

# Hyper-Aether: An AI-Driven VM-Native Computer Architecture with VM Fabric and Ensemble Reliability Model

Hirofumi Inomata  
Hyper-Aether Laboratory  
Email: inomata@acm.org

**Abstract**—We present Hyper-Aether, an AI-native computer architecture that integrates virtualization, capability-based protection, and AI-driven system generation.

At its core lies VM Fabric, a dynamic execution substrate that organizes virtual machines into structured computational graphs.

VM Fabric enables ensemble computation as a first-class execution model, improving reliability through structured redundancy.

We further provide a mathematical framework for analyzing reliability, correlation, and cost trade-offs.

**Index Terms**—AI-native computing, virtualization, ensemble systems, VM Fabric

## I. INTRODUCTION

Modern computing systems rely on large operating systems and manually written programs, resulting in increasing complexity and limited adaptability.

Recent advances in AI suggest that software can be generated automatically. However, existing systems still depend on execution models designed for static programs.

Hyper-Aether introduces a paradigm shift: systems are dynamically constructed from user specifications.

This transforms computation from program-centric execution to structure-centric execution.

## II. ARCHITECTURE OVERVIEW

Hyper-Aether consists of:

- AI Shell
- VM Fabric
- Hardware Hypervisor
- Capability Protection

Unlike conventional systems, execution structures are generated dynamically.

## III. VM FABRIC

VM Fabric represents computation as a graph of interacting virtual machines.

$$\mathcal{F} = (V, E) \quad (1)$$

This abstraction enables parallel execution, modularity, and dynamic reconfiguration.




Fig. 1. Hyper-Aether architecture




Fig. 2. Ensemble computation in VM Fabric

## IV. EXECUTION MODEL

Execution is defined as:

$$Exec = \sum_i \mathcal{H}(v_i) \quad (2)$$

This model supports concurrency, fault isolation, and scalability.

## V. AI ENSEMBLE

Multiple VMs generate outputs:

$$o_1, \dots, o_n \quad (3)$$

These outputs are aggregated:

$$o^* = \text{Aggregate}(o_1, \dots, o_n) \quad (4)$$

## VI. RELIABILITY MODEL

$$P_{maj}(n, p) = \sum_{k=\lceil n/2 \rceil}^n \binom{n}{k} p^k (1-p)^{n-k} \quad (5)$$

If  $p > 0.5$ , ensemble execution improves reliability.




Fig. 3. Trade-off between reliability and cost. Correlation reduces the benefit of redundancy, leading to an optimal ensemble size.

## VII. COST-RELIABILITY TRADE-OFF

## VIII. DISCUSSION

VM Fabric elevates ensemble computation to a first-class architectural abstraction.

It enables explicit control over:

- reliability: improving correctness through redundancy
- cost: balancing performance and resource usage
- diversity: reducing correlation among outputs

These dimensions define a structured design space.

## IX. CONCLUSION

We introduced Hyper-Aether, an AI-native architecture based on dynamically generated execution structures.

VM Fabric enables modular, adaptive, and reliable systems.

By combining ensemble computation with formal optimization, it provides a foundation for next-generation AI-driven computing.