A Demonstration of DSN Clock Synchronization by VLBI

W. J. Hurd
Communications Systems Research Section

A prototype of a semi-real-time system for synchronizing the DSN station clocks by radio interferometry was successfully demonstrated on August 30, 1972. Synchronization accuracies as good as 100 ns rms were achieved between DSS 11 and DSS 12, both at Goldstone. The accuracy can be improved by increasing the system bandwidth until the fundamental limitations due to baseline and source position uncertainties and atmospheric effects are reached. These limitations are about 10 ns for transcontinental baselines.

I. Introduction

It is well known that the clocks at widely separated antenna ground stations can be synchronized by the techniques of very long baseline interferometry (VLBI). There are two reasons that an operational VLBI time synchronization system may be desirable for the DSN. First, accuracies an order of magnitude better than those currently attained by the moon-bounce system may be attainable with little initial investment and with operational costs that should be no higher than for the existing system. Second, VLBI may be the only operationally feasible method for achieving the 10–20 ns accuracies required for two-station tracking of deep space probes (Refs. 1, 2).

The time synchronization accuracy attainable by interferometry over very long baselines is fundamentally limited by the uncertainties in the differential time delay from the radio source to the antennas. These uncertainties, which increase with the baseline length, are caused by errors in the estimates of the source positions and antenna locations, and by the variable propagation delays in the atmosphere. It is anticipated that the antenna locations will soon be known to within about 1 m, and source position errors can be reduced to this same level by interferometry. The atmospheric effects depend on frequency in a known manner, and can be calibrated by receiving on two frequencies, say S- and X-band. The fundamental limitation of accuracy can probably thus be reduced to 10 ns or less for intercontinental baselines.

Until the fundamental limit is approached, the synchronization accuracy depends primarily on the utilized bandwidth, provided that the signal-to-noise ratio (SNR) is high enough for reliable detection. The experiment reported on here confirms the two most important results of Ref. 3: first, that reliable estimates can be achieved with a small enough amount of data (about 1 million bits) so that semi-real-time processing is feasible; and second, that the rms errors are less than 0.1 times the inverse system bandwidth, so that the fundamental limitations of accuracy can be achieved with system bandwidths of only about 10 MHz.

II. Description of Experiment

As a first step in demonstrating the feasibility of an operational system for DSN clock synchronization by VLBI, an experiment was conducted on August 30, 1972, between the 26-m-diameter antennas at DSSs 11 and 12, both at Goldstone. The experiment was implemented using a minimum of special interfacing hardware in addition to standard DSN station equipment. The data were acquired and processed following the approximate maximum-likelihood method of Ref. 3.

A simplified block diagram of the experiment is shown in Fig. 1. At each station, the received signals were demodulated in two phase quadrature channels, filtered, quantized to 1 bit, and buffered into a TCP computer. Besides the receivers, the TCP computers are the major portion of the system. The special equipment for the experiment consisted of the two-channel demodulators, the filters, limiters, and samplers, and the buffers from the sampler to the TCP computers. All of this equipment was contained in one small chassis for each station, plus cables to interface to the computers.

The experiment procedure was to initiate sampling at the same time at each receiver, according to the station master clocks, and to fill the TCP computer memories with data at the highest possible sampling rate. The computer speeds limited the data rate to 500 kbps, or 250 kbps per channel, so that the system bandwidth was limited to 250 kHz. Furthermore, the maximum number of samples which could be taken at this rate was limited by the memory sizes to approximately 320,000 bits. In an operational system, the data could be transmitted directly from the computers to JPL over the high-speed data lines and processed within a few minutes in the Network Control System (NCS) or other computers. In the experiment, however, real-time operation was not required, but instead it was desired to make a number of independent estimates of time sync using each of several radio sources. Therefore, the data were written onto magnetic tape and processed later on a Sigma 5 computer at JPL. Five different radio sources were observed, with a total of 504 batches of data taken at 10-s intervals.

III. Principal Results and System Implications

The desirability of an interferometry time sync system for the DSN depends on the ability to achieve reliable results with a reasonable amount of data. This, in turn, depends on the availability of radio sources with enough correlated flux, i.e., with enough electromagnetic flux which appears to be from an ideal point source when viewed by the long baseline interferometer. In this section, we set a standard for the required source intensity for various system configurations based on experimental and analytical results, and show that adequate sources are available to result in an operationally feasible system.

The experimental results were limited by the system parameters of two 26-m-diameter antennas with temperatures of 16.3 and 37 K, the 250-kHz bandwidth, and the 3.2×10^5 bit buffer size. For the three strongest radio sources, rms processing errors of 96, 228, and 403 ns were achieved, in close agreement with theory. The results for the weakest of these sources, with an estimated correlated flux of 4.6 fu, are most significant for two reasons: first, the estimates were reliable even though the SNR was somewhat lower than the desired minimum, and second, the results were in close agreement with theory, indicating that the theory does not break down until the SNR is reduced below this level.

Based on both the theoretical and experimental results, we conclude that a source intensity of 5.5 fu would have been adequate to reliably achieve an rms error of less than 0.4 μs, or less than one tenth of the inverse of the system bandwidth. Whenever possible, higher accuracies should be achieved by increasing the system bandwidth and not the amount of data or the SNR, both because few sources have more than 2–3 fu of correlated flux over long baselines and because increasing the amount of data is expensive in terms of buffer storage, computer time, and ground communications facility (GCF) usage.

Table 1 presents the source intensities required to achieve the reliable performance level of one-tenth of the inverse system bandwidth for various antenna sizes and receiver noise temperatures in the DSN. Two buffer sizes are considered: the 0.32 Mbits usable in the TCP computers, and 1.0 Mbit, which is a practical size if special-purpose memories are used for wider bandwidths. In utilizing Table 1, one should keep in mind that the system temperatures increase at low elevation angles, so that the required fluxes might increase by a factor of about 1.5.

The availability of known radio sources was surveyed using Ref. 4 and a computer program for mutual visibility devised by J. G. Williams of the Tracking and Orbit Determination Section. Considering sources to be jointly visible only when the altitude from both stations is 10 deg or greater, there is always at least one source

of 1.3 fu or stronger visible by the station pairs at Goldstone and Spain, Goldstone and Australia, and Spain and South Africa. Sources of 2.0 fu are available for most of the day, and sources of 3–6 fu are normally visible for at least a few hours each day. The source 3C454.3, which is sometimes as strong as 6.38 fu (Ref. 4), is visible to each of the above pairs for at least 3 h a day, but unfortunately it has at other times been observed to be considerably weaker, and similar variations occur with some of the other strong sources. It is therefore not desirable to base a system on the few strongest sources.

Considering both the source intensities required and the availability, it is safe to say that station pairs with at least one 64-m-diameter antenna can be synchronized at will to within one-tenth of the inverse system bandwidth with 1.0 Mbit of data. That is, there would be little if any operational restriction as to time of day due to lack of mutual visibility of adequate sources. Synchronization of two 26-m-diameter antennas could be accomplished with some restrictions on time of day, or by using more data. It is important to note that the amount of data used is not restricted by the high-speed buffer size, but convenience is sacrificed if it becomes necessary to fill the buffer several times, storing the data on magnetic tape between fills, and then to transmit a larger amount of data to the central computer for processing.

We conclude that a system with 1.0-Mbit buffers would be operationally feasible. It would be less restricted than the X-band moon-bounce system, for which moon visibility restricts the time of day, and even the time of year for two northern hemisphere stations.

IV. Estimation Procedure

The estimates of time synchronization error were made using the approximate maximum-likelihood method (Ref. 3). This method is distinguished from normal cross-correlation methods in that the cross products are multiplied by appropriate weighting functions before being summed or envelope detected. This weighting accounts for changes in clock offset during the measurement time, and provides an optimum method for resolving the time estimates to greater accuracy than the time between samples.

A functional diagram of the estimation procedure is shown in Fig. 2. The ith samples in the phase quadrature channels after demodulation, filtering, and limiting are denoted by X_i and Y_i for DSS 11 and by Z_i and W_i for DSS 12. These signals have cross correlations which depend on i and on ρ, τ, δ, ω, and ϕ, where

$$\rho^2 = \frac{T_c^2}{T_{11}T_{12}} \tag{1}$$

is the product of the input SNRs, and

T_c = increase in system temperatures due to *correlated* flux from the source

T_{11}, T_{12} = total system noise temperatures at DSSs 11 and 12, including total source flux, correlated or otherwise

δ = difference in path length from the source to the two stations, in seconds (often called τ_g)

τ = error in clocks, or actual time difference between first samples at the two stations

ω = stopped fringe frequency, or apparent doppler difference after demodulation, in radians per second

ϕ = stopped fringe phase

The cross correlations are the expected values of the cross products:

$$E\{X_iZ_j\} = \frac{2}{\pi}\rho a_{ij}(\tau,\delta)\cos(j\Delta + \phi) \tag{2}$$

$$E\{X_iW_j\} = \frac{2}{\pi}\rho b_{ij}(\tau,\delta)\sin(j\Delta + \phi) \tag{3}$$

$$E\{Y_iZ_j\} = \frac{2}{\pi}\rho c_{ij}(\tau,\delta)\sin(j\Delta + \phi) \tag{4}$$

$$E\{Y_iW_j\} = \frac{2}{\pi}\rho d_{ij}(\tau,\delta)\cos(j\Delta + \phi) \tag{5}$$

where $\Delta = \omega \cdot 4\ \mu s$ and it is assumed that the timing is such that the cross products are uncorrelated except for $i \approx j$. The factor $2/\pi$ arises as a result of the hard limiting.

In general, for long baselines and measurement times, the fringe rate cannot be assumed constant, and $j\Delta$ must be replaced by a phase angle $\theta(j)$, which is known by the geometry. There is no significant difference in the estimation procedure. We assume here that the fringe rate is constant, for convenience and because this is valid for the short baseline of this experiment.

The estimation procedure is to maximize an estimator function G over assumed values for τ and ω. In calculating G, the stopped fringe rate ω is first normalized by subtracting out known quantities. Thus, the frequency variable becomes

$$f = \frac{1}{2\pi}(\omega - \omega_0) \tag{6}$$

where ω_0 is the *a priori* estimate of the stopped fringe frequency. Two factors contribute to ω_0: the fringe rate as calculated from the geometry, and the difference in local oscillator frequencies, or effective receiver center frequencies, at the two stations. The frequency f is the sum of the errors due to geometry and to oscillator instabilities, and the estimate of f is the estimate of these errors.

The steps in the estimation procedure for τ and f are:

(1) Assume a value of τ, say τ_k.

(2) Form all cross products whose cross correlations are nonzero for $\tau = \tau_k$.

(3) Multiply the cross products by the cross correlations for $\tau = \tau_k$, neglecting the sinusoidal terms, i.e., form $X_i Z_j a_{ij}(\tau_k, \delta)$, etc.

(4) Assume a value for f, say f_j.

(5) Evaluate $G(\tau_k, f_j)$.

(6) Maximize $G(\tau_k, f_j)$ over the region of uncertainty in f by looping back to step 4.

(7) Maximize $G(\tau_k, f_j)$ over the region of uncertainty in τ by looping back to step 1.

(8) The estimates $\hat{\tau}$ and $\hat{f}$ of τ and f are the values of τ_k and f_j which maximize G.

The distinguishing feature of this procedure is in weighting the cross products by their assumed τ-dependence before envelope detecting. This gives a natural and optimum method for resolving the estimate of τ to greater resolution than the time between samples, and for accounting for the changes in τ over the measurement time.

V. Properties and Examples of the Estimator Function

The statistics of the estimator function have been evaluated both analytically and by simulation (Ref. 3). We summarize here some of the key statistics, and then examine graphically some typical sample functions which were observed in the experiment.

The estimation procedure is considered to be reliable when the probability is high that the estimates are in the general vicinity of the correct values of the parameters, rather than being completely extraneous. This depends on the probability distributions of G for the correct and widely erroneous values of the parameters. Once the forms of the distributions are known, the performance can be well predicted by a figure of merit which we call the signal-to-noise ratio of the estimator. It is defined as the square of the difference in the means of G for the correct and incorrect values of the parameters, to the variance of G at the correct values. When G is normalized in the natural manner, its mean is unity for widely incorrect assumed values of the parameters, and is unity also when $\rho = 0$, so that

$$R = \frac{[E\{G(\tau,f)\} - 1]^2}{\text{var}\{G(\tau,f)\}} \tag{7}$$

The estimator SNR varies approximately as ρ^2, i.e., as the product of the input SNRs, or alternatively, as the square of the source flux density. For the particular filtering and sampling method used, it is given by

$$R = \frac{1}{2}\frac{r}{1 + \frac{1}{2r}} \approx \frac{r}{2} \tag{8}$$

where

$$r = 0.267\rho^2 N \tag{9}$$

and N is the number of samples in each channel at each receiver. Since the system bandwidth is the inverse of the time between samples in one channel, N is also the system time–bandwidth product.

Estimation will be reliable whenever R exceeds about 10, because the maximum value of G will almost always occur in the vicinity of the correct values of τ and f unless the initial uncertainty in these parameters is large. For example, when the initial uncertainty in f is negligibly small, the number of independent values of G which must be calculated is approximately equal to the time uncertainty times twice the system bandwidth. For the 250-kHz bandwidth of this experiment, time uncertainties of ± 10 to ± 100 μs would require calculation of only 10 to 100 independent values of G. It can be seen from

the curves of Ref. 3 that, for these uncertainties, the results would be reliable about 98–99% of the time with $R = 10$.

The resolution of the estimates depends on the peakedness of G more than on R. An approximation to the rms error in the estimation of τ is presented in Ref. 1, and is

$$\sigma_\tau = \frac{0.79T}{\rho N^{1/2}} \tag{10}$$

where T is the time between samples in one channel, or the inverse system bandwidth. In terms of R,

$$\sigma_\tau \approx \frac{0.289T}{R^{1/2}} \tag{11}$$

so that $R = 10$ is sufficient to reduce the rms error to less than $0.1T$ as well as to result in reliable estimation.

Insight into the capabilities of the estimator function to resolve time and frequency can be gained by studying the function at high signal-to-noise ratios. Figure 3 shows a plot of an actual sample function of $G(\tau_k, f_j)$ observed for a fairly high-intensity source, 3C279, with R estimated to be 24.8. The maximum of G is 52.844 and occurs for $f_j = -0.20$, $\tau_k = 40.97$, so that these are the estimates $\hat{f}$ and $\hat{\tau}$ of f and τ. In the time domain, G is nominally symmetrical, and decreases to half its maximum in under ± 2 μs and to approximately zero in ± 4 μs. In the frequency domain, G is also nominally symmetrical about the actual value of f, although this is not apparent from the sample function because the maximum did not occur at $f_j = 0$. The measurement time of the experiment was $NT \approx 0.64$ s, and the effective bandwidth of G is slightly less than the inverse of this time. It is observed that for different f_j, the maximum of G occurs at very close to the same value of τ_k. This implies that it may be unnecessary to maximize over f_j when only estimates of τ are required, provided that the initial uncertainty in f is small compared to $1/NT$, say less than $\pm 0.1/NT$.

The performance of the estimator when the noise is significant is illustrated in Figs. 4 and 5. Each presents three sample functions from different realizations of the experiment, with the time dependence shown only for the frequency variable fixed at the nominal value, $f_j = 0$. Figure 4 is for a weaker source, 4C39.25, with R estimated to be 3.81, which is significantly below the suggested design value of 10. In one of the three cases, the maximum of G occurs near $\tau_k = 21$ μs, far removed from the true value, which is near 41 μs. Extraneous results like this occur frequently at these low SNRs. Figure 5 is for source P1127-14, with R estimated to be 8.20, which is only marginally below the design point of 10. Fairly wide variations in the maximum value of G occur at this SNR, but no extraneous maxima were observed in the 72 sets of data taken for this source.

VI. Detailed Results

A total of 504 sets of data were taken using five different radio sources, and independent estimates of the time and frequency differences at the two receivers were made for each set of data. The most important results are the means and standard deviations of the estimates of τ and f as a function of the estimator SNR, R. In order to present these results, it was necessary to estimate R from the data. The method for estimating R is presented later.

A. Joint Estimation of τ and f

Table 2 presents the results of the joint estimation of τ and f for the five sources. The statistics are based on 144 independent estimates for each of the two strongest sources, 3C273 and 3C279, and on 72 cases for the other sources. Since the true values of τ and f were not known, it was not possible to compute the actual rms errors, so the standard deviations were estimated from the data using the estimates of the means. The standard deviation of the mean estimate for one source is equal to the standard deviation of one estimate for that source, divided by the square root of the number of cases. None of the mean estimates of τ differ from the value 40.97 by more than two standard deviations. All variations in the mean estimates can thus be attributed to noise. There is no evidence to suggest any effects due to errors in source or station positions, changes in the clock sync during the experiment, or errors in data processing.

The statistics of the estimates of f cannot be attributed entirely to noise, because of local oscillator instabilities. A hydrogen maser was used for the S-band reference at DSS 12, and rubidium was used at DSS 11, so the rubidium standard contribution dominated. Both the long- and short-term stabilities are on the order of one part in 10^{11}. Errors in the nominal value of f of up to 0.1 Hz were anticipated, as were short-term variations with a standard deviation on the order of 0.01 to 0.1 Hz.

Due to noise alone, the standard deviation of f should vary as $R^{-1/2}$, provided R is about 10 or greater. This relationship was nominally satisfied for the second- and third-strongest sources, with $R = 24.8$ and 8.20, and standard deviations of 0.0946 and 0.159 Hz. For the strongest source, the frequency instability was not negligible compared to the noise. Therefore, its effect was estimated from the results for the two strongest sources, assuming the noise and instability errors to add in the mean square. The rms error due to frequency instability was estimated at 0.036 Hz, which is well within the range of uncertainty of this effect. The rms frequency estimation error due to noise is then approximately $0.468/R^{1/2}$ Hz. This relationship was also closely satisfied for the next-weakest source. An R of 10 thus results in an rms error in frequency estimate of less than 0.1 divided by the measurement time of 0.64 s, just as it results in a timing error of less than 0.1 divided by the bandwidth.

B. Estimation of τ for Fixed f

When the *a priori* uncertainty in frequency is small, τ can be estimated by maximizing G over τ only, assuming no frequency error, i.e., $f = 0$. This results in better estimates of τ than does joint estimation of τ and f when the long- and short-term frequency instabilities are very small. Before this experiment was actually performed, it was felt that the frequency stabilities would be sufficiently good to omit maximization over f, and this was confirmed in the experiment. However, the amount of long-term drift is random, and the frequency offsets in the local oscillators might have been too large on another day. It was, therefore, necessary to process the data in both manners, in order to be able to predict future performance.

The results of the estimation of τ for f fixed at zero are presented in Table 3. The theoretical rms errors in the estimation of τ are also presented, as calculated for the estimated values of R. For the three highest SNR cases, the observed and calculated rms errors were very close. For the lowest two SNRs, the observed errors were significantly higher than the calculated values. This is because the theory breaks down when R is low enough so that extraneous results occur.

The observed rms errors at low SNRs would have been still higher if the assumed region of uncertainty in τ had been greater, because there would have been more extraneous results due to noise. Throughout the experiment, the uncertainty region was assumed to be from 30 to 50 μs.

C. Comparison of Estimation Methods

In comparing the results of estimating τ jointly with f and with f fixed at zero, it is seen that there is negligible difference in the standard deviations of the estimates for the three highest SNR cases, and that all are close to theory. For the two lower SNRs, the errors are significantly higher when f is estimated instead of assumed to be zero. There are two reasons for this. First, the estimates of f are poor enough to degrade the estimate of τ. Second, more extraneous estimates occurred, because there were effectively more independent calculations of G for noise only.

D. Estimation of R, ρ, and Flux Density

For each independent case, the approximate maximum-likelihood estimate for ρ is the square root of the maximum value of G, divided by the proper normalization factor. This is the best estimate of ρ only because the maximum value of G occurs at the best estimates of τ and f. A better estimate of ρ would be obtained from a value of G at the correct values of τ and f. Therefore, since it was desired to have the overall best estimates of ρ, and hence of R, the values of ρ were estimated using the best overall estimates of τ and f. These best estimates were taken as $\tau = 40.955$ and $f = 0.0807$ Hz, the values obtained from the strongest source. The overall estimates of ρ for each source were taken as the average of the estimates of ρ for all of the cases for that source.

The estimates of R were obtained from the estimates for ρ according to Eqs. (8) and (9). To estimate the correlated fluxes, it was assumed that the system temperatures at DSSs 11 and 12 were the cold sky temperatures of 37 and 16.3 K, respectively, raised by the source total flux at the rate of 0.11 K per flux unit. Then the correlated fluxes are given by Eq. (1).

Table 3 presents the estimated flux densities, input SNRs, and estimator SNRs for the five sources.

Acknowledgements

The author acknowledges the assistance of P. F. MacDoran, J. G. Williams, and D. S. Spitzmesser, who shared their knowledge and experience in VLBI experiments, and of S. S. Brokl and the station personnel at DSSs 11 and 12, who assisted in implementing the experiment.

References

1. Hildebrand, C. E., Ondrasik, V. J., and Ransford, G. A., "Earth-Based Navigation Capabilities for Outer Planet Missions," AIAA Paper No. 72-925, AIAA/AAS Astrodynamics Conference, Palo Alto, California, September 11-12, 1972.

2. Ondrasik, V. J., Hildebrand, C. E., and Ransford, G. A., "Preliminary Evaluation of Radio Data Orbit Determination Capabilities for the Saturn Portion of a Jupiter-Saturn-Pluto 1977 Mission," in *The Deep Space Network Progress Report*, Technical Report 32-1526, Vol. X, pp. 59–75, Aug. 15, 1972.

3. Hurd, W. J., "DSN Station Clock Synchronization by Maximum Likelihood VLBI," in *The Deep Space Network Progress Report*, Technical Report 32-1526, Vol. X, pp. 82–95, Aug. 15, 1972.

4. Kellermann, K. I., *et al.*, "High Resolution Observations of Compact Radio Sources at 13 Centimeters," *Astrophys. J.*, Vol. 151, Sept. 1970, pp. 803–809.

Table 1. Source intensity required for various system parameters

Antenna diameters, m		System temperatures, K		Amount of data or buffer size, 10^6 bits	Source intensity, fu of correlated flux
26	26	17	37	0.32	5.50
26	26	17	37	1.0	3.48
26	26	17	17	0.32	3.70
26	26	17	17	1.0	2.34
64	26	17	37	0.32	2.24
64	26	17	37	1.0	1.26
64	26	17	17	0.32	1.51
64	26	17	17	1.0	0.86
64	64	17	17	0.32	0.61
64	64	17	17	1.0	0.35

Table 2. Joint estimation of τ and f

Radio source	Estimated estimator SNR	Estimation of τ		Estimation of f	
		Mean, μs	Standard deviation, μs	Mean, Hz	Standard deviation, Hz
3C273	148.	40.955	0.0956	−0.0807	0.0503
3C279	24.8	41.00	0.228	−0.0799	0.0946
P1127-14	8.20	40.95	0.403	−0.0810	0.159
DW0742+10	4.24	40.95	0.719	−0.0772	0.239
4C39.25	3.81	40.78	1.26	−0.0355	0.249

Table 3. Estimation of τ for $f = 0$

Radio source	Estimated estimator SNR	Estimation of τ		Theoretical rms error in τ, μs
		Mean, μs	Standard deviation, μs	
3C273	148.	40.955	0.0978	0.095
3C279	24.8	41.00	0.224	0.232
P1127-14	8.20	40.95	0.408	0.403
DW0742+10	4.24	40.98	0.641	0.560
4C39.25	3.81	40.92	0.952	0.591

Table 4. Estimated flux densities and estimator SNRs

Radio source	Number of cases	Total flux (Ref. 4), fu	Estimated correlated flux, fu	Estimated geo. mean input SNR	Estimated estimator SNR
3C273	144	39.	22.	0.0834	148.
3C279	144	12.2	8.1	0.0344	24.8
P1127-14	72	6.2	4.6	0.0199	8.20
DW0742+10	72	3.7	3.3	0.0145	4.24
4C39.25	72	3.8	3.1	0.0138	3.81

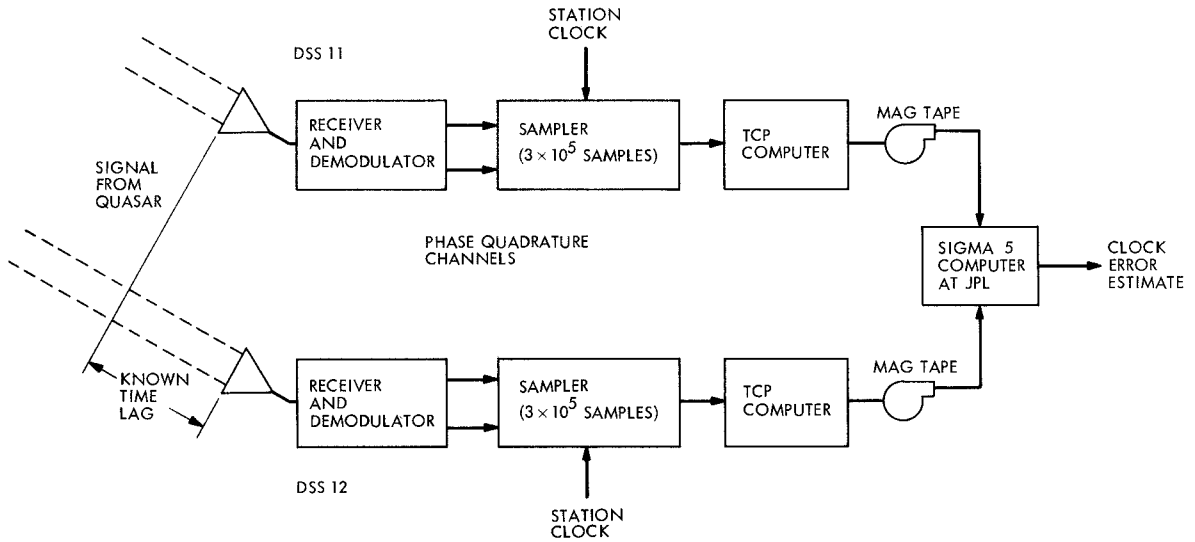

Fig. 1. VLBI time-sync demonstration block diagram

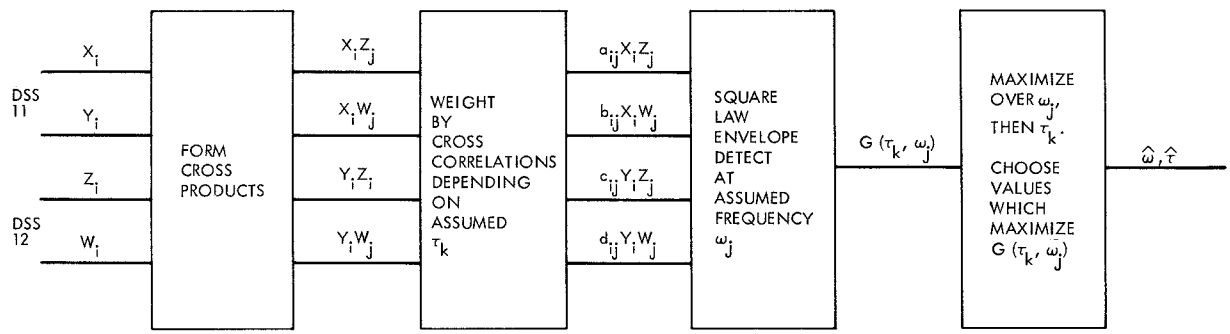

Fig. 2. Functional diagram of estimation procedure

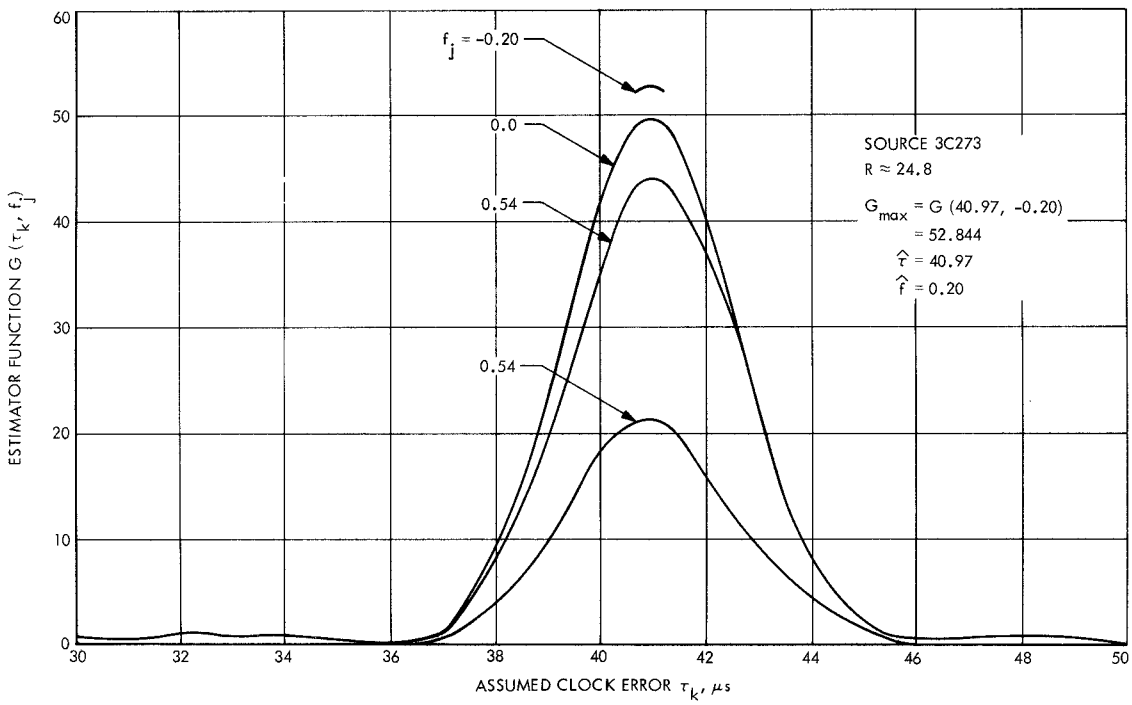

Fig. 3. An estimator sample function at a high SNR

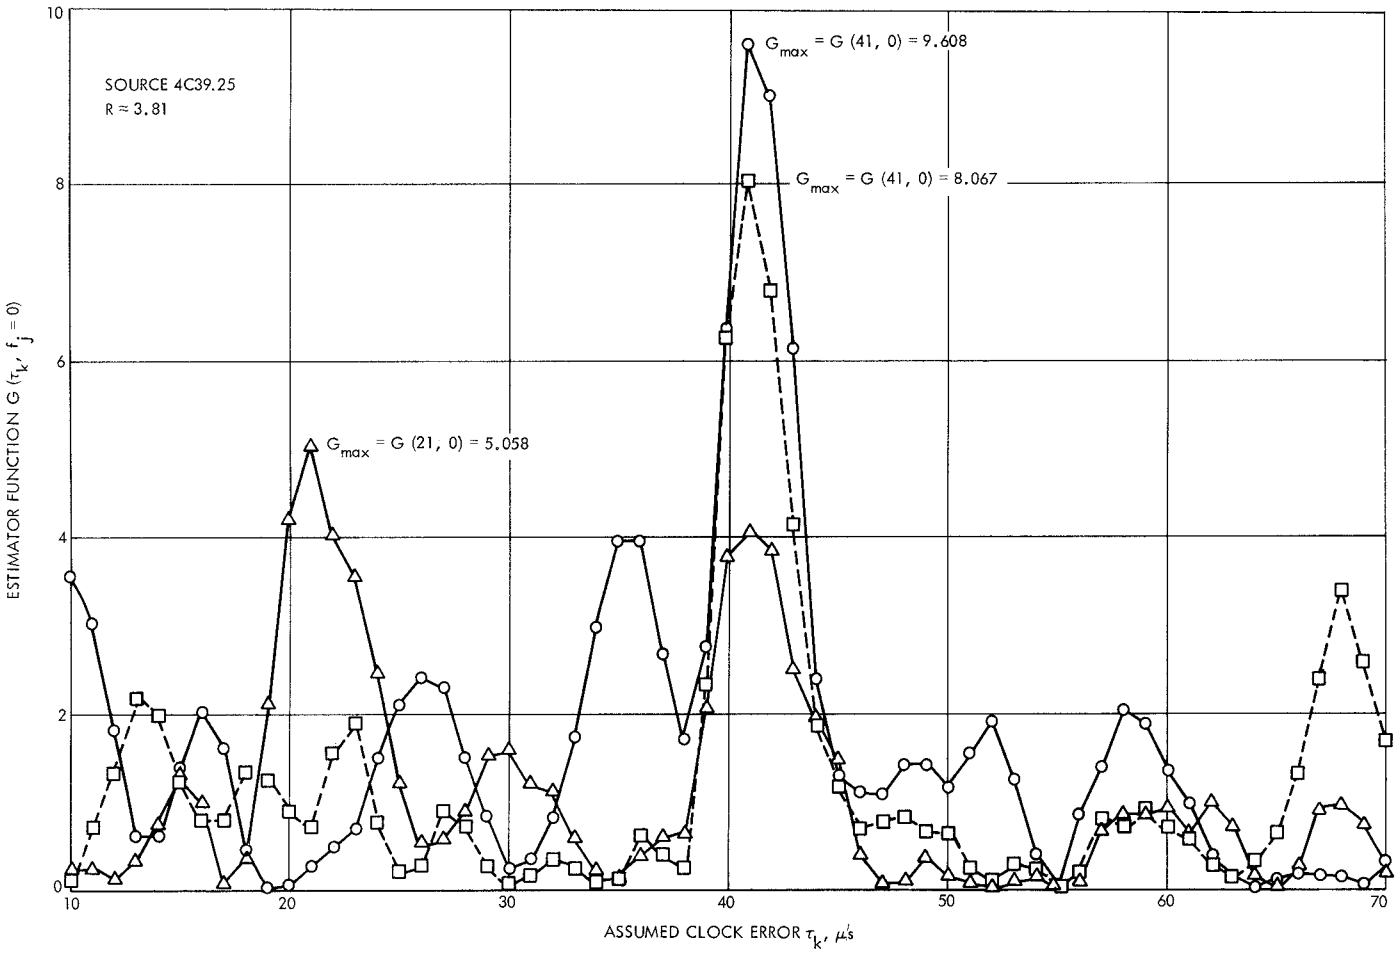

Fig. 4. Three estimator sample functions at a low SNR

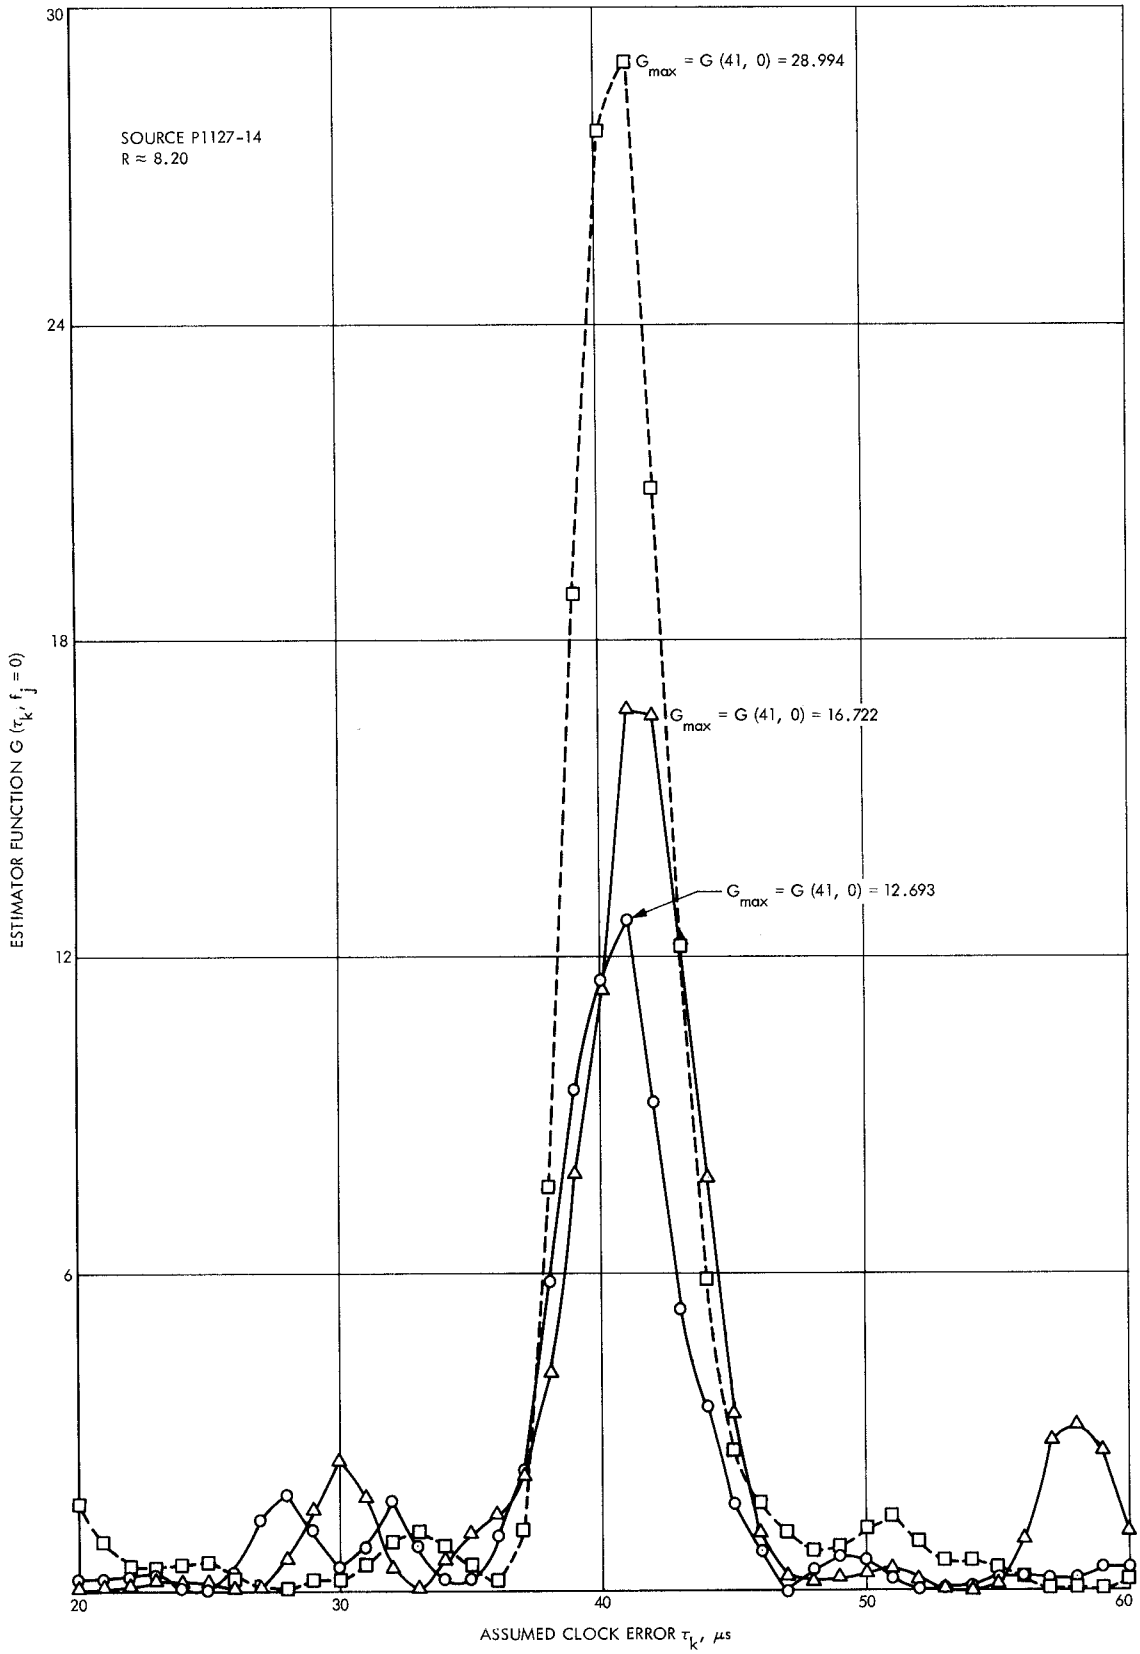

Fig. 5. Three estimator sample functions at a marginal SNR