

Time Synchronization over the Internet Using an Adaptive Frequency-Locked Loop

Judah Levine

Abstract—This paper describes the operation of an algorithm for synchronizing the time of computers using messages transmitted over packet-switched networks such as the Internet. The algorithm configures itself to realize any specified performance level at minimum cost (measured in computer cycles or network bandwidth). If the highest possible accuracy is requested, the performance will be limited by the larger of the instability of the local clock oscillator or the noise in the measurement process between the client and the server; I have obtained uncertainties of about 8 ms RMS using standard workstations and average network connections. Lower accuracy can be realized at substantially lower cost because the cost varies approximately as the inverse of the accuracy squared over a wide range of these parameters. The algorithm makes better use of scarce network bandwidth than previous methods. This improvement is realized by using a pure frequency-locked loop (rather than mixed frequency/phase locking algorithms currently proposed for the NTP) with unequal spacing between calibration cycles. The result is a cleaner separation between network noise and clock noise, which is especially important when the highest possible accuracy is desired. In addition, the algorithm is an improvement over the pure-FLL “Interlock” algorithm that I described previously because it is self configuring. In addition to supporting an explicit trade-off between cost and accuracy, the algorithm provides better performance than previous methods because it is better able to adapt itself to fluctuations in the asymmetry of the network delay. This robustness can be realized without the preliminary tuning that was necessary to realize optimum performance using previous methods.

I. INTRODUCTION

THIS PAPER discusses a method for synchronizing the clock of a client computer using messages transmitted over the Internet from a remote server. The design principles would also be appropriate for other types of connections between the client and the server, provided only that the delay through the network connecting them is symmetrical on average.

The work was motivated by two observations: the load on the network time servers operated by the National Institute of Standards and Technology (NIST) has been increasing at an average rate of 7% per month for several years with no sign that this rate of increase is slowing, and there are many indications that a significant fraction

(perhaps even a large majority) of the users of NIST digital time and frequency services do not need the millisecond-level accuracy that these services can support. Therefore, I set out to design an algorithm that would support an explicit trade-off between the synchronization accuracy that could be achieved using the procedure and the cost of realizing it (the cost is proportional to the average number of calibration requests generated by each client system). Such an explicit trade-off is not present in any of the other algorithms of which I know. As I will show subsequently, the relationship between accuracy and cost is approximately quadratic so that a modest relaxation in the accuracy requirement (to 100 ms RMS, for example) would be accompanied by a decrease of more than two orders of magnitude in the average cost of realizing it. The result would be to continue to provide adequate support to our clients while significantly extending the useful life of the existing time servers. At the same time, the underlying capability of the algorithm would not be compromised, so that those users who needed millisecond-level accuracy could achieve it using the same software.

A number of other algorithms for synchronizing clocks in this environment have been published [1]–[6]. Each of these uses a different metric for evaluating the performance of the algorithm and for deciding what is meant by optimum operation. Some algorithms, for example, are designed to guarantee an upper bound on the absolute value of the time error of the client [4], [5]. To realize this objective, they must make relatively frequent requests to servers that are close by (in a network delay sense) because the error caused by an asymmetry in the network delay is bounded by one-half of the round-trip value. Supporting this density of servers may be quite expensive in practice, especially as the Internet becomes more crowded so that the average delay between any two points increases. Other algorithms place very strong emphasis on dealing with unreliable servers or with frequent large fluctuations in the network delay or in its asymmetry [1], [3]. They realize these objectives by querying a number of servers on each calibration cycle and then by choosing the winner based on various criteria.

We argue that all algorithms with these kinds of goals are relatively expensive from the point of view of the number of servers that they require to realize their design goals. The reason is that the average operating cost (measured in terms of the number of calibration requests that are generated by a client) is driven by the desire to limit the maximum time error or to detect server failures, which should be relatively rare events. Although this machinery

Manuscript received May 22, 1998; accepted November 21, 1998. This work was supported in part by the National Science Foundation through grants NCR-9115055 and NCR-9416663 to the University of Colorado. We gratefully acknowledge this support.

The author is with JILA and Time and Frequency Division, National Institute of Standards and Technology and University of Colorado, Boulder, CO 80303.

undoubtedly improves the worst-case performance of a system synchronized in these ways, it may have little impact on the RMS time accuracy of a typical client.

Our previous Internet-based algorithm [3] and the method that we describe here use the RMS time offset of the clock in the client as the measure of performance, and they are designed to maximize the ratio of this RMS performance to the average cost (measured in terms of the load on the servers and the network). This optimization can be implemented either by finding the best performance that can be achieved at a fixed cost or by minimizing the cost for a given level of performance. As we will show subsequently, the relationship between cost and performance is not linear—incremental improvements in performance become increasingly expensive to achieve. Some levels of performance may not be achievable at any cost in the given environment because the frequency of the local clock is too unstable to support the amount of averaging that is required to cope with a noisy communications channel.

A. Synchronizing the Local Clock

All synchronization algorithms start from the same basic data—the measured time difference between the local machine and the distant server and the network portion of the round-trip delay between the two systems. (Delays in the distant server are usually not a problem. Either they are small enough to be ignored or they are measured by the server and removed by the client.) These data are processed to develop a correction to the reading of the local clock. The usual approach is to use the measured time difference after it has been corrected by subtracting one-half of the round-trip delay. This model is based on the assumption that the transmission delay through the network is symmetrical so that the one-way delay is one-half of the measured round-trip value. This corrected value may be used to discipline the local clock directly or it may be combined with similar data from other servers to detect outliers or to compute a weighted average time difference, which is then used to steer the local clock.

These steering corrections generally take two forms: time steps, which adjust the local clock by a fixed amount essentially instantaneously, and frequency steps, which adjust the effective frequency of the local clock oscillator and thereby slew the time relatively slowly. Frequency steps are usually realized by adjusting the size of the software tick—the amount added to the clock register in response to periodic interrupts from the physical clock oscillator. The frequency of the hardware oscillator itself is usually not accessible to software control.

Time steps, especially those that move the clock backwards in time, can cause unwanted side effects in other programs, and their use is usually limited to setting the clock during an initial start up of the system. Frequency steps are the preferred means of adjustment, although they can have problems too. The minimum frequency adjustment that can be applied simply is to change the size of the tick

by one least count. The resulting change in frequency is one least count per hardware interrupt, which is a fractional change of about 10^{-3} for most systems. This value is much too large. The actual frequency offset is likely to be one-tenth of this value or less, and the precision with which this offset can be determined is even greater.

In practice, frequency steering generally requires a resolution of at least 5×10^{-8} if the granularity of steering corrections is not to degrade the frequency stability of the clock oscillator [2, Fig. 1]. This level of frequency adjustment is usually implemented in two ways: either the kernel is modified to allow static frequency adjustments that are smaller than one least count per hardware interrupt to be specified or the larger, standard frequency offset is used periodically with the ratio of the on interval to the off interval determined so that the average frequency offset corresponds to the estimate derived from the time difference measurements. The first method has the advantage that the resulting time adjustment is smoother, but it requires access to the kernel source code. The second method is more general and can be applied to almost any system, but it results in a sawtooth-like variation in the time of the local clock with respect to the distant server. The amplitude of this sawtooth is a function of the system design, but it can be as large as a full tick (i.e., on the order of milliseconds) in some implementations. The first method (i.e., modifying the kernel) is the best choice if it can be implemented in the client configuration, and it is the method that we use to keep our servers on time.

The maximum frequency offset that is required is usually on the order of seconds per day, which is a fractional frequency offset of about 5×10^{-5} . The dynamic range required in the frequency steering loop is thus about 1000:1. This requirement is usually not a problem, and almost any method that can provide the required resolution can also satisfy the dynamic range requirement in steady-state operation. The maximum fractional frequency offset that is supported by the standard kernel of most systems is about 4×10^{-3} , so that it is not practical to remove time offsets of more than a second or so using this method because it would take too long. The clock is usually set initially using a single time step for this reason.

B. Effects of Measurement Noise

Individual time difference measurements are likely to have a substantial uncertainty because of the noise in the measurement process itself. This noise arises from many sources—from fluctuations in the response time of the local operating system to interrupts, from jitter in the delay through the network or the interface hardware, and from other hardware-related causes. Whatever the cause, this noise is not associated with the frequency of the clock oscillator. Using the time differences themselves to set the local clock directly is therefore not a great idea—no matter how the adjustment is performed. Doing this would convert the phase noise of the measurement process into

frequency noise in the clock. The only reason for doing this might be simplicity. A procedure that sets the clock to the time received from the server would be simple to implement and would not have to run continuously as a background or daemon process. The price for this simplicity is that the performance of the local clock oscillator is degraded.

The contributions to the noise in the measurement process vary very rapidly with time, and consecutive measurements for which the interval between them is long compared with the characteristic period of the variations will be affected by noise that is almost completely uncorrelated both with the value on the previous measurement and with the underlying frequency of the clock oscillator. The spectrum of these variations is therefore approximately white, and averaging closely spaced time interval measurements results in an estimate that converges to the underlying mean value of the time difference, provided only that the measurements are made quickly enough so that the parameters of the oscillator have not changed.

It is usually easy to satisfy both of these time requirements. The characteristic period of the hardware is on the order of microseconds or less, so that even measurements made a few seconds apart are affected by very different noise environments. On the other hand, the time dispersion caused by the fluctuations in the frequency of the oscillator is usually negligible on this time scale.

However, the averaging process cannot be continued indefinitely. It will only do the right thing as long as the spectrum of the fluctuations in the data can be characterized as predominantly white phase noise. The range of averaging times over which this is true is usually pretty small for typical computer clocks, and the nonstationary statistics of the delay in a typical Internet path further restrict the averaging times over which this is an appropriate strategy. The result is that algorithms that develop a correction based on the average measured time difference are likely to be useful only in local area networks and with a relatively short interval between calibration cycles (on the order of 1000 s or less). This argument is particularly relevant to algorithms that use a pure phase-lock loop to discipline the local clock [4], and it explains why algorithms based on frequency loops tend to achieve comparable synchronization at much lower cost [2], [3].

It is not always easy to detect the effects of too much averaging, that is averaging for a time interval that exceeds the domain in which the noise is dominated by white phase noise. The average time difference still exists in this domain, but it is no longer a stationary quantity. The synchronization loop appears to be operating normally, but it is, in fact, increasingly dominated by flicker (and, at sufficiently long averaging times by random-walk) processes [2, Fig. 9].

Algorithms that operate by stabilizing the frequency of the local clock (as opposed to its time) have an easier job, in principle. These algorithms can operate at much longer averaging times when white phase noise is no longer the main problem, and the performance is limited by the fre-

quency stability of the local clock. It turns out that clock oscillators in many computer workstations are surprisingly good. Many of them have Allan deviations of less than 10^{-7} at averaging times of 10^4 s, so that they have a free-running time stability of better than 1 ms RMS [2, Figs. 1 and 3]. One way of realizing this stability is to average the individual time differences for a time that is well within the domain in which white phase noise dominates the noise spectrum and to then switch to averaging the frequency (that is, the first difference of these time differences) as long as the noise process can still be characterized by white frequency noise. For the oscillators typically found in computer hardware, this extends the averaging interval to about 15,000 s. The details of the design will depend on the stability of the local oscillator and on the noise in the network link to the server. However, a design configured for the highest possible accuracy would involve averaging time difference measurements over a period of a few seconds combined with averaging the frequency for a period of a few hours. This strategy is not equivalent to using the same number of measurements that are equally spaced in time, even though the average load on the servers would be the same in both cases [3]. The time between calibration cycles is often called the poll interval, and the argument presented here suggests that the average poll interval does not completely specify the performance capability of the procedure—that grouping the calibration requests makes a better use of the available resources than simply uniformly spreading them out.

Frequency-locked loops are not better or more accurate than those based on phase-locking techniques—they are advantageous because they are almost always cheaper to operate for essentially the same performance. This is true because the high level of measurement noise over the Internet means that the local clock is usually more stable at short times than the distant server seen through the noisy Internet. A frequency-locked loop can better exploit this stability by longer averaging and, therefore, can be configured with longer intervals between the calibration cycles than the corresponding phase lock. Although neither design is a priori more sensitive to glitches, the longer interval between calibration cycles in a frequency loop and its correspondingly longer time constants suggest that a glitch is likely to persist for a longer time in that design before it is detected and removed.

C. Details of the Method

The algorithm design is based on the principle of separation of variance—that is it is possible to separate the contributions of the clock and the measurement process using statistical techniques. A second implicit assumption is that the variances of both the measurement process and the clock frequency can be modeled using stochastic parameters. These turn out to be good approximations to the performance of real hardware over a wide range of operating parameters, although the second assumption tends

to break down at averaging times approaching one day. To achieve this separation, the program evaluates the following two statistics after each calibration cycle has been completed:

S-1. The standard deviation of a group of closely spaced time difference measurements that are made quickly enough so that the parameters of the local clock have not changed significantly while they are being made. As we discussed previously, this requirement is easily satisfied if the ensemble of measurements is completed within a few seconds.

S-2. The error in predicting the currently observed average time difference using the previously measured value and the estimated frequency offset. This prediction is done using a simple linear relationship:

$$\hat{x}_i = x_{i-1} + \hat{y}_{i-1}\tau, \quad (1)$$

where x_i is the time difference measured at time t_i , y_i is the frequency difference between the local clock and the server estimated at the same time, and τ is the time interval from t_{i-1} to t_i . We use (1) to predict the time difference; the error in prediction on this cycle is the difference between the predicted and measured values:

$$\varepsilon_i = |x_i - \hat{x}_i|. \quad (2)$$

Both statistics are maintained in two versions: the value determined in the current calibration cycle and a sliding average over the most recent 12 hours or three calibration cycles, whichever is longer.

The first statistic is sensitive primarily to the phase noise in the measurement process. The program first compares the average value of this parameter with the desired level of performance that was specified by the operator when the software was installed. The number of measurements in each group is adjusted once per day until the two are roughly equal—the size of the group is increased if the RMS of the mean is larger than the requested accuracy and decreased if the mean is much better than necessary. Assuming that the measurement noise is approximately white phase noise, the RMS value of the mean of a group of measurements varies as the square root of the size of the group. The specified performance level, therefore, has a quadratic effect on the cost of the algorithm. In the limit, it may not be possible to achieve the desired level of performance with a reasonably sized group (less than 25 or 50 members). The program defaults to a specified maximum group size in this case. Likewise, the program will never decrease the size of a group below a defined minimum size (usually three) no matter how small the RMS becomes.

The RMS of the group of time differences sets an upper bound to the RMS performance of the synchronization loop at longer averaging times. No matter how stable its frequency is, on the average, the local clock cannot be set (in time) more accurately than this RMS value, and the subsequent algorithm can do no better than remember this adjustment without further degradation. This RMS value may be as small as 75 to 100 μ s when the server and the

client are on the same network but is more typically on the order of milliseconds for a continental length path. This limit is independent of the details of the synchronization procedure.

If the standard deviation of the current group of time differences is much larger than the average, the most likely reason is that it is due to jitter in the symmetry of the network delay during the time that the group of measurements was being made. Although the measurements are made rapidly enough that the parameters of the local and remote clocks have not changed, the same cannot be said of the network (The magnitude of the delay is not a problem because it is measured on each cycle—only its asymmetry causes trouble.) The algorithm will drop a single member of the group as an outlier if doing so will reduce the standard deviation to a more reasonable value (i.e., one that is consistent with the average over the previous observations) and will try to repeat the entire group if dropping a single value will not fix the problem. The assumption that underlies this procedure is that asymmetries in the network delay are transient effects. Small asymmetries will be averaged by multiple measurements, and large ones will be short-lived and detected as outliers.

This method will fail if the path from the client to the server has a static asymmetry, which will bias the results but will not contribute to the variance of the group of measurements. The offset caused by this asymmetry is bounded by one-half of the round-trip delay, so that, at the instant of synchronization, the time difference between the client and the server is guaranteed to be less than the sum of the measured time difference and one-half of the round-trip delay. This may be good enough for some purposes, but, as we will show, this guarantee becomes weaker with time because of stochastic frequency variations in the clock oscillator.

It is sometimes possible to detect a static asymmetry by requesting timing data from a second server, but the two data sets are likely to disagree consistently in this case, and it may not be possible to decide where the problem lies. In particular, if the local clock is synchronized initially to a server via a path that has a static asymmetry, then data from a second server received via a more symmetric path may be rejected as coming from a machine that is thought to be broken.

Under normal, steady-state operating conditions, the magnitude of the second statistic is sensitive primarily to the frequency stability of the local clock, and its average value is proportional to the value of the Allan deviation for an averaging time equal to the time interval between calibration cycles. This magnitude is affected by two parameters—the time interval between calibrations and the time constant of the frequency update loop [2, eq. (3)].

The time constant of the frequency update loop represents the optimum averaging time for the frequency estimates, assuming that the process can be analyzed statistically¹. Averaging for a shorter time interval does not optimally attenuate the white frequency noise of the oscillator. Averaging for a longer time results in an average

that exists in a formal sense but is no longer stationary because the procedure has entered the flicker and random-walk portions of the spectrum. This time constant is primarily a function of the design of the oscillator, and it needs to be determined only once when the algorithm is started for the first time. It is measured as part of the cold-start portion of the algorithm by evaluating the Allan deviation as a function of lag when the local clock is free running. This parameter is usually on the order of 10,000 to 20,000 s, so that this initial evaluation process to determine it usually takes somewhat less than 1 day. (To have confidence in the Allan deviation estimates, it is usual to compute the deviation using a time series that is at least three times longer than the longest lag that is of interest.)

Once the algorithm is running in steady state, the average value of the statistic S-2 and the desired performance level are used to adjust the interval between calibration cycles. If the average prediction error is much larger than the measurement noise (i. e., statistic S-2 \gg statistic S-1) then the prediction error is due to stochastic frequency fluctuations in the local oscillator. If the prediction error is larger than the desired level of performance, then the time between calibration cycles is too long—the frequency stability of the oscillator is not good enough to support free running for that time interval, and more frequent calibrations are necessary. If the prediction error is not larger than the desired level of performance, then there is no need to have more frequent calibrations, even though the performance would be improved by doing so.

If the prediction error is much smaller than the measurement noise, then the calibration interval is unambiguously too small—the frequency stability of the oscillator would support the same level of stability with less frequent calibrations and the extra ones are essentially being wasted. The interval between calibrations can be increased in this case with minimal impact on the performance of the loop. The only reason for not doing so would be a concern about outlier detection, and the user can specify an absolute maximum interval between calibrations if this is a concern.

The relationship between the performance of the loop and the interval between calibrations can be calculated if the noise type is known. This relationship tends to make predictions that are too optimistic when the averaging time exceeds a few hours—the performance tends to de-

grade faster than predicted as the time between calibration cycles is increased in that domain. There are two reasons for this. The first is that when the interval between calibration cycles becomes longer than the time constant of the frequency update loop, then that time constant is effectively forced to a value that is larger than optimum. In other words, the distribution of the prediction errors calculated using (2) become increasingly dominated by flicker and random-walk processes, and these errors are effectively fed back into the synchronization loop. The second reason is that at longer periods the frequency of the oscillator begins to be driven by nonstochastic effects such as fluctuations in the ambient temperature. These fluctuations are in addition to the increase in the stochastic level of the frequency variations. Temperature fluctuations in a typical office environment may pull the frequency of the oscillator by up to 1 ppm with a period ranging from several hours to 1 day. Those frequency variations are larger than predictions based on a smooth extrapolation of the Allan deviation from its value at shorter periods.

D. Parameter Update

If the statistical comparisons indicate that the current time differences are consistent with the past performance, then the current data are used to update the parameters of the procedure. These include the average values of both statistics and the average frequency of the oscillator. This average frequency is used to steer the local clock using one of the methods described previously—either it is transmitted directly to the kernel or it is used to schedule small periodic time adjustments that realize the computed offset frequency on the average. The details of these adjustment procedures are more fully discussed in [2].

E. Error Detection

A comparison of the values for each of the two statistics found on the current cycle with the corresponding average values can be used to detect many kinds of problems. We have already discussed this point previously with regard to the standard deviation of the group of time difference measurements, which can be used to detect jitter in the symmetry of the network delay. Likewise, this standard deviation is a measure of how well closely spaced time differences received from two different servers should be expected to agree.

If the tests based on the first statistic are okay, but the prediction error on this cycle is much larger than the value expected based on the average of this parameter, the software attempts to determine the cause of this apparent change in performance. (This test is significant only if the algorithm is operating in steady state. If, based on the procedures outlined previously, the interval between calibrations has just been increased, then an increase in the prediction error is to be expected, and a comparison with the average value is not appropriate until the average value

¹The primary assumption in any statistical analysis is that the data are stationary—that is that the properties of the data set have no dependence on time. This assumption usually fails primarily because of a sensitivity of the oscillator frequency to fluctuations in the ambient temperature. This sensitivity introduces diurnal and other longer-period coherent frequency variations. These fluctuations could be dealt with by measuring the local temperature and using these data to model the performance of the oscillator, but most installations do not have the hardware needed to make these temperature measurements. These temperature-induced fluctuations are important for systems that are designed to realize time with millisecond-level accuracy, and the usual method of dealing with these effects is to limit the poll interval to about 8 hours maximum. These effects are usually ignored in lower accuracy applications.

comes to equilibrium at this increased time interval.) The software conducts a number of tests in the following order:

- The simplest possibility is that the change is due to a temporary change in the symmetry of the network delay (because of congestion, for example). The program repeats the entire measurement protocol several times to see if the prediction error is improved on subsequent attempts. These repeated measurements are made to the same server. This procedure removes the discrepancy about 60% of the time after a single retry and about 80% of the time after a second one. This statistic differs in an important way from the clock filter procedure used in NTP and described in [4]. This test (and the ones described subsequently) are only activated when the initial tests based on the comparison of data from a single server with the local clock indicate that there may be a problem. The NTP procedure normally polls every server every time. This is a difference of principle as well as of implementation. The point is that the local clock should contribute to the data based on its statistics, and simply comparing the measurements made using a number of external servers does not do this in an optimum fashion.
- If the error persists then (and only then) the program queries an alternate server and uses majority voting to decide if the problem is in the first server or in the local clock. (A network problem that is common to both servers will be considered—incorrectly—as a problem in the local clock, and the second server should be chosen to minimize the probability of this happening.) Note that the time of the local clock contributes to this comparison based on its expected stability as estimated by its Allan deviation. If the two servers agree (to within the phase noise of the measurement process as estimated by S-1), then the local clock has experienced a change in its characteristics. The response to this conclusion is outlined subsequently. If the data from the second server indicate that the first server is broken (or its path delay is significantly asymmetric), then the two interchange roles, and the alternate server becomes the primary one for the next 24 hours. The program tries to switch back to the first server at that point, and does so if its data satisfy the tests described previously.
- If the problem cannot be resolved in the second step (i.e., if the time differences between both servers are not consistent with the expected prediction error and if the data from the two servers are not mutually consistent either) then either the problem cannot be assigned to one cause or our expectations based on the previous average performance are too optimistic. The operating parameters are not modified, and the program waits a short time and tries again. This can happen if the network is very congested or if the path to one or both of the servers has a significant static asymmetry. The program will use a third server in an analogous manner if one is defined. If only two servers are

specified (or if data from a third server fails to resolve the question), the program increases the average variances and goes into holdover mode. The holdover performance preserves the last estimate of the frequency offset of the local clock and continues to steer it in frequency based on that estimate. It makes no changes to its time. The clock accuracy will be limited by the flicker and random-walk frequency fluctuations of its oscillator. The time dispersion caused by these fluctuations depends on the quality of the oscillator; typical values for the Allan deviation of the frequency are on the order of 10^{-6} at 1 day, which would result in a time dispersion of 50 to 100 ms. If the problem is due to a transient degradation of the network, then it will probably disappear on the next cycle; if it is due to a more permanent degradation of its characteristics, then the synchronization procedure will slowly learn this by repeated increases in the average variances that form the basis for all of the tests that we have described.

The threshold for deciding that the current estimate of a statistic is much larger than the average value is set based on the usual compromise between the need to maintain adequate sensitivity to problems and the desire not to be overly sensitive to the statistical fluctuations in the amplitude of the noise. We have used a threshold of three times the standard deviation in all of our experiments. If the underlying statistics of the processes were Gaussian, a threshold of three times the standard deviation should cause the measurements to trigger these error tests about 1% of the time in normal operation. In fact, the data are not Gaussian, and the tests are triggered in about 10% of the calibration cycles. These outliers are often quite large so that the algorithm is not sensitive to the exact value used for the threshold. (The use of additional data about 10% of the time should be compared with other algorithms that use additional data on every poll cycle.)

F. Error Modeling

The most common cause of an error in older systems was a time step caused by a lost interrupt. These were easy to detect because the prediction error was very close to an integral number of systems ticks. This problem is usually not important in newer systems, for which a large prediction error is usually due to a significant secular change in the frequency of the clock oscillator. As we suggested earlier, the usual cause of this frequency drift is a change in the ambient temperature, although a change in line voltage may also contribute to a lesser extent. The relationship between ambient temperature and oscillator frequency is complicated because it depends both on the temperature and on its spatial gradient. Attempting to correct the oscillator frequency based on temperature measurements is too complex for most installations and is not a very profitable strategy in any case. The best strategy is to limit both the time constant for the frequency update loop and the maximum interval between calibration cycles so that these temperature-induced frequency variations are mod-

eled as deterministic effects by the frequency estimator and removed. In other words, the parameters of the loop must be short enough so that these temperature-induced effects are treated as real frequency changes rather than as stochastic effects that should be averaged without appreciably changing the loop parameters.

G. Synchronization Experiment

To test these ideas, we conducted two synchronization experiments. The first was designed to synchronize the local clock as accurately as possible and to find out what accuracy could be achieved and how expensive it would be to realize it. The second experiment was designed to synchronize the same clock so that it was correct only to the nearest second using the same servers and network path as in the first experiment. The goal was to see how much this relaxation in accuracy would save in the operating costs.

In both experiments, we synchronized the clock of a computer that was located in our laboratory in Boulder, Colorado to a server on the West Coast in Redmond, Washington and one on the East Coast in Gaithersburg, Maryland. The time of the local clock was also measured using diagnostics derived from other systems in Boulder, but these diagnostics were for information only and were not used to discipline the clock. The server in Redmond was the primary one; the one on the East Coast was only used as part of the error evaluation procedure as outlined previously.

The details of the paths between our client in Boulder and our two servers were not under our control. Although the path to Gaithersburg, Maryland had a delay that was always very nearly symmetrical, the path from Boulder to Redmond was often 42 ms longer than the return path because the outbound path was routed through Atlanta, Georgia. When this was true, the two servers seemed to have a time offset of 21 ms—an offset that was large compared with the time dispersion created by all other causes for averaging times up to several hours. As we mentioned previously, static asymmetries of this type sometimes cannot be unambiguously removed by any statistical procedure. Our solution was to add a third server located in Reston, Virginia and to use its data in a simple extension of the majority-voting algorithm described previously. The result was to exclude the server in Redmond whenever its delay became asymmetric and to make the server in Gaithersburg the primary one with the server in Reston as its backup when needed. Even when three servers are specified in the configuration file, only one is actually used most of the time—the others are used only to resolve ambiguities when a large prediction error indicates a possible problem or a change in the parameters of the system.

The data from the first high accuracy experiment are shown in the first set of figures. Fig. 1 shows the measured time difference between the local clock and UTC, and Fig. 2 shows the interval between calibration cycles for the same period. Note that the interval between calibrations is gradually increased during the first week of op-

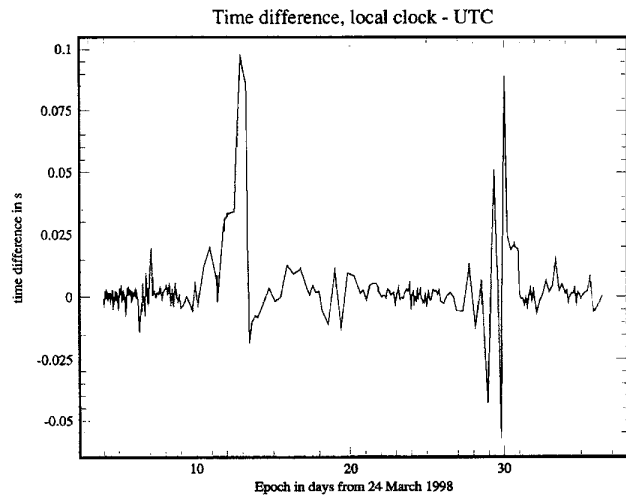


Fig. 1. The time difference in seconds between the local clock and UTC. Figs. 2, 3, and 4 use the same scale as Fig. 5 for the X axis.

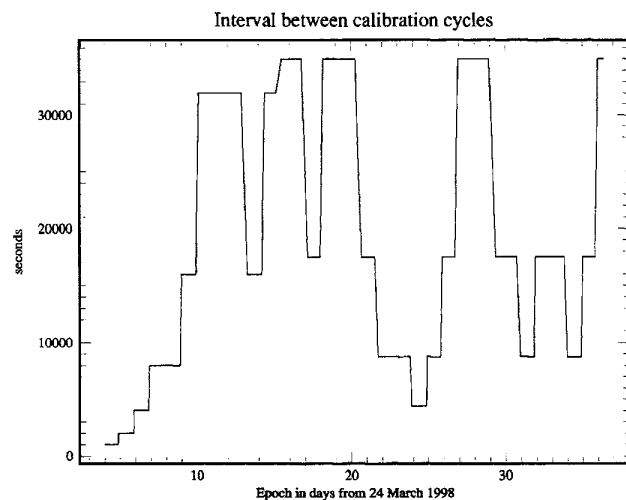


Fig. 2. The interval in seconds between calibration cycles as a function of epoch for the first experiment.

eration without significantly affecting the accuracy of the clock. This situation changes abruptly when the interval reaches 32,000 s—an interval comparable with the length of the working day—because both the local temperature and the network delay have deterministic fluctuations near this period. Although it does not happen immediately, one of these fluctuations eventually exceeds the threshold, and the algorithm is forced to decrease the time between calibration cycles to maintain the desired performance level. The same story is repeated on a smaller scale several more times until about day 28—a day that had both a large change in temperature followed by the failure of a network element when the same kind of problem is repeated. The RMS accuracy of the entire data set is about 30 ms, but most of this is due to the two large events—the RMS of the data excluding these events is about 10 ms.

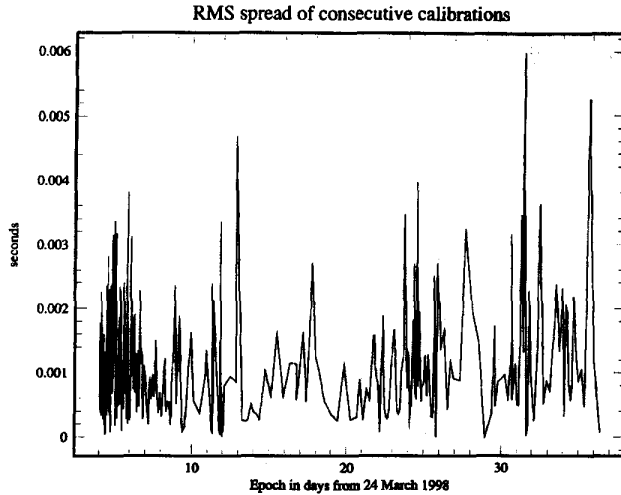


Fig. 3. The RMS spread of consecutive time difference measurements made over a short time interval.

The large fluctuations shown in Fig. 1 are primarily due to large and slowly varying network asymmetries or glitches in the local clock itself. As shown in Fig. 3, they are not due to more rapid jitter in the network asymmetry because they are not visible in the RMS spread of consecutive measurements (S-1). Based on Fig. 3, the synchronization loop should have been able to synchronize the time of the clock with an uncertainty of 1 to 2 ms RMS; the actual loop achieved this on occasion, but the overall performance, even excluding the glitches, was not that good. The reason is that the interval between calibrations was allowed to grow to values larger than the optimum time constant for the frequency averaging loop. The best performance would have been realized if we had clamped the interval between calibrations to a value less than 12,000 s—the measured optimal time constant for the frequency loop. Note that the interval between calibrations was almost always larger than this value.

Fig. 4 shows the average frequency offset of the clock as estimated by the algorithm. Diurnal effects are visible on occasion, but they are smaller than the random-walk frequency noise, which dominates the noise spectrum at this period. These frequency fluctuations have an amplitude of about 3×10^{-6} peak to peak with an irregular period of about 1 week, and they can be used to suggest an initial configuration for a synchronization loop designed to keep the clock on the machine accurate only to the nearest second.

If a static offset of about 10.5×10^{-6} were applied to the clock oscillator of this system and if the machine was then left free running with no other updates, then the time dispersion over the one month period shown in Fig. 4 would have been somewhat less than 1 s RMS. It would not be wise to allow the clock to run freely for this long of a period because there would be no means of detecting a glitch in the operation of the local clock. However, this result emphasizes the fact that keeping the clock correct

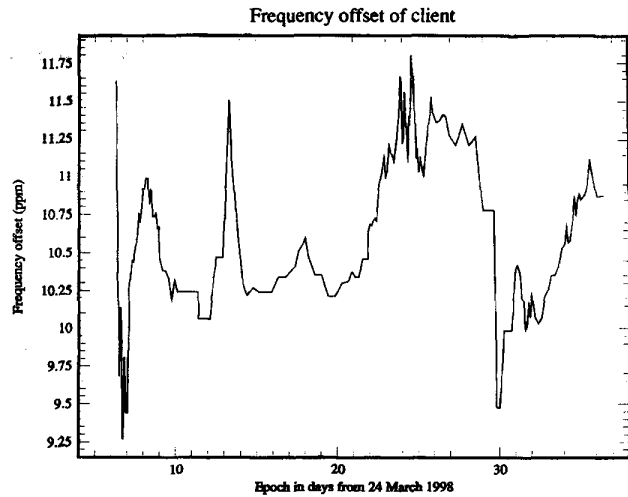


Fig. 4. The frequency offset of the local clock as estimated by the algorithm.

to the nearest second requires very little external information, although the clock in the system we were using would gain almost 1 s/day if it were run with no correction at all. This frequency offset is very stable and can be predicted with considerable accuracy even a month into the future. Even the somewhat poorer clock shown in Fig. 1 of [2], for example, would have a time dispersion of about 1 s with one update about every 18 days. In practice, the interval between calibration cycles would be set more by our concern that we detect a glitch relatively promptly than by the need to adjust the steering of the clock when it is operating normally.

We tested these ideas by configuring the algorithm to achieve an accuracy of 1 s RMS using the same servers as in the previous experiment. We set the initial interval between calibration cycles to 32,000 s, and we specified that the interval should not exceed 200,000 s. The time constant for the frequency update loop was left at 12,000 s; because the minimum interval was larger than this value, the effective time constant was three sample times as discussed previously.

It took just over 10 days for the interval between calibrations to reach the maximum specified value. The time differences during this period are shown in Fig. 5. Note that the RMS time difference increases with the interval between calibrations, but has not yet reached the specified accuracy of 1 s RMS, even with an interval between calibrations of 200,000 s. The time constant for the frequency update loop (600,000 s) is long enough to average over diurnal thermally induced frequency fluctuations, and the main contributors to the variance will be the random-walk frequency noise of the clock oscillator and the long period changes in the ambient temperature. Neither of these effects can be accurately modeled using the statistical machinery that we have developed; so the long-term performance of the clock is difficult to predict. The performance

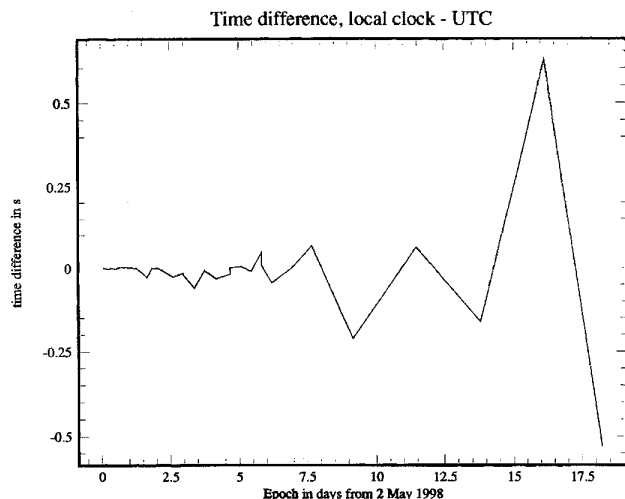


Fig. 5. The time difference in seconds between the local clock and UTC when the algorithm is configured to realize an accuracy of 1 s RMS.

to date has been substantially better than the design specification of 1 s RMS.

As we discussed previously, the highest accuracy with this particular configuration would have been achieved with an interval between calibrations of about 4,000 s because that would have resulted in three calibration cycles in a time equal to the time constant of the frequency update loop. The performance shown in Fig. 5 is about 50 times less expensive in terms of network bandwidth and server load.

II. CONCLUSIONS

We have designed an algorithm for synchronizing the clock of a computer using messages transmitted over the Internet. The algorithm uses the network and message formats of the NTP [4] and can communicate with conventional unmodified NTP servers.

The algorithm is based on our previous frequency-locked designs, but this version automatically adapts to changes in the symmetry of the network delay. It also automatically adjusts its operating parameters to realize a specified level of performance. The adjustment algorithm uses the ratio of the RMS performance to the average cost as the metric for choosing the optimal operating conditions.


In addition to synchronizing the local clock, the algorithm has a number of other diagnostic modes. It can be used to evaluate the Allan deviation of the freely running local clock or to estimate the jitter in the delay in a local network. These evaluations use the same machinery as the synchronization loop, except that the local clock is not steered in these modes.

We have conducted two experiments to illustrate the capabilities of the procedure. The first experiment was designed to synchronize a clock as accurately as possible; it achieved an accuracy of 30 ms RMS and would be capable of achieving an accuracy of better than 10 ms RMS if its configuration were changed as described in the text. This performance was realized using servers that were more than 2,000 km away, and neither the servers nor the paths to them were conditioned in any way for this work.

The second experiment was designed to synchronize a clock so that it was correct to the nearest second, a level which may be adequate for many users. The maximum interval between calibrations was set to 200,000 s. During the first two weeks of operation, the performance of the loop exceeded the specified performance level even at the longest permitted interval between calibrations.

REFERENCES

- [1] D. L. Mills, "Improved algorithms for synchronizing computer network clocks," *IEEE/ACM Trans. Networking*, vol. 3, no. 3, pp. 245–254, Jun. 1995.
- [2] J. Levine, "An algorithm to synchronize the time of a computer to universal time," *IEEE/ACM Trans. Networking*, vol. 3, no. 2, pp. 42–50, Feb. 1995.
- [3] —, "Time synchronization using the Internet," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 45, no. 2, pp. 450–460, 1998.
- [4] D. L. Mills, "Network time protocol (version 3); specification, implementation and analysis," DARPA Network Working Group, Newark, DE, Rep. RFC-1305, 1992.
- [5] K. Arvind, "Probabilistic clock synchronization in distributed systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 5, pp. 474–487, Mar. 1994.
- [6] F. Cristian, "A probabilistic approach to distributed clock synchronization," *Distrib. Comput.*, vol. 3, pp. 146–158, Jun. 1989.



Judah Levine was born in New York City in 1940. He received a B.A. degree from Yeshiva College in 1960 and a Ph.D. degree in Physics from New York University in 1966. He is currently a physicist in the Time and Frequency Division of NIST and a Fellow of JILA, a joint institute operated by the University of Colorado and NIST. He has held these positions since 1969.

His research interests include developing methods for distributing time and frequency information using both satellite and ground-based techniques with special emphasis on the statistics of time and frequency distribution.

Dr. Levine is a fellow of the American Physical Society and a member of the American Association of Physics Teachers and the American Geophysical Union.