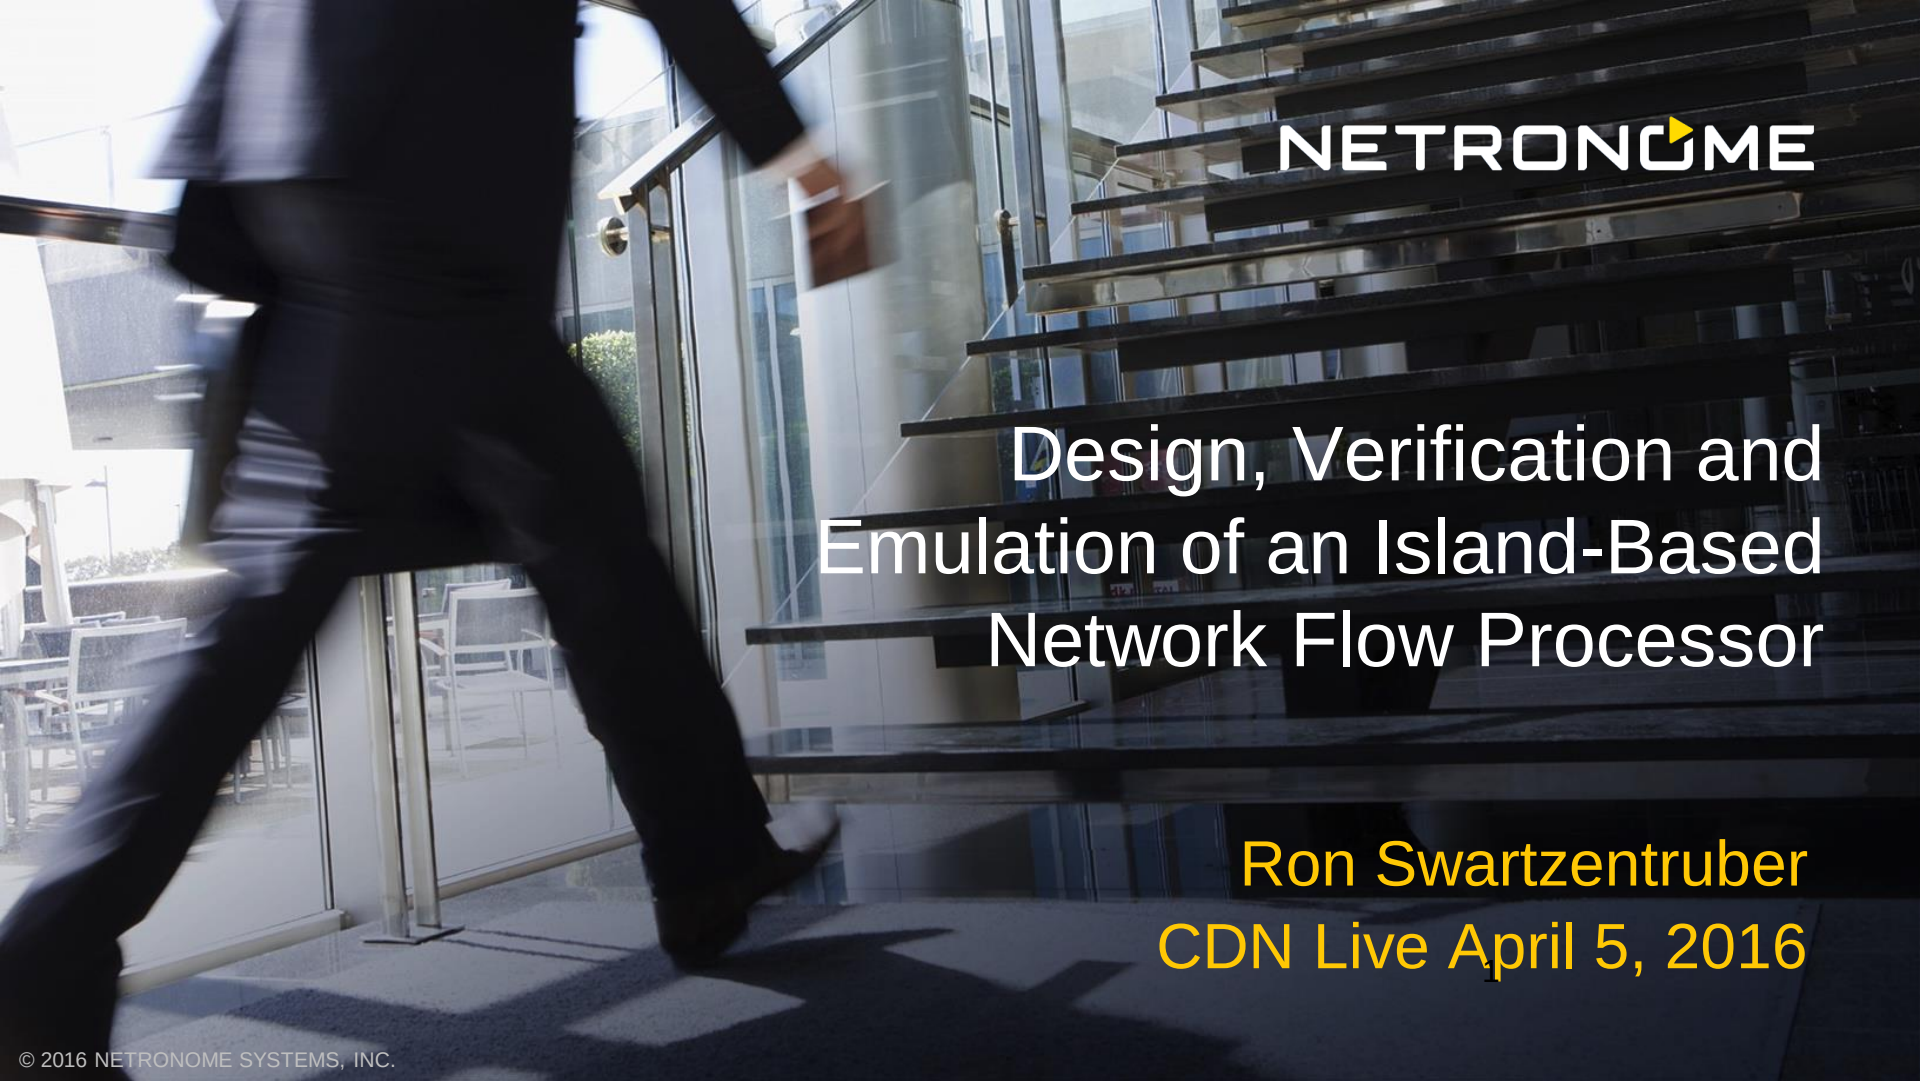

NETRONOME

Design, Verification and Emulation of an Island-Based Network Flow Processor

Ron Swartzentruber

CDN Live April 5, 2016

Problem Statements

NETRONOME

1) Design a large-scale 200Gbps Network Processor containing over 200 processors with multiple high speed I/O and large amounts of internal memory using APR blocks that can be replaced and interchanged across the floorplan
   - Allow for changes in topology later
   - Common building block floorplan saves on design time
2) Verify this very large design using a full-chip environment that can instantiate a few blocks, several blocks, or all of the blocks
   - Speed of the simulation can be determined by how many blocks are instantiated in the testbench
3) Emulate this SoC design to find potential bottlenecks and guarantee system performance
   - Enable Software applications to run pre-silicon and prove out design

Building an NFP Island

NETRONOME

- APR Blocks with a Common Footprint
- External Signal Interconnect with Fixed Pin Locations
  - Fabric Ports
  - Register Interface
  - Interrupts and Events
  - JTAG, Clocks, DFT
- Re-use I/O Timing constraints 50/50 budget
- Common Test Logic, complete for DFM

- Identical Overlay Mesh for Internal Signals and Busses for all Islands

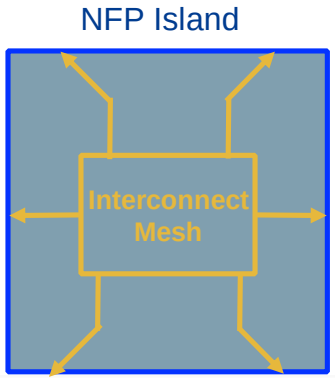

Island Block Topology

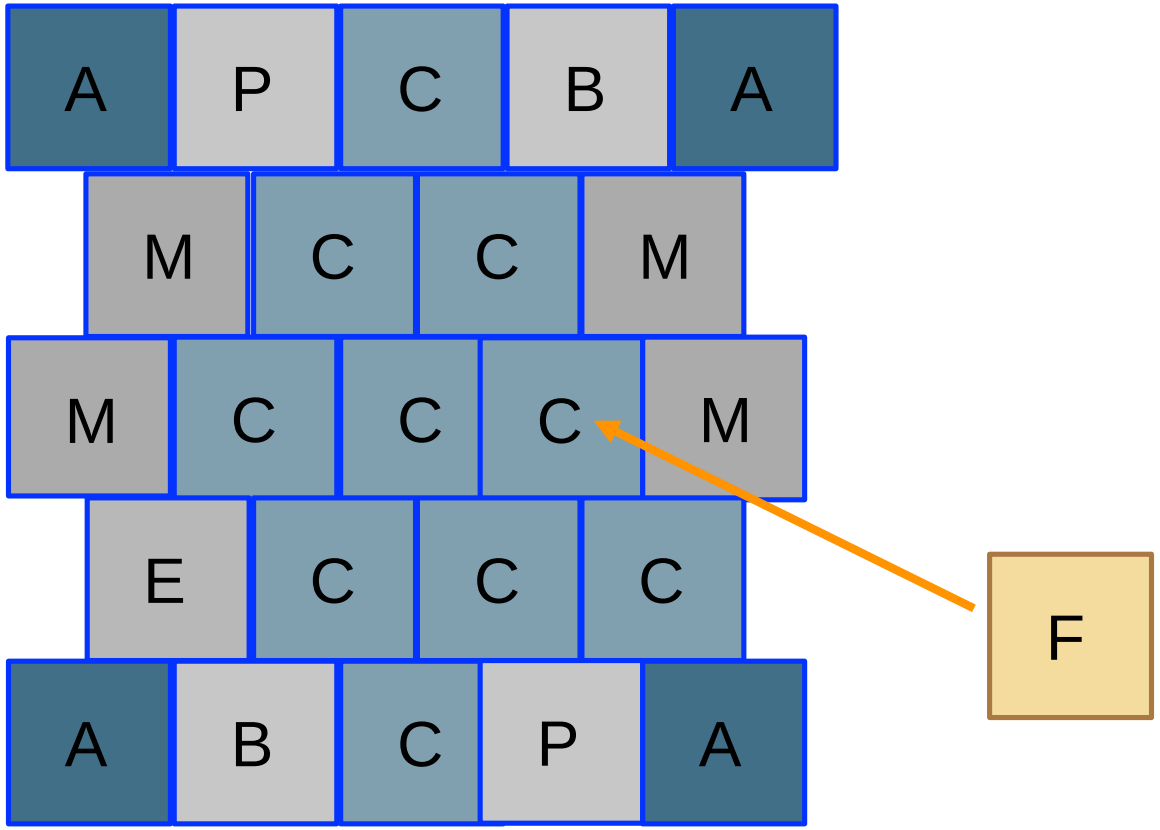

- Innovative Heterogeneous Island Architecture
- Identical Overlay Mesh
- APR Blocks Connected by Abutment
- Latency Tolerant Processing Architecture with 8 threads
- Blocks can be Easily Interchanged; Replaced and Repeated

Full Chip Test Environment

NETRONOME

- Verify the SoC using a full-chip test environment that can instantiate a few blocks, several blocks, or all of them
- Common Verify this very large design using a full-chip environment that can instantiate a few blocks, several blocks, or all of the blocks
- Speed of the simulation can be determined by how many blocks are instantiated in the test bench
- Test bench created by combining the I/O's of interest with multiple internal islands
- Python scripts create Verilog top-level module and test bench comprised of multiple blocks and common interfaces
- UVCs instantiated based on the I/Os of interest

Full Chip Verification

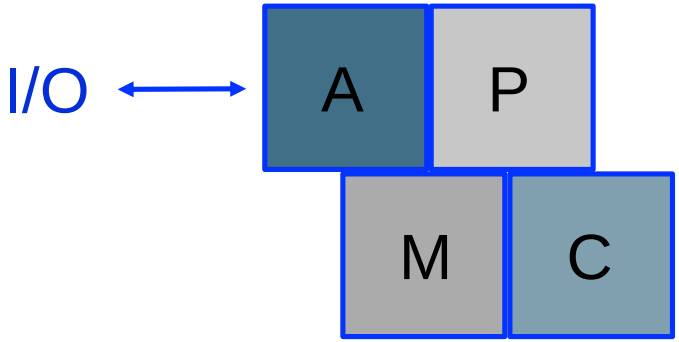

Small, Fast

10/25GbE test bench

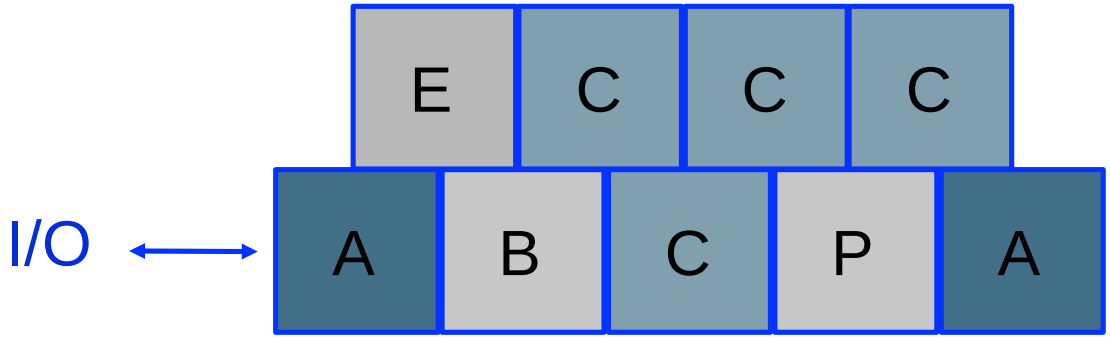

Larger,
100GbE test bench

Emulation Environment

NETRONOME

- Goal: Run full chip System Verilog simulations at 20x speed
- Run real Software applications to validate performance and find potential bottlenecks
- Test many thousands of packets in a fraction of the time as compared to simulation. Goal: 9,000x w/ speed-bridge(s)
- Create make/run environment that allows any SW engineer to test NFP application code pre-silicon
- Incorporate Cadence Palladium supported I/Os, Speedbridge, BFMs and Packet Generator
- Treat DUT as a “NIC” and connect to a VM via PCIe Speedbridge
- Software load to “NIC” via external PCIe interface

Emulation Overview

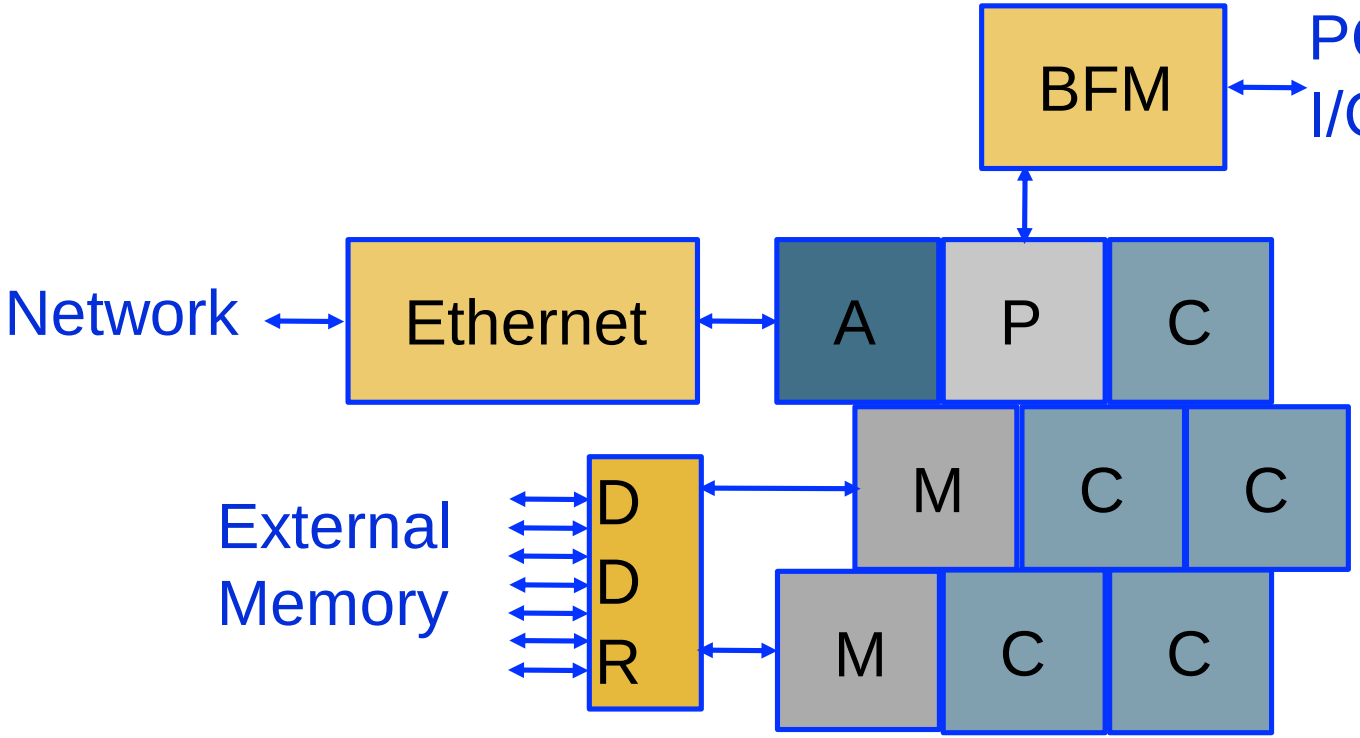

Host PCIe to Network with External Memory Testing

For use on Intelligent Server Adapters

- Six ports 40GbE (or 24x10GbE)
- 2x100GbE support
- Four PCIe Gen3 x8

Comprehensive features with LNOD 2.1

- RX/TX with SR-IOV and stateless offloads
- Extensive, flexible tunneling support (e.g. VXLAN, GRE)
- Transparent offload of OVS datapath
- Stateful flow tracking
- Stateful Load Balancing

Firmware Data Plane (blue)

External DDR3 for deeper flow tables

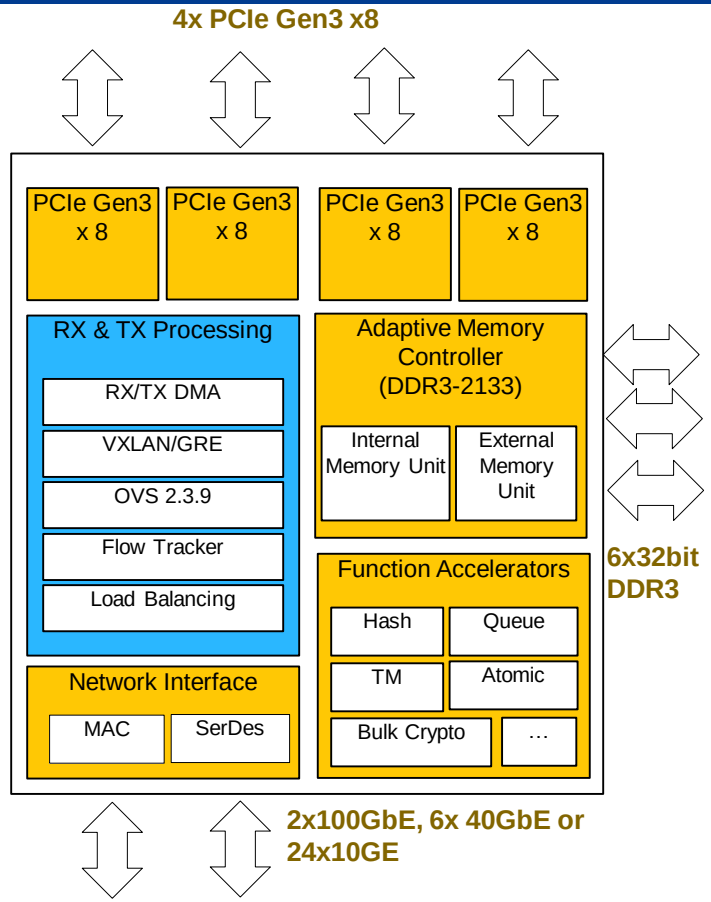

Reference Patents

NETRONOME

- US Patent Application No. 13/399,433: Staggered Island Structure in an Island-based Network Flow Processor
- US Patent Application No. 13/399,888: Island-based Network Flow Processor Integrated Circuit
- US Patent Application No. 13/399,958: Processing Resource Management in an Island-based Network Flow Processor