Benchmarking Network Processors: More than just MIPS

Nathan Pham
Performance Engineer
IBM Microelectronics

IBM®

Network Processor

- Core component of network equipment (routers, switches, firewalls, web switches, etc.)
- The goal is to be to networking products what CPUs are to PCs
- Hybrid solution that provides high performance through hardware and flexibility through software programmability
- Optimized to handle packet processing

IBM®

NP Functions

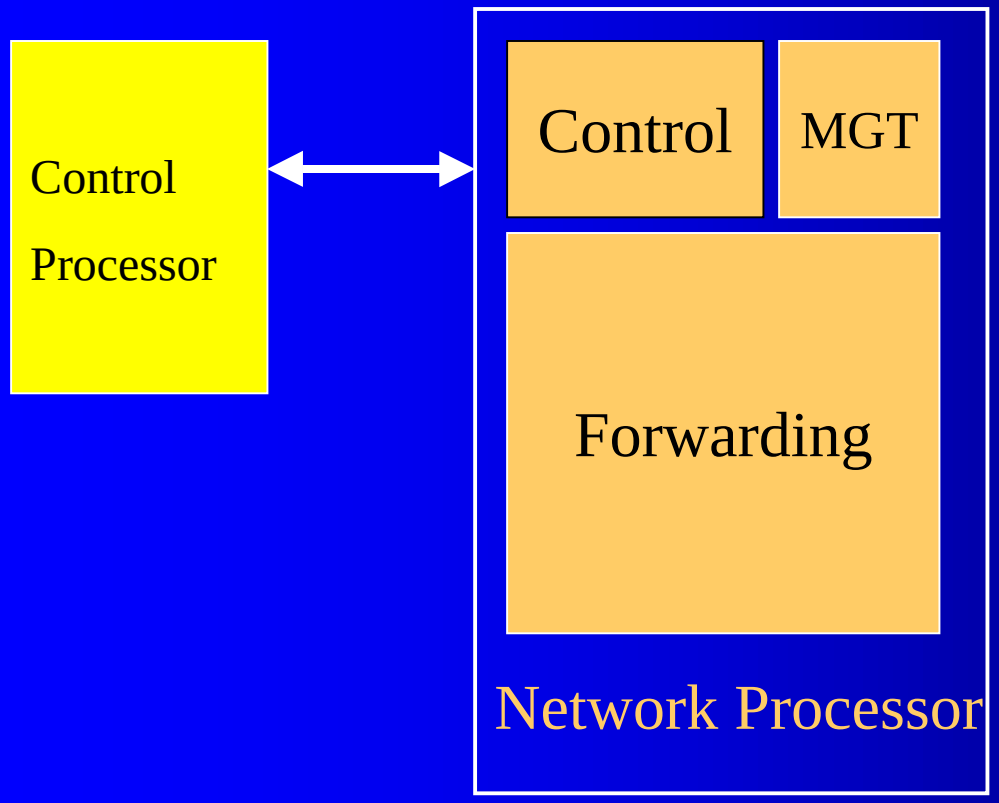

- Steady-state functions
  - Ex: frame storage, alteration and classification
- Control functions
  - Ex: routing and signaling
- Management functions
  - Ex: NP configuration and diagnostics

IBM®

NP Architecture

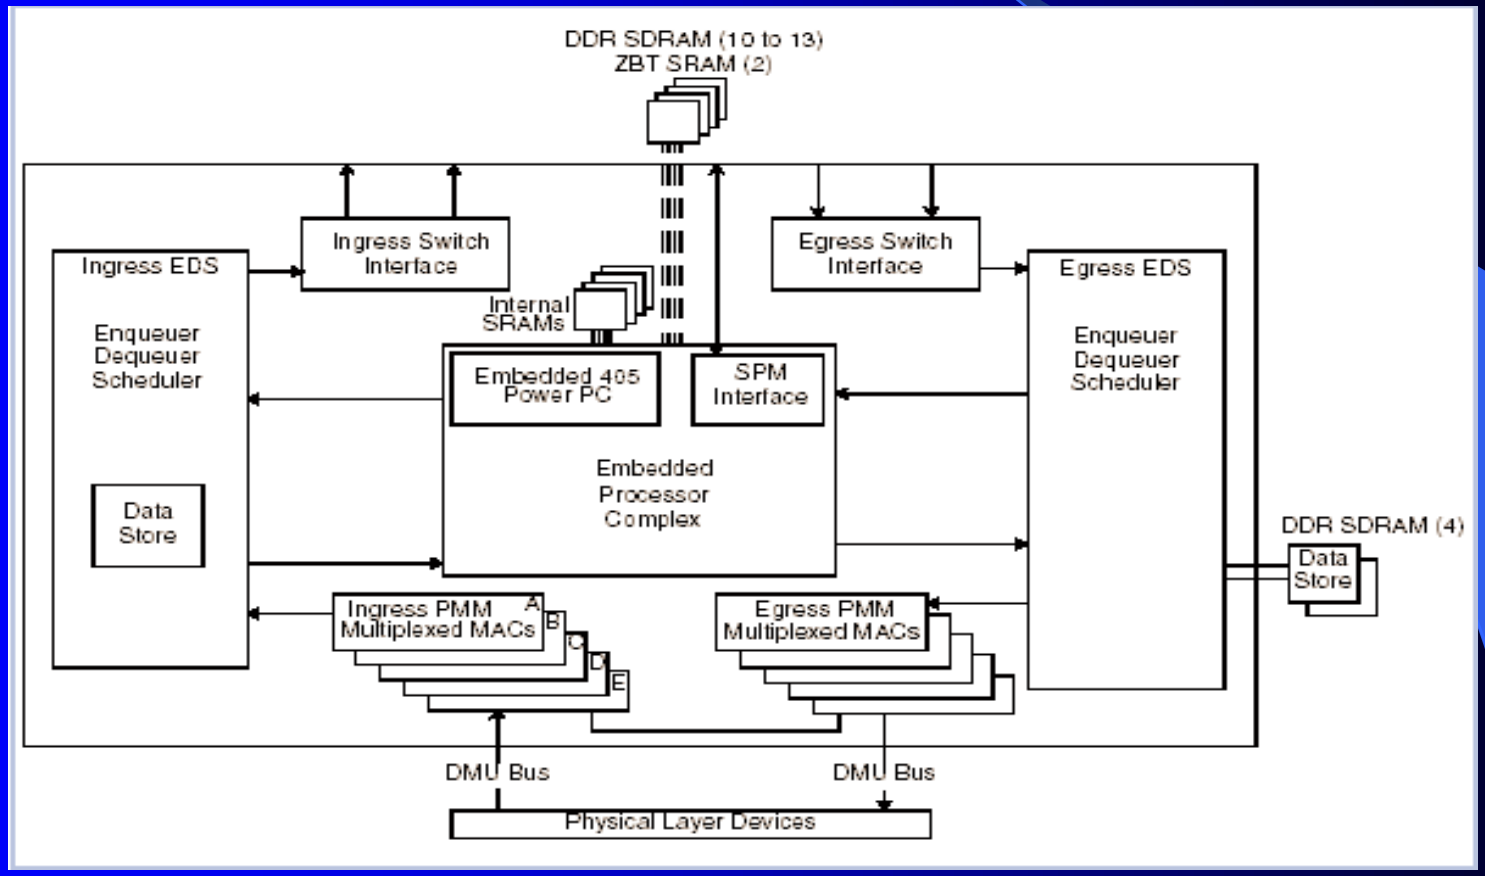

Benchmark μPs vs. NPs

- NP is designed for fast packet processing, not for general applications.
- Existing computational-intensive benchmarks for CPUs are not applicable to NPs.
- Different performance metrics

IBM®

NP Benchmarking Levels

- System level
  - Ex.: routers, firewalls, and web switches.
- Function level:
  - Ex.: IP forwarding, MPLS forwarding, QoS, etc.
- Micro operation level:
  - Ex.: LPM table lookups, 5-tuple table lookups, and CRC calculations.
- Hardware operation level:
  - Ex.: throughput/latency for accesses to memory.

IBM®

Benchmark Requirements

- Architecture independent
- Specific to the NP application domain of interest
- Meaningful performance metrics
- Realistic test environment

IBM®

IPv4 forwarding function level benchmark

- Basic IPv4 forwarding function
- Developed by Network Processing Forum (NPF) benchmark working group
- Members: Over 80 companies including IBM, Intel, Agere, EZ Chip, Vitesse, etc.
- Other benchmarks: IPv6, MPLS, and DiffServ

IBM®

IPv4 Forwarding Benchmark Setup

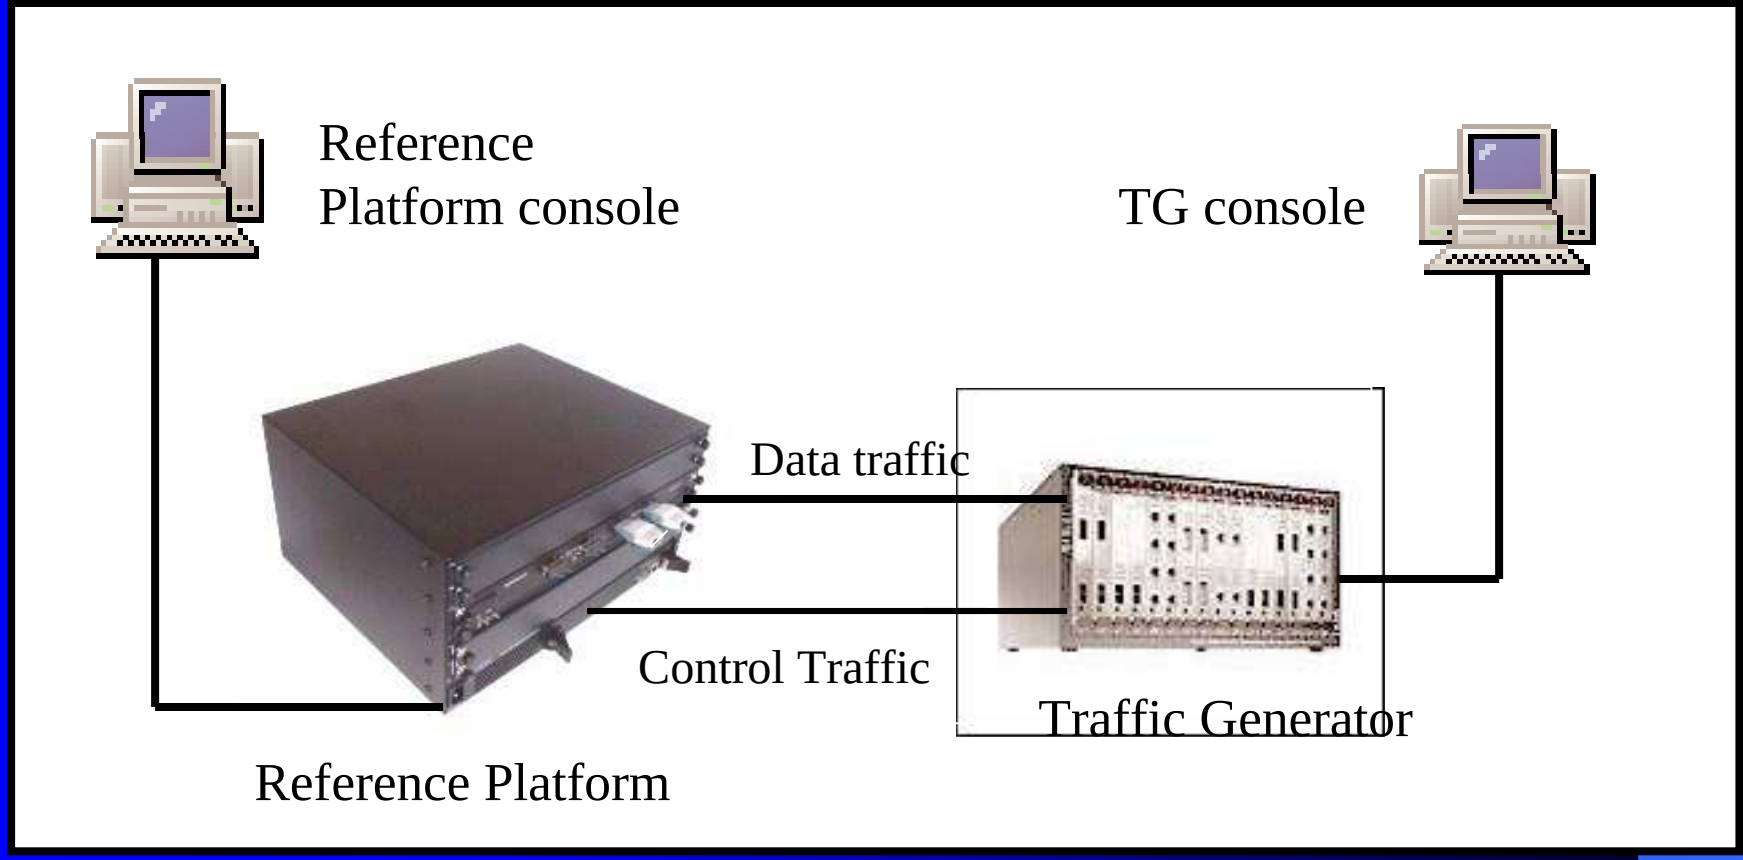

IPv4 Forwarding Benchmark Metrics – Forwarding Rate

- Max rate that frames can be forwarded
- Traffic sent to NP at max line rate
- Influencing factors:
  - Packet processing time
  - Packet rate
  - Queuing mechanism (queue size, discard mechanism)
- Reporting numbers
  - Frame rate: in Million packets per second (Mpps)
  - Bit rate: in Gigabit per second (Gbps)
  - Percent of line rate

IBM®

Metrics – Throughput Rate

- Max rate that frames can be forwarded with no frame loss
- Cannot always derive from forwarding rate
- Influencing factors:
  - Packet processing time
  - Packet rate
  - Queuing mechanism (queue size, discard mechanism)
- Reporting numbers
  - Frame rate: in Million packets per second (Mpps)
  - Bit rate: in Gigabit per second (Gbps)
  - Percent of line rate

Metrics – Latency

- Time needed to process and forward a data frame
- Sources of latency
  - Queuing delay
  - Processing time
  - Frame movement internal to NP
  - Stall time
- Influencing factors
  - Data rate
  - Frame size
  - Queuing mechanism (queue size, discard mechanism)
  - Software efficiency
  - Resources utilization efficiency

IBM®

Metrics – Loss rate

- Percent of incoming data frames dropped by NP
- Not simply the reverse of throughput rate
- Two thresholds
  - Max line rate
  - Throughput rate
- Shows NP forwarding behavior between thresholds

IBM®

Metrics – Overload forwarding rate

- Forwarding rate in extremely congested scenario
- Data frames sent at much higher than max line rate.
- Influencing factors
  - Queue size
  - Discard mechanism
  - Flow control
  - Processing time
- Shows NP forwarding behavior in stressed condition

IBM®

Metrics – Forwarding table update rate

- Most important control function for IPv4
- Max rate at which forwarding table entries can be added, updated, or deleted
- Influencing factors:
  - Control Point (CP) processor power
  - Communication channel b/t NP and CP
  - Control action processing time
- Implementation details
  - Routing protocol
  - API calls
  - Simulated packets

IBM®

Metrics – Headroom

- Excess processing power left over while concurrently keeping up with data traffic at throughput rate
- Simulate real world usage of NP
- Data and control traffic sent concurrently
- Measured by route updates/second
- Influencing factors:
  - NP processing power
  - Priority mechanism (b/t control and data traffic)

IBM®

IPv4 Forwarding Benchmark Parameters

- Routing table
  - LPM Table lookup time ~ processing time
  - Table size and structure
- Frame size
  - Smaller frame size ~ higher packet rate
  - Real world traffic profile
- IPv4 forwarding operation
  - Vanilla IPv4 forwarding
  - IPv4 with option/control
- Traffic mapping
- Run time

IBM®

Benchmark Requirements Check

Requirements	IPv4 forwarding benchmark
Architecture independent	Yes
Specific to application domain	Yes
Meaningful metrics	Yes
Realistic test environment	Yes

IBM®

IBM PowerNP NP4GS3

- 16 programmable picoprocessors provide 2128 MIPS aggregate processing capability
- Embedded PowerPC processor included
- Hardware accelerators
- Multi-threading supported
- 40 Fast Ethernet/4-Gb MACs/OC-48c/OC-48/four OC-12/sixteen OC-3
- Up to 64 NPs can be connected via switch fabric

IBM®

IBM PowerNP NP4GS3 Performance

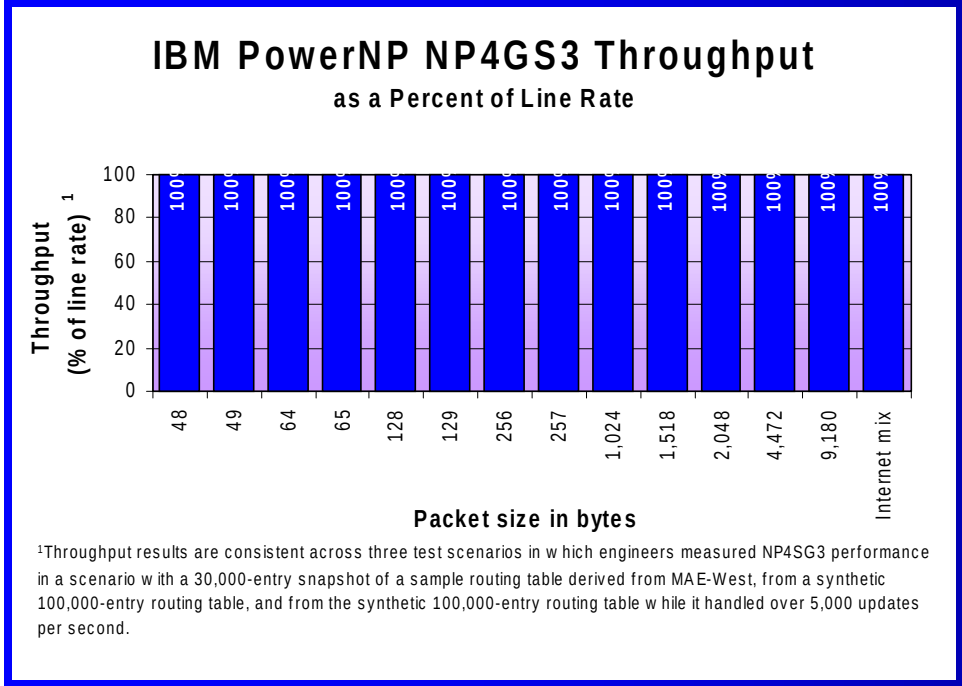

[1]Throughput results are consistent across three test scenarios in w hich engineers measured NP4SG3 performance in a scenario w ith a 30,000-entry snapshot of a sample routing table derived from MAE-West, from a synthetic 100,000-entry routing table, and from the synthetic 100,000-entry routing table w hile it handled over 5,000 updates per second.

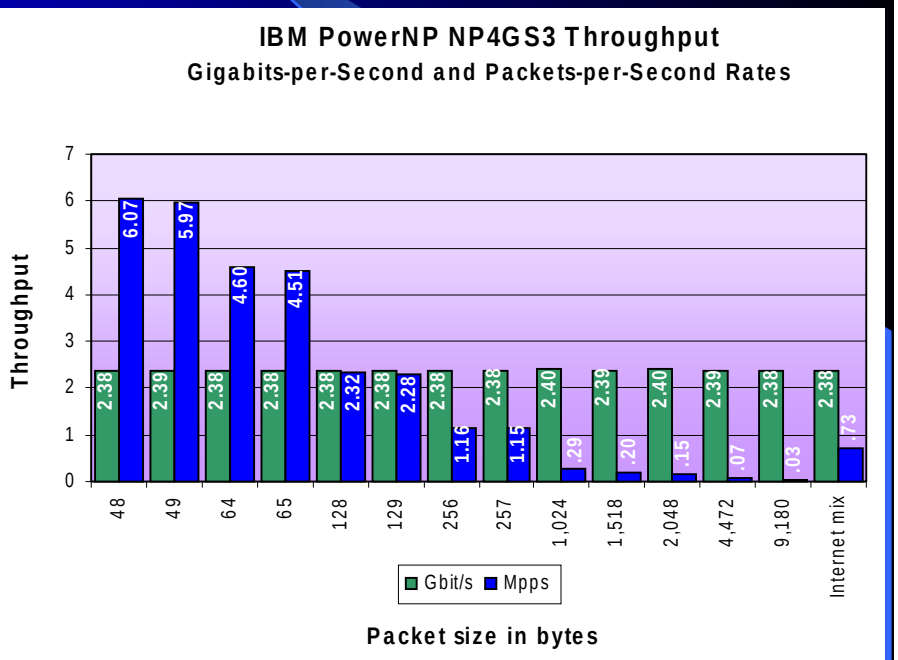

NP4GS3 Performance

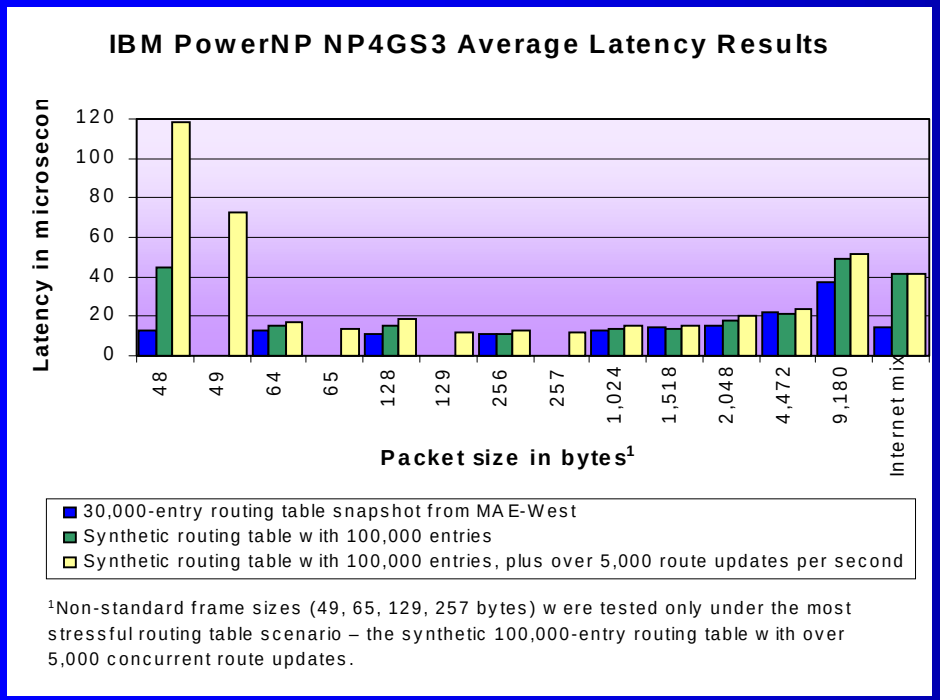

[1]Non-standard frame sizes (49, 65, 129, 257 bytes) were tested only under the most stressful routing table scenario – the synthetic 100,000-entry routing table with over 5,000 concurrent route updates.

- Delivers OC-48c wire-speed performance
- Maintains line rate with a 30,000-entry real-world routing table
- Maintains line rate with a 100,000-entry synthetic routing table
- Maintains line rate with 100,000-entry synthetic routing table while concurrently handling 7,300 routing table updates per second
- Maintains low latency in all scenarios

Summary

- Network Processor is a new and important component of modern network equipments
- Four levels of performance benchmarking the NP: system, function, micro operation, and hardware
- Benchmark should be developed based on the application NPs are used for
- IBM PowerNP NP4GS3 delivers an industry-first verified single-chip solution that can handle OC-48c IPv4 packet processing at line rate.

IBM®

Further information

- IBM Microelectronics: http://www.chips.ibm.com
  - PowerNP NP4GS3 network processor specs, documentation, etc.
  - Tolly Group report on IBM PowerNP4GS3 performance
  - NP4GS3 - MDR's processor of the year award
- Network Processing Forum http://www.npforum.org

IBM®