Cascaded Clocks Measurement and Simulation Findings

Don Chislow and George Zampetti
Telecom Solutions

1. Introduction

This paper will examine aspects related to network synchronization distribution and the cascading of timing elements. Methods of timing distribution have become a much debated topic in standards forums and among network service providers (both domestically and internationally). Essentially these concerns focus on the need to migrate their existing network synchronization plans (and capabilities) to those required for the next generation of transport technologies (namely, the Synchronous Digital Hierarchy (SDH), Synchronous Optical Networks (SONET), and Asynchronous Transfer Mode (ATM). The particular choices for synchronization distribution network architectures are now being evaluated and are demonstrating that they can indeed have a profound effect on the overall service performance levels that will be delivered to the customer. The salient aspects of these concerns reduce to: (1) identifying that the devil is in the details" of the timing element specifications and the distribution of timing information (i.e., small design choices can have a large performance impact), (2) developing a standardized method of performance verification that will yield unambiguous results, and (3) presentation of those results. Specifically, this will be done for two general cases: an ideal input, and a noisy input to a cascaded chain of slave clocks.

The method most commonly used by network providers in recent years is a master–slave or hierarchical timing configuration. An attractive feature of a hierarchical timing configuration is that existing digital transmission facilities, between digital switching nodes can be used for synchronization distribution. At the same time this will not diminish the traffic carrying capacity of a particular carrier system. Care must be exercised, however, in the choice of the primary and secondary transmission facilities and routes when designing this synchronization network because the integrity of the facility directly affects the service availability to the subscriber. Since timing error will increase with hierarchical level the objective is not to have too many levels. Further, additional levels and more complex topologies leave the network vulnerable to the formation of timing distribution loops[1]

An unfortunate consequence of a synchronous networks such as SDH or SONET is that each

---

[1]A timing loop is a configuration wherin the master or controlling timing unit is influenced or steered by a lower level timing element. This arrangement will contaminate the behavior of the higher order timing element and will result in performance degradations.

network element (NE) by definition terminates and recycles timing information. This is because each SDH NE has a clock in it to maintain the frequency tolerances required for transmission continuity. The process of synchronizing these clocks leads to a delay "breathing" phenomenon as they try to maintain lock with a reference.

In order for clocks in telecommunications equipment to maintain lock with an upstream master reference signal they are required to develop an estimate of the phase and frequency characteristics of that reference with respect to a local oscillator. The ability to calibrate the frequency of a slave clock to a network timing reference consequently becomes a critical factor because the calibration will translate directly into the slip performance during an outage. When clocks are cascaded the characteristics of the local oscillator, its control circuitry, together with noise all contribute additional phase variations that work to contaminate the estimate. So much so in fact that the error can be much greater than the drift of the oscillator and this can compromise investment in implementing a quality oscillator.

These facts serve to underscore one of the guiding principles of synchronization planning. That is, to minimize the number of slave clocks through which timing is chained.

2. Synchronization Distribution

The distribution of telecommunication network synchronization for the Plesiochronous Digital Hierarchy (PDH) is to a large extent hierarchical in plan (or master–slave arrangements). Furthermore, the topological flow of synchronization trails through a network are to a large extent influenced by the manner in which the network is configured to transport information. The reasons for doing this are essentially twofold. First, that management and administration of the synchronization follows the same path that the information payloads follow, and second as a result of this no new facilities are required to provide this capability. The first situation allows for straightforward trouble isolation and problem resolution. In the second case one can realize considerable economic efficiency. Consequently, there is strong motivation to build upon these advantages and integrate future services into this plan. The introduction of transport technologies such as: SONET, SDH and ATM present new challenges to the original synchronization distribution plan.

The basic objective of synchronization distribution is the creation of equal time scales at each location (within some time or frequency error budget. Typically, this budget is arbitrarily based on service objectives and bounded by performance and/or economic constraints). To achieve this the synchronization element[2] (SE) must follow the master not the reference. The distinction being that the master timing source dictates the performance of the network, whereas the local reference (which in many cases is a noisy representation of the master) may have deviated from it. It is therefore the function of the SE to detect any deviation from the last known satisfactory estimate of the master and disconnect from the reference before it can be misinformed by such deviations and forced in the wrong direction. This functionality requires the SE to support multiple tasks. Principally those are: (1.) minimizing the noise, and (2.)

---

[2] a synchronization element is used ubiquitously to refer either to an SETS, an SSU, a BITS, or NE clock. In general, any device that generates or re-generates timing information; colloquially, a "clock."

detecting and removing transient phase movements. The latter serves to minimize the error when the SE is in a hold–over state.

The synchronization network reference chain intended for deploying SDH follows a quasi-hierarchical approach shown in Figure 1[1,2].

The nodes are to be connected by network elements (NEs) each with internal timing elements compliant with a developing international recommendations. As illustrated, the chain should not exceed K slave clocks compliant with recommendations G.81s[3] and G.812[4]. Further, it is stated that the value of N will be limited only by the quality of timing required by the last network element in the chain typically. (Currently this number is considered to be approximately 10 maximum). It is assumed that each Synchronization Supply Unit (SSU) or Synchronous Equipment Timing Source (SETS) represents a physically distinct building location. Note also that the function of the SSU is to only provide a timing interface for the SDH elements and the immediate tributary interface NEs to which they may connect.

Note that each SDH NE that receives a G.811[5][3] traceable signal and generates a new output signal (in effect re–cycling the timing) for use by another SE in the synchronization trail represents an intervening clock to a down–stream office or NE. It was pointed out that each intervening clock will degrade the timing stability of the entire synchronization trail by some amount. Further, the number of intervening clocks that can be cascaded together will be limited not only by the stability of the regenerated and transported signals, but also by the ability of each SE to prevent clock re–arrangement phase movement, as well as other transient activity from propagating through the network. Overall, it seems reasonable that network performance will be improved if the number of intervening synchronization elements is kept to a minimum. There is however, serious interest on the part of service providers to consider more complex transport architectures such as self healing ring topologies. Such discussions have placed the number N as high as 22.

Alternatively, proposals for distribution of synchronization information within a SONET based network follow the architecture shown in Figure 1b. The difference here being that the Building Integrated Timing Supplies (BITS) serve to externally time each SONET NE that would be used to distribute timing. In fact there may not be any intervening "line timed"[4] SONET NE that is in the timing distribution path. The BITS clocks serve to provide an equal level timing interface which will prevent transient signal propagation and hold–over stability in excess of the NE itself. Moreover, the BITS clocks also supply timing to the entire office. This is not the case, however, for the SSUs in the SDH network. In many administrations that intend to implement SDH networks it is common practice to segregate timing entities within an office and dedicate them to specific service technologies, i.e., data, voice, etc. The reasons for this are primarily a combination of historical precedence (i.e., it was always done that way), and differing philosophical approaches to office operations and practices.

The control of (E1[5], or DS1) slips in the PDH and pointer adjustment activity in SDH, or SONET requires that all E1/DS1 and SDH/SONET timing devices operate at the same frequency within

---

[3]Effectively, a reference quality signal.

[4]A network element timing configuration in which the output (any direction) is determined by the input line signal.

[5]E1 is a 2.048 Mb/s international primary rate signal. DS1 is a 1.544 Mb/s American primary rate signal.

some achievable bound. This can be accomplished provided that the frequency characteristics of these synchronization elements (SEs) can be traced to a primary reference source (PRS) or G.811[5] type device. This concept is in agreement with the hierarchical master–slave timing distribution method just discussed. A pre–requisite for implementing this methodology is that all the clocks in the timing distribution chain remain synchronized to a free–running clock of equal or higher performance level. This will ensure that during failure conditions slave clocks will be able to maintain synchronization with each of the other SEs. As mentioned earlier, the process of synchronizing these clocks will lead to a breathing phenomenon as they try to maintain lock with one another.

3. Synchronization Element Clock Model

The synchronization element in this investigation is viewed as a "narrow bandwidth" digitally controlled clock. An internal block diagram illustrating the components that have been modelled is provided in Figure 2.

The classic type 2 loop filter (two integrators) that is used in the model study reflects the majority of known telecommunications NE clocks. Noise is initially introduced through the oven controlled crystal oscillator (OCXO). This represents a simple random walk–frequency modulated (RW–FM) process characteristic of crystal oscillator technology. The level of the random walk noise process is set to 5×10^{-11} at 1000 seconds. Effectively what is presented is a servo–control mechanism that will suppress low frequencies from the OCXO (by roughly 40 dB/decade) and high frequencies from the reference input (by approximately 20 dB/decade).

Peaking effects from the input reference are limited by design of the quantizing phase–detector, and represent a challenge to properly suppress in the face of introducing additional phase error by applying tighter filter bandwidths. Therefore properly managing changes in the loop gain so that corresponding increase or decrease in the system bandwidth does not allow wander to accumulate as the network size increases, or the phase error at any one particular clock is excessive is the ultimate goal.

4. Measurement Tools

The primary reason to select a particular measurement parameter is intimately related to its fundamental properties: time or frequency. Spectral characteristics are crucial because appropriate communication filters needed to be developed. Similarly, "time–difference" is important because time and phase relationships are key from the telecommunications perspective. Principally, because the general level of customer performance is related to slip activity in NE buffers. Further, the telecommunications network when studied was found to be similar to a measurement system as opposed to those timing characteristics that are indicative of a frequency standard. Because of this, the fundamental property of interest (for network measurements) becomes that of phase or time (e.g., slips), not frequency as would be the case for a frequency standard. Examples of this include network synchronization, phase–locked servo systems and time distribution systems.

In order to extract meaningful information from the data set generated from the model described in the previous section, a set of comprehensive measurement criteria must be agreed. It has been demonstrated that neither time or frequency domain characteristics alone are sufficient to properly describe the performance of clocks. It is rapidly becoming common practice to employ (at minimum) two criteria, those being, maximum time interval error (MTIE) and time variance (TVAR) type measurements to accomplish the objectives stated above. Essentially these criteria have stood the test of time and practice and have emerged as a minimal set of measurement tools.

Attributes common to each criterion are born out by the measurement process itself, that is, each can be viewed as a two stage process: a linear filter, followed by a statistical estimate. The measures outlined in the following have been shown to accurately and consistently describe real network situations. Further, they have utility as diagnostic tools. The significance of this is explained in what follows.

4.1 MTIE

MTIE is reasonably straight forward to describe. It represents the peak to peak time error (variation) between the device under test and an arbitrary reference in a given observation interval. Another interval will yield another MTIE value. Typically, one will view the error signal with increasing time, sliding the observation interval along in the process to obtain an historical record of this signal. The MTIE for the complete data set is the maximum of all the individual MTIE samples. It is noteworthy to point out that there is no filtering action that takes place in the computation of MTIE[6]. The process is simply one of peak detection and memory storage.

The purpose behind maximum time interval error as a performance measure is to constrain peak variations in network timing signals. As such MTIE is well suited to characterize network synchronization effects. Establishing an MTIE bound seeks to ensure that clock phase movement will not accumulate significantly during reference or hardware impairments. By controlling phase movement the number of slips which accumulate in the PDH network and down stream reference switching occurrences are limited. Further, by prescribing a phase slope requirement (via an MTIE specification) will help ensure that error conditions will not propagate.

This process is particularly useful to describe network transient events and slip phenomena because it directly corresponds to the peak–to–peak fill or depletion of data in network buffer stores. It should be noted however, that similar phenomena is also related to the frequency at which these movements occur. Unfortunately, MTIE does not provide information about rate of change of phase movement in a buffer or timing reference signal qualification (yet another source of error). See Figure 3.

4.2 TVAR

These difficulties call for the investigator to treat the problem in a different fashion, that is to examine the broadband $\frac{1}{f^n}$ noise characteristics of the error signal. This is accomplished

with the use of the well established Modified Allan Variance[4]. For the purposes of the telecommunications community this criterion has been modified to better represent those needs and has taken the form of a new measure known as time variance (TVAR). TVAR is the square of a 2 sample standard deviation and because of that demonstrates similar computational properties to that of the standard variance, with one notable exception: a single variance estimate is distribution dependent and will exhibit high scatter if the noise process being measured is divergent in the window of observation. Alternatively if two sample variances are calculated and then averaged until all the data are exhausted, the result is a convergent sequence.

The square root of TVAR (or time deviation TDEV) is proportional to the rms change in the mean value of the time errors averaged over an interval, τ. This is the principle difference between the TVAR and the standard variance. The standard variance computes the *rms* level of time values not average differential of those same values. It is worth mentioning that TVAR is normalized so that it will reduce to the standard deviation for a white noise phase modulated process. The utility of this type of variance is that it is extremely efficient in computing a wide band spectral density.

Thus TVAR represents the effective power output of a software–based filter the input of which is the phase modulated waveform. The filter characteristics are characterized by a bandpass filter. The upper cut–off frequency of this filter is approximately equal to the reciprocal of the observation period. The lower cut–off frequency is roughly one decade below this value and the filter response peaking occurs at about half the upper frequency value. See Figure 4.

4.3 ZTIE

Examining both MTIE and TVAR one may discover yet a third measure that may prove useful. Colloquially, this is referred to a ZTIE, or Z–transformed TIE[6]. It is an intermediate process between MTIE and TVAR because the computed value is the peak of the averaged first difference of the time sample values. ZTIE captures the peak power in a manner analogous to MTIE. In addition, an analogous bandpass filter function filter function similar to that for TVAR is employed, the center frequency of which, is controlled by the choice of averaging time, τ. The conceptual relationship requires slightly more explanation, but crudely, ZTIE provides a measure of peak power measured through a bandpass filter. See Figure 5.

By building upon the computational utility of TVAR certain efficiencies may be realized using ZTIE. Since this criterion retains peak power, which was explained to be very valuable with regard to providing network performance information, it is reasonable to conclude that to prescribe certain network performance limits that ZTIE would find utility as a performance monitoring metric. Principally because it can be calculated and acted upon while in service. Further, since the computation of ZTIE does not have to repeat the difference operation, as for TVAR, it can be thought of as representing a running snap–shot of peak network phase and frequency performance effectively marrying the attributes of MTIE and TVAR.

As shown in reference [7], ZTIE is very effective in illustrating both the short term peak jitter noise component as well as the long term frequency bias. The advantage gained using ZTIE is that in essentially one simple computation each of these characteristics may be derived. Moreover, ZTIE provides a utility as a calibration tool because if the computation are done

in–service then the frequency of the slave clock can be steered to the network reference in a very orderly fashion. This capability has direct implications for the network slip performance during reference outage periods.

Figure 5a depicts the computational concepts of each of the data gathering processes by applying the appropriate filters to the observed data.

The block averaging exhibits time dependent low pass filter characteristic the main lobe occurring at 1/T. The averaging function serves to suppress jitter and discriminate between flicker noise PM and white noise PM. This low pass filter attribute has been demonstrated to be significant in removing strong high frequency signals often found in telecommunications timing signals.The first difference function, shown next, provides suppression of the f^{-n} divergent power law noise process. This is analogous to a single pole filter producing 20 dB per decade roll–off noise rejection. This is adequate to ensure a stationary noise process for input phase noise, the dominant component of which is white FM. This function has the additional feature that it limits phase bias as well. Finally the second difference function provides an additional degree of noise suppression equivalent to a second order filter producing 40 dB per decade roll–off. This will ensure a stationary process for noise that behaves like random walk FM, which is the most divergent type of noise observed for general oscillator applications.

The advantage gained by separating these various filter processes into their constituent parts is that it becomes evident that each can convey useful information. Specifically, by examining the result of the first difference operation it can be observed that it this operation will not suppress the frequency bias term.The consequence of this is that it ZTIE correctly reflect frequency drift and transient offsets that TVAR will suppress. Moreover, because ZTIE is well behaved for non–divergent noise processes and exhibits a similar trend to that of TVAR it is much more valuable than MTIE in representing the correct trend of the noise process. Note that for conciseness the simulation results are presented for the TVAR case only. Since the simulations were run for Gaussian noise processes, the peak performance is well behaved. The MTIE and ZTIE data for this well controlled case provides little additional insight.

5. Simulation

The simulation results provided build upon earlier efforts to diagnose the statistics of telecommunication network timing performance[7]. The results enable an analysis of timing distribution topologies and permit educated speculation about the effect of degraded synchronization conditions on network performance. As was mentioned earlier these results have implications for the current asynchronous network design, but are particularly relevant to SONET, SDH and ATM network engineering (i.e., with regard to accommodation of timing signal variations and slip performance).

It is necessary to define the following model parameters to adequately define the timing signal generator at each node. For each of the cases examined the following constants apply:

1. Simulation time: 500K seconds
2. Proportional control factor: 221

3. Integral control factor: 8192

4. Damping factor: 3.0

5. Sampling time: 1.0 seconds

The phase–lock loop (PLL) parameters within the synchronization elements are representative for stable slave or transit node (OCXO) type timing devices. In the noisy environment scenario the noise signal is introduced at the first clock only.

5.1 Noise Free Linear Clock Model

Initially it is assumed that the timing reference at the source node is ideal. That is the initial phase error is zero and that phase variations are initially centered with respect to any buffer elements within the synchronization elements. The SEs are configured to reflect the cascaded timing chain as shown in Figure 1a & 1b. Figure 6 illustrates the TVAR performance for a linear chain of SEs.

The key feature that this graph makes evident is the significant growth of noise at the natural frequency of the control loop even for this over–damped situation. This noise growth under ideal input conditions is typical of telecommunication network SE designs. Notice that the short term stability remains well with in reasonable limits (i.e., minimal likelihood of pointer generation for SDH or SONET systems, and cell delay variation or loss for ATM systems). The growth of long term instability, however, will impair the holdover calibration values and make trouble isolation more difficult. This is because the actual source of the holdover value contamination may be far removed from the site at which the symptoms first begin to appear (in a network).

5.2 Noise Free Non–Linear Clock Model

The identical simulation to that in Section 5.1 is now repeated, that is, with respect to the hypothetical reference connection, but in this instance the phase (quantization) detector is deliberately modified to sample the signal in a non–linear manner. The point of emphasis here is that each SE in the chain will exhibit the same "large signal" time constant damping behavior as in the previous case, but the loop will also exhibit a wider bandwidth for low level input noise. The significance of these characteristics are borne out when the chain is subjected to actual network noise stresses. This behavior is illustrated in Figure 7a.

5.3 Discussion of Results

It is instructive to make a side–by–side comparison of the results of the linear and non–linear simulations. In Figure 8 the data are plotted in terms of decibels (dB) to emphasize the substantial difference in the growth of the phase noise as the clocks are cascaded. As seen in the figure the worst case instability is reduced by over 40 dB. This comparison dramatically illustrates the impact that a relatively small "implementation" issue such as, quantization phase error mapping, can have on overall network performance.

5.4 Verification of Simulation Results

In an effort to verify the validity of the simulation results a test fixture was designed to examine the behavior of an actual chain of 8 network synchronization elements (clocks). These clocks represented the improved design, employing the non–linear algorithms just discussed. Data were collected over a 72 hour period and compared to the previous results. Figure 9 reveals how well the simulated data compare with the actual network clocks. As shown the measured results compare favorable with the simulation model. Moreover, the figure also indicates the phase noise reduction between the linear model and the measured data, by roughly a factor of 10 for the example shown. Any small variance observed between the model simulation and the measured data is well within expected limits associated with correlated temperature effects.

5.5 Noisy Linear Clock Model

Although for many practical purposes it is common to consider synchronization signals at the point of origination as noise free, but in reality this is never really the case. Furthermore, "good" signals can arbitrarily and capriciously degrade and become unstable as a result of noise bursts, signal anomalies or related transient phenomena. Consequently, it is not difficult to appreciate that timing signal transport and signal processing along or within a chain of elements traversing a telecommunications network, will inevitably degrade the timing signal further. It is the signal emanating from the final transport NE that is the issue of concern. It is this signal that will eventually be delivered as a reference to synchronize some other clock within the subtending network(s). It is the purpose of this section to investigate the extent to which this signal can be degraded and to offer suggestions that could improve the situation.

The input signal characteristics used to evaluate the effect on the chain of cascaded clocks is described in reference 8[8]. This template which effectively serves as a worst case network noise reference was selected on the basis of an evaluation and extensive survey that was performed on existing network clocks by various participants in ANSI[6] T1X1.3. It represents the peak noise effects of timing signals from switching systems, cross–connect equipment as well as transport systems. It is noteworthy to point out that this template is based on peak performance of observed data over the short term $\leq$1000 seconds and is not intended to model actual long term network instability, but merely serve as a baseline for the development of SONET NE clock filter specifications.

The timing reference at the source node is now represented as the noise signal the characteristics of which are indicated in Figure 10. As before the initial phase error is zero and the phase of each clock is initially centered. Again the timing reference connection is the cascaded chain shown in Figure 1a & 1b. Figure 10a illustrates the TDEV performance for such a chain of SEs.

In this instance the input noise dwarfs the injected oscillator noise so much that it is no longer a factor in contributing to the short term stability. This type of noise if it does become a factor will only become apparent at very long integration times.

This figure not only describes the effect that a noisy input signal has upon the output from the

---

[6]American National Standards Institute

final clock in the reference chain, but moreover, reveals the underlying filtering and (oscillator) noise injection process that occurs between each of the clocks in the chain. This becomes evident by observing the substantial improvement in short term stability brought about by the low pass filtering of the first few clocks. This improvement soon vanishes with increasing observation time because of the peaking effect brought on by cascading the multiple filter poles. At this point the oscillator noise characteristics can no longer be filtered and due to the cascaded filter elements contribute to the injection of noise into the characteristics of the output signal. From this point on, the resultant behavior of the chain is dominated by the effects of the (cascaded) internal oscillator(s) because there is essentially no other possible source of noise in this controlled simulation environment.

Essentially, this simulation identifies the fact that these network clocks are trading–off high frequency noise rejection for low frequency noise peaking[7]. It is important to note that this is not meant to imply poor performance, but only that this is the best the performance that could be achieved. The results also correlate well with what one would expect to observe when measuring the end of a timing distribution path in a telecommunication network. It is therefore, concluded that the model adequately reflects the characteristics of a telecommunications network.

The essential feature of this comparison is the response of the non–linear implementation to the accumulated effects of large signal noise levels. It will be demonstrated in what follows that the non–linear implementation is able to maintain noise levels very near those of the original noise-free input because of a sophisticated non–linear gain control mechanism. Further, this mechanism relies on the aspect of the precise manner in which the integral and proportional gain factors are modified with respect to a previously stable condition. Hence, while there is the possibility for predictive correction the possibility for over–correction, or mis–correction is not allowed. In effect providing a "clutch" mechanism in the servo–control process.

5.6 Noisy Non–Linear Clock Model

When the same noise source is applied to the reference connection as diagramed in Figure 1, employing cascaded non–linear clocks described in Sections 3 and 5.2 the simulation results at first glance appear curious. Interestingly enough, the salient feature of Figure 10b is that the network output response is exactly the same as for the linear clock model. The short term stability is not at all affected by the input noise or influenced by the non–linear phase quantizer; moreover, the longer term instability noise values peak in the same manner as did the linear example. Perhaps surprisingly, this is precisely what one should expect to see. The significance of this result is that the response of the non–linear control loop compensates for the noisy input such that for higher noise levels the non–linear design reduces the loop gain effectively reduces the impact of the noise. Whereas, the well designed linear clock must filter the input.

In this example, however, the gain is reduced to be equivalent to that of the linear model. Further, lower noise levels will cause the gain to be increased so that the control loop will be tightly locked to the input signal (as it should) and its performance will match that of the linear model. The focus here is not that we have achieved analogous performance but the manner in which it has been done. Clearly, the two models differ in how they process the input noise. As

---

[7]This effect is also referred to as the Gibbs phenomenon.

noise levels increase either suddenly or slowly as by some network breathing phenomenon the non–linear implementation will manage its proportional and integral gain factors independently such that the performance exceeds the linear implementation. Conceptually, this must be the case because simply brute force filtering the input must approach a limit. Moreover, depending on the type of input noise the performance of the linear implementation is also limited because it is generally optimized for particular noise types, whereas the non–linear implementation manages several types handily.

It is also noteworthy to point out that in each scenario, as the number of cascaded elements increases the affect of filtering is progressive and barely noticeable. Essentially, the noise "transfer (function)" characteristics develop an increasingly steeper roll–off as would be expected from linear system filter theory. Recall that our hypothetical reference connection is a linear system for this noisy signal model. It is merely the individual synchronization elements that process the data in a non–linear fashion.

6. Summary

By examining the characteristics of the clock model described in Figure 3 and contrasting these with an understanding of network behavior that has been put forward one can postulate how the so–called boundary conditions might be manipulated to achieve an (overall) advantage at the expense of perhaps minor network effects.

Figures 10a and 10b demonstrate the results of the empirical trials undertaken to evaluate the impact of these performance trade–offs. If these figures included more detail of the higher frequency regions (times less than 256 seconds) it would be seen that the short term noise is slightly elevated in comparison to the well designed linear clock model. This, however, does not portray the correct picture because the linear approach does not adequately manage transient phenomena and the dynamic behavior of the telecommunications network. Alternatively, the non–linear design is able to make decisions necessary to manage such behavior. Notice also, that for the non–linear approach the stability values for times in excess of 1000 seconds are dramatically lower, so much so that it appears as if the size of the hypothetical network could be extended indefinitely. This of course is not practical, but the rather radical improvements brought about by implementation of the non–linear clock design hold great promise for bounding network stability in the face of ever increasing network complexity and the timing requirements that would obtain from them. In particular, the consequences of introducing SDH and SONET ring architectures pose formidable challenges to network synchronization engineering.

Finally, to reiterate the key points regarding the implementation of such non–linear clock designs as described herein: since the model has the potential to shift into a wide bandwidth mode by virtue of the input noise level and become unstable, some form of fail safe mechanism must be implemented to control this possibility. It was seen that this action can be prevented by virtue of building–in a "clutch" mechanism that prohibits the measurement decision circuitry from dwelling on a particular level too long so that changes in the control–loop gain may be gracefully managed. The control of the proportional and integral gain factors are managed to ensure stable modes of operation as the (non–linear) clock switches from wide band–width to narrow band–width based on a real time measurement of the input noise.

References

[1] ITU, Architectures of transport networks based on the synchronous digital hierarchy (SDH), G.803, Geneva, January 1993.

[2] ANSI, *"SONET Hypothetical Reference Circuit"*, T1X1.2/93–015, Liaison Statement from T1X1.2 to T1X1.3 on Network Architecture, 1993.

[3] ITU, WP6/Q21, *"Timing Characteristics of Slave Clocks Suitable for Operation of SDH Equipment"*, G.81s, SG XIII/WP6, Q21, Geneva, January 1993.

[4] ITU, WP6/Q21, *"Timing Characteristics of Slave Clocks Suitable for Operation of SDH Equipment"*, G.812, SG XIII/WP6, Q21, Geneva, January 1993.

[5] ITU, WP6/Q21, *"Timing Characteristics of Primary Reference Clocks"*, G.811, SG XIII/WP6, Q21, Geneva, January 1993.

[6] G. Zampetti, *"Synopsis of Timing Measurement Techniques used in Telecommunications"*, Telecom Solutions, Noordwijk, The Netherlands, 17–19 March 1992.

[7] G. Zampetti, *"Cascaded Clocks Measurement and Simulation Findings"*, European Frequency and Time Forum, Neuchatel, Switzerland, 16–18 March 1993.

[8] Bellcore, *"Clocks for the Synchronized Network: Common Generic Criteria,"* Technical Reference, TR–NWT–001244, Issue 1, June 1993.
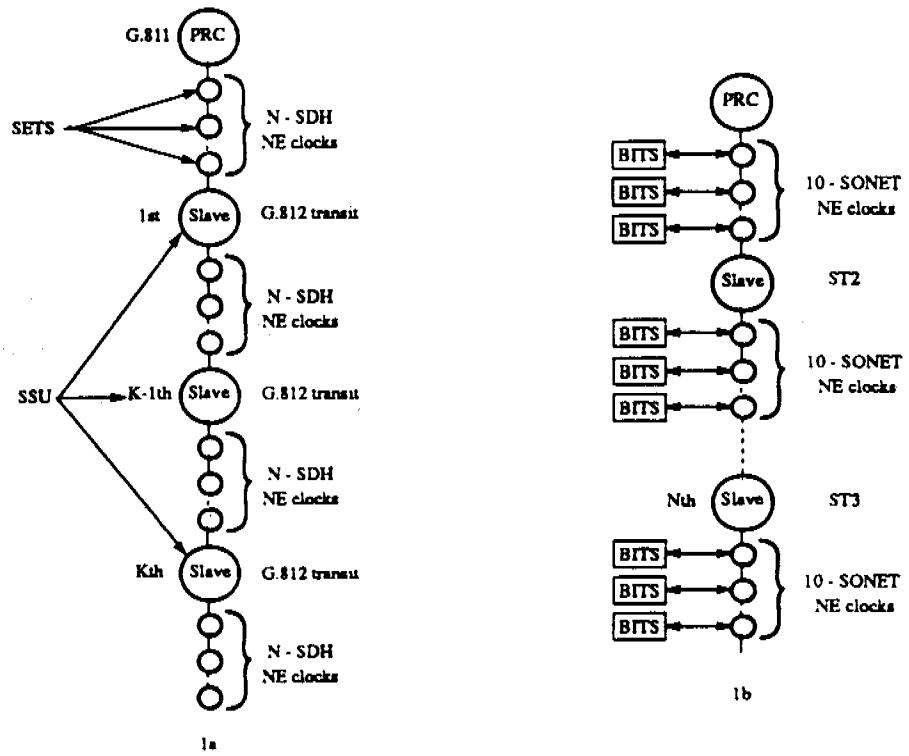
Figure 1. This is Figure 6.4/G.803 from IUT recommendation G.803, 1993
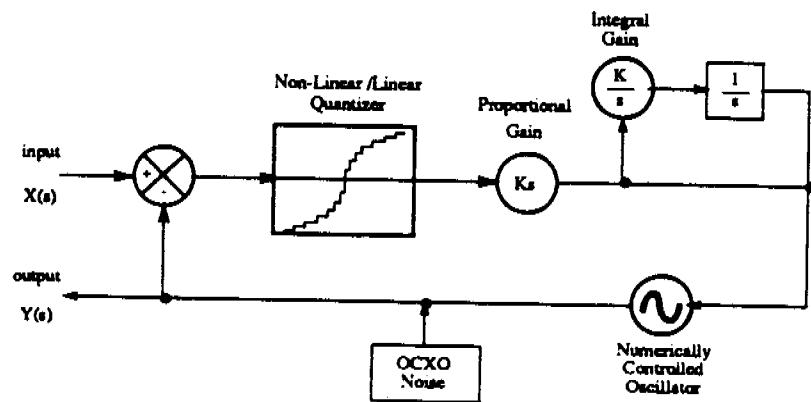
Figure 2. Generic clock architecture

161

Input Signal

MTIE
no
Filtering

Sampled Clock Signal

Amplitude

Time

Max
Peak-Peak
in window

Figure 3

Input Signal

TVAR
Filtering
N=16

Sampled Clock Signal

Amplitude

Time

Calculate
Classical
Variance

Figure 4

162

Figure 5

$$S_x(f) \quad \text{Input Phase}$$

$$|S_{avg}(f)|^2 = \left(\frac{1}{n^2}\right)\frac{1-\cos(2\pi f)}{1-\cos(2\pi(f/n))}$$

Block Averaging

$$|S_{avg\,diff}(f)|^2 = \frac{1-\cos(2\pi f)}{2}$$

$$|S_{avg\,diff}(f)|^2 = \frac{1-\cos(2\pi f)}{2}$$

First Difference
$$S_{ZTIE}(f)$$

Second Difference
$$S_{TVAR}(f)$$
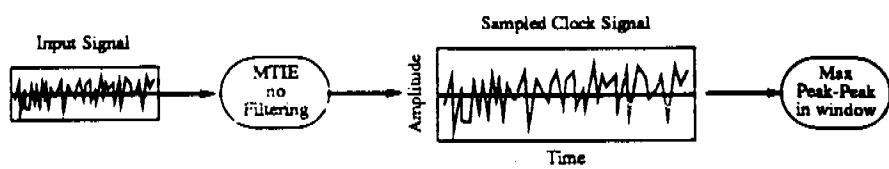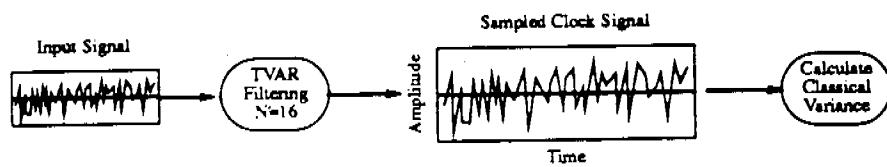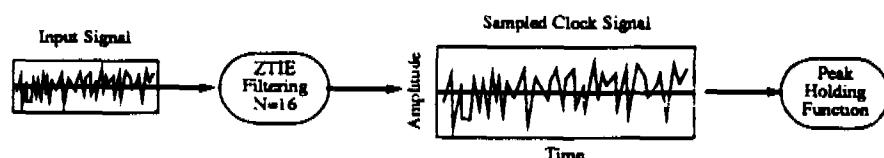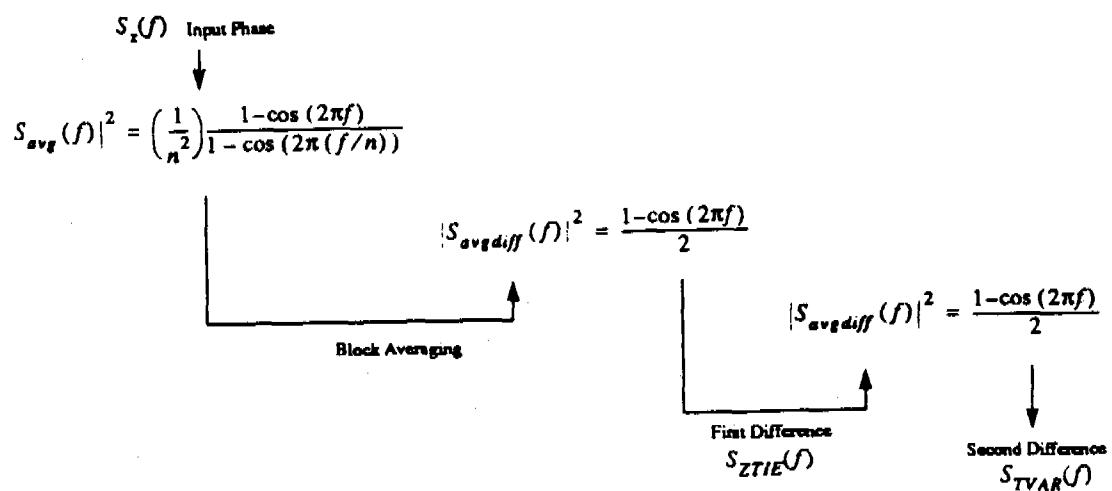
Figure 5a. Illustration of TVAR and ZTIE processing shown in the frequency domain as filter transfer functions.

163
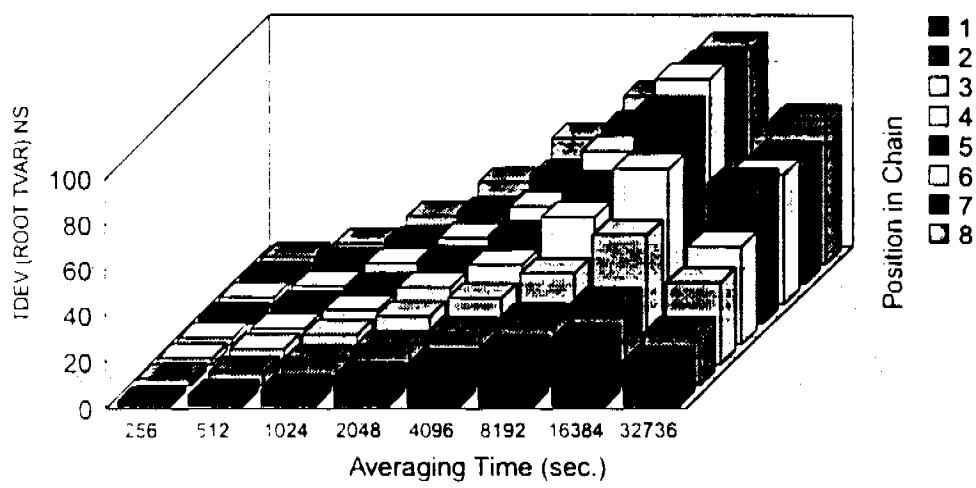
Figure 6. Linear Cascaded Clocks
Eight overdamped clocks
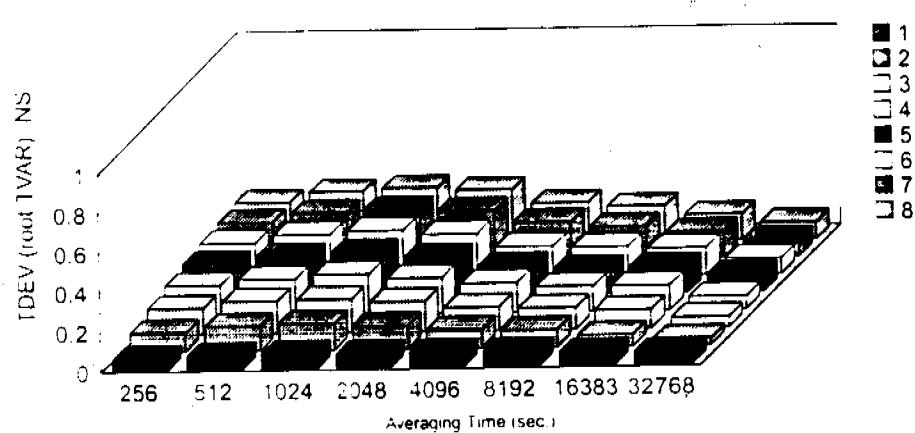
Figure 7a. Non-linear quantized cascaded clocks
Eight well designed clocks
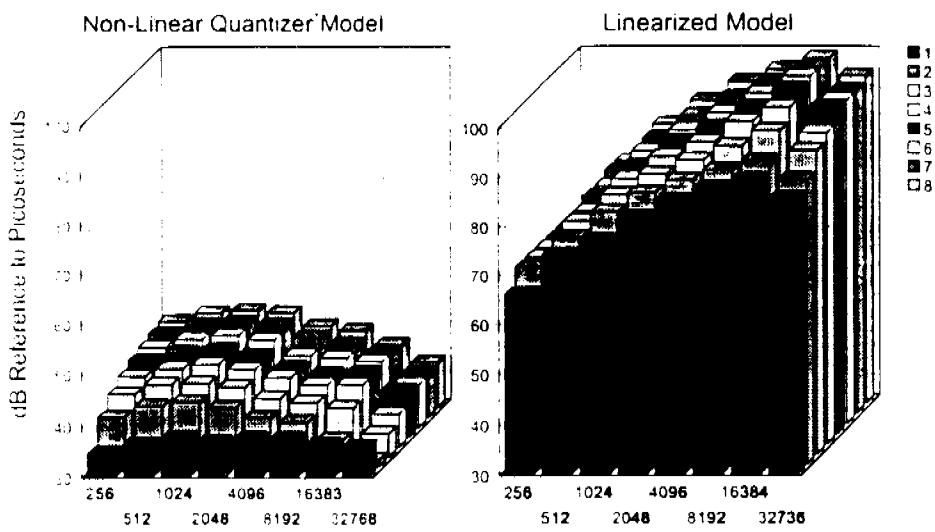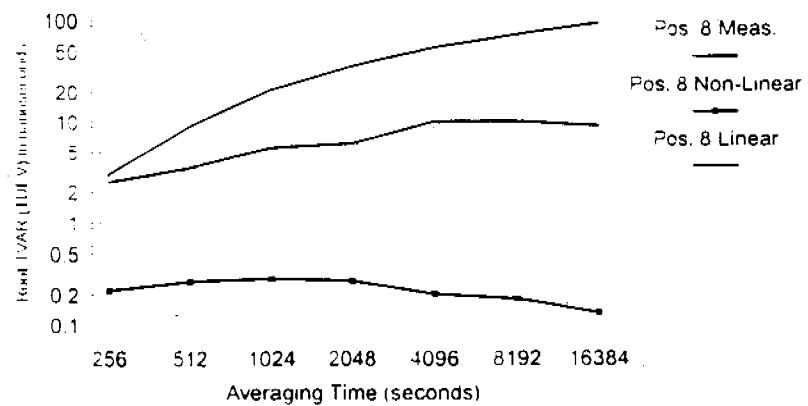
Figure 8. Comparison of Performance (Ideal Input)

Measurement data collected over a weekend.
Lab Temperature Changes (5 deg C) produce
some correlated noise accumulation in meas. data

Figure 9. Simulated vs. measured performance
Last clock in chain (ideal input to first clock)
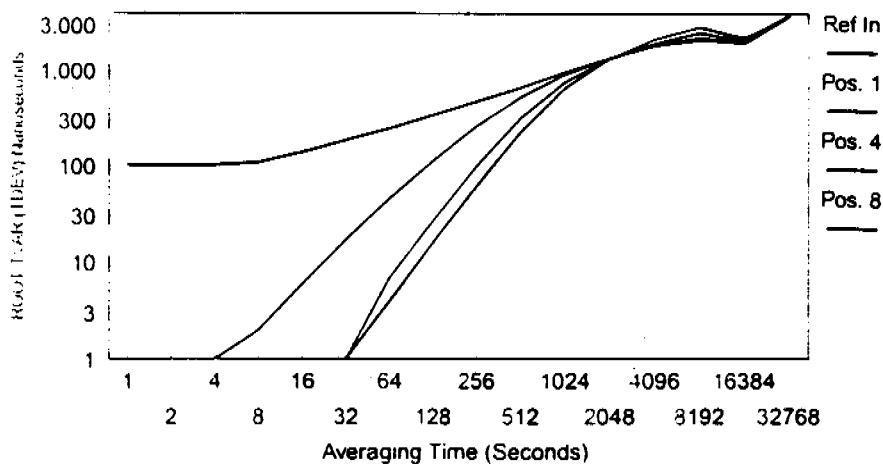
Figure 10a. Ref input noise ANSI T1.101 network noise mask linear simulation results
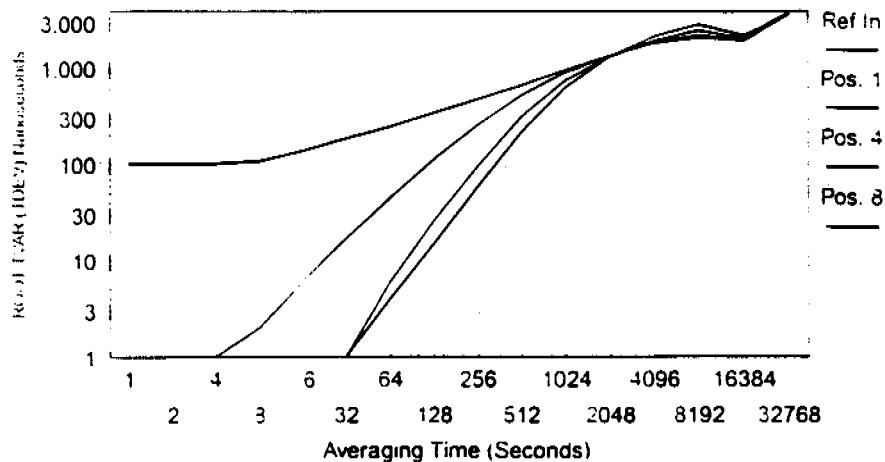

Figure 10b. Ref input noise ANSI T1.101network noise mask non-linear quant. simulation results

QUESTIONS AND ANSWERS

Marc Weiss, NIST: What would happen if we used the optical fibers? Optical fibers are extremely good for time transfer. What would happen if we bounced pulses back and forth between each network element to maintain synchronization throughout the network? How would that help some of these problems?

George Zampetti: Yes, the issue there is not technical. When I was at AT&T, we did experiments on commercial fiber; you basically have tens of ps per degree kelvin temperature variation, and that is about it. And for buried fiber, you get very good stability relative to what the network needs. The problem is that you don't have accessibility to the fiber; you have multiple vendors who are terminating that fiber. And they get paid for moving data and information, not for moving synchronization. And Japan has been pushed to try to utilize the fiber more directly. And I think that is going to have to iron itself out in standards; because, there really is sort of two camps: one camp says let's put more and more smarts and management right in the embedded network elements; others say that it can't happen, it's not manageable, let's take it out of that and make a separate overlay synch network. And so the use of the fiber is something where the issue is not technical; it is how do you get access to the pure fibers, or as close enough to it that you can achieve your goal.

David Allan, Allan's Time: I have two comments and one question. First of all, thank you for a very outstanding and very timely paper. I found it extremely interesting. One point of clarification, TVAR isn't a simple scaling from the modified two–sample variance because the multiplier tau is a variable. And when you do the Fourier transform, you see quite a different picture in Fourier space of these two functions. So it is a slightly different animal.

The question is that we are now seeing quite inexpensive GPS receivers which have claims of, even with SA one pps outputs which have RMSs in the vicinity of 40 or 50 ns, if is such inexpensive receivers could be located throughout networks, would this be useful? What is your perspective?

George Zampetti: My perspective is that that is happening for different administrations. They have already crossed some mental boundary which says GPS, with all its woes, is more stable (this may sound strange to this audience), more predictable and more manageable as a source of calibration traceability than trying to deal with its constantly changing network issues. So you are seeing GPS going in. Now what else you are seeing is acknowledgment that when we get done with this transition and we put down this GPS system, we take a step back and we ask the question, "Is this office better off for doing that?" And that is the big issue. Because, GPS is something that also has to be managed at the gate before it enters your whole central office. Like a virus gets in. And if something happens, if somebody opens up a hatch in

the roof, then all of a sudden that moves over. What type of GPS do you really have? So there is an issue here that says yes, that basically it looks good. Let's look at it operationally and let's start looking at different vendors and look at what they are doing in terms of fault tolerance and sound length; how they manage the time scale; and oh by the way, it still has to be absolutely inexpensive. That is almost a given.