

FluidCluster (fluidcluster.pt): Hardware-Accelerated FHE Netting Engine via Vectorized Zero-Sum Obfuscation and Photonic Coprocessing

FluidCluster Project Team

September 15, 2026

Abstract

Privacy-preserving clearing networks heavily rely on Trusted Execution Environments (TEEs) due to the severe performance limitations of Fully Homomorphic Encryption (FHE). However, TEE architectures introduce existential vulnerabilities to hardware side-channel attacks and fundamentally fracture synchronous smart contract composability. This paper introduces the **FluidCluster Project** (`fluidcluster.pt`), a hardware-accelerated financial clearing engine operating natively over an fhEVM. By framing multi-party debt clearing and credit limits as vectorized tensor operations under the GWEN framework, we inject structured, zero-sum camouflage loops validated via succinct Zero-Knowledge Proofs (ZKPs). To eliminate the exponential complexity bottlenecks of graph routing and Rodney Garratt’s (2019) Shapley value cost allocation (2^N), we introduce a dedicated electro-optical architecture leveraging Photonic Mach-Zehnder Interferometer (MZI) meshes. Our model guarantees complete cryptographic privacy, preserves full on-chain composability, and accelerates clearing to the physical limits of light propagation.

1 Introduction

Modern decentralized infrastructure demands a convergence of operational throughput and data privacy. In multi-party obligation netting and multilateral clearing houses (e.g., Cycle.money, Circle Arc), protecting corporate metadata—such as counterparty connections, balance fluctuations, and underlying graph topologies—is paramount.

To bypass the latency overhead of running pure mathematical schemes like Torus FHE (TFHE) on

standard electronic CPUs, prevailing systems utilize a hybrid off-chain/on-chain framework powered by hardware-isolated Trusted Execution Environments (TEEs). While operationally fast, TEE designs break the core property of decentralized applications: **synchronous, trustless composability**. Furthermore, micro-architectural exploits (e.g., timing analysis, power tracking) expose enclaves to catastrophic leakage of underlying plaintexts.

This paper presents **FluidCluster** (`fluidcluster.pt`), a paradigm shift that rejects the TEE model entirely. We demonstrate that combining a natively composable Fully Homomorphic Encryption EVM (fhEVM) with a vectorized structural architecture accelerated by Photonic Mach-Zehnder Interferometer (MZI) grids achieves sub-millisecond execution speeds while relying purely on mathematical security boundaries.

2 The Vectorized Tensor Netting Architecture

Traditional clearing routing engines process obligation networks as sparse, pointer-heavy graphs. FluidCluster formats the entire global ecosystem as structured linear algebra, optimizing the network natively inside encrypted boundaries.

2.1 The GWEN Camouflage Formulation

We utilize the mathematical primitives of the GWEN model to formalize the net bilateral obligation matrix $\mathbf{N}$ as a single, dense General Matrix Multiplication (GEMM) process:

$$\mathbf{N} = (\mathbf{G} + \mathbf{E}) \odot \mathbf{W}_{\text{scale}} \quad (1)$$

Where $\mathbf{G} \in \mathbb{R}^{N \times N}$ represents the private, real Gross transaction values submitted directly to the fhEVM as hidden ciphertexts. $\mathbf{W}_{\text{scale}}$ represents the public whole-integer scaling multiplier established off-chain to eliminate numerical truncation and division-by-zero errors.

The privacy matrix $\mathbf{E} \in \mathbb{R}^{N \times N}$ introduces pre-injected, structured, zero-sum camouflage loops generated locally by client machines before ciphertext submission. The system enforces the fundamental flow-conservation constraint:

$$\sum_{j=1}^N E_{i,j} = \sum_{j=1}^N E_{j,i} \quad \forall i \quad (2)$$

Because the rows and columns of the structural noise matrix sum to exactly zero, the fictitious volume acts as a perfect steganographic mask for public clearings, yet mathematically vaporizes upon legal state settlement, leaving real payouts unaffected.

2.2 Bilateral Risk and Credit Thresholding

To transition this from academic noise injection to institutional grade infrastructure, FluidCluster maps a parallel tensor matrix $\mathbf{C} \in \mathbb{R}^{N \times N}$ representing pre-negotiated credit boundaries:

$$\mathbf{G}_{i,j} \leq \mathbf{C}_{i,j} \quad \forall i, j \quad (3)$$

The matrix $\mathbf{C}$ forms an immutable mathematical boundary, enabling the engine to concurrently perform global liquidity optimizations and risk-tolerance assessments within the encrypted domain.

3 Off-Chain Verification via ZK-SNARKs

Computing individual flow conservation validations natively on-chain within an fhEVM creates an $\mathcal{O}(N^2)$ algorithmic explosion, incurring prohibitive gas fees.

FluidCluster avoids this by separating validation from execution. Before transferring ciphertexts, local client enclaves generate a compact, 300-byte ZK-SNARK validity proof π . The proof strictly certifies compliance with the zero-sum constraints of Matrix $\mathbf{E}$ without leaking any element values. The public EVM smart contract executes a flat $\mathcal{O}(1)$ verification routine, cutting on-chain validation overhead by over 99%.

4 Physical Acceleration via Photonic MZIs

The primary runtime barrier of Zama’s underlying TFHE framework on electronic processors is polynomial modular multiplication, heavily bottlenecked by Number Theoretic Transform (NTT) memory lookups.

4.1 Optical Polynomial Convolutions

FluidCluster bypasses electronic transistors by using optical waveguides. Because light beams inherently compute continuous Fourier Transforms as they travel and pass through optical lenses, the computational core of FHE is transformed into a zero-latency, lossless physical interaction of overlapping light waves.

4.2 Solving the 2^N Shapley Cost Allocation

Distributing the financial benefits and liquidity overhead fairly across clearing members requires calculating the marginal network contribution of each node. Implementing Rodney Garratt’s landmark (2019) Shapley Value framework introduces an NP-hard combinatorial bottleneck scaling at $\mathcal{O}(2^N)$ variations.

Photonic MZI meshes resolve this exponential scaling by encoding coalition variations as discrete spatial pathways and phase-shifts across laser beams. Utilizing Wavelength Division Multiplexing (WDM), multiple colors of light pass through the same physical mesh simultaneously. By mirroring quantum-like state superposition within a classical optical computing plane, the MZI mesh solves the 2^N combinatorial intersections in parallel, transforming a heavy algorithmic search into a deterministic sub-millisecond optical event.

5 Conclusion and Future Work

The FluidCluster Project provides a blueprint for next-generation institutional clearing houses. By executing vectorized, zero-sum masked tensor equations natively on the fhEVM and leveraging specialized Photonic MZI co-processing hardware, the system removes the structural trust trade-offs inherent in modern TEE sidecars.

Future iterations of the `fluidcluster.pt` framework will target the optimization of high-frequency Analog-to-Digital Conversion (ADC) pipelines to minimize optoelectronic transcoding latency.

This is for informational purposes only. For medical advice or diagnosis, consult a professional. AI responses may include mistakes.