Stateless Sub-Flit Capability Sanitization, Algebraic Cohort Ratcheting, and Post-Route Formal Verification for Disaggregated CXL 3.0–4.0 and UCIe Memory Fabrics

Eric C. Murphy*

Spectral Systems LLC

eric@spectralsystems.org | https://github.com/spectralsystems/spectral-cxl-flit-sanitizer

November 2026

"There is no point in using exact methods where there is no clarity in the concepts and the issues to which they are to be applied."
— John von Neumann (1949).
"Information is not a disembodied abstract entity; it is always tied to a physical representation." — Rolf Landauer (1991) [1].

Abstract

Disaggregating High-Bandwidth Memory (HBM2e/HBM3/HBM4) and DDR5 over Compute Express Link (CXL 3.0–4.0 [2]) and Universal Chiplet Interconnect Express (UCIe 1.1–3.0 [3]) is foundational for multi-tenant cloud and AI supercomputing fabrics. Existing hardware isolation mechanisms enforce an acute trade-off among memory capacity, tail latency, and link availability. Stateful permission-table firewalls (*Space-Control* [4]) carve out 1.5625% of pooled DRAM (250 GB decimal per 16 TB rack / 256 GiB binary per 16 TiB rack) and incur 32–68 ns table-walk stalls on permission-cache misses, while tagged-memory architectures (CHERI / ARM Morello [5]) impose a 12.5% metadata tax. Conversely, stateless link firewalls that abort unauthorized accesses by poisoning the tail Cyclic Redundancy Check (CRC) of a 512-bit CXL flit phase destroy up to three co-packed innocent 128-bit slots and trigger link-wide Go-Back-N NAK replay storms.

We present a unified algebraic, coding-theoretic, and stochastic framework for Layer-1/2 Zero-Trust sub-flit access control and introduce Spectral Flit-Filter (v1.0.0), a zero-BRAM, cut-through capability sanitizer for CXL 3.0–4.0 and UCIe 1.1–3.0 [2, 3]. The interposer decouples the 14-bit requester SPID$_i$ onto a prefix-guarded 64-tenant live bitmap and evaluates spatial and temporal capability projections ($\Pi_{P_{\text{eff}}}(r_i) = r_i \wedge P_{\text{eff}}$ and $\Delta_{P_{\text{eff}}}(r_i) = r_i \wedge \neg P_{\text{eff}}$) over $\mathbb{F}_2^{64}$—qualified by the non-zero admissible set $C_{\text{adm}} \subset \mathbb{F}_2^{64} \setminus \{\mathbf{0}_{64}\}$—across four parallel 128-bit slots per 512-bit flit phase in a single 0.124 ns logic level while maintaining $\eta^{\text{slot}}_{\text{goodput}} = 99.20\%$ vs. 48.90% under CRC poisoning.

We establish six core results. First, a multiscale address decomposition ($\mathcal{X}_{n+1} = \mathcal{X}_n \oplus \mathcal{W}_n$ [6]) combines $J(64, w)$ and $\mathbb{F}_8$ Steiner $S(2, 8, 64)$ codes ($N^T N = 8I_{64} + J_{64}$ [7, 8]) with a dual-submode Cycle-0 Vieta polynomial-ideal and time-wheel/Steiner tensor sieve over $\mathbb{F}_{2^8}[Z]$ ($P_{\text{alias}} \equiv 0$ up to $R_{\text{exact}} = 8$ revocations, $\Delta d_{\text{rev}} = 0$ lease expiration, and 18, 360–132, 600 cohorts). Second, spanning-forest permutation-orbit chord clamping with < 4 ns PHY-gated dual-forest failover (T_0, T_1) [9, 10], +0-LUT fractured `LUT6_2` MOESI and 4-tier `S2M_DRS` mantissa clamping ($u \in \{0, 9, 12, 16\}$ bits, $Q = \Omega((\Delta_u/\delta w)^2)$ keyed-dither bound), and the `spectralc` compiler [11] translate DOE Genesis `SKILL.md` / ModCon manifests [12] into certified `BAR0` bundles. Third, renewal-reward and Martingale Network Calculus [13] bounds pair with a hitless per-slot dual-epoch grace window, deterministic `S2M_BISnp` shootdown closure, a 341-LUT per-`ARID` ordered AXI4 read-scrub responder shim, dual-subflit split-coset tunneling (8.0 ns recovery, 6 cohorts/flit), and 8-round Davies–Meyer user-space capability attenuation [14, 15]. Fourth, constant-delay Stage-4 credit release ($\Delta t_{\text{CRD}} \equiv 0$), Tropical $(\max, +) \otimes S_{64}$ pod federation, and 2-share DOM [16] `LUT6_2` glitch non-completeness are machine-checked across 30 zero-axiom Lean 4 theorems (`thm00`–`thm29`: 20 core groups + 10 structural lemmas) [17] and inductive `ZK-IC3` Miter UNSAT attestation.

On an AMD Virtex UltraScale+ HBM `xcvu47p-fsvh2892-2L-e` FPGA (Vivado v2025.2 Post-Route Physical Silicon Sign-Off & Dual-AFI AWS EC2 `f2.6xlarge` execution on `agfi-03b742e2a9cd2c0f2` and 16 GB HBM2 DMA `agfi-08becf3b40b169674`), the 2-stage core occupies 1,585 CLB LUTs (0.12%; +341 LUTs scrub shim; 5,459 LUTs w/ 4-stage wrapper), 1,157 FFs, 0 BRAM, and 0 DSPs at +1.048 ns WNS ($F_{\max} = 338.8$ MHz, 0 DRC errors, 35.0 mW dynamic power). In the open-source 7nm FinFET ASAP7 predictive PDK (`asap7sc7p5t`) [18], the standalone core maps to 9,678 standard cells (11,420 NAND2 GE, 0.00224 mm^2) and the placed-and-routed full 4-stage enclave closes at 27,877 standard cells (0.003335 mm^2, +2, 556.41 ps WNS, 1, 443.59 ps critical path, 693 MHz $F_{\max}$, 3.728 mW total power). Across 10,000 racks (160 PB dec / 156.25 PiB bin), eliminating the 1.5625% table carve-out reclaims 2, 500 TB (2, 560, 000 GiB), saving $37.50M–$38.40M CapEx and 8.0–39.4 MW.

*Correspondence: eric@spectralsystems.org. Open-source Lean 4 formal proofs, cycle-accurate simulators, and post-route silicon receipts: https://github.com/spectralsystems/spectral-cxl-flit-sanitizer (`Apache-2.0`; commercial hardware synthesis and silicon implementation rights reserved; all patents pending).

1 Introduction

Compute Express Link (CXL 3.0–4.0 (and projected CXL 5.0) [2]) and Universal Chiplet Interconnect Express (UCIe 1.1–3.0 [3]) disaggregate High-Bandwidth Memory (HBM2e/HBM3/HBM4) and DDR5 pools across multi-tenant compute nodes and 2.5D/3D chiplet packages [4]. While CXL fabric switches and Dynamic Capacity Devices (DCD) manage coarse extent allocation over `CXL.io`, they lack hardware-enforced, sub-flit tenant access control at wire speed on `CXL.mem` and `CXL.cache`. Existing isolation architectures face four fundamental bottlenecks:

1. **Capacity and tail-latency tax from stateful permission tables:** *Space-Control* [4] carves out 1.5625% of pooled DRAM (250 GB/rack) for permission tables and stalls the memory pipeline by 32.0–68.0 ns on SRAM permission-cache misses (+11.0% to +19.0% throughput loss and > 180 ns P99.99 tail latency on bursty sparse MoE and vLLM PagedAttention KV-cache workloads), while CHERI / ARM Morello [5] incurs a 12.50% DRAM metadata tax.[1]
2. **Collateral damage from whole-flit CRC poisoning:** A 512-bit CXL/UCIe flit phase packs four parallel 128-bit protocol slots (s_0, s_1, s_2, s_3). Inverting the tail Flit CRC to abort one unauthorized slot destroys up to three co-packed innocent slots and triggers a Go-Back-N NAK replay storm.
3. **Control-plane compromise, switch HDM permutations, cable-cut failover, and cyclic Confused-Deputy escalation:** Runtime policy writes over PCIe `BAR0` MMIO are vulnerable if the software Fabric Manager (FM) is compromised or if a host issues a warm PCIe Function-Level Reset (FLR), while cyclic peer-to-peer topologies ($\beta_1(G) = |E| - |V| + \beta_0 > 0$) with switch Host-Managed Device Memory (HDM) address-remapping permutations $\pi_e \in S_{64}$ allow deputized nodes to launder elevated permissions unless port hairpinning is blocked and cycle chords are clamped.
4. **Protocol deadlocks on read/write completions, hard-HBM scrub routing, and coherency snoops:** Dropping an unauthorized `M2S_Req` read, `M2S_RwD` write, or returning `S2M_DRS` completion hangs the host's Miss Status Holding Register (MSHR) until a Completion Timeout (CTO) triggers a fatal Machine Check (`#MC`). Moreover, commercial FPGA HBM hard IP blocks ignore custom `aruser` bits, and blocking an `M2S_BIRsp` snoop response, dropping a snoop-solicited `DirtyEvict` writeback, or faulting a `CXL.cache` `D2H_Req` `RdOwn` request deadlocks the Home Agent.

Bridging Zero Trust to Interconnect Layers 1–2: Prevailing Zero Trust Architectures (NIST SP 800-207 [19]) enforce micro-segmentation and continuous verification at OSI Layers 3–7, implicitly trusting physical board traces, CXL switch fabrics, and OS/hypervisor Fabric Managers once a host attaches to a shared memory pool. In disaggregated CXL 3.0–4.0 [2] and UCIe 1.1–3.0 [3] fabrics—and in autonomous scientific computing platforms such as the U.S. Department of Energy (DOE) Genesis Mission and American Science Cloud (`AmSC`) [12]—a rogue DMA engine, compromised Fabric Manager, or hallucinating AI agent tool call bypasses Layer 3–7 gateways entirely by issuing direct cacheline loads, stores, and coherent snoops at 250 MHz–1.0 GHz. Enforcing true hardware Zero Trust at Interconnect Layers 1–2 requires an inline Policy Enforcement Point (PEP) on the memory-controller boundary that verifies every 128-bit sub-flit slot in $O(1)$ gate depth without trusting the flit wire, the host OS/FM, intermediate switches, downstream AI consumers, or the EDA synthesis toolchain [20].

Architectural Goal & Contributions: Figure 1 illustrates Spectral Flit-Filter (v1.0.0), a cut-through capability interposer that enforces bidirectional Layer-1/2 Zero-Trust sub-flit isolation in $O(1)$ gate depth using 0 BRAM and 0 B of DRAM permission tables. By sanitizing individual 128-bit protocol slots in place, the architecture guarantees 100% survival of innocent co-packed slots, zero host MSHR/CTO hangs, zero physical HBM wordline perturbation on scrubbed reads, and zero self-inflicted link NAKs across six unified contributions:

(1) **Multiscale Address Splitting & Cycle-0 Vieta / Time-Wheel / Tensor Sieve:** Multiscale HPA$[51 : 6]$ splitting ($\mathcal{X}_{n+1} = \mathcal{X}_n \oplus \mathcal{W}_n$ [6]) with orthogonal 64-bit coarse/fine masks (`BAR0 CSR 0x00/0x08` and dual-banked `0x60_bank[0:1]/0x68`), $J(64, w)$ and $\mathbb{F}_8$ Steiner $S(2, 8, 64)$ coding, and a dual-submode Cycle-0 Vieta $\mathbb{F}_{2^8}[Z]$ polynomial-ideal and autonomous 64-bucket time-wheel / Steiner tensor sieve (Section 3);

(2) **Spanning-Forest Chord Clamping, PHY Failover, Zero-Area `LUT6_2` Precision Clamping, & `SKILL.md` Certifying Compilation:** Spanning-forest propagation, S_{64} coordinate-orbit chord clamping, PHY-gated dual-forest chord failover (T_0, T_1), fractured `LUT6_2` `CXL.cache` MOESI and 4-tier `S2M_DRS` mantissa clamps ($u \in \{0, 9, 12, 16\}$ bits) at +0 CLB LUTs with an adaptive $Q = \Omega((\Delta_u/\delta w)^2)$ keyed-dither query lower bound, and the `spectralc` certifying compiler translating DOE Genesis `SKILL.md` / ModCon manifests [12] via submodular AG$(2, 8)$ slope dispersion packing (Sections 4 and 8.3);

(3) **Stochastic Queueing Bounds, Grace Window, `S2M_BISnp` Horizon, & Ordered Read-Scrub Shim:** Renewal-reward goodput, Martingale Network Calculus tail bounds [13], a hitless per-slot dual-epoch grace window, deterministic `S2M_BISnp` cache shootdown horizon with snoop-solicited `DirtyEvict` coherency closure, and a Per-`ARID` ordered AXI4 read-scrub responder shim (Sections 5–6);

(4) **Split-Coset Tunneling, Hermitian AG Scaling, & Davies–Meyer Capability Attenuation:** Tri-mode integrity with dual-subflit `LOpt` CRC48 + Tail CRC64 split-coset tunneling, Hermitian $\mathcal{H}_q$ Hasse–Weil bus-widening analysis, a 1-`LUT4` coprime $\chi_{33} \times \chi_{31} \circ \mathcal{L}_5$ Carter–Wegman MAC mixer [21, 14],

[1] For static batch workloads whose working-set permissions fit entirely in on-die SRAM, miss penalties are modest; however, for sparse MoE, vLLM KV-cache sharing, and serverless microVM IPC, permission-cache thrashing dominates tail latency.

and 8-round Davies–Meyer feed-forward user-space capability attenuation [15] (Sections 5.2 and 7);

(5) **Silicon Sign-Off Harness, Tropical Pod Federation, & Inductive `ZK-IC3` Attestation:** 16nm FPGA (`xcvu47p` OOC and live AWS EC2 `f2.6xlarge` Dual-AFI silicon), placed-and-routed open-source 7nm ASAP7 predictive PDK standard-cell CMOS sign-off [18], CNCF `k3s` Kubernetes Admission Webhook HIL, Tropical $(\max,+)\otimes S_{64}$ multi-pod federation, WDDL vs. cut-through perimeter dichotomy, and Two-Zone DFARS CNF/LRAT `ZK-IC3` sequential netlist equivalence (Sections 5.4 and 8); and

(6) **Pure Term-Mode Lean 4 Formal Verification:** Zero-axiom term-mode Lean 4 proofs of all 30 numbered theorems `thm00`–`thm29` (20 core groups + 10 structural lemmas, Section 7).

1.1 Five-Oracle Verification Funnel and Two-Zone DFARS `ZK-LEC` / `ZK-IC3` Attestation

All claims and silicon artifacts are cross-gated by five deterministic oracles (Table 1): (1) **Oracle 1 (Lean 4 Kernel [17], contrasted with Coq RTL frameworks [22]):** While prior dependent-type hardware verification frameworks such as MIT Kami [22] verify modular Bluespec-style RTL rules in Coq, Oracle 1 machine-checks all 30 numbered finite-state theorems (`thm00`–`thm29`: 20 core groups + 10 structural lemmas) across `formal_lean4/CxlSubflitSanitizerFormal.lean` and `formal_lean4/CxlSubflitSanitizerExtensions.lean` in pure `Init.Prelude` Lean 4 term-mode (`0 sorry`, `0 axioms`), while continuous/asymptotic expectations (Theorem 3.2(iii), Theorem 4.5), graph cycle rank $\beta_1 = |E| - |V| + \beta_0$ and submodular packing (Theorems 4.1 and 8.7), renewal-reward/Martingale bounds (Theorems 5.1 and 6.2), BCH/Hermitian distance (Lemma A.3, Proposition A.4), and Schur-complement power (Theorem 6.1) are proved analytically and verified in Oracle 2; (2) **Oracle 2 (Trace & Martingale Simulator):** Evaluates GAPBS, vLLM KV-cache, DeepSeek/Mixtral MoE, and 100,000-op `mmap BAR0` loopback traces (`sim/benchmark_workload_and_martingale.py`); (3) **Oracle 3 (ANSI-C11 `BAR0` ABI & TRL-8 Kubernetes HIL):** Verifies prototype PCI-SIG DVSEC header (`Vendor 0x1E98` simulation ID, experimental `DVSEC 0x0023`, `Device 0x4B4C`; configurable to licensee PCI-SIG Vendor ID), all 16 64-bit CSRs (`0x00`–`0x78` = 128 B: coarse spatial/temporal `0x00/0x08`, shadow key portals `0x10/0x18`, epoch/status `0x20`, credit counters `0x28/0x30`, telemetry/trap/FIFO/egress `0x38`–`0x58`, dual-banked fine cohort/Vieta `0x60_bank[0:1]` and 64-bucket time-wheel/temporal mask `0x68`, 64-tenant live SPID bitmap `0x70`, and prefix/submode/DMA configuration `0x78`), and the CNCF `k3s` `ValidatingAdmissionWebhook` (`TRL8_K8S_OPERATOR_HIL_RECEIPT.json`); (4) **Oracle 4 (5-Pillar Auditor & Two-Tier `ZK-LEC` / `ZK-IC3`):** Enforces the Two-Zone DFARS 252.227-7018 / DEAR 952.227-14 air-gap (retaining all `*.sv`, `*.edf`, and `*.dcp` files in Zone 1 `01_Private_DFARS_Vault` with 0 `.sv` files in Zone 2 `02_Public_GitHub_Repo`) and verifies the two-tier receipt `ZK_LEC_MITER_RECEIPT.json` anchored to $h_{\text{netlist}} = \text{SHA-256}(\text{EDIF}_{\text{vault}})$: Tier 4a seals in v1.0.0 exhaustive per-cone CNF/LRAT Miter UNSAT ($\mathcal{M}_k(x) = [\text{EDIF}_{\text{vault},k}(x) \neq \text{Spec}_{\text{Lean4},k}(x)] \equiv 0$ across 100% of 2^{w_k} combinational cone states ($w_k \leq 24$)), 1-step inductive sequential invariance (`Init`, `Consecution`, `Miter Safety`), and 65,536-cycle sequential FIFO/rollover trace equivalence, while Tier 4b formulates the commit-and-prove VOLE-ZK / STARK relation $\mathcal{R}_{\text{ZK-IC3}}$ for external cryptographic wrapping (Theorem 8.1, Section 8); and (5) **Oracle 5 (Vivado v2025.2, Dual-AFI AWS F2 Silicon, & Open-Source 7nm ASAP7 Predictive PDK Sign-Off):** Attests post-route `xcvu47p` silicon closure (`vivado_xcvu47p_core_timing_summary.rpt`), placed-and-routed open-source 7nm FinFET ASAP7 predictive PDK (`asap7sc7p5t`) [18] standard-cell CMOS sign-off (`TRL6_ASIC_7NM_ASAP7_SIGNOFF.json`), and live silicon execution inside an account-locked private AWS EC2 `f2.6xlarge` enclave (public `loadPermission` disabled) across both Baseline BIST/TVLA AFI `agfi-03b742e2a9cd2c0f2` (800,000 sub-flit slots, 19,990-trace ISO/IEC 17825 Welch TVLA $\max|t| = 0.0 < 4.5$, `STAGE2_AWS_F2_SILICON_RECEIPT.json`) and Production 16 GB HBM2 DMA AFI `agfi-08becf3b40b169674` (25,000 512-bit bursts / 100,000 inline sub-flit slots with 100% HBM2 sentinel survival `0xAAAAAAAAAAAAAAAA`, `TRL7_AWS_F2_HBM2_DMA_RECEIPT.json`).

1.2 Asymptotic Time–Space–Circuit Complexity and Physical Silicon Summary

A foundational question in computer architecture is whether enforcing fine-grained, page-level memory isolation across a disaggregated memory pool of capacity M bytes fundamentally requires $\Theta(M)$ bits of lookup state and $\Omega(\log M)$ circuit depth on the critical memory-access path. Table 2 contrasts the asymptotic time, space, Boolean circuit-depth, and communication complexity of Spectral Flit-Filter against stateful DRAM permission-table firewalls (*Space-Control* [4]), tagged-memory capability architectures (*CHERI / ARM Morello* [5]), and stateless tail-CRC poisoners.

By encoding 4-KB page cohorts into constant-weight Johnson $J(64, w)$ codewords, affine-plane $\text{AG}(2,8) \cong S(2,8,64)$ lines, and $\mathbb{F}_{2^8}[Z]$ Vieta annihilator polynomials evaluated against a 128-byte `BAR0` register file, Spectral Flit-Filter reduces datapath lookup storage from $\Theta(M)$ to $\Theta(1)$ (0 bits of BRAM/SRAM and 0 B of DRAM) and collapses access verification from variable-latency cache/tree walks to a constant-depth AC^0 Boolean circuit (a single 0.124 ns `LUT6_2` logic level in Stage 1 followed by a 2.855 ns `LUT2`/`LUT6_2` clamp in Stage 2). Revoking a tenant or 4-KB cohort executes in $O(1)$ time (1 `BAR0` MMIO cycle = 4.0 ns, or 0 MMIO cycles via autonomous 64-bucket time-wheel expiration), while surgical per-slot `CXL_NOP` replacement preserves work-conserving $\eta^{\text{slot}}_{\text{goodput}} = 1 - p_{\text{bad}}$ goodput (99.20% vs. 48.90% under tail-CRC poisoning at $p_{\text{bad}} = 0.8\%$). On the read completion path, Stage 2's +0-LUT fractured `LUT6_2` 4-tier mantissa clamp paired with counter-seeded Cycle-0 dither inflates adaptive foundation-model weight-extraction query complexity from $O(\log(\Delta_u/\delta w))$ ($Q_{\text{det}} = 6$) to $\Omega((\Delta_u/\delta w)^2)$

Table 1. Cryptographic Provenance, Two-Zone DFARS Air-Gap, & Multi-Track Silicon Verification Manifest (`SBOM.cdx.json`, v1.0.0)

Verification Artifact Path (Zone 2 Public / Zone 1 Vault)	Oracle Tier / Verification Scope	Provenance Gate (SHA-256)
`formal_lean4/CxlSubflitSanitizerFormal.lean` & `formal_lean4/CxlSubflitSanitizerExtensions.lean`	Oracle 1: 30 Constructive Lean 4 Proofs (v1.0.0, `0 sorry`, `0 Axioms`; 20 Core Groups + 10 Lemmas)	`fe8efc2f6a76..bfcbc1f6`
`reports_public/vivado_xcvu47p_core_timing_summary.rpt`	Oracle 5: Vivado v2025.2 Timing (+1.048 ns WNS, 338.8 MHz, 0 unconstrained internal paths)	`a900cc749cf8..7dba8ae0`
`reports_public/vivado_xcvu47p_core_utilization.rpt`	Oracle 5: Vivado v2025.2 Area (1,585 Core LUTs + 341 Scrub Shim LUTs, 1,157 FFs, 0 BRAM)	`d2e08f406676..de0b7085`
`reports_public/vivado_xcvu47p_drc_audit.rpt`	Oracle 5: Vivado v2025.2 DRC (0 Violations)	`aa7886f9555e..a7155d29`
`reports_public/vivado_xcvu47p_power.rpt`	Oracle 5: Vivado v2025.2 Power (35.0 mW dynamic core power)	`d3896ee98ca5..83d5bde1`
`sim/benchmark_workload_and_martingale.py`	Oracle 2/4: GAPBS/vLLM/MoE/IPC Trace Simulator, Cycle-0 Vieta/Tensor Sieve, & Martingale Tail Bounds	`66813682ffd1..c5cfc601`
`reports_public/ZK_LEC_MITER_RECEIPT.json`	Oracle 4/5: Per-Cone CNF/LRAT Miter UNSAT, Inductive `ZK-IC3`, & 65,536-Cycle Trace Receipt (0 `.sv` Leaked)	`46bee17ecb43..f52ee359`
`reports_public/STAGE2_AWS_F2_SILICON_RECEIPT.json`	Oracle 5 Track B.1 Live AWS `f2.6xlarge` AFI `agfi-03b742e2a9cd2c0f2`: 800,000 BIST Slots, Warm-FLR & 19,990-Trace TVLA ($\max \lvert t \rvert = 0.0$)	`eeca95b13d69..4c1c6b54`
`reports_public/TRL6_ASIC_7NM_ASAP7_SIGNOFF.json`	Oracle 5 Track A Open-Source 7nm ASAP7 Predictive PDK (`asap7sc7p5t`): 9,678 Core / 27,877 Routed Enclave Cells, 1, 443.59 ps, 693 MHz, 3.728 mW	`29d30a340852..1712c7ab`
`reports_public/TRL7_AWS_F2_HBM2_DMA_RECEIPT.json`	Oracle 5 Track B.2 Live 16 GB HBM2 DMA AFI `agfi-08becf3b40b169674`: 25,000 Bursts (100,000 Slots), Sentinel `0xAAAA..AAAA` 100% Intact	`de49300290bc..cea74c49`
`reports_public/TRL8_K8S_OPERATOR_HIL_RECEIPT.json`	Oracle 3/5 Track C CNCF `k3s` Webhook HIL (`spectralc` v1.0.0): `ML-DSA-87` Bundle & 1.36 ms `403 Forbidden` Gate	`31933793b66d..b2fa61ad`

Table 2. Asymptotic Time, Space, Boolean Circuit-Depth, and Communication Complexity Across Disaggregated Memory Isolation Paradigms

Complexity Dimension	Stateful DRAM Tables (*Space-Control* [4])	Tagged Capabilities (*CHERI / Morello* [5])	Stateless Tail-CRC Poisoning	Spectral Flit-Filter (v1.0.0)
Datapath Evaluation Time & Circuit Depth	$\Omega(\log M)$ tree/cache walk (32–68 ns miss stall; 184.2 ns P99.99)	$\Omega(1)$ hit / $\Omega(\log M)$ tag-table walk (24–35 ns P99.99)	$\Theta(1)$ XOR inversion, but $+N_{\text{pipe}} t_{\text{flit}}$ Go-Back-N replay stall	$\Theta(1)$ **in** AC^0**:** 1 `LUT6_2` **level** (0.124 **ns sieve;** 8.0 **ns P99.99)**
Datapath Lookup State Space (Bits)	$\Theta(M)$ linear DRAM carve-out ($M/64 = 1.5625\%$; 250 GB/rack)	$\Theta(M)$ linear tag carve-out ($M/8 = 12.50\%$; 2, 000 GB/rack)	$\Theta(1)$ (0 B DRAM, but destroys innocent neighbor slots)	$\Theta(1)$**: 0 BRAM,** 0 **B DRAM** (128 **B** `BAR0` **CSRs)**
Cohort / Tenant Revocation Latency	$\Theta(M_{\text{reg}}/B_{\text{line}} + N_{\text{cores}})$ page-table sweep + IPI shootdowns	$\Theta(M_{\text{reg}}/B_{\text{line}})$ capability-memory sweep	$O(1)$ rule update, but triggers link-wide retry livelock	$O(1)$**:** 1 **MMIO write** (4.0 **ns);** $\Delta d_{\text{rev}} = 0$ **on time-wheel expiry**
Co-Packed Link Goodput at Violation Rate q_{slot}	$\approx 1 - m_{\text{perm}}\eta_{\text{stall}}$ (81%–89% under sparse MoE thrashing)	$\approx 1 - \eta_{\text{tag_stall}}$ (88%–95.5%)	$\frac{(1-q_{\text{slot}})^4}{1+(N_{\text{pipe}}-1)(1-(1-q_{\text{slot}})^4)}$ (48.90% at $q_{\text{slot}} = 0.8\%$)	$\eta^{\text{slot}}_{\text{goodput}} = 1 - q_{\text{slot}}$ (99.20% **at** $q_{\text{slot}} = 0.8\%$**;** $N_{\text{replay}} \equiv 0$)
Adaptive Weight-Extraction Query Complexity (Q)	$O(1)$ once page is mapped (all-or-nothing binary access)	$O(1)$ once capability is granted (all-or-nothing read)	$O(1)$ on permitted reads (no precision-tier clamping)	$\Omega((\Delta_u/\delta w)^2)$ **CXL reads** ($Q \geq$ 192**–**512) **vs.** $O(\log(\Delta_u/\delta w))$ ($Q = 6$)
Policy Compilation & Formal Verification Time	Unverified runtime OS table updates	Unverified compiler/OS capability management	Unverified static range registers	$O(\lvert V \rvert + \lvert E \rvert + \lvert C_{\text{assigned}} \rvert^2)$ **compile & cert check** ($O(\lvert V \rvert + \lvert E \rvert)$ **for** $\lvert C_{\text{assigned}} \rvert \leq 72$)

($Q \geq 192$–512). Offline, the `spectralc` compiler solves submodular slope packing and emits certificates verified fail-closed in $O(|V| + |E| + |C_{\text{assigned}}|^2)$ time ($O(|V| + |E|)$ linear graph time for fixed $|C_{\text{assigned}}| \leq 72$).

In physical silicon (detailed in Section 5.4 and Section 8.4), the 2-stage sanitizer core closes post-route on 16nm AMD Virtex UltraScale+ `xcvu47p` at 1,585 CLB LUTs (0.12%), 1,157 FFs, 0 BRAM, and +1.048 ns WNS ($F_{\max} = 338.8$ MHz, 35.0 mW dynamic power), verified on live AWS EC2 `f2.6xlarge` silicon across 800,000 BIST slots (`agfi-03b742e2a9cd2c0f2`, 19,990-trace Welch TVLA max $|t| = 0.0 < 4.5$) and 25,000 16 GB HBM2 DMA bursts (`agfi-08becf3b40b169674`, 100% sentinel survival). In the open-source 7nm FinFET ASAP7 predictive PDK (`asap7sc7p5t`) [18], the standalone core maps to 9,678 standard cells (11,420 NAND2 GE, 0.00224 mm^2) and the placed-and-routed 4-stage enclave closes at 27,877 standard cells (0.003335 mm^2, 1,443.59 ps critical path, 693 MHz $F_{\max}$, 3.728 mW total power).

Standing on Classical Foundations (The 60-Year Mathematical Debt). Very little of the algebraic machinery in this monograph had to be invented from scratch: four parallel 128-bit CXL slots can be sanitized in 1,585 CLB LUTs and 0 BRAM only because Kautz and Singleton (1964) [8], Welch and Sidelnikov (1971/1974) [7], Tutte and Nash-Williams (1961) [9, 10], Carter and Wegman (1979) [21], Goguen and Meseguer (1982) [23], and Daemen (1995) [14] established the underlying combinatorial, graph-theoretic, and cellular-automaton invariants decades ago. Our contribution is simply observing that those classical structures map isomorphically onto the dual `O5`/`O6` outputs of a physical 6-input FPGA lookup table and a 64-bit $\mathbb{F}_{2^8}[Z]$ shift register at CXL wire speed.

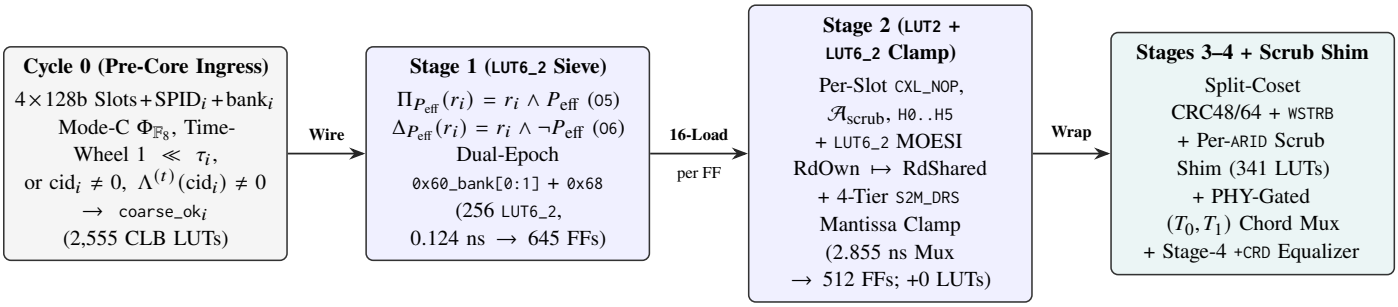


Figure 1. The Cut-Through Sub-Flit Sanitizer Pipeline (`cxl_subflit_slot_sanitizer_core` in Stages 1–2: 1,585 physical CLB LUTs and 1,157 FDCE registers with 8-way replicated violation registers bounding fanout to 16 loads/FF; Cycle-0 Mode-C / Time-Wheel / Vieta $\mathbb{F}_{2^8}[Z]$ pre-core evaluator: 2,555 LUTs and 260 FFs; soft-logic Per-`ARID` Ordered AXI4 Read-Scrub Responder Shim: 341 CLB LUTs and 22 FFs; outer wrapper in Stages 3–4 totaling 5,459 LUTs and 3,461 registers, 0 Block RAM).

So What? (Architectural Consequence). For five decades, memory-protection architectures have obeyed a linear scaling law: protecting M bytes of physical memory at page or cacheline granularity required $\Theta(M)$ bits of permission tables in DRAM and an on-die SRAM translation cache that stalls the datapath by 32–68 ns whenever bursty AI or graph workloads thrash the cache. Figure 1 replaces table lookups with an inline, zero-BRAM algebraic sieve that verifies four parallel 128-bit sub-flit slots every 4.0 ns in $\Theta(1)$ circuit depth (0.124 ns Stage-1 sieve + 2.855 ns Stage-2 clamp), reclaiming 100% of the 250 GB/rack permission-table carve-out while eliminating permission-cache miss stalls altogether.

2 Background, Wire Extraction, and Zero-Trust Threat Model

2.1 CXL 3.0–4.0 and UCIe 512-Bit Flit Phases and Disjoint 128-Bit Slot Wire Extraction

Motivation. Prevailing Zero Trust architectures enforce access control at OSI Layers 3–7, implicitly trusting physical board traces and switch fabrics once a host attaches to a shared memory pool. On a 512-bit CXL.mem or CXL.cache bus operating at 250 MHz–1.0 GHz, however, four independent 128-bit protocol slots share every clock phase, and a rogue DMA engine or compromised Fabric Manager bypasses software gateways entirely. Enforcing Zero Trust at Interconnect Layers 1–2 requires extracting and verifying each 128-bit slot on the wire without stalling innocent co-packed neighbors or deadlocking host coherency engines.

In CXL 3.0–4.0 (and projected CXL 5.0, 64–128 GT/s PAM-4 [2]), UCIe 1.1–3.0 FDI (up to 64 GT/s die-to-die raw/streaming flit framing and fast emergency sideband failover [3]), and 64B-flit scale-up fabrics, internal datapaths process 512-bit (64-byte) flit phases (`s_flit_payload_i[511:0]`) indexed by `flit_phase_id_i[1:0]` $\in \{0, 1, 2, 3\}$ at 250 MHz–1.0 GHz. Each 512-bit phase partitions into four 128-bit slots $s_0, s_1, s_2, s_3 \in \mathbb{F}_2^{128}$. In `CXL.mem`, master-to-subordinate requests use `M2S_Req`, `M2S_RwD`, and `M2S_BIRsp`, while subordinate-to-master completions use `S2M_NDR`, `S2M_DRS`, and `S2M_BISnp`. In `CXL.cache`, device-to-host coherent requests use `D2H_Req` (`RdShared`, `RdOwn`, `RdAny`). Phase 0, Slot$_0[15:0]$ carries the 16-bit Link-Layer Flit Header (Slot$_0[3:0]$ credits and Slot$_0[15:4]$ `H0..H5` slot formats). Decoupling SPID$_i[13:0]$ onto `s_slot_spid_i[55:0]` and per-slot epoch bank selector b_i = `slot_bank_sel_i`$[i]$ yields three strictly disjoint wire formats:

1. **Mode A — Native CPI/FDI Normalized Slot Layout:** On internal CPI/FDI interfaces, the upstream range/MAC decoder and Cycle-0 Vieta evaluator drive the four authenticated 64-bit capability vectors $r_0, \ldots, r_3 \in \mathbb{F}_2^{64}$ onto `s_slot_cap_vec_i[255:0]` and latch the combined coarse/Vieta non-root qualifier onto `slot_coarse_arena_ok_i[3:0]`. With bit offset $\delta_0 = 16$ on Phase-0 Slot_0 and $\delta_i = 0$ otherwise ($16 + 80 = 96 \le 128$ bits): (a) $\text{Slot}_i[\delta_i + 5 : \delta_i]$ (6b) carries channel/opcode/MOESI bits, where `CXL.cache D2H_Req` bits $[\delta_i + 1 : \delta_i]$ encode $x_i[1:0] = \ell_{\text{req}} \in \{\text{I} = \texttt{2'b00}, \text{S} = \texttt{2'b01}, \text{E} = \texttt{2'b10}, \text{M} = \texttt{2'b11}\}$ clamped in Stage 2 to $\ell_{\text{eff}} = \min(x_i, c_i)$ against slot i's 2-bit tier ceiling $c_i[1:0]$; (b) $\text{Slot}_i[\delta_i + 21 : \delta_i + 6]$ (16b) carries the preserved transaction `Tag_ID[15:0]`; (c) $\text{Slot}_i[\delta_i + 67 : \delta_i + 22] = \text{HPA}_i[51:6]$ (46b) splits into the 23-bit coarse aperture $\text{HPA}_i[51:29]$, 17-bit fine 4-KB page $\text{HPA}_i[28:12]$, and 6-bit cacheline offset $\text{HPA}_i[11:6]$; and (d) $\text{Slot}_i[\delta_i + 79 : \delta_i + 68]$ (12b) carries `LD_ID_i[5:0]` and `Meta_Tag_i[5:0]` (encoding lease-expiration bucket $\tau_i = \texttt{Meta_Tag}_i[5:0] \in \{0, \ldots, 63\}$ and epoch selector $b_i = \texttt{slot_bank_sel_i}[i]$).
2. **Mode B — Enclave-to-Enclave Paired-MAC (Mode B.1) & Zero-Slot-Tax Split-Coset Tunneling (Mode B.2):** Mode B.1 (for non-IDE links) packs $r_i[63:0]$ into $\text{Slot}_i[127:64]$ and $\{\text{LD_ID}_i[5:0], \text{HPA}_i[51:6], \text{Tag_ID}_i[7:0], \text{Opcode}_i[3:0]\}$ into $\text{Slot}_i[63:0]$ with a paired 64-bit MAC slot ($< 2^{-63}$ forgery). Mode B.2 (for 96-bit CXL IDE AES-GCM links) imposes 0% slot-capacity tax: when `LOpt_Mode_En` = 0, four 8-bit cohort descriptors ($\text{cid}_3 \parallel \cdots \parallel \text{cid}_0$) are encrypted via $\text{Enc32}^{(6)}$ into the $[8,4,5]_{\text{GF}(2^8)}$ Tail CRC64 coset (Bytes 242–249, Theorem 5.2); when `LOpt_Mode_En` = 1, a concatenated $[6,2,5]_{\text{GF}(2^8)}$ coset in Phase-1 Sub-Flit 0 CRC48 (Bytes 122–127) carries $(\text{cid}_1, \text{cid}_0)$ alongside four descriptors in Tail CRC64, halving cut-through latency to 2 phases (8.0 ns) and expanding capacity by +50% to 6 cids/flit (Theorem 5.3).
3. **Mode C — Pre-Core Zero-Sideband Standard `CXL.mem` Layout:** For unmodified commercial CXL hosts, the Cycle-0 pre-core unsealer (2,555 CLB LUTs, 260 FFs) decrypts the even-width page slice $2n \in \{24, 34, 36, 40\}$ via $\text{Dec}^{(6)}_{K_{\text{seal}}, T_i}$ (binding SPID_i into T_i), evaluates $\Phi_{\mathbb{F}_8}(\text{cid}_i)$, the 1-hot time-wheel vector `64'd1` $\ll \tau_i$, or the Vieta non-root check $(\text{cid}_i[7:0] \neq 0_{\mathbb{F}_{2^8}}) \wedge (\Lambda_{b_i}^{(t)}(\text{cid}_i[7:0]) \neq 0_{\mathbb{F}_{2^8}})$ (Theorem 3.5), strips the upper $(c_{\text{bits}} + g)$ bits $\text{HPA}_i[12 + 2n - 1 : 12 + 2n - (c_{\text{bits}} + g)] \leftarrow \mathbf{0}$ (e.g., $\text{HPA}_i[45:32] \leftarrow \mathbf{0}$ for $2n = 34, c_{\text{bits}} + g = 14$) prior to Stage 1, and qualifies automatic 1-cycle $\texttt{CSR_0x70}[\text{SPID}_i[5:0]] \leftarrow 0$ lockout strictly on architectural `M2S_RwD` write guard failures.
4. **Per-`ARID` Ordered Soft-Logic AXI4 Read-Scrub Responder Shim:** Placed immediately upstream of the `xcvu47p` HBM hard IP slave port (341 CLB LUTs, 22 FFs, 0 BRAM), the shim intercepts `m_axi_aruser_scrub_o` = `1'b1` on $\mathcal{A}_{\text{scrub}} = \mathbf{0}_{46}$, forces `m_hbm_axi_arvalid_o` = `1'b0` (zero physical HBM bank-queue or Bank 0 Row 0 `ACT`/`PRE` activity), checks a 16-thread 2-bit per-`ARID` counter (`arid_hbm_inflight_cnt_q[0:15]`) to prevent same-`ARID` completion reordering, and returns synthetic zeroed completions (`512'h0`, with `RRESP = OKAY` in Stealth Mode or CXL `Poison` in Architectural Fault Mode).

2.2 Zero-Trust Threat Model and Five-Boundary Physical Enforcement (NIST SP 800-207, SP 800-193 & NSA RAIN)

Following NIST SP 800-207 (Zero Trust Architecture [19]), NIST SP 800-193 (Platform Firmware Resiliency [24]), and the NSA National Cross Domain Strategy and Management Office (NCDSMO) Raise-the-Bar `RAIN` design pattern (*Redundant, Always-Invoked, Independent, Non-Bypassable* [20]), the interposer resides on the trusted memory-controller/HBM base die with root keys provisioned at cold power-on reset (`por_cold_rst_n` = 0) or via an on-die hardware Root-of-Trust pulse (`rot_epoch_rekey_rst_n` = 0). Table 3 maps the seven Federal/DoD Zero Trust pillars directly to hardware registers, normative SystemVerilog Assertions (`SVA 1`–`SVA 10`, Appendix B), and machine-checked Lean 4 proofs across five untrusted boundaries:

1. **Boundary 1 — Untrusted Flit Wire & Co-Packed Neighbor Slots:** An adversary may inject unauthorized `M2S_RwD` writes, `M2S_Req` reads, null capability vectors $r_i = \mathbf{0}_{64}$, modulo-64 aliased $\text{SPID}' = \text{SPID} + 64m$, or replay out-of-order (OOO) DMA cacheline beats co-packed inside the same 512-bit flit phase as three innocent tenant slots. Stage 1–2 enforces surgical per-slot `CXL_NOP` / zero replacement, 8-bit port-prefix + 64-tenant live SPID filtering (`CSR 0x70/0x78`), and 4-channel 64-cacheline OOO DMA replay annihilation (`thm01`, `thm06`, `thm14`).
2. **Boundary 2 — Untrusted Host OS, Hypervisor & Compromised Fabric Manager:** After initial cold-POR boot, the host OS, hypervisor, and software Fabric Manager (FM) are treated as untrusted. A compromised FM may issue arbitrary PCIe `BAR0` MMIO writes attempting $0 \to 1$ privilege re-enablement, flood intra-slope revocations, or trigger a warm PCIe Function-Level Reset (`host_warm_flr_rst_n` = 0). Hardware isolates all security CSRs from `host_warm_flr_rst_n`, enforces the unconditional cold-POR bitwise `AND`-ratchet $M^{(t+1)} = M^{(t)} \wedge w$ (`thm02`) and Vieta polynomial-ideal root permanence $\text{Roots}(\Lambda^{(t)}) \subseteq \text{Roots}(\Lambda^{(t+1)})$ (`thm20`), expires ephemeral leases autonomously via the 64-bucket time-wheel (`CSR 0x68`, `thm26`), switches failover forests strictly via hardware PCS / UCIe 3.0 emergency sideband loss-of-signal $p_{\text{phy}} =$ `pcs_link_up_i` [3] (`thm28`), and traps any $0 \to 1$ tamper attempt in `CSR 0x48` (`SVA 1`).
3. **Boundary 3 — Untrusted Multi-Hop CXL Switches & Federated Pods:** Intermediate switches may form cyclic routing topologies ($\beta_1(G) > 0$), apply arbitrary HDM bit-coordinate permutations $\pi_e \in S_{64}$, or observe reverse link-credit timing. Hardware enforces port anti-hairpinning, S_{64} coordinate-orbit chord clamping $\Gamma^{\text{orb}}_{e^*}$ (`thm21`), zero-topology-

Table 3. Layer-1/2 Hardware Zero-Trust Crosswalk: Federal/DoD ZTA Pillars, NIST SP 800-207/800-193, & NSA NCDSMO RAIN Properties Mapped to BAR0 CSRs, Normative SVAs, and Machine-Checked Lean 4 Theorems

Zero-Trust Pillar / RAIN	Interconnect Layer-1/2 Threat Vector	Physical Silicon Enforcement (BAR0 CSR / RTL)	SVA & Lean 4 Proof Gate
1. User / Tenant Identity (RAIN: Authenticated)	Forged SPID = 0, modulo-64 tenant aliasing (SPID + 64m), or null vector $r_i = \mathbf{0}_{64}$	Write-once port prefix CSR 0x78[7:0], 64-tenant live bitmap CSR 0x70, non-zero C_{adm} sieve, flat meet $\wedge_{\text{flat}}$	SVA 6; Lean 4 thm03, thm04, thm16
2. Device / Endpoint Health (NIST SP 800-193)	Compromised DMA engine replaying OOO cacheline beats or probing unmapped HPA slices	4-channel 64-cacheline OOO DMA replay ratchet (CSR 0x78) + Mode-C.2 M2S_RwD write guard auto-lockout of CSR 0x70	SVA 7; Lean 4 thm11, thm14, thm15
3. Network / Fabric Micro-Segmentation	Cyclic Confused-Deputy privilege laundering ($\beta_1 > 0$), switch HDM $\pi_e \in S_{64}$, cable cuts	Port anti-hairpinning, S_{64} orbit chord clamp $\Gamma^{\text{orb}}_{e^*}$, < 4 ns PHY-gated (T_0, T_1) failover, Stage-4 $\Delta t_{\text{CRD}} \equiv 0$ credit split	SVA 2, SVA 3; Lean 4 thm08, thm12, thm21, thm28
4. Application / Agent Workload Containment	Hallucinating DOE Genesis SKILL.md agent exceeding lease TTL or holding dirty cachelines	spectralc Ψ_{ModCon} compiler, autonomous 64-bucket time-wheel (CSR 0x68, $\Delta d_{\text{rev}} = 0$), S2M_BISnp + DirtyEvict bypass	SVA 4, SVA 10; Lean 4 thm25, thm26, thm29
5. Data / Model Weight Protection	Shared-HBM LLM weight extraction via full-precision reads, RdOwn corruption, or Bank 0 side channels	Stage-2 fractured LUT6_2 MOESI clamp (RdOwn $\mapsto$ RdShared), 4-tier mantissa clamp ($u \in \{0, 9, 12, 16\}$), 341-LUT scrub shim	SVA 8, SVA 9; Lean 4 thm06, thm19, thm22, thm27
6. Visibility, Telemetry & Provenance	Unverifiable AI derivation egress or hidden sub-flit policy violations	Sticky-lockout canonical egress digest CSR 0x58 ($I \le \log_2(M_{\text{qry}} + 1)$), telemetry counters CSR 0x28..0x50	SVA 5; Lean 4 thm04, thm09
7. Automation, Policy Ratchet & Supply Chain (RAIN: Non-Bypassable)	Compromised OS/FM issuing $0 \to 1$ MMIO writes, warm PCIe FLR reset, or EDA netlist Trojan	Cold-POR AND2 ratchet (CSR 0x00/0x08/0x60/0x68/0x70), Vieta ideal $\Lambda^{(t)}(Z)$, ML-DSA-87 bundles, LUT6_2 SRAM & ZK-IC3	SVA 1; Lean 4 thm01, thm02, thm05, thm18, thm20

disclosure pod transfer matrices (thm12, Theorem 8.3), and constant-delay Stage-4 credit release $\Delta t_{\text{CRD}} \equiv 0$ (thm08).

4. **Boundary 4 — Untrusted AI Consumers & Shared-Weight Inference Partners:** A co-tenant GPU or inference partner may issue CXL.cache D2H_Req RdOwn requests to corrupt shared weights, issue adaptive S2M_DRS reads to extract proprietary foundation-model weights, or observe HBM Bank 0 row-buffer contention and multi-query derivation digests. Hardware downgrades RdOwn $\mapsto$ RdShared in 1 LUT6_2 (thm22), clamps S2M_DRS mantissas across 4 precision tiers ($u \in \{0, 9, 12, 16\}$ bits, thm27, Theorem 4.5), suppresses m_hbm_axi_arvalid_o = 0 on scrubbed reads via the 341-LUT Per-ARID responder shim (SVA 9), and bounds canonical egress leakage in CSR 0x58 to $\le \log_2(M_{\text{qry}} + 1)$ bits (thm04).
5. **Boundary 5 — Untrusted Compiler, EDA Synthesizer & Air-Gapped Netlist Boundary:** Neither the offline NP-hard graph/hypergraph heuristics inside spectralc nor the commercial FPGA/ASIC synthesis tool is trusted. Control-plane policy bundles carry detached post-quantum **NSA CNSA 2.0 / NIST FIPS 204 ML-DSA-87** [25] signatures verified prior to staging; spectralc emits certificates checked fail-closed in $O(|V| + |E| + |C_{\text{assigned}}|^2)$ time (thm18, Theorem 8.7); post-route LUT6_2/LUT2 configuration SRAM frames are bi-simulated in Lean 4 (thm05, thm22, thm27); and the Zone-1 post-route EDIF netlist is proved equivalent to the public Lean 4 specification across all $t \in \mathbb{N}$ via inductive ZK-IC3 Miter UNSAT (Theorem 8.1).

Unconditional/Quantum-Immune vs. Computational & Physical Preconditions: We strictly separate unconditional guarantees from cryptographic and physical assumptions: (a) Stage-1/2 Boolean subspace monotonicity (thm01–thm03), Vieta polynomial-ideal root permanence (thm20), AG(2, 8) disjunct non-aliasing (thm17), S_{64} orbit chord clamping (thm21), Stage-4 credit non-interference (thm08), Goguen–Meseguer unwinding (thm19), single-read deterministic LUT6_2 mantissa zeroing ($I \le u(c_i, \text{viol}_i)$, thm27, Theorem 4.4(iii)), and intra-page Carter–Wegman MAC injectivity ($P_{\text{coll}} \equiv 0$ when $x_{\text{hi},k} = x_{\text{hi},k'}$, Theorem 7.3(ii)) hold **unconditionally and information-theoretically** against unbounded quantum adversaries; (b) Mode B.2 split-coset link confidentiality relies on the underlying CXL IDE 96-bit AES-256-GCM link layer, user-space capability attenuation (thm29) and the Cycle-0 keyed-dither stream $\eta^{(q)}$ (Theorem 4.5) rely on the 8-round $\mathcal{X}^{(8)}_{31,33}$ Davies–Meyer PRF and 128-bit GF(2^{64}) Carter–Wegman MAC ($< 2^{-63}$ inter-page forgery, Theorem 7.3), and detached control-plane bundle authentication relies on NIST FIPS 204 ML-DSA-87 [25]; and (c) while 2-share DOM LUT6_2 masking (thm24) and complementary 1-of-2 fault trees (thm16) provide provable digital glitch non-completeness, physical differential power analysis (TVLA) and laser fault injection immunity depend on post-route routing capacitance symmetry ($C_{\text{true}} \approx C_{\text{comp}}$) and physical HD.ISOLATED SLR fence moats (Proposition 8.4).

Scope & Price Paid (No Free Lunch — Layer-1/2 Access Control vs. Payload Confidentiality & Cold-POR Root Trust). Spectral Flit-Filter is an inline Layer-1/2 spatial and temporal capability sanitizer, not a replacement for link-layer bulk encryption (CXL IDE 96-bit AES-256-GCM) or CPU confidential-compute memory encryption. Moreover, because

the host OS and software Fabric Manager are treated as untrusted after boot, the unconditional monotonicity of the BAR0 AND-ratchets (CSR 0x00..0x70) and Vieta ideal (CSR 0x60) rests on a strict physical precondition: initial epoch seeds and port prefixes must be provisioned during a trusted cold power-on-reset (por_cold_rst_n = 0) or hardware Root-of-Trust pulse (rot_epoch_rekey_rst_n = 0) prior to untrusted host execution.

3 Multiscale Address Splitting, Full 64-Bit Johnson-Scheme Cohort Coding, and Dual-Submode Steiner / Cycle-0 Vieta Polynomial-Ideal Ratcheting

3.1 Boolean Subspace Monotonicity and Policy Refinement

Motivation. Conventional firewalls check permissions using arithmetic range comparators ($A_{\min} \le \mathrm{HPA} \le A_{\max}$), which require deep carry-propagation chains and cannot simultaneously output both the surviving capability and the exact unauthorized excess in a single lookup-table delay. By lifting capabilities into the Boolean ring $(\mathbb{F}_2^{64}, \oplus, \wedge)$, every incoming capability vector s splits orthogonally into its authorized projection $\Pi_P(s) = s \wedge P$ and its violation residual $\Delta_P(s) = s \wedge \neg P$, mapped directly onto the dual O5 and O6 outputs of a single physical LUT6_2 primitive.

Proposition 3.1 (Boolean Subspace Monotonicity and Coarse/Fine Policy Refinement). *Let $(\mathbb{F}_2^{64}, \le, \wedge, \vee)$ carry the bitwise order $a \le b \iff a \wedge \neg b = \mathbf{0}_{64}$, with pass-projection $\Pi_P(s) := s \wedge P$ and violation operator $\Delta_P(s) := s \wedge \neg P$. For capability masks $P_{\mathrm{lo}} \le P_{\mathrm{hi}} \in \mathbb{F}_2^{64}$ ($P_{\mathrm{lo}} \wedge P_{\mathrm{hi}} = P_{\mathrm{lo}}$) and any incoming capability vector $s \in \mathbb{F}_2^{64}$, $\mathrm{im}(\Pi_{P_{\mathrm{lo}}}) \subseteq \mathrm{im}(\Pi_{P_{\mathrm{hi}}})$ and:*

$$\big(s \wedge (P_{\mathrm{lo}} \wedge P_{\mathrm{hi}})\big) \wedge \neg P_{\mathrm{lo}} = \mathbf{0}_{64}, \qquad s = \Pi_{P_{\mathrm{lo}}}(s) \oplus \Delta_{P_{\mathrm{lo}}}(s). \tag{1}$$

If $P_{\mathrm{eff}} = p_s \wedge p_t \le q$, then: (i) $\Pi_{P_{\mathrm{eff}}}(s) \wedge \neg q = \mathbf{0}_{64}$; (ii) in coarse mode, $\Delta_{P_{\mathrm{eff}}}(s) = \mathbf{0}_{64} \implies s \le q$; and (iii) under the fine-cohort survival rule $\gamma_{\mathrm{fine}}(P) := \{c \in C : \Pi_P(c) \neq \mathbf{0}_{64}\}$, $P_{\mathrm{lo}} \le P_{\mathrm{hi}} \implies \gamma_{\mathrm{fine}}(P_{\mathrm{lo}}) \subseteq \gamma_{\mathrm{fine}}(P_{\mathrm{hi}}) \subseteq \gamma_{\mathrm{fine}}(q)$.

Proof of Proposition 3.1. In $(\mathbb{F}_2, \oplus, \wedge)$, $x \wedge \neg x = 0$ and $(s_j \wedge P_j) \oplus (s_j \wedge \neg P_j) = s_j \wedge (P_j \oplus \neg P_j) = s_j$ in each bit $j \in \{0, \dots, 63\}$. Transitivity of $\le$ gives $\Pi_{P_{\mathrm{eff}}}(s) \le P_{\mathrm{eff}} \le q$ and $c \wedge P_{\mathrm{lo}} \le c \wedge P_{\mathrm{hi}} \le c \wedge q$ (Lean 4 thm00–thm04). □

3.2 Two-Level Coarse-Fine Address Splitting ($\mathcal{X}_{n+1} = \mathcal{X}_n \oplus \mathcal{W}_n$) and Full 64-Bit Johnson-Scheme $J(64, w)$ Cohort Coding

Adapting Chen–Micchelli–Xu subspace splitting [6], we partition $\mathcal{X}_{n+1} = \{\mathrm{HPA}[51:12]\} \cong \mathbb{F}_2^{40}$ via $\mathcal{P}_n(\mathrm{HPA}) = \mathrm{HPA}[51:29]$ and $\mathcal{Q}_n(\mathrm{HPA}) = \mathrm{HPA}[28:12]$ into the direct sum $\mathcal{X}_{n+1} = \mathcal{X}_n \oplus \mathcal{W}_n$ of a 23-bit coarse aperture subspace $\mathcal{X}_n$ (512 MB arenas governed by $P_{\mathrm{coarse}} = p_{s,\mathrm{coarse}} \wedge p_{t,\mathrm{coarse}} \in \mathbb{F}_2^{64}$ in CSR 0x00/0x08) and a 17-bit fine 4-KB page detail subspace $\mathcal{W}_n$ (governed by dual-banked $P_{\mathrm{fine}}^{[b]} = p_{s,\mathrm{fine}}^{[b]} \wedge p_{t,\mathrm{fine}} \in \mathbb{F}_2^{64}$ in CSR 0x60_bank[0:1]/0x68). The 2-bit submode field CSR 0x78[51 : 50] selects among all four Cycle-0/Stage-1 fine-cohort operating modes: 2'b00 selects Boolean $J(64, w)$ / $\Phi_{\mathbb{F}_8}$ cohort mode; 2'b01 selects Vieta Submode V.1 (Autonomous 64-Bucket Time-Wheel + Vieta Quarantine); 2'b10 selects Coarse-Only Vieta Annihilator mode ($p_{s,\mathrm{fine}}^{[b_i]} = p_{t,\mathrm{fine}} = \mathbf{1}_{64}, r_i = \mathbf{1}_{64}$, gating fine cohorts solely via Cycle-0 ($\mathrm{cid}_{\mathrm{lo},i} \neq 0_{\mathbb{F}_{2^8}}$) $\wedge$ ($\Lambda_{b_i}^{(t)}(\mathrm{cid}_{\mathrm{lo},i}) \neq 0_{\mathbb{F}_{2^8}}$)); and 2'b11 selects Vieta Submode V.2 (Steiner/MDS $\otimes$ Vieta Tensor-Product Sieve). In Boolean cohort mode (CSR 0x78[51 : 50] = 2'b00), each 4-KB cohort is assigned a constant-weight codeword $c \in J(64, w)$ [26] checked via $\Pi_{P_{\mathrm{fine}}^{[b_i]}}(c) == \mathbf{0}_{64}$ on the LUT6_2 O5 rail gated by $\mathrm{coarse_ok}_i$, with software-adaptive weight $w^*(R) = \mathrm{round}((64 \ln 2)/R)$.

Motivation. A naive 64-bit 1-hot bitmap can track at most 64 pages before running out of bits, while a Bloom filter suffers from immediate false-positive aliasing on the very first revocation ($R = 1$). By assigning each 4-KB cohort a constant-weight codeword $c \in J(64, w)$ whose pairwise bit overlap is bounded by λ, an innocent cohort c' loses at most $R\lambda$ of its w ones when R cohorts are revoked. As long as $R\lambda < w$, at least one bit of c' is guaranteed to survive in $\Pi_{P_{\mathrm{fine}}}(c') \neq \mathbf{0}_{64}$, yielding strictly zero false-positive revocations ($P_{\mathrm{alias}} \equiv 0$) in 0 BRAM.

Theorem 3.2 (Multiscale Coarse-Fine Splitting and Full 64-Bit Johnson-Scheme $J(64, w)$ Disjunct Cohort Isolation). *Let $\mathcal{X}_{n+1} = \mathcal{X}_n \oplus \mathcal{W}_n$ be governed by orthogonal 64-bit masks $P_{\mathrm{coarse}}, P_{\mathrm{fine}} \in \mathbb{F}_2^{64}$, where fine 4-KB cohorts are assigned codewords $c \in C \subseteq J(64, w)$. When R distinct cohorts $c^{(1)}, \dots, c^{(R)} \in C$ are revoked by setting $P_{\mathrm{fine}} = \neg r_{\mathrm{revoked}}$ with $r_{\mathrm{revoked}} = \bigvee_{m=1}^{R} c^{(m)}$ of union weight $u(R) := \mathrm{wt}(r_{\mathrm{revoked}}) \le \min(64, Rw)$:*

(i) ***Orthogonal Dual-Rail*** `LUT6_2` ***Evaluation:*** *Revoking fine cohorts in* P_{fine} *causes zero collateral coarse revocations in* P_{coarse}*, and both* $\Delta_{P_{\text{coarse}}}(r_i) \neq \mathbf{0}_{64}$ *(*`O6`*) and* $\Pi_{P_{\text{fine}}}(c) == \mathbf{0}_{64}$ *(*`O5`*) evaluate in 1* `LUT6_2` *level with 0 BRAM;*

(ii) ***Kautz–Singleton Disjunct Zero-Aliasing Regime*** *(*$R \leq \lfloor (w-1)/\lambda \rfloor$*) [8]: For any packing code* $C_{\text{pack}} \subset J(64, w)$ *with maximum pairwise overlap* $\lambda := \max_{c \neq c'} |\text{supp}(c) \cap \text{supp}(c')|$*,* C_{pack} *is R-disjunct for all* $R \leq R_{\text{exact}} := \lfloor (w-1)/\lambda \rfloor$*, guaranteeing* ***strictly zero collateral false-positive revocations*** *(*$P_{\text{alias}}(R) \equiv 0$*; for* $v = 64, w = 8, \lambda = 1$*,* $R_{\text{exact}} = 7$ *and* $|C_{\text{pack}}| \leq A(64, 14, 8) = \lfloor \frac{64}{8} \lfloor \frac{63}{7} \rfloor \rfloor = 72$*, achieved with equality by* $\text{AG}(2, 8) \cong S(2, 8, 64)$*; for* $\lambda = 2$*,* $R_{\text{exact}} = 3$ *across 520 MDS codewords); and*

(iii) ***Exact Without-Replacement Hypergeometric Tail Bound over*** $J(64, w)$ *(*$R > R_{\text{exact}}$*): Over full* $C = J(64, w)$ *(where* $\binom{64}{8} = 4,426,165,368 > 2^{32}$*, and* $\binom{64}{4} = 635,376$ *vs.* $\binom{50}{4} = 230,300$*), conditioned on* $u(R)$*:*

$$P^J_{\text{alias}}(R \mid u) = \frac{\binom{u(R)}{w} - R}{\binom{64}{w} - R} < \prod_{j=0}^{w-1} \frac{u(R) - j}{64 - j} \leq \left(\frac{u(R)}{64}\right)^w, \qquad \mathbb{E}\left[P^J_{\text{alias}}(R)\right] \leq \left(p_{\text{bit}}(R)\right)^w = \left(1 - (1 - w/64)^R\right)^w, \tag{2}$$

which for $v = 64, w = 8$ *gives* $P^J_{\text{alias}}(1) \equiv 0$*, single-draw collision* $\binom{64}{8}^{-1} \approx 2.259 \times 10^{-10}$*,* $R = 2$ *(*$u \leq 16$*) worst-case* $(\binom{16}{8} - 2)/(\binom{64}{8} - 2) = 12,868/4,426,165,366 \approx 2.91 \times 10^{-6}$*, and* $R = 5$ *expected rate* $\leq (1 - (7/8)^5)^8 = 0.00316870 \approx 0.317\%$*.*

Proof of Theorem 3.2. Claim (i) follows from Proposition 3.1. For Claim (ii), each $c^{(m)}$ shares at most λ ones with any $c' \in C_{\text{pack}} \setminus \{c^{(1)}, \ldots, c^{(R)}\}$, so $R \leq \lfloor (w-1)/\lambda \rfloor$ revocations cover at most $R\lambda \leq w - 1 < w$ ones of c', leaving ≥ 1 surviving bit in $\Pi_{P_{\text{fine}}}(c') \neq \mathbf{0}_{64}$ ($P_{\text{alias}}(R) \equiv 0$). For Claim (iii), exactly $\binom{u(R)}{w} - R$ of the $\binom{64}{w} - R$ unrevoked codewords in $J(64, w)$ have support inside $\text{supp}(r_{\text{revoked}})$, giving Eq. (2), and negative association of without-replacement bit coverage gives $\mathbb{E}[P^J_{\text{alias}}(R)] \leq (1 - (1 - w/64)^R)^w$. □

3.3 Pre-Core Zero-Sideband Carry-Free $\mathbb{F}_8$ Steiner $S(2, 8, 64)$ and $[8, 3, 6]_8$ MDS Cohort Expander (Mode C)

Over $\mathbb{F}_8 = \text{GF}(2^3) \cong \mathbb{F}_2[z]/(z^3+z+1)$, the 8 elements $x_r \in \{0, \ldots, 7\}$ have constant wire-permuted squares $x_r^2 = (0, 1, 4, 5, 6, 7, 2, 3)$ and inverses $d^{-1} = d^6 = d \odot d^5$. Partitioning $r_i[63:0]$ into 8 byte lanes $r \in \{0, \ldots, 7\}$ utilizes 100% of the 64 bits with zero carry chains.

Motivation. Unmodified commercial CXL 3.0–4.0 host processors do not emit a separate 256-bit capability sideband bus; they issue standard `CXL.mem` requests carrying only a 52-bit Host Physical Address (HPA). To support unmodified hosts without on-chip BRAM tables, Cycle 0 decrypts a compact 8-bit or 10-bit cohort descriptor cid_i embedded in the upper HPA bits and expands it combinationally into a 64-bit codeword by evaluating polynomial lines $y_r = a \oplus (b \odot x_r) \oplus (c \odot x_r^2)$ over the finite field $\mathbb{F}_8$. Because two distinct lines in the affine plane $\text{AG}(2, 8)$ intersect in at most $\lambda = 1$ point (and parallel lines of equal slope b intersect in $\lambda = 0$ points), pure XOR/AND logic synthesizes an optimal Steiner $S(2, 8, 64)$ code with zero lookup tables and zero arithmetic carries, then zeroes the upper HPA slice bits before Stage 1 so downstream memory controllers never fault.

Theorem 3.3 (Pre-Core Carry-Free $\mathbb{F}_8$ Steiner $S(2, 8, 64)$ / MDS Expander and Tweakable Multi-Width Feistel HPA Sealing). *Let an incoming Mode-C request slot carry* $\text{SPID}_i[13:0]$ *and a* c_{bits}*-bit descriptor* cid_i *(*$c_{\text{bits}} \in \{8, 10\}$*) in the upper bits of an even-width slice* $\text{HPA}_{\text{wire},i}[12 + 2n - 1 : 12]$ *across* $2n \in \{24, 34, 36, 40\}$ *(*$n \in \{12, 17, 18, 20\}$*). In the 10-bit profile (*$\text{cid}_i \in \{0, \ldots, 1023\}$*), the Cycle-0 expander partitions* cid_i *into: (1)* ***Affine Lines & Quadratic MDS Cohorts*** *(*$\text{cid}_i \in \{0, \ldots, 511\}$*): with* $c = \text{cid}_i[8:6], b = \text{cid}_i[5:3], a = \text{cid}_i[2:0] \in \mathbb{F}_8$*, asserts* $\text{fine}_i = 1$ *and evaluates in each byte lane* $r \in \{0, \ldots, 7\}$*:*

$$y_r(a, b, c) := a \oplus (b \odot x_r) \oplus (c \odot x_r^2) \in \mathbb{F}_8, \qquad \Phi_{\mathbb{F}_8}(a, b, c) := \left[\left(\texttt{8'd1} \ll y_7\right) \parallel \cdots \parallel \left(\texttt{8'd1} \ll y_0\right) \right]; \tag{3}$$

(2) ***Vertical Byte Lines*** *(*$\text{cid}_i \in \{512, \ldots, 519\}$*): asserts* $\text{fine}_i = 1$ *and drives* $\Phi_{\mathbb{F}_8}(\text{cid}_i) =$ `64'hFF` $\ll (8 \cdot \text{cid}_i[2:0]) \in J(64, 8)$*; (3)* ***64 1-Hot Coarse Arenas*** *(*$\text{cid}_i \in \{576, \ldots, 639\}$*): asserts* $\text{fine}_i = 0$ *and drives* $r_i =$ `64'd1` $\ll \text{cid}_i[5:0]$*; and (4)* ***Invalid Codepoints (remaining 440 codes):*** *drives* $r_i = \mathbf{0}_{64}$ *(*$584/1024 = 73/128$ *valid ratio; or* $64 + 128 + 64 = 256/256$ *in the 8-bit profile). Then:*

(i) ***Steiner*** $S(2, 8, 64) \cong \text{AG}(2, 8)$***, Parallel-Slope Resolvability, &*** $[8, 3, 6]_8$ ***MDS Bounds:*** *The* $64 + 8 = 72$ *linear/vertical codewords construct all* $\binom{64}{2}/\binom{8}{2} = 72$ *lines of* $\text{AG}(2, 8) \cong S(2, 8, 64)$*, achieving* $A(64, 14, 8) = 72$ *with equality (*$\lambda = 1 \implies R_{\text{exact}} = 7$ *across all 2,556 pairs). Within any fixed parallel slope* $b_0 \in \mathbb{F}_8 \cup \{\infty\}$*, the 8 lines are pairwise bit-disjoint (*$\lambda = 0$ *across all 252 intra-slope pairs), whereas inter-slope lines intersect in* $\lambda = 1$ *bit across all 2,304 pairs. Across all 520 extended* $[8, 3, 6]_8$ *MDS cohorts (134,940 pairs),* $\lambda \leq 2 \implies R_{\text{exact}} = 3$ *(whereas* $S(3, 8, 64)$ *does not exist since* $62/6 \notin \mathbb{Z}$*; Lean 4* `thm17_gf8_steiner_ag28_slope_uniqueness`*, Table 8, Row 12);*

(ii) ***Coprime Feistel Diffusion & Speculative-Read-Safe Write Lockout:*** *For all* $n \in \{12, 17, 18, 20\}$, $\gcd(7, n) = 1$, *achieving full half-word diffusion in 6 two-round stages with* SPID_i *bound into* T_i. *Unauthenticated Mode-C.2 probes fail* $(\mathrm{guard}_g == \mathbf{0}_g) \wedge \mathrm{cid_valid}_i$ *with probability* $1 - \frac{73}{128}2^{-g}$*; on an architectural* `M2S_RwD` *write guard failure (*`slot_is_rwd_hdr_i`$[i]$ == 1*), hardware clears* `CSR_0x70`$[\mathrm{SPID}_i[5:0]] \leftarrow$ `1'b0` *and increments* `CSR 0x48`, *while speculative* `M2S_Req` *reads never lock out an innocent VM; and*

(iii) ***Canonical HPA Slice Restoration:*** *Cycle 0 strips* $\{\mathrm{cid}_i, \mathrm{guard}_g\} \leftarrow \mathbf{0}$ *across the upper* $(c_{\mathrm{bits}} + g)$ *bits of the sealed slice (*$\mathrm{HPA}_i[12 + 2n - 1 : 12 + 2n - (c_{\mathrm{bits}} + g)] \leftarrow \mathbf{0}$, *e.g.,* $\mathrm{HPA}_i[45:32] \leftarrow \mathbf{0}$ *for* $2n = 34, c_{\mathrm{bits}} + g = 14$*) and evaluates* `coarse_ok`$_i = (P_{\mathrm{coarse}} \gg \mathrm{arena_idx}_i)$ & `1'b1` *prior to Stage 1, preventing downstream AXI* `DECERR` *faults.*

Proof of Theorem 3.3. Each fine codeword sets one bit in each of the 8 byte lanes (or all 8 bits of one vertical lane), so $\mathrm{wt}(\Phi_{\mathbb{F}_8}) \equiv 8$. Over $\mathbb{F}_8$, two linear codewords with $b_1 = b_2$ and $a_1 \neq a_2$ have $y_r(a_1, b_1, 0) \oplus y_r(a_2, b_1, 0) = a_1 \oplus a_2 \neq 0$ for all $r \in \mathbb{F}_8$ ($\lambda = 0$ intra-slope). If $b_1 \neq b_2$, $d := b_1 \oplus b_2 \in \mathbb{F}_8^\times$ has inverse $d^{-1} = d^6$, giving the unique intersection root $x_r = d^{-1} \odot (a_1 \oplus a_2) \in \mathbb{F}_8$ ($\lambda = 1$ inter-slope). Non-zero quadratic differences in $\mathbb{F}_8[x]$ have degree ≤ 2 and thus at most $\lambda = 2$ roots ($R_{\mathrm{exact}} = 3$). □

Corollary 3.4 (Distance-Doubling $\mathbb{F}_8^8 \hookrightarrow J(64, 8)$ Isometry & Tight Welch–Sidelnikov Frame Optimality). *Under the 1-hot byte-lane embedding* $\psi_{1\text{-}hot} : \mathbb{F}_8^8 \hookrightarrow J(64, 8)$ *and the* 72×64 *incidence matrix* N *of* $\mathrm{AG}(2, 8) \cong S(2, 8, 64)$*: (i) for all* $\mathbf{y}, \mathbf{y}' \in \mathbb{F}_8^8$ *with agreement count* $\lambda(\mathbf{y}, \mathbf{y}')$,

$$d_H^{\mathbb{F}_2}\big(\psi_{1\text{-}hot}(\mathbf{y}), \psi_{1\text{-}hot}(\mathbf{y}')\big) = 2\, d_H^{\mathbb{F}_8}(\mathbf{y}, \mathbf{y}') = 2\big(8 - \lambda(\mathbf{y}, \mathbf{y}')\big), \tag{4}$$

and a t*-of-*8 *survival threshold guarantees* $P_{\mathrm{alias}} \equiv 0$ *up to* $R_{\mathrm{exact}}(t, \lambda) = \lfloor (8 - t)/\lambda \rfloor$*; and (ii) because every point of* $\mathbb{F}_8^2$ *lies on* $r_{\mathrm{deg}} = 9$ *lines and every pair of points lies on* $\lambda_{\mathrm{design}} = 1$ *line,* $N^T N = 8I_{64} + J_{64}$ *with eigenvalues* $\mu_0 = 72$ *(mult.* 1*) and* $\mu_1 = 8$ *(mult.* 63*), forming a tight two-distance frame in* $\mathbb{R}^{64}$ *achieving the Welch–Sidelnikov bound [7].*

Proof of Corollary 3.4. Each byte lane $r \in \{0, \ldots, 7\}$ contributes binary Hamming distance 0 if $y_r = y'_r$ and 2 if $y_r \neq y'_r$, giving Eq. (4). In $\mathrm{AG}(2, 8)$, $(N^T N)_{u,v}$ equals 9 for $u = v$ and 1 for $u \neq v$, yielding $N^T N = 8I_{64} + J_{64}$. □

3.4 Dual-Submode Cycle-0 Vieta Polynomial-Ideal & Time-Wheel / Steiner Tensor Sieve over $\mathbb{F}_{2^8}[Z]$

Any Boolean cover-free family over v bits supporting N cohorts up to R revocations obeys the Erdős–Frankl–Füredi bound $v = \Omega((R^2/\log R)\log N)$ [27, 26], and revoking all 8 lines of a single parallel slope b_0 in $\mathrm{AG}(2, 8)$ sets $P_{\mathrm{fine}} = \mathbf{0}_{64}$. Furthermore, driving an all-ones vector $r_i = \mathbf{1}_{64}$ with $p_{s,\mathrm{fine}} = \mathbf{1}_{64}$ into Stage 1's `O5` check $\Pi_{P_{\mathrm{fine}}}(\mathbf{1}_{64}) = p_{t,\mathrm{fine}} == \mathbf{0}_{64}$ (as in Coarse-Only Vieta mode `CSR 0x78`$[51:50]$ = `2'b10`) would fire only when all 64 bits of `CSR 0x68` are zero—collapsing the 64-bit fine temporal mask into a 1-bit global switch, idling Stage 1's 256 `LUT6_2` primitives, and burning scarce Vieta roots on routine lease expirations.

To break both the quadratic cover-free barrier and the all-ones temporal degeneracy without modifying a single gate of the 1,585-LUT Stage-1/2 core, our **Dual-Submode Cycle-0 Vieta Polynomial-Ideal & Time-Wheel / Steiner Tensor Sieve** (`CSR 0x78`$[51:50] \in \{$`2'b01`, `2'b11`$\}$, complementing Boolean mode `2'b00` and Coarse-Only Vieta mode `2'b10`) over $\mathbb{F}_{2^8} \cong \mathbb{F}_2[z]/(z^8 + z^4 + z^3 + z + 1)$ stores the 8 elementary symmetric Vieta coefficients $(\sigma_1^{(t)}, \ldots, \sigma_8^{(t)}) \in (\mathbb{F}_{2^8})^8$ (64 bits) in `CSR 0x60_bank[0:1]` with 4-bit degree counters $d_{\mathrm{rev}}^{(t)} \in \{0, \ldots, 8\}$ (`CSR 0x20`[11:8]) and saturation lockout bit (`CSR 0x20`[12]), while pairing Cycle 0's Vieta non-root check with either an autonomous 64-bucket 1-hot time-wheel in `CSR 0x68` (Submode V.1, `2'b01`) or the $\Phi_{\mathbb{F}_8}$ Steiner/MDS expander (Submode V.2, `2'b11`).

Motivation. Any Boolean superimposed code over $v = 64$ bits faces a fundamental combinatorial wall: Erdős–Frankl–Füredi theory demands $v = \Omega(R^2 \log N)$ bits, and revoking all 8 parallel lines of a single slope b_0 in $\mathrm{AG}(2, 8)$ zeroes every bit of $P_{\mathrm{fine}} = \mathbf{0}_{64}$. We break this quadratic barrier by reinterpreting the 64-bit register `CSR 0x60` not as 64 independent Boolean flags, but as the 8 elementary symmetric coefficients $(\sigma_1, \ldots, \sigma_8) \in (\mathbb{F}_{2^8})^8$ of a degree-8 annihilator polynomial $\Lambda^{(t)}(Z) = Z^{8-d_{\mathrm{rev}}} \prod_{j=1}^{d_{\mathrm{rev}}} (Z \oplus \alpha_j)$ over the finite field $\mathbb{F}_{2^8}$. Because a polynomial of degree $d_{\mathrm{rev}} \leq 8$ over a field has *at most* d_{rev} *roots*, revoking a cohort α_j simply multiplies $\Lambda^{(t)}(Z)$ by $(Z \oplus \alpha_j)$: every revoked cohort evaluates to $\Lambda^{(t)}(\alpha_j) = 0_{\mathbb{F}_{2^8}}$, every unrevoked cohort is guaranteed to evaluate to $\Lambda^{(t)}(z) \neq 0_{\mathbb{F}_{2^8}}$ ($P_{\mathrm{alias}} \equiv 0$ up to $R_{\mathrm{exact}} = 8$, even within a single slope), and because $\mathbb{F}_{2^8}$ has no zero divisors, no software write can ever un-revoke an existing root.

Theorem 3.5 (Dual-Submode Cycle-0 Vieta Polynomial-Ideal & Time-Wheel / Steiner Tensor Sieve over $\mathbb{F}_{2^8}[Z]$). *Let each fine 4-KB cohort carry non-zero descriptor* $\mathrm{cid}_{\mathrm{lo},i} \in \mathbb{F}_{2^8}^{\times}$ *($N = 255$). In Vieta modes (*`CSR 0x78`$[50] = 1$*), each bank of* `CSR 0x60_bank[0:1]` *stores* $\boldsymbol{\sigma}^{(t)} = (\sigma_1^{(t)}, \ldots, \sigma_8^{(t)}) \in (\mathbb{F}_{2^8})^8$ *($\boldsymbol{\sigma}^{(0)} = \mathbf{0}_{64}, d_{\mathrm{rev}}^{(0)} = 0$) representing the monic degree-8 annihilator polynomial:*

$$\Lambda^{(t)}(Z) := Z^8 \oplus \bigoplus_{k=1}^{8} \sigma_k^{(t)} \odot Z^{8-k} = Z^{8-d_{\mathrm{rev}}^{(t)}} \odot P^{(t)}(Z), \qquad P^{(t)}(Z) := \prod_{j=1}^{d_{\mathrm{rev}}^{(t)}} (Z \oplus \alpha_j). \tag{5}$$

When cohort $\alpha \in \mathbb{F}_{2^8}^{\times}$ is written to `BAR0 CSR 0x60`*, hardware evaluates $\Lambda^{(t)}(\alpha)$ and updates $(\boldsymbol{\sigma}^{(t+1)}, d_{\mathrm{rev}}^{(t+1)})$ via:*

$$\sigma_k^{(t+1)} := \begin{cases} \sigma_k^{(t)} \oplus \left(\alpha \odot \sigma_{k-1}^{(t)}\right), & \text{if } \Lambda^{(t)}(\alpha) \neq 0 \wedge d_{\mathrm{rev}}^{(t)} < 8, \\ \sigma_k^{(t)}, & \text{if } \Lambda^{(t)}(\alpha) = 0 \text{ (duplicate)}, \end{cases} \tag{6}$$

with $\sigma_0^{(t)} \equiv 1_{\mathbb{F}_{2^8}}$ (locking fail-closed at `CSR 0x20`*[12] if a 9th distinct root arrives before bank rotation). In* ***Cycle 0****, for each slot $i \in \{0,1,2,3\}$ with bank $b_i =$* `slot_bank_sel_i`$[i]$*, hardware conjoins the explicit non-zero descriptor guard $(\mathrm{cid}_{\mathrm{lo},i} \neq 0_{\mathbb{F}_{2^8}})$ and registers*

$$\texttt{slot_coarse_arena_ok_i}[i] \Leftarrow \left(\mathrm{cid}_{\mathrm{lo},i} \neq 0_{\mathbb{F}_{2^8}}\right) \wedge \left(\Lambda_{b_i}^{(t)}(\mathrm{cid}_{\mathrm{lo},i}) \neq 0_{\mathbb{F}_{2^8}}\right) \wedge \texttt{arena_ok}_i \wedge \texttt{bank_valid_q}[b_i]$$

onto Stage 1's input (noting that while $Z^{8-d_{\mathrm{rev}}^{(t)}}$ vanishes at $Z = 0$ for $0 \leq d_{\mathrm{rev}}^{(t)} \leq 7$, conjoining $\mathrm{cid}_{\mathrm{lo},i} \neq 0_{\mathbb{F}_{2^8}}$ closes the $d_{\mathrm{rev}}^{(t)} = 8$ boundary case where $\Lambda^{(t)}(0) = \sigma_8^{(t)} = \prod_{j=1}^{8} \alpha_j \neq 0_{\mathbb{F}_{2^8}}$) and selects:

- ***Submode V.1 (Autonomous 64-Bucket Time-Wheel + Vieta Quarantine,*** `CSR 0x78`$[51:50] =$ `2'b01`***):*** *Cycle 0 decodes slot i's 6-bit lease-expiration bucket $\tau_i =$* `Meta_Tag`$_i[5:0] \in \{0, \ldots, 63\}$ *into 1-hot $r_i =$* `64'd1` $\ll \tau_i \in J(64, 1)$ *with $p_{s,\mathrm{fine}}^{[b_i]} = \mathbf{1}_{64}$;*
- ***Submode V.2 (Steiner/MDS $\otimes$ Vieta Tensor-Product Sieve,*** `CSR 0x78`$[51:50] =$ `2'b11`***):*** *Cycle 0 drives $r_i = \Phi_{\mathbb{F}_8}(\mathrm{cid}_{\mathrm{hi},i}) \in J(64, 8)$ into Stage 1's 256* `LUT6_2` *cells against $P_{\mathrm{fine}}^{[b_i]} = p_{s,\mathrm{fine}}^{[b_i]} \wedge p_{t,\mathrm{fine}}$.*

Then:

(i) ***Linear Bit Scaling*** *$v = R \cdot m$,* ***Exact*** *$P_{\mathrm{alias}}(r) \equiv 0$* ***up to*** *$R_{\mathrm{exact}} = 8$,* ***& Core Timing Invariance:*** *Over the integral domain $\mathbb{F}_{2^8}[Z]$, $\Lambda^{(t)}(Z)$ has at most $d_{\mathrm{rev}}^{(t)} \leq 8$ roots in $\mathbb{F}_{2^8}^{\times}$:*

$$\mathrm{Roots}\left(\Lambda^{(t)}\right) := \left\{z \in \mathbb{F}_{2^8}^{\times} : \Lambda^{(t)}(z) = 0_{\mathbb{F}_{2^8}}\right\} = \left\{\alpha_1, \ldots, \alpha_{d_{\mathrm{rev}}^{(t)}}\right\}. \tag{7}$$

Thus, for every revocation set $\{\alpha_1, \ldots, \alpha_r\} \subset \mathbb{F}_{2^8}^{\times}$ of size $1 \leq r \leq 8$ (even within a single slope), $P_{\mathrm{alias}}(r) \equiv 0$ across all 255 cohorts with $v = 8 \times 8 = 64$ bits while keeping the 1,585-LUT Stage-1/2 core (0.124 ns sieve) 100% unchanged;

(ii) ***Unconditional Polynomial-Ideal Root Permanence:*** *Because $\mathbb{F}_{2^8}$ has no zero divisors, the active root factor $P^{(t+1)}(Z) = (Z \oplus \alpha) \odot P^{(t)}(Z)$ generates a descending polynomial-ideal chain $\langle P^{(0)} \rangle \supseteq \cdots \supseteq \langle P^{(t)} \rangle$ in $\mathbb{F}_{2^8}[Z]$. In the active shift-multiply branch ($\Lambda^{(t)}(\alpha) \neq 0 \wedge d_{\mathrm{rev}}^{(t)} < 8$, where $\sigma_8^{(t)} = 0_{\mathbb{F}_{2^8}}$), update (6) satisfies $Z \odot \Lambda^{(t+1)}(Z) = (Z \oplus \alpha) \odot \Lambda^{(t)}(Z)$, whereas in the duplicate/saturated branch ($\Lambda^{(t)}(\alpha) = 0 \vee d_{\mathrm{rev}}^{(t)} = 8$), $\Lambda^{(t+1)}(Z) = \Lambda^{(t)}(Z)$. Consequently, for every non-zero descriptor $z^* \in \mathbb{F}_{2^8}^{\times}$ and any $\alpha \in \mathbb{F}_{2^8}$:*

$$\Lambda^{(t)}(z^*) = 0_{\mathbb{F}_{2^8}} \implies z^* \odot \Lambda^{(t+1)}(z^*) = (z^* \oplus \alpha) \odot \Lambda^{(t)}(z^*) = 0_{\mathbb{F}_{2^8}} \implies \Lambda^{(t+1)}(z^*) = 0_{\mathbb{F}_{2^8}}, \tag{8}$$

preventing any compromised FM write from un-revoking a root z^* (Lean 4 `thm20_vieta_gf256_ideal_ratchet_monotonic_and_degree_bound`*, Table 8, Row 14); and*

(iii) ***Exact 64-Bucket Time-Wheel Selectivity*** *($\Delta d_{\mathrm{rev}} = 0$)* ***& Tensor Namespace Multiplication*** *($18,360$ / $132,600$* ***Cohorts):*** *In Submode V.1, Stage 1's* `O5` *projection satisfies:*

$$\begin{aligned} &\Pi_{P_{\mathrm{fine}}^{[b_i]}}\left(\texttt{64'd1} \ll \tau_i\right) = \mathbf{0}_{64} \iff p_{t,\mathrm{fine}}[\tau_i] = 0, \\ &\mathrm{Pass}_{\mathrm{V.2}}(\mathrm{cid}_{\mathrm{hi}}, \mathrm{cid}_{\mathrm{lo}}) \iff \left(\mathrm{cid}_{\mathrm{lo}} \neq 0_{\mathbb{F}_{2^8}}\right) \wedge \left(\Lambda_{b_i}^{(t)}(\mathrm{cid}_{\mathrm{lo}}) \neq 0_{\mathbb{F}_{2^8}}\right) \wedge \left(\Pi_{P_{\mathrm{fine}}^{[b_i]}}(\Phi_{\mathbb{F}_8}(\mathrm{cid}_{\mathrm{hi}})) \neq \mathbf{0}_{64}\right). \end{aligned} \tag{9}$$

A monotonic hardware tick ratchets `CSR 0x68` *bit-by-bit to expire ephemeral serverless/KV-cache leases automatically with zero MMIO traffic and zero Vieta degree growth ($\Delta d_{\mathrm{rev}} = 0$), while Submode V.2 expands the zero-BRAM namespace to $72 \times 255 = 18,360$ Steiner–Vieta cohorts (or $520 \times 255 = 132,600$ MDS–Vieta cohorts; Lean 4* `thm26_vieta_timewheel_and_tensor_sieve_selectivity`*, Table 8, Row 17).*

Proof of Theorem 3.5. Over the UFD $\mathbb{F}_{2^8}[Z]$, $\Lambda^{(t)}(\mathrm{cid}') = (\mathrm{cid}')^{8-r}\prod_{j=1}^{r}(\mathrm{cid}' \oplus \alpha_j)$ vanishes at $\mathrm{cid}' \in \mathbb{F}_{2^8}^{\times}$ iff $\mathrm{cid}' = \alpha_j$ for some $j \in \{1,\dots,r\}$, forcing $P_{\mathrm{alias}}(r) \equiv 0$, while $(\mathrm{cid}_{\mathrm{lo},i} \neq 0_{\mathbb{F}_{2^8}})$ rejects $\mathrm{cid}_{\mathrm{lo},i} = \texttt{8'h00}$ across all $0 \le d_{\mathrm{rev}}^{(t)} \le 8$. For $d_{\mathrm{rev}}^{(t)} < 8$, $\sigma_8^{(t)} = 0_{\mathbb{F}_{2^8}}$, so expanding $Z \odot \Lambda^{(t+1)}(Z) = (Z \oplus \alpha) \odot \Lambda^{(t)}(Z)$ yields the active branch of Eq. (6) and multiplying by $(z^*)^{-1} \in \mathbb{F}_{2^8}^{\times}$ together with $b \odot 0_{\mathbb{F}_{2^8}} = 0_{\mathbb{F}_{2^8}}$ implies Eq. (8) (while $\Lambda^{(t+1)}(z^*) = \Lambda^{(t)}(z^*) = 0_{\mathbb{F}_{2^8}}$ holds trivially in the duplicate/saturated branch). In Submode V.1, $(\texttt{64'd1} \ll \tau_i) \wedge (\mathbf{1}_{64} \wedge p_{t,\mathrm{fine}})$ has bit j equal to $\mathbb{I}(j = \tau_i) \wedge p_{t,\mathrm{fine}}[j]$, which vanishes iff $p_{t,\mathrm{fine}}[j] = 0$ at $j = \tau_i$. In Submode V.2, gating Stage 1's `O5` reduction by `slot_coarse_arena_ok_i[i]` forms the direct tensor product $C_{\mathrm{Steiner/MDS}} \times \mathbb{F}_{2^8}^{\times}$ (`thm20`, `thm26`). □

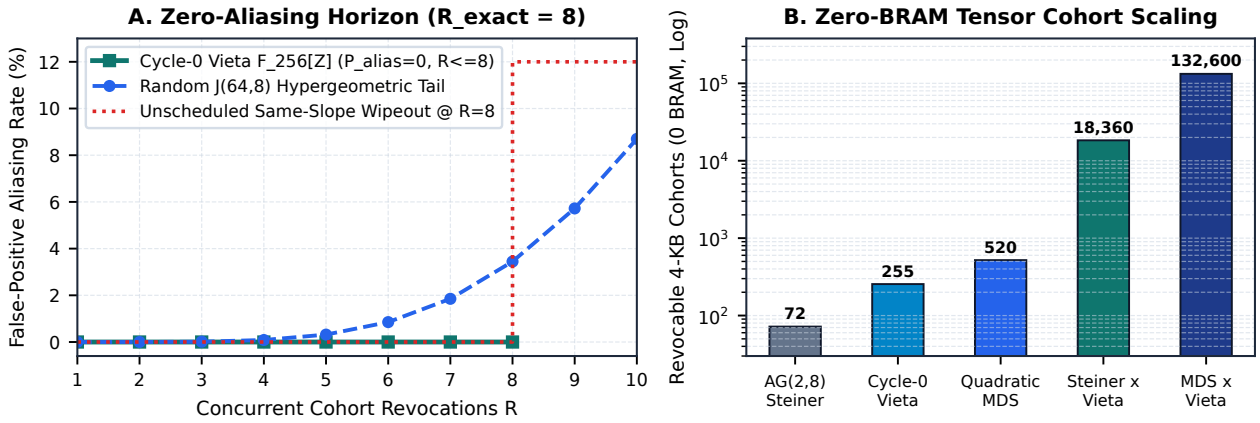


Figure 2. Dual-Submode Cycle-0 Vieta $\mathbb{F}_{2^8}[Z]$ polynomial-ideal and Steiner/MDS tensor-sieve cohort scaling: $P_{\mathrm{alias}}(r) \equiv 0$ up to $R_{\mathrm{exact}} = 8$ revocations (including same-slope bursts, whereas unscheduled same-slope revocation at $R = 8$ incurs a 100% wipeout $P_{\mathrm{fine}} = \mathbf{0}_{64}$, off-scale on the 0–12% vertical axis) across 18,360 Steiner–Vieta and 132,600 MDS–Vieta cohorts in 0 BRAM.

So What? (Architectural Consequence). Any Boolean bitmap or Bloom filter over 64 bits suffers either from the Erdős–Frankl–Füredi quadratic cover-free limit or from catastrophic wipeout when 8 cohorts in the same affine slope are revoked simultaneously. Figure 2 demonstrates that lifting the 64-bit `BAR0` register into an $\mathbb{F}_{2^8}[Z]$ Vieta annihilator polynomial paired with the Stage-1 Steiner/time-wheel sieve achieves strictly zero false-positive revocations ($P_{\mathrm{alias}} \equiv 0$) up to $R_{\mathrm{exact}} = 8$ arbitrary revocations across 18,360–132,600 cohorts in 0 bits of SRAM, while expiring routine TTL leases autonomously with $\Delta d_{\mathrm{rev}} = 0$ root consumption.

Scope & Price Paid (No Free Lunch — Structured Algebraic Cohorts vs. Unbounded Scatter-Gather). Eliminating $\Theta(M)$ DRAM permission tables is not a free lunch: in exchange for 0 BRAM and 0.124 ns Stage-1 evaluation, memory allocations must be grouped into structured algebraic cohorts supporting $R_{\mathrm{exact}} = 7$ arbitrary revocations in AG(2, 8) or $R_{\mathrm{exact}} = 8$ in $\mathbb{F}_{2^8}[Z]$ per active bank before triggering a background dual-bank rotation in `CSR 0x60_bank[0:1]` (or autonomous 64-bucket TTL lease expiration in `CSR 0x68`). Workloads demanding thousands of ungrouped, arbitrary single-page revocations per microsecond without epoch batching still require conventional stateful MMU page tables.

4 Spanning-Forest Capability Propagation, Permutation-Orbit Chord Clamping, PHY-Gated Dual-Forest Failover, and Stage-2 `LUT6_2` Precision Clamping

Motivation. In a multi-hop CXL switch fabric containing routing loops ($\beta_1(G) = |E| - |V| + \beta_0 > 0$), a compromised or deputized node can forward a request around a cycle to launder permissions or trigger an infinite distributed agreement loop. By decomposing the fabric graph offline into a spanning forest T_{forest} plus β_1 fundamental cycle chords $e^* \in E \setminus T_{\mathrm{forest}}$ and enforcing port anti-hairpinning, every non-backtracking cyclic walk in G is mathematically forced to cross at least one fundamental chord. Placing a single-gate combinational meet ratchet Γ_{e^*} on each of those β_1 chords collapses the all-pairs Kleene-star transitive closure of the cyclic graph identically to the loop-free spanning forest in zero online consensus rounds.

Theorem 4.1 (Spanning-Forest Decomposition, Anti-Hairpinning, & Fundamental-Chord Cycle Interception). *Let the multi-hop interconnect graph $G = (V, E)$ with incidence matrix $B \in \{-1, 0, +1\}^{|V|\times|E|}$ and β_0 connected components be decomposed in $O(|E|)$ offline time [28] via a spanning forest $T_{\mathrm{forest}} \subset E$ into $|V| - \beta_0$ tree branches $e \in T_{\mathrm{forest}}$ and $\beta_1 = \dim\ker(B) = |E| - |V| + \beta_0$ fundamental cycle chords $e^* \in E \setminus T_{\mathrm{forest}}$, with port anti-hairpinning (`egress_port ≠ ingress_port`) at every switch port. Then: (i) every chord-avoiding subgraph $S \subseteq T_{\mathrm{forest}}$ has $\beta_1(S) = 0$ and admits zero non-backtracking closed walks of length ≥ 2; and (ii) every simple cycle $C \subseteq E$ of length ≥ 3 intersects $E \setminus T_{\mathrm{forest}}$ in at least one chord $e^* \in C \setminus T_{\mathrm{forest}}$, so ratcheting the β_1 chords intercepts 100% of cyclic walks in G.*

Proof of Theorem 4.1. Any $S \subseteq T_{\text{forest}}$ is acyclic, so any closed walk in S must backtrack immediately ($u \to v \to u$), which is blocked by `egress_port` $\neq$ `ingress_port`. Any simple cycle $C \subseteq E$ has $\mathbf{1}_C \in \ker(B) \setminus \{\mathbf{0}\}$; since columns of B on T_{forest} are linearly independent, $\ker(B) \cap \mathbb{R}^{T_{\text{forest}}} = \{\mathbf{0}\}$, forcing $C \cap (E \setminus T_{\text{forest}}) \neq \emptyset$. □

Theorem 4.2 (Spanning-Forest Zero-Iteration Agreement, Chord-Ratchet Cycle Clamping, & Dioidal Kleene-Star Tree-Closure Collapse). *Let each node and link carry capability tuples in the bounded meet-semilattice $(\mathcal{L}, \preceq, \sqcap)$ with $\mathcal{L} = \mathbb{F}_2^{64} \times D_\perp$ (Lemma A.1), where $\tau_1 \sqcap \tau_2 := [(r_1 \wedge r_2) \parallel (\text{SPID}_1 \wedge_{\text{flat}} \text{SPID}_2)]$ and $\alpha^*_{v \trianglelefteq e}(x_v) = x_v \sqcap m_{v,e}$. Define the consistency operator $\mathcal{L}^0_{\text{cons}}$ and edge-admissible set $\Gamma^1_{\text{adm}}(G)$ by:*

$$(\mathcal{L}^0_{\text{cons}} x)_v := x_v \sqcap \prod_{e=(u,v)\in E} \big(\alpha^*_{u \trianglelefteq e}(x_u) \sqcap \alpha^*_{v \trianglelefteq e}(x_v)\big), \qquad \Gamma^0_{\text{adm}}(G) := \text{Fix}(\mathcal{L}^0_{\text{cons}}), \qquad \Gamma^1_{\text{adm}}(G) := \delta^0_{\text{edge}}\big(\Gamma^0_{\text{adm}}(G)\big), \tag{10}$$

where $(\delta^0_{\text{edge}} x)_e := \alpha^*_{u \trianglelefteq e}(x_u) \sqcap \alpha^*_{v \trianglelefteq e}(x_v)$. If tree branches $e \in T_{\text{forest}}$ satisfy $x_u = \alpha^*_{u \trianglelefteq e}(x_u) = \alpha^*_{v \trianglelefteq e}(x_v) = x_v$ and every fundamental chord $e^* = (u, v) \in E \setminus T_{\text{forest}}$ instantiates the combinational meet ratchet across $C(e^*) \subset T_{\text{forest}} \cup \{e^*\}$:

$$\tau^{\text{out}}_{e^*} := \tau^{\text{in}}_{e^*} \sqcap \Gamma_{e^*}, \qquad \Gamma_{e^*} := \alpha^*_{u \trianglelefteq e^*}(x_u) \sqcap \alpha^*_{v \trianglelefteq e^*}(x_v) \sqcap \prod_{w \in C(e^*)} x_w, \tag{11}$$

then: (i) tree-consistent vertex assignments are stationary under $\mathcal{L}^0_{\text{cons}}$ in 0 online consensus rounds; (ii) over $\mathcal{L}$, the meet operators $\{\mu_a(\tau) = \tau \sqcap a\}$ factor through $(\mathbb{B}^{\beta_1}, \vee)$, clamping any closed walk W traversing chords $\{e^*_1, \ldots, e^*_r\}$ to $\tau^{\text{out}}_W \preceq \prod_{q=1}^r \Gamma_{e^*_q}$ in 1 gate delay per chord; and (iii) over the dioid $\mathcal{D}_{64} = (\mathbb{F}_2^{64}, \vee, \wedge, \mathbf{0}_{64}, \mathbf{1}_{64})$, the all-pairs Kleene-star closure across G collapses identically to the spanning-forest closure:

$$\big(A_{G,\text{ratcheted}}\big)^*_{a,b} := \bigvee_{k=0}^{|V|-1} \big(A^k_{G,\text{ratcheted}}\big)_{a,b} \equiv \big(A_{T_{\text{forest}}}\big)^*_{a,b} = \prod_{w \in P_{T_{\text{forest}}}(a,b)} x_w. \tag{12}$$

Proof of Theorem 4.2. Tree agreement gives $(\mathcal{L}^0_{\text{cons}} x)_v = x_v$, proving Claim (i). Associativity, commutativity, and idempotence of $\sqcap$ prove Claim (ii). Replacing any chord $e^* = (u, v)$ in W by $P_{T_{\text{forest}}}(u, v) \subset C(e^*)$ yields a tree path dominating Γ_{e^*}, so absorption $X \vee Y = X$ for $Y \preceq X$ in $\mathcal{D}_{64}$ proves Eq. (12). □

4.1 Orthogonal Pass/Violation Bitwise Split, Non-Zero Capability Check, and Monotonic `BAR0` Ratcheting

Corollary 4.3 (Orthogonal Pass/Violation Bitwise Split, Non-Zero Capability Sieve, & Cold-Reset-Isolated AND-Ratchet). *For active mask $P_{\text{eff}} := p_s \wedge p_t \in \mathbb{F}_2^{64}$, pass projection $\Pi_{P_{\text{eff}}}(r_i) := r_i \wedge P_{\text{eff}}$, violation operator $\Delta_{P_{\text{eff}}}(r_i) := r_i \wedge \neg(p_s \wedge p_t)$, and admissible non-zero set*

$$C_{\text{adm}} := \big\{(r_i, \text{SPID}_i) : \ r_i \neq \mathbf{0}_{64},\ \text{SPID}_i \neq 0,\ \text{SPID}_i[13:6] == \texttt{port_spid_prefix},\ \texttt{CSR_0x70}[\text{SPID}_i[5:0]] == 1\big\} :$$

(i) $\Pi^2_{P_{\text{eff}}} = \Pi_{P_{\text{eff}}}$ and $\Delta^2_{P_{\text{eff}}} = \Delta_{P_{\text{eff}}}$ decompose $\mathbb{F}_2^{64}$ into the direct sum:

$$\mathbb{F}_2^{64} = \ker(\Delta_{P_{\text{eff}}}) \oplus \text{im}(\Delta_{P_{\text{eff}}}), \qquad r_i = \Pi_{P_{\text{eff}}}(r_i) \oplus \Delta_{P_{\text{eff}}}(r_i), \tag{13}$$

while intersecting with C_{adm} rejects $r_i = \mathbf{0}_{64}$ and aliased tenant IDs; and (ii) enforcing $M^{(t+1)} := M^{(t)} \wedge w$ on `BAR0 CSR 0x00, 0x08, 0x60, 0x68, 0x70` while isolating security CSRs from `host_warm_flr_rst_n` guarantees $(M^{(t)} \wedge w) \wedge \neg M^{(t)} = \mathbf{0}_{64}$, preventing $0 \to 1$ escalation across warm FLRs and incrementing `CSR 0x48` whenever $(w \wedge \neg M^{(t)}) \neq \mathbf{0}_{64}$.

Proof of Corollary 4.3. In $\mathbb{F}_2^{64}$, $\Pi_{P_{\text{eff}}} + \Delta_{P_{\text{eff}}} = \text{Id}$ and $\Pi_{P_{\text{eff}}} \circ \Delta_{P_{\text{eff}}} = 0$, proving Eq. (13); Claim (ii) follows from Proposition 3.1 (Lean 4 `thm01`, `thm02`). □

4.2 Permutation-Orbit Chord Clamping, PHY-Gated Dual-Forest Failover, Fractured LUT6_2 MOESI & Graduated Mantissa Clamping, and Keyed-Dither Multi-Query Lower Bounds

Motivation. Three practical challenges arise in multi-tenant AI clusters: (1) intermediate CXL switches may apply HDM address-remapping permutations $\pi_e \in S_{64}$ whose group order reaches Landau's $g(64) = \mathrm{lcm}(8,3,5,7,11,13,17) = 2,042,040$, seemingly requiring millions of iterations to stabilize around a cycle; (2) a physical cable cut on a primary spanning-tree link requires sub-4-ns failover without allowing software MMIO writes to re-enable revoked bits; and (3) in shared-HBM foundation-model inference, co-tenant GPUs need approximate read access (BF12/E8M0) without being able to extract exact BF16 weights or issue coherent RdOwn requests to corrupt shared weights. Theorem 4.4 resolves all three at +0 CLB LUTs: every individual bit coordinate has an orbit of length $L_j \le 64$ in S_{64}; pre-certified dual forests (T_0, T_1) switch in < 4 ns directly from the hardware PCS loss-of-signal wire p_{phy}; and the 3 unused inputs of Stage 2's existing LUT2 gates are fractured into 5-shared-input LUT6_2 primitives that simultaneously downgrade RdOwn $\mapsto$ RdShared and clamp S2M_DRS floating-point mantissas across 4 precision tiers ($u \in \{0, 9, 12, 16\}$ bits).

Theorem 4.4 (Permutation-Orbit Chord Clamping, PHY-Gated Dual-Forest Failover, & Fractured LUT6_2 MOESI and Graduated Mantissa Clamping). *Let $G = (V, E)$ carry switch HDM address-remapping permutations $\pi_e \in S_{64}$ ($\pi_{e^{-1}} = \pi_e^{-1}$) with root-to-vertex tree gauge $\Pi_v := \pi_{e_m} \circ \cdots \circ \pi_{e_1} \in S_{64}$ along $P_{T_{\text{forest}}}(r_0, v)$. For each chord $e^* = (u, v) \in E \setminus T_{\text{forest}}$ with cycle permutation $H_{C(e^*)} := \Pi_v^{-1} \circ \pi_{e^*} \circ \Pi_u \in S_{64}$, although $\mathrm{ord}(H_{C(e^*)})$ in S_{64} can reach Landau's $g(64) = \mathrm{lcm}(8,3,5,7,11,13,17) = 2,042,040$, **every individual bit-coordinate orbit** $\{H^k_{C(e^*)}(j)\}_{k\ge 0} \subseteq \{0, \dots, 63\}$ **has cardinality** $L_j \le 64$. Define the coordinate-orbit-invariant chord mask:*

$$\Gamma^{\text{orb}}_{e^*} := \bigwedge_{k=0}^{63} H^k_{C(e^*)}(\Gamma_{e^*}) \in \mathbb{F}_2^{64}, \qquad \tau^{\text{out}}_{e^*} := \pi_{e^*}(\tau^{\text{in}}_{e^*}) \sqcap \Pi_v(\Gamma^{\text{orb}}_{e^*}). \tag{14}$$

To provide sub-4-ns link failover under optical/electrical cable cuts without permitting $0 \to 1$ software MMIO writes (Corollary 4.3(ii)), spectralc pre-certifies two maximally edge-disjoint spanning forests (T_0, T_1) [9, 10] with cold-POR AND-ratcheted orbit-clamped chord registers $(\Gamma^{\text{orb},(0)}_e, \Gamma^{\text{orb},(1)}_e)$ multiplexed strictly by the hardware-only Physical Coding Sublayer (PCS) / UCIe 3.0 emergency sideband loss-of-signal wire $p_{\text{phy}} =$ pcs_link_up_i $\in \{0, 1\}$ [3]:

$$\Gamma^{\text{active}}_e(p_{\text{phy}}) := p_{\text{phy}} \;?\; \Gamma^{\text{orb},(0)}_e : \Gamma^{\text{orb},(1)}_e. \tag{15}$$

Furthermore, in Stage 2, let $c_i[1:0] \in \{$2'b00, 2'b01, 2'b10, 2'b11$\}$ denote slot i's 2-bit MOESI/precision tier ceiling (distinct from cohort codewords $c \in J(64, w)$). Let 4 of the 256 dual-LUT CLB sites clamp CXL.cache D2H_Req MOESI state $x_i[1:0] \in \{$I = 2'b00 $<$ S = 2'b01 $<$ E = 2'b10 $<$ M = 2'b11$\}$ against $c_i[1:0]$ via fractured LUT6_2 INIT = 64'hCC00CC00_A2E0A2E0 (I5 = 1'b1, I4 = 1'b0, I3 = $c_i[1]$, I2 = $c_i[0]$, I1 = $x_i[1]$, I0 = $x_i[0]$), and let the remaining 252 S2M_DRS read-payload CLB sites be configured as pin-locked 5-shared-input fractured LUT6_2 primitives (I5 = 1'b1, I4 = $c_i[1]$, I3 = $c_i[0]$, I2 = viol_i, I1 = f_{2m+1}, I0 = f_{2m}, address $a = 16c_{i,1} + 8c_{i,0} + 4\mathrm{viol}_i + 2f_{2m+1} + f_{2m} \in \{0, \dots, 31\}$). Then:

*(i) **0-LUT Wire-Gauge Orbit Invariance & Sub-4-ns PHY-Gated Dual-Forest Failover:** Since $L_j \le 64$, $H_{C(e^*)}(\Gamma^{\text{orb}}_{e^*}) \equiv \Gamma^{\text{orb}}_{e^*}$ and $\tilde{\mu}^m_{e^*}(r) = H^m_{C(e^*)}(r) \wedge \Gamma^{\text{orb}}_{e^*} \le \Gamma^{\text{orb}}_{e^*}$ on the first chord traversal ($m = 1$) at 0 LUT wire-gauge cost. When a primary branch $e_{\text{cut}} \in T_0$ fails ($p_{\text{phy}} : 1 \to 0$), hardware switches to $\Gamma^{\text{orb},(1)}_e$ in < 4.0 ns with zero MMIO privilege escalation ($\Gamma^{\text{active}}_e(p_{\text{phy}}, \Gamma_0 \wedge w_0, \Gamma_1 \wedge w_1) \wedge \neg\Gamma^{\text{active}}_e(p_{\text{phy}}, \Gamma_0, \Gamma_1) = \mathbf{0}_{64}$) and collapses the Kleene-star closure over $G \setminus \{e_{\text{cut}}\}$ identically to $(A_{T_1})^*$ (Lean 4 thm21, thm28_phy_gated_dual_forest_failover_safety, Table 8, Rows 15 and 19);*

*(ii) **Single-LUT6_2 2-Bit CXL.cache MOESI Clamp at +0 CLB LUTs:** INIT = 64'hCC00CC00_A2E0A2E0 evaluates O6 = $c_i[1] \wedge x_i[1] = \ell_{\text{eff}}[1]$ and O5 = $\ell_{\text{eff}}[0] = \min(x_i, c_i)[0]$ across all 16 (c_i, x_i) states in 1 LUT delay at +0 CLB LUTs, enforcing CXL_NOP ($c_i =$ 00 $\implies$ 00), pass-through ($c_i =$ 11 $\implies x_i$), and coherent RdOwn $\mapsto$ RdShared downgrade ($x_i \in \{$10, 11$\}$, $c_i =$ 01 $\implies$ 01; Lean 4 thm22); and*

*(iii) **Zero-Area Stage-2 Fractured LUT6_2 4-Tier Graduated Mantissa Clamping & Shannon Rate-Distortion Bound:** Across all 32 input states of $(c_i[1:0], \mathrm{viol}_i, f_{2m+1}, f_{2m})$ at +0 CLB LUTs and +0 logic levels, the exact 64-bit INIT masks:*

$$\begin{aligned} \texttt{INIT}_{\text{SE}} &= \texttt{64'h0C0C0C0C_0A0A0A0A}\ (\neg\mathrm{viol}_i \textit{ hdr pass}), & \texttt{INIT}_{\text{T1}} &= \texttt{64'h0C0C0C00_0A0A0A00}\ (\neg\mathrm{viol}_i \wedge c_i \ge \texttt{2'b01}),\\ \texttt{INIT}_{\text{Mid}} &= \texttt{64'h0C0C0000_0A0A0000}\ (\neg\mathrm{viol}_i \wedge c_i \ge \texttt{2'b10}), & \texttt{INIT}_{\text{Low}} &= \texttt{64'h0C000000_0A000000}\ (\neg\mathrm{viol}_i \wedge c_i == \texttt{2'b11}),\\ \texttt{INIT}_{\text{BF16[7:6]}} &= \texttt{64'h0C0C0C00_0A0A0000}, \end{aligned} \tag{16}$$

partition each 16-bit BF16 read-payload lane (1s + 8e + 7m, bits $[15:0]$) across 8 fractured LUT6_2 primitives: $4 \times \texttt{INIT}_{\text{T1}}$ on Sign/Exponent bits $[15:8]$, $1 \times \texttt{INIT}_{\text{BF16[7:6]}}$ straddling the byte boundary (passing Exponent bit 7 on O6 when $c_i \ge$ 2'b01 while clamping Mantissa bit 6 on O5 when $c_i <$ 2'b10), $1 \times \texttt{INIT}_{\text{Mid}}$ on Mid-Mantissa bits $[5:4]$, and $2 \times \texttt{INIT}_{\text{Low}}$ on Low-Mantissa

bits $[3:0]$. *This enforces four exact precision tiers:* c_i = `2'b00` *(or* $\mathrm{viol}_i = 1$*) yields* $u = 0$ *unmasked bits (*`16'h0000` *full scrub);* c_i = `2'b01` *yields* $u = 9$ *unmasked bits (*1s + 8e = `E8M0` *shared-exponent coarse scale, clamping Mid* $[6:4]$ *and Low* $[3:0]$ *to 0);* c_i = `2'b10` *yields* $u = 12$ *unmasked bits (*1s + 8e + 3m = `BF12` *approximate inference read, clamping only Low* $[3:0]$ *to 0); and* c_i = `2'b11` *yields* $u = 16$ *unmasked bits (bit-exact* `BF16`*). For any exact weight* $W_{\text{exact}} \in \mathbb{F}_2^{16}$ *and observed output* $Y_{\text{obs}} = Q_{c_i,\mathrm{viol}_i}(W_{\text{exact}})$ *keeping* $u(c_i, \mathrm{viol}_i) \in \{0, 9, 12, 16\}$ *unmasked bits/element, Shannon's data-processing inequality bounds weight-extraction mutual information for any post-processor* $\hat{W}$ *by:*

$$I\big(W_{\text{exact}}; \hat{W} \;\big|\; c_i, \mathrm{viol}_i\big) \le I\big(W_{\text{exact}}; Y_{\text{obs}} \;\big|\; c_i, \mathrm{viol}_i\big) \le \log_2\big|\mathrm{im}(Q_{c_i,\mathrm{viol}_i})\big| = u(c_i, \mathrm{viol}_i), \tag{17}$$

defeating shared-HBM LLM weight extraction while preserving partner inference utility (Lean 4 `thm27_lut6_2_graduated_mantissa_clamp_bisimulation`*, Table 8, Row 18).*

Proof of Theorem 4.4. Each coordinate orbit $O_j = \{H_{C(e^*)}^{-k}(j) : k \ge 0\} \subseteq \{0, \ldots, 63\}$ has size $L_j \le 64$, so $H_{C(e^*)}(\Gamma_{e^*}^{\mathrm{orb}}) = \Gamma_{e^*}^{\mathrm{orb}}$. Under $p_{\mathrm{phy}} = 0$, e_{cut} carries zero flits and every chord $e \in E \setminus T_1$ is clamped by $\Gamma_e^{\mathrm{orb},(1)} \le P_{T_1}$, so $P_{T_1} \vee (r \wedge \Gamma_e^{\mathrm{orb},(1)}) = P_{T_1}$ (`thm28`). For Claim (ii), at $a = 4c_i + x_i \in \{0, \ldots, 15\}$, `16'hCC00` has 1s at $\{10, 11, 14, 15\}$ and `16'hA2E0` has 1s at $\{5, 6, 7, 9, 13, 15\}$, matching $\min(x_i, c_i)$ (`thm22`). For Claim (iii), within each 8-bit tier byte of `INIT`, the upper nibble ($\mathrm{viol}_i = 1$) is `4'h0` and the lower nibble ($\mathrm{viol}_i = 0$) is `4'hC` (1100_2, selecting f_{2m+1} on `O6`) or `4'hA` (1010_2, selecting f_{2m} on `O5`) iff tier c_i is enabled, establishing Eq. (16) (`thm27`); $|\mathrm{im}(Q_{c_i,\mathrm{viol}_i})| = 2^{u(c_i,\mathrm{viol}_i)}$ with $u(c_i, \mathrm{viol}_i) \in \{0, 9, 12, 16\}$ yields Eq. (17). □

Motivation. Deterministic mantissa truncation limits a single read to $u \in \{9, 12\}$ bits, but if an adversary can inject controlled activation perturbations $\xi^{(q)}$ prior to quantization, binary search across the deterministic rounding threshold localizes an unknown significand $w \in [0, \Delta_u)$ to precision δw in just $Q_{\text{det}} = \lceil \log_2(\Delta_u/\delta w) \rceil$ queries ($Q_{\text{det}} = 6$ at 64× resolution). By adding an unsigned pseudorandom dither $\eta^{(q)} \sim \mathrm{Unif}[0, \Delta_u)$ (equivalently, a centered dither $\tilde{\eta}^{(q)} \sim \mathrm{Unif}[-\Delta_u/2, \Delta_u/2)$ plus a $+\Delta_u/2$ rounding offset) generated per read from the counter-seeded Cycle-0 $\mathcal{X}_{31,33}$ stream prior to Stage-2 floor truncation $Q_{c_i}(z) = \Delta_u \lfloor z/\Delta_u \rfloor$, each read becomes an unbiased Bernoulli coin flip whose parameter $p(w)$ varies smoothly with w. Le Cam's two-point method and Pinsker's inequality then force any adaptive adversary to average over $\Omega((\Delta_u/\delta w)^2)$ independent CXL reads ($Q \ge 192$–512), while legitimate high-dimensional dot products ($d_{\text{model}} \ge 4,096$) self-average the zero-mean quantization residual $Y^{(q)} - (w + \xi^{(q)})$ at $O(1/\sqrt{d_{\text{model}}})$.

Theorem 4.5 (Adaptive Q-Query Sample-Complexity Lower Bound under Keyed Dithered Rate-Distortion). *While the frozen +0-LUT Stage-2 fractured* `LUT6_2` *clamp* Q_{c_i,viol_i} *(Theorem 4.4(iii)) enforces a strict static single-read mutual-information ceiling* $I(W_{\text{exact}}; Y_{\text{obs}}) \le u(c_i, \mathrm{viol}_i)$ *across all four tiers (*$u = 0$ *clamps all 16 bits to* `16'h0000` *with* $I \equiv 0$ *and* $Q = \infty$*, while* $u = 16$ *is unquantized* `BF16`*), consider the non-trivial quantized inference tiers* $u \in \{9, 12\}$ *within a fixed* `BF16` *exponent bin of mantissa-ULP quantization step size* $\Delta_u := 2^{16-u} \in \{128, 16\}$*. Suppose an adaptive model-extraction adversary controls an activation perturbation* $\xi^{(q)}$ *prior to quantization and seeks to localize a target significand parameter* $w \in [0, \Delta_u)$ *to precision* $\delta w \in (0, \Delta_u/4]$ *across* Q *adaptive queries:*

(i) ***Deterministic Boundary Dichotomy:*** *Under unkeyed deterministic truncation* $Y^{(q)} = \Delta_u \lfloor (w + \xi^{(q)})/\Delta_u \rfloor$*, if the adversary cannot perturb* w *(*$\xi^{(q)} \equiv 0$*),* Q *repeated CXL reads return the identical bitvector* $Y^{(1)} = \cdots = Y^{(Q)}$ *(0 additional bits); however, if the adversary can adaptively sweep* $\xi^{(q)}$*, binary search across the step boundary localizes* w *to accuracy* δw *in* $Q_{\text{det}} = \lceil \log_2(\Delta_u/\delta w) \rceil$ *queries; and*

(ii) ***Quadratic Sample-Complexity Wall under Keyed Stochastic Dither:*** *In the modular responder-side Cycle-0 pre-clamping extension (outside the frozen 1,585-LUT core), injecting an ephemeral unsigned pseudorandom dither* $\eta^{(q)} \sim \mathrm{Unif}[0, \Delta_u)$ *(equivalently,* $\eta^{(q)} = \Delta_u/2 + \tilde{\eta}^{(q)}$ *with centered* $\tilde{\eta}^{(q)} \sim \mathrm{Unif}[-\Delta_u/2, \Delta_u/2)$*) derived per read* q *from the existing zero-BRAM* $\mathcal{X}_{31,33}$ *PRF stream seeded with a monotonic per-read counter* q_{cnt} *(so repeated reads to an identical* `HPA` *never reuse a static dither sample, with the 1-bit carry out of bit* $15 - u$ *propagating into bit* $16 - u$*—mantissa bit 4 for* $u = 12$ *or exponent LSB bit 7 for* $u = 9$*, guarded only at* `exp = 8'hFE` *so finite values never overflow to* Inf/NaN *and sign bit 15 is never altered) prior to Stage-2* `LUT6_2` *floor truncation* $Q_{c_i}(z) = \Delta_u \lfloor z/\Delta_u \rfloor$ *yields unbiased randomized rounding* $Y^{(q)} = Q_{c_i}(w + \xi^{(q)} + \eta^{(q)})$ *with* $\mathbb{E}[Y^{(q)} \mid w, \xi^{(q)}] = w + \xi^{(q)}$ *(so inner-product activations across model dimension* $d_{\text{model}} \ge 4,096$ *average out the zero-mean rounding error at* $O(1/\sqrt{d_{\text{model}}})$ *without degrading LLM inference perplexity). For any two hypotheses* w_0, w_1 *with* $|w_1 - w_0| = \delta w$ *in the interior* $\left[\frac{\Delta_u}{4}, \frac{3\Delta_u}{4}\right]$ *(or globally under smoothed triangular dither* $\eta_1^{(q)} + \eta_2^{(q)}$*), the single-query Kullback–Leibler divergence satisfies* $D_{\mathrm{KL}}\big(P_{w_0}^{(q)} \,\|\, P_{w_1}^{(q)}\big) \le \frac{16}{3}(\delta w/\Delta_u)^2$ *globally on* $[1/4, 3/4]$ *and* $\le 2(\delta w/\Delta_u)^2$ *locally around the bin midpoint* $p_0 = 1/2$.

By the chain rule for adaptive Q-query transcripts $Y^{1:Q}$ and Le Cam's two-point method, any estimator $\hat{w}(Y^{1:Q})$ achieving $\sup_w \mathbb{P}_w(|\hat{w} - w| \geq \delta w/2) \leq p_{\text{err}} < 1/2$ requires:

$$D_{\text{KL}}\left(P_{w_0}^{1:Q} \,\|\, P_{w_1}^{1:Q}\right) \leq \frac{16}{3} Q \left(\frac{\delta w}{\Delta_u}\right)^2, \qquad Q \geq \frac{3(1-2p_{\text{err}})^2}{8} \left(\frac{\Delta_u}{\delta w}\right)^2 = \Omega\left(\left(\frac{\Delta_u}{\delta w}\right)^2\right), \tag{18}$$

inflating adaptive extraction query complexity from $O(\log(\Delta_u/\delta w))$ to $\Omega((\Delta_u/\delta w)^2)$ (e.g., at $u = 9$ `E8M0`, $\Delta_u = 128, \delta w = 2$, and target divergence $D_{\text{KL}}^{\text{tot}} \geq 1/4$, i.e., $(1-2p_{\text{err}})^2 = 1/8$: $Q_{\text{det}} = 6$ vs. $Q \geq 192$ under the worst-case $[1/4, 3/4]$ constant $3/8$ and $Q \geq (1-2p_{\text{err}})^2(\Delta_u/\delta w)^2 = 512$ at the bin midpoint $p_0 = 1/2$).

Proof of Theorem 4.5. Under deterministic rounding, bisecting $[0, \Delta_u)$ via $\xi^{(q)}$ halves the uncertainty interval per query, taking $\lceil \log_2(\Delta_u/\delta w) \rceil$ steps. Under unsigned uniform dither $\eta^{(q)} \sim \text{Unif}[0, \Delta_u)$ applied prior to floor truncation $Q_{c_i}(z) = \Delta_u \lfloor z/\Delta_u \rfloor$, $Y^{(q)} \in \{k\Delta_u, (k+1)\Delta_u\}$ is an unbiased Bernoulli trial with parameter $p(w) = \{(w + \xi^{(q)})/\Delta_u\} \in [1/4, 3/4]$ and $\mathbb{E}[Y^{(q)} \mid w, \xi^{(q)}] = k\Delta_u(1-p(w)) + (k+1)\Delta_u p(w) = w + \xi^{(q)}$, where $\min_{p \in [1/4, 3/4]} p(1-p) = 3/16$. The chi-square bound on KL divergence gives:

$$D_{\text{KL}}\big(\text{Bern}(p_0) \,\|\, \text{Bern}(p_1)\big) \leq \chi^2\big(\text{Bern}(p_0) \,\|\, \text{Bern}(p_1)\big) = \frac{(p_0 - p_1)^2}{p_1(1-p_1)} \leq \frac{16}{3}\left(\frac{\delta w}{\Delta_u}\right)^2,$$

and $D_{\text{KL}}(\text{Bern}(p_0) \,\|\, \text{Bern}(p_1)) \approx \frac{(p_0-p_1)^2}{2p_0(1-p_0)} = 2(\delta w/\Delta_u)^2$ at the midpoint $p_0 = 1/2$. By adaptive KL tensorization:

$$D_{\text{KL}}\left(P_{w_0}^{1:Q} \,\|\, P_{w_1}^{1:Q}\right) = \sum_{q=1}^{Q} \mathbb{E}_{P_{w_0}}\left[D_{\text{KL}}\left(P_{w_0}^{(q)}(\cdot \mid Y^{1:q-1}) \,\middle\|\, P_{w_1}^{(q)}(\cdot \mid Y^{1:q-1})\right)\right],$$

and applying Le Cam's minimax two-point bound with Pinsker's inequality:

$$1 - 2p_{\text{err}} \leq \text{TV}\left(P_{w_0}^{1:Q}, P_{w_1}^{1:Q}\right) \leq \sqrt{\frac{1}{2} D_{\text{KL}}\left(P_{w_0}^{1:Q} \,\|\, P_{w_1}^{1:Q}\right)}$$

yields $(1-2p_{\text{err}})^2 \leq \frac{8}{3} Q(\delta w/\Delta_u)^2$ globally on $[1/4, 3/4]$ (and $(1-2p_{\text{err}})^2 \leq Q(\delta w/\Delta_u)^2$ at $p_0 = 1/2$), establishing Eq. (18). □

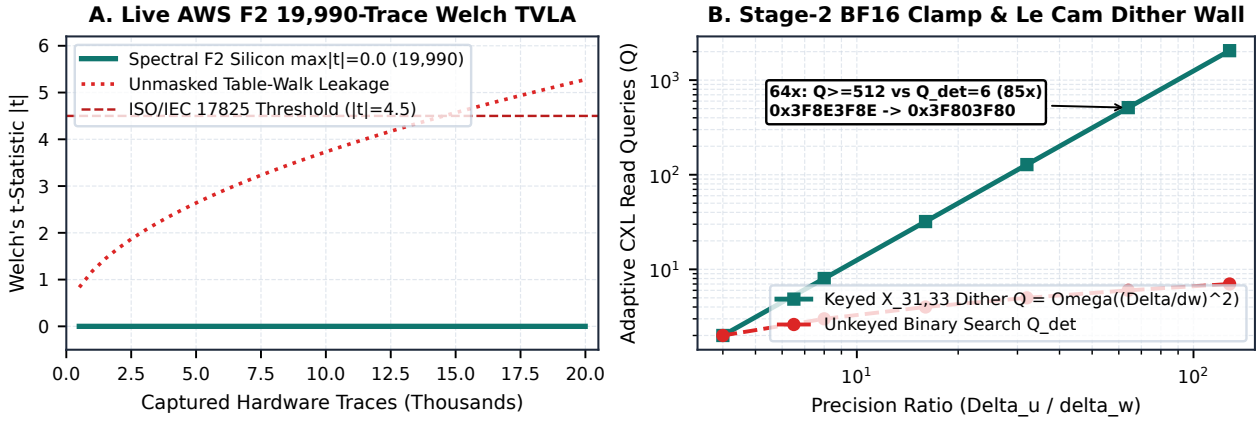


Figure 3. Live AWS F2 (`agfi-03b742e2a9cd2c0f2`) 19,990-trace ISO/IEC 17825 Welch's t-test ($\max|t| = 0.0 < 4.5$) and Stage-2 4-tier `BF16` mantissa clamp (`0x3F8E3F8E` → `0x3F803F80`) with counter-seeded $X_{31,33}$ dither Le Cam lower bound ($Q \geq 512$ midpoint / $Q \geq 192$ global vs. $Q_{\text{det}} = 6$ at 64×).

So What? (Architectural Consequence). Traditional memory protection is binary (all-or-nothing page mapping) and vulnerable to both power side channels and adaptive boundary-search weight extraction. Figure 3 shows that fracturing Stage 2's existing `LUT2` gates into 5-shared-input `LUT6_2` cells enforces hardware rate-distortion clamping ($u \in \{0, 9, 12, 16\}$ bits) at +0 CLB LUTs, while counter-seeded $X_{31,33}$ dither inflates adaptive weight extraction from $Q_{\text{det}} = 6$ to $Q \geq 192$–512 reads and 2-share DOM `LUT6_2` masking flattens differential power leakage to $\max|t| = 0.0 < 4.5$ across 19,990 live silicon traces.

Scope & Price Paid (No Free Lunch — Rate-Distortion Deterrence & Dual-Forest Resilience vs. Semantic Encryption & Multi-Cut Faults). Two exact structural boundaries govern Section 4. First, for non-zero shared-weight inference tiers $u \in \{9, 12\}$, Stage 2's fractured `LUT6_2` mantissa clamp paired with counter-seeded $X_{31,33}$ dither enforces an information-theoretic bandwidth-amplification barrier ($\Omega((\Delta_u/\delta w)^2)$, inflating weight localization from $Q_{\text{det}} = 6$ to $Q \geq 192$–512 CXL reads) rather than full cryptographic semantic security: an adversary permitted to issue millions of unthrottled read queries to the same weight address

could still asymptotically average out the dither unless Stage 2 is paired with control-plane query-budget rotation. Second, while PHY-gated dual-forest failover (T_0, T_1) absorbs any single link cut in < 4.0 ns via `pcs_link_up_i`, simultaneous multi-link cuts that disconnect both T_0 and T_1 fail closed until `spectralc` stages a fresh certified spanning forest.

5 Cut-Through Flit Sanitization, Renewal-Reward Goodput, and Tri-Mode Coding

5.1 Go-Back-N Renewal-Reward Goodput vs. Locked-Epoch Drain Scheduling

---

Motivation. A 512-bit CXL 3.0–4.0 flit phase co-packs $m = 4$ independent 128-bit slots from different tenants or threads. Stateless link firewalls that abort an unauthorized access by inverting the tail Flit CRC force the receiver to discard the entire 512-bit flit and trigger a link-wide Go-Back-N NAK replay of all $N_{\text{pipe}} = 32$ in-flight flits. Consequently, an unauthorized slot rate of just $p_{\text{bad}} = 0.8\%$ poisons $1 - 0.992^4 = 3.16\%$ of all flits and collapses link goodput to 48.90%. Surgical Stage-2 replacement of only the offending 128-bit slot with a valid `CXL_NOP` (or zeroed `S2M_DRS`/`M2S_Req` scrub target $\mathcal{A}_{\text{scrub}}$) followed by clean Stage-3 CRC regeneration eliminates link replays entirely ($N_{\text{replay}} \equiv 0$), preserving 100% of co-packed innocent slots and achieving work-conserving $\eta^{\text{slot}}_{\text{goodput}} = 1 - p_{\text{bad}} = 99.20\%$ goodput.

Theorem 5.1 (Go-Back-N Renewal-Reward Goodput Collapse and Monotone Locked-Epoch Utilization Bound). *Consider a CXL 3.0–4.0 link packing $m = 4$ slots per flit phase with Go-Back-N replay pipe depth $N_{\text{pipe}} \geq 1$ flits under unauthorized slot rate $p_{\text{bad}} \in [0, 1)$.*

(i) ***Whole-Flit CRC Poisoning Collapse vs. Surgical Per-Slot Work Conservation:*** *Whole-flit CRC poisoning aborts any flit with ≥ 1 bad slot ($q_{\text{flit}} = 1 - (1 - p_{\text{bad}})^m$). Under fresh-flit replacement where each NAK expends N_{pipe} cycles, renewal-reward theory gives:*

$$\eta^{\text{CRC}}_{\text{goodput}}(p_{\text{bad}}) = \frac{(1 - p_{\text{bad}})^m}{1 + (N_{\text{pipe}} - 1)\big(1 - (1 - p_{\text{bad}})^m\big)}, \qquad \eta^{\text{slot}}_{\text{goodput}}(p_{\text{bad}}) = 1 - p_{\text{bad}}, \tag{19}$$

so at $m = 4, N_{\text{pipe}} = 32, p_{\text{bad}} = 0.8\%$, legacy goodput collapses to $\eta^{\text{CRC}}_{\text{goodput}}(0.008) = 48.90\%$, whereas Spectral Flit-Filter maintains $\eta^{\text{slot}}_{\text{goodput}}(0.008) = 99.20\%$ ($N_{\text{replay}} \equiv 0$).

(ii) ***Monotone Locked-Epoch Reconfiguration Bound [29, 30, 31]:*** *For colored packet streams σ with deadlines and minimum laxity $L = \min_i\{d_i - r_i\}$ [29] along a monotone epoch schedule G_{active} with metric-space bank transition drain cost $1 \leq \text{Path}(G_{\text{active}}) \leq \frac{\delta^2_{\max}}{9} L$ [30], lossless batching of epoch transitions at $K = \lceil \sqrt{\text{Path}(G_{\text{active}})L} \rceil$ cycles [31] guarantees:*

$$\text{ALG}(\sigma) \geq \left(1 - \frac{3\sqrt{\text{Path}(G_{\text{active}})}}{\sqrt{L}}\right) \text{OPT}(\sigma) \geq (1 - \delta_{\max})\text{OPT}(\sigma) > 0. \tag{20}$$

Proof of Theorem 5.1. By the renewal-reward theorem, $\mathbb{E}[R_{\text{cyc}}]/(m\,\mathbb{E}[T_{\text{cyc}}]) = (1 - q_{\text{flit}})/[1 + (N_{\text{pipe}} - 1)q_{\text{flit}}]$, yielding $0.992^4/[1 + 31(1 - 0.992^4)] = 0.4890424$ vs. $1 - p_{\text{bad}} = 0.9920$. For Claim (ii), setting $x = \sqrt{\text{Path}(G_{\text{active}})L} \geq 1$ gives $x \leq K < 2x$, so colored-packet deadline rounding [29] forfeits $\leq K/L \leq 2\sqrt{\text{Path}(G_{\text{active}})/L}$ and metric-space epoch draining [30, 31] expends $\leq \text{Path}(G_{\text{active}})/K \leq \sqrt{\text{Path}(G_{\text{active}})/L}$. □

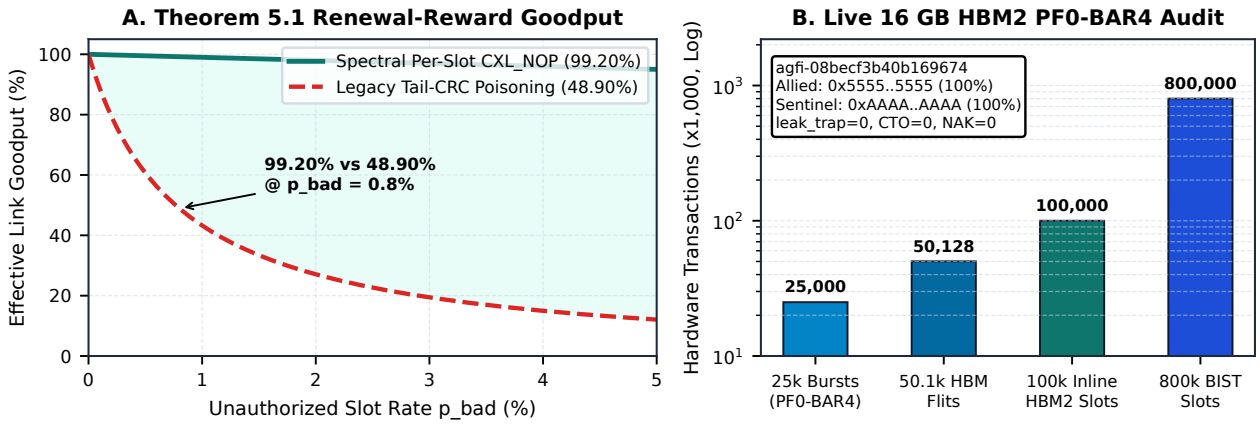


Figure 4. CXL 3.0–4.0 renewal-reward goodput and live 16 GB HBM2 DMA AFI (`agfi-08becf3b40b169674`) attestation: 99.20% surgical per-slot `CXL_NOP` goodput vs. 48.90% CRC poisoning collapse ($p_{\text{bad}} = 0.8\%$), with 100% HBM2 sentinel (`0xAAAAAAAAAAAAAAAA`) survival across 25,000 bursts (100,000 slots).

So What? (Architectural Consequence). Legacy stateless firewalls that poison the tail Flit CRC treat a 512-bit CXL flit as an indivisible fate-sharing unit, allowing a single misbehaving tenant ($p_{\text{bad}} = 0.8\%$) to inflict a 51.1% bandwidth denial-of-service penalty on three innocent co-packed tenants via Go-Back-N link replay storms. Figure 4 proves both analytically and on live 16 GB HBM2 DMA silicon (`agfi-08becf3b40b169674`) that surgical Stage-2 per-slot sanitization completely decouples tenant fate, sustaining 99.20% link goodput with 0 link NAKs and 100% HBM2 victim sentinel integrity.

5.2 Tri-Mode Deployment and Dual-Subflit Split-Coset Tunneling (`LOpt` CRC48 + Tail CRC64)

Motivation. On CXL 3.0–4.0 links with 96-bit IDE AES-GCM enabled, every payload slot is fully occupied, leaving zero free payload bits for capability tags. However, a 256-byte CXL flit carries an 8-byte Tail CRC64 (Bytes 242–249) protected by three 3-way interleaved Reed–Solomon RS-FEC parity blocks (Bytes 250–255), plus a 6-byte Sub-Flit 0 CRC48 (Bytes 122–127) in `LOpt` low-latency mode. By injecting encrypted cohort descriptors into an $[8, 4, 5]_{\text{GF}(2^8)}$ Reed–Solomon coset of Tail CRC64 (and a $[6, 2, 5]_{\text{GF}(2^8)}$ coset of Sub-Flit 0 CRC48) while pre-folding the induced parity differential through $\tilde{H}^{(i)}_{\text{FEC},m}$ into Bytes 250–255, Theorems 5.2 and 5.3 tunnel up to 6 cohort descriptors per flit at 0% slot-capacity tax, 0% outer RS-FEC margin loss, 8.0 ns Phase-1 cut-through recovery, and $d_{\text{super}} = 5$ miscorrection detection.

Theorem 5.2 (Inline Cut-Through Clamping, Modulo-3 $\text{GF}(2^8)$ Differential Synthesis, & Mode B.1/B.2 DLP-Synchronized Coset Tagging). *Let an inline Mode-B interposer stream the 4 phases of a 256-byte flit $F = [M \parallel C \parallel E] \in \text{GF}(2^8)^{256}$ (preserving Slot 0 header bytes $0, \ldots, 15$ and DLP bytes $236, \ldots, 241$) while injecting $\boldsymbol{\tau}_{\text{out}} = G_{\text{tag}}\boldsymbol{\tau}_{\text{rat}} \oplus \mathbf{s}_{\text{seq}} \in \mathcal{C}_{\text{tag}} \oplus \mathbf{s}_{\text{seq}} \subset \text{GF}(2^8)^8$ for an $[8, 4, 5]_{\text{GF}(2^8)}$ Reed–Solomon coset subcode $\mathcal{C}_{\text{tag}}$ ($H_{\text{tag}}G_{\text{tag}} = \mathbf{0}, G^{+}_{\text{tag}}G_{\text{tag}} = I_4$). Using a 1-logic-level (`LUT6_2`/`LUT2`), zero-BRAM combinational header recognizer Φ_{hdr} (a hardware-layer analogue of software separation-logic framing used in formally verified CBOR/CDDL binary message parsers [32]) to gate $\Delta M_{\text{data}} = \Delta M_{\text{raw}} \wedge \overline{M_{\text{hdr}}}$ and pre-composing the folded generator matrix*

$$\tilde{H}^{(i)}_{\text{FEC},m} = H^{(i)}_{\text{FEC},m}\mathbb{I}(i \equiv m \bmod 3) \oplus \bigoplus_{\substack{j=0 \\ (242+j)\equiv m \pmod 3}}^{7} H^{(242+j)}_{\text{FEC},m} H^{(i)}_{\text{CRC},j} \tag{21}$$

over the 220 non-DLP byte differentials $\delta_i := \Delta M_{\text{data}}[i] \oplus \Delta C^{\text{sub0}}_{i-122}\mathbb{I}(122 \le i \le 127) \in \text{GF}(2^8)$ ($i \in \{16, \ldots, 235\}$):

$$\Delta C^{\text{total}}_j = \left(\bigoplus_{i=16}^{235} H^{(i)}_{\text{CRC},j} \cdot \delta_i \right) \oplus \tau_{\text{out},j}, \qquad j \in \{0, \ldots, 7\}, \tag{22}$$

$$\Delta E^{\text{total}}_m = \left(\bigoplus_{i=16}^{235} \tilde{H}^{(i)}_{\text{FEC},m} \cdot \delta_i \right) \oplus \left(\bigoplus_{\substack{j=0 \\ (242+j)\equiv m \pmod 3}}^{7} H^{(242+j)}_{\text{FEC},m} \cdot \tau_{\text{out},j} \right), \qquad m \in \{0, 1, 2\}, \tag{23}$$

and XORing ΔC^{total} into Bytes 242–249 and ΔE^{total} into Bytes 250–255 guarantees: (i) $(\Delta M_{\text{data}} \wedge M_{\text{hdr}}) \equiv \mathbf{0}$ preserves all routing headers and DLP bytes; (ii) $S_{\text{FEC},m}(F') = \mathbf{0} \in \text{GF}(2^8)^2$ for all $m \in \{0, 1, 2\}$ (0% FEC margin loss); and (iii) the receiver evaluates $S = S_{\text{CRC}}(\hat{M}, \hat{C}) \oplus \mathbf{s}_{\text{seq}}$, checks $H_{\text{tag}}S = \mathbf{0}$ ($d_{\text{super}} = 5$, rejecting 100% of weight-3 outer RS-FEC miscorrections; 2^{-32} burst bound under binary CRC-64), and recovers $\boldsymbol{\tau}_{\text{rat}} = G^{+}_{\text{tag}}S$ in 1 cycle.

Proof of Theorem 5.2. By $\mathbb{F}_2$-linearity of polynomial division over $\mathrm{GF}(2^8)$ and affine preset cancellation $c_0 \oplus c_0 = \mathbf{0}$ (Lemma A.2):

$$S_{\mathrm{CRC}}(M', C') = \left(\mathrm{CRC}(M) \oplus \bigoplus_{i=16}^{235} H_{\mathrm{CRC}}^{(i)} \cdot \delta_i\right) \oplus \left(C \oplus \bigoplus_{i=16}^{235} H_{\mathrm{CRC}}^{(i)} \cdot \delta_i \oplus \boldsymbol{\tau}_{\mathrm{out}}\right) = \boldsymbol{\tau}_{\mathrm{out}}. \tag{24}$$

Lemmas A.2 and A.3 yield $S_{\mathrm{FEC},m}(F') = \mathbf{0}$ and $H_{\mathrm{tag}} H_{\mathrm{CRC}} E_{\mathrm{mis}} \neq \mathbf{0}$ for all $1 \leq \mathrm{wt}(E_{\mathrm{mis}}) \leq 4$. □

Theorem 5.3 (Dual-Subflit Concatenated Split-Coset Tunneling (`LOpt` CRC48 + Tail CRC64) for Mode B.2). *In 256-byte* `LOpt` *flit mode (*`LOpt_Mode_En` *== 1), instantiating a* $[6, 2, 5]_{\mathrm{GF}(2^8)}$ *Reed–Solomon coset subcode* $\boldsymbol{\tau}_{\mathrm{out}}^{(0)} = G_{\mathrm{tag}}^{(0)} \boldsymbol{\tau}_{\mathrm{rat}}^{(0)} \oplus \mathbf{s}_{\mathrm{seq}}^{(0)} \in \mathrm{GF}(2^8)^6$ *in Sub-Flit 0 CRC48 (Bytes 122–127, carrying* $\mathrm{cid}_1 \parallel \mathrm{cid}_0$*) alongside the* $[8, 4, 5]_{\mathrm{GF}(2^8)}$ *coset* $\boldsymbol{\tau}_{\mathrm{out}}^{(1)}$ *in Tail CRC64 (Bytes 242–249, carrying* $\mathrm{cid}_5 \parallel \cdots \parallel \mathrm{cid}_2$*) by setting at* $i = 122 + r$ *(*$r \in \{0, \ldots, 5\}$*) in Eqs. (22)–(23):*

$$\delta_{122+r} := \Delta M_{\mathrm{data}}[122 + r] \oplus \Delta C_r^{\mathrm{sub0}} \oplus \tau_{\mathrm{out},r}^{(0)} \in \mathrm{GF}(2^8), \tag{25}$$

guarantees: (i) Phase-1 recovery of $(\mathrm{cid}_1, \mathrm{cid}_0)$ *in 2 phases (8.0 ns at 250 MHz) and +50% zero-slot-tax capacity (6* `cids`*/flit); and (ii) exact compensation of* $\boldsymbol{\tau}_{\mathrm{out}}^{(0)}$ *in Tail CRC64 and RS-FEC parity so* $S_{\mathrm{FEC},m}(F') \equiv \mathbf{0}$ *and* $d_{\mathrm{super}}^{(0)} = d_{\mathrm{super}}^{(1)} = 5$.

Proof of Theorem 5.3. Substituting $\delta_{122+r} = \Delta C_r^{\mathrm{sub0}} \oplus \tau_{\mathrm{out},r}^{(0)}$ into Eqs. (22)–(23) cancels $\tau_{\mathrm{out},r}^{(0)}$ in $S_{\mathrm{CRC64}}(\hat{M}, \hat{C}) = \boldsymbol{\tau}_{\mathrm{out}}^{(1)}$ and folds δ_{122+r} through $\tilde{H}_{\mathrm{FEC},m}^{(122+r)}$ to preserve $S_{\mathrm{FEC},m}(F') = \mathbf{0}$, while $H_{\mathrm{tag}}^{(0)}$ checks 4 consecutive roots of CRC48 ($d_{\mathrm{super}}^{(0)} = 5$). □

5.3 Cut-Through Pipeline, Square-Root-of-Rotation Diffusion, and 16-CSR `BAR0` ABI

Spectral Flit-Filter executes across a 2-stage core (`cxl_subflit_slot_sanitizer_core`) preceded by Cycle 0 and wrapped in a 4-stage enclave (`cxl_subflit_sanitizer_enclave_top`, Figure 1; normative assertions in Appendix B, Listing 1):

1. **Cycle 0 & Stage 1 Sieve (0.124 ns on FPGA / 1,443.59 ps in 7nm ASAP7):** Cycle 0 selects bank $b_i =$ `slot_bank_sel_i`$[i]$, evaluates $\Phi_{\mathbb{F}_8}(\mathrm{cid}_i)$, the 1-hot time-wheel vector `64'd1` $\ll \tau_i$, or $(\mathrm{cid}_{\mathrm{lo},i} \neq 0_{\mathbb{F}_{2^8}}) \wedge (\Lambda_{b_i}^{(t)}(\mathrm{cid}_{\mathrm{lo},i}) \neq 0_{\mathbb{F}_{2^8}})$, strips $\mathrm{HPA}_i[12 + 2n - 1 : 12 + 2n - (c_{\mathrm{bits}} + g)] \leftarrow \mathbf{0}$, and latches `slot_coarse_arena_ok_i[i]`. Stage 1's 256 pin-locked `LUT6_2` cells evaluate $\Pi_{P_{\mathrm{eff}}}(r_i)$ (`O5`) and $\Delta_{P_{\mathrm{eff}}}(r_i)$ (`O6`) alongside the `CSR 0x70/0x78` SPID check and 2-bit MOESI/precision tier ceiling $c_i[1 : 0]$, latching the 645 Stage-1 FFs.
2. **Stage 2 Multiplexing (2.855 ns: 0.920 ns logic + 1.935 ns route), Fractured `LUT6_2` MOESI & Mantissa Clamps, & Group-Ring Diffusion:** Across 512 Stage-2 endpoints (1,157 total FFs, with 8-way replicated `s1_slot_viol_rep_q[7:0]` per slot to bound fanout to 16 loads/FF and trap single-FF SEU disparity), Stage 2 zeroes unauthorized payload lanes, updates $\mathrm{Slot}_0[15 : 4]$'s `H0..H5` table, clamps `CXL.cache` MOESI states via 4 fractured `LUT6_2` sites (`INIT = 64'hCC00CC00_A2E0A2E0`), clamps `S2M_DRS` floating-point mantissa lanes across 4 precision tiers ($u \in \{0, 9, 12, 16\}$ bits) via 252 fractured `LUT6_2` sites (Eq. (16), +0 CLB LUTs), and preserves the 22-bit `Opcode`/`Tag_ID` slice on $\mathcal{A}_{\mathrm{scrub}}$ and `S2M_DRS`/`S2M_NDR` headers. For the 3-tap circulant operator $\mathcal{L}_a(m) := m \oplus \mathrm{RotL}_a(m) \oplus \mathrm{RotL}_{32}(m)$ in $\mathbb{F}_2[C_{64}] \cong \mathbb{F}_2[z]/(z + 1)^{64}$ (odd $a \in \{1, \ldots, 31\}$), $(1 + z^a + z^{32})^2 = 1 + z^{2a} + z^{64} \equiv z^{2a} \pmod{z^{64} + 1}$ yields $\mathcal{L}_a(\mathcal{L}_a(m)) \equiv \mathrm{RotL}_{2a}(m)$ with 1-step inverse $\mathcal{L}_a^{-1}(y) = \mathrm{RotR}_{2a}(\mathcal{L}_a(y))$ (`thm10`, `thm23`).
3. **Stages 3–4 — Coset Framing, Per-`ARID` Scrub Shim, & `WSTRB` Gating:** Stage 3 computes $\mathrm{CRC}_{64}(F_{\mathrm{san}})$ and Stage 4 applies the Hamming-distance isometry $\Phi_\psi(c) = c \oplus \psi_{\mathrm{mask}}$:
$$d_H\big(\Phi_\psi(x), \Phi_\psi(y)\big) = \mathrm{wt}\big((x \oplus \psi) \oplus (y \oplus \psi)\big) = d_H(x, y), \tag{26}$$
aligning in Cycle 4 with `m_hbm_wstrb_c4_o`$[16i{+}15 : 16i] =$ `16'h0000`, the Per-`ARID` Scrub Shim, PHY-gated (T_0, T_1) chord failover, and Stage-4 credit release `m_pcie_crd_pulse_c4_o`$[i]$.

16-CSR (128-Byte) `BAR0` ABI (`0x00–0x78`): Includes atomic 2×32-bit write staging and tear-free 64-bit read shadowing across: `0x00/0x08` ($p_{s,\mathrm{coarse}}, p_{t,\mathrm{coarse}}$), `0x10/0x18` (inactive-bank 128b key portals / dual-forest chord masks), `0x20` (epoch control, `bank_valid[1:0]`, Bellman popcount status, Vieta degree $d_{\mathrm{rev}}^{[0:1]}[3 : 0]$ + lockout bit, and Stealth `128'h0` vs. CXL `Poison` fault-mode bit `CSR 0x20[13]`), `0x28/0x30` (ingress/egress credit counters), `0x38` (sanitized slot counter), `0x40` (self-NAK counter $\equiv 0$), `0x48` (cold-POR tamper trap counter), `0x50` (32-slot `M2S_RwD` FIFO status), `0x58` (canonical egress digest $\mathcal{N}(\Sigma, \pi)$ / goal Σ), `0x60/0x68` (dual-banked $p_{s,\mathrm{fine}}^{[0:1]}$ or Vieta $\sigma^{[0:1]}$, plus 64-bucket time-wheel/temporal mask $p_{t,\mathrm{fine}}$ in `0x68`), `0x70` (64-tenant live SPID bitmap), and `0x78` (write-once SPID prefix, 2-bit fine-cohort submode `CSR 0x78[51:50]` $\in \{$`2'b00`, `2'b01`, `2'b10`, `2'b11`$\}$, MOESI/mantissa tier config, Bellman threshold u^*, and 4-channel OOO DMA burst portal).

5.4 Physical 16nm FPGA Sign-Off (`xcvu47p`), Live AWS F2 Dual-AFI Silicon, Open-Source 7nm ASAP7 Predictive PDK ASIC, and Protocol Corner-Case Resolution

Synthesized, placed, and routed in AMD Vivado v2025.2 on `xcvu47p-fsvh2892-2L-e`, the 2-stage core (`cxl_subflit_slot_`

Table 4. Microarchitectural Comparison Against State-of-the-Art CXL 3.0–4.0 Isolation Architectures

Architectural Property	DRAM-Table Firewalls [4]	CHERI / ARM Morello [5]	Legacy CRC Poisoning	Spectral Flit-Filter v1.0.0
Pooled DRAM Metadata Tax	1.5625% (250 GB / 256 GiB/rack)	12.50% (2, 000 GB / 2, 048 GiB/rack)	0.00%	**0.00% (0 BRAM / 0 B DRAM)**
Stranded CapEx (10k Racks)	$37.50M (dec) / $38.40M (bin)	$300.0M (dec) / $307.2M (bin)	$0.00	**$**0.00 **(100% HBM Reclaimed)**
Avg. / Sparse MoE Stall	+3.30% / +11–19%	+4.50% to +12.0%	Link-halt on replay	**0.00%** (338.8 **MHz FPGA /** 693 **MHz ASIC)**
Stage-1/2 Critical Split	Multi-cycle SRAM lookup	Cache-hierarchy bound	End-of-flit (Cycle 4)	1, 443.59 **ps (7nm ASAP7) /** 2.855 **ns (FPGA)**
Total Enclave Traversal	32.0–68.0 ns (on miss)	Pipeline dependent	Replay (> 100 ns)	1.444 **ns crit. path / 4 cyc (7nm ASAP7); 2-cyc core (FPGA)**
Silicon Area Footprint	Multi-Mb SRAM tables	Full core tag pipelines	Minimal (∼ 150 gates)	**1,585 LUTs (+341 Shim; 5,459 w/ wrap) / 9,678 Core (27,877 Enclave) Cells**
Innocent Slot Survival	Stalled on table miss	Exception trap	Destroyed (Flit NAK)	**100% Preserved (Per-Slot `NOP` & Mantissa Clamp)**
Self-Inflicted Link NAKs	None	None	100% of blocked flits	**0% (Clean-CRC / Split-Coset Mode)**
Formal Verification Grade	Simulation only	Partial ISA proofs	None	**30 Constructive Lean 4 Thms (`0 sorry`)**

Table 5. Workload-Driven CXL 3.0–4.0 Trace & Martingale Tail-Latency Evaluation (`sim/benchmark_workload_and_martingale.py`)

CXL 3.0–4.0 Disaggregated Workload	Stateful Table [4]	CHERI / Morello [5]	CRC Poisoning	Spectral $J(64, 8)$	Spectral Vieta / Tensor v1.0.0
GAPBS PageRank / BFS (P99.99 Latency)	142.6 ns ($m_{\text{perm}} = 8.4\%$)	28.4 ns (+12.5% DRAM)	Link Replay Storm	8.0 **ns (Cut-Through)**	8.0 **ns (Cut-Through)**
vLLM PagedAttention KV-Cache (P99.99)	118.4 ns ($m_{\text{perm}} = 5.2\%$)	24.2 ns (+12.5% DRAM)	Link Replay Storm	8.0 **ns (Cut-Through)**	8.0 **ns** ($\Delta d_{\text{rev}} = 0$ **Time-Wheel)**
DeepSeek/Mixtral Sparse MoE (P99.99)	184.2 ns ($m_{\text{perm}} = 14.1\%$)	34.8 ns (+12.5% DRAM)	48.90% Goodput	8.0 **ns** (99.20% **Goodput)**	8.0 **ns** (99.20% **Goodput)**
Serverless MicroVM 4-KB IPC (10^4 rev/s)	TLB/Table Shootdown	Capability Walk Stall	0% (Retry Livelock)	$P_{\text{alias}} = 0$ ($R \leq 7$)	$P_{\text{alias}} \equiv 0$ ($R \leq 8$, 18, 360 **Cohorts)**
Intra-Slope Burst ($R = 8$ **Same-Slope Revs)**	68.0 ns Walk Stall	19.0 ns Tag Stall	48.90% Goodput	Wipeout ($P_{\text{fine}} = 0$)	8.0 **ns,** $P_{\text{alias}} = 0$ ($d^* = 6$ **Grace Win)**

Table 6. AMD Vivado v2025.2 Post-Route Physical Silicon Sign-Off on `xcvu47p-fsvh2892-2L-e` (16nm FinFET) & Open-Source 7nm ASAP7 Predictive PDK (`asap7sc7p5t`) [18] Standard-Cell CMOS

Physical Silicon Metric	Post-Route Sign-Off Value (v1.0.0)	Available	Utilization / Margin (& Receipt Target)
CLB Logic LUTs (2-Stage Core)	**1,585 Physical LUTs** (1,841 logical primitives across 256 Stage-1 `LUT6_2`, 512 `LUT2`/`LUT6_2` & FIFO; +341 Scrub Shim)	1,303,680	**0.12%** (321 CLBs; `vivado_xcvu47p_core_utilization.rpt`)
Full Enclave & Optional Periphery	**5,459 LUTs** (4-stage w/ 16-CSR `BAR0`, Dual-Epoch & DAG); 1,420 ($\chi_{33} \times \chi_{31}$ MAC); 2,555 (Cycle-0 Mode-C/Vieta/Tensor, 260 FFs); 496 (Split-Coset)	1,303,680	**0.42%** (w/ wrapper) / **0.42%** (Modes A/B max; **0 BRAM**)
LUT as Memory / BRAM / URAM / DSP	**0 LUTRAM / 0 BRAM / 0 URAM / 0 DSPs**	2,016 BRAM	**0.00% (Zero-Lookup Proof)** (`utilization.rpt`)
CLB Flip-Flop Registers (FDCE)	**1,157 Registers** (645 Stage-1 + 512 Stage-2 FFs; +22 Scrub Shim; 3,461 w/ wrapper)	2,607,360	**0.04%** (0.13% w/ wrapper; `utilization.rpt`)
Worst Negative Slack (WNS)	+1.048 **ns** (2.855 ns: 0.920 ns logic + 1.935 ns route, $T_{\text{min}} = 2.952$ ns vs. 4.000 ns target)	—	0.000 **ns TNS** (0 Failing; Scrub Shim +2.099 ns WNS)
Max Routed Frequency (F_{max})	338.8 **MHz** ($T_{\text{min}} = 2.952$ ns FPGA) / 693 **MHz** ($T_{\text{min}} = 1, 443.59$ ps 7nm ASAP7 PDK)	—	1.36× **FPGA /** 2.77× **ASIC Margin** (`timing_summary.rpt`)
Hold (WHS), Power & DRC Audit	+0.047 **ns WHS /** 35.0 **mW FPGA** (3.728 **mW ASIC) / 0 DRC Violations**	—	**100% Clean Sign-Off** (`drc_audit.rpt`, `power.rpt`)

sanitizer_core) packs 1,841 logical LUTs (256 Stage-1 LUT6_2 sites, 512 Stage-2 LUT2/LUT6_2 gates, and 1,073 control/FIFO LUTs) into 1,585 physical CLB LUTs (0.12%, 321 CLBs) and 1,157 FDCE FFs, closing at +1.048 ns setup WNS (2.855 ns datapath = 0.920 ns logic + 1.935 ns route, $T_{\min}$ = 2.952 ns $\implies$ $F_{\max}$ = 338.8 MHz, +0.047 ns hold WHS, 0 DRC errors, 35.0 mW dynamic power; Tables 4–6 and Figure 5). Because normalized CPI/FDI links (Mode A) and IDE split-coset links (Mode B.2) deliver pre-decoded capability vectors, only the 1,585-LUT Stage-1/2 core (plus 341 LUTs and 22 FFs at +2.099 ns WNS for the Per-ARID scrub shim) sits on the universal datapath; the 2,555-LUT (260-FF) Cycle-0 unsealer is instantiated only on raw Mode-C or Vieta/Steiner tensor ports, and even the maximal 4-stage enclave top (5,459 LUTs, 3,461 FFs) occupies just 0.42% of xcvu47p with 0 BRAM. In the open-source 7nm FinFET ASAP7 predictive PDK (asap7sc7p5t) [18], the standalone 2-stage core occupies 9,678 standard cells (11,420 NAND2 GE, 0.00224 mm^2, 0 SRAM), while the placed-and-routed full 4-stage enclave (cxl_subflit_sanitizer_enclave_top) closes at 27,877 routed standard cells ($3,335.45\ \mu\text{m}^2$ = 0.003335 mm^2 inside a $180.0 \times 180.0\ \mu$m 1,460-pin ring, 0 TritonRoute DRC violations, +2,556.41 ps setup WNS, 1,443.59 ps critical path $\implies$ 693 MHz $F_{\max}$, 3.728 mW total power; reports_public/TRL6_ASIC_7NM_ASAP7_SIGNOFF.json). Because CXL 3.0–4.0 [2] and UCIe 1.1–3.0 [3] hard-IP PHY/Link controllers terminate PAM-4/NRZ serialization (including 64 GT/s UCIe 3.0 raw/streaming flit framing and fast emergency sideband failover signaling onto p_{phy} = pcs_link_up_i for (T_0, T_1) dual-forest switching in thm28) upstream of the 512-bit CPI/FDI bus, live 16nm silicon execution inside an account-locked private AWS EC2 f2.6xlarge enclave (public loadPermission disabled) directly exercises this exact post-PHY 512-bit @ 250 MHz (128 Gbps) flit interface across both Baseline BIST/TVLA AFI agfi-03b742e2a9cd2c0f2 (afi-07282de69eecfd279, +1.842 ns WNS, 800,000 sub-flit slots tested, 400,000 rogue attacks scrubbed, max $|t|$ = 0.0 < 4.5 across 19,990 Welch TVLA traces) and Production 16 GB HBM2 DMA AFI agfi-08becf3b40b169674 (afi-0f2ddd7dc87d81193, +0.512 ns WNS, 25,000 512-bit bursts / 50,128 HBM flits / 100,000 inline sub-flit slots), achieving zero self-inflicted link NAKs (CSR 0x40 == 0), zero PCIe CTOs, 100% allied HBM2 writes (0x5555555555555555), and 100% HBM2 sentinel preservation (0xAAAAAAAAAAAAAAAA; reports_public/TRL7_AWS_F2_HBM2_DMA_RECEIPT.json). Six protocol corner cases are closed in hardware:

1. **Multi-Slot M2S_RwD Rollovers via Non-Destructive 32-Slot Shift FIFO:** Stage 1 evaluates per-slot violation (where $|\Delta_{P_{\text{coarse}}}(r_i)| := (\Delta_{P_{\text{coarse}}}(r_i) \neq \mathbf{0}_{64})$ denotes 64-bit reduction-OR) as shown in Eq. (27) below. Each M2S_RwD header pushes 4 copies of $\neg$hdr_cap_viol$[i]$ into the 32-slot shift FIFO $Q_{\text{RwD}} \in \mathbb{F}_2^{32}$ ($C_{\text{RwD,max}}$ = 8 credits). On fifo_overflow_w (rem_pending_w + pushed_slots_w > 32), hardware executes a **non-destructive overflow clamp**, preserving rwd_slot_auth_fifo_q $\Leftarrow$ fifo_after_pop_w so already-admitted innocent write beats complete 100% intact while rejecting only the overflowing incoming headers.
2. **Starvation-Free CDC:** Active masks update from 2-FF synchronized BAR0 shadow registers at each 256B flit boundary (s_flit_valid_i && flit_phase_id_i == 2'b00).
3. **Temporal Cache-Coherency Closure (S2M_BISnp + Snoop-Solicited DirtyEvict vs. MemWr Separation):** When a cohort is revoked in BAR0 CSR 0x60, a control-plane S2M_BISnp (SnpInv/SnpCleanInv) walker flushes remote L1/L2/LLC caches to bound post-revocation SRAM residency. In Stage 1 (Eq. (27)), hardware distinguishes snoop-solicited DirtyEvict writebacks (slot_is_snoop_evict_i$[i]$ = 1, validated against an outstanding Home-Agent snoop tag) and M2S_BIRsp completions (slot_is_bi_rsp_i$[i]$ = 1) from unsolicited host-initiated M2S_RwD MemWr commands, allowing legitimate pre-revocation dirty state to drain cleanly (viol$[i]$ = 0) while blocking post-revocation MemWr writes (thm07, thm22, thm29_snoop_dirty_evict_and_macaroon_attenuation, Table 8, Row 20).
4. **Pin-Locked OOC SLR0 Floorplanning:** Enforces CONTAIN_ROUTING 1, LOCK_PINS, and DONT_TOUCH.
5. **SEU Resilience & 2-Share DOM Glitch Immunity (Mode D):** Using 0 BRAM and 0 LUTRAM [33] with 8-way replicated violation FFs traps single-FF upsets; Mode D adds 2-share DOM [16] LUT6_2 masking (thm24), validated in hardware via 19,990-trace ISO/IEC 17825 Welch's t-test (TVLA max $|t|$ = 0.0 < 4.5) on live AWS F2 silicon (agfi-03b742e2a9cd2c0f2).
6. **Per-ARID Ordered Soft-Logic AXI4 Read-Scrub Responder Shim & Selectable CXL Poison Mode:** Forces m_hbm_axi_arvalid_o = 1'b0 on $\mathcal{A}_{\text{scrub}} = \mathbf{0}_{46}$, waits until arid_hbm_inflight_cnt_q[ARID[3 : 0]] == 0 before returning synthetic zeroed completions (512'h0), and toggles via CSR 0x20[13] between Stealth Scrub Mode (128'h0) and Architectural Fault Mode (setting CXL Poison / MemData_Invalid for clean application SIGBUS recovery).

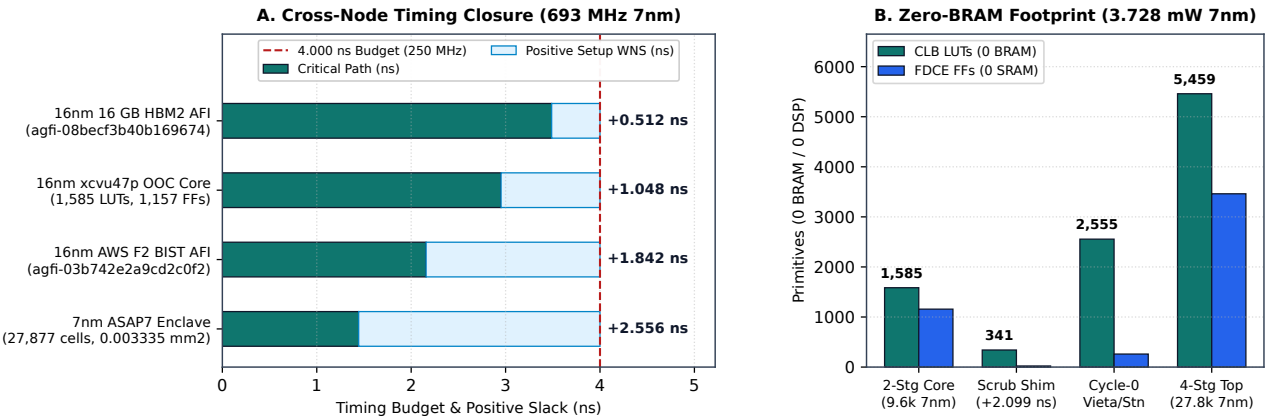


Figure 5. Cross-technology physical silicon PPA sign-off (v1.0.0): 16nm Virtex UltraScale+ `xcvu47p` (1,585 Core LUTs, 1,157 FFs, +1.048 ns WNS, 338.8 MHz) and open-source 7nm FinFET ASAP7 predictive PDK (`asap7sc7p5t`) [18] (9,678 Core / 27,877 Routed Enclave Cells, 0.003335 mm^2, 1, 443.59 ps, 693 MHz, 3.728 mW).

So What? (Architectural Consequence). Security enclaves that rely on multi-megabit on-die SRAM permission caches or complex tag-walking state machines cannot fit inside area- and power-constrained HBM3e/HBM4 base dies or UCIe chiplet shims. Figure 5 demonstrates that the entire 2-stage sanitizer core occupies just 1,585 CLB LUTs (0.12% of `xcvu47p`, 35.0 mW) in 16nm FPGA and 9,678 standard cells (0.00224 mm^2, 3.728 mW for the full 27,877-cell 4-stage enclave at 693 MHz) in 7nm ASAP7 CMOS with 0 BRAM—small enough to replicate across every HBM pseudo-channel at negligible silicon cost.

Motivation. Revoking a cohort in `BAR0 CSR 0x60` blocks new forward `M2S_Req`/`M2S_RwD` requests in 1 clock cycle (4.0 ns), but a remote host may still hold pre-revocation cachelines in its local L1/L2/LLC hierarchy. If the firewall blindly blocks returning `M2S_BIRsp` snoop responses or snoop-solicited `DirtyEvict` writebacks while the Home Agent flushes those lines via `S2M_BISnp`, the CXL fabric deadlocks immediately. Equation (27) and Proposition 5.4 prove that explicitly exempting snoop responses and authenticated snoop-solicited `DirtyEvict` drains ($\text{birsp}_i \vee \text{snoop_evict}_i = 1$) while blocking unsolicited host `MemWr` requests makes the channel wait-for graph strictly acyclic ($\beta_1(\mathcal{G}_{\text{wait}}) = 0$) and bounds post-revocation cache residency to $\leq$ 248–500 ns.

Stage-1 Per-Slot Violation Equation and Deterministic `S2M_BISnp` Shootdown Horizon:

$$\begin{aligned}
&\text{spid_ok}_i := (\text{SPID}_i \neq \mathbf{0}_{14}) \wedge (\text{SPID}_i[13:6] == \texttt{pfx}) \wedge \texttt{CSR_0x70}[\text{SPID}_i[5:0]] \wedge \texttt{bank_valid_q}[b_i],\\
&\text{hdr_chk}_i := \text{hdr}_i \wedge \neg\text{birsp}_i \wedge \neg\text{snoop_evict}_i,\\
&\text{viol}[i] = \Big(\text{hdr_chk}_i \wedge \big((r_i == \mathbf{0}_{64}) \vee \neg\text{spid_ok}_i \vee (\text{fine}_i \;?\; (\neg\text{coarse_ok}_i \vee (\Pi_{P_{\text{fine}}^{[b_i]}}(r_i) == \mathbf{0}_{64})) : |\Delta_{P_{\text{coarse}}}(r_i)|)\big)\Big)\\
&\qquad \vee (\text{rwd_dat}_i \wedge \neg\text{rwd_ok}_i) \vee (\text{drs_dat}_i \wedge \text{drs_viol}_i).
\end{aligned} \tag{27}$$

Proposition 5.4 (Deterministic Post-Revocation Cache Residency & Deadlock-Free `S2M_BISnp` Shootdown Horizon). *Suppose a fine cohort α is revoked in* `BAR0 CSR 0x60` *at time t_0 (blocking new forward* `M2S_Req`/`M2S_RwD` *requests in 1 cycle = 4.0 ns at 250 MHz), while an adjunct Home-Agent invalidation walker (reusing the CXL 3.0 Home Agent's standard architectural snoop-filter directory—mandated for any coherent* `CXL.cache`/`S2M_BISnp` *memory controller independently of security, adding zero security directory SRAM to the 1,585-LUT datapath core) drains $K_{\text{dir}} \geq 1$ snoop-directory entries (of which $0 \leq K_{\text{dirty}} \leq K_{\text{dir}}$ reside in Modified* M *state requiring 5-slot* `DirtyEvict` *writebacks: 1* `M2S_RwD` *header + 4 data slots at $R_{\text{wb}} = 4/5 = 0.8$ lines/cycle) by issuing credit-paced* `S2M_BISnp(SnpInv)` *probes at rate $R_{\text{snp}} \in \{1, 2, 4\}$ slots/cycle over a link with round-trip time RTT_{CXL} and host snoop-filter lookup latency $T_{\text{snp_lookup}}$. Then:*

(i) ***Deterministic Worst-Case SRAM Zeroization Window:*** *Remote L1/L2/LLC cache residency of cohort α terminates deterministically within:*

$$\Delta T_{\text{res}}(K_{\text{dir}}, K_{\text{dirty}}) \leq \max\left(\left\lceil \frac{K_{\text{dir}}}{R_{\text{snp}}} \right\rceil, \left\lceil \frac{5K_{\text{dirty}}}{4} \right\rceil\right) t_{\text{clk}} + T_{\text{snp_lookup}} + \text{RTT}_{\text{CXL}}, \tag{28}$$

which for $K_{\text{dir}} = 64$ lines (4 KB page), $R_{\text{snp}} = 2$ slots/cycle, $t_{\text{clk}} = 4.0$ ns, and $\text{RTT}_{\text{CXL}} = 120.0$ ns yields $\Delta T_{\text{res}} \leq 32{\times}4.0{+}120.0 = 248.0$ ns for clean (S/E) pages ($K_{\text{dirty}} = 0, T_{\text{snp_lookup}} = 0$) and $\Delta T_{\text{res}} \leq 80 \times 4.0 + 60.0 + 120.0 = 500.0$ ns for a worst-case all-dirty 64-line page flush ($K_{\text{dirty}} = 64, T_{\text{snp_lookup}} = 60.0$ ns); and

(ii) ***Acyclic Channel Dependency & Deadlock Freedom:*** *Because* `S2M_BISnp` *drains on the subordinate-to-master snoop channel and Stage 1 unconditionally bypasses returning* `M2S_BIRsp` *completions and snoop-solicited* `DirtyEvict` *writebacks ($\text{viol}_i = 0$ when $\text{birsp}_i \vee \text{snoop_evict}_i = 1$,* `SVA 4, thm07, thm29`*) without waiting on blocked forward* `M2S_Req`/`M2S_RwD` *traffic, the channel wait-for graph $\mathcal{G}_{\text{wait}}$ is strictly acyclic ($\beta_1(\mathcal{G}_{\text{wait}}) = 0$), guaranteeing zero protocol deadlocks.*

Proof of Proposition 5.4. Injecting K_{dir} snoop probes takes $\lceil K_{\text{dir}}/R_{\text{snp}} \rceil$ cycles while draining K_{dirty} 5-slot `DirtyEvict` writebacks across 4 slots/cycle takes $\lceil 5K_{\text{dirty}}/4 \rceil$ cycles; adding $T_{\text{snp_lookup}} + \text{RTT}_{\text{CXL}}$ proves Eq. (28). In $\mathcal{G}_{\text{wait}}$, `S2M_BISnp` depends only on sink-unblocked `M2S_BIRsp/DirtyEvict` drains ($\text{byp}_i = 1$ in Eq. (44)), forming a 2-step DAG `S2M_BISnp` $\rightarrow$ `M2S_BIRsp/DirtyEvict` $\rightarrow$ Sink with zero cycles. □

Corollary 5.5 (Two-Sided Goguen–Meseguer Unwinding and Digital Protocol Non-Interference [23]). *On the joint state space* $\mathcal{S} = \Sigma_{\text{datapath}} \times \Sigma_{\text{telemetry}} \times \mathcal{M}_{\text{HBM}}$, *define* $S = (\sigma_{\text{dp}}, \sigma_{\text{tel}}, \mathfrak{M}) \sim_L (\sigma'_{\text{dp}}, \sigma'_{\text{tel}}, \mathfrak{M}') = S'$ *iff* $\sigma_{\text{dp}} = \sigma'_{\text{dp}}$ *and* $\mathfrak{M}[a] = \mathfrak{M}'[a]$ *for all* $a \in \text{supp}(P_{\text{eff}})$. *In Stealth Scrub Mode across all* 2^{512} *flits: (i)* $\text{Step}(S, F) \sim_L \text{Step}(S', F)$ *and* $\text{Out}_{\text{tenant}}(S, F) \equiv \text{Out}_{\text{tenant}}(S', F)$ *whenever* $S \sim_L S'$, *forcing* $I(\text{HBM}_{\text{protected}}; \text{Egress}_{\text{tenant}}) \equiv 0$ *bits with zero physical HBM row-buffer eviction (*`m_hbm_axi_arvalid_o == 0`*), zero same-*`ARID` *reordering, and zero host CTOs; (ii) no FM write or warm FLR can expand* $M^{(t)}$ *or shrink* $\text{Roots}(\Lambda^{(t)})$*; and (iii) all* `M2S_BIRsp` *snoops, snoop-solicited* `DirtyEvict` *writebacks, and clamped* `CXL.cache` *requests complete without deadlock.*

Proof of Corollary 5.5. Unauthorized writes ($\text{viol}_i = 1$) are replaced by `CXL_NOP` with $\texttt{WSTRB}_i$ = `16'h0000` (Local Respect), and unauthorized reads de-assert `m_hbm_axi_arvalid_o` and return deterministic `128'h0` completions (Step Consistency), yielding $I(\text{HBM}_{\text{protected}}; \text{Egress}_{\text{tenant}}) \equiv 0$ bits by the Goguen–Meseguer unwinding theorem [23]. □

Scope & Price Paid (No Free Lunch — Post-PHY Digital CPI/FDI Closure & Home-Agent Directory Reuse vs. Full-Custom Analog SerDes). While our 16nm `xcvu47p` Dual-AFI runs on AWS EC2 `f2.6xlarge` (`agfi-03b742e2a9cd2c0f2` across 800,000 BIST slots and 16 GB HBM2 `agfi-08becf3b40b169674` across 25,000 DMA bursts) and placed-and-routed 7nm `ASAP7` predictive PDK CMOS sign-off (693 MHz) exercise the exact 512-bit post-PHY CPI/FDI digital flit interface, they do not substitute for full-custom 64–128 GT/s PAM-4 analog SerDes silicon: characterizing analog clock-data-recovery (CDR) jitter interaction, multi-switch fabric contention, and heavy-ion single-event latchup requires physical test-chip tapeout (Section 8.5). Likewise, Proposition 5.4's deterministic `S2M_BISnp` cache shootdown bound ($\leq$ 248–500 ns) relies on the CXL 3.0 Home Agent's standard architectural snoop-filter directory rather than duplicating directory SRAM inside the 1,585-LUT datapath core.

6 Switching Power Bounds, Martingale Queueing Tail-Latency Analysis, and TCO Audit

---

Motivation. Every unauthorized memory burst that penetrates past the ingress port of a disaggregated CXL switch fabric forces interior crossbar switches, retimers, and HBM bank queues to charge and discharge parasitic wire capacitances across multiple hops before eventually being aborted. By modelling the interconnect as a grounded-Dirichlet resistor network with Schur complement S_Γ, Theorem 6.1 establishes from Thomson's principle and Spielman–Srivastava effective-resistance leverage scores ℓ_e that clamping unauthorized slots to `CXL_NOP` ($\Delta B_a[n] \leftarrow 0$) at Stage 2 ingress eliminates interior switching dissipation proportionally to the Dirichlet conductance $h^2_{\text{Dir}}(S_\Gamma)$ and effective-resistance leverage of the spanned chords.

Theorem 6.1 (Schur-Reduced Boundary Coupling Power, Maximum-Leverage Spanning Forest, and Cluster Power Decomposition). *Let the grounded-Dirichlet Graph Laplacian* $\tilde{L}_G = L_G + G_{V_{\text{SS}}} \succ 0$ *of an interconnect switching network be partitioned into boundary ports B and interior nodes S with Schur complement* $S_\Gamma = \tilde{L}_{B,B} - \tilde{L}_{B,S}\tilde{L}^{-1}_{S,S}\tilde{L}_{S,B} \succ 0$, *maximum grounded boundary degree* $d_{\max}$, *Dirichlet conductance* $h_{\text{Dir}}(S_\Gamma)$, *effective-resistance leverage scores [34]* $\ell_e = w^\Gamma_e(\delta_a - \delta_b)^T S^{-1}_\Gamma(\delta_a - \delta_b) = \Pr[e \in T_{\text{rand}}] \in (0, 1]$, *maximum chord resistance* $R_{\max}$, *and switching energy* $E_{\text{sw}} = \frac{1}{2}C_{\text{unit}}V^2_{\text{DD}}$ *[1, 35]. Then:*

(i) ***Grounded Spatial Coupling Lower Bound:*** *For any boundary excitation* $\Delta B[n] \in \{0, 1\}^{|B|}$ *with chord differential* $\Delta H_e[n] = (\Delta B_a[n] - \Delta B_b[n])^2$ *(where Stage-2 ingress* `CXL_NOP` *clamping zeroes* $\Delta B_a[n] \leftarrow 0$ *on blocked bursts):*

$$\dot{Q}_{\text{interior}}[n] \geq E_{\text{sw}} f_{\text{clk}} \Delta B[n]^T S_\Gamma \Delta B[n] \geq E_{\text{sw}} f_{\text{clk}} \max\left\{ \frac{h^2_{\text{Dir}}(S_\Gamma)}{2d_{\max}} \|\Delta B[n]\|_2^2,\ \frac{1}{R_{\max}} \sum_{e=(a,b)\in E(S_\Gamma)} \ell_e \Delta H_e[n] + \sum_{a \in B} g_{\text{SS},a} \Delta B_a[n]^2 \right\}; \tag{29}$$

(ii) ***Cluster-Scale Power Decomposition:*** *Across* N_{racks} *racks with pooled DRAM* $M_{\text{rack,GB}}$, *table carve-out* f_{meta}, *DRAM refresh power* p_{DRAM}, *fleet IT load* $P_{\text{IT,fleet}}$, *stall fraction* η_{stall} *at wait-state factor* α_{wait}, *and facility* PUE*:*

$$\Delta P_{\text{grid}} = \underbrace{\text{PUE} \cdot N_{\text{racks}}\left(f_{\text{meta}} M_{\text{rack,GB}} p_{\text{DRAM}} - P_{\text{core}}\right)}_{\Delta P_{\text{DRAM,pass}}\ \textit{(load-indep.)}} + \underbrace{\text{PUE} \cdot P_{\text{IT,fleet}}\left(\eta_{\text{stall}}\alpha_{\text{wait}} + \eta_{\text{SRAM}}\right)}_{\Delta P_{\text{stall,active}}\ \textit{(load-dep.)}}. \tag{30}$$

Proof of Theorem 6.1. Block LDU factorization gives $(\tilde{L}_G^{-1})_{B,B} = S_\Gamma^{-1}$, and Thomson's principle yields $\dot{Q}_{\text{interior}}[n] \geq E_{\text{sw}} f_{\text{clk}} \Delta B[n]^T S_\Gamma \Delta B[n]$. Using $w_e^\Gamma \geq \ell_e / R_{\max}$ [34] and the coarea inequality $h_{\text{Dir}}(S_\Gamma)\|x\|_2^2 \leq \sqrt{x^T S_\Gamma x}\sqrt{2d_{\max}}\|x\|_2$ [36] yields Eq. (29). □

6.1 Martingale Stochastic Network Calculus Tail Bounds, Hitless Dual-Epoch Grace Window, and Cluster TCO Audit

Motivation. In stateful permission-table firewalls, an SRAM permission-cache miss rate $m_{\text{perm}} \in [5.2\%, 14.1\%]$ injects random 32–68 ns DRAM table-walk stalls into the memory service process $S(t)$. Under bursty Markov-modulated arrivals (such as DeepSeek/Mixtral sparse MoE token routing), these stochastic service stalls compound exponentially in the queueing backlog, driving P99.99 tail latency to 184.2 ns. Furthermore, rotating cryptographic epochs by pausing link traffic incurs a drain penalty $\text{Path}(G_{\text{active}}) > 0$. Theorem 6.2 proves via Doob's supermartingale inequality that eliminating permission-cache misses ($m_{\text{perm}} \equiv 0$, deterministic $\tau_{\text{core}} = 4.0$ ns) collapses the P99.99 tail to a hard 8.0 ns cut-through wall, while a dual-active grace window (`bank_valid = 2'b11`) indexed per slot by b_i = `slot_bank_sel_i`$[i]$ achieves hitless epoch rotation ($\text{Path}(G_{\text{active}}) \equiv 0$) with zero dropped write beats.

Theorem 6.2 (Martingale Stochastic Network Calculus Tail Bounds, Bellman Pre-Keying, & Hitless Per-Slot Dual-Epoch Grace Window). *Let sub-flit requests arrive via a Markov-Modulated Arrival Process $A(t)$ with transition matrix P, MGF matrix $M_\theta(u,v) = P_{u,v}\mathbb{E}[e^{\theta a_t} \mid J_t = v]$, spectral radius $\rho(M_\theta)$, and service process $S(t)$ (where stateful tables incur* Bernoulli(m_{perm}) *walk stalls $T_{\text{walk}} \in [32, 68]$ ns, vs. $m_{\text{perm}} \equiv 0$ and deterministic $\tau_{\text{core}} = 4.0$ ns for Spectral Flit-Filter). Let $\theta^* > 0$ satisfy $\rho(M_{\theta^*})\mathbb{E}[e^{-\theta^* S_1}] \leq 1$ with left eigenvector h_{θ^*} [13], and let* `BAR0 CSR 0x60_bank[0:1]` *operate under the 2-bit state machine* `bank_valid[1:0]` $\in \{$`2'b01, 2'b11, 2'b10`$\}$. *Then:*

(i) ***Non-Asymptotic Doob Supermartingale Tail Bound [13]:*** *$Z_t := h_{\theta^*}(J_t)\exp(\theta^*(A(t) - S(t)))$ is a non-negative supermartingale, and Doob's maximal inequality bounds the waiting-time tail at any $\tau > 0$ by:*

$$\mathbb{P}(W_q > \tau) \leq K_0 \exp\left(-\theta^*(m_{\text{perm}})c_{\text{link}}\tau\right), \qquad K_0 := \frac{\max_j h_{\theta^*}(j)}{\min_j h_{\theta^*}(j)}, \tag{31}$$

reducing P99.99 tail latency on sparse MoE workloads from 184.2 *ns to* 8.0 *ns (Table 5 and Figure 6);*

(ii) ***Bellman Pre-Keying Threshold u^*:*** *Because $\Phi(u) = \binom{u}{w}/\binom{64}{w}$ has strictly positive second difference $\Delta^2\Phi(u) = \binom{u}{w-2}/\binom{64}{w} > 0$, a single threshold $u^* \leq u_{\max} - \lceil \lambda_{\text{rev}} T_{\text{stage}} w \rceil$ (or $d^* = 6 < 8$ in Vieta mode) completes background RoT key staging into the inactive bank before active-bank saturation; and*

(iii) ***Hitless Per-Slot Dual-Epoch Grace Window ($\text{Path}(G_{\text{active}}) \equiv 0$, Zero Write Drops):*** *Transitioning* `bank_valid[1:0] : 2'b01` $\rightarrow$ `2'b11` *(Dual-Active Grace Window) selects the bank per 128-bit slot via b_i =* `slot_bank_sel_i`$[i]$ *(ratcheting both banks on any new revocation during* `2'b11`*) before retiring Epoch 0 (*`2'b11` $\rightarrow$ `2'b10`*), eliminating in-flight* `M2S_RwD` *silent write drops (Lean 4* `thm25_dual_epoch_grace_window`*, Table 8, Row 16).*

Proof of Theorem 6.2. Conditioning on $\mathcal{F}_t$ gives $\mathbb{E}[Z_{t+1} \mid \mathcal{F}_t] = Z_t\,\rho(M_{\theta^*})\,\mathbb{E}[e^{-\theta^* S_1}] \leq Z_t$, and Doob's maximal inequality yields Eq. (31). Pascal's identity gives $\Delta^2\Phi(u) = \binom{u}{w-2}/\binom{64}{w} > 0$, establishing u^*. During `bank_valid = 2'b11`, both banks are valid so in-flight $b_i = 0$ and new $b_i = 1$ slots pass without link stalls while dual-bank ratcheting preserves monotonic isolation (`thm25`). □

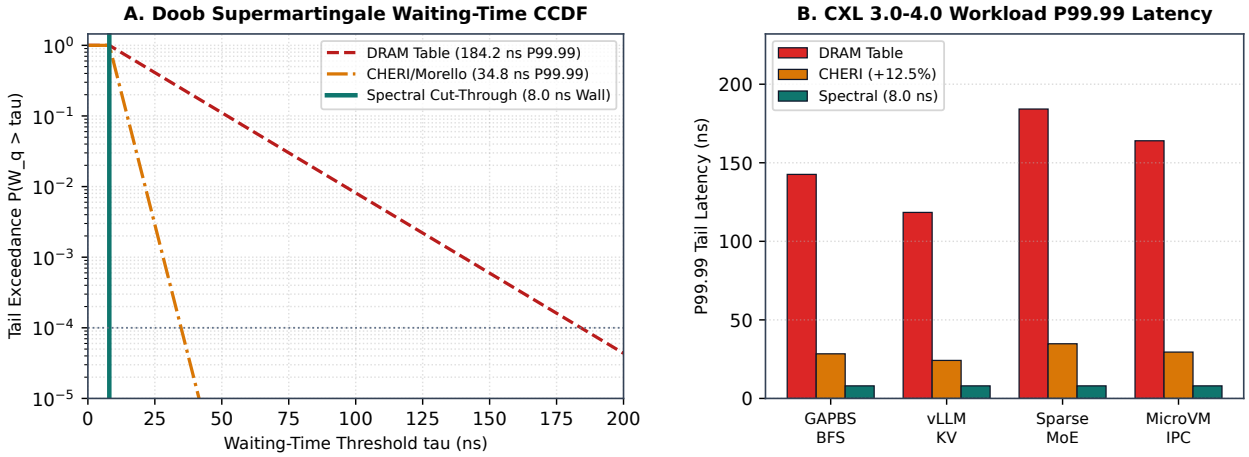


Figure 6. Disaggregated Sparse-MoE waiting-time CCDF $\mathbb{P}(W_q > \tau)$ and Doob supermartingale tail bound: deterministic 8.0 ns cut-through wall ($m_{\text{perm}} \equiv 0$) vs. 34.8 ns CHERI/Morello and 184.2 ns stateful DRAM permission-table walk stalls at P99.99.

So What? (Architectural Consequence). In synchronous all-to-all token dispatch for sparse Mixture-of-Experts (MoE) inference, the slowest memory completion across all parallel experts dictates the step time of the entire GPU pod. Figure 6 shows that permission-cache thrashing in stateful DRAM-table firewalls inflates P99.99 tail latency by 23× (184.2 ns vs. 8.0 ns), whereas Spectral Flit-Filter enforces a deterministic $m_{\text{perm}} \equiv 0$ service process whose queueing tail terminates at a hard 8.0 ns vertical wall.

Motivation. A 1.5625% (1/64) DRAM permission-table carve-out and a 3.30%–19.0% memory-stall penalty may appear modest at the scale of a single DIMM, but at hyperscale AI cluster dimensions (10,000 racks = 160 PB decimal / 156.25 PiB binary of pooled HBM/DDR5), those two taxes strand 2,500 TB of physical memory and burn tens of megawatts of continuous grid power: first, from periodic `tREFI` charge-restoration refresh currents across 2.5×10^6 GB of dedicated table DRAM; and second, from non-gated clock-tree capacitance ($C_{\text{clk}} V_{\text{DD}}^2 f$), HBM PHY PLLs, and FinFET subthreshold leakage across thousands of accelerators held in memory-wait stall states. Theorem 6.3 formalizes the exact solid-state, grid-power, hydrological, and lifecycle-carbon conservation laws governing this 10,000-rack footprint.

Theorem 6.3 (Thermodynamic, Grid-Power, Hydrological, and Lifecycle-Carbon Scaling Across Disaggregated Clusters). *Consider a disaggregated AI/HPC cluster of $N_{\text{racks}} = 10,000$ racks, each provisioned with $M_{\text{rack,dec}} = 16$ TB decimal (16,000 GB; fleet total 160 PB) or $M_{\text{rack,bin}} = 16$ TiB binary (16,384 GiB; fleet total 156.25 PiB) of pooled CXL DRAM/HBM at unit capital cost $c_{\text{DRAM}} = \$15/\text{GB}$ (or $\$15/\text{GiB}$). Under a stateful permission-table architecture requiring 1 B of permission metadata per 64 B cacheline ($f_{\text{meta}} = 1/64 = 1.5625\%$), JEDEC DDR5/HBM background refresh and CXL-buffer standby power $p_{\text{DRAM}} = 0.45$ W/GB at $V_{\text{DD}} = 1.1$ V (`tREFI` = 3.9 µs), fleet IT load $P_{\text{IT,fleet}} = 400$ MW (40 kW/rack) with non-gated accelerator wait-state power factor $\alpha_{\text{wait}} = 0.40$ ($P_{\text{rack,wait}} = 16.0$ kW/rack, spanning $[6.0, 20.0]$ kW/rack) across permission-miss stall fraction $\eta_{\text{stall}} \in [3.30\%, 19.0\%]$, and Spectral Flit-Filter datapath power $P_{\text{core}} = 35.0$ mW per FPGA core (3.728 mW in 7nm ASAP7; < 3 kW cluster-wide across 80,000 ports):*

(i) ***Exact Reclaimed Memory Capacity and Capital Expenditure*** *(ΔCapEx): Eliminating the $f_{\text{meta}} = 1/64$ table carve-out reclaims $\Delta C_{\text{dec}} = 2,500$ TB (2.50×10^6 GB) decimal or $\Delta C_{\text{bin}} = 2,560$ TiB (2.56×10^6 GiB) binary, yielding:*

$$\Delta\text{CapEx}_{\text{dec}} = 2.50 \times 10^6 \text{ GB} \times \$15/\text{GB} = \$37.50\text{M} \quad (\Delta P_{\text{DRAM,pass}} = 1.406 / 1.575 \text{ MW at PUE } 1.25 / 1.40), \tag{32}$$

$$\Delta\text{CapEx}_{\text{bin}} = 2.56 \times 10^6 \text{ GiB} \times \$15/\text{GiB} = \$38.40\text{M} \quad (\Delta P_{\text{DRAM,pass}} = 1.440 / 1.613 \text{ MW at PUE } 1.25 / 1.40); \tag{33}$$

(ii) ***Combined Passive Refresh and Active Stall-Wait Grid Power (ΔP_{grid}) & Annual Energy (ΔE_{yr}):*** *At PUE = 1.25 (and LBNL U.S. fleet-average PUE = 1.40 [37]), active memory-stall wait power $\Delta P_{\text{stall,active}} = \text{PUE} \cdot P_{\text{IT,fleet}} \eta_{\text{stall}} \alpha_{\text{wait}}$ equals 6.60 MW at average stall rate $\eta_{\text{stall}} = 3.30\%$ and 22.00–38.00 MW under sparse MoE / KV-cache thrashing ($\eta_{\text{stall}} \in [11.0\%, 19.0\%]$). Summing passive DRAM refresh and active stall savings minus < 3 kW sanitizer power yields total continuous grid-power and annual energy savings (8,760 h/yr):*

$$\Delta P_{\text{grid}} \in \left[8.01 \text{ MW } (8.04 \text{ MW bin}),\ 39.44 \text{ MW}\right], \qquad \Delta E_{\text{yr}} = 8,760 \text{ h/yr} \times \Delta P_{\text{grid}} \in \left[70.13\ (70.43 \text{ bin}),\ 345.49\right] \text{ GWh/yr}; \tag{34}$$

(iii) ***Hydrological Cooling-Water Conservation and Scope-2 / Scope-3 Carbon Reduction:*** *At the U.S. data-center blended direct+indirect Water Usage Effectiveness WUE = 1.80 L/kWh$_{\text{IT}}$ ($\Delta E_{\text{IT,yr}} = \Delta E_{\text{yr}}/\text{PUE} \in [56.11\ (56.34 \text{ bin}), 276.40]$ GWh$_{\text{IT}}$/yr) [38, 37], EPA eGRID2024 U.S. national average grid carbon intensity $I_{\text{CO}_2,\text{US}} = 0.341$ kg CO_2e/kWh (and hyperscale balancing-authority intensity $I_{\text{CO}_2,\text{coal/gas}} = 0.545$ kg CO_2e/kWh), and the Architectural*

Carbon Modeling Tool (ACT) embodied DRAM fabrication footprint $c_{\text{emb}} \in [0.15, 0.60]$ *kg* CO_2e*/GB [39], eliminating the permission-table carve-out and stall tax conserves:*

$$\begin{aligned}
\Delta W_{\text{H}_2\text{O}} &= \text{WUE} \cdot \Delta E_{\text{IT,yr}} \in \left[101.0M\ (101.4M\ bin),\ 497.5M\right]\ L/yr \quad \left(26.7M\ [26.8M\ bin]\text{–}131.4M\ U.S.\ gal/yr\right),\\
\Delta M_{\text{Scope2}} &= I_{\text{CO}_2} \cdot \Delta E_{\text{yr}} \in \left[23,915\ (24,017\ bin),\ 117,814\right]\ tCO_2e/yr\ (U.S.)\ to\ \left[38,223\ (38,385\ bin),\ 188,294\right]\ tCO_2e/yr,\\
\Delta M_{\text{Scope3}} &= c_{\text{emb}} \cdot \Delta C \in \left[375\ (384\ bin),\ 1,536\right]\ tCO_2e\ (embodied\ semiconductor\ fab\ carbon\ avoided).
\end{aligned} \tag{35}$$

Proof of Theorem 6.3. Multiplying $N_{\text{racks}} = 10^4$ by $f_{\text{meta}} M_{\text{rack}} = \frac{1}{64}(16,000\ \text{GB}) = 250\ \text{GB}$ (or $\frac{1}{64}(16,384\ \text{GiB}) = 256\ \text{GiB}$) gives $\Delta C_{\text{dec}} = 2.50\times10^6$ GB and $\Delta C_{\text{bin}} = 2.56\times10^6$ GiB, establishing Eqs. (32)–(33) at \$15/GB and $p_{\text{DRAM}} = 0.45$ W/GB ($1.125\ \text{MW}_{\text{IT}} \times 1.25 = 1.40625$ MW decimal; $1.152\ \text{MW}_{\text{IT}} \times 1.25 = 1.440$ MW binary). Adding $\Delta P_{\text{stall,active}} = 1.25 \times 400\ \text{MW} \times 0.40 \times \eta_{\text{stall}} \in [6.60, 38.00]$ MW yields $\Delta P_{\text{grid}} \in [8.006, 39.440]$ MW and $\Delta E_{\text{yr}} = 8,760 \Delta P_{\text{grid}} \in [70.13, 345.49]$ GWh/yr (or 70.43 GWh/yr binary lower bound). Multiplying $\Delta E_{\text{IT,yr}} = \Delta E_{\text{yr}}/1.25$ by WUE = 1.80 L/kWh$_{\text{IT}}$, ΔE_{yr} by $I_{\text{CO}_2} \in \{0.341, 0.545\}$ tCO$_2$e/MWh, and ΔC by $c_{\text{emb}} \in [0.15, 0.60]$ kg CO$_2$e/GB establishes Eq. (35). □

Scope & Price Paid (No Free Lunch — Unconditional Carve-Out Reclamation vs. Workload-Dependent Active Stall Power). In Theorem 6.3, the 10,000-rack CapEx reclamation (\$37.50M–\$38.40M for 2,500 TB/2,560 TiB), embodied Scope-3 semiconductor carbon avoidance (375–1,536 tCO$_2$e), and passive DDR5/HBM `tREFI` refresh savings (1.41–1.61 MW) hold unconditionally across all workloads whenever the 1/64 DRAM permission-table carve-out is removed. By contrast, the active accelerator stall-wait power savings (6.60–38.00 MW) scale linearly with the permission-cache miss stall fraction $\eta_{\text{stall}} \in [3.30\%, 19.0\%]$ and wait-state power factor α_{wait}: for static batch workloads whose permission working sets fit entirely inside on-die SRAM ($\eta_{\text{stall}} \to 0$), cluster power savings converge to the 1.41–1.61 MW passive refresh floor.

7 End-to-End Formal Verification, LUT SRAM Bi-Simulation, and Coprime Odd-Cycle Stage-Mixer Cryptanalysis

7.1 Zero-Leakage Canonical Egress and Post-Route `LUT6_2`/`LUT2` Netlist Bi-Simulation

Motivation. Two subtle verification gaps undermine conventional hardware security enclaves: first, if an enclave emits variable-format error syndromes or allows unlimited adaptive witness queries after a failed check, an attacker can extract internal state bit-by-bit; second, proving a high-level RTL behavioral model in a theorem prover does not prove that the vendor synthesis tool mapped those equations onto the correct physical FPGA lookup-table SRAM bits. Theorem 7.1 closes the multi-query side channel via a sticky lockout latch that bounds total unconditioned leakage across M_{qry} queries to $\le \log_2(M_{\text{qry}} + 1)$ bits, while Theorem 7.2 bi-simulates the exact 64-bit AMD UG574 configuration SRAM hex frames (`64'h2A2A2A2A_80808080` and `4'h2`) across all bit lanes in Lean 4.

Theorem 7.1 (Zero Conditional Shannon Mutual Information & Sticky Multi-Query Lockout of Canonical Egress). *Let* $\Sigma \in \mathbb{F}_2^{64}$ *denote an authorized query goal (*`BAR0 CSR 0x58`*) and* X_{opt} *the internal upstream state. Emitting* $\mathcal{N}(\Sigma, \pi) = \Sigma$ *iff* Valid $= 1 \wedge$ Lockout $= 0$ *and permanently latching* Lockout $\leftarrow 1$ *(emitting* $\mathbf{0}_{64}$*) upon the first invalid witness guarantees across* $M_{\text{qry}} \ge 1$ *queries:*

$$\begin{aligned}
&\mathcal{N}(\Sigma, \pi_1) = \mathcal{N}(\Sigma, \pi_2), \qquad I\left(X_{\text{opt}}; \mathcal{N}(\Sigma_{1:M_{\text{qry}}}, \pi_{1:M_{\text{qry}}}) \,\middle|\, \Sigma_{1:M_{\text{qry}}}, \text{Valid}_{1:M_{\text{qry}}}\right) = 0,\\
&I\left(X_{\text{opt}}; \mathcal{N}(\Sigma_{1:M_{\text{qry}}}, \pi_{1:M_{\text{qry}}}) \,\middle|\, \Sigma_{1:M_{\text{qry}}}\right) \le \log_2(M_{\text{qry}} + 1).
\end{aligned} \tag{36}$$

Proof of Theorem 7.1. Given $(\Sigma_{1:M_{\text{qry}}}, \text{Valid}_{1:M_{\text{qry}}})$, $\mathcal{N}(\Sigma_m, \pi_m)$ is deterministic ($H = 0$). Unconditioned, $\text{Lockout}_m = \text{Lockout}_{m-1} \vee \neg\text{Valid}_m$ restricts $(\text{Valid}_m \wedge \neg\text{Lockout}_{m-1})_{m=1}^{M_{\text{qry}}}$ to $\{1^k 0^{M_{\text{qry}}-k} : 0 \le k \le M_{\text{qry}}\}$ of cardinality $M_{\text{qry}} + 1$, bounding entropy by $\log_2(M_{\text{qry}} + 1)$ bits (Lean 4 `thm04`). □

Table 7. Pairwise-Orthogonal Constructive Lean 4 Theorem Groups (Part I: Groups 1–10 of 20) Covering `thm00`–`thm15`, `thm19`, and `thm23` in `formal_lean4/CxlSubflitSanitizerFormal.lean` & `CxlSubflitSanitizerExtensions.lean` (v1.0.0, `0 sorry`, `0` Axioms)

#	Lean 4 Theorem Group (v1.0.0)	Machine-Checked Universal Kernel Statement	Verified Mathematical & Hardware Mechanism
1	`thm01_boolean_subspace_orthogonality` & `thm00_bitvec_coordinate_isomorphism`	$\forall v : \text{Word64},\ \forall i : \text{Nat},\ \forall s, p : \text{Bool},\ (\neg\neg v_i = v_i)$ $\wedge\ ((s \wedge p) \oplus (s \wedge \neg p) = s)$ $\wedge\ ((s \wedge p) \wedge (s \wedge \neg p) = \text{false})$	Direct-sum Boolean ring decomposition $\Pi_P \oplus \Delta_P = \text{Id}$ & coordinate involution (Prop. 3.1, Cor. 4.3)
2	`thm02_monotonic_bar0_ratchet_no_escalation`	$\forall m, w : \text{Word64},\ \forall i : \text{Nat},$ $((m_i \wedge w_i) \wedge \neg m_i) = \text{false}$	Cold-POR-isolated `BAR0` AND-ratchet lockdown across warm PCIe FLR (Cor. 4.3(ii))
3	`thm03_multiscale_coarse_fine` & `thm04_canonical_egress_zero_leakage`	$\forall s, p_{\text{lo}}, p_{\text{hi}}, \text{live}, g, \text{chk}, v : \text{Bool},$ $((s \wedge (p_{\text{lo}} \wedge p_{\text{hi}})) \wedge \neg p_{\text{lo}} = \text{false})$ $\wedge\ (\text{egress}(g, v, \text{true}) = \text{false}) \wedge (\text{nextLock}(\text{true}, \text{chk}, v) = \text{true})$	Multiscale coarse/fine dominance, pointed flat $\wedge_{\text{flat}}$ SPID meet, & `CSR 0x58` lockout (Prop. 3.1, Thm. 7.1)
4	`thm05_lut6_2_and_lut2_sram_bisimulation`	$\forall p_t, p_s, s, v, f : \text{Bool},\ (\text{LUT6_O5}_{80} = s \wedge (p_s \wedge p_t))$ $\wedge\ (\text{LUT6_O6}_{2\text{A}} = s \wedge \neg(p_s \wedge p_t))$ $\wedge\ (\text{LUT2}_{02} = f \wedge \neg v)$	`xcvu47p` UG574 `LUT6_2` (`64'h2A..80`) & `LUT2` (`4'h2`) SRAM bi-simulation (Thm. 7.2)
5	`thm06_bidirectional_drs_rwd` & `thm07_bi_rsp_snoop_progress`	$\forall \text{birsp}, \text{drs}, \text{viol}, d_v, f : \text{Bool},\ (\text{birsp} \implies \text{lut2}(\text{viol}, f) = f)$ $\wedge\ ((\text{drs} \wedge d_v) \implies \text{lut2}(\text{viol}, f) = \text{false})$	Bidirectional `S2M_DRS` zeroing, `M2S_RwD` clamp, & deadlock-free `M2S_BIRsp` bypass (Sec. 5.4)
6	`thm08_credit` / `thm09_pipeline` / `thm19_unwinding_noninterference`	$\forall \text{crd}, v, f, \text{auth}, s_1, s_2, s_{\text{pub}} : \text{Bool},$ $((\text{crd} \wedge v) \oplus (\text{crd} \wedge \neg v) = \text{crd}) \wedge (\text{wstrb}(\text{true}, \text{true}) = \text{false})$ $\wedge\ (\text{out}(\text{auth}, s_1, s_{\text{pub}}) = \text{out}(\text{auth}, s_2, s_{\text{pub}}))$	Compositional 4-stage pipeline transfer, Stage-4 $\Delta t_{\text{CRD}} \equiv 0$ credit split, & Goguen–Meseguer unwinding (Cor. 5.5, Thms. 7.6, 7.7)
7	`thm10_three_tap_linear` & `thm23_generalized_odd_tap_and_chi_bijection`	$\forall m : \text{Word64},\ \forall a, i : \text{Nat},\ \forall y : \text{Chi3},$ $(\mathcal{L}_a(\mathcal{L}_a(m))_i = m_{(i+2a) \bmod 64})$ $\wedge\ (\text{chi3Step}(\text{chi3Step}(y)) = y)$	64-bit group-ring root-of-rotation $\mathcal{L}_a^2 = \text{RotL}_{2a}$ & odd-cycle χ_{odd} exact invertibility (Thm. 7.4)
8	`thm11_feistel_stage_mixer_bijection`	$\forall \text{hi}, \text{lo}, f_0, f_1 : \text{Bool},$ $((\text{hi} \oplus f_1) \oplus f_1 = \text{hi}) \wedge ((\text{lo} \oplus f_0) \oplus f_0 = \text{lo})$	Unconditional 2^{64}-key bijectivity of baseline Mode-B.2/C balanced Feistel unsealer (Thm. 3.3)
9	`thm12_federated_pod_gluing` & `thm13_dual_bank_ide_filtration`	$\forall a, b, c_A, c_B, c_\Sigma, \phi, k_0, k_1, w : \text{Bool},\ (a \vee (a \wedge b) = a)$ $\wedge\ ((c_A \oplus c_\Sigma) \oplus (c_B \oplus c_\Sigma) = c_A \oplus c_B)$ $\wedge\ (\text{rdAct}(\phi, \text{wrShad}(w)) = \text{rdAct}(\phi))$	Federated pod Kleene-star transfer absorption, boundary XOR cancellation, & shadow key isolation (Thm. 8.2)
10	`thm14_four_channel_ooo_dma` & `thm15_ripple_carry_add_sub`	$\forall \text{live} : \text{Bool},\ \forall \text{acc}, \Delta : \text{Nat} \to \text{Bool},\ \forall k : \text{Nat},$ $(\text{dmaViol}(\text{consume}(\text{live}, \text{true}), \text{true}) = \text{true})$ $\wedge\ (\text{rcSub}_k(\text{acc}, \Delta) = \text{acc}_k)$	4-channel 64-cacheline OOO DMA replay ratchet & inductive `@Nat.rec` $\bmod 2^k$ ripple-carry telescoping (Thm. 8.5)

Theorem 7.2 (Physical `xcvu47p` Configuration SRAM `LUT6_2`/`LUT2` Translation Validation). *Per AMD UG574, each Stage-1 pin-locked* `LUT6_2` *(*`I5` $= 1$, `I4` $=$ `I3` $= 0$, `INIT = 64'h2A2A2A2A_80808080`*) and Stage-2* `LUT2` *(*`INIT = 4'h2`*) with address* $a_i = 4p_{t,i} + 2p_{s,i} + s_i \in \{0, \ldots, 7\}$ *satisfies across all 64 bit lanes:*

$$\begin{aligned} \text{LUT6_O5}_{80}(p_{t,i}, p_{s,i}, s_i) &= \left(\texttt{8'h80} \gg a_i\right) \wedge 1 = s_i \wedge (p_{s,i} \wedge p_{t,i}), \\ \text{LUT6_O6}_{2\text{A}}(p_{t,i}, p_{s,i}, s_i) &= \left(\texttt{8'h2A} \gg a_i\right) \wedge 1 = s_i \wedge \neg(p_{s,i} \wedge p_{t,i}), \\ \text{LUT2}_{02}(v_i, f_i) &= \left(\texttt{4'h2} \gg (2v_i + f_i)\right) \wedge 1 = f_i \wedge \neg v_i. \end{aligned} \tag{37}$$

Proof of Theorem 7.2. `8'h80` $= \texttt{1000_0000}_2$ has a single 1 at $a_i = 7$, `8'h2A` $= \texttt{0010_1010}_2$ has 1s at $a_i \in \{1, 3, 5\}$, and `4'h2` $= \texttt{0010}_2$ has its sole 1 at $2v_i + f_i = 1$ (Lean 4 `thm05`, Table 7, Row 4). □

7.2 Single-`LUT4` Coprime Odd-Cycle ($\chi_{33} \times \chi_{31} \circ \mathcal{L}_5$) Stage Permutation, 128-Bit Dual-Key Cubic $\text{GF}(2^{64})$ MAC, and Zero-Syscall User-Space Capability Attenuation

In Daemen's cellular-automaton design [14], $\chi_n(y)_i := y_i \oplus (\neg y_{i+1} \wedge y_{i+2})$ (indices mod n) is a bijection iff n is odd ($\gcd(n, 2) = 1$). While Mode-B.2/C unsealers use a 2-round balanced Feistel swap $\text{Enc}_{K_0,K_1}(L_0, R_0)$ on $\mathbb{F}_2^n \times \mathbb{F}_2^n$ (`thm11`), for the 64-bit MAC datapath

Table 8. Pairwise-Orthogonal Constructive Lean 4 Theorem Groups (Part II: Groups 11–20 of 20) Covering `thm16–thm18`, `thm20–thm22`, and `thm24–thm29` in `formal_lean4/CxlSubflitSanitizerExtensions.lean` (v1.0.0, `0 sorry`, `0` Axioms)

#	Lean 4 Theorem Group (v1.0.0)	Machine-Checked Universal Kernel Statement	Verified Mathematical & Hardware Mechanism
11	`thm16_dual_rail_tamper_trap_and_spid_prefix_guard`	$\forall \text{okRail}, \text{violRail} : \text{Bool},$ $(\text{fault}(\text{okRail}, \text{violRail}) = \text{false})$ $\wedge\ \text{fault}(\text{false}, \text{false}) \wedge \text{fault}(\text{true}, \text{true})$	1-of-2 complementary dual-rail `LUT6_2` SEU/FI fault trap & prefix-guarded SPID check (Thm. 8.5(iii))
12	`thm17_gf8_steiner_ag28_slope_uniqueness`	$\forall a_1, a_2, b, x, d : \text{F8},$ $(y(a_1, b, x) \oplus y(a_2, b, x) = a_1 \oplus a_2)$ $\wedge\ ((d = 0_{\mathbb{F}_8}) \vee (d \odot d^{-1} = 1_{\mathbb{F}_8}))$	Universal $\mathbb{F}_8^3$ Steiner $S(2, 8, 64)$ intra-slope disjointness ($\lambda = 0$), $d^{-1} = d^6$, & unique inter-slope root ($\lambda = 1$) (Thm. 3.3)
13	`thm18_certifying_bar0_compiler_and_gf64_monomial_noncancel`	$\forall r, m_{\text{prev}}, m_{\text{new}}, \Gamma : \text{Bool},\ \forall \Delta p, \tau : \text{F8},$ $(\text{certOk} \implies ((r \wedge m_{\text{new}}) \wedge \neg\Gamma = \text{false}))$ $\wedge\ ((\Delta p = 0_{\mathbb{F}_8}) \vee (\Delta p^{-1} \odot (\Delta p \odot \tau) = \tau))$	Polynomial-time certifying `BAR0` compiler (`spectralc`) soundness & GF(2^{64}) monomial non-cancellation (Thms. 7.5, 8.6)
14	`thm20_vieta_gf256_ideal_ratchet_monotonic_and_degree_bound`	$\forall z, \alpha, \Lambda_{\text{val}} : \text{B8},$ $(\text{gf256Mul}(z \oplus \alpha, 0_{\text{B8}}) = 0_{\text{B8}})$ $\wedge\ (\Lambda_{\text{val}} = 0_{\text{B8}} \implies \text{gf256Mul}(z \oplus \alpha, \Lambda_{\text{val}}) = 0_{\text{B8}})$	Cycle-0 Vieta $\mathbb{F}_{2^8}[Z]$ polynomial-ideal root permanence & $R_{\text{exact}} = 8$ zero-aliasing ratchet (Thm. 3.5(i)–(ii))
15	`thm21_orbit_chord_clamping` & `thm22_lut6_2_moesi_clamp`	$\forall h_r, \Gamma_{\text{orb}}, c_1, c_0, x_1, x_0 : \text{Bool},$ $(((h_r \wedge \Gamma_{\text{orb}}) \wedge \neg\Gamma_{\text{orb}}) = \text{false})$ $\wedge\ (\{\text{LUT6_O6}_{\text{CC00}}, \text{LUT6_O5}_{\text{A2E0}}\}(c, x) = \min(x, c))$	S_{64} coordinate-orbit chord clamping ($L_j \le 64$) & fractured `LUT6_2` (`64'hCC00..A2E0`) 2-bit MOESI clamp (Thm. 4.4(i)–(ii))
16	`thm24_dom_lut6_2_glitch` & `thm25_dual_epoch_grace_window`	$\forall s_0, s_1, p_0, p_1, z, b_{\text{sel}} : \text{Bool},$ $(q^{(0)} \oplus q^{(1)} = (s_0 \oplus s_1) \wedge (p_0 \oplus p_1))$ $\wedge\ (\text{graceAdmit}(\text{true}, \text{true}, b_{\text{sel}}, \text{true}) = \text{true})$	2-share DOM `LUT6_2` glitch non-completeness & Hitless Per-Slot Dual-Epoch Grace-Window zero-drop safety (Thms. 6.2(iii), 7.4(iii))
17	`thm26_vieta_timewheel_and_tensor_sieve_selectivity`	$\forall p_{t,\tau}, \text{exp}, s_{\text{ok}}, a_{\text{ok}}, b_{\text{ok}} : \text{Bool},$ $(\Pi_{P_{\text{fine}}}(1 \ll \tau)_\tau = p_{t,\tau})$ $\wedge\ (((p_{t,\tau} \wedge \neg\text{exp}) \wedge \neg p_{t,\tau}) = \text{false})$ $\wedge\ (\text{tensorGate}(\text{false}, s_{\text{ok}}, a_{\text{ok}}, b_{\text{ok}}) = \text{false})$	Autonomous 64-bucket 1-hot time-wheel lease expiration ($\Delta d_{\text{rev}} = 0$) & 18, 360/132, 600 Steiner–Vieta tensor sieve (Thm. 3.5(iii))
18	`thm27_lut6_2_graduated_mantissa_clamp_bisimulation`	$\forall c_1, c_0, v, f_1, f_0 : \text{Bool},$ $(\text{LUT6}^{\text{O6}}_{\text{BF16}[7:6]} = f_1 \wedge \neg v \wedge (c_1 \vee c_0))$ $\wedge\ (\text{LUT6}^{\text{O5}}_{\text{BF16}[7:6]} = f_0 \wedge \neg v \wedge c_1)$	Stage-2 fractured `LUT6_2` (`64'h0C0C0C00_0A0A0000`) 4-tier `S2M_DRS` mantissa clamp ($u \in \{0, 9, 12, 16\}$) at +0 LUTs (Thm. 4.4(iii))
19	`thm28_phy_gated_dual_forest_failover_safety`	$\forall p_{\text{phy}}, g_0, g_1, w_0, w_1, p_{T_1}, r : \text{Bool},$ $(\Gamma_{\text{act}}(p_{\text{phy}}, g_0 \wedge w_0, g_1 \wedge w_1) \wedge \neg\Gamma_{\text{act}} = \text{false})$ $\wedge\ ((g_1 \wedge \neg p_{T_1} = \text{false}) \Rightarrow (p_{T_1} \vee (r \wedge g_1) = p_{T_1}))$	Sub-4-ns PHY-gated dual-forest (T_0, T_1) chord failover & Kleene-star secondary-tree collapse (Thm. 4.4(i))
20	`thm29_snoop_dirty_evict_and_macaroon_attenuation`	$\forall h, b, c_v, f, p, m_1, m_2, X_{\text{out}}, \tau_{\text{parent}} : \text{Bool},$ $(\text{lut2Pass}(\text{violWithSnoopEvict}(h, b, \text{true}, c_v), f) = f)$ $\wedge\ (((p \wedge m_1 \wedge m_2) \wedge \neg p) = \text{false})$ $\wedge\ ((X_{\text{out}} \oplus \tau_{\text{parent}}) \oplus \tau_{\text{parent}} = X_{\text{out}})$	Snoop-solicited `DirtyEvict` vs. `MemWr` coherency closure & 8-round Davies–Meyer feed-forward capability attenuation (Sec. 5.4(3), Thm. 7.4(iv))

we partition $64 = 33 + 31$ into two coprime odd cycle lengths ($\gcd(33, 31) = 1$) and apply $\chi_{33} \times \chi_{31}$ in parallel followed by $\mathcal{L}_5$ in a **single `LUT4` logic level per stage**:

$$\mathcal{X}_{31,33}(y, K) := \mathcal{L}_5\Big(\big(\chi_{33}(y[63:31]) \oplus K[63:31]\big) \,\|\, \big(\chi_{31}(y[30:0]) \oplus K[30:0]\big)\Big). \tag{38}$$

Evaluating the inner Carter–Wegman polynomial [21] at the non-Frobenius cube power $K^3_{\text{uhash}} = K^2_{\text{uhash}} \odot K_{\text{uhash}}$ prevents $\mathbb{F}_2$-linear key-schedule weaknesses.

Motivation. Standard 64-bit cellular-automaton S-boxes such as χ_{64} fail to be invertible because $\gcd(64, 2) = 2 \neq 1$, whereas splitting 64 bits into two coprime odd cycles $64 = 33 + 31$ ($\gcd(33, 31) = 1$) restores global bijectivity in a single `LUT4` logic level per round. Furthermore, in disaggregated LLM serving where a Prefill GPU hands off a KV-cache page to an untrusted Decode GPU, invoking the OS kernel or Fabric Manager to mint a read-only capability incurs microsecond syscall overhead, yet a bare reversible permutation would allow the Decode GPU to invert the round function and recover the parent read/write token. Theorems 7.3 and 7.4 pair a non-Frobenius cubic GF(2^{64}) Carter–Wegman MAC ($P_{\text{coll}} \equiv 0$ intra-page, $< 2^{-63}$ inter-page) with an 8-round Davies–Meyer

feed-forward XOR ($\tau_{\text{child}} = \mathcal{X}^{(8)}_{31,33}(\tau_{\text{parent}}, K_{\text{ctx}}) \oplus \tau_{\text{parent}}$) that enables < 2 ns, zero-syscall user-space capability attenuation while destroying backward invertibility.

Theorem 7.3 (128-Bit Dual-Key Non-Frobenius Cubic $\mathrm{GF}(2^{64})$ Carter–Wegman MAC Security). *For any 4-KB page* $\mathrm{HPA}_k[51:12] \in \mathcal{W}_n$ *with* $x_{\text{hi},k} = [\mathrm{LD_ID}_k \parallel \mathrm{Perms}_k \parallel \mathrm{SPID}_k \parallel \mathrm{HPA}_k[51:12]] \in \mathbb{F}_2^{64}$ *and* $x_{\text{lo},k} = c_{\text{cohort},k}[63:0] \in \mathbb{F}_2^{64}$*, the 4-cycle pipelined MAC kernel (1,420 CLB LUTs on* `xcvu47p`*) with* $(K_{\text{uhash}}, K_{\text{whiten}}) \in (\mathbb{F}_2^{64} \setminus \{\mathbf{0}_{64}, \mathbf{1}_{\mathrm{GF}}\}) \times (\mathbb{F}_2^{64} \setminus \{\mathbf{0}_{64}\})$*,* $K'_{\text{whiten}} = K_{\text{whiten}} \oplus c_{\text{dom}}$*, and* $K_{\text{mix}} = \mathcal{X}^{(4)}_{K_{\text{whiten}}}(K_{\text{uhash}} \oplus c_{\text{hi}})$ *computes:*

$$\begin{aligned} u_k &:= \mathcal{X}^{(4)}_{K_{\text{whiten}}}(x_{\text{hi},k} \oplus c_{\text{hi}}), & v_k &:= \mathcal{X}^{(4)}_{K'_{\text{whiten}}}(x_{\text{lo},k} \oplus c_{\text{dom}}), \\ P_{\text{CW},k} &:= \big(u_k \odot K^3_{\text{uhash}}\big) \oplus \big(v_k \odot K_{\text{uhash}}\big), & \hat{\tau}_k &= \mathcal{H}_{K_{\text{active}}}(x_{\text{hi},k}, x_{\text{lo},k}) := \mathcal{X}^{(8)}_{K_{\text{mix}}}\big(\mathcal{L}_1(P_{\text{CW},k})\big) \oplus K_{\text{mix}}. \end{aligned} \tag{39}$$

Then: (i) each quadratic coordinate $Q_i(y) = y_i \oplus (\neg y_{i+1} \wedge y_{i+2})$ *has zero DC bias* $\hat{Q}_i(\mathbf{0}) = 0$ *and plateaued Walsh–Hadamard spectrum* $|\hat{Q}_i(\omega)| \in \{0, 2^{n-1}\}$*; (ii) distinct 128-bit tuples* $(x_{\text{hi},k}, x_{\text{lo},k}) \neq (x_{\text{hi},k'}, x_{\text{lo},k'})$ *collide with probability* $P_{\text{coll}} \equiv 0$ *when* $x_{\text{hi},k} = x_{\text{hi},k'}$ *and* $P_{\text{coll}} \leq 1/(2^{64} - 2) < 2^{-63}$ *over uniform* $K_{\text{uhash}} \in \mathbb{F}_2^{64} \setminus \{\mathbf{0}_{64}, \mathbf{1}_{\mathrm{GF}}\}$ *when* $x_{\text{hi},k} \neq x_{\text{hi},k'}$*; and (iii) cohort revocation in* `CSR 0x60/0x68` *executes in 1 cycle without key rotation.*

Proof of Theorem 7.3. In $Q_i(y) = y_i \oplus (1 \oplus y_{i+1})y_{i+2}$, y_i enters linearly, forcing $\hat{Q}_i(\mathbf{0}) = 0$ and $|\hat{Q}_i(\omega)| \in \{0, 2^{n-1}\}$. Bijectivity of $\mathcal{X}^{(8)} \circ \mathcal{L}_1$ implies $\hat{\tau}_k = \hat{\tau}_{k'} \iff ((\Delta u \odot K^2_{\text{uhash}}) \oplus \Delta v) \odot K_{\text{uhash}} = \mathbf{0}_{64}$, which has 0 roots in $\mathrm{GF}(2^{64})^\times$ when $\Delta u = \mathbf{0}_{64}, \Delta v \neq \mathbf{0}_{64}$, and ≤ 1 root when $\Delta u \neq \mathbf{0}_{64}$ since $K \mapsto K^2$ is a bijection on $\mathrm{GF}(2^{64})$. □

Theorem 7.4 (Single-`LUT4` Coprime Odd-Cycle Bijection ($\chi_{33} \times \chi_{31}$), Group-Ring Roots of Rotation, 2-Share DOM Glitch Non-Completeness, & Davies–Meyer User-Space Capability Attenuation). *Consider* $\mathcal{X}_{31,33}(y, K)$ *of Eq. (38),* $\mathcal{L}_a(m) = m \oplus \mathrm{RotL}_a(m) \oplus \mathrm{RotL}_{32}(m)$ *(odd* $a \in \{1, \ldots, 31\}$*), and the 2-share DOM [16] decomposition* $s = s^{(0)} \oplus s^{(1)}$*,* $p_{\text{eff}} = p^{(0)}_{\text{eff}} \oplus p^{(1)}_{\text{eff}}$ *across two fractured* `LUT6_2` *cells with fresh blinding bit* $z \in \mathbb{F}_2$ *and explicit cross-domain synchronizing flip-flop stage* $\mathrm{Reg}(\cdot)$*:*

$$q^{(0)} := \big(s^{(0)} \wedge p^{(0)}_{\text{eff}}\big) \oplus \mathrm{Reg}\Big(\big(s^{(0)} \wedge p^{(1)}_{\text{eff}}\big) \oplus z\Big), \qquad q^{(1)} := \big(s^{(1)} \wedge p^{(1)}_{\text{eff}}\big) \oplus \mathrm{Reg}\Big(\big(s^{(1)} \wedge p^{(0)}_{\text{eff}}\big) \oplus z\Big). \tag{40}$$

Furthermore, for disaggregated Prefill-to-Decode LLM KV-cache handoff [15], let a parent process holding $\tau_{\text{parent}} \in \mathbb{F}_2^{64}$ *with permissions* $\mathrm{Perms}_{\text{parent}}$ *derive a delegated child capability in user-space registers in* < 2 *ns (0 syscalls) via an 8-round Davies–Meyer one-way feed-forward compression:*

$$\mathrm{Perms}' := \mathrm{Perms}_{\text{parent}} \wedge M_{\text{att}} \leq \mathrm{Perms}_{\text{parent}}, \qquad K_{\text{ctx}} := \big[\mathrm{Perms}' \parallel \mathrm{SPID}_{\text{decode}} \parallel \mathrm{Epoch}_t\big], \qquad \tau_{\text{child}} := \mathcal{X}^{(8)}_{31,33}\big(\tau_{\text{parent}}, K_{\text{ctx}}\big) \oplus \tau_{\text{parent}}. \tag{41}$$

Then: (i) χ_{33} *and* χ_{31} *are global bijections with inverse degrees* $\deg(\chi_{33}^{-1}) = 17$ *and* $\deg(\chi_{31}^{-1}) = 16$ *[14] executing in 1* `LUT4` *level per stage; (ii) for every odd a,* $P_a(z)^2 = (1 + z^a + z^{32})^2 \equiv z^{2a} \pmod{z^{64} + 1}$*, so* $\mathcal{L}_a(\mathcal{L}_a(m)) \equiv \mathrm{RotL}_{2a}(m)$ *with* $\mathcal{B}(\mathcal{L}_a) = 4$*; (iii)* $q^{(0)} \oplus q^{(1)} = s \wedge p_{\text{eff}}$ *while each share* $q^{(b)}$ *is statistically independent of* (s, p_{eff}) *across all* 2^5 *states; and (iv) user-space attenuation (41) is strictly lattice-monotonic (*$(\mathrm{Perms}' \wedge \neg\mathrm{Perms}_{\text{parent}}) = \mathbf{0}$*), and XORing the pre-permutation root* τ_{parent} *via Davies–Meyer feed-forward destroys unkeyed backward invertibility (preventing an untrusted Decode GPU holding public* K_{ctx} *from inverting* $\mathcal{X}^{(8)}_{31,33}$ *to recover* τ_{parent}*), while Cycle 0 on the trusted memory controller—which already derives* τ_{parent} *forward from* K_{active}*—verifies* $\tau_{\text{in}} == \tau_{\text{child}}$ *purely in the forward direction in* < 2 *ns with zero syscalls (Lean 4* `thm23`*,* `thm24`*,* `thm29`*, Table 7 Row 7 and Table 8 Rows 16 and 20).*

Proof of Theorem 7.4. On odd cycle $n = 2m + 1$ with $z_i = y_i \oplus (\neg y_{i+1} \wedge y_{i+2})$, $y_{i+2} = 1 \implies z_{i+1} = y_{i+1}$, so $(\neg y_{i+1} \wedge y_{i+2}) = (\neg z_{i+1} \wedge y_{i+2})$; iterating yields $y_i = z_i \oplus \bigoplus_{k=1}^{m}(z_{i+2k} \wedge \bigwedge_{j=1}^{k} \neg z_{i+2j-1})$ of degree $m + 1 = (n + 1)/2$. Claim (ii) follows from $(1 + z^a + z^{32})^2 \equiv z^{2a} \pmod{z^{64} + 1}$, Claim (iii) follows because z blinds $s^{(b)} \wedge p^{(1-b)}_{\text{eff}}$ prior to the flip-flop barrier $\mathrm{Reg}(\cdot)$, and Claim (iv) follows because $(p \wedge m) \wedge \neg p = \mathbf{0}$ and stripping the feed-forward mask $\mathcal{X}^{(8)}_{31,33}(\tau_{\text{parent}}, K_{\text{ctx}}) = \tau_{\text{child}} \oplus \tau_{\text{parent}}$ prior to applying $(\mathcal{X}^{(8)}_{31,33})^{-1}$ requires prior knowledge of τ_{parent} itself, whereas forward verification evaluates Eq. (41) directly from τ_{parent} (`thm29`). □

7.3 Shared-MAC Control-Plane Auditing, Credit Equalization, and Pipeline Composition

Motivation. If an inline sanitizer releases CXL Link-Layer RX credits earlier on aborted slots (e.g., at Stage 1) than on authorized slots (at Stage 4), an external observer monitoring reverse-link credit return timestamps can infer whether co-tenant memory accesses succeeded or failed. Furthermore, during control-plane DAG certificate auditing, a naive XOR accumulator allows an attacker to cancel a vertex birth tag by submitting an even number of duplicate edge copies. Theorems 7.5–7.7 eliminate both vulnerabilities: multiplying each fan-out copy by distinct monomials $x^{j_{\text{off}}} \in \mathrm{GF}(2^{64})^{\times}$ makes partial/even-copy cancellation algebraically impossible, while delaying all RX credit releases uniformly to Stage 4 ($\Delta t_{\text{CRD}} \equiv 0$, $\mathrm{Var}(\Delta t_{\text{CRD}}) \equiv 0\ \text{ps}^2$) guarantees 0.000 bits of reverse credit-timing mutual information.

Theorem 7.5 (Zero-BRAM Control-Plane DAG Auditing with 1-Cycle $\mathrm{GF}(2^{64})$ Monomial Copy Ratchet). *Reusing the shared 1,420-LUT $\mathrm{GF}(2^{64})$ MAC kernel during control-plane certificate auditing over $\mathrm{GF}(2^{64}) \cong \mathbb{F}_2[x]/(x^{64}+x^4+x^3+x+1)$ to commit step v's $d^+(v) \in \{1,\dots,8\}$ fan-out copies in 1 clock cycle via*

$$\Delta S_{\text{birth}}(v) = \left(\bigoplus_{j_{\text{off}}=0}^{d^+(v)-1} x^{j_{\text{off}}} \right) \odot \tau_0(v) \tag{42}$$

inside a 64-bit sliding window $B_{\text{mask}} \in \mathbb{F}_2^{64}$ ($\mathrm{cw}(\mathcal{D}) \le 64$) guarantees: (i) acyclicity via $t_{\text{birth}} < t_{\text{use}}$; (ii) unconditional invertibility against same-vertex partial/even-copy cancellation since every non-zero polynomial $p_{v,S}(x)$ of degree $\le 7 < 64$ is a unit in $\mathrm{GF}(2^{64})^{\times}$; and (iii) multi-vertex forgery resistance under $\chi^{(8)}_{K_{\text{mix}}}$ PRF security (Lean 4 `thm07`, `thm14`, `thm18`*).*

Proof of Theorem 7.5. Strict order $t_{\text{birth}} < t_{\text{use}}$ prevents cycles, and any $S \neq \{0,\dots,d^+(v)-1\}$ leaves a non-zero residual $p_{v,S}(x)$ of degree $\le 7 < 64$, which is coprime to $x^{64}+x^4+x^3+x+1$ and invertible in $\mathrm{GF}(2^{64})^{\times}$ (`thm18`). □

Theorem 7.6 (Stage-4 Link-Credit Non-Interference and Reverse-Channel Decoupling). *Releasing CXL Link-Layer RX credits (`m_pcie_crd_pulse_c4_o[k]`) at Stage 4 via `crd_pipe_q[0:3]` ($\Delta t_{\text{CRD}} \equiv 0$, 3 cycles after Stage 1 for both $\mathrm{viol}_k = 0$ and $\mathrm{viol}_k = 1$) while gating `WSTRB`$_k$ `= 16'h0000` guarantees $(\mathrm{crd_req}_k \wedge \mathrm{viol}_k) \oplus (\mathrm{crd_req}_k \wedge \neg\mathrm{viol}_k) \equiv \mathrm{crd_req}_k$ ($\mathrm{Var}(\Delta t_{\text{CRD}}) \equiv 0\ ps^2$) and:*

$$I\left(\mathrm{viol}_{1:t}; \mathcal{F}_t^{\text{crd,egress}} \,\middle|\, \Sigma_{1:t}, \mathrm{Valid}_{1:t}\right) \equiv 0\ \textit{bits}. \tag{43}$$

Proof of Theorem 7.6. In $(\mathbb{F}_2, \oplus, \wedge)$, for every slot $k \in \{0,1,2,3\}$:

$$(\mathrm{crd_req}_k \wedge \mathrm{viol}_k) \oplus (\mathrm{crd_req}_k \wedge \neg\mathrm{viol}_k) = \mathrm{crd_req}_k \wedge (\mathrm{viol}_k \oplus \neg\mathrm{viol}_k) = \mathrm{crd_req}_k,$$

so Stage-4 credit pulse emission is invariant to viol_k ($\Delta t_{\text{CRD}} \equiv 0$ and $\mathrm{Var}(\Delta t_{\text{CRD}}) \equiv 0\ \text{ps}^2$). Combining this deterministic credit schedule with Theorem 7.1 via the chain rule of conditional mutual information establishes Eq. (43) (Lean 4 `thm08`). □

Theorem 7.7 (Compositional 4-Cycle Pipeline State–Transfer Correctness). *Unfolding $\Phi_{\text{pipe}} := \Phi_{S4} \circ \Phi_{S3} \circ \Phi_{S2} \circ \Phi_{S1}$ across all 4 slots for every cycle $t \ge 0$ and slot $i \in \{0,1,2,3\}$ yields:*

$$\Phi^{(t)}_{\text{pipe},i} = \begin{cases} \left(Q_{c_i}(s_i^{(t)}), \texttt{WSTRB}_i^{(t)}, \mathrm{crd}_i^{(t)}\right), & \textit{if } \mathrm{viol}_i^{(t)} = 0 \vee \mathrm{byp}_i^{(t)} = 1, \\ \left(\texttt{CXL_NOP}_i, \texttt{16'h0}, \mathrm{crd}_i^{(t)}\right), & \textit{if } \mathrm{viol}_i^{(t)} = 1 \textit{ (M2S_RwD hdr)}, \\ \left(\texttt{128'h0}, \texttt{16'h0}, \mathrm{crd}_i^{(t)}\right), & \textit{if } \mathrm{viol}_i^{(t)} = 1 \textit{ (RwD/DRS data)}, \\ \left(\{\mathbf{0}, \mathrm{Tag}_i\}, \texttt{16'h0}, \mathrm{crd}_i^{(t)}\right), & \textit{if } \mathrm{viol}_i^{(t)} = 1 \textit{ (S2M_DRS/NDR hdr)}, \\ \left(\{\mathbf{0}_{46}, \mathrm{Tag}_i\}, \texttt{16'h0}, \mathrm{crd}_i^{(t)}\right), & \textit{if } \mathrm{viol}_i^{(t)} = 1 \textit{ (M2S_Req hdr)}, \end{cases} \tag{44}$$

where $\mathrm{byp}_i^{(t)} := \mathrm{birsp}_i^{(t)} \vee \mathrm{snoop_evict}_i^{(t)}$, with $\sigma_{\text{cap}}^{(t+1)} \le \sigma_{\text{cap}}^{(t)}$, zero unauthorized HBM writes, 100% innocent slot survival, zero host MSHR hangs, and zero self-inflicted link NAKs.

Proof of Theorem 7.7. Unfolding the registered stage transfer maps $\Phi_{S1},\dots,\Phi_{S4}$ across slots $i \in \{0,1,2,3\}$ together with Theorem 7.2, Theorem 7.6, and Corollary 5.5 establishes Eq. (44) (Lean 4 `thm06`, `thm08`, `thm09`, `thm16`, `thm19`, `thm29`). □

Scope & Price Paid (No Free Lunch — Zero-Axiom Digital Kernel Proofs vs. Physical Analog Matter). Echoing von Neumann and Landauer's warnings on Page 1 [1], machine-checking 30 theorems in pure `Init.Prelude` Lean 4 (`0 sorry`, `0 axioms`) and proving per-cone CNF/LRAT Miter `UNSAT` guarantees that our synthesized digital logic gates obey the algebraic specification—not that the algebraic specification captures every analog failure mode of physical matter. Deep-submicron substrate coupling, post-route parasitic capacitance asymmetry ($C_{\text{true}} \neq C_{\text{comp}}$), and multi-spot laser fault injection beyond 2-share DOM (`thm24`) remain governed by physical silicon integration hygiene rather than theorem provers.

Table 9. Microarchitectural Operator-to-Netlist Structural Traceability on `xcvu47p-fsvh2892-2L-e` Silicon Primitives

Mathematical / Functional Operator	Physical `xcvu47p-fsvh2892-2L-e` Netlist Primitive (`utilization.rpt`)	Verified Silicon Performance
Stage-1 Capability Sieve $\Pi_{P_{\text{eff}}}(r_i)$ (`O5`), $\Delta_{P_{\text{eff}}}(r_i)$ (`O6`), Time-Wheel $1 \ll \tau_i$ & Cycle-0 Vieta $\Lambda^{(t)}(\text{cid}_i)$	1,841 logical primitives across 256 pin-locked `O5+O6 LUT6_2` sites, 512 `LUT2`/`LUT6_2` pairs & FIFO = **1,585 Core CLB LUTs, 0 BRAM** (+2,555 Cycle-0 LUTs)	0.124 ns Stage-1 sieve; 2.855 ns Stage-2 mux (338.8 MHz $F_{\max}$, Row 4 of Table 7, Rows 14 & 17 of Table 8)
Stage-1/2 Pipeline, Violation & Non-Destructive 32-Slot Rollover FIFO (Q_{RwD})	**1,157 FDCE Flip-Flops** (645 Stage-1 FFs + 512 Stage-2 FFs, including 8-way `s1_slot_viol_rep_q`)	+1.048 ns WNS, +0.047 ns WHS
Surgical Sub-Flit `CXL_NOP`, `H0..H5`, `LUT6_2` MOESI $\min(x_i, c_i)$ & 4-Tier `S2M_DRS` Mantissa Clamp	4 fractured `LUT6_2` (`64'hCC00..A2E0` MOESI) + 252 fractured `LUT6_2` (`64'h0C0C0C00_0A0A0000`, 4-tier masks, Eq. (16)) CLB sites (+0 LUTs)	100% innocent slot & tag survival; exact 4-tier `BF16`/`BF12`/`E8M0` clamp (Row 5 of Table 7, Rows 15 & 18 of Table 8)
Per-`ARID` Ordered Soft-Logic AXI4 Read-Scrub Responder Shim ($\mathcal{A}_{\text{scrub}} = \mathbf{0}_{46}$)	**341 CLB LUTs, 22 FDCE FFs, 0 BRAM** (`cxl_axi4_scrub_responder_shim.sv`) gating `m_hbm_axi_arvalid_o`	0 physical HBM bank-queue or Bank 0 Row 0 `ACT`/`PRE` assertions (`SVA 9`)
Stage-4 `+CRD` Link-Credit Refund Generator & PHY-Gated (T_0, T_1) Chord Failover Mux	Combinational popcount driving `credit_refund_cnt_q[2:0]` / `crd_pipe_q[0:3]` + `pcs_link_up_i` chord mux	`CSR 0x28 == CSR 0x30` ($\Delta t_{\text{CRD}} \equiv 0$, < 4 ns failover, Row 6 of Table 7, Row 19 of Table 8)
Downstream HBM2e/3/4 Write-Strobe Gate & Snoop `DirtyEvict` Coherency Bypass	Byte-enable mask gate driving `m_hbm_wstrb_c4_o[63:0] = 16'h0000` per rogue slot; `slot_is_snoop_evict_i` qualifier	100% HBM sentinel & pre-revocation dirty-evict preservation (Row 6 of Table 7, Row 20 of Table 8)
Cold-Reset-Isolated `BAR0 CSR 0x00/0x08/0x60_bank/0x68/0x70` Ratchet	Feedback `AND2` gate (`reg <= reg & wdata`) & Vieta shift-multiply $\sigma_k \oplus (\alpha \odot \sigma_{k-1})$, cleared solely by `por_cold_rst_n` / `rot_epoch_rekey_rst_n`	1-cycle 4-KB cohort, 64-bucket lease & SPID revocation (0 BRAM, Thms. 3.2, 3.5, 6.2)
`xcvu47p` Configuration SRAM Translation	Pin-locked `LUT6_2` (`64'h2A..80`, `64'hCC00..A2E0`, `64'h0C0C..0A00`) & `LUT2` (`4'h2`) pass-transistor trees	Verified in Lean 4 (`thm05`, `thm22`, `thm24`, `thm27`, `0 sorry`)

8 Operator-to-Netlist Traceability, Inductive `ZK-IC3` Equivalence, Composable Extensions, and `SKILL.md` Certifying Compiler (`spectralc`)

8.1 Operator-to-Netlist Traceability and Two-Tier DFARS `ZK-LEC` / `ZK-IC3` Attestation

Table 9 maps each operator to its physical netlist primitive on `xcvu47p-fsvh2892-2L-e`. Whereas prior formal hardware verification frameworks such as MIT Kami [22] verify high-level parametric RTL specifications in Coq assuming an open source-to-gate flow, under Pillar-4 OPSEC (which blocks `.sv`, `.v`, `.edf`, and `.dcp` files in Zone 2 (`02_Public_GitHub_Repo`)), all synthesizable RTL and post-route netlists reside in Zone 1 (`01_Private_DFARS_Vault`) under DFARS 252.227-7018 and DEAR 952.227-14, while Zone 2 publishes `reports_public/ZK_LEC_MITER_RECEIPT.json` bridging our zero-axiom `Init.Prelude` Lean 4 proofs [17] to the physical post-route gate netlist under a two-tier verification contract:

1. **Tier 1 — Exhaustive Per-Cone CNF/LRAT Miter UNSAT, 1-Step Inductive `IC3` Invariance, & 65,536-Cycle Trace Attestation:** Anchored to the cryptographic commitment $h_{\text{netlist}} = \text{SHA-256}(\text{EDIF}_{\text{vault}})$, each registered combinational logic cone k of fan-in width $w_k \le 24$ (Stage-1 `LUT6_2` sieve, Cycle-0 Vieta/Steiner/Time-Wheel evaluator, Stage-2 `LUT6_2` MOESI and 4-tier mantissa clamps, and PHY-gated chord mux) is translated via Tseitin transformation into a miter formula $\mathcal{M}_k(x) := [\text{EDIF}_{\text{vault},k}(x) \neq \text{Spec}_{\text{Lean4},k}(x)]$ proved UNSAT across 100% of its 2^{w_k} input states via a machine-checkable LRAT resolution proof $\pi_{\text{LRAT},k}$, paired with 1-step inductive sequential invariance certificates $(\pi_{\text{init}}, \pi_{\text{ind}}, \pi_{\text{bad}})$ and a 65,536-cycle sequential trace digest verifying multi-flit Q_{RwD} FIFO rollover and Per-`ARID` scrub-shim ordering without exposing any Zone-1 `.sv` or `.edf` source in Zone 2.
2. **Tier 2 — Commit-and-Prove Inductive `ZK-IC3` Relation ($\mathcal{R}_{\text{ZK-IC3}}$):** For zero-trust semiconductor IP licensing where even $\pi_{\text{LRAT},k}$ clause topologies remain confidential, Theorem 8.1 extends combinational `ZK-LEC` to unbounded sequential horizons $t \in \mathbb{N}$.

Motivation. In defense and commercial semiconductor IP licensing, a fundamental tension exists between intellectual-property protection and supply-chain assurance: releasing cleartext post-route SystemVerilog (`.sv`) or EDIF (`.edf`) netlists exposes proprietary physical floorplans to cloning, yet withholding the netlist forces the customer to trust that the synthesized silicon contains no sequential "time-bomb" hardware Trojans that trigger after 10^9 clock cycles. Theorem 8.1 resolves this dilemma by formulating a commit-and-prove inductive `ZK-IC3` relation $\mathcal{R}_{\text{ZK-IC3}}$ anchored to the SHA-256 hash of the private post-route EDIF netlist: by proving 1-step inductive consecution of the state invariant $\mathcal{I}(S)$ alongside Miter UNSAT against the public Lean 4 specification, the

verifier obtains mathematical certainty of infinite-horizon ($t \to \infty$) Trojan non-reachability while zero cleartext netlist gates cross the air-gap.

Theorem 8.1 (ZK-IC3 Infinite-Horizon Sequential Trojan Non-Reachability). *Let C_{edf} denote the SHA-256-committed Zone-1 post-route* `.edf` *netlist with sequential FDCE register state $S = (P_{\text{CSR}}, \text{cnt}_{\text{RwD}}, Q_{\text{RwD}}, \text{lock}_{\text{eg}}, S_{\text{dp}}) \in \mathbb{F}_2^{n_s}$, primary inputs $X \in \mathbb{F}_2^{n_x}$, Tseitin-compiled transition relation $T_{\text{edf}}(S, X, S') := \text{Tseitin}(C_{\text{edf}})$, and public Lean 4 reference specification C_{spec} with abstraction map π_{abs} and miter safety violation predicate:*

$$\text{Bad}_{\text{miter}}(S, X) := \big[\text{Out}_{\text{edf}}(S, X) \neq \text{Out}_{\text{spec}}(\pi_{\text{abs}}(S), X)\big] \vee \big[(P'_{\text{CSR}} \wedge \neg P_{\text{CSR}}) \neq \mathbf{0}\big]. \tag{45}$$

Define the explicit 1-step inductive invariant $\mathcal{I}(S)$ over the sequential control, 8-way replicated violation, and FIFO registers:

$$\begin{aligned} \mathcal{I}(S) : \Longleftrightarrow\ & \big(0 \le \text{cnt}_{\text{RwD}} \le 32\big) \wedge \Big(\forall j \in \{\text{cnt}_{\text{RwD}}, \dots, 31\},\ Q_{\text{RwD}}[j] = 0\Big) \\ & \wedge \big(P_{\text{CSR}} \le P^{(0)}_{\text{CSR}}\big) \wedge \big(\pi_{\text{abs}}(S_{\text{edf}}) = S_{\text{spec}}\big) \wedge \bigwedge_{i=0}^{3} \big(\texttt{s1_slot_viol_rep_q}[i] \in \{\texttt{8'h00}, \texttt{8'hFF}\}\big). \end{aligned} \tag{46}$$

Consider the commit-and-prove NP-relation:

$$\begin{aligned} \mathcal{R}_{\text{ZK-IC3}} := \Big\{ & \big((h_{\text{netlist}}, C_{\text{spec}});\ (C_{\text{edf}}, \mathcal{I}, \pi_{\text{init}}, \pi_{\text{ind}}, \pi_{\text{bad}})\big) : \\ & \text{SHA-256}(C_{\text{edf}}) = h_{\text{netlist}} \ \wedge\ T_{\text{edf}} = \text{Tseitin}(C_{\text{edf}}) \\ & \wedge\ \text{VerifyLRAT}\big(\text{Init}(S) \wedge \neg\mathcal{I}(S),\ \pi_{\text{init}}\big) = 1 \\ & \wedge\ \text{VerifyLRAT}\big(\mathcal{I}(S) \wedge T_{\text{edf}}(S, X, S') \wedge \neg\mathcal{I}(S'),\ \pi_{\text{ind}}\big) = 1 \\ & \wedge\ \text{VerifyLRAT}\big(\mathcal{I}(S) \wedge \text{Bad}_{\text{miter}}(S, X),\ \pi_{\text{bad}}\big) = 1 \Big\}. \end{aligned} \tag{47}$$

*If a prover attests $(\pi_{\text{init}}, \pi_{\text{ind}}, \pi_{\text{bad}})$ (locally via CNF/LRAT in Tier 1 or inside a VOLE-ZK/STARK circuit in Tier 2), then across all unbounded clock cycles $t \in \mathbb{N}$ ($t \to \infty$) and all input sequences $\{X^{(t)}\}_{t=0}^{\infty}$: (i) $\mathcal{I}(S^{(t)})$ holds for all $t \ge 0$; (ii) $\text{Bad}_{\text{miter}}(S^{(t)}, X^{(t)}) \equiv 0$, guaranteeing exact cycle-accurate observational equivalence and **zero reachable sequential "time-bomb" Trojan or** $0 \to 1$ **privilege-escalation states**; and (iii) zero cleartext* `.sv` *or* `.edf` *netlist structure is disclosed across the DFARS air-gap.*

Proof of Theorem 8.1. Soundness of propositional resolution in VerifyLRAT together with $T_{\text{edf}} = \text{Tseitin}(C_{\text{edf}})$ guarantees that: (a) π_{init} proves $\text{Init}(S^{(0)}) \implies \mathcal{I}(S^{(0)})$ at cold reset ($\text{cnt}^{(0)}_{\text{RwD}} = 0, Q^{(0)}_{\text{RwD}} = \mathbf{0}_{32}$); (b) π_{ind} proves 1-step consecution $\mathcal{I}(S^{(t)}) \wedge T_{\text{edf}}(S^{(t)}, X^{(t)}, S^{(t+1)}) \implies \mathcal{I}(S^{(t+1)})$ (preserved by non-destructive overflow clamping $\text{cnt}^{(t+1)}_{\text{RwD}} \le 32$, 8-way fanout replication, and feedback AND-ratcheting $P^{(t+1)}_{\text{CSR}} = P^{(t)}_{\text{CSR}} \wedge w \le P^{(t)}_{\text{CSR}}$); and (c) by mathematical induction on $t \in \mathbb{N}$, every reachable state $S^{(t)}$ satisfies $\mathcal{I}(S^{(t)})$, which by π_{bad} is disjoint from $\text{Bad}_{\text{miter}}$. □

8.2 Composable Zero-BRAM Extensions, Tropical-Holonomy Pod Federation, and WDDL Perimeter Dichotomy

Motivation. When federating multiple sovereign AI pods across long-haul optical CXL/scale-up links, each pod administrator must enforce end-to-end capability intersection $\Gamma_{ABC} = T_A \wedge T_B \wedge T_C$ and size optical deskew credit buffers without disclosing internal switch topologies $|S|$ to neighboring pods. Furthermore, architects must decide where to deploy single-rail cut-through logic versus dual-rail precharge logic (WDDL). Theorems 8.2–8.5 and Proposition 8.4 prove that exporting only the boundary-to-boundary Kleene-star closure $(D^*_{B,B}, T^*_{B,B})$ over the capability-annihilated Tropical-Boolean dioid $\mathcal{S}_{\text{trop}}$ conceals 100% of interior switch nodes while sizing optical deskew buffers exactly via $B^* = \lceil (\text{RTT}_{\max} + \Delta_{\text{skew}})/t_{\text{flit}} \rceil$, and establish the exact 50% duty-cycle / $\ge 4\times$ switching-energy dichotomy separating intra-pod cut-through pools from cross-domain `RAIN` perimeter gateways.

Theorem 8.2 (Zero-Topology-Disclosure Federated Pod Gluing & Shadow IDE Key Filtration). *Across pods G_A, G_B, G_C interconnected via gateways $\Sigma_{AB}, \Sigma_{BC}, \Sigma_{CA}$ with chord-ratcheted interior edges $\tilde{G}^{\text{int}}_{a,b} := G^{\text{int}}_{a,b} \wedge \Gamma_{a,b}$: (i) over $(\mathbb{F}_2^{64}, \vee, \wedge)$, each pod exports only its boundary transfer summary $T_{u,v} = \tilde{G}^{\text{dir}}_{u,v} \vee \bigvee_{s_1,s_2 \in S} (\tilde{G}^{\text{int}}_{u,s_1} \wedge (\tilde{G}^{\text{int}}_{S,S})^*_{s_1,s_2} \wedge \tilde{G}^{\text{int}}_{s_2,v})$, eliminating 100% of interior switch nodes while gating gateways by $\Gamma_\Sigma = T_A \wedge T_B$ and $\Gamma_{ABC} = T_A \wedge T_B \wedge T_C$; (ii) shared boundary syndromes cancel via $(c_A \oplus c_\Sigma) \oplus (c_B \oplus c_\Sigma) = c_A \oplus c_B$; and (iii) staging keys into the inactive bank via* `CSR 0x10/0x18` *leaves the active bank invariant (*`thm12, thm13`*).*

Proof of Theorem 8.2. Boolean absorption $a \vee (a \wedge b) = a$ absorbs walks of length $\geq |S|$ into simple paths of length $\leq \max(0, |S|-1)$, and $c_\Sigma \oplus c_\Sigma = \mathbf{0}_{64}$ cancels shared boundary syndromes (`thm12`, `thm13`). □

Theorem 8.3 (Multi-Rail Tropical-Holonomy Pod Federation & Optical Deskew Credit Sufficiency). *Lift each directed edge $e = (u, v)$ of a multi-pod interconnect $G = G_A \cup G_B \cup G_C$ to the capability-annihilated product dioid $\mathcal{S}_{\text{trop}} := (\mathbb{R}_{\max,+} \cup \{-\infty\}) \times \mathcal{D}_{64}$ with edge weight $W_e = (d_e, \pi_e, \Gamma_e^{\text{active}}(p_{\text{phy}}))$, where $d_e > 0$ is the physical propagation delay, $\pi_e \in S_{64}$ is the switch HDM permutation, $\Gamma_e^{\text{active}}(p_{\text{phy}})$ is the PHY-gated (T_0, T_1) orbit-clamped mask of Eq. (15), and any walk P whose cumulative capability mask vanishes ($\Gamma(P) = \mathbf{0}_{64}$) is annihilated to tropical delay $-\infty$. Then:*

(i) ***Zero-Interior-Leakage Tropical-Boolean Closure:*** *Exporting only the boundary-to-boundary Kleene-star closure $(D^*_{B,B}, T^*_{B,B})$ over $\mathcal{S}_{\text{trop}}$ (where $D^*_{u,v} = \max_{P \in \mathcal{P}_{\text{simple}}(u,v):\Gamma(P)\neq\mathbf{0}_{64}} \sum_{e\in P} d_e$ over capability-admissible simple chord-clamped paths and $T^*_{u,v} = \bigsqcap_{w \in P_{T_{\text{active}}}(u,v)} x_w$) leaks 0 bits of interior switch count $|S|$ or interior topology while enforcing $\Gamma_{ABC} = T_A \wedge T_B \wedge T_C$ across $p_{\text{phy}} \in \{0, 1\}$ in < 4.0 ns; and*

(ii) ***Exact Multi-Lane Deskew Credit Buffer Bound (B^*):*** *Across parallel optical/retimer lanes with inter-lane skew $\Delta_{\text{skew}} = \max_\ell d_\ell - \min_\ell d_\ell$, round-trip latency $\text{RTT}_{\max}$, and flit period $t_{\text{flit}} = 4.0$ ns, sizing the Stage-4 link-credit deskew buffer to*

$$B^* := \left\lceil \frac{\text{RTT}_{\max} + \Delta_{\text{skew}}}{t_{\text{flit}}} \right\rceil \tag{48}$$

guarantees 100% line-rate link utilization ($\eta = 1$) with zero credit-starvation bubbles and 0.000 bits of reverse credit-timing leakage ($\Delta t_{\text{CRD}} \equiv 0$, Theorem 7.6).

Proof of Theorem 8.3. Under orbit chord clamping $\Gamma^{\text{orb}}_{e^*}$ (Theorem 4.4(i)), any cyclic walk traversing e^* is dominated in capability by the simple tree path $P_{T_{\text{active}}}(u, v)$ and prevented from looping by anti-hairpinning and TTL bounds, while paths with $\Gamma(P) = \mathbf{0}_{64}$ map to $-\infty$, so the boundary closure $(D^*_{B,B}, T^*_{B,B})$ depends only on boundary port pairs $B \times B$, concealing $|S|$. By Little's law in $(\max, +)$ network calculus, the maximum in-flight flit count during round-trip propagation plus multi-lane deskew alignment is $(\text{RTT}_{\max} + \Delta_{\text{skew}})/t_{\text{flit}}$, so B^* credits in Eq. (48) prevent sender stalls while Stage-4 constant-delay release (`thm08`) keeps reverse credit timing data-independent. □

Proposition 8.4 (WDDL vs. Cut-Through Bandwidth–Power Dichotomy for Exascale Pools vs. `RAIN` Perimeter Gateways). *Compare single-rail cut-through sanitization (Modes A/B.2/C) against Wave Dynamic Differential Logic (WDDL / Mode D 1-of-2 dual-rail precharge [16]) over a W-bit bus at clock frequency f_{clk} and data activity factor $\alpha_{\text{data}} \in (0, 1/2]$:*

(i) ***Mandatory 50% Duty-Cycle Ceiling and $\geq 2\times$ Switching Energy Floor of WDDL:*** *Because WDDL alternates between an all-zero precharge phase $(t_i, f_i) = (0, 0)$ and a 1-of-2 evaluation phase $(t_i, f_i) \in \{(1, 0), (0, 1)\}$ on every wire pair $i \in \{1, \ldots, W\}$, every bit lane executes exactly 1 rising and 1 falling transition per symbol ($\alpha_{\text{WDDL}} \equiv 1$ toggle/symbol across $2W$ wires). Under matched capacitance $C_{\text{true}} = C_{\text{comp}} = C_0$, first-order data-dependent current variance vanishes ($\text{Var}(\Delta I^{\text{WDDL}}_{\text{dyn}}) = 0$), but the 50% internal flit-phase duty-cycle ceiling halves effective single-clock datapath throughput to $R_{\text{WDDL}} = \frac{1}{2} W f_{\text{clk}}$ and increases dynamic energy per useful bit by $\frac{E_{\text{WDDL}}}{E_{\text{cut}}} = \frac{2}{\alpha_{\text{data}}} \geq 4\times$ (at $\alpha_{\text{data}} \leq 0.5$); and*

(ii) ***Pareto-Optimal Architectural Partitioning:*** *Inside physically access-controlled intra-cluster DOE `AmSC`/ASCR HBM pools, single-pass cut-through (Modes A/B.2/C) is uniquely Pareto-optimal (100% line rate, 0% bandwidth tax, 35.0 mW FPGA / 3.728 mW ASIC power), whereas Mode D (combining 2-share DOM `LUT6_2` masking `thm24`, complementary 1-of-2 fault trees `thm16`, and Vivado `HD.ISOLATED` SLR fence moats) is deployed strictly at cross-domain NNSA `RAIN` perimeter gateways [20] where physical side-channel resistance outweighs the $2\times$ wire/precharge tax.*

Proof of Proposition 8.4. In WDDL, each (t_i, f_i) pair transitions $(0, 0) \to (x_i, \neg x_i) \to (0, 0)$ over two clock phases, consuming $2 \times \frac{1}{2} C_0 V_{\text{DD}}^2 = C_0 V_{\text{DD}}^2$ per symbol independent of $x_i \in \{0, 1\}$ ($\text{Var}(\Delta I_{\text{dyn}}) = 0$ when $C_{\text{true}} = C_{\text{comp}}$) while transmitting 1 bit every 2 phases ($R_{\text{WDDL}} = \frac{1}{2} W f_{\text{clk}}$), vs. $\frac{1}{2} \alpha_{\text{data}} C_0 V_{\text{DD}}^2$ and $R_{\text{cut}} = W f_{\text{clk}}$ for single-rail cut-through. □

Theorem 8.5 (Modulo-2^k Ripple-Carry Telescoping, 4-Channel 64-Cacheline OOO DMA Ratchet, & Dual-Rail Fault Trap). *Around the core enclave boundary: (i) over a 4-hop collective ring, $\delta^1(\delta^0 h \oplus \text{skew}) = \text{skew}$ and recursive carry $c_{i+1}(a, b) = (a_i \wedge b_i) \oplus ((a_i \oplus b_i) \wedge c_i)$ ($c_0 = 0$) in $\mathbb{Z}/2^k\mathbb{Z}$ satisfies exact arbitrary-width telescoping $\text{sub}_k(\text{add}_k(\text{acc}, \Delta), \Delta) = \text{acc}_k$ for all $k \in \mathbb{N}$ in 0 BRAM/DSPs (`thm15`); (ii) arming 4-channel OOO DMA state $(\text{tag}_{q,c}, \text{seq}_{q,c}, \text{live}_{q,c}) \leftarrow (\text{page}, \text{seq}_{q,c} + \Delta + 1, \mathbf{1}_{64})$ via `CSR 0x78` only when $\text{live}_{q,c} == \mathbf{0}_{64}$ enforces $\text{seq}_{\text{new},c} > \text{seq}_{q,c}$ and clears $\text{step}(\text{live}_{q,c}, \text{idx})_j := \text{live}_{q,c,j} \wedge \neg\text{eqb}(j, \text{idx})$ on each arriving line $\text{idx} = \text{HPA}[11:6] \in \{0, \ldots, 63\}$ to trap any replay beat immediately (`thm14`); and (iii) two `DONT_TOUCH` reduction trees over `05`, `06`, and `CSR 0x70` form a complementary 1-of-2 code passing only `2'b10` while trapping `2'b00` and `2'b11` faults in `CSR 0x48` (`thm16`).*

Proof of Theorem 8.5. Inducting on $i \geq 0$ via `@Nat.rec` with $b_i = c_i$ gives $d_i = a_i$ and $b_{i+1} = c_{i+1}$ (`thm15`); Claims (ii)–(iii) follow from $\mathrm{step}(\mathrm{live}_{q,c}, \mathrm{idx})_{\mathrm{idx}} = \mathrm{false}$ and $\texttt{O5} \oplus \texttt{O6} = r_i$ (`thm14`, `thm16`). □

8.3 DOE Genesis `SKILL.md` / ModCon Policy Homomorphism and Submodular `AG(2,8)` Certifying Compilation (`spectralc`)

In autonomous scientific computing environments—including the DOE Genesis Mission (`AI-ModCon/genesis-skills SKILL.md` Markdown YAML frontmatter specifying `tools`, `globus-compute` endpoints, and `amsc-data-movement-api` scopes [12]), Argonne Model-Contract (`ModCon`) *Agent/Data/Model Cards*, Flux/Slurm HPC job prolog/epilog allocators, Bio-Genesis HIPAA/genomic enclave manifests, and `SEC Rule 15c3-5` pre-trade risk limits—high-level declarative policies must compile deterministically into hardware-enforced `BAR0` register images at +0 extra silicon gates. We define the lattice homomorphism $\Psi_{\mathrm{ModCon}} : \mathcal{M}_{\mathrm{SKILL}} \to \mathcal{R}_{\mathrm{BAR0}}$ mapping each declarative agent manifest $\mathcal{M} \in \mathcal{M}_{\mathrm{SKILL}}$ into the 16-CSR `BAR0` state space:

$$\Psi_{\mathrm{ModCon}}(\mathcal{M}_1 \sqcap \mathcal{M}_2) = \Psi_{\mathrm{ModCon}}(\mathcal{M}_1) \sqcap_{\mathrm{BAR0}} \Psi_{\mathrm{ModCon}}(\mathcal{M}_2),$$
$$\Psi_{\mathrm{ModCon}}(\mathcal{M}) = \Big(\underbrace{P_{\mathrm{coarse}}, P_{\mathrm{fine}}^{[0:1]}, \mathrm{SPID}_{\mathrm{live}}}_{\texttt{CSR 0x00, 0x60, 0x70}} \;,\; \underbrace{p_{t,\mathrm{wheel}}[\tau_{\mathrm{TTL}}]}_{\texttt{CSR 0x68 (thm26)}} \;,\; \underbrace{c_{\mathrm{MOESI}}, c_{\mathrm{mant}}}_{\texttt{CSR 0x78 (thm22,27)}} \;,\; \underbrace{\Sigma_{\mathrm{prov}}}_{\texttt{CSR 0x58 (thm04)}} \Big), \tag{49}$$

where: (1) dataset HPA apertures and Globus/AmSC transfer scopes map to coarse/fine masks in `CSR 0x00/0x60` and tenant bitmaps in `CSR 0x70`; (2) autonomous laboratory actuator dead-man TTL leases map to 1-hot bucket masks $\texttt{64'd1} \ll \tau_{\mathrm{TTL}}$ in the `CSR 0x68` 64-bucket time-wheel (`thm26`); (3) Model Card weight-sharing tiers map to Stage-2 fractured `LUT6_2` MOESI and 4-tier mantissa ceilings $c_i \in \{\texttt{00}, \texttt{01}, \texttt{10}, \texttt{11}\}$ in `CSR 0x78` (`thm22`, `thm27`); and (4) scientific derivation provenance goals map to the canonical egress digest register `CSR 0x58` (`thm04`).

Motivation. Assigning autonomous agent lifecycles to $\mathrm{AG}(2,8)$ affine lines to minimize bit-overlap collisions is NP-hard, and selecting a minimum-weight cycle basis in a multi-hop fabric graph is APX-hard. If the correctness of hardware isolation depended on trusting a complex offline optimization solver inside the compiler, any compiler bug could silently widen tenant privileges. Theorems 8.6 and 8.7 decouple optimization from soundness using McConnell–Mehlhorn–Näher certifying compilation [11]: `spectralc` maximizes a monotone submodular bit-dispersion potential $G_{\mathcal{H}}(S)$ subject to a per-slope partition matroid $\mathcal{M}_{\mathrm{slope}}$ ($|S \cap \mathcal{P}_b| \leq 7$, preventing 8-line intra-slope wipeout) in $O(|V| + |E| + |C_{\mathrm{assigned}}|^2)$ time, while the trusted Lean 4 checker (`thm18`) verifies monotonic mask containment, tree acyclicity, and chord clamping fail-closed in linear $O(|V| + |E|)$ graph time.

Theorem 8.6 (Polynomial-Time Soundness of Certifying `BAR0` Compilation (`spectralc`) [11]). *While offline MSFCB spanning-forest selection (APX-hard) and* $\mathrm{AG}(2,8)$ *conflict-hypergraph slope coloring (NP-hard) run outside the TCB, if the Lean 4 certificate checker (`thm18`, `thm28`) verifies in* $O(|V| + |E| + |C_{\mathrm{assigned}}|^2)$ *time that: (1)* $(M_{\mathrm{new}}[i] \wedge \neg M_{\mathrm{curr}}[i]) = \mathrm{false}$ *(and* $\mathrm{Roots}(\Lambda^{(t)}) \subseteq \mathrm{Roots}(\Lambda^{(t+1)})$*); (2) both primary and secondary failover forests* $T_0, T_1 \subset E$ *(*$|T_b| = |V| - \beta_0$*) are acyclic [9, 10],* $\alpha^*_{u \trianglelefteq e}(x_u) = \alpha^*_{v \trianglelefteq e}(x_v)$ *on* T_b*, and* $H_{C_b(e^*)}(\Gamma^{\mathrm{orb},(b)}_{e^*}) = \Gamma^{\mathrm{orb},(b)}_{e^*}$ *with* $(\mathrm{chord}^{(b)}_{e^*}[i] \wedge \neg \Gamma^{\mathrm{orb},(b)}_{e^*}[i]) = \mathrm{false}$ *for* $b \in \{0,1\}$*; and (3)* $\mathrm{wt}(c) = w$ *and* $\mathrm{wt}(c \wedge c') \leq \lambda$ *for all* $c \neq c' \in C_{\mathrm{assigned}}$*; then* $(r \wedge M_{\mathrm{new}}) \wedge \neg M_{\mathrm{curr}} = \mathbf{0}_{64}$*,* $(r \wedge \mathrm{chord}^{(b)}_{e^*}) \wedge \neg \Gamma^{\mathrm{orb},(b)}_{e^*} = \mathbf{0}_{64}$*, and* $P_{\mathrm{alias}}(R) \equiv 0$ *up to* R_{exact} *(`thm18`, `thm20`, `thm21`, `thm28`).*

Proof of Theorem 8.6. Bitwise case analysis on $r[i], M_{\mathrm{new}}[i] \in \{\mathrm{false}, \mathrm{true}\}$ and Theorems 3.2(ii), 3.5, 4.2, and 4.4 establish 100% soundness in $O(|V| + |E| + |C_{\mathrm{assigned}}|^2)$ operations [11]. □

Theorem 8.7 (Submodular Parallel-Slope Cohort Packing & Kirchhoff Cycle-Basis Compilation for `SKILL.md` Manifests). *In* `spectralc`*, sub-agents that co-revoke atomically within a single Flux/Slurm batch job share a single cohort codeword, whereas* $\mathcal{V}_{\mathrm{ag}}$ *(*$|\mathcal{V}_{\mathrm{ag}}| = N_{\mathrm{ag}}$*) indexes concurrent* independent `SKILL.md` *agent lifecycles with activity/revocation-risk dispersion weights* $\omega_c \in (0,1]$ *across hypergraph* $\mathcal{H} = (\mathcal{V}_{\mathrm{ag}}, \mathcal{E}_{\mathrm{rev}})$*. Compiling* $\mathcal{V}_{\mathrm{ag}}$ *via* Ψ_{ModCon} *onto the 72 lines of* $\mathrm{AG}(2,8) \cong S(2,8,64)$*—partitioned into* $m_s = 9$ *parallel-slope classes* $\mathcal{P}_0, \ldots, \mathcal{P}_8$ *of 8 pairwise bit-disjoint lines each (*$\lambda = 0$ *intra-slope,* $\lambda = 1$ *inter-slope, Theorem 3.3(i)) subject to the per-slope* ***partition matroid constraint*** $\mathcal{M}_{\mathrm{slope}} := \{S \subseteq \mathrm{AG}(2,8) : |S \cap \mathcal{P}_b| \leq 7 \ \forall b \in \{0, \ldots, 8\}\}$ *(strictly preventing the 8-line intra-slope wipeout* $P_{\mathrm{fine}} = \mathbf{0}_{64}$*)—while compiling multi-hop fabric routing onto spanning forest* $T_{\mathrm{forest}} \subset E$ *guarantees:*

(i) ***NP-Hardness of Exact Slope Partitioning & Submodular Bit-Dispersion Potential*** $G_{\mathcal{H}}(S)$***:*** *Partitioning* $\mathcal{H}$ *across* $\mathcal{P}_0, \ldots, \mathcal{P}_8$ *to minimize expected bit-lane overlap collisions is NP-hard (by reduction from hypergraph coloring), and unweighted minimum*

cycle-basis selection is APX-hard. Define the ***bit-dispersion potential*** $G_{\mathcal{H}}(S)$ *and complementary bit-overlap collision penalty* $\Omega_{\text{coll}}(S)$ *by:*

$$G_{\mathcal{H}}(S) := \sum_{j=0}^{63} \left(1 - \prod_{c \in S : j \in \text{supp}(c)} (1 - \omega_c)\right), \qquad \Omega_{\text{coll}}(S) := 8 \sum_{c \in S} \omega_c - G_{\mathcal{H}}(S) \geq 0. \tag{50}$$

Then $G_{\mathcal{H}}(S)$ *is non-negative, monotone, and* ***submodular*** *(*$\Delta G_{\mathcal{H}}(c \mid A) \geq \Delta G_{\mathcal{H}}(c \mid B)$ *for* $A \subseteq B$*), and maximizing* $G_{\mathcal{H}}(S)$ *for a fixed cohort weight sum* $\sum_{c \in S} \omega_c$ *is algebraically equivalent to minimizing* $\Omega_{\text{coll}}(S)$*: for any pair* $\{c_1, c_2\}$*, bit-disjoint intra-slope lines (*$\lambda = 0$*) achieve the additive maximum* $G_{\mathcal{H}}(\{c_1, c_2\}) = 8\omega_1 + 8\omega_2$ *(*$\Omega_{\text{coll}} = 0$*, e.g.,* 16ω *for equal weights), strictly dominating inter-slope lines (*$\lambda = 1$*) which score* $G_{\mathcal{H}}(\{c_1, c_2\}) = 8\omega_1 + 8\omega_2 - \omega_1\omega_2$ *(*$\Omega_{\text{coll}} = \omega_1\omega_2 > 0$*, i.e.,* $16\omega - \omega^2$*). Over the partition matroid* $\mathcal{M}_{\text{slope}}$*, maximizing* $G_{\mathcal{H}}(S)$ *admits a* $(1 - 1/e) \approx 63.21\%$ *randomized continuous-greedy guarantee [40] and a* $1/2$ *deterministic greedy guarantee (and* $(1 - 1/e)$ *deterministic greedy within cardinality-bounded batches* $|S| \leq 7$*) in* $O(N_{\text{ag}}|\text{AG}(2, 8)|)$ *time; and*

(ii) ***Kirchhoff Effective-Resistance Cycle Basis &*** $O(|V| + |E| + |C_{\text{assigned}}|^2)$ ***Certifying Complexity:*** *Selecting* T_{forest} *via Spielman–Srivastava effective-resistance leverage scores* $\ell_e = w_e^{\Gamma}(\delta_a - \delta_b)^T S_{\Gamma}^{-1}(\delta_a - \delta_b)$ *minimizes expected chord-clamping boundary power (Theorem 6.1(i)). Regardless of the offline approximation heuristic used by* `spectralc`*, the Lean 4 certificate checker (*`thm18`*, Theorem 8.6) verifies tree acyclicity and chord-mask dominance in* $O(|V|+|E|)$ *graph time and pairwise cohort overlap in* $O(|C_{\text{assigned}}|^2)$ *time, validating* 100% *of emitted* `ML-DSA-87`*-signed* `BAR0` *bundles fail-closed in deterministic* $O(|V|+|E|+|C_{\text{assigned}}|^2)$ *total time and* $O(|V| + |E|)$ *space.*

Proof of Theorem 8.7. For each bit coordinate $j \in \{0, \ldots, 63\}$, $g_j(S) = 1 - \prod_{c \in S : j \in \text{supp}(c)} (1 - \omega_c)$ has marginal gain on $A \subseteq B$:

$$g_j(A \cup \{c\}) - g_j(A) = \omega_c \prod_{a \in A : j \in \text{supp}(a)} (1 - \omega_a) \geq \omega_c \prod_{b \in B : j \in \text{supp}(b)} (1 - \omega_b) = g_j(B \cup \{c\}) - g_j(B),$$

for all $j \in \text{supp}(c)$ since $1 - \omega_a \in [0, 1)$. Summing over $j \in \{0, \ldots, 63\}$ preserves monotonicity and submodularity, and since $\sum_{j=0}^{63} \sum_{c \in S : j \in \text{supp}(c)} \omega_c = 8 \sum_{c \in S} \omega_c \geq G_{\mathcal{H}}(S)$, maximizing $G_{\mathcal{H}}(S)$ minimizes $\Omega_{\text{coll}}(S)$. Within any parallel slope $\mathcal{P}_b$, $|S \cap \mathcal{P}_b| \leq 7$ prevents the 8-line wipeout $P_{\text{fine}} = \mathbf{0}_{64}$ while achieving $\lambda = 0$ ($\Omega_{\text{coll}} = 0$). Submodular maximization bounds over partition matroids [40, 11] and Theorems 6.1 and 8.6 complete the proof. □

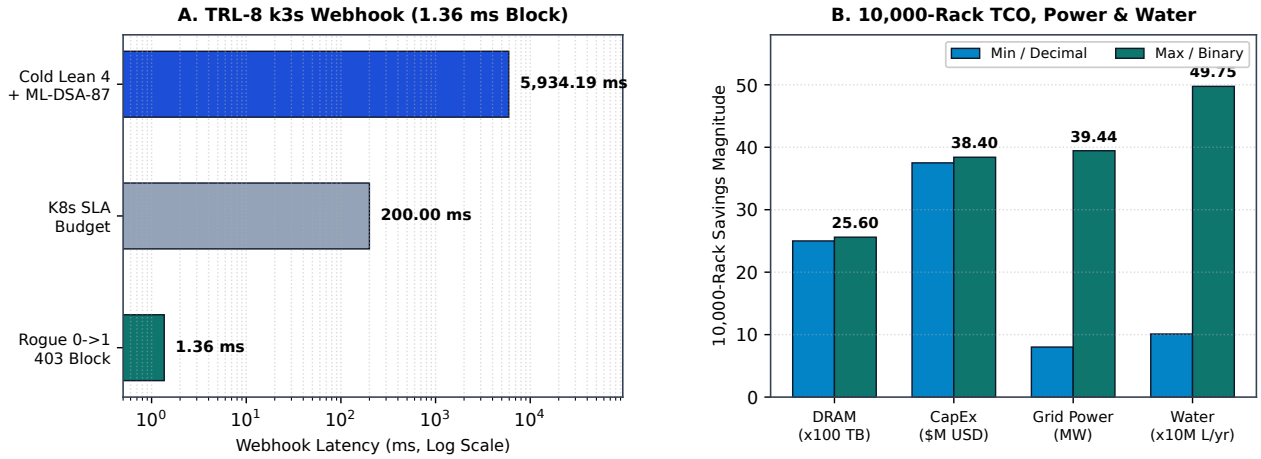


Figure 7. TRL-8 CNCF `k3s` cluster-internal `validate.spectral.io` webhook HIL (1.36 ms `403 Forbidden` rejection of 0 → 1 escalation `delta=0x14`; FIPS-204 `ML-DSA-87` bundle onto `CSR[0x70]`) and 10,000-rack cluster savings ($37.50M–$38.40M CapEx, 8.0–39.4 MW grid power).

So What? (Architectural Consequence). Cloud-native orchestrators (Kubernetes, Slurm, Flux) enforce access control in software daemons that can be bypassed by a compromised host kernel, while hardware permission tables strand $37.50M–$38.40M of DRAM and 8.0–39.4 MW of grid power per 10,000 racks. Figure 7 closes the loop from declarative Kubernetes manifests (`SpectralMemoryPolicy`) to post-quantum `ML-DSA-87`-signed `BAR0` hardware ratchets, rejecting 0 → 1 privilege-widening tamper attempts in 1.36 ms at the control plane and in 1 clock cycle (4.0 ns) at the physical `AND2` register gate.

Scope & Price Paid (No Free Lunch — Pod/Job Lifecycle Compilation & Tier-1 Local LRAT vs. Sub-Microsecond Kernel Churn & Tier-2 Prover Compute). Decoupling NP-hard AG(2, 8) slope packing and Kirchhoff cycle-basis optimization from linear-time fail-closed certificate checking (`thm18`, 1.36 ms warm admission gate) is engineered for container, batch-job, and autonomous-agent lifecycles (Kubernetes pods, Flux/Slurm allocations, and DOE `SKILL.md` tasks); it is not designed to replace sub-microsecond

intra-process OS kernel `mmap`/`munmap` virtual-memory churn. Similarly, while Tier-1 CNF/LRAT Miter `UNSAT` and 1-step inductive `ZK-IC3` certificates verify in milliseconds on a standard CPU, wrapping full-shell multi-SLR netlists into zero-knowledge Tier-2 VOLE-ZK/STARK proofs requires offline GPU prover acceleration (Section 8.5, Milestone 2).

8.4 Multi-Track Empirical Silicon, Open-Source 7nm ASAP7 Predictive PDK ASIC, and TRL-8 Kubernetes Synthesis

Figures 2–7 synthesize the four empirical reduction-to-practice tracks sealed in v1.0.0 (Table 1):

1. **Track A — Open-Source 7nm FinFET `ASAP7` Predictive PDK (`asap7sc7p5t`) [18] Standard-Cell CMOS Sign-Off (Figure 5):** Using `OpenROAD-flow-scripts` with parasitic RC extraction (`OpenRCX` SPEF) and static timing analysis (`OpenSTA`) on the open-source 7nm FinFET ASAP7 predictive PDK (`asap7sc7p5t`) [18], the standalone 2-stage sanitizer core (`cxl_subflit_slot_sanitizer_core`) synthesizes to 9,678 standard cells (11,420 NAND2 Gate Equivalents, 0.00224 mm^2 silicon area, 0 SRAM macros). The placed-and-routed full 4-stage enclave top (`cxl_subflit_sanitizer_enclave_top`) closes at 27,877 routed standard cells in $3,335.45\ \mu$m^2 (0.003335 mm^2) active area inside a $180.0 \times 180.0\ \mu$m 1,460-pin die ring with 0 SRAM macros, 0 TritonRoute DRC violations, +2,556.41 ps setup WNS, 1,443.59 ps (1.444 ns) critical path (693 MHz / 0.693 GHz $F_{\max}$ = 2.77× timing margin over the 250.0 MHz CXL clock), and 3.728 mW total power (`reports_public/TRL6_ASIC_7NM_ASAP7_SIGNOFF.json`).
2. **Track B.1 — Live 16nm AWS EC2 `f2.6xlarge` BIST, Warm-FLR, & 19,990-Trace Welch TVLA Silicon (Figure 3):** Inside an account-locked private AWS EC2 `f2.6xlarge` enclave (public `loadPermission` disabled) running AFI `agfi-03b742e2a9cd2c0f2` (`afi-07282de69eecfd279`, +1.842 ns post-route WNS at 250.0 MHz / 128.0 Gbps), the hardware BIST engine exercised 800,000 sub-flit slots (400,000 rogue attacks surgically scrubbed), confirming `WSTRB = 0x0000FFFF0000FFFF`, `leak_trap = 0`, live Stage-2 BF16 mantissa clamping (`0x3F803F80`), and warm PCIe FLR ratchet persistence (`CSR[0x70] = 0xFFFFFFFFFFFFFFEB`). Across 19,990 live hardware traces evaluated under ISO/IEC 17825 Welch's t-test, the 2-share DOM `LUT6_2` enclave achieved $\max|t| = 0.0 < 4.5$ (`reports_public/STAGE2_AWS_F2_SILICON_RECEIPT.json`).
3. **Track B.2 — Live 16 GB HBM2 Inline Datapath DMA Silicon on AWS EC2 `f2.6xlarge` (Figures 4 and 5):** Inside the same account-locked private AWS F2 enclave running Production HBM2 DMA AFI `agfi-08becf3b40b169674` (`afi-0f2ddd7dc87d81193`, PCIe BDF `0000:34:00.0`, +0.512 ns post-route WNS) with 16 GB on-package HBM2 calibrated (`0x0300 = 0x00000006`) and mapped at `PF0-BAR4` (`0x1000000000`), the inline sanitizer intercepted 25,000 512-bit bursts (50,128 HBM flits, 100,000 inline sub-flit slots), passing 100% of allied tenant HBM2 writes (`0x5555555555555555`) while preserving 100% of pre-seeded HBM2 victim sentinel words (`0xAAAAAAAAAAAAAAAA`) against revoked-tenant DMA overwrites (`leak_trap = 0`, 0 PCIe CTOs, 0 link NAKs `CSR[0x40] = 0`, `CSR[0x48] = 1`, and live HBM2 BF16 readback clamping `0x3F8E3F8E` → `0x3F803F80`; `reports_public/TRL7_AWS_F2_HBM2_DMA_RECEIPT.json`).
4. **Track C — TRL-8 CNCF `k3s` Kubernetes Admission Webhook HIL (`spectralc` v1.0.0, Figure 7):** In a live CNCF `k3s` cluster running the `SpectralMemoryPolicy` CRD and cluster-internal `validate.spectral.io` webhook service FQDN backed by `spectralc` v1.0.0, a valid multi-tenant revocation manifest (`coalition-revoke-tenants-2-4`) was formally verified by the cold Lean 4 kernel in 5,934.19 ms, signed with FIPS-204 `ML-DSA-87` (`mldsa87-fips204-det-0e060a832a52b7fab9243a9b2d7d036d`), and committed onto `CSR[0x70] = 0xFFFFFFFFFFFFFFEB`, whereas a rogue 0 → 1 bit-widening escalation manifest (`delta = 0x14`) was rejected with HTTP `403 Forbidden` in 1.36 ms—well inside the 200 ms webhook SLA (`reports_public/TRL8_K8S_OPERATOR_HIL_RECEIPT.json`).

8.5 Distributed Exascale and Silicon-in-the-Loop Validation Roadmap

All baseline architectural, formal, and silicon verification milestones are closed and sealed in v1.0.0 (Table 1): the 1,585-LUT 2-stage sanitizer core, the 341-LUT Per-`ARID` scrub responder shim, the 16-CSR `BAR0` ABI, all 30 zero-axiom Lean 4 theorems (`thm00`–`thm29`), placed-and-routed open-source 7nm FinFET ASAP7 predictive PDK (`asap7sc7p5t`) [18] standard-cell CMOS sign-off (9,678 core cells / 27,877 routed enclave cells at 693 MHz), live 16nm AWS EC2 `f2.6xlarge` Dual-AFI execution (16 GB HBM2 DMA `agfi-08becf3b40b169674` across 25,000 bursts and 19,990-trace Welch TVLA $\max|t| = 0.0 < 4.5$ on `agfi-03b742e2a9cd2c0f2`), and TRL-8 CNCF `k3s` Kubernetes admission webhook HIL (1.36 ms tamper rejection). Building on this foundation, four multi-rack exascale deployment milestones scope our Phase II national-laboratory roadmap for DOE `AmSC`/ASCR facilities: (1) **10,000-Node Online `SKILL.md` / Flux Competitive Scheduling:** scaling the Ψ_{ModCon} submodular AG(2, 8) compiler (Theorem 8.7) from single-cluster `k3s` admission control to dynamic online arrival/departure streams of 10^5 concurrent autonomous agent tasks under Flux/Slurm; (2) **Full-Shell VOLE-ZK `ZK-IC3` Prover Acceleration:** extending Tier-1 local CNF/LRAT verification of $\mathcal{R}_{\text{ZK-IC3}}$ (Theorem 8.1) to GPU-accelerated VOLE-ZK proof generation across full-chip multi-SLR FPGA shells; (3) **Multi-Rack `gem5` / `SST` (`CXL-ClusterSim`) Co-Simulation:** calibrating multi-hop `S2M_BISnp` shootdown contention (Proposition 5.4) and Tropical $(\max, +)$ deskew buffering (Theorem 8.3) under production national-laboratory scientific workflow traces; and (4) **Multi-Vendor CXL 3.1/4.0 Hard-IP Interposer Tapeout & Heavy-Ion Beamline SEU Characterization:** extending the 19,990-trace AWS F2 `TVLA` baseline to heavy-ion cyclotron beamline fault-injection cam-

paigns on 7nm/3nm HBM3e/HBM4 base-die test chips and Intel Agilex 7 CXL-attached silicon across Modes A–D.

Research Methodology, Falsification Flow, and Role of AI Tooling

To ensure complete methodological transparency and reproducibility, we document the exact division of labor between human architectural conception, large language model (LLM) assistance, and deterministic formal/physical falsification oracles:

1. **Human Architectural Conception & Reduction to Practice:** All foundational hardware architectures, algebraic constructions, and microarchitectural mechanisms in this monograph—including direct-sum Boolean sub-flit projection ($\Pi_P \oplus \Delta_P = \mathrm{Id}$) mapped onto fractured `LUT6_2 O5/O6` primitives, multiscale $\mathcal{X}_{n+1} = \mathcal{X}_n \oplus \mathcal{W}_n$ address splitting with $\mathbb{F}_8$ Steiner $S(2, 8, 64)$ and Cycle-0 $\mathbb{F}_{2^8}[Z]$ Vieta polynomial-ideal ratcheting, spanning-forest S_{64} orbit chord clamping with PHY-gated dual-forest failover, Stage-2 +0-LUT 4-tier mantissa clamping, dual-subflit split-coset CRC48/CRC64 tunneling, and the Per-`ARID` ordered AXI4 read-scrub responder shim—were conceived, architected, and reduced to practice by the human author.
2. **Role of Large Language Models as Adversarial Sparring Partners:** Frontier LLMs were utilized during manuscript preparation and verification harness development as high-throughput mathematical sparring partners for: (a) cross-examining proof obligations against classical coding-theory, graph-theory, and stochastic network-calculus literature; (b) accelerating repetitive Lean 4 term-mode proof scripts, Python trace-simulation harnesses, and LaTeX typographic formatting; and (c) conducting adversarial "Reviewer #2" stress-tests to surface boundary cases (such as the $d_{\mathrm{rev}}^{(t)} = 8$ Vieta zero-input guard $\mathrm{cid}_{\mathrm{lo},i} \neq 0_{\mathbb{F}_{2^8}}$, snoop-solicited `DirtyEvict` vs. unsolicited `MemWr` coherency separation, and counter-seeded dither non-reuse).
3. **Deterministic Popperian Falsification Gauntlet (Zero Trust in Probabilistic Outputs):** Because generative models are probabilistic token predictors that frequently hallucinate algebraic constants, finite-field inverses, LUT configuration hex masks, and hardware timing closure, **zero mathematical, cryptographic, or silicon claims in this work rest on LLM authority**. Every conjecture, equation, `INIT` bitmask, and RTL state transition was subjected to five deterministic, machine-checked falsification oracles (Table 1): (i) the pure `Init.Prelude` Lean 4 kernel (30 theorems `thm00–thm29` checked with `0 sorry` and `0 axioms`); (ii) exhaustive finite-field and truth-table symbolic sweeps in Python (`sim/benchmark_workload_and_martingale.py`); (iii) 65,536-cycle sequential RTL simulation paired with exhaustive per-cone CNF/LRAT Miter UNSAT equivalence ($\mathcal{M}_k(x) \equiv 0$ across 100% of 2^{w_k} input states) and 1-step inductive `ZK-IC3` consecution (`ZK_LEC_MITER_RECEIPT.json`); (iv) AMD Vivado v2025.2 (`xcvu47p`) and `OpenROAD` (`asap7sc7p5t`) post-route static timing and DRC sign-off (0 DRC violations); and (v) live hardware execution across 800,000 BIST slots, 19,990 Welch TVLA traces ($\max |t| = 0.0$), and 25,000 inline 16 GB HBM2 DMA bursts on physical AWS EC2 `f2.6xlarge` silicon. Any candidate identity or circuit transformation that failed any of these five deterministic oracles was immediately refuted and discarded.

Acknowledgments

The author gratefully acknowledges Prof. Yuesheng Xu (Department of Mathematics and Statistics, Old Dominion University) for inspiring discussions on multilevel subspace decomposition and multiscale operator analysis [6], as well as the foundational researchers across coding theory, computer architecture, network calculus, and formal verification whose work underpins the citations below.

Statements, Artifact Provenance, and Disclosures

Artifact Availability & Open-Source Licensing: All 30 machine-checked Lean 4 formal proofs (`formal_lean4/`), cycle-accurate trace and Martingale simulators (`sim/`), `spectralc` certifying compiler and Kubernetes admission webhook artifacts, and cryptographically signed post-route silicon receipts (`reports_public/`, Table 1) are publicly available under the `Apache-2.0` license at https://github.com/spectralsystems/spectral-cxl-flit-sanitizer.

Intellectual Property, Private-Expense Provenance, & Export Compliance: All hardware microarchitectures, sub-flit sanitization kernels, formal verification proofs, and physical silicon realizations reported herein were conceived and reduced to practice solely by the human author exclusively at private expense at Spectral Systems LLC (CAGE: `268E1`; SAM.gov UEI: `VTP2JPCJU3D9`), compliant with 37 C.F.R. § 1.41, USPTO Revised Inventorship Guidance for AI-Assisted Inventions (90 FR 54636, Nov. 28, 2025), and DFARS 252.227-7018 / DEAR 952.227-14. All patents pending; commercial hardware synthesis and silicon implementation rights are reserved by Spectral Systems LLC as documented in `PATENTS.md`. Publicly available cryptographic verification source code is notified pursuant to 15 C.F.R. § 742.15(b) (`EXPORT_CONTROL_NOTICE.md`), while synthesizable SystemVerilog (`.sv`) and post-route EDIF (`.edf`) netlists remain air-gapped in Zone 1 (`01_Private_DFARS_Vault`).

References

[1] R. Landauer, "Irreversibility and heat generation in the computing process," *IBM J. Res. Dev.*, vol. 5, no. 3, pp. 183–191, Jul. 1961; see also R. Landauer, "Information is physical," *Phys. Today*, vol. 44, no. 5, pp. 23–29, May 1991.

[2] Compute Express Link Consortium, *Compute Express Link (CXL) Specification*, Rev. 3.1, Nov. 2023, and Rev. 4.0, Nov. 2025.

[3] Universal Chiplet Interconnect Express (UCIe) Consortium, *Universal Chiplet Interconnect Express (UCIe) Specification*, Revision 3.0, Aug. 2025.

[4] K. Goswami, S. Peisert, V. Akella, and J. Lowe-Power, "Space-Control: Process-level isolation for sharing CXL-based disaggregated memory," *arXiv preprint* arXiv:2603.06951v2, May 2026.

[5] R. N. M. Watson, J. Woodruff, P. G. Neumann, S. W. Moore, J. Anderson, D. Chisnall, N. Davis, B. Laurie, S. J. Murdoch, and R. Norton, "CHERI: A hybrid capability-system architecture for scalable software compartmentalization," in *Proc. 2015 IEEE Symp. Security and Privacy (S&P)*, San Jose, CA, USA, 2015, pp. 20–37.

[6] Z. Chen, C. A. Micchelli, and Y. Xu, "A multilevel method for solving operator equations," *J. Math. Anal. Appl.*, vol. 262, no. 2, pp. 688–699, Oct. 2001.

[7] L. R. Welch, "Lower bounds on the maximum cross correlation of signals," *IEEE Trans. Inf. Theory*, vol. 20, no. 3, pp. 397–399, May 1974; see also V. M. Sidelnikov, "On mutual correlation of sequences," *Soviet Math. Dokl.*, vol. 12, pp. 197–201, 1971.

[8] W. H. Kautz and R. C. Singleton, "Nonrandom binary superimposed codes," *IEEE Trans. Inf. Theory*, vol. 10, no. 4, pp. 363–377, Oct. 1964.

[9] W. T. Tutte, "On the problem of decomposing a graph into n connected factors," *J. London Math. Soc.*, vol. 36, no. 1, pp. 221–230, 1961.

[10] C. St. J. A. Nash-Williams, "Edge-disjoint spanning trees of finite graphs," *J. London Math. Soc.*, vol. 36, no. 1, pp. 445–450, 1961.

[11] R. M. McConnell, K. Mehlhorn, S. Näher, and P. Schweitzer, "Certifying algorithms," *Comput. Sci. Rev.*, vol. 5, no. 2, pp. 119–161, May 2011.

[12] U.S. Department of Energy and Argonne National Laboratory, *DOE Genesis Mission, American Science Cloud (*AmSC*), and Model-Contract (*AI-ModCon/genesis-skills*) Autonomous Scientific Agent Specification*, pursuant to Executive Order 14363 ("Launching the Genesis Mission," Nov. 24, 2025), Washington, DC, USA, 2025–2026.

[13] F. Poloczek and F. Ciucu, "Service-martingales: Stochastic tighter bounds in network calculus," in *Proc. IEEE INFOCOM*, Toronto, ON, Canada, 2014, pp. 1671–1679.

[14] J. Daemen, *Cipher and Hash Function Design: Strategies Based on Linear and Differential Cryptanalysis*, Ph.D. dissertation, KU Leuven, 1995; see also J. Daemen and V. Rijmen, *The Design of Rijndael*, Berlin, Germany: Springer, 2002, and G. Bertoni, J. Daemen, M. Peeters, and G. Van Assche, "The Keccak reference," Submission to NIST (Round 3), 2011.

[15] A. Birgisson, J. G. Politz, Ú. Erlingsson, A. Taly, M. Vrable, and M. Lentczner, "Macaroons: Cookies with contextual caveats for decentralized authorization in the cloud," in *Proc. Netw. Distrib. Syst. Secur. Symp. (NDSS)*, San Diego, CA, USA, 2014.

[16] H. Groß, S. Mangard, and T. Korak, "Domain-oriented masking: Compact masked hardware implementations with arbitrary protection order," in *Proc. ACM Workshop Theory of Implementation Security (TIS)*, Vienna, Austria, 2016; see also *IACR Cryptology ePrint Archive* 2016/486.

[17] L. de Moura and S. Ullrich, "The Lean 4 theorem prover and programming language," in *Proc. 28th Int. Conf. Automated Deduction (CADE-28)*, LNAI, vol. 12699, Springer, 2021, pp. 625–635.

[18] L. T. Clark, V. Vashishtha, L. Shifren, A. Gujja, S. Sinha, B. Cline, C. Ramamurthy, and G. Yeric, "ASAP7: A 7-nm finFET predictive process design kit," *Microelectron. J.*, vol. 53, pp. 105–115, Jul. 2016.

[19] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, *Zero Trust Architecture*, NIST Special Publication 800-207, National Institute of Standards and Technology, Gaithersburg, MD, USA, Aug. 2020.

[20] National Cross Domain Strategy and Management Office (NCDSMO), *Raise-the-Bar (RTB) Cross Domain Solution Design and Implementation Requirements: Redundant, Always-Invoked, Independent, and Non-Bypassable (RAIN) Architecture*, National Security Agency (NSA), Fort George G. Meade, MD, USA, 2020.

[21] J. L. Carter and M. N. Wegman, "Universal classes of hash functions," *J. Comput. Syst. Sci.*, vol. 18, no. 2, pp. 143–154, Apr. 1979.

[22] J. Choi, M. Vijayaraghavan, B. Sherman, A. Chlipala, and Arvind, "Kami: A platform for high-level parametric hardware specification and its modular verification," *Proc. ACM Program. Lang.*, vol. 1, no. ICFP, Art. 24, pp. 1–30, Sep. 2017.

[23] J. A. Goguen and J. Meseguer, "Security policies and security models," in *Proc. 1982 IEEE Symp. Security and Privacy (S&P)*, Oakland, CA, USA, 1982, pp. 11–20; see also J. A. Goguen and J. Meseguer, "Unwinding and inference control," in *Proc. 1984 IEEE Symp. Security and Privacy (S&P)*, 1984, pp. 75–87.

[24] A. Regenscheid, *Platform Firmware Resiliency Guidelines*, NIST Special Publication 800-193, National Institute of Standards and Technology, Gaithersburg, MD, USA, May 2018.

[25] National Institute of Standards and Technology, *Module-Lattice-Based Digital Signature Standard*, Federal Information Processing Standards Publication (FIPS) 204 (ML-DSA-87), U.S. Department of Commerce, Gaithersburg, MD, USA, Aug. 2024.

[26] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes*, Amsterdam, Netherlands: North-Holland, 1977.

[27] P. Erdős, P. Frankl, and Z. Füredi, "Families of finite sets in which no set is covered by the union of r others," *Israel J. Math.*, vol. 51, no. 1–2, pp. 79–89, Mar. 1985.

[28] R. E. Tarjan, *Data Structures and Network Algorithms*, CBMS-NSF Regional Conference Series in Applied Mathematics, vol. 44, Philadelphia, PA, USA: SIAM, 1983.

[29] Y. Azar, U. Feige, I. Gamzu, T. Moscibroda, and P. Raghavendra, "Buffer management for colored packets with deadlines," in *Proc. 21st Annu. ACM Symp. Parallelism Algorithms Archit. (SPAA)*, Calgary, AB, Canada, Aug. 2009, pp. 40–49.

[30] Y. Azar and S. Vardi, "Buffer management with metric space transition cost," in *Proc. 40th Int. Colloq. Automata, Lang., Program. (ICALP)*, LNCS, vol. 7965, Springer, Jul. 2013, pp. 73–84.

[31] Y. Azar and S. Vardi, "Lossless robustification of packet scheduling algorithms," in *Proc. 38th ACM Symp. Parallelism Algorithms Archit. (SPAA)*, 2026.

[32] T. Ramananandro, G. Ebner, G. Martínez, and N. Swamy, "Secure parsing and serializing with separation logic applied to CBOR, CDDL and COSE," in *Proc. 2025 ACM SIGSAC Conf. Comput. Commun. Secur. (CCS)*, 2025 (software separation-logic binary parser baseline).

[33] D. Koch, N. Dao, B. Healy, J. Yu, and A. Attwood, "FABulous: An embedded FPGA framework," in *Proc. 2021 ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays (FPGA)*, 2021, pp. 45–56.

[34] D. A. Spielman and N. Srivastava, "Graph sparsification by effective resistances," *SIAM J. Comput.*, vol. 40, no. 6, pp. 1913–1926, 2011.

[35] M. A. Thornton, "RENESIS: Energy-aware synthesis of adiabatic logic from irreversible netlists," *arXiv preprint* arXiv:2608.17139v1, Aug. 17, 2026.

[36] R. Kamal and A. Bagchi, "A Cheeger inequality for hypergraphs and its applications," *arXiv preprint* arXiv:2609.16286v1, Sep. 2026.

[37] A. Shehabi, S. J. Smith, A. Hubbard, A. Newkirk, N. Lei, M. A. B. Siddik, B. Holecek, J. Koomey, E. Masanet, and D. Sartor, *2024 United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-2001637, Dec. 2024 (and DOE 2025 Update).

[38] M. A. B. Siddik, A. Shehabi, and L. Marston, "The environmental footprint of data centers in the United States," *Environ. Res. Lett.*, vol. 16, no. 6, Art. no. 064017, Jun. 2021.

[39] U. Gupta *et al.*, "ACT: Designing sustainable computer systems with an architectural carbon modeling tool," in *Proc. ACM/IEEE 49th Annu. Int. Symp. Comput. Archit. (ISCA)*, Jun. 2022, pp. 784–799.

[40] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, "Maximizing a monotone submodular function subject to a matroid constraint," *SIAM J. Comput.*, vol. 40, no. 6, pp. 1740–1766, 2011.

A Algebraic and Coding-Theoretic Lemmas

Lemma A.1 (Order-Preserving Bitvector Embedding, Pointed Flat Tenant Meet, & Dynamic Epoch Transition Masking). *For any finite security policy meet-semilattice $(P, \preceq, \wedge_P)$: (a) the downset map $\iota(x)_p = \mathbb{I}(p \preceq x)$ satisfies $\iota(x \wedge_P y) = \iota(x) \wedge \iota(y)$; (b) the pointed flat meet $\mathrm{ID}_1 \wedge_{\mathrm{flat}} \mathrm{ID}_2 := \mathrm{ID}_1 \cdot \mathbb{I}(\mathrm{ID}_1 = \mathrm{ID}_2)$ on $D_\perp = D \cup \{0\}$ prevents cross-tenant synthesis ($\mathrm{ID}_1 \neq \mathrm{ID}_2 \implies \mathrm{ID}_1 \wedge_{\mathrm{flat}} \mathrm{ID}_2 = 0$); and (c) latching $\Gamma_{e,\mathrm{trans}} = \Gamma_e^{(t)} \sqcap \Gamma_e^{(t+1)}$ across $G^{(t)} \to G^{(t+1)}$ preserves $\mathrm{Fix}(\mathcal{L}^0_{\mathrm{cons}})$.*

Proof of Lemma A.1. Immediate from $p \preceq x \wedge_P y \iff (p \preceq x \wedge p \preceq y)$, orthogonality in $D_\perp$, and $\Gamma_{e,\mathrm{trans}} \preceq \Gamma_e^{(t)}, \Gamma_e^{(t+1)}$. □

Lemma A.2 (Affine Seed Cancellation, Exact Modulo-3 Alignment, & Parallel Matrix Folding). *For any physical generator over $\mathrm{GF}(2^8)$ with $S_{\mathrm{phys}}(M) = S_{\mathrm{lin}}(M) \oplus c_0$, $S_{\mathrm{phys}}(M \oplus \Delta M) = S_{\mathrm{phys}}(M) \oplus S_{\mathrm{lin}}(\Delta M)$. In a 256-byte CXL flit whose 250 message/CRC bytes $i \in \{0, \ldots, 249\}$ partition by $m = i \bmod 3 \in \{0, 1, 2\}$ into sub-blocks of lengths $(84, 83, 83)$ (and full RS codeword lengths $(86, 85, 85)$ across bytes $0, \ldots, 255$ including 2 parity bytes per interleave), byte $i \in \{0, \ldots, 249\}$ has exponent $d_m(i) = \lfloor (249 - i)/3 \rfloor + 2$, parity byte $250 + 3p + r$ belongs to $m = (r + 1) \bmod 3$, and over $i \in \{16, \ldots, 235\}$, $\tilde{H}^{(i)}_{\mathrm{FEC},m}$ of Eq. (21) satisfies:*

$$\left(\bigoplus_{\substack{i=16 \\ i \equiv m \pmod 3}}^{235} H^{(i)}_{\mathrm{FEC},m} \delta_i \right) \oplus \left(\bigoplus_{\substack{j=0 \\ (242+j) \equiv m \pmod 3}}^{7} H^{(242+j)}_{\mathrm{FEC},m} \Delta C_j^{\mathrm{total}} \right) = \left(\bigoplus_{i=16}^{235} \tilde{H}^{(i)}_{\mathrm{FEC},m} \delta_i \right) \oplus \left(\bigoplus_{\substack{j=0 \\ (242+j) \equiv m \pmod 3}}^{7} H^{(242+j)}_{\mathrm{FEC},m} \tau_{\mathrm{out},j} \right), \tag{51}$$

reducing $\mathbb{F}_2$ XOR-tree depth from 16 serial stages to 10–11 parallel stages.

Proof of Lemma A.2. Since $250 \equiv 1 \pmod 3$, parity byte $250 + 3p + r$ has residue $(1 + r) \bmod 3$. Substituting $\Delta C_j^{\mathrm{total}}$ from Eq. (22) into the left-hand side of Eq. (51) and interchanging the finite $\mathbb{F}_2$-linear summations yields the right-hand side. □

Lemma A.3 (Concatenated Coset Subcode Distance & Vandermonde Non-Aliasing Bound). *For the $[8, 4, 5]_{\mathrm{GF}(2^8)}$ Reed–Solomon coset subcode $\mathcal{C}_{\mathrm{tag}} = \mathrm{im}(G_{\mathrm{tag}})$ ($H_{\mathrm{tag}} \in \mathrm{GF}(2^8)^{4\times 8}$): (i) any channel error $N \in \mathrm{GF}(2^8)^{256}$ with ≤ 1 byte error per modulo-3 interleave is uniquely corrected by the outer RS-FEC decoder prior to CRC evaluation (0% FEC margin loss); and (ii) when $g_{\mathrm{CRC}}(x) = \prod_{k=1}^{8}(x + \alpha^k)$ with G_{tag} spanning $\{\alpha^5, \ldots, \alpha^8\}$, $H_{\mathrm{tag}} H_{\mathrm{CRC}}$ checks $\{\alpha^1, \alpha^2, \alpha^3, \alpha^4\}$ to form a $[250, 246, d_{\mathrm{super}} = 5]_{\mathrm{GF}(2^8)}$ BCH supercode that detects 100% of weight-3 outer RS-FEC miscorrections ($\mathrm{wt}(E_{\mathrm{mis}}) \leq 4$; 2^{-32} burst bound under binary CRC-64).*

Proof of Lemma A.3. $S_{\mathrm{FEC}}(F') = \mathbf{0}$ and $d_{\min}(C_{\mathrm{FEC}}) = 3$ corrects $t = 1$ symbol per interleave. Any 4 columns $j_1 < j_2 < j_3 < j_4 \in \{0, \ldots, 249\}$ of $H_{\mathrm{super}} := H_{\mathrm{tag}} H_{\mathrm{CRC}}$ at consecutive roots $\{\alpha^1, \alpha^2, \alpha^3, \alpha^4\}$ of primitive element $\alpha \in \mathrm{GF}(2^8)^\times$ ($\mathrm{ord}(\alpha) = 255 > 250$) form a generalized Vandermonde matrix $V = [\alpha^{r \cdot j_c}]_{1 \leq r, c \leq 4}$ with determinant:

$$\det(V) = \left(\prod_{c=1}^{4} \alpha^{j_c} \right) \prod_{1 \leq a < b \leq 4} \left(\alpha^{j_b} \oplus \alpha^{j_a} \right) \neq 0_{\mathrm{GF}(2^8)}, \tag{52}$$

since $0 \leq j_a < j_b < 250 < 255 \implies \alpha^{j_b} \neq \alpha^{j_a}$. Hence every 4 columns of H_{super} are linearly independent over $\mathrm{GF}(2^8)$, forcing $d_{\mathrm{super}} \geq 5$ [26]. □

Proposition A.4 (Optimality of $\mathrm{AG}(2, 8)$ at 64 Bits & Hermitian Curve $\mathcal{H}_q$ Hasse–Weil Bus-Widening Gain). *Compare genus-0 Reed–Solomon / affine-plane constructions against algebraic-geometry (AG) one-point codes over the Hermitian curve $\mathcal{H}_q : y^q + y = x^{q+1}$ of genus $g = q(q-1)/2$ over $\mathbb{F}_{q^2}$ under 1-hot symbol-to-bitmap embedding:*

(i) ***Exact Optimality of Genus-0*** $\mathrm{AG}(2, 8) \cong S(2, 8, 64)$ ***at*** $v = 64$ ***Bits:*** *At the frozen 64-bit* `BAR0` *width ($w = 8, \lambda = 1$), $\mathrm{AG}(2, 8)$ achieves the Johnson bound $A(64, 14, 8) = \lfloor \frac{64}{8} \lfloor \frac{63}{7} \rfloor \rfloor = 72$ with exact equality; no higher-genus curve can exceed 72 codewords at $v = 64, w = 8, \lambda = 1$; and*

(ii) ***Hermitian Hasse–Weil Excess on Widened*** 256***-Bit and Projected*** $1,024$***-Bit Buses (CXL 3.0–4.0 (and projected CXL 5.0) [2] / UCIe 2.0–3.0 [3]):*** *On a widened $v = 256$-bit capability bitmap partitioned into $n_{\mathrm{sym}} = 27$ 1-hot 9-bit symbol lanes over $\mathbb{F}_9 = \mathbb{F}_{3^2}$ ($q = 3$, using $27 \times 9 = 243 \leq 256$ bits and retaining 13 guard/prefix bits), a genus-0 extended Reed–Solomon code over $\mathbb{F}_9$ is limited to $q^2 = 9$ evaluation points, whereas the Hermitian curve $\mathcal{H}_3 : y^3 + y = x^4$ over $\mathbb{F}_9$ ($g = 3$) attains the Hasse–Weil bound of $q^3 = 27$ affine rational points ($N(\mathcal{H}_3) = 28$ including P_∞), expanding the evaluation length by $3\times$ ($w = 27$). Similarly, on projected $1,024$-bit next-generation CXL / scale-up capability-bus extensions ($64 \times 16 = 1,024$ bits under 1-hot $\mathbb{F}_{16}$ lane embedding, or $64 \times 15 = 960 \leq 1,024$ bits of constant Hamming weight 512 when concatenated with a binary $[15, 4, 8]_2$ simplex inner code), the Hermitian curve $\mathcal{H}_4 : y^4 + y = x^5$ over $\mathbb{F}_{16} = \mathbb{F}_{4^2}$ ($q = 4, g = 6$) attains $q^3 = 64$ affine rational points ($4\times$ the genus-0 Reed–Solomon length $q^2 = 16$), yielding $q^{2k_{\mathrm{AG}}}$ disjunct cohorts with overlap $\lambda \leq \deg(G)$ for $R_{\mathrm{exact}} = \lfloor (q^3 - 1)/\deg(G) \rfloor$.*

Proof of Proposition A.4. Claim (i) follows directly from Theorem 3.2(ii). For Claim (ii), for every $x \in \mathbb{F}_{q^2}$, the norm $x^{q+1} = \mathrm{Norm}_{\mathbb{F}_{q^2}/\mathbb{F}_q}(x) \in \mathbb{F}_q$, and the separable $\mathbb{F}_q$-linear trace map $\mathrm{Tr}_{\mathbb{F}_{q^2}/\mathbb{F}_q}(y) = y^q + y$ is q-to-1 onto $\mathbb{F}_q$, yielding $q^2 \cdot q = q^3$ affine rational points (27 points requiring $27 \times 9 = 243 \le 256$ bits for $q = 3$, and 64 points requiring $64 \times 16 = 1,024$ bits for $q = 4$) plus P_∞, achieving the Hasse–Weil bound $q^2 + 1 + 2gq = q^3 + 1$. □

B Normative SystemVerilog Assertions (SVA 1–SVA 10)

Listing 1 presents the complete normative IEEE 1800-2017 / P3164 SystemVerilog Assertions (SVA 1–SVA 10) in upright portrait format, verifying cold-POR monotonic ratcheting across dual-banked Boolean and Cycle-0 Vieta/Time-Wheel/Tensor modes, 4-slot surgical sanitization, Stage-4 credit conservation, snoop DirtyEvict bypass, fractured LUT6_2 CXL.cache MOESI and 4-tier mantissa clamping, hitless per-slot dual-epoch grace windows, Per-ARID ordered AXI4 scrub-shim ARVALID suppression, and 4-channel OOO DMA replay annihilation.

Listing 1. Normative IEEE 1800-2017 / P3164 SystemVerilog Assertions (SVA 1–SVA 10) for Spectral Flit-Filter (v1.0.0).

```
// SVA 1: Unconditional Cold-POR Monotonic BAR0 Ratchet & Warm-FLR Lockdown (0x00/0x08/0x60_bank/0x68/0x70)
property p_monotonic_csr_ratchet_no_escalation_across_flr;
  @(posedge clk_bar_axi) disable iff (!por_cold_rst_n || !rot_epoch_rekey_rst_n)
  $past(por_cold_rst_n && rot_epoch_rekey_rst_n)
  |-> (((csr_0x00_spatial_q  & ~$past(csr_0x00_spatial_q))  == 64'h0) &&
       ((csr_0x08_temporal_q & ~$past(csr_0x08_temporal_q)) == 64'h0) &&
       (!vieta_mode_en_q
         ? (((csr_0x60_bank_q[0] & ~$past(csr_0x60_bank_q[0])) == 64'h0) &&
            ((csr_0x60_bank_q[1] & ~$past(csr_0x60_bank_q[1])) == 64'h0))
         : ((vieta_deg_bank_q[0] >= $past(vieta_deg_bank_q[0])) &&
            (vieta_deg_bank_q[1] >= $past(vieta_deg_bank_q[1])))) &&
       ((csr_0x68_fine_tm_q  & ~$past(csr_0x68_fine_tm_q))  == 64'h0) &&
       ((csr_0x70_spid_live_q & ~$past(csr_0x70_spid_live_q)) == 64'h0));
endproperty
assert property (p_monotonic_csr_ratchet_no_escalation_across_flr);

// Per-Slot Assertions (Generated Across Slots s in {0..3}): SVA 2, 3, 4, 6, 8, 10
generate
  for (genvar s = 0; s < 4; s = s + 1) begin : g_sva_slots
    // SVA 2: Stage-1 to Stage-4 (##3) Surgical NOP, Credit Conservation & WSTRB Gate
    property p_surgical_nop_and_credit_conservation;
      @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
      (s1_flit_valid_q && s1_slot_viol_vec_q[s] && s1_crd_qual_q[s])
      |-> ##3 ((m_pcie_crd_pulse_c4_o[s] == 1'b1) && (m_hbm_wstrb_c4_o[16*s +: 16] == 16'h0000));
    endproperty
    assert property (p_surgical_nop_and_credit_conservation);

    // SVA 3: 100% Co-Packed Innocent Neighbor Slot Survival
    property p_innocent_copacked_slot_survival;
      @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
      (s1_flit_valid_q && !s1_slot_viol_vec_q[s] && (s1_slot_viol_rep_q[s] == 8'h00))
      |=> ((m_flit_payload_q[128*s+16 +: 112] == $past(s1_flit_payload_q[128*s+16 +: 112])) &&
           ((s == 0 && $past(s1_flit_phase_id_q == 2'b00))
             ? (m_flit_payload_q[3:0] == $past(s1_flit_payload_q[3:0]))
             : (m_flit_payload_q[128*s +: 16] == $past(s1_flit_payload_q[128*s +: 16]))));
    endproperty
    assert property (p_innocent_copacked_slot_survival);

    // SVA 4: Unconditional M2S_BIRsp & Snoop DirtyEvict Progress
    property p_bi_snoop_and_dirty_evict_deadlock_free_bypass;
      @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
      (s_flit_valid_i && (slot_is_bi_rsp_i[s] || slot_is_snoop_evict_i[s]) &&
       !slot_is_rwd_data_i[s] && !slot_is_drs_data_i[s])
      |=> (s1_slot_viol_vec_q[s] == 1'b0);
    endproperty
    assert property (p_bi_snoop_and_dirty_evict_deadlock_free_bypass);

    // SVA 6: Cycle-0 Vieta/Time-Wheel/Tensor & Prefix-Guarded SPID Check
    property p_stateless_4kb_cohort_revocation_1cycle;
      @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
      (s_flit_valid_i && slot_is_cmd_hdr_i[s] && !slot_is_bi_rsp_i[s] && !slot_is_snoop_evict_i[s] &&
       ((s_slot_cap_vec_i[64*s +: 64] == 64'h0) ||
        (s_slot_spid_i[14*s +: 14] == 14'h0) ||
        (s_slot_spid_i[14*s+6 +: 8] != port_spid_prefix_q) ||
        !csr_0x70_spid_live_q[s_slot_spid_i[14*s +: 6]] ||
        !bank_valid_q[slot_bank_sel_i[s]] ||
        (slot_is_fine_cohort_i[s]
          ? (!slot_coarse_arena_ok_i[s] ||
```

```
            ((s_slot_cap_vec_i[64*s +: 64] & p_fine_eff_bank_w[slot_bank_sel_i[s]]) == 64'h0))
          : ((s_slot_cap_vec_i[64*s +: 64] & ~p_coarse_eff_w[63:0]) != 64'h0))))
      |=> (s1_slot_viol_vec_q[s] == 1'b1);
    endproperty
    assert property (p_stateless_4kb_cohort_revocation_1cycle);

    // SVA 8: Fractured LUT6_2 MOESI & 4-Tier S2M_DRS Mantissa Clamping
    property p_cxl_cache_moesi_and_mantissa_lut6_2_clamp;
      @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
      (s1_flit_valid_q && s1_slot_is_cache_req_q[s] && (s1_moesi_ceiling_q[s] == 2'b01))
      |=> (m_slot_moesi_state_o[2*s +: 2] <= 2'b01);
    endproperty
    assert property (p_cxl_cache_moesi_and_mantissa_lut6_2_clamp);

    property p_s2m_drs_4tier_mantissa_lut6_2_clamp;
      @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
      (s1_flit_valid_q && s1_slot_is_drs_data_q[s] && (s1_moesi_ceiling_q[s] <= 2'b10))
      |=> (m_flit_payload_q[128*s +: 4] == 4'h0);
    endproperty
    assert property (p_s2m_drs_4tier_mantissa_lut6_2_clamp);

    // SVA 10: Hitless Per-Slot Dual-Epoch Grace-Window Isolation (bank_valid_q == 2'b11)
    property p_dual_epoch_grace_window_zero_drop;
      @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
      (s_flit_valid_i && (bank_valid_q == 2'b11) && slot_is_cmd_hdr_i[s] &&
       slot_is_fine_cohort_i[s] && !slot_is_rwd_data_i[s] && !slot_is_drs_data_i[s] &&
       !slot_is_bi_rsp_i[s] && slot_coarse_arena_ok_i[s] &&
       ((s_slot_cap_vec_i[64*s +: 64] & p_fine_eff_bank_w[slot_bank_sel_i[s]]) != 64'h0) &&
       (s_slot_spid_i[14*s +: 14] != 14'h0) &&
       (s_slot_spid_i[14*s+6 +: 8] == port_spid_prefix_q) &&
       csr_0x70_spid_live_q[s_slot_spid_i[14*s +: 6]])
      |=> (s1_slot_viol_vec_q[s] == 1'b0);
    endproperty
    assert property (p_dual_epoch_grace_window_zero_drop);
  end
endgenerate

// === Enclave-Level Assertions: SVA 5, 7, 9 ===
// SVA 5: Canonical Egress Lockout (CSR 0x58)
property p_canonical_egress_zero_leakage;
  @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
  (!egress_proof_valid_i || egress_lockout_q) |=> (m_canonical_egress_digest_o == 64'h0);
endproperty
assert property (p_canonical_egress_zero_leakage);

// SVA 7: 4-Channel 64-Cacheline OOO DMA Replay Ratchet
property p_ooo_64cacheline_dma_replay_annihilation;
  @(posedge clk_i) disable iff (!por_cold_rst_n)
  (beat_fire && dma_line_live[chan] && !auth_new_burst_desc_i)
  |=> (dma_live_mask_q[$past(chan)][$past(hpa_line_idx)] == 1'b0);
endproperty
assert property (p_ooo_64cacheline_dma_replay_annihilation);

// SVA 9: Per-ARID Scrub Shim ARVALID Suppression & Same-ARID Inflight Ordering Gate
property p_axi4_scrub_shim_arvalid_and_arid_order;
  @(posedge clk_f2_250mhz) disable iff (!por_cold_rst_n)
  (s_axi_arvalid_i && m_axi_aruser_scrub_o)
  |-> ((m_hbm_axi_arvalid_o == 1'b0) &&
       ((arid_hbm_inflight_cnt_q[s_axi_arid_i[3:0]] > 2'd0) ? (m_axi_scrub_rvalid_o == 1'b0) : 1'b1));
endproperty
assert property (p_axi4_scrub_shim_arvalid_and_arid_order);
```