Next Generation Very Large Array Memo #79

Synthesis of Optimal Estimators for Water Vapor Radiometry

Nathan Towne

June 17, 2020

Abstract

In ngVLA operation at higher frequencies, fast switching between the science target and a calibrator for delay and phase calibration would seriously reduce time on target. Because delay in atmospheric transit times is dominated by water vapor, water vapor radiometry (WVR) is to be used to circumvent this problem by inferring water vapor content, and hence delay, from sky brightness spectra at frequencies around the 22 GHz water vapor line. Because delay is computed from these spectra simply by integrating the spectra with a weight function, the question arises as to the weight function that most accurately and precisely estimates the delay. Such a synthesis is based on 1) how the atmosphere and the instrument map water vapor in the atmosphere to instrument spectra, and 2) the statistics of fluctuations in the atmosphere and confounding fluctuations in the instrument. The former is addressed by atmospheric modeling augmented by weather observations and instrument response measurements, while the latter by Kolmogorov's theory of scaling of fluctuations (augmented by atmospheric measurements) and bench and in-the-field measurements of instrument noise. With these constraints a 'best' weight function is computed. Once done, with the model in hand one can perform sensitivity analyses that constrain atmospheric and instrument parameters in such a way that requirements can be formulated. This memo outlines how most of these steps can be done.

1 Introduction

Water vapor in the atmosphere has a strong effect on the air's index of refraction and hence impacts radiation's time of flight. This effect is largely nondispersive at ngVLA frequencies and is quantified by the Smith-Weintraub equation (Sec. 3) — an integral of the water-vapor density along a line of sight. At the same time, the lossy part of the water-vapor and other components' dielectric constants absorb and emit radiation according to an ordinary differential equation along a path (Chandrasekhar 1960, Sec. 6, p. 9, Eq. 46), an equation that can also be derived through thermodynamic arguments. The frequency dependence of this absorption and emission, specifically at frequencies near water-vapor's 22 GHz resonance, results in radiation spectra at the Earth's surface

that reflects the amount and distribution of water vapor. Water vapor radiometry (WVR) for radio interferometry and remote sensing attempts to use this information to determine the total delay along a line of sight. The general idea of an estimator for delay is to compute a weight function that maps receiver spectral channels to an estimate of delay in a manner that minimizes the uncertainty of the estimate (Nikolic 2009). The statistical spread of delay estimates from a receiver and the estimator's weight function reflects the statistics of water-vapor fluctuations as well as the statistics of receiver noise. An optimized estimator must take into account these statistics, the physics of how water vapor along the line of sight results in spectra on the ground, and the receiver's response to these spectra. These fluctuations are systematically propagated as errors through the atmosphere and receiver response functions to the delay estimate. It is the knowledge of all these factors that allows calculation of a weight function that minimizes the spread of delay error.

This memo outlines how to go about quantifying the atmosphere and the instrument sufficiently to compute useful estimators. Some parts of the problem, such as the response and noise statistics of the instrument, are readily measured on the bench or in the field. The evolution of radiation along the line of sight is nearly so in the sense that spectra are easily computed given real or imagined temperature and composition profiles. More difficult is the choice of reference atmosphere – estimates of temperature and composition profiles about which fluctuations are thought to occur – which impacts how fluctuations in composition (and temperature) map to fluctuations of radiation spectra on the ground. The most difficult problem is the statistics of the atmospheric fluctuations themselves, and their height dependence, about which there has been a great deal of study. In this memo a few similar solutions based on (Carilli and Holdaway 1997) and their 90 minute measurement of the root phase structure function (Sec. 6.2) are constructed, which serve as placeholders for other more realistic and varied models.

A block diagram of the atmosphere/receiver/estimator model illustrating the information flow is shown in Fig. 1. In that diagram, the water-vapor profile refers to water-vapor density as a function of position along the line of sight. The Smith-Weintraub block refers to the water-vapor part of the Smith-Weintraub integral that computes the non-dispersive delay due to water vapor (Sec. 3). The atmospheric-response block is about how water-vapor fluctuations map to sky-brightness spectra as seen on the ground (Sec. 4.1). The channel-response block refers to the spectral response functions of the channels of the instrument, which are measured on the bench. Receiver noise is inserted at the '+' block on the right. And the weight function computes an estimate of the non-dispersive delay from the noisy intensities provided by the receiver's channels.

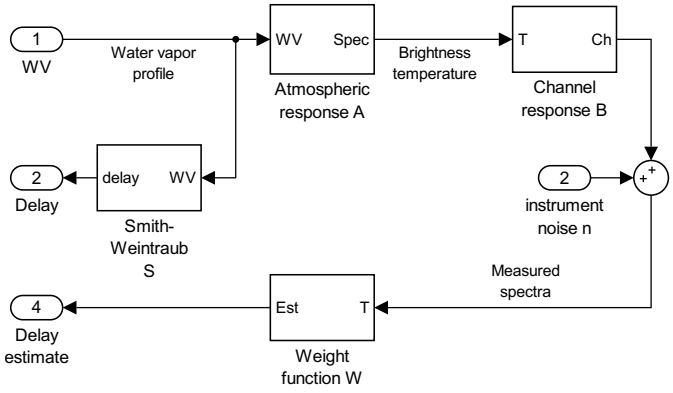

Figure 1: Diagram of the linearized atmosphere/radiometer/estimator system that is the basis of this memo.

The model is linearized in the following sense. In an interferometer, fast switching is used to periodically calibrate visibilities. In between calibrations, change in the atmosphere at each antenna leads to variation of temperature and water-vapor density along the line of sight, variations that are to be corrected with detection of incremental change of the sky brightness temperature as measured by the WVR receiver. These changes are small compared to the absolute sky brightness and a linearized model is adequate for inferring change of delay since the last source calibration. A linearized model allows more accurate formulation of the problem and simpler construction of the estimator.

A feature of the model shown in Fig. 1 is that the model atmosphere is a continuous function of position along the line of sight. There is no definition of layers of atmosphere, within which temperature, pressure and composition are constant. Perturbation of a layer means imposition of a change of infinite intensity and zero thickness, i.e., delta functions. As a linearized model the synthesis does not need complete knowledge of the distributions and correlations obeyed by the atmosphere and instrument. Instead, only covariance matrices are required, albeit functions of continuous variables. The covariance matrix obeyed by the receiver channels is directly measured on the bench and/or in situ. The covariance matrix obeyed by the water vapor is much more complicated, however, and is expected to vary widely depending on weather conditions. The root phase structure function measurement of (Carilli and Holdaway 1997) were well modeled by Kolmogorov turbulence (Kolmogorov 1991) and a few fitted parameters. That success provides hope that a general model of fluctuations and correlations can be usefully modeled by a few parameters that vary the model to current conditions.

The reference atmosphere provides a nominal composition and temperature along the line of sight (Sec. 6.1) about which variations of temperature and composition occur. It is an important part of the model because it impacts how emission by water vapor higher in the atmosphere is attenuated by water vapor lower in the atmosphere, which in turn emits its own radiation. There are a few tools available for constructing reference atmospheres, including surface conditions (temperature, pressure, and humidity), twice daily balloon soundings at five sites in Arizona and New Mexico ringing the EVLA site (conducted by the National Weather Service), parameterized standard atmospheres, and retrievals from satellite data. With a reference atmosphere one calculates in detail how emission by water vapor at every level along the line of sight affects sky brightness spectra on the ground (Sec. 4.1). The reference-atmosphere's temperature dependence also informs how water vapor becomes delay via the Smith-Weintraub equation.

The channel-response block of Fig. 1 contains the bench-measured frequency response functions of the instrument, one for each spectral bin (channel).

The weight function is one product of the synthesis, which provides delay corrections to the visibilities between source calibrations as described above. This synthesis is performed in post processing along with other calibration tasks. A second product is requirements traceability of technical specifications. Given any collection of specifications of the components in the receiver, the synthesis predicts the achievable precision of delay measurements, which, together with the delay precision requirement, constrains choices of technical requirements in a requirements flow down. Thus the synthesis directly maps design choices, such as the magnitude of receiver gain fluctuations and integration time, back to the system requirements. So these estimators are an important tool in design phases of projects. WVR for ngVLA is not an exception.

The Compact Water Vapor Radiometer (CWVR) Pathfinder project has developed four five-channel radiometers that were installed in EVLA antennas at the outputs of K-band front ends. The purpose of the project is the development of hardware and algorithmic experience applicable

to the ngVLA. Although the scope of this memo does not include the specifics of the CWVRs, this memo is part of the project, and the CWVRs serve as a test bed for ideas in this document. Thus some results from this memo are applied numerically to the CWVRs, particularly in Sec. 6 on sensitivity analysis where rough performance figures and parameter requirements are derived.

In the next section the core of the synthesis is discussed abstractly.

2 Estimator Synthesis in Operator Form

This section performs the synthesis of the estimator in a more generic operator form. Operators represent blocks of Fig. 1; what they do to physical quantities are described in detail in later sections. The reader should bear in mind that some operator products consist of integrals (such as BA), and other of finite sums (such as W acting on the receiver channels).

The Smith-Weintraub operator S and its integral realization ordinarily are expressed in terms of partial pressures, but can easily be converted to mass density via the ideal gas law. Let the transformation be represented generically by the operator S, and $\rho_{\rm wv}$ be the source (mass density) vector (or function) and d be the resulting delay. Then

$$d = S\rho_{\rm wv} \tag{1}$$

Let A be the operator representing atmospheric behavior, whose realization will be worked out in detail in Sec. 4. Let B be a vector of channel response functions mapping spectra (continuous functions of frequency) to channel intensities (before receiver noise is added). Let $C = BA$. The vector W contains the weights and v is the instrument noise. Then the estimator's output is

$$d_{\rm est} = W(C\rho_{\rm wv} + v) \tag{2}$$

The error in the estimated delay is simply $d - d_{\rm est}$, and the expectation value of the square error with respect to time, weather conditions, etc., is

$$\delta d^2 \equiv \langle (d - d_{\rm est})^2 \rangle \tag{3}$$

where the angle brackets ($\langle\rangle$) are temporal and/or spatial averages, depending on the context. Substituting Eq. (1) and Eq. (2) into Eq. (3),

$$\begin{aligned} \delta d^2 &= \langle (S\rho_{\rm wv} - W(C\rho_{\rm wv} + v))^T (S\rho_{\rm wv} - W(C\rho + v)) \rangle \\ &= \langle (\rho_{\rm wv}^T S^T - (\rho_{\rm wv}^T C^T + v^T) W^T)(S\rho_{\rm wv} - W(C\rho_{\rm wv} + v)) \rangle \\ &= {\rm Tr}(\mathbb{R}(S^T - C^T W^T)(S - WC)) + {\rm Tr}(\mathbb{V} W^T W) \end{aligned}$$

where the 'T' means transpose, 'Tr' is the trace operator, $\mathbb{R}$ is the covariance matrix of $\rho_{\rm wv}$,

$$\mathbb{R} = \langle \rho_{\rm wv} \rho_{\rm wv}^T \rangle \tag{4}$$

and $\mathbb{V}$ is the covariance matrix of instrument noise v.

$$\mathbb{V} = \langle v v^T \rangle \tag{5}$$

The result is

$$\delta d^2 = S\mathbb{R}S^T - 2S\mathbb{R}C^T W^T + W(C\mathbb{R}C^T + \mathbb{V})W^T \tag{6}$$

The trace operator is dropped because the terms are all scalars. The $S\mathbb{R}S^T$ term is the variance of the delay fluctuations to be tracked by the estimator. Recall that $\delta\rho_{\rm wv}$ refers to the change in water-vapor density since the last source calibration, and so has (hypothetically) zero mean everywhere along the line of sight. $\delta\rho_{\rm wv}$ and v are assumed uncorrelated and v also has zero mean.

So for every choice of weight function W, δd^2 is the statistical variance of the error of the delay estimate given by $d_{\rm est}$, given the statistics of $\rho_{\rm wv}$ ($\mathbb{R}$) and receiver noise ($\mathbb{V}$). To minimize this uncertainty, regard Eq. (6) as a variational problem where W is varied to find an extremum of δd^2, computed by differentiating Eq. (6) by W

$$0 = -2S\mathbb{R}C^T + 2W(C\mathbb{R}C^T + \mathbb{V}) \tag{7}$$

and then solving for the optimal weight function $W_{\rm opt}$

$$W_{\rm opt} = S\mathbb{R}C^T(C\mathbb{R}C^T + \mathbb{V})^{-1} \tag{8}$$

From the weight function the optimum value of the delay-error variance is computed by substituting Eq. (8) into Eq. (6).

$$\begin{aligned} \delta d^2_{\rm opt} &= S\mathbb{R}S^T - 2S\mathbb{R}C^T W^T_{\rm opt} + S\mathbb{R}C^T W^T_{\rm opt} \\ &= S\mathbb{R}S^T - S\mathbb{R}C^T W^T_{\rm opt} \end{aligned} \tag{9}$$

It can also be represented by

$$\delta d^2_{\rm opt} = S\mathbb{R}S^T - S\mathbb{R}C^T(C\mathbb{R}C^T + \mathbb{V})^{-1}C\mathbb{R}S^T \tag{10}$$

and

$$\delta d^2_{\rm opt} = S\mathbb{R}S^T - W_{\rm opt}(C\mathbb{R}C^T + \mathbb{V})W^T_{\rm opt} \tag{11}$$

Expressed in this operator form, these results are relatively simple and can be applied to a variety of problems. Each application has its own interpretation of the operators and vector spaces, and the user determines whether there is benefit in the use of the estimator. For context, two examples of the application of these concepts – one simple and one complex — are provided.

- The simple example is a noisy communications channel, a one-dimensional problem where noise is added to a signal: $s_n = s + n$. In the spectral domain we seek a filter $W_{\rm opt}$ that when applied to s_n best estimates s. In this problem the operators S and C above are the identity operator, and $\mathbb{R}$ and $\mathbb{V}$ are spectral densities of the signal s and noise n at any frequency, respectively. Applying Eq. (8) we have

  $$W_{\rm opt} = \frac{\mathbb{R}}{\mathbb{R} + \mathbb{V}} \tag{12}$$

  This is the classic matched filter, where the filter is matched to the statistics of both the signal and the noise. From this result we see that the filter of Eq. (8) can be regarded as a generalized matched filter, applicable to systems with many degrees of freedom and non-trivial statistics.

- The complex example is an alternate formulation of the drizzle algorithm (Fruchter and Hook 2001). The purpose of this algorithm is to resample at higher resolution collections of under-sampled images of a (usually astronomical) field of view. The algorithm estimates the intensity of each resampled pixel based on the intensities of scattered overlapping pixels among the original images — overlapping in the sense of the point spread function as well as pixel areas.

The identification of operators in Eq. (8) is interestingly straightforward. The operator S corresponds to an integral of the sky intensity over the resampled-pixel's area, which we regard as the "true" intensity of the resampled pixel, an estimate of which is to be the result of the synthesis. Corresponding to the operator A is the collection of point spread functions, each of which maps sky brightness to its image's focal plane via a (possibly slowly varying) convolution. Corresponding to operator B is the collection of pixelations, each of which is associated with how that image's image sensor maps focal-plane intensities to pixel intensities. v is pixel noise, a combination of shot noise and readout noise in digital sensors. $\mathbb{R} = \mathbb{R}(\Omega_1, \Omega_2)$ is the covariance matrix of the sky brightness between points Ω_1 and Ω_2 in the sky, which is modeled as homogeneous and isotropic, although it would ordinarily be customized to the stellar and extended-object statistics of the region in which the images reside. $\mathbb{V} = \langle v\, v^T \rangle$ is the covariance matrix of pixel noise v, which we expect is diagonal or nearly so. Finally, $W_{\rm opt}$ is the collection of weights applied to the pixels of the original images that optimally estimates the resampled-pixel's intensity. In this interpretation, $W_{\rm opt}$ applied to the original images estimate a single resampled pixel, and the resampled image is computed this way one pixel at a time.

The reader can see that the complexity of the synthesis itself when applied to more complicated problems is primarily in interpreting the operators of Eq. (8) and Eq. (9) appropriately, the calculation of the various operator products, and forging a practical algorithm from these ideas. The remainder of this memo explicates these points to the case of WVR.

3 Smith-Weintraub Equation

The Smith-Weintraub equation relates the index of refraction (non-dispersive part) to total pressure and water-vapor partial pressure (Smith and Weintraub 1953, Sec. 3).

$$n - 1 = \alpha \frac{k_{\rm B}\rho_{\rm T}}{m_d} + \left(\frac{\beta}{m_{\rm wv}} - \frac{\alpha}{m_d} \right) k_{\rm B}\rho_{\rm wv} + \gamma \frac{k_{\rm B}\rho_{\rm wv}}{m_{\rm wv} T_{\rm a}} \tag{13}$$

where $\rho_{\rm T}$ is air density, m_d is the mean molecular mass of dry air, $\rho_{\rm wv}$ is water-vapor density, $m_{\rm wv}$ is the molecular mass of water, $T_{\rm a}$ is air temperature, $k_{\rm B}$ is the Boltzmann constant, and

$$\alpha = 7.76 \times 10^{-7} {\rm K/Pa} \tag{14}$$
$$\beta = 6.48 \times 10^{-7} {\rm K/Pa} \tag{15}$$
$$\gamma = 3.776 \times 10^{-3} {\rm K^2/Pa} \tag{16}$$

The first term of Eq. (13) is due to air compression and rarefication due to convection, which is not observable with radiometry in K band. So that term is not considered in this memo. The remaining terms are linear in $\rho_{\rm wv}$ and so contribute to the linear Smith-Weintraub operator S. The realization of S mapping $\rho_{\rm wv}$ to delay is the integral

$$\begin{aligned} {\rm delay} &= S\rho_{\rm wv} \\ &= \int_{s_{\rm surf}}^{\infty} ds\, S_{\rm ker}(s)\, \rho_{\rm wv}(s) \qquad (17) \\ &= \int_{s_{\rm surf}}^{\infty} ds\, k_{\rm B} \left(\frac{\beta}{m_{\rm wv}} - \frac{\alpha}{m_d} + \frac{\gamma}{m_{\rm wv} T_{\rm a}(s)} \right) \rho_{\rm wv}(s) \qquad (18) \end{aligned}$$

where $S_{\rm ker}$ is the kernel of the operator, and referred to later.

In the spirit of the linearized model of the atmosphere/receiver/estimator system, S is applied to perturbations of $\rho_{\rm wv}$ of either sign, and not just the positive-only total water-vapor density.

4 Response to Water Vapor

There are two parts to how water vapor affects readings from a radiometer. The first is how water vapor impacts sky brightness spectra on the ground (Sec. 4.1 and Sec. 4.2), and the second is the spectral response of the receiver (Sec. 4.3).

4.1 Evolution of Sky Brightness Temperature

The radiation intensity, expressed as brightness-temperature T, is regarded as a smooth function of position s along the line of sight, as well as of frequency. Coefficients of absorption a and emission e are assumed the same, and reflection is zero. (These assumptions fail when there are larger water droplets.) Then the differential equation (Chandrasekhar 1960, Sec. 6, p. 9, Eq. 46) describing the evolution of the brightness temperature along the line of sight is

$$\frac{dT}{ds}(s,f) = e(s,f)\,(T_{\rm a}(s) - T(s,f)) \tag{19}$$

The initial condition can be

$$T(s=\infty,f) = T_{\rm CMB} \tag{20}$$

where $T_{\rm CMB}$ is the 2.7K cosmic microwave background. Eq. (20) can be modified if there is other significant (possibly frequency dependent) emission in the field.

This equation is formulated along a single direction (1-d) because, in the case of the VLA antennas, the beam in K band is narrow enough that at height there can be little decorrelation of atmospheric temperature and composition across the beam.

In this context, the emissivity function e takes into account the 22 GHz water-vapor line (Lorentzian part), the water vapor continuum excess (Sutton and Hueckstaedt 1996, Eq. 9), water ice and liquid water, and the tail of the 60 GHz O_2 line. The water vapor emissivity can be given by the Van Vleck-Weisskopf line shape (Tahmoush and Rogers 2000, Secn. 4), the HITRAN database (Gordon et al. 2017), or other source.

With a reference atmosphere thought to approximate temperature and composition along the line of sight at the time of a source calibration, and derived from available weather data, a reference emissivity $e_0(s,f)$ is derived, and a reference sky-brightness temperature $T_{\rm ref}(s,f)$ is integrated through Eq. (19).

$$\frac{dT_{\rm ref}}{ds} = -e_{\rm ref}\,(T_{\rm a} - T_{\rm ref}) \tag{21}$$

$$\delta T_{\rm ref}(s=\infty) = T_{\rm CMB} \tag{22}$$

where the s and f dependence has been suppressed. We define for later use the quantity

$$\Delta T = T_{\rm a} - T_{\rm ref} \tag{23}$$

Fig. 2 shows a solution for $T_{\rm ref}$ given a simple reference atmosphere.

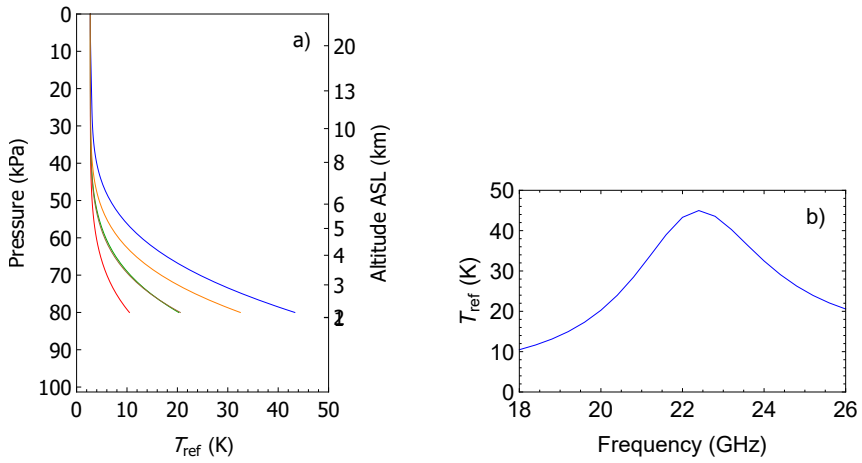

Figure 2: $T_{\rm ref}$ from a simple reference atmosphere with 2 km surface elevation, 300K surface temperature, 80 kPa surface pressure, and 50% surface relative humidity exponentially decreasing with altitude with a scale height of 4 km. Plot a) shows $T_{\rm ref}$'s pressure/altitude dependence at 18, 20, 22, 24, and 26 GHz. Traces at 20 and 26 GHz appear as one trace. Plot b) shows the power spectrum at ground level. The Van Vleck-Weisskopf line shape (Tahmoush and Rogers 2000, Sec. 4) plus the GR continuum excess (Sutton and Hueckstaedt 1996, Eq. 9) is used.

Pressure p instead of position along the line of sight s is often used as the independent variable in atmospheric calculations because the part of the atmosphere most active is at very low altitudes, while pressure spans a wide range at those low altitudes. Plots of quantities against pressure, for example, show much more clearly what is going on due to its stretching of the lower atmosphere than plots against altitude. This practice is followed here for the same reason.

Pressure at any altitude supports the weight of the air above that altitude. If h is altitude above sea level (ASL) and r is altitude from Earth's center, then the pressure p obeys

$$\frac{d}{dh}(4\pi r^2 p) = -4\pi r^2\,\rho g \tag{24}$$

where p is pressure, ρ (rho) is air density, and g is acceleration due to gravity. Using the ideal gas law, we relate ρ to pressure.

$$p = \rho\frac{k_{\rm B}T}{\overline{m}} \tag{25}$$

where $\overline{m}$ is the local mean molecular mass of the air and $k_{\rm B}$ is the Boltzmann constant. Eq. (24) becomes

$$\frac{dp}{dh} = -p\left(\frac{\overline{m}g}{k_{\rm B}T_{\rm a}} + \frac{2}{r}\right) \tag{26}$$

dp/ds is related to dp/dh through the zenith angle θ of the line of sight in the obvious way.

$$\frac{dp}{dh} = \frac{dp}{ds}\frac{1}{\cos\theta} \tag{27}$$

In terms of pressure, Eq. (21) becomes

$$\frac{dT_{\rm ref}}{dp} = \Delta T\, e_{\rm ref}\,\frac{1}{p}\,\frac{1}{\overline{m}g/k_{\rm B}T_{\rm a} + 2/r} \tag{28}$$

4.2 Incremental Brightness Temperature

As was discussed earlier, periodic source calibrations provide both a reference delay, and a sky-temperature reference of Eq. (21) for use when computing delay corrections. Spectra measured later have changed, and the incremental portions δT are available for computing incremental delay corrections. Corresponding to these measured δT are the hypothetical incremental water vapor $\delta\rho(s)$ along the line of sight responsible for δT. The path from $\delta\rho$ is through the water-vapor line model to the emissivity $\delta e = e - e_{\rm ref}$, through Eq. (19) to $\delta T = T - T_{\rm ref}$.

To streamline this process, we use the linearized form of Eq. (19) to compute the evolution of δT directly. Machine error is also limited by not taking differences of solutions. In Eq. (19), we make the substitutions $T \to T_{\rm ref} + \delta T$ and $e \to e_{\rm ref} + \delta e$. The terms to zeroth order cancel due to Eq. (21), and the second-order terms are dropped. This leaves the first-order terms

$$\frac{d\delta T}{dp} = (\Delta T\,\delta e + e_{\rm ref}\,\delta T)\,\frac{ds}{dp} \tag{29}$$

$$\delta T(p=0) = 0 \tag{30}$$

where ΔT is given by Eq. (23). Higher-order terms are dropped because of the assumption of linearity discussed in the introduction.

Using the integrating factor $\mu(p)$,

$$\log\mu(p) = -\int_{p_{\rm surf}}^{p} dp\; e_{\rm ref}\,\frac{ds}{dp} \tag{31}$$

the solution to Eq. (30) at $p_{\rm surf}$ is

$$\delta T(p_{\rm surf}) = \int_0^{p_{\rm surf}} dp\;\mu\,\Delta T\,\delta e\,\frac{ds}{dp} \tag{32}$$

In principle $\delta\rho$ can be any function, but the estimator needs a function mapping a local perturbation of ρ at a point p_0 to δT at the surface $p = p_{\rm surf}$. To do this take

$$\delta\rho_{\rm wv}(p) = \delta(p - p_0) \tag{33}$$

where δ on the right hand side is the Dirac delta function – a unit perturbation in this context. By integrating Eq. (32) across p_0,

$$A(p_{\rm surf}; p_0) = \Delta T\,\mu\,\frac{\partial\,\delta e}{\partial\rho_{\rm wv}}\,\frac{ds}{dp} \tag{34}$$

where the right hand side is evaluated at $p = p_0$.

Because the emissivity e is ordinarily expressed as a function of the water-vapor partial pressure $p_{\rm wv}$, there is utility in expressing $\partial e/\partial\rho_{\rm wv}$ as

$$\frac{\partial e}{\partial\rho_{\rm wv}} = \frac{\partial e}{\partial p_{\rm wv}} \times \frac{dp_{\rm wv}}{d\rho_{\rm wv}} \tag{35}$$

$$= \frac{\partial e}{\partial p_{\rm wv}} \times \frac{k_{\rm B}T_{\rm a}}{m_{\rm wv}} \tag{36}$$

where $m_{\rm wv}$ is the molecular mass of water. The ideal gas law was used. Combining Eq. (34) with Eq. (26), Eq. (27), and Eq. (36),

$$A(p_{\rm surf}; p) = -\Delta T\;\mu\;\frac{\partial e}{\partial p_{\rm wv}}\;\frac{k_{\rm B}T_{\rm a}}{m_{\rm wv}}\;\frac{1}{p}\;\frac{1}{\overline{m}g/k_{\rm B}T_{\rm a} + 2/r}\;\frac{1}{\cos\theta} \tag{37}$$

where the right hand side is evaluated at pressure p. A in Eq. (37) is the function that maps local perturbations of water vapor along the line of sight to incremental brightness spectra at the surface. A's dependence on $p_{\rm surf}$ is exclusively through the integration factor μ. Note also that the unit of δT is temperature per mass density.

The spectral intensity at ground level due to any incremental distribution $\delta\rho$ of water vapor along the line of sight is computed by the integral

$$\delta T(f) = \int_0^\infty dp\ A(p_{\rm surf}, p)\ \delta\rho(p) \tag{38}$$

4.3 Instrument Response

As was mentioned earlier, a detailed description of how the atmospheric temperature and composition maps channel intensities is required to synthesize an estimator. Half of this description, the behavior of the atmosphere, was developed in the previous section. The second half, operator B, is how sky spectra at ground level excite channel intensities, which, for the purpose of this memo, means integrals over frequency involving the frequency response functions of the channels of the instrument.

$$b_k = \int df\ \delta T(f) B_k(f) \tag{39}$$

where b_k is the kth channel's intensity, and B_k is the response function of that channel.

The net response of the atmosphere/instrument is the operator product C of A and B:

$$C_k(p) = \int df\ B_k(f)\, A(f; p_{\rm surf}, p) \tag{40}$$

where the frequency dependence of A is now made explicit.

In ngVLA, the receiver is expected to be digital with many (perhaps 256) channels. With this many channels, the channels are narrow, and the receiver's response function varies little over the bandwidth of an individual channel. An approximation of Eq. (40) that will no doubt be employed is

$$C_k(p) = B_k(f_k)\, \delta f\, A(f_k; p_{\rm surf}, p) \tag{41}$$

where f_k is the center frequency of the kth channel, and δf is an effective bandwidth of the channels, assuming they are all the same. Balancing aliasing against total bandwidth is set by maximizing performance through Eq. (8).

5 Evaluation of the $W_{\rm opt}$ and $\delta d^2_{\rm opt}$

Figure 3 is a diagram of computational flow of the estimator synthesis. The major steps towards the evaluation of Eq. (8) and Eq. (9) are the operator products $S\mathbb{R}S^T$, $S\mathbb{R}C^T$, and $C\mathbb{R}C^T$. The Smith-Weintraub operator S is an analytic function of pressure p. The covariance matrix $\mathbb{R}$ is similarly assumed an analytic function of two pressure variables via a useful atmospheric model. The atmospheric operator A is a numerical function of a continuous pressure variable and a discrete frequency variable as described in Sec. 4.2. So frequency is sampled across the band. The number of samples should be large enough that it samples the receiver's band-pass response function adequately. If the receiver has many narrow frequency channels, then this requirement may translate to a single sample per channel.

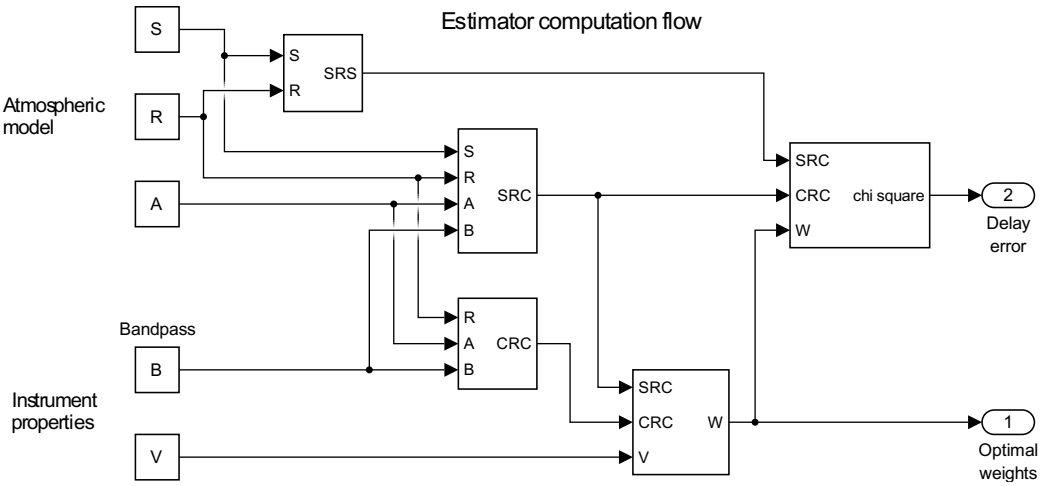

Figure 3: Flow of the computation of the optimal weights $W_{\rm opt}$ and delay error $\delta d^2_{\rm opt}$. As described above, the operators S, R, and A are the Smith-Weintraub operator, the turbulence model, and the atmospheric radiation model, respectively. These operators, along with $S\mathbb{R}S^T$ and the reference atmosphere constitute the atmospheric model. The operators B and $\mathbb{V}$ are the band-pass response functions and the noise covariance matrix of the receiver, respectively. The evaluation of the operator products in the three blocks is described in more detail in the text. The remaining two blocks are simply matrix evaluations corresponding to Eq. (8) and Eq. (9).

With these assumptions, the three operator products are directly evaluated. The first, $S\mathbb{R}S^T$, is a scalar evaluated as an integral over two path variables.

$$S\mathbb{R}S^T = \int ds_1\, ds_2\, S_{\rm ker}(s_1)\, \mathbb{R}(s_1, s_2)\, S_{\rm ker}(s_2) \tag{42}$$

where $S_{\rm ker}$ on the right hand side is the kernel of the Smith-Weintraub operator of Eq. (13). As was mentioned earlier, this term is the raw delay variance to be tracked by the estimator. The second, $S\mathbb{R}C^T$, is a one-dimensional array, each element of which is an integral of two variables and sum over frequency. The value for the lth receiver channel is

$$S\mathbb{R}C^T|_l = \Sigma_k \int ds_1\, ds_2\, S_{\rm ker}(s_1)\, \mathbb{R}(s_1, s_2)\, A(f_k; s_2)\, B_l(f_k) \tag{43}$$

The third, $C\mathbb{R}C^T$, is a two-dimensional array, each element of which is an integral of two variables and sum over two frequency indices.

$$C\mathbb{R}C^T|_{l_1\, l_2} = \Sigma_{k_1\, k_2} \int ds_1\, ds_2\, B_{l_1}(f_{k_1})\, A(f_{k_1}; s_1)\, \mathbb{R}(s_1, s_2)\, A(f_{k_2}; s_2)\, B_{l_2}(f_{k_2}) \tag{44}$$

Evaluation of these integrals is not trivial. Using Mathematica, some tuning of the accuracy and precision parameters of the numerical integrator is necessary to have practical computation times, and at the same time have sufficient precision following the numerical cancelation of the terms of Eq. (9). The reference atmosphere will in general not be analytic, but will instead be derived from piecewise continuous functions from standard temperature profiles or balloon soundings. Convergence of the integrals with such models is slower.

The remainder of the computation consists of finite-dimensional matrix arithmetic.

The operators $S\mathbb{R}A^T$ and $A\mathbb{R}A^T$ may alternatively be computed, the advantage being to separate atmospheric physics from receiver properties in B. The practical problem with this scheme is the computation time required when the number of frequencies sampled is large, particularly with $A\mathbb{R}A^T$ where an integral must be evaluated every frequency pair.

6 Atmospheric Model

The atmospheric model consists of the reference atmosphere and the turbulence model. The reference atmosphere is about the temperature and composition profiles as a function of altitude. Together these determine the frequency dependence of emission and absorption through any straight path to the surface, and consequently the spectrum of radiation at the surface and the operator A. The turbulence model describes fluctuations of the index of refraction and water-vapor concentration along the path, and correlations between points on the path and the operator $\mathbb{R}$. These quantities are integral to estimator synthesis as described in Sec. 2.

6.1 Reference Atmosphere

The reference atmosphere contains model (or actual) temperature and composition profiles of the air along the line of sight. It is important for multiple reasons. First, the emissivity function in general depends on temperature, total pressure, and water-vapor partial pressure. Second, these parameters and the emissivity function together describe how air lower in the atmosphere shadows emission from higher in the atmosphere. Finally, water vapor preferentially absorbs at the center of the emission line, further suppressing the narrower emission from higher in the atmosphere where there is less pressure broadening of the line. These details are necessary for computing the line shape of emission along the line of sight accurately to accurately account for spectra seen at ground level, which in turn allow accurate determination of total delay through the weight function.

Resources available for constructing a reference atmosphere tuned to current conditions are sparse balloon data, satellite data, weather forecasts, and on-site ground measurements. Balloon data are twice daily and spatially sparse. Temperature profiles among the five balloon launch sites in Arizona and New Mexico, track each other well above the tropopause, and to a lesser degree below. Figure 4 shows temperature and humidity profiles from five balloon soundings surrounding the VLA site. The fact that sky brightness spectra at ground level are not particularly sensitive to the temperature profile (Butler 1999), suggest that the temperature profiles can be modeled sufficiently well with balloon soundings or a standard temperature profile tapered to surface readings.

Accurately modeling water vapor is much more critical. There is also a degree of consistency of the water-vapor profiles among the five launch sites (Fig. 4 right), even during the monsoon. The example of Fig. 4 suggests modeling the water-vapor density as an exponential with scale height fitted to the average among the launch sites, and overall scale set to the water-vapor density at the surface. With such a choice, the differences among the traces would be interpreted simply as local fluctuations to be measured by the receiver and estimator through the incremental sky-brightness temperature measured by the receiver.

Given profiles of temperature and water-vapor composition, pressure as a function of altitude is computed from Eq. (26) with the initial condition being surface pressure.

This cursory glance must be further interpolated to other times of day, and its applicability to other times with varying weather conditions must be established. Satellite retrievals with their

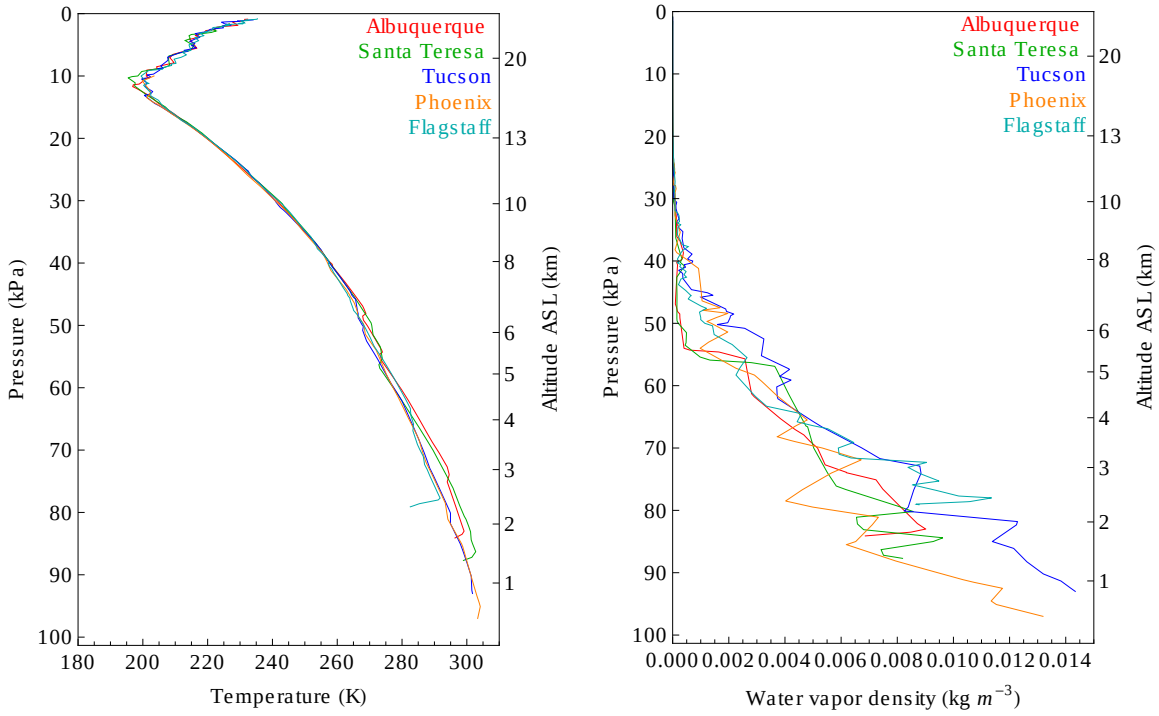

Figure 4: Temperature (left) and humidity (right) of the relatively humid atmosphere of 2018-Jul-22 at noon, from five balloon soundings in Arizona and New Mexico.

higher spatial and temporal resolutions may allow construction of better reference atmospheres, although they are less accurate than balloon data.

Assessing the sensitivity of the delay error variance to errors in the reference atmosphere is discussed in Sec. 7.1.

6.2 Turbulence Model $\mathbb{R}$

Is section Sec. 3 was discussed the Smith-Weintraub operator that maps variations of the water vapor along the line of sight to delay, which happens through the variations δn of the index of refraction n of the air. Delay (times the speed of light) is the integral of these variations along the line of sight of the antenna. Two antennas constituting a baseline see different delays, but delays that are correlated due to the correlations of δn between points along the two lines of sight, and within each line of sight. The root-mean-square difference of the delays as a function of baseline is the root-phase-structure-function $r(b)$ (RPSF). Measurements of the RPSF are accounted for by a model that assumes a layer of the atmosphere of thickness m that is uniformly turbulent, and Kolmogorov scaling of the fluctuations within that layer, valid over a wide range of spatial separations,

$$\langle(\delta n(s_1) - \delta n(s_2))^2\rangle \sim D_n^2 |s_1 - s_2|^{2/3} \tag{45}$$

where D_n is a constant and s_1 and s_2 are points within the layer. These assumptions result in an analytic form for the RPSF

$$\begin{aligned} r(b)^2 &= \langle(\delta l(x) - \delta l(x+b))^2\rangle \\ &= \frac{1}{20}D_n^2\left(15b^{8/3} - 9m^{8/3} - (15b^2 - 9m^2)(b^2+m^2)^{1/3} + 16b^{2/3}m^2\,{}_2F_1\left(\frac{1}{2},\frac{2}{3};\frac{3}{2};-\frac{m^2}{b^2}\right)\right) \end{aligned} \quad (46)$$

where b is baseline length and ${}_2F_1$ is a hypergeometric function. This formula is valid for baselines smaller than a value at which there is seen to be decorrelation. In this regime, the value of D_n is determined from the fit of the data to Eq. (46). In the regime in which there is decorrelation, the value of the RPSF is due to integrated fluctuations along the two lines of sight added in quadrature, which is evaluated from Eq. (46) applied to a single line of sight.

$$r_{\text{decorr}}^2 = 2m^2\left(\langle\delta n^2\rangle + \frac{9}{40}m^{2/3}D_n^2\right) \quad (47)$$

Knowing D_n and r_{decorr}, the third unknown δn is determined from Eq. (47).

Estimator synthesis needs to know the water-vapor covariance matrix $\mathbb{R}$

$$\mathbb{R}(s_1, s_2) = \langle\delta\rho_{\text{wv}}(s_1)\delta\rho_{\text{wv}}(s_2)\rangle \quad (48)$$

δn is connected to water-vapor fluctuations $\delta\rho$ through the kernel of the Smith-Weintraub equation Eq. (65).

$$\delta n = S_{\text{ker}}\delta\rho_{\text{wv}} \quad (49)$$

Here we make the added assumption that S_{ker} is constant in the turbulent layer. Knowing $\langle\delta n^2\rangle$, and employing the identity

$$\langle\delta n(s_1)\delta n(s_2)\rangle = \langle\delta n^2\rangle - \frac{1}{2}\langle(\delta n(s_1) - \delta n(s_2))^2\rangle \quad (50)$$

we have that

$$\begin{aligned} \mathbb{R}(s_1, s_2) &= \frac{1}{S_{\text{ker}}(s_1)}\langle\delta n(s_1)\delta n(s_2)\rangle\frac{1}{S_{\text{ker}}(s_2)} \\ &= \frac{1}{S_{\text{ker}}(s_1)}\left(\langle\delta n^2\rangle - \frac{1}{2}\langle(\delta n(s_1) - \delta n(s_2))^2\rangle\right)\frac{1}{S_{\text{ker}}(s_2)} \\ &= \frac{1}{S_{\text{ker}}(s_1)}\left(\langle\delta n^2\rangle - \frac{1}{2}D_n^2|s_1 - s_2|^{2/3}\right)\frac{1}{S_{\text{ker}}(s_2)} \end{aligned} \quad (51)$$

The measurement of $r(b)$ in (Carilli and Holdaway 1997) provides us with values of D_n and $\langle\delta n^2\rangle$ given in Fig. 5.

The assumptions in this derivation are all not universally valid. Kolmogorov scaling often fails; there can be multiple boundary layers at different levels, or none at times of the day; S_{ker} is not constant and is a poor approximation for thick turbulent layers; turbulence is not constant in the turbulent layer, and is influenced by high-level winds; and water-vapor density typically drops rapidly with altitude. Studies of horizontal communications links show that Kolmogorov or non-Kolmogorov scaling varies with altitude, with elevated values at an inversion and near the ground. As was mentioned, the turbulent peak near the inversion is influenced by winds at higher altitudes. In this picture, correlations between points along slanted paths are particularly poorly constrained. Simulations of turbulence such as in (Stirling et al. 2008) shed some light on this problem.

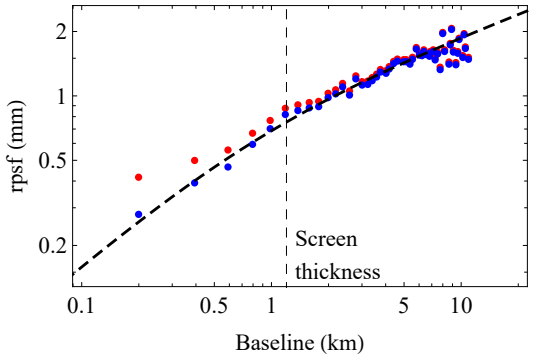

Figure 5: Fit of model root phase structure function Eq. (46) (dashed line) to the data of (Carilli and Holdaway 1997), in a unit of distance. Red dots are the raw data points and blue dots have a receiver-noise term subtracted in quadrature. The fitted parameters are screen thickness $m = 1.2$ km, receiver noise $= 0.031$ mm, the structure parameter $D_n = 0.77 \times 10^{-8}$ mm$^{-1/3}$, and $\sqrt{\langle \delta n^2 \rangle} = 0.9$ ppm.

So with this simplified model and the success of Fig. 5, we can calculate $\mathbb{R}$ through Eq. (47), Eq. (45), Eq. (49), and Eq. (48). This solution provides a model of $\mathbb{R}$ that reproduces a measured root phase structure function and can be applied to estimator synthesis. While it does not capture the full complexity of atmospheric turbulence, it is a useful starting point for design studies constraining receiver requirements, if not for calibration of visibilities. A better grasp of atmospheric turbulence is necessary in the longer term.

7 Sensitivity Analysis

There are inevitable errors in the atmospheric model and other parameters that are part of the synthesis of an estimator, and it is essential to have an idea to what degree those errors affect the delay corrections applied to a measurement set. One way to address this problem is to compare the tracking precision without the error using Eq. (9) with the tracking precision computed from Eq. (6) when the weight function contains the error (Eq. (6)). This allows us to get a sense of how much inaccuracy can be tolerated. This is what is done in Sec. 7.1, where the atmospheric model is varied and the tracking precision is compared.

When parameters enter into the synthesis analytically, such as the parameter γ (Eq. (16)) in the Smith-Weintraub operator, then a Taylor expansion of the tracking precision with respect to those parameters about their nominal values can be computed. The zeroth-order term of that expansion is δd^2_{opt}. Because of the structure of Eq. (10), the first-order term is zero, and the next higher term is the second-order term that is a quadratic form in the parameters' deviations about their nominal values. A concave quadratic form will degrade performance, which is expected of model inaccuracies, a convex quadratic form will do the opposite, and a saddle-shaped quadratic form does both. Inaccuracies of T_{cal}, gain stability, and the Smith-Weintraub operator are analyzed with this method in sections 7.2, 7.3, and 7.4, respectively. Although not done here, one can comprehensively compute a single quadratic form in all the parameters considered as a means to identify interaction among those parameters, and not just among closely related parameters. Spectral decomposition of sensitivity matrices identifies the most sensitive parameters by the eigenvalues and eigenvectors, and how the parameters interact through each eigenvector. Parameters that have little impact

on performance can be identified, allowing simplification of the estimator model and/or physical receiver.

Numerical results are given below, where the variations are compared with a baseline atmospheric model given in Sec. A.1, and with characteristics of one of the CWVRs installed in the EVLA, given in Sec. A.2.

7.1 Errors in the atmospheric model

Two alternate atmospheric models were devised to illustrate testing the sensitivity of the synthesized estimators to variations of the atmosphere. One model increases the surface temperature, tapering the temperature to a standard profile at 9 km AGL, and the second that increases the surface relative humidity with the same scale height. Variations of the turbulence model can also be explored, but for brevity that is not done here.

Fig. 6 shows the results of the comparisons. The three rows and three columns correspond to the three models, where the optimal weight function of the column is inserted into the model of the row. So the blue traces show the optimal performance of the model of the row (the blue traces in the rows are the same across columns). There is no comparison in the graphs on the diagonal. The red traces in the off-diagonal graphs show the non-optimal performance.

Note that, like atmospheric modeling, assessment of sensitivities is a large problem, one that can only be touched upon here. But this assessment does serve as the start of a framework for assessing how errors in an atmospheric model degrade estimator performance.

7.2 Errors in T_{cal}

Errors in T_{cal} are in effect gain errors among the receiver's channels that affect the magnitude of T_{sys} returned by the receiver. They can also be regarded as errors in the channel response functions B. You can imagine that if T_{cal} is the modeled T_{cal} and T'_{cal} is the actual, we have a dimensionless diagonal error matrix

$$\delta_{T_{\mathrm{cal}}} = \mathrm{diag}(T'_{\mathrm{cal}}/T_{\mathrm{cal}}) \tag{52}$$

where 'diag' means the diagonal matrix formed from the array $T'_{\mathrm{cal}}/T_{\mathrm{cal}}$ (one T_{cal} per receiver channel). This error matrix affects the error-less C (or B) by

$$C' = \delta_{T_{\mathrm{cal}}} C \tag{53}$$

Because the receiver noise is measured in terms of noise temperature, which has its root in T_{cal}, $\mathbb{V}$ is also affected by δT_{cal}.

$$\mathbb{V}' = \delta_{T_{\mathrm{cal}}} \mathbb{V} \, \delta_{T_{\mathrm{cal}}} \tag{54}$$

So the tracking precision when there is an error in T_{cal} changes Eq. (6) to

$$\delta d^2 = S\mathbb{R}S^T - 2S\mathbb{R}C^T\delta_{T_{\mathrm{cal}}} W_{\mathrm{opt}}^T + W_{\mathrm{opt}}(\delta_{T_{\mathrm{cal}}} C\mathbb{R}C^T\delta_{T_{\mathrm{cal}}} + \delta_{T_{\mathrm{cal}}} \mathbb{V} \, \delta_{T_{\mathrm{cal}}})W_{\mathrm{opt}}^T \tag{55}$$

But this is equivalent to

$$\delta d^2 = S\mathbb{R}S^T - 2S\mathbb{R}C^T W_{\mathrm{opt}}^{T\prime} + W'_{\mathrm{opt}}(C\mathbb{R}C^T + \mathbb{V})W_{\mathrm{opt}}^{T\prime} \tag{56}$$

where we have an altered weight function

$$W'_{\mathrm{opt}} = W_{\mathrm{opt}} \, \delta_{T_{\mathrm{cal}}} \tag{57}$$

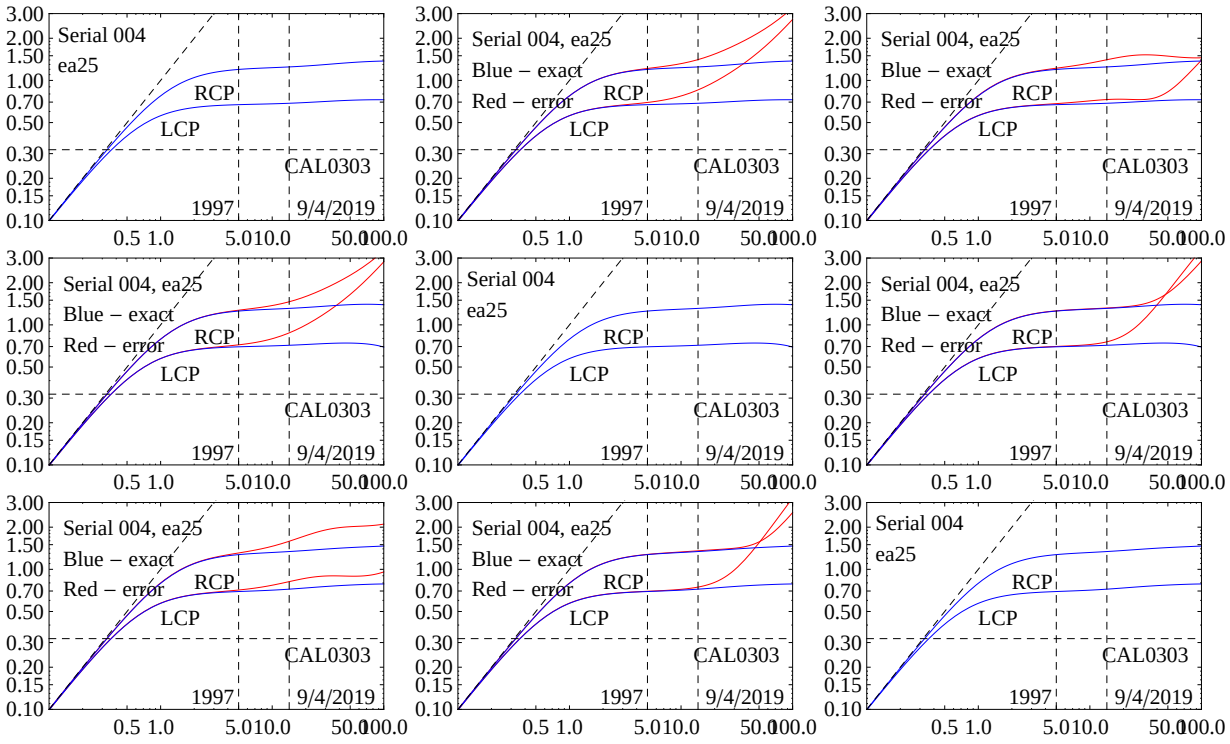

Figure 6: Comparison of optimal estimator performance with performance using a non-optimal weight function computed from an alternate atmospheric model. Rows correspond to the 'true' atmospheric models, while the columns correspond to the 'in-error' atmospheric model from which the weight function is computed. Blue traces show the optimal performance (optimal weight function for the model), while the red traces show the performance with the non-optimal weight function. The baseline atmospheric model of Sec. A.1 corresponds to the first row/column, the second row/column varies the surface temperature (300K instead of 280K), and the third row/column varies the surface relative humidity (75% instead of 50%). Although not labeled, the horizontal and vertical axes are turbulence intensity and tracking precision, respectively, both in ps. The annotations are the same as in Fig. 12.

Because the altered weight function $W'_{\rm opt}$ is not the optimal weight function for $S\mathbb{R}S^T$, $S\mathbb{R}C^T$, $C\mathbb{R}C^T$ and $\mathbb{V}$, this weight function degrades the performance of the estimator, i.e., $\delta d^2 > \delta d^2_{\rm opt}$. So Eq. (56) provides a route to a requirement on the accuracy of $T_{\rm cal}$ by limiting the impact of errors on δd^2.

A more direct requirement on the errors in $T_{\rm cal}$ comes from noting that Eq. (56) is $\delta d^2_{\rm opt}$ plus a quadratic form in the fractional errors in the $T_{\rm cal}$s. So let the array x be the fractional error

$$x = \frac{T'_{\rm cal}}{T_{\rm cal}} - 1 \tag{58}$$

Then

$$\delta d^2 = \delta d^2_{\rm opt} + x\, Q\, x^T \tag{59}$$

where Q is a symmetric matrix. It is readily evaluated from $S\mathbb{R}C^T$, $C\mathbb{R}C^T$, and $\mathbb{V}$. In the case of the atmospheric model of Fig. 12 and LCP, it is

$$Q = \begin{pmatrix} 0.00073 & -0.0372 & 0.00143 & 0.119 & -0.0375 \\ -0.0372 & 2.06 & -0.0809 & -6.79 & 2.09 \\ 0.00143 & -0.0809 & 0.00324 & 0.274 & -0.0834 \\ 0.119 & -6.79 & 0.274 & 23.4 & -7.08 \\ -0.0375 & 2.09 & -0.0834 & -7.08 & 2.19 \end{pmatrix} \text{mm}^2 \tag{60}$$

By inspection, the fourth channel requires the tightest tolerance on its fractional $T_{\rm cal}$, followed by the second and fifth. In the case of the fourth channel, the inaccuracy at which the quadratic-form term equals the CAL0303 (Hales 2019, Req. CAL0303) requirement is $x_4 = 2\%$. But note a few points about this simple calculation.

- $\delta d^2_{\rm opt}$ currently meets requirements in neither EVLA nor ngVLA. Because both $\delta d^2_{\rm opt}$ and Q depend on the receiver noise figures in $\mathbb{V}$, calculation of tolerances on $T_{\rm cal}$ accuracies cannot be definitively done before the final $\mathbb{V}$ is known and shown capable of meeting the requirement.
- If the impact of the $T_{\rm cal}$ inaccuracy were to be limited in an error budget to, say, 50% of a requirement (added in quadrature), then the accuracy requirement would be further tightened.
- The off-diagonal elements can make such estimates inaccurate. For example the second and fifth channels have matrix elements all roughly 2, implying that the contribution from those two channels provides a term $2\,\text{mm}^2 \times (x_2 + x_4)^2$, significantly different than calculating from the diagonal elements of Q separately. These observations can be formulated comprehensively in terms of a spectral decomposition of Q – one term per eigenvalue.

  In fact, Q has a dominant eigenvalue, which means that the quadratic form in the errors x can be closely approximated by a single term consisting of the square of a weighted sum of the elements of x, where the weight function is the dominant eigenvector times the square root of its eigenvalue. This is a route to a simpler approximate statement of a tolerance on the errors, one that quickly and reliably identifies the errors that matter most.

7.3 Gain Stability and Receiver Noise Temperature

As previously noted, the estimator maps changes in the receiver channel intensities between source calibrations to line-of-sight delay changes, which are written to the observation's calibration data.

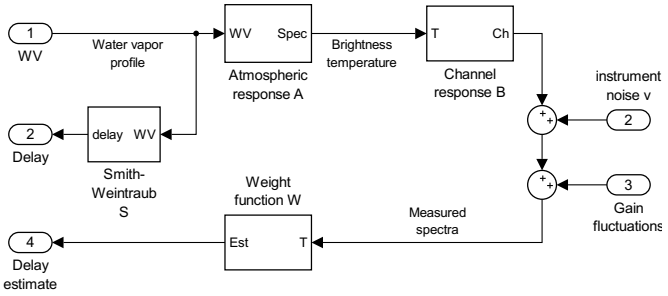

Figure 7: Diagram of the linearized atmosphere/radiometer/estimator system with gain fluctuations δg added to the receiver noise.

When there is gain drift in the receiver, drift of the intensity of the larger non-fluctuating component of $T_{\rm sys}$ due to gain drift cannot be distinguished from intensity fluctuations of $T_{\rm sys}$, for the most part. So it is essential that gain fluctuations be limited by the delay precision requirement.

Gain fluctuations can be incorporated into Sec. 2 by adding them to Fig. 1, as shown in Fig. 7. In that figure, δg is an array of fractional gain fluctuations, one for each receiver channel, each of which is scaled by its value of $T_{\rm sys}$. So for the purpose of the remainder of this section, regard $T_{\rm sys}$ as a matrix whose diagonal consists of the receiver's $T_{\rm sys}$. Then the receiver noise including gain fluctuations is

$$v' = v + T_{\rm sys}\,\mathbb{G} \tag{61}$$

where v' is the receiver noise with gain fluctuations. The covariance matrix $\mathbb{V}'$ is

$$\mathbb{V}' = \mathbb{V} + T_{\rm sys}\,\mathbb{G}\,T_{\rm sys} + \text{cross-correlation terms} \tag{62}$$

where

$$\mathbb{G} = \langle \delta g\, \delta g^T \rangle \tag{63}$$

It is reasonably assumed that δg is not correlated with $\rho_{\rm wv}$, but δg and v' likely have cross correlation. With the definition of $\mathbb{V}'$ in Eq. (62), the equations of Sec. 2 hold. But now the impact of measurements or models of $\mathbb{G}$ on tracking precision can be assessed by computing $\delta d^2_{\rm opt}$ with and without $\mathbb{G}$. The simplest plausible model is that the fractional gain fluctuations are independent of frequency and channel. Then for the CWVRs

$$\mathbb{G} = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 \end{pmatrix} \langle \delta g^2 \rangle \tag{64}$$

where δg is now regarded as a scalar gain fluctuation. With the receiver noise temperature equal to 150K for ngVLA WVR, the gain-fluctuation requirement of the receiver suggested by the figure is of order 0.01%. A cold receiver would reduce the slopes of the traces.

The interaction between gain stability and receiver temperature is illustrated in Fig. 9. In that figure a model sky temperature and very low receiver noise is assumed.

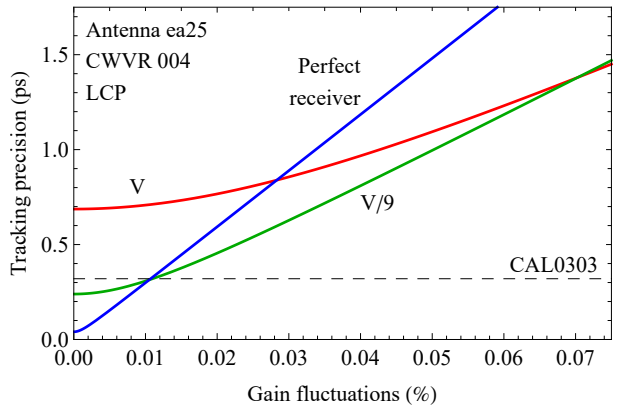

Figure 8: Tracking precision with varying intensity of gain fluctuations. The red and green traces are computed with the LCP $\mathbb{V}$ scaled by the indicated amounts. The atmospheric model is that of Fig. 12, and a model of $T_{\rm sys}$ that peaks at 200K for a room-temperature receiver is assumed. The blue trace assumes the noise of a perfect receiver given an integration time of 0.85 s, bandwidths of 1.5 and 1.0 GHz among the channels, and a receiver noise temperature of 150K.

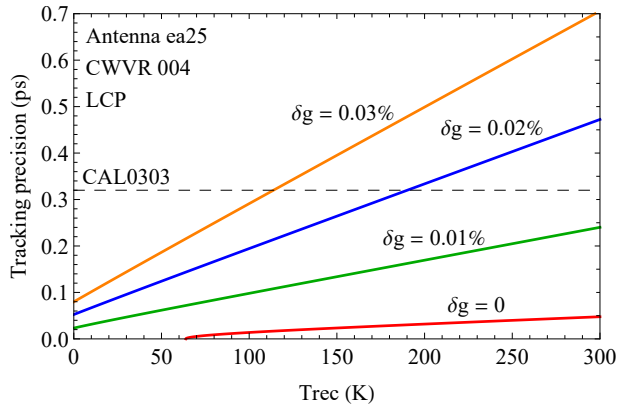

Figure 9: Tracking precison as a function of receiver noise temperature for four values of gain stability. Receiver response functions are taken from CWVR 4, and noise characteristics of an ideal receiver given the bandwidths and two-second integration time are entered into $\mathbb{V}$. A sky temperature peaking at 50K is also assumed.

7.4 Inaccuracy in the Smith-Weintraub Operator S

The Smith-Weintraub operator of Eq. (18) is linearly parameterized by two parameters, $\beta/m_{\rm wv} - \alpha/m_{\rm d}$ and $\gamma/m_{\rm wv}$, which are only approximately known. Errors in these parameters affect tracking precision by computing the weight function $W_{\rm opt}$ from Eq. (8) using an operator S' not equal to the actual operator S. This section computes the tracking precision that results from this error. The method used is to insert the non-optimal $W'_{\rm opt}$ computed with S' into Eq. (6), resulting in a tracking precision δd^2 degraded by the error. Because S is linear in the two parameters, δd^2 is the optimal δd^2 plus a quadratic form in the two parameters $\beta/m_{\rm wv} - \alpha/m_{\rm d}$ and $\gamma/m_{\rm wv}$. Like in Sec. 7.3, we will compute the (2x2) kernel of the quadratic form and extract accuracy requirements on the parameters.

Let $\delta S = S' - S$ so that $S' = S + \delta S$. Similarly, from Eq. (8) let $W'_{\rm opt} = W_{\rm opt} + \delta W$. From Eq. (6) we have

$$\begin{aligned}
\delta d^2 &= S\mathbb{R}S^T - 2S\mathbb{R}C^T W^{T\prime} + W'(C\mathbb{R}C^T + \mathbb{V})W^{T\prime} \\
&= S\mathbb{R}S^T - 2S\mathbb{R}C^T W^{T\prime} + S'\,\mathbb{R}\,C\,W^{T\prime} \\
&= S\mathbb{R}S^T - 2S\mathbb{R}C^T (W^T + \delta W^T) + (S + \delta S)\,\mathbb{R}\,C\,(W^T + \delta W^T)
\end{aligned}$$

With cancellations, the result

$$\delta d^2 = \delta d^2_{\rm opt} + \delta S\,\mathbb{R}C(C\mathbb{R}C^T + \mathbb{V})C\mathbb{R}\,\delta S^T \tag{65}$$

Given a parameterization of $\delta S(x)$ where S is linear in x, $x = (x_1, x_2, ...)$, and that $\delta S(x = 0) = 0$, then the second term of Eq. (65) is a quadratic form in x, where the kernel Q has matrix elements

$$Q_{ij} = \partial_i(\delta S)\,\mathbb{R}C^T(C\mathbb{R}C^T + \mathbb{V})C\mathbb{R}\,\partial_j(\delta S^T) \tag{66}$$

and

$$\delta d^2 = \delta d^2_{\rm opt} + x\,Q\,x^T \tag{67}$$

Then applying this result to WVR through Eq. (43),

$$\partial_i(\delta S)\,\mathbb{R}C|_l = \Sigma_k \int ds_1\,ds_2\,\partial_i(S_{\rm ker})(s_1)\,\mathbb{R}(s_1, s_2)\,A(f_k; s_2)\,B_l(f_k) \tag{68}$$

Because δS depends on two parameters, Q is of dimension two. In the case of CWVR, the atmospheric model of Fig. 12 and 14.3 ps delay fluctuations (the intensity of the observation of Sec. A.3) result in the following kernel for fractional error:

$$Q = \begin{pmatrix} 0.0026 & 0.22 \\ 0.22 & 18 \end{pmatrix} \text{mm}^2 \tag{69}$$

One can perform the spectral decomposition mentioned earlier, but by inspection it is seen that the error is dominated by the second (x_γ corresponding to Eq. (16)) parameter. Because the CAL0303 requirement for ngVLA is 0.32 ps, which corresponds to 0.1 mm, the accuracy requirement is 0.6%, which is about the accuracy that is given in (Smith and Weintraub 1953) for their 3.75×10^{-3} value. A slightly different value is given in (Stirling et al. 2008), and it is not clear as to its source and uncertainty.

$$x_\gamma \leq 0.1\text{mm}/\sqrt{18\,\text{mm}^2} = 2.3\% \tag{70}$$

This requirement is not as stringent as the accuracy to which γ is known. The requirement for the accuracy of the first parameter, $\beta/m_{\rm wv} - \alpha/m_{\rm d}$, is approximately 100%, reflecting the fact that $S_{\rm ker}$ is dominated by the γ term.

8 Discussion

This section discusses several other issues.

8.1 The Roll of Estimators in Receiver Design

This section emphasizes the importance of estimators in the design of receivers, beyond what was discussed in Sec. 6. When calibrating science data, the weight function Eq. (8) is the primary result of this memo. It optimally maps receiver data to delay in the sense that it minimized the uncertainty of the delay estimate. In this way visibilities can be better calibrated between source calibrations, allowing longer intervals between source calibrations. Source calibrations are still needed, but with more frequent delay corrections provided by Eq. (8), phase errors remain flat for longer intervals allowing more time on target.

In the design phase of these WVR projects the objective is instead to map broader delay precision and accuracy requirements to more detailed technical requirements. There, only the delay variance Eq. (9) is (directly) relevant. An optimal estimator approach allows rigor in this process when taking into account water vapor fluctuations and correlations along the line of sight, given a parameterized atmospheric model. The optimal estimator allows comparisons for varying atmospheric characteristics, such as boundary-layer thickness, water-vapor fluctuation intensities ($\langle \delta\rho^2 \rangle$); and reference-atmosphere characteristics, such as surface conditions, temperature, humidity, and water/water ice profiles. The formalism of Sec. 2 can further be used to propagate errors in the reference atmosphere to the expectation value of delay-error variance and explore its parameter space. It is true that atmospheric properties are highly uncertain and variable, and estimator synthesis for calibration can potentially be a case of garbage in-garbage out. Even so, real-world data are available, albeit for very limited range of conditions, and Eq. (9) provides a route in the design phase to employ parameterized atmospheres to explore how the parameter space impacts the precision with which delay can be estimated.

Instrumental variation can also be efficiently explored to see how it affects delay precision. Bench-measured intensities and correlations of receiver noise are an integral part of the synthesis (the radiometer equation is not assumed). The synthesis allows comparison of varying channel count, a smaller number of analog channels vs many Fourier channels in logic, varying total bandwidth, and schemes where non-contiguous segments of bandwidth are individually spectrally analyzed. Channel response functions can be varied, and the superior statistics of overlapping channel response functions can be assessed and exploited as is done with image sensors (Wikipedia 2019). The tradeoff between system noise temperature and integration time can be assessed, allowing for decisions about system cooling. Gaps between and overlaps of channels, and missing or noisy channels are automatically accounted for optimally through B and $\mathbb{V}$. The performance of non-optimal weight functions can be compared through Eq. (6). Sensitivity analysis can further constrain instrument properties, such as how system temperature impacts the gain-stability requirement. And finally, WVR-antenna beam width and offset from the main antenna can be entered into the mix of parameters constrained by the precision requirement.

In this way constraints on receiver parameters are revealed and more informed choices can be made. If a requirement cannot be met, it will be known early in the design phase. Requirements traceability at the receiver level is defined by this synthesis — from the various receiver parameters mentioned in the previous paragraph back to the delay precision requirement. For these reasons it is important when developing these receivers to synthesize these estimators early in the design

phase of the project and allow the results to drive the design.

8.2 Liquid Water and Water Ice

Not discussed up to this point, liquid water and water ice in clouds have f^2 frequency dependence to their emissivity, while contributing much less to delay. To account for clouds, the atmospheric model must include absorption and emission from clouds, requiring that liquid water and water ice have their place in the synthesis. This place can be on the same footing as water vapor, i.e., they can have their own distributions along the line of sight, distributions whose emission is shadowed by absorption lower in the atmosphere, and that shadows emission from higher in the atmosphere.

From (Thompson, Moran, and Swenson 1986, Eq. 13.115), absorption per unit length by liquid water in clouds is approximated by

$$\text{absorption coefficient} = \eta\, \rho_{\text{liq}}\, e^{(T_0 - T_{\text{a}})/T_1} \left(\frac{f}{c}\right)^2 \tag{71}$$

where $\eta = 10^3 \text{m}^4/\text{kg}$, ρ_{liq} is the density of liquid water, $T_0 = 291\text{K}$, and $T_1 = 36.60\text{K}$. This term is a component in the emissivity e in Eq. (19). The f^2 frequency dependence is identical to that of the water-vapor continuum excess (Sutton and Hueckstaedt 1996, Eq. 9), which is an integral part of the water-vapor line shape. Estimator synthesis must have both components in the atmospheric operator A for the estimator to correctly account for spectra and distinguish water vapor from liquid water (and water ice and other aerosols).

8.3 WVR Antennas with Wider Beams

The WVR pathfinder project for ngVLA uses the beam of the VLA antenna and its K-band front end to gather and amplify the water-vapor signal. The VLA 25 m dish in K band provides a very narrow beam of a couple of arc minutes plus the dish diameter, narrow enough that there is little decorrelation of water vapor across the beam. This is why the action of the atmospheric operator A need only be a one-dimensional integral and why estimator synthesis is as tractable as it is.

In the ngVLA a small WVR antenna attached to the perimeter of each main antenna is contemplated. The larger angular size of the beam of this antenna samples a larger volume surrounding and crossing the main-antenna beam; consequently there is necessarily a degree of decorrelation of water-vapor fluctuations across the WVR beam and between the two beams (Nikolic, Hills, and Richer 2007). For this reason there is an additional source of tracking error using a separate WVR antenna compared with the CWVRs on EVLA antennas.

It is possible to formulate the estimator synthesis in a natural way such that the integrals for $S\mathbb{R}S^T$, $S\mathbb{R}C^T$, and $C\mathbb{R}C^T$ are over the entire volumes of the main and WVR antennas' beam volumes and weighted by their beam patterns. In this way the tracking error takes into account the different delay fluctuations in the beams and the weight function correctly accounts for them. There is a cost in that much greater computation time is required to compute these integrals. A brute-force calculation of the difference between the delays seen by two antennas (Towne 2019b) showed rather modest error for a one-meter WVR dish size placed at the edge of an 18 m dish. Thus it seems unnecessary to perform the full-beam synthesis and that one-dimensional synthesis is sufficient.

Nevertheless, this is a tool available to the designer for cases where the discrepancy between the main and WVR beams is not so small. Also note that these calculations are only as good as

the atmospheric model, and that the complexity of turbulence in the atmosphere add a degree of uncertainty to any conclusions.

8.4 The Dry Component

The dry component of delay fluctuations refers to convectively driven density fluctuations, which are not observable directly with the 22 GHz water-vapor line. The magnitude is significant to ngVLA, and some means to deal with it beyond interpolation between source calibrations may be necessary. Boundary-layer simulation have shown correlation between fluctuations of water vapor and density fluctuations (Stirling et al. 2008), which varies with height, particularly in relation to the inversion. Knowledge, or models of, of how this correlation varies might may be useful for partly accounting for the dry component.

In this section it is outlined how such a model can be entered into the estimator formalism of Sec. 2. Let δn_d be the unknown fluctuation of the index of refraction due to density fluctuations. δn_d enters delay fluctuations in Eq. (2) as follows:

$$\delta d = S\rho_{\rm wv} + \int ds\, \delta n_d \tag{72}$$

where the integral is along the line of sight. The estimate $\delta d_{\rm est}$ in Eq. (2) is unchanged. The function space of atmospheric fluctuations has been expanded from just $\delta\rho_{\rm wv}$ to the pair $(\delta\rho_{\rm wv}, \delta d_n)$, on which the operator

$$S' = (S, \int ds) \tag{73}$$

acts to return the delay

$$\delta d = S'(\rho_{\rm wv}, \delta n_d) \tag{74}$$

as given in Eq. (72). The covariance matrix $\mathbb{R}$ of Eq. (4) expands as well to

$$\mathbb{R}'(s_1, s_2) = \left(\begin{array}{cc} \langle \delta\rho_{\rm wv}(s_1)\delta\rho_{\rm wv}(s_2)\rangle & \langle \delta\rho_{\rm wv}(s_1)\delta n_d(s_2)\rangle \\ \langle \delta n_d(s_1)\delta\rho_{\rm wv}(s_2)\rangle & \langle \delta n_d(s_1)\delta n_d(s_2)\rangle \end{array} \right)_\Sigma \tag{75}$$

where Σ indicates symmetrization with respect to s_1 and s_2. The model of the dry component fluctuations and correlations are inserted into the δn_d elements of Eq. (75). The atmospheric matrix A, and consequently C, are not sensitive to δd_n, and so the kernel $C\mathbb{R}C^T + \mathbb{V}$ is unaffected, although it is understood that it (and C) maps the δd_n part of the function space to zero. The new optimal weight function is the primed Eq. (8)

$$W'_{\rm opt} = S'\mathbb{R}'C^T(C\mathbb{R}C^T + \mathbb{V})^{-1} \tag{76}$$

and the optimal delay tracking precision is the primed Eq. (9)

$$\delta d^{2\,\prime}_{\rm opt} = S'\mathbb{R}S^{T\,\prime} - S'\mathbb{R}'C^T W^{T\,\prime}_{\rm opt} \tag{77}$$

Note that the when varying the dry-compoment model, the original estimator need not be recomputed.

9 Conclusions

- The estimator formalism provides a weight function that minimizes tracking precision by taking into account atmospheric physics and receiver characteristics. The water-vapor line shape as a function of altitude, statistical models of turbulence, a model of atmospheric temperature and composition profiles, receiver noise statistics, and receiver response functions are all utilized.
- Sensitivity analysis utilizing the estimator allows performance comparison with errors in elements of the models the estimator is built upon. For example, a gain-stability and $T_{\rm cal}$-accuracy are readily computed from the estimator. The methodologies outlined in Sec. 6 provide a general framework for evaluating sensitivities and formulating technical requirements.
- Noise in the CWVRs is dominated by fluctuations in the system not related to rf amplifier noise, although gain fluctuations may play a role. This noise must be lowered to meet a tracking-precision requirement appropriate to the EVLA (and ngVLA). There are a lot of sources of noise and errors that together limit of tracking precision. It is a challenge to identify all that matter and control or mitigate them sufficiently well to reach the very high degree of cancellation of ambient delay fluctuations required of the estimator as set forth in the ngVLA project requirements.
- The atmospheric model, particularly the turbulence model, is a weak link in estimator synthesis, although it is not known how or if it will limit performance in the EVLA and/or ngVLA.

10 Acknowledgements

Thanks to Alan Erickson, other digital-group members, Craig Hennies, Bryan Butler, Chris Hales and the front-end group.

A Appendices

A.1 Baseline Atmospheric Model

The atmospheric model is largely hypothetical, consisting of a hypothetical reference atmosphere and a turbulence model based on (Carilli and Holdaway 1997) but scaled for varying intensity.

For the purpose of a test case, the reference atmosphere has the temperature profile shown in Fig. 10. It resembles the standard atmosphere published by NASA in 1962, but with the surface temperature tapered to 280K. Relative humidity is given by an exponential with 4 km scale height and value of 50% at the surface.

As was discussed in Fig. 5, (Carilli and Holdaway 1997) motivates a baseline turbulence model with a uniform structure parameter over a turbulent-layer thickness, and constant fluctuation intensity ($\sqrt{\langle \delta n^2 \rangle}$). This model has a fairly calm intensity $\sqrt{S\mathbb{R}S^T}/c = 5$ ps. Other intensities are obtained by simply scaling $\mathbb{R}$, in effect scaling D_n and $\sqrt{\langle \delta n^2 \rangle}$ identically.

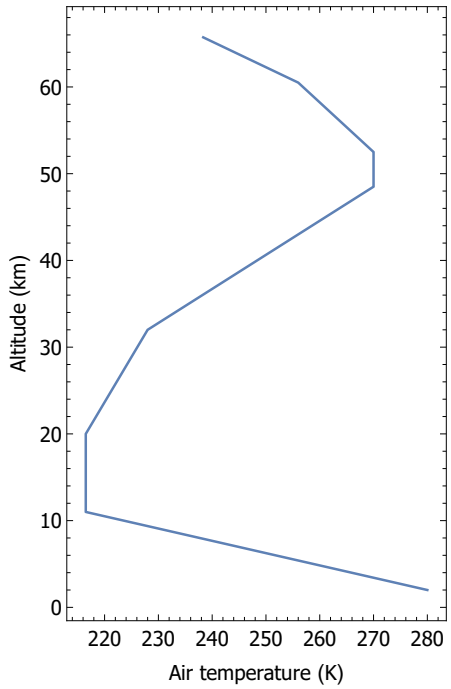

Figure 10: Reference atmosphere temperature profile that roughly follows the 1962 standard atmosphere published by NASA.

A.2 Receiver Characteristics

Figure 11 shows the frequency responses of the five channels of a CWVR.

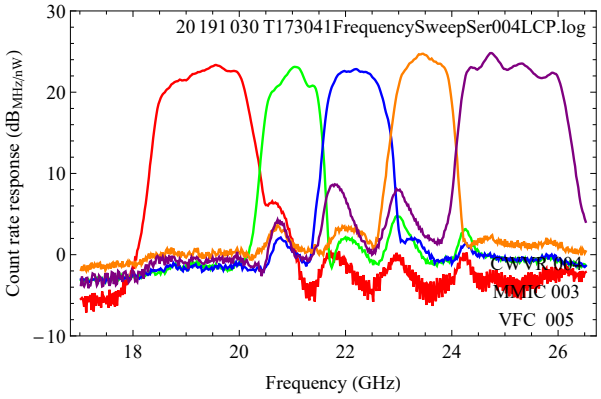

Figure 11: Frequency response of the CWVR in ea25, LCP.

The receiver noise covariance matrices were bench measured on 2019-10-25. The LCP value is

$$\mathbb{V}_{\text{LCP}} = \begin{pmatrix} 0.0136 & 0.0146 & 0.0136 & 0.0093 & 0.0098 \\ 0.0146 & 0.0185 & 0.0172 & 0.0118 & 0.0116 \\ 0.0136 & 0.0172 & 0.0173 & 0.0116 & 0.0114 \\ 0.0093 & 0.0118 & 0.0116 & 0.0085 & 0.0084 \\ 0.0098 & 0.0116 & 0.0114 & 0.0084 & 0.0091 \end{pmatrix} \text{K}^2 \tag{78}$$

and RCP is

$$\mathbb{V}_{\mathrm{RCP}} = \begin{pmatrix} 0.0659 & 0.0951 & 0.0912 & 0.0687 & 0.0571 \\ 0.0951 & 0.1412 & 0.1358 & 0.1026 & 0.0851 \\ 0.0912 & 0.1358 & 0.1321 & 0.0996 & 0.0828 \\ 0.0687 & 0.1026 & 0.0996 & 0.0764 & 0.0636 \\ 0.0571 & 0.0851 & 0.0828 & 0.0636 & 0.0545 \end{pmatrix} \mathrm{K}^2 \tag{79}$$

The difference between the two polarizations proves useful for illustrating how estimator performance varies with receiver noise. For comparison, the radiometer equation predicts, assuming $T_{\mathrm{sys}} = 50\mathrm{K}$ and the integration time is 0.85 s (Towne 2019a), of order $3 \times 10^{-6}\mathrm{K}^2$ (BW $= \sim 1$ GHz). Figure 12 shows the estimator performance given these receiver characteristics and the atmospheric model of Appendix A. The horizontal axis is the intensity of the turbulence, which is varied by simply scaling $\mathbb{R}$ in $S\mathbb{R}C^T$, $S\mathbb{R}C^T$, and $C\mathbb{R}C^T$.

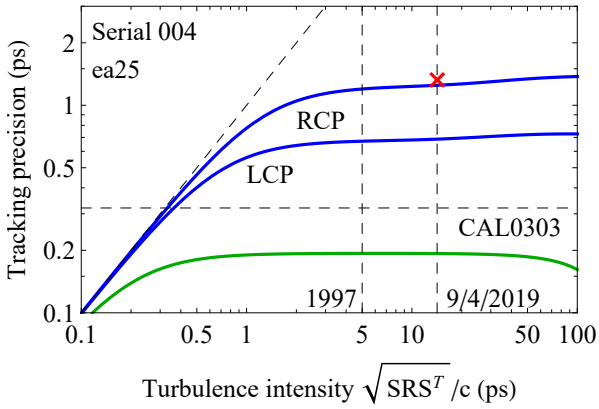

Figure 12: Tracking precision in antenna 25 of both RCP and LCP. The reference atmosphere is that described in Sec. 6.1, the turbulence model has structure parameter and fluctuation intensity scaled from those estimated from (Carilli and Holdaway 1997). The left vertical dashed line is the intensity of delay fluctuations estimated from (Carilli and Holdaway 1997), the right vertical dashed line is the intensity of delay fluctuations taken from a one-hour observation by V. Dhawan on 9/4/2019, and the horizontal dashed line marked CAL0303 is the tracking precision requirement from (Hales 2019, CAL0303). The diagonal dashed line is untracked precision. The red 'x' is the tracking precision from the observation of Sec. A.3, and the green trace assumes noise of a perfect receiver given the channel bandwidths and integration time.

But note that a measurement of $\mathbb{V}$ in situ differed from that measured on the bench. It is not known whether this difference is due to differing noise in the environments, if there are methodological problems with the measurements, or if the measurements were not sufficiently long to sample the fluctuations accurately. So just measuring receiver noise has proved problematic. The intensity of the noise swamps amplifier noise.

A.3 Calibrator Observation of 9/4/2019

The 4 Sep 2019 observation of a calibrator can serve as a rough measurement of the tracking precision. This measurement compares the difference between the calibrated phase difference between two antennas with the difference between a weighted sum of the channels of the two CWVRs, shown in Fig. 13. The antennas of that figure were separated by ~ 10 km, well into the decorrelated regime where the single-antenna delay can be taken from the figure by scaling by $1/\sqrt{2}$.

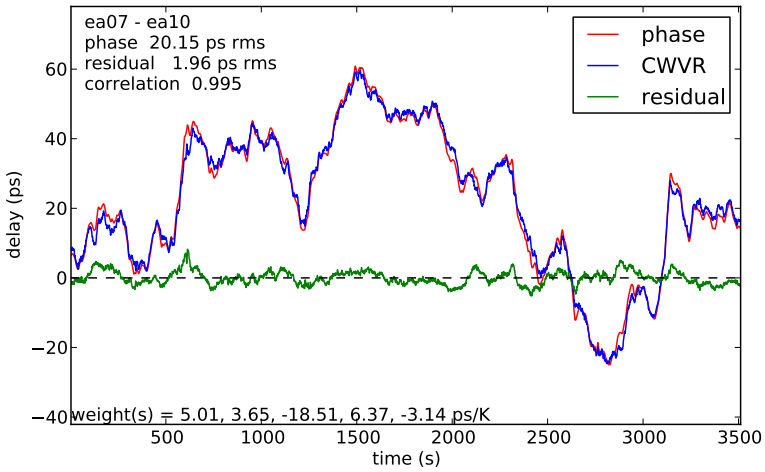

Figure 13: The calibrated phase difference between two antennas from an observation of a calibrator on 4 Sep 2019 (red) and the difference between the same antennas of an optimally weighted sum of the channels of CWVR channel intensities. The green trace is the residual. The antenna separation at the time was 10 km.

This observation is a poor comparison with the assumptions of Fig. 12 for a number of reasons, one being that at that time all the CWVRs had not had the F318 logic update for the pulse-counting problem, and had no power-supply replacement for the temperature-regulation electronics. Another is that the choice of weight function is unfair in that it is the one that minimizes the difference between the phase and weighted CWVR intensities, instead of being determined independently of the data. Nevertheless, the factor of order two between them is encouraging.

References

Butler, Bryan (1999). "Some Issues for Water Vapor Radiometry at the VLA". In: *VLA Scientific Memo 177*.

Carilli, Chris and M. A. Holdaway (1997). "Application of Fast Switching Phase Calibration at mm Wavelengths on 33 km Baselines". In: *VLA Scientific Memo No. 173*.

Chandrasekhar, S. (1960). *Radiative Transfer*. Dover Books on Intermediate and Advanced Mathematics. Dover Publications. ISBN: 9780486605906. URL: `https://books.google.com/books?id=CK3HDRwCT5YC`.

Fruchter, A. S. and R. N. Hook (Oct. 2001). "Drizzle: A Method for the Linear Reconstruction of Undersampled Images". In: *Space Telescope Science Institute Preprint*.

Gill, Ajay, Robert J. Selina, and Bryan J. Butler (2018). "A Study of the Compact Water Vapor Radiometer for Phase Calibration of the Karl G. Janksy Very Large Array". In: *Journal of Astronomical Instrumentation*.

Gordon, I. E. et al. (Dec. 2017). "The HITRAN 2016 molecular spectroscopic database". In: *jqsrt* 203, pp. 3–69. DOI: `10.1016/j.jqsrt.2017.06.038`.

Hales, Chris (May 2019). "Calibration Strategy and Requirements". In: 020.22.00.00.00-0001-REQ.

Kolmogorov, Andrei Nikolaevich (1991). "The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers". In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 434.1890, pp. 9–13.

Nikolic, B. (Mar. 2009). "Inference of Coefficients for Use in Phase Correction I". In: *arXiv e-prints*. arXiv: 0903.1179 [astro-ph.IM].

Nikolic, B., R. E. Hills, and J. S. Richer (2007). "Limits on Phase Correction Performance Due to Differences Between Astronomical and Water-Vapour Radiometer Beams". In: *Alma Memo No. 573*.

Smith, Ernest K. Jr and Stanley Weintraub (Sept. 1953). "The Constants in the Equation for Atmospheric Refractive Index at Radio Frequencies". In: *Proceedings of the IRE* 50, pp. 1035 –1037. DOI: 10.1109/JRPROC.1953.274297.

Stirling, Alison et al. (2008). "Turbulence simulations of dry and wet phase fluctuations at Chajnantor. Part I: The daytime convective boundary layer". In: *ALMA Memo 517*.

Sutton, E. C. and R. M. Hueckstaedt (Nov. 1996). "Radiometric monitoring of atmospheric water vapor as it pertains to phase correction in millimeter interferometry." In: *Astronomy and Astrophysics Supplement Series* 119, pp. 559–567.

Tahmoush, David A. and Alan E. E. Rogers (2000). "Correcting atmospheric path variations in millimeter wavelength very long baseline interferometry using a scanning water vapor spectrometer". In: *Radio Science* 35.5.

Thompson, A. R., J. M. Moran, and G. W. Swenson (1986). *Interferometry and synthesis in radio astronomy*.

Towne, Nathan (2019a). "CWVR Monitor and Control". In: *Digital Group Memo*.

— (2019b). "Precision With Which a WVR Antenna Can Measure Main Antenna Delay". In: *Digital Group Memo*.

Wikipedia (2019). "Color filter array". In: *https://en.wikipedia.org/wiki/Color_filter_array*.