

Development of a Picosecond Time-of-Flight Readout System at the Fermilab MTest Beamline

Richmond Yeung

May 2026

1 Introduction and Background

1.1 Motivation for Picosecond Timing

There are three main motivations for achieving picosecond timing resolution:

- < 10 ps Time-of-Flight Positron Emission Tomography (ToF PET)
- Probing of secondary and tertiary vertices in collider experiments
- Particle identification at the > 15 GeV scale

The goal of 10 picosecond resolution for PET has been a long-standing objective, as it enables significantly improved image resolution while reducing radiation dosage by orders of magnitude [1]. This has important implications for patient safety, particularly in pediatric imaging, and could allow radioactive tracers to be distributed more widely rather than prepared on-site.

In high-energy collider experiments, precise reconstruction of secondary and tertiary vertices is essential for identifying short-lived particles such as heavy-flavor hadrons and long-lived exotic states. These particles decay at displaced locations relative to the primary interaction point. Picosecond-level timing allows tracks to be correlated in both space and time, reducing combinatorial background and improving reconstruction accuracy.

The primary application of this work is particle identification at the ~ 15 GeV scale using Time-of-Flight techniques at Fermilab's MTest beamline. Such capabilities are also relevant to experiments including LHCb, the Electron-Ion Collider (EIC), and rare decay searches such as KOTO.

1.2 Time-of-Flight (ToF)

Time-of-Flight (ToF) systems determine a particle's velocity by measuring the difference in arrival time between two spatially separated detectors.




Figure 1: Example Time-of-Flight setup using multiple detectors

For a relativistic particle with velocity $\beta = v/c$ (taking $c = 1$), the time of flight over a distance d is:

$$\Delta t = \frac{d}{\beta} \quad (1)$$

$$\Delta t = d \sqrt{1 + \frac{m^2}{p^2}} = \frac{dE}{p} \quad (2)$$

where m is the particle mass, p its momentum, and E its energy. Using the relation $E^2 = p^2 + m^2$, the particle mass can be inferred:

$$\Delta t^2 = \frac{d^2(m^2 + p^2)}{p^2} \quad (3)$$

$$\Rightarrow m = p \sqrt{\left(\frac{\Delta t}{d}\right)^2 - 1} \quad (4)$$

Thus, combining time-of-flight measurements with momentum information enables particle identification.

1.3 Pion-Kaon Separation




Figure 2: Pion-kaon separation using time-of-flight; contours indicate required timing resolution [2]

Applying the ToF formalism to particles with known masses, such as pions and kaons, allows calculation of their expected time-of-flight difference. The contours in Figure 2 show the timing resolution required to distinguish between the two species at different momenta.

At higher momenta, the time difference becomes extremely small, requiring timing resolution at the picosecond level. Achieving ~ 1 ps resolution would enable particle identification at significantly higher energies than currently possible.

1.4 Timing Resolution and Readout Electronics

The overall timing resolution of a detector system is often limited by the readout electronics. Even if individual detectors achieve few-picosecond resolution, the electronics introduce additional uncertainty that cannot be reduced by averaging over multiple detectors.

Waveform sampling electronics are commonly used to digitize detector signals. These systems function similarly to high-speed oscilloscopes, capturing the shape of the electrical pulse for precise timing extraction.

The timing resolution of a waveform sampler is given by [3]:

$$\Delta t = \frac{\Delta u}{U} \cdot \frac{1}{\sqrt{3f_s f_{3dB}}} \quad (5)$$

where Δu is the voltage uncertainty, U is the signal amplitude, f_s is the sampling rate, and f_{3dB} is the analog bandwidth.

This relationship highlights the importance of both high sampling rate and large bandwidth in achieving picosecond timing precision.

The goal of this project is to develop a low-power, high-channel-count waveform sampling ASIC (PSEC6) with sub-picosecond timing resolution. This device is intended for deployment in Time-of-Flight particle identification systems at Fermilab's MTest beamline, with potential applications in a wide range of high-energy physics experiments.

2 PSEC6 System Design

The operating principle of the ASIC is simple - sample the detector pulse very quickly, then digitize the signal slowly. The advantage of this is Analog-to-Digital converters (ADCs) have a speed and resolution tradeoff. Slower digitization allows for a better resolution of the voltage sampled.

2.1 Sample and Hold - Switched Capacitor Circuits

"Sample and Hold" is achieved using a Switched-Capacitor Array (SCA), demonstrated by the figure below.




Figure 3: An illustration of a Switched-Capacitor Array

For each SCA, a single capacitor is closed at a time, storing the charge for a small interval of the pulse, essentially taking a slice of the pulse before the voltage stored by that slice is digitized by an ADC.

PSEC6's SCA runs at 10.24 GS/s.

2.1.1 Interleaved SCA

PSEC6 maximizes its sampling rate by interleaving multiple SCAs.




Figure 4: Interleaved Switched Capacitor Array across four phases of the sampling clock [4]

By delaying the sampling by a set amount of the clock phase, the chip samples a different portion the same signal using four different SCAs at the same time, quadrupling the sampling rate to 40 GS/s.

2.2 Performance Block Diagram

PSEC6 operates on three main components: The Sampling Block (SCA and its related components), Digital Block (which acts as the controls, with a programmable trigger), and the Phase Locked Loop (the clock source).




Figure 5: Functional Block Diagram of PSEC6. Adapted from [5]

2.3 Sampling Clock

Operating a low-jitter oscillator (to generate the wave) on an integrated circuit is a challenge, as they inject significant and unpredictable phase noise into itself, while process and temperature variations - the latter of which can happen in real time, also alters the frequency of the clock - causing it to drift.




Figure 6: An Example LC based Oscillator [6] - Capacitance and Inductance are temperature dependent, and the Oscillator's frequency is dependent on the value of LC. Our oscillator also includes a digital band - Additional capacitors that can be turned on and off by an external command to manually adjust the oscillator's frequency

The solution is to tie the chip's oscillator to a stable (and low jitter) external reference of a lower frequency. This is called a Phase Locked Loop (PLL).

2.3.1 Dual-edged triggered flip flops

Distributing a 10.24 GHz clock signal around an entire chip requires a lot of power, and the simulation takes a much longer time which makes development inefficient. As a result of this, we opted to divide the 10.24 GHz clock signal by two and distribute a 5.12GHz clock instead.

In order to recover the sampling rate, we devised a scheme to use dual-edged triggered flip flops, meaning that the clock signal (a square wave, where the rising edge acts as a tick) is utilized for both its rising and falling edges - doubling the overall sampling rate to bring each SCA back to 10.24 GSamples/s.




Figure 7: Schematic of a dual edge triggered flip flop [4], as well as simple illustration of its sampling (left: rising edge only, right: sampling on both edges)

2.4 Phase Locked Loop

Recall from figure 6 that the oscillator's output frequency is dependent on the value of LC. Using a voltage controlled capacitor (varactor), we can alter the value of LC by changing the value of the capacitor - allowing for the frequency control and forming a Voltage-Controlled Oscillator (VCO).




Figure 8: A block diagram of a Charge pump based PLL

In order to compare the VCO's output frequency to the external reference, the signal must be divided down. Flip-flops are used to achieve this: it detects a rising edge of the square wave, allowing it to pass, then requiring another rising edge to reset the flip flop before the next rising edge is allowed to pass.

The Phase-Frequency detector (PFD) compares the divided clock to the reference, and outputs high when one is faster than the other, which controls a charge pump that increases and lowers the VCO's control voltage based on the PFD's response.




Figure 9: Operating principle of a phase detector [6]

Lastly, the resistor and two capacitor ensure that the loop dynamics are stable: in essence, so that the control voltage doesn't jump around, is able to settle, and minimizes the size of ripples that directly cause jitter.

3 Methods

3.1 PLL Design and Implementation

PSEC6, like many other ASICs are designed using the Cadence suite, which comes with design tools, simulation capabilities, and a parasitics (e.g. capacitance between metal layers) modeling.

Through one of our collaborators at Fermilab, we were provided with a free-running VCO that we used in PSEC5's August 2024 submission. However, through our testing and re-runs of the simulations, it became clear that the PSEC5 clock dividers are poor and needed designing.

In order to streamline the PLL development, I wrote a mathematical model of a Phase Locked Loop and all of its components using a hardware description language (HDL) called VerilogA. This allowed the modular replacement of ideal VerilogA components with transistor-level schematics one at a time to see whether the model still managed to lock.

3.2 PLL components

Creating a phase-frequency detector (PFD) did not present too large of a challenge. A typical PFD consists of two flip-flops and a reset, which was relatively straightforward to implement using standard cells provided by TSMC. We also implemented an "active high reset path", in order to turn off the PFD, along with other components of the PLL for testing in case the PLL does fail - allowing us to fall back to a free running VCO.




Figure 10: Phase Frequency detector with active high reset path [6] [8]

As for the charge pump, we went with a simple nmos-only current steering charge pump with cascoded branches and complementary switching. This works sufficiently well, but could be improved for a future design (to produce a more stable VTune) as the design does present some charge sharing and current mismatch issues.




Figure 11: Current Steering Charge Pump

3.3 Clock Divider Chain

The VCO produces a differential output: both a "positive" and "negative" signal with equal magnitudes but opposite polarity. The "real" signal being the difference between the positive and the negative outputs.

Differential pairs come with great advantages: since the signal swings by less, it allows for much faster switching times. Furthermore, since the two conductors carry equal but opposite signals, their magnetic fields cancel each other out, resulting in much lower electromagnetic radiation and less cross talk on nearby signal lines. Moreover, it allows for very significant "common-mode" noise immunity from the power supply, since the noise would appear on both traces and the difference of the two would eliminate said noise.

However, PSEC5's design left the negative node open, and utilizes a large inductor as the only source of buffering/driving for the first stage. This makes it incredibly susceptible to supply noise, and the lopsided output likely causes further problems to the signal's integrity. Furthermore, during the re-simulation, it was found that the "fast" divider that converts the frequency from 10.24 to 5.12 GHz was very susceptible to failing (by missing edges of the clock) when there were any power drops, this called for a complete overhaul of the divider architecture.




Figure 12: Block diagram depicting PSEC6's divider architecture

For PSEC6, we utilize a common-mode-logic (CML) buffer into a fast divider that remains in CML before being converted for sampling control or further division. This way, we avoid a lopsided output, ensure proper drive for the sampling clock, while retaining the advantages of the differential pair with a more robust first divider.

4 PLL Dynamics and Stability

A charge-pump phase-locked loop is a feedback system whose small-signal behavior reduces, to first order, to that of a damped driven oscillator. The time-domain equations recover a familiar second-order ODE, while the Laplace-domain treatment exposes a pole-zero structure that directly determines stability. Both views are developed below.

4.1 PLL Dynamics

Around its free-running frequency ω_0 , the voltage-controlled oscillator is modeled as a linear element whose instantaneous frequency depends on the tuning voltage:

A linear VCO:

$$\frac{d\phi_{out}}{dt} \approx K_{VCO} \cdot V_{tune} + \omega_0 \quad (6)$$

where K_{VCO} is the tuning gain (rad/s/V). The loop filter is a series RC, which combines an integrator (the capacitor) with a stabilizing zero (the resistor) whose role becomes clear in the stability analysis below.

And a loop filter with Impedance (Z):

$$Z(s) = R + \frac{1}{sC} \quad (7)$$

The phase-frequency detector compares the reference against the *divided* VCO output, where N is the feedback divider ratio that allows the VCO to run at N times the reference frequency:

The Phase detector compares the reference with the *divided* VCO output:

$$\phi_e = \phi_{ref} - \frac{\phi_{out}}{N} \quad (8)$$

Differentiating the phase-error definition and substituting the VCO law converts this into a relation between phase-error rate and control voltage:

Taking the derivative of the phase error:

$$\frac{d\phi_e}{dt} = \frac{d\phi_{ref}}{dt} - \frac{1}{N} \frac{d\phi_{out}}{dt} = \omega_{ref} - \frac{\omega_0}{N} - \frac{K_{VCO}}{N} V_{ctrl} \quad (9)$$

The charge pump converts the phase error into a current I_p that is integrated by the capacitor and resistively dropped across R , producing the control voltage seen by the VCO:

From the loop filter, the control voltage takes:

$$V_{tune} = \frac{I_p}{2\pi} \left(R \phi_e + \frac{1}{C} \int \phi_e dt \right) \quad (10)$$

4.2 Damped, Driven Oscillator

Substituting the control-voltage expression back into the phase-error rate equation and grouping constants exposes the loop's natural dynamics. It is convenient to introduce the steady-state frequency offset $\Delta\omega$ and

a single composite loop gain K :

Define $\Delta\omega = \omega_{ref} - \frac{\omega_0}{N}$ and the loop gain $K = \frac{I_{cp}K_{VCO}}{2\pi N}$:

$$\frac{d\phi_e}{dt} = \Delta\omega - KR\phi_e - \frac{K}{C} \int \phi_e dt \quad (11)$$

Differentiating once more eliminates the integral and reveals the second-order character of the loop:

$$\frac{d^2\phi_e}{dt^2} + KR\frac{d\phi_e}{dt} + \frac{K}{C}\phi_e = \frac{d(\Delta\omega)}{dt} \quad (12)$$

This is structurally identical to a mass-spring-damper system driven by an external force, so the standard machinery for second-order oscillators applies directly:

A mechanical 2nd order oscillator can be described as:

$$\frac{d^2x}{dt^2} + 2\zeta\omega_n\frac{dx}{dt} + \omega_n^2x = f(t) \quad (13)$$

Matching coefficients term-by-term identifies the loop's natural frequency and damping ratio. Notably, ω_n is set primarily by the capacitor and the loop gain, while ζ is dominated by the resistor:

A direct comparison gives the natural frequency and damping factor:

$$\omega_n = \sqrt{\frac{K}{C}} = \sqrt{\frac{I_{cp}K_{VCO}}{2\pi NC}} \quad (14)$$

$$\zeta = \frac{KR}{2\omega_n} = \frac{R}{2} \sqrt{\frac{I_{cp}K_{VCO}C}{2\pi N}} \quad (15)$$

Where $\zeta < 1$ is underdamped, $\zeta > 1$ is overdamped, and $\zeta \approx 0.707$ is critically damped. The damping ratio governs the qualitative shape of the lock transient. Figure 13 overlays simulated VCO control-voltage transients across a sweep of loop-filter component values, spanning a strongly underdamped response (cyan, with deep undershoot before recovering) through progressively better-damped traces that settle quickly to their steady-state value. When ζ is pushed very low, as in Figure 14, the loop rings for tens of microseconds before locking — in practice an unacceptable lock time and a major contributor to reference spurs in the output spectrum.

4.3 Stability

The same loop is naturally analyzed in the Laplace domain, where the integrators become $1/s$ factors and the stability question becomes one of pole-zero placement. The phase comparison carries over directly:

The Laplace domain representation:

$$\Phi_e(s) = \Phi_{ref}(s) - \frac{\Phi_{out}(s)}{N} \quad (16)$$

Cascading the forward-path blocks — the PFD generating phase error, the charge pump converting it to current, the filter shaping that current into a voltage, and the VCO integrating frequency into phase — yields the output phase as a function of the phase error:

The forward path (PFD $\rightarrow$ Charge Pump $\rightarrow$ Loop Filter $\rightarrow$ VCO) gives:

$$\Phi_{out}(s) = \frac{K_{VCO}}{s} \cdot Z(s) \cdot \frac{I_{cp}}{2\pi} \cdot \Phi_e(s) \quad (17)$$

Substituting the loop-filter impedance produces the open-loop transfer function. Two key features stand out:




Figure 13: Simulated VCO control-voltage transients for a sweep of loop-filter R and C values. Traces span heavily underdamped (cyan) through nearly critically damped responses, illustrating how ζ shapes overshoot and settling time.

the double pole at the origin (one from the capacitor, one from the VCO's frequency-to-phase integration) and the zero at $s = -1/(RC)$ contributed by the resistor:

With $Z(s) = R + \frac{1}{sC}$, we get the open loop transfer function:

$$G(s) = \frac{K_{VCO} I_{cp}}{2\pi N} \cdot \frac{sRC + 1}{s^2 C} = K \cdot \frac{sRC + 1}{s^2 C} \quad (18)$$

Closing the loop with unity feedback (after accounting for the divider) gives:

And a closed-loop transfer function of:

$$H(s) = \frac{\Phi_{out}}{\Phi_{ref}} = \frac{N \cdot G(s)}{1 + G(s)} \quad (19)$$

The role of the resistor is now apparent. Without it, $G(s)$ would reduce to $K/(s^2 C)$, contributing -180° of phase shift at every frequency — the Nyquist contour, sketched in Figure 15(a), encircles the critical point $(-1, 0)$ and the closed loop is unstable for any non-zero gain. Adding R introduces a zero that lifts the phase back toward -90° before the loop crosses unity gain, deforming the contour as in Figure 15(b) so that it no longer wraps the critical point. The placement of this zero relative to ω_n is exactly what sets the phase margin, and it is directly equivalent to the damping ratio ζ derived in the previous subsection: a low- R design is both lightly damped in the time domain and short on phase margin in the frequency domain. Hence in a charge-pump PLL, the loop-filter resistor is essential to make the loop stable in the first place.

It is also reasonable to view phase margin as amount of leeway before negative feedback becomes positive feedback - causing the loop to drive itself infinitely and completely failing.




Figure 14: Strongly underdamped ($\zeta \ll 1$) loop response. The control voltage rings for tens of microseconds before settling, demonstrating why low damping is avoided in practical PLL design despite its faster initial rise.

4.4 PLL parameters

In order to minimize jitter, we want to keep reference spurs (ripples) to a minimum. Ripples are caused by any size mismatch between the up and down pulses in the charge pump, as well as any small spike caused from the PFD firing before the reset signal turns it off (leakage signal).

Recalling the block diagram in Figure 8, the size of the second capacitor (closer to the VCO) directly decreases the size of the ripples in the control voltage. Increasing the damping factor (by increasing the loop filter capacitance and resistance) also reduces the ringing, as the leakage signal does not swing the control voltage as much during the time it is active.

This gives us three different ways to reduce the size of the ripple, which directly contributes to jitter:

1. Increase the size of C2
2. Balance the current mismatch
3. Raise the damping factor

However, when raising the size of C2, the loop phase peak on the bode plot - Figure 16 - drifts to the left very quickly, devastating the phase margin. In order to ensure that we fabricate a stable and working ASIC, we had to recover the 45 degree phase margin by drastically increasing the size of C1 and decreasing the size of the loop resistor.

Making such changes decreases the overall system's damping factor, but the size of C2 places a larger constraint on the size of the ripples. We ended up choosing the following loop parameters:

- $R = 800\Omega$
- $C_1 = 120.2 \text{ pF}$
- $C_2 = 12 \text{ pF}$

Which gives us a ripple size of 12 mV.




Figure 15: Pole-zero map and corresponding Nyquist contour for the PLL open-loop response [6]. (a) With only the two poles at the origin (no series resistor), the contour encircles $(-1, 0)$ and the closed loop is unstable. (b) Adding the resistor introduces a zero on the negative real axis, which deforms the contour away from the critical point and stabilizes the loop.




Figure 16: Bode plot for a PLL with $R = 800\Omega$, $C_1 = 120.2 \text{ pF}$, $C_2 = 12 \text{ pF}$. The phase margin is determined from the bode plot as the loop phase at the frequency where the loop gain crosses 0.

For our purposes, we can live with an underdamped PLL as long as the overall jitter is satisfactory as the sampling can occur well after the PLL locks and the chip is calibrated for the run.

To be extra safe, we designed the loop filter in a way such that its resistance is tunable.




Figure 17: Schematic of a Tunable resistor - using pass gates to activate low-resistance paths around each series resistor.

4.4.1 Layout and parasitics

Being satisfied with the result, we carried on with the full layout of the PLL block, but encountered the loop filter once again produced new problems: Given the allotted size for the PLL of 400microns by 500 microns, the loop filter we proposed was not going to fit due to each of the chosen "mim" capacitors requiring an unpopulated region around its core.

Furthermore, re-running the initial layout of the loop filter, the parasitics cause the performance of the PLL to massively fall. The size of the ripples were 41 mV, giving us triple the jitter.

With this result, we decided that a smaller loop filter of $R_1 = 750\Omega$, $C_1 = 64.2pF$ and $C_2 = 8pF$ is the only reasonable compromise, with the expectation that we would still obtain a smaller ripple as the smaller spacing would reduce the parasitics enough to generate a sensible amount of jitter. This leaves us with:

- $\zeta = 0.173$
- $\omega_n = 1.15$ MHz
- $\omega_z = 3.31$ MHz
- $\phi_m = 50.2^\circ$

The ripple size of this loop filter was $\approx 35mV$, confirming our hypothesis.

5 Results

5.1 PLL Performance

5.1.1 Schematic level simulation

At the schematic level, a peak-to-peak jitter of approximately 600 fs was measured across 81,000 clock periods. This was achieved using an EyeDiagram - which interleaves the signal for a defined period (97.6

picoseconds for us) across the entire simulation. Using this, the peak-to-peak jitter is just the maximum horizontal thickness of the curve.



Figure 18: Eye diagram used to estimate jitter

This approximately corresponds to an RMS jitter of:

$$J_{RMS} \approx \frac{J_{pp}}{2\sqrt{2\ln(N)}} = 63 \text{ fs} \quad (20)$$

Which gave us ample leeway before the parasitics kick in.

5.1.2 Extracted layout performance

After designing a schematic, one has to produce a physical layout to be fabricated by the foundry. Beyond just the logic and connectivity of the devices, a layout specifies its physical location and placement, as well as the thickness of the connectivity.




Figure 19: Physical layout of the Phase Frequency Detector [8]. Each colour on the screen represents a different metal or oxide layer

TSMC also defines a rigorous set of design rules to ensure that a design is feasible for its lithography machines. A few examples of these involve no corners smaller than $0.04\mu\text{m}$, no islands of smaller than $0.24\mu\text{m}^2$, no traces of the same layer within a trace-width apart from each other etc.

Once a layout passes the design rules check (DRC), it moves onto a layout-versus-schematic (LVS) test - which compares the connectivity between the layout and schematic. This ensures that there are no shorts nor open nets so that the devices can operate as intended. Lastly, with a device passing LVS - parasitics extraction can take place, where the software adds dummy capacitors and resistors between all the traces to simulate the layout, allowing the designer to run the simulation again to see whether the performance is still acceptable.

Due to time constraints, we were only able to repeat the Eye-diagram measurement for an extracted layout for the full PLL without the VCO prior to the chip's submission. We did include dummy devices in the VCO. Yielding a RMS jitter of 193.7 fs , with the extra jitter likely being caused by the parasitics of the giant loop filter block, as well as a slower performance for the PFD, where the longer dead zone adds a contribution.




Figure 20: Convergence of the control voltage from the component-level parasitics simulation [8]. Here, the parasitics modelling were applied to each block individually, but not as a collective. Here, we see ripples at around 34 mV.

Since the submission, we have acquired lock on the PLL at 9.088 GHz - from a 152 MHz reference multiplied by 64 using the same digital band on the VCO recall Figure 6. This shows that including the VCO parasitics, the device has slowed down sufficiently such that the 160 MHz reference is no longer within the PLL’s lock acquisition. The fully extracted PLL at 9.088 GHz recorded a jitter of 550 fs in the typical regime. This is fully expected, as the VCO is always going to be the dominant source of the jitter.

In the meantime, adjusting the digital band of the VCO has allowed for the VCO to acquire a frequency above 11 GHz, meaning the desired 40GS/s is still acquirable, and the PLL should still lock at that setting. Furthermore, it is likely that the jitter will actually decrease by a small amount with a faster VCO, meaning it is reasonable to place an upper bound to the jitter at 550 fs.

Overall, we produced a PLL with a component-level post-extraction RMS jitter of around **550 fs**.

Figure of Merit	Value
Average Total Power	$P_{\text{typ}} = 15.7 \text{ mW}$ $P_{\text{max}} = 76.6 \text{ mW}$
RMS Jitter	$J_{\text{typ}} = 550 \text{ fs}$
Duty Cycle of Sampling Clock	46.3 – 54.8%
Charge Pump Current	$I_{\text{cp}} = 200 \text{ } \mu\text{A}$
Characteristic Damping Time Scale	$\tau = 1 - 2 \text{ } \mu\text{s}$
Phase Margin	50.13°

Table 1: Performance of PSEC6’s 10.24 GHz Phase Locked Loop




Figure 21: Physical layout of the full Phase Locked Loop. The large block (in a box) on the top right is the VCO, with the ring being the inductor. The yellow blocks in the staircase formation are the two capacitors for the loop filter. The bottom right pale-blue block are the decoupling capacitors for the VCO power supply, ensuring that it does not experience sudden drops in voltage. The entire divider chain sits just below the VCO - where the VDDA symbol sits.

There was significant interest from Fermilab for a sub-picosecond jitter 10 GHz Phase Locked Loop, since an accurate digital clock has a wide range of applications in electronics such as synchronizing signals, frequency synthesis, modulation and more. Despite being one of the most commonly used circuits in electrical engineering, no PLL has been designed for High Energy Physics (HEP) applications at 10 GHz in 65nm CMOS (currently the most widely used process for HEP due to its robustness and performance/cost tradeoff).

This design is still low power, and the jitter performance should suffice for PSEC6's application, and can be further improved in known regimes.

5.2 System-Level Performance

On the 15th of April 2026, we submitted PSEC6 for fabrication and several design decisions were made in the interest of time and getting the design over the line. For example, we opted against designing a feedback based charge pump in favor of spending more time ensuring the PLL was stable. Furthermore, a full simulation of the Phase Locked Loop is still underway, and was not done ahead of the chip's submission.

However, what we do have is a full chip that completely passes the LVS test - which was not the case for PSEC5's submission, meaning that we expect to see much fewer failure modes during testing.




Figure 22: Physical layout of the full PSEC6 ASIC. The PLL sits in the middle, flanked by 4 sampling channels on each side. Pads connecting wirebonds to the external packaging surrounds the entire chip

In terms of performance, with the main figure of merit being timing resolution - we recall this equation:

$$\Delta t = \frac{\Delta u}{U} \cdot \frac{1}{\sqrt{3}f_s f_{3dB}} \quad (21)$$

where Δu is the voltage uncertainty, U is the signal amplitude, f_s is the sampling rate, and f_{3dB} is the analog bandwidth.

In designing a Phase Locked Loop to discipline the free-running VCO, we reduce the sampling-time uncertainty associated with clock jitter. For a finite pulse slope, this timing uncertainty appears as an effective voltage uncertainty in the sampled waveform and therefore contributes to the timing-resolution budget. As you may recall from PSEC6's block diagram Figure 5, the current design relies on using external fast Analog-to-Digital converters.

This was a conscious decision as designing on-chip ADCs will be an immense undertaking, and it is more important to get a working chip under the effective operating principle first. Once this has been achieved, we can advertise PSEC6 to other experiments such as ones at the EIC, LHCb and more with a known data-rate requirement allowing us to design ADCs of the appropriate specifications after-the-fact.

However, the achieved 550 fs RMS jitter should translate to around 3.1 ps of peak-to-peak jitter across our 6.4ns fast-sampling buffer. This should be more than good enough, as most of the jitter is deterministic (for example, the ripples caused by the reference clock being compared with the divided clock at 160 MHz) and can be calibrated away. Modern SCA timing calibration techniques involves feeding in a known sinusoid and fitting, demonstrating a suppression of 10-20x on the deterministic component [3].

5.3 Overview of PSEC6's performance and Possible improvements

PSEC6's sampling rate of 40 GSamples/s upgrades on PSEC4's already impressive 17 GS/s, while achieving an analog bandwidth of 4.9 GHz compared to PSEC4's 1.6 GHz. It is also designed around achieving a multi-hit capability to capture sequential signals from a detector. This is achieved by having a long "slow buffer" per channel, which provides a coarse sample of the detector signal over 204.8 nanoseconds, with "fast buffers" triggering in time for the signal's rising edge.

Table 2: Performance specifications for PSEC6

Process	65nm TSMC
Signal to Noise	1000
Sampling Rate	40GSa/s (5GSa/s)
Fast Buffer Length	6.4ns
Slow Buffer Length	204.8ns
Analog Bandwidth	4.9 GHz
Channels/Chip	8
Max. Readout Rate	40KHz

However, besides introducing an on-chip ADC, there are a few other possible design changes that could improve the performance or functionality for the next generation of PSEC.

5.3.1 Down Sampling

PSEC6 was designed with LAPPD or super-fast PMTs in mind, where timing is paramount. This is why a great emphasis is placed on sampling the rising edge quickly and accurately. However, other experiments may have less of a need for timing resolution on the rising edge, but could find PSEC6's long sampling buffer to be incredibly useful with a desire to see the entire waveform.

This calls for the need to run the SCAs at a lower frequency, which should not be difficult as the SCA operates at the frequency of the clock it is being supplied. Since the divider chain works on many blocks of divide-by-twos, it should be quite straightforward to adjust the clock logic to supply a lower-frequency clock to the SCAs to allow the fast buffer to cover the entire waveform of the user's choosing.

5.3.2 Gate-switched Charge Pump with Feedback

As previously mentioned, to reduce the charge pump current mismatch, a better charge pump could be designed.




Figure 23: Schematic of a feedback-based charge pump [7], requiring a small amplifier

A smaller mismatch would reduce the size of the ripples on the VCO's control voltage, giving us a PLL with even lower jitter.

5.3.3 Larger MOMCAP Loop Filter

From Figure 21, we can see many squares of "mimcaps" of 2 pF each. The pale blue decoupling capacitor is 5 pF large, and is comprised of MOMCaps (utilizing only layers 1-2) instead of mims.

MOMcaps are known to have better spatial efficiency than MIMCaps but with lower accuracy, it is also possible to place MOMcaps underneath mimcaps, meaning the original goal of 130 pF total is doable with some redesigning. We opted against doing this due to the time constraint of the submission - making implementation or testing the MIMCaps (as part of a loop filter instead of decoupling) unfeasible at the time.

5.3.4 Bootstrap Switches

The sampling switches on PSEC6 were designed to minimize "switching noise", which comes as a small spike to the signal as the switch turns on and off. This noise will be present in the SCA's capture, which directly affects the voltage measurement Δu . However, the bootstrap switch - which has the smallest switching noise (but larger physical area/power consumption) is something that can be tested for a future submission.




Figure 24: Example schematic for a bootstrap switch [9]

5.3.5 Dual-edged triggered flip flops at 10 GHz

Lastly, future work should test whether PSEC6's dual-edge-triggered flip-flop can operate at a frequency of 10 GHz rather than 5GHz - thus enabling 80 Gigasamples per second.

References

- [1] Domurat-Sousa, K., et al. "Low-Dose TOF-PET Based on Surface Electron Production in Dielectric Laminar MCPs." *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1057, 168676, 2023.
- [2] Angelico, E. "Development of LAPPDs for a Precision ToF System at the Fermilab Testbeam Facility." *Doctoral Dissertation, University of Chicago*, 2020.
- [3] Stricker-Shaver, D., et al. "Novel Calibration Method for Switched Capacitor Arrays Enables Time Measurements with Sub-Picosecond Resolution." *IEEE Transactions on Nuclear Science* 61, no. 6, 3607–3617, 2014.
- [4] Park, J., et al. "Design of an 8-Channel 40 GS/s 20 mW/Ch Waveform Sampling ASIC in 65 nm CMOS." *16th Pisa meeting on Advanced Detectors, Isola d'Elba*, 2024
- [5] Park, J., et al. "Design of an 8-Channel 40 GS/s 20 mW/Ch Waveform Sampling ASIC in 65 nm CMOS." *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1072, 170165, 2025.
- [6] Razavi, B. *Design of Analog CMOS Integrated Circuits*, 2nd ed. McGraw-Hill, New York, 2016.
- [7] Razavi, B. *Design of CMOS Phase-Locked Loops: From Circuit Level to Architecture Level*. Cambridge University Press, Cambridge, 2020.
- [8] Datta, A. "10.24 GHz PLL for PSEC6, a 40 GSa/s Waveform Sampling ASIC in 65 nm CMOS." *APS Global Physics Summit*, March 2026.
- [9] Razavi, B. "The Bootstrapped Switch." *IEEE Solid-State Circuits Magazine* 7, no. 3, 12–15, 2015.