

SuperNIC: An FPGA-Based, Cloud Oriented SmartNIC



Presentation by Augustin Scanlon

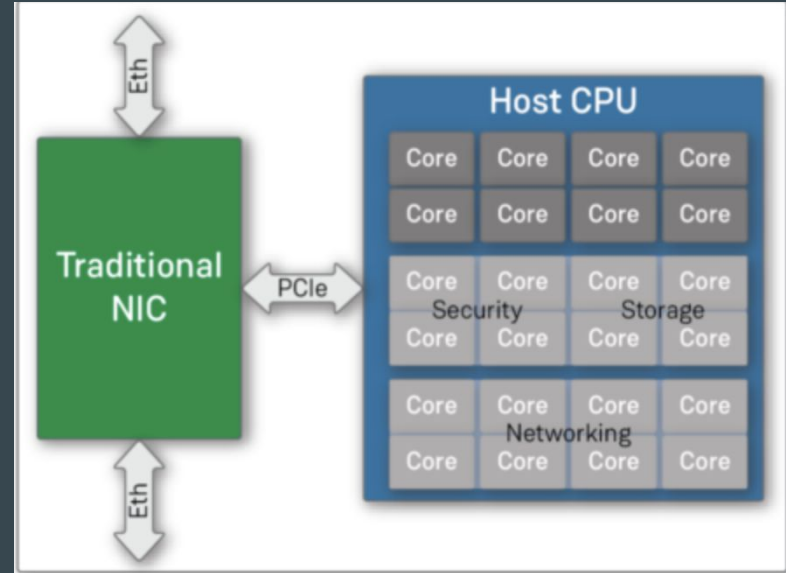
Section 1: The Evolution of Network Interface Cards (NICs) and Data Center Challenges

Introduction to NICs

- Brief context review
- Journey back to 1998
- Challenge: Computers didn't have built-in network functionality
 - Example: Upgrading not feasible due to cost
 - Solution: Buy a NIC from RadioShack, install it, set up drivers, and connect to the internet

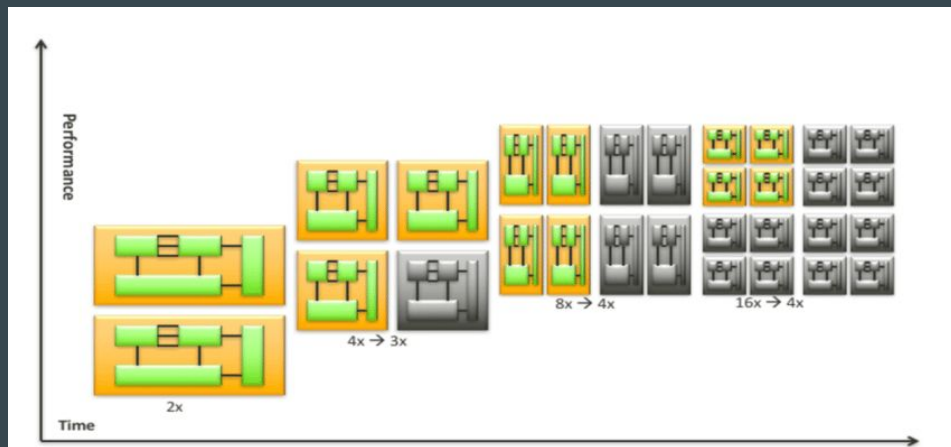
CPU and NIC Interaction (Old Model)

- CPU
 - Responsible for managing the network stack, assembling/disassembling packets
 - Performed bulk of network-related tasks
- NIC's:
 - Focused on physical network connection, data formatting, and error checking
 - CPU handled most of the processing



Inevitable Performance Slowdown

- For a while, this model worked well
- Technological forces:
 - **Moore's Law:** Slowing down of CPU performance improvements
 - **Dennard Scaling:** Insufficient to meet increasing network demands (see Dark Silicon)
- **Impact:**
 - Network speeds increased to 100Gbps, 200Gbps, etc.
 - CPUs couldn't keep up with both applications and network tasks
 - Real-time processing challenges for data centers



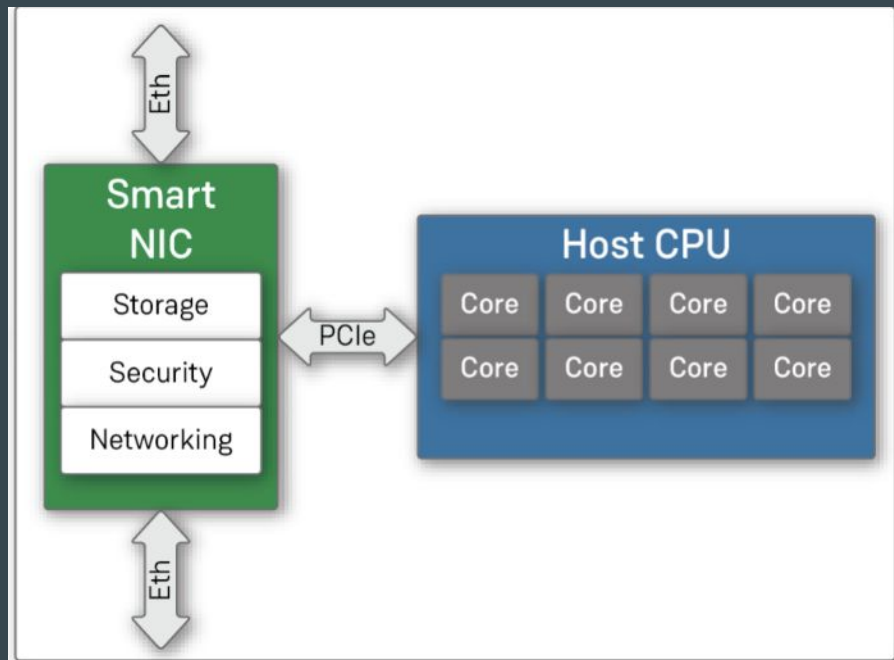
Challenges in Modern Data Centers

- **Multi-Tenancy:**
 - Supporting multiple tenants with varying network tasks in real-time
 - Flexible, scalable solutions required
- **Scalability:**
 - Cloud workloads fluctuate, data centers must scale dynamically
 - Traditional CPUs and older network solutions struggle to meet these demands

Section 2: What are SmartNICs?

Introduction to SmartNICs

- **Solution:** SmartNICs, particularly FPGA-based SmartNICs like SuperNIC
- **Why SmartNICs?**
 - Offload network tasks from the CPU
 - Composed of FPGAs, ASICs, or ARM cores



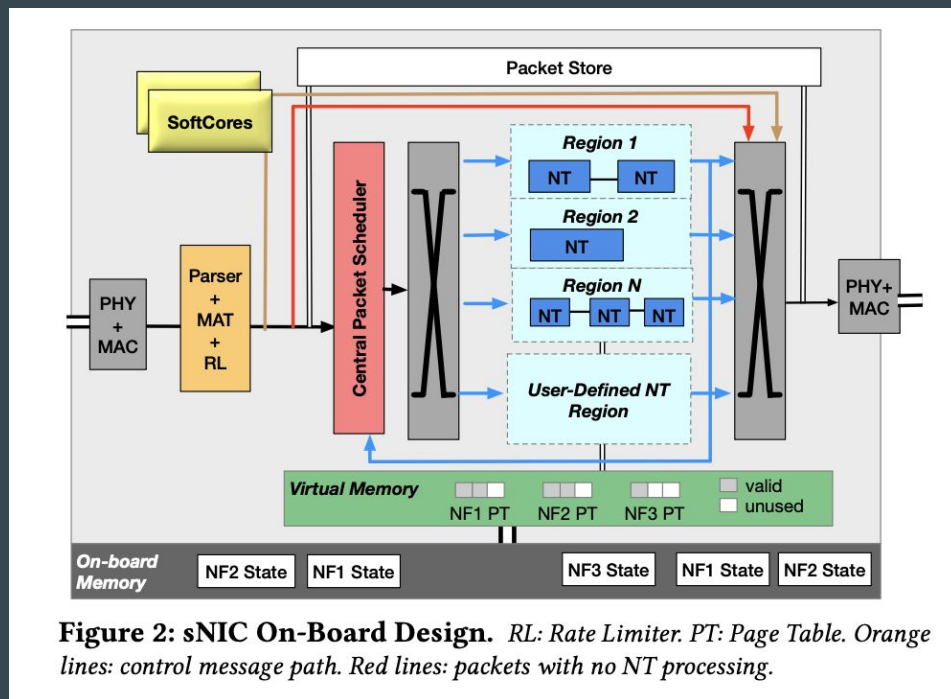
General SmartNIC Function and Benefits

- **Key Benefits of SmartNICs**
 - **Task Offloading:** Packet processing, encryption, and firewall management
 - **Programmability:** Especially when FPGA-based, allows customization of network operations for specific workloads in data centers.
 - **Performance Efficiency:** Reduce latency and increase data throughput.
- **Types of SmartNICs**
 - **ARM-based SmartNICs:** ARM cores for network processing, limitations at high data rates (e.g., 100Gbps).
 - **ASIC-based SmartNICs:** Excellent performance for fixed functions but lack flexibility.
 - **FPGA-based SmartNICs:** Balance programmability and high-speed processing, ideal for dynamic, cloud-based workloads.

Section 3: SuperNIC Architecture and Design

Overview of SuperNIC

- **FPGA-based SmartNIC platform**
 - Key objective: Efficiently offload network functionalities in multi-tenant environments
 - **Multi-tenancy:** Ensures tasks are fairly distributed across multiple users
 - Virtual chains of network tasks mapped onto FPGA regions



Detailed Design of FPGA Regions

- **FPGA Regions:**
 - Programmable hardware for executing network tasks (NTs)
 - Real-time dynamic reconfiguration to adapt to tenant needs
 - Virtual NT chains (sequence of NTs) map onto FPGA regions for processing
 - Reconfiguration for tasks: Regions reallocated based on load and tenant requirements

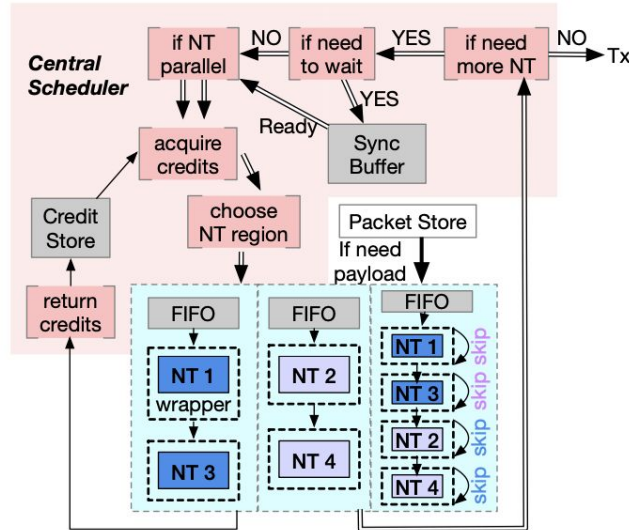


Figure 3: sNIC Packet Scheduler and NT Region Design.

Double arrows, single arrows, and thick arrows represent packet headers, credits, and packet payload.

Task Management with Directed Acyclic Graphs (DAGs)

- **DAGs for organizing tasks:**
 - DAGs arrange tasks in a flowchart-like structure (no loops)
 - **Parallel Processing:** Tasks executed simultaneously to reduce latency
- **DAG Reconfiguration:**
 - Dynamic reconfiguration of DAGs for changing workload demands
 - Task skipping allows SuperNIC to bypass unnecessary steps, optimizing processing




Figure 1: Different Ways of Mapping a DAG to Physical Chains. Each blue box is a region. Dash arrows represent NT skipping.

Schedule and Resource Allocation

- **Packet Scheduling Mechanism:**
 - Central scheduler assigns incoming packets to FPGA regions
 - Scheduler steps back once assignment is made, allowing FPGA to handle tasks
- **Fair Resource Sharing:**
 - Space-sharing and time-sharing balance FPGA resources between tenants
 - **Time-sharing:** Ensures no tenant monopolizes resources

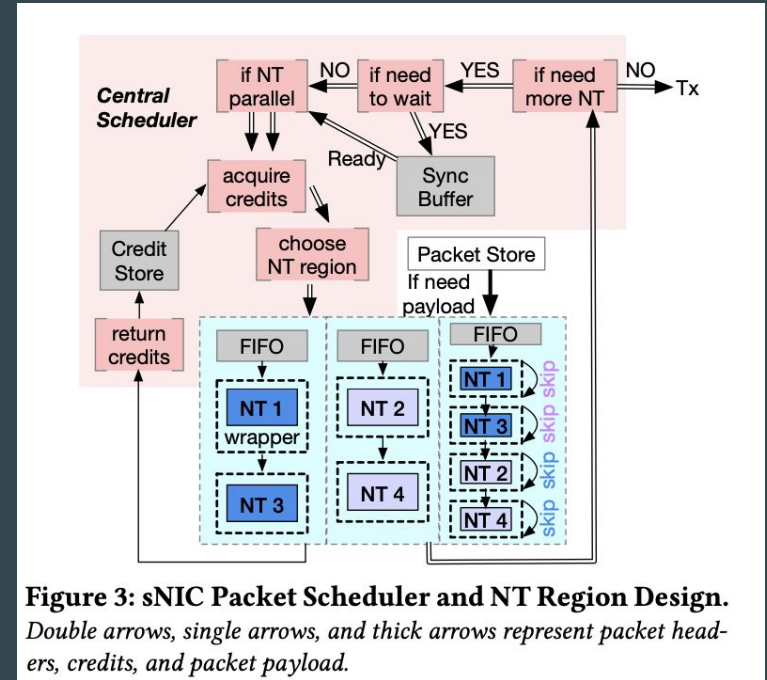


Figure 3: sNIC Packet Scheduler and NT Region Design.
Double arrows, single arrows, and thick arrows represent packet headers, credits, and packet payload.

Task Optimization and NT Skipping

- **NT Skipping:**
 - Skips unnecessary tasks (e.g., encryption) to improve efficiency
 - **Real-time Skipping:** Based on traffic requirements
- **Parallelism:**
 - **DAG Parallelism:** Executes different parts of the DAG simultaneously
 - **Instance Parallelism:** Multiple instances of the same DAG handle packets concurrently
 - **FPGA Optimization:** Efficient space utilization by consolidating tasks into FPGA regions

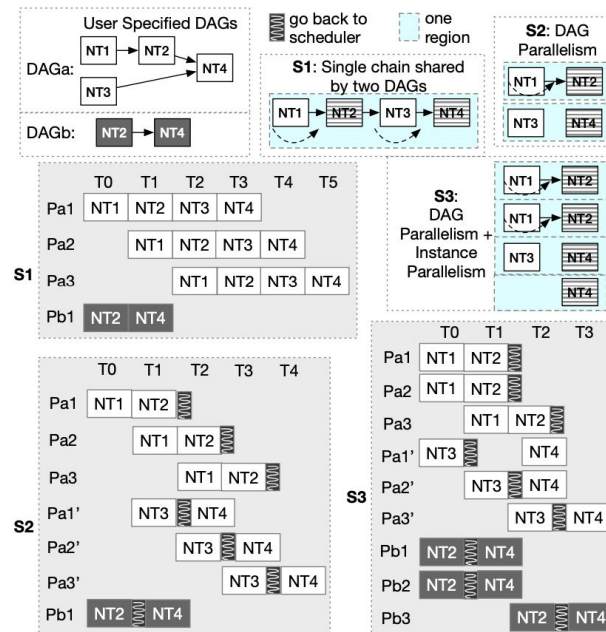


Figure 4: sNIC NT Pipeline. Two deployed DAGs, *a* and *b*. *S1*, *S2*, and *S3* are three ways of executing them. Pa_i/Pb_i refer to the *i*th packet targeting the first/second DAG, P'_i refers to a forked packet. T_i refers to a time unit in the timeline.

Section 4: SuperNIC Performance Evaluation

Benchmark Results

- **Throughput:** Supports up to **100Gbps** with only **196ns** scheduling overhead
- **Latency:**
 - Reduces network task DAG latency by **40%** compared to PANIC.
 - Adds **1.3 microseconds** to packet latency due to third-party PHY/MAC modules
 - Scheduler itself contributes just **196ns** of delay
- **FPGA Utilization:**
 - Improves FPGA resource utilization by up to **3.83x** compared to PANIC.




Figure 7: Throughput with different credits.




Figure 8: Dummy NT Chain Latency.

Comparison with Other SmartNIC Solutions

- **ASIC-based SmartNICs:**
 - Fixed-function design limits flexibility and adaptability.
 - Optimized for specific tasks but lacks programmability and scalability.
- **PANIC:**
 - Higher latency and less efficient compared to SuperNIC.
 - SuperNIC provides better resource scalability and dynamic task management.
 - More efficient due to virtual NT chain mapping.

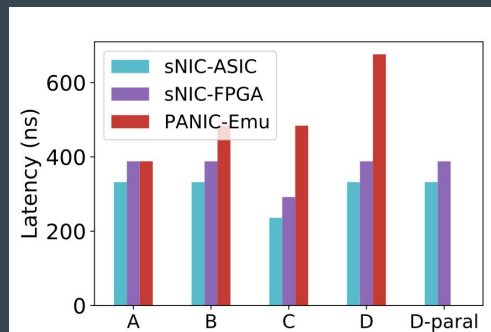


Figure 12: Real NT DAG Latency. *sNIC-ASIC shows projected ASIC performance.*

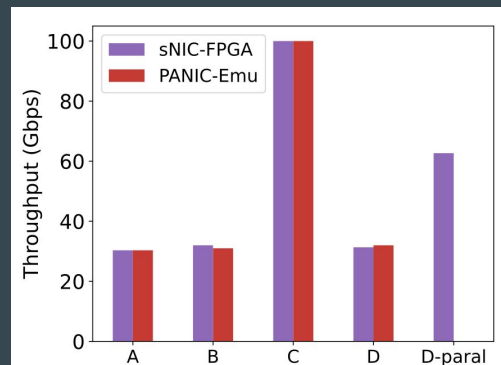


Figure 13: Real NT DAG Throughput.

SuperNIC's Scalability and Future Potential

- **Scalability:**
 - Ideal for dynamic cloud environments due to flexible FPGA reconfiguration.
 - Dynamically allocates resources based on workload.
- **Parallelism:**
 - **DAG Parallelism:** Executes tasks within the DAG simultaneously.
 - **Instance Parallelism:** Runs multiple instances of the same DAG concurrently for greater throughput.
- **Future Improvements:**
 - Potential refinements in scheduling overhead and FPGA resource allocation efficiency.

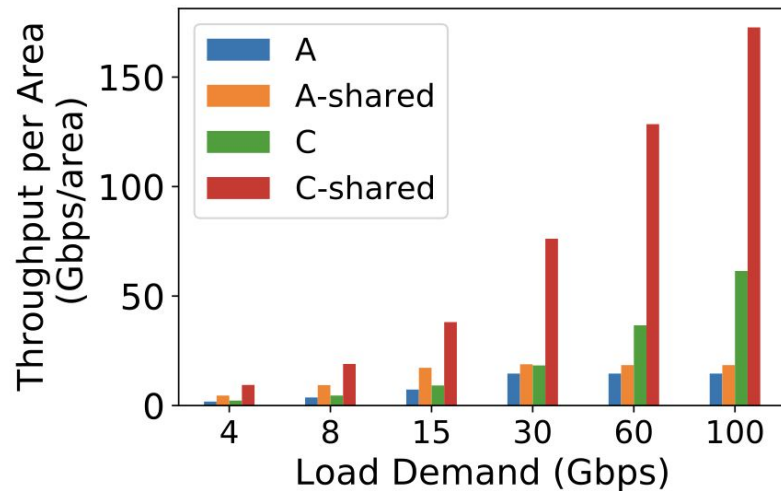


Figure 14: NT sharing *A and C (foreground) sharing chain D (background).*

Works Cited

- Casey. “What Is the Difference Between SmartNIC and NIC?” Fibermall.Com (blog), September 25, 2023. <https://www.fibermall.com/blog/difference-between-smartnic-and-nic.htm>.
- Lin, Will, Yizhou Shan, Ryan Kosta, Arvind Krishnamurthy, and Yiyang Zhang. “SuperNIC: An FPGA-Based, Cloud-Oriented SmartNIC.” In Proceedings of the 2024 ACM/SIGDA International Symposium on Field Programmable Gate Arrays, 130–41. Monterey CA USA: ACM, 2024. <https://doi.org/10.1145/3626202.3637564>.
- Lu, Chien-Ping. “AI, Native Supercomputing and the Revival of Moore’s Law.” APSIPA Transactions on Signal and Information Processing 6, no. 1 (2017). <https://doi.org/10.1017/ATSIP.2017.9>.