

1999 Joint Meeting EFTF - IEEE IFCS

HYDROGEN MASER ENSEMBLE PERFORMANCE AND CHARACTERIZATION OF FREQUENCY STANDARDS

Thomas E. Parker

National Institute of Standards and Technology
Time and Frequency Division
325 Broadway, Boulder, CO 80303, USA

ABSTRACT

The performance of a post-processed clock ensemble dominated by active, cavity-tuned, hydrogen masers has been evaluated. The ensemble has an average frequency drift of less than $\pm 3 \times 10^{-15}$ per year, and an Allan deviation less than 1×10^{-15} for τ between 4 hours and 100 days. This ensemble has been used to help characterize the frequency stability of a number of commercially available, high performance frequency standards.

1. INTRODUCTION

A complete evaluation of the frequency stability of commercial, high-performance cesium beam and hydrogen maser frequency standards is very difficult to perform because there is no single standard consistently available that is better in all aspects of stability. Therefore, the National Institute of Standards and Technology (NIST) has developed a post-processed ensemble of high performance standards that is intended to serve as a stable frequency reference in both short and long term.

The ensemble, referred to as AT1E, consists of five active, cavity-tuned, hydrogen masers [1] and four high-performance cesium beam standards [2]. The AT1 algorithm [3], [4] is used in a post-processed mode, and the maximum weight of any standard in the scale is limited to $\sim 30\%$. We also make use of modeled frequency drift parameters to handle the relatively high drift rates of the hydrogen masers. Post-processing helps in determining the drift parameters and in handling unusual events, such as equipment failures or environmental disruptions. We have not yet found it necessary to model the frequency drift of the cesium standards. Though four cesium standards are present in the ensemble, their total weight, both short and long term, is less than 5%. Thus they play a relatively small role in the performance of the scale. AT1E has been calculated since April 27, 1997 (MJD 50565), which is the earliest point in time from which at least four masers have always been available. The goal for AT1E is to provide continuity and to be comparable to, or more stable than any of the individual standards in the

U. S. Government work not protected by U. S. copyright

ensemble. Its purpose is to serve as a convenient frequency reference for evaluating the stability of various frequency standards such as commercial devices, or even the short and medium-term stability of primary standards. In Section 2 we discuss the performance of the ensemble, and in Section 3 we use the ensemble to help characterize the frequency stability of the high-performance, commercial frequency standards that are in the ensemble.

2. ENSEMBLE PERFORMANCE

The long-term stability of AT1E is evaluated with data from primary frequency standards such as CS1 and CS2 at the Physikalisch-Technische Bundesanstalt (PTB) and NIST-7. AT1E is not steered directly, but uses information from primary standards to model the long-term frequency drift of the hydrogen masers. Figure 1 shows the fractional frequency difference of AT1E relative to three primary standards for the entire period of its calculation. Data for CS1 and CS2 were obtained from publications of the Bureau International des Poids et Mesures (BIPM), including Circular T and the Annual Report of the BIPM Time Section. All of the primary standards have uncertainties near 1×10^{-14} . For convenience the initial frequency offset of AT1E was arbitrarily set to 0.

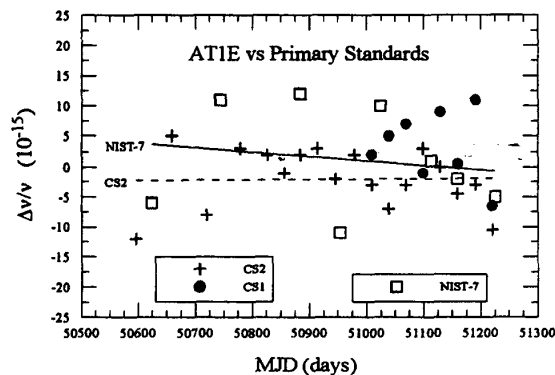


Figure 1. Frequency of AT1E relative to primary frequency standards.

Our primary interest is in the frequency stability of AT1E. The solid line indicates the average frequency drift of AT1E relative to NIST-7, and its slope is $-2.6(\pm 6) \times 10^{-15}$ per year. The slope was obtained by the method of linear regression. The dashed line indicates the average frequency drift of AT1E relative to CS2, which is $+0.2(\pm 2) \times 10^{-15}$ per year. There is not yet enough CS1 data from Circular T to make a meaningful estimate of frequency drift. The conclusion drawn from the data in Fig. 1 is that the drift modeling has been successful with the average frequency drift of AT1E being less than $\pm 3 \times 10^{-15}$ per year. (The lowest overall individual maser drift rate is $+1.2 \times 10^{-14}$ per year.)

The stability of the primary frequency standard data is not good enough to provide meaningful information on the stability of AT1E for time intervals less than about one year. Changes in maser drift rates that occur over intervals shorter than one year can only be estimated by comparison to the ensemble itself. This requires some care since the masers being evaluated are also in the ensemble. However, since AT1E is a post-processed scale the modeled drift rates can be iterated to provide an overall low drift rate. More accurate evaluation of short-term changes in maser drift rates will become possible when data from cesium fountain standards is available at regular intervals.

Figure 2 shows Allan deviation plots as calculated by two different methods. The curves with triangles and diamonds are for AT1E and were obtained from three-corner-hat calculations using TA(USNO) and EAL, and TA(USNO) and TA(AMC), respectively. EAL is a free atomic scale calculated by the BIPM, TA(USNO) is the independent local atomic time scale calculated at the U.S. Naval Observatory in Washington DC, USA, and TA(AMC) is the independent local atomic time scale calculated at the Alternate Master Clock in Colorado Springs, CO, USA. Data for EAL, TA(USNO), and TA(AMC) were obtained from Circular T and the Annual Report of the BIPM Time Section. The solid circle curve is a comparison of maser 2 with AT1E

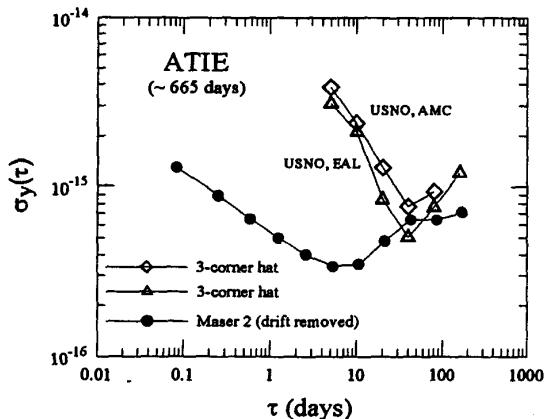


Figure 2. Allan deviation of AT1E.

using internal NIST measurements. The short-term stability (τ less than ~ 40 days) of the three-corner-hat curves is dominated by common-view GPS noise [5] and therefore does not represent the true performance of AT1E.

A better estimate for the short-term stability of AT1E is obtained from the curve with solid circles. The weight of maser 2 in AT1E is approximately 30 %, which means that the actual ensemble noise is about 65 % of that indicated by the solid circle curve [6]. The data for this curve has not been corrected for this correlation since the short and long-term weights of maser 2 are not the same, and furthermore, the weights are not constant with time. Therefore, the curve should be considered only an approximation of the true behavior of AT1E since a precise correction would be almost impossible to perform. Frequency drift was removed for the solid circle curve because of the relatively high drift rate of maser 2. No frequency drift was removed for the three-corner-hat curves. (The average drift of AT1E relative to EAL is only -4.3×10^{-15} per year.)

There is risk involved in evaluating the stability of an ensemble by comparing it to one of the individual standards in the ensemble, as in the solid circle curve. Some correlation exists between the scale and the standard, as determined by the weight of the standard. This can be approximately accounted for as pointed out above. However, there is also the risk of other unknown common-mode phenomena causing all of the standard frequencies to behave in a similar fashion. Although we have put considerable effort into identifying and eliminating such potential common-mode problems [7], it is always useful to make comparisons to remote, and independent frequency references. That is why the three-corner hat curves are shown. Unfortunately, the noise inherent in a GPS common-view comparison dominates in the short-term as seen in Fig. 2. Therefore, a truly independent evaluation of the short-term stability of AT1E cannot be made until better frequency comparison technology, such as carrier phase GPS, is available. However, the Allan deviation for the three curves is in good agreement for τ longer than about 40 days, indicating that it is a reasonable measure of the long-term stability of AT1E. Table 1 summarizes the estimated performance of AT1E.

Table 1 AT1E Performance

$\sigma_y(\tau = 0.1 \text{ day})$	$\sim 1 \times 10^{-15}$
$\sigma_y(\tau = 1 \text{ day})$	$\sim 4 \times 10^{-16}$
$\sigma_y(\tau = 10 \text{ days})$	$\sim 2.5 \times 10^{-16}$
$\sigma_y(\tau = 100 \text{ days})$	$\sim 8 \times 10^{-16}$
Frequency Drift	$\leq \pm 3 \times 10^{-15}$ per year

3. FREQUENCY STANDARD PERFORMANCE

There are five, active, cavity-tuned hydrogen masers at the NIST facility in Boulder, CO, USA, and all are contained in individual environmental chambers to control temperature and humidity [7]. Figure 3 shows the long-term fractional frequency drift of the five hydrogen masers relative to AT1E for the full 695 days that AT1E has been calculated. Initial frequency values have been arbitrarily offset for viewing convenience. In fact, the maser frequencies are distributed over a range of 8×10^{-12} . The frequency drift clearly differs from maser to maser, and the newer ones generally have higher drift rates. (The masers are numbered in the order in which they were purchased.) There also is a tendency for the larger drift rates to decrease with time. The slopes of the curves in Fig. 3 range from $+3.2 \times 10^{-13}$ per year (early in the curve for maser 3) to -3.3×10^{-14} per year for maser 4.

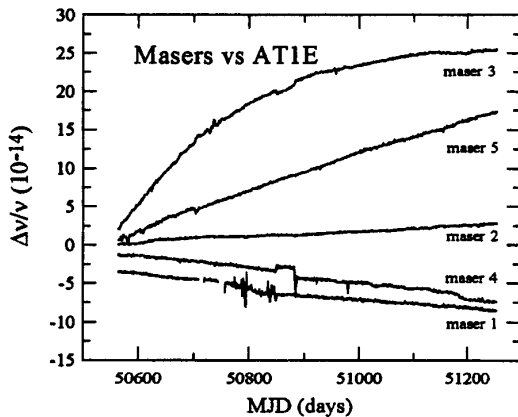


Figure 3. Frequency drift of hydrogen masers.

Nearly all of the transient structure observed in the curves can be identified with environmental testing, failures of environmental chambers, or the moving of masers from one chamber to another. The gaps in the maser 1 data between MJD 50700 and MJD 50800 occurred when the maser was out of service. Maser 2 has been virtually undisturbed for the full time and shows the best overall performance.

The most interesting unexplained event in Fig. 3 started on MJD 50850 when masers 3 and 4 both exhibited frequency steps at the same time. The frequency of maser 4 went up 8×10^{-15} and that of maser 3 went down by 5×10^{-15} . The monitored environmental data (temperature, relative humidity, barometric pressure, power line voltage, and vertical magnetic field strength) showed no obvious cause. Thirty four days later both masers exhibited opposite frequency steps when the maser 4 environmental chamber failed. Note that the two masers are in different chambers and none of the other masers showed any unexplained behavior.

At this time we still have no satisfactory explanation for the occurrence.

Figure 4 shows TOTAL deviation (interpreted the same as the Allan deviation) plots for three representative masers compared to AT1E. Allan deviation curves for the other two masers fall within the range of the curves in Fig. 4. Upward corrections to the data ranging from 10 % to 20 % should in principle be made since these masers are all in AT1E. However, as mentioned earlier, these corrections are difficult to make precisely, so no corrections have been applied. A mean second-difference frequency drift has been removed in all cases. For clarity confidence limits are only shown on one curve.

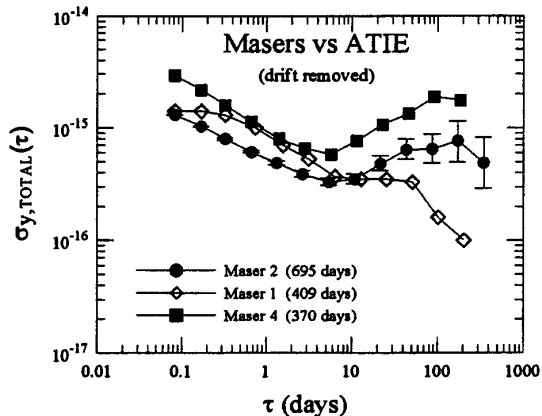


Figure 4. Allan (TOTAL) deviation of hydrogen masers.

Maser 2 has been undisturbed for the full period that AT1E has been calculated so the data in Fig. 4 is for a period of 695 days. For masers 1 and 4 the data covers only the most recent 409 and 370 days respectively. The stability curves vary significantly from maser to maser, not only in level, but also in shape. However, all of the masers show frequency stabilities that drop into the 10^{-16} range between one and ten days. All of the masers except maser 1 show a rise in the TOTAL deviation beyond about ten days. This rise appears in most cases to be caused by changes in the frequency drift rate, although occasionally it can be traced to environmental stability problems. The drop beyond 100 days is caused by the frequency drift removal process.

Figure 5 shows TOTAL deviation plots for two of the cesium standards. Again the standards are numbered in the order in which they were purchased. No frequency drift has been removed for these plots, and no correlation correction is needed because of the low weight of the clocks. The data in Fig. 5 show that the noise of the cesium standards is white frequency out to about 80 days where the level is in the low to mid 10^{-15} range. Beyond 80 days the level either flickers out or begins to rise. The Allan deviation curves of cesium standards 2 and 3 fall between the two curves in Fig. 5.

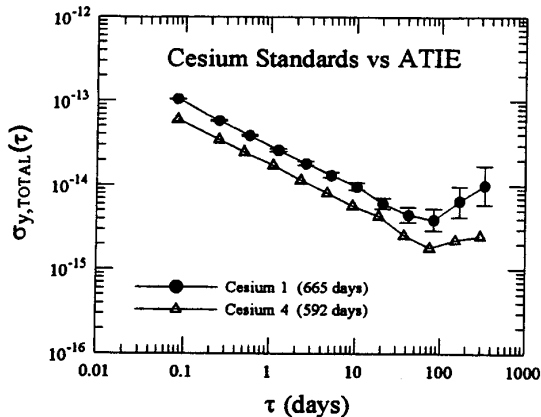


Figure 5. Allan (TOTAL) deviation of cesium frequency standards.

Frequency drift rates of the cesium standards relative to ATIE are generally low and are summarized in Table 2, along with the maser drift rates. The drift rates for cesium standards 2 and 3 were calculated for the entire interval that ATIE has been calculated, while for cesium standards 1 and 4 the period is slightly less. Cesium 4 is the newest device and has not been in operation for the full length of ATIE operation. Cesium standard 1, on the other hand, is the oldest device and it ran out of cesium in March, 1999. It became very noisy during its last 25 days, and the frequency of cesium 1 also exhibited a slight upward trend for the last 100 days before it got noisy. This has resulted in it having the highest overall cesium drift rate.

Table 2. Frequency Drift Rates Relative to ATIE

Standard Number	Freq. Drift Rate
Maser 1	-2.0×10^{-14} per year
Maser 2	$+1.6 \times 10^{-14}$ per year
Maser 3	$+2.8 \times 10^{-14}$ per year
Maser 4	-3.3×10^{-14} per year
Maser 5	$+7.5 \times 10^{-14}$ per year
Cesium 1	$+1.6 \times 10^{-14}$ per year
Cesium 2	$+5.8 \times 10^{-15}$ per year
Cesium 3	-4.4×10^{-15} per year
Cesium 4	$+2.9 \times 10^{-15}$ per year

The maser drift rates in Table 2 apply only to the most recent 300 days. They are generally higher than those of the cesium standards, with only maser 2 having a drift rate equal to the worst cesium.

4. SUMMARY

A hydrogen maser ensemble has demonstrated a frequency drift rate less than $\pm 3 \times 10^{-15}$ per year, and a minimum Allan deviation near 2×10^{-16} at $\tau \approx 10$ days. This ensemble has proven to be very useful in evaluating the frequency stability characteristics of the best commercial frequency standards and also in the evaluation of stability characteristics of primary frequency standards [8].

5. ACKNOWLEDGEMENTS

The author acknowledges the valuable assistance of Jim Gray and Trudi Pepler in maintaining the frequency standards in the ensemble and Judah Levine for creating the computer program for ATIE.

6. REFERENCES

- [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.
- [2] Hewlett-Packard model HP5071A with high performance tube.
- [3] P. Tavella and C. Thomas, "Comparative Study of Time Scale Algorithms," *Metrologia*, vol. 28, pp. 57-63, 1991.
- [4] M. Weiss and T. Weissert, "AT2, A new Time Scale Algorithm: AT1 Plus Frequency Variance," *Metrologia*, vol. 28, pp. 65-74, 1991.
- [5] T.E. Parker, D.A. Howe and M. Weiss, "Accurate Frequency Comparisons at the 1×10^{-15} Level," in Proc. 1998 IEEE International. Freq. Control Symp., pp. 265-272, 1998.
- [6] P. Tavella and C. Thomas, "Study of the Clock-Ensemble Correlation in ALGOS Using Real Data," in Proc. 5th European Frequency and Time Forum, pp. 435-441, 1991.
- [7] T.E. Parker, "Influence of Environmental Factors on Hydrogen Maser Frequency Stability," in Proc. 12th European Frequency and Time Forum, pp. 126-131, 1998.
- [8] S.R. Jefferts, D.M. Meekhof, F. Levi, W.D. Lee, F.L. Walls, L.W. Hollberg and T.E. Parker, "NIST Cesium Fountain Preliminary Accuracy Evaluation," Joint Meeting of the 13th European Frequency and Time Forum and 1999 IEEE International Freq. Control Symp. (in this Proceedings).