Reimagining SerDes for Scalable AI: Architectures, Bottlenecks, and Breakthroughs

Phani Suresh Paladugu

Synopsys

Conf42.com Kube Native 2025

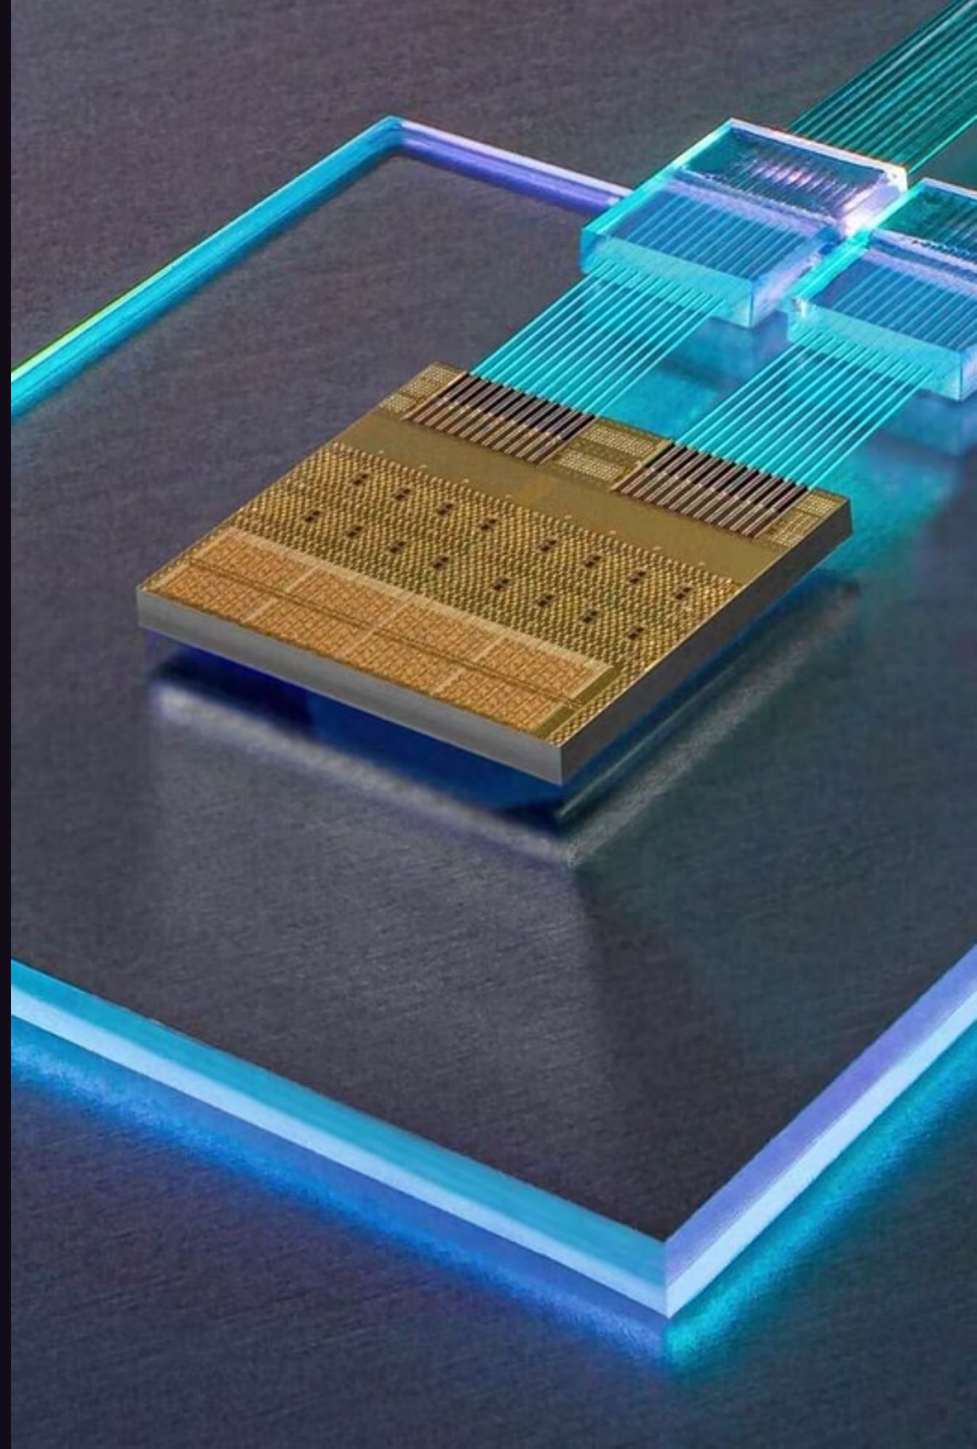

Session Agenda

1

AI Workload Impact on Interconnect Design

How training and inference workloads shape SerDes requirements

2

Signal Integrity at Multi-Gigabit Rates

Challenges at 112G/224G PAM4+ environments

3

SerDes Architecture Trade-offs

Balancing power, throughput, design approach, and specialization

4

Future Directions in SerDes for AI

Emerging Technologies

The AI Compute Explosion

Modern AI training models show exponential growth in compute demands:

- GPT-4: **~1.8 quintillion FLOPs** for training
- Model sizes doubling every **3 to 4 months** since 2018
- Training times rising to **weeks or months** despite parallelization
- Memory bandwidth requirements increasing **20-30% annually**

This growth trajectory places extreme pressure on interconnect, memory technologies to deliver more bandwidth at lower latency and power.

SerDes in the AI Era: Conflicting Demands

Ultra-high Data Rates

112G PAM4 becoming baseline requirements, with 224G PAM4 on horizon

Power Constraints

Target efficiency of < 4-5 pJ/bit while supporting higher frequencies

Signal Integrity Challenges

Increasing losses, reflections, crosstalk in copper channels at multi-GHz

AI Workload Patterns

Bursty traffic, asymmetric bandwidth needs, collective operations

These conflicting requirements create a design paradox that demands innovative approaches beyond traditional SerDes architectures.

AI Workload Impact on Interconnect Design

Training Workloads

- All-to-all communication patterns for parameter distribution
- Sustained high-bandwidth data movement during backpropagation
- High sensitivity to latency spikes that can stall pipeline
- Rising bisection bandwidth demands for model parallelism

Inference Workloads

- Bursty traffic patterns with variable load intensity
- Often memory-bound rather than compute-bound
- Stricter tail latency requirements (e.g., <5ms for real-time applications)
- Asymmetric bandwidth needs (more reads than writes)

These distinct workload characteristics require specialized SerDes design approaches that align with the communication patterns of modern AI systems.

Signal Integrity at Multi-Gigabit Rates

Key Challenges at 112G/224G

- Channel loss exceeding 40-45dB at Nyquist frequency
- Increasing inter-symbol interference (ISI)
- Reflections from impedance discontinuities
- Crosstalk from adjacent channels
- Stringent jitter budgets (< 0.1 UI)
- Amplitude noise sensitivity in PAM4 (3× worse than NRZ)

As we move toward 224G and beyond, traditional equalization techniques reach fundamental limits, requiring novel approaches to maintain signal integrity.

Innovation Frontiers in SerDes Design

DSP-Augmented PAM4

Advanced DSP techniques including FFE (Feed-Forward Equalization) with 7+ taps, DFE (Decision Feedback Equalization) with 15+ taps, and sophisticated CTLE (Continuous Time Linear Equalization) designs enable robust signal recovery.

Adaptive Equalization

Real-time adaptation algorithms that continuously optimize equalization parameters based on channel conditions, utilizing techniques like LMS (Least Mean Squares) and sign-sign algorithms.

Machine Learning for SerDes

ML-based calibration and adaptation that can predict optimal SerDes parameters based on channel characteristics, reducing time-to-lock by 40-60% compared to traditional methods.

Advanced CDR & Coding Techniques

Next-Generation CDR Approaches

- Digital Bang-Bang Phase Detectors with enhanced phase resolution (< 1ps)
- Multi-phase sampling for improved jitter tolerance
- Frequency-domain CDR techniques resistant to SSC (Spread Spectrum Clocking)
- Hybrid analog-digital CDR architectures balancing power and performance

Forward Error Correction

- Reed-Solomon FEC providing 6-7dB coding gain
- Low-Density Parity-Check (LDPC) codes with iterative decoding
- Tailored coding schemes for AI traffic patterns

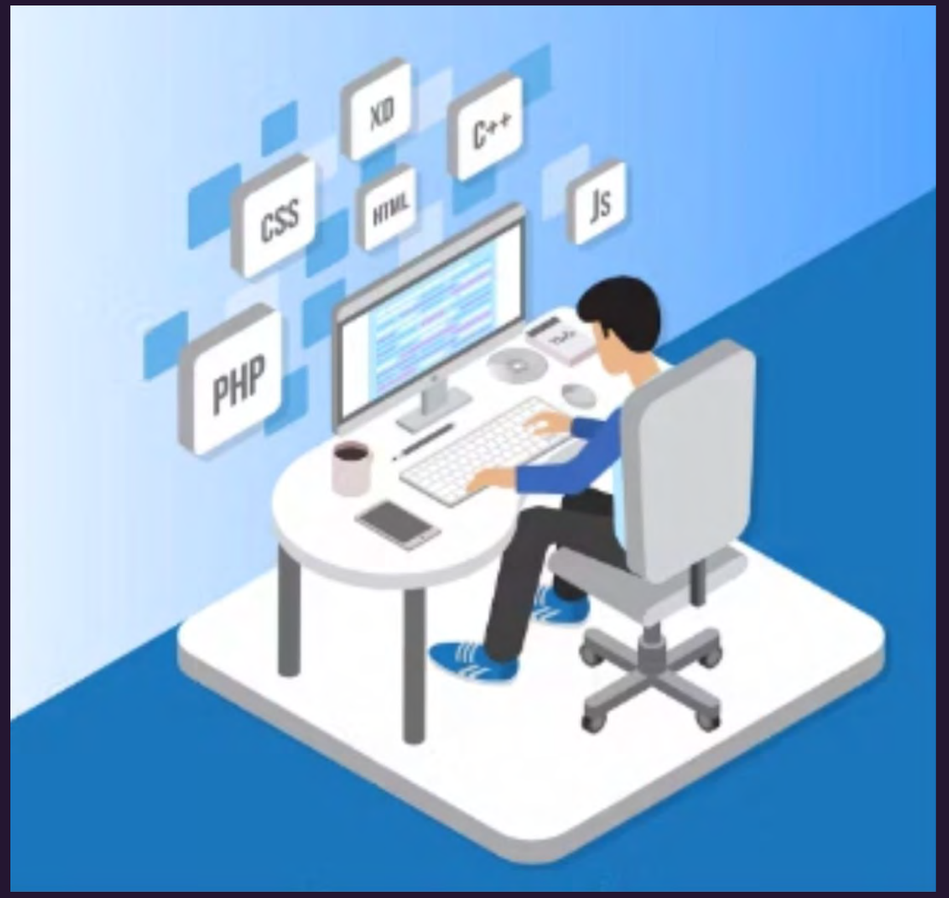

Architectural Trade-offs in SerDes Design

Power vs. Throughput

Modern designs target <4-5 pJ/bit at 112G/224G PAM4 rates. Achieving both requires architectural innovation in:

- Analog front-end simplification
- Process node optimization (5nm/3nm/2nm)
- Adaptive power scaling based on workload

Analog vs. Digital

Increasing shift toward digital-dominant designs:

- 70-80% digital logic in modern SerDes
- Enables process scaling benefits
- Facilitates adaptation and calibration
- Allows more sophisticated signal processing

Reconfigurability

Flexible SerDes architectures supporting:

- Multiple protocols
- Adaptive rate negotiation
- Power/performance operating points
- Forward compatibility with emerging standards

The Power Challenge

Power efficiency is now the primary constraint in SerDes design for AI applications. With power budgets of 5-20W per chip dedicated to I/O, innovations focus on:

Circuit Techniques

- Supply Voltage Scaling (0.7-0.8V)
- Adaptive biasing schemes
- Clock gating on inactive lanes

Architecture Optimizations

- Workload-aware power states
- Power islands to save power
- Simplified analog front-ends

System-Level Approaches

- Link utilization monitoring
- Dynamic frequency scaling
- Thermal-aware Floorplan/Placement

Holistic Solution Strategies

Workload Analysis

Traffic pattern profiling across various AI models identified 70% of communication in collective operations (all-reduce, all-gather).

Signal and Power Integrity

Build end to end models including package, PCB, connector, cable for system level simulations and come up with package design constraints.

Architecture

Implemented heterogeneous SerDes design: Power and Latency optimized Ethernet 112G/224G and PCIe Gen7/6, CXL, UALink Solutions for Scale out and Scale up networks.

Physical Design

Optimized solutions for NS and EW with 4 lane and 8 lane configurations to meet multiple lane requirements for High Performance Computing applications. Optimized bump placements to ease the system design

Future Directions in SerDes for AI

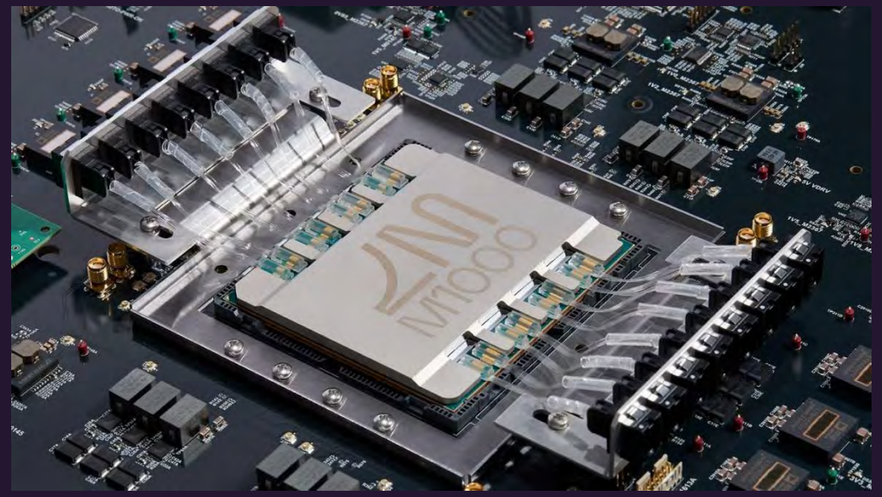

Emerging Technologies

- **448G ethernet solutions:** To Address next generation scale out requirements with optimized performance, power, area
- **Silicon Photonics Integration:** Hybrid electrical-optical SerDes enabling longer reach, higher bandwidth interconnects for AI
- **Co-Packaged Optics:** Moves optics closer to the switch/AI accelerator die to reduce electrical channel loss, eliminating electrical SerDes for long-reach communications
- **SerDes for Advanced Packaging:** High density, low latency SerDes for stacking AI accelerators, HBM and Chiplets in 3D.

These technologies promise 5-8× improvements in bandwidth density while reducing power by 2-3× compared to current electrical SerDes.

Key Takeaways

1

AI workloads fundamentally reshape SerDes requirements

Understanding traffic patterns, bandwidth asymmetry, and collective operations is essential for optimized design.

2

Power efficiency and latency are key design constraints

Future designs must achieve <4 pJ/bit while delivering 224G+ data rates to meet AI scaling demands.

3

Signal integrity requires increasingly sophisticated approaches

DSP-augmented designs, adaptive equalization, and ML-based tuning are becoming standard rather than optional.

4

Heterogeneous SerDes architectures are the future

AI SoCs will incorporate specialized SerDes optimized for different interfaces and traffic patterns rather than one-size-fits-all solutions.

Thank You