OVLBI-ES MEMO NO. 71

DETAILS OF POST-PASS PROCESSING OF TIME CORRECTION DATA

Larry R. D'Addario
24 April 1997

1. Introduction

This document describes the processing required to generate the correlator time correction file ("DeltaT file") [1] from measurements made during a tracking pass. It is specific to the NRAO OVLBI earth station at Green Bank, although many of the details are also applicable to the other existing earth station designs.

Many aspects of this have been considered in earlier documents. Design considerations for time transfer systems were discussed in [2]. General analysis of the errors involved in the time transfer process and of the prospects for their correction was given in [3]. Algorithms for estimating and removing the non-reciprocity caused by the troposphere and ionosphere were given in [4,5]. The specific effects of errors in the earth station clock were analyzed in [6]. The design of the critical hardware and firmware for the NRAO station was described in [7]. Application of the time corrections at a correlator was discussed in [8].

Here I attempt to gather together all the relevant information on the NRAO station, leading to specific formulas for computing the values needed in the DeltaT file. We adopt the view that each earth station should provide the most accurate timing that can be derived from the information available to it. Therefore, I attempt to correct for all known non-ideal behavior, even if its effect is usually small. I show here that we can then expect an accuracy of a few ps in the variation of timing error during a tracking pass, with a constant offset in absolute timing of a few μs; the latter is dominated by the error in the predicted orbit, without which it would be the order of 50 ns.

2. Classification Of Effects

Table 1 lists various effects that influence the time transfer process, and that would result in timing errors (incorrect values of the time correction, Δt) if not accounted for. They are classified according to whether they are internal or external to the earth station, and according to what aspect of the time transfer process they influence. I consider three ways in which non-ideal behavior of the system can enter the process: (1) by affecting the initialization event, when the tape-time generating clock (in the formatter) is set, and thus causing an error in the part of Δt that remains constant during a tracking pass (shown as type "Δt_0" in Table 1); (2) by causing errors in Δt at other than the initialization event ("Δt"); and (3) by causing errors in the time stamps on the Δt values ("t_i"). Whether an effect influences the constant or variable part of Δt does not imply anything about whether the effect itself is constant or variable. For example, the decoder UTC clock is used only at initialization, so its error can affect only the constant part of Δt; but it might be different from one pass to the next. Conversely, some hardware delays can in principle affect each value of Δt (and are so classified in Table 1), but in practice are taken to be constant.

Here, as in most earlier documents, a "tracking pass" means a continuous period of two-way communication with the spacecraft. If the link is briefly lost and then re-acquired, it may be necessary to have a new clock setting event; in such a case, the subsequent period is considered a new tracking pass. In accordance with the specification [1], this requires that a new DELTA_T table be started in the DeltaT file.

Some of the effects are shown later to be extremely small in magnitude. These are marked as type "0" in the table.

Each of these effects is discussed in detail in the following sections.

3. Effects Internal To The Earth Station

A simplified block diagram of the NRAO Green Bank station is given in Figure 1. I assume that readers are generally familiar with the system configuration, so many relevent details will not be discussed here. But it will be helpful to review some points closely related to the time transfer process.

The station includes several "clocks," as follows.

(a) The most precise and stable measure of station time is maintained in hardware at the antenna vertex. We'll call this the "DSP clock" because its lowest frequency components are maintained by a Digital Signal Processor in the Two Way Timing (TWT) Control Module. Reference signals from the hydrogen maser are

Table 1: Phenomena Affecting The Time Corrections

Phenomonon	Symbol	Type
Effects Internal to Earth Station		
Station UTC clock error at decoder	$\delta t_{\rm dec}$	Δt_0
Tape clock (formatter) setting error	$\delta t_{\rm fmtr}$	Δt_0
Signal processing delay, ant. thru dec.	$\tau_{\rm sig}$	Δt_0
Maser cable length variation	$\Delta\tau_{\rm maser}$	Δt
Station clock error at TWT DSP	$\delta t_{\rm DSP}$	$\Delta t, t_i$
Uplink coarse synthesizer delay	$\tau_{\rm syn,u}$	Δt
Uplink fine synthesizer delay	$\tau_{{\rm DDS}u}$	Δt
Transmitter PLL acceleration error	$\tau_{\rm pll,u}$	$\Delta t, 0$
Downlink coarse synthesizer delay	$\tau_{\rm syn,d}$	Δt
Downlink fine synthesizer delay	$\tau_{{\rm DDS}d}$	Δt
Demodulator PLL acceleration error	$\tau_{\rm pll,d}$	$\Delta t, 0$
Phase processing delay	τ_ϕ	t_i
Effects External To Earth Station		
Absolute downlink delay (setting event)	$\tau_{\rm link}(t_0)$	Δt_0
Troposphere non-reciprocity	$\Delta\tau_{\rm trop}$	Δt
Ionosphere non-reciprocity	$\Delta\tau_{\rm ion}$	Δt
Spacecraft turn-around delay	$\tau_{\rm sc}$	Δt
Spacecraft PLL acceleration error	$\tau_{\rm pll,sc}$	$\Delta t, 0$
Intrinsically Non-correctable Effects		
Geometrical error (see [3])	τ_g	Δt

brought directly here on optical fiber, and various timing signals are generated from them, ranging from 1 Hz to 500 MHz. Variation in the length of the optical fiber is measured to a precision of < 0.03 ps and an accuracy of < 1 ps. After correcting for this and for the effects of maser frequency offset since initialization, the absolute UTC is maintained by this hardware with a stability of < 1 ps/day, an absolute accuracy of < 50 ns, and an ambiguity period of 1 s. Better absolute accuracy should be possible by laboratory calibration of some hardware delays, but this is not expected to be necessary. The 1 s ambiguity is eliminated by having the station computer read results from the hardware at least once per second.

(b) The 10 MHz maser reference is brought from the antenna vertex to the equipment building on a coaxial cable approximately 30 m long. There it is used to drive a counter inside the decoder module in such a way that the counter provides a measure of station UTC with a resolution of 10 ns and an ambiguity of > 32 days (48 b). We'll call this the "decoder clock." It is subject to the same errors as the DSP clock plus the delay in the (unmonitored) coax cable and a setting uncertainty of about 1 LSB or 10 ns. Its absolute accuracy is checked against GPS time in a dedicated GPS receiver having a precision of about 30 ns and a short-term accuracy of about 200 ns peak. Better knowledge of the absolute time error is expected to be possible by long-term comparison with GPS, or by common-view techniques. It can be synchronized to the GPS receiver under manual control, but this should be done only rarely. By avoiding any disturbances to the time-determining circuitry (all of which is on uninterruptable power), we hope to keep the offset between station UTC and true UTC stable to a few ns per month. Here "station UTC" means the reading on the decoder clock after correction for the measured maser cable length variation and integrated maser frequency error.

(c) The demodulated bit streams from the spacecraft provide measures of "satellite time." These include the recovered data clock (64 MHz for VSOP, 72 or 36 or 18 MHz for RadioAstron) and the frame synchronization word periodically embedded in the data (200 Hz for VSOP, 400 or 200 or 100 Hz for RadioAstron). After signal acquisition and synchronization, the decoder uses these to synthesize 32 MHz ("S32M") and 1 Hz ("S1") clock signals for the formatter; it does this by arbitrarily selecting the beginning of one downlink frame as occurring on the S1 boundary. The formatter then contains a counter driven by S32M that provides a measure of satellite time; the content of this counter is called "tape time" and is used to write time codes to the tape headers. Note that satellite time and hence tape time include any sampling-time error on the

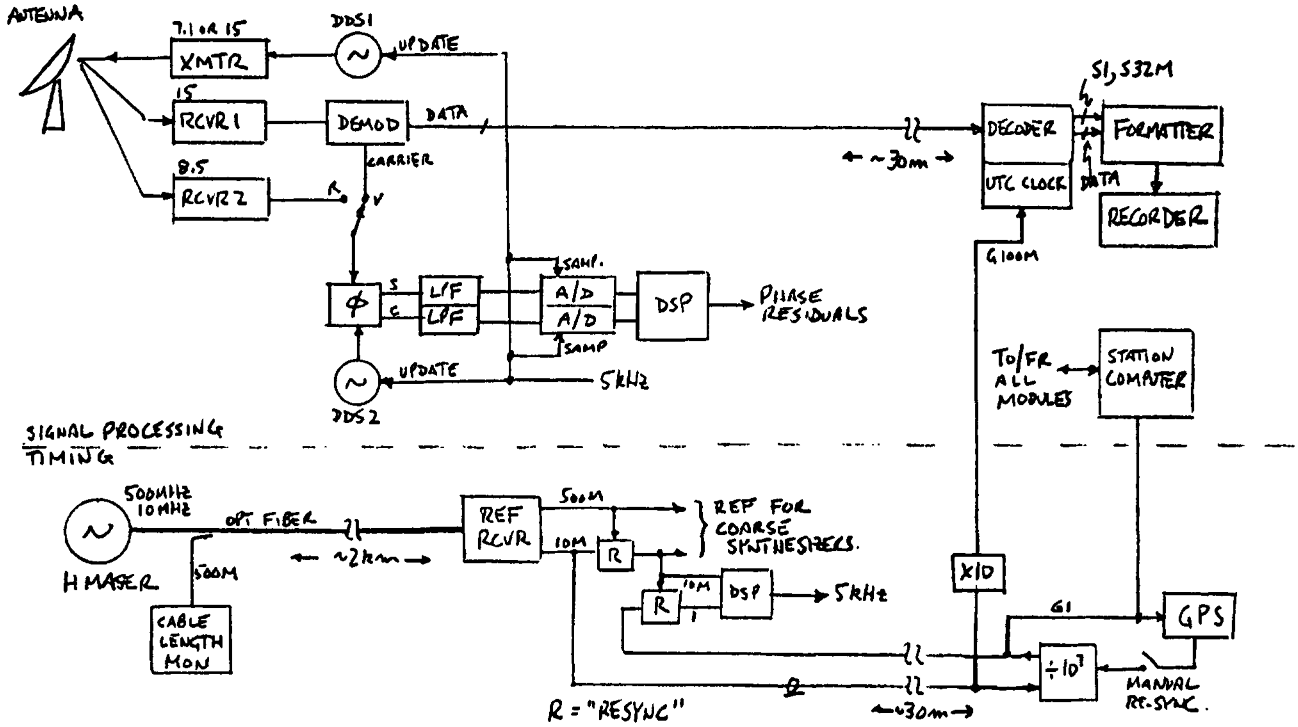

Figure 1: Simplified block diagram of the NRAO OVLBI Earth Station.

spacecraft (which nominally, but not exactly, runs at the same rate as UTC due to the uplink compensation), and it also includes the full Doppler shift of the downlink. Therefore, even if it is initialized to be approximately equal to station UTC, it will diverge from station UTC over the course of a tracking pass. The VLBA formatter allows initialization of the counter only on the S1 boundary, and then only to a value representing an integer second. We normally initialize it at the beginning of a pass to the nearest second to station UTC; this can produce a setting error of ±0.5s. An arbitrarily large setting error is allowed, as long as it is precisely known and included in the final Δt values.

With the above background, we now consider each of the items in Table 1 that is internal to the earth station.

3.1 Decoder clock error and formatter setting error

The decoder clock is used only during time initialization at the beginning of a tracking pass, and therefore it affects only the constant part of Δt. Let $t_{\rm dec}(t_0)$ be the decoder clock reading at the clock setting event, let $\delta t_{\rm dec}(t_0)$ be the clock error (difference from true UTC) at that time, and let $t_{\rm fmtr}$ be the value to which the formatter's tape clock is then initialized. These terms result in a net contribution to Δt_0 of $t_{\rm dec}(t_0) + \delta t_{\rm dec}(t_0) - t_{\rm fmtr}$.

We maintain knowledge of $\delta t_{\rm dec}$ as accurately as we can, and keep the results in our own data base in Green Bank. The primary basis of this knowledge is long-term monitoring against GPS by means of the dedicated GPS receiver. By making many measurements over several days, the effects of receiver errors, satellite geometry, and selective availability are largely averaged out. Our results can be checked and supplemented by similar measurements using other GPS receivers in Green Bank checking other clocks driven by the same hydrogen maser. In this way, knowledge of the maser's frequency error (long-term drift rate) is expected to be excellent. As a minimum, we will maintain our best estimate of δt_{dec} in the form of a 2-term polynomial, with offset and slope at a given epoch updated approximately once per day. Future upgrades that should be pursued include: more precise checks of clock offset by common-view GPS with NIST or USNO; checks of clock offsets by VLBI, requiring configuration of the station as a VLBI telescope; and longer-term fitting to higher order polynomials.

All possible care should be taken in determining and maintaining our knowledge of δt_{dec}. For example, the measured difference from GPS time should be corrected for UTC-GPS, as published by USNO, before determining the long-term trend. Known system effects on the clock reading, such as the measured variation in the optical cable length, should also be removed and handled separately; this is further discussed in 3.3 and 5.3 below. (Diurnal and faster variations in the latter are probably much smaller than other errors, but long-term drifts could affect the results if not accounted for.)

3.2 Signal processing delay

The remainder of the constant part of Δt consists of converting the UTC at the station to the UTC on the satellite. For this we must subtract from Δt the link delay (considered later) plus the signal delay in the station, from the antenna reference point up through the point in the formatter hardware where the clock reading gets associated with a data bit. The latter consists of $\tau_{sig} = \tau_{rcvr1} + \tau_{cable} + \tau_{dec}$. The receiver consists of the input optics, the front end, the downconverter, and the demodulator. There is a geometrical delay through the optics of approximately 20 ns (6.0 m), and cabling throughout the receiver of about 4.3 ns (1.0 m). Additional delay is caused by filtering in the downconverter, 10 ns; and in the demodulator, 47 ns. This is a total of 81 ns, which we assume to be constant. This value is just a rough estimate (say, ±20ns), and it may be refined later. Finally, the demodulator introduces a delay of about half the downlink clock period in recovering the data bits; so

$$\tau_{rcvr1} = (81 \pm 20\text{ns}) + 0.5/f_c$$

where f_c is the downlink clock frequency (64 MHz for VSOP, 72 or 36 or 18 MHz for RadioAstron). Next, the cable from the demodulator (at the vertex) to the decoder (at the equipment building) is about 30m long with a velocity factor of 0.78, giving a delay of $\tau_{cable} = 128\,\text{ns}$. Again, we assume that this is constant.

Finally, the decoder introduces a slight delay, as follows: each downlink frame, the value of the UTC counter is latched shortly after detection of each sync word; one of these values will be $t_{dec}(t_0)$. But the data sample that will be labeled with t_{fmtr} is the one corresponding to a particular bit of the header which arrived at a slightly different time. (The header bits replace signal samples on the donwlink, and are replaced by pseudo-random noise on the tape.) So, let N_1 be the number of downlink clock cycles between the arrival of the last sync bit at the decoder input and the latching of the UTC counter (a function of the decoder's internal design); and let N_2 be the number of clock cycles between the labeled bit and the last sync bit. This gives

$$\tau_{dec} = (N_1 + N_2)/f_c.$$

From the decoder design and laboratory measurments [12], $N_1 \approx 28$ for VSOP and 33 for RadioAstron, and $N_2 = -9$ for VSOP and -10 for RadioAstron. Since the clock's 100 MHz reference is asynchronous with the downlink, the delay is subject to an uncertainty of ±1 cycle of the reference, or ±10 ns.

In summary, we have

$$\tau_{sig} = (209 \pm 25\text{ns}) + (N_1 + N_2 + 0.5)/f_c,$$

where all values are subject to refinement from measurements.

People who are familiar with the detailed designs of the decoder and/or the formatter may know that each includes some buffering of the data, introducing additional physical delays. The decoder buffers several downlink frames, resulting in delays of many milliseconds. However, these delays have no logical significance because the associated clock signals (S32M and S1) are synchronized with the data streams and are delayed by the same amounts. A time tag has been logically attached to a data sample when the decoder clock reading is latched, so subsequent delays need not be taken into account; their effect is similar to that of the much bigger delay involved in shipping the tape to a correlator.

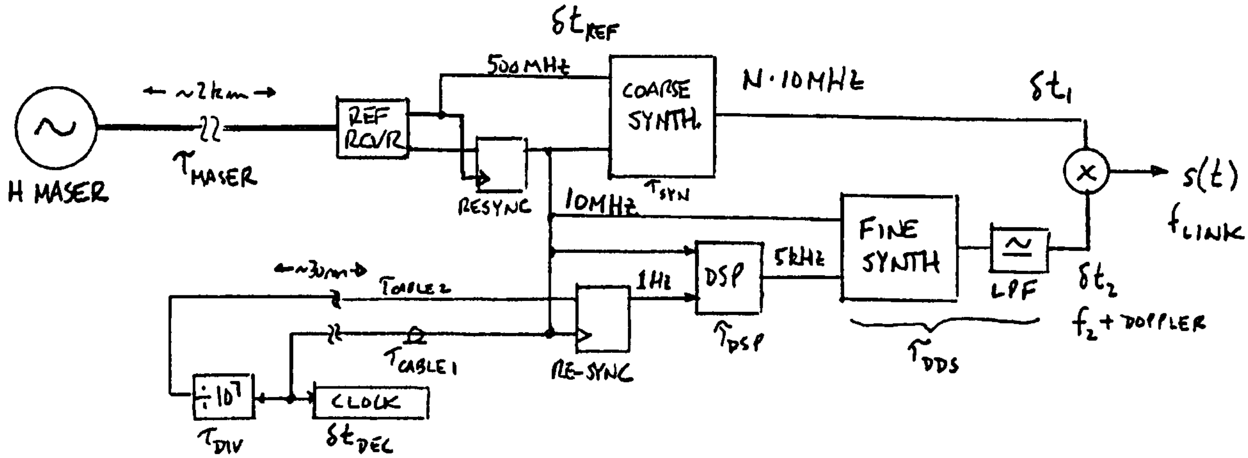

Figure 2: Scheme for synthesizing uplink or predicted downlink signal. Timing errors at location x are denoted δt_x and delays through component y are denoted τ_y.

3.3 Maser cable length variation

The buried optical fiber cable is approximately 2 km long, but its actual length and absolute delay are not accurately known. It transmits reference signals at 500 MHz and 10 MHz. The 500 MHz signal is returned to the maser room on a separate but identical fiber in the same cable, and the round-trip phase change is monitored. Assuming reciprocity, this gives the cable delay with an ambiguity of 1.0 ns. The resolution and noise of the measurement system is around .03 ps, and the absolute accuracy is believed to be well under 1 ps. For details of this system's design and testing, see [9]. Experience shows that the typical delay variation is 20 ps diurnally, 80 ps with synoptic weather systems, and 300 to 500 ps seasonally; thus, it always remains within the ambiguity limit. Although these numbers are typical, we occasionally see delay changes of 10 ps in 5 min.

The cable length affects the accuracy of both the decoder clock and the DSP clock. The effect on the decoder clock is small compared with other effects, but it should be taken into account in analyzing long-term trends, as discussed above. The effect on the DSP clock is more significant because this clock drives the generation of microwave signals for the two-way timing link to the spacecraft. A derivation of this was given in approximate form in [6]; in the following section I give a more complete formulation for the NRAO station.

3.4 Timing errors in uplink and downlink synthesizers

Figure 2 shows the essential features of the circuitry for synthesizing the uplink signal $s_{\rm up}(t)$ or the predicted downlink signal $s_{\rm dn}(t)$ from the maser references. In each case, the synthesis is separated into two parts: a "coarse" signal is effectively created at a harmonic of 10 MHz near the desired frequency; and a "fine" signal is generated in a direct digital synthesizer (DDS). Only the latter is phase-steered as a function of time during the tracking pass. The instantaneous phase of the final signal (at the microwave link frequency) is the sum of that of the coarse and fine parts. Various details of the implementation are different for the uplink and downlink signals; these details are not important to the present discussion so they will not be covered here.

Recall that the maser transmits reference signals at 10 MHz and 500 MHz. The received 10 MHz is "resynchronized" by forcing its zero-crossings to coincide with zero-crossings of the 500 MHz reference; further, it is guaranteed by design that the difference in delay for transmission of the two references is always much less than 1/(500 MHz). Thus, only the variations in the transmission delay of the 500 MHz reference are important; it is the master. This delay is measured in real time. Next, a 1 Hz reference is generated by dividing the 10 MHz by 10^7. This is actually done in the equipment building, requiring 2-way transmission on cables. However, the resulting 1 Hz signal is resynchronized to the 10 MHz at the vertex, and the cable delay variation is much less than 1/(10 MHz); thus, the cables have no effect on the stability and the 500 MHz signal remains the master of all three timing references. The 1 Hz reference is also used to drive the decoder clock at the equipment building, so the cables do produce an offset between the decoder clock and the vertex timing.

Let $\omega = N\omega_{10} + \omega_{\mathrm{DDS}}$ be the nominal link frequency, where $\omega_{10} = 2\pi(10\,\mathrm{MHz})$. Then the desired synthesized signal (either uplink or downlink) is

$$s_{\mathrm{ideal}}(t) = \cos\omega[t \pm \hat{T}(t)]$$

where t is true UTC (coordinate time) at the station and $\hat{T}$ is the predicted link delay; the plus sign applies to the uplink case and the minus sign to the downlink. Then the actual signal may be represented as

$$s(t) = \cos[N\omega_{10}(t+\delta t_1) + \omega_{\mathrm{DDS}}(t+\delta t_2) \pm \omega\hat{T}(t+\delta t_2)]$$

where δt_1 is the effective timing error of the coarse synthesizer (whose various contributions are considered later); and δt_2 is the effective timing error in the fine synthesizer, including the DSP clock error, delays in the DDS's electronics, and the delay in the filter following the DDS. From this it can be seen that the actual signal is not quite expressible as the desired signal with a time offset; this is a result of the splitting of the synthesis into coarse and fine parts, with each having a different timing error (not considered in [6]). After some algebra, the last equation can be put into a form involving only the link frequency ω:

$$s(t) = \cos\omega[t + \delta t \pm \hat{T}(t+\delta t_2)],$$

where $\delta t = (N\omega_{10}\delta t_1 + \omega_2\delta t_2)/(N\omega_{10} + \omega_2)$ is the frequency-weighted average of the two timing errors.

Using the last expression, adding subscripts u and d to denote the uplink and downlink, and following the same steps as in [6] eventually shows that the measured timing residual will be

$$\tau_{\mathrm{meas}} = T_u - \hat{T}_u(t + \delta t_{2u} - T_u - T_d) + T_d - \hat{T}_d(t + \delta t_{2d}) + \delta t_d - \delta t_u + \tau_{\mathrm{rcvr2}} + \tau_{\mathrm{xmtr}}$$

where T_u, T_d are the actual uplink and downlink delays. The last two terms are needed to refer the measurement to the station reference location for which the predicted link delays $\hat{T}_{u,d}$ were calculated. By convention, the standard reference location is the intersection of the antenna's rotation axes. No physical signal exists there, but we must account for the propagation times of physical signals to or from the reference point. Then τ_{rcvr2} is the delay from the reference point to the phase detector, and τ_{xmtr} is the delay from the uplink synthesizers to the reference point.

The preceding equation can be compared to the ideal residual (in the absence of timing errors), which is

$$\tau_{\mathrm{ideal}} = T_u - \hat{T}_u(t - T_u - T_d) + T_d - \hat{T}_d(t).$$

Following the same reasoning as in [6], the apparent value of Δt is

$$\Delta t_{\mathrm{apparent}} = \Delta t_0 + \tau_{\mathrm{meas}}/2$$

but the correct value is

$$\begin{aligned}\Delta t &= \Delta t_0 + \tau_{\mathrm{ideal}}/2 - \delta t_u - [\hat{T}_u(t - T_u - T_d + \delta t_{2u}) - \hat{T}_u(t - T_u - T_d)] \\ &\approx \Delta t_0 + \tau_{\mathrm{ideal}}/2 - \delta t_u - \delta t_{2u}\dot{\hat{T}}_u(t - T_u - T_d)\end{aligned}$$

where the last two terms are the amount by which the uplink was delayed from the ideal timing. Substituting from the preceding three equations gives

$$\Delta t = \Delta t_0 + \tau_{\rm meas}/2 - [\delta t_u + \delta t_d + \tau_{\rm rcvr2} + \tau_{\rm xmtr} + \delta t_{2u}\dot{T}_u(t - T_u - T_d) - \delta t_{2d}\dot{T}_d(t)]/2$$

where $\dot{T}_u, \dot{T}_d$ are the derivatives of the predicted delays $\hat{T}_u, \hat{T}_d$ with respect to coordinate time t. (The derivative terms were explicitly neglected in [6].) Recall that $\delta t_{2d}, dt_{2u}$ are the clock errors of the fine synthesizers. If these errors are nearly equal for the uplink and downlink synthesizers, then the derivative terms will nearly cancel. Otherwise, they introduce a time-varying correction to Δt even if the clock errors themselves are constant.

The analysis in this section has so far been generic, and should apply to most earth station implementations. We now consider, for the Green Bank station, the origin of the timing errors $\delta t_1, \delta t_2$; their magnitudes; and how well known they are. Referring again to Figure 2, let

$$\delta t_1 = \delta t_{\rm ref} - \tau_{\rm syn}$$

and

$$\delta t_2 = \delta t_{\rm ref} - \tau_{\rm DSP} - \tau_{\rm DDS}$$

where $\delta t_{\rm ref}$ is the effective timing error in the vertex reference signals (500 MHz, 10 MHz, and 1 Hz) with respect to true coordinate time UTC; for example, it can be taken to be the difference between a positive zero-crossing of the 1 Hz reference and the nearest integer UTC second mark. (Recall that this must also be a zero crossing of the 10 MHz and 500 MHz references, since all are synchronized in hardware. Very small and constant delays within this hardware are neglected.) Then $\tau_{\rm syn}$ is the effective delay in the coarse synthesizer's electronics; $\tau_{\rm DSP}$ is the delay in the DSP's firmware (discussed later); and $\tau_{\rm DDS}$ is the effective delay in the fine synthesizer's electronics.*

These terms can be further broken down as follows.

Let

$$\delta t_{\rm ref} = \delta t_{\rm clk} + \Delta\tau_{\rm maser} - \tau_{\rm cable1} - \tau_{\rm div} + K_1 T_{10}$$

where $\delta t_{\rm clk}$ is the decoder clock error after removing the change in maser cable delay since some reference epoch, $\Delta\tau_{\rm maser} = \tau_{\rm maser}(t) - \tau_{\rm maser}(t_{\rm ref})$; $\tau_{\rm cable1}$ is the delay in the 10 MHz antenna cable; $\tau_{\rm div}$ is the delay in the divide-by-10^7 box; K_1 is an integer; and $T_{10} = 1/(10\,{\rm MHz})$ is the period of the 10 MHz reference. The last term accounts for the 1 Hz antenna cable's delay, in that the signal is re-synchronized to the 10 MHz reference upon receipt at the vertex (we think that $K_1 = 1$ or 2, but this needs to be checked).

The last equation relates the decoder clock error, of which we already have some knowledge, to the vertex electronics timing. Notice that the cable delay and the divider delay do not actually affect the vertex timing, but are included because they affect the decoder clock. In practice, these terms should be very stable; our procedure for determining the decoder clock error, since it involves comparisons with GPS over many days, ensures that small variations in the cable or divider delays will average out. Also, we have by definition excluded the measured maser cable delay from $\delta t_{\rm clk}$. Therefore, we can treat the cable and divider delays as constant (estimated at 243 ns) and the variation in $\delta t_{\rm clk}$ is purely intrinsic to the maser (being the integrated effect of its absolute frequency error). The only other time-variable term in $\delta t_{\rm ref}$ is the maser cable, which is measured.

Next, $\tau_{\rm syn}$ arises entirely in microwave electronics; its value is unknown, but it is believed to be very small (a few ns) and very stable, arising mostly in short cables. Although it may be different for different link frequencies and may vary slightly with environmental stresses like temperature changes, we treat it as a constant for any one satellite.

Next, $\tau_{\rm DSP}$ is a firmware delay in the digital signal processor. A critical feature of the design is that this processor's clock is the 10 MHz reference, rather than an independent crystal. Because of this, firmware delays are precisely controlled and known; they are always an integral number of 10 MHz cycles. In this

* Throughout these notes, variables like δt_x represent timing errors such that $t_{\rm apparent} = t + \delta t_x$; and variables like τ_y represent delays (generally positive) in hardware.

case, the 1 Hz reference causes an interrupt to the processor, and this is used to synchronize the processor-generated 5 kHz signal that updates both DDSs. The delay is equal to the interrupt latency plus several cycles of code. (In firmware version 2.23, it is 40 cycles or $\tau_{DSP} = 4000\,$nsec.) The very same signal is used to sample the residual phase, as discussed later.

Next, τ_{DDS} represents the latency in each DDS chip, plus the delay in the following low pass filter. The DDS clock is precisely the 10 MHz reference multiplied by 10 (i.e., 100 MHz), and the latency (from receipt of an update pulse from the DSP to producing the new output frequency) is a fixed number of clock cycles (for the AD9850 chip now used, it is 17 cycles or 170 nsec). The filter is a 5-pole Chebychev with $-3\,$dB point at 42.5 MHz; its typical delay is 120 ns, but this is somewhat a function of the DDS frequency. For any one satellite, the frequency varies only by the Doppler shift, which is a small fraction of the filter bandwidth; we therefore treat this delay as a constant for each satellite. Nevertheless, it is fairly large and its uncertainty is the largest contribution to error in our knowledge of the absolute time error δt_2.

Finally, τ_{xmtr} consists of any delay in the phase locked loop that combines the signals from the coarse and fine synthesizers (shown below to be negligible) plus the delay in a short cable (about 1 m, or 4.3 nsec) and in the free-space optics to the reference point (about 20 nsec).

3.5 PLL acceleration errors

The synthesis processes for both the uplink and downlink involve phase locked loops. Each loop is a second order servo with high d.c. gain, so it can maintain essentially zero phase error for any frequency offset, provided that it is in lock. However, such loops have phase errors proportional to the rate of change of frequency (or the second derivative of the desired phase). The error is given by $\Delta\theta = \dot{\omega}/\omega_n^2$, where ω_n is the loop natural frequency [10]. The maximum frequency rate for VSOP is 450 kHz/s, and the smallest loop natural frequency in our system is 5 kHz, giving a worst-case phase error of .003 rad. There are three PLLs in the system, two at microwave frequencies and one near 600 MHz; for the latter, this phase error corresponds to a timing error of 0.7 ps. I conclude that in all cases the PLL acceleration error is negligible, so we make no attempt to correct for it.

3.6 Phase processing delay

The measurement of residual delay τ is accomplished in an analog phase detector whose outputs are proportional to $\sin\omega_d\tau$ and $\cos\omega_d\tau$. The latter signals are processed by analog low pass filters (1.8 kHz at $-3\,$dB) and are then sampled at 5 kHz by the same DSP-generated signal used to update the DDSs. They are then converted to units of phase, cycle slips are resolved, and the result is processed by a digital FIR low pass filter (5 Hz cutoff) and simultaneously decimated to 10 Hz sampling rate. Finally, during each UTC second the 10 samples from the preceding second are recorded by the station computer and labeled with the (integer) time at the end of the preceding second. Each such sample will be used to compute a corresponding value of Δt. The DeltaT file spec [1] requires that the latter values be labeled with the time of arrival of the corresponding signal at the antenna reference point. To do this, we must account for all of the processing delays. This is done in Appendix A.

The dominant delay is in the digital filter. This filter is implemented in the DSP, whose clock is the maser-derived 10 MHz reference. This processor (ADSP2101) also executes exactly one instruction per clock, so by counting instructions it is possible to know precisely the value of this delay. It therefore contributes little to the *error* in the total delay. The latter is dominated by knowledge of the delay through the analog LPF. Nevertheless, the total estimated error is $1\,\mu$s, which is well under the $7\,\mu$s accuracy requirement calculated in [8].

4. Effects External To The Earth Station

4.1 Absolute downlink delay.

A large component of Δt_0, the constant part of Δt, is the estimated propagation delay from satellite to earth station at the clock setting event. This is simply the sum of the estimated vacuum delay (from the predicted orbit) and the excess delay due to the medium (troposphere and ionosphere): $\hat{R}(t_0)/c + \tau_{trop}(t_0) + \tau_{ion}(t_0)$. The error in this estimate will almost always be dominated by the error in the first term, since the predicted orbit will normally have errors of several hundred meters (up to 1000 m), or about $1\,\mu$s. The total excess path length through the troposphere is about 84 ns at 5 deg elevation and sea level, and the total delay through the ionosphere at the frequencies of interest is considerably less. Thus, even if the medium terms are neglected entirely, there is little effect on the error.

4.2 Medium non-reciprocity

The effects involved are discussed in detail in [3], and algorithms for computing corrections are discussed in [4] and [5]. Here I emphasize only that the dominant effect is the motion of the earth station–satellite vector through the medium during the round-trip time, so that the two directions traverse different media; of much less importance is the difference in delay through the same medium at the different uplink and downlink frequencies, due to ionospheric dispersion. Thus, the effect is primarily geometrical and is largest near the horizon. The correction to Δt is half of the delay difference; this can be as large as 350 ps for the troposphere and several hundred ps for the ionosphere at VSOP's frequencies (see [4] and [5]), and rapidly changing.

For the Green Bank station, we will compute and apply these corrections using spatial models of the troposphere and ionosphere. For each value of Δt, the uplink and downlink delays are computed separately; half of their difference is included in Δt. The troposphere model is a stratified extrapolation of meteorological measurements made in real time near the station, as detailed in [5]. The ionosphere model is not yet finalized, but will probably be based on recent global measurements rather than local real-time data.

4.3 Spacecraft turn-around delay

The uplink signal is received at the spacecraft's link antenna (a small paraboloid), sent by coax and waveguide to the transponder, and used to lock a PLL. The locked oscillator is used to synthesize the downlink signal, which is then sent by a similar path back to the link antenna for radiation. As long as this round-trip is reciprocal with respect to the phase detector in the transponder, it has no effect on Δt and the corrected time refers to the signal at the transponder's phase detector. The measured residual delay will include effects from the fact that the link antenna is not at the spacecraft's center of gravity (to which the predicted orbit refers), and that the total delay through the spacecraft is different from the free-space delay between the antenna to the CG. Note that this is time-varying, since the projection of the antenna–CG vector onto the link direction changes along the orbit, as the link antenna tracks the earth station.

But there is likely to be some non-reciprocity because the transmission lines for VSOP and RadioAstron are each at least 2 m long and somewhat dispersive; and because the uplink receiver and downlink transmitter have different filters and hence different delays. If the non-reciprocity is perfectly stable, then its consequence is a fixed offset in time; this should be small (maybe a few ns), so it should be of no significance. Since we have no data to the contrary about either satellite, we must for now assume that the non-reciprocity is stable. We therefore plan to make no corrections for it.

4.4 Spacecraft PLL acceleration error

Each spacecraft uses a second order PLL with high d.c. gain, just like those in the earth station. In each case the loop bandwidth is about 1 kHz. By calculations similar to those in 3.5, it can be shown that this will introduce negligible acceleration error, especially considering that the phase acceleration will normally be held to very low values by the uplink compensation.

5. Computation Of Values Needed In DeltaT File

Based on all the results previously derived, I summarize here the expressions for the various values needed in the DeltaT file.

5.1 Constant part of Δt

$$\begin{aligned}\Delta t_0 = {} & t_{\mathrm{dec}}(t_0) + \delta t_{\mathrm{dec}}(t_0) - t_{\mathrm{fmtr}} && \text{(clock)} \\ & - (\tau_{\mathrm{rcvr1}} + \tau_{\mathrm{cable}} + \tau_{\mathrm{dec}}) && \text{(station hardware)} \\ & - [\hat{T}_d(t_0) + \tau_{\mathrm{trop,dn}}(t_0) + \tau_{\mathrm{ion,dn}}(t_0)] && \text{(link).}\end{aligned}$$

The instantaneous clock error is $\delta t_{\mathrm{dec}}(t_0) = \delta t_{\mathrm{clk}}(t_0) + \Delta\tau_{\mathrm{maser}}$, where $\delta t_{\mathrm{clk}}(t_0)$ is the model derived from long-term clock monitoring; the change in maser cable length has been excluded from this model and is added back at each epoch for which the clock correction is needed. In the present context the maser cable length is a very small effect, but it is important in the time correction values, to be considered shortly. By far

Table 2: Numerical Values of Parameters, VSOP Case

Symbol	Description	Value
τ_{xmtr}	Uplink transmitter delay	24 ns
τ_{rcvr1}	Downlink receiver delay, data link	81 ns
τ_{rcvr2}	Downlink receiver delay, timing link	135 ns
τ_{cable}	Downlink signal antenna cable delay	128 ns
τ_{dec}	Decoder delay	302.3 ns
τ_{cable1}	10 MHz antenna cable delay	143 ns
τ_{div}	1 Hz divider delay	100 ns
K_1	10 MHz re-sync delay, cycles	2
τ_{DSP}	DSP synchronization latency	4000 ns
$\omega_u/2\pi$	Nominal uplink frequency	15300 MHz
N_u	Uplink coarse multiplier	1550
$\omega_{2u}/2\pi$	Uplink nominal DDS frequency	-200 MHz
$\omega_d/2\pi$	Nominal downlink frequency	14200 MHz
N_d	Downlink coarse multiplier	1421
$\omega_{2d}/2\pi$	Downlink nominal DDS frequency	-10 MHz
$\tau_{synu,d}$	Coarse synthesizer delays (assumed)	0 ns
$\tau_{DDSu,d}$	Fine synthesizer delays (both)	195 ns
τ_{pd}	Phase detector filtering delay	1673.4828 ms

the largest error in Δt_0 (up to a few μs) is that in $\hat{T}_d$, the (vacuum) link delay calculated from the predicted orbit. The station hardware delay is assumed constant for a given satellite, but the decoder component τ_{dec} varies significantly according to the downlink data format. The total hardware delay is believed known to within 20 ns now, and this may be refined later.

5.2 Time of the clock setting event

The DeltaT file needs to include the time of clock setting referred to the signal arrival time at the earth station antenna:

$$t_0 = t_{dec} + \delta t_{clk} + \Delta\tau_{maser} - (\tau_{rcvr1} + \tau_{cable} + \tau_{dec}).$$

Note that t_0 is the value known as "GND_TIME" in the file specification [1], and that t_{fmtr} is the value known as "TAPETIME."

5.3 Time correction values

$$\Delta t(t) = \Delta t_0 + [\tau(t) - \tau(t_0)]/2 + \Delta t_{medium}(t) - \Delta t_{medium}(t_0)$$

where the clock corrections have been included in τ via

$$\tau = \tau_{meas} - [\delta t_d + \delta t_u + \tau_{rcvr2} + \tau_{xmtr} + \dot{T}_u(t - T_u - T_d)\delta t_{2u} - \dot{T}_d(t)\delta t_{2d}],$$

with

$$\delta t_{d,u} = (N_{d,u}\omega_{10}\delta t_{1d,1u} + \omega_{2d,2u}\delta t_{2d,2u})/\omega_{d,u}$$
$$\delta t_{1d,1u} = \delta t_{ref} - \tau_{synd,u}$$
$$\delta t_{2d,2u} = \delta t_{ref} - \tau_{DSP} - \tau_{DDSd,u}$$
$$\delta t_{ref} = \delta t_{clk} - \Delta\tau_{maser} - \tau_{cable1} - \tau_{div} + K_1 T_{10};$$

and where

$$\Delta t_{medium} = (\tau_{trop,up} - \tau_{trop,dn} + \tau_{ion,up} - \tau_{ion,dn})/2.$$

Our current best estimates of the numerical values of all the parameters, evaluated for VSOP, are given in Table 2. Quantities not covered there are determined from real-time measurements or models. Note that some frequencies are negative.

The estimates in Table 2 result in values of δt_1, δt_2, and δt that are actually the same for both uplink and downlink. Most of the terms involved (δt_{ref} and τ_{DSP}) are indeed common to both uplink and downlink in our implementation. But two terms, τ_{syn} and τ_{DDS}, arise in physically separate hardware and could be different. As a practical matter, we neglect τ_{syn} completely; it should be small (at most 1 cycle of a fixed microwave frequency of 7 to 15.5 GHz, < 142 ps) and constant (subject to small variations with equipment temperature) for a given satellite. But τ_{DDS} is larger and more variable. As discussed in section 3.4, it consists of a digital latency of 170 ns, identical for uplink and downlink, plus the delay in an analog filter. The latter is theoretically 20 to 31 ns, a function of the DDS frequency setting; that setting is substantially different for uplink and downlink, and varies slightly because of Doppler tracking. These considerations allow some simplification in the expression for the corrected residual delay:

$$\begin{aligned}\tau = \tau_{meas} - 2\delta t_{ref} - \tau_{rcvr2} - \tau_{xmtr} - (\omega_{2d}/\omega_d + \omega_{2u}/\omega_u)\tau_{DSP} \\ - [\dot{T}_u(t - T_u - T_d) - \dot{T}_d(t)](\delta t_{ref} - \tau_{DSP}) \\ + [\dot{T}_u(t - T_u - T_d)\tau_{DDSu} - \dot{T}_d(t)\tau_{DDSd}].\end{aligned}$$

Finally, the following arguments allow the terms involving $\dot{T}_d, \dot{T}_u$ to be neglected for the Green Bank station. The maximum clock error δt_{ref} will be kept below a few μs by operating procedures, and the maximum difference between uplink and downlink delay rates is approximately $\ddot{T} \cdot (T_u + T_d)|_{max} \approx 3 \times 10^{-10}$ (for VSOP, less for RadioAstron); this makes the first bracketed term less than 2×10^{-15} s, which is completely negligible. For the final bracketed term, note that $|\dot{T}_u \approx \dot{T}_d| < 3 \times 10^{-5}$, and that the maximum difference $|\tau_{DDSu} - \tau_{DDSd}|$ is about 11 ns (see discussion in preceding paragraph). Then the last bracketed term is at most 0.33 ps. Although this is time varying and correctable, it is small enough to neglect at the maximum observing frequency of 22 GHz. If the difference between the uplink and downlink clock errors were substantially larger, a correction would be required.

We are thus left with time-varying terms in Δt due to:

- τ_{meas}, residual delay, measured in real time;
- δt_{clk}, clock error, from long-term clock monitoring;
- $\Delta\tau_{maser}$, maser cable length variation, measured in real time;
- $\tau_{trop,up} - \tau_{trop,dn}$, troposphere non-reciprocity, from model based on real time local measurements; and
- $\tau_{ion,up} - \tau_{ion,dn}$, ionosphere non-reciprocity, from model based on recent global measurements.

5.4 Time stamp on first Δt value

This is the "UTC_DATA" value required in the header of the DELTA_T table of the file [1].

$$t_1 = t_{phr} + \delta t_{ref} + \tau_{DSP} - (\tau_{rcvr1} + \tau_{pd} + 1.0\,\text{s})$$

where t_{phr} is the time recorded by the real time software in the phase residuals file (".phr" file) for the measurement corresponding to the first Δt value; τ_{pd} is the phase detector filtering delay; and the other terms have been considered earlier. The terms in parentheses represent the phase processing delay, τ_ϕ. The 1.0 s term is a software artifact, but is exact. The largest term is the digital filter delay (part of τ_{pd}, 1.6806000 s in firmware version 2.23), but it is known very accurately. Additional details and numerical values are given in Appendix A.

In practice, δt_{ref} can be neglected here because we can always keep it less than a few μs by careful maintenance of the clock. It was shown in [8] that time stamp errors cause a worst-case correlator timing error of .07 ps/μs, so time stamp errors of a few μs are easily acceptable. Therefore, the DeltaT file time stamps are related to the phase residuals file time stamps by an additive constant. This constant is rather large (> 2.6 s) in our implementation, but it is believed to be stable (±130 ps) and known very accurately (±1 μs).

5.5 Correspondence between Δt time stamps and tape time (TAPETIME table)

The specification [1] allows for an optional TAPETIME table, where each entry consists of two numbers: a sample of earth station UTC, and the tape time (formatter clock reading) of the data sample received at the antenna reference location at that time. These values may be calculated as follows.

$$\begin{aligned}t_{ground} &= t_{dec} + \delta t_{clk}(t) + \Delta\tau_{maser} - \tau_{rcvr1} - \tau_{cable} - \tau_{dec} \\ t_{tape} &= t_{fmtr} + [C_f(t) - C_f(t_0)]/f_f\end{aligned}$$

where C_f is the reading of the decoder's frame counter and f_f is the downlink frame rate (e.g., 200 Hz for VSOP). As before, t_{fmtr} is the formatter's clock setting at the initialization epoch t_0.

6. Dropouts And Data Blanking

During a planned tracking pass, it is possible that problems may occur in maintaining the spacecraft timing. There may be a temporary loss of signal to or from the spacecraft; there may be a "glitch" or brief malfunction of the earth station equipment; or there may be other, similar events. Here I distinguish three sharply different cases of such disruptions.

(a) Degraded data. If analysis of monitor data shows that the accuracy of the measured timing residual τ_{meas} or of any of the corrections is degraded, so that the expected accuracy is not likely to be obtained, the DeltaT file will nevertheless contain our best estimate of the time corrections. For example, failure of the weather station will cause us to revert to simplified atmospheric model that does not depend on real-time data. There is no provision for flagging of degraded data within the DeltaT file, but information about the occurance of anomalies is available in the Performance Log file [11].

(b) Bad data. Complete loss of signal from the satellite, as well as certain other failures, means that no useful estimate of Δt can be provided. During such periods, the samples of Δt in the file will be replaced by a special number (the IEEE representation of $-\infty$) to indicate that no data is available, as provided in the specification. For the NRAO station, the decision to "blank" some results in this manner is made by the offline software on the basis of monitor data from the station log, using algorithms to be specified elsewhere.

If a dropout is of sufficiently short duration, and if re-acquisition of signal occurs smoothly, then it should be possible (very nearly) to recover the timing that would have been obtained if the dropout had not occurred. This is because the station hardware is designed to avoid any ambiguities in timekeeping, and because the oscillators representing satellite time will "flywheel" at nearly the correct frequency in the absence of a signal. Upon re-synchronization of the decoder, the timing may have changed by an integral number of bits, but this number is known. (A change of less than one downlink frame is handled in hardware, with no effect on tape time. A change of a few frames is detectable from the frame count in the header.) Upon re-locking of the demodulator to the downlink carrier, an ambiguity of 1/4 cycle of downlink phase occurs (17.6 ps at 14.2 GHz), but this fairly small so we will ignore it. For the NRAO station, it is the job of the real-time software to decide whether such a recovery will be possible, using algorithms to be specified elsewhere; if so, then it is the job of the offline software to determine any additional corrections needed in Δt, and to include these in the DeltaT file.

(c) Re-initialization. If a dropout is too long, then recovery of the pre-dropout timing may not be possible. It is the job of the real-time software to decide whether this is the case, and if so to re-initialize the timing automatically. This means that the formatter (tape) clock will be reset, just as if we were starting a new tracking pass. The constant timing offset Δt_0 will be re-determined, using the predicted orbit at the new clock setting event. Generally, this means that the new timing has a large and unknown offset from the old timing. Accordingly, the offline software will record this in the DeltaT file by starting a new DELTA_T table.

The sharp distinction between cases (b) and (c) is this: in (b) we have information that allows the old and new timing to be connected; in (c) we do not. Usually case (b) will include a period of bad (blanked) timing corrections between the old and new timing, but this is not guaranteed. (For example, there could be a loss of the data link but not of the timing link, so good values of Δt are available continuously but a frame slip has occurred when the data link is recovered.) Case (b) may or may not produce a discontinuity between the new timing and a smooth extrapolation of the old timing, depending on whether or not a frame slip has occurred. In case (c), no attempt should be made to connect old and new timing; correlators should handle it as if a new tracking pass has started, just as if we had a handoff to a new tracking station.

Acknowledgment

This document benefited considerably from discussions with Anthony Minter, who pointed out various errors and omissions in an earlier version.

REFERENCES

[1] L. D'Addario and G. Langston, "Time Corrections File Interface," NRAO Specification A34300N004E, 95/04/19.

[2] L. D'Addario, "Time synchronization in orbiting VLBI." *IEEE Trans. Instr. & Meas.*, 40:584-590, 1991.

[3] L. D'Addario, "Post-Pass Processing of Two-Way Timing Measurements (revision of OVLBI-ES Memo 38)," OVLBI-ES Memo 64, 96/07/08.

[4] A. Minter, "Ionospheric and Tropospheric Corrections for the DeltaT File," OVLBI-ES Memo 62, 96/05/14.

[5] A. Minter, "Estimating the Tropospheric Correction for DeltaT: Data from Weather Station Measurements," OVLBI-ES Memo 63, 96/05/25.

[6] L. D'Addario, "Effects of Earth Station Clock Error on Correlator Time Corrections," OVLBI-ES Memo 65, 96/08/10.

[7] L. D'Addario, "Real-Time Processing of Downlink Phase Measurements," OVLBI-ES Memo 25, 92/03/11.

[8] L. D'Addario, "Correlating Orbiting Radio Telescope Data With Time Corrections," OVLBI-ES Memo 55, 95/03/22.

[9] L. D'Addario, "Reference Transmission System: Design and Test Results," OVLBI-ES Memo 48, 94/07/02.

[10] F. Gardiner, *Phase Lock Techniques*. New York:Wiley, 1979.

[11] L. D'Addario, "Spacecraft/Station Performance Log Interface," NRAO Specification A34300N008B, 95/10/06.

[12] L. D'Addario and R. Escoffier, "Measurement of the delay through the Decoder." [ftp.gb.nrao.edu:ovlbi/doc/decDelay.txt]

APPENDIX A: CALCULATION OF THE PHASE PROCESSING DELAY

After the downlink wavefront arrives at the earth station's reference location (defined to be the point of intersection of the azimuth and elevation axes of the antenna), it is processed in a series of steps that culminate in either (a) recording on the wideband tape, or (b) measurement of the carrier phase relative to the prediction of the orbit file, for the data and timing downlinks, respectively. In each case, the final result is recorded with a time stamp that differs from the time of arrival at the reference location. Here we consider the residual phase measurement.

Processing	Delay	Stability, ns
Antenna optics:		
Vertex to focus = 5.08m		
Vertex to subr to feed (est) = 8.0m*		
Vertex to axis int (est) = 2.0m*; net=6.0m	20 ns	0.0001
Front End electronics:		
Bandwidth 2 GHz	0.5 ns	0.001
Cables to bin, 1m*, v/c=0.77	4.3 ns	0.001
Downconverter electronics:		
Amplifiers and mixers, about 500 MHz BW	2.0 ns	0.004
Filter, 550 MHz BW, 6 poles	5.5 ns	0.008
Second conversion (IRM):		
IF filter, 10 MHz* bandwidth	100.0 ns	0.01
Phase detector:		
Analog filter, 1768 Hz LP, 8 pole Butterworth	282.8 μs	100
Digital FIR filter, 5 Hz LP, delay=8403 samples at 5000 samp/s	1680.6 ms	1
Digital FIR filter, firmware time offset	−7.4 ms	1
Software delay: 1st item in record to timestamp	1000.0 ms	0
Timekeeping (phase detector sampling to master station time):		
Cable, vertex to eqpt bldg (10MHz) 30m* 0.7*	−143 ns	0.020
Logic delays in 1Hz divider (est)	−100 ns	0.1
Recapture by 10MHz at vertex	+200 ns	0
DSP firmware (ver. 2.23) synchronization latency	−4000 ns	0

* value needs to be checked.

Overall, the available data indicates that we have, for the phase measurement:

Total delay	2.6734788893 s
Estimated uncertainty	0.000001
Estimated rms stability	0.00000013

APPENDIX B: SOFTWARE NOTES

1. The constant part of $\delta t_{2d,u}$ can be largely eliminated by compensating for it in OrbitCo. Each line of the orbit file would then be computed for a true UTC slightly *later* than its time stamp by an amount $\tau_{\rm DSP} + \tau_{\rm DDS} = 4290\,{\rm ns}$.

2. The long-term clock error model can be easily maintained in an ASCII file where each line contains reference epoch $t_{\rm ref}$ followed by coefficients of the Taylor series expansion of $\delta t_{\rm clk}$ about $t - t_{\rm ref}$. Two coefficients (offset and rate) should be sufficient if the model is updated every few days, but client programs should allow for the possibility of several more terms whose coefficients have default values of zero. Resetting the decoder clock (which should happen only rarely) or unplanned glitches will force a new line in the file. In the absence of such events, successive lines should represent a continuous function.

3. Considering the accuracy of our hardware implementation, and assuming adequate stability, it has been shown in this report that the corrected residual delay reduces to

$$\tau = \tau_{\rm meas} - 2(\delta t_{\rm clk} - \Delta\tau_{\rm maser}) + {\rm const.}$$

However, in order to provide flexibility in case actual errors are found to be larger than expected, the offline program for producing the DeltaT file should include provisions for:

(a) different values of the "constant" term each run;

(b) additional terms proportional to the coarse synthesizer delays $\tau_{{\rm syn}u,d}$, which might later be found to have a known dependence on temperature in Rack V (or some other parameter) and hence be correctable; and

(c) additional terms $\dot{T}_u \tau_{{\rm DDS}u}(\omega_{2u}) - \dot{T}_d \tau_{{\rm DDS}d}(\omega_{2d})$.

For (b) and (c), I suggest providing subroutines to return $\tau_{{\rm syn}u}$, $\tau_{{\rm syn}d}$, $\tau_{{\rm DDS}u}(\omega_{2u})$, and $\tau_{{\rm DDS}d}(\omega_{2d})$. Initially, these can be dummy functions, always returning zero; later they can be made more sophisticated if necessary.