Advanced Algorithm and System Development for Cassini Radio Science Tropospheric Calibration

S. J. Keihm and K. A. Marsh
Microwave, Lidar, and Interferometer Technology Section

An important error source for the calibration of tropospheric delay variations at microwave frequencies is the "retrieval error" that is due to the uncertainty in conversion from observables (sky brightness temperatures and surface measurements) to path delay. A large database of vertical water-vapor density profiles (from lidar measurements) and vertical temperature profiles (from a radio acoustic sounding system (RASS)) has been used to quantify the expected retrieval error as a function of time scale, instrumental configuration, and measurement accuracy. One key result is that a three-channel (e.g., 22.2-, 23.8-, and 31.4-GHz) water vapor radiometer (WVR) and a two-channel microwave temperature profiler (MTP) are nearly optimal and are needed in order to meet the accuracy goals for the Cassini Gravitational Wave Search Experiment (GWE). A second key result is that model-based Bayesian inversion techniques provide substantially better accuracy than do statistical retrieval methods. The estimated retrieval error for the measurement of path delay variations (using Bayesian methods and the above instrumental configuration with targeted stability) is a factor >6 smaller than the total GWE tropospheric calibration requirement at time scales >5000 s. At a 1000-s time scale, the error is a factor of ∼2 smaller than the GWE requirement. However, the retrieval error is approximately equal to the GWE requirement at a time scale of 100 s.

I. Introduction

The planned Cassini Gravitational Wave Search Experiment (GWE) requires very accurate calibration of fluctuations in the line-of-sight, water vapor–induced microwave path delay over ∼10-hour time intervals. Significant components of the error budget for such a tropospheric calibration measurement include the effects of dry (nonvapor) fluctuations, radiometer calibration and stability, pointing, beam matching and offset between the tropospheric calibration (tropo-cal) radiometer and the DSN antenna, vapor absorption model accuracy, and retrieval algorithm accuracy. This article describes developments in the area of retrieval algorithm techniques and the radiometer calibration and stability requirements for meeting the Cassini tropo-cal specifications.

As used here, "retrieval error" is defined as the error in estimated wet path delay that results solely from the nonunique conversion of observable measurements (primarily microwave brightness temperatures) to the delay value, i.e., when all other error sources, including modeling and instrument errors, are assumed

to be zero. In the results that will follow, however, retrieval algorithms are tested using plausible ranges of (nonzero) observable noise. Thus, the resultant retrieval errors will reflect the combined effects of algorithm error and instrument stability.

For application to the Cassini GWE tropo-cal effort, the distinction between retrieval accuracy and precision is crucial. In the sidereal Doppler tracking of the Cassini spacecraft, the *variations* in atmospheric delay over minute-to-multihour time scales are the critical quantities to be measured. Uncalibrated constant offsets and linear trends will not compromise the success of the GWE. Thus, the GWE tropo-cal requirements and goals are expressed in terms of the Allan standard deviation (ASD), which provides a measure of the nonlinear delay variations as a function of time interval Δt [1]:

$$ASD(\Delta t) = \frac{\left\{\langle[\delta(t+2\Delta t) - 2\delta(t+\Delta t) + \delta(t)]^2\rangle\right\}^{1/2}}{\Delta t\sqrt{2}} \tag{1}$$

in s/s, where $\langle\rangle$ denotes ensemble averaging and $\delta(t)$ is the delay in seconds at time t.

The emphasis on measuring delay variations, rather than absolute delay, requires a modification to the normal method (e.g., [2]) of evaluating retrieval algorithms. For example, the predicted performance of statistical retrieval algorithms is usually given as an rms residual from a zero-bias, multilinear regression of observables versus path delay, including the effects of observable noise. The regression is (usually) derived from a multiyear radiosonde archive of computed path delays and observables. When the observable noise is negligible, the residual for path delay retrieval reflects the nonuniqueness of the inversion solution. Variations in temperature conditions and the height distribution of the vapor contribute to the retrieval scatter. When, however, the derived algorithm is applied to subsets of the observable archive (stratified, for example, by temperature conditions), the errors appear more as offsets, with reduced scatter due to the restricted range of conditions. We expect that retrieval algorithms applied to short time-interval conditions would also exhibit this character, with retrieval error scatter determined by the variability in the temperature and vapor-height distribution conditions over the time interval of interest.

The major difficulty in evaluating the expected reduction in retrieval error scatter over time intervals from 100 to 10,000 s (the range of interest for the Cassini GWE) has been the absence of reliable data characterizing profile changes in vapor and temperature over intervals of minutes to hours. Radiosondes are routinely obtained every 12 hours and, thus, cannot provide a useful testbed for evaluating the variations in retrieval errors over Cassini GWE time scales. To remedy this problem, we have obtained from Yong Han at the Environmental Testing Laboratory of the National Oceanic and Atmospheric Administration (NOAA/ETL) in Boulder, Colorado, a data set that provides 2-minute time series of both RASS-derived temperature and lidar-derived water vapor profiles over ∼10-hour nighttime intervals from Coffeeville, Kansas. The lidar data appear to be the best yet published, exhibiting excellent agreement with colocated radiosonde data [3]. The profile data are from November and December 1991, with clear conditions and path delay levels close to the expected range for Goldstone in winter. From the 2-minute Coffeeville profile data, time series of path delay, radiometric brightness temperatures, and surface meteorology observables have been computed and used as a testbed for assessing the accuracy and precision of candidate retrieval algorithms.

In Section II of this article, the testbed of lidar plus RASS profile data will be described. In Section III, the candidate algorithms considered to date will be described, including both statistical and Bayesian techniques. In Section IV, we outline the simulation procedures by which each algorithm was evaluated using the testbed data. Sections V and VI contain first-stage and second-stage results in terms of the optimum retrieval method, required instrumentation, frequency selection, and radiometer calibration requirements. In Section VII, the ASD performance of a proposed prototype Cassini tropo-cal system is evaluated as a function of radiometer stability and compared with the GWE requirements. Section VIII contains conclusions and plans for future work.

II. Testbed of Profile Time Series Data

The Coffeeville data received from NOAA/ETL consisted of 11 nights of temperature and vapor-density profile data at 75-m vertical spacing and 2-min time intervals. The temperature profile data from 0.35 km to ∼2 km were derived from RASS measurements taken at 15-min intervals with a vertical resolution of 150 m. Temperatures below the lowest RASS range gate (0.35 km) were obtained by interpolation using surface meteorology data. Linear interpolation was used to produce the final 2-min profiles with 75-m vertical spacing. Above 2 km, research quality radiosonde data, obtained at 3- to 12-hour intervals, were used to complete the temperature profiles (with considerable interpolation smoothing in the time domain).

The vapor density profile data were provided by the Raman lidar measurements, obtained at 1-min intervals (and averaged to produce the 2-min data) with 75-m vertical resolution. Absolute calibration was obtained by a single least-squares regression to the vapor profile data obtained from 41 coincident radiosonde measurements over a 3-week period. The estimated random error of the lidar vapor density measurements (due to instrument noise) ranged from approximately 0.3 to 1.0 percent for the lowest 2 km to approximately 10 percent at ∼8 km. Near the 8-km range limit, the lidar-derived vapor densities often exceeded saturation for the corresponding temperature. For these data, and all heights above 8 km, we have applied a correction, assigning all vapor densities to be equivalent to a 20-percent relative humidity condition. For heights below the lowest lidar range gate (185 m), vapor densities were obtained by interpolation with the surface meteorology data.

From the 11 nighttime data sets, we have selected 4 thus far for algorithm simulation studies. This selection was based primarily on data completeness and the time variability of the profile data. (A few of the data sets contained significant time gaps, and one night's data had no time variation in the temperature data.) The four selected November 1991 data sets have been designated days 12, 14, 15, and 16, corresponding to the truncated file names transmitted from NOAA/ETL (not the exact days of the month). The path delay and path-delay fluctuation levels at ∼1-hour time scales for days 14, 15, and 16 are comparable to typical Goldstone winter conditions. The day-12 path delay and fluctuation levels are comparable to the highest levels observed for Goldstone winter conditions [4]. Temperature and vapor-density time series, at selected heights, are shown for two of the data sets in Figs. 1 and 2. The striking features are the large fractional excursions in vapor density on hour time scales or less, evidence of the degree and scale of tropospheric turbulence.

Figure 3 shows the zenith path-delay time series computed from the days-12 and -15 profile data, along with the calculated ASD fluctuation statistics. The day-15 path delay is typical of a Goldstone winter night, while the day-12 data are representative of the wettest winter delay conditions expected at Goldstone. The range of ASD values, typical of active winter Goldstone conditions [4], are approximately 5 to 50 times larger than the Cassini requirements and provide an excellent testbed for evaluating candidate tropo-cal instrumentation and retrieval algorithms.

III. Description of Retrieval Algorithms

For the purposes of this article, the term "retrieval algorithm" refers to a specific inversion technique applied to a selected set of instrument observables. The two inversion techniques considered are the linear statistical method and the model-based Bayesian method. The candidate instrumentation includes 2-, 3-, and 11-channel water vapor radiometers (WVRs); a 2-channel, elevation-scanning microwave temperature profiler (MTP); and surface measurements of temperature, pressure, and vapor density. As used here, "case number" refers to a selected subset of the candidate instrumentation and a specific assignment of measurement noise values. Also as used here, "noise" refers to the random error of a single measurement. Perhaps a better term would be stability since we are using noise as a measure of an instrument's precision over all time scales of the retrieval interval. Unless specified otherwise, all assigned instrument errors are assumed to be Gaussian-distributed white noise.

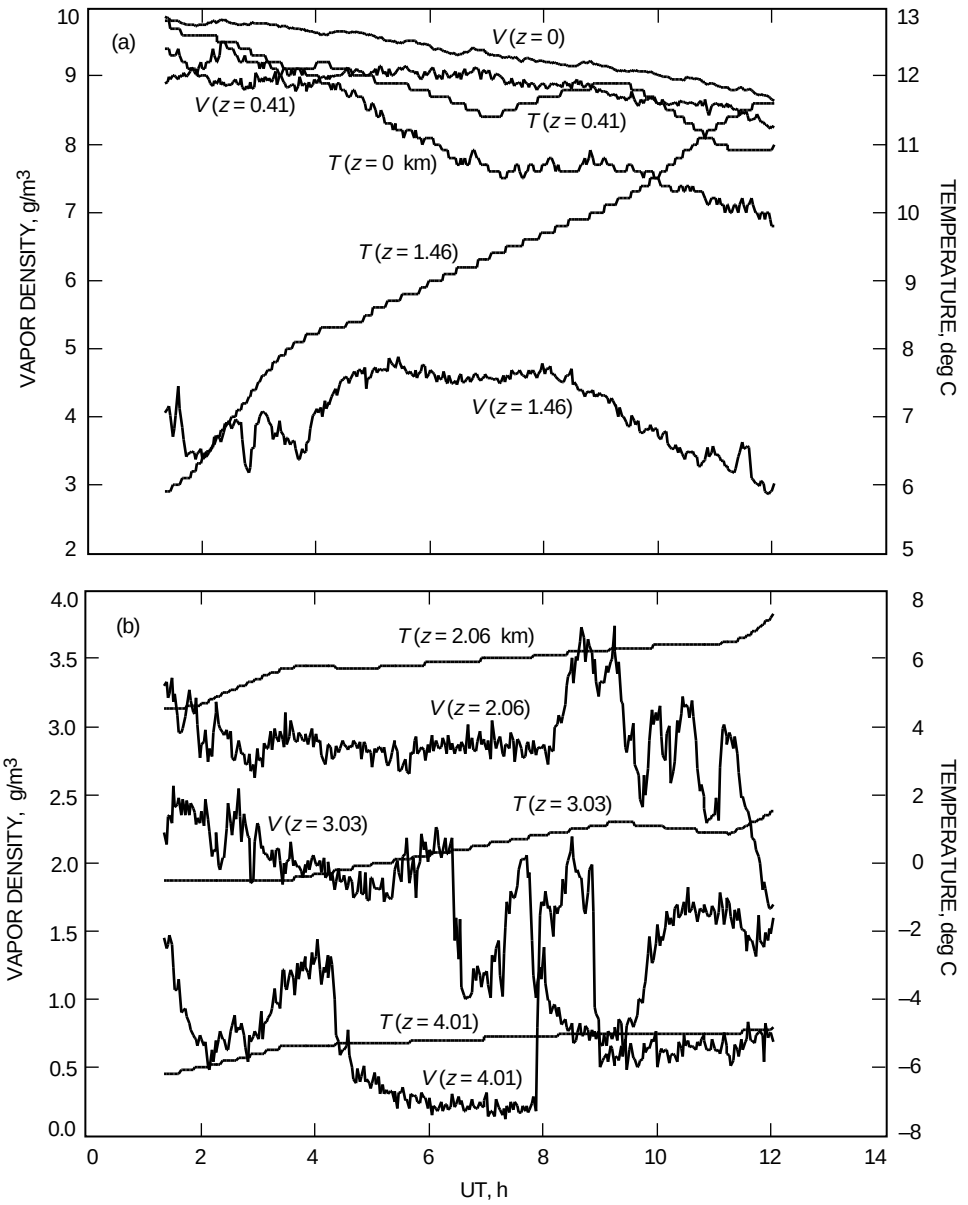

Fig. 1. Time series of temperature and vapor density from the lidar and RASS data obtained for day 12 at Coffeeville at heights (a) < 2 km and (b) > 2 km.

A. Statistical Algorithms

All statistical algorithms were generated using a 3-year (1990–1992) radiosonde archive from Topeka, Kansas. Topeka is located ∼160 miles due north of Coffeeville and has nearly the same elevation (340 m at Topeka; 290 m at Coffeeville) and similar humidity conditions. The statistical algorithms assume that the desired parameter (in our case, zenith wet path delay) can be expressed as a linear combination of the instrument observables:

$$PD = PD_{\text{mean}} + \Sigma\, C_i(OB_i - OBMEAN_i) \tag{2}$$

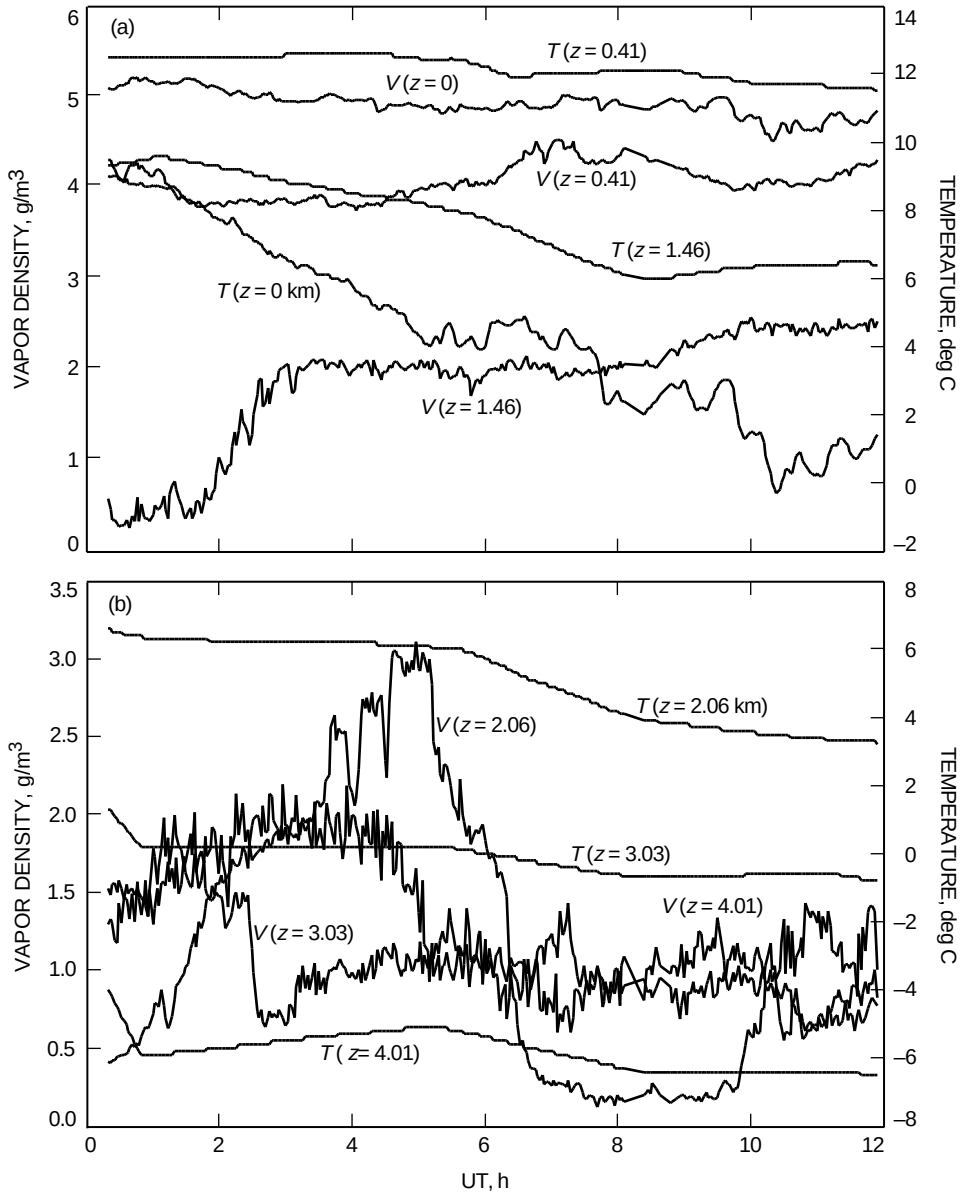

Fig. 2. Time series of temperature and vapor density from the lidar and RASS data obtained for day 15 at Coffeeville at heights (a) < 2 km and (b) > 2 km.

where PD_{mean} is the archive average path delay, OB_i is the measured value of the ith observable (computed from the profile data), $OBMEAN_i$ is the computed archive average for the ith observable, and the C_i are the linear retrieval coefficients generated by regression of the path delay versus observables-plus-noise archive computed from the radiosonde database. The generated retrieval coefficients depend on the assigned observable noise values because the error vector is added to the observables before regressing with path delay. Thus, for low-noise cases, the coefficients approach values that would be produced by a simple multilinear regression of computed path delays versus computed observables, and the predicted retrieval error approaches the fit's rms residual. For high-noise cases, the coefficients approach zero as the observable information content diminishes, retrieved path delay values remain close to the archive average, and the predicted retrieval error approaches the archive standard deviation in path delay.

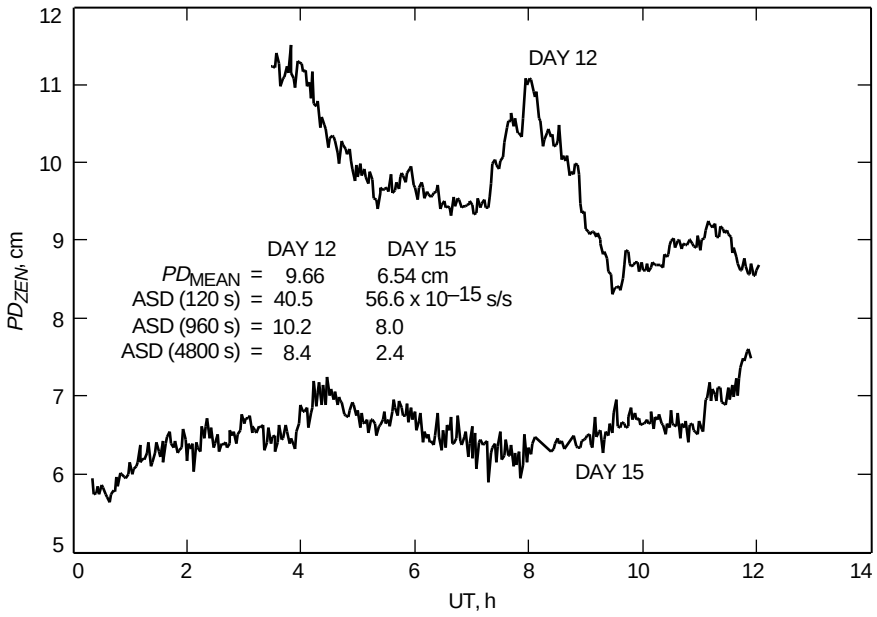

Fig. 3. Zenith wet path delay time series computed from the lidar and RASS measurements at Coffeeville for the day-12 and day-15 intervals.

B. Bayesian Algorithms

The Bayesian method employs model-based inversion techniques to solve the radiative transfer equation for discrete vapor and temperature profiles given the array of observable measurements. The profiles are varied iteratively, with theoretical observables calculated each step, until a most-probable profile solution is obtained. From this most-probable profile solution, the final retrieved path delay is computed directly. The criteria for a most-probable solution include constraints based on a priori covariance statistics of the profiles as well as the residuals between measured and computed observables. The covariance statistics are derived from the same Topeka radiosonde archive used for the statistical algorithms and provide a means for downweighting a profile solution that varies too greatly from the archive mean. In qualitative terms, the most-probable profile solution is that which minimizes the residuals between measured and computed observables while best conforming to the constraints of the a priori statistics.

Observable errors are taken into account by adjusting the relative weights assigned to the residual and statistical constraints. Similarly to the statistical algorithms, higher observable errors lead to increased weight for the a priori archive and less weight for the measurements.

IV. Retrieval Simulation Procedures

For each of the four selected lidar-plus-RASS nighttime profile data sets, 2-min time series of path delay and all candidate observables were computed by numerical solution of the radiative transfer equation using our current nominal model of atmospheric absorption and refractivity. All candidate observables used in the first two stages of testing are listed in Table 1 along with the range of associated Gaussian errors that have been considered.

The range of WVR brightness temperature errors spans worst-to-best cases in the sense that poorly monitored gain variations can easily result in 0.5-K drifts over multihour periods, while 0.01 K represents the multihour stability goal of the advanced WVR. Current WVRs, with careful gain drift monitoring, have demonstrated stability at the 0.1-K level. The stability of current generation ground-based MTPs is less certain, but the 0.1- to 1.0-K range probably brackets the actual performance. The

Table 1. Candidate observables and associated errors.

Observable no.	Description	Frequency, GHz	Elevation angle, deg	Error range, K
1	WVR TB[a]	20.0	90	0.01–0.50
2	WVR TB	20.4	90	0.01–0.50
3	WVR TB	20.7	90	0.01–0.50
4	WVR TB	21.2	90	0.01–0.50
5	WVR TB	21.7	90	0.01–0.50
6	WVR TB	22.235	90	0.01–0.50
7	WVR TB	22.7	90	0.01–0.50
8	WVR TB	23.2	90	0.01–0.50
9	WVR TB	23.6	90	0.01–0.50
10	WVR TB	24.0	90	0.01–0.50
11	WVR TB	31.4	90	0.01–0.50
12	MTP TB	54.4	90	0.10–1.00
13	MTP TB	54.4	42	0.10–1.00
14	MTP TB	54.4	30	0.10–1.00
15	MTP TB	57.97	90	0.10–1.00
16	MTP TB	57.97	42	0.10–1.00
17	MTP TB	57.97	30	0.10–1.00
18	Surface temperature	—	—	0.2 K
19	Surface pressure	—	—	1.0 mb
20	Surface vapor density	—	—	0.5 g/m^3

[a] TB = brightness temperature.

single-point estimates of surface meteorology observable precision are values that are commonly cited and that we assume are attainable with carefully calibrated off-the-shelf technology.

Using the computed "truth" time series of observables as algorithm inputs, retrieval simulations were performed for individual cases by selecting an observable subset, adding Gaussian noise of the specified level to each 2-minute observable array, and then retrieving the time series of zenith path delay over the 10- to 12-hour interval of the selected lidar-based data set. The time series of retrieval error (retrieved-minus-truth path delay) was then computed, and measures of performance, including bias, scatter, and error ASD, were calculated. Evaluations were then made comparing algorithm techniques (Bayesian versus statistical) and different observable subsets, and assessing the effects of measurement stability. When comparing the retrieval-error ASD results with the Cassini requirements and goals, it is important to note that the first-stage retrieval simulations were performed for zenith path delay only, an overly optimistic comparison in that the retrieval errors will scale with air mass and the actual Cassini tracking will be very nearly sidereal. This issue is addressed in Section VI.

V. First-Stage Results

The initial selection of test cases was intended to provide insights on the following main issues:

(1) When, if at all, are Bayesian techniques clearly superior to statistical ones for path delay retrieval?

(2) What improvements can be obtained by adding channels to the standard three-frequency WVR?

(3) How much improvement is afforded by adding MTP and surface meteorology observables?

In the first stage of simulations, 24 cases were selected to investigate the above issues. These cases are listed in Table 2 in terms of observable selection and assigned errors.

Table 2. First-stage simulation case description.

Case no.	Observable	Error				
		WVR, K	MTP, K	T_s, K	P_s, mb	V_s, g/m^3
1	3-channel WVR	0.50	—	—	—	—
2	3-channel WVR	0.10	—	—	—	—
3	3-channel WVR	0.03	—	—	—	—
4	3-channel WVR	0.01	—	—	—	—
5	11-channel WVR	0.50	—	—	—	—
6	11-channel WVR	0.10	—	—	—	—
7	11-channel WVR	0.03	—	—	—	—
8	11-channel WVR	0.01	—	—	—	—
9	3-channel WVR $+T_s$	0.01	—	0.2	—	—
10	3-channel WVR $+P_s$	0.01	—	—	1.0	—
11	3-channel WVR $+V_s$	0.01	—	—	—	0.5
12	3-channel WVR $+T_s + P_s + V_s$	0.01	—	0.2	1.0	0.5
13	11-channel WVR $+T_s$	0.01	—	0.2	—	—
14	11-channel WVR $+P_s$	0.01	—	—	1.0	—
15	11-channel WVR $+V_s$	0.01	—	—	—	0.5
16	11-channel WVR $+T_s + P_s + V_s$	0.01	—	0.2	1.0	0.5
17	3-channel WVR $+MTP + T_s$	0.01	1.0	0.2	—	—
18	3-channel WVR $+MTP + T_s$	0.01	0.5	0.2	—	—
19	3-channel WVR $+MTP + T_s$	0.01	0.1	0.2	—	—
20	3-channel WVR $+MTP + T_s + P_s + V_s$	0.01	0.1	0.2	1.0	0.5
21	11-channel WVR + MTP $+T_s$	0.01	1.0	0.2	—	—
22	11-channel WVR + MTP $+T_s$	0.01	0.5	0.2	—	—
23	11-channel WVR + MTP $+T_s$	0.01	0.1	0.2	—	—
24	11-channel WVR + MTP $+T_s + P_s + V_s$	0.01	0.1	0.2	1.0	0.5

Note that "3-channel WVR" refers to the standard JPL WVR J-unit operating frequencies of 20.7, 22.2, and 31.4 GHz, while "11-channel WVR" refers to the full ensemble of WVR frequencies listed in Table 1. For all cases that include MTP measurements, a two-channel (54.4- and 57.97-GHz) elevation-scanning system is simulated with sky brightness temperature measurements at 30-, 42-, and 90-deg elevations. In an earlier unpublished study, this frequency and elevation combination has been determined to be near optimum for temperature profile retrievals for heights up to 5 to 7 km.

Retrievals for observable set case numbers 1 through 24 were simulated for the testbed data sets designated days 12, 14, 15, and 16 using both the Bayesian and statistical methods. Results are shown in Figs. 4 through 7 in terms of retrieval error bias and ASD. The most significant result in terms of bias

(Fig. 4) is that Bayesian methods applied to advanced WVR-plus-MTP observable sets (cases 17 through 24) always reduce the retrieved path delay bias to a level of ∼0.01 cm or less, even for MTP errors as large as 1.0 K. For cases without MTP observables (1 through 16), the biases have little or no correlation with WVR errors. This result suggests that the nonlinear model-based algorithms, when temperature profiling observables are included, effectively eliminate offsets due to pervasive conditions that depart from the statistical archive average.

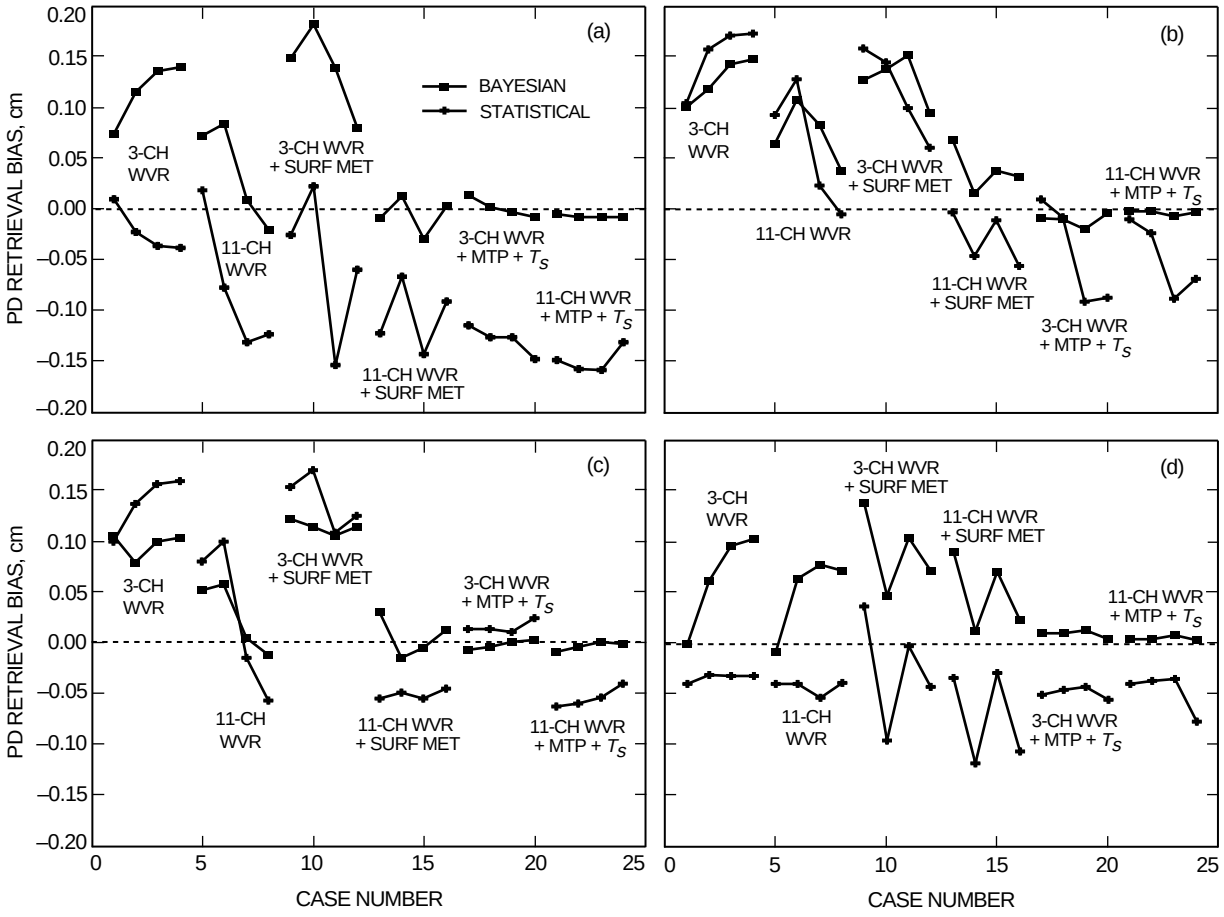

Fig. 4. A comparison of Bayesian and statistical path delay retrievals in terms of bias for the Coffeeville testbed data. The case numbers refer to specific selections of observables and observable errors (see Table 2 and the text for a more detailed description): (a) day 12, (b) day 14, (c) day 15, and (d) day 16.

Figures 5 through 7 show the retrieval error ASD for three time scales that nearly span the GWE intervals of interest. Also indicated on the plots are the Cassini ASD requirements (CR) and goals (CG)[1] and annotated information describing the observables and observable errors for each case number. (For further clarification, refer to Table 2.) Note that the Cassini specifications shown in Figs. 5 through 7 have been scaled downward by a $1/\sqrt{2}$ factor to account for the two-way tracking required by Cassini. (For two-way tracking, the ASD of the retrieval error increases by a factor of approximately $\sqrt{2}$, assuming the uplink and downlink retrieval errors are uncorrelated. This will be the case as long as the Cassini signal round-trip travel time is greater than the ASD time scale of interest.)

[1] *Cassini Radio Science Ground System D-Level and Cost Review* (internal document), Jet Propulsion Laboratory, Pasadena, California, February 27, 1995.

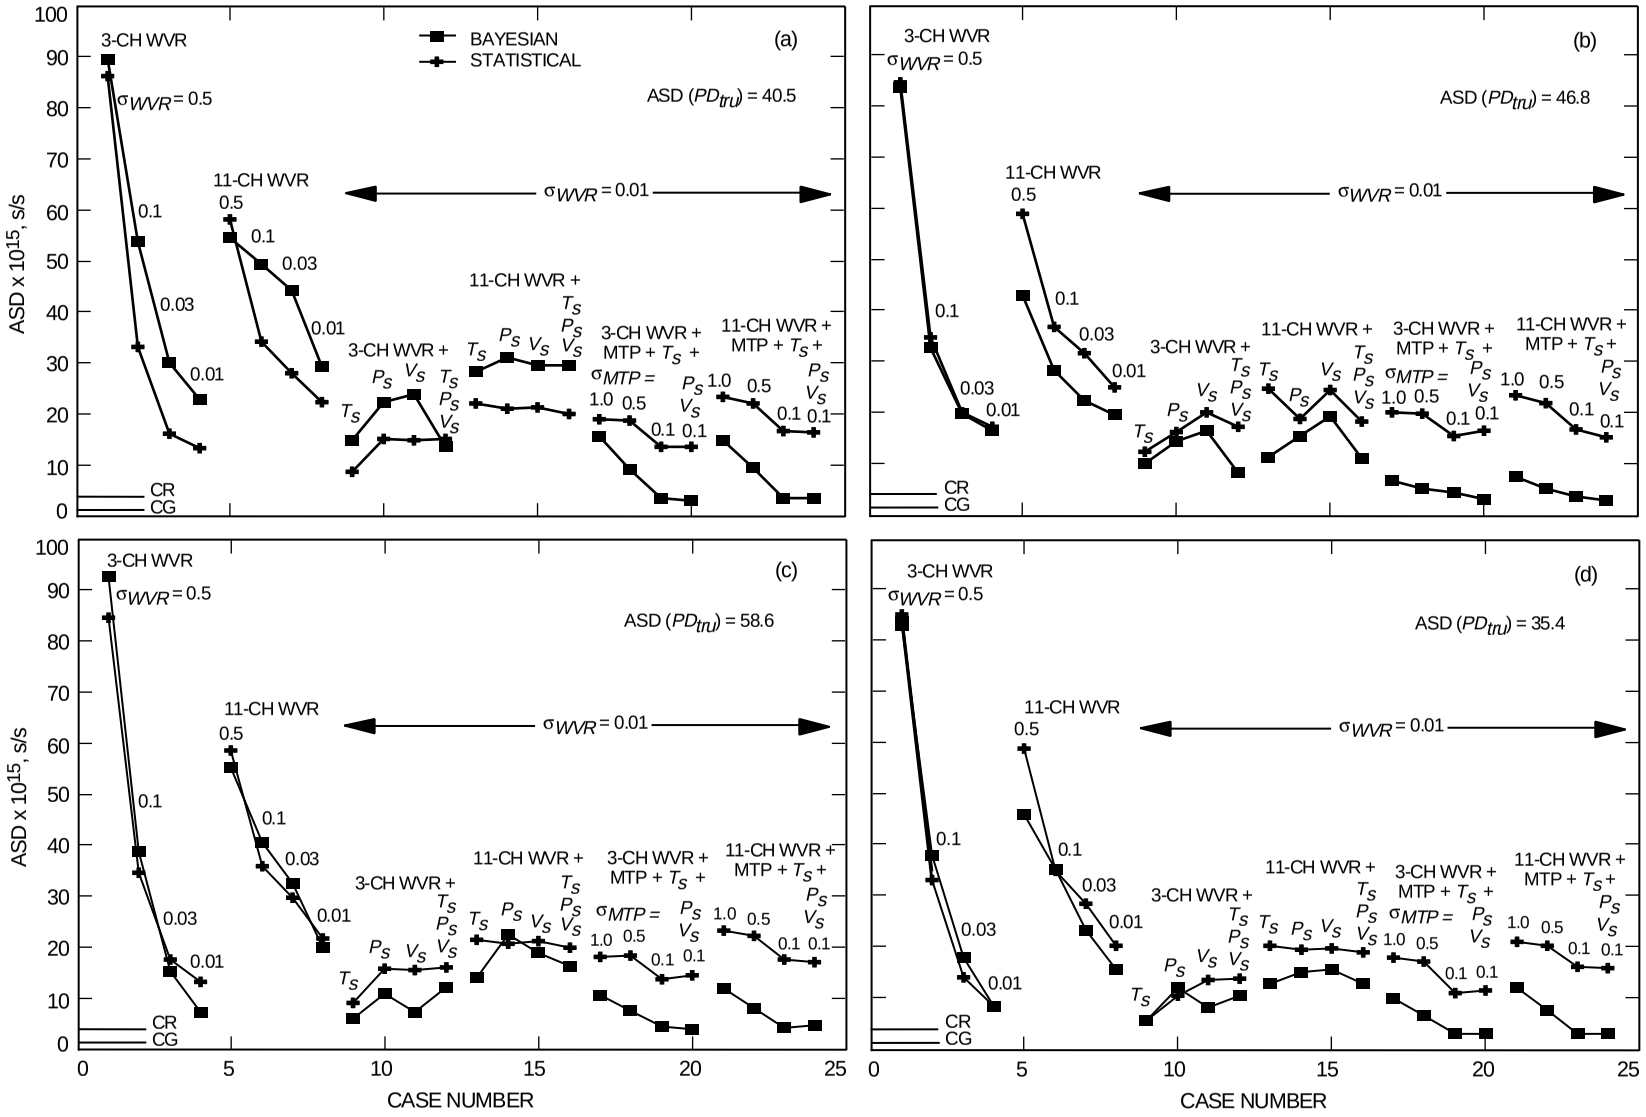

Fig. 5. A comparison of Bayesian and statistical path delay retrievals in terms of ASD at a 120-s time interval for the Coffeeville testbed data. The case numbers refer to specific selections of observables and observable errors (see Table 2 and the text for a more detailed description): (a) day 12, (b) day 14, (c) day 15, and (d) day 16.

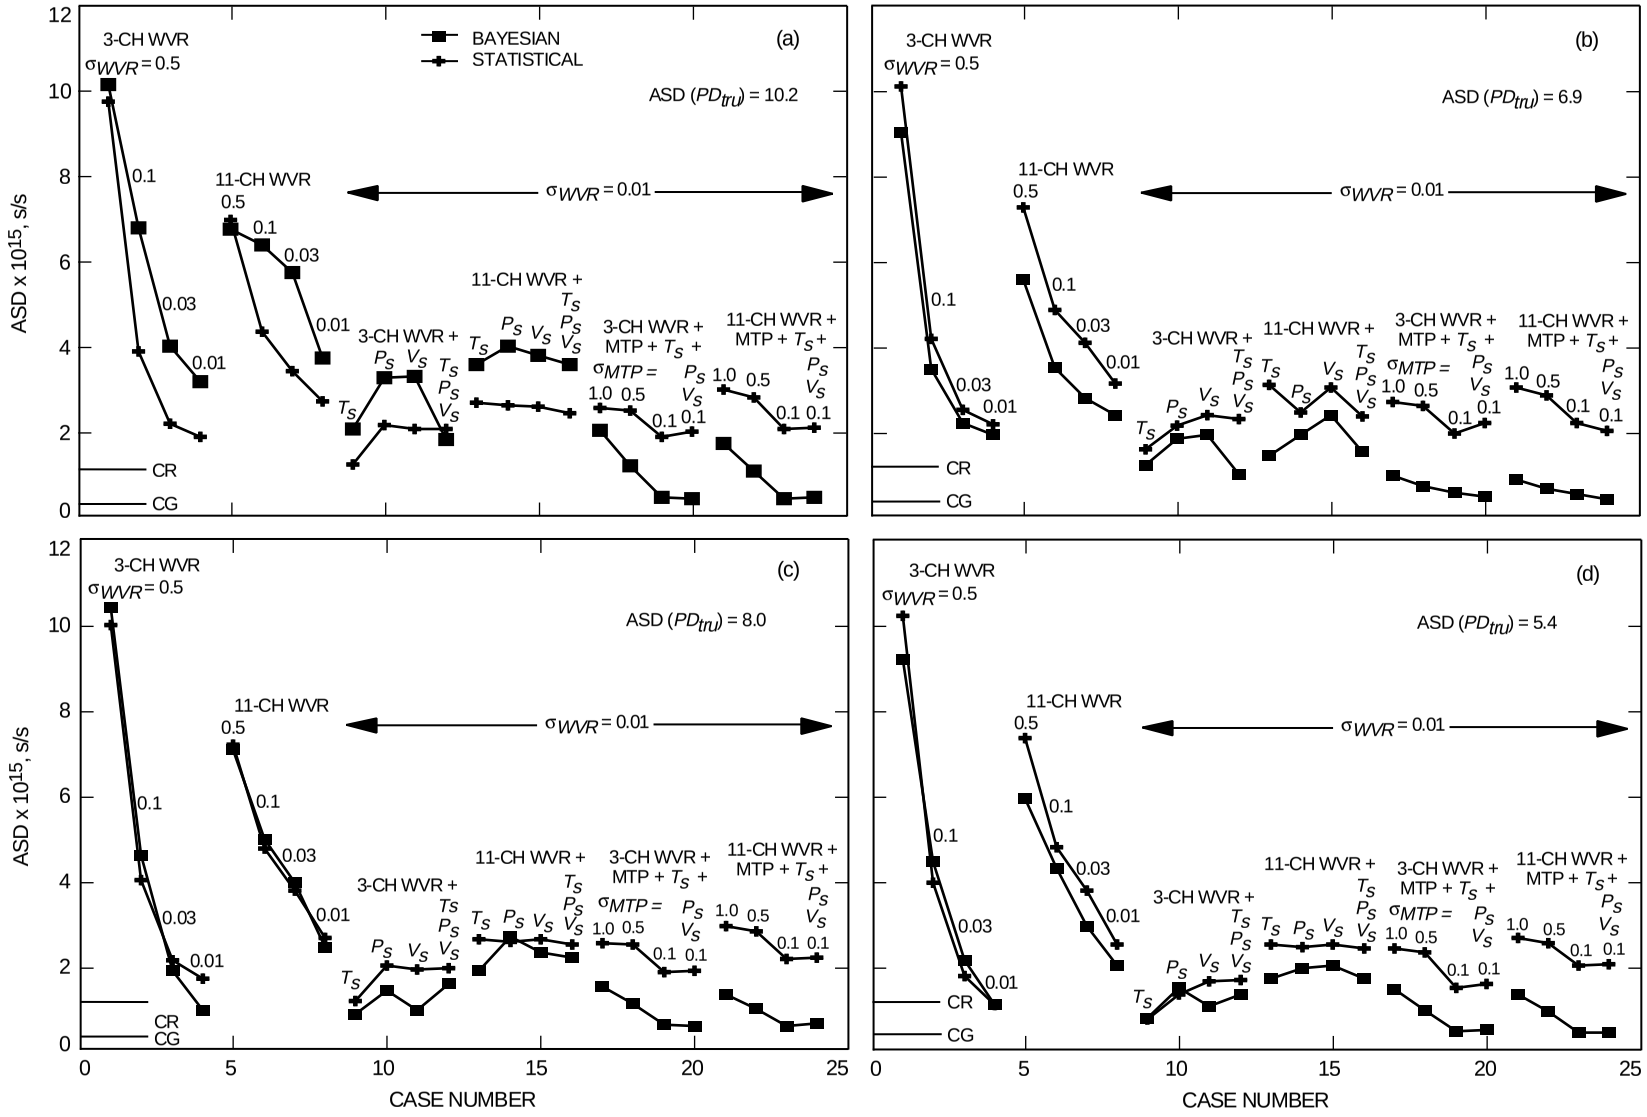

Fig. 6. A comparison of Bayesian and statistical path delay retrievals in terms of ASD at a 960-s time interval for the Coffeeville testbed data. The case numbers refer to specific selections of observables and observable errors (see Table 2 and the text for a more detailed description): (a) day 12, (b) day 14, (c) day 15, and (d) day 16.

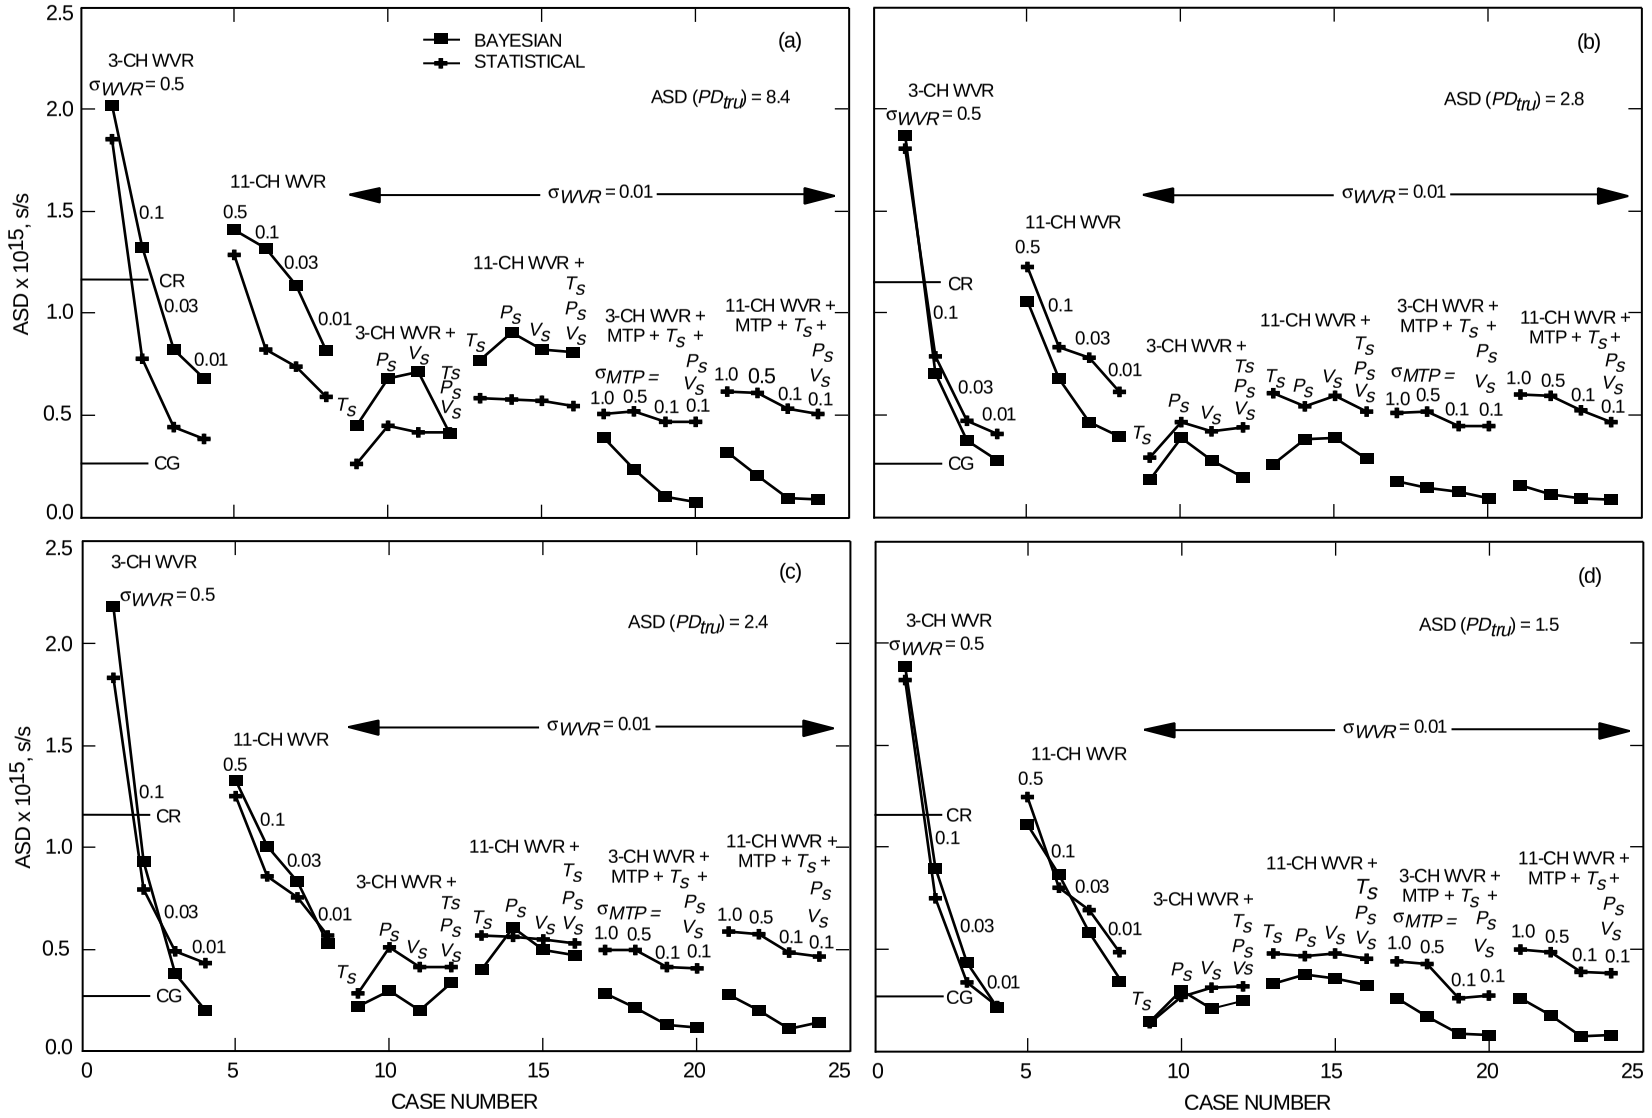

Fig. 7. A comparison of Bayesian and statistical path delay retrievals in terms of ASD at a 4800-s time interval for the Coffeeville testbed data. The case numbers refer to specific selections of observables and observable errors (see Table 2 and the text for a more detailed description): (a) day 12, (b) day 14, (c) day 15, and (d) day 16.

Based on the results shown in Figs. 5 through 7, and addressing the three issues outlined above, we draw the following conclusions:

(1) In terms of ASD, Bayesian retrievals are clearly superior to statistical methods when temperature profile information (MTP) is available (cases 17 through 24). At the shorter time scales (Δt = 120 s, Fig. 5), the Cassini requirements can only be approached by utilizing Bayesian techniques with precise WVR and MTP observables. At the longer time scales (Figs. 6 and 7), the Cassini goals can be met or exceeded using the Bayesian techniques with the precise WVR-plus-MTP observable set. When only WVR or WVR-plus-surface meteorology observables are available (cases 1 through 16), the results are mixed. For the testbed data sets of days 14, 15, and 16, the Bayesian results are generally as good or better than the statistical. For day 12, the statistical retrievals are superior for non-MTP cases.

(2) The addition of channels in the 20- to 24-GHz band to the standard three-channel WVR produces no significant retrieval improvement in the ASD domain, with or without MTP and surface meteorology observables, when the WVR errors are less than 0.1 K. In fact, for WVR precision approaching the goal of the advanced WVR (0.01 to 0.03 K), the ASD performance is actually significantly worse for the 11-channel system relative to the standard 3-channel WVR when no MTP observables are used. (Compare cases 5 through 8 with 1 through 4 and cases 13 through 16 with 9 through 12.) This behavior occurs for both the Bayesian and statistical retrievals at all time scales, and we have no definitive explanation for it at this time. One contributing factor appears to be the effect of increasing system noise with the addition of WVR channels. Theoretically, this should not occur. If the covariance matrices are correctly calculated, the addition of redundant or low information observables should not increase the retrieval noise. However, examination of the statistical algorithms' retrieval coefficients reveals the opposite to be occurring. From Eq. (2), it is straightforward to derive the retrieval noise that results from independent observable noise:

$$\Delta PD = \left[\Sigma(C_i[\Delta OB_i])^2\right]^{1/2} \tag{3}$$

where ΔPD is the retrieval noise and ΔOB_i are the observable errors. Because the observable errors are considered random, the ΔPD values are easily converted into ASD values using Eq. (1). The result is that the system noise is the dominant component of the statistical algorithm ASD values shown in Figs. 5 through 7 and is responsible for the degradation in performance going from the 3-channel to the 11-channel WVR observable sets. This result suggests that, for low WVR observable noise (<0.1 K) and strongly information-redundant observables, the correlations of observables with path delay are not being accurately calculated, due either to the inadequacy of the Topeka data archive for resolving small differences in the effects of path delay on observables or errors that are amplified in the inversion of a near-singular regression matrix.

We have also investigated the performance of a two-channel WVR for the Cassini tropocal system to evaluate trade-offs between performance and the design simplification and cost reductions afforded by a two-channel (versus three-channel) radiometer. The issue was addressed using WVR-plus-MTP-plus-T_s observable sets including two-channel (20.7-, 31.4-GHz and 22.2-, 31.4-GHz) and three-channel (20.7-, 22.2-, and 31.4-GHz) WVRs with wide ranges of WVR and MTP errors. The simulations were performed using only the Bayesian algorithm on testbed days 12 and 15, which had the largest and most variable path delay. The results indicated that for all tested combinations of WVR and MTP errors the three-channel WVR system produces better ASD performance. The

fractional improvement generally increases as both the WVR and MTP errors decrease. For the low range of errors considered (<0.03 K for the WVR and <0.3 K for the MTP), the three-channel reduction in the retrieval error ASD (relative to the best two-channel result) ranges approximately from 20 to 50 percent for the three time scales of 120, 960, and 4800 seconds. Based on these results, we concluded that a three-channel WVR is the optimum choice for the Cassini tropo-cal system.

(3) With advanced WVR precision (0.01 K) and a Bayesian algorithm, the addition of MTP (and the surface temperature) observables produces dramatic improvement in the ASD domain of retrieval performance (cases 17 through 24). The performance depends strongly on the MTP observable errors, improving by a factor of approximately 2 to 4 times as the MTP errors are reduced from 1.0 to 0.1 K. For the best-case MTP observables, with 0.1-K precision (cases 19, 20, 23, and 24), Cassini requirements for near-zenith tracking are met or exceeded at all time scales. These results appear to be insensitive to the variation in fluctuation levels seen in the four testbed data sets.

The addition of surface pressure and surface vapor density as observables does not appear to significantly improve performance (e.g., compare cases 19 and 20). However, considering the relatively low cost of deployment and other uses, such as monitoring of the dry delay component and calibration checks on the microwave instrumentation, a surface meteorology package, including temperature, pressure, and humidity sensors, should be included in the Cassini tropo-cal system design.

VI. Second-Stage Results

Analysis of retrieval simulations for the initial 24 candidate sets of observables and observable errors led to three main conclusions: (1) The advanced WVR for Cassini tropo-cal needs no additional channels beyond the standard three-channel WVR system; (2) inclusion of an MTP or equivalent temperature profiling system is required to attain optimal tropospheric calibration; and (3) Bayesian retrieval techniques are required to maximize performance for a tropo-cal system that includes advanced WVR precision and MTP observables.

In the second stage of simulations, three remaining issues were addressed using the lidar testbed data set: (1) optimization of the three-frequency WVR channel selection, (2) the effects of sidereal line-of-sight tracking, and (3) the effects of WVR calibration offsets.

A. Optimization of the Three-Frequency WVR Channel Selection

We have assumed to this point that the frequencies of the three-channel WVR include the vapor resonance line center (22.235 GHz), the lower-frequency "hinge" point (20.7 GHz, see below), and a cloud-sensitive channel (31.4 GHz) identical to the J-series WVRs designed at JPL. From previous studies (e.g., [2]), these frequencies have been determined to be near-optimum for WVR-only path-delay retrieval accuracy in both clear and cloudy weather. For the Cassini GWE, which requires unprecedented precision in tropospheric calibration, and is expected to primarily utilize cloud-free data, the 31.4-GHz channel may not be optimum, and we decided to investigate possible improvements afforded by replacing the 31.4-GHz channel with a lower frequency near the vapor line.

To address the frequency selection issue, we ran simulations for numerous three-frequency combinations with the following results. Line center (22.235 GHz) and a frequency within ∼0.5 GHz of the absorption line hinge points are required to achieve optimal retrieval performance. The line center provides maximum signal-to-noise ratio in terms of water vapor sensitivity. The hinge points, near 20.7 and 23.8 GHz, are least sensitive to the height distribution of the vapor due to pressure broadening effects. There is considerable latitude, however, in the choice of the third frequency for clear conditions. This is shown in Fig. 8. The

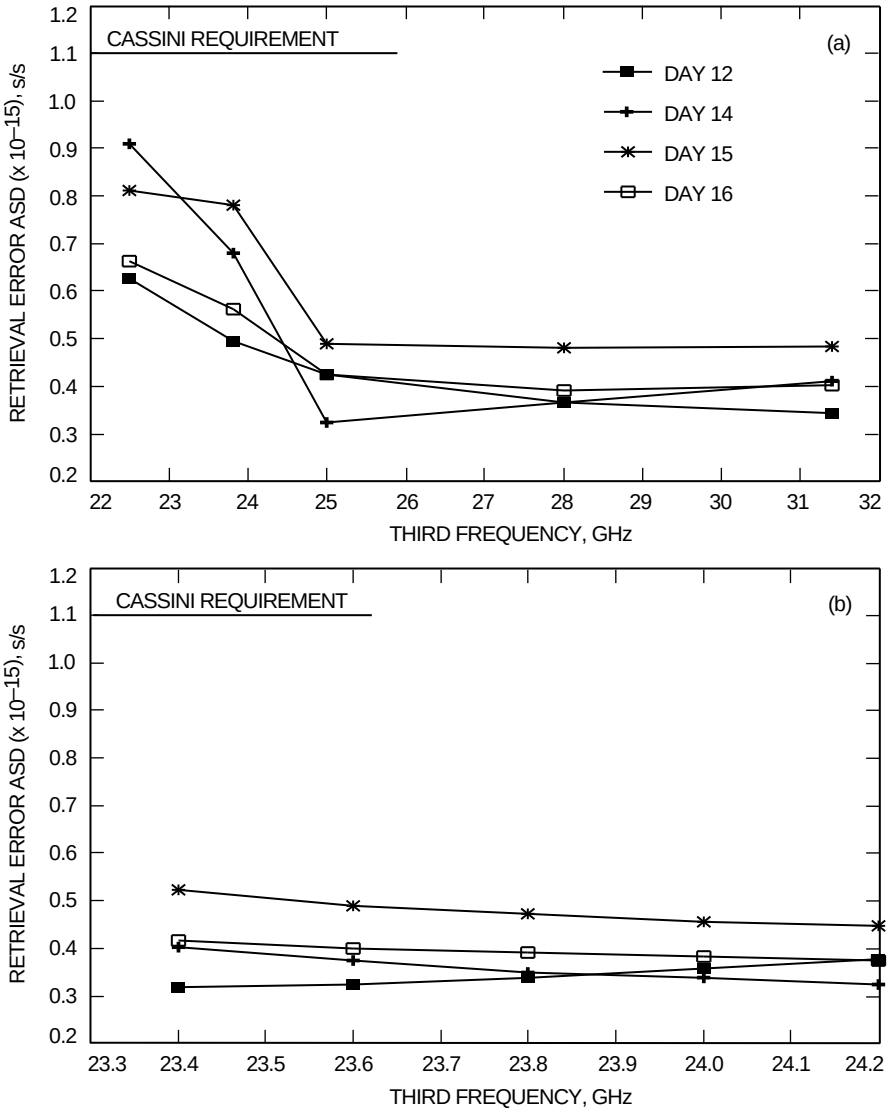

Fig. 8. The dependence of the retrieval error ASD at a 960-s time scale on the choice of a third WVR frequency for the 4 days of Coffeeville testbed data, where the first two frequencies are (a) 20.7 and 22.235 GHz and (b) 22.235 and 31.4 GHz.

upper panel displays the retrieval error ASD (at a 960-s time scale) versus the choice of the third WVR frequency for the four nights of lidar testbed simulations. For this case, line center and the 20.7-GHz hinge point have been chosen as the primary vapor-sensing channels. Note that there is no significant performance impact for selected third frequencies in the range of 25 to 31.4 GHz and that performance actually deteriorates if the third frequency is chosen closer to the 22.235-GHz vapor line center. These results also indicate that there is no significant advantage in clear weather of replacing the 31.4-GHz channel with a lower frequency. Thus, based on the expectation that the Cassini GWE will desire a "best-effort" performance during cloudy conditions, it is recommended that the 31.4-GHz channel be retained as the third frequency in the advanced WVR design.

The lower panel of Fig. 8 shows the sensitivity of ASD performance as the hinge frequency is varied over a 0.8-GHz range surrounding 23.8 GHz. Here we have assumed the line center and the 31.4-GHz cloud-sensing channel as the two given frequencies. The essentially flat result indicates that the hinge

point frequency can be varied by at least ±0.4 GHz without significantly degrading the ASD retrieval performance. Results are very similar at the 120- and 4800-s time scales and suggest that we will have sufficient flexibility in frequency selection to avoid potential problems, such as radio frequency interference (RFI), that could restrict our channel options. Assuming no restrictions, our WVR frequency recommendation would be 22.235, 23.8, and 31.4 GHz. (The 23.8-GHz hinge point yields slightly better ASD performance than the 20.7-GHz channel.) We will hereafter refer to this WVR frequency combination, coupled with the MTP and surface meteorology observables, as the nominal Cassini GWE tropospheric calibration system.

B. The Effects of Sidereal Line-of-Sight Tracking

Path delay and path-delay fluctuation levels increase with air mass. All simulations to this point have assumed constant zenith observations, for which the line-of-sight delay errors are minimized. For sidereal tracking of Cassini, line-of-sight calibration will be required for all spacecraft positions above 20-deg elevation, and the question remains of how the retrieval errors will scale with air mass. It is not necessary to simulate the continuous elevation variation of sidereal tracking to address this issue. Because the effective speed of sidereal tracking through the troposphere is much less than the wind speed, the retrieval error ASD results will not differ significantly from the case of fixed-elevation observations obtained at the midpoint elevation of any Cassini tracking interval. To investigate the worst-case scenario, we have repeated the lidar testbed simulations using the nominal GWE calibration system, but with the WVR pointing fixed at 20-deg elevation, the maximum air mass encountered during Cassini tracking. The result of the fixed 20-deg simulation is that the ASD of the line-of-sight retrieval error increases by a factor of three relative to the fixed zenith case, the same factor as the air mass ratio between 20-deg elevation and zenith. Thus, we have confirmed in the ASD domain that the retrieval error will scale linearly with air mass, and the tropospheric calibration system performance relative to Cassini requirements must be evaluated accordingly. The results shown in Figs. 5 through 7 are representative of near-zenith tracking. For elevation angles above 60 deg, corrections to the predicted ASD performance shown in Figs. 5 through 7 will be less than 15 percent. At the lowest tracking elevations (20 to 40 deg), the predicted retrieval error ASD values will increase by a factor of from 1.5 to 3.0 relative to the results shown in Figs. 5 through 7. The implications for meeting the Cassini GWE requirements at all time scales and elevation angles are discussed further in Section VII.

C. The Effects of WVR Calibration Offsets

The radiometer instrument errors to this point have been assumed to be purely stochastic with zero bias. The question remains of the effect of realistic instrument calibration offsets, in particular offsets that vary with WVR channel. To evaluate the effects of offsets, we have repeated the retrieval simulations for days 12 and 15, using the nominal calibration system (22.235-, 23.8-, and 31.4-GHz WVR plus MTP plus surface meteorology), but with 1-K offsets added to one, two, or all three WVR channel observables. For all cases, significant path-delay retrieval biases (0.1 to 0.5 cm) occurred. As expected, the case for which all three channel biases were set to 1 K produced no significant degradation in ASD performance; i.e., the retrieval error did not vary significantly with time. However, for cases in which only one or two of the WVR channels contained 1-K offsets, ASD performance degraded by approximately 40 to 60 percent, dependent both on the testbed data set (day 12 or 15) and on which of the channels included the offset. Thus, *relative* WVR channel offsets can be expected to degrade the tropo-cal performance, and efforts should be made to ensure that the actual relative offsets are limited to a level less than ∼0.5 K.

We also simulated relative offsets in the MTP measurements and found the effects to be much less than those for comparable WVR offsets. The worst case examined assumed a +1 K offset for the 54.4-GHz channel and a −1 K offset for the 57.97-GHz channel, producing an approximate 10-percent degradation in ASD performance over the 960- and 4800-second time scales. MTP offset effects at the 120-second time scale were insignificant for all cases considered in which relative channel offsets did not exceed 2 K.

VII. Nominal System Performance Versus Radiometer Stability

Thus far the retrieval simulations have been used to establish a prototype tropospheric calibration system consisting of a three-channel (22.235-, 23.8-, and 31.4-GHz) WVR, a two-channel (54.4- and 57.97-GHz) MTP, and a three-component surface meteorology package. If the targeted radiometer stability goals (0.01 K for the advanced WVR and 0.1 K for the MTP) are met, the nominal tropo-cal system will meet or exceed the GWE requirements at all time scales for the near-zenith (approximately >60-deg elevation) portion of Cassini sidereal tracking. To assess the degradation of performance in the event that the radiometric stability goals cannot be met, we have repeated the simulations using the nominal tropo-cal system for wide ranges of assumed WVR and MTP stability that span the targeted and worst-case performance range. The simulations were performed on the day-12 and day-15 testbed profiles, assuming a constant line-of-sight elevation angle of 45 deg, an intermediate value for Cassini tracking intended to be representative of mean air mass conditions for the GWE. The results for day 12 are shown in Fig. 9 as contour plots of two-way retrieval error ASD versus WVR and MTP stability. Results for day 15 differed from day 12 by less than 20 percent at all time scales. The values designated "CR" and "CG" in the plots are the Cassini requirements and Cassini goals, respectively, for two-way tropospheric calibration. Note that the contours represent the net effect of the retrieval algorithm plus instrument stability error components.

From the results shown in Fig. 9, we draw the following conclusions regarding the impact of radiometer stability on the performance of the prototype Cassini tropo-cal system:

(1) At the 120-s time scale [Fig. 9(a)], the estimated ASD performance for targeted radiometer stability (0.01 and 0.1 K for the WVR and MTP) is approximately 8×10^{-15} s/s, approximately 30-percent higher than the Cassini requirement. Any significant degradation in radiometer stability would lead rapidly to unacceptable performance. However, at this short time scale, it is expected that the stated radiometer stability goals will be most easily achieved, and very possibly surpassed, producing tropospheric calibration performance very close to the GWE requirements.

(2) For intermediate time scales [960 s, Fig. 9(b)], the Cassini requirement is matched with radiometer errors approximately 2 to 3 times as large as the best case considered (0.01 K for the WVR and 0.1 K for the MTP). We note that, while the stated design goal for the advanced WVR is 0.01-K stability, no comparable analysis has yet been made of MTP stability limitations. Our WVR experience suggests that 0.1-K MTP stability should be attainable, but further analysis is needed of the existing ground-based MTP instruments.

(3) For the largest time scales relevant to the GWE [>4000 s, Fig. 9(c)], the retrieval-plus-stability error component should make no significant contribution to the Cassini ASD requirement. In fact, if we are within a factor of three in meeting the radiometer stability targets, the GWE *goal* at $\Delta t >$ 1-hour time scales would be comfortably met for nearly all segments of the Cassini sidereal track.

VIII. Conclusions and Future Work

Simulations of path delay retrieval over approximate 10-hour intervals, using real data for the time variation of atmospheric conditions, indicate that a tropo-cal system consisting of a three-channel WVR, two-channel MTP, and surface meteorology measurements is required to meet the Cassini GWE specifications. A model-based Bayesian inversion algorithm is also required to minimize retrieval errors in the ASD domain. Recommended channels for the advanced WVR are 22.2, 23.8, and 31.4 GHz, with the additional note that any frequency in the range from 23.4 to 24.2 GHz could replace the 23.8-GHz channel with no significant performance impact.

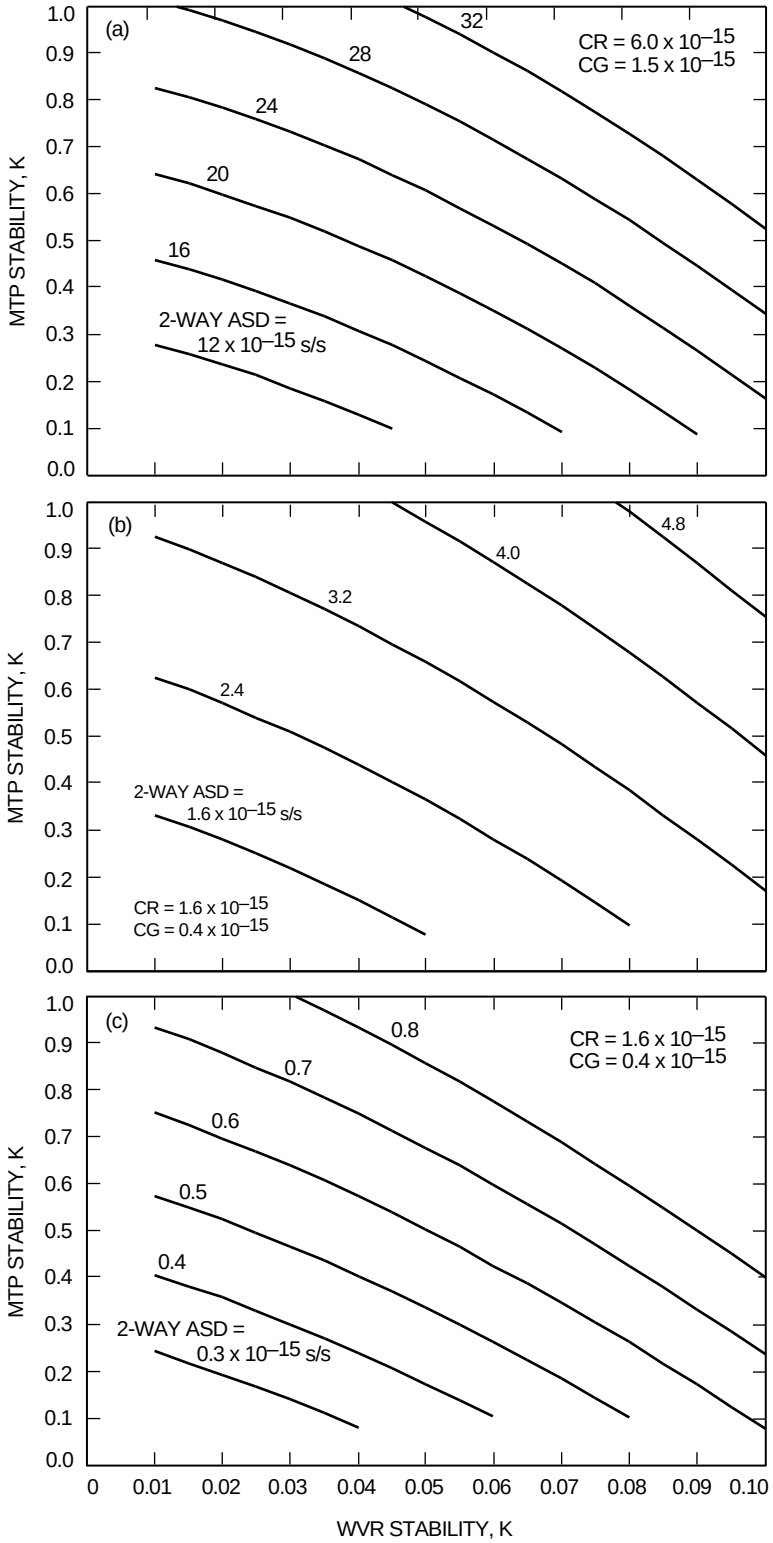

Fig. 9. Contour plots of the ASD of the two-way retrieval error at a 45-deg elevation versus WVR and MTP instrument stability for the day-12 testbed data at a *dt* of (a) 120 s, (b) 960 s, and (c) 4800 s.

If the targeted radiometer stability goals (0.01 K for the WVR and 0.1 K for the MTP) are achieved, the simulations indicate that the proposed tropospheric calibration system will meet or surpass the Cassini GWE requirements at almost all relevant time scales and elevation angles. At the shortest time scale (∼100 s), the retrieval algorithm error component alone will approximately match the GWE requirement, with only slight improvement afforded by improvements in radiometer stability. At the intermediate (∼1000-s) time scale, the targeted radiometer stabilities can be relaxed by a factor of approximately three and still meet the Cassini requirements. For hour time scales and longer, the retrieval algorithm plus radiometer stability error components will not be significant relative to the GWE requirements. At the ∼5000-s time scale, the specified Cassini goal (0.4×10^{-15} s/s ASD) could be met with WVR and MTP stabilities of 0.03 and 0.3 K, respectively.

The stated tropospheric calibration system performance would be significantly degraded by *relative* WVR channel offsets of ∼1-K magnitude. For example, simulations performed with a +1-K offset in the 23.8-GHz WVR channel, and a zero offset in all other radiometer channels, produced a 60-percent degradation in the day-12 ASD performance and a 40-percent degradation in the day-15 performance. Offsets common to all three WVR channels produced no significant degradation in ASD performance. Similar analyses of the effects of MTP offsets showed much smaller effects, with an approximate 10-percent ASD performance degradation in the extreme case of a 2-K relative offset between the two MTP channels.

Recent laboratory measurements of the advanced WVR prototype stability conducted by Dr. Alan Tanner of JPL's Microwave, Lidar, and Interferometer Technology Section suggest that the 0.01-K goal could be surpassed at <1-hour time scales. At a 1000-s time scale, 0.005-K performance is typical. At shorter time scales, the stability becomes thermal noise limited, with a precision level of ∼0.003 K for a 100-s integration time. Dr. Tanner also estimates that planned calibration tests for the advanced WVR will constrain channel calibration relative offsets to the 0.1- to 0.2-K level. An important near-future test for the advanced WVR will be to determine if the stability goals can be achieved outside the controlled laboratory environment.

The stability of current generation MTPs over 10-hour time scales needs to be assessed to determine if precision levels approaching 0.1 K are attainable. Alternative temperature profiling instrumentation (such as RASS) should also be considered, evaluating trade-offs between cost, reliability, and stability. The assumed precision of surface meteorology instrumentation should also be investigated. Of particular importance is determining the availability of temperature and pressure sensors with precision at the 0.2-K and 1-mb level, respectively.

Acknowledgments

The authors thank R. Linfield and M. J. Mahoney for their very useful comments in reviews of preliminary drafts of this article.

References

[1] D. W. Allan, "Statistics of Atomic Frequency Standards," *Proc. IEEE*, vol. 54, no. 2, pp. 221–230, 1966.

[2] B. L. Gary, S. J. Keihm, and M. A. Janssen, "Optimum Strategies and Performance for the Remote Sensing of Path Delay Using Ground-Based Microwave Radiometers," *IEEE Trans. Geosci. Rem. Sensing*, vol. GE-23, pp. 479–484, 1985.

[3] Y. Han, J. B. Snider, E. R. Westwater, S. H. Melfi, and R. A. Ferrare, "Observations of Water Vapor by Ground-Based Microwave Radiometers and Raman Lidar," *JGR*, vol. 99, pp. 18,695–18,702, 1994.

[4] S. J. Keihm, "Water Vapor Radiometer Measurements of the Tropospheric Delay Fluctuations at Goldstone Over a Full Year," *The Telecommunications and Data Acquisition Progress Report 42-122, April–June 1995*, Jet Propulsion Laboratory, Pasadena, California, pp. 1–11, August 15, 1995.
http://tda.jpl.nasa.gov/tda/progress_report/42-122/122J.pdf