Symmetricom®

Stochastic Model Estimation of Network Time Variance

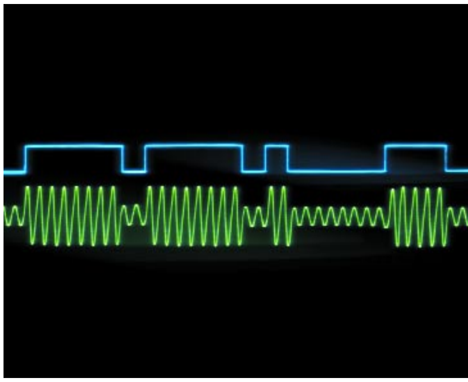

WHITE PAPER

Stochastic Model Estimation of Network Time Variance

Introduction

This white paper presents a study of clock drift in computers and other networked devices.

- The first section, Clocks, Time, and Drift, provides a non-technical introduction to clock architecture, distribution of standard time, drift in oscillators, and efforts to improve clock precision.

- The second section, An Equation for Clock Drift, shows how the calculation for modeling drift in networks was established.

- The third section, Monte Carlo Calculations, presents the statistical projections and results of the study.

- The study concludes that clock error can be significant, on the order of several minutes per month, even in relatively small networks.

Clocks, Time, and Drift

A common understanding of the following three points is required in order for measured time to be a useful and commonly understood standard:
- Frequency
- Units of measure
- Zero-point

Clock Architecture, Oscillators, and Counters

Clocks provide the first two aspects of time - frequency and units of measure. The basic architecture of every clock consists of an oscillator and a counter. The oscillator generates a consistent frequency, while the counter accumulates the frequency's 'beats' and translates them into a common time unit. In a pendulum clock, for example, the pendulum is the oscillator. The gears inside the clock serve as the counter. Every time the pendulum completes a defined number of cycles, the gears move the minute hand forward one unit. The essential property of all oscillators is that they behave in a stable and predictable way. This stability determines how well two or more clocks remain synchronized with each other over time. For this reason, clocks have used a variety of phenomena and materials that display stable and predictable cyclical behavior. Examples include the Earth's rotation and orbit, dripping water, pendulums, piezo-electric quartz crystals, cesium, rubidium, and hydrogen masers.

While oscillators vary in their stability and predictability, counters are practically perfect in that they are consistently able to transform oscillations into units of time without error. The challenge with clock counters is setting them to the right time, or zero-point, so they agree with other clocks.

Standard Time

Standard time provides the common zero-point by which to set clocks. Several time standards exist. The most common, Coordinated Universal Time (UTC) is kept by national time laboratories using atomic clocks. The time kept by these labs agrees to within several nanoseconds.

The accuracy of a clock can be measured by comparing it with UTC. To the extent that it agrees with UTC, we know that it agrees with other clocks. UTC is distributed by the Internet, telephone, radio, and GPS satellite.

Piezo-Electric Oscillators and Drift

Most computers and networking equipment today use inexpensive piezo-electric quartz oscillators designed to vibrate at approximately 32,768 times per second (Hz). This frequency is determined by the size, cut, and orientation of the quartz crystal. As with most commodity parts, there is a certain degree of tolerance for variations in performance. The resulting oscillators tend to run faster or slower than 32,768 Hz. Furthermore, oscillators change with age and are affected by environmental variables, such as mechanical vibration, magnetic fields, and especially temperature. Consequently, clocks that use quartz oscillators can drift up to several seconds per day. Over time, this drift adds up and becomes significant.

Interestingly, wristwatches drift less because the stable body temperature of the wearer regulates the oscillator better than the changing temperature inside a PC.

Improving Clock Precision and Accuracy

Improvements to clock precision and accuracy have centered on reducing environmental effects and on correcting drift when it occurs.

Oven-controlled quartz oscillators (OCXOs) maintain the oscillator in a chamber at a fixed temperature, well above the ambient temperature range, where it can operate at a specified frequency. By shielding the oscillator from environmental temperature changes, OCXOs maintain their intended frequency much better than conventional quartz oscillators.

In a process known as disciplining, high precision clocks calculate their drift relative to standard time and adjust the frequency-regulating mechanism to cancel-out future drift. This is similar to using the adjustment dial on a mechanical clock.

Since the early 1950s, a new class of clocks based on the atomic resonance of specific elements has emerged. Fundamentally, atomic clocks work by bombarding elemental ions with energy sources such as light or microwaves, and detecting the frequency at which those ions absorb the energy source. An early clock using molecules of ammonium based on this approach was developed in 1949. Subsequently, an atomic clock using cesium was developed that kept time to within one second over twenty million years. This is so accurate that, in 1967, the resonance frequency of cesium, which is 9,192,631,770 Hz, was adopted as the definition of one second, replacing the previous standard based on astronomical measurements.

Two other types of atomic clocks deserve mention: rubidium, and hydrogen-maser. Rubidium-based clocks are less accurate than cesium clocks, but can be used in applications such as telecommunications and GPS satellites because of their relatively compact size, low cost, and low power consumption. Hydrogen maser-based clocks currently have the highest short and medium-term stability of any of the established clock types.

Efforts to improve clock precision and accuracy continues today.

An Equation for Clock Drift

The goal of this study was to develop a simple expression for clock drift. Starting with the fundamental definition of frequency:

$$f := \frac{d\phi}{dt}$$

where

- f... frequency
- ϕ... phase
- t... time

We rearrange and integrate each side of the equation over time:

$$\phi := \int f(t)\,dt$$

For an oscillator.

$$f := f_{nom} + \Delta f_0 + a\cdot(t - t_0) + \Delta f_n(t) + \Delta f_e(t)$$

$$\phi := \phi_{nom} + \Delta\phi$$

where

- t_0... starting time
- ϕ_{nom}... nominal, time varying phase
- $\Delta\phi$... phase, or time error
- f_{nom}... nominal frequency
- Δf_0... initial frequency error
- a... aging rate
- Δf_n... short-term frequency instability (noise) term
- Δf_e... environmental term

High-order frequency terms are neglected. This provides a good approximation for all common oscillators.

$$\phi_{nom}(t) + \Delta\phi(t) - \left[\phi_{nom}(t_0) + \Delta\phi(t_0)\right] := \int_{t_0}^{t}$$

$$\left[f_{nom} + \Delta f_0 + a\cdot(\tau - t_0) + \Delta f_n\cdot(\tau) + \Delta f_e\cdot(\tau)\right]d\tau$$

Since $\phi_{nom}(t) - \phi_{nom}(t_0) = f_{nom}$ _ $(t-t_0)$

Short-term frequency variation is neglected in these calculations. Short-term frequency variation has zero mean, and does not lead to accumulated time errors. In other words, the positive and negative parts cancel each other out, on average. The amplitude of these short-term variations is small enough that they do not cause the clock to accelerate or decelerate erratically.

The environmental term, primarily due to temperature, can be quite significant. For our simulations, we assume that the oscillators experience diurnal temperature variation such as the 24-hour heating-cooling cycle in most office buildings. Since our calculations are for clock drift over many days, the temperature effects were modeled with an effective average constant frequency offset. Inclusion of an oscillatory temperature would not significantly affect the results of this study. Changes to weekend and holiday temperature cycles were excluded from this study as well.

Neglecting the noise and environmental terms:

$$\Delta\phi(t) := \Delta\phi(t_0) + \int_{t_0}^{t} \Delta f_0 + a\cdot(\tau - t_0)\,d\tau$$

The resulting equation for clock drift provides a nice simple result for a single clock:

$$\Delta\phi(t) := \Delta\phi(t_0) + \Delta f_0\cdot(t - t_0) + \frac{a}{2}\cdot(t - t_0)^2$$

However, to determine the effects of drift for multiple networked clocks, we need to perform Monte Carlo probability simulations.

Monte Carlo Calculations

Monte Carlo calculations are used to statistically model probability using large sample populations to generate data. In this study, random numbers were used to evaluate the time error evolution in the clock population. We choose Df0 from a gaussian (normal) distribution, typical for cheap oscillators such as found in PCs:

Mean = 0, _ = $2x10^{-5} * f_{nom}$

The aging rate, a, is gaussian, distributed with mean = 0, _ = $2x10^{-6} * f_{nom}$

Note the following:

- To better show the rate at which clocks drift apart, the clocks in this study are synchronized to within one second of each other at **t =0**. In a real-world sample, the initial time error (the amount of time the clocks are off by) would most likely be larger.

- Note that time-error for clocks grows quadratically, accelerating over time. For cheap PC clocks, on a scale of weeks the linear term is stronger than the quadratic term, so individual clocks wander off at an almost constant rate.

- Individual clock properties are uncorrelated. This may not be exactly true for clocks in the same building, with similar temperature cycles, but correlation effects are expected to be small.

- Clock parameters are normally distributed. The variances of the constants in the clock drift equation (initial time error, initial frequency error, aging rate) are all inputs to the model.

- The typical sample size used in these simulations is approximately 100,000 clocks (e.g., 1000 simulations of 100 clocks). Some of the texture of the data is due to the finite size of the sample. Smother results could be achieved using larger samples.

- The spread in clock errors for a given size network is the average spread determined by simulating a number of networks of the same size.

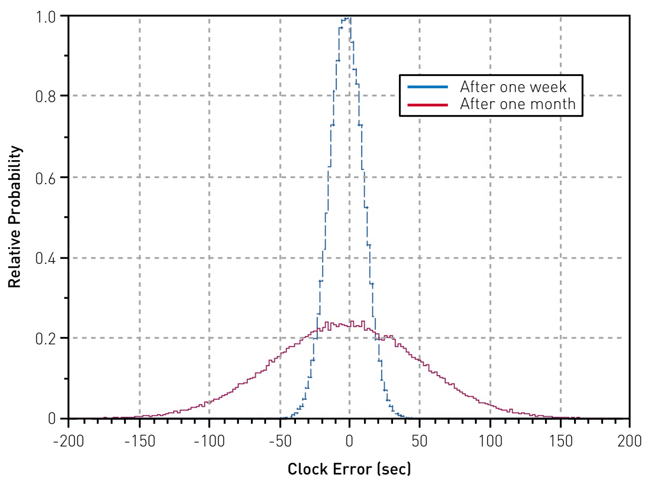

FIG.1 Relative Probability of Clock Error for one week and one month. The number of clocks in the network is 100.

- The spread in clock time grows almost linearly as a function of time, due to the dominance of the linear term in the clock drift equation.

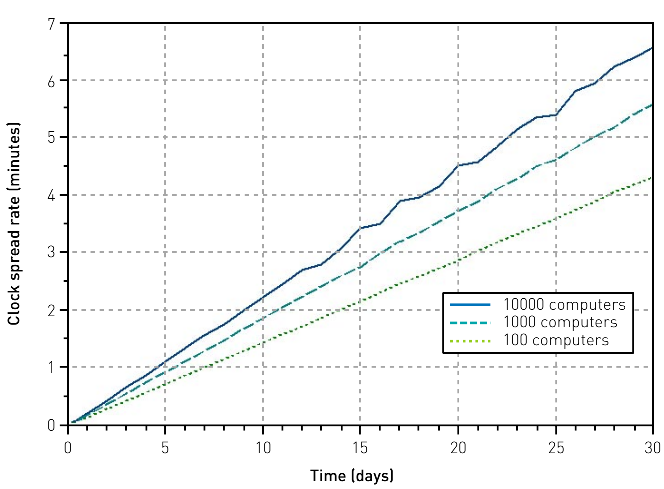

FIG.2 Clock Spread over 30 days.

- Clock spread grows logarithmically with respect to the size of the network. The problem of clock drift is significant, even in smaller networks, gradually leveling off as the network size increases.

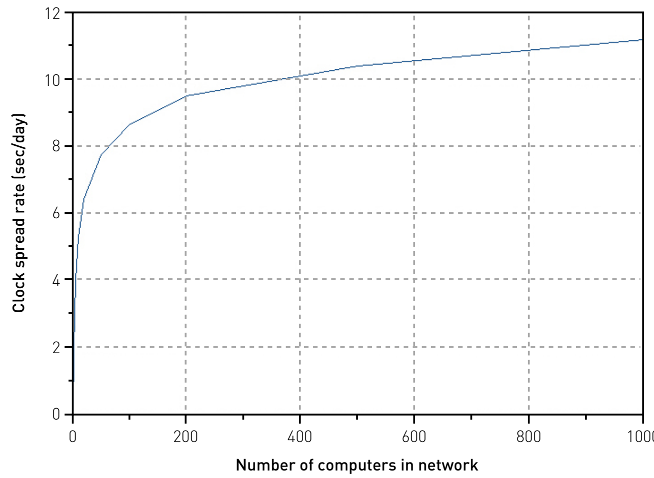

FIG.3 Rate of Clock Spread for the Number of Computers in a network.

To conclude, the above simulations clearly demonstrate that, whether the network contains 100 or 1000 clocks, time spread will exceed one minute in less than one week.

References

From Sundials to Atomic Clocks, J. Jespersen, J. Fitz-Randolph, 1999

Theory of Oscillators, A. A. Andronov, A. A. Vitt, and S. E. Khaikin, Dover, 1966.

NIST Technical Note 1337: Characterization of Clocks and Oscillators, Edited by D. B. Sullivan, D. W. Allan, D. A. Howe, F. L. Walls, 1990.

http://www.ofc.com/features.htm, Technical Notes and Feature Articles, Corning Frequency Control Inc.

http://aa.usno.navy.mil/faq/docs/UT.html, What is Universal Time?, United States Naval Observatory

http://www.wenzel.com/oscillators.html, Crystal Oscillator Information, Wenzel Associates Inc.

Resources

The calculations were performed on a PC using O-Matrix mathematical calculation software (http://www.omatrix.com/)

SYMMETRICOM, INC.
TT&M DIVISION
3750 Westwind Boulevard
Santa Rosa, California
95403
tel : 707.528.1230
fax : 707.527.6640
ntp-systems@symmetricom.com
www.ntp-systems.com

SYMMETRICOM, INC.
CORPORATE
2300 Orchard Parkway
San Jose, California
95131-1017
tel : 408.433.0910
fax : 408.428.7896
info@symmetricom.com
www.symmetricom.com