Q-Bahn: Scalable and Deployable QoS for the Wired/Wireless Internet

Kihong Park

Network Systems Lab
Department of Computer Sciences
Purdue University

PURDUE UNIVERSITY

Network Systems Lab Projects

Workload Sensitive Traffic Control	Scalable Internet QoS	Network Security

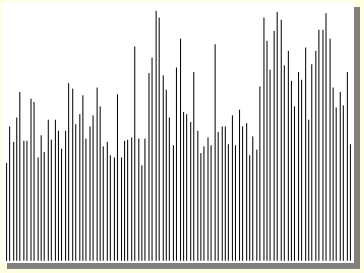

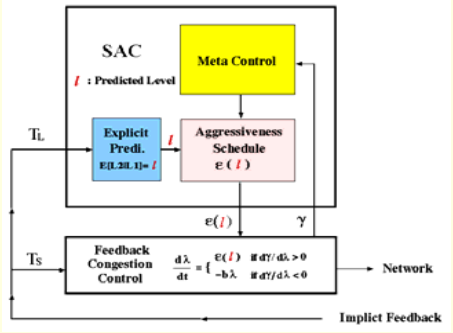

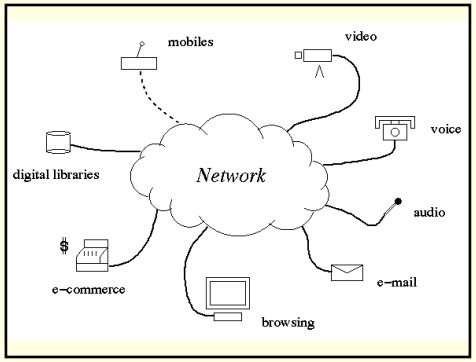

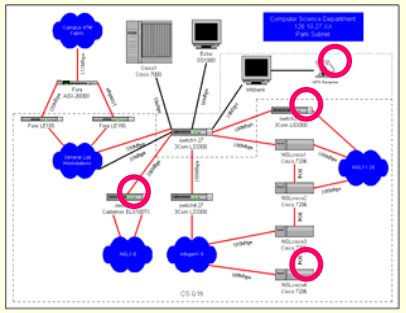

DDoS & Worm

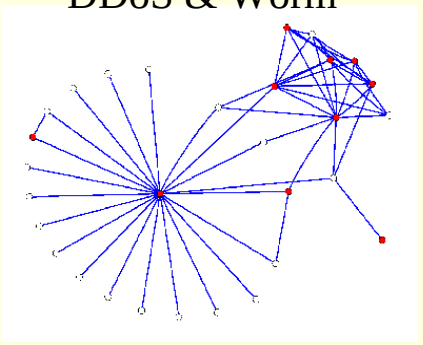

Outline

- NGN challenges

- Q-Bahn approach
  - Foundations
  - Design features
  - Implementation

- Q-Bahn Demo

Challenges: Three Types

- Functional
  - QoS
  - Security
  - Fault-tolerance

- Performance
  - Scalability
  - Deployability

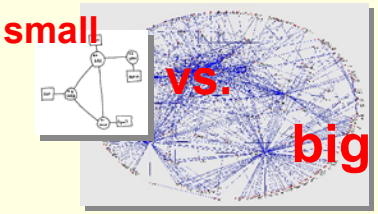

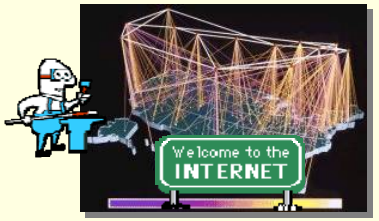

- Organizational
  - Policy barrier
  - Business model

Perspective: End-to-End QoS

End System

- legacy app
- legacy OS
- server, PC, handheld
- CPU ➔
- bw, power, …

Local Access

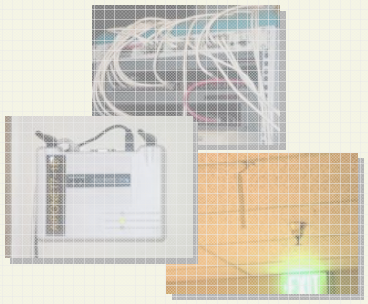

- wireless
- WLAN
- mobility
- access control

Intra-Domain

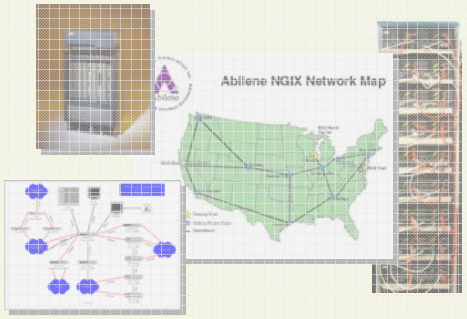

- access router
- core router
- admission control

Inter-Domain

- policy

Perspective: End-to-End QoS

End System

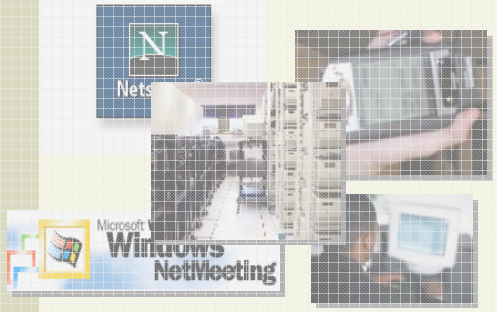

- legacy app
- legacy OS
- server, PC, handheld
- CPU
- bw, power, …

Local Access

- wireless
- WLAN
- mobility
- access control

Intra-Domain

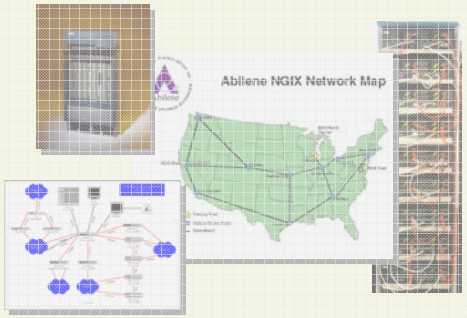

- access router
- core router
- admission control

Inter-Domain

- policy

Perspective: End-to-End QoS

End System

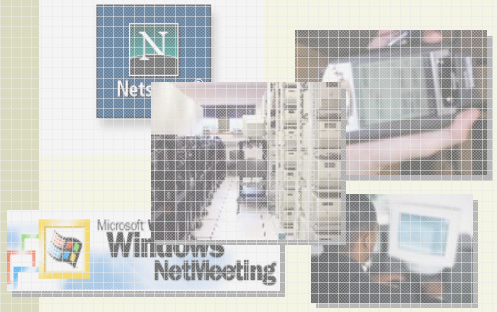

- legacy app
- legacy OS
- server, PC, handheld
- CPU
- bw, power, …

Local Access

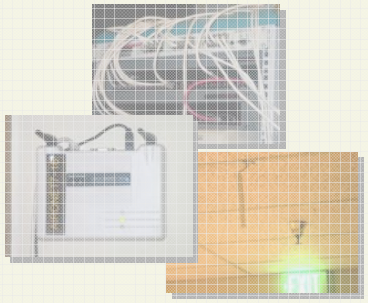

- wireless
- WLAN
- mobility
- access control

Intra-Domain

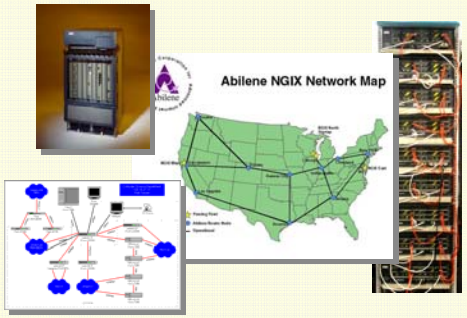

- access router
- core router
- admission control

Inter-Domain

- policy

Perspective: End-to-End QoS

End System

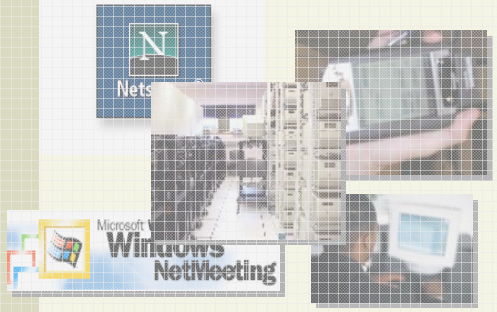

- legacy app
- legacy OS
- server, PC, handheld
- CPU
- bw, power, …

Local Access

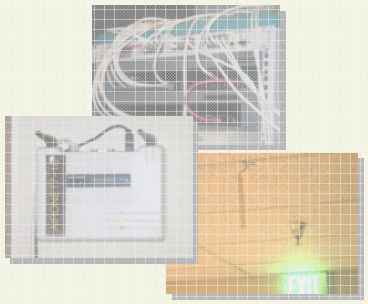

- wireless
- WLAN
- mobility
- access control

Intra-Domain

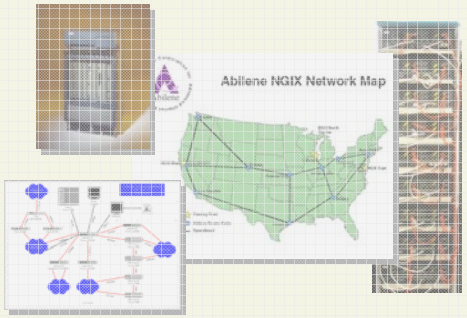

- access router
- core router
- admission control

Inter-Domain

- policy

Perspective: End-to-End QoS

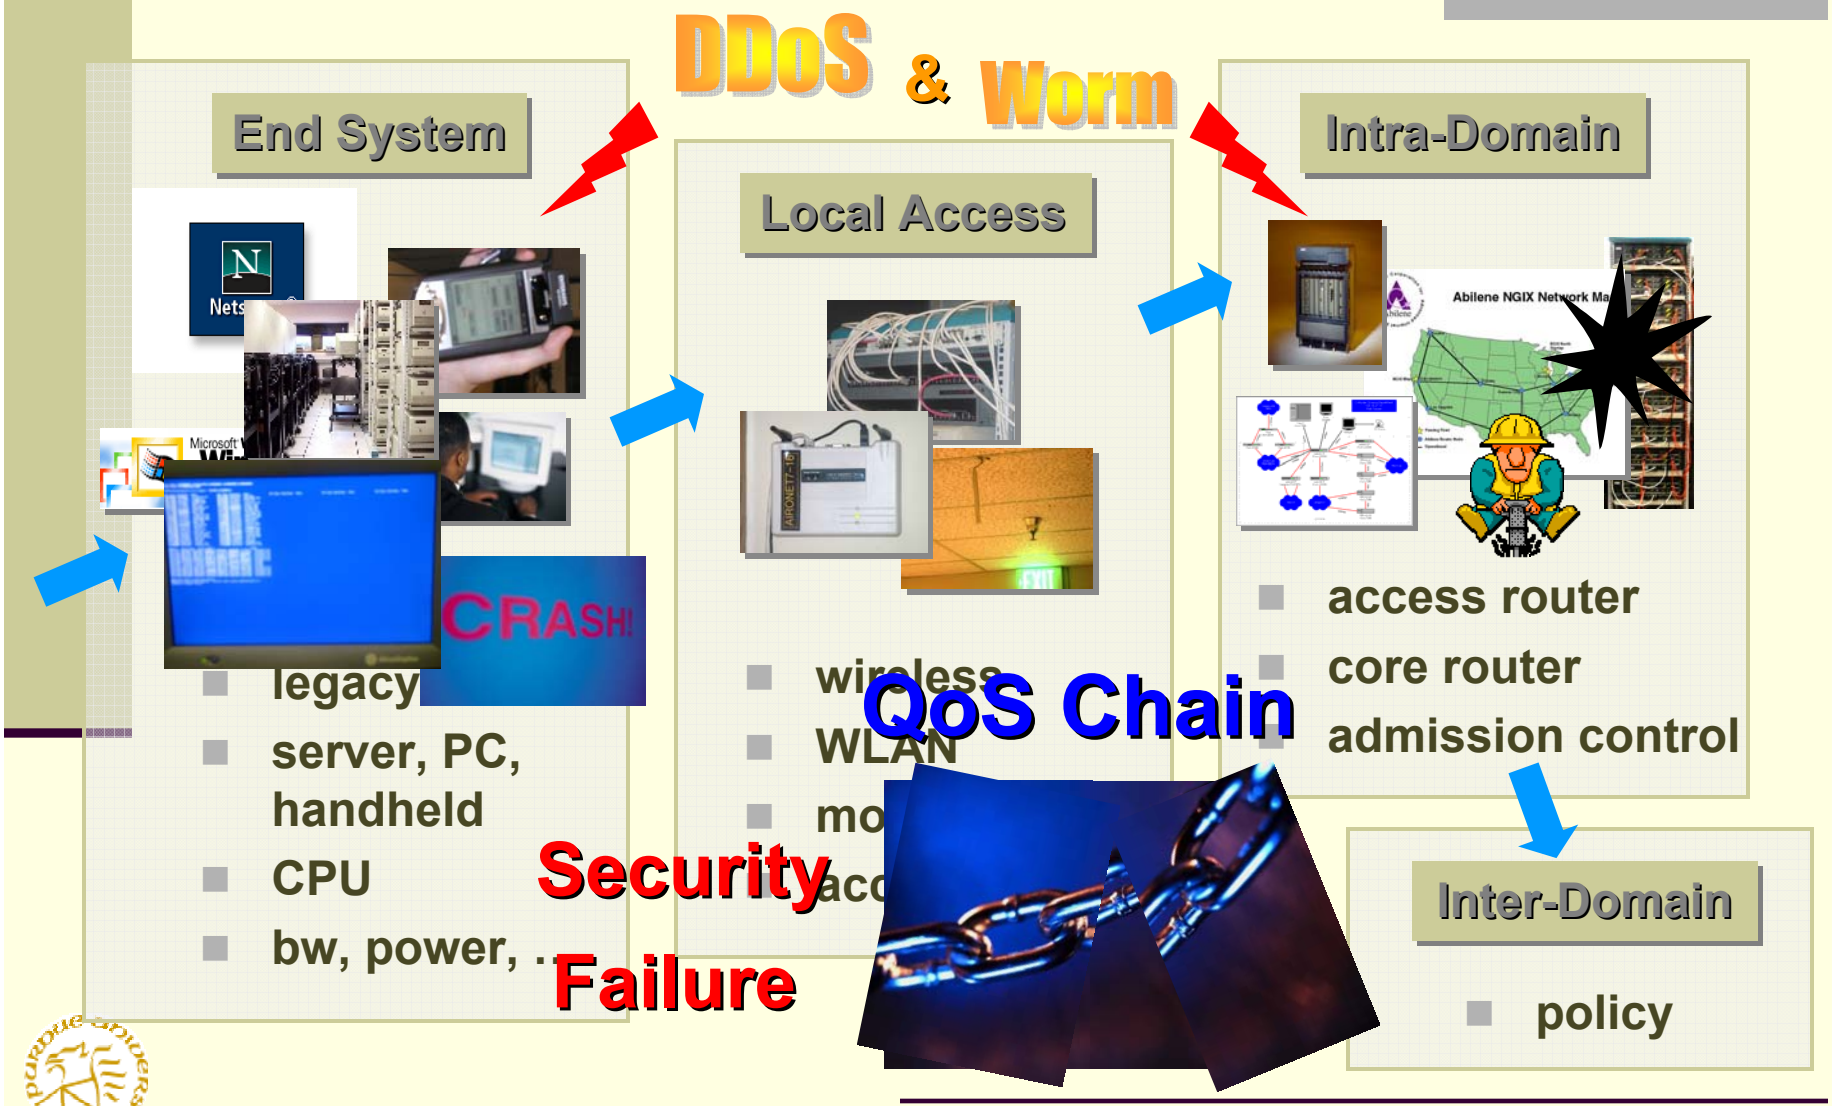

Challenges: Three Types

- Functional
  - QoS
  - Security
  - Fault-tolerance
- Performance
  - Scalability
  - Deployability
- Organizational
  - Policy barrier
  - Business model

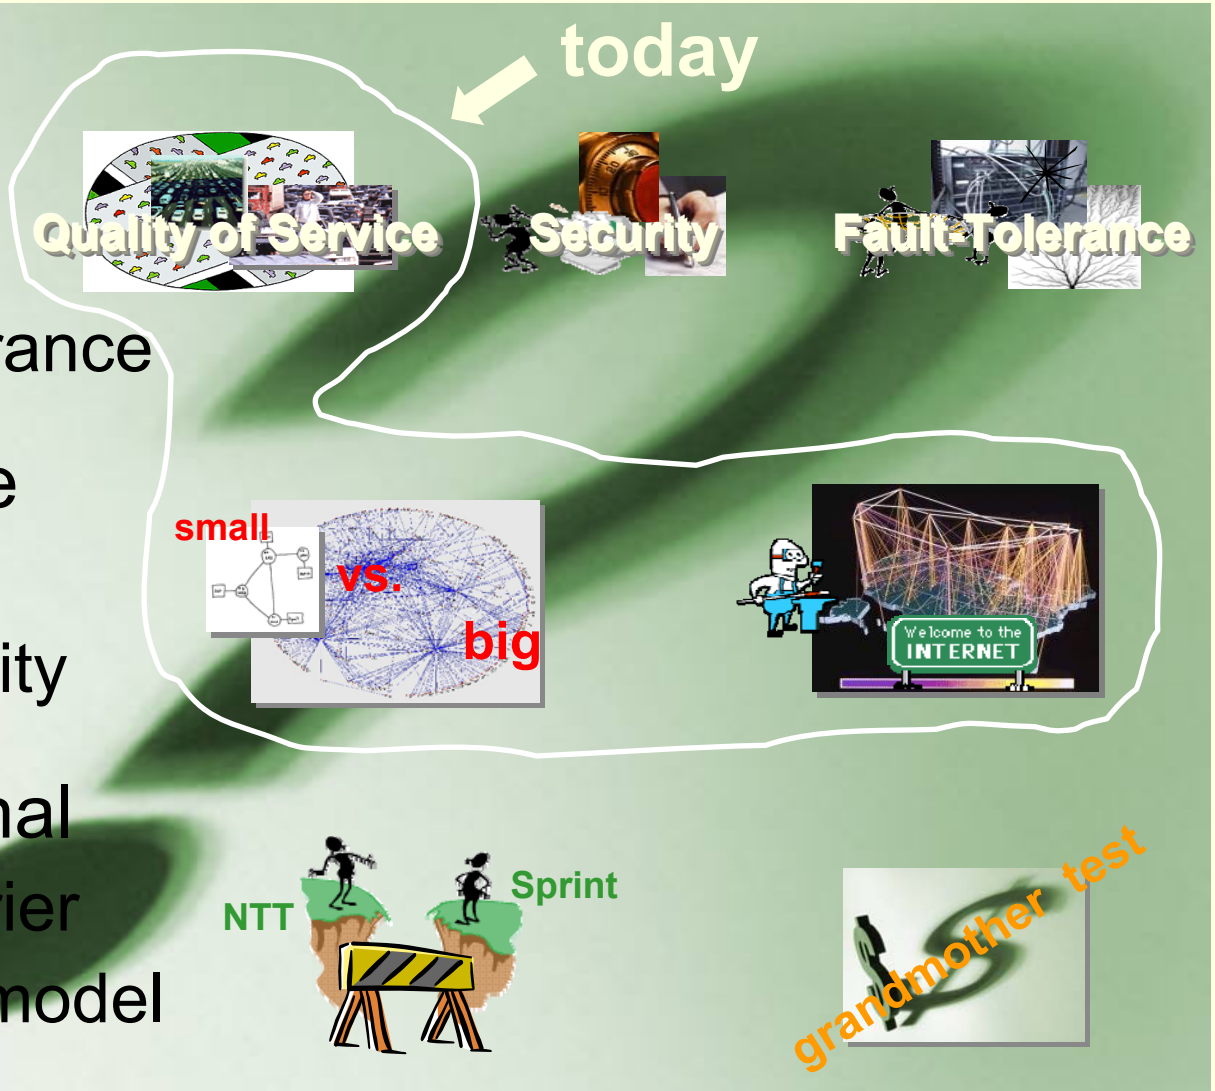

Q-Bahn Approach: Objectives

- Scalability
  - → reservation-less service
    - aggregate-flow scheduling
    - end-to-end QoS control
    - system efficiency
- Deployability
  - → legacy application support
    - backward compatibility
    - extensibility
    - business model

Q-Bahn: Intra-domain enterprise QoS system

Cisco router
Windows/Linux
Test bed

Value Added Service Provisioning

- Q-Bahn Approach:

  → endow value added capabilities to legacy apps

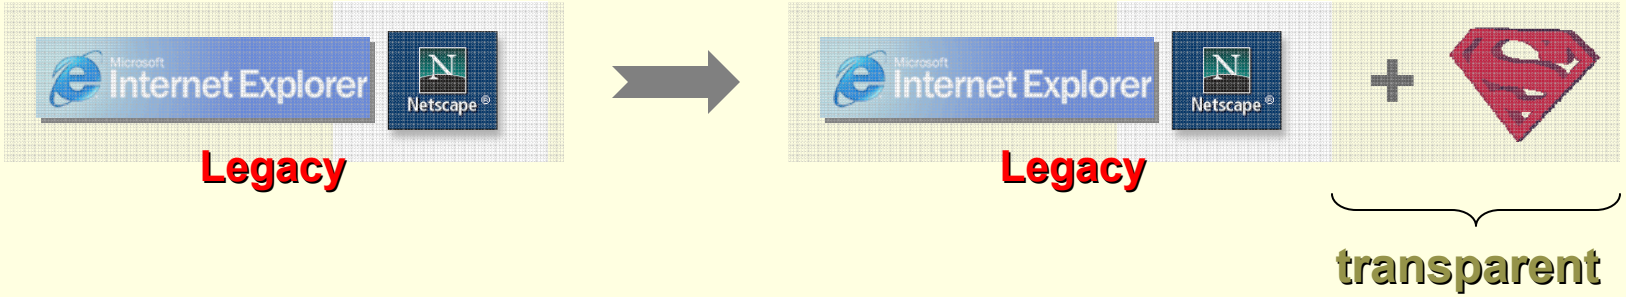

- Foundation for advanced networked services

  → enabling technology for next generation services

Q-Bahn Features

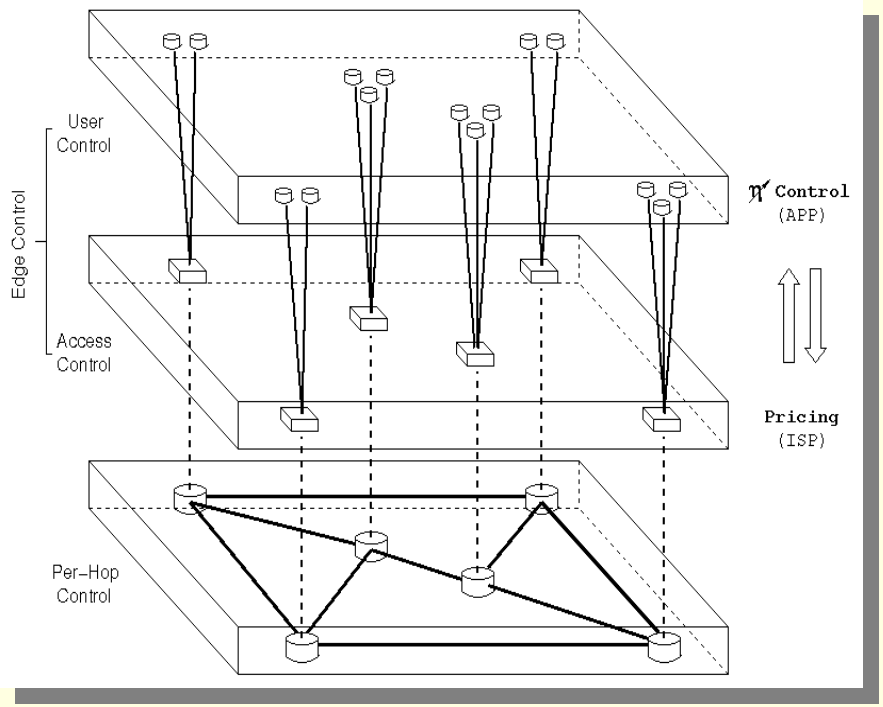

Theory + Systems

- System building
- Workload generation
- Aggregate-flow scheduling
- Game theory
- Traffic control
- Business model

Q-Bahn Features

→ **legacy app & OS**

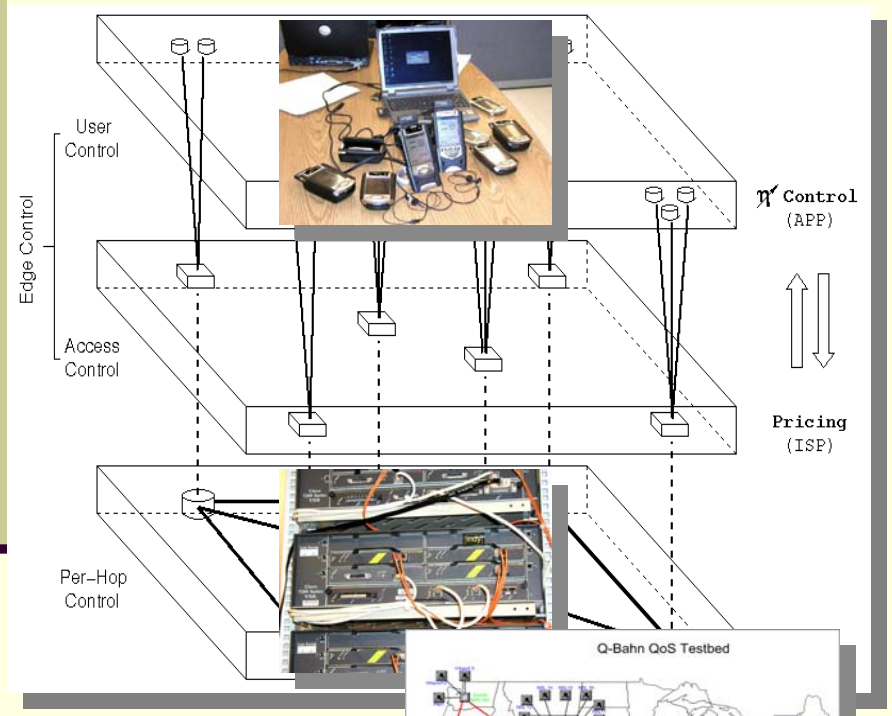

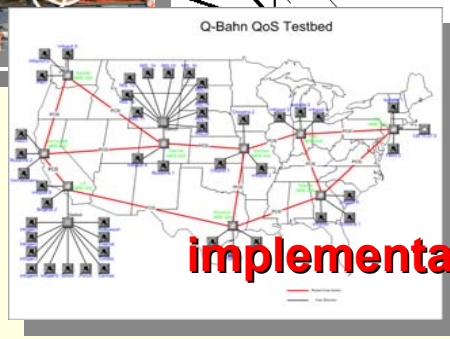

- System building
- Workload generation
- Aggregate-flow scheduling
- Game theory
- Traffic control
- Business model

implementation & testbed benchmarking

Q-Bahn Features

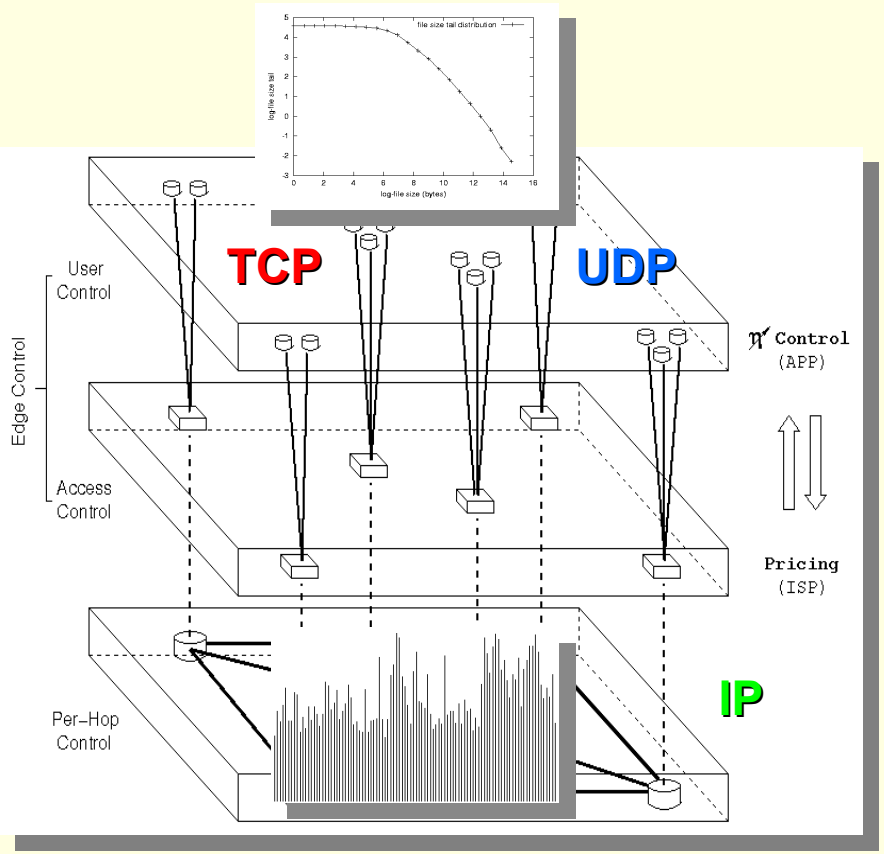

→ **heavy-tailed, self-similar traffic**

- System building
- Work ation
- Aggregate-flow scheduling
- Game
- Traffic control
- Busin

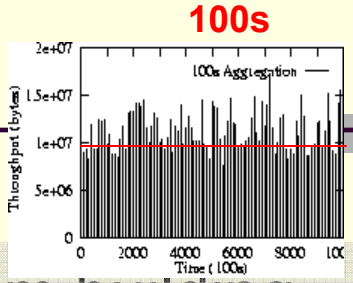

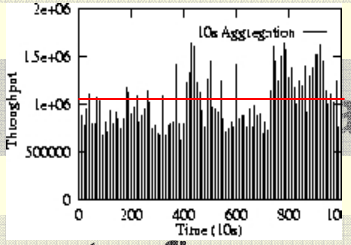

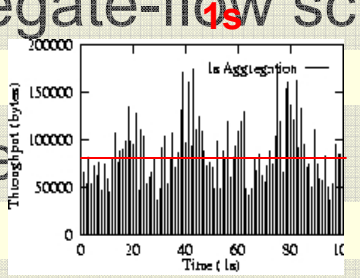

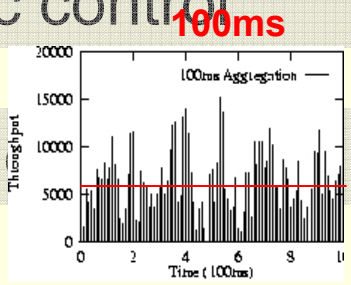

Q-Bahn Features

effective aggregation

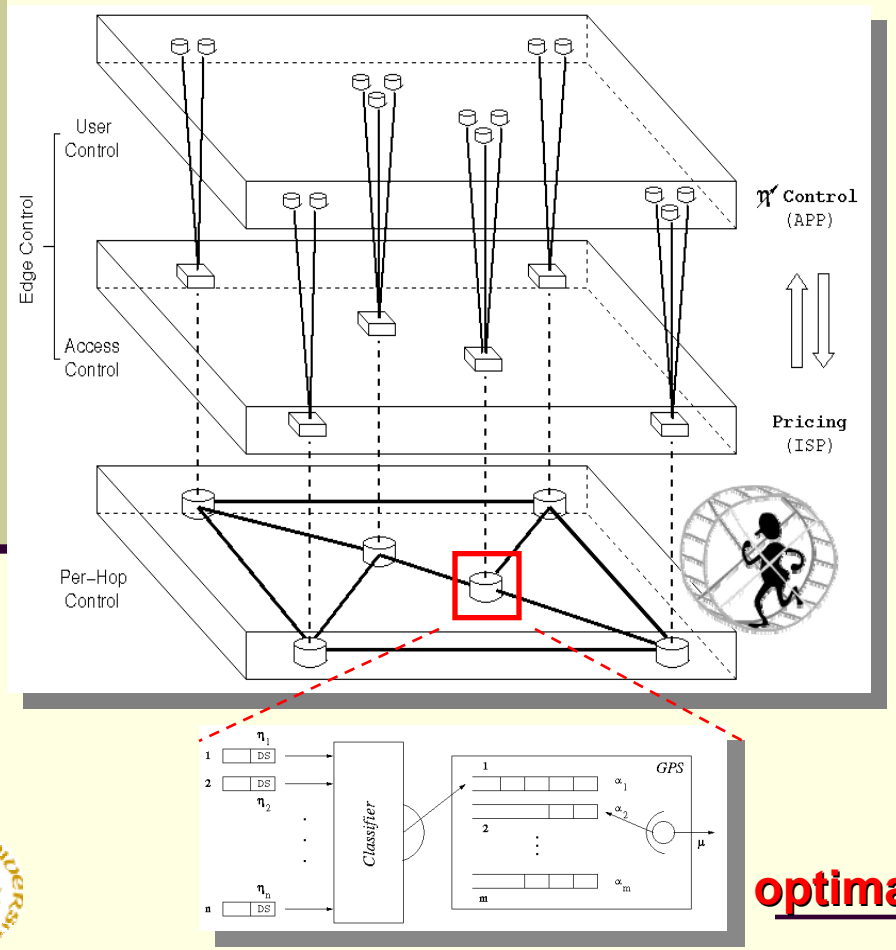

- System building
- Workload generation
- Aggregate-flow scheduling
- Game theory
- Traffic control
- Business model

optimal PHB design

Q-Bahn Features

vs.

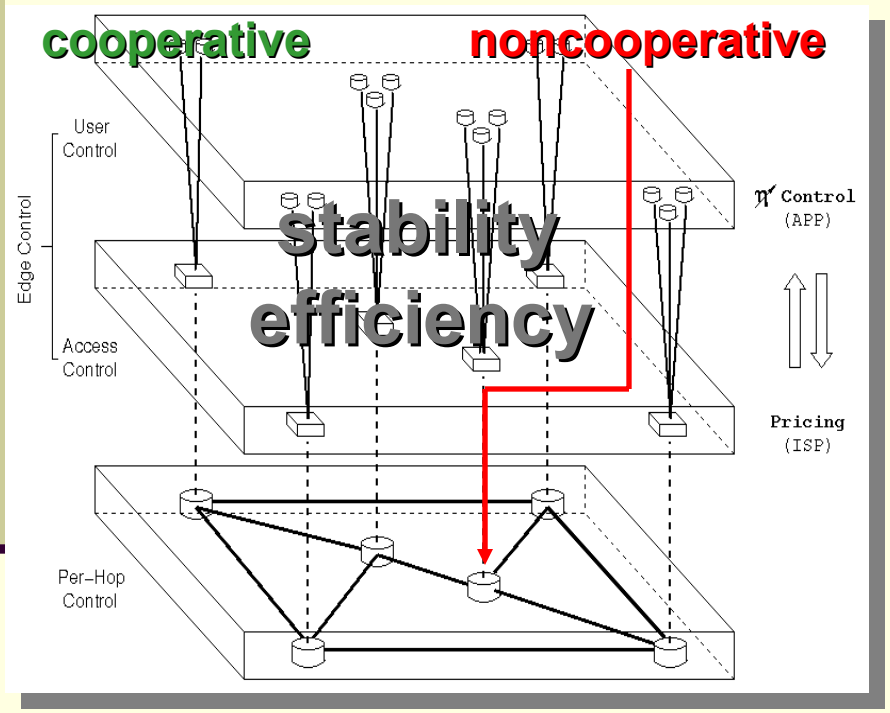

- System building
- Workload generation
- Aggregate-flow scheduling
- Game theory
- Traffic control
- Business model

- **scheduling**
- **pricing**

Q-Bahn Features

closed- & open-loop label control

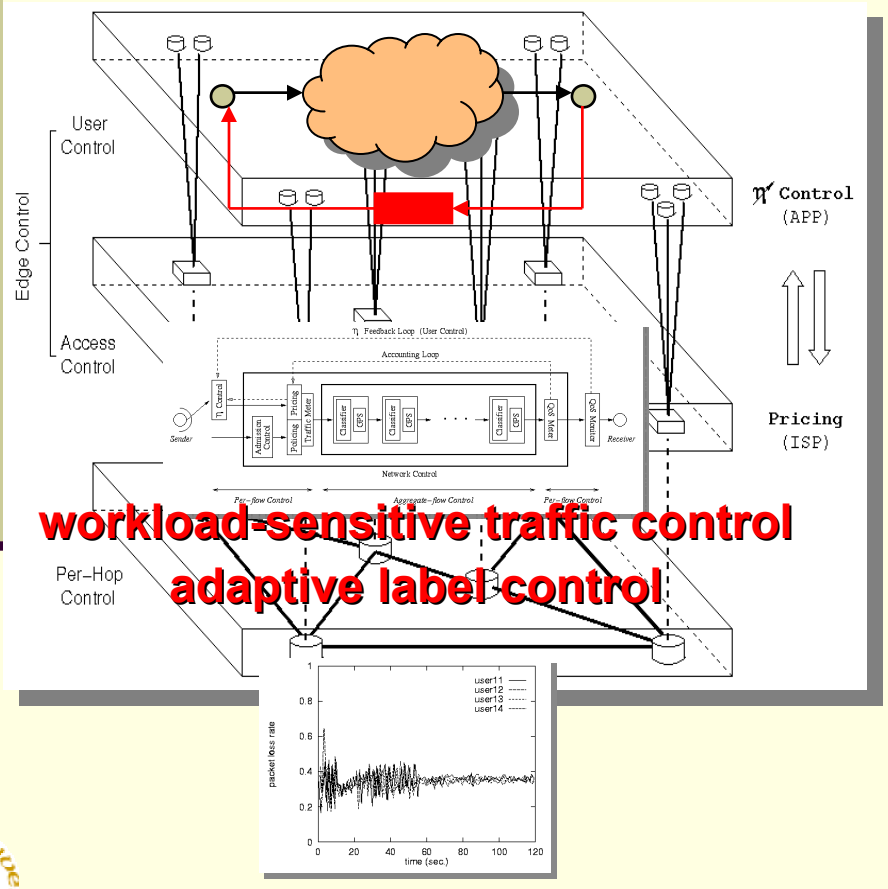

- System building
- Workload generation
- Aggregate-flow scheduling
- Game theory
- Traffic control
- Business model

Q-Bahn Features

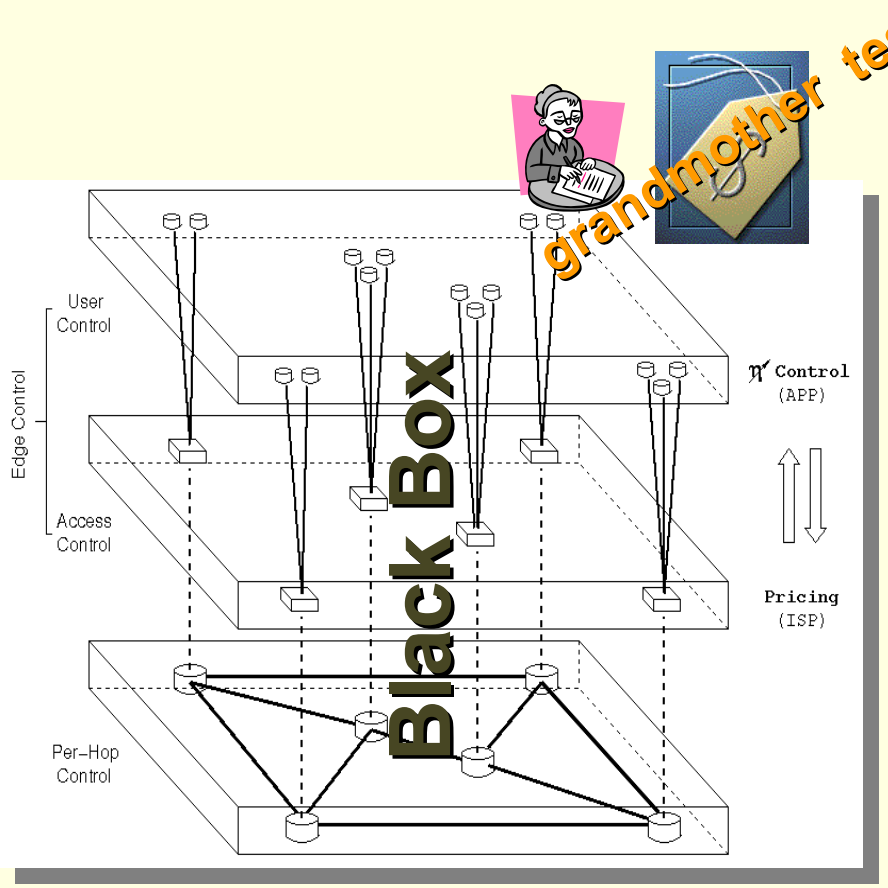

- System building
- Workload generation
- Aggregate-flow scheduling
- Game theory
- Traffic control
- Business model

System Building & Benchmarking

Router/Switch || End System || Testbed

Router QoS Control

- Optimal PHB implementation

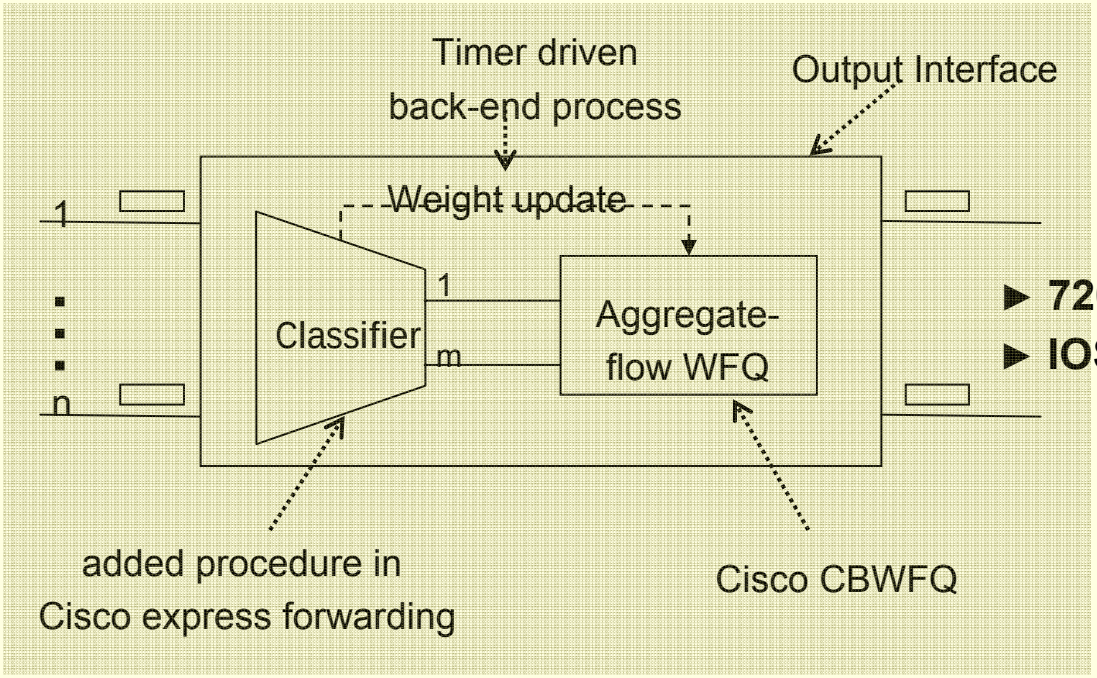

- **7200 series backbone routers**
- **IOS 12.2**

- Joint work with F. Baker, S. Kweon, G. Reitsma

QoS Switching Performance

- QoS space
  - ► Cisco 7206 VXR NP400
  - ► IOS 12.2 purdue-phb

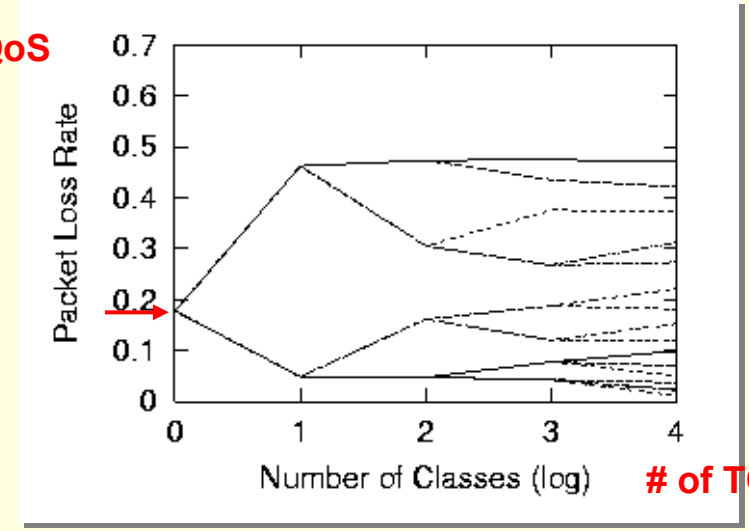

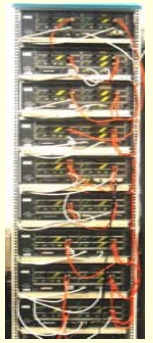

- Footprint
  - ► processing overhead
  - ► optimal vs. EF, AF, FIFO

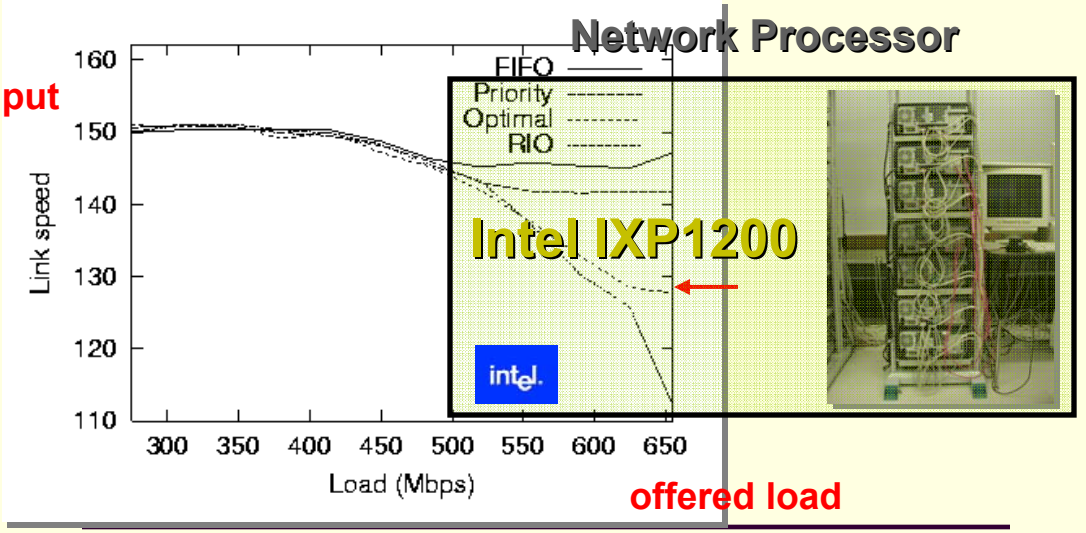

End-to-End QoS Control

- End system: host
  - Adaptive label control
  - Mark IP TOS field to achieve target QoS
    - → e.g., 4-bit TOS field for 16 classes

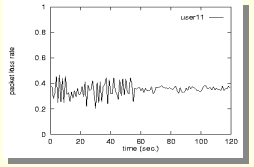

- Key feature: end-to-end QoS control
  - Open-loop & closed-loop control
  - Admission control: access control & accounting
  - Unified QoS currency

End-to-End QoS Control

- Integrated QoS control
  → single QoS currency

- CPU scheduling
- End-to-end QoS control
- QoS routing
- IP packet scheduling
- Wireless MAC

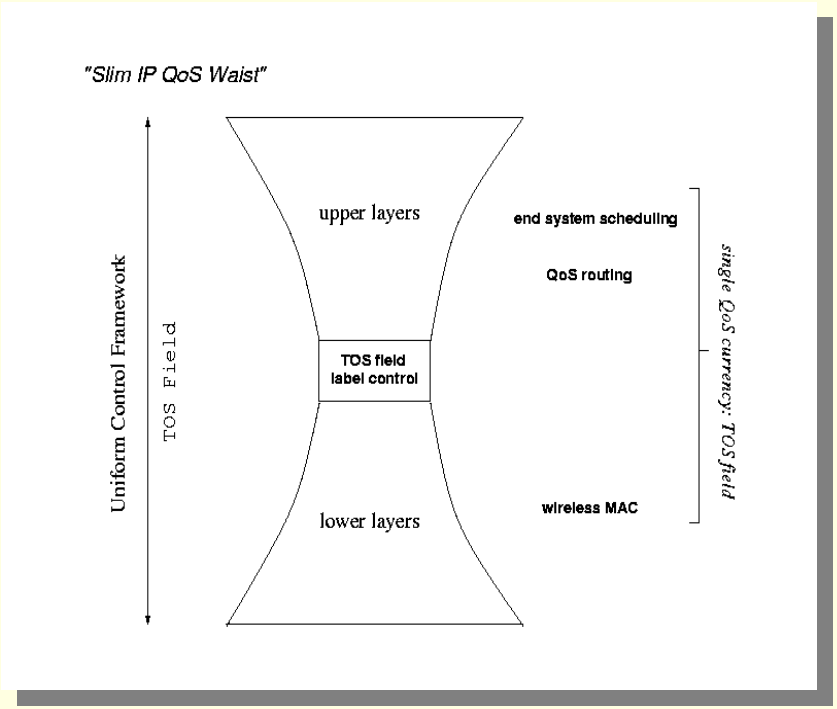

End System QoS Control

- Legacy application QoS support

  → deployability: Q-Port

  - Q-Interface
  - Q-Manager
  - Q-Driver
  - Q-Policy
  - Q-Measure

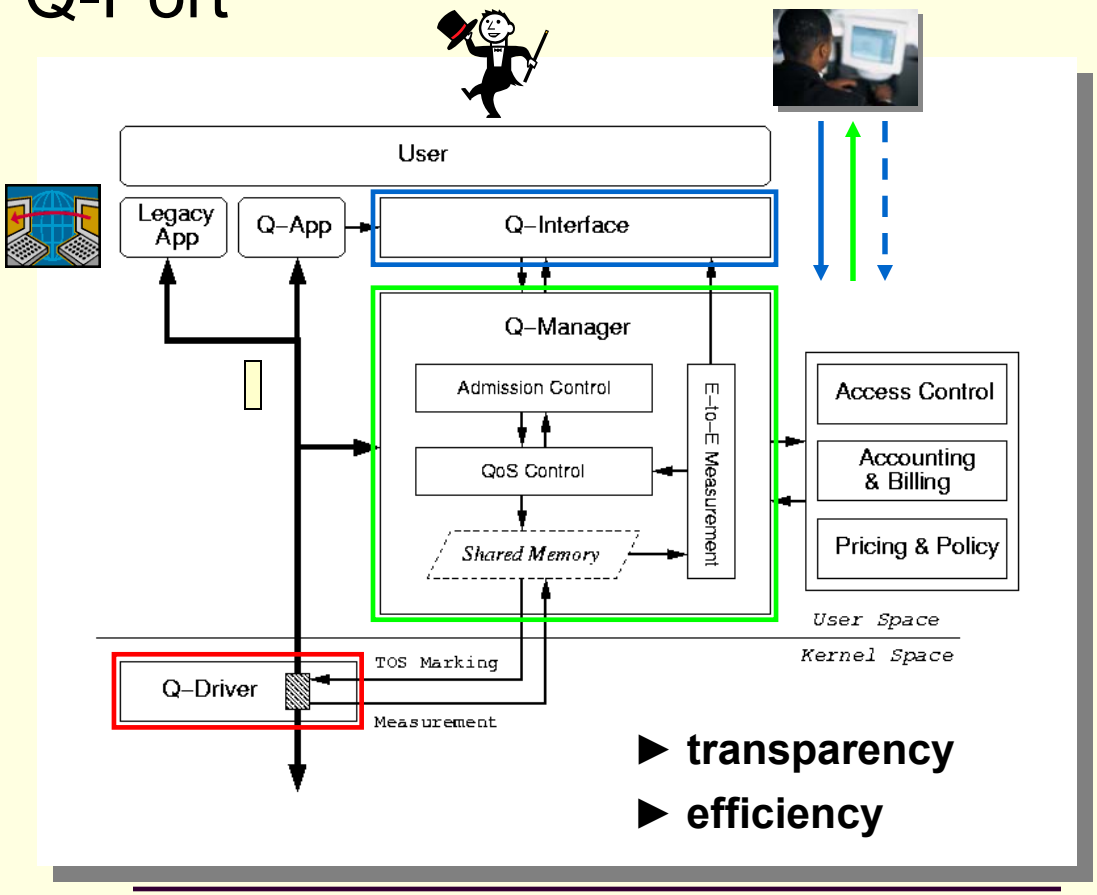

End System QoS Control

- Q-Port implementation design
  - Windows XP and CE
    - → Q-Driver installation in NDIS

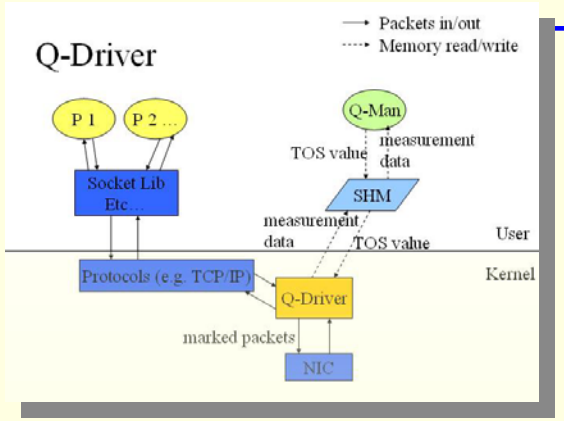

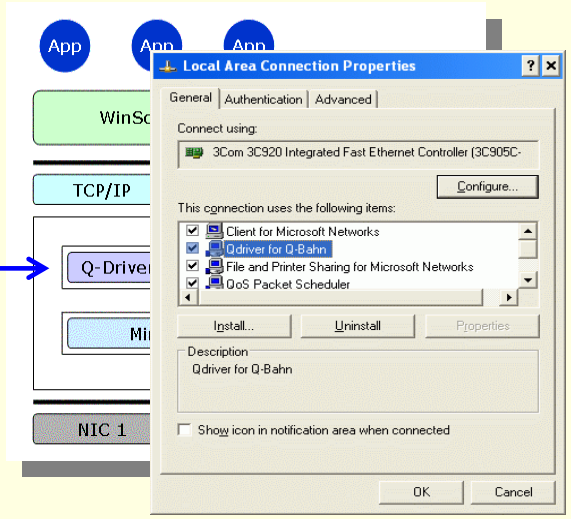

  - Linux
    - → netfilter (Linux 2.4+)
    - → dynamically loadable Q-Driver

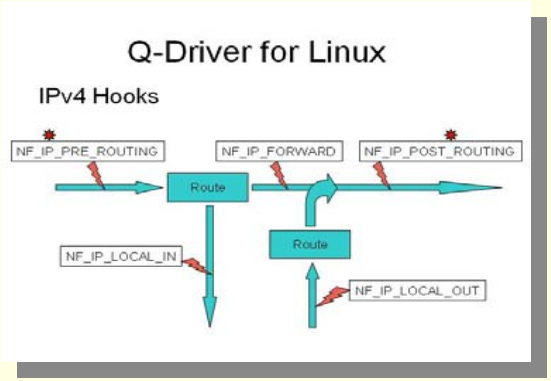

Q-Bahn Testbed

- Physical system: network core
  - 9-node IP-over-SONET backbone
  - Cisco 7206 VXR routers

Q-Bahn Backbone

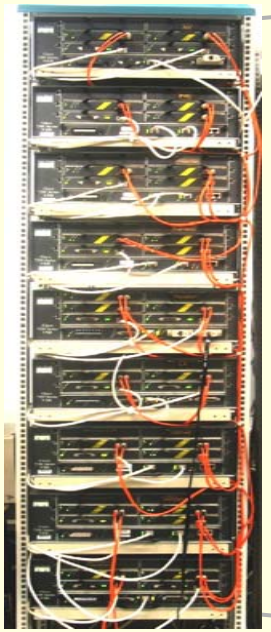

Abilene/Internet2 Connectivity

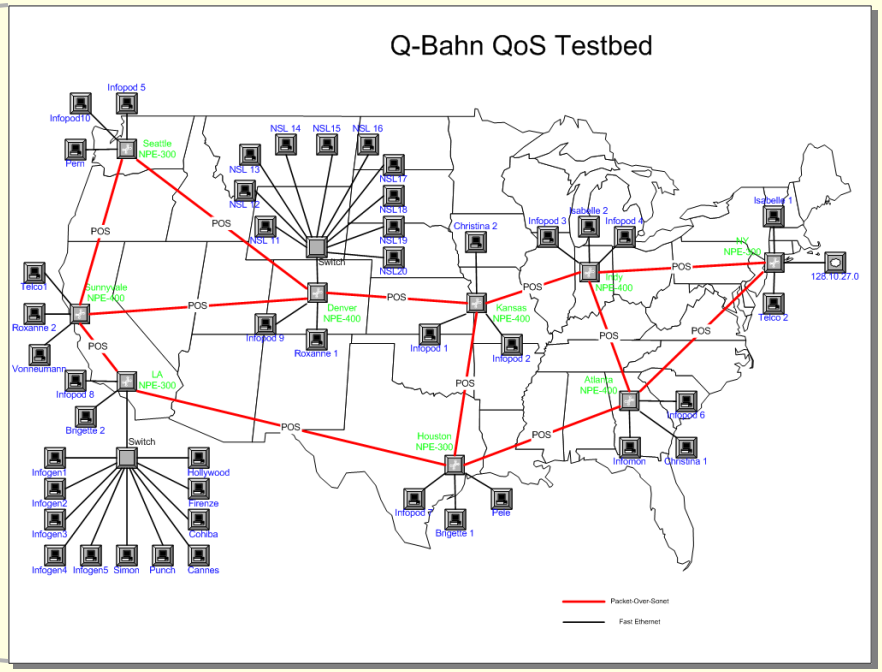

Q-Bahn Testbed

- Physical system: network core
  - Cisco 7206VXR routers: custom IOS
  - purdue-phb: implements optimal aggregate-flow scheduling

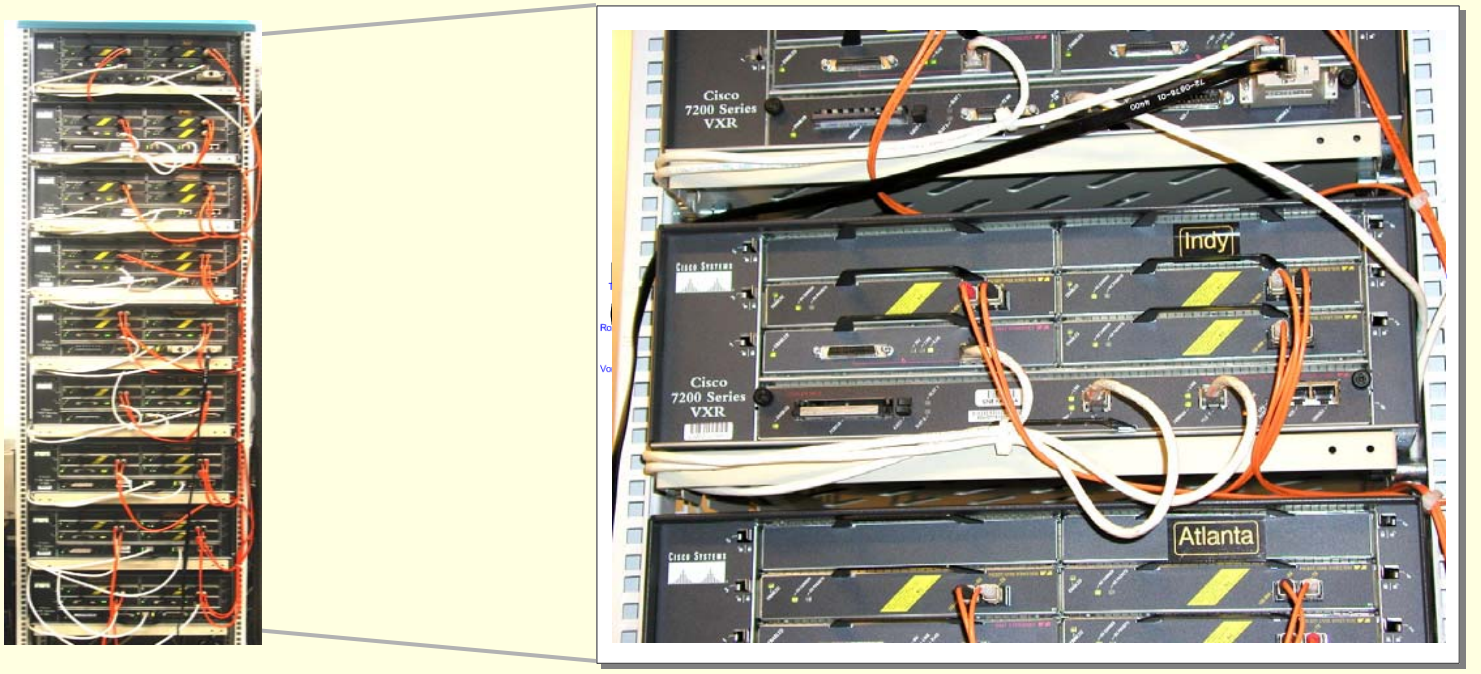

Q-Bahn Testbed

- Physical system: end system
  - Workstations, PCs, labtops, handhelds running Linux and Windows
  - Transparent end system QoS support: Q-Port

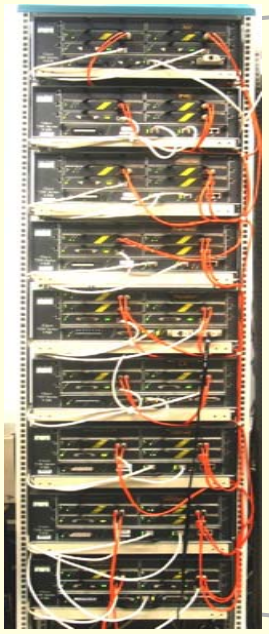

Experiment: Demo

- Application domain: real-time cable TV CDN
  → also VoD CDN, VoIP and teleconferencing

- Legacy application
  - Client: NetMeeting, OpenPhone
  - Server: OpenMCU

► **multi-threaded**

- Legacy OS
  - Windows XP
  - Linux 2.4+ with netfilter

Experiment: Demo

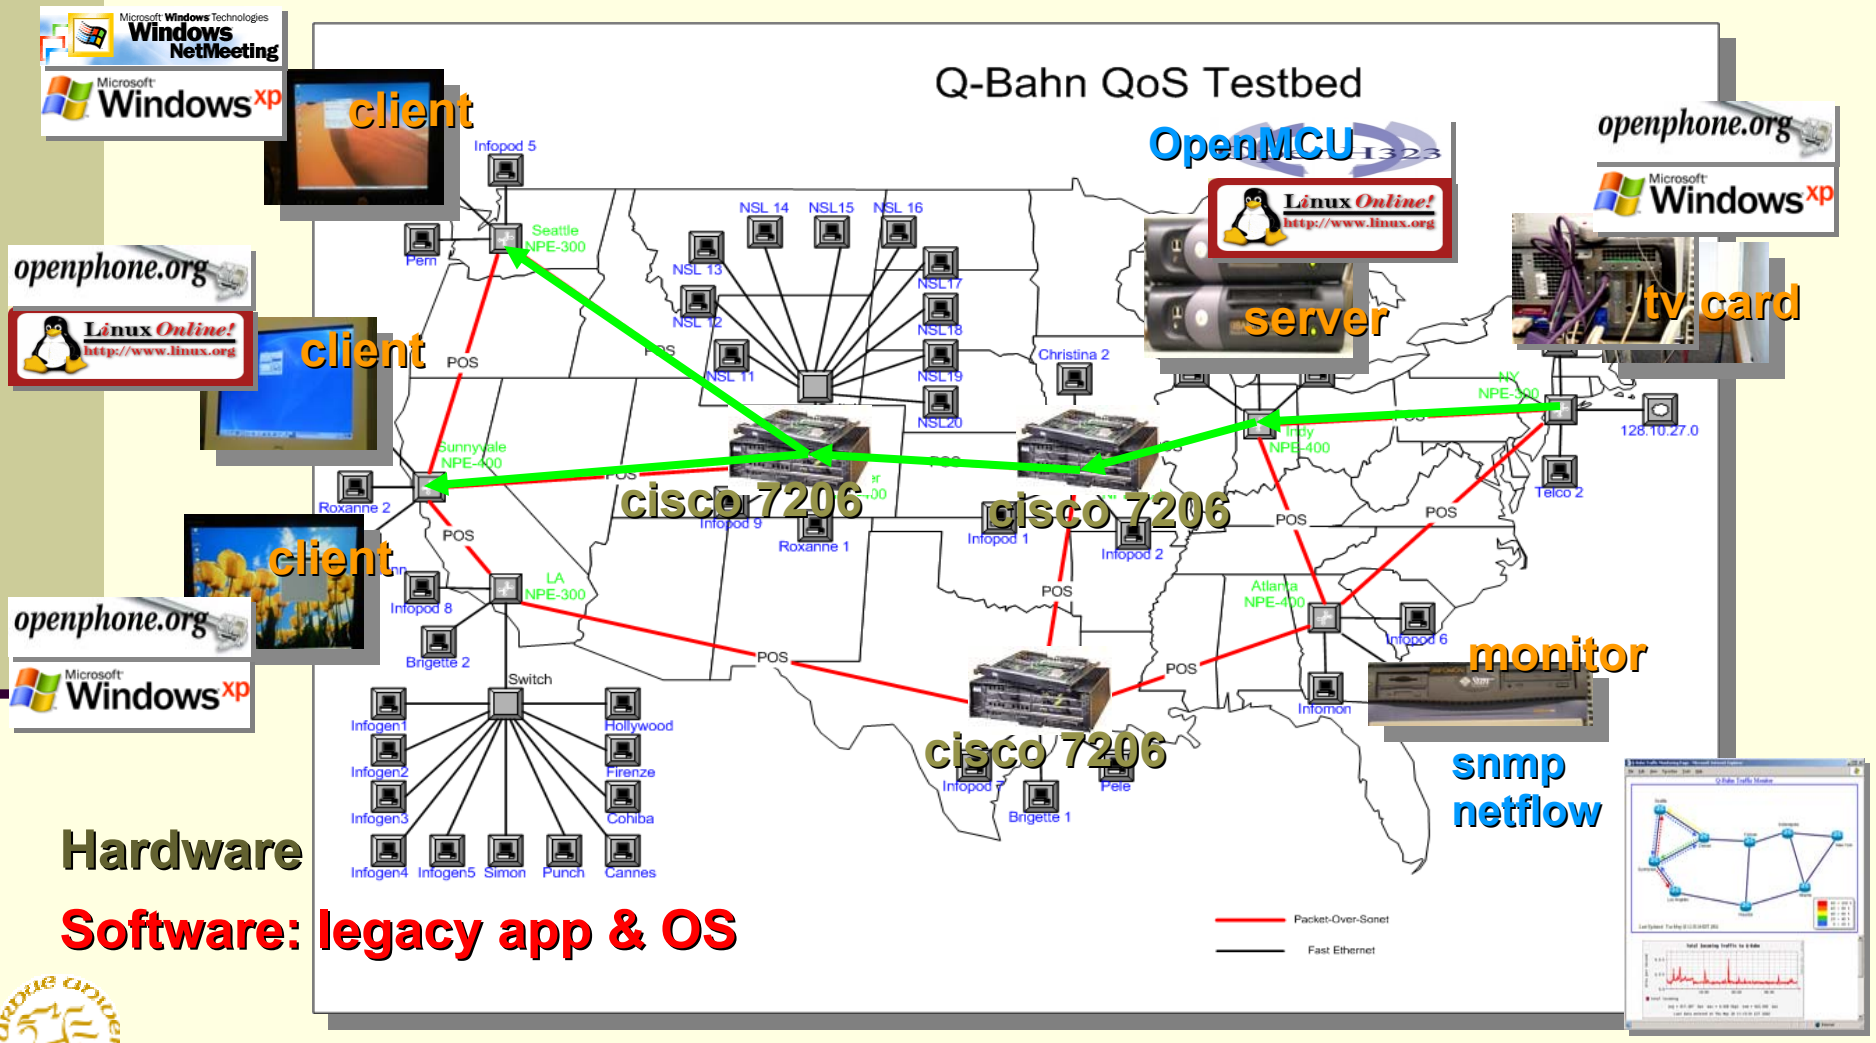

Hardware

Software: legacy app & OS

Experiment: Demo

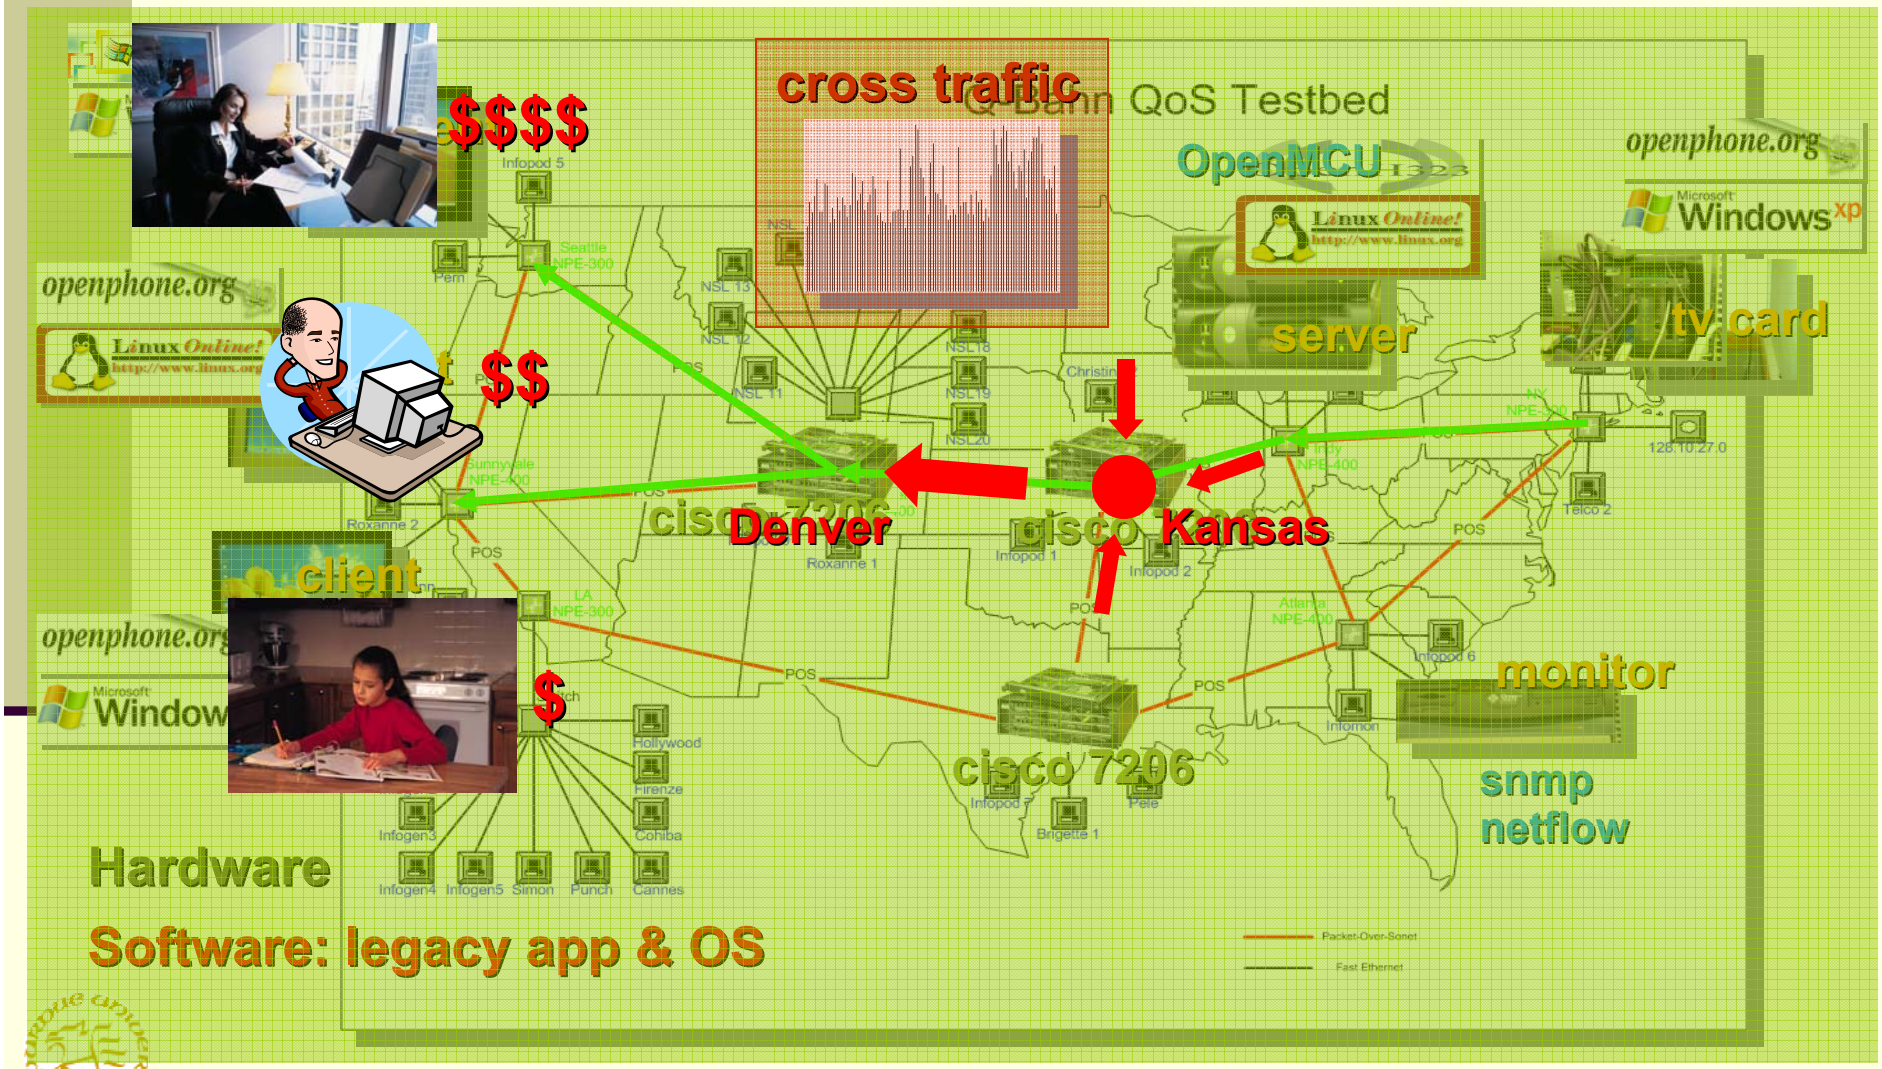

Experiment: Demo

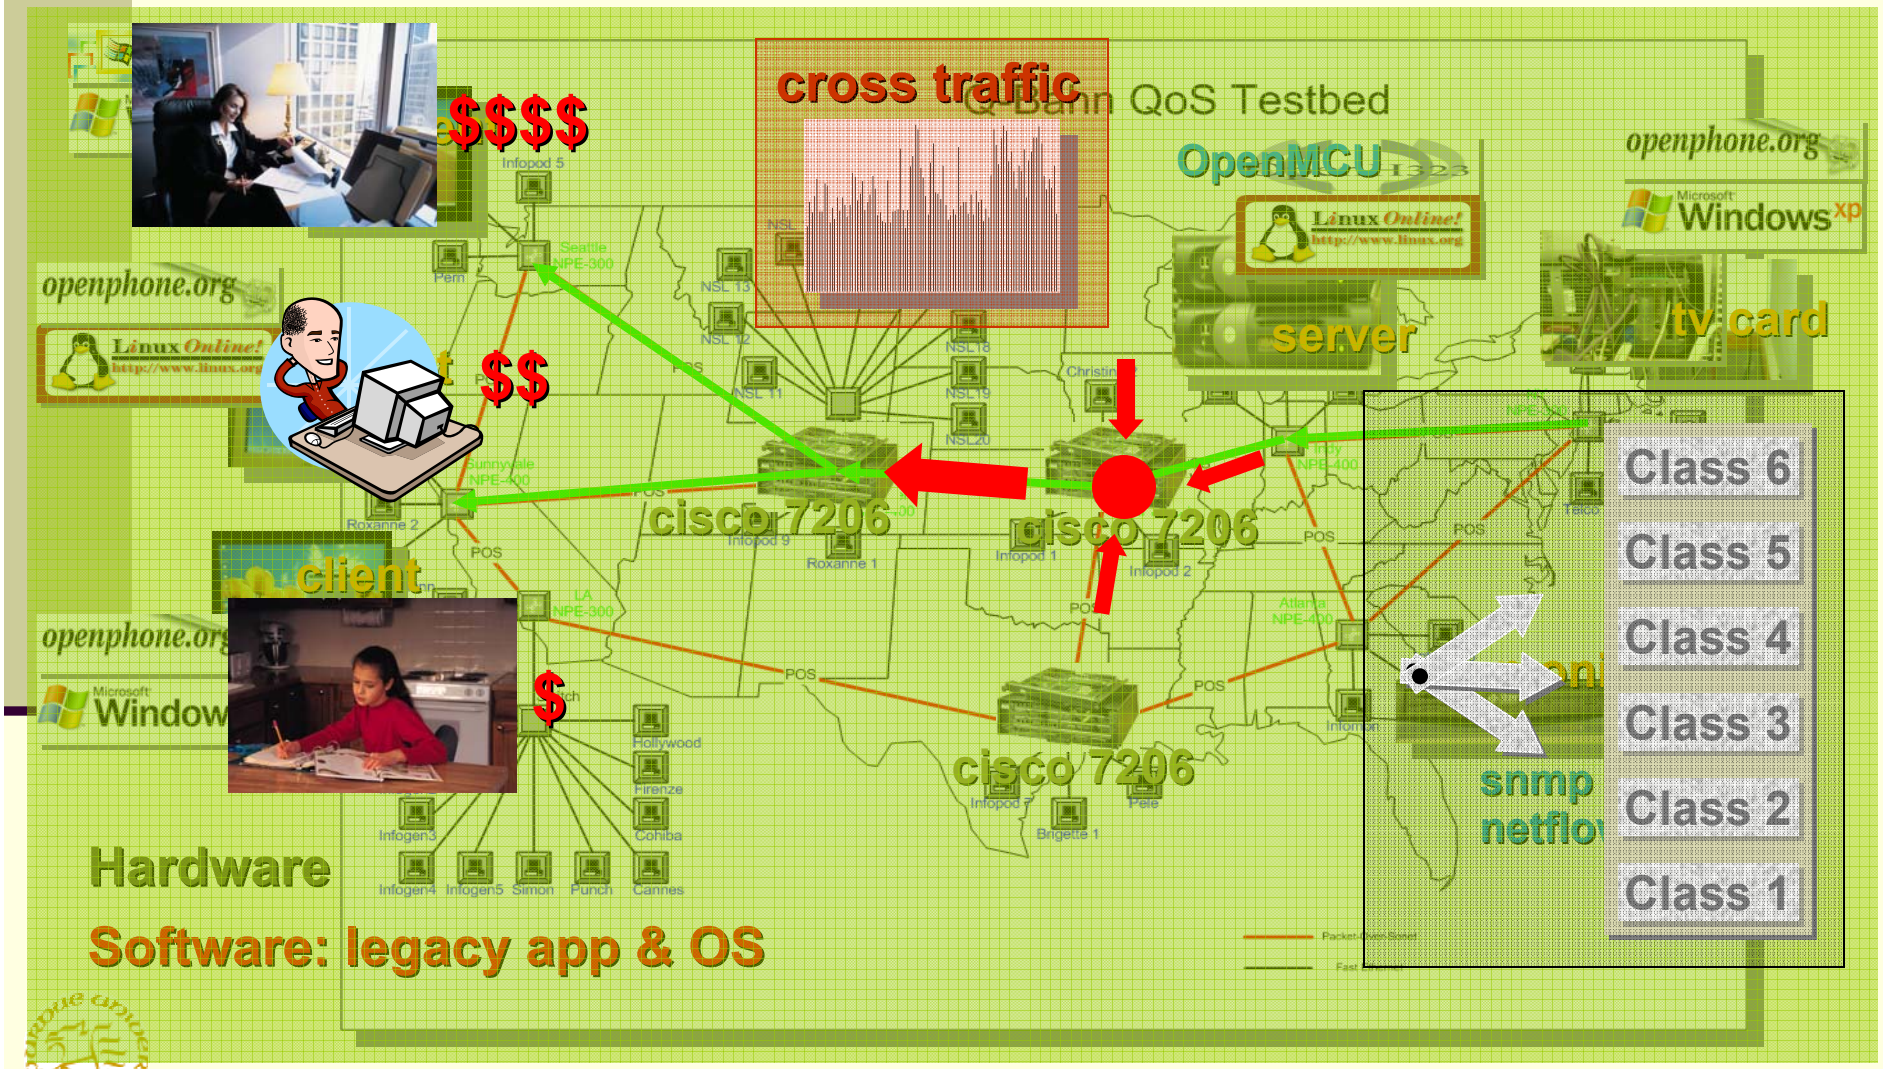

Experiment: Demo

VoIP & Teleconferencing

VoD, CDN, Web Server

client

client

server

monitor

snmp
netflow

cisco 7206

cisco 7206

wireless & 4G

Hardware

Software: legacy app & OS

Wireless Extension

- Physical system: wireless & mobile
  - 6-AP Enterasys RoamAbout 802.11b WLAN

Floor 1

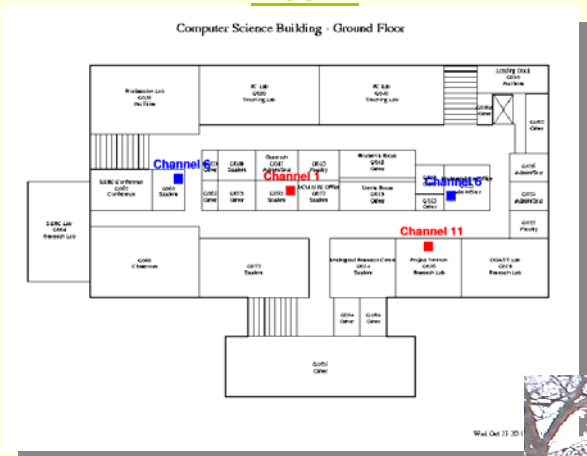

Floor 2

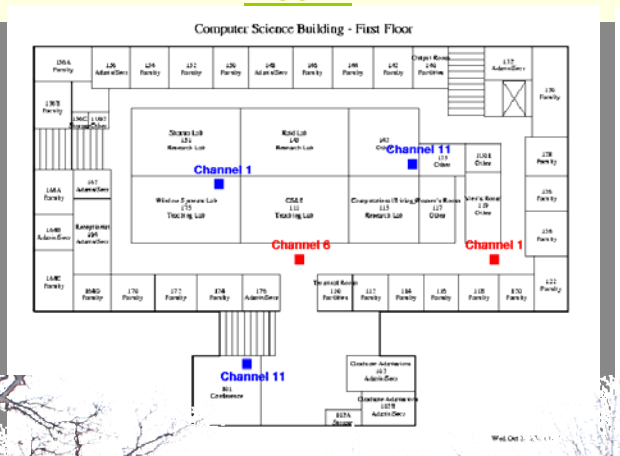

Floor 3

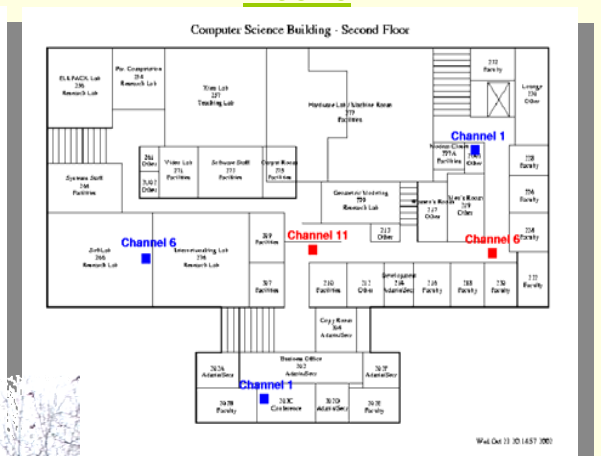

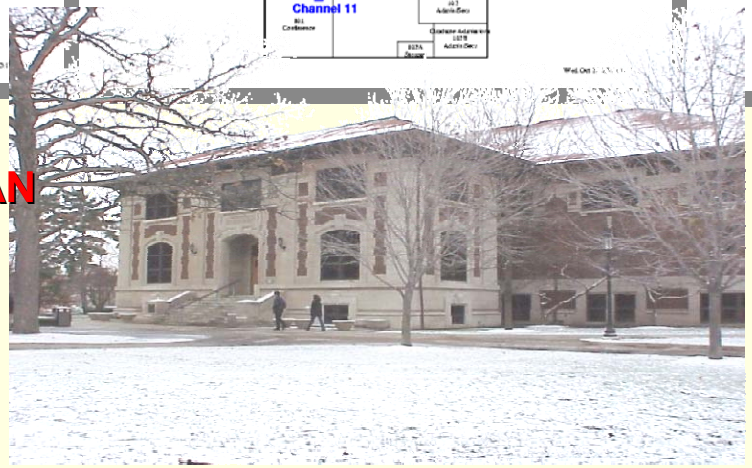

- **Network Systems Lab WLAN**
- **CS Dept. WLAN**

Wireless Extension

- Physical system: wireless & mobile
  - Mobiles: pocket PCs, laptops

Floor 1 Floor 2 Floor 3

Wireless Experiment: Demo

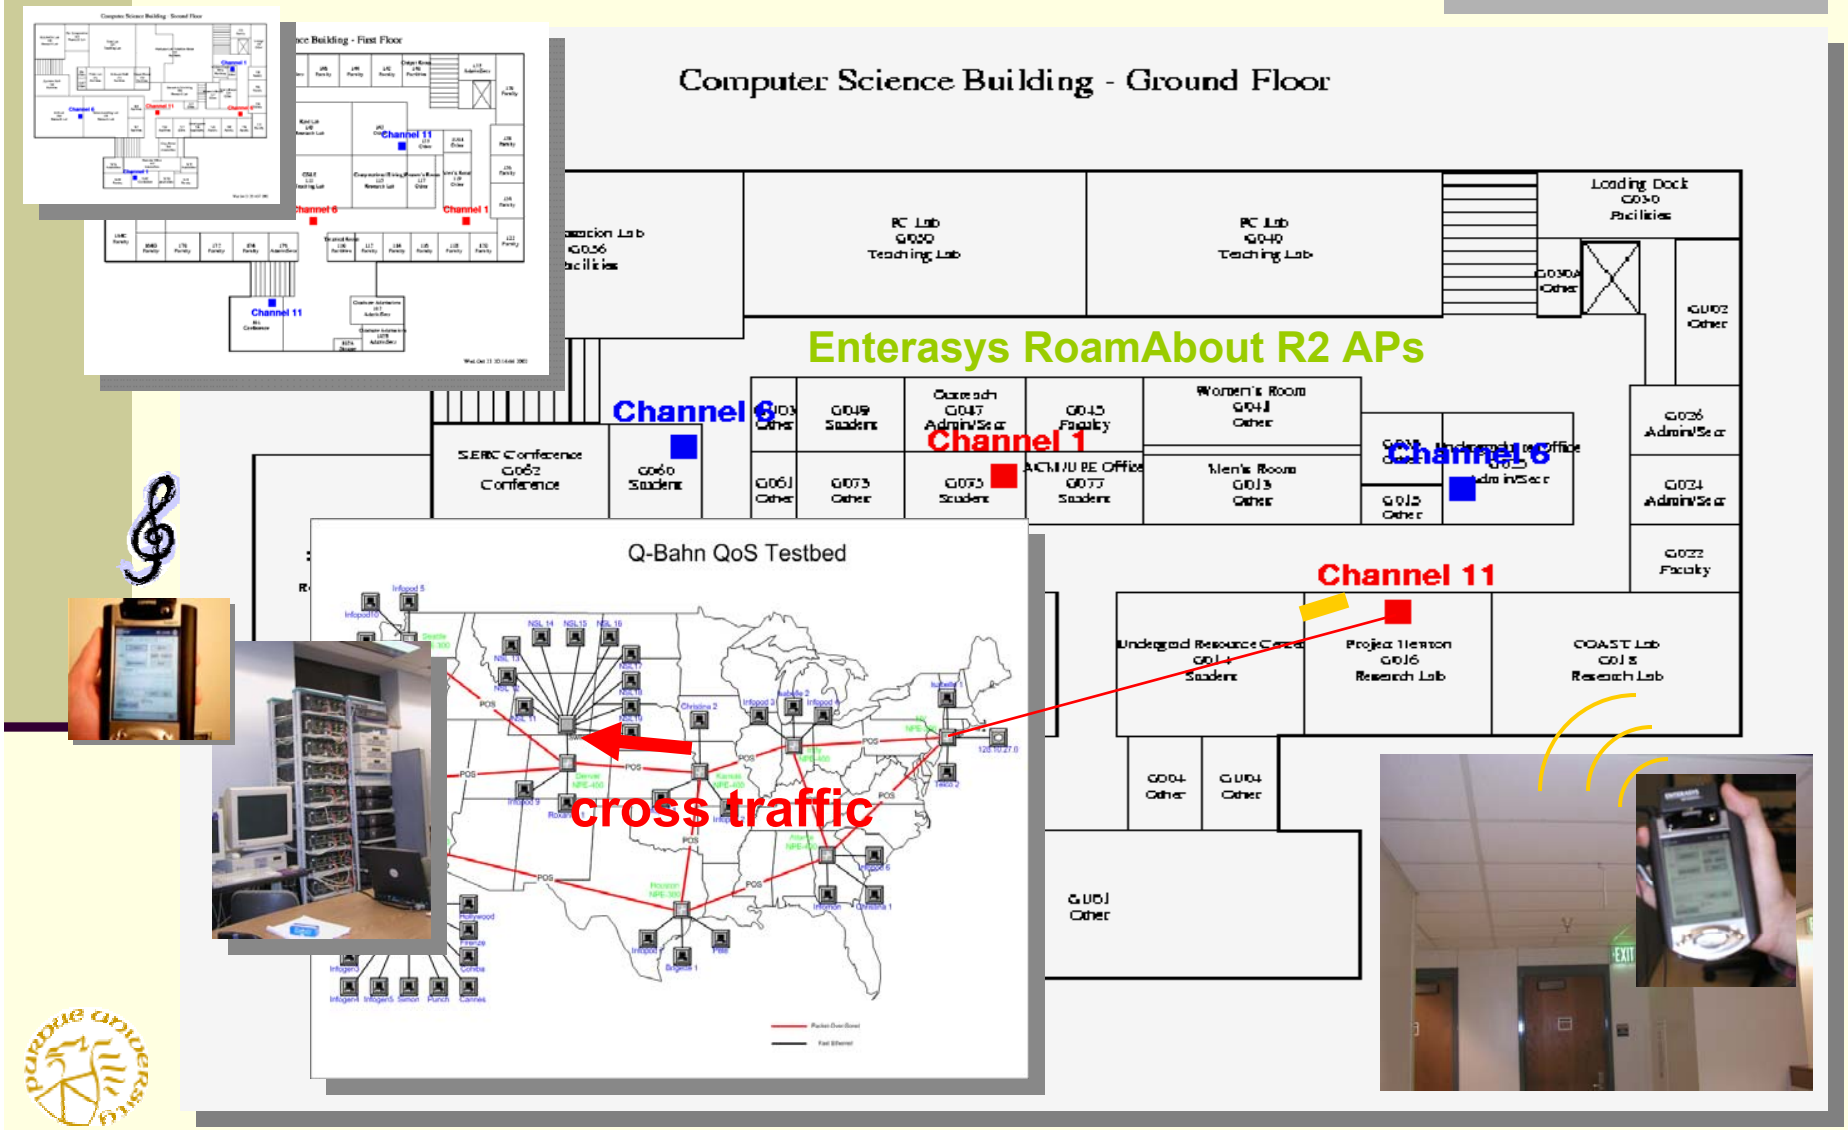

Wireless Experiment: Demo

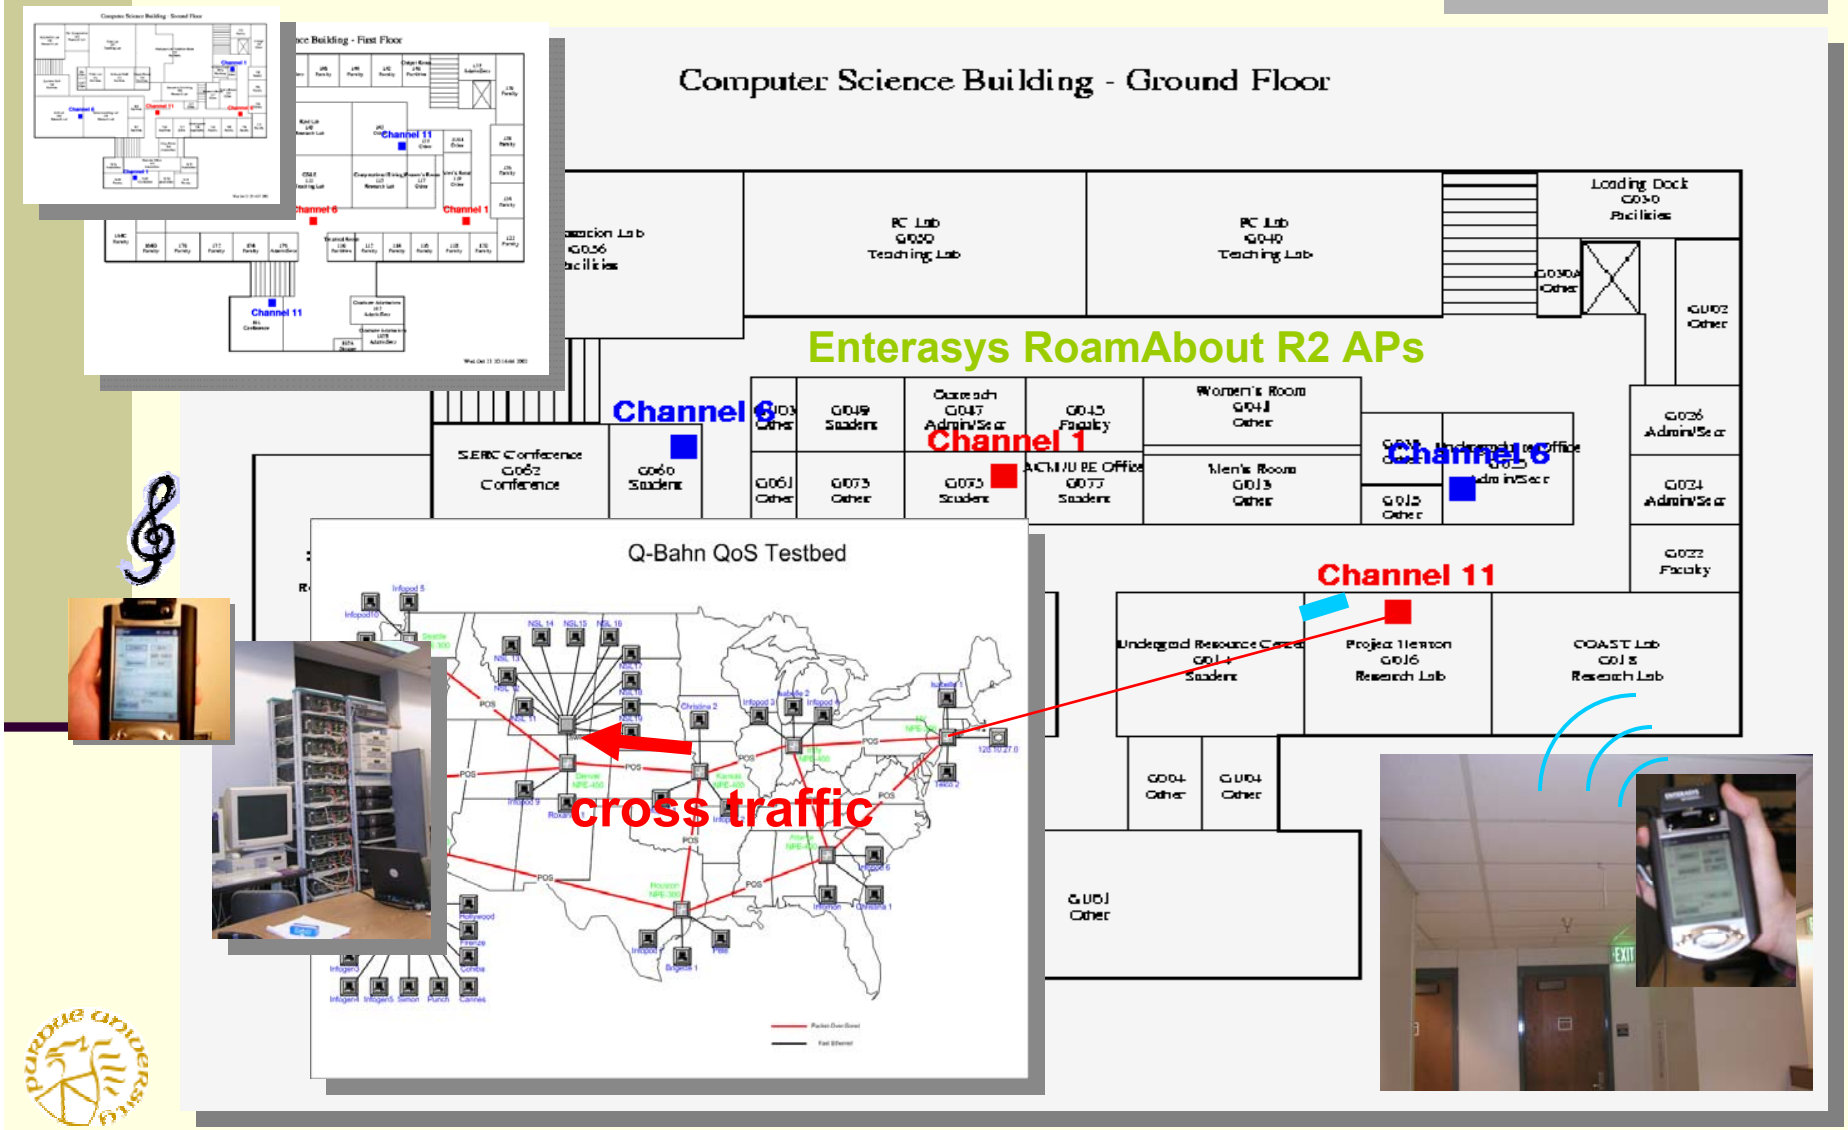

Wireless Experiment: Demo

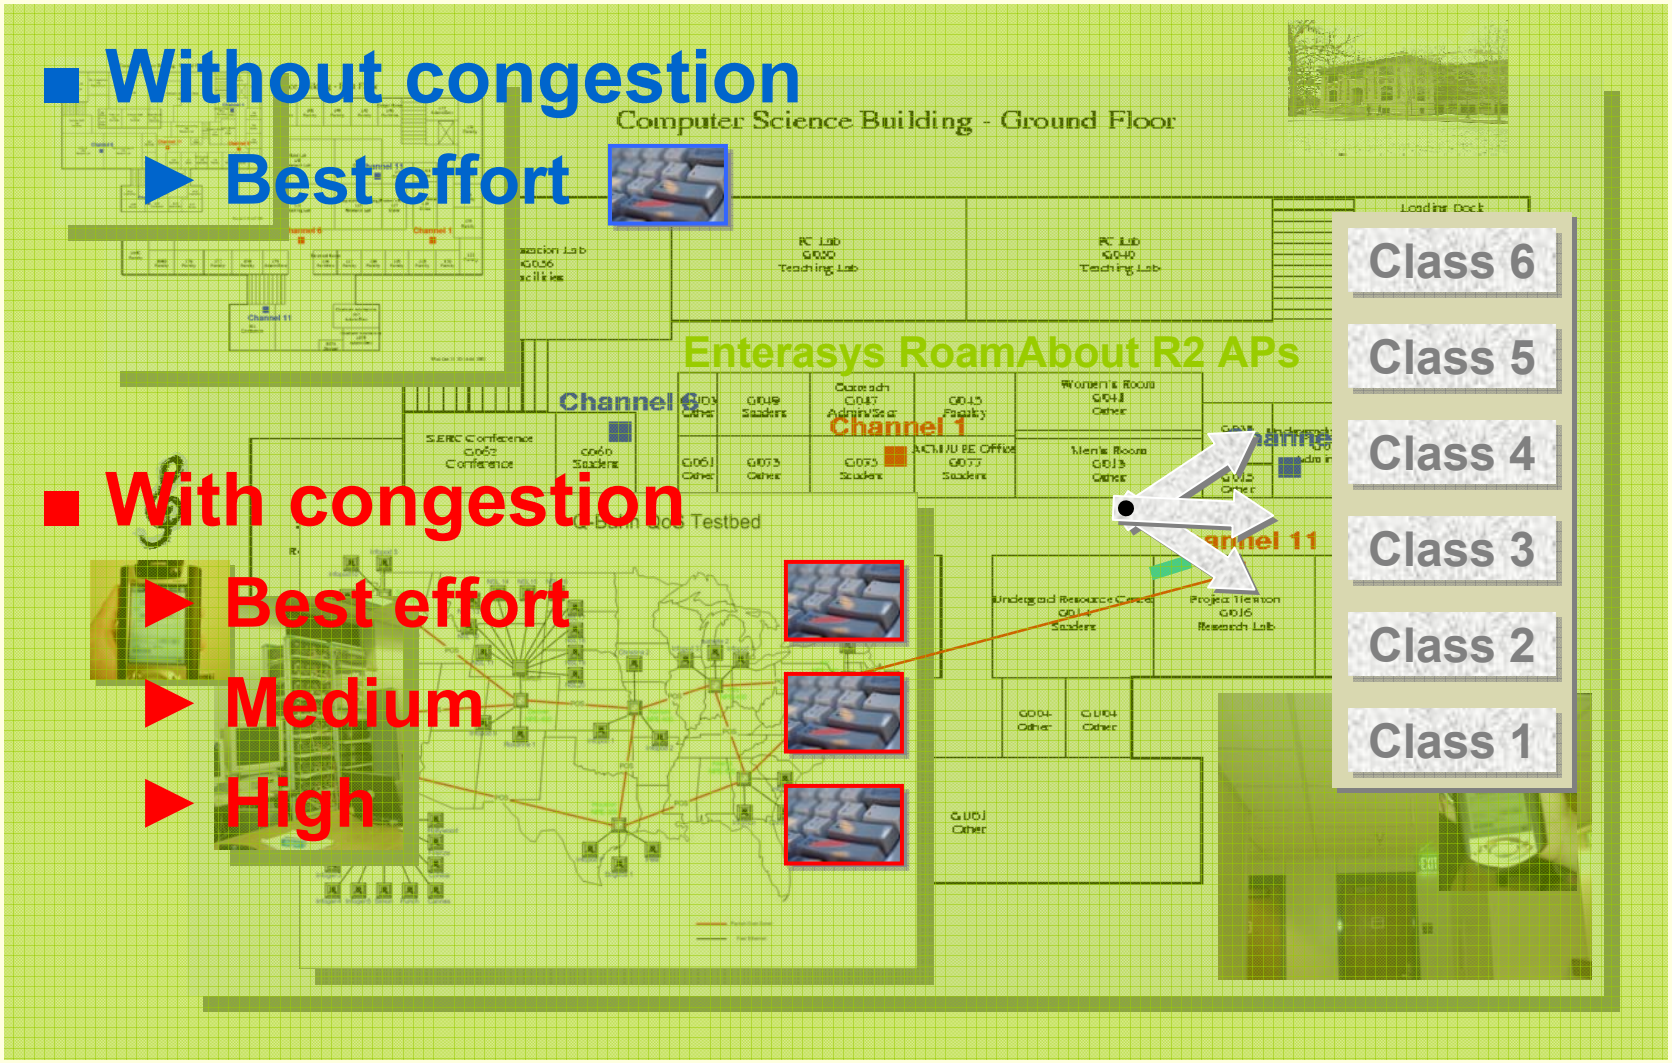

Benchmark Results: Structural

- Dynamic workload process

# of sessions

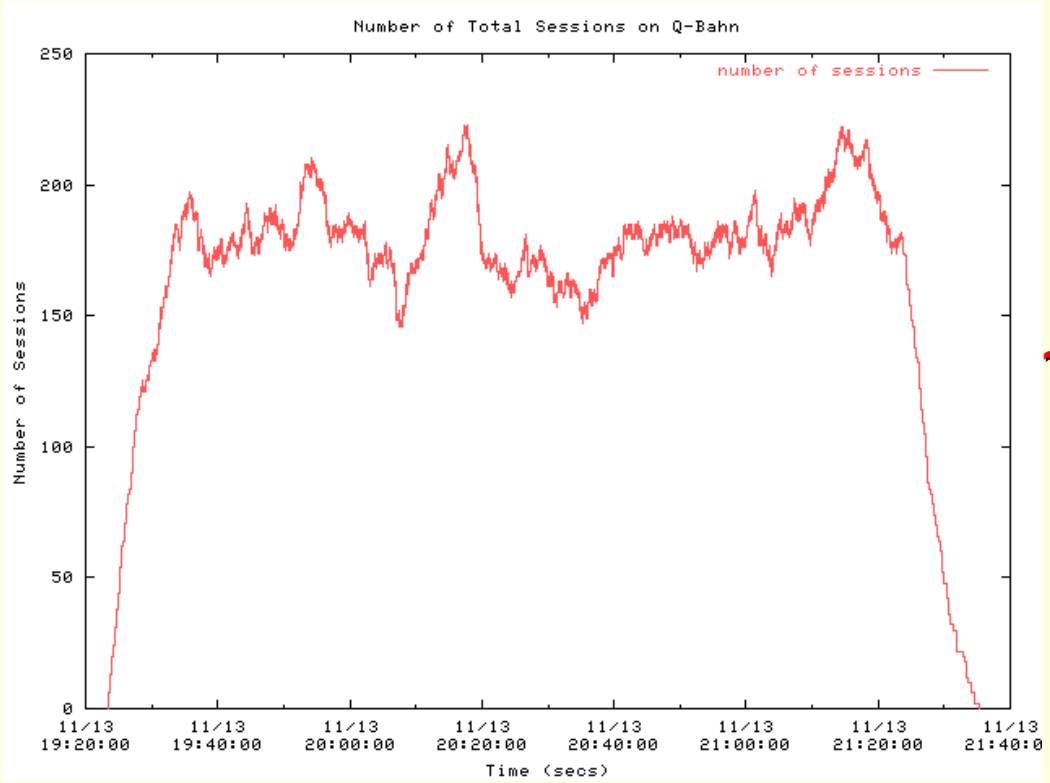

~2,000 session

time

→ bursty arrivals: 11/13/02, 7:20pm-9:40pm

Benchmark Results: Structural

- Performance: TOS field value 2 vs. 7

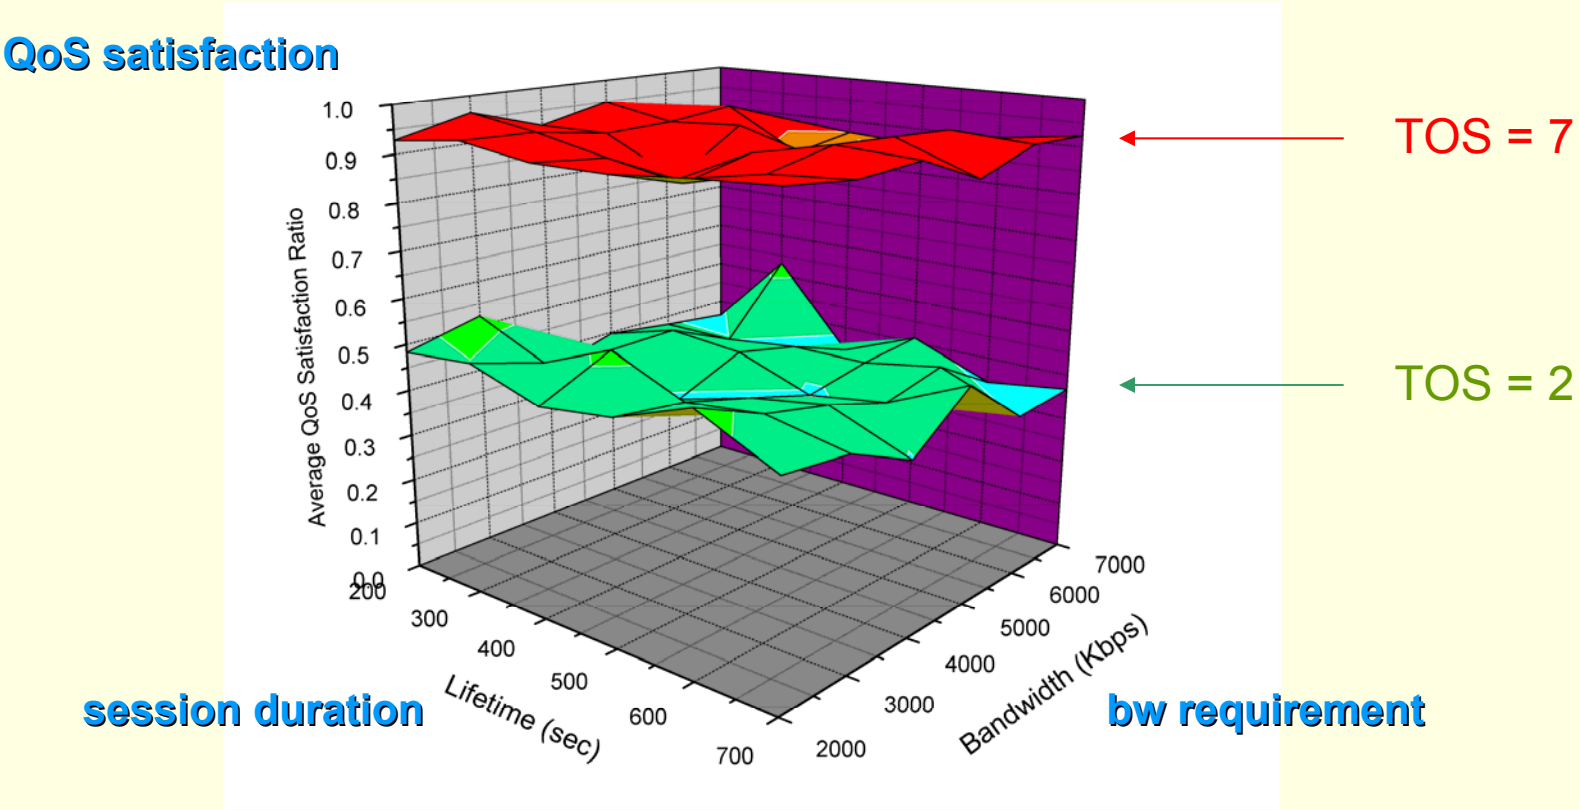

→ robust w.r.t. heterogenous workloads

Benchmark Results: Structural

- Workload: TCP file transfer

  → 80%+ of Internet traffic is HTTP traffic

completion time

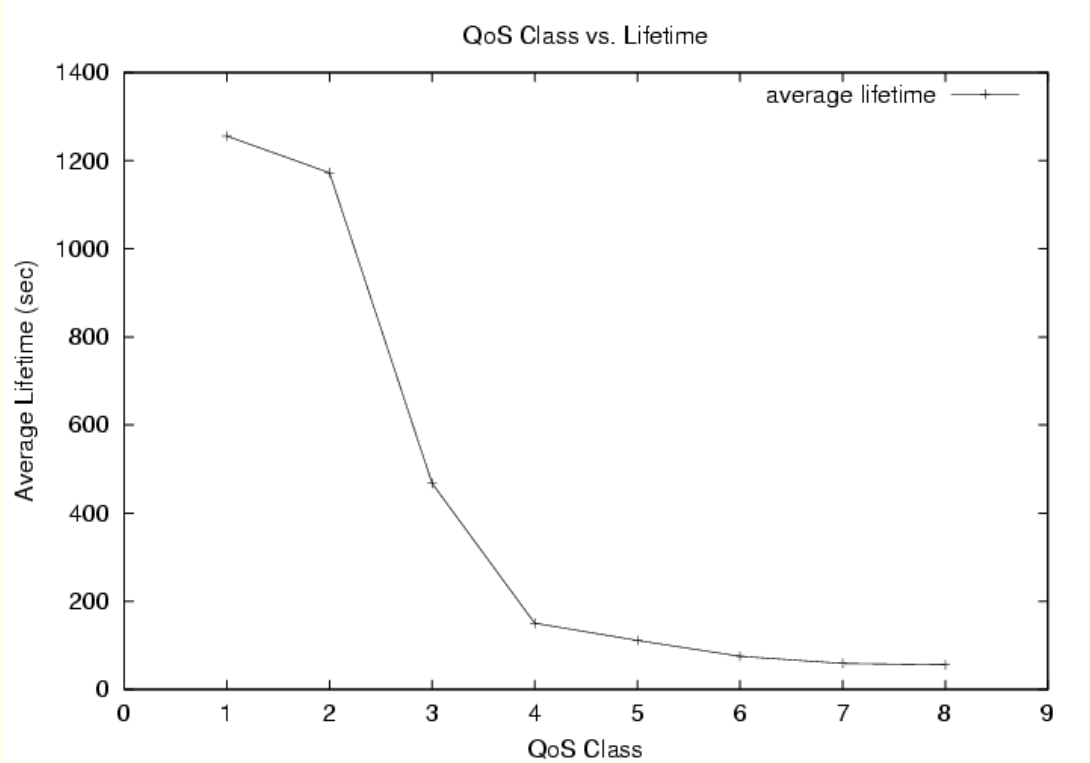

service class

Other Application Domains

- Functional integration
  → new services

- Enterprise systems
  → Xerox document services

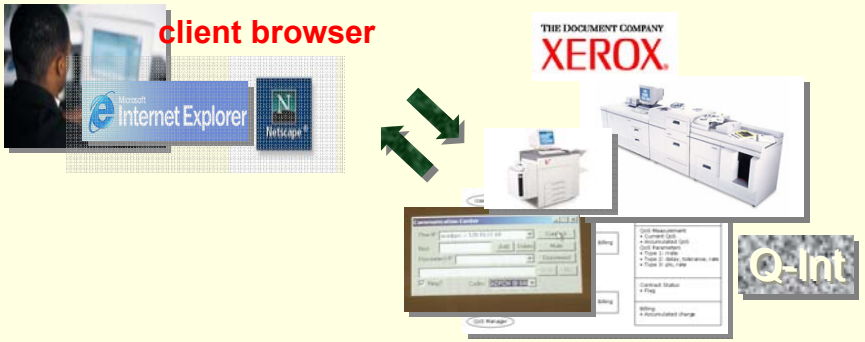

- University services
  → PURDUE UNIVERSITY intra-domain network mgt.

► **ResNet**

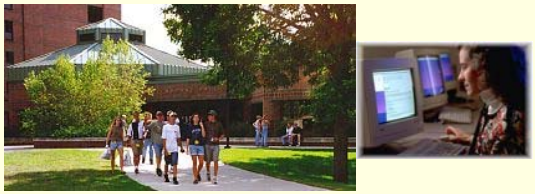

Further Info & Acknowledgment

- Contact
  - E-mail: park@cs.purdue.edu
  - http://www.cs.purdue.edu/nsl

- Supported by
  - NSF
  - DARPA
  - CERIAS, Xerox, Intel, ETRI