

A Clock Stabilization System for CHIME/FRB Outriggers

J. Mena-Parra¹ , C. Leung^{1,2} , S. Cary^{1,3} , K. W. Masui^{1,2} , J. F. Kaczmarek⁴ , M. Amiri⁵ , K. Bandura^{6,7} , P. J. Boyle^{8,9} , T. Cassanelli^{10,11} , J.-F. Cliche⁸ , M. Dobbs^{8,9} , V. M. Kaspi^{8,9} , T. L. Landecker¹² , A. Lanman^{8,9} , and J. L. Sievers^{9,13,14} 

CHIME/FRB Collaboration

¹ MIT Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA
jdmena@mit.edu

² Department of Physics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

³ Department of Astronomy, Wellesley College, 106 Central Street, Wellesley, MA 02481, USA

⁴ National Research Council Canada, Herzberg Astronomy and Astrophysics Research Centre, Dominion Radio Astrophysical Observatory, PO Box 248, Penticton, British Columbia, V2A 6J9 Canada

⁵ Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T 1Z1 Canada

⁶ Lane Department of Computer Science and Electrical Engineering, 1220 Evansdale Drive, PO Box 6109, Morgantown, WV 26506, USA

⁷ Center for Gravitational Waves and Cosmology, West Virginia University, Chestnut Ridge Research Building, Morgantown, WV 26505, USA

⁸ Department of Physics, McGill University, 3600 rue University, Montréal, QC H3A 2T8, Canada

⁹ McGill Space Institute, McGill University, 3550 rue University, Montréal, QC H3A 2A7, Canada

¹⁰ Dunlap Institute for Astronomy & Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada

¹¹ David A. Dunlap Department of Astronomy & Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada

¹² Dominion Radio Astrophysical Observatory, Herzberg Research Centre for Astronomy and Astrophysics, National Research Council Canada, PO Box 248, Penticton, BC V2A 6J9, Canada

¹³ Perimeter Institute for Theoretical Physics, 31 Caroline Street N, Waterloo, ON N2L 2YB, Canada

¹⁴ School of Chemistry and Physics, University of KwaZulu-Natal, 238 Mazisi Kunene Road, Glenwood, Durban, 4041, South Africa

Received 2021 October 1; accepted 2021 November 12; published 2022 January 7

Abstract

The Canadian Hydrogen Intensity Mapping Experiment (CHIME) has emerged as the prime telescope for detecting fast radio bursts (FRBs). CHIME/FRB Outriggers will be a dedicated very-long-baseline interferometry (VLBI) instrument consisting of outrigger telescopes at continental baselines working with CHIME and its specialized real-time transient-search backend (CHIME/FRB) to detect and localize FRBs with 50 mas precision. In this paper, we present a minimally invasive clock stabilization system that effectively transfers the CHIME digital backend reference clock from its original GPS-disciplined ovenized crystal oscillator to a passive hydrogen maser. This enables us to combine the long-term stability and absolute time tagging of the GPS clock with the short- and intermediate-term stability of the maser to reduce the clock timing errors between VLBI calibration observations. We validate the system with VLBI-style observations of Cygnus A over a 400 m baseline between CHIME and the CHIME Pathfinder, demonstrating agreement between sky-based and maser-based timing measurements at the 30 ps rms level on timescales ranging from one minute to up to nine days, and meeting the stability requirements for CHIME/FRB Outriggers. In addition, we present an alternate reference clock solution for outrigger stations that lack the infrastructure to support a passive hydrogen maser.

Unified Astronomy Thesaurus concepts: Radio astronomy (1338); Radio transient sources (2008); Radio pulsars (1353); Astronomical instrumentation (799); Very long baseline interferometry (1769)

1. Introduction

Fast radio bursts (FRBs; Lorimer et al. 2007) are transient pulses of radio light observed out to cosmological distances; both their origins and emission mechanisms remain unclear. Even though thousands of FRB events occur over the full sky every day (Bhandari et al. 2018; The CHIME/FRB Collaboration 2021), their detection with traditional radio telescopes is challenging due to the randomly occurring nature of the majority of bursts.

With its unique design optimized for rapid wide-field observations and a powerful real-time transient-search engine (CHIME/FRB; CHIME/FRB Collaboration et al. 2018), the

Canadian Hydrogen Intensity Mapping Experiment (CHIME¹⁵; CHIME Scientific Collaboration et al. 2021 (in preparation)) has become the leading facility for the detection of FRBs, detecting over 500 FRBs (The CHIME/FRB Collaboration 2021) and 18 new repeating sources (CHIME/FRB Collaboration et al. 2019a, 2019b; Fonseca et al. 2020) in its first year of full operation. Such an unprecedented sample of events with a single survey has enabled detailed studies of statistical properties of the FRB population, such as fluence distribution and sky rate, scattering time, dispersion measure distribution, spatial distribution, burst morphology, and correlations with a large-scale structure (Chawla et al. 2021; Josephy et al. 2021; Rafie-Ravandi et al. 2021; Pleunis et al. 2021; The CHIME/FRB Collaboration 2021).

However, except for FRBs with low dispersion measure (Michilli et al. 2021; Bhardwaj et al. 2021a, 2021b), CHIME/FRB's arcminute localization precision is insufficient to



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

¹⁵ <https://chime-experiment.ca>

localize these bursts to their host galaxies, which is crucial to understanding their nature and unlocking their potential as probes of the intergalactic medium and large-scale structure. To overcome this limitation, the CHIME/FRB collaboration is currently developing CHIME/FRB Outriggers, a program to deploy CHIME-like outrigger telescopes at continental baseline distances. CHIME and the outriggers will form a dedicated very-long-baseline interferometry (VLBI) network capable of detecting hundreds of FRBs each year with subarcsecond localization precision in near real time, allowing for the unique identification of FRB galaxy hosts and source environments.

Because VLBI localizes sources by precisely measuring the difference in the arrival time of astronomical signals between independent telescopes across far-separated sites, it is critical to use very stable local reference signals (i.e., clocks) that allow the synchronization of VLBI stations without losing coherence during observations and between calibrations. This is particularly important for stationary telescopes like CHIME and outrigger stations that can only be calibrated when a bright radio source transits through their field of view. The superior stability performance of hydrogen masers on short and intermediate timescales makes them the preferred option for VLBI applications (Matthews et al. 2018; Event Horizon Telescope Collaboration et al. 2019; Schmittberger & Scherer 2020). Here, we present a hardware and software clock stabilization solution for the CHIME telescope that effectively transfers the reference clock from its original GPS-disciplined crystal oscillator to a passive hydrogen maser during VLBI observations, meeting the timing requirements for FRB VLBI with CHIME/FRB Outriggers. Furthermore, this system can be implemented without interrupting CHIME’s current observational campaign and without modifications to the correlator or the data-analysis pipelines for cosmology and radio transient science.

The paper is organized as follows. Section 2 describes the features of the CHIME instrument that are relevant to its use as a VLBI station in CHIME/FRB Outriggers. Section 3 discusses the CHIME/FRB Outriggers clock stability requirements for FRB VLBI. Section 4 describes the hardware and software of the stabilization system that transfers CHIME’s reference clock to a passive hydrogen maser. Section 5 shows the results of the suite of tests that validate the clock stabilization system with VLBI-style observations between CHIME and the CHIME Pathfinder (an early small-scale prototype of CHIME recently outfitted as an outrigger test bed; Bandura et al. 2014; Leung et al. 2021). Section 6 presents an alternate clock solution for outrigger stations that do not have the infrastructure to support a hydrogen maser. Section 7 presents the conclusions.

2. Instrument Overview

A detailed description of the CHIME instrument and the CHIME/FRB project is presented in CHIME Scientific Collaboration et al. (2021, in preparation) and CHIME/FRB Collaboration et al. (2018). In this section, we give a brief introduction to these systems focused on the features that are relevant for FRB VLBI. We also give an overview of CHIME/FRB Outriggers.

2.1. CHIME and CHIME/FRB

CHIME is a hybrid cylindrical transit interferometer located at the Dominion Radio Astrophysical Observatory (DRAO) near Penticton, B.C., Canada. It consists of four $20\text{ m} \times 100\text{ m}$ cylindrical reflectors oriented north–south and instrumented with a total of 1024 dual-polarization feeds and low-noise receivers operating in the 400–800 MHz band. The cylinders are fixed with no moving parts, so CHIME operates as a drift-scan instrument that surveys the northern half of the sky every day with an instantaneous field of view of $\sim 120^\circ$ north–south by 2.5° – 1.3° east–west.

Although CHIME’s design was driven by its primary scientific goal to probe the nature of dark energy by mapping the large-scale structure of neutral hydrogen in the universe across the redshift range $0.8 \leq z \leq 2.5$, its combination of high sensitivity and large field of view also make it an excellent instrument to study the radio transient sky. Thus, in its final stages of commissioning, the CHIME correlator was upgraded with additional hardware and software backends to perform additional real-time data processing operations for pulsar timing and FRB science.

The correlator (Bandura et al. 2016a, 2016b; Denman et al. 2020) is an FX design (temporal Fourier transform before spatial cross-multiplication of data), where the F engine digitizes the 2048 analog inputs at 800 MSPS and separates the 400 MHz input bandwidth into 1024 frequency channels with a 390 kHz spectral resolution. The F engine also implements the corner-turn network that rearranges the complex-valued channelized data (also known as “baseband”) before sending it to the X-engine that computes a variety of data products for the different real-time scientific backends: interferometric visibilities for the hydrogen intensity mapping backend (CHIME Scientific Collaboration et al. 2021 (in preparation)), dual-polarization tracking voltage beams for the pulsar monitoring backend (CHIME/Pulsar Collaboration et al. 2021), and high-frequency resolution power beams for the 21 cm absorption systems backend (Yu et al. 2014) and for the CHIME/FRB backend that is triggered by highly dispersed radio transients to search for FRBs in real time (CHIME/FRB Collaboration et al. 2018). Additionally, a $\sim 36\text{ s}$ long memory buffer in the X-engine stores baseband data (2.56 μs time resolution, 390 kHz spectral resolution, and 4-bit real + 4-bit imaginary bit depth for the 2048 correlator inputs) that can be saved to disk when the CHIME/FRB search pipeline detects an FRB candidate, enabling polarization and high-time resolution analysis of FRB events, as well as subarcminute localization precision (Michilli et al. 2021). Eventually it will also enable VLBI localization with CHIME/FRB Outriggers.

2.2. CHIME/FRB Outriggers

The scientific goal of CHIME/FRB Outriggers is to provide 50 mas localization for nearly all CHIME-detected FRBs with subhour latency. This angular resolution is sufficient to determine galaxy hosts and source environments and is well matched to current best optical follow-up observations. To this end, the CHIME/FRB collaboration is currently building outrigger telescopes at distances ranging from hundreds to several thousands of kilometers from DRAO. The outriggers will be small-scale versions of CHIME, each with about one eighth of CHIME’s collecting area, the same field of view, and

tilted such that they monitor the same region of the sky as CHIME.

In contrast to traditional VLBI that is typically performed for known targets with small fields of view and manageable data rates, the random nature of most FRBs requires the real-time processing of massive data rates in order to detect and localize these events in blind searches with wide fields of view. The baseband data rate of CHIME is 6.6 Tbit s^{-1} , while that of each outrigger station will be 0.8 Tbit s^{-1} . Since such high data rates cannot be continuously saved, the outriggers will adopt the triggered FRB VLBI approach demonstrated in Leung et al. (2021), where each station buffers its local baseband data in memory and only writes it to disk upon receipt of a trigger from the CHIME/FRB real-time search pipeline over internet links. The local data of each station is then transmitted to a central facility where the signals are correlated together such that the outriggers operate with CHIME as an interferometric instrument with the angular resolution of a telescope with an aperture of thousands of kilometers.

3. Clock Stability Requirements

Accurate timing is critical for VLBI since the localization of radio sources is ultimately derived from the relative time of arrival of signals at the telescope stations. By synthesizing the available frequency channels, it is possible to obtain a statistical precision on the measured delay given by (Rogers 1970)

$$\sigma_{\tau}^{\text{stat}} = \frac{1}{2\pi \cdot \text{SNR} \cdot \text{BW}_{\text{eff}}}, \quad (1)$$

where SNR is the signal-to-noise ratio of the VLBI event, and BW_{eff} is the effective bandwidth. For the CHIME/FRB detection threshold¹⁶ and bandwidth (BW), this corresponds to

$$\sigma_{\tau}^{\text{stat}} \approx \frac{184 \text{ ps}}{\left(\frac{\text{SNR}}{7.5}\right) \left(\frac{\text{BW}}{400 \text{ MHz}}\right)}. \quad (2)$$

For a VLBI baseline b , a delay precision σ_{τ} corresponds to a statistical localization uncertainty

$$\sigma_{\theta} \approx \frac{c}{b} \sigma_{\tau}, \quad (3)$$

which gives $\sigma_{\theta}^{\text{stat}} \lesssim 11 \text{ mas}$ for a 1000 km baseline. However, the (relative) delay measured by the interferometer includes not only the geometric delay (which ultimately provides the source localization) but also additional contributions that need to be accounted for, such as propagation through the troposphere and ionosphere, baseline errors, drift between clocks of different stations (clock timing errors), and other instrumental delays. In practice, the localization uncertainty of CHIME/FRB Outriggers will be limited by systematic errors due to uncompensated delay contributions and particularly by errors in the determination of the dispersive delay due to the ionosphere. Although the large observation bandwidth of the instrument

¹⁶ The SNR in VLBI is related to the SNR at CHIME as $\text{SNR}/\text{SNR}_{\text{CH}} = \sqrt{2A_0}/A_{\text{CH}} = 1/2$, where A_{CH} and A_0 are the collecting areas of CHIME and the outrigger, respectively, and the factor of $\sqrt{2}$ comes from the difference in the detailed noise statistics of a cross-correlation compared to an auto-correlation (Masui et al. 2015). While the CHIME/FRB real-time detection pipeline has a detection threshold of ~ 10 (Michilli et al. 2021), the SNR rises by $\sim 50\%$ through the more detailed analysis of the saved baseband data. As such, we take the floor on the CHIME detection SNR to be $\text{SNR}_{\text{CH}} = 15$.

helps to mitigate this effect, it still represents the most important challenge for system stability at CHIME frequencies.

Our simulations indicate that we can reliably localize FRB events to 50 mas, which, for $b = 1000 \text{ km}$, corresponds to a delay error budget of $\sigma_{\tau} \approx 800 \text{ ps}$. Anticipating that the ionosphere will be the main contributor to delay errors, the clock timing error specification¹⁷ has been set to $\sigma_{\tau}^{\text{clk}} \lesssim 200 \text{ ps}$.

Note that for blind FRB searches, this specification must be met at all times. Indeed, it may not always be possible to find a calibrator immediately after the detection of an FRB for phase referencing. An additional complication is that stationary telescopes like CHIME and the outriggers observe the sky as it transits through their field of view and thus cannot slew toward favorable calibrators. Therefore, it is especially important to have a reference clock that is reliable on the timescales required to connect an FRB detection to a calibrator observation, potentially hours later.

Although the list of steady radio sources potentially suitable for calibrating low-frequency VLBI arrays with $\gtrsim 1000 \text{ km}$ baselines has significantly increased thanks to the ongoing LOw Frequency ARray (LOFAR) Long-Baseline Calibrator Survey (LBCS; Moldón et al. 2015; Jackson et al. 2016), during its initial stages CHIME/FRB Outriggers will adopt a more conservative strategy relying mainly on bright pulsars for calibration. Pulsars are compact, can be separated from the steady radio background in the time domain, and are sufficiently abundant to be used as the primary sources for phase referencing. Accordingly, the backend of each outrigger will also have the ability to form tracking baseband beams for pulsar analysis and calibration. Recently, Cassanelli et al. (2021) demonstrated the potential of pulsars as calibrators for CHIME/FRB Outriggers through the triggered VLBI detection of an FRB over a $\sim 3000 \text{ km}$ baseline between CHIME and the Algonquin Radio Observatory 10 m Telescope using PSR B0531+21 in the Crab Nebula for phase referencing.

We estimate that an FRB detection can be phase-referenced to an SNR $\gtrsim 15$ pulsar¹⁸ within less than $\sim 10^3 \text{ s}$. This timescale and the relative clock timing error specification set the clock stability (Allan deviation) requirement to

$$\sigma_y(10^3 \text{ s}) \lesssim 2 \cdot 10^{-13}. \quad (4)$$

As explained in Appendix [appendix](#), the Allan deviation is a measure of stability commonly used in precision clocks and oscillators, with $\Delta t \cdot \sigma_y(\Delta t)$ roughly representing the rms of clock timing errors a time Δt after calibration. The specification in Equation (4) is typically obtained with hydrogen masers. However, in Section 6, we show that even frequency references that initially do not meet this requirement can be used as reference clocks for the outriggers by interpolating timing solutions between calibrators.

4. Clock Stabilization System

In this section, we discuss the considerations that led to the current clock stabilization solution for CHIME, as well as the hardware and data analysis for the maser signal.

¹⁷ This is not a hard upper limit but rather a reasonable reference value that represents our goal to keep the clock timing errors well below the 800 ps total timing error budget.

¹⁸ This represents the VLBI SNR after coherent addition of the pulses within a single pulsar transit.

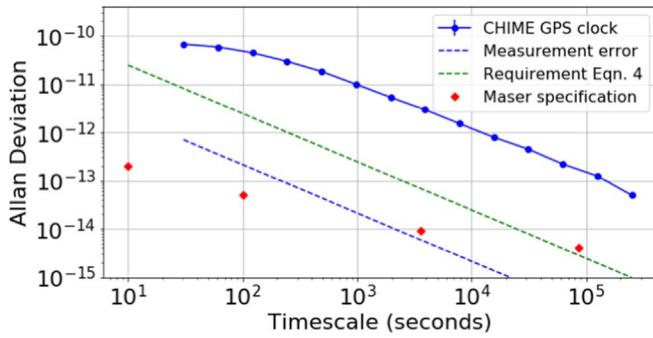


Figure 1. Allan deviation of the CHIME GPS clock and the DRAO hydrogen maser. Blue: Allan deviation of the CHIME GPS clock as measured with the clock stabilization system described in Section 4. A total of 10 days of raw ADC data at 30 s cadence were collected for the measurement. Dashed blue: expected measurement error contribution to the Allan deviation obtained from simulations of uncorrelated but time-dependent errors in the range ~ 4 – 20 ps rms (the range observed in the measured delays). Dashed green: stability requirement from Equation (4) assuming white noise delay errors. Red: manufacturer-specified Allan deviation of the DRAO maser. The CHIME GPS clock does not meet the stability requirements for FRB VLBI, but the DRAO maser does (Equation (4)).

4.1. Hardware/Software Considerations

The CHIME F engine is implemented using the ICE hardware, firmware, and software framework (Bandura et al. 2016a). It consists of field programmable gate array (FPGA)-based motherboards specialized to perform the data acquisition and channelization of the 2048 CHIME analog inputs. The ICE motherboards are packaged in eight crates with custom backplanes that implement the networking engine that reorganizes and sends the baseband data to a dedicated graphics processing unit (GPU) cluster that performs the X-engine operations. The outriggers will also use an ICE-based F engine.

The data acquisition and signal processing of the F engine are driven by a single 10 MHz clock signal provided by a Spectrum Instruments TM-4D global positioning system (GPS)-disciplined ovenized crystal oscillator. The GPS module also generates an Inter-Range Instrumentation Group (IRIG)-B time-code signal internally synchronized to the clock and that is used by the correlator to time stamp the data. A low-jitter distribution system sends the clock and time signals to each ICE backplane and motherboard and is ultimately used to generate the analog-to-digital converter (ADC) sampling clocks. The time stamping process is implicit: The F engine uses the IRIG-B signal to synchronize the start of the data acquisition to an integer second (up to the 10 ns resolution of the IRIG-B decoder in the FPGA firmware), and it also tags each data frame with a frame counter value. The X-engine time stamps the data by calculating the offset from the start time based on the frame counter value and assuming a fixed 2.56 μ s baseband sampling time.

Figure 1 shows the Allan deviation of the CHIME GPS clock (blue line) as measured with the clock stabilization system described in Sections 4.2 and 4.3. The GPS-disciplined crystal oscillator, being locked to a GPS time reference determined by a vast network of atomic clocks, will eventually surpass the stability performance of a single hydrogen maser on very long timescales ($\Delta t \gtrsim 10^6$ s). On intermediate and long timescales ($\Delta t \sim 10^3$ – 10^5 s), including the ones of interest for CHIME/FRB Outiggers, the CHIME clock stability is dominated by white delay noise ($\sigma_y(\Delta t) \propto 1/\Delta t$) corresponding to ~ 6 ns rms

timing errors. While the coherence of this frequency standard is sufficient for CHIME’s operations as a connected interferometer and for all its backends, the high precision needed for FRB VLBI requires the development of a more stable clock system.

As a continuously tracking global navigation satellite system station and as part of the Western Canada Deformation Array (WCDA; Chen et al. 1996) and the Canadian Active Control System (CACS¹⁹; Duval et al. 1996), DRAO is equipped with an atomic frequency standard consisting of a T4Science pH Maser 1008 passive hydrogen maser owned and operated by Natural Resources Canada (NRCan). The maser is installed in a seismic vault at the DRAO site and has a primary output of 5 MHz (sine wave). It is directly connected to a low-noise distribution amplifier in the same rack that serves as electrical isolation for the maser and also derives multiple copies of the 10 MHz reference signal (sine wave). NRCan has approved the use of two of those signals for CHIME-related operations.

The manufacturer-specified Allan deviation of the DRAO maser is shown in Figure 1 (red points). The maser clearly exceeds the stability requirements for FRB VLBI with CHIME (Equation (4)). Some of the outrigger sites will also have access to hydrogen maser frequency references with similar performance.

Although in principle the ICE system can be operated with an IRIG-B time signal that is not phase-locked to the 10 MHz reference clock (such as an independent maser), an important restriction to using the maser as the master clock for the CHIME correlator is the fact that both the F-engine and X-engine software and the scientific data analysis pipelines were developed on the premise that the clock and IRIG-B signals are synced (e.g., to time stamp the data), something that cannot be guaranteed if the two signals are generated by independent systems (the maser and the GPS receiver). Even if the relative drift between clock and time signals could be tracked, both the correlator and data analysis pipelines would need to be updated to implement this change. As such, we had to develop a clock stabilization system that did not impact the normal operations of CHIME, its existing real-time backends, and the other scientific teams.

4.2. Maser Signal Conditioning

The clock stabilization system designed for FRB VLBI with CHIME keeps the current GPS-disciplined crystal oscillator as the master clock and instead feeds the maser signal to one of the ICE ADC daughter boards so it is digitized by the crystal-oscillator-driven F engine. The data are processed to monitor the variations in the phase of the sampled maser signal, which correspond to variations in the relative delay between the maser and the master clock. By using this information to correct phase variations in the baseband data recorded at the time of an FRB detection, the system effectively transfers the reference clock from the GPS-disciplined oscillator to the more stable maser signal during FRB observations. As shown in Figure 1 (dashed blue line), the noise penalty associated with the clock transfer operation is essentially white ($\sigma_y(\Delta t) \propto 1/\Delta t$) on timescales relevant for FRB VLBI and small ($\lesssim 20$ ps) compared to the total clock timing error budget.

A block diagram of the maser signal path is shown in Figure 2. The first point of access to the 10 MHz maser signal is

¹⁹ <http://cgrsc.ca/resources/geodetic-control-networks/canadian-active-control-system-cacs>

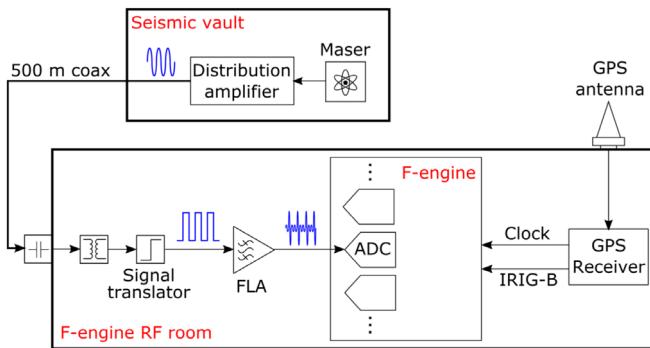


Figure 2. Maser signal path. The 10 MHz maser signal is transported through \sim 500 m of buried coaxial cable from the seismic vault to one of the CHIME F-engine RF huts. There, the maser signal is conditioned to a waveform that can be digitized by the CHIME F engine (see Section 4.2 for details).

the low-noise distribution amplifier within the seismic vault. From there, the signal is transported through \sim 500 m of buried coaxial cable to one of the two radio-frequency (RF)-shielded huts that house the CHIME F engine. We use the same type (LMR-400) of low-loss coaxial cable used in the CHIME analog receivers and whose thermal susceptibility has been extensively tested in the field. At the RF hut, the cable interfaces with a ground block, and the signal is then carried inside the RF room using standard SMA cables where it is connected to an isolation transformer to refer the next stages to the F-engine crate ground. One complication in the digitization of the maser signal is that the ADC daughter boards that specialize the ICE system for CHIME have a bandpass transfer function that strongly attenuates signals below \sim 100 MHz. For this reason, instead of feeding the maser signal directly to an ADC daughter board, the signal is used to drive a low-noise sine-to-square wave signal translator that generates 10 MHz harmonics well into the CHIME band. The output of the translator is then filtered to the CHIME band using the same band-defining filter amplifier (FLA) used in the CHIME receivers (Bandura et al. 2014, CHIME Scientific Collaboration et al. 2021 (in preparation)). Finally, the FLA is connected directly to one of the correlator inputs where it is digitized at 800 MSPS with an 8-bit ADC.

4.3. Clock Stabilization Pipeline

The FPGA within each ICE motherboard processes the data from its digitizers using the custom CHIME F-engine firmware (for details, see Bandura et al. 2016a, 2016b). Briefly, the raw ADC data from each input are passed to the frequency channelizer module as frames of 2048 8-bit samples. The channelizer forms the baseband data by splitting the 400 MHz input bandwidth into 1024 frequency channels, each truncated to a 4-bit real + 4-bit imaginary complex number. Additionally, a probe submodule within the channelizer can be configured to periodically capture a subset of the raw ADC data that is separately saved and typically used in CHIME for system monitoring.

By default, the CHIME F-engine software pipeline saves one raw ADC frame (2.56 μ s of data) from each input every 30 s, but this cadence can be modified before starting a data acquisition. The clock stabilization pipeline extracts the raw ADC frames from the maser input, Fourier transforms each frame via a Fast Fourier Transform (FFT), and separates the frequency channels corresponding to the harmonics of the

10 MHz signal in the CHIME band. The quality of each harmonic is assessed based on its signal-to-quantization-noise ratio and its susceptibility to spurious aliased harmonics (relevant for harmonics near the edges of the CHIME band). Low-quality harmonics are discarded.

Since the ADC that digitizes the maser signal uses the GPS clock as the reference for sampling, the variations in the delay of the sampled maser signal represent the delay variations of the GPS clock with respect to the maser, the latter of which is more stable on short and intermediate timescales. These delay variations $\Delta\tau(t)$ will induce phase variations $\Delta\phi(t, \nu)$ in the maser harmonics of the form

$$\Delta\phi(t, \nu) = 2\pi\nu\Delta\tau(t), \quad (5)$$

where ν is the harmonic frequency.

Since we are interested in the GPS clock delay variations relative to the delay at the time of VLBI calibration, the phase of the maser harmonics is initially referenced to the phase of a frame close to calibration time. Then for each frame, a line described by Equation (5) is fit to the phase as a function of harmonic frequency to recover the GPS clock delay (relative to the maser) as a function of time. The dominant component of the recovered delay $\Delta\tau(t)$ is a slow linear drift as a function of time (\sim 50 ns day $^{-1}$) corresponding to a constant offset of the maser frequency from 10 MHz. This linear trend is removed from $\Delta\tau(t)$ since it is due to the maser frequency calibration and not due to the instability of the GPS clock.

The captured raw ADC frames are only a small fraction of the available CHIME data; thus, the times at which GPS clock delay measurements are available are not necessarily aligned with the times of a calibration observation or an FRB detection. This means that the GPS clock delay time stream must be interpolated in order to find the clock contribution to the total delay measured in a VLBI observation. We use linear interpolation to find the GPS clock delay at any arbitrary time, a method that is motivated by the short-timescale behavior of the clock delay variations. Figure 3 shows a few examples of the behavior of the GPS clock delay on timescales of a few seconds as measured by the clock stabilization system. On these timescales, the timing variations are dominated by the tuning jitter generated by the algorithm that disciplines the crystal oscillator. In essence, the algorithm works by counting the number of clock cycles between successive GPS receiver pulse-per-second (PPS) pulses and adjusting the crystal's temperature to ensure 10 million counts between pulses. The size of the temperature tuning steps is progressively reduced as the crystal oscillator frequency approaches 10 MHz. As shown in Figure 3, this discipline procedure gives a characteristic triangle-wave shape to the tuning jitter, and although in a perfectly-tuned oscillator the transitions should occur every second, in practice we observe that they can take longer. Thus, as long as the GPS clock delay is sampled at cadences below \sim 500 ms we can track the tuning jitter features and a linear interpolation provides a good approximation to the true delay at any time.

As the current version of the F-engine control software only allows saving raw ADC data for all the correlator inputs at the time, raw ADC data at cadences below 10 s cannot be saved during normal telescope operations and are restricted to operations during times scheduled for hardware maintenance and software upgrades. A modification to the F-engine control software is ongoing to allow saving fast-cadence raw ADC data for the maser input while keeping the default cadence for the

remaining correlator inputs, a change that does not impact the normal operations of the correlator and the data-analysis pipelines.

It is also possible to process baseband data directly to extract the maser signal and measure the GPS clock delay variations. The operation is very similar to that of raw ADC data, except that the maser data have already been transformed to the frequency domain by the F engine. Since in this case most maser harmonics do not lie exactly in the center of a frequency channel, the pipeline selects the closest F-engine frequency channel. Then for each selected channel, it performs an additional channelization by using an FFT along the time domain to isolate the harmonic frequency.

Although working with baseband data is logically convenient in cases where we need to test the performance of the clock stabilization system (see Section 5), during regular operations the current system is designed to work mostly with raw ADC data. This is mainly because a baseband dump for an FRB event is typically collected in ~ 100 ms segments at different times for each frequency channel in order to account for the dispersion delay of the transient, with a total event duration lasting tens of seconds (Michilli et al. 2021). This leaves only a few megahertz of bandwidth available at any particular instant, making the monitoring of clock delay variations more challenging. Furthermore, when using baseband dumps we still need to rely on the continuously saved raw ADC data to track and correct the long-timescale linear drift of the maser.

5. Validation of the Clock Stabilization System

We tested the reliability of clock stabilization system by installing it in the Pathfinder telescope and comparing maser-based measurements of the CHIME–Pathfinder relative clock drift to independent measurements obtained from VLBI-style observations of steady radio sources.

5.1. The Pathfinder as an Outrigger

The Pathfinder is presented in detail in Bandura et al. (2014). It is a small-scale prototype of CHIME with identical design and field of view, and it has the same collecting area of the outriggers under construction. The telescope is located ~ 400 m from CHIME and was constructed before CHIME as a test bed for technology development. With the same correlator architecture as CHIME, the Pathfinder operates as an independent connected interferometer with its own GPS-disciplined clock. Recently, Leung et al. (2021) repurposed the Pathfinder as an outrigger to demonstrate the feasibility of triggered FRB VLBI for CHIME/FRB Outriggers. The Pathfinder correlator is now equipped with a custom baseband data recorder capable of processing one quarter of the CHIME band and programmed to write its local baseband data to disk upon receipt of a trigger from CHIME/FRB. We also connected the additional copy of the maser signal from the seismic vault to one of the Pathfinder correlator inputs using a signal path identical to that of CHIME and shown in Figure 2 (except for the transport cable which is longer for the Pathfinder setup).

5.2. Comparison to Interferometric Observations

The clock stabilization system measures the delay variations of the CHIME and Pathfinder clocks by processing the maser

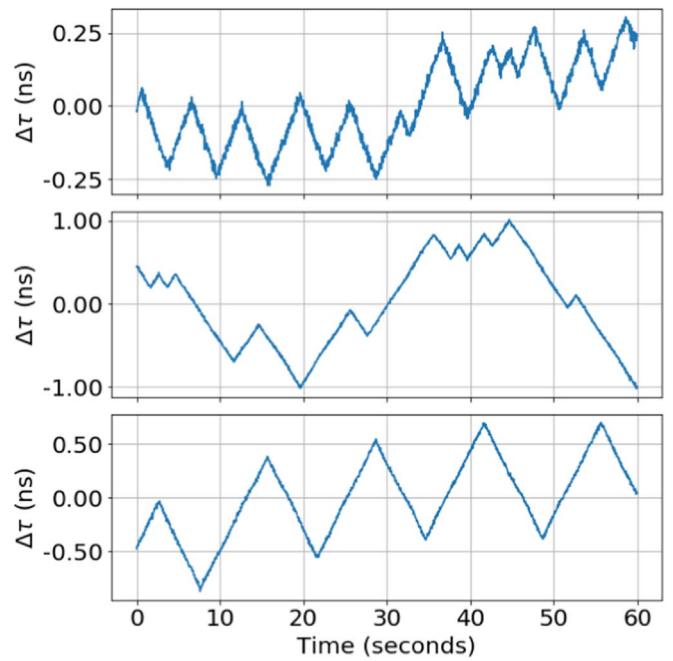


Figure 3. Three examples of the behavior of the GPS clock delay on timescales of a few seconds as measured by the clock stabilization system with respect to the DRAO maser. The raw ADC data cadence for this measurements was 40 ms. The measurement errors are in the range ~ 2 –13 ps. The characteristic triangle-wave pattern is due to the algorithm that disciplines the crystal oscillator in the GPS unit. The algorithm works by counting the number of clock cycles between successive GPS PPS pulses and adjusting the crystal’s temperature to ensure 10 million counts between pulses. The size of the temperature tuning steps changes depending on the tuning history of the oscillator.

data from each telescope, as described in Section 4. The delays from each clock are then interpolated to the observation times so the relative clock drift can be tracked over time.²⁰ An independent way to measure the CHIME–Pathfinder relative clock delay is to interferometrically track a known point source over time using both telescopes and their independently running backends. If we properly account for all other contributions to the measured interferometric delay (geometric, ionosphere, etc.) as the source transits through the field of view of the two telescopes, then any residual delay should correspond to the relative drift between the two clocks. If the clock stabilization system is robust, its measurements should agree very closely with the interferometric measurement, which we use as a standard.

5.2.1. Short-timescale Test

We used Cygnus A (henceforth referred to as CygA) for the VLBI-style observations since it is the brightest radio source seen by CHIME that is unresolved on a CHIME–Pathfinder baseline. For the first test, we programmed the CHIME/FRB backend to trigger short baseband dumps simultaneously for CHIME and the Pathfinder during a single CygA transit. In this way, we collected seven 10 ms-long baseband dumps, spaced by a minute, while the source was in the field of view. The observation was carried out in 2020 November during a day scheduled for instrument maintenance, so we were also able to

²⁰ If the maser data comes from simultaneous baseband dumps instead of raw ADC samples, then clock delays from the two telescopes can be directly compared without interpolation.

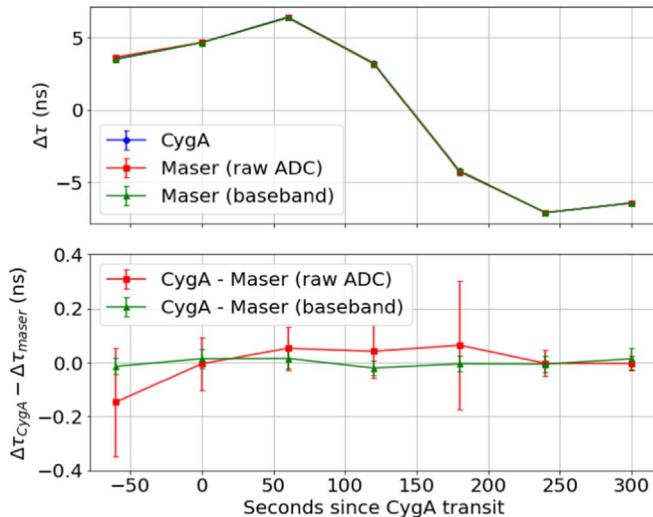


Figure 4. Comparison of the CHIME–Pathfinder relative clock delay inferred via the clock stabilization system and interferometric observations from a single transit of CygA. Top: relative clock delay (in ns) as a function of time, as inferred from CygA baseband data (blue), raw ADC maser data (red) and maser baseband data (green). Bottom: difference between sky-based and maser-based measurements of the relative clock delay (raw ADC maser data in red, baseband maser data in green), demonstrating agreement between the two methods. The large error bars for all but the last point in the raw ADC data analysis (red) are due to current limitations of the Pathfinder raw ADC acquisition system (see Section 5.2.1 for details). The error bar in the last red point of the top plot (~ 14 ps) is representative of the expected accuracy of the clock stabilization system using raw ADC data.

collect raw ADC maser data at 200 ms cadence with the two telescopes.

Since the telescopes are co-located, they experience a common ionosphere, suppressing relative ionospheric fluctuations (we see no evidence for these in our observations). Thus, the residual delay in the interferometric visibility after accounting for the geometric contribution gives a measurement of the relative clock drift. The residual interferometric delays are calculated using a procedure identical to that described in Leung et al. (2021). In summary, each telescope is internally calibrated (to measure the directional response of each antenna in the telescope array) using a separate observation of a bright point source (CHIME Scientific Collaboration et al. 2021 (in preparation)). Then, for each telescope and baseband dump, the data are coherently summed over antennas to form a phased-array voltage beam toward CygA. The beamformed data from CHIME and the Pathfinder are then cross-correlated on a frequency-by-frequency basis to form the complex visibility. The phase of each visibility is compensated for the geometric delay. The variance of each visibility is found empirically by splitting each baseband dump into short time segments, computing the visibility for each segment, and calculating the variance over segments as a function of frequency. The seven visibilities are phase-referenced to that of the first baseband dump since we are only interested in changes of the relative clock delay. Finally, the wideband fringe-fitting procedure to find residual delay performs a least-mean-squares fit of a complex exponential with linear phase and frequency-dependent amplitude to the measured visibilities.

The top panel of Figure 4 shows the resulting comparison between the CHIME–Pathfinder relative clock delays calculated from interferometric observations (blue) and those found through the clock stabilization system (from raw ADC data in

red, from baseband data in green). The 1σ error bars are too small to be visible in the plot but they are in the range ~ 17 – 22 ps for CygA measurements, ~ 14 – 238 ps for raw ADC maser data measurements, and ~ 18 – 33 ps for baseband maser data measurements. Measurements with the clock stabilization system show excellent agreement with the sky. This is further highlighted in the bottom panel of Figure 4 that shows the difference between sky-based and maser-based measurements of the relative clock delay (raw ADC maser data in red, baseband maser data in green).

The large error bars for all but the last point in the raw ADC data analysis (red) are dominated by the error in the measurements of the maser delay at the Pathfinder. These can be traced back to current limitations of the Pathfinder raw data acquisition system, which occasionally drops packets when we collect raw ADC data at fast cadence for the reasons explained in Section 4.3. This limitation will be solved in the next upgrade of the F-engine control software. For the observation times that fell within sections of missing Pathfinder raw ADC data (the longest of which was ~ 80 s), the delay values were obtained by performing a smoothing spline interpolation based on the available measurements. To estimate the uncertainty in the delay values obtained with this method, we analyzed a segment of the delay time stream for which there were no gaps due to dropped packets, ~ 10 minutes before the sky observations. The errors were found using a procedure similar to the one used in Section 6 to evaluate the performance of alternate reference clocks, where we introduce artificial gaps in the delay time stream and analyze the statistics of an ensemble of interpolation residuals. Only the raw ADC delays (red) for the last observation time could be measured using the default interpolation for both telescopes (see Section 4.3). The uncertainty for this measurement is ~ 14 ps, and is representative of the expected accuracy of the clock stabilization system using raw ADC data.

5.2.2. Long Timescale Test

To test the performance of the clock stabilization system on long timescales we collected five CygA baseband dumps, each 10 ms in duration, spaced one minute apart, for nine days in a row for a total of 45 delay measurements. The observations were carried out in 2021 April during normal CHIME operations so we relied on baseband maser data for delay measurements with the clock stabilization system for the reasons explained in Section 4.3. Both interferometric and maser-based delays are calculated using the same procedure as in the short-timescale test described in Section 5.2.1, with the visibilities phase-referenced to one of the observations on the fifth day.

The results of the long timescale test are shown in Figure 5. The interferometric measurements agree with the maser-based measurements at the ~ 30 ps rms level, demonstrating that after correction with the clock stabilization system the resulting reference clock is stable over timescales of more than a week, and that the signal chain used to inject the maser signal into the correlator is not a limitation for the system’s performance in CHIME/FRB Outiggers.

6. A Reference Clock for Outrigger Stations without a Maser

The clock stabilization system allows us to inject external reference clock signals into radio telescopes that share the

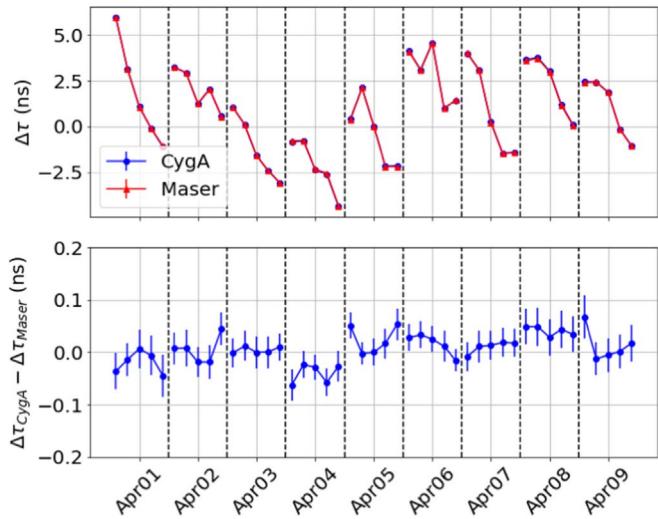


Figure 5. Top: comparison of the CHIME-Pathfinder relative clock delay inferred via the clock stabilization system (red) and interferometric observations (blue) of from multiple transits of CygA. For each transit, we made five measurements of the relative clock delay, spaced by one minute, for nine days in a row. Bottom: difference between sky-based and maser-based measurements of the relative clock delay. The two methods show excellent agreement on short (minute) and long (many days) timescales. This indicates that the clock stabilization system we have implemented can track clock delay variations with better than ~ 30 ps rms level precision.

CHIME correlator architecture. In addition to its use in the new outriggers, the system enables reference clocks to be swapped out at existing telescopes like CHIME and legacy systems like the Pathfinder without making major changes to the software framework or existing scientific backends, while expanding the telescopes' capabilities to include VLBI.

Most outriggers will also have access to hydrogen maser frequency references that can be used in the same way as CHIME (Section 4) to meet the stability requirements for FRB VLBI. However, we still need to address the possibility that certain outrigger stations may be built at locations (e.g., greenfield land) that will lack the infrastructure to support a hydrogen maser. In this scenario, alternate reference signals (e.g., from rubidium microwave oscillators) can also be injected into the correlator to track and compensate for GPS clock drifts. The performance of these oscillators is inferior to that of a hydrogen maser, but they are still more stable than the CHIME GPS clock on the timescales relevant for FRB VLBI. They are also less expensive and more readily available than a maser. Even if these frequency references can potentially be used directly as the correlator master clock since they typically come in units that can provide GPS disciplining as well as absolute time, it is still desirable to use them separately as free-running clocks for short and intermediate timescale observations. Not only are they inherently more stable than the primary CHIME clock, but they are not subject to short-timescale tuning jitter when not locked to GPS.

Equation (4) provides a convenient way to determine whether an off-the-shelf clock meets the requirements for FRB VLBI by simply reading the $\sigma_y(10^3 \text{ s})$ value from the unit's data sheet. Passive hydrogen masers meet and exceed this requirement. However, this specification was derived from Equation (A5), which assumes uncorrelated clock timing errors and perfect calibration measurements. In practice, the actual clock timing errors will depend on aspects not necessarily captured by this equation including the detailed statistics of the

delay variations, the methods used to estimate them, the timing and accuracy of the calibration measurements, and the technique that we use to inject the clock to our system. These aspects become relevant when the frequency standard does not clearly exceed the specification in Equation (4).

As part of the implementation of the clock stabilization system, Cary et al. (2021) developed a software package with methods to determine the suitability of precision clocks for VLBI with transit telescopes like CHIME/FRB Outriggers. These methods take into account the details of the noise processes that determine the stability of the clocks and simulate realistic timing calibration scenarios. The basic input to the software is a time stream that represents the delay variations as a function of time of the clock under test. The delay time stream data can be either from measurements or from simulations; in the latter case, the software provides tools to generate time streams described by combinations of power-law noise processes commonly observed in precision clocks and oscillators including white phase modulation noise, white frequency modulation noise, flicker frequency modulation noise, and random walk frequency modulation noise (Allan 1987). Similarly, the software provides tools to generate delay time streams from a set of Allan deviation measurements, which is convenient for evaluating the performance of a clock based on its manufacturer specifications. In this case, it is assumed that the delay variations are described by a combination of power-law noise processes where the weight of each noise component is found by fitting the Allan variance data to a model consisting of a linear combination of the Allan variances of the previously described noise processes.

A calibrator is parameterized by its observing time, number of clock timing measurements, and SNR per transit. For example, for a calibrator at the equator, the observing time with CHIME is ~ 6 minute, and with a ~ 2 minute integration time we would have three delay measurements per transit. The SNR determines the uncertainty in the calibration delay measurements (see Equation (1)). Given a timescale Δt_{cal} that represents the maximum expected time separation between calibrators, the method masks a random Δt_{cal} -long section of the delay time stream, interpolates using a best-fit function determined from the available calibration measurements at each end of the masked section, and keeps the interpolation residuals. The process is repeated a configurable number of times to obtain a statistical ensemble of interpolation residual time streams, each of length²¹ $\sim \Delta t_{\text{cal}}$. As the default metric of the stability of the clock at Δt_{cal} timescales, the method uses the largest value of the ensemble standard deviation in the interval $[0, \Delta t_{\text{cal}}]$. Other metrics of performance are available. Since throughout the paper we have used the convention that Δt represents the time between a calibration and an observation, this metric can be interpreted as an estimate of the largest clock timing rms error for Δt up to $\sim \Delta t_{\text{cal}}/2$.

Different interpolation methods are available including linear fit, smoothing spline, and nearest available calibrator. The fitting weights are determined by the calibrator SNR and the level of the noise added by the timing stabilization system.

As a candidate for outriggers without a maser, we evaluated the performance of the EndRun Technologies Meridian II US-Rb rubidium oscillator. The unit was installed in the Pathfinder RF room and connected to a separate input of the correlator so

²¹ In practice, there is an additional overhead equivalent to one integration.

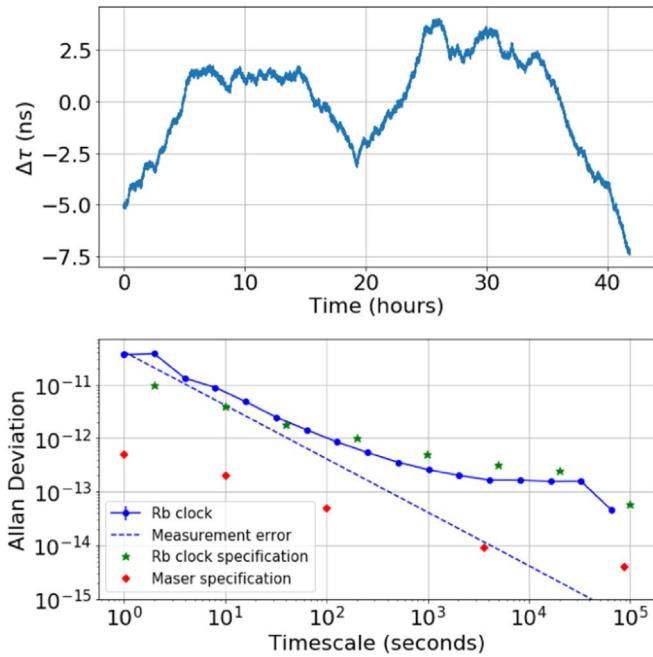


Figure 6. Top: measured delay variations of the rubidium oscillator tested as a candidate reference clock for outriggers without a maser. Bottom: measured Allan deviation of the Rb clock (blue), measurement error (dashed blue), and manufacturer-specified Allan deviation of the Rb clock (green points) and the DRAO maser (red points). The measured Allan deviation of the Rb clock is consistent with the specification at intermediate and long timescales. At short timescales, the noise of the timing stabilization system dominates the performance, but it is still small ($\sim 10\text{--}30$ ps) compared to the clock timing error budget. This confirms that the hardware of the clock stabilization system is not a limitation for clock performance in CHIME/FRB Outriggers.

we could test its performance against the DRAO maser under conditions comparable to those of a typical outrigger. A signal conditioning chain identical to the maser was used (signal translator + FLA). We collected ~ 42 hr of raw ADC data at 1 s cadence for both the maser and the Rb clock and used the clock stabilization pipeline to extract the clock delay variations relative to the maser. The top panel of Figure 6 shows the measured delay variations after removing the slow linear drift component due to the maser (see Section 4.3). The bottom panel shows the corresponding Allan deviation of the Rb clock (blue), the measurement error (dashed blue), and the manufacturer-specified Allan deviation of the Rb clock (green points) and the DRAO maser (red points). For $\Delta t \lesssim 40$ s the variations are dominated by the noise associated to the timing stabilization system, which is in the range $\sim 10\text{--}30$ ps and is small compared to the ~ 200 ps clock timing error budget. At longer timescales, the measured Allan deviation of the Rb oscillator is consistent with the manufacturer specification. These results are also consistent with direct Allan deviation measurements of the Rb oscillator performed with a phase noise analyzer, confirming that the hardware of the stabilization system is not a limitation for clock performance in CHIME/FRB Outriggers.

Note that if we rely only on the manufacturer-specified Allan deviation, the Rb clock does not meet the requirement in Equation (4). This justifies a more detailed analysis of the clock performance to determine whether it can still be used as a frequency reference for the outriggers without a maser.

The measured delay time stream was analyzed with the software package described above and in Cary et al. (2021) to

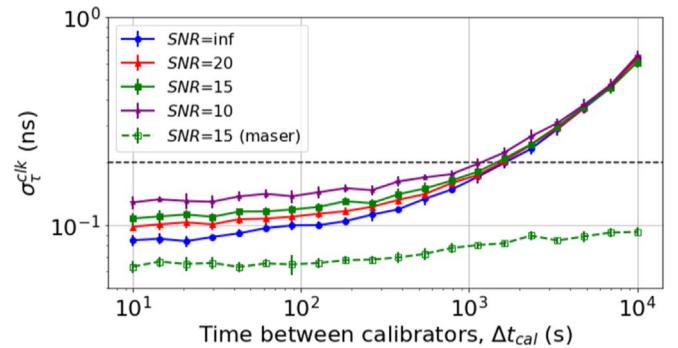


Figure 7. Projected clock errors, $\sigma_{\tau}^{\text{clk}}$, of the Rb clock as a function of the time between calibrators Δt_{cal} from measured delay variations and simulations of realistic timing calibration scenarios. This metric represents an estimate of the largest clock timing error for Δt up to $\sim \Delta t_{\text{cal}}/2$ (see Section 6 for details). The dashed black horizontal line represents $\sigma_{\tau}^{\text{clk}} = 200$ ps. Even in the most conservative scenario where we assume that all the calibrators have $\text{SNR} = 15$ (solid green), the Rb clock timing errors stay below 225 ps up to $\Delta t \sim 10^3$ s by interpolating between timing solutions, meeting the requirements for FRB VLBI with CHIME/FRB Outriggers.

evaluate the expected performance of the Rb clock under different calibration conditions. The results are shown in Figure 7. For the analysis, we assumed that the calibrator observing time was 9 minutes (roughly the value at CHIME’s zenith) with two integrations per observation. The best performance is obtained with linear interpolation between calibrators. We tested calibrator SNRs of 10 (purple), 15 (green), 20 (red), and ∞ (blue), with the latter representing the case where the clock performance is not limited by calibration errors. For comparison, we also show the projected clock timing errors for the DRAO maser using synthetic data generated from the manufacturer-specified Allan deviation with $\text{SNR} = 15$ (dashed green). Figure 7 shows that, even in the most conservative scenario where we assume that all the calibrators have $\text{SNR} = 15$, the Rb clock timing errors stay below 225 ps up to $\Delta t \sim 10^3$ s by interpolating between timing solutions (solid green). This clock timing error is still well below the total timing budget of ~ 800 ps, leaving enough room to handle ionospheric delay errors.

7. Conclusions

We developed a clock stabilization system for CHIME/FRB Outriggers that allows synchronization of CHIME and outrigger stations at the ~ 200 ps level on short and long timescales. This meets the requirements for 50 mas localization of FRBs detected with the CHIME/FRB real-time pipeline. Our proof-of-principle clock transfer has demonstrated that a variety of different data products can be used for precise time transfer from an external reference clocks into data acquisition backends using the ICE framework. This method is minimally invasive to existing telescopes like CHIME and the Pathfinder, expanding the capabilities of these instruments to include VLBI without impacting their existing scientific backends. It also allows for increased flexibility and modularity for future systems such as those at CHIME outriggers. For outriggers that do not have the infrastructure to support a hydrogen maser, we demonstrated that it is still possible to meet the required clock stability specification by using alternate reference clocks and interpolating timing solutions between calibrations.

We acknowledge that CHIME is located on the traditional, ancestral, and unceded territory of the Syilx/Okanagan people.

We are grateful to the staff of the Dominion Radio Astrophysical Observatory, which is operated by the National Research Council Canada. CHIME is funded by a grant from the Canada Foundation for Innovation (CFI) 2012 Leading Edge Fund (Project 31170) and by contributions from the provinces of British Columbia, Québec, and Ontario. The CHIME/FRB Project is funded by a grant from the CFI 2015 Innovation Fund (Project 33213) and by contributions from the provinces of British Columbia and Québec, and by the Dunlap Institute for Astronomy and Astrophysics at the University of Toronto. Additional support was provided by the Canadian Institute for Advanced Research (CIFAR), McGill University and the McGill Space Institute via the Trottier Family Foundation, and the University of British Columbia. The CHIME/FRB Outriggers program is funded by the Gordon and Betty Moore Foundation and by a National Science Foundation (NSF) grant (2008031). FRB research at MIT is supported by an NSF grant (2008031). FRB research at WVU is supported by an NSF grant (2006548, 2018490). J.M.P. is a Kavli Fellow. C.L. was supported by the U.S. Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. V.M.K. holds the Lorne Trottier Chair in Astrophysics & Cosmology and a Distinguished James McGill Professorship and receives support from an NSERC Discovery Grant and Herzberg Award, from an R. Howard Webster Foundation Fellowship from CIFAR, and from the FRQNT Centre de Recherche en Astrophysique du Québec. J.L.S. acknowledges support from the Canada 150 programme.

Facility: CHIME

Software: AllanTools (Wallin et al. 2018)

Appendix

Clock stability and Allan Deviation

On timescales of a second and larger, clock stabilities are usually quoted in terms of the Allan variance, or its square-root the Allan deviation (Allan 1966). In this appendix, we show how the clock stability required for FRB VLBI relates to the Allan deviation.

Let the “true” time be t , and the difference in seconds between what our clock reads and the true time be $x(t)$. In general, we will not know what the clock timing error $x(t)$ is for all time, but for CHIME/FRB Outriggers we will measure it whenever we do an on-sky calibration. If we calibrate at two points in time, we naturally would like to predict the timing error at the midpoint between the two calibrations. In particular, if the calibration times are t and $t + 2 \cdot \Delta t$, then we want to estimate $x(t + \Delta t)$ given $x(t)$ and $x(t + 2 \cdot \Delta t)$. The simplest assumption that we can make is that our clock runs at a constant rate between t and $t + 2 \cdot \Delta t$, in which case our prediction for $x(t + \Delta t)$ will be $[x(t) + x(t + 2 \cdot \Delta t)]/2$. The variance of the timing error estimate halfway between the calibrations will then be:

$$\begin{aligned} & \left\langle \left[x(t + \Delta t) - \frac{x(t) + x(t + 2 \cdot \Delta t)}{2} \right]^2 \right\rangle \\ &= \frac{1}{4} \langle [x(t) - 2x(t + \Delta t) + x(t + 2 \cdot \Delta t)]^2 \rangle. \end{aligned} \quad (\text{A1})$$

The Allan variance is defined to be

$$\begin{aligned} \sigma_y^2(\Delta t) &\equiv \frac{1}{2 \cdot \Delta t^2} \\ &\langle [x(t) - 2x(t + \Delta t) + x(t + 2 \cdot \Delta t)]^2 \rangle, \end{aligned} \quad (\text{A2})$$

so we can now relate our timing error variance halfway between the calibrations directly to the Allan variance:

$$\begin{aligned} & \left\langle \left[x(t + \Delta t) - \frac{x(t) + x(t + 2 \cdot \Delta t)}{2} \right]^2 \right\rangle \\ &= \frac{\Delta t^2}{2} \sigma_y^2(\Delta t). \end{aligned} \quad (\text{A3})$$

In general, and up to factors of order unity, the variance of the timing error will be determined by the product $\Delta t^2 \cdot \sigma_y^2(\Delta t)$, so

$$\sigma_x(\Delta t) \approx \Delta t \cdot \sigma_y(\Delta t) \quad (\text{A4})$$

is typically used as an approximation to the standard deviation of the clock timing error after a time Δt (Barnes et al. 1971; Kartaschoff 1979; Rogers & Moran 1981). In practice, the actual errors will depend on the details of the prediction algorithm and the noise processes that dominate the stability of the clock. For the particular case of white timing noise, it follows from Equation (A2) that

$$\sigma_y^{\text{WN}}(\Delta t) = \sqrt{3} \frac{\sigma_x}{\Delta t}. \quad (\text{A5})$$

From Equation (A2), the Allan deviation is dimensionless, so it is telling us the fractional uncertainty in our clock between calibrations. Quantitatively, if our pulsar calibrations are separated by $2 \cdot \Delta t = 2000$ s, and the limit on the *differential* timing residual at the midpoint is 200 ps rms (Section 3), then our Allan deviation requirement is

$$\sigma_y(10^3 \text{ s}) \approx \sqrt{\frac{3}{2}} \cdot \frac{2 \cdot 10^{-10} \text{ s}}{10^3 \text{ s}} \approx 2 \cdot 10^{-13}. \quad (\text{A6})$$

ORCID iDs

- J. Mena-Parra  <https://orcid.org/0000-0002-0772-9326>
- C. Leung  <https://orcid.org/0000-0002-4209-7408>
- S. Cary  <https://orcid.org/0000-0003-1860-1632>
- K. W. Masui  <https://orcid.org/0000-0002-4279-6946>
- J. F. Kaczmarek  <https://orcid.org/0000-0003-4810-7803>
- M. Amiri  <https://orcid.org/0000-0001-6523-9029>
- K. Bandura  <https://orcid.org/0000-0003-3772-2798>
- P. J. Boyle  <https://orcid.org/0000-0001-8537-9299>
- T. Cassanelli  <https://orcid.org/0000-0003-2047-5276>
- J.-F. Cliche  <https://orcid.org/0000-0001-6509-8430>
- M. Dobbs  <https://orcid.org/0000-0001-7166-6422>
- V. M. Kaspi  <https://orcid.org/0000-0001-9345-0307>
- T. L. Landecker  <https://orcid.org/0000-0003-1455-2546>
- A. Lanman  <https://orcid.org/0000-0003-2116-3573>
- J. L. Sievers  <https://orcid.org/0000-0001-6903-5074>

References

- Allan, D. 1966, in Proc. IEEE, 54, 221
- Allan, D. 1987, IEEE Trans. Ultrason. Ferroelectr. Freq. Control, 34, 647
- Bandura, K., Bender, A. N., Cliche, J. F., et al. 2016a, JAI, 5, 1641005

Bandura, K., Cliche, J. F., Dobbs, M. A., et al. 2016b, *JAI*, **5**, 1641004

Bandura, K., Addison, G. E., Amiri, M., et al. 2014, *Proc. SPIE*, **9145**, 738

Barnes, J. A., Chi, A. R., Cutler, L. S., et al. 1971, *IEEE Trans. Instrum. Meas.*, **IM-20**, 105

Bhandari, S., Keane, E. F., Barr, E. D., et al. 2018, *MNRAS*, **475**, 1427

Bhardwaj, M., Gaensler, B. M., Kaspi, V. M., et al. 2021a, *ApJL*, **910**, L18

Bhardwaj, M., Kirichenko, A. Y., Michilli, D., et al. 2021b, *ApJL*, **919**, L24

Cary, S., Mena-Parra, J., Leung, C., et al. 2021, *RNAAS*, **5**, 216

Cassanelli, T., Leung, C., Rahman, M., et al. 2021, arXiv:2107.05659

Chawla, P., Kaspi, V. M., Ransom, S. M., et al. 2021, arXiv:2107.10858

Chen, X., Langley, R. B., & Dragert, H. 1996, in *GPS Trends in Precise Terrestrial, Airborne, and Spaceborne Applications*, ed. G. Beutler et al. (Berlin: Springer), 70

CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2018, *ApJ*, **863**, 48

CHIME/FRB Collaboration, Amiri, M., Bandura, K., et al. 2019a, *Natur*, **566**, 235

CHIME/FRB Collaboration, Andersen, B. C., Bandura, K., et al. 2019b, *ApJL*, **885**, L24

CHIME/Pulsar Collaboration 2021, in preparation

Denman, N., Renard, A., Vanderlinde, K., et al. 2020, *JAI*, **9**, 2050014

Duval, R., Héroux, P., & Beck, N. 1996, in *GIS Conf.*, Vol. 53, 7

Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., et al. 2019, *ApJL*, **875**, L2

Fonseca, E., Andersen, B. C., Bhardwaj, M., et al. 2020, *ApJL*, **891**, L6

Jackson, N., Tagore, A., Deller, A., et al. 2016, *A&A*, **595**, A86

Josephy, A., Chawla, P., Curtin, A. P., et al. 2021, *ApJ*, **923**, 2

Kartaschoff, P. 1979, *IEEE Trans. Instrum. Meas.*, **28**, 193

Leung, C., Mena-Parra, J., Masui, K., et al. 2021, *AJ*, **161**, 81

Lorimer, D. R., Bailes, M., McLaughlin, M. A., Narkevic, D. J., & Crawford, F. 2007, *Sci*, **318**, 777

Masui, K., Amiri, M., Connor, L., et al. 2015, *A&C*, **12**, 181

Matthews, L. D., Crew, G. B., Doebleman, S. S., et al. 2018, *PASP*, **130**, 015002

Michilli, D., Masui, K. W., Mckinven, R., et al. 2021, *ApJ*, **910**, 147

Moldón, J., Deller, A. T., Wucknitz, O., et al. 2015, *A&A*, **574**, A73

Pleunis, Z., Good, D. C., Kaspi, V. M., et al. 2021, *ApJ*, **923**, 1

Rafiei-Ravandi, M., Smith, K. M., Li, D., et al. 2021, *ApJ*, **922**, 42

Rogers, A. E. E. 1970, *RaSc*, **5**, 1239

Rogers, A. E. E., & Moran, J. M. 1981, *IEEE Trans. Instrum. Meas.*, **IM-30**, 283

Schmittberger, B. L., & Scherer, D. R. 2020, arXiv:2004.09987

The CHIME/FRB Collaboration, Amiri, M., et al. 2021, arXiv:2106.04352

Wallin, A. E. E., Price, D. C., Carson, C. G., & Meynadier, F. 2018, *allantools: Allan deviation calculation, Astrophysics Source Code Library*, ascl:1804.021

Yu, H.-R., Zhang, T.-J., & Pen, U.-L. 2014, *PhRvL*, **113**, 041303