any relevant Lyapunov exponents and their corresponding modes will then have to be exponents and modes of the routine—not of Eq. (1). Following the completion of each analysis these modes will acquire zero amplitude at all land sites.

Although we shall refer to the specific procedure that we have finally adopted as “breeding,” it departs somewhat from the spirit of the original breeding procedure. Instead of introducing perturbations at frequent intervals and allowing each one to mature before being put to use, we introduce a small but otherwise arbitrary perturbation z_{jn} at the single preinitial time $n = -L$ and then apply the observation–analysis–forecasting routine to both Z_{jn} and $Z_{jn} + z_{jn}$ throughout the 90-day preinitial period and the 5-yr test period. Also, following Toth and Kalnay (1993), to keep the perturbation from eventually becoming large or possibly from decaying to zero, we normalize it at each step; that is, we multiply the j values z_{jn} by a factor that will reduce, or possibly increase, the sum of their squares to its initial value. At

each 6-h step n we then target the site j where z_{jn} assumes its maximum absolute value.

In a variant of the procedure, which we may call multiple breeding, we introduce an ensemble of K perturbations z_{jnk} , with $k = 1, \dots, K$, at the preinitial time $-L$. We then apply the routine to the unperturbed analysis Z_{jn} and to each perturbed analysis $Z_{jn} + z_{jnk}$, and at each step n we normalize each perturbation and target the site j where $\sum z_{jnk}^2$ assumes its maximum value.

Simple breeding, to which multiple breeding reduces when $K = 1$, yields considerably better analyses and forecasts than random selection, but, especially when $K \geq 8$, multiple breeding is a marked improvement over simple breeding. Figure 8, produced with $K = 15$ and with the format of Figs. 5–7, summarizes the results. The improvement over Fig. 6 is quite evident. In the western continental and eastern oceanic regions, the areas of mediocre and poor forecasts have been displaced downward—toward longer range—by a full day or more from their positions in Fig. 6. We conclude that operationally attainable procedures that are superior to random selection do exist and that the multiple breeding procedure is one of them.

In a somewhat similar scheme, which we call “multiple replication” (or just replication if $K = 1$), we make explicit use of the fact that the observations are imperfect. Again we introduce K small perturbations at time




FIG. 8. As in Fig. 6 except that the site of the supplementary observation is determined by the multiple-breeding strategy, with $K = 15$.

$-L$ but, then whenever we assimilate an observation Y_{jn} into an unperturbed analysis Z_{jn} , we assimilate modified observations $Y_{jn} + y_{jnk}$ into modified analyses $Z_{jn} + z_{jnk}$ where the new “observational errors” y_{jnk} are chosen randomly from the same distribution as the original observational errors a_{jn} . We omit the normalization process. Effectively, each ensemble member is an attempt to replicate what the analysis might have been if the original observational errors had been different, and ideally we should have added the new errors to the true solution X_{jn} , but of course this is not known to the forecaster. Multiple replication is similar to the Observation System Simulation Experiment (OSSE) procedure in use at the Atmospheric Environment Service in Canada (Houtekamer and Derome 1995).

We find that multiple replication outperforms any other procedure that we have investigated. Figure 9, in the usual format, shows the results, again with $K = 15$. Additional downward displacement is evident.

Whereas the superiority of multiple replication over multiple breeding in this model is something verifiable by computation, a proposed explanation as to why multiple replication should be superior is bound to be rather subjective. Forgetting about such concepts as Lyapunov exponents and modes, we simply remark here that an analysis is an estimate of the truth, whereas the ensemble members in the multiple replication scheme are addi-

tional estimates of the truth, with about the same likelihood of being good ones. It seems reasonable that the errors in one estimate are most likely to be large at sites where this estimate differs most from the remaining ones. In the multiple-breeding procedure, since the perturbations are kept small, the separate ensemble members are effectively the *same* estimate of the truth, and our simple argument no longer holds.

6. Further experiments

Other procedures that we have sought to exploit are the use of a form of “singular vectors”—infinitesimal perturbations whose growth during specified limited periods, rather than infinite periods, is greatest (Farrell 1989; Mureau et al. 1993; Buizza and Palmer 1995; Palmer et al. 1998)—and single-site perturbations that most greatly influence specified land sites at specified future times, and are most economically found through an adjoint procedure (Cacuci 1981; Hall et al. 1982). In the former case the targeted site is the one where the perturbation in question assumes its greatest amplitude; in the latter it is simply the site of the perturbation.

Use of the perturbation where, relative to the amplitude itself, the instantaneous time derivative of the amplitude is largest did little better than random selection. Use of the perturbation that would amplify most greatly




FIG. 9. As in Fig. 6 except that the site of the supplementary observation is determined by the multiple-replication strategy, with $K = 15$.

in the few days following the present was also unproductive. When instead we used the less conventional procedure of determining the perturbation that, when introduced a few days *before* the present, would amplify the most during the period up to the present, the results were rather good. In determining the amplification of the perturbation z_{jn} during this period, we assimilated the intervening observations into the perturbed analysis $Z_{jn} + z_{jn}$ as well as into Z_{jn} —something that we could not do when determining amplification forward from the present time. When the perturbation was introduced 10 days before the present, a lead time that seemed to produce optimal results, the improvement was not as good as that produced by multiple replication, but superior to multiple breeding.

The results using the different procedures were in some instances so close that one might question whether 5 years is long enough to distinguish adequately between their merits. Accordingly, for random selection, multiple breeding with $K = 15$, the use of singular vectors introduced 10 days before the present, and multiple replication with $K = 15$, we have extended the experiments to 25 years. In Table 1 we present the root-mean-square forecast errors for the westernmost land site and for the average of all land sites, at ranges of 1, 3, 6, and 10 days, for five consecutive 5-yr periods. Thus,

for each strategy, the tabulated results fall into eight columns, each containing five numbers.

Although there is considerable spread among the numbers within some of the columns, the conclusions are fairly clear. There is no overlap between the numbers in any column for the obviously poorest strategy, random selection, and those in the corresponding column for any other strategy. Likewise, except for the 10-day range, there is no overlap between numbers for the best strategy, multiple replication, and those for any other strategy. The intermediate strategies are hardly distinguishable at the longer ranges, but at shorter ranges the singular-vector procedure has an edge over multiple breeding, despite some overlap in the numbers.

We might add that in a real-world test nobody would wait 25 years, or even 5 years, before settling upon a strategy. One would more likely adopt a strategy, at least temporarily, even before its apparent superiority could pass a statistical significance test.

When we wish to determine the sensitivity of *each* variable X_j in Eq. (1) at some future time to a small modification of a *single* variable X_i at the present time, we can integrate Eq. (1) twice, once without and once with, the value of X_i initially modified. This is, in fact, what we did in constructing Fig. 3, which shows the sensitivity of each variable, at various ranges, to an

TABLE 1. Rms errors for five consecutive 5-yr periods (numbered 1–5), in hundredths of a unit, at the westernmost land site (*W*) and averaged over all land sites (*A*), at ranges of 1, 3, 6, and 10 days, when the routine observations are supplemented by single-targeted observations, with the sites selected randomly (*R*) by multiple breeding (*MB*) with $K = 15$, by singular vectors (*SV*), or by multiple replication (*MR*) with $K = 15$. Note that an error that has reached saturation would appear as 510.

Strategy	Range (<i>W</i>)				Range (<i>A</i>)			
	1	3	6	10	1	3	6	10
R 1	340	421	467	488	103	213	343	453
2	302	410	458	470	98	206	342	447
3	296	360	424	465	93	190	329	440
4	291	378	434	487	92	199	334	447
5	324	413	446	477	99	207	342	453
MB 1	221	318	402	457	73	166	303	425
2	238	323	387	441	75	164	302	422
3	237	310	363	424	73	160	287	410
4	209	293	393	456	69	167	304	426
5	192	310	367	429	65	152	290	419
SV 1	167	281	368	452	59	147	287	417
2	211	309	395	453	71	164	301	425
3	164	271	362	421	57	144	284	407
4	213	298	384	441	69	162	296	419
5	182	289	370	427	65	150	284	412
MR 1	123	213	345	419	47	125	268	402
2	132	242	354	440	47	131	273	411
3	128	227	321	402	47	125	262	393
4	145	255	359	445	51	136	282	413
5	126	226	337	421	47	128	272	406

initial modification of X_{10} . Alternatively, we can derive from Eq. (1) the equation

$$dx_j/dt = X_{j-1}x_{j+1} - x_j + (X_{j+1} - X_{j-2})x_{j-1} - X_{j-1}x_{j-2} \quad (7)$$

governing small departures x_j from X_j , and integrate Eq. (7) once, with only the i th variable x_i differing from zero initially. Since the coefficients X_{j-1} , etc., in (7) are time dependent, we must also integrate Eq. (1) once to determine them.

If instead we are interested in the sensitivity of a *single* variable at some future time to separate modifications of *each* variable at the present, we can integrate Eq. (7) many times, obtaining much unneeded information in the process, but it is far more economical to integrate the adjoint to Eq. (7)—the equation in which, for each pair (i, j) , the value of $\partial(dx_j/dt)/\partial x_i$ equals the value that $\partial(dx_i/dt)/\partial x_j$ possesses in Eq. (7)—a single time. The adjoint is readily seen to be

$$dx_j/dt = X_{j-2}x_{j-1} - x_j + (X_{j+2} - X_{j-1})x_{j+1} - X_{j-1}x_{j+2}. \quad (8)$$

Again the coefficients are time dependent and must be found by integrating Eq. (1).

Particularly when we are interested in good forecasts at a specific land site at a specific range, it would appear reasonable to target the oceanic site to which the land site is most sensitive. We have done just this, using the adjoint procedure to find the sensitivity, 3 and 7 days in advance, at site 21, the westernmost land site, to each oceanic site. The results have been disappointing. With the 7-day sensitivity, the performance is virtually in-

distinguishable from that of random selection. With the 3-day sensitivity, there is great improvement in the 1-day forecast at site 21, offset, however, by poorer forecasts in the east, and there is no improvement at all in the 3-day forecast, where the greatest gain might have been anticipated. At longer ranges the forecasts are actually poorer than those produced by random targeting.

Two possible reasons for the failure of the procedure suggest themselves. First, the sensitivity to an infinitesimal modification, which is what the adjoint procedure determines, need not be the same as the sensitivity to the sometimes rather large change that corrects the analysis at the targeted site. Investigation shows that this is generally the case for sensitivities at ranges beyond 5 days but not at shorter ranges. Second, the chosen land site may tend to be most sensitive to the same oceanic site, or to a small group of sites, many times in a row, so that sites where the analysis has been more or less corrected continue to be targeted. This proves to be the case for sensitivities at the shorter ranges.

We have attempted to devise procedures that combine the adjoint and breeding procedures, but we have simply obtained results intermediate to those yielded by the two procedures separately. Noting that modifying an analysis through an adjoint procedure has produced marked improvements in short-range forecasts with large models (e.g., Rabier et al. 1996), we feel that there should be a way to put the adjoint procedure to good use with our model, but have yet to discover it. We are thus forced to face the possibility that what works best in our model may not do so well in the real world, and vice versa.

TABLE 2. Five-year rms errors, in hundredths of a unit, averaged over all sites, at ranges of 0, 1, 3, 6, and 10 days, when there are no routine observations and M targeted observations, with the sites selected randomly (R), by multiple breeding (MB) with $K = 15$ or by multiple replication (MR) with $K = 15$.

Strategy	M	Range				
		0	1	3	6	10
R	1	425	447	472	486	500
	2	347	385	436	471	493
	3	264	308	382	439	481
	4	160	202	290	383	455
	5	102	136	218	329	429
	6	64	91	165	286	404
	7	44	64	135	259	385
	8	33	50	112	237	376
	9	29	42	99	224	365
MB	1	449	470	485	498	503
	2	359	404	445	476	497
	3	117	151	217	317	418
	4	28	44	105	232	370
	5	26	43	110	241	385
MR	1	371	418	461	479	495
	2	259	327	399	449	453
	3	42	67	145	273	396
	4	27	42	101	230	373
	5	23	36	88	211	357

7. Further comments and conclusions

There are many more things that one might do with the present model. One of the simplest refinements would be to let the rates of dissipation and external forcing over the ocean differ from those over land. Possibly there would be some unexpected changes in predictability at coastal land sites.

If we are inclined to be pessimistic, we might ask what should be done if our observation system, far from being enhanced by a few observations of opportunity, were to be subjected to a general downsizing. Perhaps we could perform a study somewhat like the present one, possibly with just as many ramifications, to seek to determine where or when we should sacrifice some observations, if our analyses and forecasts are to suffer the least.

Proceeding to another scenario, we might ask what would happen if all of our observations were of the adaptive or targeted type. How many of these would be needed, for example, to compensate for a complete absence of a regular observing network? This question is readily answered, given the targeting strategy to be used, in the context of the present model.

Ocean and land sites have now lost their distinction, and at any range, with the strategies that we have pursued, the long-term-average forecast errors should be the same at all sites. In Table 2 we show these errors at ranges of 0 (the analysis), 1, 3, 6, and 10 days, for the indicated numbers of targeted observations per 6 h, when our strategies are random selection, multiple breeding with $K = 15$, and multiple replication with $K = 15$.

We see that with the latter strategies, four targeted

observations every 6 hours lead to analysis errors only slightly exceeding the prescribed observational error, 0.2 units, whereas the subsequent error growth as the forecast range increases is not too much greater than that demanded by the two-day doubling. Perhaps somewhat surprisingly, a fifth supplementary observation seems to be largely redundant. Random targeting can similarly reduce the errors, but only if 9 or 10 sites are targeted. In attempting to extend these results to the real world, we should note that four targeted sites in the model are equivalent to perhaps 20% of the observations that we now assimilate. Also, had we used a larger value of F , four sites would not have sufficed.

More generally, the model ought to prove ideal for numerous other purposes, not necessarily related to site selection. These can range from pilot studies for specific meteorological undertakings to investigations of chaotic dynamical systems known to have moderately high fractional dimensionality.

Returning to the present study, we first recall that our assimilation scheme, where we merely introduce observed values into the analysis, is far from optimal. Improvements are not hard to come by.

At the eastern land sites, the errors in the first guess result mainly from the amplification of errors in the observations 6 h earlier, although the model error is partly responsible. If, as assumed, the errors in successive observations are independent, some linear combination of a first guess and an observation should on the average provide a better analysis than the first guess or the observation alone.

Let the mean-square observation and analysis errors be e^2 and f^2 , respectively; e but not f is unspecified.

The mean-square first-guess error should then be $\alpha^2 f^2$, where, if the error-doubling time is about 2 days, or eight times steps, α is about 1.09. If the observations and first guesses are given weightings $1 - k$ and k in the linear combination, the new mean-square analysis error, which should equal the old one, f^2 , should be $(1 - k)^2 e^2 + k^2 \alpha^2 f^2$. Equating these values, solving for f/e , and choosing k to minimize f/e , we find that $k = \alpha^{-2}$, or about 0.84, while $f/e = (1 - k)^{1/2}$, or about 0.40. At the western land stations, where the first-guess error is due largely to propagation of larger errors from oceanic sites, α should be larger, whence k should be smaller.

We have applied the new assimilation procedure to the random-selection, multiple-breeding, and multiple-replication strategies, letting $k = 0$ at sites 21–23 and $k = 0.84$ elsewhere. We find in each case a marked general improvement in the forecasts. Typically it takes a day or so longer for an error to reach a given magnitude. However, there is little change in the relative merits of the various schemes.

Independently of this empirically based assimilation procedure, which may be considered a rudimentary form of 3D-VAR (reducing to 1D for this model), we can introduce a dynamically based procedure, in the spirit of 4D-VAR, that makes explicit use of Eq. (1). Specifically, from the way that the “weather” has been varying at coastal land sites, we can estimate what the weather must have been just offshore to produce these variations.

If $Z_{1,n,1}$, the first guess at site 1, is to be replaced as an analysis by a new estimate $Z_{1,n}$, the first guess $Z_{40,n,1}$, at site 40 should be replaced by a new guess $Z_{40,n}$, where, to a fair approximation, since $\partial(dX_{40}/dt)/\partial X_1 = X_{39}$ according to Eq. (1),

$$Z_{40,n} = Z_{40,n,1} + Z_{39,n,1}(Z_{1,n} - Z_{1,n,1})\Delta t. \quad (9)$$

Ideally $Z_{40,n}$ should equal the observation $Y_{40,n}$. Equating these quantities, we can solve for $Z_{1,n}$. The value will not be exact because of the observation and model errors and because an instantaneous time derivative does not exactly determine a previous 6-h change.

The trouble with this procedure becomes evident when we note that we must divide by $Z_{39,n,1}$, and we might find ourselves trying to divide by zero. While this is extremely unlikely, it is not at all unlikely that we would be dividing by something very small, thus obtaining far too large a value for $Z_{1,n}$. We can circumvent this difficulty and still improve our analyses significantly by requiring that $Z_{1,n}$ should not differ too greatly from $Z_{1,n,1}$. This we do by choosing $Z_{1,n}$ to minimize the quantity $(Z_{40,n} - Y_{40,n})^2 + k^2(Z_{1,n} - Z_{1,n,1})^2$, where k is a weighting factor, somewhat arbitrarily chosen to equal 0.5. In a like manner, we can reestimate $Z_{20,n}$ so that the new guess $Z_{22,n}$ is close to $Y_{22,n}$, after which we can reestimate $Z_{19,n}$ to make $Z_{21,n}$ close to $Y_{21,n}$.

With the random-selection and multiple-breeding strategies we find marked improvements in the shorter-

range forecasts at the farther-west land stations; these progress eastward as the forecast range increases. With the multiple-replication strategy the improvements are less pronounced, but this strategy remains superior to multiple breeding.

Combining the two new assimilation schemes retains the separate advantages of each. More refined schemes ought to improve the forecasts still more, but, in any event, the purpose of the model experiments is to evaluate the *relative* merits of the procedures rather than to produce the best simulated forecasts that the model will allow.

Next, an assumption of a complete absence of observations over the ocean, or over any other sufficiently extensive area, is not realistic. Actual oceanic areas are deficient mainly in observations of those quantities that are not too well estimated by satellite. It would be possible, even with a small model, to introduce other assumptions as to where observations are and are not regularly or perhaps sporadically made. Although we have not attempted to do this, we suspect that again the scoring positions of the strategies—multiple replication first, singular vectors next, etc.—would not change, but we must again face the possibility that some adjoint procedure would become a contender.

In any case, there is a real question as to whether the scoring positions would be the same if the strategies were applied to the real atmosphere, or to a more comprehensive model. Our model lacks certain important atmospheric properties, including room for smaller-scale structure, and a vertical and second horizontal dimension. It would certainly be possible to enlarge the model by adding these dimensions, retaining the same sort of linear and nonlinear terms, although this would greatly reduce its computational advantage over more conventional models. However, the *behavior* of our model also lacks some perhaps crucial features. We have produced nothing analogous to rapidly or explosively deepening cyclones, which we might have expected to encounter and which suggest themselves as likely sites for supplementary observations. Indeed, if the *errors* in a cyclone or some other restricted region are growing explosively, a succession of adaptive observations, even though confined to this region, might still not be redundant. We have not created strong horizontal gradients and, in fact, have not even attained the spatial continuity through which an observation at one point might tell us a good deal about conditions at neighboring points. These objections can presumably be overcome by going to a more *realistic* model—preferably something at least as sophisticated as a multilayer quasi-geostrophic model.

It is also quite unrealistic, incidentally, to assume that an airplane—a likely observing platform—will simply fly to a distant target and then release a dropsonde. Presumably, for relatively little additional expense, it can gather data of some sort along its entire route. With a three-dimensional model, we can at least consider the

problem of selecting optimum routes rather than optimum single targets—a far more daunting task than we have faced so far.

We cannot say with any assurance that, with a more comprehensive model, schemes designed to resemble the ones that we have been studying will retain their relative merits, nor whether some scheme that cannot even be formulated in the context of our model might outperform them. We feel strongly, however, that any investigators undertaking studies with newer models should, no matter what strategies they may be testing, bear in mind the principal finding of this study and at least compare their strategies with one that attempts to target the locations where the first-guess errors are largest.

Acknowledgments. We have received some very helpful comments from several reviewers. Our work has been supported by the Atmospheric Sciences Section of the National Science Foundation, under Grant 9318422-ATM of the Climate Dynamics Program and Grant 9634239-ATM of the Large-scale Dynamics Program.

REFERENCES

- Buizza, R., and T. N. Palmer, 1995: The singular-vector structure of the atmospheric general circulation. *J. Atmos. Sci.*, **52**, 1434–1456.
- Cacuci, D. G., 1981: Sensitivity theory for nonlinear systems. I: Nonlinear functional analysis approach. *J. Math. Phys.*, **22**, 2794–2808.
- Farrell, B. F., 1989: Optimal excitation of baroclinic waves. *J. Atmos. Sci.*, **46**, 1193–1206.
- Hall, M. C. G., D. G. Cacuci, and M. E. Schlesinger, 1982: Sensitivity analysis of a radiative-convective model by the adjoint method. *J. Atmos. Sci.*, **39**, 2038–2050.
- Houtekamer, P. L., and J. Derome, 1995: Methods for ensemble prediction. *Mon. Wea. Rev.*, **123**, 2181–2196.
- Joly, A., D. Jorgensen, M. A. Shapiro, A. Thorpe, P. Bessemoulin, K. A. Browning, J.-P. Cammas, J.-P. Chalon, S. A. Clough, K. A. Emanuel, L. Eymard, R. Gall, P. H. Hildebrand, R. H. Langland, Y. Lemaître, P. Lynch, J. A. Moore, P. O. G. Persson, C. Snyder, and R. M. Wakimoto, 1997: Definition of the Fronts and Atlantic Storm-Tracks Experiment (FASTEX): Scientific objectives and experimental design. *Bull. Amer. Meteor. Soc.*, **78**, 1917–1940.
- Kaplan, J. L., and J. A. Yorke, 1979: Chaotic behavior of multidimensional difference equations. *Lecture Notes in Mathematics*, H.-O. Peitgen and H.-O. Waters, Eds., Springer-Verlag, 204–227.
- Lorenç, A. C., 1981: A global three-dimensional multivariate statistical interpolation scheme. *Mon. Wea. Rev.*, **109**, 701–721.
- Lorenz, E. N., 1982: Atmospheric predictability experiments with a large numerical model. *Tellus*, **34**, 505–513.
- , 1996: Predictability: A problem partly solved. *Proc. Seminar on Predictability*, Vol. 1, ECMWF, Reading, Berkshire, UK, 1–18.
- Mureau, R., F. Molteni, and T. N. Palmer, 1993: Ensemble prediction using dynamically conditioned perturbations. *Quart. J. Roy. Meteor. Soc.*, **119**, 299–323.
- Palmer, T. N., R. Gelaro, J. Barkmeijer, and R. Buizza, 1998: Singular vectors, metrics, and adaptive observations. *J. Atmos. Sci.*, **55**, in press.
- Parrish, D. F., and J. C. Derber, 1992: The National Meteorological Center's spectral statistical-interpolation analysis system. *Mon. Wea. Rev.*, **120**, 1747–1763.
- Rabier, F., E. Klinker, P. Courtier, and A. Hollingsworth, 1996: Sensitivity of forecast errors to initial conditions. *Quart. J. Roy. Meteor. Soc.*, **122**, 121–150.
- Simmons, A. J., R. Mureau, and T. Petrolagiadis, 1995: Error growth estimates of predictability from the ECMWF forecasting system. *Quart. J. Roy. Meteor. Soc.*, **121**, 1739–1771.
- Snyder, C., 1996: Summary of an informal workshop on adaptive observations and FASTEX. *Bull. Amer. Meteor. Soc.*, **77**, 953–961.
- Thépaut, J.-N., and P. Courtier, 1991: Four-dimensional variational data assimilation using the adjoint of a multilevel primitive-equation model. *Quart. J. Roy. Meteor. Soc.*, **117**, 1225–1254.
- Toth, Z., and E. Kalnay, 1993: Ensemble forecasting at NMC: The generation of perturbations. *Bull. Amer. Meteor. Soc.*, **74**, 2317–2330.