Low Power Network Processor Design Using Clock Gating

Yan Luo, Jia Yu, Jun Yang and Laxmi Bhuyan

Computer Science and Engineering
University of California Riverside

Design Automation Conference
Anaheim, CA, June 16, 2005

Packet Processing in the Future Internet

Future Internet

More packets
&
Complex packet processing

ASIC

General-Purpose Processors

Network Processor

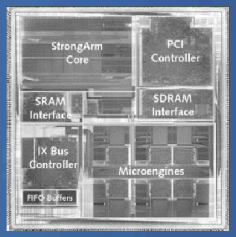

- High processing power
- Support wire speed
- Programmable
- Scalable
- Optimized for network applications
- ...

Typical Network Processor Architecture

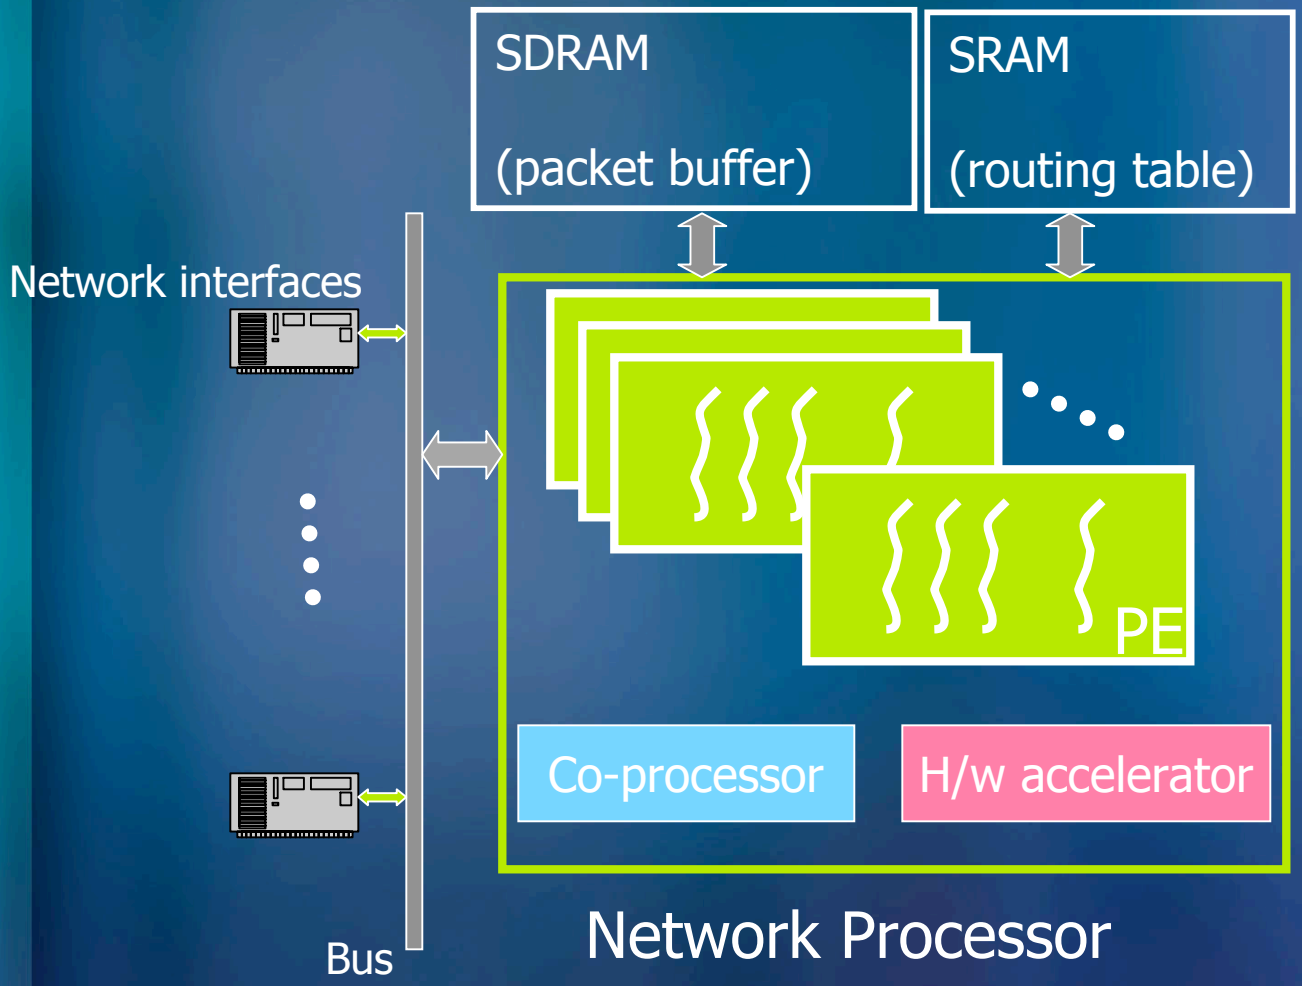

IXP1200: 6 MEs
IXP2400: 8 MEs
IXP2800: 16 MEs

Network Processor Research Overview

- Performance has been the primary interest
  - Throughput
  - Latency
  - Packet loss ratio
- Power efficiency of network processors is becoming a big concern

ST200 edge router from Laurel Networks : 8 NP boards
Each NP board consumes: **95~150W**
2 chassis in a router
The total power - **2700W**
[*Reference: Laurel Networks ST series router data sheet*]

Real-time Traffic Varies Greatly

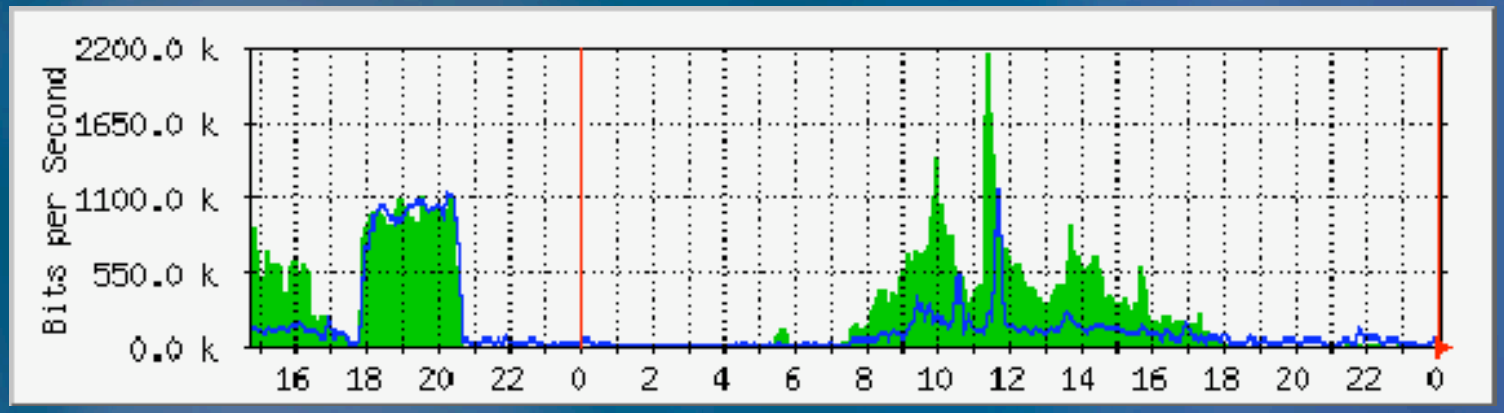

- Shutdown unnecessary PEs, re-activate PEs when needed
- Clock gating retains PE instructions

Indicators of Gating/Activating PEs

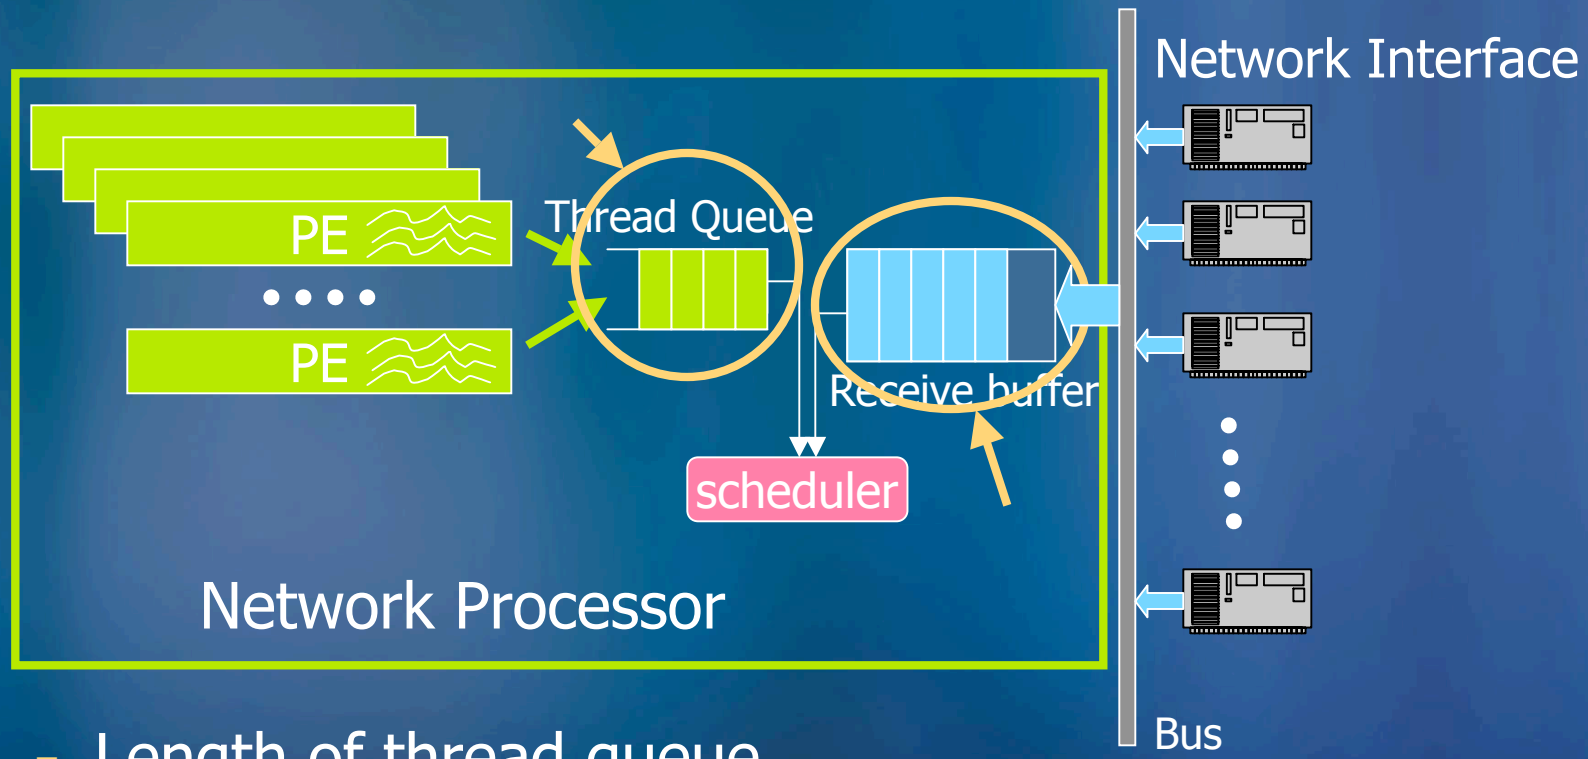

- Length of thread queue
- Fullness of internal buffers

Control Logic to Clock-gate/Activate PE

If (thread_queue_length > T)
  increment counter;

If (counter exceeds threshold)
  { turn-off-a-PE;
    decrement threshold }

If (buffer is full)
  { turn-on-a-PE;
    increment threshold }

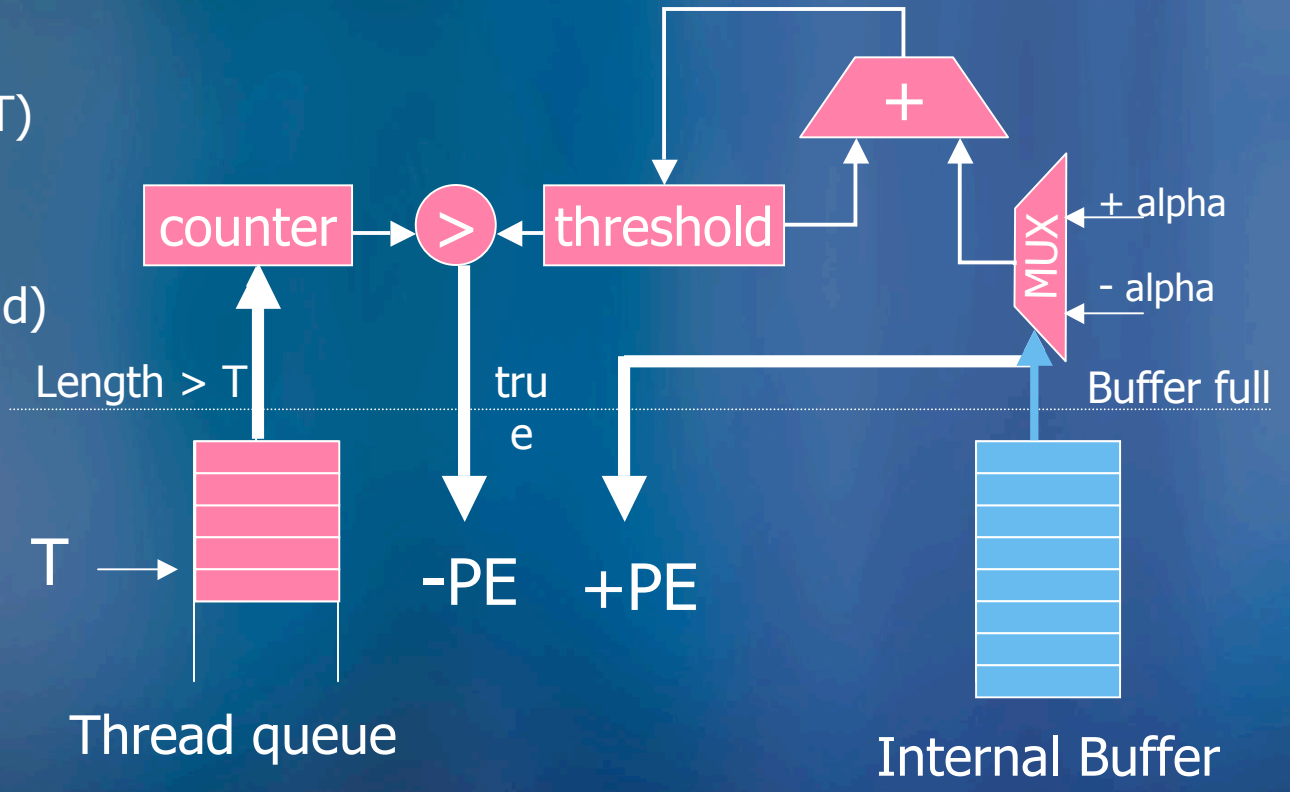

Challenges of Clock Gating PEs

- Terminating threads safely
  - Threads request memory resources
  - Stop unfinished threads result in resource leakage
- Reschedule packets to avoid "orphan" ports
  - Static thread-port mapping prohibits shutting down PEs
  - Dynamically assign packets to any waiting threads
- Avoid "extra" packet loss
  - Burst packet arrival can overflow internal buffer
  - Use a small extra buffer space to handle burst

Dynamic Port Mapping and Extra Buffer

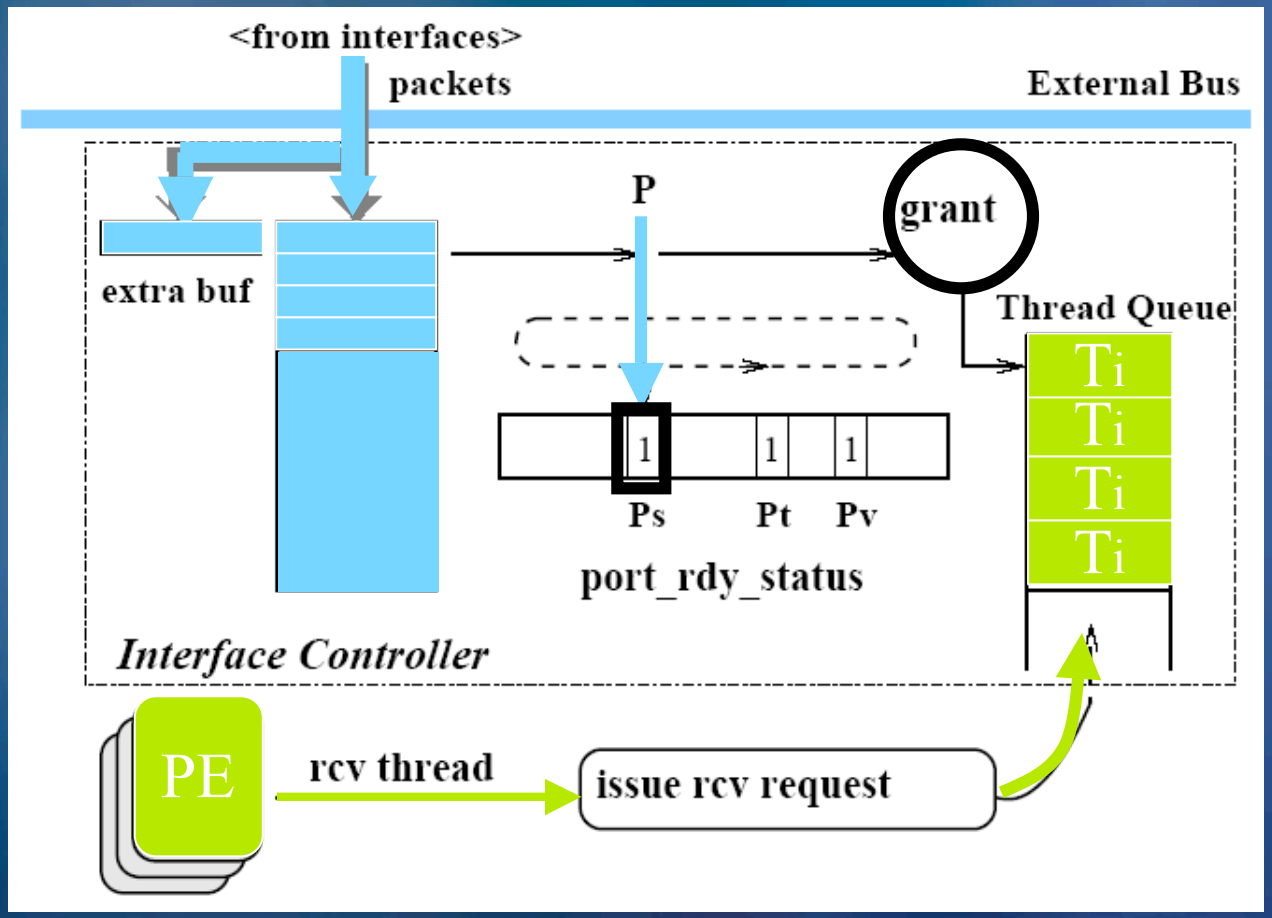

Putting It All Together

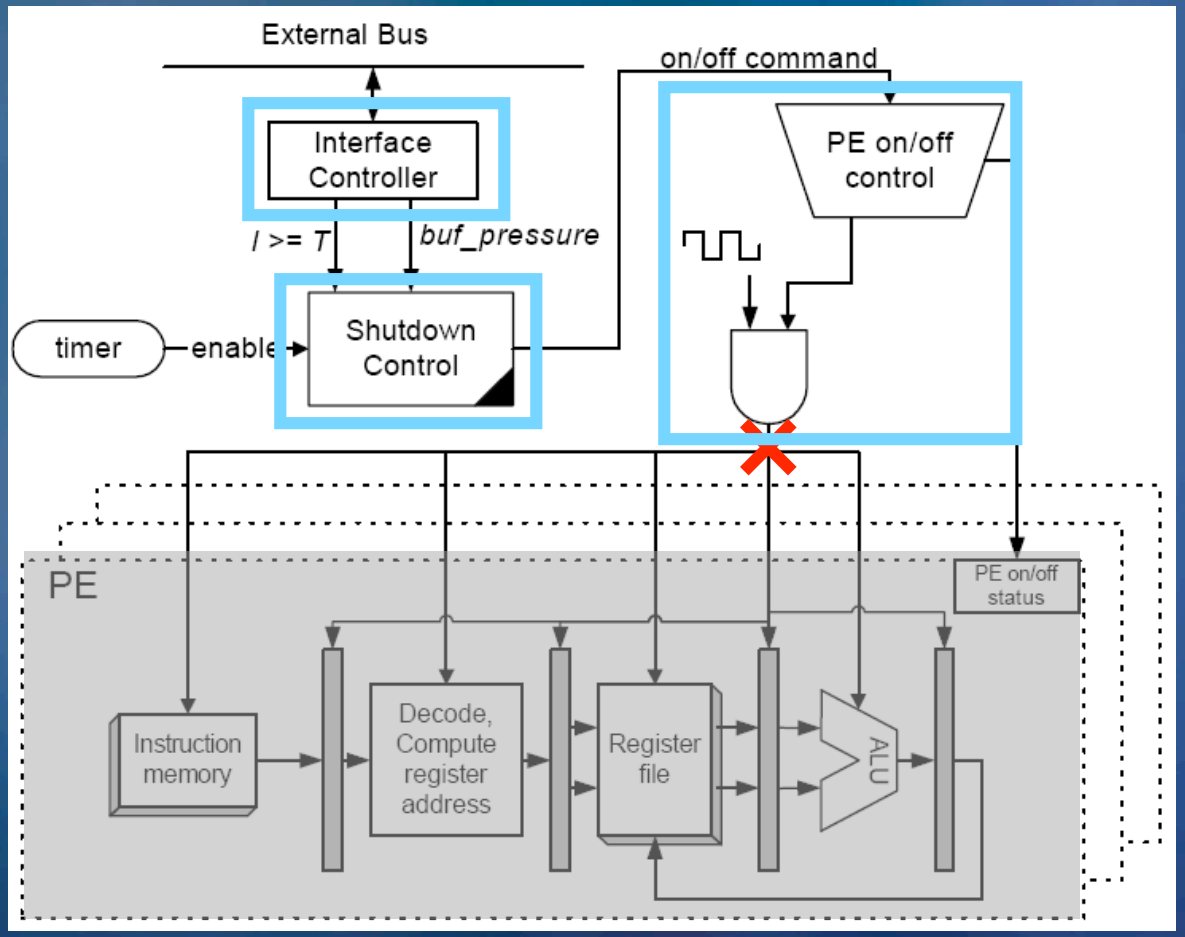

Performance Evaluation Setup

- NePSim – performance and power evaluation tool
- Augment clock power onto NePSim
- Benchmarks
  - Ipfwdr : IPv4 forwarding
  - Nat: Network address translation
  - Url: URL pattern matching
  - Md4: computation of 128b message digest
- Packet traces from NLANR

Power Saving

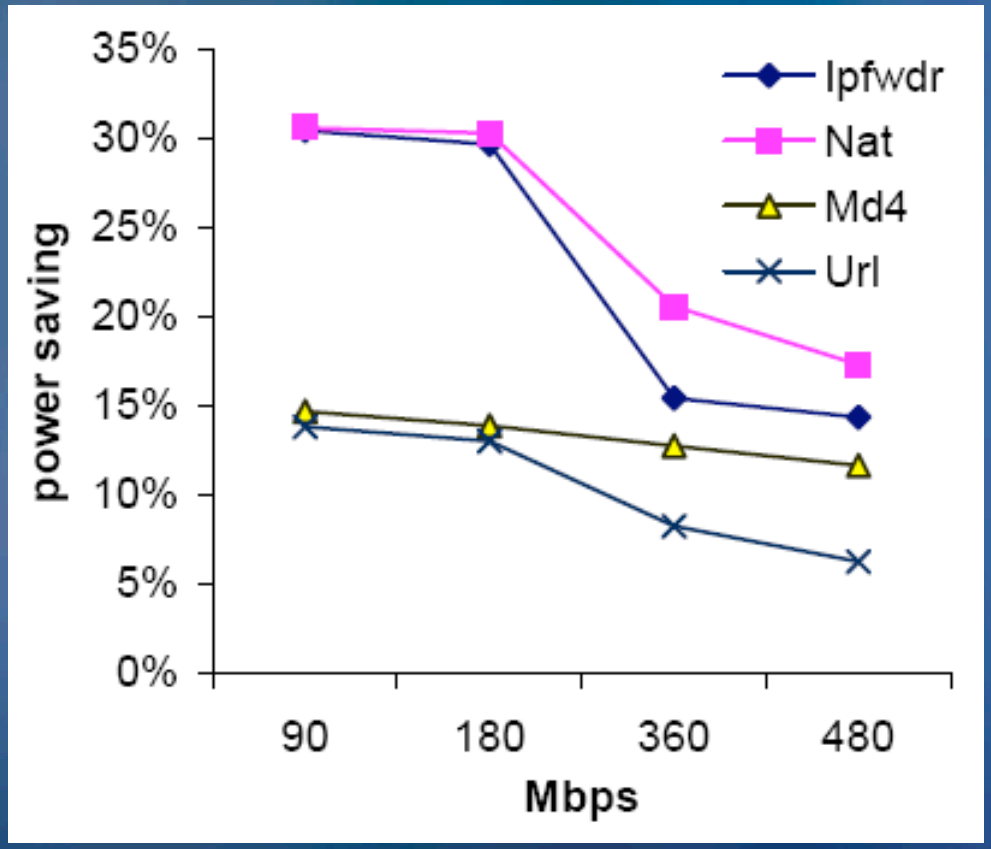

<4% reduction on system throughput

Conclusion

- Power consumption of NPs has become a big concern
- Variation of traffic gives opportunities to save power
- Shutdown/activate PEs dynamically based on traffic load
- Thread queue and internal buffer are effective parameters of PE shutdown/activation policy
- Consider PE termination procedure, PE-port mapping and burst arrival of packets
- Obtain up to 30% power saving with <4% throughput reduction