Impact of New High Stability Frequency Standards on the Performance of the NIST AT1 Time Scale

Thomas E. Parker, *Fellow, IEEE,* and Judah Levine

***Abstract*—The recent addition of new commercial high stability frequency standards to the National Institute of Standards and Technology (NIST) real time AT1 time scale has resulted in significant improvements in the performance of the scale. The frequency stability of the scale at one day has increased by a factor of 2 to 4×10^{-15} and the stability at 100 days has improved to approximately 1×10^{-15}. As a result UTC (NIST) has been kept within 50 ns of UTC (Coordinated Universal Time) for the last year. Further improvements are anticipated as more high performance clocks are added to the AT1 ensemble.**

I. INTRODUCTION

THE AT1 TIME SCALE at the National Institute of Standards and Technology (NIST) operates in real time and consists of three main components: an ensemble of 8 to 10 commercially available frequency standards, the associated time interval measurement hardware, and the AT1 algorithm [1], [2]. An ensemble of clocks is used because it increases the reliability of the time scale and, in principle, has better frequency stability than any individual clock. The algorithm is used to optimally combine the measured clock difference data and to calculate the scale time. The AT1 time scale is used as a reference for the generation of the UTC (NIST) and also can be used as a reference for research on primary frequency standards. Currently, there are two hydrogen masers and seven cesium frequency standards in AT1, and all are kept in environmental chambers that control temperature and relative humidity.

In the last few years a new generation of commercial high performance frequency standards which offers significantly improved frequency stability over earlier devices has become available [3], [4]. As these new standards are incorporated into the clock ensemble, it is fully expected that the time scale performance will improve [5]. However, because AT1 is a real-time time scale, we have been very conservative about making major changes in AT1. We report here the increased frequency stability of the AT1 time scale resulting from the addition of these new clock technologies [6].

AT1 is currently in a transition phase and is presently dominated by two cavity-tuned hydrogen masers[1] and three new high performance cesium frequency standards (Hewlett-Packard model HP5061B with high performance tube). The masers have excellent short-term frequency stability ($\sigma_y(\tau = 1 \text{ day}) \approx 1 \times 10^{-15}$), while the cesium frequency standards offer outstanding long-term stability. $\sigma_y(\tau = 1 \text{ day})$ for the cesium standards is $\sim 2.5 \times 10^{-14}$ and $\sigma_y(\tau = 100 \text{ days}) \approx 3.5 \times 10^{-15}$. The short-term weight of a clock in the scale is inversely proportional to $\sigma_y^2(\tau = 1 \text{ day})$, but the maximum short-term weight of any individual clock in AT1 is administratively limited to 30%. Therefore, the two masers have only 60% of the short-term weight in the scale. As a result the time scale is in an unusual state. Except for frequency drift, the two masers exhibit better frequency stability than the scale. This state will continue to exist until four masers are present in the scale.

II. NEW HIGH PERFORMANCE FREQUENCY STANDARDS

Before examining the performance of AT1, it is useful to illustrate, in a qualitative sense, the improved frequency stability of the new commercial high performance frequency standards that are now operating at NIST. As a reference we will use one of our hydrogen masers[1], which will be referred to as M1. Fig. 1 shows the normalized frequency difference between M1 and the primary frequency standard NIST-7 [7] for a recent 500-day period. Days are given in Modified Julian Dates (MJD). The maser has a fractional frequency stability at one day of about 1×10^{-15} and, as shown in Fig. 1, a very constant average frequency drift rate of -1.1×10^{-16}/day. This constant and relatively low drift rate makes M1 a very useful reference. The error bars in Fig. 1 indicate the 1σ uncertainty of NIST-7. Data comparing the frequency of M1 to the International Atomic Time (TAI) as generated by the Bureau International des Poids et Mesures (BIPM) shows that the drift rate of -1.1×10^{-16}/day also extends fairly accurately back to MJD 49400.

Fig. 2 shows the normalized frequency difference as a function of time between one of our older generation high

Manuscript received July 15, 1996; accepted May 9, 1997.

The authors are with National Institute of Standards and Technology, Time and Frequency Division, 847.5, Boulder, CO 80303 (e-mail: tparker@bldrdoc.gov).

---

[1] Sigma Tau Standards series 2000 cavity tuned hydrogen masers. Commercial equipment has been identified since it would be impossible to duplicate the results without this information. Such identification does not imply recommendation or endorsement by the National Institute of Standards and Technology.

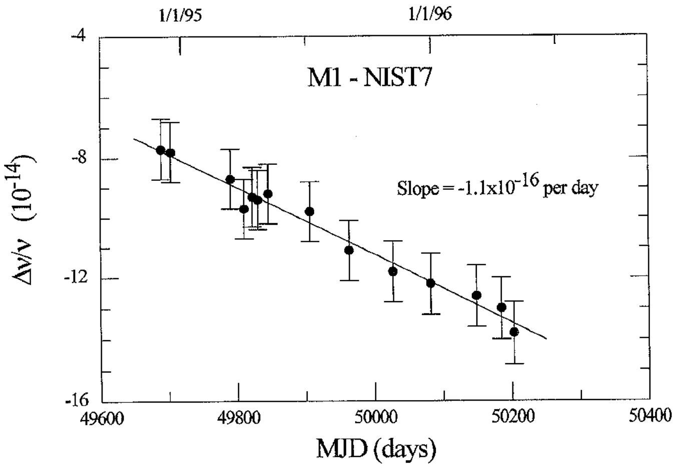

Fig. 1. Normalized frequency difference between maser M1 and NIST-7.

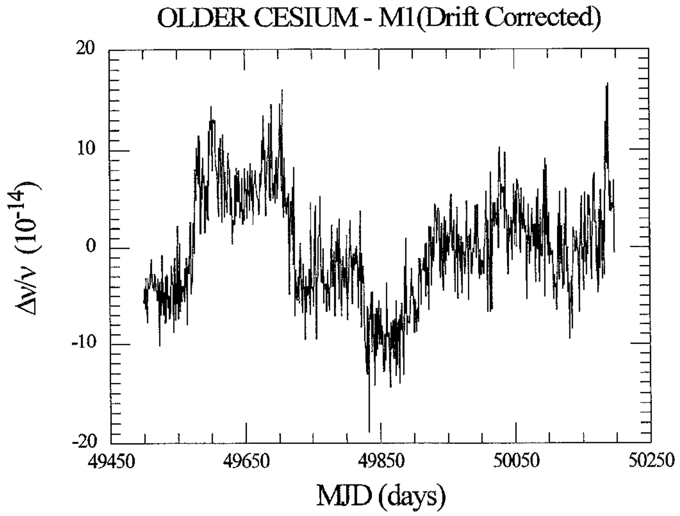

Fig. 2. Normalized frequency difference between an older generation high performance cesium standard and maser M1.

performance cesium beam frequency standards (Hewlett-Packard model HP5061B with high performance tube) and maser M1 for a 700-day period. The mean frequency difference and the average frequency drift of the maser (-1.1×10^{-16}/day) have been removed and the data are plotted as 24-hour averages. This frequency data was obtained from the first difference of the clock data. Virtually all of the frequency variations present in the data (some as large as 1×10^{-13}) are from the cesium standard. $\sigma_y(\tau)$, at τ equals 1, 10, and 100 days, is 3.0×10^{-14}, 1.8×10^{-14}, and 3.8×10^{-14}, respectively, for the data in Fig. 2. Until early 1994 this type of cesium standard provided the highest performance clocks in the AT1 ensemble. For convenient comparison, each minor tick on the y axis represents 1×10^{-14} in all of the frequency versus time plots in this paper.

Fig. 3 is a similar plot for a new generation high performance cesium standard (Hewlett-Packard model HP5071A with high performance tube). There is a small reduction in the level of the day-to-day frequency fluctuations as compared to the older standard in Fig. 2, but the long-term frequency stability is dramatically improved. $\sigma_y(\tau)$, at τ equals 1, 10, and 100 days, is 2.2×10^{-14}, 8.0×10^{-15}, and 3.8×10^{-15}, respectively, for the data in Fig. 3. Three of these standards were added to the AT1 ensemble in 1994.

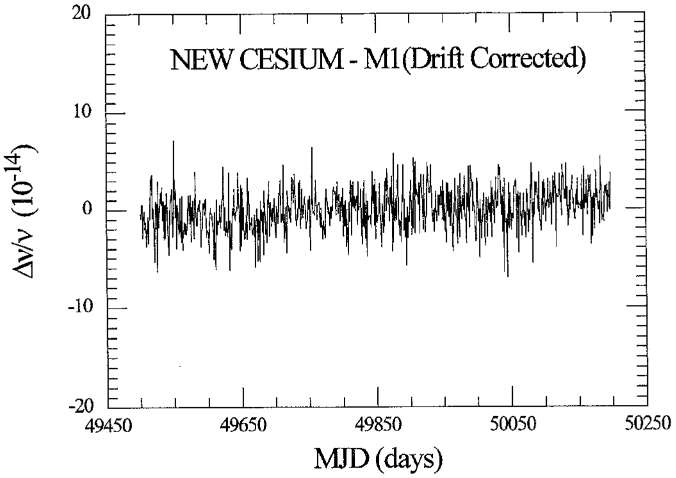

Fig. 3. Normalized frequency difference between a new generation high performance cesium standard and maser M1.

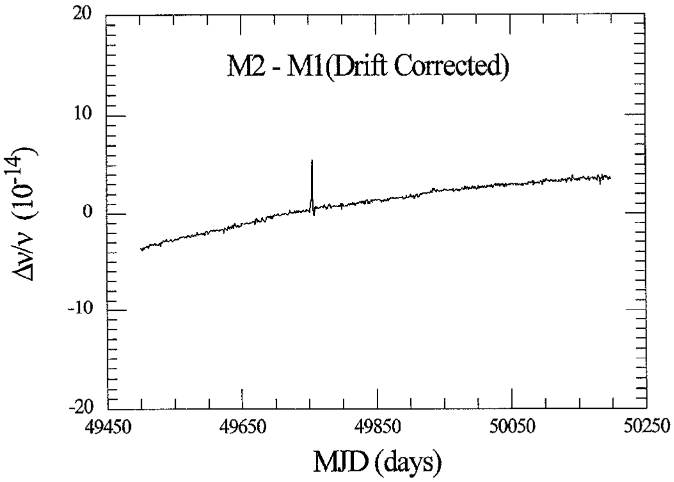

Fig. 4. Normalized frequency difference between masers M2 and M1.

Fig. 4 shows the normalized frequency difference between M1 and another maser, M2, of the same type. The scale of Fig. 4 is the same as that of Figs. 2 and 3. Fig. 4 clearly illustrates the main frequency stability characteristics of these hydrogen masers. They have very low short-term frequency fluctuations and, in contrast, significant long-term frequency drift. The day-to-day fluctuations observable in Fig. 4 are a combination of those from both masers since they have comparable noise levels. However, the frequency drift is predominantly that of M2. The frequency spike that occurs around MJD 49755 is from M1 and was caused by the failure of its environmental chamber. $\sigma_y(\tau)$, at τ equals 1, 10, and 100 days, is 2.1×10^{-15}, 1.2×10^{-15}, and 7.0×10^{-15}, respectively, for the data in Fig. 4. Maser M1 was added to AT1 in 1994 and M2 was added in 1995.

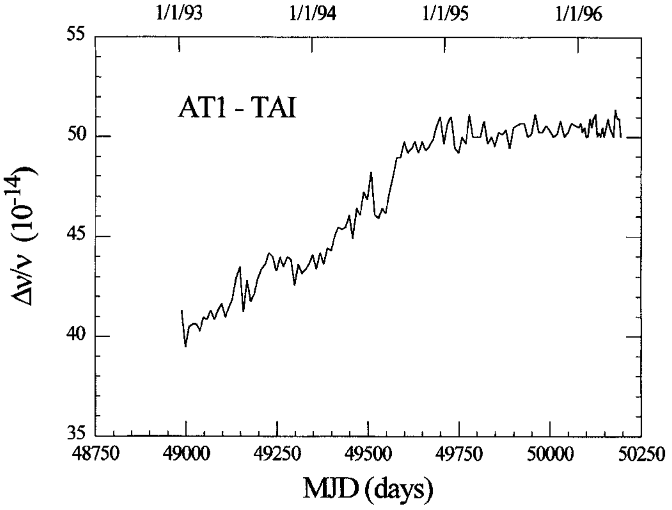

Fig. 5. Normalized frequency difference between the NIST AT1 and TAI.

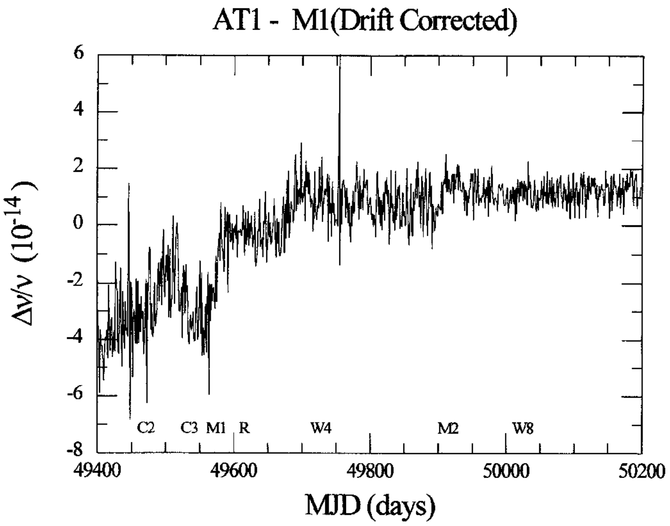

Fig. 6. Normalized frequency difference between AT1 and M1.

III. AT1 Time Scale Performance

The normalized frequency difference between the NIST AT1 time scale and TAI for the last 3 1/2 years is shown in Fig. 5. Each data point represents a 10-day average frequency difference, except for data after January 1, 1996, where the interval has decreased to 5 days. The most striking feature of the data in Fig. 5 is the decrease in frequency drift of AT1 from approximately $+1.3 \times 10^{-16}$/day to approximately $+1.2 \times 10^{-17}$/day that occurred near the end of 1994. There is also a decrease in random frequency fluctuations starting at about the same time, but this is not so obvious from the data in Fig. 5. This aspect will be discussed in more detail later. The improved frequency stability of AT1 follows closely the introduction of new high performance frequency standards and also some changes in the parameters of the algorithm. The frequency of AT1 is not steered (it is a free running time scale), so the large fractional frequency difference of 5×10^{-13} between AT1 and TAI is not an issue.

The data in Fig. 5 clearly show the long-term frequency characteristics of AT1, but the 10 or 5 day data interval of TAI is not useful for investigating the performance of AT1 at time intervals as small as 1 day. Therefore, we will again use maser M1 as a frequency reference. With excellent short-term stability and a constant drift rate, M1 makes a very good reference for evaluating the performance of AT1. The use of M1 as a reference is possible only because the scale is in a transition phase in which one or two individual clocks are more stable than the scale. This does provide, however, an opportunity to document in great detail the improvements in the scale as changes are made.

Fig. 6 shows the normalized frequency difference between AT1 and M1 for the 800-day period from MJD 49400 (February 17, 1994) to MJD 50200 (April 27, 1996). As with previous figures the mean frequency difference and average drift of the maser have been removed so that essentially only the drift of AT1 remains. During this entire period the scale was significantly noisier than M1, so that virtually all of the detail shown comes from AT1. Only the frequency spikes near MJD's 49445 and 49755 are due to M1. This figure contains information similar to that of Fig. 5, but now more short-term detail is visible since each data point represents a 26-hour average rather than a 10- or 5-day period. The decrease in drift rate starting near MJD 49700 is clearly visible and a decrease in day-to-day fluctuations of approximately a factor of 2 is also evident between the first 200 days and the last 200 days.

The notations along the bottom edge of Fig. 6 indicate when major milestones occurred in the evolution of AT1. C2 and C3 indicate when the second and third of the new high performance cesium frequency standards entered the ensemble. The first new high performance cesium standard entered the scale before the start time in Fig. 6. M1 and M2 indicate when masers M1 and M2 entered the scale. The frequency drift of the masers is not currently modeled in AT1. All new clocks are initially put into the scale at low weight, and it can take from 30 to 100 days for a clock to reach full weight in the ensemble. Even though M1 was in the scale for most of the period shown in Fig. 6, the level of correlation between the frequency of AT1 and M1 is small because the short-term weight of M1 in the scale is limited to 30%. This limits the correlation to no more than 17% [8], and, in fact, it is smaller than that since the scale is at least 3 times noisier than M1 (for $\tau = 1$ to 10 days) at all times during the period shown in Fig. 6. The analysis of [8] assumes the scale is always quieter than the clock.

The "R" near MJD 49600 indicates where the reset criterion in AT1 was changed. Prior to MJD 49597, any time interval measurement on a clock that gave a time deviation from the ensemble that was more than 3 standard deviations from the predicted value was declared erroneous. The clock was then given zero weight for that measurement cycle and the predicted value was used to reset the

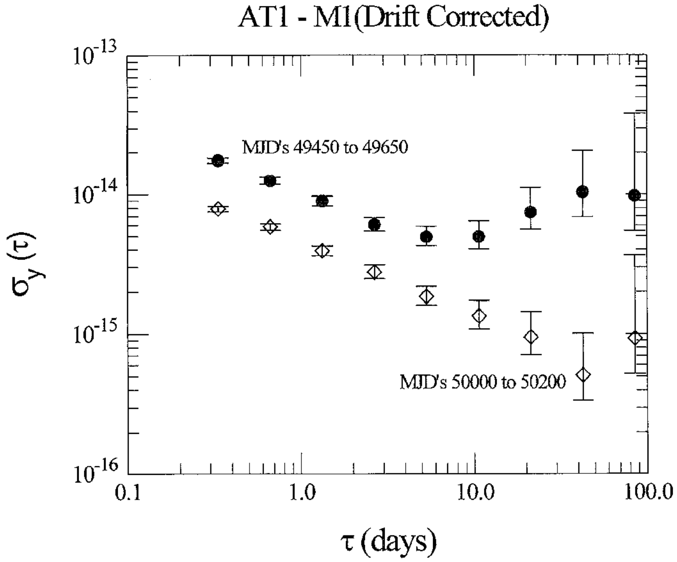

Fig. 7. Allan deviations for the data in Fig. 6. Solid circles are from an early 200-day period, and the diamonds are for the last 200 days. The error bars are 3σ.

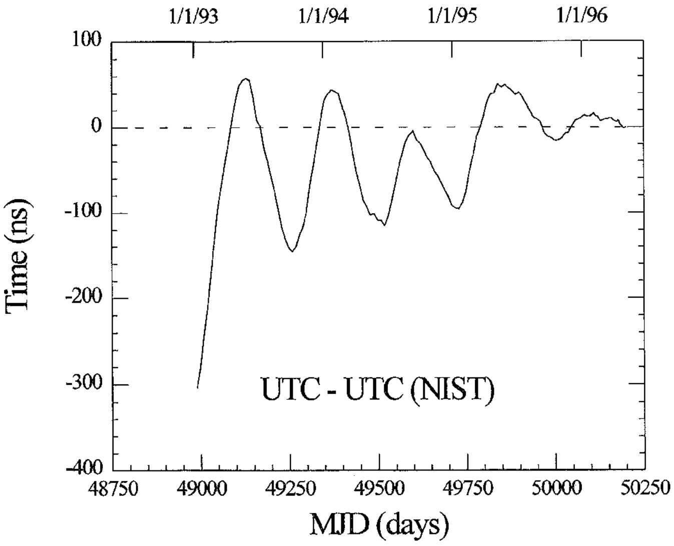

Fig. 8. Difference between UTC and UTC (NIST).

clock. After 49597 the reset algorithm was changed so that a clock is gradually deweighted for errors in the range of 3 to 4 standard deviations. The weight is set to zero only for errors larger than 4 standard deviations [9]. For the period 49400 to 49600, the drift rate of AT1 with the old reset parameters is about $+1.5 \times 10^{-16}$/day. For the same period using an experimental version of AT1 (see Section IV), in which the only change is the new reset criterion, the drift rate is reduced to about 1.1×10^{-16}/day. However, a much smaller difference in drift rates is observed for the two reset algorithms when a later period is examined in which the scale is dominated by the new high performance clocks. This is not surprising because there are fewer resets with the new clocks.

The notations W4 and W8 indicate where the averaging times for the prediction rate (frequency) [2] of the three new high performance cesium standards were increased from 2 days to 4 and 8 days, respectively. These changes serve to increase the long-term weights of these clocks. With these changes the five high performance clocks (three cesiums and two masers) have 99% of the long-term weight in the scale. The weights are divided approximately equally among these five clocks.

Fig. 7 gives a more quantitative evaluation of the improvement in the frequency stability of AT1. The Allan, or two-sample deviations shown in Fig. 7 are calculated from the data in Fig. 6 over two different time periods. The solid circles are for the 200-day period from 49450 to 49650. Here the masers have had virtually no impact yet and the long-term stability of the scale was still heavily influenced by the older generation cesium standards. The diamonds are for the period 50000 to 50200 and represent a period where all of the new clocks are at full weight and all of the changes to the algorithm parameters are in effect. For both of these calculations the drift of the maser has been removed, but not the drift of the scale. The impact of the two masers for τ less than about 5 days is clearly evident. The individual masers are about 10 times quieter than the cesium-based scale, so one would expect a little better than a factor of two reduction in scale noise with the introduction of two masers limited to 30% weight each. This is exactly what is seen in Fig. 7. At τ greater than 10 days, the improvement in the scale is even more dramatic. Here the new high performance clocks (both cesium and maser) have displaced the older cesiums entirely and the long-term stability has improved by more than a factor of 10.

With the improved frequency stability of AT1 as demonstrated in Figs. 6 and 7, one would also expect to realize better performance in generating UTC (NIST). UTC (NIST) is a steered version of AT1 in which fractional frequency steps are inserted that are no larger than 2.3×10^{-14} (2 ns/day) at a rate no more often than once per month. Fig. 8 shows the difference between UTC and UTC (NIST) for the last 3 1/2 years; it also exhibits a dramatic improvement. UTC (NIST) has been within 50 ns of UTC for the last year, and within 100 ns for about the last 2 years.

IV. Experimental Versions of AT1

The data in Figs. 6 and 7 suggest that the new high performance clocks, as well as the associated adjustments to the parameters of AT1, are responsible for the improved performance of AT1. However, this hypothesis can be further tested by using an experimental version of AT1 that is not run in real-time. This allows the same clock data to be used in calculating alternative versions of AT1 using different administrative decisions (when clocks enter or leave the scale, clock weights, number of clocks, etc.). Also different versions of the AT1 algorithm can be used (new reset criterion or old). Fig. 9 shows the normalized frequency

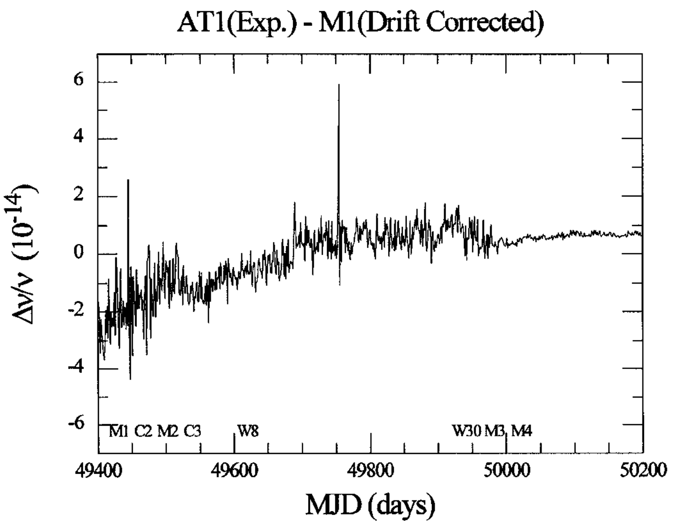

Fig. 9. Normalized frequency difference between the experimental AT1 and M1.

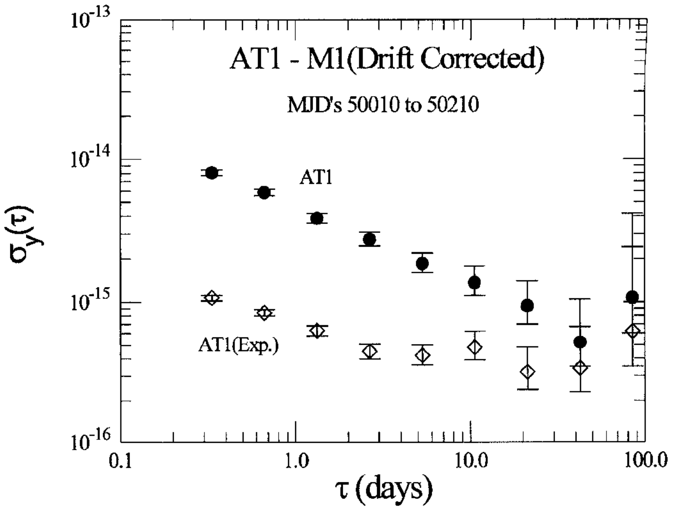

Fig. 10. Allan deviations for a recent 200-day period from the real time AT1 (solid circles) and the experimental version of AT1 (diamonds). The error bars are 3σ.

difference for one such experimental time scale compared to M1. This data covers the same time period as Fig. 6, but several changes have been made. First, the experimental version of AT1 uses the new reset algorithm from the beginning. Also, masers M1 and M2 have been put into the scale earlier and the new high performance cesiums go in initially with higher long-term weight (though at the same time). Eventually two additional masers, M3 and M4, are put into the scale. The maximum short-term weight of any clock remains at 30%. A drift parameter was used in this scale only for M3. Comparing Fig. 9 with Fig. 6, we see that the overall frequency drift is reduced by about a factor of 2. However, there is still significant drift in Fig. 9 for the period 49600 to 49700 even after both M1 and M2 are in the scale and the averaging times of the three new high performance cesiums are increased to 8 days (W8). It is not clear why the drift rate decreased after 49700 because no major changes were made at that time. The large frequency fluctuations in Fig. 6 near 49550 and 49700 are significantly reduced in Fig. 9, and a dramatic reduction in short-term noise occurs after MJD 50000 when M3 and M4 have entered the scale. Though not everything is fully understood, the earlier introduction of the new high performance clocks and the associated parameter changes in the experimental version of AT1 have clearly resulted in better scale performance.

Fig. 10 is a plot of Allan deviations for a recent 200-day period of the real time AT1 (see Fig. 6) and also the same period for an experimental version of AT1 nearly identical to that shown in Fig. 9. This is a period with all four masers in the experimental scale and a dramatic reduction in short-term noise is clearly seen. However, the improvement at τ greater than 10 days is not nearly as large. This is entirely as expected. The third and fourth masers result in a more normal scale operation in which each of the masers is below, or just at, the administrative limit of 30% weight. Thus the short-term stability of the scale is actually better than that of any of the individual masers. This, however, means that the diamonds in Fig. 10 are more representative of M1 than the scale. Also there is correlation between AT1 (Exp.) and M1 that reduces the $\sigma_y(\tau)$ values by about 17% [8]. The less dramatic improvement at τ greater than 10 days occurs because the new high performance cesium standards and masers M1 and M2 already dominate the long-term stability of the real time AT1 scale. Therefore, the presence of the third and fourth masers is less significant.

The data in Fig. 10 show that, with a mixture of high performance cesium standards and cavity-tuned hydrogen masers, a scale with $\sigma_y(\tau)$ less than 1×10^{-15} for τ greater than 1 day and less than 100 days is possible. However, considerable care will be required in adapting the AT1 algorithm to a mixture of clocks with such different noise characteristics [5]. Because the hydrogen masers have significantly lower noise levels in the short term, they will dominate the scale in the short term. However, it is desirable to have the cesium standards dominate in the long term because of their smaller frequency drift rates. Achieving this balance may require further adjustments to AT1 since AT1 was not designed for such a mixture of clocks. Adjustments are not only required for inserting and maintaining a mixture of cesiums and hydrogen masers, but the scale must also be able to handle the sudden loss of a maser from the scale. A fifth maser has recently been acquired and a time scale with five masers could exhibit a flicker floor close to 3×10^{-16} at a few days.

Another issue that needs to be addressed is the environmental sensitivities of the masers. Is it possible that the frequencies of all of the masers have some common mode fluctuation that is not visible when comparing two masers? We do not know of any such problems, but not all environmental parameters have been characterized at the low 10^{-16} level.

V. CONCLUSIONS

The availability of new high performance cesium beam and hydrogen maser frequency standards has resulted in a substantial improvement in the frequency stability of the NIST AT1 time scale. This has resulted in UTC (NIST) being kept within 50 ns of UTC over the last year. Further improvements will occur as additional high performance clocks are added to the scale, but changes will have to be made in the AT1 algorithm to handle the very different frequency stability characteristics of the cesium frequency standards and the hydrogen masers.

ACKNOWLEDGMENTS

The authors acknowledge the valuable assistance of Jim Gray, Trudi Peppler, and Marc Weiss in maintaining and improving the AT1 time scale.

REFERENCES

[1] P. Tavella and C. Thomas, "Comparative study of time scale algorithms," *Metrologia*, vol. 28, pp. 57–63, 1991.

[2] M. Weiss and T. Weissert, "AT2, A new time scale algorithm: AT1 plus frequency variance," *Metrologia*, vol. 28, pp. 65–74, 1991.

[3] H. B. Owings, P. A. Koppang, C. C. MacMillan, and H. E. Peters, "Experimental frequency and phase stability of the hydrogen maser standard output as affected by cavity auto-tuning," in *Proc. 1992 IEEE Freq. Contr. Symp.*, pp. 92–103.

[4] J. A. Kusters and J. L. Johnson, "A new cesium beam frequency standard performance data," in *Proc. 1992 IEEE Freq. Contr. Symp.*, pp. 143–150.

[5] L. A. Breakiron, "Timescale algorithms combining cesium clocks and hydrogen masers," in *Proc. 23rd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting*, 1991, NASA Conference Publication 3159, Pasadena, CA, pp. 297–305.

[6] T. E. Parker and J. Levine, "Recent improvements in the performance of the NIST AT1 time scale," *Proc. 1996 IEEE International Frequency Control Symposium*, pp. 1131–1136.

[7] W. D. Lee, J. H. Shirley, J. P. Lowe, and R. E. Drullinger, "The accuracy evaluation of NIST-7," *IEEE Trans. Instrum. Meas.*, vol. IM-44, no. 2, pp. 120–123, 1995.

[8] P. Tavella and C. Thomas, "Study of the clock-ensemble correlation in ALGOS using real data," in *Proc. 5th European Frequency and Time Forum*, Neuchatel, Switzerland, 1991, pp. 435-441.

[9] M. A. Weiss and T. P. Weissert, "Sifting through nine years of NIST clock data with TA2," *Metrologia*, vol. 31, pp. 9–19, 1994.

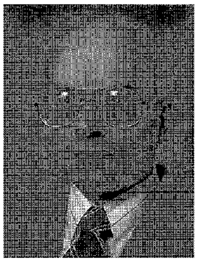

Thomas E. Parker (M'79–SM'86–F'94) was born in Natrona Heights, PA, on September 17, 1945. He received his B.S. in physics from Allegheny College in 1967. He received his M.S. in 1969 and his Ph.D. in 1973, both in physics, from Purdue University.

In August 1973, Dr. Parker joined the Professional Staff of the Raytheon Research Division, Lexington, MA. Initially, his work was primarily related to the development of improved temperature-stable surface acoustic wave materials. From 1977, Dr. Parker was responsible for the development of high performance surface acoustic wave (SAW) oscillator technology at the Research Division, including the "All Quartz Package" for SAW devices. His primary interest was frequency stability, with an emphasis on $1/f$ noise, vibration sensitivity, and long-term frequency stability. In June 1994 Dr. Parker joined the Time and Frequency Division of the National Institute of Standards and Technology in Boulder, CO. He is the group leader for the Time Scale and Coordination Group; his interests include improved time scales and time transfer technology.

Dr. Parker is a Fellow of the IEEE, and a member of Sigma Xi and Sigma Pi Sigma. He has served as an elected member of the Administrative Committee of the IEEE Ultrasonics, Ferroelectrics, and Frequency Control Society (1988–1990) and is currently Chair of the Frequency Control Standing Committee of the UFFC-S. He has served on the Technical Program Committees of both the Ultrasonics and the Frequency Control Symposia, and was the Technical Program Chair for the Frequency Control Symposium in 1990 and 1991. Dr. Parker is currently the Associate Editor for Frequency Control—Acoustics of the UFFC-S Transactions.

He received the 1988 Outstanding Transactions Paper Award from the IEEE Ultrasonics, Ferroelectrics, and Frequency Control Society as a co-author of two papers which appeared in the May 1988 and November 1988 issues of the Transactions. Dr. Parker was the recipient of a Thomas L. Phillips "Excelence in Technology Award" from Raytheon in 1992. In 1994 he received the W. G. Cady Award presented by the IEEE International Frequency Control Symposium. Dr. Parker is the 1996 Distinguished Lecturer of the IEEE Ultrasonics, Ferroelectrics, and Frequency Control Society.

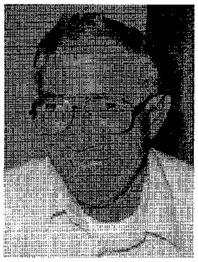

Judah Levine was born in New York City in 1940. He received a Ph.D. degree in Physics from New York University in 1966. He is currently a physicist in the Time and Frequency Division of the National Institute of Standards and Technology (NIST) in Boulder, Colorado, and is also a Fellow of the Joint Institute for Laboratory Astrophysics, which is operated jointly by NIST and the University of Colorado at Boulder.

He has worked on time scale algorithms at NIST for about 25 years. His most recent work involves developing the new time-scale algorithm described in this paper. He is also working on methods for distributing time and frequency information using digital networks such as the Internet, and on ways of improving satellite-based time and frequency distribution.

Dr. Levine is a Fellow of the American Physical Society and is a member of the American Geophysical Union and the IEEE Computer Society.