Directory-Based Coherency at Scale: From 8 Cores to 256 Agents

A SignatureIP technical white paper on scaling cache coherency from small clusters to large heterogeneous SoCs

SignatureIP Engineering Team
Center of Excellence for Compute and I/O Subsystem IP

***Abstract*, Cache coherency at small agent counts is solved by broadcast snooping, which works well up to roughly eight coherent agents but does not scale. Snoop traffic grows as the square of agent count, response bandwidth grows linearly, and aggressive broadcast variants reach a practical ceiling around 128 cores. Every production SoC at modern scale is directory-based. The design problem is how to scale the directory itself: sparse and coarse-vector storage formats, direct-transfer features (DMT, DCT, DWT in AMBA CHI Rev E.b), and home-node distribution are the three mechanisms that let the same protocol scale from 8 to 256 agents. SignatureIP's C-NOC is the production-grade implementation, with a UCIe die-to-die interface in the current product for the chiplet path; multi-die coherency extensions are on the roadmap. The portfolio handles the system problems that come with scale: NC-NOC, Proxy Cache, ATC, AXI2CHI, CHI2AXI, PCIe Controller, CXL Controller, CXL Subsystem, AXI4-LLL, Ethernet IP, and the Inoculator topology generator.**

***Index Terms*, AMBA CHI, directory-based coherency, snoop filter, sparse directory, coarse-vector directory, home-node distribution, DMT, DCT, DWT, RN-F, HN-F, UCIe, multi-die, CHI Rev E.b, scaling.**

I. EXECUTIVE SUMMARY

Cache coherency at small agent counts is solved by broadcast snooping. Every coherent request reaches every cache, every cache reports its state, and the protocol resolves the transaction from the response set. The mechanism is simple and works well up to roughly eight coherent agents on a shared bus. Past that point it breaks. Snoop traffic grows as the square of the agent count, response bandwidth at each cache grows linearly, and the energy cost grows with both. Aggressive broadcast variants with probe filtering stretch the practical ceiling to roughly 128 cores on some workloads, but the ceiling exists. Every production SoC at modern scale is directory-based.

The interesting design problem is not whether to use a directory; it is how to scale the directory itself. A naive bit-vector directory that stores one presence bit per agent per cache line is intractable at large agent counts. Sparse and coarse-vector formats cut storage by orders of magnitude. Direct-transfer features (DMT, DCT, DWT in AMBA CHI Rev E.b) remove the home node from the data path on the common case. Home-node distribution turns a single ordering point into many, each owning a slice of the address space and independently sized for its slice. These three mechanisms together are what let the same protocol scale from 8 agents to hundreds.

SignatureIP's C-NOC is the production-grade implementation of this approach. It implements AMBA CHI Rev E.b with directory-based snoop filtering at the home node, DMT/DCT/DWT enabled on the data path, parameterized home-node distribution, and a UCIe die-to-die interface in the current product for the chiplet path. Around C-NOC, the rest of the portfolio handles the system problems that come with high agent counts: NC-NOC for non-coherent traffic, Proxy Cache for I/O coherency without full CHI, ATC for address translation, AXI2CHI and CHI2AXI for migration and interop, PCIe Controller for standardized off-die IO transport, CXL Controller for off-die coherent attach, AXI4-LLL for ultra-low latency C2C connectivity, Inoculator for topology generation, and Ethernet IP as a representative non-coherent master. Multi-die coherency extensions across UCIe links and ISO 26262 functional-safety support are on the SignatureIP roadmap.

II. WHY BROADCAST SNOOPING BREAKS

A broadcast-snoop fabric answers the coherency question by sending every request to every cache. The mechanism is correct at any agent count. What does not survive past a small handful of agents is the cost model.

Snoop traffic per transaction scales linearly with agent count, because every request is delivered to every cache. Aggregate snoop traffic on the fabric therefore scales as the square of the agent count, because every cache also issues requests. The response bandwidth required at each cache scales linearly, because every cache has to answer every transaction. The fabric link bandwidth, the directory storage (if any), and the energy cost all grow with one of these factors.

At eight agents the cost is tractable. At thirty-two it is uncomfortable. At one hundred and twenty-eight, broadcast snooping with aggressive probe filtering still works on some workloads but produces a snoop traffic floor that occupies a meaningful fraction of the coherent fabric's bandwidth budget [1]. Beyond that, the mechanism does not generalize.

As one illustrative measurement, on a representative dual-socket server-class system, memory bandwidth utilization has been reported to saturate near 53 percent of theoretical peak under heavy coherency traffic, against roughly 82 percent without it, a close-to-30-point efficiency gap [2]. Inter-core communication latency on the same system ranges from approximately 36 ns when cores share an L3 segment to approximately 137 ns when they sit in different sockets, a 3.8x topology penalty [2]. The specific numbers are one system's; the direction and magnitude are typical of the regime.

The conclusion is structural. SoCs above modest agent counts have to use a directory.

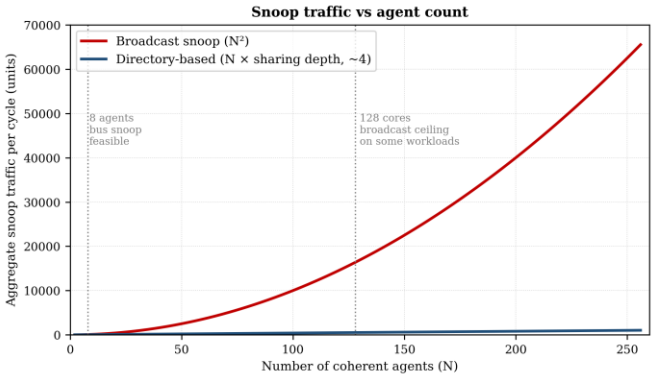

Fig. 1. Snoop-traffic scaling on a broadcast fabric grows as the square of the coherent agent count; a directory fabric grows linearly. The 128-core ceiling on broadcast variants reflects a practical limit reached on some workloads.

III. HOW DIRECTORY-BASED COHERENCY WORKS

A. The Home Node and the Directory

Directory-based coherency enforces consistency at a home node. In AMBA CHI Rev E.b (Issue E), the fully coherent home node is HN-F. The home node owns the directory for a range of cache lines and arbitrates every transaction to that range. On a coherent read, the home node consults the directory, determines which agent (if any) holds the line, issues a snoop to that agent (and only that agent), collects the response, and forwards data to the requester [3].

The directory entry per line carries three pieces: the line state (Invalid, Shared, Modified, Owned, Exclusive, or protocol-specific variants), the sharer list, and any per-line metadata the protocol needs (for example, the ECC tag in safety-grade variants). The state and metadata are small; the sharer list is what scales with agent count.

B. Sparse and Coarse-Vector Directory Formats

A full bit-vector directory uses one bit per agent per cache line. At 256 agents that is 32 bytes per line, plus the state and metadata. Per gigabyte of directory-covered memory, with 64-byte cache lines, the directory itself occupies 500 megabytes of on-die SRAM. That is not tractable.

Sparse directory formats cut storage two ways. Coarse-vector formats group agents into bins (say, eight bins of 32 agents each) and track presence per bin instead of per agent, trading false-positive snoops for storage savings. Pointer-based formats store only the agents that actually hold the line (typically one or two pointers) and fall back to broadcast for the rare lines with wider sharing. Production fabrics use a mix of these strategies, with the directory sized for the expected sharing depth of the workload rather than the worst case [4].

Fig. 2. Directory storage per format for 256 agents at 64-byte cache lines: full bit-vector uses 500 MB per GB; a coarse-vector with 8 bins of 32 uses 16 MB per GB with false-positive snoops; a pointer-based format uses 32 MB per GB with a broadcast fallback for wide sharing.

C. Direct Transfer: DMT, DCT, DWT

The home node is on the control path. It does not have to be on the data path. CHI separates the two so the data path is shortened to the minimum number of hops the topology allows.

Direct Memory Transfer (DMT) applies on a directory miss. The slave node, typically memory, sends data directly to the requester instead of going through the home node first. Direct Cache Transfer (DCT) applies on a directory hit. The home node tells the snooped cache to send the line directly to the requester. Direct Write Transfer (DWT) applies the same principle on the write-back path. In each case the control path keeps coherency state correct while the data path bypasses the home node entirely [5].

These three features together are the single largest latency saving available to a CHI fabric without redesigning the protocol. They are also the largest determinant of how the fabric scales: without them, the home node becomes the data-path bottleneck above modest agent counts.

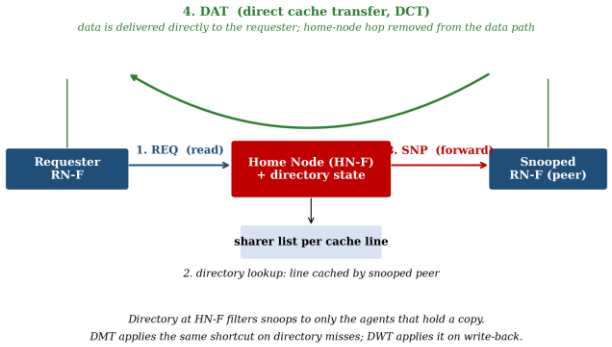

Fig. 3. DMT and DCT keep the home node on the control path while the data path bypasses it. The requester reaches memory (DMT) or a peer cache (DCT) in one hop instead of two.

D. Home-Node Distribution

A single home node serializes all transactions to its address range. Above modest agent counts that becomes the bottleneck. Distributing home nodes across the fabric, with each home node owning a slice of the address space, turns one ordering point into many.

The distribution is constrained by two requirements. The address-map design has to balance directory load across the home nodes so no single home node becomes hot under representative workloads. The snoop-routing logic has to deliver a snoop from any home node to any sharing agent with bounded latency. Production-grade fabrics generally distribute home nodes along the fabric's natural topology (typically a mesh or a hierarchical NoC), with the directory partitioned by address hashing.

IV. SCALING FROM 8 TO 256 AGENTS

The implementation choices change as the agent count rises. The protocol stays the same. What changes is the directory format, the home-node count, the snoop filter accuracy, and the topology that delivers traffic between them.

A. At 8 Agents

Broadcast snooping is feasible. A single home node, if used, is not under bandwidth pressure. Directory storage is modest with any format. The fabric is typically a bus or a small crossbar. The verification surface is tractable. SignatureIP C-NOC at this scale is configured with a single home node and direct-transfer features enabled; the directory format is overkill for the sharing depth but cheap in absolute terms.

B. At 32 Agents

Broadcast snooping is on the edge. Snoop traffic grows by 16x relative to the 8-agent case and starts to compete with data traffic for fabric link bandwidth. Directory-based coherency becomes the better answer. A single home node may still serve, or the SoC may distribute to two or four. Sparse directory formats are now worth the design effort. The fabric is usually a small mesh or a flattened crossbar.

C. At 128 Agents

Broadcast snooping is no longer viable. Aggressive probe filtering can stretch it on specific workloads, but the snoop traffic floor consumes too much of the fabric budget and the verification surface for the broadcast case becomes unbounded [1]. Directory-based coherency is mandatory. Home nodes are distributed; eight to sixteen is typical, partitioned by address hashing. Sparse directory formats with coarse-vector grouping cut storage to a tractable fraction of on-die SRAM. DMT/DCT/DWT keep data off the home-node path on the common case. The fabric is a mesh, sized for the bandwidth between any home-node pair and any RN-F.

D. At 256 Agents

The directory is the dominant design problem. Sparse format choice (coarse-vector versus pointer-based) becomes workload-dependent: workloads with deep sharing favor coarse-vector at the cost of false-positive snoops; workloads with narrow sharing favor pointer-based at the cost of falling back to broadcast on the rare wide-sharing lines. Home-node count is typically sixteen to thirty-two, distributed across the fabric topology. DMT/DCT/DWT are non-negotiable; without them, the home-node hop saturates the link budget. Snoop routing across the fabric has to be bounded under contention, which puts QoS class design in the critical path.

Multi-die topologies become attractive at this scale. A monolithic 256-agent SoC at the reticle limit can be split across two or four dies connected by UCIe links, with each die hosting its own slice of the directory. The protocol does not change. The directory placement and snoop routing do.

SignatureIP Solution, C-NOC. C-NOC implements AMBA CHI Rev E.b with directory-based snoop filtering at HN-F, sparse and coarse-vector directory formats configurable per address range, and DMT, DCT, and DWT enabled on the data path. Home-node count is parameterized; production deployments range from one home node at small scale to multiple distributed home nodes at large scale. The same RTL covers the range from 8 to 256 agents without protocol change; what changes is the configuration generated against the system specification.

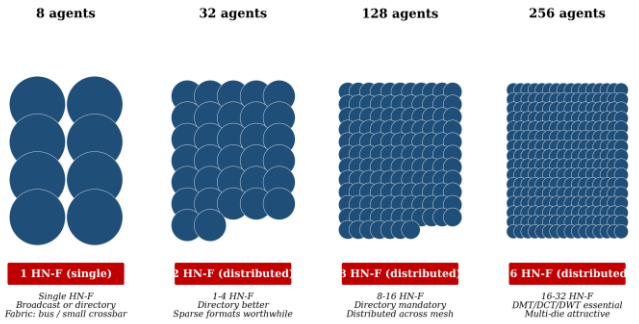

Fig. 4. Representative C-NOC configurations at 8, 32, 128, and 256 agents. Home-node count, directory format, and fabric topology scale with the system specification; the protocol stays constant.

V. DEPLOYED IN PRACTICE

The case for directory-based coherency at scale is not theoretical. Both hyperscaler CPU silicon and production AI accelerators have converged on the architectural pattern. Automotive ADAS silicon adds an ISO 26262 dimension to the same pattern. This section catalogs the published evidence.

A. Hyperscaler CPUs at Scale

The three highest-volume hyperscaler CPU families today all describe directory or directory-equivalent coherency in their published architecture material. The naming differs across vendors; the mechanism is the same.

Ampere's 192-core AmpereOne uses, in the vendor's own product brief language, “a coherent mesh interconnect with 64 distributed home nodes and directory-based snoop filters” [6]. The snoop-filter capacity is sized to cover all 384 megabytes of L2 cache on the chip. The fabric is an 8 by 9 mesh, with certain quad-core clusters co-located with home nodes that carry directory slices and a slice of system-level cache. Past the 128-core mark this is the canonical example of distributed-home-node directory coherency in a commercial Arm server CPU.

Intel's Xeon 6 (Granite Rapids and Sierra Forest) uses the same architectural pattern under different vendor terminology. Each mesh stop is paired with a Caching/Home Agent (CHA) that contains a slice of L3 cache and an inclusive sparse-directory snoop filter. Intel's 6985P-C exposes 120 CHA instances running at 2.2 gigahertz, providing 480 megabytes of total L3 across the chip on a mesh that spans three EMIB-stitched compute dies [7]. The CHA is, in CHI terminology, a fully coherent home node distributed across a mesh.

AMD's 5th-generation EPYC family (Turin, up to 192 Zen 5c cores) uses a centralized rather than distributed directory: a single probe filter resides in the I/O Die and tracks all L3 cache slices across all CCDs. AMD's own 4th- and 5th-Gen EPYC architecture whitepapers state verbatim that “AMD maintains coherency across all L3 cache (and memory) and uses a probe filter to track all the L3 cache slices” [8]. The probe filter is the directory; the placement is different from Ampere and Intel. All three architectures filter snoops against a tracked sharer state rather than broadcasting them.

These three exhibits together establish that at hyperscaler scale the directory is not a research idea or a niche choice. It is the consensus mechanism across both Arm and x86 commercial silicon past roughly one hundred coherent agents.

B. AI Accelerators with Coherent Multi-Die Attach

AI accelerators present a sharper version of the same problem. The coherent agent population now includes CPU cores, GPU compute dies, and accelerator chiplets in the same address space, often across a die boundary.

The cleanest published example of directory-based coherency at multi-die scale is AMD's Instinct MI300A APU, which combines three Zen 4 CCDs, six CDNA3 XCDs, and four I/O dies in a single coherent address space behind 128 gigabytes of HBM3. AMD's CDNA 3 architecture description, summarized in Chips and Cheese's deep dive, names “Coherent Masters” at the XCDs and “Coherent Slaves” with probe filters co-located with Infinity Cache slices at each memory controller, and states that “cache coherency across multiple XCDs is maintained by the Infinity Cache's snoop filter, which efficiently resolves most inter-XCD coherent requests without disturbing the heavily utilized L2 caches” [9]. The pattern is a distributed directory at the memory-controller boundary, across a chiplet topology. MI300A is the compute substrate of the El Capitan exascale system.

NVIDIA's Grace Hopper Superchip is the cleanest “AMBA CHI in production at chip-to-chip granularity” disclosure available. The NVLink-C2C link between Grace CPU and Hopper GPU, in NVIDIA's own architecture blog, “supports the Arm AMBA Coherent Hub Interface (AMBA CHI) architecture... supports cache coherency at the 64-byte granularity, snoop filtering, and different cache models with data forwarding” [10]. The Grace CPU caches GPU-resident memory at cache-line granularity and the GPU caches CPU-resident memory, with hardware-enforced coherency rather than software-managed page migration. This is the architectural picture Section VII (The Multi-Die Frontier) describes, made concrete in a shipping product.

A useful counterpoint to these coherent designs lives in the dataflow-accelerator camp. Cerebras WSE-3 (approximately nine hundred thousand cores) explicitly forgoes caches and a memory hierarchy; inter-PE communication is via tagged wavelet packets routed across a 2D mesh, with each PE holding

48 kilobytes of local SRAM and no global coherency. Google's TPU v5 and v6 use software-managed VMEM scratchpads and explicit DMA scheduled by the XLA compiler, again without inter-core coherency. The lesson is not that coherence is optional, but that the design space splits cleanly: workloads with statically analyzable data flow (large-batch training, fixed-graph inference) can be served by compiler-scheduled message passing; everything else (general-purpose compute, control planes, dynamic graphs, heterogeneous SoCs with mixed traffic) needs a coherent fabric, and at scale that fabric is directory-based.

Even AI-first silicon that is mostly non-coherent maintains a coherent CPU island for control. Tenstorrent's Blackhole exposes up to 140 Tensix tiles on the full die (120 on the shipping p150 card) on a non-coherent NoC and a 16-core SiFive x280 cluster that is coherent within itself, sharing 64 megabytes of L2 and 32 gigabytes of coherent GDDR6 over a dedicated interconnect separate from the Tensix mesh. The architecture acknowledges what every general-purpose SoC has to acknowledge: the control plane needs coherent shared memory, and at meaningful agent counts that means a directory.

C. Automotive ADAS at ASIL-D

The automotive ADAS case raises the bar from "coherent" to "coherent and provably safe." Three shipping or imminent automotive SoCs publish enough detail on both fronts to ground the safety chapter.

Mobileye's EyeQ family of vision SoCs is the strongest single piece of evidence in the corpus for the directory-plus-ASIL combination. Arteris's press release with Mobileye attribution describes the Ncore cache-coherent interconnect as "the on-chip communications backbone of Mobileye's next-generation ISO 26262 ASIL B(D) capable next generation EyeQ system-on-chip devices" and quotes Mobileye: "We chose the Arteris Ncore cache coherent interconnect because of its unique proxy caches and their ability to underpin high-performance, low power, cache coherent clusters" [11]. The architectural mechanism is, per Arteris, a configurable proxy-cache design that "optimizes system bandwidth by eliminating communications to a centralized cache directory", a distributed-directory coherent NoC carrying an explicit ISO 26262 ASIL claim.

Renesas's next-generation R-Car platform (Gen 5, X5H) couples a coherent CPU cluster, Cortex-R52 lockstep real-time cores, and AI accelerators across a fabric that, per Arteris, "supports ISO 26262 ASIL D requirements at the system level, with functional safety support provided through the Arteris FlexNoC resilience option" [12]. The architectural lesson for the white paper is that ASIL-D at the system level is being achieved through a NoC IP with a dedicated resilience package, not as a property of the CPU IP alone.

NVIDIA's Drive Orin is the most-shipped autonomous-driving SoC today and is useful as a primary-source nuance. The Hot Chips 34 architecture deck states verbatim that Orin is "FuSA ASIL-B Chip, ASIL-D Systematic" with the ASIL-D random-fault claim limited to a separate Safety Island built from four Cortex-R52 lockstep pairs delivering up to 10K ASIL-D DMIPS [13]. The trade-press summary that loosely calls Orin "ASIL-D" omits the distinction. The white paper's framing, that ASIL-D today is typically achieved on an isolated lockstep island rather than across the full coherent agent population, is consistent with Orin's published architecture.

The pattern across these three exhibits is that the industry currently partitions the safety case. A coherent NoC carries the perception and planning traffic at ASIL-B(D) or system-level ASIL-D, and a separate lockstep safety island carries the small set of cycles that must be element-level ASIL-D random-fault certified. Extending element-level ASIL-D across the full coherent agent population, rather than around it, is the open problem the next decade of ADAS silicon will work on, and is the target of the functional-safety roadmap described in Section VIII.

VI. HETEROGENEITY INSIDE THE COHERENT DOMAIN

Modern SoCs are not homogeneous CPU clusters. A compute SoC at the upper end carries multiple CPU clusters, GPU compute units, neural accelerators, video and image engines, DMA engines, network controllers, storage controllers, and a long tail of peripherals. Coherency is expected across some subset of them, not all.

CHI defines this as a spectrum. Fully coherent agents (RN-F) hold cacheable copies of memory, are snooped, and respond to snoops. I/O coherent agents (RN-I) can read coherent memory and may snoop into other caches, but they are not themselves snooped, because they do not retain modified data over long horizons. Non-coherent agents bypass the coherent fabric entirely and rely on software to flush caches at the right boundaries.

The right level for each block depends on how it uses memory. CPU clusters and tightly coupled accelerators need full coherency. DMA engines and network controllers benefit from I/O coherency. Boot code, configuration registers, and trace buffers are non-coherent and stay that way. Picking the wrong level wastes bandwidth, area, or correctness, depending on the direction of the error.

At 256 agents, the spectrum is what keeps the directory tractable. If every agent were RN-F, the directory would have to track 256 potential sharers per line. In practice the count of fully coherent agents that actually share working sets is much smaller, and the I/O coherent and non-coherent agents are kept off the directory entirely. Heterogeneity in the coherent domain also raises a subtler failure mode: false sharing, where unrelated variables placed on the same cache line generate spurious coherency traffic and can silently erode performance across CPU clusters and tightly coupled accelerators. Recent work at MICRO 2024 demonstrates that hardware assistance from the coherence protocol itself can detect and repair false sharing on the fly [14], which is easier to build on a directory fabric where sharing state is already tracked than on a broadcast fabric where it is not.

SignatureIP Solution, Proxy Cache. Most non-CPU IP on a real SoC is not CHI native. DMA engines, network controllers, storage controllers, and most accelerator blocks expose AXI4 or AXI4-Stream. Proxy Cache is a CHI agent on the C-NOC side, an AXI4 endpoint on the NC-NOC side, and a small local cache that absorbs repeated reads of shared lines. Non-coherent agents reach coherent memory through Proxy Cache without redesign and without inflating the directory. The agents see a coherent memory view; the directory sees one well-behaved CHI participant.

SignatureIP Solution, NC-NOC. I/O traffic that does not need coherency does not have to ride on the coherent fabric. NC-NOC is the high-performance non-coherent fabric carrying AXI4, AXI3, AXI4-Lite, AXI4-Stream, AHB, and APB traffic on its own network, with configurable QoS, address interleaving, traffic regulation, and Q-Channel and P-Channel power-domain control. Coherent bandwidth on C-NOC stays available for traffic that actually requires coherency.

VII. THE MULTI-DIE FRONTIER

Chiplet-based SoCs raise the agent count beyond the reticle limit of a single die. UCIe (2.0 and later; version 2.0 ratified August 2024, version 3.0 released August 2025 with 48/64 GT/s and 3D packaging support) is the converged die-to-die transport standard [15]. A monolithic 256-agent SoC at the reticle limit can be split across two or four dies, with each die hosting its own slice of the directory and UCIe links carrying coherency traffic between them.

The coherency protocol does not change. What changes is that every snoop, response, and forward that previously crossed an on-die wire now crosses a packaged interconnect with higher latency, lower bandwidth per pin, and a non-zero bit error rate. Directory placement, snoop filtering, and home-node distribution become first-order design decisions rather than implementation details. An architecture that performs well monolithically can lose its margin entirely on the same workload across two dies if these decisions were not made with the chiplet path in mind.

SignatureIP Solution, C-NOC UCIe interface. C-NOC today supports a UCIe interface for die-to-die communication, so the physical and link path for chiplet integration is in the current product, not a future capability. The remaining piece, coherency extensions across the die boundary (directory distribution and snoop routing that turn UCIe links into a coherent multi-die fabric), is on the SignatureIP roadmap as a future release. The architectural foundation in the current monolithic C-NOC, including directory placement and home-node distribution, is being designed against that path. Customers committing to C-NOC today are positioned to extend into chiplet topologies without reworking the coherency protocol and without adding the die-to-die transport, because the transport is already there.

Off-die coherent attach for accelerators and pooled memory follows the same partition. CXL 3.x on the Flex Bus carries CXL.io (PCIe-equivalent), CXL.cache (coherent device caches), and CXL.mem (memory expansion). The CXL Controller bridges a CHI fabric to the Flex Bus, appearing to C-NOC as both RN-F (when the off-die device caches host memory) and SN-F (when host transactions address device-attached memory). At 256 agents this extends the directory's logical reach beyond the package.

SignatureIP Solution, CXL Controller. CXL Controller covers CXL 3.x across CXL.io, CXL.cache, and CXL.mem; it is the bridge between a CHI-coherent fabric and the Flex Bus, appearing to C-NOC as both RN-F and SN-F, with a bias-tracker flush path for Type 2 HDM-DB devices. CXL Subsystem is the pre-integrated drop-in subsystem (sig_cxl_system_top) that packages the CXL Controller with the AXI-SFI bridge, CPI-CHI bridge, ATU, ATC, and Flex Bus PHY interface in one wrapper.

VIII. SAFETY AND COMPLIANCE AT SCALE

The coherent fabric is a safety-critical data path. At small agent counts the safety case is tractable. At 256 agents the coherent fabric has more state, more buffers, and more pipeline stages than a small cluster, and every one of them is a potential propagation point for random hardware faults.

Protection at scale needs SECDED ECC on data buses, parity on address and control buses, end-to-end coverage across pipeline stages, and lockstep on critical state machines. The directory itself is a high-value target for ECC, because a flipped sharer bit can produce wrong-data results that are not caught by the data-path protection. Freedom from interference (FFI) becomes harder to demonstrate because the proof scales with the number of initiator pairs, not just the number of initiators.

For ADAS and autonomous driving systems, the safety case is the design driver. Sensor fusion, perception, planning, and actuation pipelines run across CPU clusters, AI accelerators, and high-bandwidth I/O on the same SoC, with deterministic

worst-case latency budgets measured in microseconds and a safety classification that does not tolerate silent data corruption. Cache coherency in these systems is not a performance optimization; it is the mechanism that lets the perception stack and the control stack agree on what the vehicle just saw, with bounded latency, every cycle. A coherency miss with unbounded latency on a perception-to-actuation path is a safety issue.

Functional safety has to cover both data at rest and data in motion. Data at rest lives in on-chip SRAMs (cache lines, directory entries, buffers, register files) and in off-chip DRAM. Data in motion lives on the buses, bridges, and pipeline stages of the interconnect itself. A safety-grade coherent fabric needs both, with the protection coverage carried end-to-end across every hop from initiator to target. For applications that demand very high reliability or run in harsh environments such as space systems and avionics, DECTED (double-error-correct, triple-error-detect) is replacing SECDED on the protected paths because of the higher rate of multi-bit upsets caused by radiation and aging.

Functional safety across the SignatureIP portfolio is on the SignatureIP roadmap, not in the current product release. C-NOC, NC-NOC, Proxy Cache, ATC, AXI2CHI, CHI2AXI, Ethernet IP, PCIe Controller, CXL Controller, CXL Subsystem, and AXI4-LLL today are production-grade for compute, networking, and storage SoCs, but they do not ship with an ISO 26262 [16] certification package, FMEDA, or formal safety case. ASIL-B through ASIL-D readiness across the portfolio is targeted as a future release, with the protection mechanisms described above being designed into the architecture against that target. Customers evaluating SignatureIP IP for ASIL-classified SoCs should treat functional-safety support as a forward-looking capability and engage with SignatureIP on the certification path.

IX. THE SIGNATUREIP ARCHITECTURE FOR SCALE

The architecture SignatureIP ships is built around directory-based coherency on C-NOC, with the supporting portfolio handling the system problems that come with high agent counts. The same architecture covers the range from 8 to 256 agents, with the configuration changing per system rather than the IP set.

C-NOC is the coherent fabric. It implements AMBA CHI Rev E.b end to end, with directory-based snoop filtering at the home node, sparse and coarse-vector directory formats per address range, DMT/DCT/DWT on the data path, and parameterized home-node distribution. It is the fabric for CPU clusters, GPU compute, and accelerators that need full coherency. The CHI channel separation, the directory at the home node, and the direct-transfer features together let the fabric carry sustained coherent traffic without saturating any single link, which is the failure mode older bus-based and crossbar coherent fabrics reach first under modern workloads. The UCIe die-to-die interface is in the current product; multi-die coherency extensions are on the roadmap.

NC-NOC is the non-coherent fabric. It carries AXI4, AXI3, AXI4-Lite, AXI4-Stream, AHB, and APB traffic at high throughput, with QoS, address interleaving, traffic regulation, performance counters, and Q-Channel and P-Channel power-domain control built in. At 256 agents it carries the traffic that would otherwise saturate the coherent fabric.

Proxy Cache and ATC sit at the boundary. Proxy Cache bridges non-coherent agents into the coherent domain through a CHI-speaking intermediary, absorbing repeated reads and keeping the directory from tracking non-coherent agents directly. ATC provides IOMMU-style address translation for DMA-capable masters on NC-NOC, holding recently used translations close to the requester so DMA latency stays tolerable in virtualized systems.

AXI2CHI and CHI2AXI are the bidirectional protocol bridges. AXI2CHI lets AXI-native initiators reach C-NOC without redesign. CHI2AXI lets a CHI-native fabric reach existing AXI subordinates without forcing them to migrate.

PCIe Controller, CXL Controller, CXL Subsystem, and AXI4-LLL are the off-die extension. PCIe Controller is the standardized non-coherent off-die transport for I/O, storage, and accelerator attach. CXL Controller is the coherent off-die bridge that extends the CHI fabric onto the Flex Bus, appearing to C-NOC as both RN-F and SN-F. CXL Subsystem is the pre-integrated drop-in version. AXI4-LLL is the link-layer IP for photonic and copper serial interconnect, with integrated Reed-Solomon FEC for sub-microsecond CPU-to-remote-memory latency where retransmission overhead would break the latency budget.

SignatureIP Solution, Inoculator. NC-NOC and C-NOC are configurable along enough axes (agent count, traffic profile, protocol mix, address map, QoS class, home-node count, directory format) that hand-configuration at 256 agents is not viable. Inoculator is SignatureIP's cloud-based topology generation tool. It takes a system specification and produces a validated NC-NOC and C-NOC topology configuration. Current scope covers NC-NOC and C-NOC; coverage for the rest of the portfolio is on the Inoculator roadmap and will be added incrementally. Inoculator is a productivity and correctness tool. It is not part of the ECC, parity, or lockstep path, and it is not a functional safety, fault injection, or error-protection block.

Ethernet IP is the representative non-coherent master that exercises the integration end to end. It runs at packet-rate throughput, sits on NC-NOC, reaches coherent memory

through Proxy Cache, and is the worked example of how a non-coherent master participates in system coherency without inflating the directory.

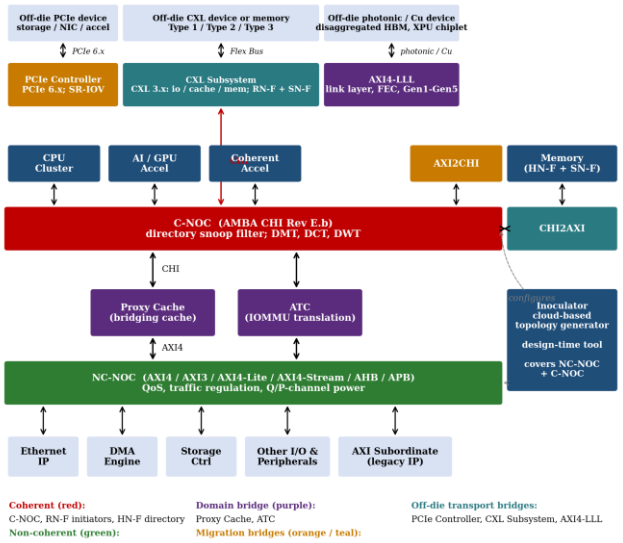

Fig. 5. SignatureIP portfolio in one SoC context: C-NOC as the coherent fabric with a UCIe die-to-die interface; NC-NOC as the non-coherent fabric; Proxy Cache and ATC at the boundary; AXI2CHI and CHI2AXI as the protocol bridges; PCIe Controller, CXL Controller, CXL Subsystem, and AXI4-LLL for off-die transport; Ethernet IP as the worked-example non-coherent master; Inoculator generates the topology.

X. CONCLUSION

Directory-based coherency is the only mechanism that scales from small CPU clusters to large heterogeneous SoCs. Broadcast snooping survives at small agent counts but does not generalize past roughly 128 cores under any variant. The directory is what makes scaling possible: it filters snoops to only the agents that hold a copy, keeps storage tractable through sparse and coarse-vector formats, and distributes home nodes across the fabric so no single ordering point becomes the bottleneck. Direct-transfer features (DMT, DCT, DWT in CHI Rev E.b) remove the home node from the data path on the common case, which is the largest single latency saving available without redesigning the protocol.

The implementation choices change as the agent count rises. At 8 agents broadcast is feasible, the directory is overkill but cheap, and the fabric is small. At 32 the directory becomes the better answer. At 128 it becomes mandatory and home nodes are distributed across the fabric. At 256 the directory format choice (coarse-vector versus pointer-based) becomes workload-dependent, the home-node count climbs into the tens, and DMT/DCT/DWT are non-negotiable. Multi-die topologies become attractive at the upper end of this range and are the next frontier.

For SoC architects evaluating coherent interconnect IP for agent counts in this range, the criteria are concrete: AMBA CHI Rev E.b conformance with DMT, DCT, and DWT enabled; a directory-based home node with sparse and coarse-vector format support; parameterized home-node distribution sized for the agent count; channel separation that lets snoop, response, and data traffic schedule independently; a bridging cache for non-coherent agents that need I/O coherency; address translation close to the DMA masters; bidirectional AXI bridges that preserve ordering; off-die coherent and non-coherent attach paths; a UCIe die-to-die interface for chiplet integration; a topology generator that produces a validated configuration from a system specification; and a roadmap that extends into multi-die coherency without protocol rework.

SignatureIP's portfolio is designed to meet these criteria today: C-NOC for the coherent traffic, NC-NOC for the non-coherent traffic, Proxy Cache and ATC at the boundary, AXI2CHI and CHI2AXI for migration and interop, PCIe Controller for standardized off-die IO transport, CXL Controller for off-die coherent attach, AXI4-LLL for ultra-low latency C2C connectivity, Ethernet IP as the worked-example non-coherent master, and Inoculator as the cloud-based topology generator covering NC-NOC and C-NOC today. The UCIe die-to-die interface on C-NOC is in the current product. Multi-die coherency extensions across UCIe links and full functional-safety support across the portfolio are on the SignatureIP roadmap.

REFERENCES

[1] EECG University of Toronto, snoop-filtering reference. [Online]. Available: https://www.eecg.utoronto.ca/~moshovos/filter/doku.php?id=source-based_snoop_filters

[2] R. K. Somarouthu, "Demystifying Cache Coherency in Modern Multiprocessor Systems," *European Journal of Computer Science and Information Technology*, vol. 13, no. 37, 2025.

[3] Arm Limited, *AMBA CHI Architecture Specification* (IHI0050), latest revision. [Online]. Available: https://developer.arm.com/documentation/ihi0050/latest/

[4] Arteris, *Snoop Filters in Coherent NoC Systems: Types, Trade-Offs, and Directory-Based Design*. [Online]. Available: https://www.arteris.com/learn/snoop-filter/

[5] Arm Limited, *Direct Memory Transfer, Direct Cache Transfer, PrefetchTgt, and Direct Write Transfer*, Arm Developer Documentation (102407). [Online]. Available: https://developer.arm.com/documentation/102407/latest/

[6] Ampere Computing, *AmpereOne Family Product Brief*. [Online]. Available: https://amperecomputing.com/briefs/ampereone-family-product-brief

[7] C. Lam, "A Look into Intel Xeon 6's Memory Subsystem," *Chips and Cheese*, 2024. [Online]. Available: https://chipsandcheese.com/p/a-look-into-intel-xeon-6s-memory

[8] Advanced Micro Devices, *5th Gen AMD EPYC Processor Architecture White Paper*, 2024. [Online]. Available: https://exertisenterprise.com/wp-content/uploads/2024/11/5th-gen-amd-epyc-processor-architecture-white-paper.pdf

[9] C. Lam, "AMD's CDNA 3 Compute Architecture," *Chips and Cheese*, 2024. [Online]. Available: https://chipsandcheese.com/p/amds-cdna-3-compute-architecture

[10] NVIDIA Corporation, "NVIDIA Grace Hopper Superchip Architecture In-Depth," *NVIDIA Developer Blog*, 2023. [Online]. Available: https://developer.nvidia.com/blog/nvidia-grace-hopper-superchip-architecture-in-depth/

[11] Arteris, *Arteris IP Ncore and FlexNoC Interconnects and Resilience Packages Licensed by Mobileye for AI-Powered EyeQ Chips*, press release. [Online]. Available: https://www.arteris.com/press-releases/mobileye-eyeq-arteris-ip-ai-autonomous-driving/

[12] Arteris, *Arteris Network on Chip IP Deployed in Renesas Next-Gen R-Car Automotive Technology*, press release, March 2026. [Online]. Available: https://www.globenewswire.com/news-release/2026/03/24/3261293/0/en/Arteris-Network-on-Chip-IP-Deployed-in-Renesas-Next-Gen-R-Car-Automotive-Technology.html

[13] M. Ditty, *NVIDIA Orin System-on-Chip*, Hot Chips 34, August 2022. [Online]. Available: https://hc34.hotchips.org/assets/program/conference/day2/ADAS%20and%20Grace/HC2022.NVIDIA.Mike_Ditty.v6.pdf

[14] T. Liu et al., "Leveraging Cache Coherence to Detect and Repair False Sharing On-the-fly," in *Proc. IEEE/ACM MICRO*, 2024. [Online]. Available: https://www.cse.iitk.ac.in/users/swarnendu/files/papers/micro24.pdf

[15] UCIe Consortium, *Universal Chiplet Interconnect Express (UCIe) 2.0 and 3.0 Specifications*, August 2024 (2.0) and August 2025 (3.0). [Online]. Available: https://www.uciexpress.org/specifications

[16] ISO, *ISO 26262-1:2018 Road vehicles: Functional safety*. [Online]. Available: https://www.iso.org/standard/68383.html

About SignatureIP

SignatureIP is a provider of semiconductor IP solutions for chiplet and multi-die design across today's advanced SoC architectures. SignatureIP's network-on-chip (NoC) interconnect IP and high-speed interface IP — spanning PCIe Gen 6, CXL 3.0, and coherent and non-coherent NoC — enable system architects to build higher-performance, lower-power designs with faster time to silicon. SignatureIP's IP portfolio allows its customers to focus on product differentiation while SignatureIP handles the complexity of on-chip and die-to-die communication. Learn more at www.signatureip.ai.