# Hardware Acceleration of ZK-STARK Number Theoretic Transforms (NTT) via Pierpont-Embedded Hydra Cores

**Ender UYGUN-Creative Commons Attribution-NonCommercial 4.0 International**
*Independent Hardware Architect & Researcher-10.5281/zenodo.18652122*

February 2026

**Abstract**

**Abstract.** Zero-Knowledge Scalable Transparent Arguments of Knowledge (ZK-STARKs) are critical for blockchain scalability but suffer from high proof-generation latency due to the heavy use of Number Theoretic Transforms (NTT) over finite fields. This technical report proposes a specialized arithmetic unit tailored for Pierpont primes ($p = 2^a 3^b + 1$). By leveraging the "Hydra" determinant-locked architecture, we replace generic modular multipliers with deterministic shift-and-add logic. This approach not only accelerates the NTT stages essential for STARK provers but also minimizes dynamic power consumption. Furthermore, we identify a direct application of this low-power logic in Neuromorphic Spiking Neural Networks (SNN), establishing a unified architecture for both privacy-preserving proofs and edge-AI efficiency.

---

## 1 1. The Bottleneck: NTT in ZK-Rollups

Layer-2 blockchain scaling solutions (ZK-Rollups) rely on STARKs to verify massive transaction batches. The computational bottleneck of a STARK prover is the generation of execution traces and polynomial commitments, which are dominated (over 70%) by Number Theoretic Transforms (NTT).

Current hardware accelerators (FPGA/GPU) use generic prime fields, requiring expensive Montgomery multiplication circuits. This results in significant "Area-Power-Product" costs and limits the throughput of decentralized provers.

## 2 2. The Solution: Pierpont-Hydra Architecture

We introduce a domain-specific architecture that exploits the arithmetic properties of Pierpont primes ($p = 2^a 3^b + 1$). These primes are inherently "FFT-friendly," allowing for highly efficient Cooley-Tukey decomposition.

### 2.1 2.1. Determinant-Locked Arithmetic

Instead of treating the modulus $p$ as a random integer, the Hydra core utilizes the algebraic rigidity of the ring structure.

- **Mechanism:** The "Determinant Lock" ($ad - bc = 1$) allows modular reduction to be computed via bit-wise shifts (for the $2^a$ component) and simplified sparse multiplication (for the $3^b$ component).
- **Hardware Benefit:** In ASIC design, bit-shifts are implemented as fixed wiring (zero logic gates). This drastically reduces the silicon area compared to generic 64-bit or 256-bit multipliers.

## 3 3. Performance Implications for STARKs

Based on preliminary architectural modeling and software simulations of the Hydra core [1]:

1. **Throughput:** By eliminating pipeline stalls associated with iterative division, NTT butterfly operations can achieve single-cycle throughput.
2. **Latency:** We project a **10x-20x reduction** in modular reduction latency for Pierpont-based fields compared to generic prime fields.
3. **Prover Efficiency:** This acceleration directly translates to faster block finality on ZK-Rollup networks (e.g., StarkNet, Polygon Miden).

## 4 4. Extended Application: Neuromorphic & IoT (Dual-Use)

*This section establishes the architectural applicability of the Hydra core to energy-constrained Edge AI systems.*

The central feature of the Hydra architecture—**Multiplication-Free Arithmetic**—addresses the primary constraint of Spiking Neural Networks (SNN) and Neuromorphic processors: Energy per Operation (J/Op).

- **Event-Driven Gating:** The determinant-lock logic acts as a passive gate that only triggers dynamic power consumption when a valid algebraic "event" occurs.
- **IoT Inference:** For edge devices running encrypted inference (FHE) or lightweight authentication (ZK), the Hydra core offers a unified silicon footprint that handles both security and compute tasks with minimal battery drain.

## 5 5. Conclusion

The generic approach to modular arithmetic is reaching its scaling limits. By adopting structure-aware architectures like Pierpont-Hydra, we can unlock the next generation of performance for

ZK-STARKs. Simultaneously, this logic provides a foundational IP block for ultra-low-power neuromorphic computing.

We invite collaboration with ZK-hardware teams and neuromorphic chip designers to implement this architecture on FPGA prototypes.

## References


[1] E. Uygun, "The Hydra-FHE Architecture: Accelerating Homomorphic Large-Integer Arithmetic," *Zenodo*, 2026.

[2] E. Uygun, "A Determinant-Locked (Hydra) Cofactor Completion Method for Pierpont Numbers," *Zenodo*, 2026.

[3] Ben-Sasson, E., et al., "Scalable, Transparent, and Post-Quantum Secure Computational Integrity (STARKs)," 2018.