Tellus (2001), 53A, 192–214
Printed in UK. All rights reserved

TELLUS
ISSN 0280–6495

Moisture retrievals from simulated zenith delay "observations" and their impact on short-range precipitation forecasts

By MANUEL S. F. V. DE PONDECA[1]* and XIAOLEI ZOU[2], [1]*National Center for Atmospheric Research, 3450 Mitchell Lane, Boulder, Colorado 80303, USA;* [2]*Department of Meteorology, Florida State University 404 Love Building, Tallahassee, Florida 32306, USA*

(Manuscript received 1 September 1999; in final form 13 September 2000)

ABSTRACT

The feasibility of assimilating the GPS total zenith delay into atmospheric models is investigated within the framework of the "Observing System Simulation Experiment." The total zenith delay is made up of two terms: one is proportional to the pressure at the site of the GPS ground-based receiver and the other to the overlying amount of water vapor. Using the MM5 mesoscale model and its adjoint, a set of 4-dimensional variational (4DVAR) experiments is performed. Results from the assimilation of simulated precipitable water observations are used as the benchmark. The model domain covers Southern California. The observations are simulated with a 10 km horizontal resolution model that includes full physics, while a 20-km resolution and a less comprehensive physics package are used in the 4DVAR experiments. Both, the 10-km and 20-km models employ the same set of 15 vertical levels. Moisture fields retrieved from the total zenith delay are found to compare very well with those retrieved from the precipitable water. Verified against the observations, the vertically integrated moisture is found to be very accurate. An overall improvement is also achieved in the vertical profiles of the moisture fields. The use of the so-called background term and model initialization are shown to greatly reduce the negative impact that the sole assimilation of the total zenith delay can have on the pressure field and integrated water vapor. The adverse effect stems from the poor resolution of the topography needed to evaluate the model pressure at the GPS sites. The analysis increments of all model fields are found to be similar to the counterparts obtained from the assimilation of the precipitable water. The same is true for the short-range precipitation forecasts initiated from the 4DVAR-optimal initial conditions.

1. Introduction

The remote sensing of the atmosphere through the Global Positioning System (GPS) is a state-of-the-art technique that estimates the atmospheric moisture content with a high spatial and temporal resolution (Bevis et al., 1992). Soon, such estimates will be available to the numerical weather prediction community on a routine, real time basis. Through data assimilation, operational centers will in particular increase the ability to forecast storm development and precipitation.

The GPS comprises a constellation of 24 non-stationary satellites put in orbit primarily for geodetic and high precision positioning (Leick, 1990; Dixon, 1991; Ware and Businger 1995). Their usage for meteorological purposes implies only relatively low additional costs (Bevis et al., 1992). These satellites transmit microwave radio signals into the atmosphere, which are picked up either by ground-based receivers or by receivers

* Corresponding author.
e-mail: pondeca@bamboo.met.fsu.edu

carried on low Earth orbit satellites (Businger et al., 1996; Ware et al., 1996). Because the propagation medium differs from a vacuum, the radio signals experience a bending and retardation before reaching the receiver. The total zenith delay (TZD) is the observation of central importance in this paper. It is the delay associated with the (electrically) *neutral* atmosphere as the radio signals propagate along the zenith direction. Estimates of the TZD are derived from signal phase measurements at the ground-based receivers.

As it will become clear shortly, the TZD comprises two terms: the hydrostatic zenith delay (HZD) and the wet zenith delay (WZD). The former represents a rescaling of the pressure at the site of the receiver, while the WZD is proportional to the local precipitable water (PW). It is common to convert the TZD into the meteorologically more appealing PW. This requires an estimate of the pressure at the site of the receiver and weighted vertical mean temperature. In dense GPS networks, however, pressure measurements may not be available for a large number of GPS sites. The question thus arises whether the sole information of TZD can still be useful to the numerical weather prediction community. The purpose of this work is to explore this issue. Through an idealized case study, we use an adjoint based variational method to retrieve moisture fields given either the TZD or the PW. We use the moisture fields retrieved from the PW as a benchmark to evaluate the quality of the fields retrieved from the TZD. In addition, we assess the impact of the retrieved fields on short-range precipitation forecasts. The retrieval experiments make use of the MM5 adjoint modeling system (Zou et al., 1995; Zou et al., 1997). Our proof-of-concept is based on the "Observing System Simulation Experiment (OSSE)," which means that the "observations" are generated by the model itself. An OSSE-setup provides a convenient way of gaining insight into basic theoretical aspects of the problem at hand, as one need not deal with the issue of erroneous (incompatible) observations. In future work, we will present results from the 4DVAR assimilation of real TZD data stemming from a network of ground-based GPS receivers located in Southern California. We note that the results of Kuo et al. (1996) and Xiao et al. (1999) justify our use of the PW derived fields as the benchmark. Using real and simulated observations, these authors demonstrate the feasibility of assimilating PW data into atmospheric models. Their results show an effective recovery of the vertical structure of water vapor and an improvement in the quality of the moisture analysis. Furthermore, when added to the information on wind and temperature, the assimilation of PW data and surface humidity or rainfall leads to significant improvement in short-range precipitation forecasts.

Section 2 reviews briefly the concepts of total, hydrostatic and wet zenith delay. Section 3 contains the model description and overview of the synoptic situation. A brief summary of the method of the 4-dimensional variational data assimilation is found in Section 4. Section 5 presents details on the evaluation of the station pressure and vertical integrals that enter the definition of the model TZD and PW. Section 6 describes the experiment design. The numerical results are presented in Section 7, and the discussion and conclusions in Section 8.

2. The concepts of total, hydrostatic and wet zenith delay

As shown in the literature (Saastamoinen, 1972; Davis et al., 1985; Bevis et al., 1992; Janssen, 1993), it is convenient to separate the TZD into the hydrostatic zenith delay (HZD) and the wet zenith delay (WZD). The latter is caused by the permanent dipole moment of water vapor only, which makes the expression "wet delay" appropriate. The HZD is caused by the induced dipole moments, mainly in N2 and O2, but also in water vapor.

The hydrostatic zenith delay is given in mm by

$$\mathrm{HZD} = (2.2779 \pm 0.0024)\frac{p_s}{f(\lambda, H)}, \tag{1}$$

where p_s is the pressure in hPa at the site of the GPS antenna. H is the height of the antenna above the ellipsoid in kilometers, and

$$f(\lambda, H) = 1 - 0.00266 \cos 2\lambda - 0.00028H \tag{2}$$

is a correction factor that accounts for variations of the acceleration due to gravity with latitude λ and height H.

The wet zenith delay is calculated as

$$\mathrm{WZD} = 10^{-6}\left[\int \frac{e}{T}\left(k_2' + \frac{k_3}{T}\right) \mathrm{d}z\right], \tag{3}$$

where z is the geometrical coordinate along the zenith path, $k_2' = (17 \pm 10)\ \mathrm{K\ hPa^{-1}}$ and $k_3 = (3.776 \pm 0.004) \times 10^5\ \mathrm{K^2\ hPa^{-1}}$. In (3), e is the water vapor pressure in hPa, and T is the temperature (K). The WZD possesses the units of z.

One can show that the WZD is proportional to the precipitable water

$$\mathrm{PW} = \frac{1}{\rho_\mathrm{w}} \int_H^\infty \rho_\mathrm{v}(z)\ \mathrm{d}z, \tag{4}$$

where ρ_w and $\rho_\mathrm{v}(z)$ are the density of liquid water and of water vapor, respectively (Davis et al., 1985). The non-dimensional constant of proportionality is expressed in terms of a suitably defined column mean temperature, the density of water and the constants k_2' and k_3. When expressed in the same units of length, the WZD is approximately 6.4 times the PW (Bevis et al., 1992). This relationship is accurate to better than 1.5% in our case study. Typical values of the HZD at sea level lie around 2.3 m. Clearly, in the absence of rapid pressure changes, as those associated with the passage of a low-level cyclone, the HZD remains fairly constant with time. In contrast, the WZD is a smaller, highly variable quantity in both space and time (Elgered et al., 1991). Values of the WZD range approximately from 1 cm in arid regions to 40 cm in humid regions.

3. Model description and brief synoptic overview

3.1. The model

We use the MM5 Adjoint Modeling System described by Zou et al. (1997). The MM5 is a primitive equation, finite-difference based non-hydrostatic mesoscale model (Anthes and Warner, 1978; Dudhia, 1993; Grell et al., 1994). It uses a sigma vertical coordinate defined in terms of a time independent reference pressure (Dudhia, 1993; Grell et al., 1994). The model includes moisture and planetary boundary layer processes. In our study, the state vector comprises all grid point values of the u, v and w components of the wind field, the temperature, T, perturbation pressure p' (actual pressure minus the reference pressure), and specific humidity, q. The MM5 Adjoint Modeling System also includes the tangent linear and adjoint versions of the MM5 as well as a 4-dimensional variational data assimilation package.

The geographical region targeted in this study is centered around the Los Angeles Basin and stretches roughly from 112°W to 125°W in longitude and from 30°N to 38°N in latitude (Fig. 1). A network of ground-based GPS receivers is under construction in this region. In a future paper, we will discuss the results from the assimilation of TZD data from that network.

3.2. Synoptic overview

The synoptic situation features the passage of a cold front across Southern California between the 9th and the 11th of December 1996. At 1200 UTC 9 December 1996, the system was approaching the coast from the Pacific Ocean (not shown). For 0000 UTC 10 December 1996, Fig. 1 shows the approximate position of the frontal system as obtained from the NCEP global analysis. By 1200 UTC 10 December 1996, much of the front had moved well inland (not shown). Heavy rainfall was registered ahead of this frontal system. Our case study focuses on the period from 1200 UTC 9 December 1996 to 0000 UTC 10 December 1996.

4. Four-dimensional variational (4DVAR) data assimilation

In 4DVAR data assimilation we minimize a cost function J that measures the misfit between the available observations $\{\boldsymbol{y}\}$ and the model forecast analogues:

$$\begin{aligned} J(\boldsymbol{x}_0) &= (\boldsymbol{x}_0 - \boldsymbol{x}_\mathrm{B})^\mathrm{T} \mathscr{B}^{-1} (\boldsymbol{x}_0 - \boldsymbol{x}_\mathrm{B}) \\ &\quad + \sum_n^N \sum_{r=n\mathrm{i}}^{r=n\mathrm{f}} [\boldsymbol{H}^n(\boldsymbol{x}_{t_r}) - \boldsymbol{y}^n_{t_r}]^\mathrm{T} \mathscr{O}^{n\,-1}_{t_r} \\ &\quad \times [\boldsymbol{H}^n(\boldsymbol{x}_{t_r}) - \boldsymbol{y}^n_{t_r}], \end{aligned} \tag{5}$$

where $\boldsymbol{x}_{t_r}$ is the model state vector at time level t_r. In general, we have N types of vectors of observations, $\boldsymbol{y}^n$. The observations are collected at the discrete times t_r, from $t_{n\mathrm{i}}$ to $t_{n\mathrm{f}}$. $\boldsymbol{H}^n$ is the model analogue to $\boldsymbol{y}^n$, and its elements may be nonlinear functionals of the state vector. The first term on the rhs of (5) is the background term. Here, $\boldsymbol{x}_\mathrm{B}$ is

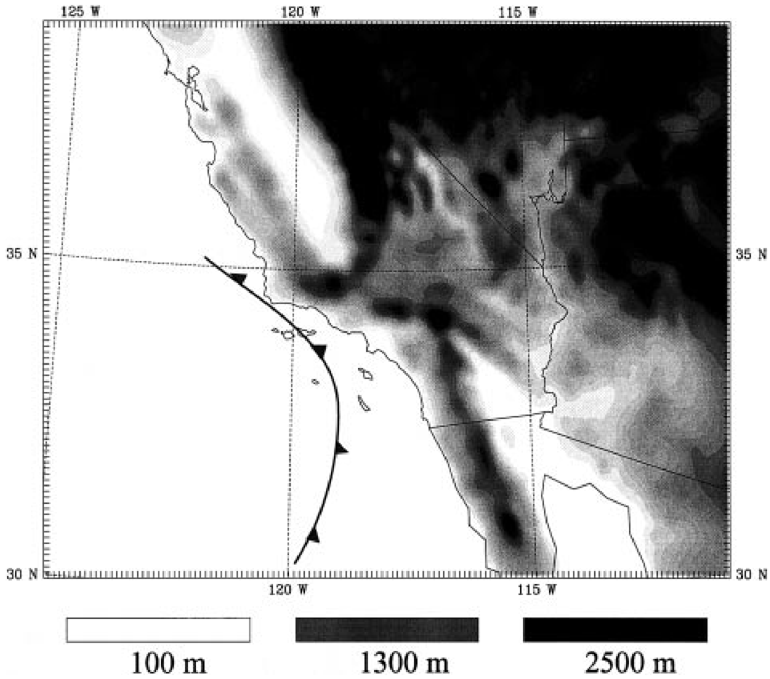

Fig. 1. The 10-km topography used to obtain the GPS station elevations. Boxes below panel illustrate levels of shading associated with 3 different values of the field. Values vary from 0 to roughly 3300 m. Also indicated is the position of the cold front for 0000 UTC 10 December 1995.

the background state vector and $\boldsymbol{x}_0$ the model initial condition. The former introduces a priori information on the atmospheric state. The matrices $\mathscr{B}^{-1}$ and $\mathscr{O}_{t_r}^{n^{-1}}$ are the inverses of the error covariance matrices of the a priori information and the observations, respectively. In the simplest approach, these matrices are taken to be diagonal. Such an approach neglects spatial correlations in the error distribution. We note that the quadratic form chosen for J is consistent with the assumption of Gaussian error distributions (Bennett, 1992). Also, (5) assumes a perfect model framework. We refer the reader to Ide et al. (1997) for an alternative way of writing eq. (5). The uniqueness of the model solution $\boldsymbol{x}_{t_r}$ renders the cost function a sole functional of the initial condition $\boldsymbol{x}_0$. The MM5 Adjoint Modeling System employs the iterative limited-memory Newtonian method after Liu and Nocedal (1989) to find the minimum of $J(\boldsymbol{x}_0)$. The algorithm requires the renewed computation of the gradient of J with respect to $\boldsymbol{x}_0$ at successive iteration levels. That gradient is obtained by integrating the adjoint model backwards in time from the largest value of t_{nf} to t_0. In the course of the integration, appropriate forcing terms are added to the adjoint variables at those time levels when observations are available (Talagrand and Courtier, 1987; Navon et al., 1992).

In this work, special emphasis is put on simulations without the background term. For lack of a better terminology, we call such simulations "(data) retrieval experiments." The exclusion of the background term serves well our purpose of studying the impact of the GPS data alone on the model initial condition. We denote by "optimal initial condition" the initial condition attained through a retrieval or a full 4DVAR experiment.

5. Evaluation of the model PW and TZD

The 4DVAR method requires computation of model-equivalent TZD and PW at every iteration level. The input variables are the model state vector and station elevations. The latter determine the lower boundaries of the vertical integrals of eqs. (3) and (4). They also enter the calculation of the TZD in two additional ways.

1. The correction factor of eq. (2) is a function of the station elevation. However, this is only a weak dependence. For instance, for $\lambda = 0$ rad, HZD changes by less than 0.1% as one goes from a site with $H = 0$ km to a site with $H = 3$ km.

2. The HZD is proportional to the station pressure, which depends strongly on the station elevation.

One finds 3 distinct cases when comparing the station elevation with the model topography.

1. *The station elevation is higher than the model topography*

In this case, we interpolate the station pressure from the pressure at the two sigma levels immediately below and above the station. The interpolation assumes an exponential dependence for $p(z)$, i.e., pressure as a function of the geometric height. We note that in the (non-hydrostatic) MM5 model the heights of each sigma level are fixed in time. The upper-limit of integration in eqs. (3) and (4) is the height of the upper-most sigma level, $z_{\max}$. For clarity, we note that the upper-most sigma level is a 50 hPa constant pressure surface. The lower-limit of integration is the station elevation z_{station}.

2. *The station elevation is lower than the model topography*

In this situation, we interpolate the station pressure from the model surface and sea level pressure values. Consistent with the methodology adopted to diagnose the sea level pressure in the MM5 model, again the interpolation assumes an exponential law for $p(z)$. The vertical integrations are from $z_{\max}$ to z_{station}. In the integration, when z is less than the model topography we take the model fields to be constant. The values are those from the lowest sigma level.

3. *The station elevation coincides with the model topography*

This is a trivial case. The station pressure is the model surface pressure and the vertical integrations are from $z_{\max}$ to the model topography.

While the method adopted to evaluate the model station pressure seems intuitively correct and is supported by simple analytical models of the atmosphere, the definition of the lower boundary in the vertical integrations is a disputable issue. We discuss this issue further in Section 8.

6. Experiment design

6.1. *The nature of the observations*

The observations are the TZD, PW and surface humidity simulated with a 10-km resolution model. The mesh of that model is $97 \times 133 \times 15$, i.e., 97 grid points in the north-south direction, 133 in the east-west direction, and 15 vertical levels. The inclusion of surface humidity in our experiments follows Kuo et al. (1996). These authors find the retrieval of the moisture fields to be significantly improved by adding surface humidity to the information of PW. We simulate the observations with a model configuration that features a bulk-planetary boundary layer, the Grell cumulus parameterization scheme (useful for grid sizes of 10 to 20 km) and explicit moisture physics. The assimilating model has a 20-km horizontal resolution and a mesh of $49 \times 67 \times 15$. It includes a bulk-planetary boundary layer, the Kuo cumulus parameterization scheme (more useful for grid sizes bigger than 30 km), and no explicit moisture physics. We supply observations at every grid point. Such an inter-site distance is in good agreement with the average distance in the main cluster of GPS sites from the Southern California GPS Integrated Network. The total number of stations is clearly much larger in our idealized case study. We use such a large number of stations in order to minimize some of the shortcomings that arise from the fact that a retrieval experiment, as defined in this work, represents an under-determined problem. Also, because our study is primarily intended to assess how well the direct assimilation of the TZD fares compared with the assimilation of the PW, it would appear that the exact number of stations is of minor importance. The observations are directly computed at the co-location points of the 20-km model grid. This is done by interpolating the relevant 10-km model fields to these locations and applying relations (1), (3) and (4). We supply observations of PW and TZD every 15 min, which is a realistic time frequency. Indeed, GPS estimates of PW and TZD from the Southern California Integrated Network are available every 5 min. Observations of surface humidity are prescribed every hour.

We present results from two retrieval experiments and one "quasi-4DVAR" assimilation experiment (Table 1). The observations are the PW and surface humidity in experiment 1, and the TZD and surface humidity in experiments 2. In addition to the TZD and surface humidity, experiment 3 includes the a priori information on pressure. In all simulations, the minimization runs for 30 iterations. The time window is always 3 h, starting 1200 UTC 9 December 1996. We focus on model initial conditions generated for this time and associated short-range precipitation forecasts. Comparisons of the simulated with the actual-observed precipitation patterns are, however, beyond the scope of this work. Such comparisons will be presented in a future paper for a different case study.

Fig. 1 shows the 10-km resolution topography from which the station elevations are obtained. It is apparent from this figure that Southern California is a region of complex topography. The terrain values range from 0 to roughly 3300 m. The field of the difference between the 20-km model topography and the station elevations (not shown) contains isolated values with magnitudes as large as 900 m. Differences of this magnitude can occur also when real GPS data are assimilated into a much higher resolution model.

6.2. *Choice of the guess initial condition (IC)*

The minimization process starts from a prescribed "guess IC." We obtain the guess IC by applying a 5-point smoother (filter) four times to the u, v, T, w and p' fields of the 20-km resolution analysis for 1200 UTC 9 December 1996. The smoother simply replaces the value of the field at a given grid point with the average from the four nearest grid points. The analysis is obtained as part of the MM5 package and consists of a coarse $2.5^\circ \times 2.5^\circ$ NCEP analysis improved with an objective analysis of rawinsonde, surface station and ship observations. We choose the specific humidity (q) of that guess IC in a more elaborate way: Starting from the analysis, we integrate the 20-km model for 48 h and smooth the resulting field 10 times. Appropriate to our proof of concept, this procedure yields a guess moisture field significantly different from the observations at the initial time. For 1200 UTC 9 December 1996, Fig. 2a shows the PW observations. We recall that this is the pattern of the 10-km analysis interpolated into the 20-km grid. Fig. 2b is the PW of the guess IC. The differences between both panels are obvious. The difference pattern (not shown) is dominated by large-scale features. It also reveals that, with the exception of the southwestern portion of the domain, the PW of the guess IC is always larger than the observations. The correlation coefficient (cc) between the fields of Figs. 2a, b is only 85.93%, and the rms error is 0.59 cm. According to the approximate rule given in Section 2, this rms error is equivalent to a rms error of 3.78 cm in the WZD and TZD. The latter value is more than six times larger than the typical GPS-TZD measurement errors, which hardly ever exceed 0.6 cm. Vertical structures of the moisture fields are discussed at a later stage.

Our choice for the guess IC requires further consideration on a different aspect: while the exceptionally degraded moisture field fits well with our proof-of-concept, it remains to demonstrate that the smoothing applied to the remaining fields results in a perturbation (with respect to the analysis) which projects significantly onto non-decaying directions of phase space. Results from model integrations (not shown) reveal a rapid growth of the perturbation during the first 12 h. The e-folding time evaluated from an approximate domain averaged total energy is 1.02 days. This value is comparable to typical e-folding times of perturbation growth in strongly unstable background flows. At the initial time, the perturbation is dominated by small-scale features and its amplitude is scattered mainly over land. Later, the perturbation attains a more defined shape, with amplitude over the Pacific Ocean as well. Also, an increase in the spatial scale takes place during the initial hours of the integration. Baroclinicity is present throughout the 12-h period. A rigorous 4DVAR assimilation experiment would require showing that the guess IC is consistent with the background error distribution. However, this

Table 1. *Description of the 3 experiments performed in this work*

Experiment	PW	TZD	q_{sfc}	$p'(t_0)$
1	yes	no	yes	no
2	no	yes	yes	no
3	no	yes	yes	yes

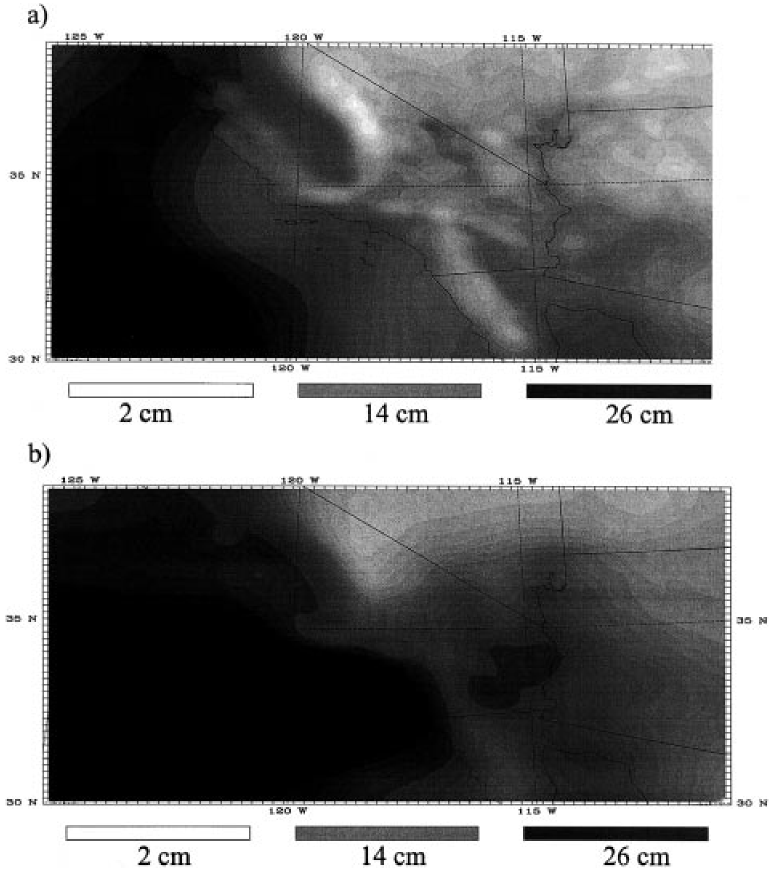

Fig. 2. The precipitable water at 1200 UTC 9 December 1996. Panel (a) displays the observations and (b) the guess field. Values vary from 5 to 38 cm in panel (a) and from 5 to 36 cm in panel (b). Boxes below panels illustrate levels of shading associated with 3 different values of the fields.

appears unnecessary here, as we are mostly concerned with (moisture) retrieval experiments.

7. Numerical results

7.1. *The integrated moisture retrieved from the PW observations*

This is experiment 1 of Table 1. The error covariance matrix of the PW observations is a constant diagonal matrix with the entry of $(1/30)\,\mathrm{cm}^{-2}$ a value borrowed from Kuo et al. (1996). For the surface humidity, the error covariance matrix is diagonal with entries equal to the squares of three-hour forecast differences. Fig. 3 shows the evolution of the cost-function and modulus of its gradient with the iteration number. The cost-function undergoes a monotonic decrease of almost three orders of magnitude, while the modulus of its gradient decreases by two orders. The pattern of the PW of the optimal IC is very similar to the observations (Fig. 2a), and therefore is not displayed here. It correlates at 99.60% with the observations, and the rms error is only 7.37×10^{-2} cm. The largest error at any one grid point amounts to less than 3%. Compared with the values obtained for the guess field (cc $= 85.93\%$ and rmse $= 0.59$ cm), we conclude that the improvement in the PW field is nearly perfect. Our results on the accuracy of the integrated moisture content agree qualitatively well with those of Kuo et al. (1996). We use the results

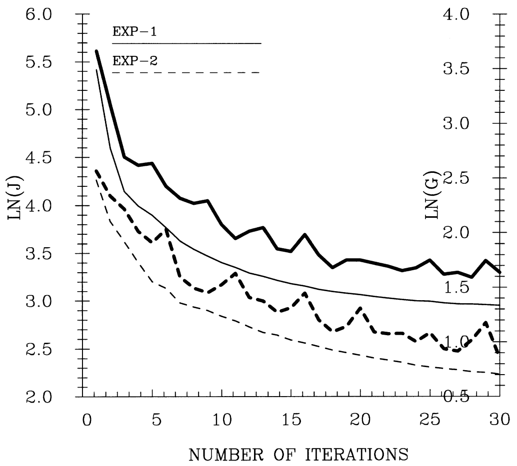

Fig. 3. Base 10 logarithm of the numerical values of the cost function (light lines) and the modulus G of its gradient (darker lines) at the different iteration levels. Results for experiments 1 and 2.

obtained in this subsection as a benchmark to evaluate how the TZD observations fare in a similar retrieval setup.

7.2. *The integrated moisture retrieved from the TZD observations*

In experiment 2, the observations are the TZD and the surface humidity. The error covariance matrix of the TZD is a constant diagonal matrix with the unique entry of $(1/320)\,\mathrm{m}^2$. This value roughly corresponds to the square of the inverse of the maximum 3 h forecast difference found at any one station location. Fig. 4 shows the TZD observations at the initial time. The overall pattern resembles the 10 km resolution surface pressure interpolated into the 20 km model grid. Clearly, this is because the HZD is the dominant term in the TZD.

The evolution of the cost function and magnitude of its gradient are shown in Fig. 3. The same figure also displays results for the benchmark experiment. We recognize that both quantities drop somewhat faster in experiment 2. We note that, due to the different nature of the observations, the comparison is solely based on the slopes of the curves. It is tempting to speculate that the transition from PW to TZD observations may improve the conditioning of the Hessian matrices that enter the minimization algorithm. However, investigations along these lines surpass the main goal of this work.

The optimal IC reveals a remarkable improvement in the integrated water vapor. The PW (not shown) correlates at 99.72% with the observations (Fig. 2a), and the rms error is 6.55×10^{-2} cm. These values compare very well with those from the benchmark simulation (99.60% and 7.37×10^{-2} cm, respectively). A rms error of 6.55×10^{-2} cm in the GPS-PW translates into a 0.42 cm rms error in the TZD. Hence, the retrieval experiment brings the initially large rms error (3.78 cm) to a value comparable to the GPS-TZD measurement error. Fig. 5 shows the difference between the PW of experiment 2 and that of experiment 1. This pattern further confirms that the integrated moisture retrieved from the TZD is comparable to that retrieved from the PW data. The shading in the figure covers a relatively small area. Furthermore, the maximum absolute value in the difference field amounts to 0.21 cm, which is only 1% of the PW observation at that specific location. As discussed later in Subsection 7.5, the difference field of Fig. 5 results mainly from errors in the dry component of the TZD.

7.3. *The vertical profiles of the retrieved moisture fields*

It is important to assess the impact of the retrieval experiments on the vertical distribution of moisture. Despite the lack of a background term, it is reasonable to expect some improvement in the vertical profiles. This is because we supply information on the integrated water vapor at several time levels, while using the forecast model dynamics as a constraint. The results are displayed in Fig. 6. They are for the ten lowest sigma levels, which represent the medium and lower troposphere. The upper-most level ($\sigma = 0.550$) and lower-most ($\sigma = 0.980$) approximate the 570 hPa and 930 hPa pressure levels, respectively. The solid line shows the discrepancies between q^{obs} and q^{guess}. The largest rms errors occur in the lowest model levels. The values exceed $2.00\ \mathrm{g\,kg^{-1}}$ in the lowest two levels. The smallest value is $1.12\ \mathrm{g\,kg^{-1}}$ and occurs in the sixth level. However, for this level, the vertical profile of the correlation coefficient indicates the pattern of q^{guess} has little resemblance with the observations. The correlation coefficient

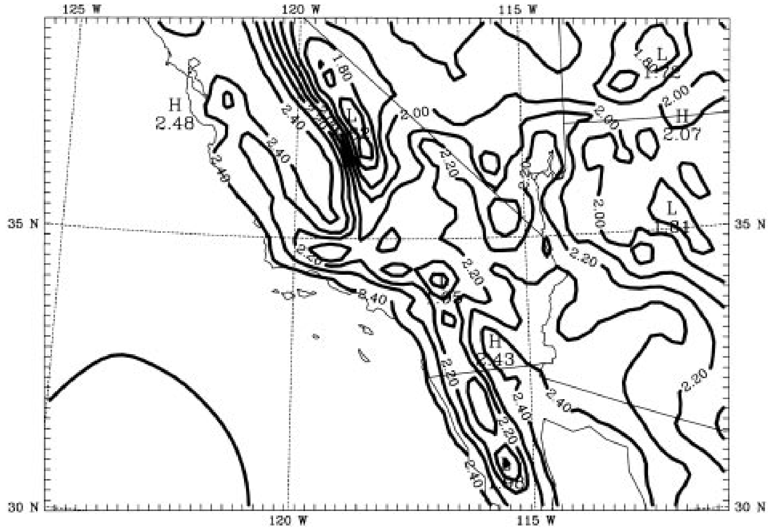

Fig. 4. The observations of TZD for 1200 UTC 9 December 1996. Values in meters. Contour interval is 0.1 m.

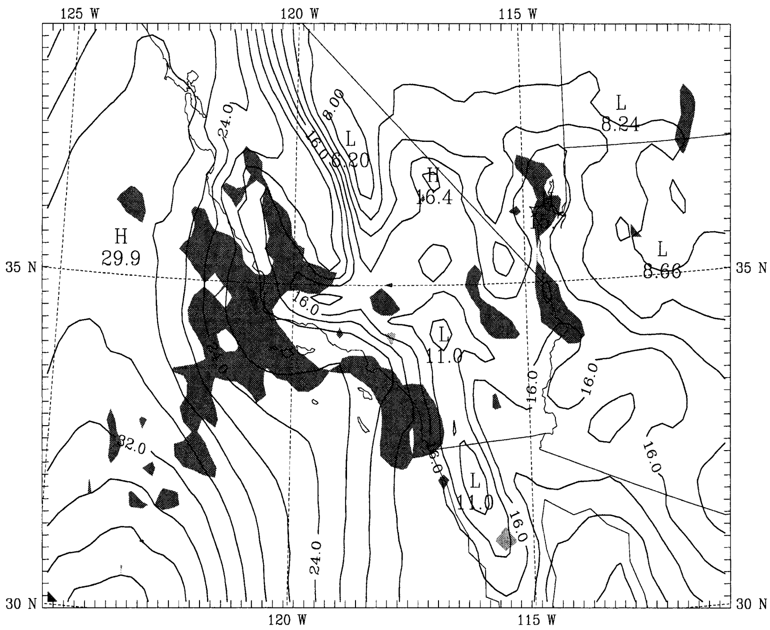

Fig. 5. Field of the difference between the PW of the optimal IC of experiment 2 and that of experiment 1 (shaded areas). Light shading for negative values and darker for positive values. Field values from −2.1 to 2.0 mm. Contour lines represent the PW of the optimal IC of experiment 1. Contour interval is 2 mm.

has its minimum value of 46.77%! The relatively small values of the rms errors found in the upper levels stem from the higher concentration of moisture in the lower troposphere. The average rms error (ARMSE) is 1.64 g kg^{-1}. The correlation coefficient is 92.19% in the lowest two levels and less than 90.00% in the remaining (Fig. 6). It drops quite drastically in the highest levels, so

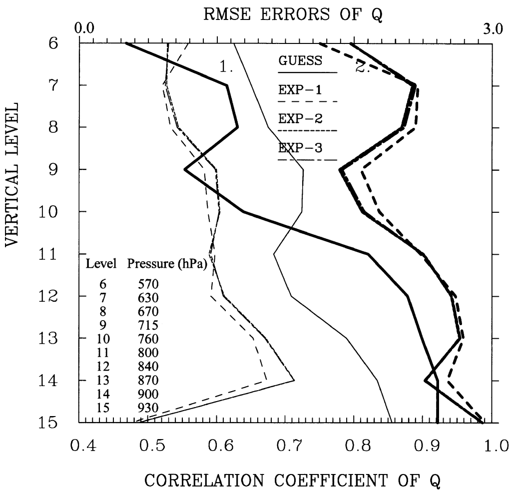

Fig. 6. Vertical profiles of the rms errors (light lines) and correlation coefficients (darker lines) of the specific humidity. Values of the rms errors are given in g/kg. All calculations performed with respect to the observations for 1200 UTC 9 December 1996. Mean pressure values at each vertical level also indicated. Note the rear overlapping of the profiles for experiments 2 and 3.

that the average correlation coefficient (ACC) is only 73.46%. We see from Fig. 6 that both experiment 1 and 2 produce an improvement of the moisture field at all vertical levels. For experiment 1, the ARMSE and ACC are 8.86×10^{-1} g kg^{-1} and 89.10%, respectively. A tremendous improvement is achieved in the lowest level, where the rms error is 3.83×10^{-1} g kg^{-1} and the correlation coefficient is 99.06%. Clearly, this is in large part due to the assimilation of the surface humidity. The improvement in the second-lowest level is only marginal. The rms error is 1.37 g kg^{-1} and the correlation coefficient is 93.58%.

This points to the need for more information on the vertical profile of the moisture field if we are to obtain a better moisture analysis in the optimal IC. These results are in accordance with the work of Kuo et al. (1996), who find the assimilation of PW data to have little effect on the nonbias errors of the specific humidity in some of the verifying sounding stations. We note that the specific levels where the largest errors occur depend on the choice of the guess IC. The results for experiment 2 lie only slightly below those from the benchmark simulation. With the exception of the sixth and eleventh levels, the rms errors are slightly larger in experiment 2. The largest deviation between experiments 2 and 1 occurs in the second-lowest level and amounts to approximately 0.3 g kg^{-1}. The ARMSE is now 9.94×10^{-1} g kg^{-1}, i.e., only slightly larger than in the benchmark simulation (8.86×10^{-1} g kg^{-1}). With the exception of the second-lowest level, the correlation coefficients also show an improvement with respect to the guess field. However, the values are smaller than those for the benchmark simulation. The ACC is now 83.67%, while it was 89.10% in the PW-experiment. For the second-lowest level, the small value of the correlation coefficient (85.42%) indicates a depreciation of the moisture field with respect to the guess IC. This is the level where the benchmark simulation also yields somewhat large errors in the moisture structure. We note that Fig. 6 also displays results from experiment 3, which are discussed later in Subsection 7.5.2.

Fig. 7 shows mean vertical profiles of the rms errors associated with the three different types of GPS stations described in Section 5. There are 827 stations of type 1 (receiver above model topography), 1075 of type 2 (receiver below model topography) and 1266 of type 3 (receiver at the model surface). The figure shows that both experiment 1 and 2 reduce significantly the errors contained in the guess IC. For all three station types, this is always true below the sixth model level ($\approx$ 570 hPa pressure level). We recall that most of the water vapor is concentrated below the 500 hPa level. Also, and important to our proof-of-concept, the results from experiment 2 are comparable to those from the benchmark simulation.

7.4. *Impact on the remaining model variables*

Owing to the model dynamics, perturbations to a given predictive variable are usually accompanied by changes in other variables. However, we do not expect our retrieval experiments to yield significant improvements in the u, v, T and p'-fields. This is because we are essentially assimilating water vapor, which to some degree can be viewed as a passive tracer. Therefore, rather than comparing the new (optimal) fields with the observations, we present the analysis increments. The latter are the differences between the integration from the optimal IC and that from the guess IC. The terminology applies when the differences are

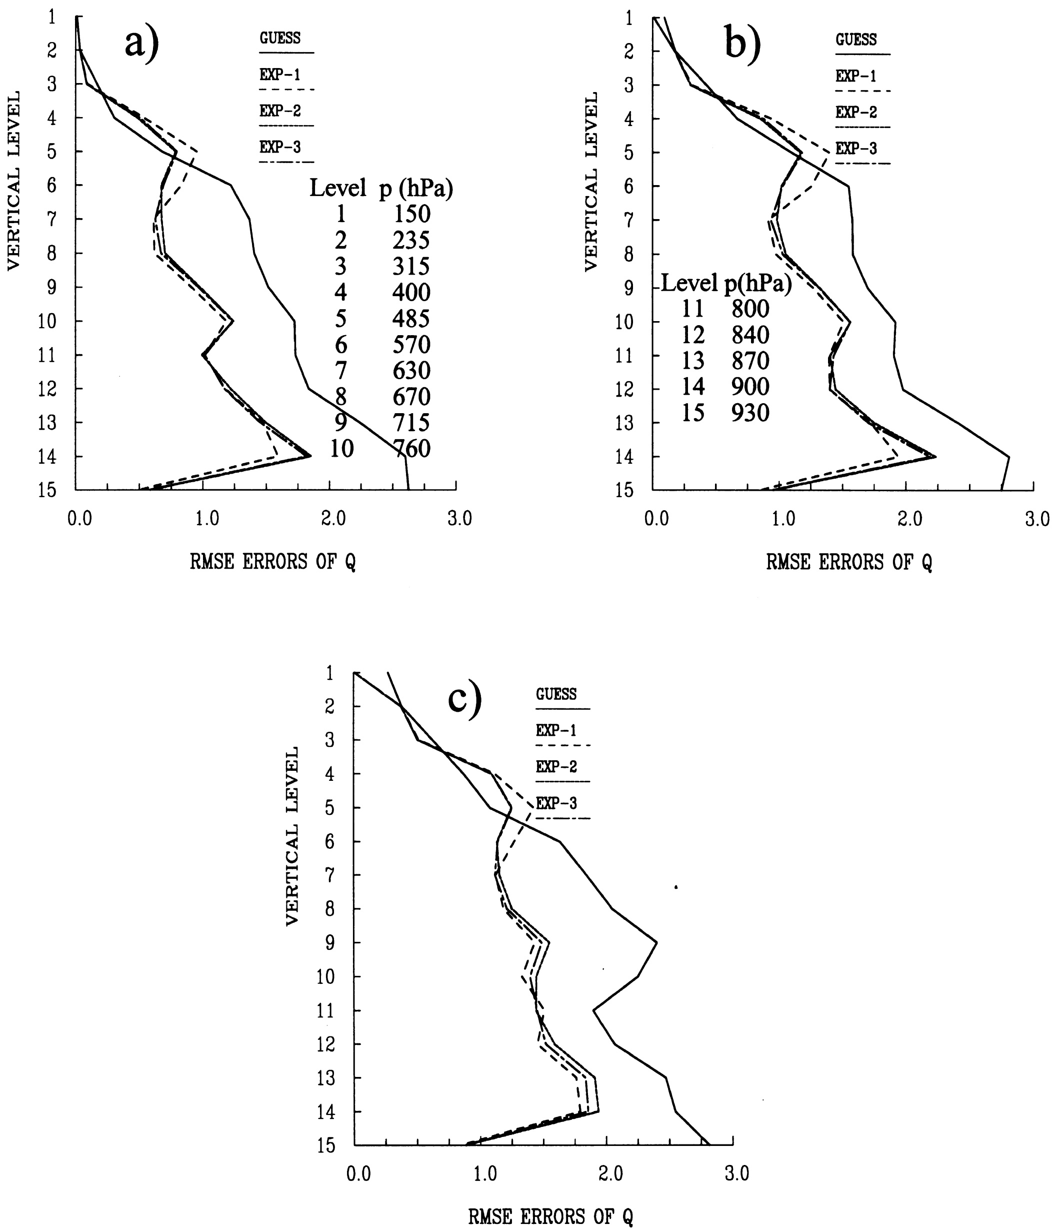

Fig. 7. Mean profiles of the rmse of the specific humidity for all stations with elevations (a) higher than the model topography, (b) lower than the model topography and (c) equal to the model topography. Values in g/kg. Mean pressure values at each vertical level indicated in panels (a) and (b).

taken within the time window. For brevity, we discuss the u-component of the wind field only. Our main conclusions are valid also for the v, w and T fields. As we will see, the p'-field requires a more careful discussion. Fig. 8 shows the analysis increments of the u-field for experiments 1 and 2. They are calculated at the beginning and end of the minimization time window. The patterns are

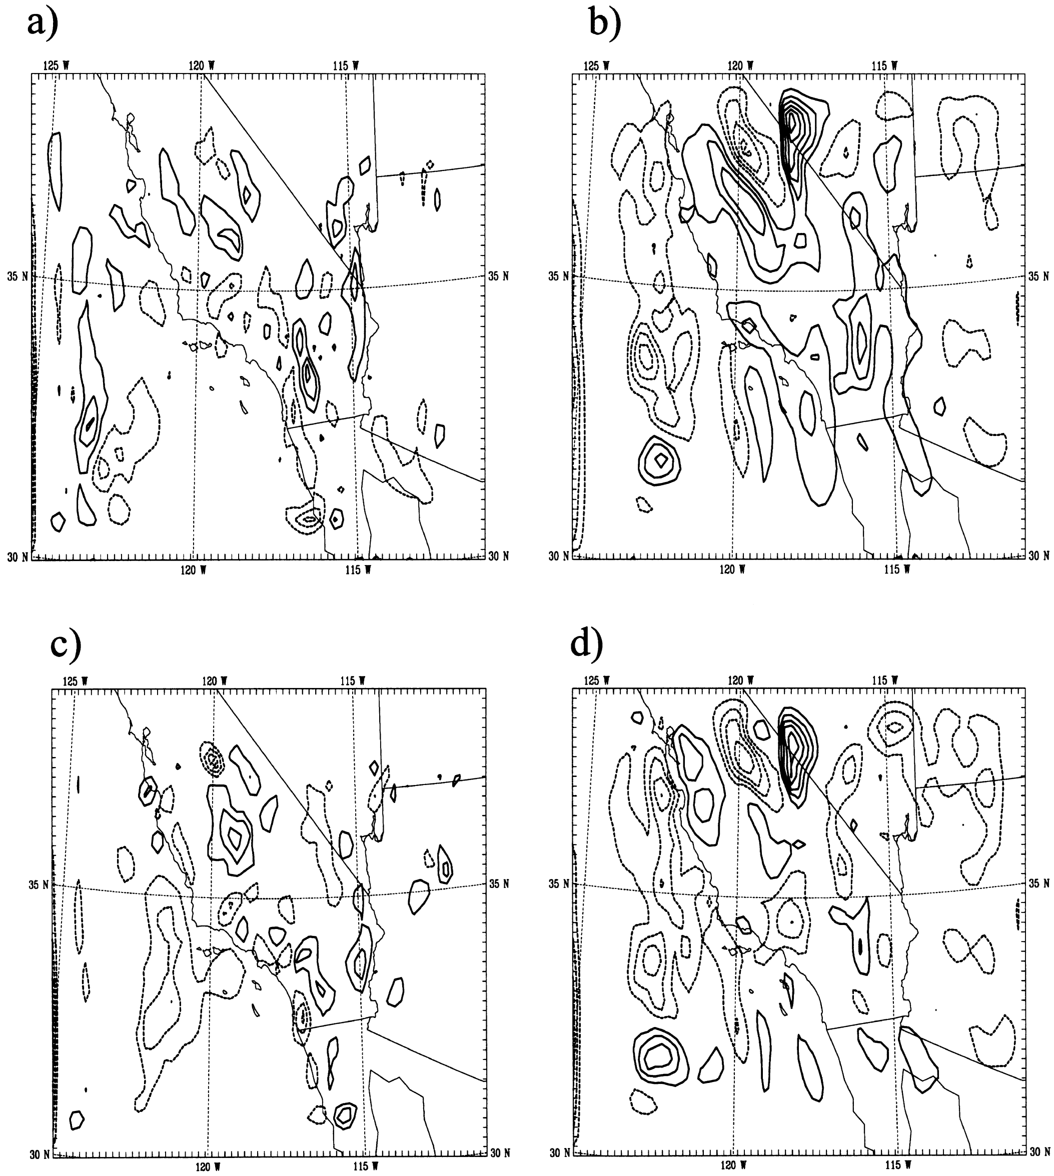

Fig. 8. Analysis increments of the u-field for experiment 1 (panels (a) and (b)) and experiment 2 (panels (c) and (d)). Panels (a) and (c) are the initial increments and (b) and (d) the final increments. Contour interval is 0.5 m/s for the initial time and 1 m/s for the final time. Zero line omitted and dashed lines for negative values. Numerical value of the lowest contour level is the negative of the contour interval. Plots are for the 4th-lowest model level which approximates the 840 hPa pressure level.

for $\sigma = 0.870$, which approximates the 840 hPa pressure level. This is the approximate climatological position of the low-level jet, which is instrumental in most cases of winter season rainfall in Southern California. We see striking similarities in the overall patterns both at the initial and final time. Also, most of the maxima are comparable in magnitude and occur at approximately the same geographical regions. The similarities are even more pronounced at the end of the time window. This suggests that the small deviations found in the optimal IC of experiment 2 project mainly onto decaying directions of phase space. Similar figures were produced for all model variables at every vertical level. Unsatisfactory results are found only for the initial increments of p' in the two-lowest model levels in experiment 2. The patterns contain excessive noise! Fig. 9 displays the lower-level p' increments for experiments 1 and 2. We note the use of different contour intervals in panels (a) and (b) of this figure. For the TZD-experiment, the initial increments represent a negative pattern over most of the domain. However, this is not associated with a biased station pressure operator, as the analysis increments do not correlate with the topographic differences (10-km minus 20-km resolution topographies). The justification for the negative pattern is subtle: as mentioned in Section 6, the PW of the guess IC is larger than the observations for most of the domain. It is also important to note that the guess pressure field is fairly close to the observations. In the lowest two model levels, the largest absolute value of the difference at any one grid point is only 0.66 hPa. Therefore, just as with the PW, the TZD of the guess IC is larger than the observations. We thus expect the minimization procedure to yield an overall reduction of the TZD of the guess IC. Because the TZD-operator does not contain any "bogus" term constraining the changes in p', the "optimal TZD" is attained through a reduction of both the PW and the model station pressure. In sharp contrast with the small maxima around 0.5 hPa found for the benchmark experiment, experiment 2 yields initial maxima with magnitudes as high as 2.85 hPa. By the end of the time window, however, the analysis increments in both experiments are comparable. A closer look at the integration from the optimal IC of experiment 2 (not shown) reveals the only non-decaying structure is the small focus of positive values initially around 36°N, 118°W. This is also the location of the largest initial increments in experiment 1. Interestingly, we also note that the pattern of Fig. 9b correlates well with the difference field of Fig. 5. This is further discussed in Subsection 7.5.2.

Plots of the gradient of the response function $(\nabla J)_{x_0}$ were made for various iteration levels (not shown). They reveal that, though the details in the fields may be different, the geographical regions of largest magnitudes are fairly the same in experiments 1 and 2. Just as with the analysis increments, the only exception is found in the p'-component of the gradient at the lowest model levels. This result suggests the minimization proceeds in a comparable way in both experiments.

7.5. *Improving the retrieval setup*

7.5.1. Model initialization. As we mentioned, there are increased similarities in the analysis increments of experiments 1 and 2 at the end of the assimilation time window. This suggests the noise found in the p'-field of experiment 2 projects strongly onto high frequency gravity-inertia waves. If this is the case, model initialization should help improve the p'-pattern. To test our hypothesis, we use the digital filter of Lynch and Huang (1992). In a simple-minded way, one would "digitally" remove higher frequencies from a signal $f(t)$ by multiplying its Fourier transform $[F(\omega)]$ with the Heaviside function $H(\omega)$ and finally performing the inverse Fourier transform on the product. The Heaviside function $H(\omega)$ is defined through a cutoff frequency ω_c. The result from the inverse Fourier transform is the desired filtered signal $f^*(t)$. The above steps are equivalent to performing the convolution of $f(t)$ with $h(t)$, where $h(t)$ is the inverse Fourier transform of $H(\omega)$. In our application, the $f(t)$'s represent finite time series of the model predictive variables. Gibbs oscillations due to truncation are reduced by multiplying the Fourier components of the Heaviside function with a Lanczos window. The implementation of the filter involves short-term adiabatic forward and backward model integrations from the originally noisy IC. The filtered IC is a linear combination of the model outputs from both integrations at regular time intervals.

We apply the digital filter to the optimal IC of experiment 2. The cutoff period is 3 h. Following Lynch and Huang (1992), we use the mean absolute tendency of the lowest-level p' as a measure of the high frequency noise. We find the filtering to reduce the initial noise by more than 90%. Furthermore, the reduced noise level remains fairly constant with time (figure not shown). Fig. 9c shows the new initial analysis increment of the

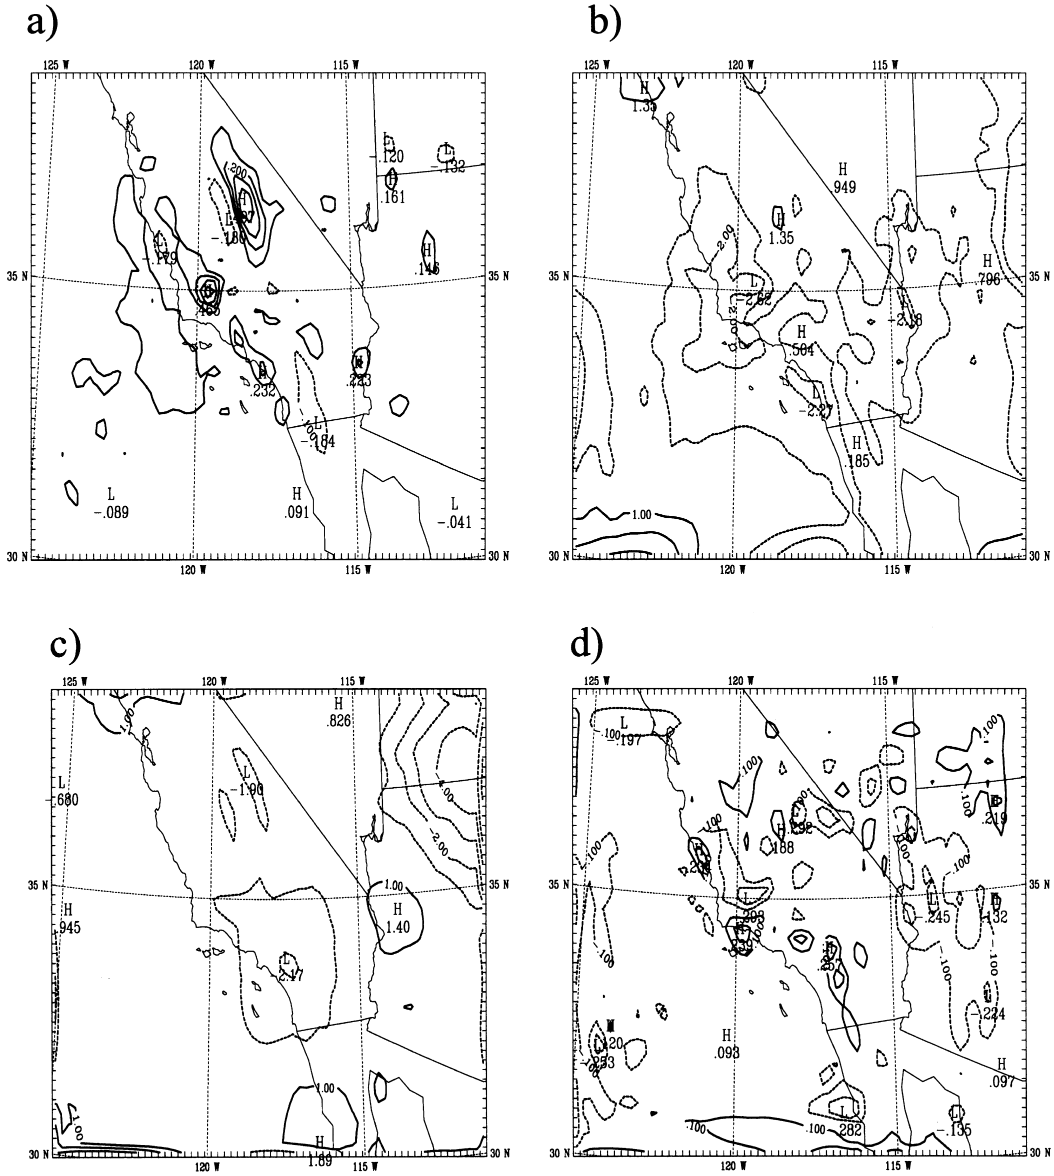

Fig. 9. Initial analysis increments of the lowest-level perturbation pressure for (a) experiment 1, (b) experiment 2, (c) experiment 2 with application of the digital filter and (d) experiment 3. Contour interval is 0.1 hPa in panels (a) and (d) and 1 hPa in panels (b) and (c). Zero line omitted and dashed lines for negative values.

lower-level p'. It is computed as the filtered optimal IC minus the guess IC. Compared with Fig. 9b, we see a reduction of the size of the large-scale negative pattern. The absolute values of the minima, however, are still somewhat large. On a more negative note, we see new contour lines on the northeastern corner of the domain. Because the initialization should in fact be incorporated into the nonlinear and adjoint models, we believe our results may be further improved.

7.5.2. The rôle of the background term. To further our understanding of the negative impact of the TZD data on the p'-field, we recall that the TZD comprises two terms: a moisture term and a term proportional to the station pressure (Section 2). As we will see shortly, we can ascribe the noise in the p'-field to shortcomings in modeling the station pressure with the 20-km model.

Our station pressure operator yields average errors of 1.0 to 1.5 hPa. To show this, we compute differences between the values from the operator and the true values of station pressure. The former are obtained from a 20-km model integration, while the latter are the surface pressure values forecast with the 10-km model. The station elevations that enter the operator are the 10-km terrain values at the co-location points of the coarser grid. Table 2 shows results for an integration period of 3 h, with verification every 30 min. In both integrations, the initial condition is the analysis for 1200 UTC 9 December 1996. The model options coincide with those used to generate the observations (Section 6). It seems appropriate to proceed in this way for two reasons.

Table 2. *Maximum and minimum differences between the model station pressure and the "true" values*

Time (min)	I1	J1	I2	J2	Min value (hPa)	Max value (hPa)
000	35	31	37	34	−2.78	+2.76
030	39	31	36	49	−1.07	+1.74
060	39	31	36	49	−2.47	+1.40
090	39	31	39	57	−2.17	+1.18
120	39	31	36	49	−1.24	+1.40
150	36	32	39	32	−1.03	+1.08
180	36	32	39	32	−0.88	+1.16

Only stations at least 10 grid points (180km) away from the lateral boundaries are considered. The "true" values are the surface pressure values forecast with the 10-km resolution model and interpolated to the 20-km resolution grid. The model station pressure is computed from the 20-km model state vector using the procedure described in Section 5. The same physical options were used to run the 10-km and 20-km models. In both runs, the IC is the analysis for 12000 UTC 9 December 1996. The grid point location (I,J) where these values occur are also indicated. (I) stands for the north-south position in the mesoscale grid and (J) for the east-west location.

1. Both the 20-km and the 10-km resolution analysis are actually created from the same data set.

2. The patterns of the sea level pressure computed from both initial conditions are very similar (not shown). The maximum relative difference at any one grid point is less than 0.1%. This gives us confidence that the operator is essentially accounting for the different topographies. For clarity, we note that the sea level pressure is diagnosed from the three-dimensional fields of p', T and q.

Shown in Table 2 are the minimum and maximum values of the pressure difference (observation minus model value) found at any location in the domain. Only stations at least 10 grid points (180 km) away from the lateral boundaries are considered. During the first minutes of integration, the discrepancies between the model and the true values can be as large as almost 3.0 hPa. This is likely associated with the method used to form the initial condition at the various grid resolutions and the dynamical adjustment present in the early stages of the model integration. Subsequently, the discrepancies decrease by more than 1.0 hPa.

A full 4DVAR problem is over-determined, and its solution is the least squares fit of the model trajectory to the observations. This suggests the use of additional information on the p'-field to compensate for shortcomings in modeling the station pressure. The rôle of the additional information is to mask the errors originated from the assimilation of the TZD. Because the premise of this work is the absence of pressure readings at the GPS sites, the additional information must come from the background term only. We repeat the TZD-assimilation experiment with the difference that eq. (5) includes a priori information on p'- at all model grid points. The a priori information is the 20-km analysis of p' for 1200 UTC 9 December 1996. The covariance matrix is again diagonal with entries equal to the squares of three-hour forecast differences. This simulation constitutes experiment 3 (Table 1). Fig. 9d shows the initial analysis increments of the lower-level p' for this experiment. The pattern now compares favorably with the benchmark (Fig. 9a), in the sense that the magnitudes of the increments are quite small. The PW of the optimal IC (not shown) is again very similar to the observations. The rms error between that pattern and the observations

(Fig. 2a) is 6.17×10^{-2} cm, and the correlation coefficient is 99.74%. Fig. 10 shows the difference between the PW of experiment 3 and that of experiment 1. Compared with Fig. 5, we see a dramatic reduction of the size of the pattern. We recall that Fig. 5 represents the difference between the PW of experiment 2 and that of experiment 1. This result shows that the positive pattern of Fig. 5 is essentially "forced" by the pressure perturbations. This does not come as a surprise, as the pressure gradient is a forcing term in the momentum equations. The wind field, on its turn, determines the water vapor distribution. We recall that water vapor is to some extent a passive tracer. For experiment 3, the vertical profiles of the rms errors and correlation coefficients of the q-field are shown in Fig. 6. Also, the mean errors in the vertical profiles of the specific humidity are displayed in Fig. 7. The results in both figures are nearly identical to those from experiment 2.

7.6. Impact on short-range precipitation forecasts

We perform 12-h model integrations to assess the impact of the retrieved moisture fields on short-range precipitation forecasts. The model options are those used to generate the observations. In particular, we use the Grell cumulus parameterization scheme. Since the results are derived from a 20-km model, it is convenient to define the "observed precipitation" through a 20-km model integration from the initial observations. Alternatively, one could define it as the 10-km precipitation interpolated into the 20-km grid. However, due to marked topographic differences, the coarse model integrations would always miss some of the details in the observed precipitation. This is because orographic enhancement of precipitation is a trademark of winter storms over Southern California. Fig. 11 shows the observed 3-h accumulated precipitation ending 1500, 1800, 2100 UTC 9 December 1996 and 0000 UTC 10 December 1996. We see a rainband moving eastward at approximately 15 m s^{-1}. The 3-h rainfall verified at 2100 UTC 9 December shows a well-defined northwest-southeast tilting band with maxima near the coast line. During the following three hours, most of the rainband moves into the land mass. Figs. 12, 13 and 14 display the results of the integration from the guess IC, and the

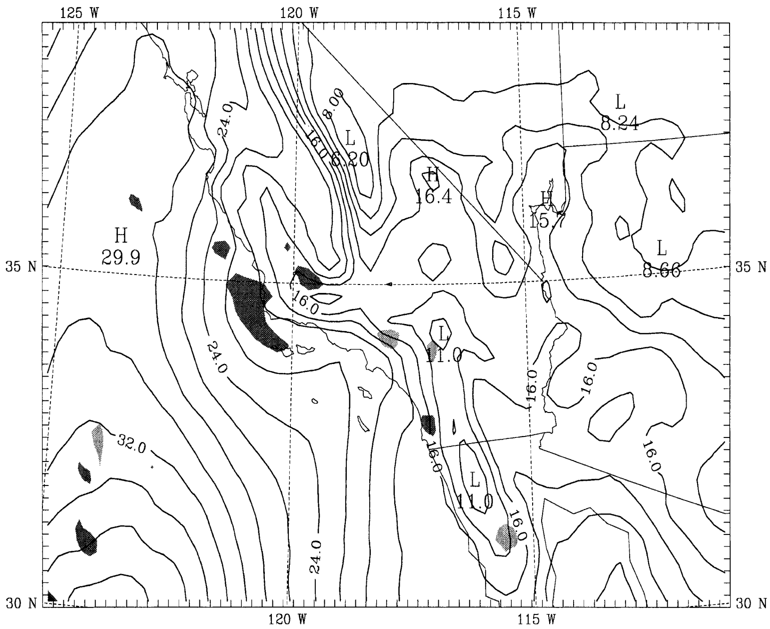

Fig. 10. Field of the difference between the PW of the optimal IC of experiment 3 and that of experiment 1 (shaded areas). Light shading for negative values and darker for positive values. Field values from -2.0 to 1.1 mm. Contour lines represent the PW of the optimal IC of experiment 1. Contour interval is 2 mm.

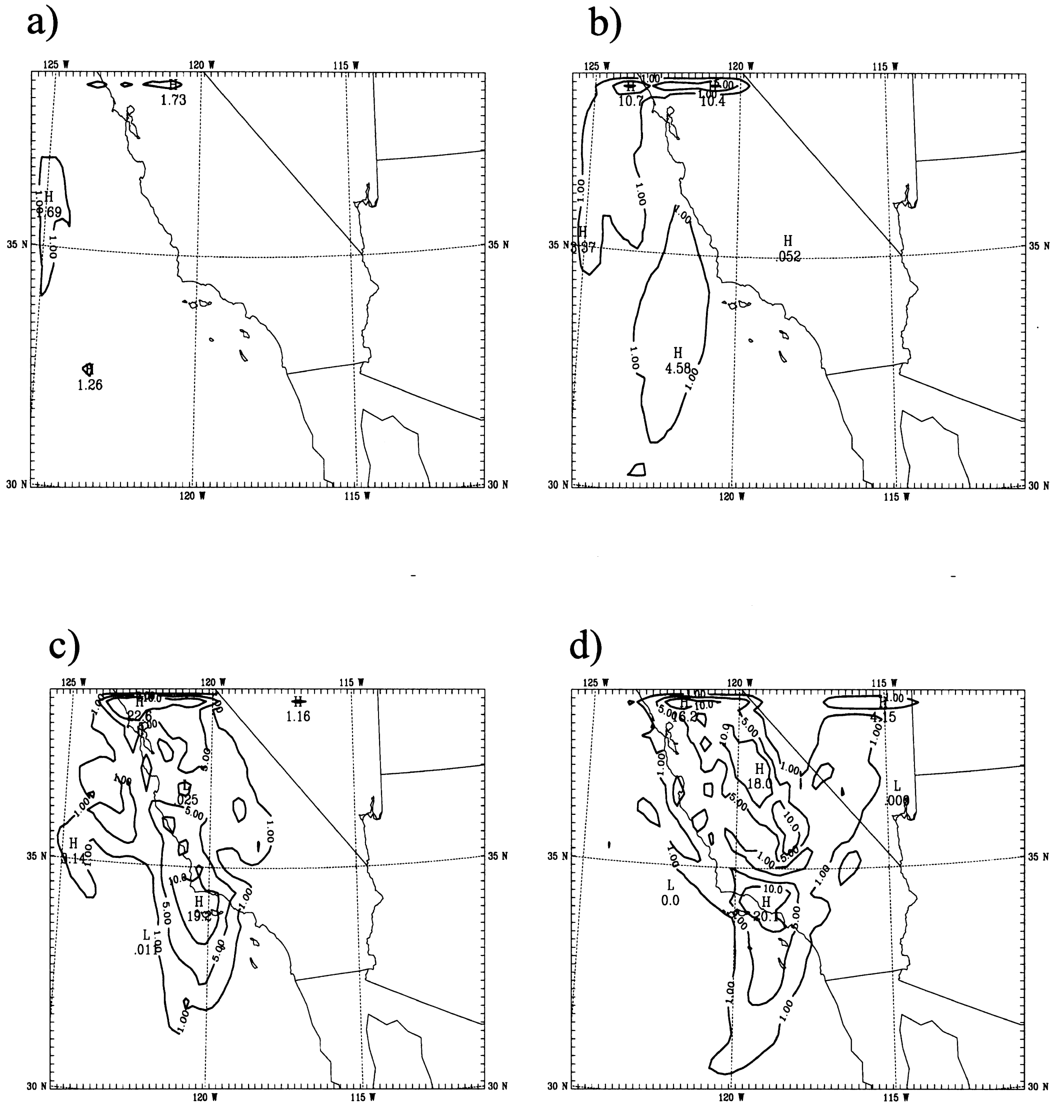

Fig. 11. The 3-h accumulated precipitation ending (a) 1500, (b) 1800, (c) 2100 UTC 9 December 1996 and (d) 0000 UTC 10 December 1996. Patterns obtained from a 20-km model integration starting with the IC of the observations at 1200 UTC 9 December 1996. Contour interval is 5.0 mm and lowest contour level is 1.0 mm.

optimal IC's of experiments 1 and 2, respectively. We recall that the initial flow fields of the guess and optimal IC's differ from the observed flow by small amplitude perturbations only. This is a desirable aspect, as it allows isolation of the relative impact of the different moisture structures on the precipitation forecast (Kuo, et al. 1996). Compared with Fig. 11, the integration from the guess IC misplaces the rainband. The precipitation verified at 1800 and 2100 UTC clearly shows the rainband is too far to the east. Also, the guess IC exaggerates the amount of precipitation. We see values in excess of 1 cm h^{-1} especially in the 3-h rainfall verified at 2100

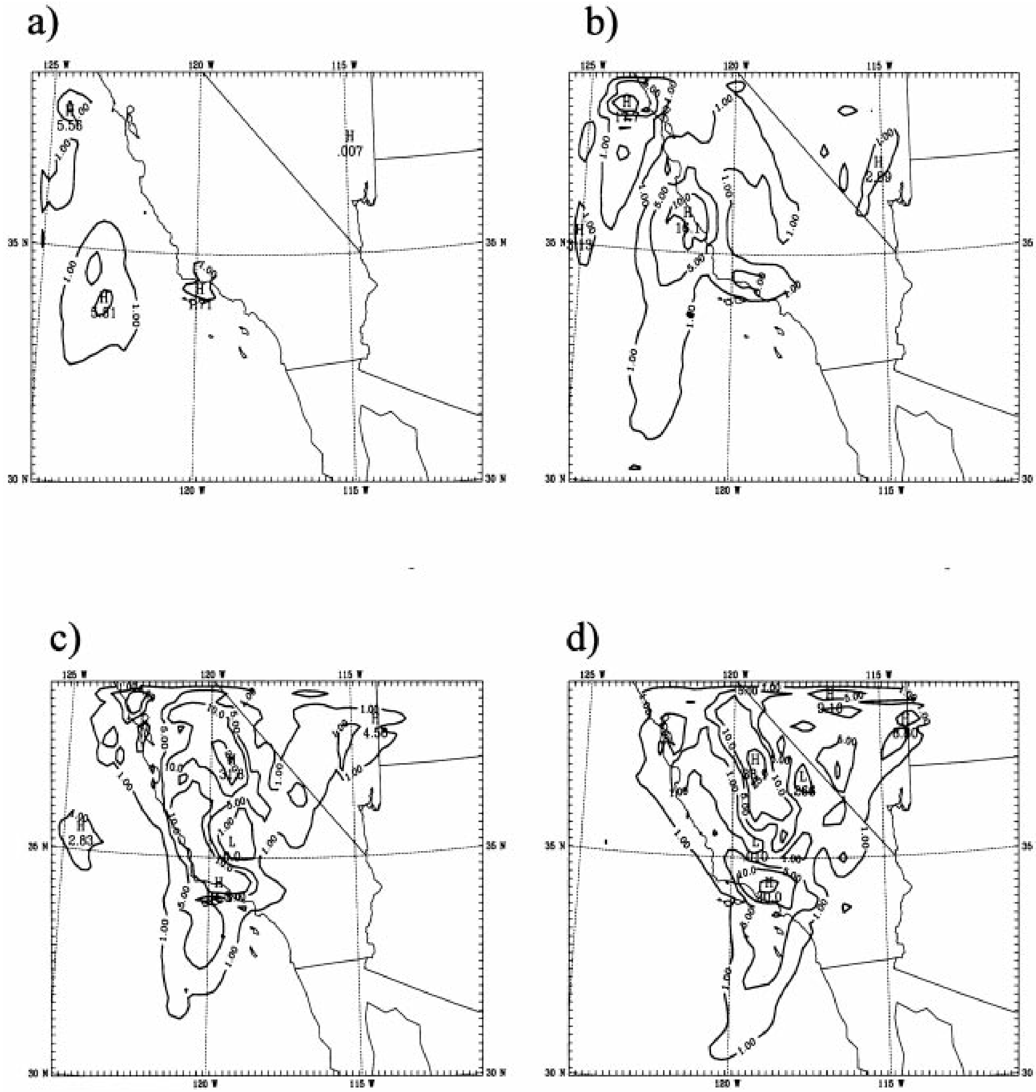

Fig. 12. As in Fig. 11, except for the integration from the guess IC.

UTC. Some of the discrepancies are due to the early orographic precipitation enhancement that occurs in the integration from the guess IC. We note that, unlike the observations, the PW of the guess IC has its maximum very close to the land mass (Fig. 2). The similarities between Figs. 13 and 14 and the observed rainfall (Fig. 11) are rather obvious. The integrations from the optimal IC's of experiments 1 and 2 successfully place the rainband at the correct location. Also, the maxima in the rainbands are very similar and comparable to those from the benchmark. These results are very encouraging for two reasons.

1. The PW and TZD contain no information on the vertical structure of moisture and yet their assimilation leads to remarkable improvements in short-range precipitation forecasts. This

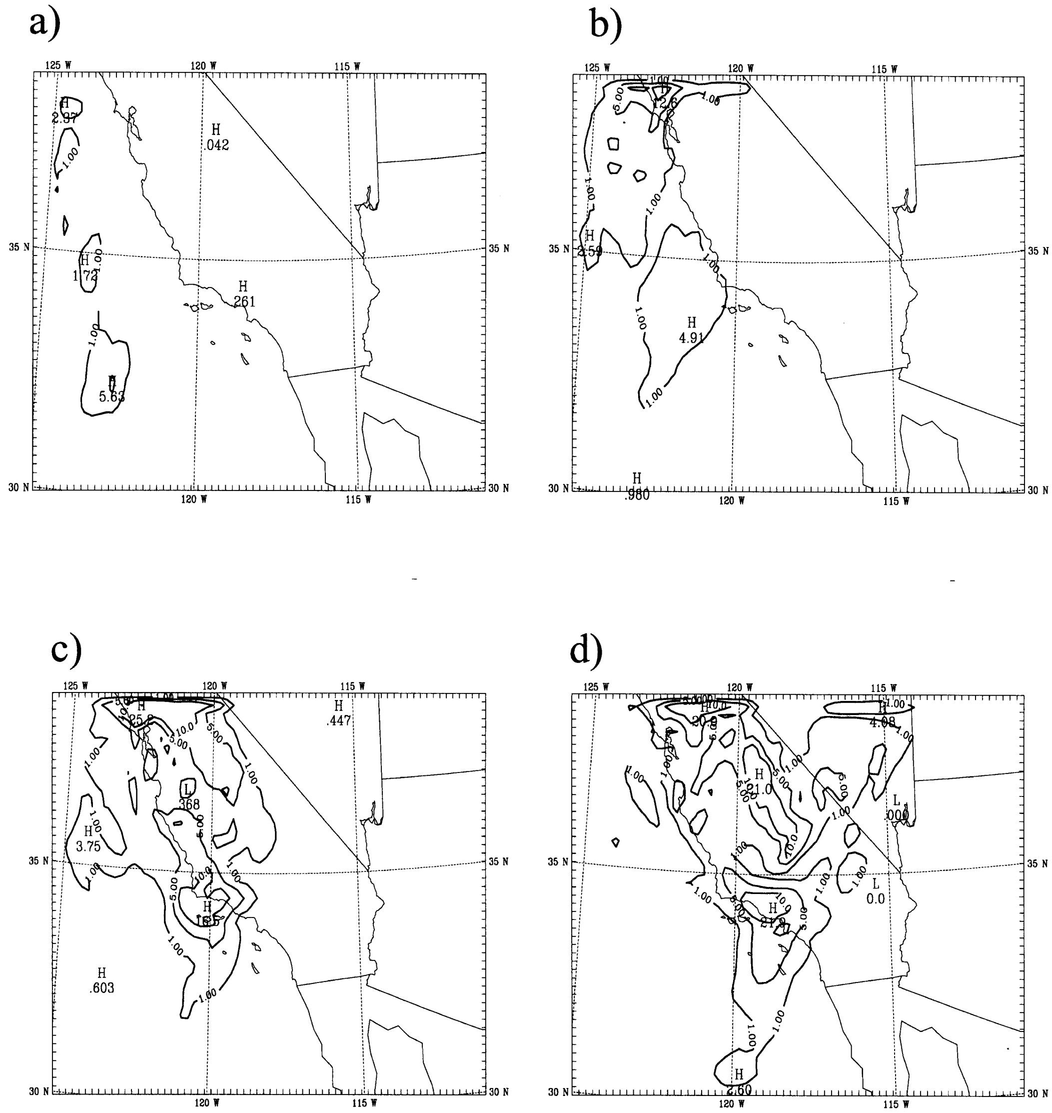

Fig. 13. As in Fig. 11, except for the integration from the optimal IC of experiment 1.

is even more encouraging in view of the somewhat large errors found in the second lowest-level q-patterns of the optimal IC's of experiments 1 and 2.

2. The results suggest that the direct assimilation of the TZD has a similar impact on short range precipitation forecasts as that attained by assimilating the PW.

8. Summary and discussion

TZD estimates are readily available from phase signal measurements provided by ground-based GPS receivers. Converting the TZD into the PW requires the availability of meteorological instruments at the GPS sites. It is natural to inquire about the feasibility of directly assimilating the

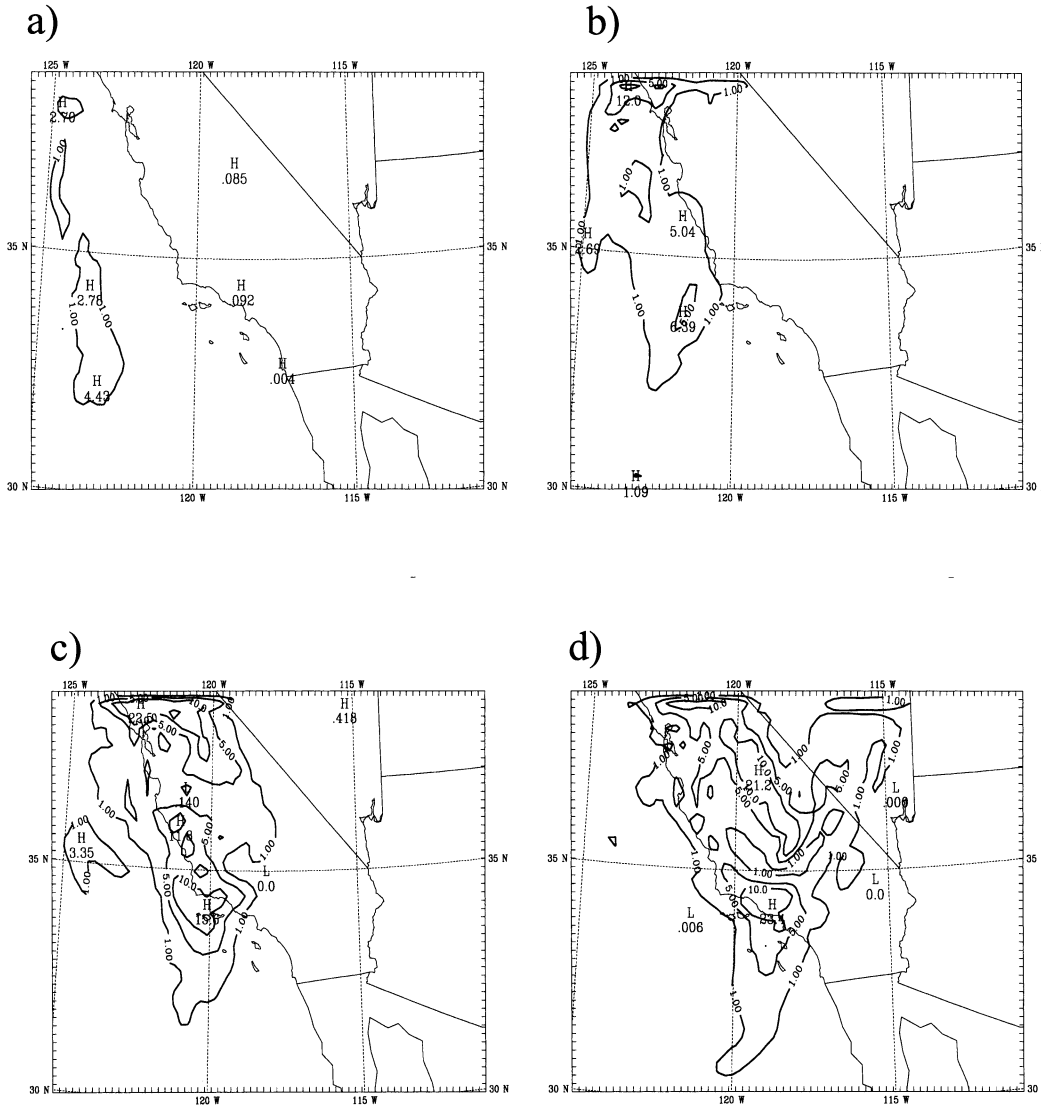

Fig. 14. As in Fig. 11, except for the integration from the optimal IC of experiment 2.

TZD into atmospheric models. We have used an OSSE-setup and a 4DVAR method to tackle this issue. The results from the assimilation of simulated PW observations were used as the benchmark. The geographical area targeted was Southern California, where a network of ground-based GPS receivers is under construction. The main results are as follows.

1. The convergence in the TZD-minimization experiment is slightly faster than in the benchmark experiment. However, it is premature to draw conclusions regarding the possible "pre-conditioning" of the cost-function associated with the TZD data.

2. The assimilation of the TZD allows the retrieval of the vertically integrated water vapor pattern with a very high accuracy, comparable to that obtained through the assimilation of the PW itself.

3. During the minimization process, the gradi-

ent $(\nabla J)_{x_0}$ possesses its largest magnitudes in the same geographical areas as in the benchmark experiment. The same is true for the analysis increments. The only exceptions were the initial increment of p' at the two lowest-levels. The fields contained excessive noise in the TZD-experiment.

4. The optimal initial condition can contain local errors as large as 3 hPa in the lowest model levels when TZD alone is assimilated. This is particularly true for orographically complex regions. Such errors result from the often drastic height differences between the model topography and the GPS station elevations. The field of the difference between the model topography (20-km topography) and the station elevations (10-km topography) contains values ranging from -300 m to 900 m. Differences this large may seem at first exaggerated, thus reflecting a poor proof-of-concept experimental set-up. However, this is not the case. Indeed, compared with the true elevations of the GPS antennas from the Southern California Integrated GPS network, a 6-km model contains height differences of the same order of magnitude. In order for the assimilation of the TZD to be equivalent in accuracy to the assimilation of the PW, the model should simulate the station pressure with a barometer accuracy. When available, measurements of the station pressure at the GPS sites are made with an accuracy better than 0.4 hPa. In orographically complex regions, it may prove very difficult to achieve such an accuracy with the model resolutions affordable with present-day computers. This is even so if the method of data assimilation uses a 4DVAR algorithm, which is computationally rather expensive. Our methodology to evaluate the station pressure is supported by theoretical models of the atmosphere and yields average errors of the order of 1.0 to 1.5 hPa. Errors of this magnitude can alone account for half of the measurement error of the TZD. The use of a digital filter on the optimal IC was shown to reduce these errors substantially. This proves these errors project strongly onto gravity-inertia waves. Also, and most importantly, the inclusion of background information of pressure lead to an overall improvement of the results. The additional information essentially "masks" the errors associated with the evaluation of the station pressure. This was also shown to lead to an improved field of integrated water vapor.

5. In our case study, the assimilation of the TZD leads to remarkable improvements in 3- to 12-h precipitation forecasts. The results are comparable to those attained through the assimilation of the PW.

The study suggests that a full 4DVAR assimilation of the TZD yields moisture fields quantitatively close to those obtained through the assimilation of PW data. By "full" it is meant a set-up that includes the background term and model initialization. Moreover, the quality of our results is likely to improve in geographical regions where topography plays a less important rôle. Further investigation, however, may be necessary: one should study the sensitivity of the results to the incorporation of more realistic error covariance matrices for both the GPS observations and the background term. The definition of the lower integration boundaries in eqs. (3) and (4) also deserves further attention. We repeated the PW-experiment but with the lower boundary always coincident with the model topographic height. This approach is based on the assumption that the model dynamics itself adjusts all vertically integrated quantities. In other words, no further topographic related "refinements" are necessary when verifying against the observations. We found the results to differ very little from those of the original PW-experiment. In particular, the second-lowest level shows errors in the specific humidity field very similar to those found in the original simulation. A study with some statistical value should, however, be performed to aid in justifying the choice of the lower integration boundary. For example, it is conceivable that different seasons and depths of the marine boundary layer in coastal Southern California require different treatments of the lower integration boundary. Here, the approach of Cucurull et al. (2000) may prove more useful. These authors move the entire model profile from the model topography to the height of each GPS station with preservation of the stability properties of the planetary boundary layer. Finally, we note that our experiment set-up would reasonably correspond to a real situation: surface moisture data are often available with hourly periodicity and GPS observations with a temporal frequency of 5 to 15 min. Also, for our domain size, a time window of 3 h roughly corresponds to the maximum permissible window, if the simulations are to be carried out within an acceptable period of time. To make the point, we note that the J90 Cray CPU time for the PW-experiment is 21.57 h. Ongoing code vectorization will, in the future, save computing time.

9. Acknowledgments

This work was sponsored by the JPL-NASA project "Retrieval, Assimilation and Modeling of Atmospheric Water Vapor from Ground-and Space-Based GPS Networks: Investigation of the Global and Regional Hydrological Cycles." The authors acknowledge gratefully the stimulating discussions held with Dr. Yoaz Bar-Sever, Mr. Yong-Run Guo and Drs. Ying-Hwa (Bill) Kuo, Chris Rocken, Francois Vandenberghe, Steven Marcus, Mike Chin and Jean Dickey during the course of this work. Thanks are also due to two anonymous reviewers for their valuable comments.

REFERENCES

Anthes, R. A. and Warner, T. T. 1978. Development of hydrodynamic models suitable for air pollution and other mesometeorological studies. *Mon. Wea. Rev.* **106**, 1045–1078.

Bennett, A. F. 1992. *Inverse methods in physical oceanography*. Cambridge University Press, Port Chester, NY, 346 pp.

Bevis, M., Businger, S., Herring, T. A., Rocken, C., Anthes, R. A. and Ware, R. H. 1992. GPS Meteorology: remote sensing of atmospheric water vapor using the Global Positioning System. *J. Geophys. Res.* **97**, 15,787–15,801.

Businger, S., Chiswell, R. S., Bevis, M., Duan, J., Anthes, R. A., Rocken, C., Ware, R. H., Exner, M., VanHove, T. and Solheim, F. S. 1996. The promise of GPS atmospheric monitoring. *Bull. Am. Meteorol. Soc.* **77**, 5–18.

Cucurull, L., Navascues, B., Ruffini, G., Elosegui, P., Rius, A. and Vila, J. 2000. On the use of GPS to validate NWP systems: the HIRLAM model. *J. Atmos. Oceanic Technol.* **17**, 773–787.

Davis, J. L., Herring, T. A. Shapiro, I. I., Rogers, A. E. and Elgered, G. 1985. Geodesy by radio interferometry: effects of atmospheric modeling errors on estimates of baseline length. *Radio Sci.* **20**, 1593–1607.

Dixon, T. 1991. An introduction to the global positioning system and some tectonic applications. *Rev. Geophys.* **29**, 249–276.

Dudhia, J. 1993. A nonhydrostatic version of the Penn State-NCAR mesoscale model: validation test and simulation of an Atlantic cyclone and cold front. *Mon. Wea. Rev.* **121**,1493–1513.

Elgered, G., Davis, J. L., Herring, T. A. and Shapiro, I. I. 1991. Geodesy by radio interferometry: water vapor radiometry for estimation of the wet delay. *J. Geophys. Res.* **96**, 6541–6555.

Gelb, A. 1974. *Applied optimal estimation*, 374 pp. MIT Press, Cambridge, Mass.

Grell, G. A., Dudhia, J. and Stauffer, D. R. 1994. *A description of the 5th-generation Penn State/NCAR mesoscale model (MM5)*. NCAR Technical Note, NCAR/TN-398+STR, National Center for Atmospheric Research, Boulder, CO, 138 pp.

Ide, K., Courtier, P., Ghil, M. and Lorenc, A. C. 1997. Unified notation for data assimilation: operational, sequential and variational. *J. Meteor. Soc. Japan.* **75**, 181–189.

Janssen, M. A. 1993. *Atmospheric remote sensing by microwave radiometry*. John Wiley, New York, 592 pp.

Kållberg, P., Ed., 1989. *The HIRLAM level 1 system. Documentation manual*, 160 pp. [Available from SMHI, S 60176 Norrköping, Sweden.]

Kuo, Y.-H., Zou, X. and Guo, Y.-R. 1996. Variational assimilation of precipitable water using a nonhydrostatic mesoscale adjoint model. Part I: Moisture retrieval and sensitivity experiments. *Mon. Wea. Rev.* **124**, 122–147.

Leick, A. 1990. *GPS satellite surveying*. John Wiley, New York, 352 pp.

Liu, D. C. and Nocedal, J. 1989. On the limited memory BFGS method for large scale optimization. *Math. Program.* **45**, 503–528.

Lynch, P. and Huang, X.-Y. 1992. Initialization of the HIRLAM model using a digital filter. *Mon. Wea. Rev.* **120**, 1019–1034.

Navon, I. M., Zou, X., Derber, J. and Sela, J. 1992. Variational data assimilation with an adiabatic version of the NMC spectral model. *Mon. Wea. Rev.* **120**, 1433–1446.

Saastamoinen, J. 1972. Atmospheric correction for the troposphere and stratosphere in radio ranging of satellites. The use of artificial satellites for geodesy, *Geophys. Monogr., Ser.* **15**, AGU, 247 pp.

Talagrand, O. and Courtier, P. 1987. Variational assimilation of meteorological observations with the adjoint vorticity equation (I). *Quart. J. R. Met. Soc.* **113**, 1311–1328.

Ware, R. and Businger, S. 1995. Global positioning for geosciences research. *Eos. Trans. Amer. Geophys. Union* **76**, 187.

Ware, R. et al. 1996. GPS sounding of the atmosphere from low Earth orbit: preliminary results. *Bull. Amer. Meteorol. Soc.* **77**, 19–40.

Xiao, Q., Zou, X. and Kuo, Y.-H. 1999. Incorporating the SSM/I derived precipitable water vapor and rain rate into a numerical model: a case study for ERICA IOP-4 cyclone. *Mon. Wea. Rev.* **128**, 87–108.

Zou, X., Kuo, Y.-H. and Guo, Y.-R. 1995. Assimilation of atmospheric refractivity using a nonhydrostatic adjoint model. *Mon. Wea. Rev.* **123**, 2229–2249.

Zou, X., Vandenberghe, F., Pondeca, M. and Kuo, Y.-H. 1997. *Introduction to adjoint techniques and the MM5 adjoint modeling system.* NCAR Technical Note, NCAR/TN-435-STR, 110 pp.