Single-Column Data Assimilation for the Atmospheric Radiation Measurement (ARM) Program

J.-F. Louis
Atmospheric and Environmental Research, Inc.
Cambridge, MA 02139

Why assimilate ARM data?

The main purpose of the ARM program is to provide the necessary data to develop, test and validate the parameterization of clouds and of their interactions with the radiation field, and the computation of radiative transfer in climate models. For the most part, however, the ARM observations will be imperfect, incomplete, redundant, indirect, and unrepresentative. This is unavoidable, despite the best efforts at equipping the Cloud and Radiation Testbed (CART) site with the best instruments. To understand these limitations, we must consider the structure of a climate model and the observation constraints.

The basic prognostic variables of any climate model are atmospheric temperature, horizontal wind components and humidity, and some surface variables: surface pressure, snow amount, soil temperature, etc. These variables are defined at a set of grid points (or, equivalently, as a set of spectral coefficients) at a small number of vertical levels. The grid boxes generally are a few hundred kilometers on the side, and hundreds to thousands of meters thick. The radiation scheme and cloud parameterization therefore use as input quantities averaged over the model grid box, and produce average fluxes.

By contrast, most of the ARM observations at individual sites will be obtained in a relatively small area compared to the scale of one grid box of a climate model. This is certainly true of radiosonde data, which sample extremely small volumes, both in time and space, but it also applies to most other observation techniques. ARM observations will thus be unrepresentative, meaning that they will be affected by scales of motion that do not exist in climate models. This question of scales and representativity of the data may be particularly serious for the radiative processes, which are highly nonlinear.

Besides fairly infrequent radiosondes, in situ observations of the atmosphere over the ARM sites will not normally be available because of cost considerations. Instead, the ARM plan calls mainly for remote sensing instrumentation. This means that we are not measuring temperature or moisture directly, but indirectly through their effect on electromagnetic radiation.

Although different instruments are available and will be used to measure the basic variables of climate models—temperature, humidity and winds—these measurements are not sufficient to completely validate cloud parameterizations or radiation schemes. Many quantities that would be needed for this purpose will be unavailable, except possibly during campaigns or special observing periods. Vertical profiles of radiative fluxes, cloud droplet distribution, aerosol distribution, turbulent fluxes within clouds, or optical properties of ice crystals are but a few examples. In many respects, the data will thus be incomplete, but, on the other hand, some quantities will be measured or inferred by several different instruments or methods, with different error characteristics.

We need a way to reconcile conflicting observations. This problem can be illustrated by comparing temperature observations made with two different radiosondes (Figure 1). One is the ARM data on 29 October and the other one the Oklahoma City (OKC) sounding of the National Weather Service (NWS) 4 hours later. One obvious difference is the resolution. In the NWS sounding we only have the data at the mandatory levels. The questions are "Which profile is more appropriate for a climate model?" and "Can the two sets of data be reconciled?"

Finally, measurements are obviously not perfect. Instrument errors can be minimized but not entirely avoided. In addition, the effects (on the measurements) of the scales that are

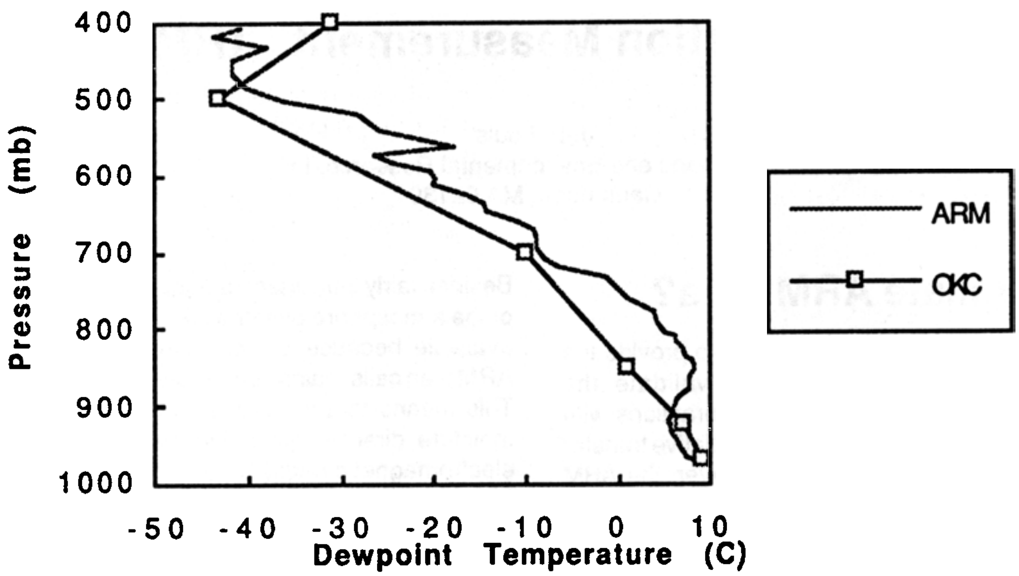

Figure 1. Dewpoint temperature profile at Oklahoma City (OKC) and at the ARM site, 4 hours apart.

not represented by the climate models cannot be distinguished easily from actual observation errors, except that they may be spatially correlated. Furthermore, when remote sensing is used, the retrieval methods introduce errors which cannot be avoided. In addition, remote sensing measurements may have exotic error characteristics.

Inaccuracy, incompleteness, redundancy, unrepresentativeness, and indirect measurements are all problems with which the weather forecasting community has been struggling for years. Various techniques of data assimilation have been developed to deal with them. Common to all the modern techniques of data assimilation is the idea that all the information available about the atmosphere should be used, and that includes not only the observations, but also our knowledge of the physical processes, which can be expressed in a model. The model is used to ensure time continuity and to constrain the analysis to be consistent with the physics of the atmosphere and representative of the desired scales. Estimates of the expected errors of the various data sources can also be used to combine the observations and the model estimate in an optimal way. Finally, the model can provide estimates of quantities that are not observed. While these estimates are not necessarily "the truth," they are compatible both with the other observations and with the model constraints.

The most advanced methods of data assimilation make use of the variational principle, which consists of adjusting some parameters of the model to minimize the difference between the observations and the model simulation. One of the major advantages of this method is that it is quite easy to include indirect measurements, as long as the quantity measured can be simulated by the model. For example, it is not necessary to use a retrieval method to turn the radiances of remotely sensed observations into temperature profiles. The radiances can be used directly because the model can estimate the radiances corresponding to its own temperature profile. Surface fluxes, precipitation, or cloudiness observations can also be used.

The variational method can be expensive because it requires an iterative minimum search, which involves running the model many times. In the ARM context, though, we are interested in data assimilation at a single site. Therefore, we can use a single-column model, which makes the variational method quite feasible, even with modest computer resources.

Our purpose is twofold: to explore the variational data assimilation technique, which has not yet been used in an operational context, and to provide to the ARM community a tool to turn the observations into the measurements that are needed to develop and validate climate model components.

Variational Data Assimilation

I introduced the concept of variational data assimilation and adjoint method at the 1992 Science Team meeting. I summarize it here briefly.

A model is used to simulate the evolution of the atmosphere during an assimilation period. The assimilation period should be long enough to include enough observations to constrain the model, but short enough that the evolution of a small perturbation can be described by the linear tangent model (LTM)[a]. A period of 12 to 24 hours seems to be reasonable.

As the model is run, we compute a cost function, which is essentially a measure of simulation errors. Typically, it is the weighted sum of the squared differences between observations and model output, the weights being the inverse of the expected errors. Other constraints can be introduced in the cost function as will be seen later.

The method then consists in adjusting some model parameters (called control variables) to minimize the cost function. At the minimum, the model simulation becomes the "analysis." We have different choices of control variables. They could be model physical parameters that are not well known, the initial state of the simulation, some nudging terms in the model equations that account for the deficiencies of the model, or a combination of all of these. We have been experimenting with the different possibilities.

Initially we will concentrate on the model's physical parameters in order to find the optimum set of parameters for the CART site. Once these parameters are set, we tend to prefer using nudging terms as control variables. Derber (1989) found that this resulted in a better analysis; it also has the advantage of resulting in a continuous description of the atmospheric state, without jumps at the beginning of each assimilation period.

In the minimization process, we need to compute the gradient of the cost function with respect to the control variables. This is used to determine the direction and size of the step in the minimum search. When the model is nonlinear, the gradient cannot be written analytically, but it can be computed by integrating the adjoint model, which, in a discrete model, is the transpose of the linear tangent model (Hoffman et al. 1992).

The ALFA model

The model used in this work is an extension of the AER Local Forecast and Assimilation (ALFA) model, which we started developing at Atmospheric and Environmental Research (AER) with the goal of doing local forecasting (Louis et al. 1990). It is a single column model that computes all the physics in the ground and atmosphere and takes the horizontal derivatives needed for the advection terms from a large scale forecast (or analysis). In forecasting mode, the model would be used to compute how the local conditions modulate the large scale flow predicted by a global or regional model. The adjoint technique is used to optimize the model parameters.

Much of the work in the first 2 years of our ARM contract has been to incorporate a sophisticated radiation scheme into the ALFA model. The scheme chosen is the one developed by Toon et al. (1989). It makes use of a generalized two-stream approximation and is designed for vertically inhomogeneous, multiple-scattering atmospheres. We have modified it to allow for fractional cloudiness and have written its adjoint. We have also written the adjoint of the convection scheme (Anthes et al. 1982) and the stratiform precipitation scheme.

An example of the results obtained with the Toon scheme is given in Figures 2 and 3, which show the computed solar and downward IR fluxes for 8-9 July 1992 at Oklahoma City. During the day, the model computes about 50% cloudiness; at the end of the second night, it computes 100% of low clouds.

(a) The linear tangent model is the forecast model linearized around the actual trajectory followed by the forecast model in the phase space.

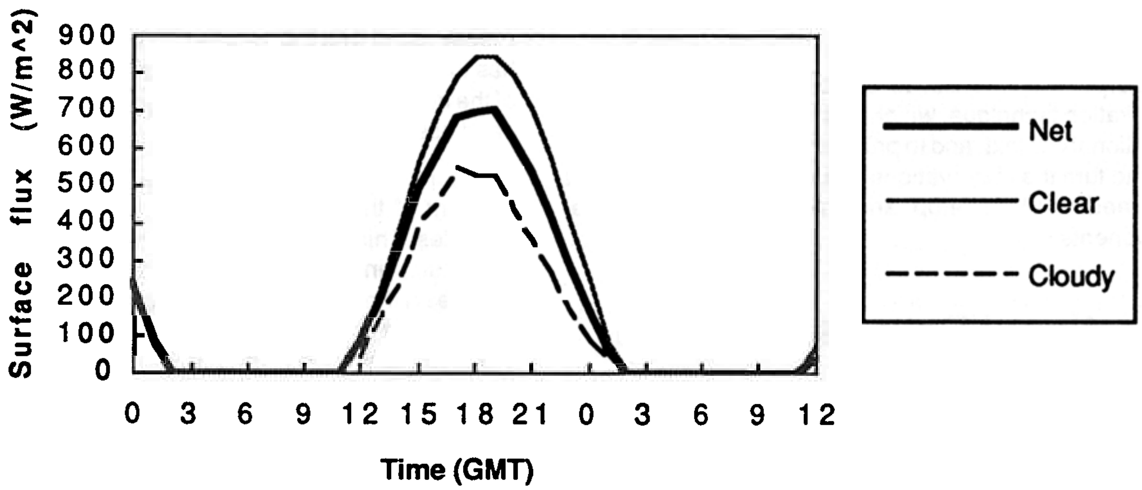

Figure 2. Computed solar fluxes for the OKC station, 8-9 July 1992. Clear and cloudy computations are shown separately.

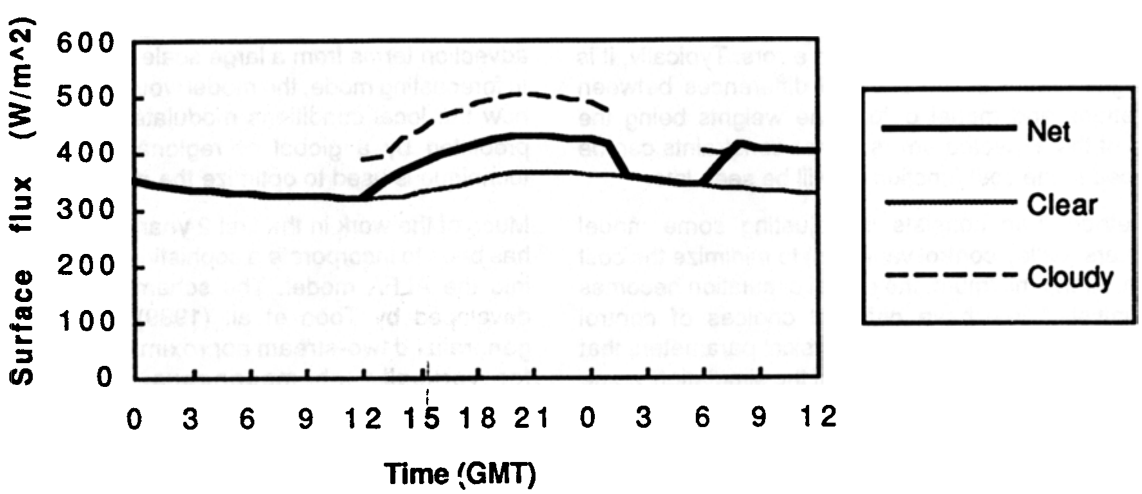

Figure 3. Computed downward infrared fluxes for the OKC station, 8-9 July 1992. Clear and cloudy computations are shown separately.

Note that this kind of output can be compared to radiometric measurements and can therefore be used in computing the cost function in data assimilation. That is something that no other data assimilation method can do easily. In fact, much effort has been put into developing temperature retrieval techniques for satellite data because the operational data assimilation schemes could not use the radiances as input data.

Although writing the adjoint of a model is not terribly difficult, and somewhat mechanical once the linear tangent model has been written, debugging it is quite a challenge. To test the LTM, we estimate the gradients by finite difference (by slightly perturbing the input of the original model) and compare them to the output of the LTM. The adjoint can be tested in two ways: 1) ensure that the scalar product of the input of the LTM by the output of the adjoint is equal to the scalar product of the input of the adjoint by the output of the LTM, or 2) compare the gradient computed by the adjoint with that computed by the LTM. These should agree to within machine precision. Unfortunately both of these tests are global in nature; that is, they test the entire code at once. If a discrepancy is found, there is no indication where the error might be!

Data Assimilation Tests

We are now doing an extensive set of tests of both the model optimization and data assimilation. So far, we have only used NWS data for Oklahoma City, but we will soon start using the ARM data as well.

We show here a couple of examples. In both figures we show the surface temperature observations for 2 days; our ALFA forecasts before optimization, which we also call first guess; the ALFA analysis for the first day; and the resulting forecast for the second day. The first guess forecast is performed with what we think may be "reasonable" physical parameters for the Oklahoma site, with the analysis of the NWS as initial condition.

In Figure 4, we do an optimization of the model parameters, using data for the first day. Although the figure shows only surface temperature, all the available NWS data are used, i.e., surface temperature, dewpoint temperature, winds every hour, and sounding profiles (mandatory levels only) every 12 hours. A 24-hour forecast with the new parameters is then performed from the state at the end of the optimization period.

In Figure 5 we use the Derber nudging algorithm to assimilate data during the first day. We have used constant nudging: at each time step we add constant terms to all the tendency equations. These terms are different for all the variables and also depend on height. They are our control variables. They are all zero at the start of the assimilation procedure. The iterative process is stopped when the forecast error during the first day is minimum. Again, a 24-hour forecast is performed from the end of the assimilation period.

Conclusions

After struggling (somewhat longer than expected) with writing the adjoint of the Toon radiation scheme, we are now at the point where we are beginning to get results of model optimization and data assimilation using the adjoint method. The first results are very encouraging, as can be seen in the figures shown here. We are confident that the variational data assimilation method, using a single-column model, will prove to be a powerful tool for data fusion and data assimilation.

A lot of work remains to be done. Optimization of the model parameters needs to be done with a much longer series of data, to cover different meteorological situations. The Derber nudging method will require considerable tuning, especially in defining the vertical profiles of the nudging terms. Up to now we have let them adjust freely, but that creates a problem when observations are available at only a few levels. A smoothness constraint should probably be enforced.

So far, we have also chosen fairly simple situations, avoiding convective cases. It is not known yet whether the kind of thresholds involved in the convection will create convergence problems in the minimization. Finally we need to develop what might be called "observation simulators," i.e., algorithms to create output similar to the observed quantities, for as many of the ARM instruments as possible.

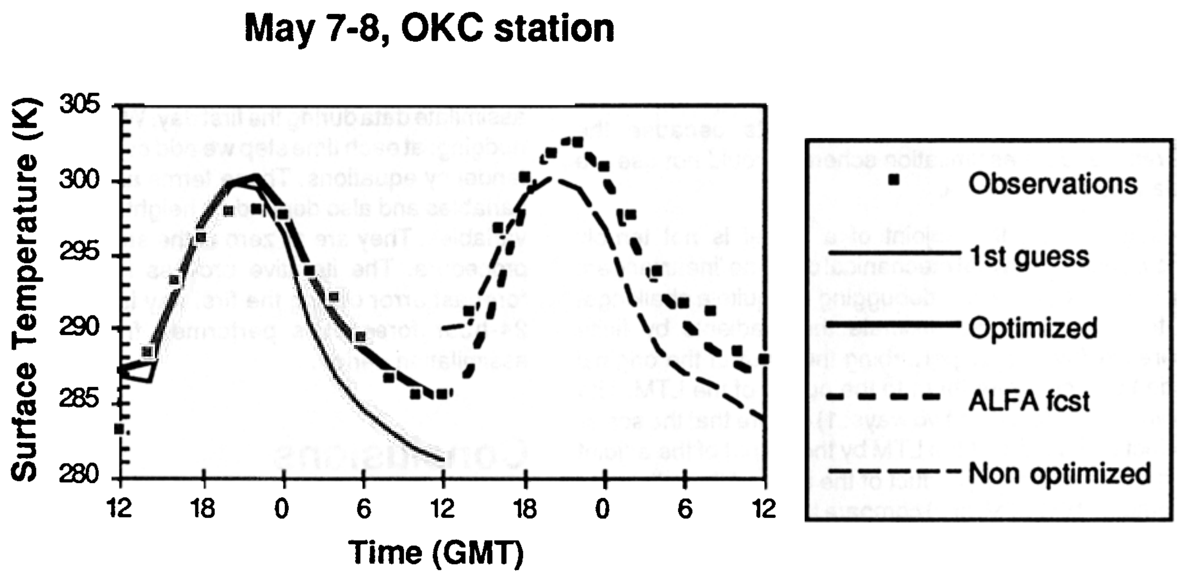

Figure 4. Example of model parameter optimization for Oklahoma City. The optimization is performed during the first 24 hours.

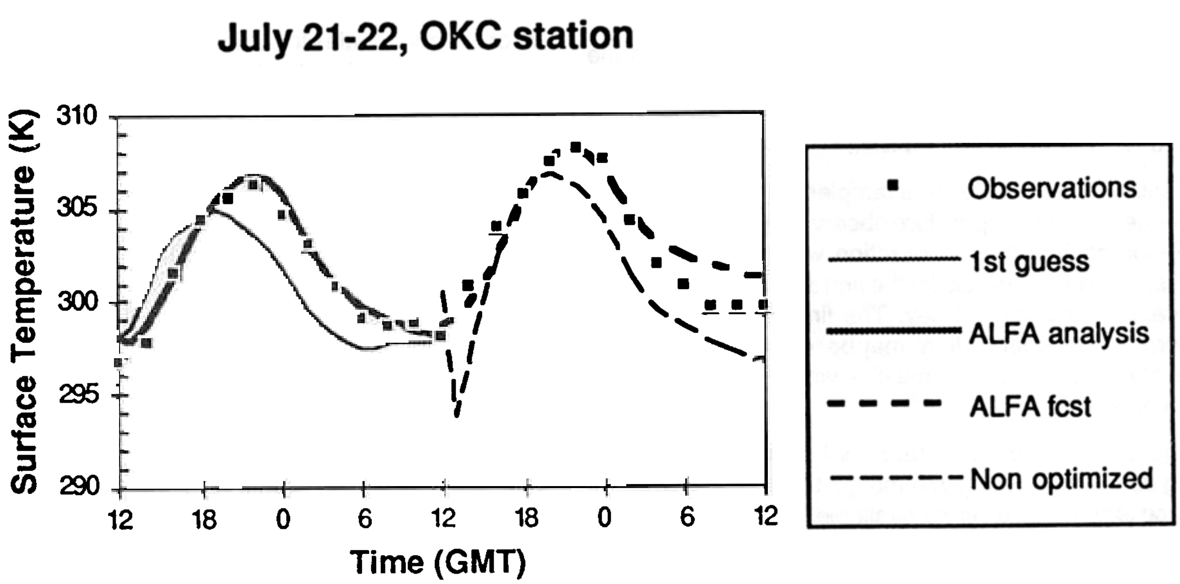

Figure 5. Example of Derber nudging assimilation at Oklahoma City. The assimilation is performed during the first 24 hours.

References

Anthes, R. A., Y.-H. Kuo, S. G. Benjamin, and Y.-F. Li. 1982. The evolution of the mesoscale environment of severe local storms: Preliminary modeling results. *Mon. Wea. Rev.* **110**:1187-1213.

Derber, J. C. 1989. A variational continuous assimilation technique. *Mon. Wea. Rev.* **117**:2437-2446.

Hoffman, R. N., J.-F. Louis, and T. Nehrkorn. 1992. A method for implementing adjoint calculations in the discrete case. Technical Memorandum No. 184. European Centre for Medium-Range Weather Forecasts, Reading, United Kingdom.

Louis, J.-F., R. N. Hoffman, and T. Nehrkorn. 1990. Optimizing a local weather forecast model. NSF Contract No. ISI-8960592. Available from Atmospheric and Environmental Research, Cambridge, MA 02139.

Toon, O. B., C. P. McKay, T. P. Ackerman, and K. Santhanam. 1989. Rapid calculation of radiative heating rates and photodissociation rates in inhomogeneous multiple scattering atmospheres. *J. Geophys. Res.* **94**:16287-16301.