

THE ALGORITHM USED TO REALIZE UTC(NIST)

Judah Levine and Tom Parker
 Time and Frequency Division
 National Institute of Standards and Technology
 Boulder, Colorado 80305
 USA

Abstract - The National Institute of Standards and Technology (NIST) maintains a real-time estimate of Coordinated Universal time (UTC), which is called UTC(NIST). UTC(NIST) is realized by steering the output of a hydrogen maser to minimize the difference between UTC(NIST) and UTC, a time scale computed by the International Bureau of Weights and Measures (BIPM). The difference between UTC(NIST) and UTC is published in the BIPM monthly Circular T. The reference for UTC(NIST) is AT1, a free-running time scale computed at NIST in real-time from an ensemble of cesium clocks and hydrogen masers maintained at the NIST laboratory in Boulder. The steering that implements UTC(NIST) is realized using a commercial phase stepper that produces an output that can be offset either in time or in frequency from the input signal from the maser. The current algorithm, which has been used for many years, maximizes the frequency smoothness of UTC(NIST) at the expense of its time accuracy. Since both AT1 and UTC have improved significantly in the last few years as a result of better clock hardware, this algorithm may no longer be optimum. We are investigating other choices which may be able to preserve the desirable goal of frequency smoothness while at the same time improving the RMS time accuracy.

Keywords: Time scale, UTC

THE NIST CLOCK ENSEMBLE AND AT1

The National Institute of Standards and Technology maintains a clock ensemble consisting of about 10 commercial cesium clocks and hydrogen masers. The times of the clocks in the ensemble are measured with respect to one of them (which is designated as the working standard) every 2 hours, and the weighted average of these data is used to construct a time scale called AT1. The ensemble algorithm automatically assigns a weight to each clock based on its past performance; there are also administrative limits imposed on this process to prevent any one clock from having a relative weight of more than 30%. (The working standard is chosen for its stability and longevity – it is not necessarily the “master clock” in the ensemble and its time is generally not the same as UTC(NIST) or the ensemble average time.)

THE DEFINITION AND REALIZATION OF UTC(NIST)

The 5 MHz output of one of the hydrogen masers in the scale is passed through a computer-controlled phase stepper, and the 1 pps output of this phase stepper is the real-time realization of UTC(NIST). In order to steer the phase stepper, its 5 MHz output signal, which has a known and stable relationship to its 1 pps output, is measured with

respect to the working standard as if it were an ordinary clock in the ensemble, except that these data are assigned a weight of 0 in computing the ensemble average time. As a result of this process, the time difference between the output of the phase stepper and the ensemble average time is computed by the time-scale algorithm exactly as for any other clock. The phase stepper is controlled so that these measured time differences are steered towards UTC based on the data published in the BIPM Circular T. Specifically, the phase stepper is controlled so that the measured time difference between its output and the ensemble average time is equal to the value computed from the equation

$$x(t) = x_k + y_k(t - t_k), \quad (1)$$

where x_k and y_k are the specified time offset and rate offset, respectively, between UTC(NIST) and AT1 for the period starting at epoch t_k , and $x(t)$ is the measured time difference between the output of the phase stepper and the ensemble average time measured at epoch t , which is at or after the epoch t_k . (The AT1 ensemble time does not include leap seconds, so that the value in eq. 1 is modulo 1 s.)

For many years, the time offset and rate offset parameters in eq. 1 were changed only at 0 UTC on the first day of each month. In every case, the parameters were computed so that the time computed from eq. 1 was continuous at the month boundary and only the steering frequency was changed. Therefore, except at these monthly transitions, the frequency stability of UTC(NIST) was identical to the frequency stability of AT1, apart from a small amount of white phase noise introduced by the phase stepper. In order to minimize this white phase noise, phase stepper commands were computed and transmitted to the hardware every 12 minutes, and the magnitude of a time-step command was limited to 10 ps under normal operating conditions. (This limit is easily maintained, since the free-running frequency stability of the maser results in a time dispersion over 12 minutes of much less than this value.) The monthly changes to the offset frequency in eq. 1 were limited to ± 2 ns/day, which corresponds to a maximum fractional frequency step of about $\pm 2.3 \times 10^{-14}$ at the monthly boundary. Although changes of this magnitude were quite common in the 1980s, this maximum frequency adjustment is rarely used at present – the last frequency adjustment of this magnitude was in April, 1994.

The rate offsets (the values of y_k) that have been used to define UTC(NIST) since January, 1988 are shown in the following figure. In each case, the corresponding time offset, x_k , is computed so that there will be no discontinuity at the monthly boundary in the time difference computed using eq. 1. (The long-term increase in the steering

frequency shown in fig. 1 is a reversal of a nearly monotonic decrease in the steering frequency that started in the mid 1980s. The minimum in the steering frequency was reached in mid-1997, shortly before the start of the period shown in the figure.)

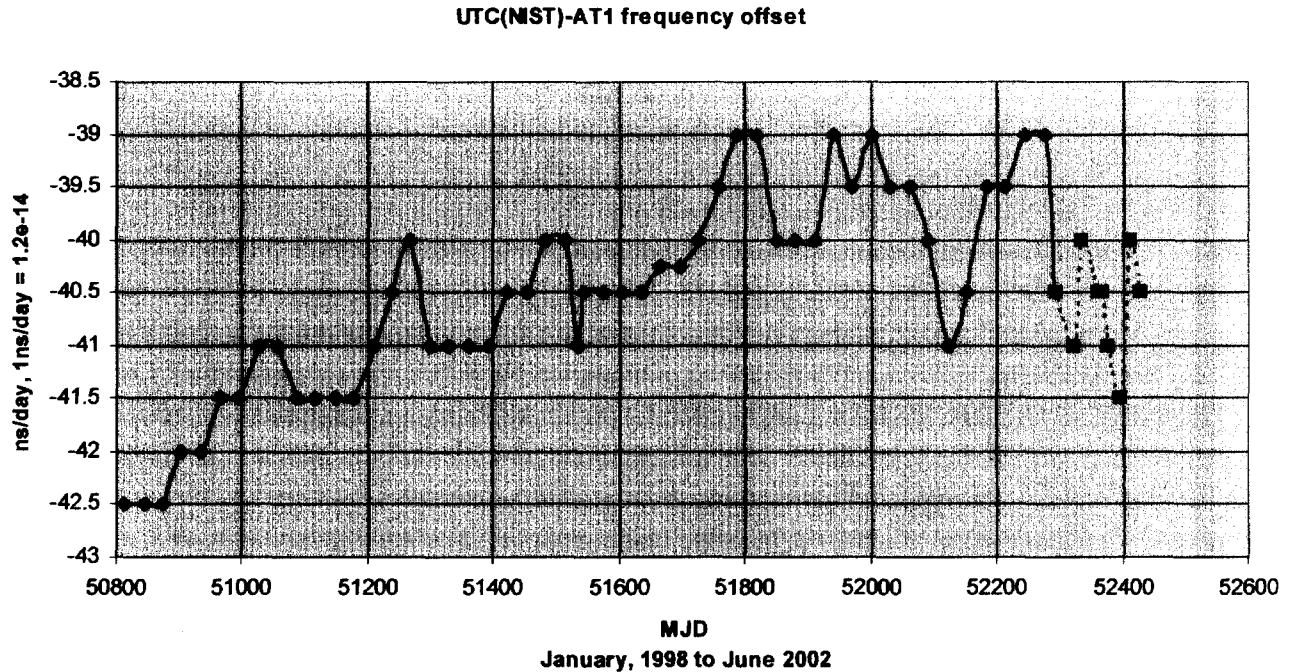


Fig. 1. The monthly rate offsets that have been used to compute UTC(NIST) from the free-running time scale AT1 using eq. 1. The values connected by dotted lines near the end of the plot show the recent corrections computed using the method described in this paper.

THE FREQUENCY STABILITY OF UTC

Except for leap seconds, UTC is identical to International Atomic Time (TAI). The TAI time scale, in turn, is computed from two contributions: (1) the weighted average of about 250 cesium clocks and hydrogen masers located at various National Metrology Institutes and timing laboratories, and (2) data from a much smaller number of primary frequency standards, which are used to apply small frequency steering corrections to TAI. The size and timing of these adjustments is determined by the BIPM. Since these corrections are applied relatively infrequently, the stability of TAI for averaging times of a few months or less is essentially the same as the stability of the ensemble average of cesium clocks and hydrogen masers, which is called EAL. The stability of EAL is given in the BIPM Circular T as the sum of three contributions: a white frequency noise whose magnitude is $6 \times 10^{-15} / \tau^{1/2}$, a flicker component whose magnitude is 6×10^{-16} , and a random walk frequency noise whose magnitude is $1.6 \times 10^{-16} \tau^{1/2}$. In each case, τ is a time interval measured in days.

The retrospective nature of TAI and the mechanical delays in the publication of Circular T by the BIPM mean that eq. 1 must always be an extrapolation. For the last few years, the Circular T for any month is usually published on or about the 15th of the following month. Because of the fact that the NIST procedure was to change the steering parameters only at the start of any month, the Circular T data for any month, which was received on the 15th of the following month, would not have any effect on the steering until the 1st day of the next month. The maximum period of the required extrapolation was therefore 2 months, and small changes in the publication date of Circular T do not have any affect on this value.

Using the parameters from Circular T defined above, the stability of TAI for averaging times of 30-60 days is about 1.5×10^{-15} . This averaging period is in a transition zone in which no one noise type dominates the Allan variance. For an intermediate averaging time of 40 days, for example, the amplitudes of white FM, flicker FM and random walk FM are about 0.95×10^{-15} , 0.6×10^{-15} , and 1×10^{-15} ,

respectively. The time dispersion over a 40 day extrapolation due to this frequency instability would be of order $\tau \sigma_y(\tau) \approx 5$ ns RMS, and this is probably the best that any real-time realization of UTC can provide, given the current stability of TAI and the retrospective nature of its computation. (This assumes that UTC(k), the realization of UTC by laboratory “k,” is statistically independent of UTC itself. This is discussed in more detail below.)

THE FREQUENCY STABILITY OF AT1 AND UTC(NIST)

The UTC(NIST) steering equations are defined with respect to our free-running time scale AT1, and no matter what algorithm we choose to use to realize UTC(NIST), the time dispersion between UTC(NIST) and UTC between the monthly computations of TAI and publications of Circular T can never be better than the frequency stability of AT1 with respect to TAI over this same period. Although the clocks that form the ensemble that defines AT1 are also reported to the BIPM for the construction of EAL, TAI, and UTC, the NIST contribution to the BIPM scales is quite small (of order 2%), so that AT1 and TAI can be considered as essentially statistically independent of each other. (This is probably true of most of the laboratories that contribute clock data to the BIPM with the exception of the US Naval Observatory (USNO). The large number of clocks at the USNO that also contribute to the BIPM time scales means that any USNO time scale is likely to have a significant correlation with EAL, and therefore with TAI and UTC.)

The stability of AT1 with respect to TAI can be evaluated in a number of different ways. Using the BIPM data for TAI-AT1 (AT1 is called TA(NIST) by the BIPM) for the last several years, the stability of AT1 with respect to TAI is about 1.3×10^{-15} for an averaging time of 40 days and is not greater than about 1.7×10^{-15} over the range in averaging times from 30-60 days. A “3-corner hat” evaluation using data from NIST, PTB, and USNO gives a somewhat poorer stability estimate of about 1.8×10^{-15} for AT1 over the same range of averaging times. It is not obvious how to combine this estimate with the stability of TAI itself. The most conservative estimate would simply sum the two estimates; another alternative would be to sum them in quadrature, and still a third possibility would be simply to take the larger of the two as the estimate of the stability of the difference. A stability of 2×10^{-15} is a reasonable estimate of the stability of AT1 with respect to TAI, and that this value also represents a reasonable estimate for the stability of UTC(NIST) with respect to UTC between the monthly computations of TAI. Using this value for the frequency stability, the time dispersion over 40 days between UTC and UTC(NIST) would be about $\tau \sigma_y(\tau) \approx 7$ ns RMS. The exact value of this estimate would depend on details of the noise spectrum, but this approximate calculation is certainly correct to within a factor of 2. This value is roughly the same as the time dispersion of UTC

itself that was estimated above. Therefore, improving the statistical performance of AT1 (by adding more clocks, for example) probably would not result in a significant improvement in the RMS difference between UTC(NIST) and UTC until the number of clocks that were common to AT1 and EAL was so large that UTC(NIST) and UTC were no longer statistically independent.

THE HISTORICAL STABILITY OF UTC(NIST)

The values of UTC-UTC(NIST) as computed by the BIPM are shown in the following figure.

The Allan deviation of these data is about 3.4×10^{-15} for an averaging time of 40 days, and the RMS value of all of the data is 16 ns. This value of the Allan deviation is almost a factor of 2 larger than we would have expected based on the previous discussion, suggesting that the current steering algorithm is not optimum. At least one obvious problem is that the extrapolation interval of 60 days that we currently use is longer than necessary, and could be decreased if we were willing to apply steering changes to UTC(NIST) in the middle of a month when we receive Circular T instead of waiting until the start of the next month to make use of these data.

We also note that the variance is dominated by a quasi-periodic effect with a period of about 200 days. It is possible that this is a seasonal variation driven by long-term changes in temperature; it may also be a result of the time constant in the parameter estimation of the current steering algorithm, which is about this same period. If the quasi-periodic variation really is driven by something like seasonal temperature variations, and if we could model this variation using a static or slowly-varying admittance to externally measured temperature fluctuations, then using that strategy to steer UTC(NIST) would be better than our current ideas, which are based on purely stochastic models. Although a correlation with temperature has been suspected for a long time, it has never been possible to identify a temperature record that can act as a robust predictor of the observed variation in the frequency of the NIST clock ensemble.

TESTS OF STEERING ALGORITHMS USING HISTORICAL DATA

The current steering algorithm is limited by three constraints: (1) that time steps will never be used, (2) that changes to the parameters in the steering equation (eq. 1) will be made only at the start of a month, and (3) that the maximum rate change at the start of any month will be limited to ± 2 ns/day. We have used historical Circular T data to evaluate the performance of a number of different steering algorithms, which relax (or completely abandon) these constraints.

UTC-UTC(NIST)




Fig. 2. UTC-UTC(NIST) as a function of time as computed by the BIPM. The plot starts at MJD 50800 (18 December 1997) and continues through the most recent data available as of May, 2002. The points connected by a dotted line near the end of the figure show the first four months of data using the steering method discussed in this paper.

Although we would never steer using a true time step, we have tested an approximation of this idea in which we removed the time offset between UTC(NIST) and UTC(NIST) by applying a relatively large frequency step (up to ± 3 ns/day) for a relatively short periods of time (5 days) starting on the first MJD ending in 4 or 9 after the new Circular T would have been received. The time offset was taken to be the value of UTC-UTC(NIST) on the last MJD of the report. (This is the “fast” steering test in the following table and figure.) This approach will clearly degrade the frequency stability of UTC(NIST) at short periods, since these steering corrections are much larger than the free-running stability of AT1. On the other hand, since this process tightly couples UTC and UTC(NIST), the Allan deviation of the difference between these two time scales at long periods will be a minimum with this approach. However, this does not necessarily mean that the stability of UTC(NIST) has been improved in any absolute sense – for averaging times of a few months the process may have simply replaced the instabilities of AT1 and UTC(NIST) with the roughly comparable, but uncorrelated, instabilities of TAI and UTC.

We also tried a number of variations of this idea which used smaller frequency adjustments (± 1 ns/day or less) to remove the time offset over a longer period of time ranging from 2 weeks to 2 months. (These are identified as the

“moderate” and “slow” steering tests, respectively, in the following table and figure.) While all of these adjustments are more aggressive than the current algorithm, the average frequency changes due to the “moderate” and “slow” steering algorithms are now comparable to the free-running stability of AT1 itself, so that they improve the time accuracy but have a much smaller impact on the frequency stability. (A frequency adjustment of 1 ns/day for 14 days would result in an effective time step of 14 ns relative to a scale that did not have that adjustment. When averaged over a two-month interval, this time step is nearly invisible to most users because it is comparable to the time dispersion over the same averaging time due to the frequency instability of the scale itself. However, a user whose clock had very good short term stability could detect this step with a few days of averaging. Even with GPS common-view, the noise in the time transfer process would mask the time dispersion due to this frequency offset at shorter periods.) The time differences UTC-UTC(NIST) that would have resulted from each of these steering algorithms are characterized in the following table, which shows the RMS value of the difference, and the maximum and minimum values over the interval of the experiment. The corresponding values for the current algorithm are shown for comparison.

TABLE 1
Parameters of the steering experiments
Using Circular T data from MJD 51200 – MJD 52200

Experiment	UTC-UTC(NIST)	
	RMS	MAX, MIN
Fast	8 ns	+25 ns, -17 ns
Moderate	10 ns	+24 ns, -23 ns
Slow	12 ns	+25 ns, -24 ns
Current	15 ns	+40 ns, -27 ns

Using the same historical data set, we also computed the Allan deviations of UTC-UTC(NIST) that would have resulted using the various steering procedures. The results are presented in fig. 3, which also shows the Allan deviation of the current procedure for comparison. In all of these cases, the more aggressive steering reduces the long-term Allan deviation of the difference UTC - UTC(NIST) relative to the deviation obtained using the current algorithm. (For all of the algorithms we tested, this

improvement in the Allan deviation at long time intervals is the chief improvement over the current method.) Although the more aggressive steering also improves the RMS value of the difference between UTC and UTC(NIST), the improvement is not very dramatic (See table 1, above). However, the more aggressive steering does reduce the maximum and minimum values of this difference quite significantly, because the current algorithm responds very slowly and with a significant time delay to a frequency change. This delay and slow response is especially noticeable in the response of the current method to the random-walk frequency noise in AT1, especially near an inflection point of the quasi-periodic variation shown in fig 2. The resultant time dispersion can grow to a relatively large value before it will be removed. (The large offset in fig. 2 near MJD 52200 is an example of this effect. There have been others in the past, although this is the largest one shown in this data set.)

Allan Deviations of Steering Tests




Figure 3. Allan deviations of the differences UTC-UTC(NIST) for the steering tests described in the text. The curve “UTC-UTC(NIST)” gives the Allan deviation for the current steering algorithm. The other algorithms are described in the text.

DISCUSSION AND CONCLUSIONS

Although the “fast” steering algorithm realizes the best improvement in the RMS value of UTC-UTC(NIST) of all of the algorithms we have tried, it degrades the frequency

stability of UTC(NIST) by a significant amount (see fig. 3), and we consider this unacceptable at this time. We are currently experimenting with defining the steering parameters for UTC(NIST) using a variation of the “moderate” algorithm. From the results in table 1, this

algorithm improved the RMS time difference nearly as much as the “fast” algorithm but did not have any noticeable effect on the short-term frequency stability compared to our current algorithm. This version of the moderate steering algorithm uses up to two frequency adjustments per month: one at the start of the month and a second optional one that is applied shortly after Circular T is received, but not necessarily on a day whose MJD value ends in 4 or 9. We do not think that limiting the insertion of a steering correction in this way is necessary, since the BIPM observes

UTC(NIST) nearly continuously using GPS common view and two-way satellite time transfer. Both steering adjustments will be limited to ± 1.5 ns/day for the present, with the expectation that the maximum frequency adjustment will not be used very often, and that smaller corrections of ± 1 ns/day or less will be used in general. This new algorithm has been in use for about 4 months (since February, 2002), and the values of UTC-UTC(NIST) for this period are shown connected by a dotted line near the end of the data in fig. 2. It is too soon to draw any conclusions about this test.