Programmable Network Processors -- an opportunity for proactive networking

Alan Crouch
Director, Network Platforms
Intel Labs
alan.crouch@intel.com
May 13, 2002

NOSSDAV  **2002**

The 12th International Workshop on Network and Operating Systems Support for Digital Audio and Video

OUTLINE

- The Environment
- What is a network processor?
- Intel's Vision and Contribution
- Technical Challenges
- Summary

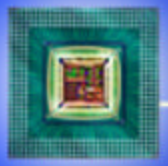

Communications Industry Transition

- **Vertically integrated communication vendors moving to building block approach**
  - **Programmable switching platform**
  - **Control/Service logic moves to standard computing platforms**
- **Implications**
  - **Multiple, independent services on same switching platform**
  - **Allows rapid, dynamic innovation in services**
  - **Customer, service or VPN-specific resource guarantees possible**
  - **Example: SoftSwitch Initiative**

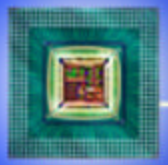

Lesson from the past: Computer Industry Transition

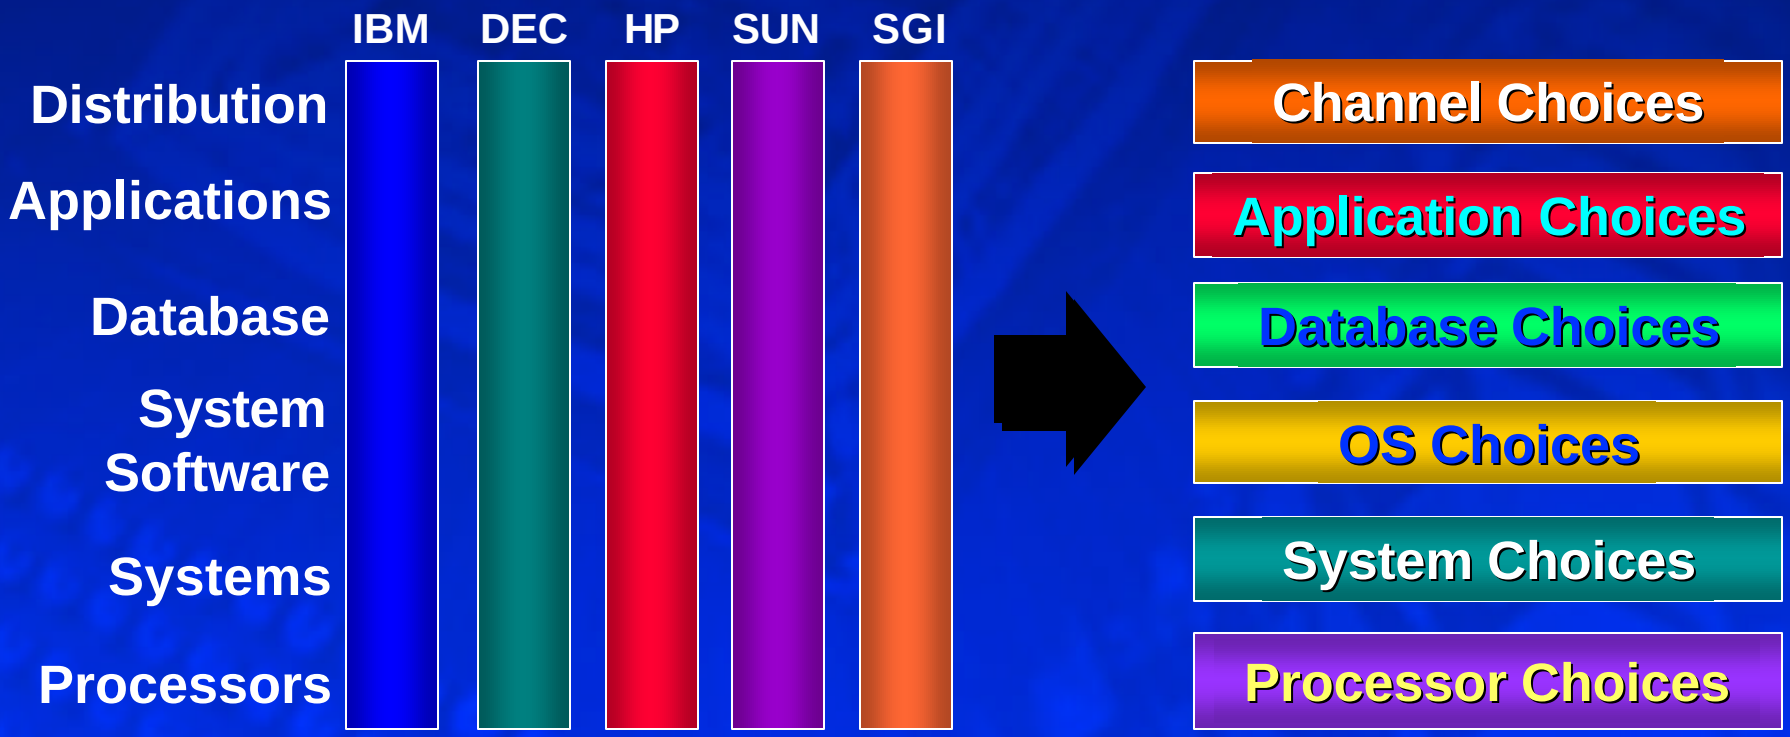

"Old Computer Industry"

"New Computer Industry"

Business Value Proposition for Systems Vendors

- ASICs giving way to programmable Silicon
  - Time-to-Market, Price/performance, Flexibility reasons
  - Efficient suppliers: IBM, AMCC, Intel, Motorola, …
- Value moving to services and applications
  - L3/L4 switches already a commodity
  - End2End services require integration across multiple technologies and infrastructures
  - Service providers want quick ways of deploying new services
- New services via middleware infrastructure
  - E.g., policy, security infrastructure, directories, JAVA*/JINI*, CORBA*/DCOM*, …

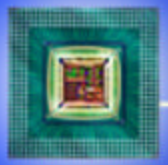

*Other names and brands may be claimed as the property of others

Where Network Services Live or “edge” is everywhere

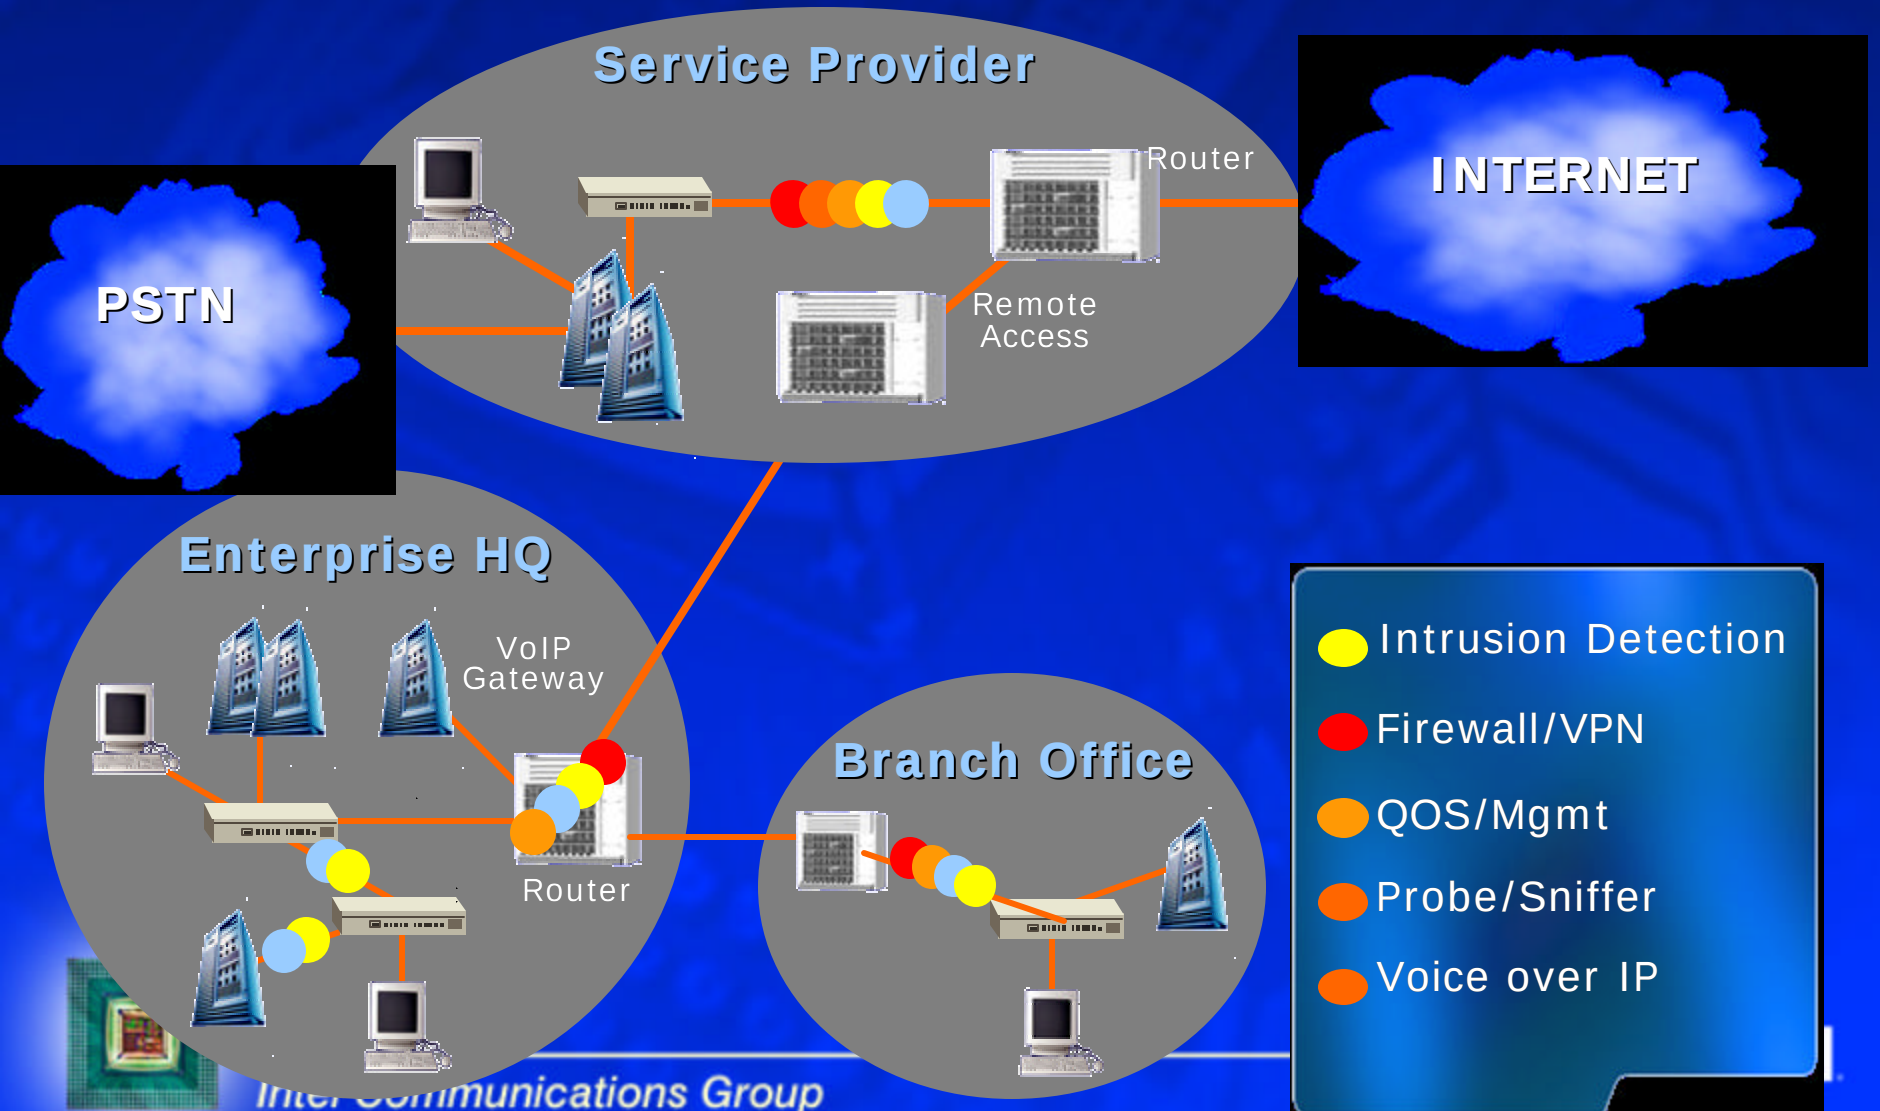

Changing notion of “network edge”

- New services deployed throughout the nodes of the network
  - VPNs, intrusion detection
  - QoS-enabled Video streaming and distribution
- Edge services require:
  - More per-packet processing
  - Flexibility to upgrade nodes in the field
- IP-based networks are no longer passive packet forwarders
  - Must proactively respond to network dynamics to ensure high-availability and security.

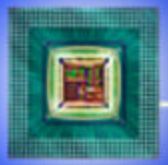

Services Increase the Processing Load

By Increasing Packet Touch

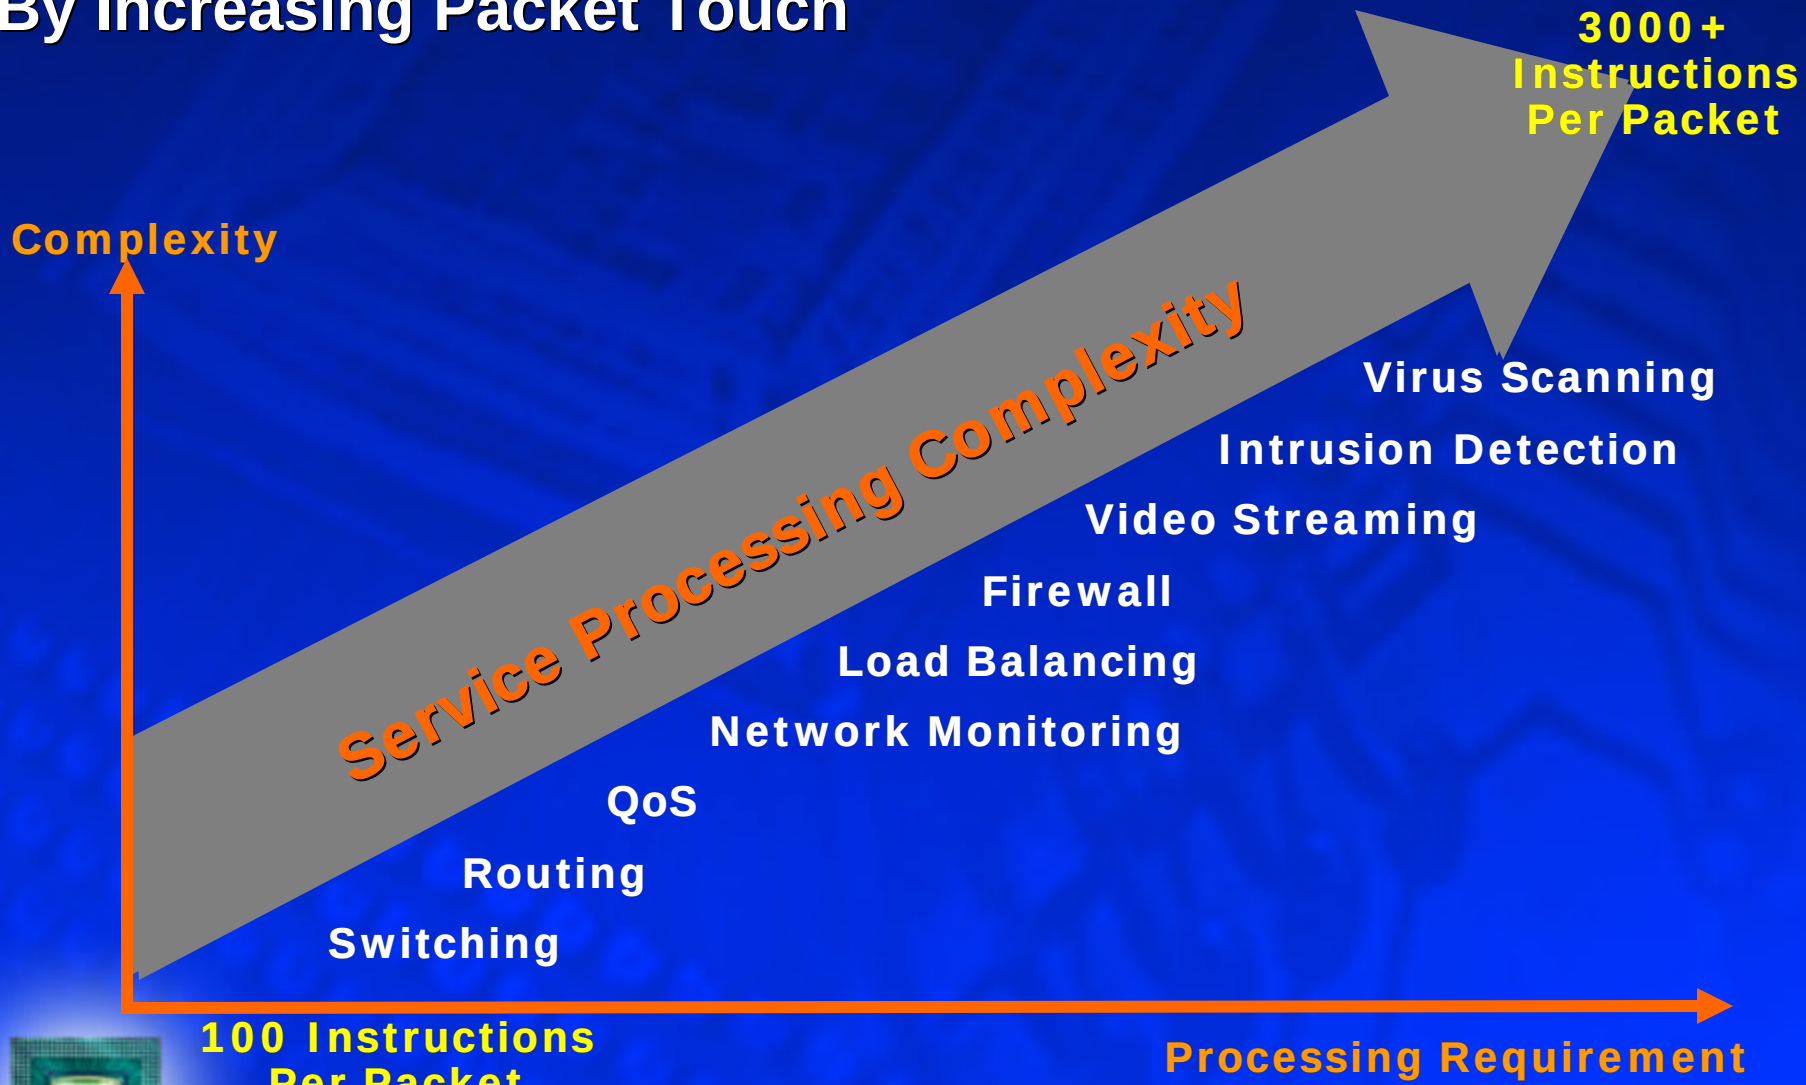

More Detailed Categorization of Service Processing Requirements

- ***The characterization is really multi-dimensional (still simplified):***

Service	Data touch	State touch	Compute Requirements
Switching	Low	Low/Med	Low/Med
Routing	Low	Low/Med	Low/Med
QoS	Low/Med	Low/Med	Low/Med
Stateful Firewall	Low/med	Low/Med	Low-High
Proxy Firewall	Med/high	Med	Med
Load Balancing	Med	Med/High	Low/Med
CB Load Balance	High	Med/High	Low/Med
VPN	High	Med	High
Virus Detection	High	High	High
IDS	High	High	High

* Crypto processing needs a special processor.

Network processors to the rescue ….

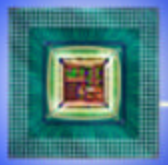

OUTLINE

- The Environment
- What is a network processor?
- Intel's Vision and Contribution
- Technical Challenges
- Summary

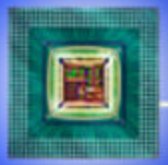

What is a Network Processor (NPU)?

- **Flexible Programming Capability**
  - Programmability, with rich packet processing features
- **Capable of L2-L4+ data processing at wire rates**
- **Highly integrated Processor**
  - Microcoded and/or hardwired acceleration engines
  - Memory subsystem controller(s)
  - Closely coupled with media (integrated or attached)
  - Good data flow management, internal comm's support, and partitioning to enable acceleration
  - Memory latency hiding techniques includes memory controllers

Highly integrated compute, memory and media resources to tackle packet processing problems close to the wire.

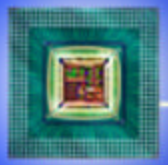

Problem Spaces Targeted by NP's

- CPE, Edge/Access, and Core applications
  - Encapsulation/Decapsulation to fabric/backplane
  - Switching/Routing (L2/L3/L4)
  - Cell/packet conversion
- L4-L7 applications; content and/or flow-based
- Intrusion Detection (IDS) and Proactive monitoring of networks
  - Particularly challenging
  - needs processing of many state elements in parallel,
    - unlike most other networking apps which are more single-path per packet/cell

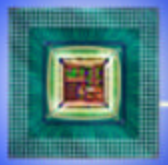

Why not just use a GHz+ Pentium?

- Network processing is basically a dataflow problem
  - Almost every new packet requires new state
  - uP caches do not hide latency
- Time budgets per cell/packet at high wire speeds are low
  - OC-48 arrival rate = 160ns, OC-192 arrival rate = 35ns, OC-768 arrival rate = 8ns

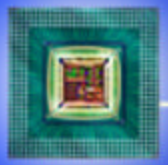

NPUs can provide acceleration

- Keep stalls on memory and cache to a minimum, and provide specialized compute capabilities
- Offload high-touch portions of applications
  - Header parsing, checksums/CRCs, RegEx string search
- Offload latency-intensive portions to reduce stall time
  - pointer-chasing in hash table lookups, fetch of candidate portion of packet for header parsing
- Offload compute-intensive portions with specialized engines
  - Crypto computation, ATM CRC, packet classification
- Provide efficient system management
  - Buffer management, descriptor management, communications among units, timers, queues, etc.

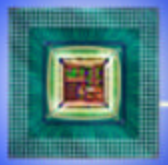

OUTLINE

- The Environment
- What is a network processor?
- Intel's Vision and Contribution
- Technical Challenges
- Summary

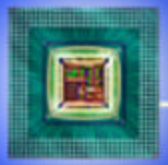

Intel Vision: rich offering of modular networking building blocks

Key Industry Transitions

- Ethernet
  - 10G emergence in metro
  - Provide building blocks at every level
    - Optics, framers, network processors, switch fabrics
- Processing in the Network
  - IXP family of network processors
- End2End view
  - Clients (LAN, wireless, cable modem, xDSL, etc)
  - Enterprise (SAN, switch fabrics)
  - Metro/Core
    - Building blocks for access/edge and core routers

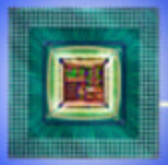

IXP Network Processors

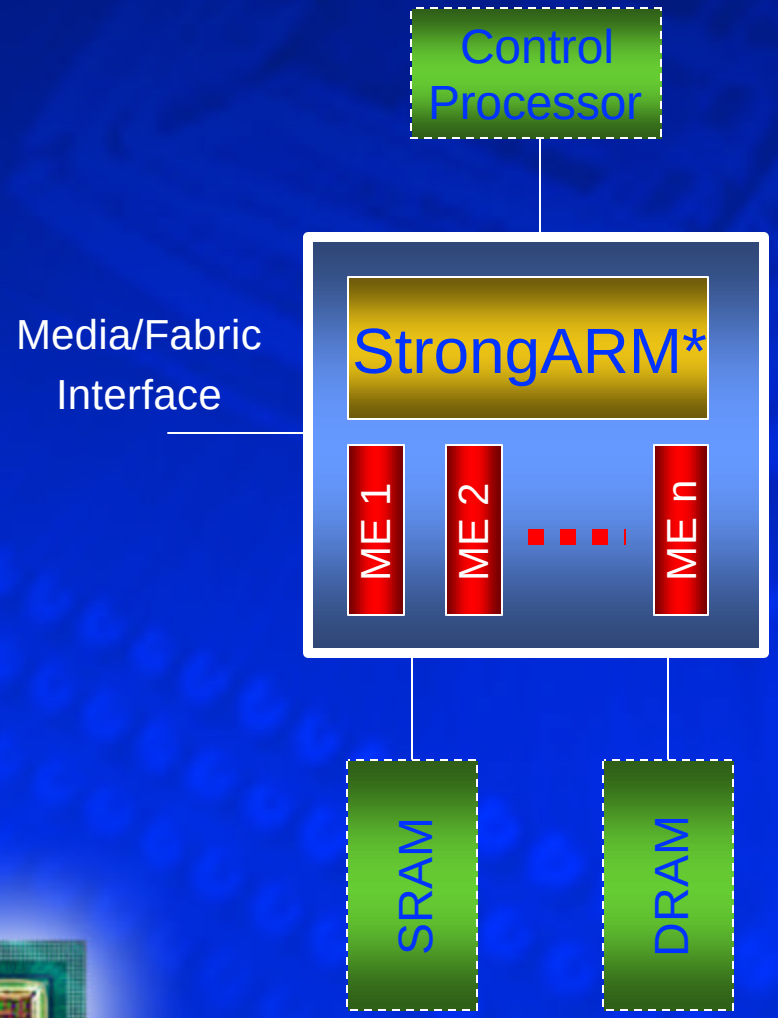

- **Microengines**
  - **RISC processors optimized for packet processing**
  - **Hardware support for multi-threading**
- **Embedded StrongARM/Xscale™**
  - **Runs embedded OS and handles exception tasks**

IXP Network Processor Family - Multi-Processor Architecture

Parallel Micro-Engines

- Programmable 32-Bit RISC Processors (232 MHz->1.4GHz)
  - Independent Control Store
  - Local I/O Transfer Store
- Equal access to all resources
  - SRAM & SDRAM Memory
  - IX Bus & PCI Peripherals
- Non-Blocking Internal Buses
  - > 2X External Bandwidth
  - Eases programming

Parallel Operation Supports Very High Throughput

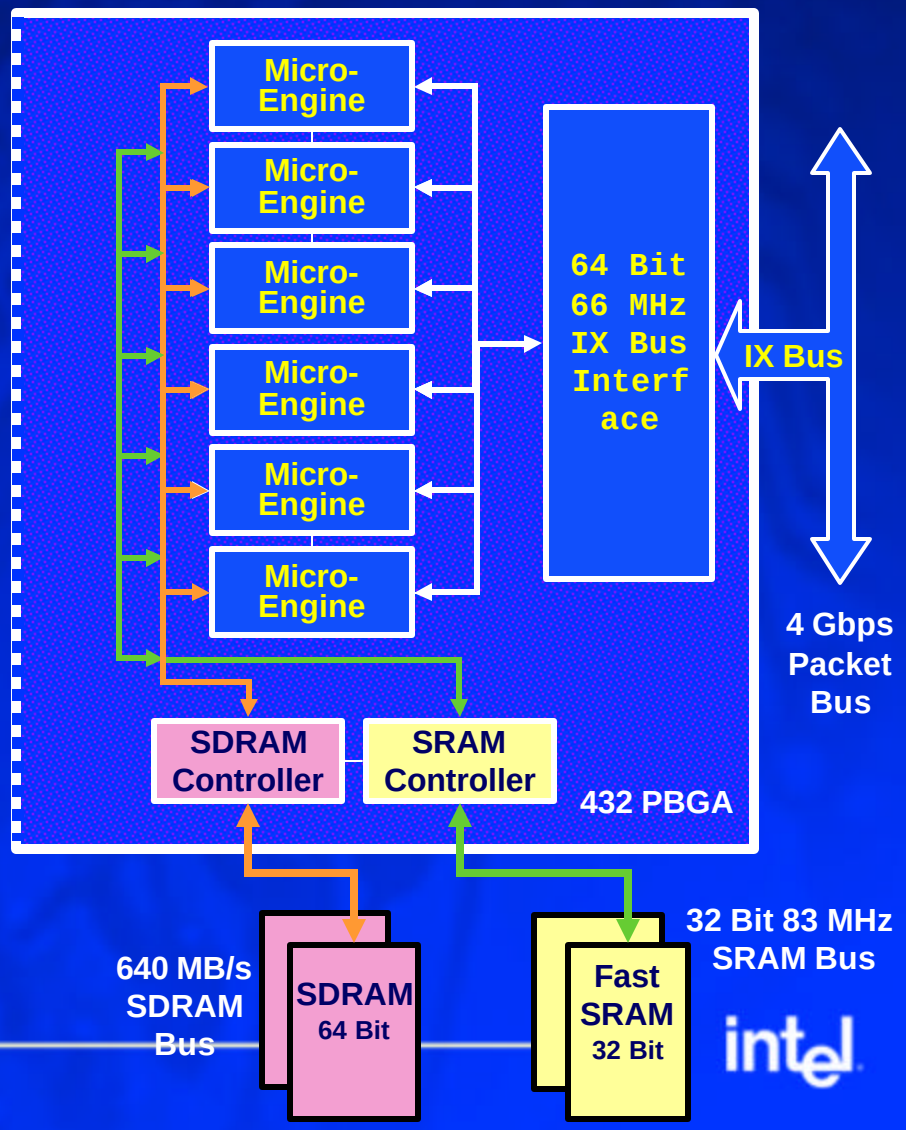

Current Offering: IXP12xx Network Processor

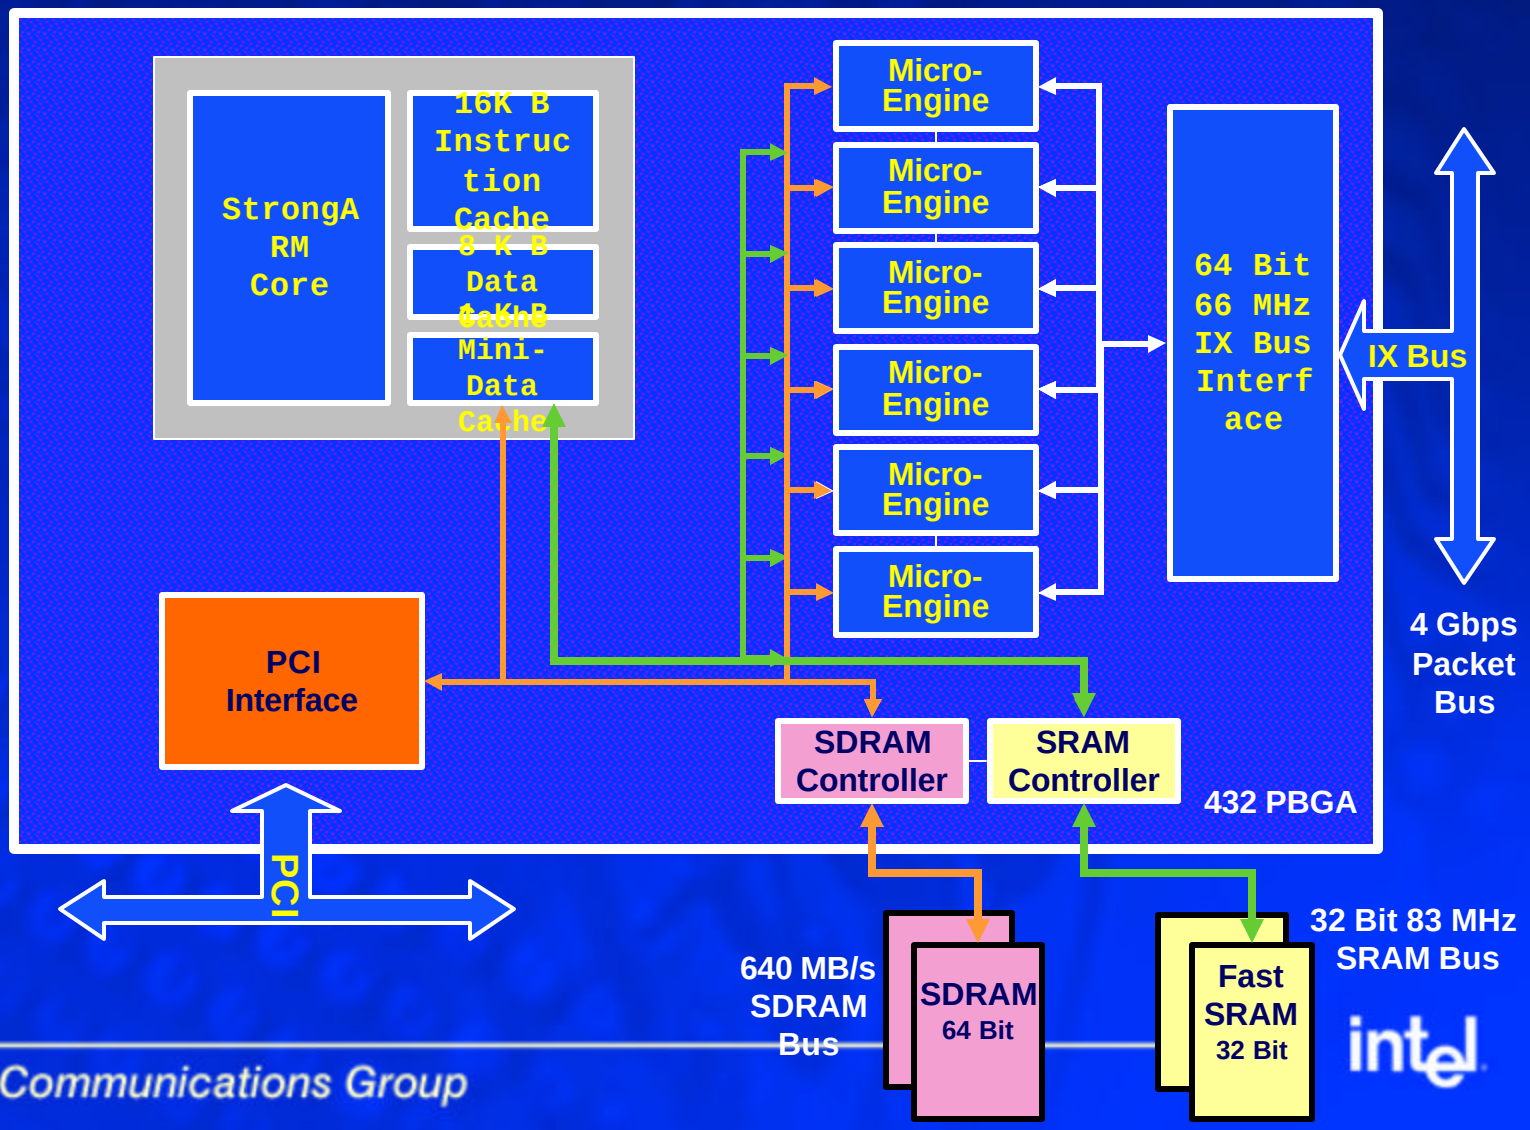

Next Generation: IXP2800

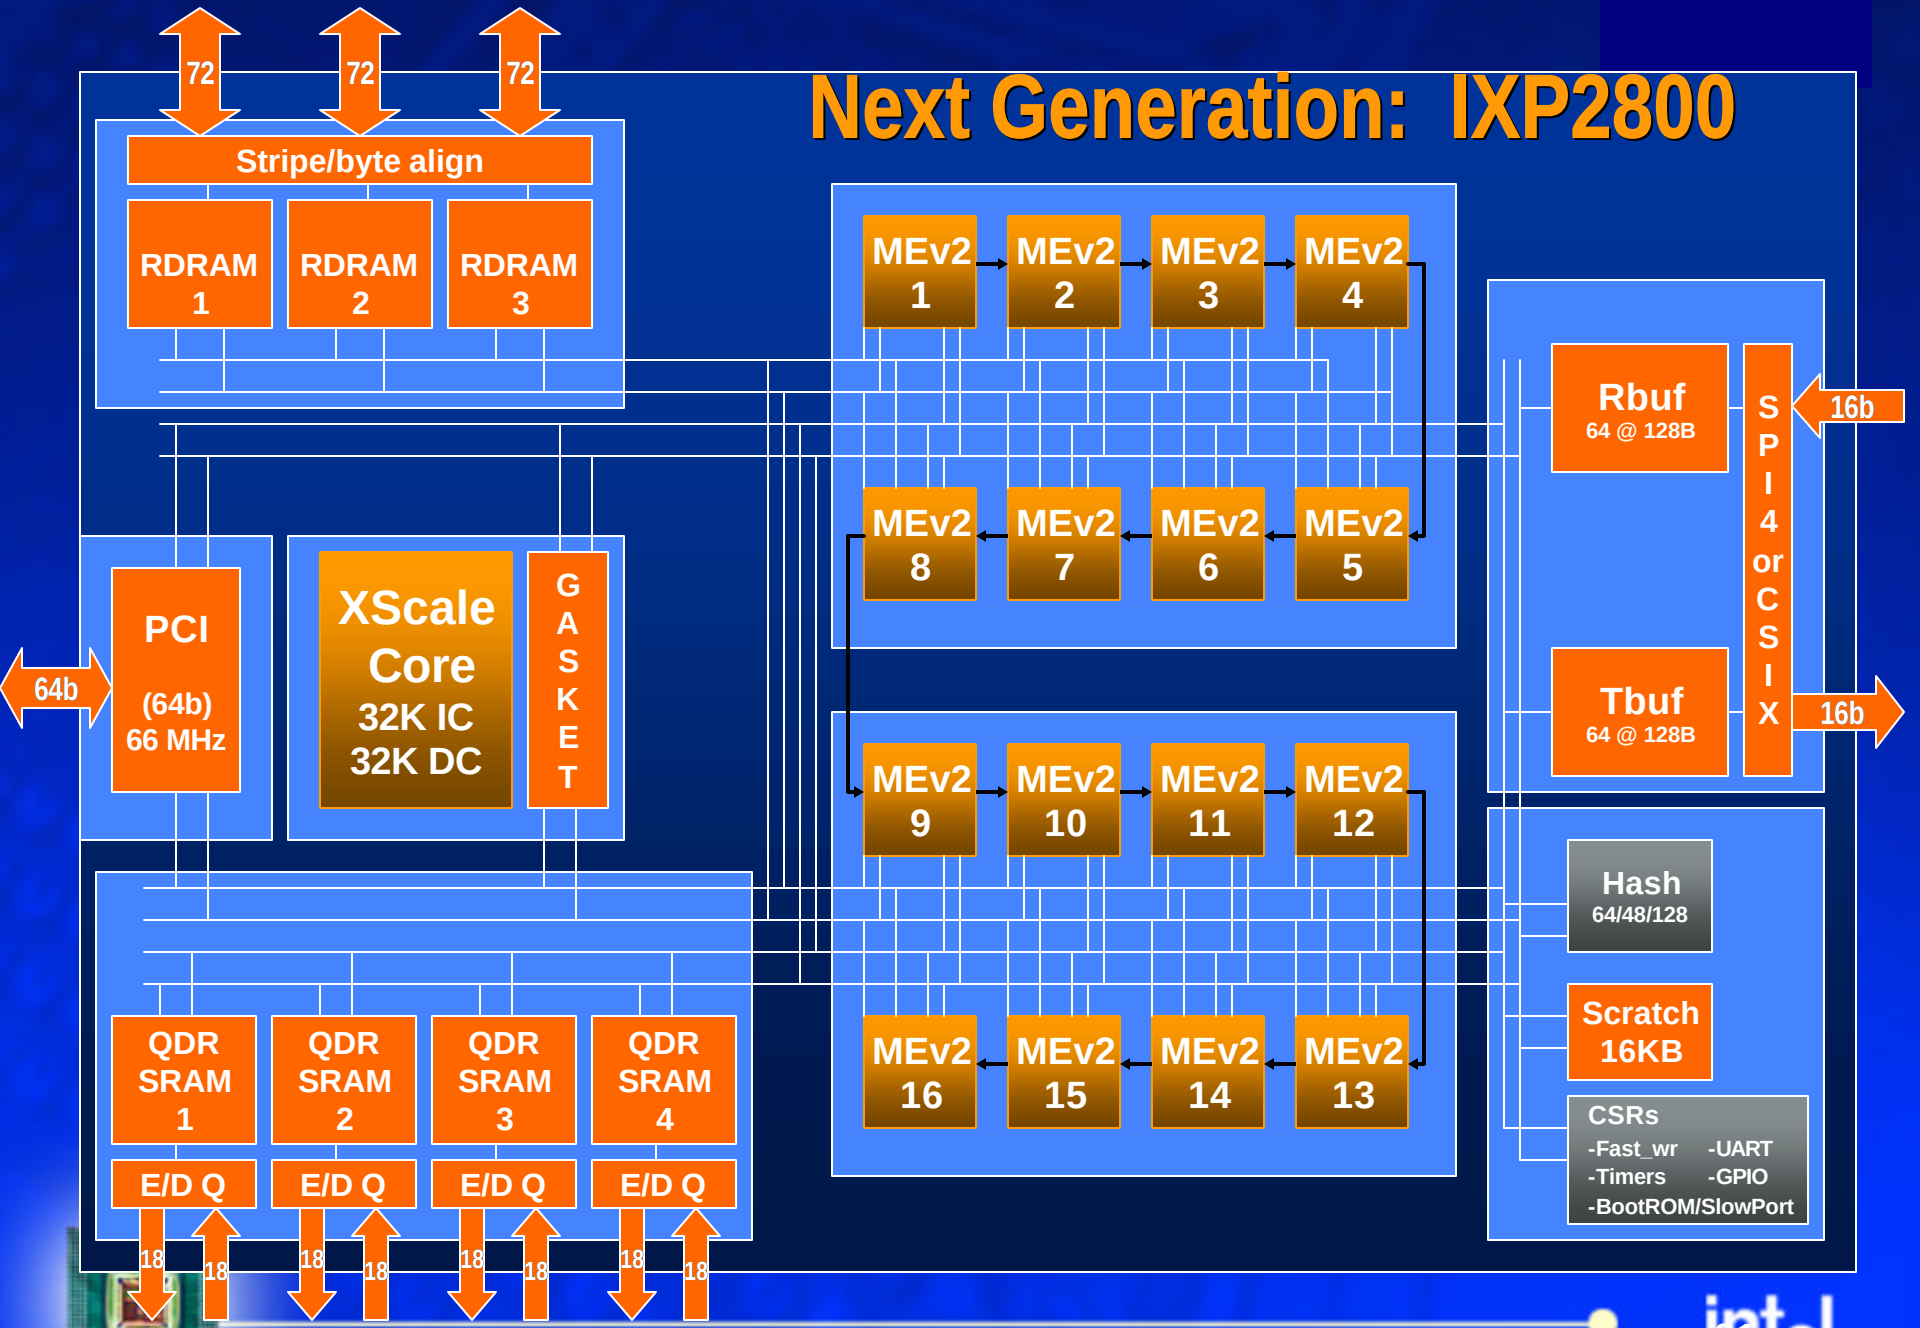

2nd Generation Hardware Features: *Performance and Ease of Programming with a Mixed Pipeline*

- **Ring Buffers**
  - **Allows for Context and Functional Pipeline to talk to each other without horrendous synchronous timing issues**
  - **Use the best type of pipeline for the job**
- **Next Neighbor Registers and Signaling**
  - **Allows for single cycle transfer of context to the next logical micro-engine to dramatically improve performance**
  - **Simple, easy transfer of state**
- **Distributed data caching within each micro-engine**
  - **Allows for all threads to keep processing even when multiple threads are accessing the same data**
- **1.4Ghz initial micro-engine clock speed**
  - **Speed is good!**

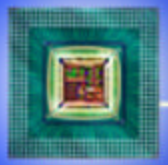

intel

Why Hardware Multi-Threading?

- Ensures efficient use of micro-engine pipeline
  - Simulation results indicate ~75% of cycles spent waiting for memory
- If single threaded, the equivalent performance would require 24 engines
  - Huge silicon savings
- Keeps micro-engine pipeline active during otherwise idle periods due to memory latency

Four Threads Executing on a Single Microengine

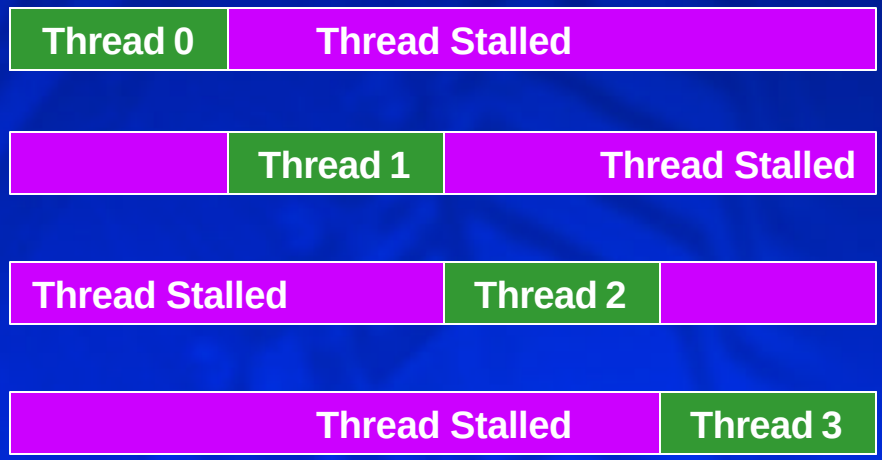

Going to 16 engines, 8 Threads each in IXP2800 (128 HW threads!)

IXA Portability Framework

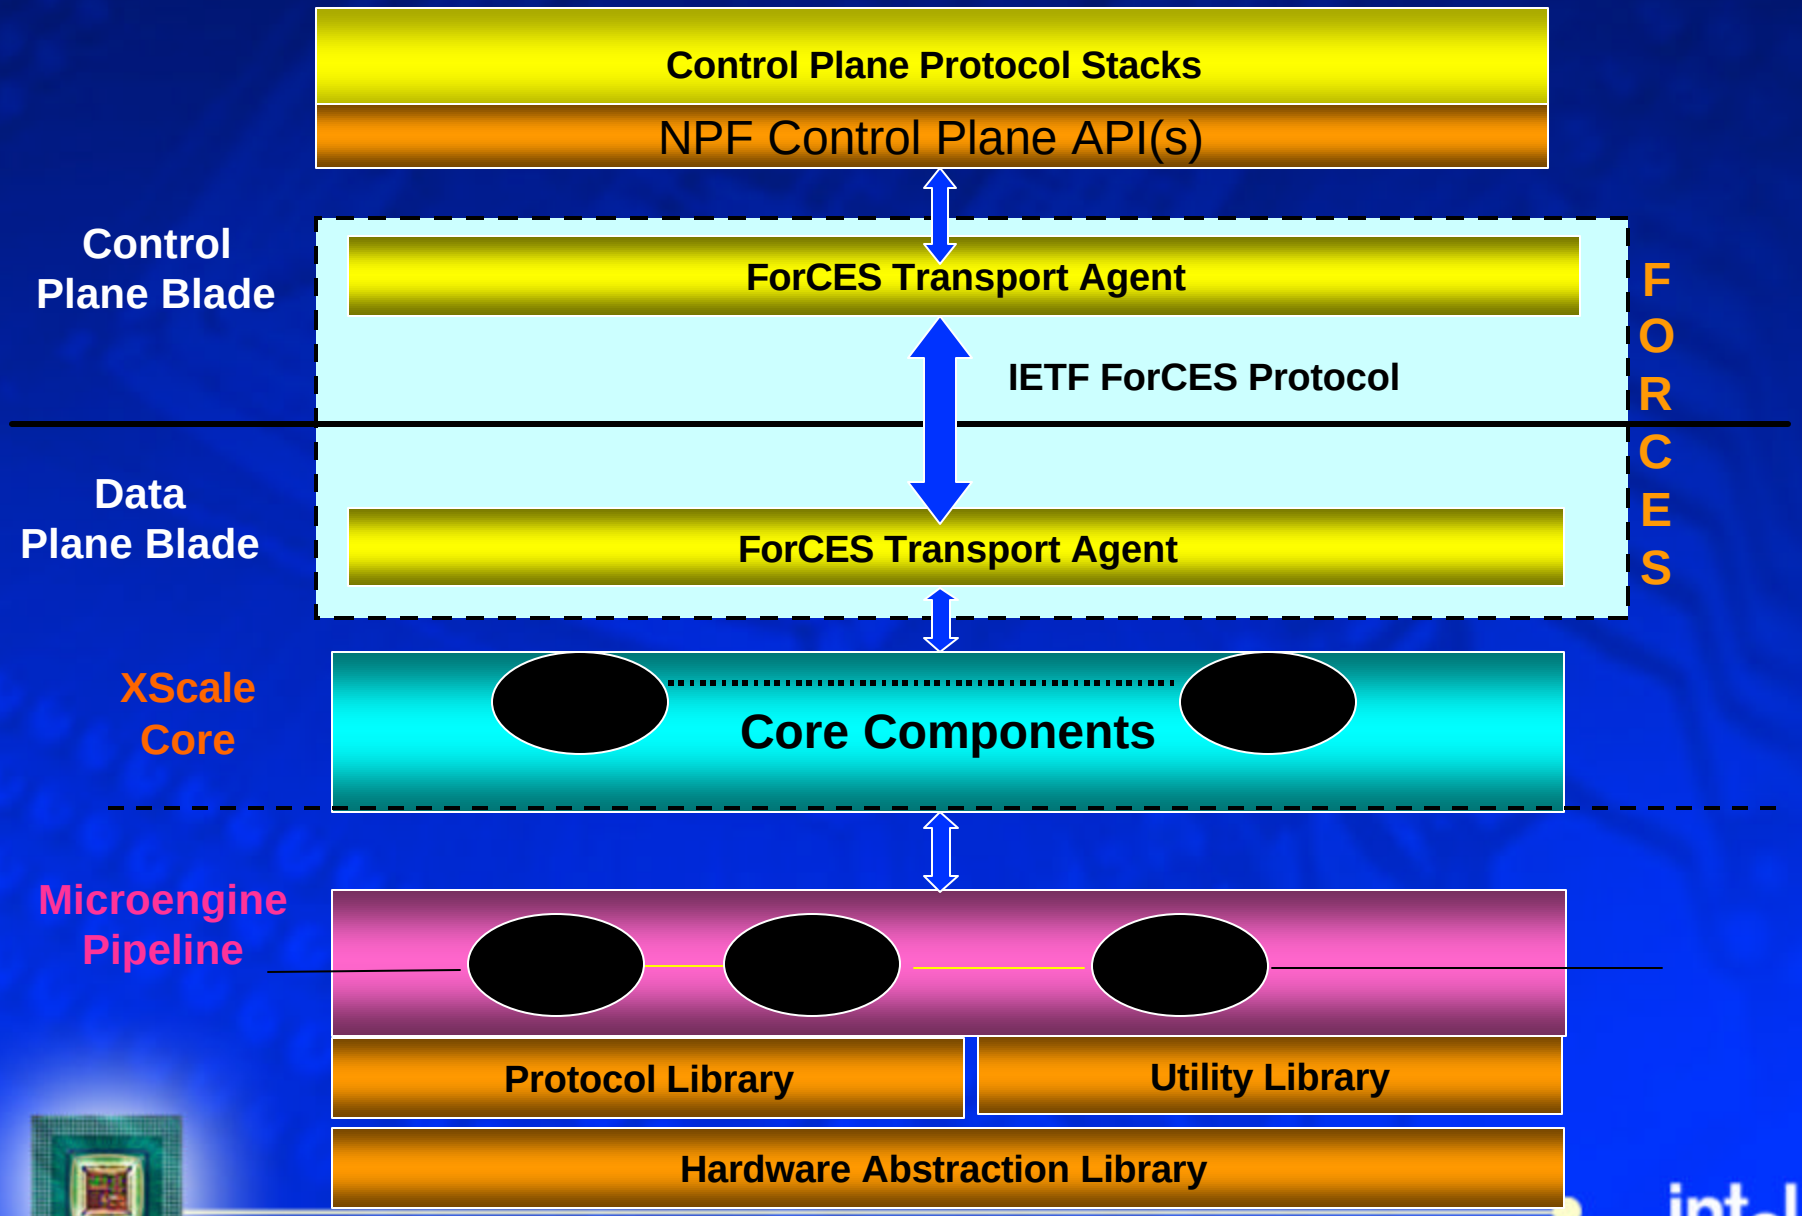

What is IETF ForCES?

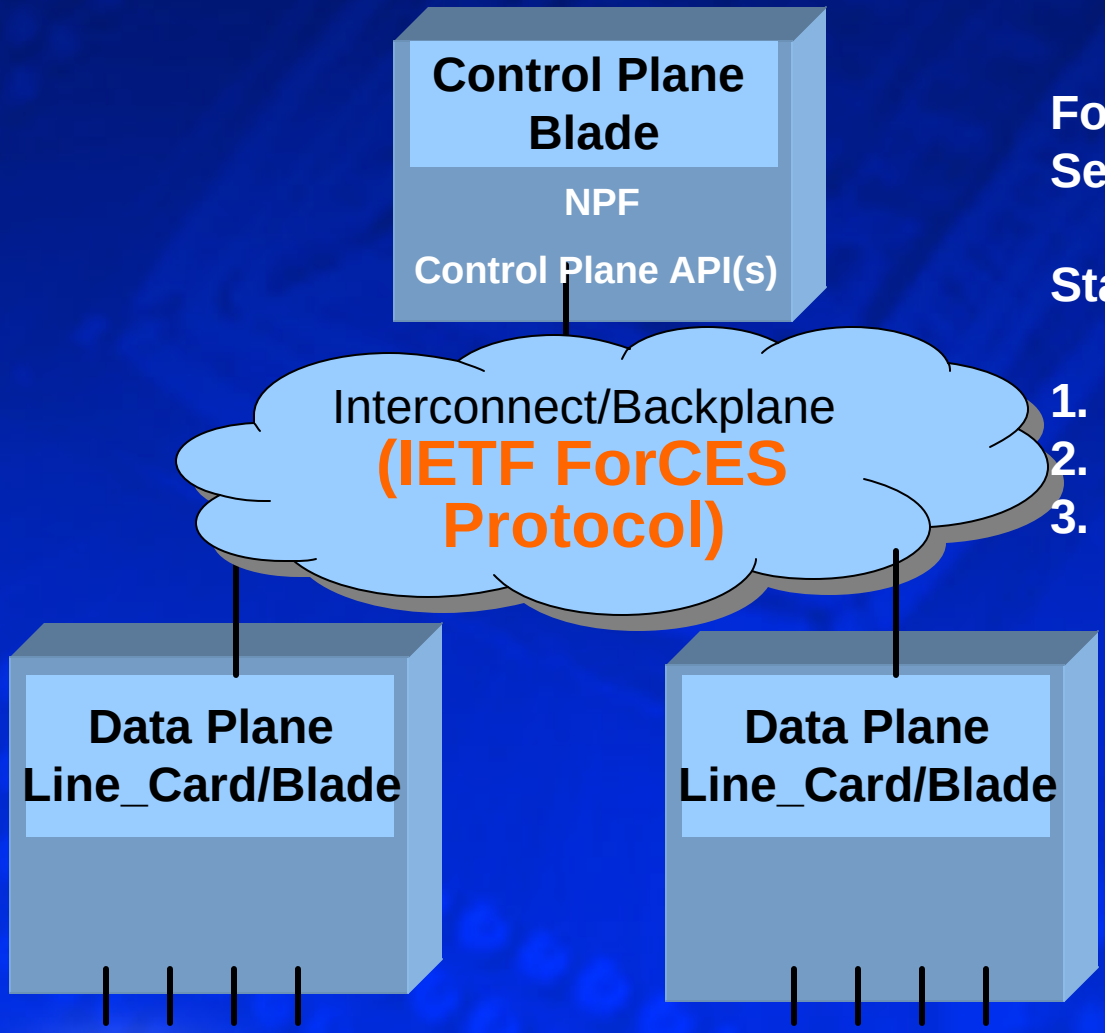

Forwarding and Control Element Separation for IP Networks

Standards track work in three steps (expected milestones):

1. **Internet draft -- current**
2. **Draft standard – 2002/early 2003**
3. **Standard – late 2003**

Join us: www.sstanamera.com/~forces

Next IETF meeting: July 14-19, Yokahama, Japan

IXA University Program

- Enable universities around the world to incorporate IXA-based components into research and curriculums:
- Three primary goals:
  1. Encourage the use of the IXA in both research and curriculum.
  2. Create a community of IXA-related research universities
  3. Give students the hands-on experience and training in building and experimenting with real-world network applications using IXA.

For More Information:

www.intel.com/research/university/comm/index.htm

OUTLINE

- The Environment
- What is a network processor?
- Intel's Vision and Contribution
- Technical Challenges
- Summary

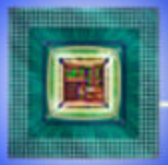

Programmability brings flexibility and potential for change

- Migrating Internet to a robust infrastructure when ….
  - Fragile and vulnerable IP networks vs. robust voice networks
  - Inflexible ASICs have limited capabilities in the fast path with little insight into network health
  - Only insufficient, stale & unreliable statistics provided through old SNMP MIB-II & RMON polling-based mechanisms
- … The world is moving
  - DoS and DDoS Attacks are on the rise
  - Network worms are devastating networks (eg. Code Red)
  - Highly transactional mission critical applications demand robust networks
  - Streaming applications, SLAs and QoS
  - Dynamic Traffic Engineering

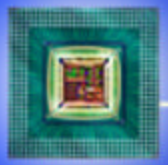

Example Opportunity

- **Programmable NPUs for advanced video streaming platforms**
  - **Handle a variety of streams**
  - **QoS and Quality adaptation over mobile and wireless**
  - **Efficient overlays to achieve multipoint streaming independent of underlying infrastructure**
  - **Dynamic adaptation to attacks, failures**

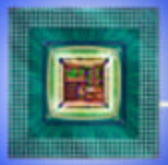

How Network processors can help

- **Programmable and secure infrastructure**
  - **Robust, verifiable code on the network nodes**
  - **Verifiable transactions among nodes**
  - ***Watchdog* and *bloodhound* nodes distributed in the infrastructure**
    - **A cooperative model based on information sharing**
    - **Watch for performance anomalies and trigger reactions**
      - Route changes, install new label-switched paths, etc
    - **Watch for behavior anomalies in traffic and nodes**
      - Squelch the traffic closer to the entry
      - “replace” misbehaving nodes with alternate paths

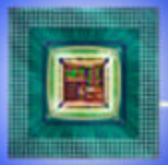

Example: Proactive, robust video streaming infrastructure

Traffic can be redirected across uncontested paths, or marked with high QoS

Service Verification E2E

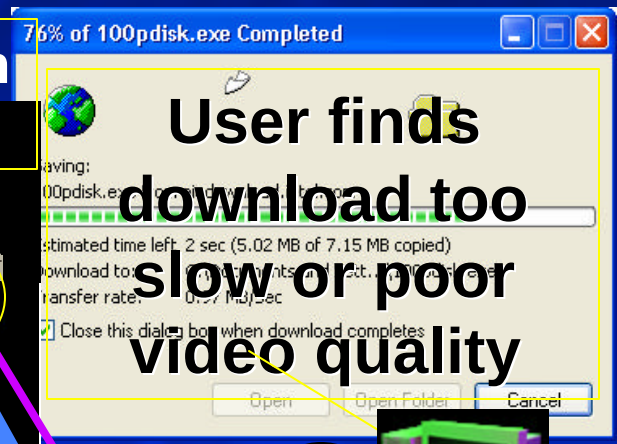

User finds download too slow or poor video quality

Web Server state distinguishes mission critical apps & users

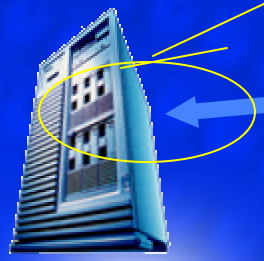

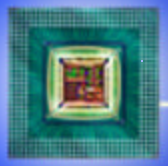

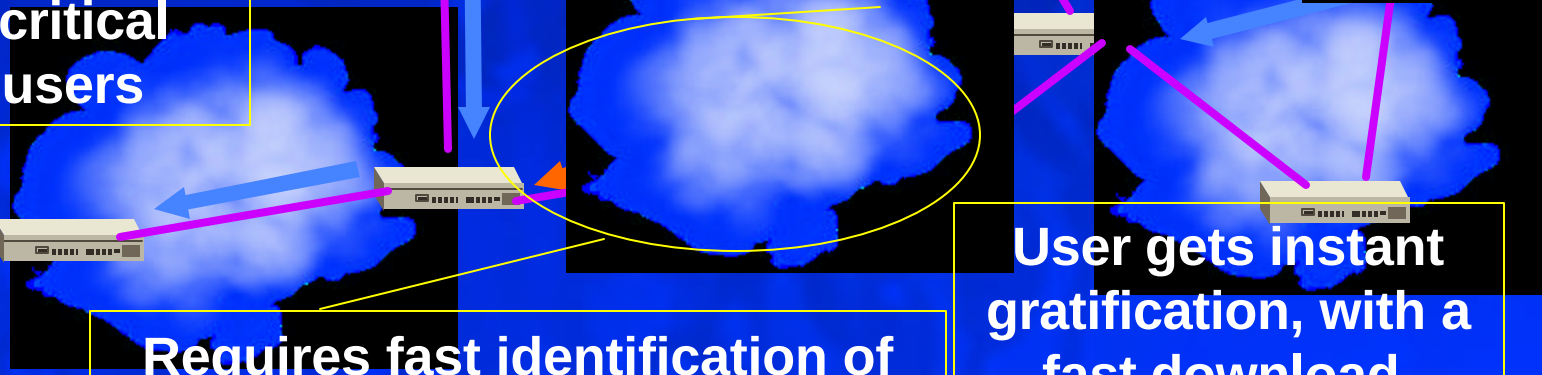

Requires fast identification of mission critical flows correlated with server state

User gets instant gratification, with a fast download.

Server Is Targeted for a Distributed TCP SYN Flood Attack

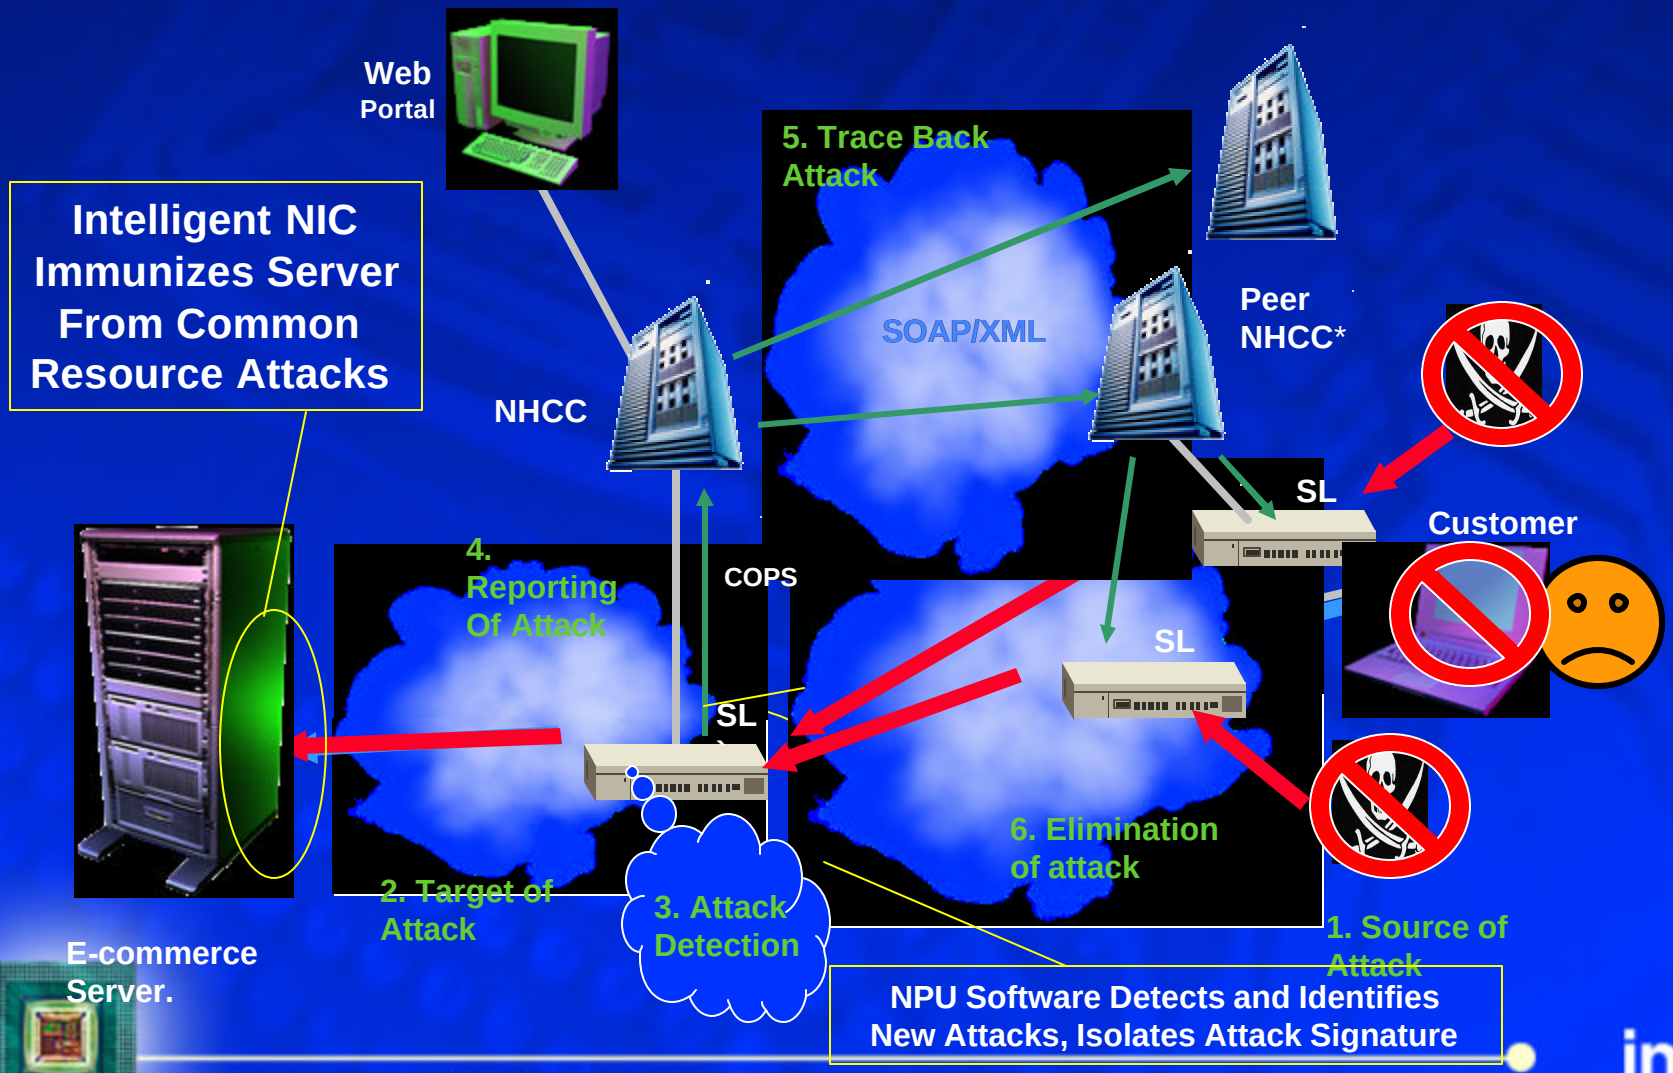

OUTLINE

- Overall Vision
- IXP Architecture
- Software Support
- Technical Challenges
- Summary

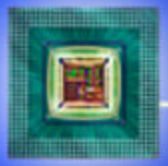

Summary

“Edge” is everywhere

- Services being deployed at all network nodes
- Dynamic/Rapid service deployment requires data-plane programmability

Network Processors evolving to meet the challenge

- Packet Processing: more instructions per packet
- Evolving functionality: Big advancements this year

Opportunity: Proactive Networks programmable platform

- Quality of Service, Media Processing, Anomoly Detection and Correction
- Rapid/Dynamic deployment: Add new functionality on the fly

Join us in Proactive Networks research efforts

www.intel.com/research/university/comm/index.htm

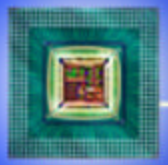

intel

BACKUP

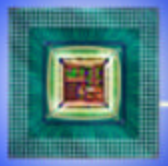

IXP1200 Programming Book

By Erik J. Johnson and Aaron R. Kunze
List Price: $49.95
Softcover, 320 Pages
ISBN: 097128878X

Available at Intel Press Booth, Amazon.com, fatbrain.com, and Barnes&Noble.com

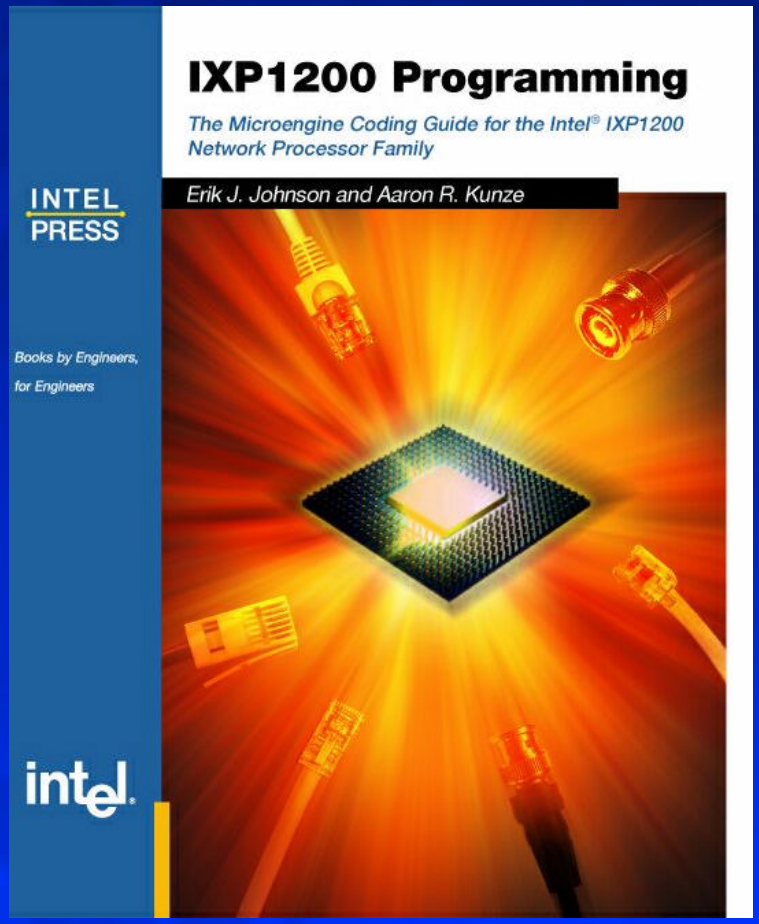

- **Topics include:**
  - **Microblocks and microACEs**
  - **Microengine C**
  - **Receiving, processing and transmitting packets**
  - **Data structure and algorithm design**

- **Ordering: www.intel.com/intelpress/ixp1200**
- **Chapter 2 online: http://www.intel.com/intelpress/ixp1200/descriptions.htm**