Technical Note: P-ONE Timing Requirements

Nathan Whitehorn, Michigan State University

December 21, 2023

1 Introduction and Requirements

P-ONE's primary scientific objective is identification of neutrino point sources. Sensitivity to these sources scales roughly as $N_{\text{signal}}/\sqrt{N_{\text{bkg}}}$ in the area around the source. Shrinking angular resolution by a factor of x does not affect the signal from the source, but does reduce the area of interest by a factor of x^2, thus reducing N_{bkg} by x^2. Because point-source sensitivity is proportional to $\sqrt{N_{\text{bkg}}}$, and $N_{\text{bkg}} \propto x^2$, we expect P-ONE's primary science objective to improve linearly with angular resolution—at least so long as the sources in question are bright enough to make the rate of interactions in the detector non-zero.

Angular resolution, in turn, is determined primarily by the timing of photons in P-ONE. All reconstruction algorithms are fundamentally concerned with matching the timing of photon arrival in P-ONE's PMTs with expectations from a particle on a particular trajectory—the best-matching trajectory is the reconstruction result. Because of this, all other things being equal (similar path through the detector, etc.), timing variance is directly related to angular resolution variance, and for a fixed detector geometry and energy spectrum, are essentially proportional. You can see this by considering a simple line-fit algorithm, which compares photon arrival times to expectations by minimizing χ^2:

$$\chi^2 = \sum_i \frac{\left(t_i - \hat{t}_i(x, y, z, t_0, \theta, \phi)\right)^2}{\sigma_{t,i}^2} \tag{1}$$

Here, the resulting uncertainty in any of the track variables involved in predicting t_i will just be linearly proportional to σ_t. More complicated reconstructions do other things but work with the same information and so have the same statistical scaling.

2 Contributions to Timing Uncertainty

The timing uncertainty is the product of a number of different factors in the detector:

1. The photomultiplier transit-time spread, σ_{tts}

2. The timing reconstruction uncertainty from waveform digitization, σ_{digi}
3. The optical-module time synchronization uncertainty, σ_{sync}
4. The propagation time spread from scattering and dispersion, σ_{prop}
5. The uncertainty in the mean PMT transit time, σ_{tt}
6. The uncertainty in the optical-module position, $\sigma_{\mathrm{geo}}/(c/n_g)$

In any individual optical module—I will assume that different optical modules are uncorrelated—these combine in quadrature as follows, taking into account their individual scalings with the number of photons N detected at that module (scalings explained in each subsection):

$$\sigma_t = \frac{\sigma_{\mathrm{tts}}}{\sqrt{N}} \oplus \frac{\sigma_{\mathrm{digi}}}{N} \oplus \sigma_{\mathrm{sync}} \oplus \frac{\sigma_{\mathrm{prop}}}{\sqrt{N}} \oplus \sigma_{\mathrm{tt}} \oplus \frac{\sigma_{\mathrm{geo}}}{(c/n_g)} \tag{2}$$

Rearranging terms to share scalings with N, we obtain:

$$\sigma_t = \frac{\sigma_{\mathrm{digi}}}{N} \oplus \frac{\sigma_{\mathrm{tts}} \oplus \sigma_{\mathrm{prop}}}{\sqrt{N}} \oplus \left(\sigma_{\mathrm{sync}} \oplus \sigma_{\mathrm{tt}} \oplus \frac{\sigma_{\mathrm{geo}}}{(c/n_g)}\right) \tag{3}$$

Within any of these three groups, the scalings can be neglected for optimization purposes: if $\sigma_{\mathrm{tts}} < \sigma_{\mathrm{prop}}$ at one N, its contribution to the total uncertainty σ_t will be lower at any N. As such, it makes sense to think of this as composed, fundamentally, of three uncertainties: one related to the digitization electronics, one related to the phototubes and the propagation medium, and one related to the calibration.

As an additional guide for optimization, we can consider the second term, $\frac{\sigma_{\mathrm{tts}} \oplus \sigma_{\mathrm{prop}}}{\sqrt{N}}$, to be fundamentally irreducible. The medium is largely what it is—though σ_{prop} is nearly linear in the detector-module spacing—and the phototubes we are using have the lowest transit-time spread available. The rest of the terms are related to our design choices for the modules and array. The goal for optimization of σ_t, then, is just to make sure that this second term is the dominant one at all relevant N:

$$\frac{\sigma_{\mathrm{digi}}}{N} \oplus \left(\sigma_{\mathrm{sync}} \oplus \sigma_{\mathrm{tt}} \oplus \frac{\sigma_{\mathrm{geo}}}{(c/n_g)}\right) \ll \frac{\sigma_{\mathrm{tts}} \oplus \sigma_{\mathrm{prop}}}{\sqrt{N}} \tag{4}$$

We can break this apart further into two requirements by considering the scalings of the two pieces on the left. The first term matters only for small N, so we can consider this for one photon:

$$\sigma_{\mathrm{digi}} \ll \sigma_{\mathrm{tts}} \oplus \sigma_{\mathrm{prop}} \tag{5}$$

For the second term, we can set a requirement based on our per-PMT dynamic range requirement[1], which is 200 photons:

[1]For module-wide quantities like σ_{geo} and σ_{sync}, this should probably be smaller by a factor of $\sqrt{N_{\mathrm{PMTs}}} = 4$.

$$\sigma_{\text{sync}} \oplus \sigma_{\text{tt}} \oplus \frac{\sigma_{\text{geo}}}{(c/n_g)} \ll \frac{\sigma_{\text{tts}} \oplus \sigma_{\text{prop}}}{\sqrt{200}} \tag{6}$$

Given that reconstruction studies by Christian Haack and Jean-Pierre Twagirayezu that neglect all terms except σ_{tts} and σ_{prop} show that reconstruction performance improves with lower σ_{tts} with the baseline P-ONE geometry down to $\sim 1\,\text{ns}$, and our $\sigma_{\text{tts}} = 1.3\,\text{ns}$, we can approximate the combination of $\sigma_{\text{tts}} \oplus \sigma_{\text{prop}}$ as just $\sigma_{\text{tts}} = 1.3\,\text{ns}$. That leads to the following numerical constraints:

$$\boxed{\sigma_{\text{digi}} \ll 1.3\,\text{ns}} \tag{7}$$

$$\boxed{\sigma_{\text{sync}} \oplus \sigma_{\text{tt}} \oplus \frac{\sigma_{\text{geo}}}{(c/n_g)} \ll 0.1\,\text{ns}} \tag{8}$$

2.1 Photomultiplier transit-time spread

The timing of the arrival of the electron cascade at the PMT anode varies depending on propagation path through the dynodes, the electric field uniformity at the photocathode, and other random effects. These behave like a random, near-Gaussian shift in arrival time added to each photon independently; for our PMTs, this adds a 1.3 ns uncertainty [1]. Because the ensemble of photons controls the reconstruction uncertainties, the random shifts add incoherently, resulting in a lower overall uncertainty on photon timing in the presence of multiple detected photons of $1.3\,\text{ns}/\sqrt{N}$.

2.2 Timing reconstruction uncertainty

Some uncertainty is derived from the reconstruction of the photon time from the digitized waveform [2]. This uncertainty is controlled by three things:

- The sampling frequency of the digitizer
- The frequency content of the pulse
- The signal-to-noise ratio of the pulse

The easiest thing to understand is the scaling of the uncertainty with the signal-to-noise ratio (SNR), which is proportional to N. The noise source is largely uncorrelated[2] thermal and other electronic Gaussian noise added to the true measured voltage in each sample. Because there is no intrinsic scale, the way there is with a Poissonian process, linearly increasing the pulse amplitude is equivalent to linearly decreasing the uncertainty from the noise. A linear decrease in bin-to-bin uncertainty then propagates into a linear decrease in the

[2]The noise spectrum from our ADC is actually slightly blue, but this is a second-order effect

uncertainty on the fit parameters (t_0, amplitude) for the pulse, exactly as it did for point sources in the introductory section. As such, we expect σ_{digi} to scale as $1/N$. (For a fuller proof, see any proof of the Shannon-Hartley information theorem.)

The dependence of the base uncertainty on the other parameters is more complex. Somewhat counterintuitively, the sampling frequency is not dominant and contributes only indirectly.

To start with, consider a short pulse being sampled at a rate f such that the pulse is much shorter than $1/f$. What sets the uncertainty on the pulse time? Because the ADC is sampling the input signal, if the pulse is entirely located between two sample times, it will simply be missed; if it *does* get seen, it is known that the pulse was high at a particular time, so the uncertainty is set by the width of the pulse Δt and *not* by the frequency f, which just sets the rate at which the pulse gets detected.

We can make it so that the pulse is always detected by spreading it out in time, so that it is wider than the sampling interval $1/f$—this is the key function of the pre-ADC shaping electronics in most experiments. Note that now the uncertainty on the pulse is related to $1/f$, but only because we chose to make it that way. By broadening it in a particular way, however, we can not only guarantee we don't miss the pulse but also push the uncertainty below not only $1/f$ but also Δt. The way this works is best understood in the frequency domain, since that is the frame in which most of the standard sampling-theory math is written and in which the shaping electronics generally work. Per the Nyquist theorem, we can reconstruct the amplitude *and phase* of every frequency mode in a signal below $f/2$ (the "Nyquist limit") perfectly given an infinite-amplitude-resolution set of samples at frequency f; this means that we can reconstruct any function with power purely below the Nyquist limit.

This only applies if there are no frequencies above $f/2$—what creates the timing ambiguity with the short square pulse is the presence of higher frequencies in the sharp rising and falling edges that interfere ("alias") with the set of modes in the first Nyquist zone, below $f/2$. Applying a low-pass filter at $f/2$ creates a slower pulse that nonetheless can be resolved more precisely in time—one way to think about this is to see that you now have more data points (samples) and can now use shape information to pin down t_0, much the way that you can determine the mean of a Gaussian in a histogram to less than a bin width so long as the Gaussian itself is significantly wider than one bin.

In fact, per the Nyquist theorem, a pulse thus slowed down can be reconstructed *perfectly* in the absence of noise, with $\sigma_{\text{digi}} = 0$ *irrespective of the sampling frequency* f; the normal procedure for ADC systems is thus to apply such a low-pass filter (an "anti-aliasing" filter) before the ADC inputs.

In the presence of noise, measurement of the pulse time will be limited because the shape of the pulse cannot be measured precisely—we saw how noise causes a degradation in resolution at the beginning of this section. This is where the sampling frequency couples indirectly to the resolution. The total pulse SNR is a function of the integrated SNR of the pulse across all frequencies in the system. Applying a low-pass filter at $f/2$ removes some of the modes of the signal, lowering its overall power and thus SNR without lowering the noise, which is largely internal to the ADC and thus unaffected. If the filter were made very low, you can imagine smearing the pulse out so widely that it disappears.

For a single pulse, what the ADC sample frequency controls, then, is how much of the pulse survives the anti-aliasing filter built to match it and in that way indirectly affects σ_{digi} via the SNR: a fast ADC sees a more localized pulse with the same noise per unit time, so has higher SNR in that bin[3]. Since the SNR is then actually set by the first two things in the list (the sample frequency and the Nyquist-limited version of the input pulse), in addition to the ADC noise, we can think of the time resolution σ_{digi} as being entirely determined by pulse SNR.

The exact dependence of σ_{digi} is best computed numerically, but a simple analytic approach closely matches the detailed study performed by Christian Haack. The intrinsic pulse shape from the PMTs is approximately Gaussian with a 10–90% rise time of 2.9 ns [1], which roughly corresponds to a Gaussian with $\sigma = 1.4\,\text{ns}$. In Fourier space, this in turn corresponds to a Gaussian of width $\frac{1}{2\pi}\sigma^{-1} = 114\,\text{MHz}$, quite close—and not by accident—to our ADC Nyquist frequency of 105 MHz. Assuming white noise and using the fact that SNR adds in quadrature across frequencies, we can integrate the square of this Gaussian below 105 MHz to get the important incidental result that we lose very little SNR by our choice of filter and ADC sample rate, retaining 90% of the SNR we would have had using an infinite-sample-rate ADC with the same noise level and PMTs. Given this limited effect on the pulse shape by the Nyquist filter, we will assume the pulses stay Gaussian and neglect the small filter-induced pulse asymmetry in the remainder of this section.

To evaluate timing performance, we need to have $\sigma_{\text{digi}} \ll 1.3\,\text{ns}$ for all pulses with height we care about. For trigger-efficiency reasons, we separately target being able to detect pulses down to a quarter of the single-photo-electron pulse-height distribution peak (the "SPE peak", hereafter "1 PE"). Assuming the pulse shape from the previous paragraph, we require $\text{SNR} \gg (1.4\,\text{ns})/(1.3\,\text{ns}) \gg 1.07$ for all pulses of interest, which implies we need $\text{SNR} > 1.1$ for a quarter-PE pulse if we choose that point to saturate the limit; the SNR for an average pulse would then be at least 4.4, well below saturation of the uncertainty budget in the average case. This gives the timing requirement:

$$\boxed{\text{SNR}\,(1\,\text{PE}) \geq 4.4} \tag{9}$$

The trigger system separately sets a requirement of $\text{SNR} \geq 4$ for any pulse of interest to be able to distinguish it from noise. Given the quarter-PE level-of-interest threshold, the trigger system then requires $\text{SNR}\,(1\,\text{PE}) \geq 16$, easily meeting the timing criterion above. Our other requirement on SNR comes from the dynamic range requirement. Given a typical lab-measured ADC noise of 0.6 counts and 90% of the ADC's 4096-count dynamic range on the side of zero containing the pulse, the maximum SNR of any pulse is 6100. The project dynamic-range requirement to be able to measure a 200 PE pulse thus requires

[3] A faster ADC also helps disentangle multiple nearby pulses.

that SNR (1 PE) < 6100/200 < 31. Together, these impose a requirement on pulse SNR (a function of the fixed ADC noise, PMT gain, and analog front-end transimpedance) of:

$$16 < \mathrm{SNR}\,(1\,\mathrm{PE}) < 31 \tag{10}$$

Our timing requirement (Eq. 9) is satisfied easily by any value in this range allowed by the trigger and dynamic-range requirements.

2.3 Time synchronization uncertainty

The optical modules report photon times to the data acquisition and trigger system in TAI (UTC without leap seconds). This—or any other common time base—requires that the modules have measurement of absolute time. Errors in this timebase create a global time-shift for the module, contributing a per-module systematic uncertainty σ_{sync} to σ_t.

Timestamps for the modules are created using three inputs:

- A common reference clock for the array
- A stream of synchronization markers at 100 Hz
- A coarse absolute time reference delivered by IP networking using PTP

The goal of the timing synchronization system is to associate every tick of the common clock, when received at the module's ADC, with a particular array-global time that must be known, per Eq. 8, to $\ll 100\,\mathrm{ps}$. As the ADC attached to the phototubes samples at a frequency f_{ADC} that is derived from the common clock, this problem can be solved simply by declaring that the time of the i-th ADC sample is $t_i = t_0 + i/f_{\mathrm{ADC}}$, where t_0 is the time at which the ADC started. The accuracy of t_i is then dependent on the accuracy of t_0, which is derived from a combination of the 100 Hz synchronization marker stream and the PTP time, and the accuracy of the array-global reference clock.

The determination of t_0 relies on the synchronization markers. When DAQ start is requested, the ADC aligns its internal sampling clock, and so the time of the first sample, to the next delivered synchronization marker (used here as a JESD204B SYSREF signal). Determination of the absolute time of that marker relies on PTP: given a sufficiently low rate of synchronization markers and a known phase of the markers with respect to PTP time, it is possible to unambiguously determine the time of a given synchronization marker. For example, if you know the markers are emitted every 10 ms on 10 ms boundaries, detection of a synchronization mark at 8:01:42.031 implies the true time is 8:01:42.030, with whatever high precision the markers have. Effectively, the role of the synchronization marker is to enhance the precision of standard network time protocols and the rate of markers is set low (here, 100 Hz) to reflect the uncertainty in network time.

For this to work, there are two key requirements: the phase of the synchronization markers ϕ_{sync} must be known relative to the phase of the absolute time

ϕ_{PTP} to much less than one period of the synchronization clock ($1/f_{\text{sync}}$) and the uncertainty of network time σ_{PTP} must be much less than one period of the synchronization clock. If either uncertainty became larger, the association with a TAI time would become ambiguous and the pulse at 8:01:42.031 could be associated with a variety of possible times (42.020, 42.040, or 42.036 if the phase is not known). For the 100 Hz markers used in P-ONE and order-of-magnitude headroom, this translates into the requirement:

$$\begin{aligned} \sigma_{\text{PTP}} &< 1\,\text{ms} \\ |\phi_{\text{sync}} - \phi_{\text{PTP}}| &< 1\,\text{ms} \end{aligned} \tag{11}$$

Other physics goals may require tighter uncertainties on ϕ_{sync}. For example, multimessenger supernova science requires microsecond-scale timing, amending the requirements above to the following, where module and root subscripts distinguish the accuracies of the PTP time at the locations of the modules and root clock, respectively:

$$\boxed{\begin{aligned} \sigma_{\text{PTP,module}} &< 1\,\text{ms} \\ \sigma_{\text{PTP,root}} &\lesssim 1\,\mu\text{s} \\ |\phi_{\text{sync}} - \phi_{\text{PTP,root}}| &\lesssim 1\,\mu\text{s} \end{aligned}} \tag{12}$$

This sets a further constraint on the accuracy of the array-global clock. Because the ADC samples have times determined by the time at which the ADC started and the frequency, errors in the array-global clock frequency will become accumulating errors in the ADC sample time. This cannot exceed the globally desired synchronization with PTP (second requirement above) without harming science output and cannot exceed the earlier, 1 ms, requirement without losing synchronization between the ADC and subsystems that take time references later on, in particular the LED flasher timing system. For a day-length period between ADC (and DAQ) restarts, this sets the requirement that the integrated error on cycle counting cannot exceed the level above after a day. For 1 μs uncertainties after a day, this sets the requirement that the maximum allowable shift in the array clock frequency[4] from its nominal value is:

$$\boxed{\left|\frac{\Delta f}{f_{\text{nominal}}}\right| < \frac{1\,\mu\text{s}}{1\,\text{day}} < 10^{-11}} \tag{13}$$

Without the supernova constraint, this degrades to the 1 ms set by the internal synchronization requirement, yielding $|\Delta f/f| < 10^{-8}$.

All the requirements on the root clock are met by standard PTP-disciplined ovenized crystal oscillators (OCXO).

In addition to the requirements on the root clock, the time distribution system needs to have low uncertainties consistent with the overall budget on

[4]These requirements on the root clock stability apply *both* to the high-speed array-global reference clock *and* to the slower, 100 Hz synchronization clock. It is quoted here as one requirement on the basis that the 100 Hz synchronization clock and the high-speed array-global clock are presumably derived from the same source.

σ_t. Assuming that all clocks in the distribution system are locked to the master clock, the long-term frequency stability criterion on the root clock above will be met automatically by the distributed clocks. The more important requirement is how well known the arrival time of the synchronization markers is at the optical modules. This contributes directly to the calibration portion of σ_t since it is a global uncertainty in all time stamps from the module, giving a constraint of:

$$\boxed{\sigma_{\text{sync}} \ll 100\,\text{ps}} \tag{14}$$

2.4 Scattering and dispersion: σ_{prop}

As Cherenkov photons travel from their emission sites to the phototubes, they are delayed by two processes: scattering and dispersion. These both increase with distance from the source.

Scattering is caused by trajectory changes from collisions of photons with non-uniformities (air bubbles, sediment) in the water column, and takes the form of a random walk in three dimensions, adding an average delay proportional to the square root of the number of steps, which $\propto \sqrt{\Delta r/\alpha_s}$, where α_s is the scattering length. In addition, because of variability on the path, it adds uncertainty proportional to the same factor. Because α_s is very long in sea water, this is a small effect for P-ONE, though it is the dominant contributor not just to σ_{prop} but σ_t for IceCube—this is why we are building P-ONE.

Dispersion is caused by the unknown wavelength of the detected photons at the PMTs and variations in the propagation speed c/n_g at different wavelengths. This adds a deterministic, linear delay with propagation distance that looks like jitter because the wavelengths of the detected photons are not known and so which delay applies to which PMT hit is likewise unknown. For P-ONE, scattering is so small that this is the dominant contribution to σ_{prop}.

This uncertainty, because it operates on photons independently, averages down with the number of photons, contributing $\sigma_{\text{prop}}/\sqrt{N}$ to the total uncertainty σ_t. Although it is variable with distance and so is different for each event, unlike the other uncertainties in this report, Monte Carlo studies (see intro) suggest that in the modules that dominate the information used by reconstructions, the uncertainty from dispersion, so long as the dispersion itself is known, is subdominant to that from the PMT transit-time spread.

2.5 PMT mean transit time uncertainty: σ_{tt}

The time at which photons are recorded by the DAQ as detected is the sum of their arrival time at the photocathode, the time spent crossing the phototube, the propagation delay of the signal from the PMT to the ADC, any internal delays in the ADC, and any constant module-specific offsets in the timing synchronization. The photon-to-photon jitter in this process is accounted for in σ_{tts}, but the mean delay is an important calibration constant with an uncertainty that also contributes to σ_t. Because these factors matter only as a sum, and because their measurement involves the same sum, we will lump all of these pieces together as the "PMT transit time", even though that is only part of the delay described.

Measurement of the transit time relies on an external light source that flashes at a known time, then observing the arrival time in the DAQ of the light from the source. All the random factors (scalings with $1/N$ and $1/\sqrt{N}$) can be eliminated by sufficient integration time. Timing of the flash at the appropriate level of precision is a complicated problem that is not treated here. With a single bright light source for the entire detector, only relative arrival times matter and the absolute flash time does not need to be known; with multiple sources, the relative flash times of the sources must be known precisely, and so almost any approach involves cross-calibration of multiple OMs to the same flash and a step-by-step walk around the detector. The total uncertainty is constrained by Eq. 8 to:

$$\boxed{\sigma_{\rm tt} \ll 100\,\mathrm{ps}} \tag{15}$$

2.6 Position uncertainty

Optical-module position uncertainty appears in reconstructions as an uncertainty on the timing prediction: if the module is closer to the source than you think, light comes early; farther, late.[5] If the position, determined by some combination of acoustic and optical triangulation and accelerometer-based dead reckoning, is not known to a level $\sigma_{\rm geo}$, it will add an uncertainty $\sigma_{\rm geo}/(c/n_g)$ to the timing. Assuming this is uncorrelated between optical modules, it will add coherently (no scaling with N) within a module, but, like $\sigma_{\rm sync}$, will add incoherently across optical modules. Assuming $n_g = 1.35$, keeping $\sigma_{\rm geo}$ from dominating the 100 ps calibration-timing uncertainty budget (Eq. 8) requires:

$$\boxed{\sigma_{\rm geo} \ll 2.2\,\mathrm{cm}} \tag{16}$$

Although this is presented as a module constraint, this level of geometry uncertainty is about a quarter the width of one of our phototubes. This makes it important to know both the module position as a whole and the orientation of the module ψ to a high enough precision to resolve the position of the photocathode to this uncertainty on the module:

$$\boxed{\sigma_\psi \ll \arcsin\left(\frac{\sigma_{\rm geo}}{r_{\rm module}}\right) \ll 5^\circ} \tag{17}$$

It additionally sets a requirement on the uncertainty of n_g, which is simply the ratio of $\sigma_{\rm geo}$ and the typical propagation length of light $r_{\rm prop}$. Using the typical attenuation length of 30 m as the propagation length, we need to know n_g to:

[5] To second order, position shifts also affect the anticipated dispersion and scattering.

$$\boxed{\sigma_n/n_g \ll \frac{\sigma_{\text{geo}}}{30\,\text{m}} \ll 0.1\%} \tag{18}$$

Insofar as n_g is determined by flashers, measurement of n_g at this precision is implied by meeting the requirement on σ_{tt} and on meeting the requirement on σ_{geo} by a non-optical method. Alternately, module positions and n_g can be measured using an optical method that breaks the degeneracies between the mean PMT transit time, n_g, and the module position such as use of multiple, tightly-synchronized ($\ll 100\,\text{ps}$) beacons at well-known positions ($\ll \sigma_{\text{geo}}$). n_g is likely to be static, or nearly so, at this level of precision.

Although mostly out of scope for this document, measurement of the geometry imposes somewhat stricter requirements. For optical or acoustic position measurement, the propagation-speed (n_g for light or c_s for sound) uncertainty must be reduced by such a method to:

$$\sigma_n/n_g \,\text{or}\, \sigma_c/c_s \ll \frac{\sigma_{\text{geo}}}{|\Delta r_{\text{beacon}}|} \tag{19}$$

Given typical 1 km spacings from beacons to modules, calibration of the module position at the level from Eq. 16 requires either the optical or acoustic propagation speed to be known to:

$$\sigma_n/n_g \,\text{or}\, \sigma_c/c_s \ll 2 \times 10^{-5} \tag{20}$$

At this level, the propagation speed is likely to be time-variable—dependent on salinity, temperature, and pressure—and to be a product of the calibration measurement, likely using a multiple-beacon technique as described above. Because precision at this level is not required for physics operations (Eq. 18), snapshot measurements taken only during calibration measurements are sufficient and only n_g or c_s needs to be constrained at this tighter level, depending on which technique is in use.

Returning to direct requirements on timing, use of acoustic calibration also requires that the acoustic pulse time be measured at the module to within $\sigma_{\text{geo}}/c_s \ll (2.2\,\text{cm})\,/\,(1500\,\text{m/s}) \ll 15\,\mu\text{s}$, which is trivially met by the other module timing requirements described in this document.

References

[1] Hamamatsu Photonics. R14374 Datasheet. URL: https://www.hamamatsu.com/jp/en/product/optical-sensors/pmt/pmt_tube-alone/head-on-type/R14374.html.

[2] IceCube Collaboration. Energy reconstruction methods in the IceCube neutrino telescope. *Journal of Instrumentation*, 9(3):P03009, March 2014. arXiv:1311.4767, doi:10.1088/1748-0221/9/03/P03009.