Noise analysis of unevenly spaced time series data

C. Hackman and T. E. Parker

Abstract. Two-way satellite time and frequency transfer (TWSTFT) data are typically recorded on Monday, Wednesday and Friday. This produces an unevenly spaced time series on which it is difficult to perform an accurate two-sample variance analysis. We have investigated the effect of uneven data spacing on the computation of $\sigma_x(\tau)$ [1-2]. Evenly spaced simulated data sets were generated for noise processes ranging from white phase modulation to random walk frequency modulation. $\sigma_x(\tau)$ was then calculated for each noise type. Data were subsequently removed from each simulated data set to create two unevenly spaced sets with average intervals of 2,8 and 3,6 days. These sparse sets correspond to typical TWSTFT data patterns. $\sigma_x(\tau)$ was then calculated for each sparse data set using two different approaches. First, the missing data points were replaced by linear interpolation and $\sigma_x(\tau)$ calculated from this now full data set. The second approach ignored the fact that the data were unevenly spaced and calculated $\sigma_x(\tau)$ as if the data were equally spaced with average spacing of 2,8 or 3,6 days. The impact of uneven data spacing on the results of these two approaches is significant and is discussed. Finally, techniques are presented for correcting errors caused by uneven data spacing in simulated TWSTFT data sets, and the appropriate technique is applied to a real data set.

1. Introduction

Data points obtained from an experiment are often not evenly spaced. In this paper, we examine the application of $\sigma_x(\tau)$ [1-2] to the unevenly spaced time-series data (τ is the data spacing) obtained from two-way satellite time and frequency transfer (TWSTFT). We do so by using $\sigma_x(\tau)$ with both evenly and unevenly spaced data of known power-law noise type and magnitude. The noise types examined are white phase modulation (WHPM), flicker phase modulation (FLPM), white frequency modulation (WHFM), flicker frequency modulation (FLFM), and random walk frequency modulation (RWFM) [3].

$\sigma_x^2(\tau)$ is not the classical variance of x, but rather a two-sample variance which can be related to the modified Allan variance, mod $\sigma_y^2(\tau)$, by the equation $\sigma_x^2(\tau) = \tau^2 \operatorname{mod} \sigma_y^2(\tau)/3$ [1-2]. Two-sample variances are preferred to classical variances in analysing the frequency or time stability of oscillators because it is common for oscillators to display non-white frequency noise at long averaging times. For a white noise process, the classical variance converges on a fixed value as the number of data samples increases. The uncertainty of that value also decreases. However, if the noise is not white, the classical variance does not converge as the number of data samples increases. In fact, the value obtained for the classical variance will depend on the number of samples used. This problem can be avoided by using one of the family of two-sample variances (the Allan variance, $\sigma_y^2(\tau)$; the modified Allan variance, mod $\sigma_y^2(\tau)$; and the time variance, $\sigma_x^2(\tau)$). These estimators of frequency or time stability converge for white and non-white processes toward a fixed value for a given averaging time even as the sample size increases. Furthermore, by observing the slopes present in a log-log plot of $\sigma_y(\tau)$, mod $\sigma_y(\tau)$, or $\sigma_x(\tau)$ versus the averaging time τ, one can determine which power-law noise processes dominate at which averaging times [2-5]. The Allan variance $\sigma_y^2(\tau)$ can be used to determine the presence of RWFM, FLFM and WHFM, but cannot distinguish between WHPM and FLPM. For this reason, the modified Allan variance mod $\sigma_y^2(\tau)$ was devised: it can distinguish between all five of the power-law noise processes. The time variance $\sigma_x^2(\tau)$ is simply a rescaling of the modified Allan variance and is more useful in analysing systems in which white or flicker phase noise is of primary interest. This is the case for a time transfer system such as TWSTFT. In this paper we analyse the use of the time deviation $\sigma_x(\tau)$, which is the square root of the time variance.

Vernotte et al. [6] studied the analysis of noise and drift in unevenly spaced pulsar data. However, the data obtained from pulsar studies are much more sparse in time than are the data obtained from TWSTFT. For example, Vernotte et al. [6] generated a sequence of 8192 evenly spaced data points with known noise type and then removed points so that only 167 of the initial 8192 points remained. They then tried to use these 167 points to deduce the noise type of the underlying 8192-point sequence. In TWSTFT, the task is less daunting: time transfers are typically measured

C. Hackman and T. E. Parker: National Institute of Standards and Technology, Time and Frequency Division, 325 Broadway, Boulder, CO 80303, USA.

on Monday, Wednesday and Friday, so that, in a perfect world, we would have a data density of three data points present out of a possible seven.

This paper is not intended to be a rigorous treatment of how to calculate $\sigma_x(\tau)$ in all possible cases of unevenly spaced data. Rather, our purpose is to suggest methods and corrections which may be applied to data such as those produced by TWSTFT in order to obtain a more accurate assessment of the underlying time stability and noise type.

The National Institute of Standards and Technology (NIST) regularly performs time transfers with several laboratories in North America and Europe. Two of these laboratories are the United States Naval Observatory (USNO) in Washington, D.C., and the Van Swinden Laboratorium (VSL) in Delft, the Netherlands. The transfers with the USNO occur on Monday, Wednesday and Friday. The transfers with the VSL occurred on Monday and Wednesday for modified Julian dates (MJDs) 49387 to 49533 (4 February 1994 to 30 June 1994). Since then, they have occurred on Mondays, Wednesdays and Fridays. We use the NIST-USNO and the NIST-VSL data obtained from MJDs 49387 to 49770 (4 February 1994 to 22 February 1995) as our data-spacing templates for this experiment. We use two different data-spacing templates because, among all of the time transfers that the NIST performed over the period MJD 49387 to 49770, the transfers that occurred with the greatest frequency were those with the USNO. Among those that occurred with lesser frequency were those with the VSL.

2. Method of evaluation

We evaluated the use of $\sigma_x(\tau)$ with unevenly spaced data having the five different power-law noise types: WHPM, FLPM, WHFM, FLFM and RWFM. Ten data files were generated for each noise type. The WHPM, WHFM and RWFM files were generated by using integration in combination with the random-number generator of the VAX FORTRAN version 5.0 math library. The FLPM and FLFM files were generated according to the algorithm of Kasdin and Walter [7] and used the random number generator RAN1 of Press et al. [8]. The only difference between the ten files of a given noise type was the seed number used to generate the data file. All ten data files of each noise type had 384 evenly spaced data points spaced one day apart, and the data points were assigned time tags of MJDs 49387 through 49770. In the next step, we removed data points from each file so that the remaining data points were aligned with the data points obtained from the NIST-USNO TWSTFT for MJDs 49387 to 49770. This produced a file containing 137 unevenly spaced data points. The missing data points were then filled in by performing linear interpolation between the remaining data points. After this last step, there are once again 384 evenly spaced data points.

This process was also performed using the data spacing obtained from the NIST-VSL time transfers for MJDs 49387 to 49770. Therefore, for each seed number of each noise type, we finally had five data files:

File Type 1: The originally generated 384 evenly spaced data points with known noise type and magnitude.

File Type 2: A data file of 137 data points spaced as in the NIST-USNO time transfers of MJDs 49387 to 49770. This file is obtained by removing the appropriate data points from File #1. The average spacing (see below) is 2,816 days.

File Type 3: File #2 with the missing data points filled in by linear interpolation.

File Type 4: A data file of 108 data points spaced as in the NIST-VSL time transfers of MJDs 49387 to 49770. This file, like File #2, is obtained by removing points from File #1. The average spacing (see below) is 3,579 days.

File Type 5: File #4 with the missing data points filled in by linear interpolation.

Having created all fifty files for a given noise type, we then performed a $\sigma_x(\tau)$ analysis of each file. For the data files with even spacing (File Types 1, 3 and 5 above) we computed $\sigma_x(m\tau_{0,\mathrm{evn}}$; $m = 1$, 2, 4, 8, 16, 32, 64, 128; $\tau_{0,\mathrm{evn}} = 1$ day) in the usual fashion [1-2]. For the files with unevenly spaced data (File Types 2 and 4) we computed $\sigma_x(\tau)$ by treating the adjacent data points as if they were evenly spaced with $\tau_{0,\mathrm{avg}}$ calculated as follows:

$$\tau_{0,\mathrm{avg}} = (\mathrm{MJD}_{\mathrm{last}} - \mathrm{MJD}_{\mathrm{first}})/(N-1), \qquad (1)$$

where $\mathrm{MJD}_{\mathrm{first}}$ and $\mathrm{MJD}_{\mathrm{last}}$ are the time tags for the first and last data points, and N is the number of data points. For File Type 2, $\tau_{0,\mathrm{avg}} = 2{,}816$ days, and for File Type 4, $\tau_{0,\mathrm{avg}} = 3{,}579$ days. In both of these latter cases, we computed $\sigma_x(n\tau_{0,\mathrm{avg}})$ for n equal to 1, 2, 4, 8, 16 and 32.

Having obtained $\sigma_x(\tau)$ versus τ for all fifty files, we then computed the average values of $\sigma_x(\tau)$ for each file type. For example, we had, for File Type 3, ten sets of $\sigma_x(\tau)$ versus τ, where $\tau = 1$, 2, 4, 8, 16, 32, 64 and 128 days. The average value of $\sigma_x(\tau = 1$ day) was calculated from the ten values of $\sigma_x(\tau = 1$ day). We repeated this process for $\sigma_x(\tau = 2$ days), $\sigma_x(\tau = 4$ days), ... up to $\sigma_x(\tau = 128$ days). Therefore, for each power-law noise type, we finally have five plots of $\sigma_x(\tau)$ versus τ:

(1) Average $\sigma_x(\tau = 1$, 2, 4, 8, 16, 32, 64 and 128 days) for File Type 1, i.e. the files with known noise type. This plot shows the "correct" values for $\sigma_x(\tau)$.

(2) Average $\sigma_x(\tau = 2{,}816,\ 5{,}632,\ 11{,}264,\ 22{,}528,\ 45{,}056$ and $90{,}112$ days$)$ for File Type 2. This represents the results we obtain by using unevenly spaced data with the NIST-USNO distribution.

(3) Average $\sigma_x(\tau = 1,\ 2,\ 4,\ 8,\ 16,\ 32,\ 64$ and 128 days$)$ for File Type 3. This represents the results we obtain by taking unevenly spaced data with the NIST-USNO distribution, performing linear interpolation to make an evenly spaced data file, and then performing the $\sigma_x(\tau)$ analysis.

(4) Average $\sigma_x(\tau = 3{,}579,\ 7{,}158,\ 14{,}316,\ 28{,}632,\ 57{,}264$ and $114{,}528$ days$)$ for File Type 4. This represents the results we obtain by using unevenly spaced data with the NIST-VSL distribution.

(5) Average $\sigma_x(\tau = 1,\ 2,\ 4,\ 8,\ 16,\ 32,\ 64$ and 128 days$)$ for File Type 5. This represents the results we obtain by taking unevenly spaced data with the NIST-VSL distribution, performing linear interpolation to make an evenly spaced data file, and then performing the $\sigma_x(\tau)$ analysis.

Finally, for each average value of $\sigma_x(\tau)$ for File Types 2 to 5, we computed a "correction factor". The correction factor is defined as:

$$\text{correction factor}\ (\sigma_x(\tau)_{\text{File Type}\,j}) = \operatorname{avg}\sigma_x(\tau)_{\text{File Type}\,1}/\operatorname{avg}\sigma_x(\tau)_{\text{File Type}\,j}. \tag{2}$$

In other words, multiplying the $\sigma_x(\tau)$ values obtained using File Type j by the correction factors for File Type j produces the correct values for $\sigma_x(\tau)$ as given by File Type 1. Because the τ values for File Types 2 and 4 do not match the τ values for File Type 1, various types of interpolation were used to obtain the correction factors for these two file types. The details of obtaining the correction factors for the different noise types and file types are discussed in the next section.

3. Results

Figures 1 to 5 show the results obtained for the noise types WHPM, FLPM, WHFM, FLFM and RWFM, respectively. As mentioned above, each of the points shown corresponds to the mean of ten values. The standard deviation of each set of ten values was also computed, but, for visual clarity, error bars indicating ± 1 standard deviation are shown only on the File Type 1 (i.e. correct) values. Approximately the same size error bars should be applied to each of the file type curves.

Figure 1 shows the results obtained for white PM noise. There are several important points here. First of all, File Types 3 and 5 (interpolating unevenly spaced data to form evenly spaced data) yield values of $\sigma_x(\tau)$ which are much too small when τ is less than the $\tau_{0,\text{avg}}$ of the corresponding unevenly spaced data set. On the

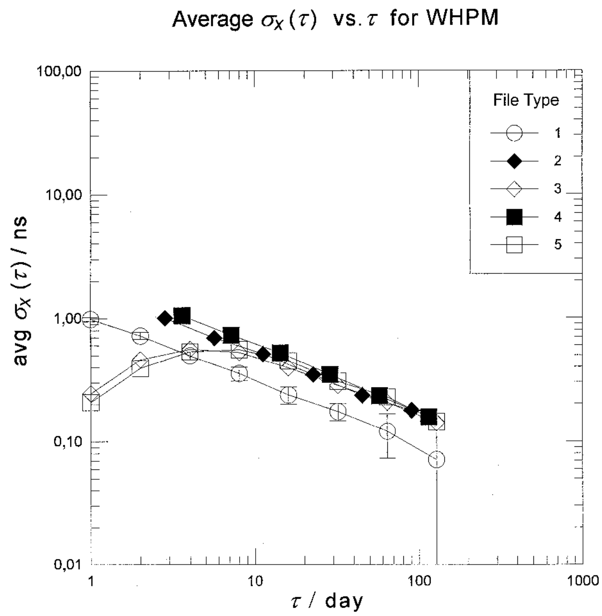

Figure 1. The average values of $\sigma_x(\tau)$ obtained from simulated WHPM data. File Type 1 indicates the correct values obtained from the original evenly spaced simulated data. The error bars on the File Type 1 data are the standard deviations associated with each of the average File Type 1 values. File Type 2 and File Type 3 show the results obtained when some of the original data points are deleted, thus forming an average data spacing of 2,816 days, and then the remaining points analysed two different ways. File Type 4 and File Type 5 indicate results obtained when data are decimated to produce an average data spacing of 3,579 days. For visual clarity, the error bars are not shown for File Types 2 to 5. However, the sizes of the missing error bars are approximately the same as those shown for File Type 1.

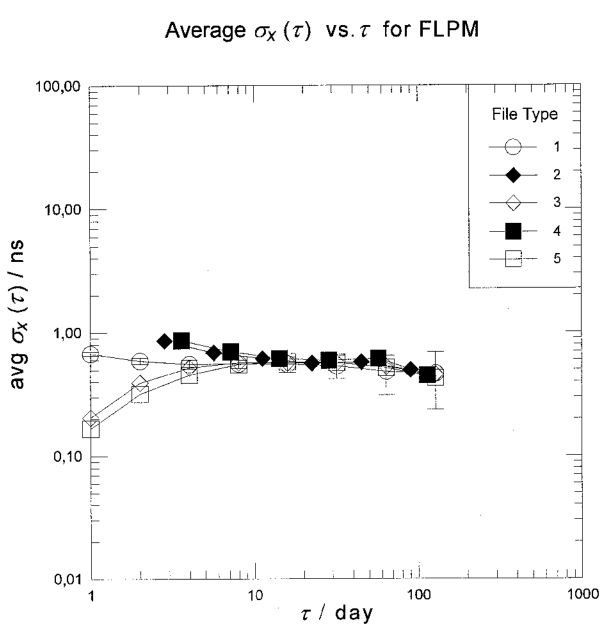

Figure 2. The average values of $\sigma_X(\tau)$ obtained from simulated FLPM data. See Figure 1 caption for more details.

other hand, File Types 2 and 4 (the unevenly spaced data) yield $\sigma_x(\tau)$ values which have the –1/2 slope appropriate to white PM [1], but which are consistently too high. In fact, for $\tau \geq 8$ days, all of the methods used converge to yield approximately the same too-large values for $\sigma_x(\tau)$.

For File Types 2 and 4, the correction factor is in theory constant for all values of τ and can be expressed as:

$$\text{correction factor (WHPM)} = (\tau_{0,\,\text{evn}}/\tau_{0,\,\text{avg}})^{1/2}. \quad (3)$$

It is easily understood why the calculation of $\sigma_x(\tau)$ from unevenly spaced data with an underlying WHPM process yields the correct slope but values of $\sigma_x(\tau)$ that are uniformly too high. With WHPM noise each data point in the time series is independent of all others. Therefore, the point-to-point fluctations are independent of the time between the points and

$$\begin{aligned}\sigma_x(m\,\tau_{0,\,\text{avg, File Type 2}})_{\text{File Type 2}} &= \sigma_x(m\,\tau_{0,\,\text{avg, File Type 4}})_{\text{File Type 4}}\\ &= \sigma_x(m\,\tau_{0,\,\text{evn}})_{\text{File Type 1}},\end{aligned}$$

where m is the same value throughout this equation. Because this is true, and because the log-log plot for the File Type 1 data has a –1/2 slope, the results obtained from the File Type 2 and File Type 4 data will also have the correct –1/2 slope.

Figure 2 shows the flicker PM results. Once again, File Types 3 and 5 yield values of $\sigma_x(\tau)$ which are too small at short averaging times. Also, the lower-τ values of $\sigma_x(\tau)$ for File Types 2 and 4 are again too high. However, the results obtained from all file types converge toward the correct value as τ increases. Similar results are obtained for white FM (Figure 3) and flicker FM (Figure 4).

Figure 5 shows the RWFM results. Here, the use of interpolated data (File Types 3 and 5) provides virtually the same results as the originally generated data file (File Type 1) and the use of unevenly spaced data (File Types 2 and 4) provides values of $\sigma_x(\tau)$ which are too

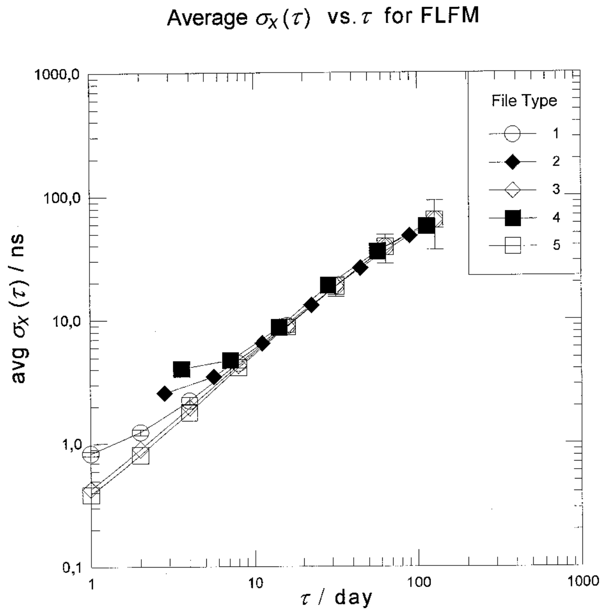

Figure 4. The average values of $\sigma_X(\tau)$ obtained from simulated FLFM data. See Figure 1 caption for more details.

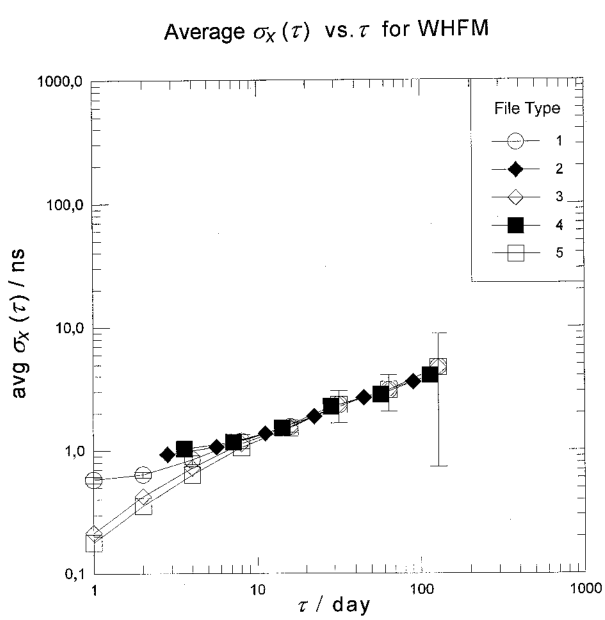

Figure 3. The average values of $\sigma_x(\tau)$ obtained from simulated WHFM data. See Figure 1 caption for more details.

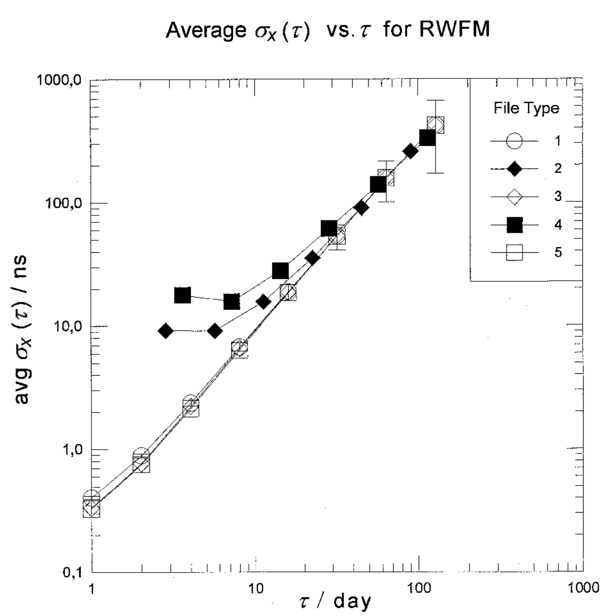

Figure 5. The average values of $\sigma_X(\tau)$ obtained from simulated RWFM data. See Figure 1 caption for more details.

Table 1. Correction factors for File Types 3 and 5.

τ/days	WHPM		FLPM		WHFM		FLFM		RWFM	
	USNO	VSL	USNO	VSL	USNO	VSL	USNO	VSL	USNO	VSL
1	4,02	4,74	3,28	4,02	2,71	3,24	1,94	2,15	1,22	1,23
2	1,56	1,83	1,51	1,85	1,51	1,80	1,38	1,53	1,16	1,18
4	0,88	0,93	1,07	1,22	1,16	1,32	1,13	1,24	1,07	1,11
8	0,67	0,64	0,98	1,01	1,05	1,13	1,04	1,08	1,02	1,06
16	0,60	0,53	0,98	0,96	1,01	1,02	1,01	1,02	1,01	1,02
32	0,61	0,57	0,95	0,94	1,00	1,00	1,00	1,01	1,00	1,00
64	0,58	0,53	0,91	0,93	1,00	1,00	1,00	1,00	1,00	1,00
128	0,50	0,49	1,02	1,08	1,00	1,00	1,00	1,00	1,00	1,00

large at small τ. In fact, as we progress from the WHPM process to the low-frequency-dominated noise processes (e.g. RWFM) [3], the use of linear interpolation to fill in missing data points becomes an increasingly better approximation of the truth. For smaller values of τ, using the unevenly spaced data becomes an increasingly worse approximation of the truth. As we progress from FLPM to RWFM, the results obtained using all methods converge on the correct value as τ increases.

From the results shown in Figures 1 to 5 we have computed correction factors. Table 1 shows the correction factors obtained from the file types (3 and 5) which have evenly spaced data. These correction factors were obtained by simply taking the ratio $\sigma_x(m\tau_{0,\mathrm{evn}})_{\text{File Type 1}}/\sigma_x(m\tau_{0,\mathrm{evn}})_{\text{File Type 3 or 5}}$ where $\tau_{0,\mathrm{evn}} = 1$ day. Tables 2 to 3 show the correction factors obtained for the file types (2 and 4) with unevenly spaced data. Because the averaging times for the unevenly spaced files (e.g. 2,816, 5,632, …, etc. days for File Type 2) do not match the averaging times for File Type 1 (1, 2, 4, …, etc. days), we cannot simply take a ratio of two values to obtain the correction factor. Generally, interpolation of some sort is required. For white PM, we performed linear fits to the File Type 1, 2 and 4 data and then used the linear fits to compute the correction values. Note that the correction factors for WHPM shown in Tables 2 and 3 all fall within 10% of the values calculated from (3). For flicker PM, we interpolated between points as needed: for example, to get the correction factor for File Type 2, $\tau = 2{,}816$ days, we used

$$\text{correction factor } (2,816 \text{ days}) = \\ = \frac{\sigma_x(\tau = 2,816\text{d})_{\text{File Type 1, interp}}}{\sigma_x(\tau = 2,816\text{d})_{\text{File Type 2}},} \tag{4}$$

where the quantity in the numerator is obtained by performing linear interpolation between $\sigma_x(\tau = 2$ days$)$ and $\sigma_x(\tau = 4$ days$)$. For white FM, flicker FM and random walk FM we performed a cubic fit to each of the curves and then used the fits to compute the correction factors.

Table 2. Correction factors for File Type 2.

τ/days	WHPM	FLPM	WHFM	FLFM	RWFM
2,816	0,60	0,67	0,80	0,64	0,16
4	0,59	0,71	0,87	0,76	0,27
5,632	0,59	0,81	0,92	0,86	0,41
8	0,58	0,85	0,95	0,94	0,57
11,264	0,58	0,90	0,97	0,99	0,71
16	0,57	0,94	0,99	1,01	0,84
22,528	0,57	0,97	1,00	1,01	0,92
32	0,56	0,94	1,00	1,01	0,97
45,056	0,56	0,89	1,02	1,01	1,00
64	0,55	0,88	1,04	1,02	1,00
90,112	0,55	0,95	1,08	1,05	1,00

Table 3. Correction factors for File Type 4.

τ/days	WHPM	FLPM	WHFM	FLFM	RWFM
3,579	0,49	0,65	0,80	0,50	0,12
4	0,49	0,66	0,82	0,56	0,14
7,158	0,50	0,80	0,91	0,81	0,33
8	0,50	0,81	0,92	0,85	0,38
14,316	0,50	0,91	0,97	0,94	0,59
16	0,50	0,91	0,98	0,95	0,63
28,632	0,51	0,91	1,00	0,93	0,78
32	0,51	0,90	1,01	0,92	0,80
57,264	0,52	0,80	1,03	0,92	0,90
64	0,52	0,81	1,04	0,92	0,93
114,528	0,52	1,04	1,09	1,06	1,10

4. Discussion

There is, unfortunately, no way to apply these results blindly. The user will need to have an idea of what sort of noise types make sense in the context of the measurement. Initially, one log $\sigma_x(\tau)$ versus log (τ) plot using the original set of unevenly spaced data and one log $\sigma_x(\tau)$ versus log (τ) plot using a full data set formed by linear interpolation should be constructed.

At medium-to-large averaging times (in our experiments, $\tau \geq 8$ days), almost all methods, in their uncorrected state, provide the correct slope for the log $\sigma_x(\tau)$ versus log (τ) plot. For WHPM, the unevenly spaced data give the correct slope at all values of τ. Thus, the user can determine which power-law noise process dominates at medium-to-large averaging times. (The exception to this rule occurs when RWFM predominates, and the unevenly spaced data are used to make the log $\sigma_x(\tau)$ versus log (τ) plot. In this case, the slope of the plot is slow in converging to the correct +3/2 value.) The more difficult part arises when the value of m in $\tau = m\tau_0$ is small. It is here that we see the largest effects of not having an evenly spaced data set. In addition, in this regime the noise process which dominates a measurement often changes from one type to another.

If data are recorded on Monday, Wednesday and Friday, it will be impossible to get a reliable estimate of $\sigma_x(\tau = 1\ \text{day})$ – that information simply is not available. One can, however, make a fair estimate of $\sigma_x(\tau = 2\ \text{days})$ in this case because Monday-Wednesday and Wednesday-Friday are each two-day intervals. To be completely safe, one could avoid stating values of $\sigma_x(\tau)$ for $\tau < \tau_{0,\text{avg}}$. Finally, in this analysis, the ratio of the data length (384 days) to $\tau_{0,\text{avg}}$ (2,816 and 3,579 days) was always greater than 100; therefore, it may not be appropriate to use these results with short, sparse data sets.

If there is only one, known, noise type present, then the correction factors shown in Tables 1 to 3 can be applied. Unless a user has exactly the same average data spacing as we did, some interpolation may be needed in order to use the correction factors. Fortunately, the values of most of the correction factors are not strongly dependent on the average spacing for the range of spacing that was examined. If the noise type is not known, one could begin by deciding into which of the two following categories the experiment fits: (i) the results contain only measurement noise; or (ii) the results contain measurement noise and clock behaviour.

Examples of the former category are common-clock or closure TWSTFT experiments. An example of the latter category is performing TWSTFT between two remotely located clocks. We examine each of these situations below.

4.1 Measurement noise

If the results should only contain measurement noise, then the noise type will most likely be white PM or flicker PM. Fortunately, as Figure 1 shows, if WHPM is the dominant noise type, the log $\sigma_x(\tau)$ versus log (τ) plot for the unevenly spaced data will have a clear –1/2 slope and it will be obvious that the WHPM corrections should be applied. This method was used in [9]. Similarly, if the log $\sigma_x(\tau)$ versus log (τ) plot has zero slope at large τ (Figure 2), then apply the FLPM corrections. In this case it is important to be certain that the noise type at large τ has been correctly ascertained because, if the noise type is FLPM, the corrections which are applied at large τ are fairly small. If the noise type is WHPM, the corrections which are applied at large τ are relatively large.

4.2 Combination of clock noise and measurement noise

If the experiment measures clock behaviour (or some other quantity which is characterized by a low-frequency-dominated noise type), then the situation becomes more complicated because the results will contain a mixture of noise types – the noise type associated with the measurement and the noise type(s) associated with the behaviour of the clocks under study. We have evaluated various analysis techniques and have arrived at the following recommendations which combine ease of use with acceptable accuracy.

First, examine the $\sigma_x(\tau)$ plots for evidence of measurement noise (WHPM, FLPM). The simplest way to see if there is any measurement noise is to look at the $\sigma_x(\tau)$ plot of the interpolated data set in the region where τ is small to medium. As Figures 1 to 3 show, for WHPM, FLPM and WHFM, the $\sigma_x(\tau)$ plot of the interpolated data will curve down as τ decreases to approach $\tau = 1$ day. In the case of FLFM, the $\sigma_x(\tau)$ plot of the interpolated data makes a straight line as τ decreases. In the case of RWFM, the $\sigma_x(\tau)$ plot curves up slightly as τ decreases. Therefore, if a downward curve is present at small τ and if there is evidence of a flat transition area at medium τ, there is probably significant measurement noise present.

If there indeed is measurement noise mixed in with the long-term noise, we suggest the following procedure (hereafter called the "hybrid method"): compute $\tau_{0,\text{avg}}$ from the unevenly spaced data and then simply use the $\sigma_x(\tau)$ values obtained from the interpolated data for $\tau > \tau_{0,\text{avg}}$. Then, estimate $\sigma_x(m\tau_{0,\text{evn}})$, where $m\tau_{0,\text{evn}}$ is the largest integral multiple of $\tau_{0,\text{evn}}$ that is less than $\tau_{0,\text{avg}}$, as follows:

(a) Using the values of log $\sigma_x(\tau = \tau_{0,\text{avg}})$ and log $\sigma_x(\tau = 2\tau_{0,\text{avg}})$ obtained from the unevenly spaced data, perform a linear extrapolation to smaller τ to obtain an estimate for log $\sigma_x(\tau = m\tau_{0,\text{evn}})$ for the unevenly spaced data set.

(b) Compute the average of log $\sigma_x(\tau = m\tau_{0,\text{evn}})$ obtained from (a) and log $\sigma_x(\tau = m\tau_{0,\text{evn}})$ obtained from the interpolated data set.

(c) Use this average value as an estimate of the correct value of log $\sigma_x(\tau = m\tau_{0,\text{evn}})$. From this obtain $\sigma_x(\tau = m\tau_{0,\text{evn}})$.

For example, the NIST-USNO data have $\tau_{0,\text{avg}} = 2{,}816$ days. Therefore, to obtain values of $\sigma_x(4\ \text{days} \le 128\ \text{days})$ we would use the $\sigma_x(\tau)$ values obtained from the interpolated data. To estimate $\sigma_x(\tau = 2\ \text{days})$ we would use the three steps outlined above. Further examples of this process are presented below.

This technique works because, for typical clock noise types (WHFM, FLFM, RWFM), the uncorrected values obtained from the interpolated data set are a pretty good estimate of the true values for medium to long averaging times. For measurement noise types WHPM, FLPM and WHFM, at small values of τ, taking the average of the *logarithm* of $\sigma_x(\tau)$s associated with the interpolated and the unevenly spaced data sets yields an acceptable estimate of the true value of $\sigma_x(\tau)$. Note that this averaging technique works very poorly at small τ for RWFM (Figure 5) and not so well at small τ for FLFM (Figure 4). Therefore, this technique should not be used if it looks as if clock noise dominates all the way down to small values of τ.

If inspection of the $\sigma_x(\tau)$ plots reveals no hint of measurement noise (i.e. it appears that clock noise dominates even at small τ), then determine the noise type from the large-τ values of $\sigma_x(\tau)$ and apply the appropriate correction factors from Table 1 to the $\sigma_x(\tau)$ values obtained from the interpolated data set.

We now give three examples of the analysis of mixed simulated noise types, ranging from situations in which the measurement noise dominates out to medium τ to situations in which the measurement noise is quickly overwhelmed by clock behaviour. In combination 1 (Figures 6a and 6b), we see a case in which inspection of the initial $\sigma_x(\tau)$ plots (Figure 6a) reveals obvious signs of the presence of both measurement and clock noise. The average data spacing is 2,816 days. As Figure 6b shows, using the hybrid method provides very good estimates of the correct values of $\sigma_x(\tau)$: the largest error is only 10% of the true $\sigma_x(\tau)$. In addition, we do not need to know precisely what types of noise are present (in this case, WHPM and WHFM) in order to arrive at the final estimates for $\sigma_x(\tau)$. Finally, note that we do not attempt to obtain a value for $\tau = 1$ day. We obtain the estimate of $\sigma_x(\tau = 2$ days) as follows: for the unevenly spaced data, log ($\sigma_x(\tau$ = 2,816 days), in nanoseconds) = 0,113, and log ($\sigma_x(\tau$ = 5,632 days)) = 0,031. By performing a linear extrapolation from these two points, we can estimate that log ($\sigma_x(\tau$ = 2 days)) = 0,154. From the interpolated data, we obtain log ($\sigma_x(\tau$ = 2 days)) = –0,222. By taking the average of these two values of log ($\sigma_x(\tau$ = 2 days)), we obtain our final estimate that log ($\sigma_x(\tau$ = 2 days)) = –0,034.

In combination 2, we again see signs of both measurement noise and clock noise in the initial $\sigma_x(\tau)$ plots (Figure 7a). The average data spacing for combinations 2 and 3 (see below) is 3,008 days. As Figure 7b shows, the hybrid method again provides a good estimate of the correct values for this combination of WHPM and FLFM.

In combination 3, it is difficult to tell if there is any measurement noise present. The $\sigma_x(\tau)$ plot of the interpolated data set does exhibit a very faint downward curve as τ decreases towards 1 day, but other than that, it looks like FLFM (Figure 8a). We have used both the hybrid technique and the simple application of the FLFM corrections (Table 1). As Figure 8b shows, the FLFM corrections work marginally better. As it turns out, the true $\sigma_x(\tau)$ curve shows clear evidence of measurement noise (WHPM) only at $\tau = 1$ day – a time interval about which we can gain no information from the sparse ($\tau_{0,\mathrm{avg}} = 3{,}008$ days) data set.

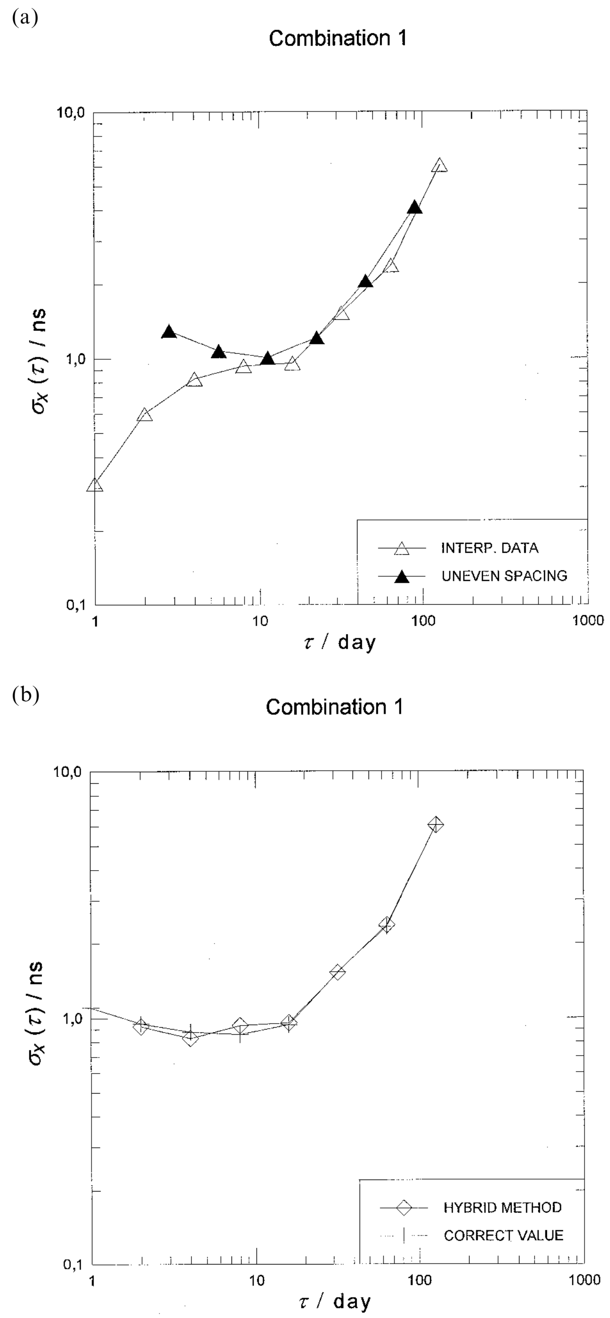

Figure 6. Figure 6a shows the uncorrected $\sigma_x(\tau)$ values obtained from a sparse data set with a mixture (combination 1) of noise types. Figure 6b shows the corrected values of $\sigma_x(\tau)$ obtained using the hybrid method (see text) and the values obtained from the original, evenly spaced data set.

4.3 Experimental data

To demonstrate the usefulness of the hybrid method on real rather than simulated TWSTFT data, we have analysed data from a series of TWSTFT measurements made between masers at the NIST and the USNO. Hydrogen masers were used in order to minimize the clock noise. Details of these measurements are given

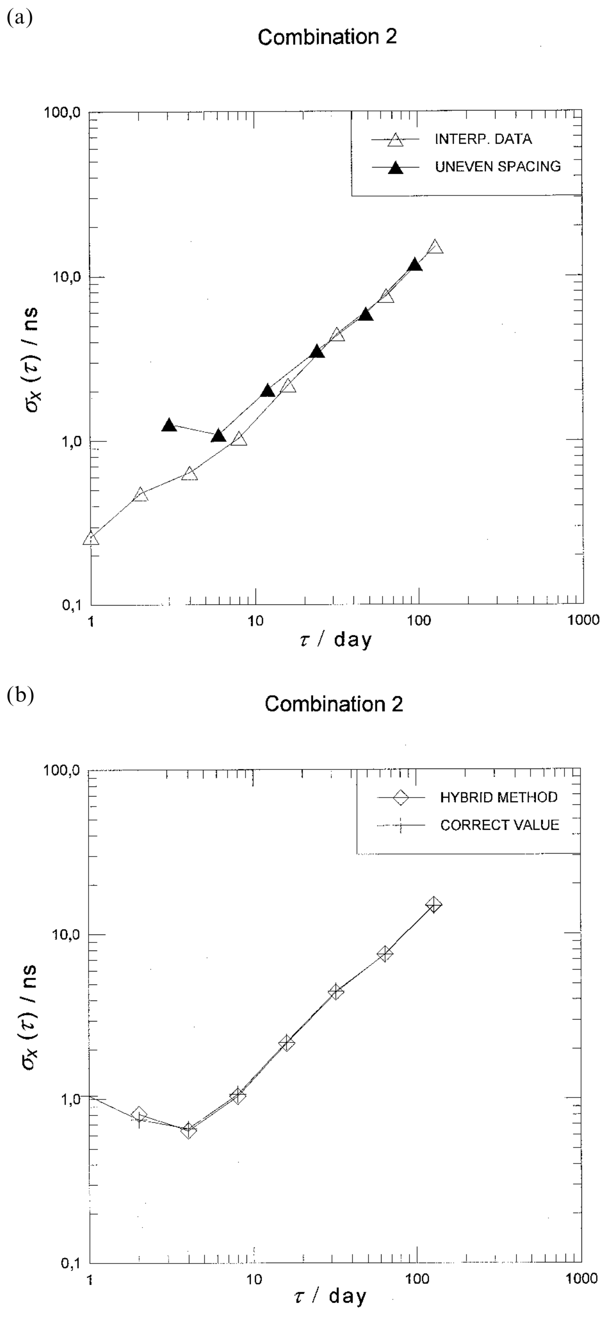

Figure 7. These figures (7a and 7b) are similar to those of Figure 6. However, a different mixture of noise types (combination 2) is used.

in [9]. Over a 129 day period, 5 minute time transfers were made on Mondays, Wednesdays and Fridays. However, during one three-week subset of these measurements, time transfers were also performed on Tuesdays, Thursdays, Saturdays and Sundays, resulting in the formation of a 21 day set of once-per-day time transfers.

Figure 9 shows the two sets of $\sigma_x(\tau)$ versus τ computed from these data. The squares in Figure 9 represent the values of $\sigma_x(\tau)$ calculated from the

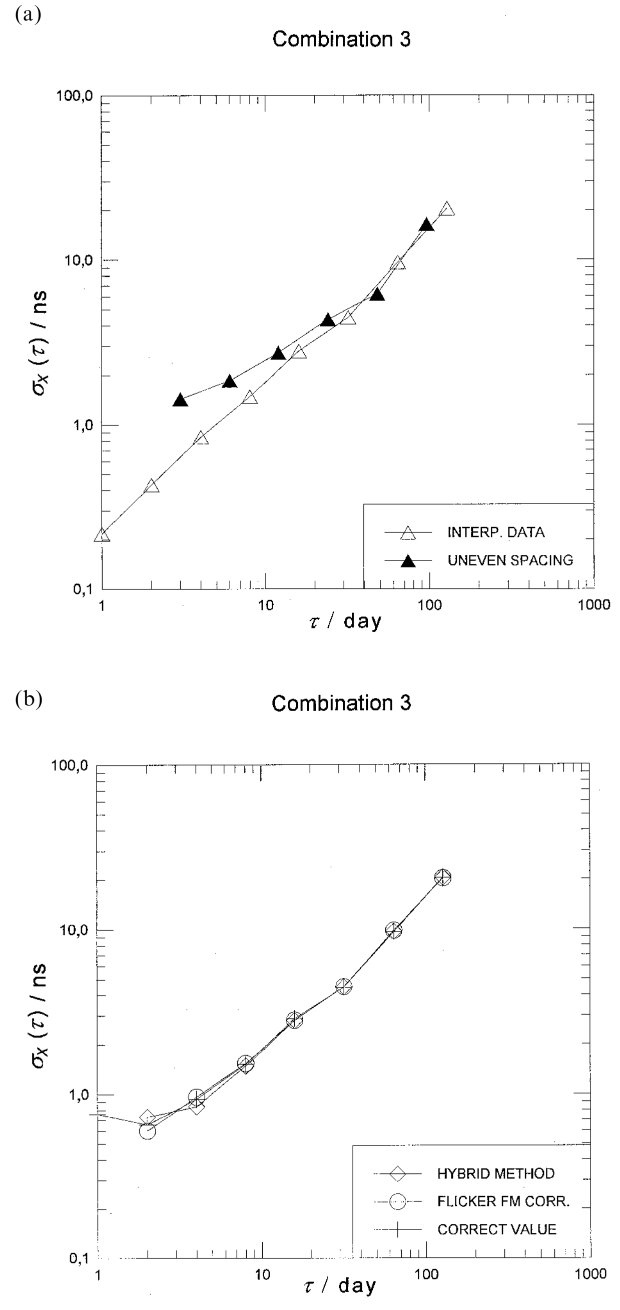

Figure 8. Figure 8a is similar to Figures 6a and 7a; however, a different mixture of noise types (combination 3) is used. Figure 8b shows the results of applying two types of corrections: those obtained using the hybrid method and those obtained using FLFM corrections. Again, the values as obtained from the original, evenly spaced data set are also shown in Figure 8b.

21 day data set in which data were taken every day. Here the values of τ range from 1 day to 4 days. Since these data were evenly spaced, no special techniques were required to calculate $\sigma_x(\tau)$. The diamonds represent the values of $\sigma_x(\tau)$ calculated from a data set consisting only of the Monday, Wednesday and

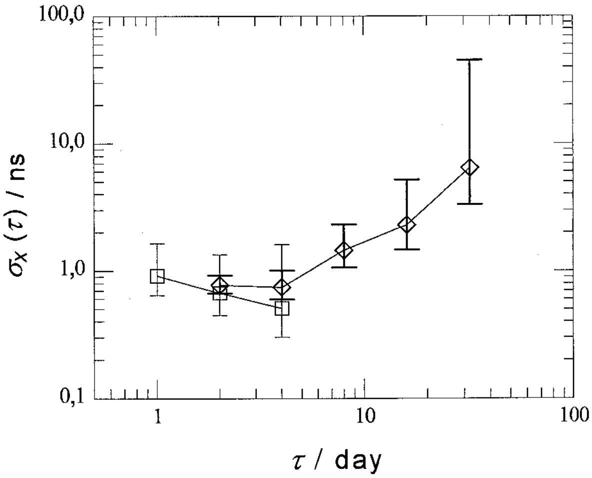

Figure 9. The values of $\sigma_x(\tau)$ for a set of real TWSTFT data obtained between masers at the NIST and the USNO. The squares represent the stability of a set of 21 consecutive daily time transfers. The diamonds represent the stability of a set of time transfers taken approximately three times per week over the course of 129 days. The hybrid method (Section 4.2) was used to compute the latter values of $\sigma_x(\tau)$.

Friday time transfers taken over the 129 days. This set, which represents a typical TWSTFT data pattern, contains 48 measurements and thus has an average data interval of 2,7 days. The hybrid method, as described in Section 4.2, was used to calculate these values of $\sigma_x(\tau)$ for τ values ranging from 2 days to 32 days. The error bars for the diamonds and the squares represent 95% confidence limits [10].

Unlike the simulated data in Section 4.2, we do not know the true noise levels in the experimental data of Figure 9. For τ greater than about 5 days the observed noise is close to FLFM and is consistent with the maser noise. In this region the interpolated TWSTFT data should be a good measure of the true noise level, as observed in Section 4.2. However, it is at the smaller values of τ ($\tau \leq 5$ days) that the effect of the uneven data spacing is most prominent. Here, the best we can do is to use the 21 days of consecutive data to estimate the true noise levels at small values of τ. It is clear from these data that the dominant noise process at small τ values is either WHPM or FLPM and that this is definitely not clock noise [9] but rather the noise of the TWSTFT process. The $\sigma_x(\tau)$ values in Figure 9 calculated using the hybrid method are in good agreement at small τ with the values of $\sigma_x(\tau)$ calculated from the 21 days of consecutive data. The hybrid method values fall well within the confidence limits of the $\sigma_x(\tau)$ values from the evenly spaced 21 day data set.

One important point to note is that for real experimental data there may be a large frequency offset between the clocks being compared. For evenly spaced data this offset is not relevant and can be ignored. However, for unevenly spaced data this is not true and the mean frequency offset must be removed before $\sigma_x(\tau)$ is calculated. A fractional frequency offset of $1{,}1 \times 10^{-14}$ was removed before using the hybrid method in Figure 9.

5. Conclusions

We have used two TWSTFT time series data sets to investigate the impact of unevenly spaced data on the calculation of $\sigma_x(\tau)$. We have analysed simulated data sets that have had points removed to match the typical Monday-Wednesday-Friday TWSTFT data patterns. $\sigma_x(\tau)$ was calculated from these sparse data sets using two techniques. One technique involves analysing the sparse data as if they were evenly spaced with an average time interval, and the second uses interpolated data to recreate an evenly spaced data set. Correction factors for both approaches have been calculated for noise processes ranging from WHPM to RWFM. For all of the noise processes except WHPM, the values of $\sigma_x(\tau)$ calculated with either of the two approaches converge on the correct values at large τ. However, significant errors may be introduced for small τ. Finally, we suggest techniques for estimating correct values of $\sigma_x(\tau)$ in situations where the type of noise is unknown or where more than one noise type is present, and apply these techniques to a set of real TWSTFT data.

Acknowledgements. The authors wish to thank Judah Levine, Don Sullivan, Matt Young (all from the National Institute of Standards and Technology), and Jim DeYoung (United States Naval Observatory) for their useful comments concerning this manuscript. We would also like to acknowledge the cooperation of Angela McKinley, Bill Powell, and Phu Mai of the United States Naval Observatory in the NIST-USNO time transfer experiment.

References

1. Allan D. W., Weiss M. A., Jesperson J. L., In *Proc. 45th Ann. Frequency Control Symposium*, 1991, 667-678.
2. Thomas C., Wolf P., Tavella P., Time Scales, *BIPM Monographie 94/1*, Sèvres, Bureau International des Poids et Mesures, 1994, 49-51.
3. Allan D. W., *IEEE Trans. Ultrasonics, Ferroelectrics, and Frequency Control*, 1987, **UFFC-34**, 647-654.
4. Allan D. W., Barnes J. A., In *Proc. 35th Ann. Frequency Control Symposium*, 1981, 470-474.

5. Stein S., In *Precision Frequency Control*, Vol. 2 (Edited by E. A. Gerber and A. Ballato), New York, Academic Press, 1985, 191-232, 399-416.
6. Vernotte F., Zalamasky G., Lantz E., In *Proc. 25th Ann. Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, 1993, 379-388.
7. Kasdin N. J., Walter T., In *Proc. 1992 IEEE Frequency Control Symposium*, 1992, 274-283.
8. Press W. H., Flannery B. P., Teukolsky S. A., Vetterling W. T., *Numerical Recipes in C*, 1st ed., New York, Cambridge University Press, 1988, 735 p.
9. Hackman C., Jefferts S. R., Parker T., In *Proc. 1995 IEEE Frequency Control Symposium*, 1995, 275-281.
10. Weiss M. A., Walls F. L., Greenhall C., Walter T., In *Proc. 9th European Time and Freq. Forum*, 1995, 153-165.

Received on 20 September 1995 and in revised form on 20 February 1996.