Design and performance of a high-stability water vapor radiometer

Alan B. Tanner and A. Lance Riley

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, USA

[1]   The design of two new high-stability microwave water vapor radiometers is presented along with a performance evaluation. The radiometers operate next to a spacecraft tracking station at NASA's Goldstone facility in California, where they will be used to calibrate tropospheric path delay fluctuations during an upcoming gravity-wave search experiment (GWE) involving the Cassini spacecraft. Observing frequencies of the radiometers are 22.2, 23.8, and 31.4 GHz, and the antenna beam width is 1°. The instruments are room temperature Dicke radiometers with additive noise injection for gain calibration. Design highlights include: (1) a practical temperature control system capable of stabilizing the entire receiver to a few millikelvin from day to night; (2) redundant noise diode injection circuits with 30 ppm RF power stability; and (3) a voice coil actuated waveguide vane attenuator which is used as a high-performance Dicke switch. Performance of the radiometers is evaluated from intercomparisons of the two radiometers and from continuous tip curve calibrations spanning nearly 1 year. Structure function analysis of the intercomparison data indicates that the brightness temperature stability of these radiometers is better than 0.01 K on 1000–10,000 s timescales. Analysis of tip curve calibrations indicates RMS errors of ∼0.05 K on 30-day timescales and 0.15 K on 1-year timescales.   *INDEX TERMS:* 6994 Radio Science: Instruments and techniques; 5757 Planetology: Fluid Planets: Remote sensing; 0360 Atmospheric Composition and Structure: Transmission and scattering of radiation; *KEYWORDS:* microwave radiometer, water vapor, radiometer design

Citation: Tanner, A. B., and A. L. Riley, Design and performance of a high-stability water vapor radiometer, *Radio Sci.*, *38*(3), 8050, doi:10.1029/2002RS002673, 2003.

1. Introduction

[2]   Water vapor adds a variable delay to radio signals passing through Earth's troposphere that can limit the phase stability and coherence time for very long baseline interferometry (VLBI) or for experiments that involve precise spacecraft tracking. One technique to correct these errors involves measuring the water vapor along a line of sight with a water vapor radiometer (WVR). WVRs measure the thermal radio emissions of water vapor near a resonance at 22 GHz, and the radio path delay is inferred from these data using models and estimates of the atmospheric temperature and pressure. There are a variety of errors associated with this technique, including model errors and uncertainty in the vertical distribution of temperature and water vapor, but past experiments have been largely limited by the accuracy and stability of the WVR itself [*Linfield et al.*, 1996].

[3]   At JPL, an upcoming series of gravitational-wave search experiments (GWE) involving the Cassini spacecraft will require the most sensitive measure of the water vapor path delay to date. Starting in December of 2001, and during the spacecraft's cruise to Saturn, the Cassini GWE will be conducted from a specially equipped ground station at NASA's Goldstone Deep Space Communication Complex (GDSCC) in California. The principle GWE measurements will be of very small perturbations (potentially less than ∼1 mm) in the spacecraft tracking data on timescales of 1000 to 10,000 s. Of the various errors associated with this experiment (including those of the time standards, mechanical stability of the ground station and spacecraft, etc.), it is the tropospheric path delay corrections that will likely determine the detection threshold for the GWE [*Tinto and Armstrong*, 1998]. When translated into radiometric specifications, Cassini GWE goals call for a WVR that is stable to better than 0.01 K in brightness temperature on 1000 to 10,000 s timescales. This represents about a ten-fold improvement over the demonstrated perform-
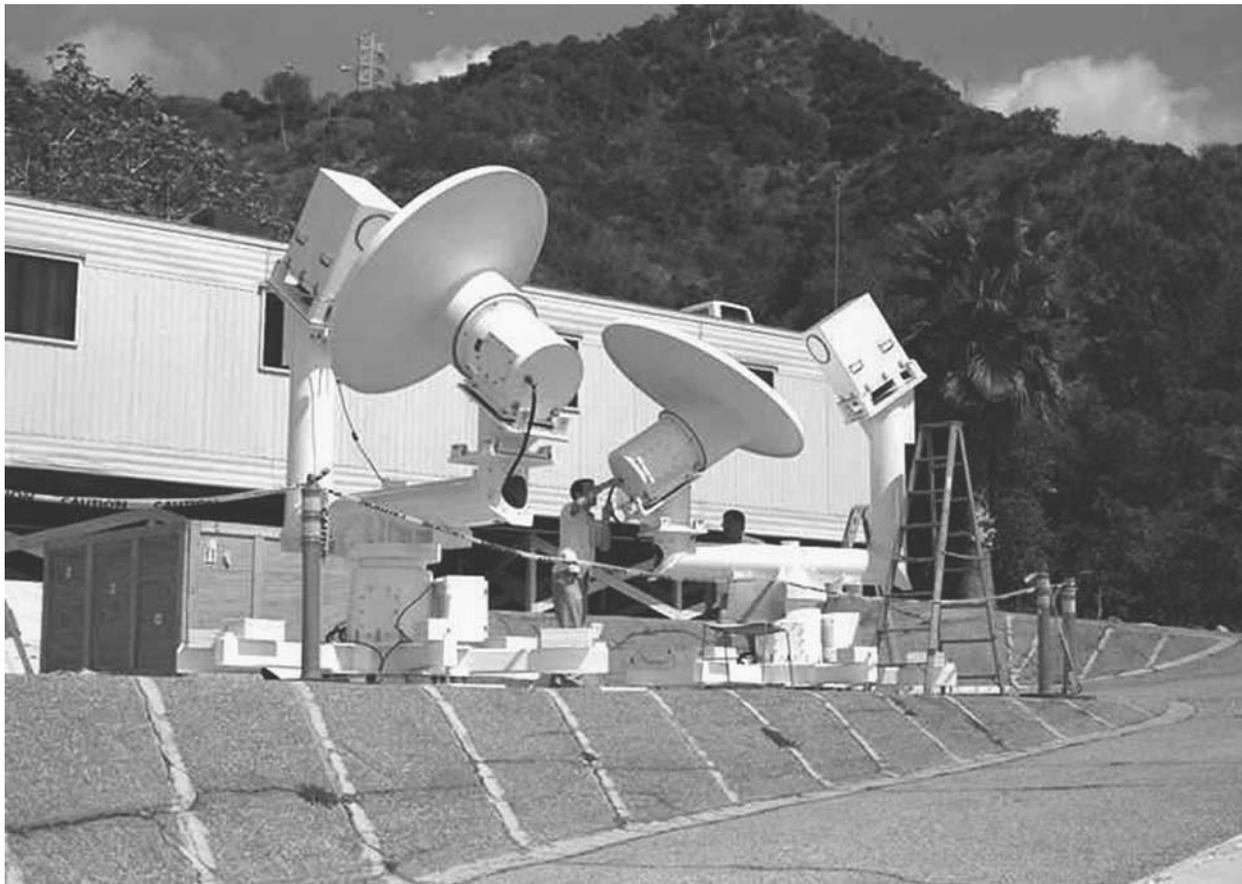
Figure 1. Two AWVRs during assembly at JPL. See color version of this figure at back of this issue.

ance of any previous WVR, and has motivated the development of the radiometers described in this paper.

[4] This paper presents the design and performance evaluation of two WVRs which were developed for the Cassini GWE. An overview of the completed systems (dubbed the ''advanced'' water vapor radiometers (AWVRs)) is presented in section 2. These instruments are based on laboratory tests and prototype work, which were previously reported [*Tanner*, 1998]. The laboratory work produced a number of innovations, which are reviewed in section 3. Section 4 briefly describes the internal calibration algorithm, and section 5 presents a performance evaluation of the radiometers based on side-by-side comparisons and long-term calibration data spanning nearly 1 year.

2. System Overview

[5] The two AWVR units, designated by serial numbers A1 and A2, are located approximately 40 m to the

southeast and southwest, respectively, of the much larger tracking antenna, designated DSS-25, where the Cassini GWE will be conducted. During a spacecraft track both AWVRs will be pointed in the same direction as the DSS-25 antenna. During off hours the AWVRs perform continuous sweeps in elevation as part of a "tip curve" procedure that is the basis for the system calibration. Two identical radiometers have been built to ensure reliability and to provide a crosscheck of the system performance.

[6] Figure 1 shows the two AWVRs during assembly at JPL. Each AWVR consists of the radiometer box, the supporting pedestal, and the offset parabolic reflector. The overall height of the system is about 3 m, and the reflector size is about 1.6 by 1.9 m. The radiometer box contains the antenna feedhorn (hidden behind the circular radome on the face of the enclosure), RF electronics, data system, and temperature control systems. The pedestal includes the azimuth positioner at the base of the structure and the elevation positioner, which rotates the reflector on a focal axis that is inclined by 35° from

Table 1. Short List of AWVR Specifications

Frequencies, GHz	Bandwidth, MHz	NEDT, $1000-10,000$ s, K	Calibration Uncertainty, K	Antenna Beamwidth, deg	Pointing Uncertainty, deg	1000 s Pointing "Jitter," deg
22.2 23.8 31.4	600	0.01	0.5	1.0	0.05	0.005

horizontal. This 35° inclination is complementary to a 55° offset between the focal axis and the main beam of the antenna. Rotation of the elevation positioner to other angles results in a conical sweep of the antenna beam about the focal axis, and requires a coordinate transformation to convert between positioner angles and the true azimuth and elevation, as referenced to the local horizon. In practice these conversions are calculated by the controlling computer and are transparent to the user.

[7] A short list of specifications for the AWVR is summarized in Table 1. The frequency selections are based on work by *Keihm and Marsh* [1996, 1998] to minimize the uncertainty in the path delay inversion algorithm. The antenna beam width is designed to minimize the mismatch with the 34 m diameter tracking antenna at a range of 2 km, which is approximately the midpoint of the tropospheric water vapor distribution (a 1° beam projects a 34 m spot diameter at 2 km range, which roughly matches the near-field pattern of the DSS-25 tracking antenna) [*Linfield and Wilcox*, 1993; *Linfield*, 1998]. The pointing requirements of Table 1 were derived from the sensitivity of the measurement to air mass, which increases rapidly (as the secant of elevation angle) as the antenna beam approaches the horizon.

[8] A priority of the antenna design was to minimize far-sidelobe contamination. The feedhorn provides a Gaussian illumination with a very low main reflector edge illumination of about −25 dB, and very low spill-over of less than 0.3%. These requirements were derived from simulations of worst-case noise temperature contamination during sidereal antenna movements, and on the requirement that fluctuations not exceed 0.01 K on the 1000 s timescale of interest to the Cassini GWE. Subsequent tests of the completed antenna and radiometers have verified that there is no evidence of antenna backlobes at a detectable threshold of 0.01 K during repeated azimuth scans of the AWVR pedestals.

3. Radiometer Design

3.1. Temperature Control

[9] The AWVR "box" of Figure 1 incorporates a number of design features that relate to the stability of the system. First and foremost is the temperature control. Figure 2 gives a cross section view of the AWVR interior to illustrate the design. The AWVR uses two thermo-

electric cooler (TEC) circuits and a double insulated enclosure to achieve very tight temperature control of the RF components. The "main" TEC is attached to the base plate of the airtight enclosure, and serves as the heat exchanger to the outside air. This TEC can heat or cool the instrument under the control of a computer. Fans are used to circulate air from the main TEC throughout the main enclosure and around the interior RF enclosure, as illustrated with arrows in Figure 2. The enclosure incorporates air ducts to guide the air over the top of the interior RF enclosure before returning to the main TEC.
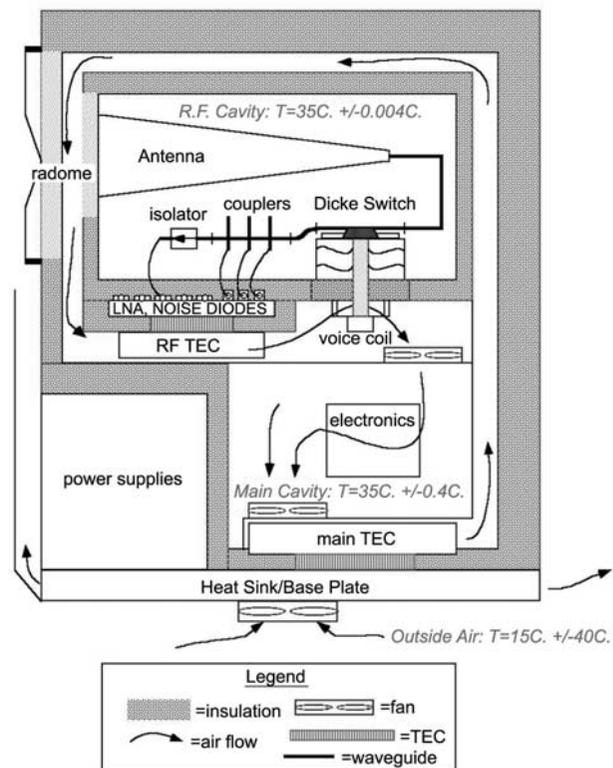

Figure 2. AWVR RF "box within a box" cross-section. Two stages of temperature control and a circulating air system keep the RF component temperatures stable to within ±0.004°C. The voice coil actuated Dicke switch is also shown in cross section. See color version of this figure at back of this issue.

The computer, which is embedded in the electronics of Figure 2, reads temperatures from thermistors distributed throughout the radiometer. A weighted average of these temperatures is applied to a software control loop to maintain the average temperature surrounding the interior RF enclosure. The target temperature for this main control loop is 35 C. Diurnally changing gradients among the thermistors surrounding the RF enclosure are typically ±0.4 C. Also, because of the turbulent airflow within the main cavity, random fluctuations of ±0.1 C are common. These errors are dampened by the insulation of the interior RF enclosure. Within the RF cavity, the design objective is to minimize all heat sources and heat sinks so that RF components can passively settle to a single reference temperature of 35 C. No fans are used inside the RF cavity because these only add heat and turbulence. All active RF components, which include amplifiers and noise diodes, are isolated on a plate that is cooled by a smaller secondary TEC to the same 35 C. Local gradients among the active components are about 0.1 C, but these gradients are constant with time. Insulation is packed around these components to prevent radiant and convective heat flow into the RF cavity. By this arrangement, temperatures within the RF cavity can be stabilized to ±0.004 C. The calibrated accuracy of the temperature sensors is approximately ±0.05 C, based on tests conducted in the JPL Standards Laboratory.

3.2. Humidity Control

[10] Along with temperature control, humidity control within the radiometer is also needed to maintain stability of the RF circuits. This sensitivity was observed in laboratory tests which showed that the forward coupling of a waveguide directional coupler could be changed significantly with a moderate change in humidity (about 200 ppm RF power changes were observed for about 20% change in relative humidity). These errors were tied to changes in the dielectric of the air within the waveguide that affected the standing waves along the coupled arm of the waveguide. To control the humidity of the AWVR the enclosure is sealed and a nitrogen purge is applied for at least one day to remove any moisture. It is important to purge the radiometer because moisture will otherwise tend to migrate into the insulation near the exterior walls of the radiometer at night, and then back into the RF cavity when the walls heat up in the daytime. In the field, if no nitrogen purge is available, air can be vented to the outside via a silica desiccator. Venting is needed to allow for barometric changes.

3.3. Radome

[11] The radome of the AWVR provides a moisture and thermal barrier to the antenna feedhorn, as shown in

Figures 1 and 2. Expanded polystyrene is used for the thermal insulation of the radome in two layers; one on the RF box of Figure 2, and one on the exterior enclosure. To reduce RF reflections all surfaces of the foam are cut with a pattern of v-grooves (not visible in the figures). The exterior moisture barrier consists of a thin sheet of Teflon. As shown in cross-section in Figure 2, this Teflon sheet is also modified with an impressed cone shape. This cone shape is needed to disperse reflected power away from the feedhorn aperture.

3.4. Dicke Switch

[12] Figure 2 also shows the basic layout of the AWVRs' waveguide circuits following the horn antenna. The output of the horn connects to WR-34 waveguide to the right of the figure (drawn as a heavy black line), which then feeds a Dicke switch below. The Dicke switch in this case is a unique mechanical voice coil actuated vane attenuator, which is also shown in cross section in the figure. This switch provides the reference black body temperature to the radiometer by periodically inserting a vane attenuator into a broad-wall slotted waveguide. The voice coil actuator moves the vane up and down approximately 0.7cm at a rate of about 3 Hz with a settling time of about 10 ms. The voice coil itself is at the base of the assembly outside of the RF cavity (for cooling) and is connected to the vane by an insulating fiberglass and foam tube. This tube is held by loudspeaker "spiders" that guide the moving assembly in a straight line without any bearings or contacting parts. The reliability of this switch is excellent: after over two years of continuous operation, or about 200 million switch cycles, there has been no evidence of degradation in either of the two AWVRs.

[13] The RF characteristics of the vane attenuator were the prime motivation to developing the actuator. This device works over the entire waveguide band, and greatly simplified the radiometer design by allowing all three channels of the AWVR to be combined into a single front-end. The insertion loss in the "on" state (i.e., when the vane is pulled out of the slot) is less than 0.3 dB; the loss in the "off" state (vane fully inserted) is more than 45 dB; and in either state of the switch the return loss is well below −30 dB. The loss of this device in the "on" state is also very reliable and stable since, in this state, it essentially is a straight piece of waveguide.

3.5. Noise Diodes

[14] Figure 2 shows a set of three cross-guide couplers following the Dicke switch. Each of the couplers connects via a coaxial cable to a corresponding noise diode on the TEC plate below. The noise diodes are stabilized in temperature and are carefully controlled in terms of bias current and duty cycle to provide a repeatable and stable
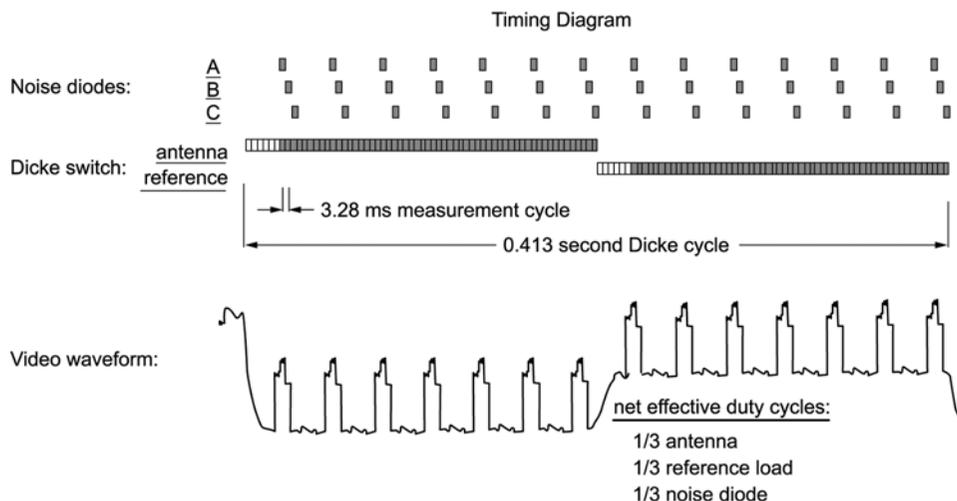
Timing Diagram

Figure 3. Timing diagram, showing the relative timing and duty cycles of the three noise diodes and Dicke switch. See color version of this figure at back of this issue.

additive noise reference to the radiometer. These noise diodes are switched on and off via their bias currents under the control of the computer, and the responses are used to monitor and stabilize the gain of the receiver. The effective noise temperatures of the noise diodes, as referenced to an equivalent source at the antenna, is nominally 300–500 K in the two AWVRs. The stability with which the output of the noise diodes can be compared to power at the antenna input is a critical determinant of the radiometer stability. The fact that there are three noise diodes is another key to the AWVR performance.

[15] One of the most difficult problems concerning the use of coupled noise diodes is the multitude of small reflections and associated standing waves that are inevitable in the interconnecting transmission lines of the radiometer front-end. The laboratory work preceding the AWVR development revealed just how problematic standing waves are, and how difficult it is to compare the power coupled from the noise diode with the power input at the antenna. A simple innovation of the AWVR in this regard is the use of redundant noise diodes. By injecting several noise diodes at different points along the transmission line behind the antenna the system can at least measure (and to a degree mitigate) variations in the noise diode versus antenna coupling. This feature is invaluable as a diagnostic tool as it provides a continuous and relevant measure of the system performance. Relative changes in the noise diode outputs are continuously monitored in the data system. On many occasions these changes have been used to quickly identify and repair some very subtle problems in the AWVR. For example, internal humidity changes, movement of the radome, and a variety of coaxial cable problems have all been identified

with this technique. The fact that the noise diodes are injected after the Dicke switch also helps by providing such comparisons in either state of the Dicke switch.

3.6. Receiver

[16] The output of the couplers in Figure 2 is connected to a ferrite isolator, which in turn passes the signal through a waveguide-to-coaxial adapter, and then to the low noise amplifier (LNA) on the RF TEC plate below. The isolator prevents the relatively poor mismatch and the cold noise temperature of the LNA from affecting the calibration circuits above. The LNA has a noise figure of approximately 3 dB from 22 GHz to 32 GHz. Following the LNA, the remainder of the receiver (not shown in the figure) is a 3-channel single-sideband direct-detect system: a frequency triplexer divides the RF signal into the three observing bands, which are again amplified, and then passed to three zero-bias Shottky detectors. All of these devices are commercial coaxial components, and all are attached to the same TEC plate at the base of the RF enclosure. The RF power level reaching each of the detectors is approximately −17 dBm. OP−37 video amplifiers then boost the detected signal to about 2 v, which is then passed to digitizer/integrators in the electronics assembly below the RF enclosure. The digitizer/integrators consist of synchronous voltage-to-frequency converters followed by digital counters, as is common in radiometer design.

3.7. Timing

[17] The operation of the Dicke switch and each of the noise diodes follows the timing diagram of Figure 3. A

simulated video waveform is also illustrated in Figure 3. The timing of the measurement modes is determined by speed limitations of the mechanical Dicke switch, and a desire to operate the noise diodes at a constant duty cycle and with short pulses to minimize thermal cycling of the devices. The sequence of measurements is controlled by an interrupt service routine in the computer. Digital hardware generates interrupts every 3 milliseconds, and the service routine reads the digitizer/integrators and programs the next calibration mode on each measurement cycle. The digitized data are buffered by the interrupt routine and then passed to a data acquisition program which sorts and averages the data by frequency and by noise diode and Dicke switch modes. The averaged data are recorded on an embedded disk drive.

4. Calibration

[18] The calibration of the AWVR involves two steps: (1) a calibration of the noise diode temperatures, and then (2) application of these noise diodes temperatures to the routine radiometer observations. The noise diode calibration occurs on about a once-per-month basis, and involves a tip curve procedure as discussed below. The noise diodes are used to transfer the calibration to the routine observations according to

$$T_B \cong T_o - \frac{C_o - C_s}{C_{NDD}} T_{ND}, \qquad (1)$$

where T_B is brightness temperature, T_o is the physical temperature of the Dicke switch reference load, C_o is the measured "count" in response to this load (the word "count" here is a reference to the counter of the digitizer, and is synonymous with "integrated detector voltage"), C_s is the "sky" count as measured during the antenna mode of the Dicke switch, C_{NDD} is a sum of the counts-deflected for all the modes of the noise diodes (as defined below) and T_{ND} is a sum of the corresponding noise diode temperatures. There are three noise diodes that are injected in the two modes of the Dicke switch for a total of six noise diode measurements. The noise diode terms of (1) expand to:

$$C_{NDD} = C_{sNDA} + C_{sNDB} + C_{sNDC} - 3C_s + C_{oNDA}$$
$$+ C_{oNDB} + C_{oNDC} - 3C_o \qquad (2)$$

and

$$T_{ND} = T_{sNDA} + T_{sNDB} + T_{sNDC} + T_{oNDA} + T_{oNDB}$$
$$+ T_{oNDC}, \qquad (3)$$

where the "s" and "o" subscripts denote sky and reference modes, respectively, the "A", "B", and "C" subscripts refer to each of three noise diodes, and C_o and

C_s are the same reference and sky counts from equation (1). Note that the "s" and "o" subscripts of the noise diode temperatures in equation (3) could be omitted by assuming $T_{sNDA} = T_{oNDA}$ (etc. for B and C). In practice we find that the sky and reference mode temperatures of a given noise diode typically differ by about ±0.1%, and that it is a useful diagnostic tool to keep track of these differences (which should not change by more than 0.005% if the system is stable).

[19] In equation (1), T_o is readily measured with a thermistor placed next to the absorbing vane of the Dicke switch and is very close to 308.14 K (=35 C). This leaves only T_{ND} in equation (1) to be calibrated by the tip curve procedure. A description of this procedure is given by *Elgered* [1993], and a more detailed analysis of common errors associated with this technique is given by *Han and Westwater* [2000]. Conceptually this procedure involves observing the sky at various elevations and then extrapolating the measured curve of observed counts versus air mass in the line of sight (= secant of elevation assuming a horizontally uniform atmosphere) to zero air mass. At zero air mass the extrapolated counts correspond to the known cosmic background temperature of 2.7 K, which establishes a cold reference to calibrate the noise diodes. At low atmospheric opacities the relation between counts (proportional to brightness temperature) and air mass (proportional to atmospheric opacity) is nearly linear. At higher opacities this relation becomes non-linear, and solutions must employ a radiative transfer model to obtain an accurate calibration. In practice, data processing typically involves an iterative loop to adjust the calibration until the measured brightness temperatures fit the opacity model. The accuracy of the solution depends on the accuracy of the radiative transfer model and the degree to which the atmosphere is horizontally uniform. In the case of the AWVR, we have found that the stability of the radiometer is better than the repeatability of a single tip curve, and we typically combine many days of repeated tip curves to form a single estimate of the noise diode temperatures. The AWVRs are programmed to automatically perform continuous tip curves any time they are not in use. Some specifics of our procedure and data processing follow:

[20] Our tip curve procedure involves continuous sweeps of the antenna between ±60° from zenith. The sweep rate is a constant 0.5° of elevation per second, and radiometric data are integrated and recorded at 5 s intervals continuously throughout the sweeps. Each sweep takes about 5 min, so several hundred sweeps are typically available on a given day. These rates were set as slow as possible to minimize wear in the scan mechanism without severely undersampling the atmosphere. Five minutes corresponds, very roughly, to the time required for water vapor to move across the field of

view (given a 10 m/s wind velocity and a 3 km field of view at 1 km altitude).

[21] Tip curve data are initially processed in six-hour segments, producing four new noise diode calibrations per day. Our algorithm starts by computing a 6-hour time series of brightness temperatures with equation (1) from an initial (best guess) noise diode calibration. These data are then converted from the known elevation angle to zenith-equivalent brightness temperatures using the radiative transfer model (as by *Elgered* [1993]). The zenith-equivalent data are then screened by removing all points for which the 31 GHz brightness temperature is either (1) greater than 30 K (indicating heavy clouds or rain) or (2) is more than 2 K higher than the lowest brightness temperature within the tip (or within 5 min). This "2 K above minimum" filter eliminates a variety of transient errors, including clouds, sun crossings, and other occasional obstructions. This step also requires an accurate initial noise diode calibration, and may be bypassed if none is available. The "clean" zenith-equivalent brightness temperatures are then cross-correlated with air mass (secant of elevation angle) for the entire six-hour segment, and the resulting correlation coefficient is used to adjust the noise diode temperatures up or down. This process is repeated until the correlation coefficient is nulled, at which point the noise diode temperatures are logged, and the processing algorithm moves on to the next six-hour segment. Note that our null-correlation approach accomplishes essentially the same goal as solutions described by *Han and Westwater* [2000] which adjust the calibration to minimize the difference between the zenith mapped brightness temperatures and their mean. We have adopted the correlation approach merely for the simplicity of the computer code, and the efficiency with which many tip curves can be processed in a single iteration. The accuracy of either approach depends on the same radiative transfer model used to map the brightness temperatures to zenith equivalent temperatures.

5. Performance Evaluation

[22] The following performance evaluation is broken into three parts: (1) "delta-T" to compare the short term noise of the AWVR with theoretical limits; (2) "stability estimates" to establish the performance on timescales of about 1000 s to 2-days based on a side-by-side comparison of the two radiometers; and (3) "tip curve" results which track the noise diode calibration history for nearly a year of operations at Goldstone.

5.1. Delta-T (Theoretical White Noise)

[23] We have computed the theoretical white noise limit, or delta-T, of the AWVR from the known detection bandwidth, B, and integration time, τ, using equations (1)–(3) and the duty cycles of Figure 3. A brief summary

of this derivation follows: We started by expanding each of the counts of equations (1) and (2) in terms of the expected receiver gain, expected noise temperatures, and a small delta term to represent the measurement error. For example, the sky counts of equation (1) were replaced with $C_s = g(T_r + T_B)(1 + \delta_s)$, where g is gain (with units of counts/Kelvin), T_r is the receiver noise temperature, and δ_s is a random variable representing the zero-mean fractional (or "normalized") error associated with the sky measurement. Similar expressions for the reference and noise diode measurements (with independent errors δ_o, δ_{osNDA}, δ_{sNDB}, etc.) were applied, resulting in a long expression for the brightness temperature and all measurement errors. By applying approximations such as $(1 + \delta_s)^{-1} \cong (1 - \delta_s)$ and $(1 + \delta_s)(1 + \delta_o) \cong (1 + \delta_s + \delta_o)$, the brightness temperature equation was reduced to a form involving a linear sum of the independent errors. The variance of the net error was then calculated using the standard radiometer formula and the fact that the cross correlations of the independent errors are zero. For example, the variance of δ_o is $3(B\tau)^{-1}$ given an RF detection bandwidth B (in Hz), integration time τ (in seconds) and a duty cycle of 1/3 (from Figure 3). Finally, the delta-T was computed as the root of the net variance, resulting in the following formula:

$$\Delta T = \sqrt{\frac{3\left(1 + 3\frac{T_o - T_B}{T_{ND}}\right)^2 (T_r + T_o)^2 + 3\left(1 - 3\frac{T_o - T_B}{T_{ND}}\right)^2 (T_r + T_B)^2 + 18\left(\frac{T_o - T_B}{T_{ND}}\right)^2 \sum_{x,y}(T_r + T_x + T_{xNDy})^2}{B\tau}},$$

(4)

where x refers to the Dicke switch mode ("o" or "s" from equation (2)), and y refers to a specific noise diode ("A", "B", or "C"). In the case of the AWVR the bandwidth, B, is about 600 MHz, and the receiver noise, T_r, is about 500 K.

[24] To measure the white noise of the two AWVRs the radiometers were operated side-by side while observing the sky at zenith. The difference between the brightness temperatures of the two radiometers at each frequency was then used to estimate the combined delta-T according to:

$$E_N(t) \equiv \sqrt{\frac{1}{2}\left\langle (D(t, t) - D(t + t, t))^2 \right\rangle}, \qquad (5)$$

where $D(t, \tau)$ is the brightness temperature difference as a function of time, t, and integration time, τ, and $\langle\rangle$ is the expectation operator. If the radiometers are perfectly stable and calibrated, then E_N should equal the root-sum-square of the two radiometers' delta-Ts.

[25] Figure 4 summarizes the combined delta-T of the A1 and A2 radiometers as computed from measurements with equation (5) (solid lines) and from the root-sum-square (RSS) of the theoretical estimates of equation (4) for the two radiometers (dashed lines). Figure 4 is a log-
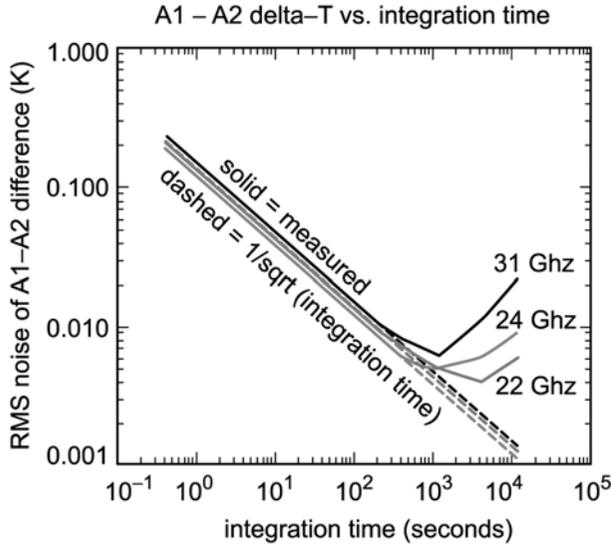
Figure 4. Theoretical (dashed) and measured (solid) delta-T from side-by-side comparisons of the two AWVRs. See color version of this figure at back of this issue.

log plot of noise versus integration time, starting at a minimum integration set by the Dicke switch cycle time of 0.413 s. As can be seen in this figure, the radiometer noise agrees very well with the theory out to an integration time of 1000 s, at which point instability of the two radiometers prevents any further noise reduction. The minima near 1000 s of the three channels are approximately 0.005–0.008 K. Since Figure 4 represents the difference of two nearly identical radiometers, these numbers can be divided by $\sqrt{2}$ to estimate that the noise of individual radiometers is approximately 0.0035 to 0.0056 K. At 1000 s timescales the Cassini requirements have been translated to a 10 mK radiometric stability requirement. From Figure 4, the AWVRs will evidently meet this requirement comfortably.

5.2. Stability Estimates From Side-By-Side Data

[26] The measured delta-Ts presented above were calculated from the difference of consecutive samples separated by the integration time of the boxcar averages. This calculation is useful for determining where the instrument stability exceeds the theoretical noise, but it is not necessarily the best measure of instrument stability at longer timescales since integration times exceeding 1000 s are of little practical interest. To measure the instrument stability on longer timescales we instead fix the integration time and vary the time separation of the A1–A2 brightness temperature difference. As a sample case for the following analysis, Figures 5 and 6 give the brightness temperature and brightness temperature differ-
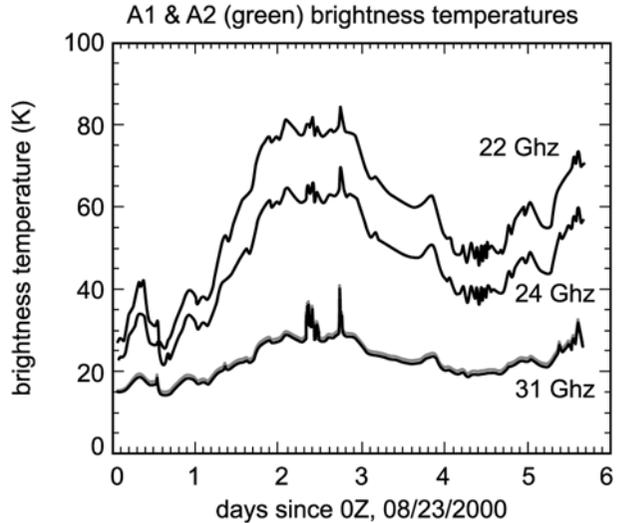

Figure 5. Brightness temperature overlay of the A1 radiometer (black) and A2 (green). Only the 31 GHz channels of A1 and A2 differ enough to discriminate between the two instruments. See color version of this figure at back of this issue.

ences, respectively, from sky measurements spanning 5 days. A boxcar integration time of 124 s applies. In Figure 5 the overlay of brightness temperatures of the two radiometers match too well to discriminate between the A1 and A2 results, but Figure 6 shows the errors clearly (including a large drift in the 31 GHz channel that will have some bearing on the following analysis). The calibrations of the AWVRs were adjusted at the start of
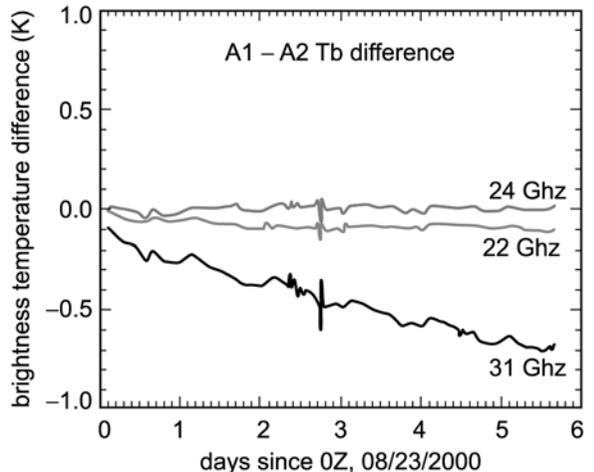

Figure 6. Difference of A1 and A2 brightness temperatures from Figure 5. Integration time is 124 s. See color version of this figure at back of this issue.
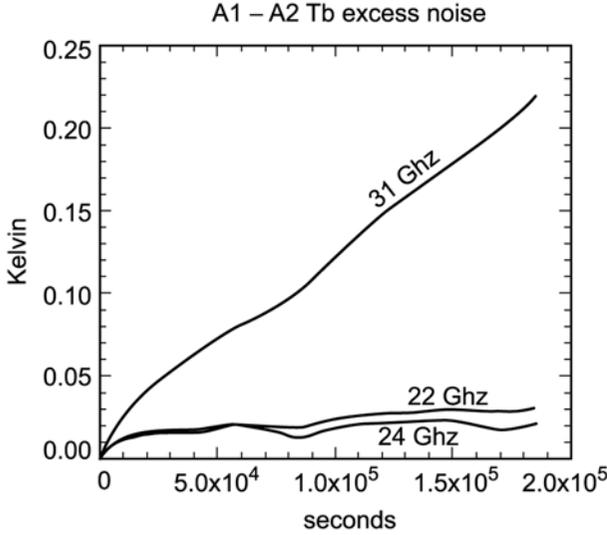
Figure 7. Excess noise structure of the A1−A2 brightness temperatures difference of Figure 6 computed with equation (6).

these measurements to null the initial difference of Figure 6, so it should be noted that these results are unrelated to the absolute accuracy of either radiometer.

[27] Figures 7 and 8 give two different estimates of the AWVR "excess" errors (meaning errors in excess of the inherent delta-T noise predicted by equation (4)) versus timescale. Figure 7 was computed from

$$E_S(\tau) \equiv \sqrt{\left\langle (D(t,\tau_0) - D(t,\tau,\tau_0))^2 \right\rangle - \left\langle (D(t,\tau_0) - D(t,\tau_0,\tau_0))^2 \right\rangle},$$

(6)

where $D(t, \tau_o)$ is the A1−A2 brightness temperature difference from Figure 6 with a fixed integration time of $\tau_o = 124$ s, and τ is the time difference corresponding to the horizontal axis of Figure 7. The right hand expectation term in equation (6) negates the white noise of the radiometer. For a sufficiently small τ_o, this subtraction renders $E_s(\tau)$ insensitive to integration time. In this case τ_0 is 124 s where, from Figure 4 and equation (5), this term is consistent with the theoretical white noise.

[28] As can be seen in Figure 7 the 22 and 24 GHz channels both exhibit a stability of less than 0.03 K excess error. The minima near 86,000 and 172,000 s (or 1 and 2 days) in the 22 and 24 GHz channels also indicate that much of the error is diurnal. The large drift in the 31 Ghz channel, on the other hand, contributes a great deal to the excess error of Figure 7.

[29] The radiometric requirements for the Cassini GWE were originally derived from requirements expressed in terms of the Allan Standard Deviation (ASD) [*Allan*, 1966] of the atmospheric path delay measurement error.

ASD is commonly used to measure the stability of clocks for applications, which are insensitive to a linear drift, as is the case with the GWE. To evaluate the stability of the AWVRs in the same context, Figure 8 plots a modified ASD, which has been calculated from the data of Figure 6 according to

$$E_A(\tau) \equiv$$

$$\sqrt{\frac{1}{2}\left\langle (D(t,\tau_0) - 2D(t,\tau,\tau_0) + D(t,2\tau,\tau_0))^2 \right\rangle - \frac{1}{2}\left\langle (D(t,\tau_0) - 2D(t,\tau_0,\tau_0) + D(t,2\tau_0,\tau_0))^2 \right\rangle}.$$

(7)

[30] Again, as in equation (6), the right-most term here negates the theoretical delta-T so that the result represents the excess error due to instrument instability. Equation (7) also differs from the conventional ASD definition in that D is in units of Kelvin (not time), and there is no division by τ (so that the result is in units of K, not s/s per convention).

[31] As can be seen in Figure 8, the definition of equation (7) has suppressed the linear drift in the 31 GHz data that produced the large error in Figure 7. Figure 8 also exhibits more pronounced minima at timescales of 1 and 2 days in all three RF channels, which again indicates that much of the remaining errors are repeatable with time of day. Overall, Figures 7 and 8 demonstrate that the stability of a radiometer is a strong function of interpretation and the timescale of interest. We also note that these data span only five days (due to availability of the instruments for this test), so one may expect sampling errors for the longer (>1 day) timescales.
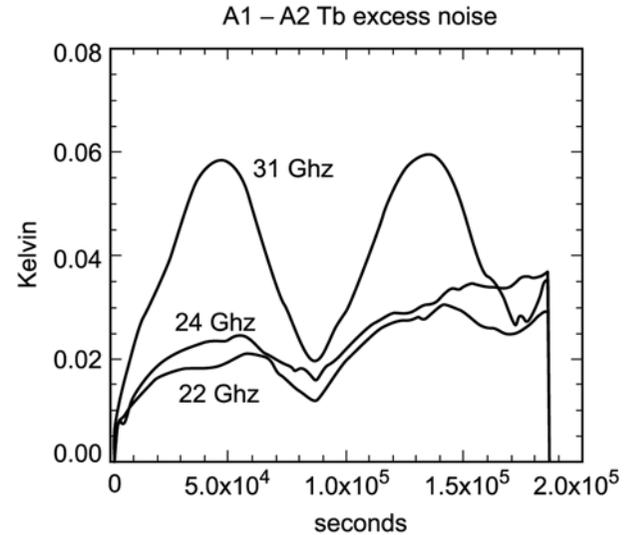

Figure 8. Excess noise of A1−A2 brightness temperatures difference from Figure 6 computed as a modified Allan standard deviation with equation (7).
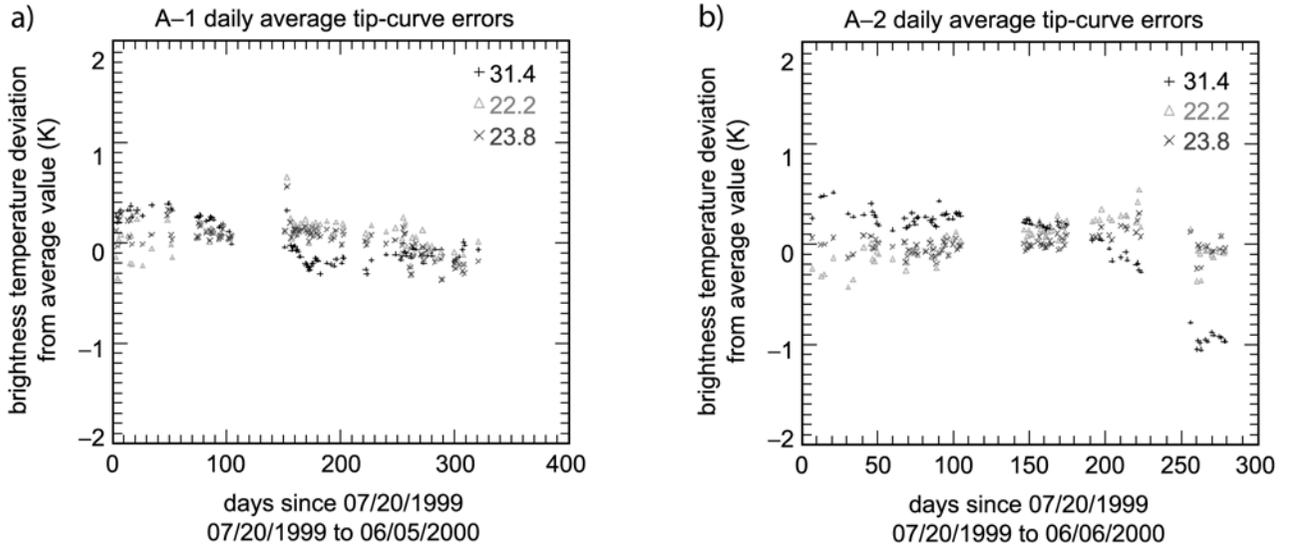
Figure 9. Difference of T_B calibrated with a one-year average of all tip curve results, and those calibrated from one-day tip curve averages for A1 (a) and A2 (b). Positive error means that the long-term calibrated AWVR overestimated the brightness temperature on a given day. See color version of this figure at back of this issue.

[32] The 31 GHz performance presented above is generally poorer than the 22 and 24 GHz channels due to known standing-wave problems. One problem is related to a standing wave between the radome and the directional couplers. The measured directivity of the couplers used for the noise diode injection is only 17 dB at 31 GHz, compared to 24 dB at 22 and 24 GHz. This is problematic because the reversed coupled signal reflects off the radome and interferes with the forward coupled signal. We have conducted tests and have confirmed that the 31 GHz channels are sensitive to movement of the radome, and we believe that thermal expansion of the enclosure and radome from night to day is partly responsible for diurnal errors observed in the AWVRs. The drift in the 31 GHz data, on the other hand, is related to a different standing-wave problem; the particular data set of Figure 5 was collected shortly after a modification to a coaxial cable in the A1 unit, and the drift in the 31 GHz channel can be linked to comparable drifts observed among the different noise diodes. Based on past experience, such errors can be caused by stresses placed on the plastic insulator after bending the coaxial cable. This drift has since stopped, based on the long-term calibration history presented in the next section.

5.3. AWVR Stability Estimates From Tip Curve Data

[33] The long-term stability of each of the AWVRs can be independently estimated from tip curve data collected over the past year at Goldstone. As discussed in section 4, the tip curve procedure is performed continuously any time the radiometers are not in use. The data collected are used to calibrate the noise diode temperatures of equation (1). The day-to-day history of the noise diode temperatures provides a long-term measure of the radiometer stability.

[34] Figure 9 summarizes the long-term history of the AWVR calibrations. Errors were estimated by subtracting brightness temperatures computed with a one-year average of noise diode temperatures from those computed with one-day averages. Only about 90 of a possible 300 days are plotted in each case, based on a selection criteria that required an entire 24 hours of cloud-free tip curve data. Poor weather, maintenance, and other experiments excluded most days. As can be seen, both AWVRs have drifted very little over nearly 1 year. Only minor long-term errors in the 31 GHz channels are evident. Otherwise, much of the short-term scatter, especially in the 22 GHz channels, can be attributed to days where high humidity degraded the tip curve accuracy. Table 2 summarizes the standard deviations of Figure 9. Also given in Table 2 are standard deviations calculated from the "best" month in which particularly dry and stable weather yielded the highest quality tip curve calibrations. These results indicate a one-month stability of about 0.05 K.

[35] Figure 10 provides estimates of the average tip curve calibration errors versus time of day. These were calculated by binning the tip curve results in 6-hour

Table 2. Brightness Temperature Errors Calculated From Day-to-Day Tip Curve Results of Figures 9 and 10

	RMS TB Errors Estimated From Tip Curves, K		
	31 GHz	22 Ghz	24 Ghz
One Year RMS From Figure 9			
A1, 94 days between 7/20/1999 and 6/5/2000	0.204	0.162	0.151
A2, 89 days between 7/20/1999 and 6/6/2000	0.431	0.185	0.093
Best Month From Figure 9			
A1, 22 days from 12/22/1999 to 1/20/2000	0.086	0.057	0.041
A2, 29 days from 12/10/1999 to 1/15/2000	0.027	0.064	0.045
Diurnal Errors From Figure 10			
A1 versus time of day	0.031	0.037	0.027
A2 versus time of day	0.057	0.039	0.031
Difference of A1−A2 versus time of day	0.072	0.035	0.022

segments sorted by time of day, and then comparing the results with the all-day averages. The standard deviation of the errors in Figure 10 is presented in the last rows of Table 2. Table 2 also presents errors computed from the difference of the A1 and A2 errors, which are in remarkably good agreement with the differential errors presented in Figure 8. Excepting the 31 GHz channel of the A2 radiometer, where we suspect a standing-wave problem, Figure 10 indicates a persistent error that is highest in the afternoon. We have recently confirmed that electrical losses in the AWVR antenna and radome can account for much of these errors. Radiometric tests of reflector samples indicate that reflector losses contribute 0.5 to 0.9 K to the antenna noise temperature. Contributions from reflector spillover have been estimated at about 0.5 K, and radome losses are about 0.3 K,

bringing the total contribution to 1.3 to 1.7 K. These values scale to approximately $0.04-0.05$ K of peak-to-peak antenna noise temperature change given an average day-to-night ambient temperature change of 10 K about a mean of 300 K, which is consistent with Figure 10. We are presently evaluating an antenna loss correction to reduce these diurnal errors.

6. Conclusion

[36] The AWVRs will meet the stability requirements of the Cassini GWE experiment. On timescales of 1000 to 10,000 s the above data demonstrates that the 0.01 K stability goal has been met, with some channels approaching 0.003 K. For a radiometer in which the observed brightness temperature is some 300 K colder
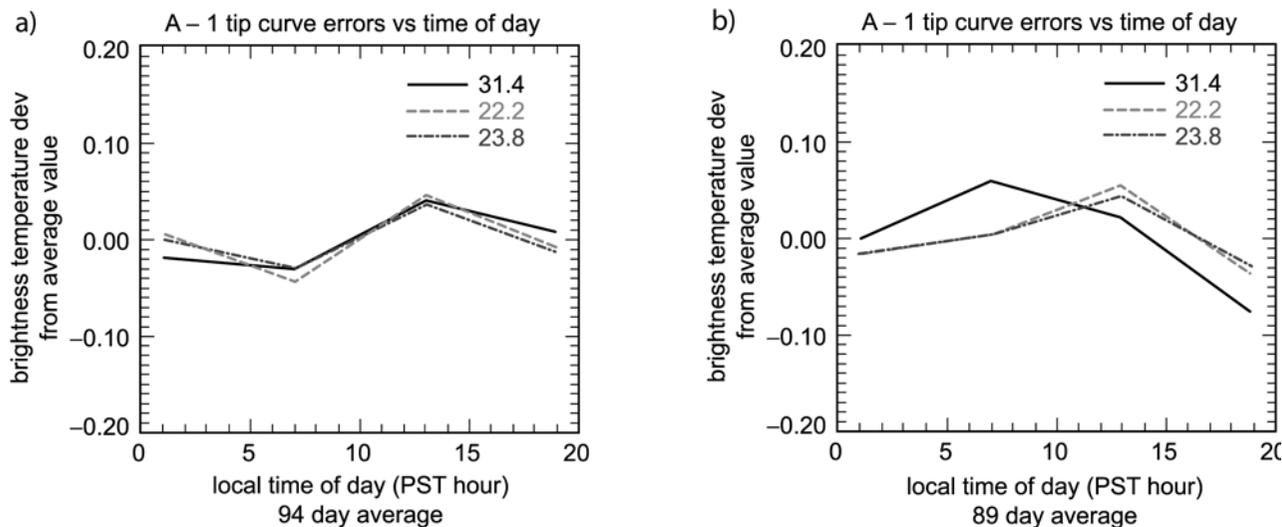

Figure 10. Average brightness temperature error versus time of day for A1 (a) and A2 (b) based on tip curve analysis. See color version of this figure at back of this issue.

than the internal (Dicke) reference temperature, a 0.003 K stability figure implies a gain stability of 10 ppm. Beyond the Cassini requirements, and well past 10,000 s timescales, the AWVRs also exhibit an excellent level of stability of about 0.02 K for day-to-day errors, and approximately 0.03–0.08 K per month.

[37] This report has focused on the stability of the radiometers, but not the absolute accuracy. In Table 1 the quoted calibration uncertainty is given as 0.5 K. This is a conservative estimate based on intercomparisons between the two AWVRs (which agree at the 0.1 K level), and errors expected from standard WVR tip curve calibrations described by *Han and Westwater* [2000]. The estimate is conservative in that at least two significant contributions to the tip cure calibration errors (finite beam width and elevation angle pointing accuracy) have been greatly reduced in the AWVR design (1° versus 5° beam width, and 0.05° versus 0.5° elevation accuracy of standard WVRs used by Han and Westwater).

References

Allan, D. W., Statistics of atomic frequency standards, *Proc. IEEE*, *54*(2), 221–230, 1966.

Elgered, G., Tropospheric radio path delay from ground-based microwave radiometry, in *Atmospheric Remote Sensing by Microwave Radiometry*, edited by M. A. Janssen, chap. 5, pp. 240–243, John Wiley, New York, 1993.

Han, Y., and E. R. Westwater, Analysis and improvement of tipping calibration for ground-based microwave radiometers, *IEEE Trans. Geosci. Remote Sens.*, *38*(3), 1260–1276, 2000.

Keihm, S. J., and K. A. Marsh, Advanced algorithm and system development for Cassini radio science tropospheric calibration, in *The Telecommunications and Data Acquisition Progress Report*, *JPL Rep. 42–127*, pp. 1–20, Jet Propulsion Lab., Pasadena, Calif., 1996.

Keihm, S. J., and K. A. Marsh, New model-based Bayesian inversion algorithm for the retrieval of wet troposphere path delay from radiometric measurements, *Radio Sci.*, *33*(2), 411–419, 1998.

Linfield, R. P., Effect of aperture averaging upon tropospheric phase fluctuations seen with a radio antenna, *Radio Sci.*, *33*(5), 1353–1359, 1998.

Linfield, R. P., and J. Z. Wilcox, Radio metric error due to mismatch and offset between a DSN antenna beam and the beam of a tropospheric calibration instrument, in *The Telecommunications and Data Acquisition Progress Report*, *JPL Rep. 42–114*, pp. 1–9, Jet Propulsion Lab., Pasadena, Calif., 1993.

Linfield, R. P., S. J. Keihm, L. P. Teitelbaum, S. J. Walter, M. J. Mahoney, R. N. Treuhaft, and L. J. Skjerve, A test of water vapor radiometer-based troposphere calibration using very long baseline interferometry observations on a 21-km baseline, *Radio Sci.*, *31*, 129–146, 1996.

Tanner, A. B., Development of a high-stability water vapor radiometer, *Radio Sci.*, *33*(2), 449–462, 1998.

Tinto, M., and J. W. Armstrong, Spacecraft Doppler tracking as a narrow-band detector of gravitational radiation, *Phys. Rev.*, *58*, 042002-1-8, 1998.

_____

A. L. Riley and A. B. Tanner, Jet Propulsion Laboratory, California Institute of Technology, MS 246-101, 4800 Oak Grove Dr., Pasadena, CA 91109, USA. (a.lance.riley@jpl.nasa.gov; alan.b.tanner@jpl.nasa.gov)
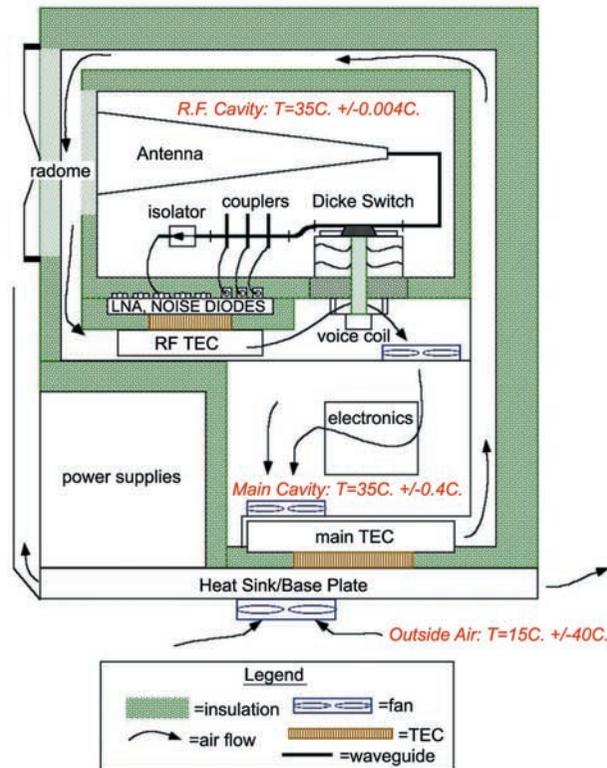
Figure 1. Two AWVRs during assembly at JPL.

Figure 2. AWVR RF "box within a box" cross-section. Two stages of temperature control and a circulating air system keep the RF component temperatures stable to within ±0.004°C. The voice coil actuated Dicke switch is also shown in cross section.
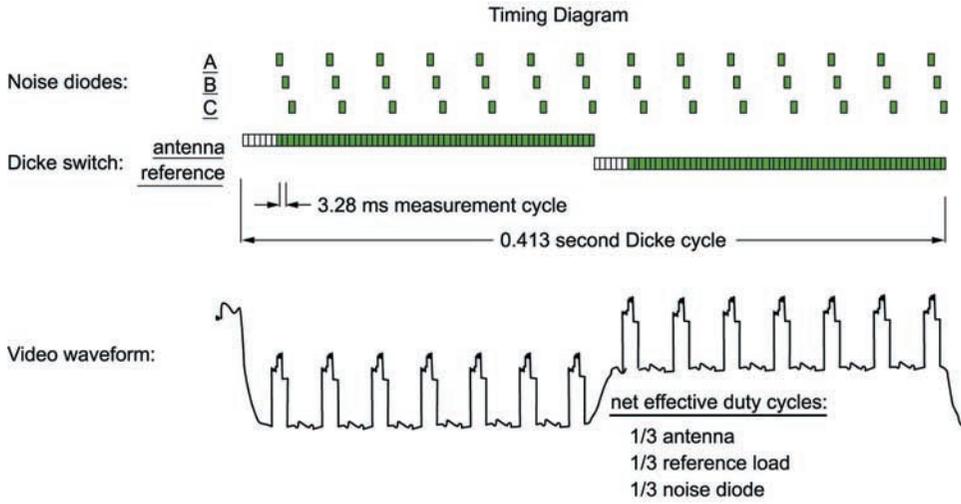
Timing Diagram

Noise diodes:
A
B
C

Dicke switch: antenna
reference

3.28 ms measurement cycle

0.413 second Dicke cycle

Video waveform:

net effective duty cycles:
1/3 antenna
1/3 reference load
1/3 noise diode

Figure 3. Timing diagram, showing the relative timing and duty cycles of the three noise diodes and Dicke switch.

A1 – A2 delta–T vs. integration time

solid = measured
dashed = 1/sqrt (integration time)

31 Ghz
24 Ghz
22 Ghz

RMS noise of A1–A2 difference (K)

integration time (seconds)
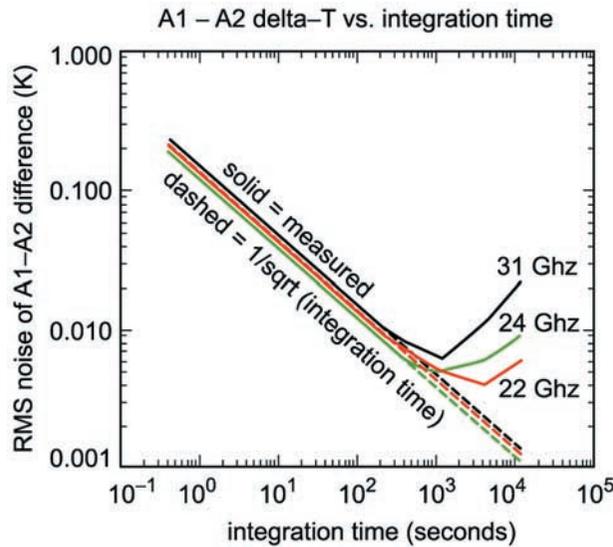
Figure 4. Theoretical (dashed) and measured (solid) delta-T from side-by-side comparisons of the two AWVRs.
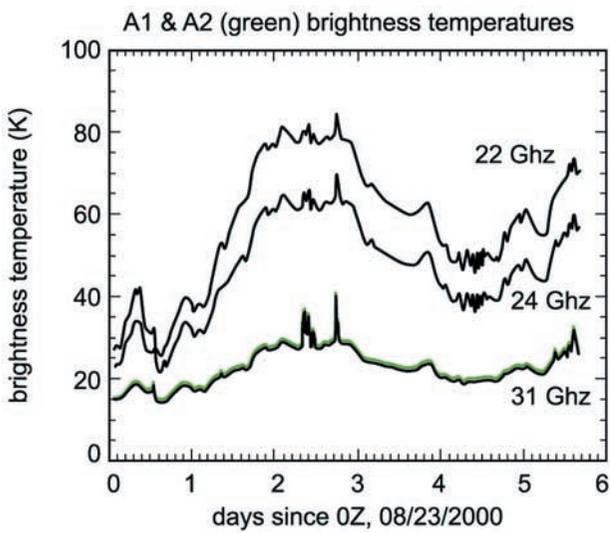
Figure 5. Brightness temperature overlay of the A1 radiometer (black) and A2 (green). Only the 31 GHz channels of A1 and A2 differ enough to discriminate between the two instruments.
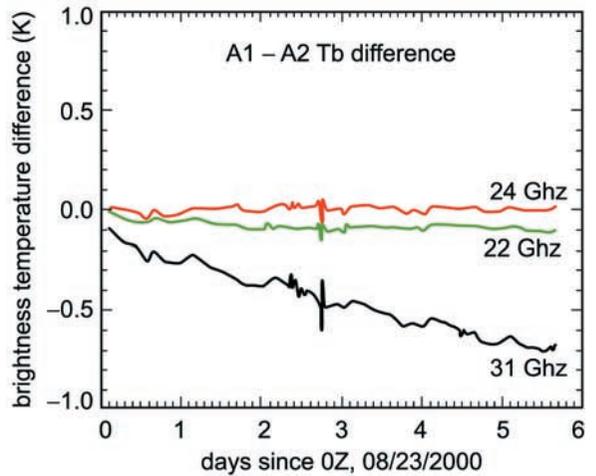
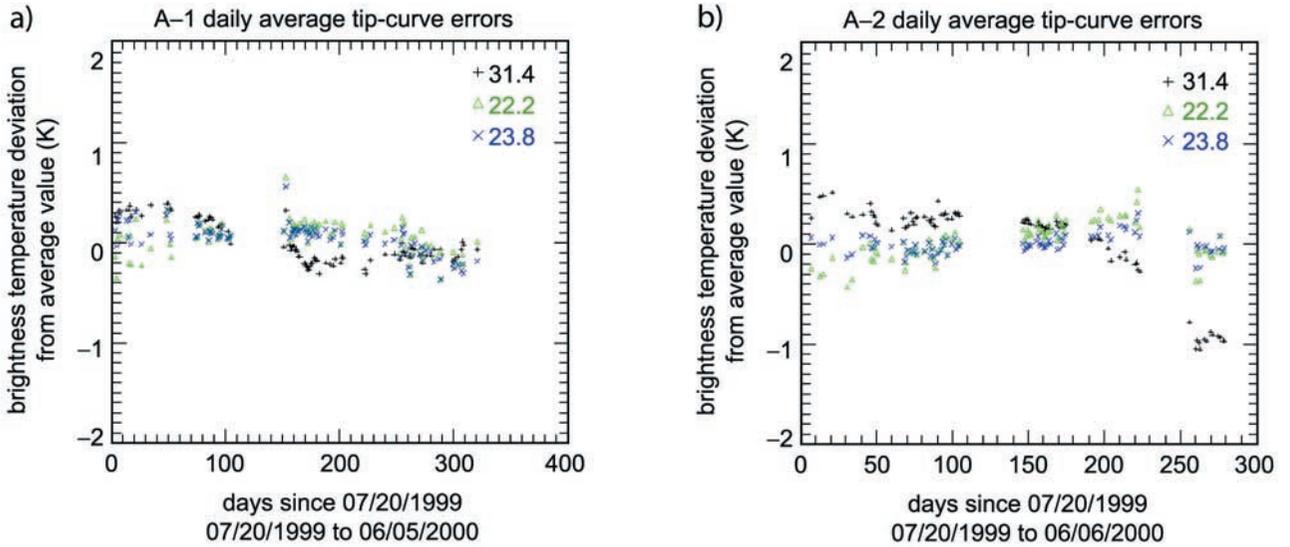**Figure 6.** Difference of A1 and A2 brightness temperatures from Figure 5. Integration time is 124 s.

Figure 9. Difference of T_B calibrated with a one-year average of all tip curve results, and those calibrated from one-day tip curve averages for A1 (a) and A2 (b). Positive error means that the long-term calibrated AWVR overestimated the brightness temperature on a given day.
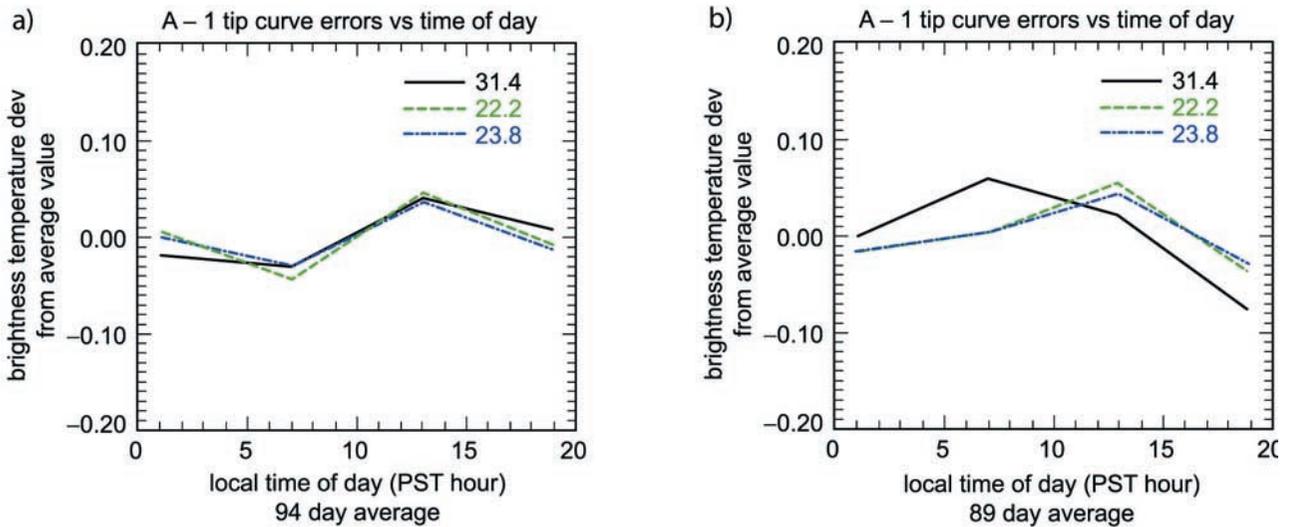

Figure 10. Average brightness temperature error versus time of day for A1 (a) and A2 (b) based on tip curve analysis.