IOP Publishing Class. Quantum Grav. **43** (2026) 065003 https://doi.org/10.1088/1361-6382/ae4708

Classical and Quantum Gravity

OPEN ACCESS

RECEIVED
15 September 2025

REVISED
9 December 2025

ACCEPTED FOR PUBLICATION
17 February 2026

PUBLISHED
17 March 2026

PAPER

A hardware testbed for LISA inter-satellite signals

Reid Ferguson[1,3,*], **Guido Mueller**[1,2,3] and **Olaf Hartwig**[1,3]

[1] Precision Interferometry and Fundamental Interactions, Max Planck Institute for Gravitational Physics, Hanover, Germany
[2] Department of Physics, University of Florida, Gainesville, FL, United States of America
[3] Faculty for Mathematics and Physics, Leibniz University Hanover, Hanover, Germany
[*] Author to whom any correspondence should be addressed.

E-mail: reid.ferguson@aei.mpg.de, guido.mueller@aei.mpg.de and olaf.hartwig@aei.mpg.de

Keywords: LISA, time delay interferometry, FPGA, testbed, gravitational waves, hardware-in-the-loop, LISA simulations

Abstract

Time-delay interferometry (TDI) is essential for the Laser Interferometer Space Antenna (LISA) to suppress laser frequency noise and extract gravitational-wave signals from inter-satellite phase measurements. While most recent studies of TDI have relied on software simulations, we present here the development and first results of a new hardware testbed that simulates the LISA inter-satellite links using radio frequency system-on-chip field-programmable gate array-based delay lines. The system digitally applies LISA-like delays, Doppler shifts, and injected gravitational-wave signals to MHz carriers with clock sidebands, and provides phasemeter-compatible outputs for TDI analysis. We demonstrate baseline performance of the delay line, including correction of differential analog-to-digital converter/digital-to-analog converter jitter, and static tests of TDI X_1 combinations with external signals. In addition, we show injection and recovery of massive black hole binary waveforms through the hardware-in-the-loop setup. These results establish a flexible platform for testing TDI, clock noise transfer, and ranging techniques in a controlled environment, and outline future steps toward a full laboratory-scale simulator of the LISA constellation.

1. Introduction

1.1. The Laser Interferometer Space Antenna (LISA) constellation

Planned for launch by the European Space Agency in 2035, the LISA is the first non-terrestrial gravitational wave (GW) observatory [1]. At its core is the measurement of the distance between test masses in free fall. This technology was proven by the LISA Pathfinder mission in 2016 [2], and the technology for free-falling test masses can be scaled up to three satellites orbiting the Sun to produce the LISA constellation. The spacecraft measure their own positions relative to the test masses floating within them, and the distance between themselves and the distant spacecraft, via highly sensitive laser interferometers. Combining these measurements provides a test mass-to-test mass measurement. These test masses can be thought of as a subset of a ring of test masses which is deformed by a passing GW, as shown in figure 1.

With the position of the test masses relative to the spacecraft handled sufficiently [2], the inter-spacecraft interferometer (ISI) is left as the remaining untested part. Its 2.5 million kilometers ($\sim$8.3 light-seconds) of length provide the sensitivity to dimensionless strain required to act as an antenna for GWs.

Controlling the arm lengths of terrestrial detectors to be nearly equal allows for the frequency noise of the laser to be decoupled from the measurement of differential phase in the interferometer arms. This is aided by the use of a complex, high-gain laser frequency control system [3, 4]. The dynamics of orbital mechanics, however, prevent such a thing from being practical in space-based observatories. One approach to stabilize the lasers would be to use the long arm itself as a reference, known as arm-locking. However, arm-locking in LISA has been shown to be neither sufficient nor necessary for the mission to succeed [5].

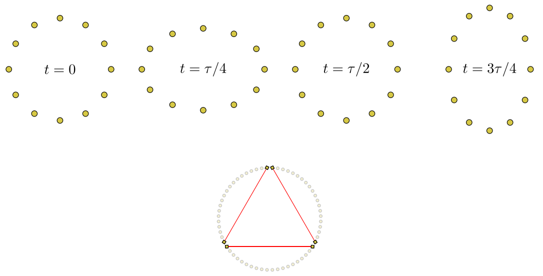

Figure 1. Effect of a $+$-polarized GW with period τ on a ring of test particles. The LISA test masses can be thought of as a subset of such a ring, shown below.

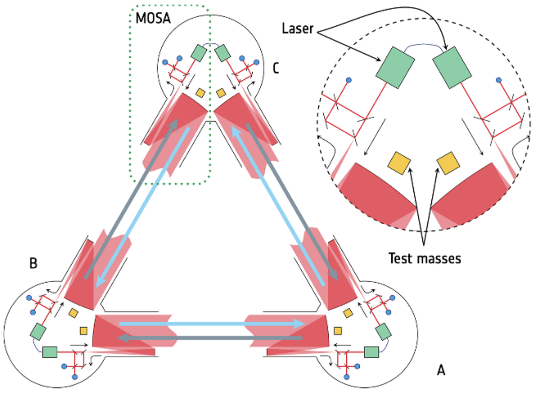

Figure 2. Simplified diagram of the LISA constellation, retrieved from the LISA Definition Study Report released by the ESA. Reproduced with permission from [1].

The constellation of satellites that make up LISA will breathe. The arm lengths of the ISI are expected to vary by up to $\sim$1.2%, and the angles of the triangular formation will therefore also vary. The spacecraft must then include movable optical sub-assemblies (MOSAs) to track these angular changes.

A reference link between the MOSAs relates the local phase measurements on one spacecraft, and each MOSA additionally has an interferometric readout of the local displacement of the test mass. The distances between two test masses along one inter-spacecraft link is measured by combining these interferometric readouts to produce the LISA observables. A basic diagram of the constellation, retrieved from the LISA Definition Study Report released by the ESA [1], is shown in figure 2.

The near-equal-arm condition that decouples the laser frequency noise from the measurements must then be recreated in post-processing using time delay interferometry (TDI) [6]. This technique shifts the measurement taken on one satellite backwards in time relative to those of the other satellites until the carrier noise common in those measurements is minimized.

In order for this to sufficiently suppress the laser carrier frequency noise, the timestamps of each measurement must be extremely precise. To achieve the required noise cancellation, the appropriate delays need to be known to the nanosecond level [6]. These delays contain not only the light travel times, but also time stamping offsets between the satellite clocks. A way to measure this so-called pseudorange, d, must be included in the inter-satellite link.

In addition to accurate timestamping, the in-band performance of these clocks is relevant as they directly couple to the inter-satellite measurement. This precision requires an in-band clock stability on the order of $10^{-13}\mathrm{s}/\sqrt{\mathrm{Hz}}$ [7]. The USO's foreseen for LISA do not reach this stability. Furthermore, the propagation of the jitters through the components of the spacecraft would add timing inaccuracies. As a consequence, the clocking jitters must be measured and corrected for.

1.2. Simulating the inter-satellite links

1.2.1. Inter-satellite link composition

A single inter-satellite link, for example between Spacecraft 1 (SC1) and Spacecraft 2 (SC2), consists of the following main elements:

SC1 generates a resonator-stabilized laser carrier at $f_{\mathrm{C12}} \approx 280\,\mathrm{THz}$, which is phase-modulated by an electro-optic modulator (EOM) using a clock-derived tone (e.g. $f_{\mathrm{SB1}} = 2.4\,\mathrm{GHz}$) from its local phasemeter. This modulation is further supplemented by a ranging code comprised of pseudo-random noise, and optionally with further out-of-band ranging tones [8]. This modulated signal is split into two paths: one is transmitted to SC2, while the other is interfered with the signal incoming from SC2.

SC2 similarly generates its own laser carrier and applies a slightly different clock sideband (e.g. $f_{\mathrm{SB2}} = 2.401\,\mathrm{GHz}$) using its local phasemeter. Both signals travel the inter-satellite distance, experiencing a time delay d_{12} ($8.3 \pm 0.1\,\mathrm{s}$), and are subject to Doppler shifts due to spacecraft motion and phase modulations from passing GWs.

The signal generated at SC2 at time $t - d_{12}$ arrives at SC1 at time t, where d_{12} is the light travel time from SC2 to SC1 (subscript convention: receiver-sender). Thus, any phase or frequency information from SC2 is delayed by d_{12} when it is received at SC1, including the modulation (such as clock sidebands) applied at SC2 at time $t - d_{12}$.

However, the GW-induced phase modulation is accumulated continuously along the light path as the signal propagates from SC2 to SC1. This means that the total GW effect measured at SC1 at time t is the integrated response of the GW along the path between SC2 at $t - d_{12}$ and SC1 at t. The GW does not simply imprint its effect at the moment of emission; instead, its influence is built up over the entire light travel time. To avoid this explicit integration, one can simulate the effect of the GW as a Doppler frequency shift, shown later in section 2.1.

While the laser carrier encodes the GW information, the GHz sidebands are needed to correct the clock jitter. Upon arrival at the distant spacecraft, the carrier-carrier beat note provides the measurement of the GW, while the sideband–sideband beat note provides a measurement of the sender's clock jitter relative to the receiver's local clock.

At the receiving spacecraft (SC1), the incoming delayed field is interfered with the local field on a photodetector. The resulting electronic signal contains the following terms of interest:

- The carrier-carrier beat note, which is sensitive to the relative motion of the spacecraft and to GW signals. This is kept within the 5–25 MHz range by feedback control.
- The sideband–sideband beat note, which is offset by 1 MHz from the carrier beat note and is used to measure the relative clock jitter between the two spacecraft.

The photodetector output is split and sent to multiple phasemeter channels. The channel targeting the carrier beat note provides the **science measurement**, while the channel targeting the sideband beat note provides the **clock measurement**. The difference between these two channels isolates the relative clock jitter, which can then be corrected in post-processing, as shown in the following derivation.

The carrier science measurement on SCi of the incoming field from SCj, after interference with the local field, can be written as

$$\mathrm{sci}_{C,ij}(t) = \Delta\omega_{ij} t + \phi_j\left(t - d_{ij}\right) - \phi_i(t) + \delta\phi_{\mathrm{GW}}(t), \tag{1}$$

where $\phi_i(t)$ and $\phi_j(t)$ are the phases of the local and remote lasers, $\delta\phi_{\mathrm{GW}}(t)$ is the accumulated GW-induced phase along the path. The $\Delta\omega_{ij}$ term accounts for the Doppler frequency shift due to the relative motion of SCi and SCj, and represents the nominal phase evolution of the carrier beat note over time. The ϕ terms encode the relative phase noise of each laser field relative to this nominal evolution.

The two sideband measurements—upper and lower—can be similarly expressed as

$$\mathrm{sci}_{\mathrm{SB}+,ij}(t) = \Delta\omega_{ij} t + \Delta\omega_{\mathrm{SB},ij} t + \phi_j\left(t - d_{ij}\right) - \phi_i(t) + \delta\phi_{\mathrm{SB,\,GW}}(t), \tag{2}$$

and

$$\mathrm{sci}_{\mathrm{SB}-,ij}(t) = \Delta\omega_{ij} t - \Delta\omega_{\mathrm{SB},ij} t + \phi_j\left(t - d_{ij}\right) - \phi_i(t) + \delta\phi_{\mathrm{SB,\,GW}}(t), \tag{3}$$

where $\Delta\omega_{\mathrm{SB},ij}$ is the frequency difference between the clocking sidebands of SCi and SCj as measured on the receiving spacecraft (i, in this case). It can be expressed as the sum of the designed frequency difference and the drift of the clock on SCj in the timebase of the clock on SCi:

$$\begin{aligned}\Delta\omega_{\mathrm{SB},ij} &= 2\pi\cdot\big(\underbrace{2.401\,\mathrm{GHz}}_{f_{\mathrm{SB},i}} - \underbrace{2.400\,\mathrm{GHz}}_{f_{\mathrm{SB},j}}\big) + \frac{\mathrm{d}}{\mathrm{d}t}\delta\phi_{\mathrm{CLK}_{ij}} \\ &= 2\pi\cdot\Delta f_{\mathrm{SB},ij} + \frac{\mathrm{d}}{\mathrm{d}t}\delta\phi_{\mathrm{CLK}_{ij}},\end{aligned} \tag{4}$$

where $\delta\phi_{\mathrm{CLK}_{ij}} = \phi_{\mathrm{CLK}_j}(t - d_{ij}) - \phi_{\mathrm{CLK}_i}(t)$ is the phase noise of the clock on SCj transmitted to SCi and compared to the phase noise of the clock on SCi.

The GW term is scaled by the slightly different clock sideband frequencies, but the difference is relatively small, and so to first order it can be written as $\delta\phi_{\mathrm{SB,\,GW}}(t) \approx \delta\phi_{\mathrm{GW}}$.

The measurement of the clocking noise between the two spacecraft is then given by the difference between upper and lower sideband measurements, which suppresses sideband readout noise:

$$\mathrm{sci}_{\mathrm{SB}+,ij} - \mathrm{sci}_{\mathrm{SB}-,ij} = 2 \cdot \Delta\omega_{\mathrm{SB},ij}. \tag{5}$$

The relative phase noise of the clocks comes from integrating this signal over time, and isolates the relative clock jitter between SCi and SCj given in equation (4), allowing for the correction of timing noise in the science measurement.

Combinating these measurements with the local test mass interferometer (TMI) readout and the reference interferometer (RFI) readout produces, for one MOSA, the combined transponder signal η_{ij}, which is used to construct the TDI combinations. In the context of this paper, where only the science interferometer is under examination, we neglect the TMI and RFI measurements[4].

In summary, the distance between satellites ensures that all information from SC2 is received at SC1 with a delay, so the measurement at SC1 at time t contains the state of SC2 at time $t - d_{12}$, *plus* the accumulated GW effect along the path. This explicit time-shifting and path integration are fundamental to the operation and simulation of the inter-satellite links, and the η_{ij} terms formalize the beat note measurements used in TDI. GW signals are primarily extracted from the carrier beat note, while the sideband beat note serves to monitor and remove timing noise.

In order to simulate these links, a hardware testbed is being developed that, in its first iteration, will implement the following:

- Generate a signal source that mimics the laser light on each spacecraft, including clock sidebands
- Maintain any clocking sidebands applied to this carrier signal
- Apply a time-variable delay to simulate inter-spacecraft distances
- Apply Doppler frequency shifting corresponding to the simulated spacecraft motion
- Inject an arbitrary GW strain into the delayed field
- Mix the delayed field with the local field to produce the relevant beat notes at the phasemeter
- Provide a phasemeter output that can be used to construct the TDI observables

1.2.2. Hardware simulations of inter-satellite links

Tests of TDI including representative inter-spacecraft delay in hardware setups have been conducted numerous times previously [9–11]. This work builds upon those, but with an eye towards integrating the delay lines into a more comprehensive hardware simulation testbed. The most challenging aspect of this task is precisely applying the appropriate delay to a signal. Optically implementing such delays is completely impractical in an Earth-bound laboratory—just one light-second is nearly 7.5 times the circumference of the planet. Instead, we rely on a hybrid approach in which the delay is applied digitally.

For this, the signal is sampled by an analog-to-digital converter (ADC), digitally delayed in an electronic memory device, and subsequently regenerated by a digital-to-analog converter (DAC) mixed with a local signal before being transmitted to a receiving phasemeter. The core component of the setup is a combination of field-programmable gate array (FPGA), as a typical computer processor would not be fast enough, and several gigabytes of high-speed random access memory (RAM) to store the signal.

The devices procured to serve as the delay lines in this experiment are ZCU208 evaluation boards manufactured by AMD Xilinx. The FPGAs on these boards is a ZU48DR Zynq UltraScale+ radio frequency system-on-chip Gen 3. This chip incorporates eight 14-bit ADCs and DACs, each capable of sampling rates ranging from 5 to 10 GHz. The system-on-chip is a PS that can run a lightweight operating system such as Linux.

The processing system (PS) has been loaded with a Linux-based Jupyter server platform called PYNQ, developed by AMD[5]. This allows communication with the programmable logic (PL) over Ethernet via memory-mapped input–outputs (MMIO). The PL is programmed to handle the data acquisition and processing, while the PS is used to control the PL and provide a user interface. Doppler

---

[4] We do not need to split our simulated spacecraft into two MOSAs to track distant satellites, so no RFI is needed, and we have no test mass, so no TMI.

[5] URL: www.pynq.io.

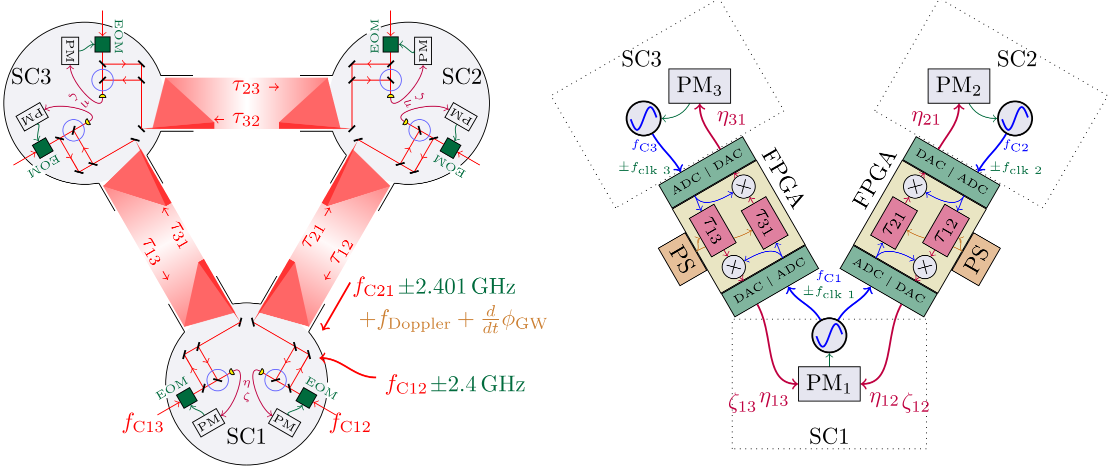

Figure 3. A side-by-side comparative diagram of the hardware simulation setup for the LISA delay line contrasted with the satellite constellation. The left-hand side shows a simplified diagram of the LISA constellation, with the inter-satellite links and their constituent frequencies and delays. The right-hand side shows a simplified diagram of the hardware simulation setup, with the components described in section 1.2.2. The PMs are phasemeters, and the EOMs are electro-optical modulators. The central beamsplitters of the inter-spacecraft interferometers are circled in blue on the left; they are simulated via digital mixers in the hardware simulation. Satellite motion is controlled and waveforms are injected via the processing system (PS). The third arm is omitted in the hardware testbed, as is the test mass-to-optical bench measurement. These will be implemented in future development.

shifts and GW phase modulations are applied via a buffer loaded by the PS and read by the PL. The PL is programmed in VHDL, a hardware description language.

The evaluation board is supplied with a CLK104 RF clocking board featuring shielded electronics that provide a stable timebase for the ADCs and DACs. It can be programmed via the PS to construct custom clocking structures. However, for the purposes of this work, it will be programmed to provide a sampling clock at a frequency of 2.048 GHz. This clock can be generated either from its internal reference oscillator or by accepting a standard 10 MHz reference tone. The CLK104 outputs a 7.68 MHz SYSREF clock tone that can be utilized to monitor the stability of the clocking system.

Additionally, the board is equipped with four 4 GB units of DDR4 RAM—8 GB for the PS, and 4 GB for the PL. The former of these can be used for on-board computation of signal offsets, as will be discussed later in section 2.1, while the latter is used entirely for the delay buffers via a series of state machines.

A diagram of the LISA constellation similar to that in figure 2 is shown on the left-hand side of figure 3, where emphasis is placed on the composition of the inter-satellite interferometer signals and how they are simulated in hardware. The right-hand side of the figure shows a simplified diagram of the hardware simulation setup, with some of the components described in the previous subsection. At this point in time, only two of the three inter-satellite link pairs are simulated, between SC1 and SC2, and between SC1 and SC3, which are required for the TDI Michelson X combinations.

The third link pair, between SC3 and SC2, is not yet implemented, as the TDI Michelson combinations require only two link pairs, and the central spacecraft can arbitrarily be called SC2 or SC3 for the Y or Z combinations, respectively. In future implementations, a third board will be used to simulate the third inter-satellite link pair. This will provide a testbed for arbitrary TDI combinations, such as the Sagnac channels, which require loops around the constellation.

The THz laser carrier signals are replaced by electronic signals of 10..60 MHz, the limits of which are determined by the analog front-end of the FPGA; the ADCs are AC-coupled, and the Nyquist limit is 64 MHz, which will be discussed in section 2.2. The local measuring phasemeters generate a tone derived from their sampling clock. This tone is used to phase-modulate the MHz carrier signal, which is then sent into the ADC of a delay line and split into two paths. One path is sent to the state machines that place it into RAM to be recovered later, and the other path is sent to a digital mixer that simulates the beamsplitter of the interferometer. The other spacecraft's signal—on the output of the converse delay—is also sent to this mixer, and its output is sent to the DAC to be synthesized and fed to the local phasemeter, building the η_{12} signal. The other η_{ij} signals are generated in a similar manner, with the other delay lines and their associated phasemeters.

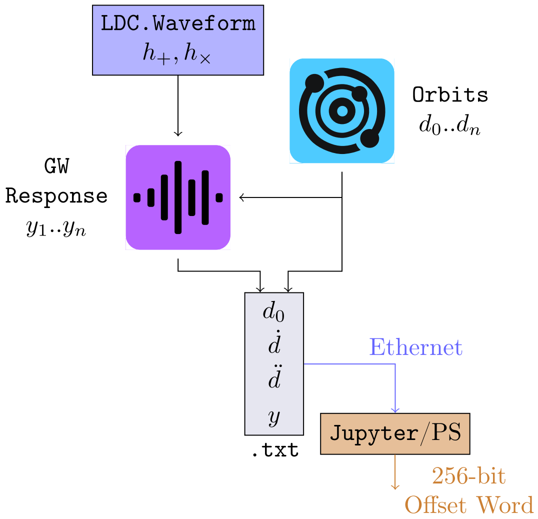

Figure 4. Inputs from orbital and waveform modules contributing to the measurable instrument response. The observable quantities d_0, $\dot{d}$, $\ddot{d}$, and y_{ij} are written to a text file to be uploaded to the PS. The initial ranges are sent for the initial delay times; everything afterwards is handled via the derivatives. Reproduced with permission from [13].

2. Delay line

2.1. Actuating the delay line

The values used to actuate the delay lines are generated via the codes `LISA GW Response`, `LISA Orbits`, and the LISA data challenge code [12–14]. The use of this codebase begins with defining a measurement length, which then gives an array of timestamps in the rest frame of the Solar System barycenter. These timestamps are used to generate a time series of simulated orbits for all three spacecraft, either Keplerian or OEMO[6], and then generating a time series of the pseudoranges $d_{ij}(t)$ and their first and second derivatives $\dot{d}_{ij}(t)$ and $\ddot{d}_{ij}(t)$ [14].

The same timestamps are given to the IMRPhenomD model [16–18] to calculate the local measurable strains, h_+ and $h_\times$, of any number of GWforms from source parameters such as ecliptic coordinates and redshift. Each of these are combined with the orbital values to calculate the interferometer response to each GW as a relative frequency shift in each arm, given by

$$y_{ij}(t_i) \approx \frac{1}{2\left(1-\hat{k}\cdot\hat{n}_{ij}(t_i)\right)} \times \left[H_{ij}\left(t_i - d_{ij}(t_i) - \frac{\hat{k}\cdot\vec{x}_j(t_i)}{c}\right) - H_{ij}\left(t_i - \frac{\hat{k}\cdot\vec{x}_i(t_i)}{c}\right)\right] \tag{6}$$

where H_{ij} is the antenna response to both polarizations of the GW, $\mathbf{x}$ are the positions of the satellites, $\hat{\mathbf{k}}$ the propagation vector of the GW, and $\hat{\mathbf{n}}$ the propagation vector of the laser light in the $j \rightarrow i$ link [13]. These are all summed together in a superposition to create a time series of relative frequency shifts to be fed to each delay line.

The proper pseudoranges, their derivatives, and the GW strain-induced frequency shifts are compiled into a text file which is then uploaded to the Jupyter server[7] running on the PS, as shown in figure 4. As there is limited memory available to the PS, timestamps are at a relatively low cadence. This is usually set to the proposed LISA final sampling rate of 4 Hz. For a measurement length of 1 day, this results in 345 600 sets of values to feed to each delay line. This results in a file size on the order of $\mathcal{O}(10)$MB for a single day of data. With 8 GB of memory available on the PS, this allows for a measurement length of up to 100 d, which is sufficient for the purposes of simulating LISA-like orbits and their associated GW responses. Higher cadences can be used with shorter measurement lengths to simulate the response to a GW with a higher frequency, but this is not necessary for the current work.

The files are uploaded and simple Python scripts run in the PS to calculate the appropriate offsets for each delay line. The GW strain responses are scaled up by the proper pseudoranges and the supposed ‘THz’ carrier frequencies to get the appropriate frequency offsets of the MHz carriers actually present in

---

[6] OEMO is the Orbit Ephemeris Message Orbits [15].

[7] This is not necessarily online; a local network functions with Ethernet.

the delay lines. The same 'THz' assumption is made for the derivatives of the proper pseudoranges in calculating the frequencies for the Doppler shifting of the carrier and the clock sidebands.

These must then be converted into binary form. For a GW, the expected frequency shift supposing the THz carrier is on the order of μHz. Since frequencies within an FPGA are represented as fractions scaled to the streaming clock,

$$\texttt{GW\_SHIFT} = \texttt{int}\left(2^B \cdot \frac{f_{\text{shift}}}{f_{\text{DSCD}}}\right)$$

for B-bit precision, this is then $\mathcal{O}(\mu\text{Hz}/\text{MHz}) \approx 1 \times 10^{-12}$. The $\log_2$ of this ratio gives the required minimum bit depth of 42. The frequency shifts, for binary efficiency's sake, are therefore represented by 64-bit signed values where any such value is fractional with $1 \rightarrow 2^{64}$.

A caveat to this is that 64 bits are not required for all the information to be sent to the PL—the top 39 bits contain no relevant information other than a flat scaling. A 64-bit integer provides, with a 128 MHz clock, approximately 7 pHz of precision in GW modulation of the delay. With an assumed upper limit of the range at 100 μHz, this dynamic range requires only 25 bits which can then be padded later to give the correct fractional value. As such, a simple 32-bit integer can be used to represent this value, with the padding back to 64 bits handled in the PL:

$$\texttt{GW\_SHIFT} = \texttt{int32}\left(2^{64} \cdot \frac{y_{ij}}{f_{\text{DSCD}}}\right)$$

in the PS, followed by

$$\texttt{PIR\_GW\_OFFSET} <= \texttt{resize}(\texttt{GW\_SHIFT, 64})$$

in the PL, where PIR is short for Phase Increment Register and the `resize` function prepends the required bits[8].

The same is done for the Doppler shifting, which has a maximum value of approximately 20 MHz. 32 bits cannot cover a sufficient dynamic range to make its offsets comparable to the GW offsets. Instead, the derivative of the Doppler shift is computed, which is at most approximately 225 mHz s^{-1}. This requires a 36-bit number to cover the dynamic range provided by 64 bits, with the top 28 bits are masked away:

$$\texttt{D\_DOPPLER\_X} = \texttt{int}\left(2^{64} \cdot \frac{\delta f_{\text{Dopp, X}}}{f_{\text{DSCD}}}\right) \,\&\, (1 \ll 36) - 1.$$

This packs each value into the appropriate bit length. While still byte-aligned in memory, once it arrives in the PL it will be unpacked on the bit level.

The PS runs on its own dedicated 100 MHz clock, and cannot be synchronized to the rest of the FPGA. This is sufficient for setting constant values—like the initial delay time or the PLL controls—in the PL, but it does mean that the offsets need to be passed through a clock domain crossing (CDC) first-in first-out (FIFO) buffer to move them into the clock domain of the data fed through the chip[9]. For a single delay line, each element in this buffer takes the form of a 256-bit word, with the structure delineated in table 1.

The first 32 bits are used for the GW phase offset. The second field is 31 bits, used for the inverse of the absolute value of the pseudorange derivative, with an additional bit used to indicate whether the pseudorange derivative is positive or negative. The next three 36-bit fields are used for the differential Doppler shifts of the carrier and upper and lower clock sidebands, and the last 32 bits are used for a timestamp to align the offset with the clock driving the data. Zeros are padded to the top to align with the 256-bit word size required by the architecture of the DMA IP core.

A 32-bit unsigned counter increments by 1 for each clock cycle of the data streaming clock domain (DSCD). When it reaches the current timestamp value (the top 32 bits of the offset word after the zero-padding), a flag is raised to read the next offset word from the FIFO buffer. The various parts of the next offset word are then unpacked and used to actuate the delay line.

---

[8] While retaining signed two's complement format.

[9] The Xilinx AXIS Data FIFO IP core is used for this with asynchronous clocks. A programmable empty flag is used to signal the PS to send the next chunk of offsets.

Table 1. Structure of the 256-bit word used to buffer the delay line actuation parameters. The first 32 bits are used for the gravitational wave phase offset, the next 31 bits are used for the inverse of the pseudorange derivative, and the next bit is used to indicate whether the pseudorange derivative is positive or negative. The next three 36-bit words are used for the Doppler shifts of the carrier and upper and lower clock sidebands, and the last 32 bits before zero-padding are used for a timestamp.

Parameter	Bit range	# Bits
δF_{GW}	0–31	32
$\lvert\dot{d}\rvert^{-1}$	32–62	31
$\mathrm{sgn}(\dot{d})$	63	1
$\dot{F}_{\mathrm{Dopp,C}}$	64–99	36
$\dot{F}_{\mathrm{Dopp,USB}}$	100–135	36
$\dot{F}_{\mathrm{Dopp,LSB}}$	136–171	36
timestamp	172–203	32
0’s	204–255	51

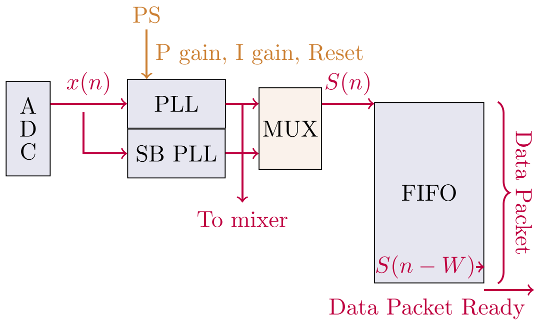

Figure 5. Simplified block diagram of initial signal processing described in section 2.2. The incoming signal is split into three copies. Two are sent two different PLLs to isolate the carrier and sideband signal components. A PS-controlled multiplexer determines which copy to send to a FIFO buffer, where W samples are grouped into a 256-bit word to match the RAM controller width. This packet is sent along with a flag to notify the command-issuing module of an incoming data packet.

2.2. Initial signal processing

A simplified block diagram of the initial signal processing is shown in figure 5. The ADCs and DACs on the FPGA are driven by a 2.048 GHz clock, sourced by the CLK104 timing board and locked by internal PLLs in one of the DAC tiles. This chosen DAC tile therefore drives the *DSCD*, as its output clock is fed to elements of the FPGA that carry the streaming data. The 14-bit ADCs average the data down 16 times, producing a 16-bit sample that is output to the DSCD at $F_{\mathrm{DSCD}} = f_{\mathrm{CLK104}}/16 = 128\,\mathrm{MHz}$. A copy of this signal can be sent to mix with the output of another delay line representing the signal coming from the other spacecraft, or directly to the DAC for noise floor comparison.

This is fed into a PLL which outputs its instantaneous accumulated phase as a 16-bit value, where $0..2\pi$ is mapped to $0..2^{16}-1$. The difference in phase between each sample and the previous one is calculated, so that the frequency of the incoming signal (scaled to a fraction of the clock frequency) is sent along. A second data pipeline concatenates this value with a 16-bit PLL output that tracks one of the local clock sidebands for timing transfer, which is crucial for reaching design sensitivity with TDI [19]. The pipeline can be selected via a multiplexer controlled by the PS.

This data sample is then packaged into a 256-bit (32-byte) word to match the bit width of the onboard RAM controller. If only a carrier is present, then this word width in samples is $W = 16$ samples, each 16 bits wide. If the concatenated pipeline is selected, then with each sample at 32 bits, $W = 8$.

The rate at which data is forwarded to the RAM is F_{DSCD}/W. For the carrier-only implementation, this is $f_{\mathrm{data}} = 128\,\mathrm{MHz}/16 = 8\,\mathrm{MHz}$. 32 bytes sent into RAM at 8 MHz is 256 MBps, and the RAM for a single delay line can easily accommodate over 10 s of delay. However, for the concatenated sideband implementation, the effective data rate is doubled. At 512 MBps, a hard maximum of 8 s is possible for the 4 GB of RAM available on one delay line using the ZCU208 board. Until further development occurs to expand the available memory of the system (e.g. by writing to an external memory unit), this implementation would need to run for shorter delay times, or the initial signal data would need to be downsampled by a factor of 2 and re-interpolated back up to 128 MHz after the delay but before the actuation. In either case, this is a limitation of this one specific platform, and not of the methodology or the technology as a whole.

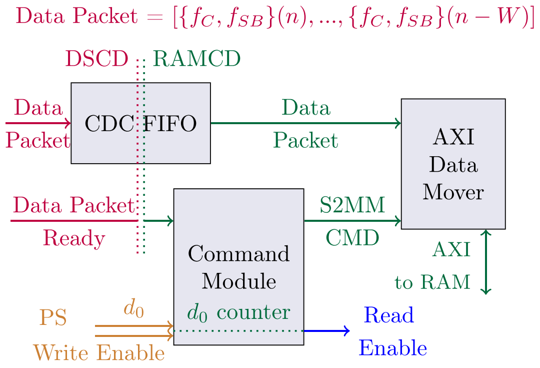

Figure 6. Simplified block diagram of the input to the delay buffer, described in section 2.3. A `ready` signal accompanies each packet of data for handshaking and proper timing. A S2MM command is generated and sent to the AXI Data Mover, which in turn sends the data packet to RAM. Sending `Write Enable` from the processing system (PS) starts a counter which goes to the programmed initial delay d_0 corresponding to the initial value of the proper pseudorange calculated from orbital simulations. Once the counter reaches this value, `Read Enable` is raised.

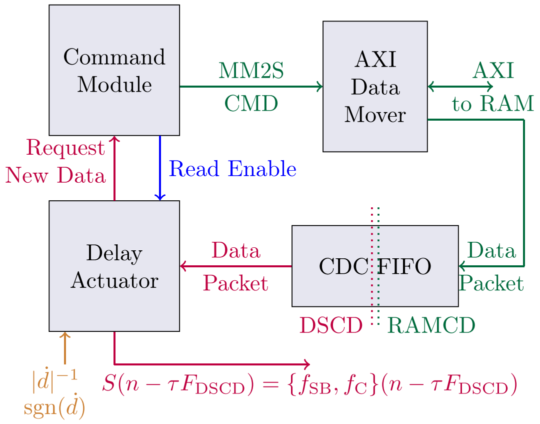

Figure 7. Block diagram of readout from RAM. After the initial delay has elapsed, the `Read Enable` flag is raised. The delay actuator module requests new data from the command module, which in turn generates a MM2S command for the RAM controller. The inverse of the proper pseudorange derivative is calculated and packed into a 31-bit integer representing a number of clock cycles in the DSCD, along with a sign bit. These are fed to the delay actuator, which changes the rate at which data is read from the RAM to sample precision—a shifting accuracy of $1/F_{\mathrm{DSCD}} = 7.8125\,\mathrm{ns}$. A CDC buffer is used to move the data from the RAMCD to the DSCD.

2.3. RAM

Bursts of these 256-bit words, along with an accompanying command, are sent to a Xilinx AXI DataMover module, which handles the interactions with the RAM controller, as shown in figure 6. The RAM chips and their controller module run on a dedicated 333.25 MHz clock, and as such, custom state machines were written to handle handshaking and flagging on either side of a CDC FIFO buffer so that the commands and word-bursts arrive with proper timing from the DSCD to the RAM clock domain (RAMCD).

No commands are sent until a `Write Enable` control signal is issued from the PS to begin writing to RAM. This control signal—and others, as discussed earlier—are written to MMIO registers from a Jupyter server along with an initial delay time. Once this command is sent, data is fetched from the CDC buffer and written to RAM when the Command state machine receives a flag that a new packet of data is ready. A counter begins to tick up, and once it reaches the value written to the `initial_delay_time` register, a flag is raised for `read` commands to be issued to the RAM controller.

Data is read from the RAM in bursts spanning a few microseconds of data, shown in figure 7. These pass through another CDC buffer back to the DSCD before filling a large FIFO buffer from which individual samples are unpacked and sent on. The state machine that handles the unpacking is used to actuate the delay times, and is fed values derived from the derivative of the proper pseudorange $\dot{d}$.

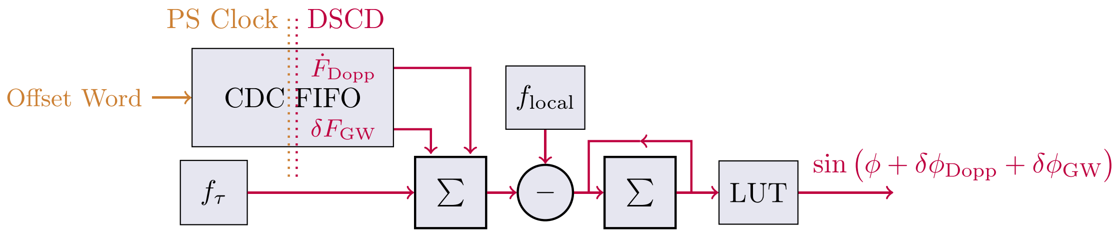

Figure 8. Reconstruction of the composite signal after retrieval from memory. Each of the delayed frequency components undergo the same treatment: Their respective Doppler and GW frequency shifts are applied; the value of the local frequency is subtracted, acting as a digital beamsplitter; the resulting frequency is accumulated into phase and fed into a LUT. The outputs of the LUTs are summed to reconstruct the full composite signal, where the $\Delta\omega$ values are the beatnote frequencies.

2.4. Simulating satellite motion

For simulating static arms, the `read` pointer of this FIFO buffer—the index of the array—increments by 1 each clock cycle. When the read pointer of this buffer reaches halfway, a flag requesting new data is raised (see figure 7) which triggers the state machines to refill the bottom half of the buffer. When the read pointer reaches the top of the buffer and loops back to the beginning, the same flag is raised and the state machines refill the top half of the buffer.

For moving satellites, the derivative of the proper pseudorange corresponding to a given delay line is generated from an orbital simulation. The absolute value of the inverse of this, $|\dot{d}^{-1}|$, is converted to a 31-bit integer. This value corresponds to the number of clock cycles to elapse before the next sample should be skipped or held an extra cycle. This is sent to the delay actuator with an extra bit indicating its sign. A counter in the delay actuator increments until this value is reached and then either skips the next sample if the sign is negative or holds the current sample for an extra clock cycle if the sign is positive. This skipping or holding of the data thereby changes, at a microscopic level, the macroscopic delay. With a 128 MHz data rate, skipping or holding for one clock cycle corresponds to a delay time shift of $1/F_{\mathrm{DSCD}} = 7.8125\,\mathrm{ns}$.

While this works for changing the travel time of the signals, this ‘motion’ does not apply the correct Doppler shifting to the signals that would mimic a LISA-like frequency shift; it would be scaled to the MHz electronic carrier frequencies, not the ∼280 THz frequencies of the LISA laser fields that they are representing. Additional frequency offsets need to be applied to the signals to simulate the Doppler shifts of the spacecraft relative to each other. As mentioned before, the correct values are pre-calculated and included in the offsets word. For each of the signal components, a 36-bit value is used to represent each of the derivatives of the proper offsets. These are integrated on-chip—accumulated each clock cycle—to give the proper frequency offsets as a 64-bit values, which are then applied to the values output by the delay actuator.

2.5. Reconstructing the signal

Some additional processing is required to reconstruct the signal with its clock tone sidebands, shown in figure 8. After the delay is applied to the frequency values and they are fed out from the delay actuator buffer (section 2.4), the value of the upper sideband is used to calculate the lower sideband frequency. The most efficient way to do this is to use the fact that the upper sideband frequency is related to the carrier frequency and the lower sideband frequency by

$$f_{\mathrm{LSB}}(n) = 2 \cdot f_{\mathrm{C}}(n) - f_{\mathrm{USB}}(n)\,. \tag{7}$$

Next, the Doppler shifts for each of these three frequencies are unpacked from the current offset word and added to their respective frequencies along with the pre-calculated GW strain response. This combined frequency can be directly sent on, or subtracted from a ′local′ undelayed signal:

$$f_{\mathrm{comb}}(n) = f_\tau(n) - f_{\mathrm{local}}(n) + \delta F_{\mathrm{Dopp}} + \delta F_{\mathrm{GW}}. \tag{8}$$

Accumulated phase values are running sums given by

$$\phi(n) = \phi(n-1) + f \tag{9}$$

for any frequency value, which is obviously a distributive operation.

The combined actuated carrier beat note to be output by the delay line is then given by

$$x_{\mathrm{C}}(n) = \sin\left(\Delta\phi_{\mathrm{C}}(n) + \delta\phi_{\mathrm{Dopp,\,C}} + \delta\phi_{\mathrm{GW,\,C}}\right). \tag{10}$$

The sideband output can be similarly regenerated with a programmable prefactor m providing a new modulation depth. Each sideband, then, is given by

$$x_{\mathrm{U/LSB}}(n) = m\sin\left(\Delta\phi_{\mathrm{U/LSB}}(n) + \delta\phi_{\mathrm{Dopp,\,U/LSB}} + \delta\phi_{\mathrm{GW,\,SB}}\right), \tag{11}$$

and the final actuated and remodulated signal is

$$x(n) = x_{\mathrm{C}}(n) + x_{\mathrm{USB}}(n) + x_{\mathrm{LSB}}(n). \tag{12}$$

Optionally, a pilot tone can be added here via direct digital synthesis (DDS) to track dispersion due to subsequent digital-to-analog or analog-to-digital interfaces (see section 3.1). The total superimposed sum of signals is then sent to the DAC.

3. Tests of the delay line

In the following sections, the *delay operators* $\mathbf{D}$ and $\mathcal{D}$ are mathematically identical. The former is used to denote physical delays present in the system, while the latter is used to denote the time shifts applied to a signal in post-processing. They are defined as

$$\mathcal{D}_{ij}f(t) = f\left(t - d_{ij}(t)\right), \tag{13}$$

where $d_{ij}(t)$ is, as mentioned earlier, the pseudorange between spacecraft i and j at time t. The delay operator $\mathcal{D}_{ij}$ is a linear operator and can be applied to any function $f(t)$, including the phase of a signal.

Delay operators can be nested to represent a chain of inter-satellite links:

$$\begin{aligned}\mathbf{D}_{ij}\mathbf{D}_{jk}f(t) &= f\left(t - d_{ij}(t) - d_{jk}\left(t - d_{ij}(t)\right)\right)\\ &= \mathbf{D}_{ijk}f(t).\end{aligned} \tag{14}$$

3.1. Baseline performance

To characterize the performance of the delay line, we use the experimental setup shown in the left side of figure 9. A signal (red) at frequency ω with phase ϕ generated by a Moku:Pro phasemeter picks up phase jitter $-\omega\delta t_{\mathrm{M}}(t)$, where δt_{M} is the timing jitter of the Moku sampling clock[10]. The delay line is programmed with a static delay τ; no offsets are applied. The signal is then sampled by the delay line board, picking up$-\omega\delta t_{\mathrm{D}}(t)$. If it is immediately sent back out of the delay line's DAC and measured (blue), these cancel out to first order. Otherwise, upon re-synthesis it picks up $+\omega\delta t_{\mathrm{D}}(t)$, but still contains the phase jitter $-\omega\delta t_{\mathrm{D}}(t-\tau)$. When digitized by the Moku phasemeter, it finally picks up $-\omega\delta t_{\mathrm{M}}(t-\tau)$ (purple). The phase of the measured (delayed) signal ϕ_τ with frequency ω is therefore given by

$$\begin{aligned}\phi_\tau(t) &= \phi(t-\tau) + \ \omega(t-\tau)\cdot\left(\delta t_{\mathrm{M}}(t-\tau) - \delta t_{\mathrm{D}}(t-\tau)\right) + \ \omega(t)\left(\delta t_{\mathrm{D}}(t) - \delta t_{\mathrm{M}}(t)\right)\\ &= \mathbf{D}_\tau\phi(t) + \ (1-\mathbf{D}_\tau)\,\omega\cdot\left(\delta t_{\mathrm{D}}(t) - \delta t_{\mathrm{M}}(t)\right),\end{aligned} \tag{15}$$

where physical electronic signal travel time (e.g. in cables) is assumed to be a constant offset to the programmed time and is thus subsumed within τ.

The lattermost term, the differential timing jitter, can be measured by generating a tone on the delay board via DDS and measuring it (green in figure 9), whereby it picks up the relevant jitters:

$$\begin{aligned}\phi_{\mathrm{DDS}} &= \omega_{\mathrm{DDS}}\cdot\left(t + \delta t_{\mathrm{D}}(t) - \delta t_{\mathrm{M}}(t)\right)\\ &= \phi_{0,\mathrm{DDS}} + \tilde{\phi}_{\mathrm{DDS}}\\ \rightarrow \phi_\tau(t) &= \mathbf{D}_\tau\phi(t) + (1-\mathbf{D}_\tau)\,\frac{\omega}{\omega_{\mathrm{DDS}}}\tilde{\phi}_{\mathrm{DDS}}(t).\end{aligned} \tag{16}$$

---

[10] These jitters, δt, couple in with a minus sign when a signal is *digitized* by an ADC and with a plus sign when it is *synthesized* by a DAC.

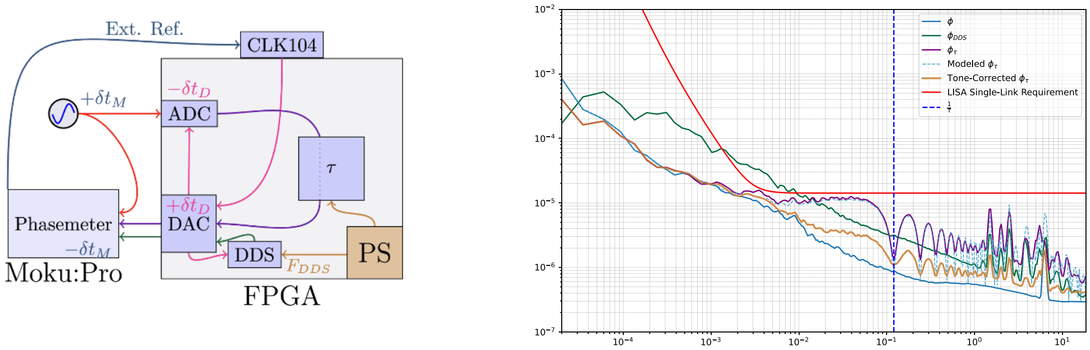

Figure 9. Left: Measurement setup for clock jitter analysis in the delay line, corresponding to equation (15). The CLK104 drives the board through one of the DAC tiles (magenta lines), and is fed an external reference by the Moku:Pro to provide a measurement with minimal differential clock noise. Timing jitter terms, δt, enter at interfaces between the digital and analog. Right: plot of phase ASD. In blue, the ASD of the measured phase of the generated signal, the green is that of the DDS signal (described in section 3.1), and the purple is that of the delayed signal. Applying equations (16) and (19) produce the bronze and teal dotted curves, respectively showing a reduction in noise and an agreement of the model to the measurement of the delayed signal. The LISA single-link phase noise requirement is shown in red for [20].

To analyze the noise in a measurement $x(t)$ and connect it to physical processes via power series in f, it is often convenient to look at its spectral density S, which is defined by

$$S[x](f) \triangleq |\tilde{X}(f)|^2, \tag{17}$$

where $\tilde{X}(f)$ is the Fourier transform of x. Theoretically, this requires infinite measurement time to have a full continuity of frequencies. The power spectral density estimate of the signal is practically given by

$$S = \lim_{T\to\infty} \frac{1}{T} |\tilde{X}_T(f)|^2, \tag{18}$$

where $\tilde{X}_T(f)$ is the Fourier transform of $x(t)$ windowed around the measurement time T.

Assuming low correlation between the measurements themselves, the spectral density of the measurement ϕ_τ is then approximated by

$$S[\phi_\tau](f) \approx S[\phi](f) + \left(\frac{\omega}{\omega_{\mathrm{DDS}}}\right)^2 4\sin^2(\pi f \tau)\, S[\phi_{\mathrm{DDS}}](f). \tag{19}$$

The $4\sin^2(\pi f \tau)$ factor comes directly from the squared Fourier transform of $1 - \mathbf{D}_\tau$. This is commonly referred to as the *delay comparison transfer function*, and frequently shows up in TDI calculations, where it is similar to how the full LISA constellation has no sensitivity at integer multiples of the light travel time [6].

Subtracting the rightmost term of equation (16) from the measurement of the delayed signal, and applying the model in equation (19) using the spectral density estimations of the DDS tone and of the original generated signal produces the plots shown in figure 9. The correction significantly reduces the noise level in the frequency band of interest (mid-mHz) and the model shows good agreement with the measurement of the delayed signal. As a benchmark (BM), we compare these results to the RMS of the design sensitivity of the LISA optical metrology system (OMS) and acceleration noise provided by [20], i.e.

$$S_{\mathrm{BM}} = \sqrt{S_{\mathrm{OMS}}^2 + 2S_{\mathrm{acc}}^2}, \tag{20}$$

where the factor of two represents the two test masses connected by the link.

In this context, the plot shows that the delay line does not add a significant amount of noise to the signal, and that this added noise can be largely accounted for and removed in post-processing. We assume that the remnant noise difference between the corrected spectrum and the undelayed spectrum is due to the perhaps false assumption that the ADC and DAC of the delay line board apply clock jitter identically—they could have differing levels of incoherent dispersion that would not cancel out. Furthermore, reaching LISA levels of performance is not the goal of this project, but rather the

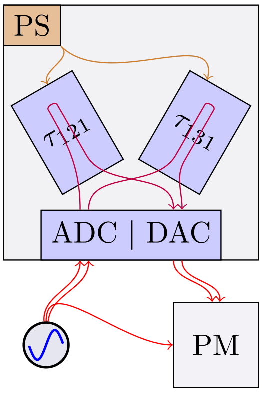

Figure 10. Measurement setup for static, single-board TDI tests. A signal generator is split into three paths. One is directly measured by a phasemeter. The other two feed a board programmed with two delay lines, each representing a round trip of different length. The PS can optionally feed offsets to these delay lines. Results of such measurements probe the current noise floor of the delay line experiments.

exploration of TDI and its ancillary required measurements. Fortunately, the noise of the tone-corrected measurement, shown in bronze in figure 9, is below the BM line—it is lower than the budget allocation for the OMS noise and the test mass acceleration noise for a single combined link between satellites, shown in red. The delay line is therefore suitable for use in the hardware testbed.

3.2. Static TDI tests

3.2.1. Readout noise floor

The next test to conduct is that of a static TDI configuration, in particular X_1. This is the TDI generation that corresponds to constant but unequal armlengths [21], and is given by

$$X_1 = (1 - \mathcal{D}_{121})(\eta_{13} + \mathcal{D}_{13}\eta_{31}) - (1 - \mathcal{D}_{131})(\eta_{12} + \mathcal{D}_{21}\eta_{21}) . \tag{21}$$

The η terms are the beatnote measurements collected on each spacecraft described in section 1.2.1. In orbit, they are naturally each only available on their corresponding spacecraft, i.e. η_{31} is only available on SC3. All η's need to be combined in post-processing. The first term is the link between spacecraft 1 and 3 transported along the round-trip path from 1 to 2 and back to 1. The second term is the link between spacecraft 1 and 2 transported along the round-trip path from 1 to 3 and back to 1. Doing so reconstructs two equivalent optical path lengths, which allows for frequency noise to be decoupled from the measurement.

As a simplified demonstrative example for testing, shown in figure 10, two delay lines were programmed into one board, each fed by the same external (and rather noisy) signal source. These can be thought of as equivalent to two separate optical paths that include a mirror reflecting each path back to the source—equivalent to the 'perfect transponder' model described in [21]. The lengths of the arms, as described above, are set to 8.1 and 8.5 s. In this model, equation (21) reduces with $\eta_{ij} + \mathcal{D}_{ij}\eta_{ji} \rightarrow \eta_{iji}$.

The internal mixer ('digital beamsplitter') is not used here; the signals are combined in post-processing to compare their relative phase evolution. Future similar measurements require the use of multiple signal generators to have different frequencies for each link, which in turn requires more phasemeter channels to fulfill the requirements of the TDI combination.

The undelayed and two delayed signal phases are measured as $\phi_1(t)$, $\phi_{121}(t)$, and $\phi_{131}(t)$, respectively. The η terms are then calculated as

$$\begin{aligned} \eta_{121}(t) &= \phi_{121}(t) - \phi_1(t) , \\ \eta_{131}(t) &= \phi_{131}(t) - \phi_1(t) , \end{aligned} \tag{22}$$

where now $\phi_{121}(t) = \mathbf{D}_{121}\phi_1(t)$ and similarly for $\phi_{131}(t)$. The η terms are then inserted into equation (21) to calculate the TDI combination X_1:

$$X_1 = (1 - \mathcal{D}_{121})\eta_{131} - (1 - \mathcal{D}_{131})\eta_{121} . \tag{23}$$

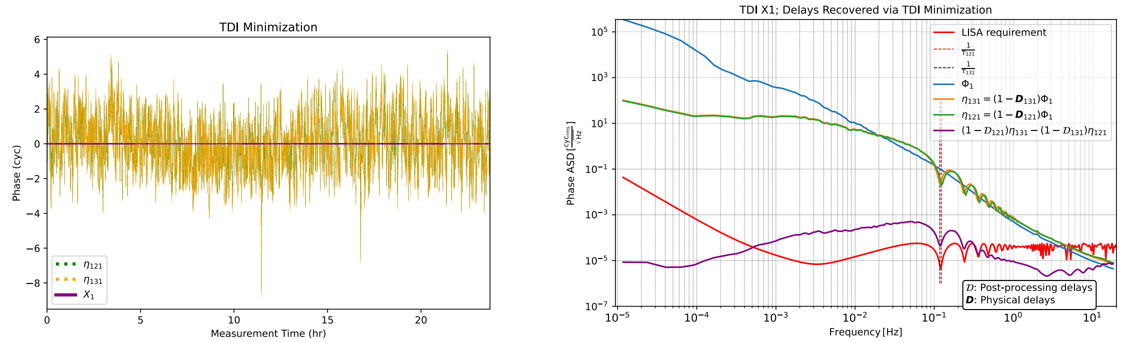

Figure 11. Results of the TDI X_1 combination, where the phase noise of the carrier signal, ϕ_1, is suppressed at the lower end of the frequency band by more than 10 orders of magnitude. The projected LISA phase noise requirement is shown in red as a benchmark, and the residual noise floor of the TDI combination is shown in purple. The individual constituent signals are shown in blue, orange, and green.

If we assume the frequencies of the constituent signals are approximately constant on the timescale of the delay, the constituent terms in this equation cancel to first order, and this measurement can be used to quantify the relative noise floor of the delay line for an external noisy signal source. The values of the pseudoranges are recovered without *a priori* knowledge by minimizing the spectral density approximation of the TDI combination calculated via Welch's method [6, 22].

The amplitude spectral density of the TDI combination is plotted in figure 11 along with those of the individual constituent signals. The projected LISA single-link phase noise requirement is show in red, and the residual noise floor of the TDI combination is shown in purple.

The TDI combination suppresses the phase noise of the carrier signal, ϕ_1, at the lower frequencies by more than 10 orders of magnitude. A large amount of noise is present across the frequency band from about 0.1 mHz to 0.25 Hz, which is unfortunately the region where LISA is most sensitive to GWs. This noise floor is likely dominated by readout noise in the measuring phasemeter coupled with errors in pseudorange estimation, but further investigation is required to determine its exact source.

Errors in the applied delays couple very strongly to the noise in the TDI combination. This can be ameliorated by increasing the sampling rate of the measuring phasemeter and by using stronger and more computationally expensive minimization routines for minimizing the TDI combination. Furthermore, a ranging tone applied to the initial generated signals would provide an accurate-enough initial estimate for the minimization routine, and future implementations of these experiments can probe the use case for such tones in the final flight configuration of LISA [8].

3.2.2. Injecting a GW signal

A time series of frequency offsets corresponding to the strain of a massive black hole binary (MBHB) merger calculated from the `IMRPhenomD` model set ([18]) is loaded into the RAM of the PS. Using the functionality described in section 2.1, these offsets are scaled by the programmed PPRs and fed into the PL to be applied to the delay lines.

Taking the expression in equation (23) and adding to the η terms the waveform strains injected into each arm gives the expected TDI measurement, shown in figure 11, as everything except the differential strain response is suppressed. This is given by

$$\mathrm{TDI}\left[\phi_{\mathrm{GW}}\right] = k \cdot c \cdot \left(d_{121} \cdot \left(1 - \mathcal{D}_{121}\right) h_{131}\left(t\right) + d_{131} \cdot \left(1 - \mathcal{D}_{131}\right) h_{121}\left(t\right)\right), \tag{24}$$

where k is the wavenumber corresponding to a 280 THz carrier and c is the speed of light.

Applying the appropriate TDI combination removes the carrier frequency noise from the data as expected, revealing the characteristic form of a binary merger in the spectral density estimation and the injected waveform in the time series, shown in figure 12. The spectral density estimation of the TDI combination is shown in the left panel, and the time series of the TDI combination is shown in the right panel. The same remnant noise floor is present in the TDI combination as in the previous test, and the successful subtraction of the injected waveform from the TDI combination suggests that the signal injection methodology is relatively noise-free.

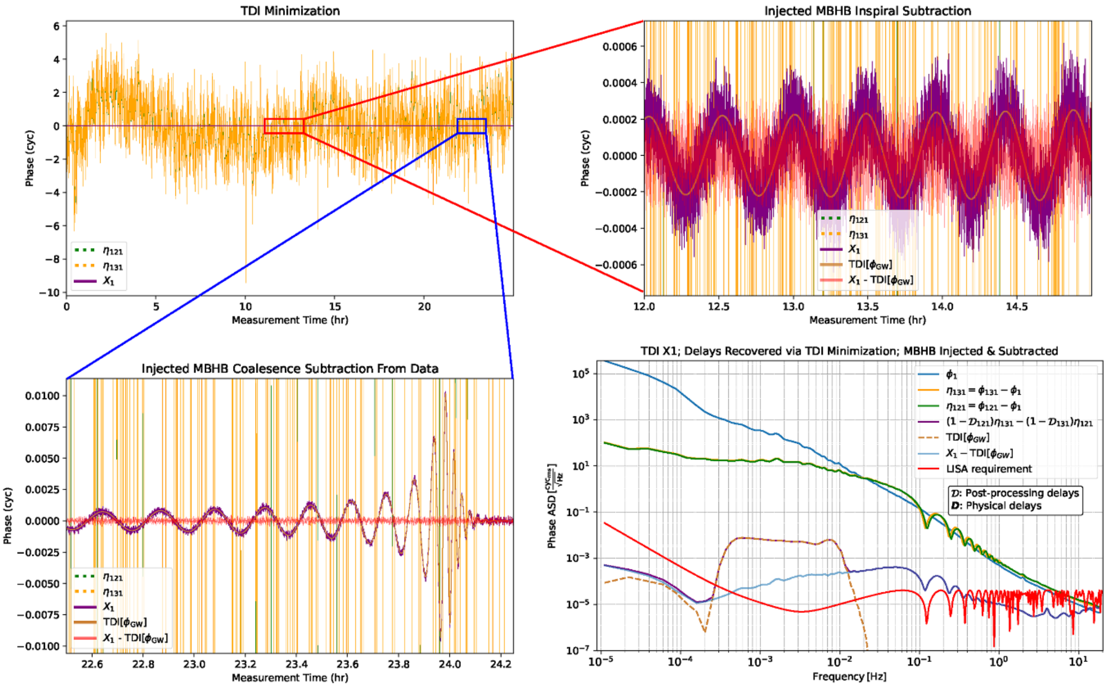

Figure 12. The appropriate TDI combination (equation (23)) suppresses the carrier frequency noise which would otherwise obscure the injected MBHB coalescence waveform. Top left: Full time series of the η composite measurements and the minimized TDI combination. Red and blue boxes indicate zoomed regions. Top right: zoom of the central region showing the inspiral revealed in the TDI combination, the injected inspiral waveform, and the subtraction of the latter from the former. Bottom left: zoom of the time around the coalescence, showing the same. Bottom right: spectral densities of these time series with the LISA benchmark scaled by the TDI X1 transfer function [20]. The residual noise floor described in the previous section is recovered, albeit with an uptick at the lowest frequency bin likely due to DC spectral leakage from the signal injection.

4. Conclusions and outlook

The delay line board has been successfully tested in a pure feedthrough measurement, where the added noise is dominated by the differential clock noise between the measuring phasemeter and the delay line board. The addition of the noise is well-modeled and can be largely removed in post-processing, should that be required, but the added noise is small enough that it does not significantly affect the measurements.

This added noise clearly scales with the relative frequency/phase noise of the signal source. The self-measurement baseline test where the Moku acted as both signal generator and phasemeter reaches the LISA single-link requirement by itself. The following tests using an external noisy signal source remain about an order of magnitude (at worst) above the LISA performance requirements. If necessary, the signal chain can be amended to include a pilot tone correction for the ADC jitter in the PL before the data is sent to memory.

The delay line board has also been successfully tested in a static TDI configuration, where the TDI combination X_1 is used to suppress the phase noise of a carrier signal by more than 10 orders of magnitude at the lowest end of the frequency band. The residual noise floor of the TDI combination is likely dominated by readout noise in the measuring phasemeter coupling with the higher frequency noise of the delayed signals, but further investigation is required to determine the exact source of this noise.

The delay line board has also been used to inject an arbitrary GW signal into the simulated inter-satellite links, which is completely obscured by the carrier frequency noise. The static TDI combination X_1 is then used to suppress this noise, revealing the characteristic form of a binary merger in the spectral density estimation and the injected waveform in the time series. The residual noise floor from the null test is recovered by subtracting the expected waveform measurement from the TDI combination of the measurements.

This residual noise floor could be arising from errors in the ranging estimation. As it currently stands, our analysis of the TDI combination involves a relatively simple routine that minimizes the power in the Welch spectrum of the TDI time series by shifting the constituent time series by constant values. It is possible that there is some non-constant delay that is unintentionally added by thermal drift in the cabling. Further measurements will include a dedicated $\sim 1\,\text{Hz}$ ranging tone, as earlier versions of this experiment implemented this to vastly improve the performance of their setup [11]. Such tests would be additionally useful in the exploration of modulation-assisted TDI ranging, proposed in [8].

The next tests to be undertaken are those of a dynamic satellite constellation, where rather than perfect round trips, each delay line is fed by a different signal source. A second board will be similarly programmed to provide the second pair of links, and the mixing of the link signals will be performed digitally in the PL. Orbital simulations will be used to generate the expected proper pseudoranges and their derivatives. This will in turn be used to generate an arbitrary waveform or superposition of waveforms, each with their own dynamic antenna response functions. The pseudorange derivatives and the projected strain will be fed into the simulated arms coherently to produce a realistic representation of the LISA system on a table top.

Such a realistic representation would also allow for the testing of clock tone transfer between the satellites, which is a crucial part of the pre-processing of the TDI observables, and has not yet been explored in a delay line setup. The clock tone transfer is necessary to ensure that the TDI observables are not contaminated by the clock noise of the measuring phasemeter. The testbed will be expanded to the stage shown in figure 3 to conduct these tests.

Further experiments are planned for these devices. The purely-electronic setup will be replaced with an electro-optic setup, where the beatnote between two cavity-stabilized lasers will be used as the primary carrier signal source. Two other lasers, each beating against one of the two stabilized lasers, will be used to generate the signals originating on the other ‘satellites’. Digital or electronic modulation will be replaced by electro-optic modulation to imprint each local phasemeter clock tone onto the respective carrier signals. The above list of tests will then be repeated with the optical setup and driven towards a robust and performant hardware testbed for LISA in the lab.

The scientific advantage of such a testbed should not be underestimated. It will allow for the testing of TDI combinations and their performance in a controlled environment, where the noise sources can be well understood and characterized. More specifically, it can show how the noise contributed by physical instruments can effect the scientific output of LISA in how it may effect the estimation of waveform parameters. It will also allow for the testing of new algorithms and techniques for TDI generation and processing, which can then be applied to the actual LISA mission.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Data access statement

The data and FPGA code supporting the findings of this paper are available from the corresponding author upon reasonable request. Simulation code and orbital datasets are based on publicly available software packages `LISA Orbits`, `LISA GW Response`, and the `LISA Data Challenge` framework.

Conflict of interest

The authors declare that they have no competing financial or non-financial interests relevant to the content of this work.

Author contributions

Reid Ferguson 0009-0006-2802-6570
Conceptualization (lead), Data curation (lead), Formal analysis (equal), Investigation (lead), Methodology (lead), Project administration (lead), Software (lead), Validation (lead), Visualization (lead), Writing – original draft (lead), Writing – review & editing (lead)

Guido Mueller 0000-0003-2262-4132
Conceptualization (supporting), Funding acquisition (lead), Project administration (supporting), Resources (supporting), Writing – review & editing (supporting)

Olaf Hartwig 0000-0003-2670-3815
Formal analysis (supporting), Methodology (supporting), Writing – review & editing (equal)

References

[1] Colpi M *et al* 2024 LISA definition study report (arXiv:2402.07571)

[2] Armano M *et al* 2018 Beyond the required LISA free-fall performance: new LISA pathfinder results down to 20μHz *Phys. Rev. Lett.* **120** 061101

[3] Meylahn F, Knust N and Willke B 2022 Stabilized laser system at 1550 nm wavelength for future gravitational-wave detectors *Phys. Rev. D* **105** 122004

[4] Mueller C L *et al* 2016 The advanced LIGO input optics *Rev. Sci. Instrum.* **87** 014502

[5] Ghosh S, Sanjuan J and Mueller G 2022 Arm locking performance with the new LISA design *Class. Quantum Grav.* **39** 115009

[6] Tinto M and Dhurandhar S V 2005 Time-delay interferometry *Living Rev. Relativ.* **8** 4

[7] Shaddock D, Ware B, Spero R and Vallisneri M 2004 Post-processed time-delay interferometry for LISA *Phys. Rev. D* **70** 081101(R)

[8] Bayle J-B, Staab M, Francis S P, Rees E R, Spero R and Heinzel G 2025 Modulation-assisted time-delay interferometric ranging for LISA *Class. Quantum Grav.* **42** 155009

[9] Thorpe J, Cruz R, Mueller G and Sankar S 2005 Time delay interferometry using the UF-LISA benchtop simulator *Class. Quantum Grav.* **22** S761–7

[10] de Vine G, Ware B, McKenzie K, Spero R E, Klipstein W M and Shaddock D A 2010 Experimental demonstration of time-delay interferometry for the laser interferometer space antenna *Phys. Rev. Lett.* **104** 211103

[11] Mitryk S J, Mueller G and Sanjuan J 2012 Hardware-based demonstration of time-delay interferometry and TDI-ranging with spacecraft motion effects *Phys. Rev. D* **86** 122006

[12] LDC Working Group 2024 LISA Data Challenge Software (LDC) Python package lisa-data-challenge, version 1.2.4, 2024, simulators and analysis tools for LISA mock datasets; please cite DOI (arXiv:2204.12142)

[13] Bayle J 2025 LISA GW Response: Python package for LISA time-domain gravitational-wave response Python package lisagwresponse, version 2.5.0, 2025, computes LISA link responses to gravitational waves and generates data files compatible with LISA Instrument and LISANode (available at: https://doi.org/10.5281/zenodo.8321733)

[14] Bayle J 2024 LISA orbits: Python package for LISA spacecraft orbit generation Python package lisaorbits, version 2.4.2, 2024, generates orbit files (state vectors, pseudoranges, proper time etc.) compatible with LISA Instrument, LISA GW Response, the LDC software, and LISANode

[15] Martens W 2021 Trajectory design for the ESA LISA mission *CEAS Space J.* **13** 409–21

[16] Pürrer M, Khan S, Ohme F, Birnholtz O and London L 2023 IMRPhenomD: phenomenological gravitational-wave model for non-precessing black-hole binaries ASCL code entry ascl:2307.019, 2023, part of LALSuite and standalone; calibrated against numerical relativity simulations

[17] Husa S, Khan S, Hannam M, Ohme F and Pürrer M 2016 Frequency-domain gravitational waves from non-precessing black-hole binaries. i. new phenomenological model imrphenomd *Phys. Rev. D* **93** 044006

[18] Khan S *et al* 2016 Frequency-domain gravitational waves from nonprecessing black-hole binaries. II. A phenomenological model for the advanced detector era *Phys. Rev. D* **93** 044007

[19] Hartwig O and Bayle J-B 2021 Clock-jitter reduction in LISA time-delay interferometry combinations *Phys. Rev. D* **103** 123027

[20] Babak S *et al* 2021 LISA sensitivity and SNR calculations (arXiv:2108.01167)

[21] Hartwig O 2021 Instrumental modelling and noise reduction algorithms for the laser interferometer space antenna *PhD Dissertation* Leibniz University Hanover

[22] Welch P D 1967 The use of fast fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms *IEEE Trans. Audio Electroacoust.* **15** 70–73