Building New Networks

Diversified Networking and the Open Network Lab

Patrick Crowley
pcrowley@wustl.edu
http://www.arl.wustl.edu/~pcrowley

Summary

- Diversified Networking: Enabling Net Innovation
- Open Network Lab
- Note
  - Joint work with Jon Turner and others at WU

Diversifying the Net

- Diversification allows many *diverse meta-networks* to co-exist within shared *substrate*.
  - » different meta-nets may differ radically
  - » no requirement that they use the same protocols
  - » enables new network architectures and services to compete
- More than just shared links.
  - » technology now makes it feasible to virtualize routers
  - » network processors and FPGAs deliver flexible packet processing
  - » virtualized switch fabrics can permit sharing with traffic isolation
- Diversification shifts the "waist of the hourglass".
  - » new diversification layer lies just below the traditional IP layer
  - » operates over multiple layer 1 and 2 network technologies
  - » allows multiple end-to-end packet delivery systems
  - » focus of diversification layer is providing resources used by meta-networks

Diversified Networking Concepts

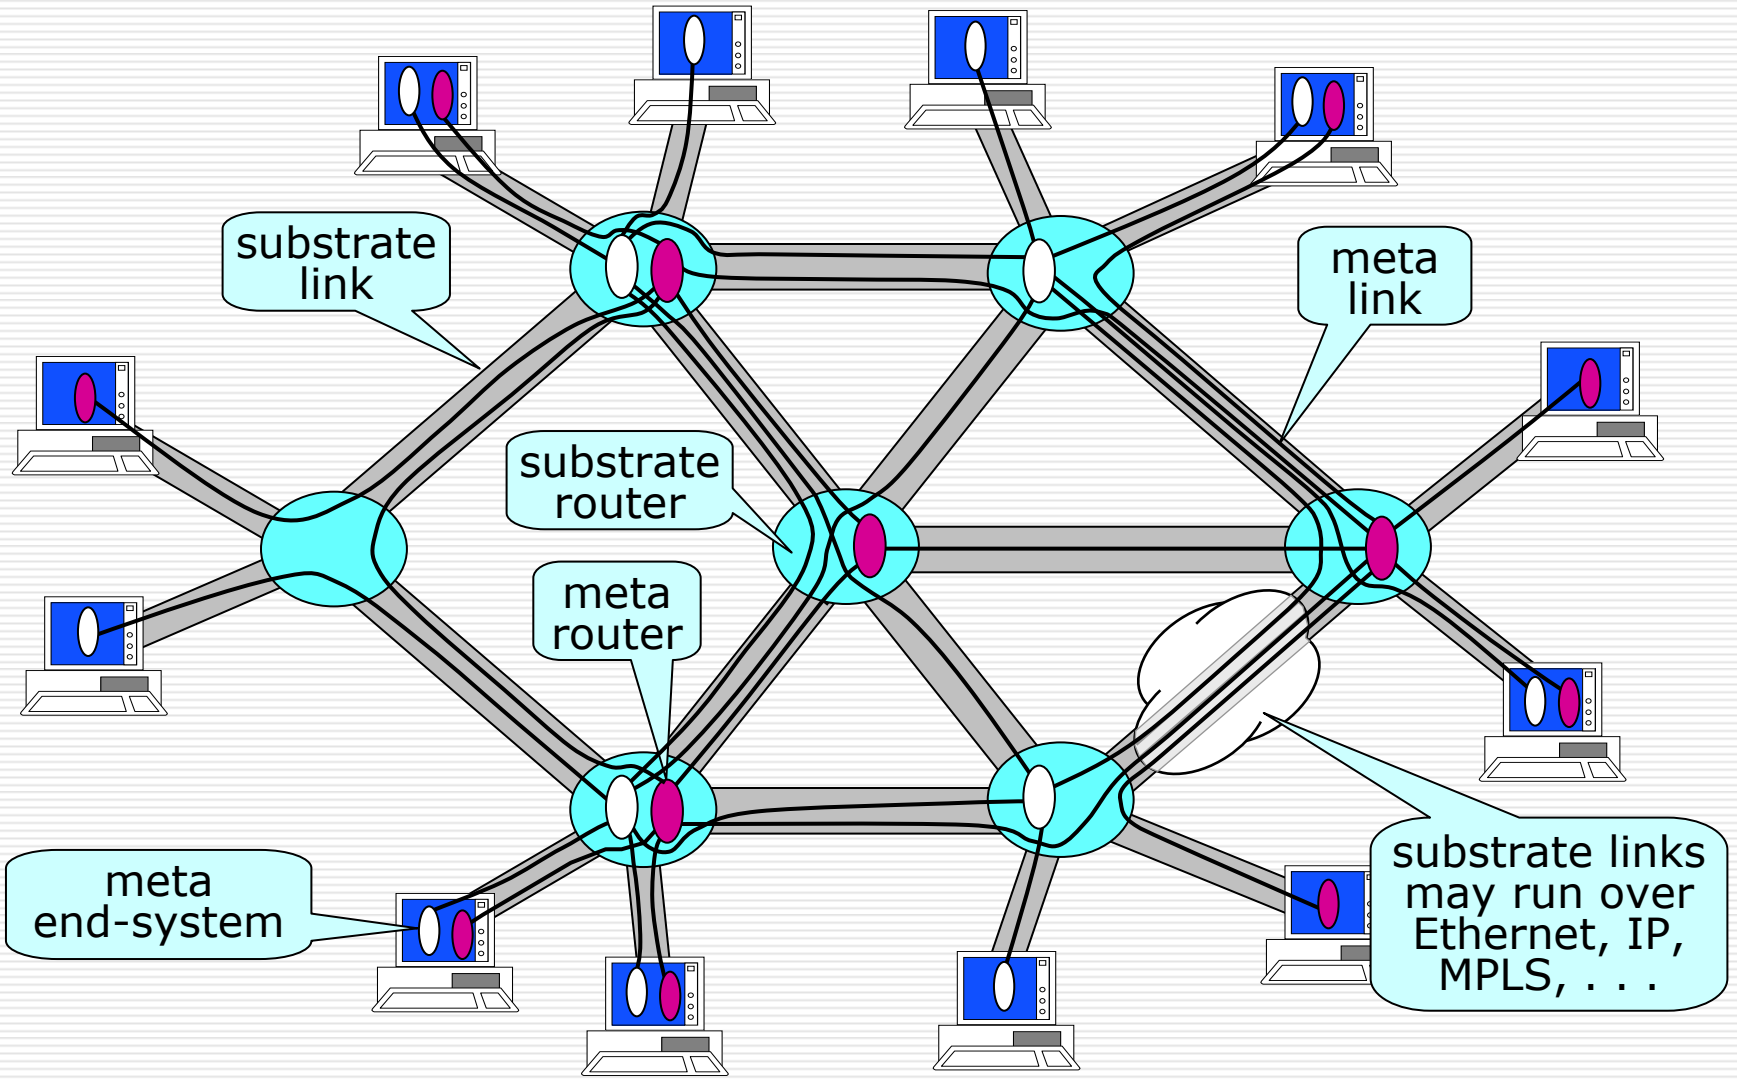

ATCA & NP-based Routers

- Development of ATCA enables use of commercial router components

- Radisys ATCA-7010
  - »10 1 Gbps links
  - »2 Intel IXP 2850 NPs
    - 17 programmable processor cores each
  - »3 high-speed RDRAMs, 3 QDR SRAMS, 1 shared 18 Mb TCAM

Radisys ATCA-7010

- Other cards
  - »Multi-core compute blades
  - »Platform FPGAs
  - »DSPs

Node Architecture

- ATCA Chassis
  - » 14 slot chassis is typical
    - Two slots for switch fabrics
    - 12 slots for processing blades
- General purpose blades.
  - » shared blades run Plab OS
    - no change to current apps
  - » also support dedicated blades
  - » use separate blade server to preserve ATCA slots for NPs
- NP blades.
  - » support dedicated PEs
    - control from Vserver on PE/GP
  - » shared PE options
    - shared NP for fast path
- Line Cards:
  - » 10 1 GE Physical Interfaces
- 10 GE fabric switch
  - » VLANs used to isolate metarouters
  - » uplinks for connecting to multiple chasses
- Good ratio of PEs to LC: 3:1
  - » 9 NP Blades, 3 LC Blades

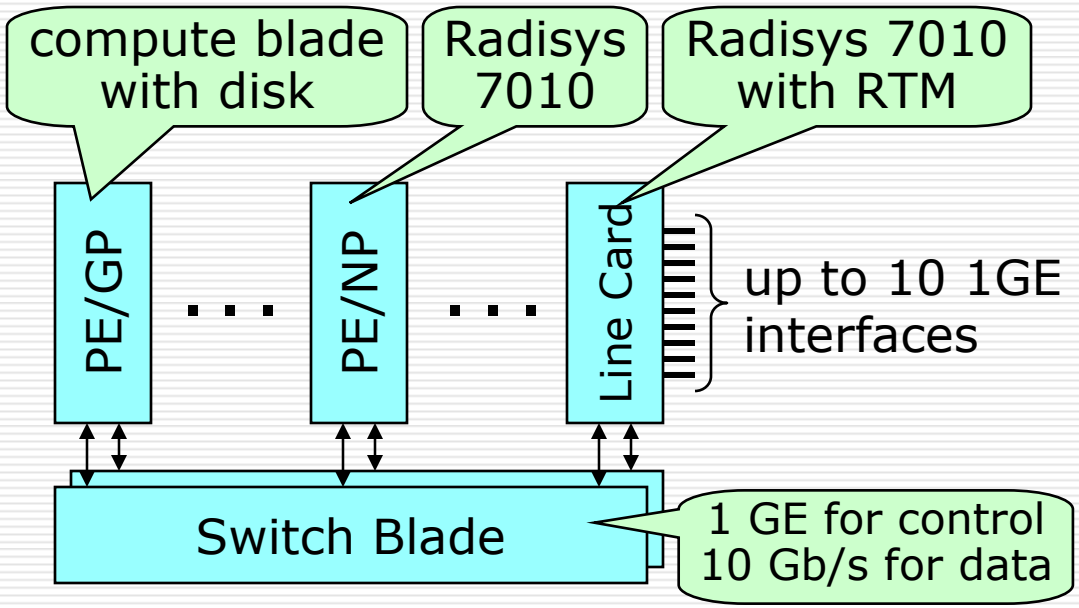

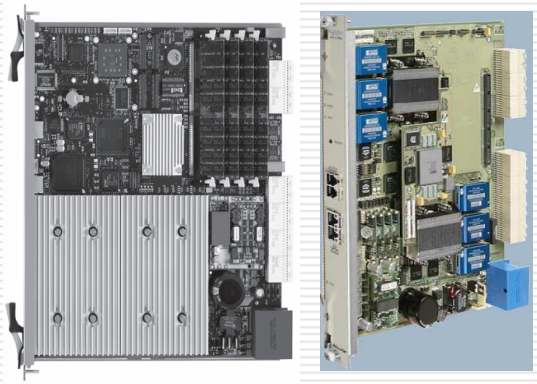

Introduction to ONL

- What is ONL?
  - remotely accessible networking lab
  - gigabit routers with configurable hardware packet forwarding and embedded processors at each port
  - routers can be remotely configured through intuitive GUI
  - extensive support for traffic monitoring/visualization
  - resource for network research community
- Why did we build ONL?
  - difficult to experiment with high performance routers
    - commercial routers are not open
    - open PC routers have limited performance & experiments using them may have limited relevance to high performance routers
  - net research community needs better experimental resources
- What can you do with ONL?
  - evaluate new and existing protocols & apps in realistic testbed
  - add new features to routers (embedded processors, hw mods)
  - mount compelling demonstrations using real-time visualization

Sample ONL Session

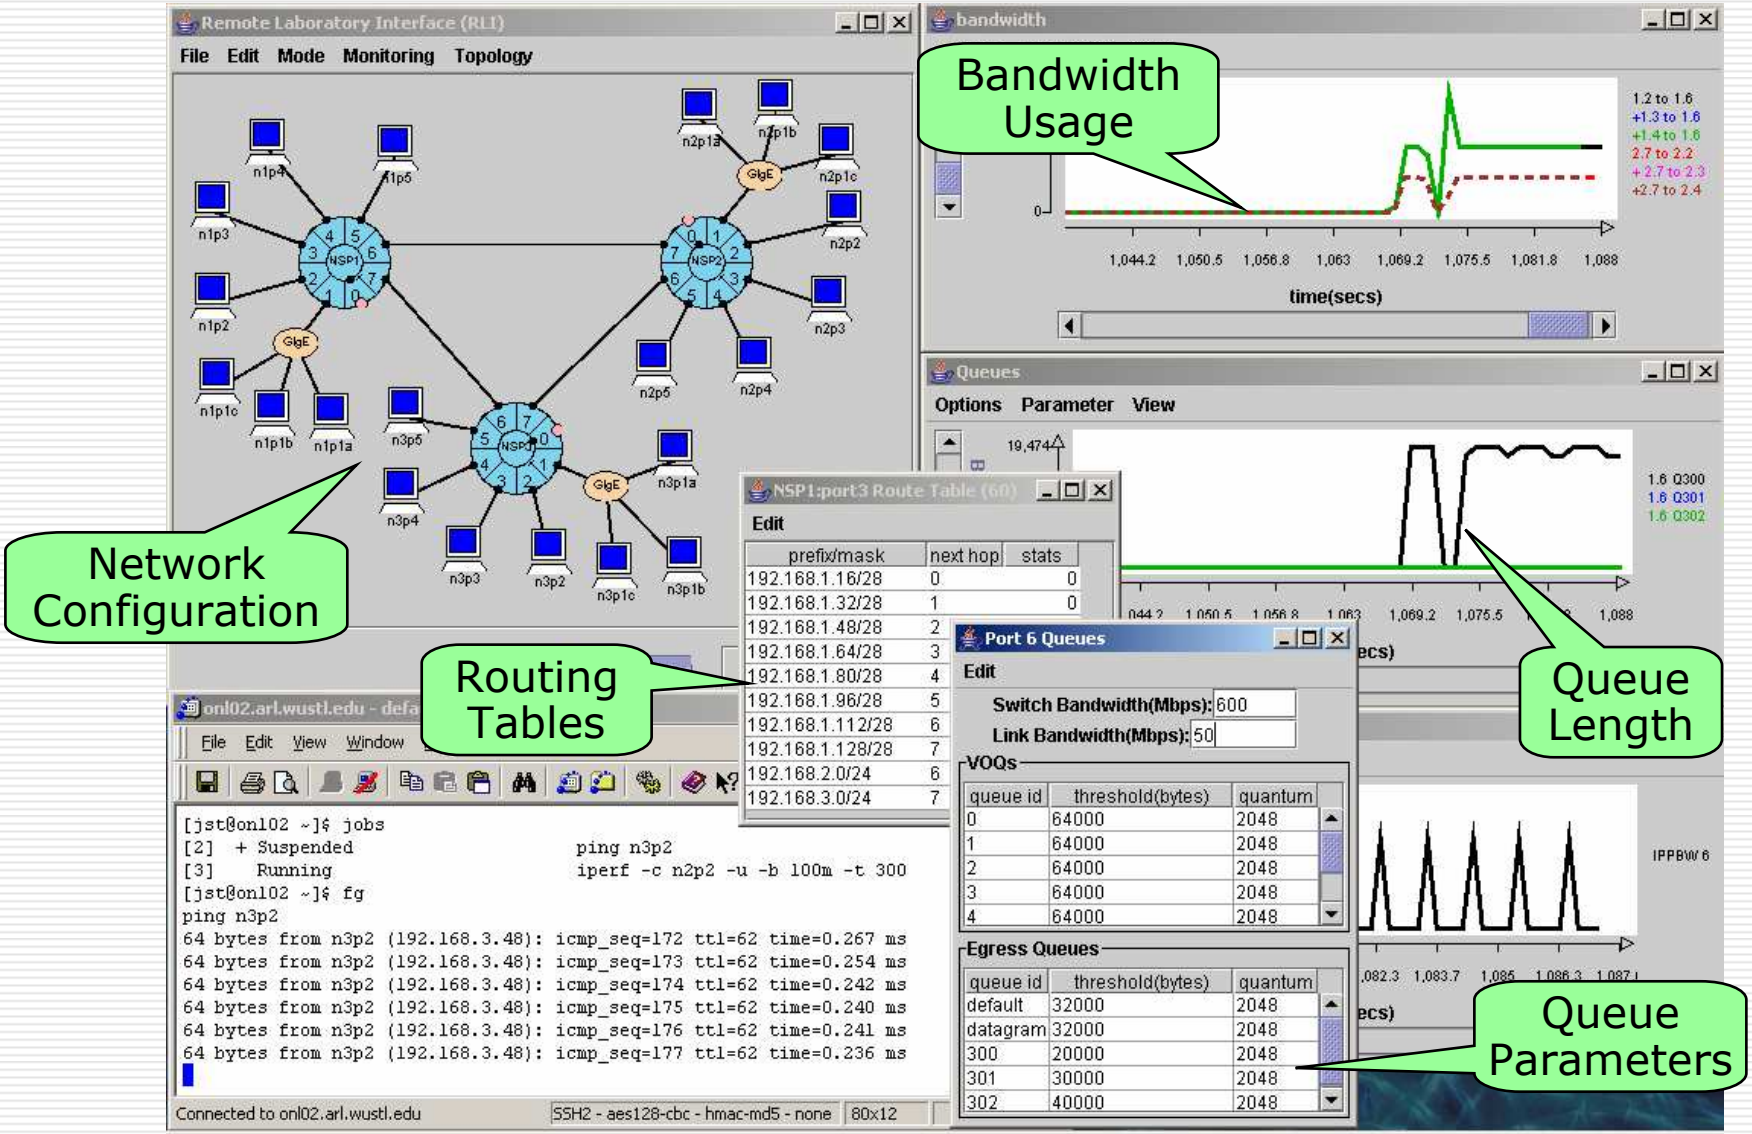

NP-Based Router Design

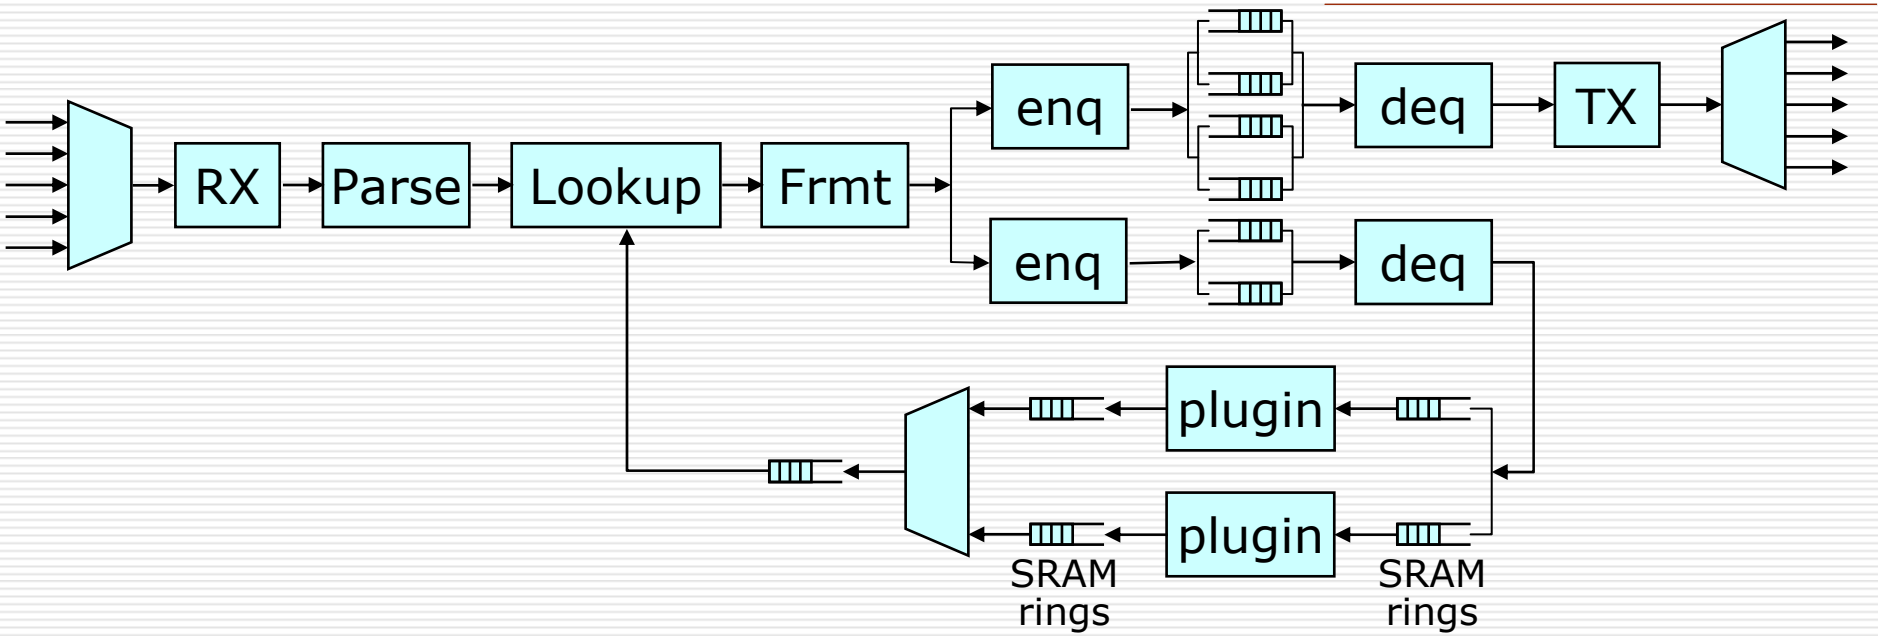

- Dual 5 port router configuration
  - »each NPU handles 5 GE ports
  - »allocate in pairs to maintain isolation between experiments
  - »"excess" micro-engines available for plugins
- Users extend via plugins, or modifying design

Getting Started w/ ONL