

Technical Note: P-ONE Inter-Module Timing

Nathan Whitehorn, Michigan State University

October 25, 2023

1 Introduction and Requirements

In order to achieve project requirements on angular resolution, we require synchronization of all ADC sampling at the 0.1 ns level or better (P-ONE Req. 1A). This, in turn, requires that all ADCs share a clock, or traceability of a clock, with a maximum deviation in frequency between modules in the array of 0.1 ns / run (10^{-15} for a one-day run) or less. It also requires that the modules be phase synchronized at the 0.1 ns level.

Like all subclass-1 JESD204B ADCs, the P-ONE PMT ADC (Analog Devices AD9083) takes an input clock in the few hundred MHz range and a single-pulse phase-synchronization signal (“SYSREF”) that sets the start of data taking. Timing is then kept by cycle-counting from the initial SYSREF pulse—that is the origin of the frequency-stability requirement above. As long as SYSREF arrives with a static phase delay relative to the system clock, ADC samples on all modules are guaranteed to be aligned to within some unknown long-term-constant delay and the small internal jitter of the ADC (the unknown constant delay sums with the unknown constant delay of the PMT mean transit time and can be calibrated out with it).

This static phase-delay requirement, and the requirement for very low phase noise on clocks for high-speed ADCs in general, normally results in the use of a dedicated, discrete PLL for generating both the ADC clock and SYSREF, as well as a locked reference clock for the SERDES unit in the FPGA that will receive the ADC data. Here, we use the Analog Devices AD9546, one of a broad class of devices built to create ADC system clocks and SYSREF pulses in a triggered way with respect to the clocks.

Synchronization of ADC sample timing then becomes a problem of knowing the time at which SYSREF was generated at each module relative to an array-global root clock to within the 0.1 ns precision requirement. This requires that the signal propagation times on all cables be known to within 0.1 ns at all times, which requires periodic in-situ measurements.

For an underwater detector, the ultimate source of time information comes from the data network (rather than GPS), using some standard like IEEE 1588 (the Precision Time Protocol). Network time, usually via some more local intermediate clock, can be used to generate a common clock and phase marker

at each module from common information (e.g. Ethernet clocks for frequency and some phase information embedded in Ethernet data).

To reduce power, system complexity, and jitter, some number of modules can also share the results of a common network-connected clock and directly share the inputs to their ADC PLLs, which is the subject of this document.¹ Because the inputs to the ADC PLLs are identical, they will maintain perfect synchronization so long as the PLLs are locked and the phase delays between PLL inputs are known.

2 Phase Delay Measurements

The ADC PLL we used (AD9546) is one of a family of devices from Analog and their competitors (notably Renesas/NEC and Silicon Labs) designed both to generate JESD204B clock and SYSREF signals and also to measure the phase skew between two clocks of identical frequency. This provides, in principle, an easy way to measure the round-trip delay between two PLLs: the one closer to the root clock (“lower”) sends a signal to the other (“up”), to which it locks and then mirrors it back to the other clock (“down”). The skew between the outbound and inbound signals on the lower clock is then equal to the round-trip delay of the link. (Mechanisms for assessing one-way delay are described in Sec. 4.)

Jitter on the PLL outputs is minimized if the reference clock frequency is high (100 kHz or higher) and such frequencies are also required by AC-coupled physical links. This presents some complications: for example, for a 125-MHz reference clock, the measured phase can only be determined within the 8 ns clock period and longer delays are not measurable. This compares poorly with the 300 ns round-trip delay even between neighboring P-ONE optical modules. To solve this problem, the AD9546, like other similar chips, supports embedding a lower-frequency clock inside a higher-frequency one using pulse-width modulation. Because the PLL locks to the rising edge of the input clocks, information can be encoded in the times of the falling edges, allowing the PLL to mark particular cycles as special (“tagging” in the AD9546 documentation); these can then be used as phase markers both for skew measurement and for triggering SYSREF generation (Fig. 1).

This solves two problems: it allows the measurement of arbitrarily-long cables without degradation in the clock frequency rate by measuring skew between the tagged samples and it allows a coarse phase marker to be unambiguously associated with a TAI time. For example, if the modules only know the absolute TAI time to 1 ms precision (typical of NTP), they would not know the correct global time associated with an ADC SYSREF pulse if the clock pulses arrive at 1 kHz or higher. By measuring time offsets relative to a slow ($\lesssim 100$ Hz) set

¹The existence of a network-connected clock meeting the requirements above (10^{-15} frequency stability relative to other network root clocks in the array over 1-day periods and, equivalently, 0.1 ns timing uncertainty relative to other network root clocks) is assumed and its details are not discussed here.

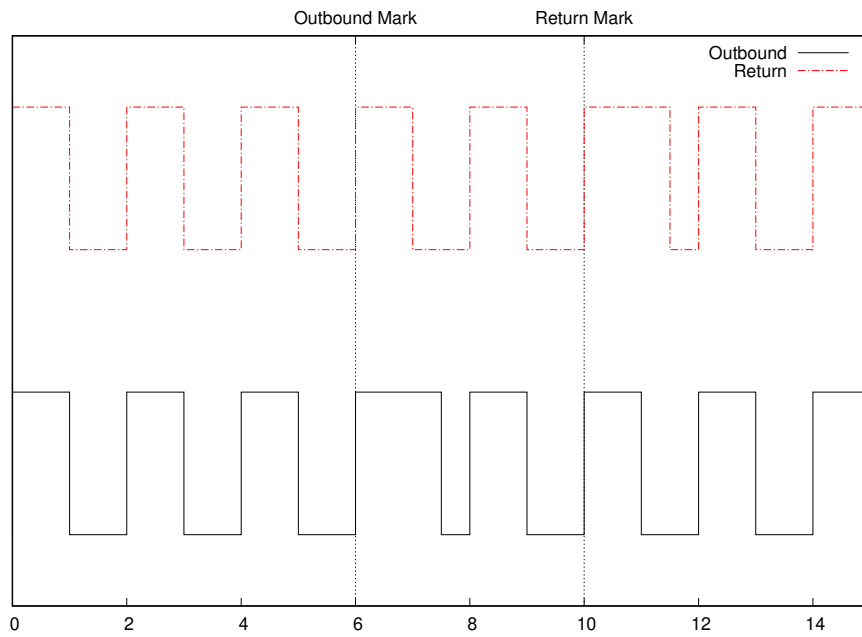


Figure 1: PWM skew measurement. The falling edge of one pulse is delayed to mark it as a synchronization event; this embeds the lower-frequency synchronization clock into the higher-frequency reference clock. The round-trip cable delay is measured as the time difference between the outbound and inbound marked pulses.

of phase-synchronization pulses, the absolute TAI time of each ADC sample—to within some microsecond-scale array-global offset from the root clock—is known.

3 Summary of System Architecture

This simple phase-measurement technique suggests a clock distribution architecture where one PLL is synchronized to a local external clock (Sec. 5) and some number of other ADC PLLs lock to its outputs. This architecture meets our requirements: it is simple, low-power, and precise simply because the ADC PLLs are directly connected and it involves no additional components except potentially physical-layer transceivers (Sec. 4).

Within this clock domain, the PLL directly connected to a network clock produces a number of copies of a frequency reference, with embedded phase markers at some interval larger than the software network time uncertainty at the modules and the maximum round-trip delay. Each ADC PLL locks directly to this reference, producing an ADC clock locked to the reference frequency and SYSREF pulses synchronously with the embedded phase markers. Propagation delays from the PLL connected to the root clock are measured by having the ADC PLL return a copy of the original signal—operating in zero-delay mode—and measuring the skew between the outbound and return signals at the root PLL (Fig. 1). Typical reference frequencies are between 200 kHz and 125 MHz, with the phase marker present at 100 Hz.

4 Physical Layer

The clock signals generated and received by the PLL are differential square-wave clock signals and can be transported by any interface between two PLLs. Two possible mechanisms are described below; these are entirely the same except for the choice of transceiver and the distances of cabling allowed as a result.

4.1 Copper M-LVDS

One possibility is to use a long-distance differential signalling method on twisted-pair cable. M-LVDS, which is designed for MHz-scale signalling over tens of meters of twisted-pair cable and allows multi-drop topologies, is the clear choice here.

Because the signalling medium is not simultaneously bidirectional, the return signal must operate on a different pair than the reference clock. This introduces the potential of asymmetric delays between the two pairs and attendant skew. Periodically reversing the direction of the return, however, allows measurement of this asymmetry. The round-trip measurement outlined above allows the measurement of the sum of the delays on the two pairs (down and up). Sending a signal up on both pairs and measuring the skew between the arrival of the two at the higher PLL allows a measurement of the difference of the delays.




Figure 2: Effect of EMI on clock signalling (green) over 50 meters of Macartney quad compared to initial signal (yellow), without signals on adjoining pairs (left) and with 10BASE-T1L Ethernet on an adjoining pair (right). The right-hand figure shows a significantly degraded signal that periodically loses timing lock.

Having both the sum and the difference then allows the measurement of the unidirectional delay.

Using the ADN4680E M-LVDS transceiver, which provides this capability in a low power envelope, dissipates approximately 300 mW per module. The high-frequency loss of the Macartney twisted-pair cable limits the reference clock frequency to no more than ~ 500 kHz, but this is sufficient to provide consistent 20-ps scale synchronization between neighboring modules in a low-EMI environment. The poor EMI rejection of the Macartney cables, however, makes this system very vulnerable to EMI (Fig. 2). In addition, high loss, even at 100-kHz-scale frequencies, limits the range of the signals to 50–100 meters. This system thus requires multiple network root clocks per line (with attendant increased power draw) and/or the use of active repeaters between modules.

The topology in a copper system is required by these considerations to involve network-attached leader clocks embedded in every few modules, with 2–4 intermediate modules sharing the signals from a given network clock and repeating the copper signal down the cable. A diagram is shown in Figure 3.

4.2 Fiber Optic

An alternative approach is to use fiber optic transport to the modules, which avoids the EMI and high-frequency loss problems of the copper cables. Frequency domain multiplexing, using bidirectional optics, also allows the up- and down-going signals to share a medium, reducing the link asymmetry and obviating the need for sum-and-difference delay measurement. Instead, the one-way delay, to within the same electronics constants that were present in the copper system, can be computed by dividing the round-trip delay by 2. This removes the EMI, propagation-distance, and loss limitations of the copper system entirely.

Common fast-Ethernet-capable bidirectional gigabit SFP modules can be

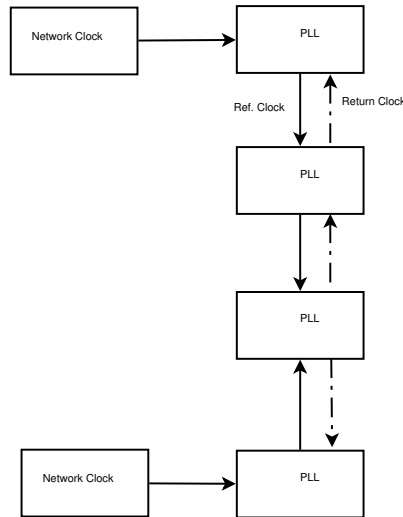


Figure 3: Topology for a system using copper transport. Clock signals are actively repeated by each PLL on modules in a chain, with network-attached root clocks present periodically (here, every 4th module).

used by using a 125 MHz reference frequency with the same 100 Hz phase marker described for the copper system, constrained so that the pulse-width distortion is $> 10\%$ to keep the signal bandwidth below that of gigabit Ethernet.

Power usage can be minimized by using a passive optical splitter from the lead clock, which enables only one SFP to be required per module. The resulting time delays are extremely stable and precisely measured ($\lesssim 0.5$ ps, Fig. 4), allowing individual leaf clocks to be measured only periodically and to have their return signals to the root clock to be time multiplexed. Because the integration period to measure phase delays is short (1 second in Fig. 4) and the need for it is intermittent (minutes to hours at most), return-signal multiplexing can be coordinated coarsely using out-of-band control signals. As with the copper transport, the measured time delays can likewise be communicated out of band from the root clock to the leaf clock. This reduces power usage to 250 mW per module using Waystream bidirectional SFPs, lower than the copper system. Typical 20-km bidirectional single-mode SFPs can drive up to 32 receivers over 1 km of fiber.

The passive splitter also reduces the number of intermediate clocks, improving overall system performance, and allows for redundancy in the leader clock by using a 2:N passive splitter with multiple inputs. A diagram is shown in Figure 6.

The elimination of the cable length considerations and use of a tree-type optical splitter suggests a topology with a single network-attached leader clock per line, located in the string junction box, and connected to a 2:20 optical splitter with one fiber terminating in each module. This arrangement, by lowering




Figure 4: Timing data from internal optical timing distribution during lab testing at MSU over a few-day period. This shows the measured delay variation from the root clock to a readout node over a 2:24 optical splitter and 220 meters of fiber. RMS uncertainties at short timescales ($\lesssim 10$ sec) are 200–300 fs; the longer timescale variation shows rough 24-hour periodicity and is believed to be due to thermal contraction and expansion of the 220 meters of fiber.




Figure 5: Photograph of a fiber timing transport system at MSU in August 2023. Shown are a network-attached root clock (top left, a Meinberg Syncbox), a test board with an AD9546 (red, center), an Opelink 2:24 Fiber-Optic Splitter (top right), and a Waystream SFP (top left). This setup produced the data in Fig. 4.

the number of clocks and transceivers, minimizes jitter and power consumption while providing redundancy in the lead clock. A photograph of such a setup is shown in Figure 5.

5 Root Clock Interface

When not being driven by the signals described here, the PLLs can synchronize to a bare frequency reference and time their phase markers against an external synchronization signal. This frequency reference can be single-ended or differential and at any frequency from 1 Hz to 500 MHz, although higher, in particular above 200 kHz, provides the lowest system jitter. The external synchronization mark can also be single-ended or differential and must occur with some known delta to TAI. It must be slower than the software network-time precision on the modules, which in practice suggests a signal in the range 1–100 Hz. A typical network/GPS clock 1PPS signal meets these requirements. This synchronization marker must be precisely synchronized (< 0.1 ns) with any other such root-clock synchronization markers generated in the array and must be coarsely synchronized ($< 1\mu s$) with TAI. The frequency reference must be long-term stable with respect to the synchronization marker, such that counting cycles of the frequency reference at its nominal frequency results in deltas relative to the synchronization marker no larger than the required accuracy of the synchronization marker.

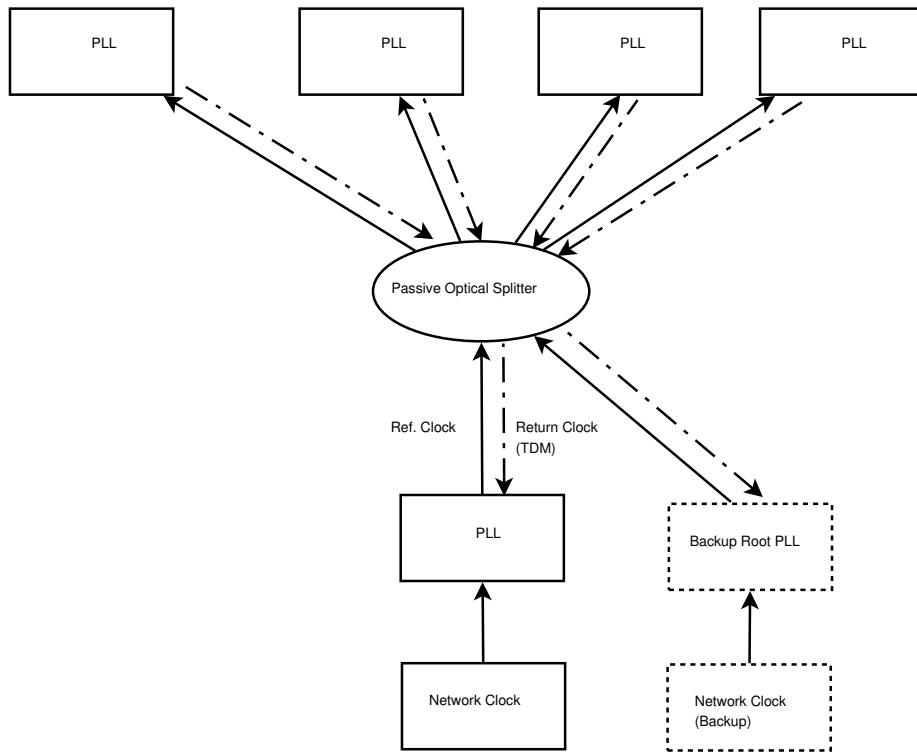


Figure 6: Topology for a system using fiber-optic transport. Clock signals are sourced by an potentially out-of-module root PLL (bottom) and then split through a passive optical splitter to each ADC PLL. Return signals for phase-delay measurement from each ADC PLL to their reference are time multiplexed through the splitter. Use of a 2:N splitter allows connection of a backup root PLL and network clock (bottom right).