Drift Correction for Active Hydrogen MASERs

Samuel J. Griffin
Frequency and Time Systems
Microchip Technology Inc.
Tuscaloosa, Alabama, USA

Dylan Meyer
Dept. of ECE
The University of Alabama
Tuscaloosa, Alabama, USA

Andrew Lemmon
Dept. of ECE
The University of Alabama
Tuscaloosa, Alabama, USA

H. Bryan Owings
Frequency and Time Systems
Microchip Technology Inc.
Tuscaloosa, Alabama, USA

***Abstract*—The drift characteristics of well-aged Hydrogen MASERs may be reasonably modeled as linear over periods greater than two weeks. There are several broad strategies used to manage drift: a) quantify and acknowledge drift in a system error budget, b) compensate using post-processing methods, c) apply a linear correction, or d) steer the MASER output frequency to a superior reference. The selection of drift management method greatly depends upon the requirements of a user's system. Users, such as radio astronomers, may de-emphasize active drift management as their primary concern is 1-to-100 second stability. These users often prefer to compensate for frequency drift after passing a predefined frequency offset threshold. This removal takes the form of a manual frequency adjustment on the MASER. However, time-scale users, who require exact characterization and compensation of drift, will need to intervene using post-processing methods to avoid tampering with a characterized MASER which may be contributing to local UTC. This paper will discuss the underlying processes that may cause linear drift, how this drift is quantified, and then give an overview of the design implications of each drift handling method. For illustration purposes, a MASER's drift will be characterized, and a linear drift correction will be applied. This illustration covers the use case when a MASER will perform as a stand-alone standard with moderate long-term stability requirements.**

***Index Terms*—Frequency Control Symposium, MASER, Frequency Standard, Drift Correction, Frequency Drift.**

I. Introduction

For frequency standards, the reduction of systemic noise is of utmost importance. Incremental to large improvements have dramatic impacts on the applications of precise time, such as navigation, scientific measurements, and communications to name a few. Consequently, continued scrutiny of frequency standard design and application is warranted.

The Active Hydrogen **M**icrowave **A**mplification by **S**timulated **E**mission of **R**adiation (MASER [1]) frequency standard is a common device in critical timing systems. Consistent with all frequency standards, Active Hydrogen MASERs have drift. Classically, this drift is modeled as linear. For practically any application, this drift will need to be managed. This drift impacts the long-term stability of the MASER.

This paper will cover the following topics in each section. Section II will discuss what is unique about Hydrogen MASERs amongst frequency standards, and how their signals are characterized. Then Section III will introduce and consider the different methods that may be used to correct frequency drift. Section IV details an example drift correction procedure.

II. Background

In service of discussions on MASER drift correction, a review of challenges faced by users and designers of frequency standards will be covered. These issues will then be highlighted for the Hydrogen MASER to provide understanding of the source and behavior of drift to be counteracted.

A. Types of non-ideal behaviors

When discussing the performance of any frequency standard, any determinations are based on the devices frequency stability (or instability). Sources of this instability arise from internal and external sources. External noise coupling into the MASER system is unavoidable, and a bulk of the systems within frequency standards aim to limit and/or compensate for these sources of instability. Time, temperature, pressure, humidity, gravity, magnetic fields- virtually any external force or event is a potential noise source. Moving toward the inner-workings of frequency standards, the interaction of external noise with the device, as well as inherent noise sources within the device are denoted as system noise. These internal instabilities may include component aging, baseline thermal noise, and shot noise. Noise also comes from active frequency compensation methods. While these methods are intended to reduce noise, they contribute to the system's noise floor either from internal control schemes or from user commands. These compensation methods include but are not limited to thermal controls, magnetic controls, and pressure controls.

Along with the discussion of external and internal origins of noise, one must consider the characteristics of the noise itself. The classic dichotomy of deterministic and stochastic noise influences the understanding of a noise source. Deterministic noise, as it is predictable, can be handled in a variety of ways. Random/stochastic noise has fewer methods of reduction, and generally contributes to the noise floor of a device. This noise floor sets an absolute limit on the improvements possible for a device's frequency stability.

As there is quite some ground to cover while analyzing a system for noise issues, some strategies must be applied to help lead any investigation. When discussing the sources of instability for any frequency standard, it is useful to define what type of system topology is being used. Does the system produce a reference signal, or does it produce an error signal when given an input signal? Defining which system topology is in use will direct discussion of the device and its uses. Many commercial atomic frequency standards, such as cesium beam

978-1-6654-3935-0/21/$31.00 ©2021 IEEE

standards, are referred to as passive standards. These standards utilize atomic transitions as a bandpass filter combined with a feedback loop to discipline the output of a quartz oscillator via an error signal. The quartz oscillator is steered to maximize the strength of the filtered signal via the error signal. In other words, when the atomic beam pulses on-frequency, the desired state of maximum current is observed [2]. Active frequency standards produce a signal from the resonance of a population of atoms, which is used as a reference [2]. Some portion of this self-sustaining signal is coupled off into a receiver system which ultimately supplies an output for use.

B. The Hydrogen MASER

The Active Hydrogen MASER maintains wide usage in precise time. Projects aimed at creating Hydrogen MASERs were intended to produce frequency standards superior to the existing commercial atomic clocks. Improvements in timing capabilities were necessitated by the proliferation of space navigation, as well as improvements in radio astronomy apparatuses and a desire to improve metrology as a practice [3]. These fields reaped the rewards of an improved frequency standard, and many still utilize improved versions of the Active Hydrogen MASER standard today [4].

The Active Hydrogen MASER is an 'active' standard because it uses hydrogen atoms to generate a reference signal. This signal is then conditioned with one or more phase-locked-loops and amplifiers and subsequently provided as an output.

Since the MASER is an active standard, the focus of any investigation of frequency instability centers on effects that influence the behavior of the hydrogen atoms, as well as the supporting electronics that interface with the signal generated by the atoms. Since the physics related to the resonant frequencies of the hydrogen atoms and the quality of the signal processing paths are the main factors determining frequency instability in a system, these are the main areas that can be improved to enhance frequency standard performance, and reduce undesired aspects such as drift.

C. Investigations of Hydrogen MASER Drift

Hydrogen MASER drift is a topic that has been investigated by several groups. These studies have varied in scope and focus. Strategies and approaches have required compilation of data from periods ranging from weeks to years. These studies typically focus on subjects such as physical causes of drift, or characterizing frequency stability over short, medium, or long periods of time. Some data sets, largely focused on observation times in the order of a month or two, have seen MASER average daily drift fit a linear model [5]. Other sets, spanning several months to several years, see average daily drift change in value one or more times over the observation period. These cases occasionally need higher order curves to accurately model the data over the entire scope of the observation [6] [7]. Environmental changes, such as failing environmental chambers or maintenance, have been postulated to induce these changes in drift rate [8]. These drift changes are causal in nature, and require the drift rate to be estimated again.

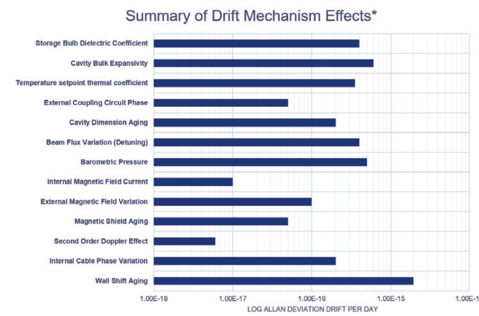

Fig. 1. Systematic Sources of Frequency Drift in Hydrogen MASERs

In Figure 1 adapted from [7] for the Microchip Active Hydrogen MASER (MHM-2020), the relative frequency instability contributions for various systemic effects are estimated. Significant sources of drift include effects such as wall shift, component aging, and cavity pulling. Wall shift is one significant source of frequency drift. Wall shift refers to the alteration in output frequency over time due to the changing nature of hydrogen atom collisions with the storage bulb wall. The walls are coated with polytetrafluoroethylene, designed to produce as near-elastic collisions with the hydrogen atoms as possible. This coating ages over time, affecting the energy transfer present in wall collisions, thus changing the output frequency of the hydrogen population [9]. Component aging is also a contributor of drift in hydrogen MASERs. Aging can result in shrinking cavity components which affect the resonant frequency of the cavity, aging of resistors controlling reference voltages which control important functions in the signal path, and aging of thermistors used to control cavity temperature causing cavity pulling [10]. Cavity pulling is the dominant noise source that has some adjust-ability in the MASER's design, as well as operationally by the users. The frequency stability of the hydrogen MASER is sensitive to changes in cavity volume in the order of one Ångstrom (1×10^{-10} m). As a result, any factors that change the volume of the cavity affect the output signal of the hydrogen MASER [10]. Significant contributions to cavity volume changes are external temperature changes, which can occur seasonally, as well as the aging of the cavity dimensions over time [10]. As stated previously, each of these phenomena affect the frequency stability of a hydrogen MASER in a quantifiable manner.

A common way of defining the frequency stability of a standard is Allan Deviation (ADEV). MASERs exhibit exceptional short and medium-term stability in relation to other commercial atomic clocks, while exhibiting greater drift in the long-term [11]. Short term Hydrogen MASER operation results in ADEV values on the order of 1×10^{-14} [9]. This occurs in the warm-up region where the MASER begins to oscillate but has not reached steady-state operation [12]. Generally, the MASER short-term performance is dominated by its local oscillator that is disciplined to the hydrogen

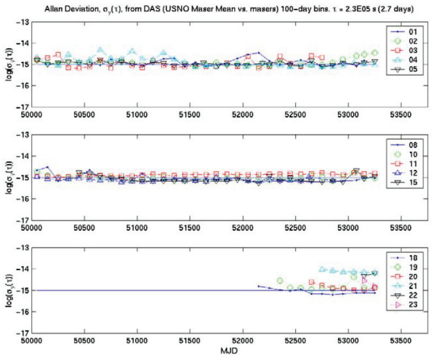

Fig. 2. Long-term ADEV measurements in 100 day time averages of USNO MASERs. Key numbers represent unique MASERs [7].

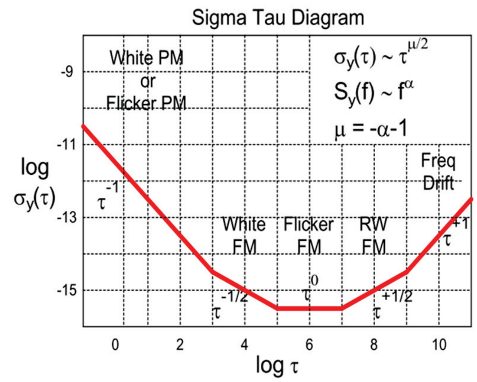

Fig. 3. ADEV graph demonstrating the characteristic curves of Power Law Noise models [16].

frequency [9]. While this performance is not cutting-edge, it does out-perform comparable passive standards [13]. As the MASER continues to operate, the stability of the output improves. In the medium term, the MASER reaches an ADEV flicker floor in the order of magnitude of 1×10^{-16} [14]. Long term ADEV values generally are on the order of 1×10^{-15} as illustrated in Figure 2 [7]. This is generally the result of the onset of deterministic frequency drift. This drift can be caused by a multitude of factors, including the previously mentioned wall shift, component aging, and cavity pulling [9] [10]. Drift correction methods will have the largest impact in this long-term measurement range.

D. Estimating Drift With Modeling

Before long term drift can be corrected, it must be quantified. Multiple models exist for frequency drift. The standard model for drift is a linear function applied to frequency data. Alternatively, a quadratic model may be used to estimate drift with a phase dataset [15].

Applying a model to a dataset involves regression. As part of this regression process, a comparison between the regression fit and the original dataset must be made by analyzing the residuals of the fit. Any non-Gaussian patterns left in the residuals indicate that the model is not fitting the data well. This is where looking at an ADEV plot is useful, and applying Power Law noise models can help identify additional noise contributions that are not imposed by linear frequency drift. See Figure 3 for the traditional representation of Power Law noise. Identifying the noise contaminants can lead to their reduction or elimination through analytic means or the reconfiguring of the testing environment to eliminate the noise source so a new dataset can be obtained. This allows for regression modeling to produce a good frequency drift estimate.

Another topic of great importance to frequency drift estimation is confidence intervals. After performing one regression on a dataset, an estimate of frequency drift is obtained. Performing a second regression on a newer set of data from the same source produces a different estimate, which is likely to be quite different from the first. Compounding the issue, using multiple models on the same dataset can provide widely varying results. This is where the concept of Chi-Squared distributions can help to resolve this issue. With the first regression, the variance on the initial estimate is extremely high (therefore the estimate confidence is low). To reduce this variance on the estimate, more estimates need to be performed. Each estimate will increase the confidence of the estimate, reducing the variance [17]. It is highly encouraged to perform scheduled drift estimations in a constant interval. In practice, one estimate can be performed over at least a 30 day period to capture frequency drift.

III. General Frequency Standard Strategies

Since Hydrogen MASER drift is well documented and the applications of MASERs are widespread and diverse; many strategies for addressing MASER drift have been adopted. For the use of this paper, one method will be referred to as post-processing. This refers to the use of data pertaining to MASER drift after the MASER was utilized to enhance calculations and measurement certainty. Post-processing is employed in radio astronomy and paper-based time scales. Another is simple regression, or linear drift correction. This approach is used in standards such as the MHM-2020. A third is adaptive correction, or periodic recalculation of regression parameters or state variables. These regression parameters or state variables are then used to correct the output of the MASER. This approach is used in state-of-the-art timescales, with both post-processed and real-time physical realizations.

1) Post-Processing

Post-processing may be used for several reasons, from the need for a frequency standard to be free-running or the desire to minimize disturbances in frequency stability [18]. Hydrogen MASERs are widely used as local oscillators or frequency references for radio astronomy observatories and arrays. One array uses a post-processing technique to better characterize the Hydrogen MASER local oscillator. This allows for increased confidence in measurement data taken by the telescope [4]. Observations conducted by the array are relatively infrequent, creating periods long enough for frequency drift to be present in the output of a Hydrogen MASER. To compensate for this change in the reference frequency, the array monitors the short and long-term health and status of each MASER. For short-term monitoring, a high-precision quartz oscillator is used to measure the performance of the MASER in a 1 second time interval. This is to give an estimate of short-term MASER performance. For long-term monitoring, a 1 PPS signal from a GPS receiver is compared to a 1 PPS signal derived from the MASER output. This is used to observe frequency drift over long periods of time, as the GPS receiver serves as a standard of superior stability in long measurement intervals. This drift rate is then used to analyze the health of the MASER while also allowing data captured over the observation period to be more accurately assessed. This increases confidence in measurements taken during the observation interval [4].

2) Simple Regression

Drift can be characterized or modelled several different ways for different oscillators. Specifically, linear drift correction consists of analyzing drift data using linear regression. The drift rate can be characterized by a linear slope describing drift per unit time, such as drift per day, or drift per month. For appropriate sampling periods, the drift present in Hydrogen MASERs can be satisfactorily modelled as linear [5]. Once this regression is satisfactorily characterized, an incremental correction can be imposed to counteract the average drift. There are several ways to impose these corrections, such as the use of an Auxiliary Output Generator (AOG) or digital synthesizers internal to the frequency standard. This method will be discussed more in depth via a practical example of linear drift correction in the MHM-2020 Hydrogen MASER in Section IV.

3) Adaptive Regression

One method of correcting drift in a frequency standard is to alter the output, or steer, a frequency standard to minimize the deviation from a superior reference or clock ensemble [19]. This can be done via alteration of the frequency standards or outside equipment that implements the desired change, such as an AOG or a phase microstepper. This is commonly seen in clock ensembles and timing labs, where a large number of frequency standards are used to create a stable and accurate reference while a single clock is used as a physical realization of the ensemble [20]. An example of this is the steering of a Hydrogen MASER master clock output through an AOG. This is done by taking the time differences between the mean timescale, generated by ensemble clocks, and the master clock. Following these measurements, these data points are used as state variables in an algorithm that adjusts the output of the master clock via an AOG. The use of the AOG allows the Hydrogen MASER to maintain a high level of frequency stability while correcting the output so it strictly adheres to the more accurate clock ensemble. This also allows for the direct output of the MASER to contribute to Universal Coordinated Time (UTC), adding several more clocks to the international timescale. [21]. Examples of these algorithms vary from classical control, Kalman filtering, or other algorithms used to weigh and measure inputs and outputs of a multi-variate system while imposing corrections. [19]. Due to the relative complexity of these strategies, these algorithms allow several different sources of error to be addressed, increasing accuracy.

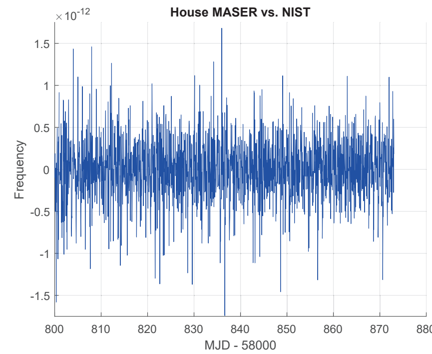

Fig. 4. Qualification of Reference

IV. Practical Drift Correction Example: MHM-2020

The Microchip Active Hydrogen MASER MHM-2020 is the latest iteration of one of the original commercial Hydrogen MASER designs. Up to this point, linear frequency drift has been externally compensated either by user command or by a secondary system on its output. With the latest design, the system can schedule a series of updates to the MASER's phase-locked loop to counteract linear drift. In order to prevent these corrections from appreciably effecting the short-term performance of the MASER, these corrections are made using the smallest step achievable by a synthesizer within the phase-locked loop and are spread across a 24 hour period.

Before this feature can be taken advantage of, an estimate of the drift rate must be obtained. Then, the MASER must be configured with a correction rate. After a period of data collection, the MASER's frequency data must be evaluated again to check the drift correction's effect.

A. Establishing a Reference

As with any analysis, it is important to ensure that a reference is performing well before using it to qualify another device. In this example, the reference is the 'House' MASER at the Microchip timing lab in Tuscaloosa, AL. The House MASER is regularly compared to UTC(NIST) using GPS common-view. See Figure 4 for the raw frequency differences between the House MASER and NIST. As this is a relatively noisy measurement comparison, the important evidence from the comparison is a lack of perceivable structure to the plot, and the House MASER is at least performing on the order of $1{\times}10^{-12}$. The House MASER is more likely performing much better than this based on historical data, less than a magnitude of 5×10^{-16}. With this comparison, there is high confidence in the accuracy of comparisons using the House MASER as a reference.

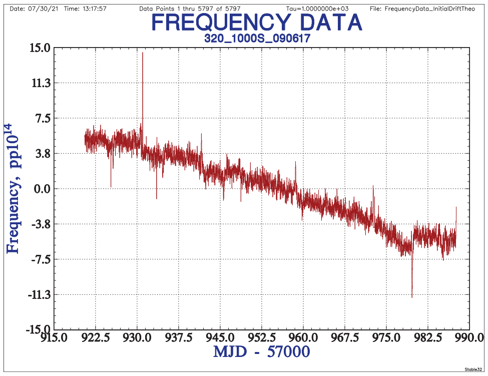

Fig. 5. Uncorrected MASER Frequency Data

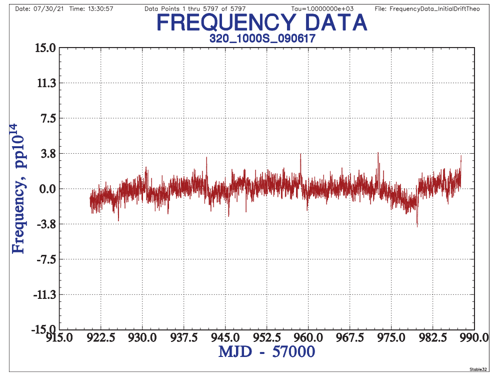

Fig. 7. MASER Frequency Data with Drift Removed Analytically

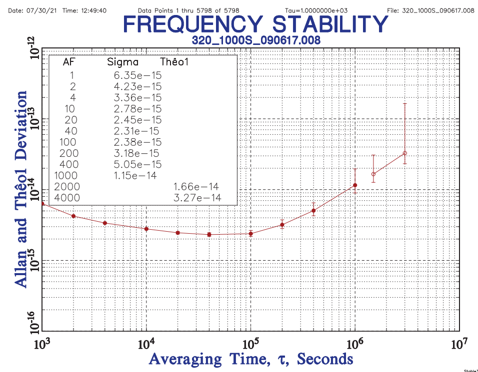

Fig. 6. Uncorrected MASER TheôH plot

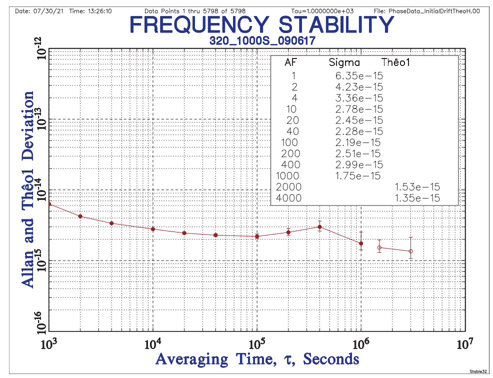

Fig. 8. MASER TheôH plot with Drift Removed Analytically

B. Baseline MASER Performance

Using the House MASER as a reference, a MASER under drift evaluation was measured. Approximately 70 days of ten-second phase measurements were collected for this device under test (DUT). This data was averaged to thousand-second data as shown in Figure 5, where a drift of $-1.943809 \times 10^{-15}$ per day was estimated. Figure 6 shows the overlapping ADEV and Theô1 frequency stability. In accordance with the earlier drift estimate, the Theô estimate shows a large decrease in stability which indicates a frequency drift.

C. Analytical Drift Removal

In order to set an expectation of what drift compensation will accomplish, the baseline frequency dataset is shown again in Figure 7 with drift analytically removed. The phase data was modeled using a quadratic function, resulting in a drift rate of $-2.249779 \times 10^{-17}$ per averaging interval.

The frequency stability plot in Figure 8 forms an expectation that the long-term stability should follow the trend of the preceding averaging intervals, improving slightly.

D. Real-Time Drift Compensation

For the real-time linear compensation of the DUT frequency drift, a correction value of 1×10^{-15} per day was chosen. This value was based on a priori knowledge of general MASER drift magnitudes. The drift estimate from Section IV-C was not used as it was much smaller than historical drift rates and a stark change in the frequency drift was desired. Over a relatively short evaluation period of approximately 80 days, the phase data of the DUT versus the House MASER was measured. The frequency conversion of this dataset is shown in Figure 9. Note a dramatic decrease in slope of the plot. Similarly, the overlapping ADEV and Theô1 frequency stability shown in Figure 10 display a marked improvement in comparison to Figure 6 for long averaging times. Additional drift estimates on new datasets captured from this DUT can be used to improve the accuracy of the correction value. Regardless, the drift was dramatically reduced, as illustrated by Figure 11, which projects the average value of the frequency data if the drift was not removed against the corrected dataset.

V. Conclusion

All frequency standards exhibit drift over their operational lifetime. In Hydrogen MASERs, this drift is particularly noticeable in long term measurement intervals. This can be attributed to systematic phenomena such as wall shift, component aging, and cavity pulling. Many applications utilizing Hydrogen MASERs necessitate a strategy to minimize this drift. These strategies vary from rudimentary methods such as post processing, to highly complex state-variable models that require a large amount of computation and instrumentation. However, for a wide variety of applications, drift can be modelled as a linear trend and corrected accordingly. Once the drift

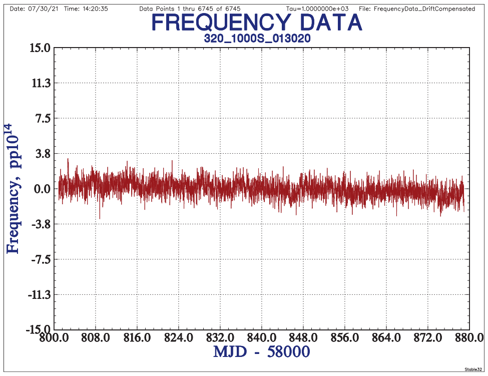

Fig. 9. Corrected MASER Frequency Drift Data

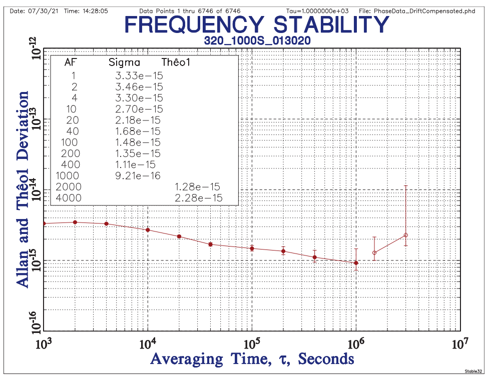

Fig. 10. Corrected MASER TheôH plot

value of the frequency standard is quantified with a sufficient degree of confidence, the drift can be compensated via a linear drift correction method. This can lead to improvements in long term stability by an order of magnitude.

ACKNOWLEDGMENT

The authors would like to thank Dr. David Howe, Dr. Richard Overstreet, and Dr. Robert Lutwak for their interest in this paper, and their discussions on the topics contained within.

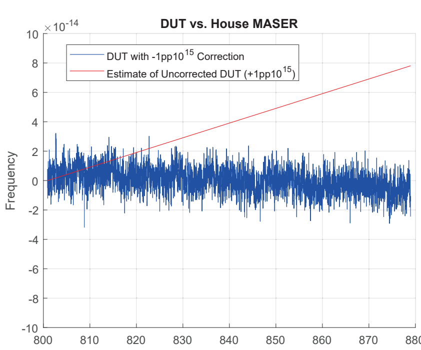

Fig. 11. MASER DUT Frequency Measurements with $-1pp10^{15}$/Day Correction

REFERENCES

[1] J. Vanier and C. Audoin, *The Quantum Physics of Atomic Frequency Standards: Volume 2*, A. E. Bailey, Ed. Philadelphia, PA: IOP PUBLISHING LTD, 1989.

[2] L. Lewis, "An Introduction to Frequency Standards," *Proceedings of the IEEE*, vol. 79, no. 7, pp. 927–935, 1991.

[3] H. Peters, T. McGunigal, and E. Johnson, "Hydrogen Standard Work at Goddard Space Flight Center," Goddard Space Flight Center, Tech. Rep., 1968.

[4] D. Ball, K. Akiyama, A. Alberdi *et al.*, "First M87 event horizon telescope results: II. array and instrumentation," *ASTROPHYSICAL JOURNAL LETTERS*, p. 28, apr 2019.

[5] H. Bryan 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 *Proceedings of the Annual Frequency Control Symposium*, 1992.

[6] T. E. Parker, S. R. Jefferts, and T. P. Heavner, "Medium-Term Frequency Stability of Hydrogen Masers as Measured by a Cesium Fountain," in *2010 IEEE International Frequency Control Symposium*, Newport Beach, CA, USA, 2010. [Online]. Available: https://ieeexplore.ieee.org/document/5556319

[7] D. Matsakis, P. Koppang, and R. M. Garvey, "THE LONG-TERM STABILITY OF THE U.S. NAVAL OBSERVATORY'S MASERS," in *36th Annual Precise Time and Time Interval (PTTI) Meeting*. Washington, D.C.: PTTI, 2004, pp. 1–11.

[8] L. A. Breakiron and D. N. Matsakis, "PERFORMANCE AND CHARACTERIZATION OF U.S. NAVAL OBSERVATORY CLOCKS," in *32nd Annual Precise Time and Time Interval (PTTI) Meeting*, 2000.

[9] H. Peters and H. Bryan Owings, "Hydrogen MASER Improvements and Future Applications," in *1997 IEEE INTERNATIONAL FREQUENCY CONTROL SYMPOSIUM*. IEEE, 1997, pp. 280–285.

[10] E. M. Mattison, "Physics of Systematic Frequency Variations in Hydrogen Masers," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 39, no. 2, pp. 250–255, 1992.

[11] L. A. Breakiron, "Timescale Algorithms Combining Cesium Clocks and Hydrogen Masers," in *Proc. of the 23rd Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, 3-5 December 1991*. Pasadena, California, USA: PTTI, 1991, pp. 297–305.

[12] J. Vanier, J. Gagnepain, W. Riley, F. Walls, and M. Granveaud, "Aging, warm-up time and retrace: important characteristics of standard frequency generators (Proposal for IEEE Standards Project P1193)," in *Proceedings of the 1992 IEEE Frequency Control Symposium*. IEEE, 1992, pp. 807–815. [Online]. Available: http://ieeexplore.ieee.org/document/269956/

[13] I. Diegel, "Hydrogen Maser Frequency Standard."

[14] T. E. Parker, J. E. Gray, and T. K. Peppler, "Environmental Sensitivities of Cavity-Tuned Hydrogen MASERs," in *Proceedings of the Annual IEEE International Frequency Control Symposium*, 1997, pp. 273–279.

[15] J. A. Barnes, "THE MEASURRMENT OF LINEAR FREQUENCY DRIFT IN OSCILLATORS," in *15th Annual PTTI Meeting*, vol. 4. PTTI, 1983, p. 32.

[16] W. J. Riley, *Handbook of Frequency Stability Analysis*. NIST, 1994, vol. 31, no. 1.

[17] D. Howe, D. Allan, and J. A. Barnes, "Properties of Signal Sources and Measurement Methods," in *Proceedings of the 35th Annual Symposium on Frequency Control*. Boulder, Colorado: NIST, may 1981, pp. 14–60.

[18] C. Thomas, P. Wolf, and P. Tavella, "Time Scales," BIPM; Queen Mary and Westfield College; Galileo Ferraris National Electrotechnical Institute, Tech. Rep., 1994. [Online]. Available: https://www.bipm.org/documents/20126/53814635/Monographie+94-1+%281994%29.pdf/c4069c93-155d-5565-7e3b-5bf5d4ff76c7

[19] M. Farina, L. Galleani, P. Tavella, and S. Bittanti, "A control theory approach to clock steering techniques," in *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 57, no. 10, oct 2010, pp. 2257–2270.

[20] T. D. Schmidt, C. Trainotti, and J. Furthner, "Comparing clock steering technique performances in simulations and measurements," in *Proceedings of the Annual Precise Time and Time Interval Systems and Applications Meeting, PTTI*, vol. 2019-Janua. Institute of Electrical and Electronics Engineers Inc., 2019, pp. 290–298.

[21] L. A. Breakiron and P. Koppang, "FREQUENCY STEERING OF HYDROGEN MASERS," United States Naval Observatory, Tech. Rep., 1996.