MULTISCALE CLOCK ENSEMBLING USING WAVELETS

Ken Senior
Naval Center for Space Technology
U.S. Naval Research Laboratory
Washington, DC 23075, USA
E-mail: ***ken.senior@nrl.navy.mil***

Don Percival
Applied Physics Laboratory
University of Washington
Seattle, WA 98105
E-mail: ***dbp@apl.washington.edu***

Abstract

There are currently many different timescale algorithms in use today ranging in application from scientific to commercial. While these algorithms differ in many respects and are sometimes tailored specifically for the intended application and mix of clocks involved, they all share the common goal of optimally combining the clock error difference observed or measured between a collection of clocks to form a reference timescale that is more stable than any of the constituent clocks. Of these algorithms, only a few are well suited to collections of highly disparate clocks. A new approach to forming timescales is presented here. This new multiscale ensemble timescale (METS) algorithm is based on a multiresolution analysis afforded by the discrete wavelet transform, is not dependent on a specific model for the clocks involved, optimally utilizes any mix of clocks both in terms of type and capability, and results in a reference timescale that is more stable than the constituent clocks over all scales (averaging intervals). The METS algorithm is presented in detail and is compared in a simulation study with a well-accepted timescale algorithm that uses a model-based Kalman approach.

INTRODUCTION

The widespread availability of high-performance clocks has led to considerable interest in the question of how to form a timescale based upon a collection of clocks. The usual purpose of a timescale is to improve upon the time kept by any individual clock and to offer protection against the failure of individual clocks. How best to form a timescale is not an easy question and has been the subject of numerous papers and five international symposia over the past 40 years. The difficulty in forming a timescale is tied to the fact that some clocks have good short-term performance (stability), but poor long-term performance, whereas the opposite is true for other clocks. Timescales are invariably based on the notion of weighting clocks by their expected performance, but a single weighting scheme has long been recognized as problematic because of differences in the inherent performance of clocks. As discussed below, certain model-based schemes that make use of the Kalman filter are capable of producing timescales with good short- and long-term stability, but these schemes are problematic for clocks that are a poor match for the presumed model and also might yield suboptimal intermediate-term stability.

In this paper, we propose a new approach to forming a timescale. Our approach is based on the concept of a multiresolution analysis (MRA) that is afforded by the discrete wavelet transform (DWT). The DWT of a time series yields a collection of wavelet coefficients, each of which captures the difference between localized averages of the time series over a particular scale (a span of time). A small (large) squared wavelet coefficient tells us that the time series under study was relatively stable (unstable) at the particular time and scale associated with the coefficient. An MRA is based upon the inverse DWT and creates a collection of new time series (the components of the MRA). Each component is associated with a particular scale. Addition of all the MRA components gives back the original series exactly. In our context, the time series of interest are time (phase) differences between clocks. The average of the squared wavelet coefficients for a collection of time differences for a particular scale can be used, in conjunction with an N-cornered-hat technique, to quantify the stability of individual clocks. The stability measures are then used to create weighted MRA components, which, when summed together, yields a timescale that has optimal stability over all scales. Our multiscale ensemble timescale (METS) algorithm does not assume any particular model for a clock and, thus, can handle clocks with a wide range of different stability properties.

The remainder of this paper is organized as follows. We first give some background on how to form a timescale from measurements that compare the time kept by pairs of clocks. We then discuss the basics of the discrete wavelet transform and the concept of a multiresolution analysis prior to introducing the METS algorithm. We present an outline of the key steps in the METS algorithm, after which we evaluate the algorithm in a simulation study that compares it to a model-based Kalman filter approach. We then state our conclusions and briefly discuss forthcoming work.

BACKGROUND

Given a set of N independent clocks, the phase (or time) error of each clock with respect to an arbitrary perfect reference clock at epoch t is labeled as $X_i(t)\ \ i=1,\ldots,N$. Such a perfect clock could only be accessed in so far as any of the individual clock errors are perfectly known. That is, given the exact errors of a clock, one can adjust the clock by that error and, hence, recover the perfect clock. However, practically speaking, one may only observe or measure the differences between a set of real clocks. From a set of N clocks, one may construct $N(N-1)/2$ possible clock differences. But, if a perfect measurement process is assumed, all clock difference information may be succinctly represented in only $N-1$ clock differences, say $X_{ir}(t)\equiv X_i(t)-X_r(t),\ i=1,\ldots,N, \mathrm{i}\neq \mathrm{r}$, where without loss of generality an arbitrary single reference clock r from among the clocks has been used to represent the differences. Note that the difference between any two clocks i and j may be determined through the simple difference, $X_{ij}(t)=X_{ir}(t)-X_{jr}(t)$. Of course, making perfect measurements of clocks is not possible and so measurements of the *N-1* clock differences will generally be made at discrete measurement epochs subject to some error,

$$Z_{ir}(t_k)=X_{ir}(t_k)+\varepsilon_{ir}(t_k), \tag{1}$$

$i = 1, \ldots, N, \mathrm{i} \neq \mathrm{r}$, $k = 1, \ldots, N_t$. It is assumed herein that all such measurement errors are mean zero and independent over the collection of clocks.

Given such a collection of such *N-1* clock difference measurements (1), the goal of a timescale algorithm is to determine or estimate as well as possible the individual clock errors. This is equivalent to constructing a new reference for the clock observations having less error than any of the constituent clock errors. Because the goal is to recover *N* quantities from *N-1* observations, the problem is rank-deficient, or ill-posed. Additional information must be specified in order to address this deficiency. Typically, this is handled in a statistical sense by imposing additional assumptions about the behavior of the collection of clocks as an ensemble that also exploit the independence of the clocks. For example, one may assume that on weighted average the ensemble errors are zero. That is, an ensemble average of the errors may be defined,

$$X_e(t_k) \equiv \sum_{i=1}^{N} w_i(t_k) X_i(t_k) = 0, \tag{2}$$

or, equivalently formed with respect to the reference clock,

$$X_{er}(t_k) \equiv \sum_{i=1}^{N} w_i(t_k) X_{ir}(t_k) = 0, \tag{3}$$

where heuristically the weights are chosen inversely to some measure of the magnitude of the errors of each of the clocks and normalized such that $\sum_{i=1}^{N} w_i(t_k) = 1$. For example, if the clocks were all equivalent in the nature of their noise, say $E[(X_i(t_k))^2] \approx \sigma^2$, then setting $w_i(t_k) = 1/N$ in (3) and noting the independence of the clocks, one would expect from the usual formula for the variance of the average of a collection of independent and identically distributed random variables with finite variance that $E[(X_e(t_k))^2] \approx \sigma^2 / N$. In other words, such a weighted average should, on statistical ensemble average, produce errors that are reduced by $1/\sqrt{N}$ as compared with each of the constituent clocks.

However, equation (3) has several practical problems. For one thing, clocks generally have deterministic offsets from one another, including time offsets, but also frequency (first derivative of time) and even drift (second derivative) offsets, as well as other deterministic effects. So, as new clocks enter the ensemble or older clocks leave it, (3) will result in deterministic changes in the ensemble as the constituents change. In order to address this shortcoming, a model of the clocks' behavior can be imposed to account for these deterministic effects and (3) correspondingly modified such that statistical assumptions about the ensemble of clocks be made with respect to the clocks' prediction errors instead of the errors directly. That is, given a model for the clock's deterministic behavior, let $X(t_k \mid t_k - \tau)$ represent the prediction of the clock to the current epoch t_k from some previous epoch $t_k - \tau$ based on the model. Then, define the ensemble by imposing the ensemble such that,

$$X_{er}(t_k) - X_{er}(t_k \mid t_k - \tau) \equiv \sum_{i=1}^{N} w_i(t_k) \left[X_{ir}(t_k) - X_{ir}(t_k \mid t_k - \tau) \right] = 0. \tag{4}$$

Here, the ensemble is defined not as the weighted average of clock errors, but is instead defined by its variations from clock predictions with respect to the model now imposed. In other words, the ensemble's variations from its predictions are the weighted average of variations of the clocks from their respective predictions. This constraint is similar to (3), but has the additional benefit that the ensemble average is not affected adversely as clocks enter or leave the ensemble.

There are still some drawbacks to constraint (4), though. In addition to having various deterministic effects, clocks' errors may also be driven by multiple and different noise processes. For example, one clock may have very large short-term variations, but remain more stable over the longer term, while other clocks may have just the opposite behavior, being very stable over short intervals (scales), but more divergent over longer intervals. In this case, a simple one-weight-per-clock constraint such as (4) will not necessarily produce an ensemble that is more stable than all constituent clocks at all scales. For example, one model that has been shown to be very useful for many clock variations is the perfect integrator model. This model consists of a clocks' phase error X, its frequency (first derivative of phase) Y, and its drift (second derivative of phase) error D, each perturbed by independent random walks, labeled RWPH, RWFR, and RWDR, i.e.,

$$\frac{d}{dt}\begin{bmatrix} X \\ Y \\ D \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}\begin{bmatrix} X \\ Y \\ D \end{bmatrix} + \begin{bmatrix} d\alpha \\ d\beta \\ d\gamma \end{bmatrix}. \tag{5}$$

In this model, a clock may have relatively small RWPH variations, while at the same time having large RWFR variations. But, with only one weight specified per clock, these independent variations cannot be constrained such that the resulting ensemble has both the smallest RWPH and RWFR (and RWDR) variations. For this reason, Stein **[1]** introduced for this specific perfect integrator model three separate weighting constraints in order to independently constrain each of the three noise types,

$$X_{er}(t_k) - X_{er}(t_k \mid t_k - \tau) \equiv \sum_{i=1}^{N} w_{1,i}(t_k)\left[X_{ir}(t_k) - X_{ir}(t_k \mid t_k - \tau)\right] = 0, \tag{6}$$

$$Y_{er}(t_k) - Y_{er}(t_k \mid t_k - \tau) \equiv \sum_{i=1}^{N} w_{2,i}(t_k)\left[Y_{ir}(t_k) - Y_{ir}(t_k \mid t_k - \tau)\right] = 0, \tag{7}$$

$$D_{er}(t_k) - D_{er}(t_k \mid t_k - \tau) \equiv \sum_{i=1}^{N} w_{3,i}(t_k)\left[D_{ir}(t_k) - D_{ir}(t_k \mid t_k - \tau)\right] = 0, \tag{8}$$

where each clock now has three weights, $w_{1,i}$, $w_{2,i}$, and $w_{3,i}$, each set inversely to the respective levels (variances) of RWPH, RWFR, and RWDR for the clock. Assuming the adequacy of this model, note that a clock's drift differs from its predictions exactly by its RWDR component. Thus,

$$E\left[\left(D_i(t_k) - D_i(t_k \mid t_k - \tau)\right)^2\right] = E\left[(d\gamma_i)^2\right] = 1/w_{3,i}.$$

Similarly, variations of the clock's frequency predictions can be shown to be exactly its RWFR component (plus its integrated RWDR). By imposing an ensemble constraint independently for each noise type, an ensemble timescale is produced with the aim of getting a small ensemble variation over all scales. Note that, because the three types of noise processes represent independent perturbations of the clock, the addition of three separate constraints (6)-(8) to the system of *N-1* clock differences is one solution to the problem of rank deficiency.

The model shown here is very easily implemented using a Kalman filter. But, in some situations, other noises may be present in a clock's error for which a Kalman filter is not the best or easiest choice to implement, as in the case of flicker noises. In the following sections, a non-model based approach to the timescale problem is presented. This wavelet-based multiscale ensembling approach can generate a multiscale ensemble timescale (METS) that is better (more stable) at all scales than any of the constituent clocks in the ensemble.

DISCRETE WAVELET TRANSFORMS

The METS algorithm described in the next section is based on decompositions of clock signals using discrete wavelet transforms (DWT). A short description of wavelets and wavelet transforms will now be presented, though the reader is directed to **[2]** and **[3]** for more comprehensive treatments.

Similar to the discrete Fourier transform (DFT), a DWT is a linear orthonormal transform that utilizes an orthogonal basis set against which a signal can be decomposed. Unlike a DFT, where the basis set consists of stretchings and scalings of sines and cosines, the basis set for a DWT consists of stretchings and scalings of small "wavelets" that have only compact support, or finite extent. A few examples of some basic wavelet shapes ("mother wavelets") from which various wavelet bases are generated are shown in Figure 1 below.

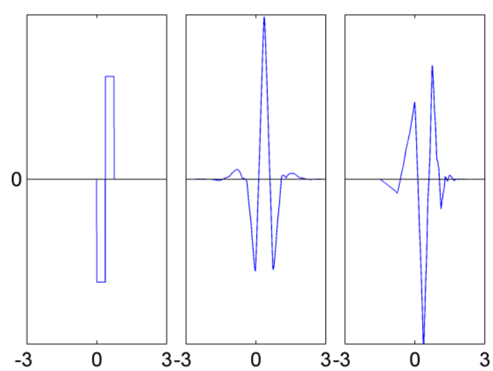

Figure 1. Three example mother wavelets, the Haar (left), the Coiflet (middle), and Daubechies (right) wavelets.

The action $\boldsymbol{\mathcal{W}}X$ of the DWT on a given discrete signal $X = X(t_k)$, $k = 1, \ldots, N_t$, is equivalent to periodized convolution of the signal against stretchings and scalings of its mother wavelet, generating as its output wavelet coefficients $\mathbf{W}$,

$$\boldsymbol{\mathcal{W}}X \equiv \mathbf{W} = \left[W_1, W_2, \ldots, W_{L_0}, V_{L_0}\right]^T \tag{9}$$

with the property that $X = \boldsymbol{W}^T \mathbf{W}$, where L_0 is the number of desired terms (levels) in the decomposition, a value which must be pre-selected in the construction of the wavelet bases and whose upper limit is tied to the sample size, $N_t = 2^{L_0}$. The above transformation yields L_0 subvectors of wavelet coefficients (each associated with a different scale) and one subvector of so-called scaling coefficients that captures large-scale averages of *X*. We can reconstruct *X* perfectly by applying the inverse wavelet transform ($\boldsymbol{W}^{-1} = \boldsymbol{W}^T$), which can be manipulated to yield an additive decomposition known as a multiresolution analysis (MRA):

$$\begin{aligned} X(t_k) &= \sum_{l=1}^{L_0} W_l^{\ T} W_l + V_{L_0}^{\ T} V_{L_0} \\ &= \sum_{l=1}^{L_0} D_l(\tau_l, t_k) + S_{L_0}(t_k) \end{aligned} \tag{10}$$

The D_l terms are referred to as the "details" of the decomposition, while the last term S_{L_0} is called the "smooth." Each of the details captures variations of the signal over that particular scale τ_l, while the smooth (based upon the scaling coefficients) captures all the remaining variation in the signal. For example, given the sampled signal *X* shown in Figure 2 below, its corresponding wavelet decomposition using the LA(8) wavelet and L_0=10 levels is shown in Figure 3.

Figure 2. An example discrete signal.

In addition to providing a scale-based decomposition, the signal the norm-preserving nature of the DWT also allows an energy decomposition of the signal as well, referred to as the wavelet variance. This variance is defined by

$$v_X^2(\tau_l) \equiv \mathrm{var}(W_l) \tag{11}$$

and it can be shown that for a very wide class of signals and for an appropriately chosen wavelet that $\sum_{l=1}^{\infty} v_X^2(\tau_l) = \mathrm{var}(X)$. One such class of processes is the fractionally differenced (FD) noises whose spectral density is given by

$$S_X(f) = \frac{\sigma_\varepsilon^2}{(4\sin^2(\pi f))^\delta} \approx \frac{\sigma_\varepsilon^2}{(2\pi f)^{2\delta}} \tag{12}$$

which are approximately equivalent to the power law processes commonly used to model clock behavior. It can be shown that the variance **Error! Reference source not found.**(11) will be well-defined for FD processes provided the underlying wavelet has sufficient "width" **[2,3]**.

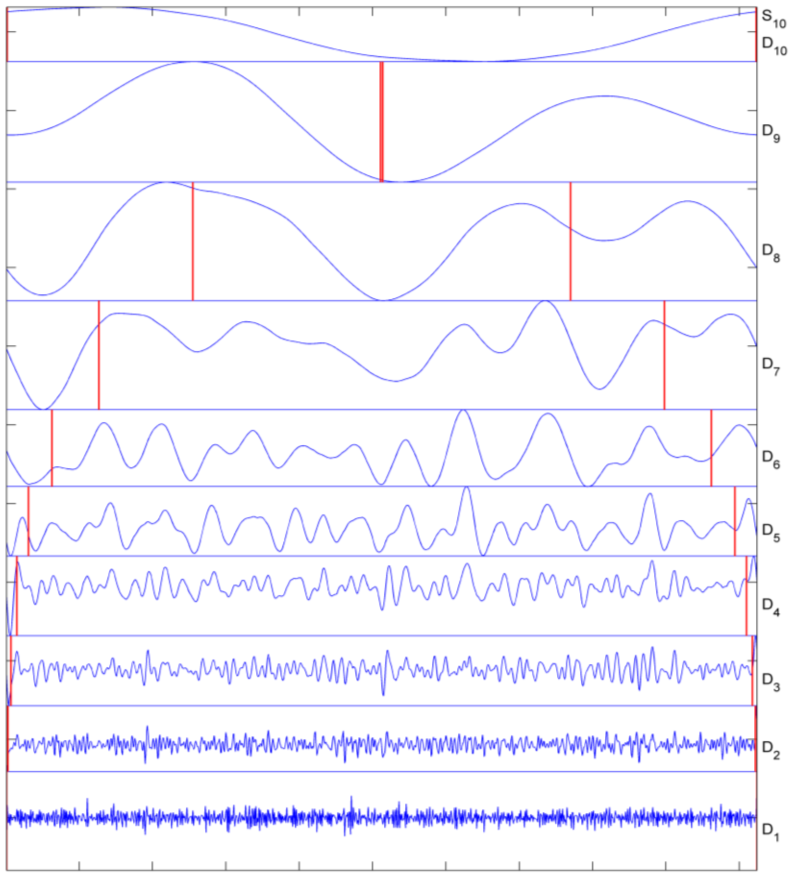

Figure 3. A wavelet decomposition of the signal shown in Figure 2 using an LA(8) mother wavelet and $L_0=10$ levels in the expansion.

METS ALGORITHM

Returning now to the timescale problem, assume given a set of clock difference measurements (1). The algorithm for generating METS can be described in four steps:

STEP 1. Apply an appropriately chosen wavelet transform separately to the collection of clock difference signals to obtain a set wavelet of wavelet coefficients for each series, as well as the clock-to-clock wavelet covariances between each series,

$$\begin{aligned} &\boldsymbol{\mathcal{W}} Z_{ir} = \mathbf{W}_{ir} = \left[W_{ir,1}, W_{ir,2}, \ldots, W_{ir,L_0}, \mathrm{V}_{ir,L_0}\right]^T, \\ &\mathrm{cov}(\mathbf{W}_{ir}, \mathbf{W}_{jr}), \end{aligned} \quad (13)$$

$i, j = 1, \ldots, N \;\; (i, j \neq r)$.

STEP 2. Next, over each scale τ_l, $l = 1,\dots,L_0$, apply N-Cornered Hat (c.f., **[4,5]**) to the *N-1* wavelet covariances for that scale in order to recover *N* wavelet variances; in other words, to recover separately the wavelet variances of each clock from the reference clock variance,

$$\left\{\mathrm{cov}(W_{ir,l}, W_{jr,l})\right\}_{\substack{i,j=1\\ i\neq r\neq j}}^{N}, \quad \xrightarrow{\text{N-HAT}} \quad \left\{v_i^2(\tau_l)\right\}_{i=1}^{N} \tag{14}$$

$l = 1,\dots,L_0$.

STEP 3. Per-clock and per-scale weights are now defined (and normalized) by setting,

$$\tilde{\alpha}_i(\tau_l) = 1/v_i^2(\tau_l), \qquad \alpha_i(\tau_l) = \tilde{\alpha}_i \Big/ \sum\nolimits_l^N \tilde{\alpha}_i \tag{15}$$

$i = 1,\dots,N$ and $l = 1,\dots,L_0$. Note that $\sum_{i=1}^{N} \alpha_i(\tau_l) = 1$ for each *l*.

STEP 4. Finally, using the weights (15), form the weighted sum of the wavelet coefficients in (13) and apply the inverse wavelet transform in order to form the METS timescale referenced to the original reference clock,

$$X_{er} \equiv \sum_{l=1}^{L_0}\sum_{i=1}^{N} \alpha_i(\tau_l) W_l^T W_{ir,l} + \sum_{i=1}^{N} \alpha_i(\tau_{L_0}) V_{L_0}^T V_{ir,L_0} \tag{16}$$

where it is noted that the weights utilized for the smooth term are the weights corresponding to the longest scale τ_{L_0}. Also, the reference clock is recovered in (16) because of linearity of the DWT and inverse DWT operators and normalization of the weights.

The only requirement necessary for the above algorithm to properly generate a METS timescale is that sufficient conditions are met such that the wavelet covariances (14) exist. However, this requirement is very minimal, since, for example, if the set of clocks have random variations that are well-modeled by finite sums of FD processes, then an appropriate wavelet may be found such that (14) is well-defined. Moreover, the class of FD processes essentially encompasses every possible noise process that has been observed to date in actual clocks. Deterministic variations including polynomial offsets of the clocks may also be admitted, provided that a wavelet is chosen with enough internal differencing operations **[3]**.

EXAMPLE SIMULATION 1

Consider a 12-clock example consisting of three classes of clocks having different timekeeping capabilities determined by random variations that are the sum of two FD noises (12) $\delta = 1, 2$ with levels given in Table 1 below. In other words, each clock is modeled approximately by a RWPH and RWFR with different levels of each noise as specified in the table. The frequency instability of 12 such simulated clocks is shown in Figure 4 below shown in the blue series, as measured using the Allan variance. Clock difference measurements were formed from the simulated clocks and the METS algorithm then applied to the difference measurements using an LA(8) wavelet and "reflection" boundary conditions. The results of the METS algorithm are also shown in Figure 4, where the frequency instabilities of the METS estimates of the clocks, as well as the instability of the resulting METS timescale as measured with the Allan variance and various wavelet variances, is shown. The METS timescale clearly performs better at all scales than any of the constituent clocks by all frequency instability measures shown. As the example illustrates, the multiple per-scale weights that are determined separately over each scale result in a timescale that minimizes variations over each scale, separately.

Table 1. Spectral density levels of FD noise for an example of three classes of clocks.

CLASS	δ	σ_{ε}
I	$\delta_1 = 1$	$\sigma_{\varepsilon_1} = 0.5$
	$\delta_2 = 2$	$\sigma_{\varepsilon_2} = 0.1$
II	$\delta_1 = 1$	$\sigma_{\varepsilon_1} = 2$
	$\delta_2 = 2$	$\sigma_{\varepsilon_2} = 1/35$
III	$\delta_1 = 1$	$\sigma_{\varepsilon_1} = 10$
	$\delta_2 = 2$	$\sigma_{\varepsilon_2} = 0.005$

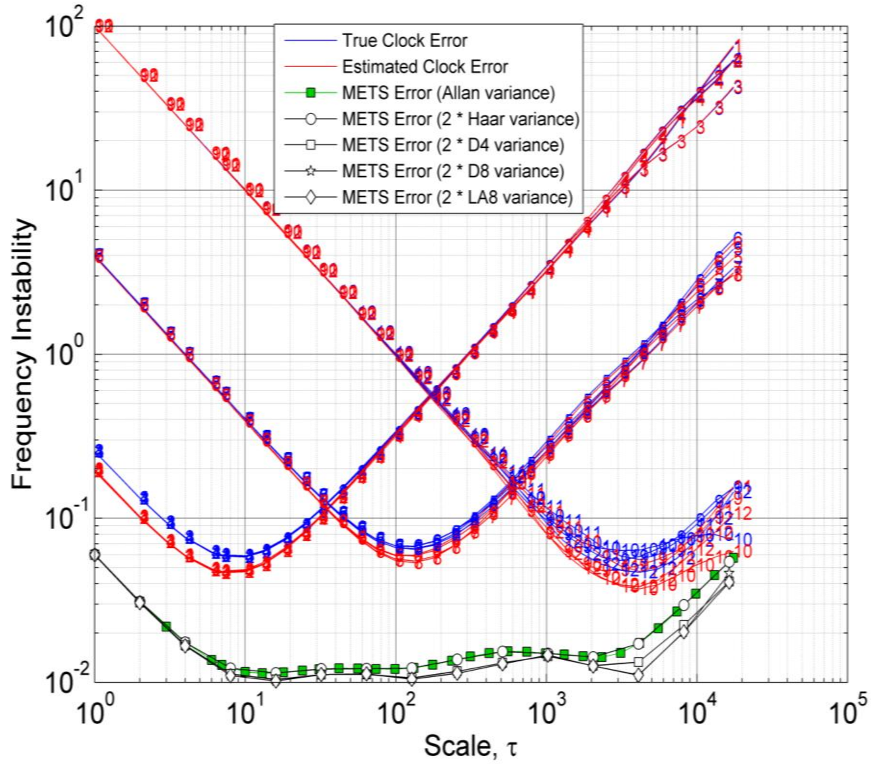

Figure 4. Frequency instability for twelve simulated clocks (blue) with FD spectral values given in Table 1 as measured with the Allan variance, three clocks from each class. The frequency instability of the resulting METS estimates of each clock are also shown (red) along with the frequency instability of the METS timescale, measured using variance wavelet variances as well as the Allan variance (green).

Note that a model-based approach such as the Kalman filter implementation of (5) can produce similar stability improvements over multiple scales by weighting each modeled noise type separately, as in the addition of constraints (6)-(8). A Monte Carlo comparison of METS with the Kalman approach using 20 simulations was conducted. Each Monte Carlo simulation consisted of 12 clocks whose errors are driven by the same FD noises in Example 1 above. For the Kalman approach, the model and constraints (5)-(8) were applied. Frequency instabilities for the simulated clocks as well as the METS and Kalman timescales, all averaged over the 20 simulations, are shown in Figure 5 below. In both cases, the resulting timescales outperform all of the constituent clocks. Although only marginally, the METS timescale appears more stable over most scales, while the Kalman-generated timescale outperforms at the longest scales. Additional analysis and simulation are required to determine these differences robustly, though.

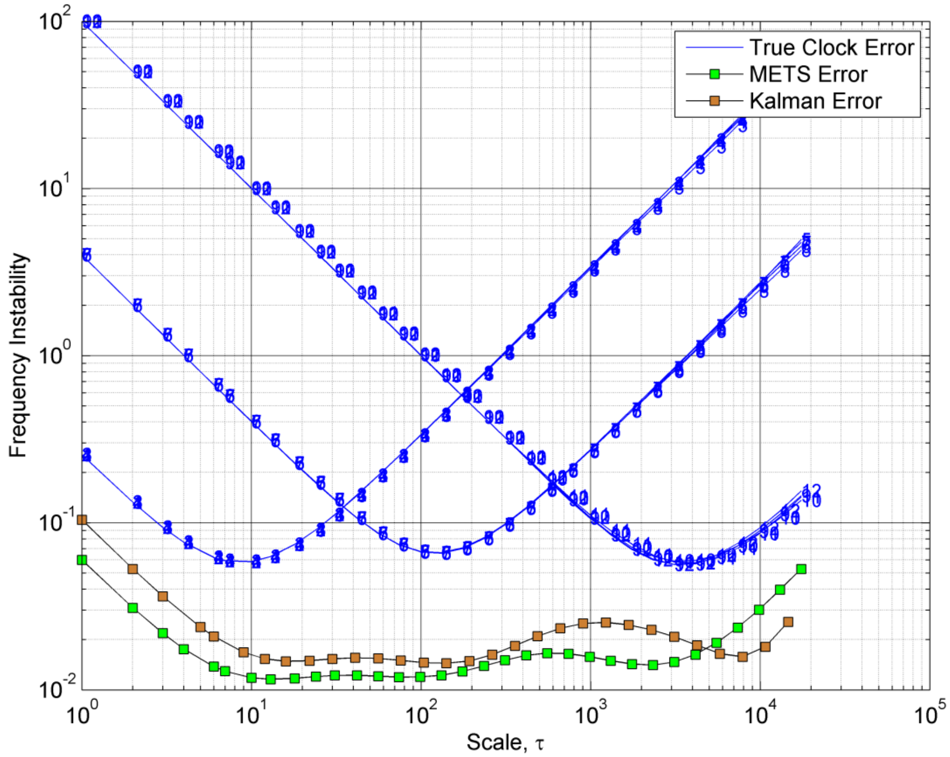

Figure 5. Monte Carlo (20 simulations) comparison of the METS algorithm with the Kalman approach. Both approaches assume the clocks' errors are driven by the FD noises specified in Table 1, while the Kalman approach assumes additionally the model (5). The plot shows the average over the 20 separate simulations of the frequency instability of 12 simulated clocks (blue), METS timescale (green), and Kalman results (orange).

EXAMPLE SIMULATION 2

As stated earlier, one major advantage of METS over a Kalman-based approach is that it may be implemented without consideration of a model for the dynamics of the clocks. Therefore, noise processes such as flicker noises, which are typically difficult to model in Kalman filters, are easily admitted in METS. Consider an example simulation consisting of six clocks whose random variations are the sum of the three FD noises (12) $\delta = 1, 3/2,$ and 2 with levels given in Table 2 below. These FD noises are each consistent with a random walk in phase (RWPH), a flicker frequency (FLFR), and a random walk in frequency (RWFR), respectively. The frequency instability of this six-clock simulation is shown below in Figure 6 below (blue series). Note from the figure that the flicker noise process dominates for more than a decade of scales. Once again, clock difference measurements were formed from the simulated clocks, followed by application of the METS algorithm using the same LA(8) wavelet and "reflection" boundary conditions as above. As Figure 6 shows, the resulting frequency instability of the METS timescale is once again better at all scales than any of the constituent clocks.

Table 2. Spectral density levels of FD noise for example 2 consisting of three noise types.

δ	σ_{ε}
$\delta_1 = 1$	$\sigma_{\varepsilon_1} = 0.5$
$\delta_2 = 3/2$	$\sigma_{\varepsilon_2} = 0.25$
$\delta_3 = 2$	$\sigma_{\varepsilon_1} = 0.005$

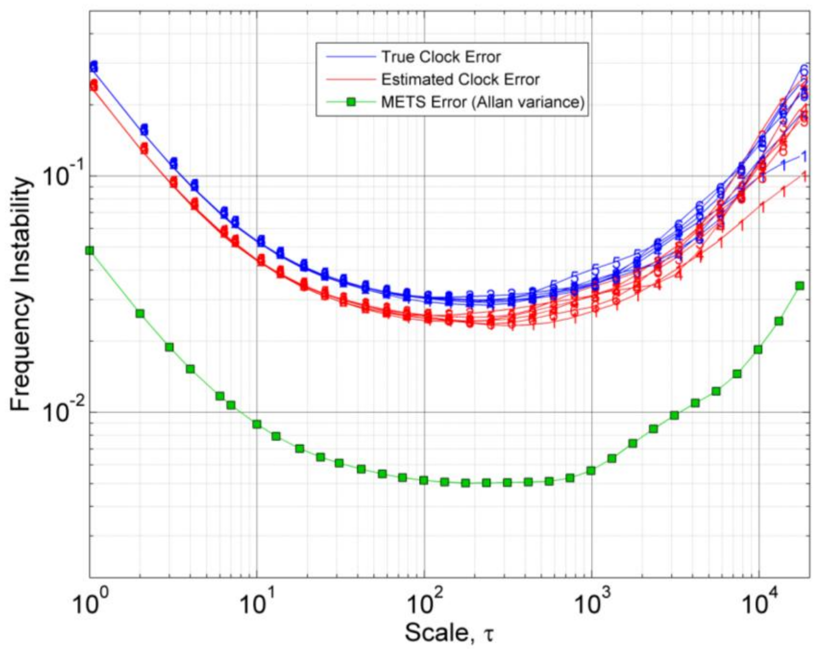

Figure 6. Frequency instability for six simulated clocks (blue) with FD spectral values given in Table 2 as measured with the Allan variance. The frequency instability of the resulting METS timescale estimates of each clock are also shown (red) along with the frequency instability of the METS timescale (green).

CONCLUSIONS

We have introduced a new algorithm – the multiscale ensemble timescale (METS) algorithm – for forming a timescale. This algorithm is based upon the multiresolution analysis afforded by the discrete wavelet transform (DWT). The basic idea is to decompose measurements of clock differences into variations at different scales (spans of time), to form an optimal timescale at each individual scale based upon the stability of the clocks at that scale (as determined by an N-cornered-hat method) and then to form an overall optimal timescale by combining the individual timescales using the inverse DWT. This approach does not make use of any specific clock model and, hence, can handle clocks that do not conform well to simple models, such as the perfect integrator model. We have demonstrated through limited simulation studies that the METS algorithm can outperform a well-known model-based timescale algorithm that is implemented in terms of a Kalman filter. While the METS algorithm is intuitively quite appealing, there are many technical details that need to be worked out before this algorithm can be used for forming timescales in practical situations. Subject to the successful outcome of this future work, the METS algorithm has the potential for providing an elegant solution to the long-standing property of how best to combine clocks with differing stability properties together to form the best possible timescale. A more detailed mathematical treatment of the algorithm presented here will be available in **[6]**.

REFERENCES

[1] S. R. Stein, 2003, *"Time Scales Demystified,"* in Proceedings of the 2003 IEEE International Frequency Control Symposium and PDA Exhibition Jointly with the 17th European Frequency and Time Forum (EFTF), 4-8 May 2003, Tampa, Florida, USA (IEEE 03CH37409C), pp. 223-227.

[2] D. B. Percival and A. T. Walden, 2002, **Wavelet Methods for Time Series Analysis** (Cambridge University Press, Cambridge, UK).

[3] D. B. Percival, 2002, *"A Tutorial on Stochastic Models and Statistical Analysis for Frequency Stability Measurements,"* presented at the 4th International Symposium on Time Scale Algorithms, 18-19 March 2002, Sèvres, France.

[4] P. Tavella and A. Premoli, 1993 *"**A** revisited three-cornered-hat method for estimating frequency standard instability,"* **IEEE Transactions on instrumentation and Measurement, IM-42,** 7-13.

[5] F. Torcaso, C. R. Ekstrom, E. A. Burt, and D. N. Matsakis, 1999, *"Estimating Frequency Stability and Cross-Correlations,"* in Proceedings of the 30th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 1-3 December 1998, Reston, Virginia USA (U.S. Naval Observatory, Washington, D.C.), pp. 69-82.

[6] D. Percival and K. Senior, 2011, *"A Wavelet-Based Multiscale Time Scale Algorithm,"* **IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control**, to be submitted.