UNIVERSITY of VIRGINIA

A Ph.D. Defense on

Quantifiable Service Differentiation for Packet Networks

Nicolas Christin
University of Virginia
Department of Computer Science
Charlottesville, VA 22904

nicolas@cs.virginia.edu

Related Publications

- J. Liebeherr and N. Christin. JoBS: Joint Buffer Management and Scheduling for Differentiated Services. In *Proceedings of IEEE/IFIP IWQoS 2001*. June 2001. Karlsruhe, Germany.
- N. Christin, J. Liebeherr, and T. Abdelzaher. A Quantitative Assured Forwarding Service. In *Proceedings of IEEE Infocom 2002*, vol. 2, pages 864-873. June 2002. New York, NY.
- J. Liebeherr and N. Christin. Rate Allocation and Buffer Management for Differentiated Services. In *Computer Networks* 40(1), pages 89-110. September 2002.
- N. Christin and J. Liebeherr. A QoS Architecture for Quantitative Service Differentiation. In *IEEE Communications Magazine* 41(6), pages 38-45. June 2003.
- N. Christin and J. Liebeherr. The QoSbox: A PC-Router for Quantitative Service Differentiation in IP Networks. Technical Report CS-2001-28. October 2001. Revised version in submission to *ACM Computer Communications Review*.
- N. Christin and J. Liebeherr. Marking Algorithms for Service Differentiation TCP Traffic. Technical Report CS-2003-04. February 2003. In submission to *Computer Communications*.

Software Dissemination

- The QoSbox. http://qosbox.cs.virginia.edu
- Simulation code included in *ns-2* as of ns-2.26
- Kernel code included in extension packages for 4.4-BSD based kernels (FreeBSD, OpenBSD, NetBSD)
  - ALTQ-3.1. http://www.csl.sony.co.jp/person/kjc/kjc/software.html
  - KAME snap-kits. http://www.kame.net
  - Inclusion in base distributions of the FreeBSD, OpenBSD, and NetBSD kernels is under consideration

Outline

- Background and thesis statement
- New service: Quantitative Assured Forwarding
- Main contributions
  - Joint Buffer Management and Scheduling (JoBS)
  - A feedback-control design
  - BSD kernel implementation (QoSbox)
  - Quantitative Assured Forwarding for TCP traffic
- Conclusions

Need for Quality-of-Service (QoS)

- What is QoS?
  - Traffic control mechanisms to differentiate performance based on network-operator or application requirements
  - Delay, loss and throughput guarantees are considered most crucial
- Why is QoS needed?
  - Different applications have different needs
  - The Internet is now a commercial network → service differentiation can create incentives for differentiated pricing
- Where is QoS needed?
  - Core(s) vs. edge(s)

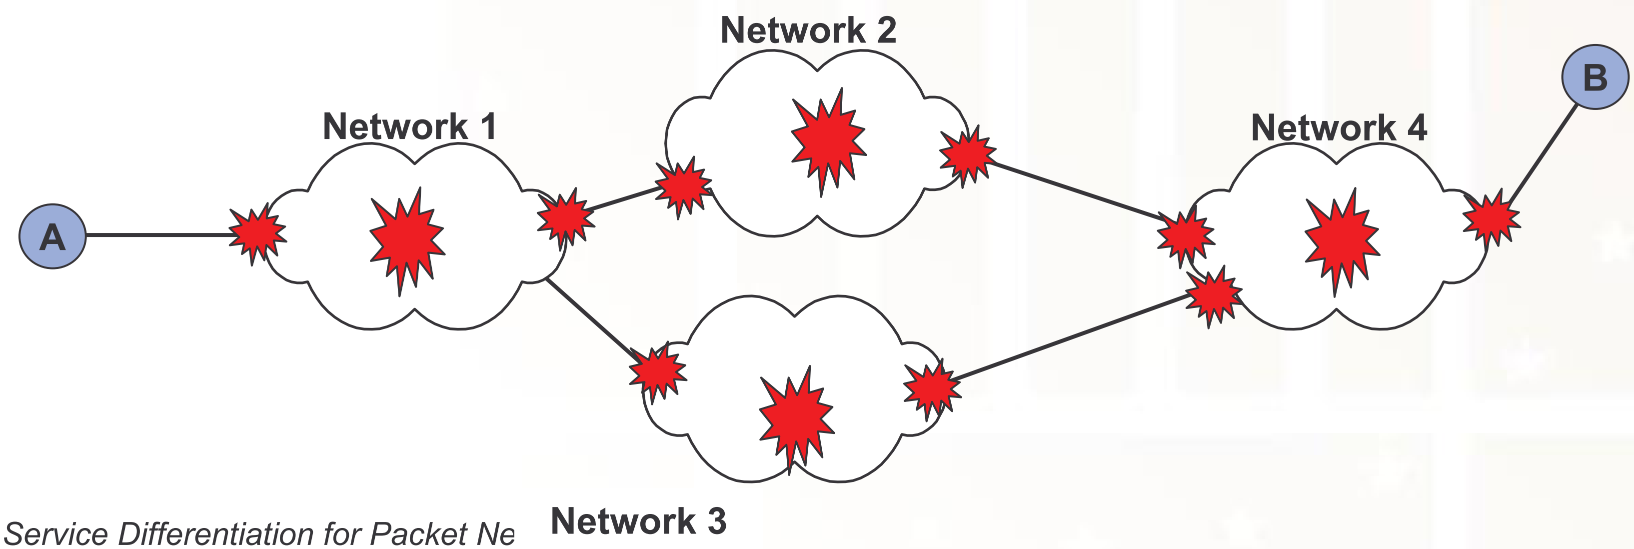

Granularity of Service Guarantees

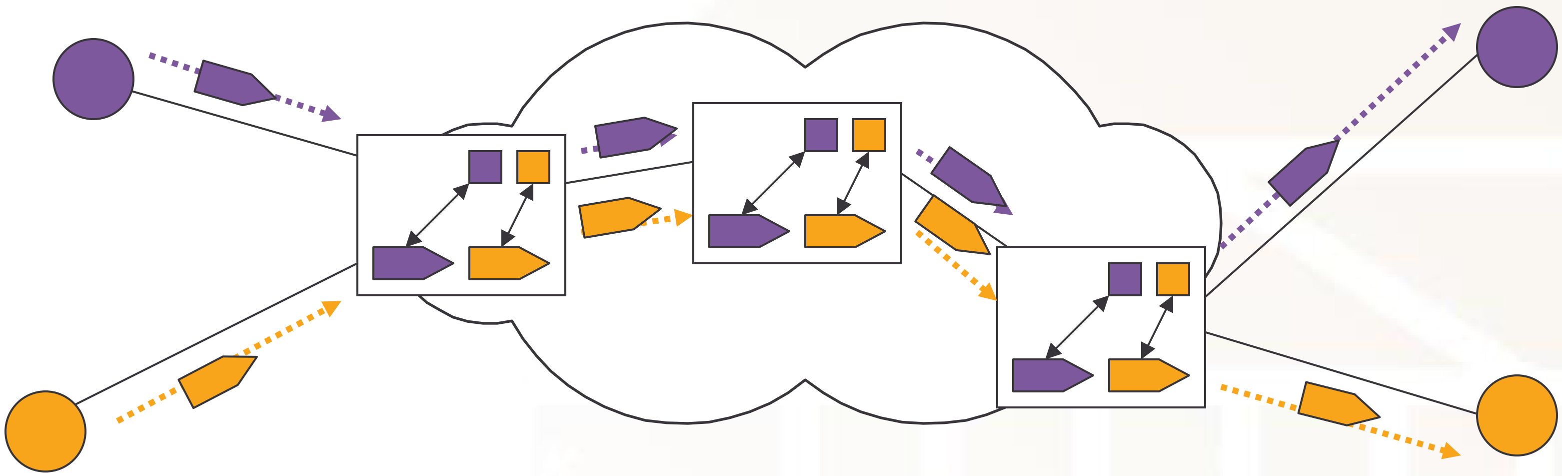

- Per-flow guarantees
  - Require per-flow reservations in the network
  - Require per-flow classification at routers

Granularity of Service Guarantees

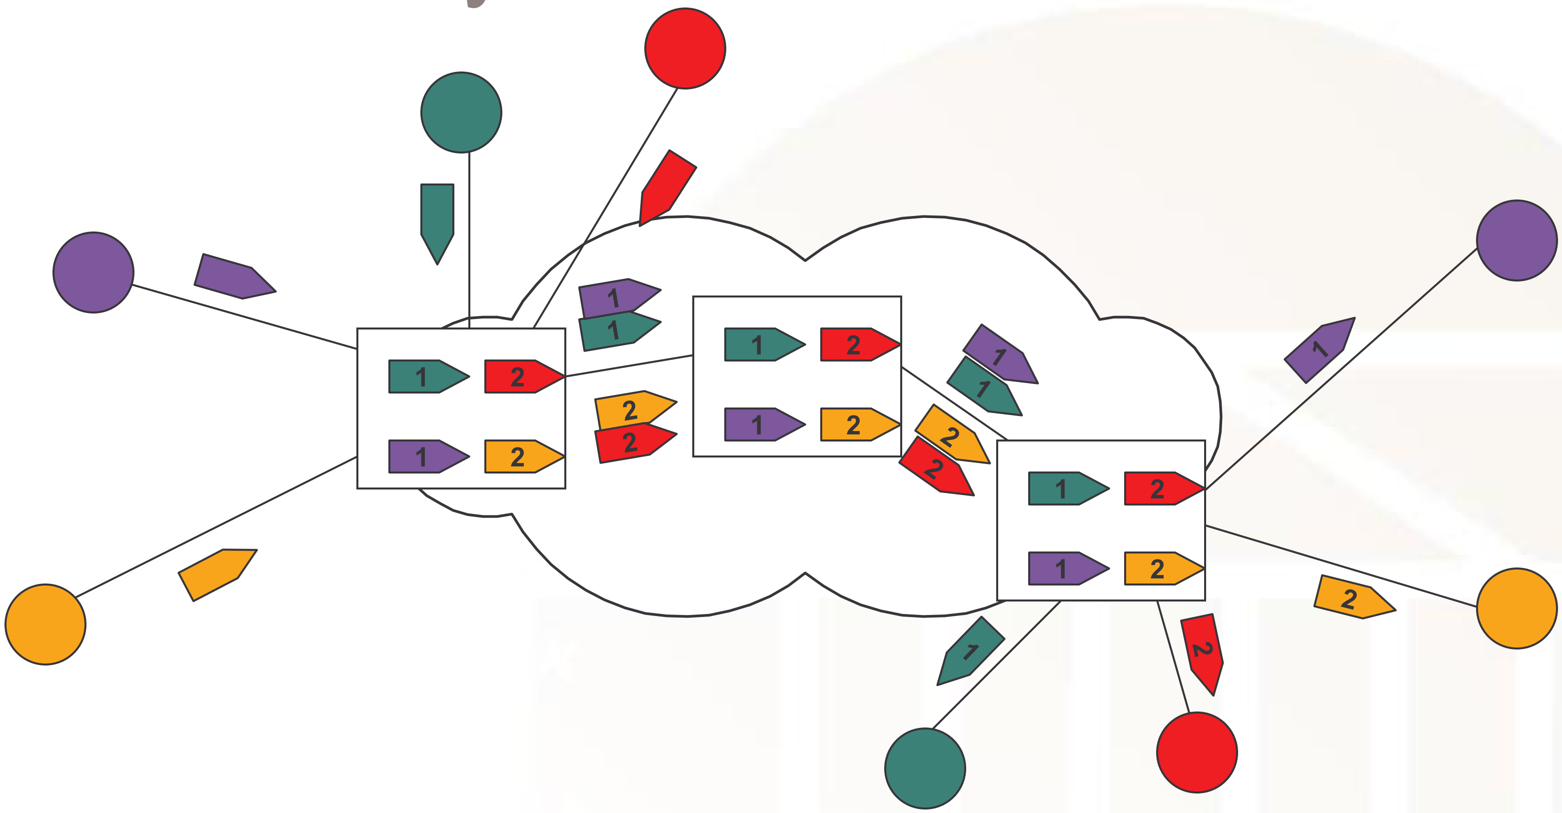

- **Per-class guarantees**
  - Bundle traffic flows with similar service requirements into a small number of aggregates (classes)
  - Per-flow reservations are not needed
  - Provide service differentiation to aggregates
  - Per-class guarantees do not immediately translate into per-flow guarantees

Strength of Service Guarantees

- Absolute guarantees
  - Generally require reservations of router resources
  - Example:
    - *Class-1 loss rate ≤ 1 %*
    - *Flow-2 throughput ≥ 2 Mbps*

- Relative guarantees
  - Generally do not require reservation of resources

  1. Qualitative guarantees
     - *Class-1 loss rate ≤ Class-2 loss rate*
     - *Class-1 delay ≤ Class-2 delay*
  2. Proportional guarantees
     - *Class-2 delay / Class-1 delay ≈ 2*
     - *Flow-3 throughput / Flow-2 throughput ≈ 5*

Problem and Context

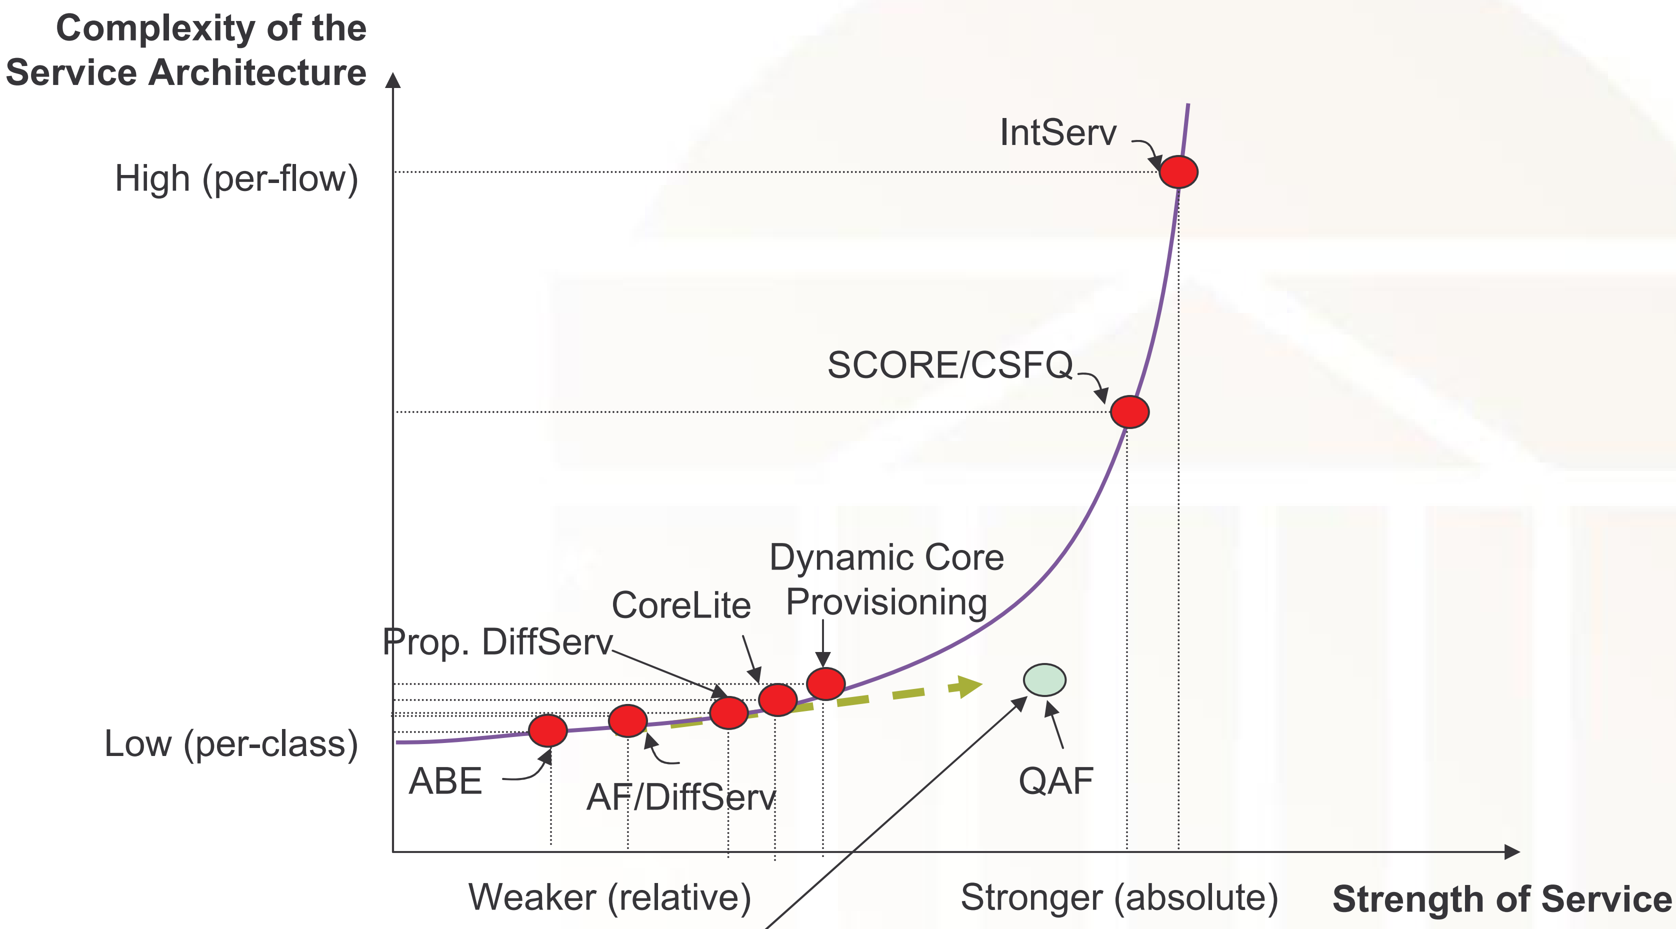

The scope of class-based service guarantees can be significantly enhanced by using appropriate buffer management, scheduling and feedback capabilities of the network

Related Work

- **Proportional Delay and Loss Differentiation** (Dovrolis, PhD thesis, 2001)
  - Proportional guarantees
- **CoreLite** (Barghavan et al., 2000)
  - Absolute and proportional delay guarantees, no loss guarantees
- **SCORE/CSFQ/DPS** (Stoica, PhD thesis, 2001)
  - Strong guarantees, but per-flow classification at access points
- **ABE Service** (Hurley, PhD thesis, 2002)
  - Two classes (one with delay bounds, the other with lower losses)
- **Dynamic Core Provisioning** (Liao, PhD thesis, 2002)
  - Throughput and loss differentiation, no absolute guarantees on delays

Quantitative Assured Forwarding

- Objective: Strongest possible class-based service without reservations
- Service guarantees of QAF:
  - **Proportional and absolute** per-class guarantees for loss and delay
  - **Lower bounds on the throughputs of classes**
- QAF can be used to express all existing class-based architectures

- Characteristics:
  - Guarantees provided on a per-hop, per-class basis
  - No admission control, no signaling, no traffic conditioning
  - No per-flow operations
  - Important: Service guarantees may need to be temporarily relaxed

Deployment of QAF service

Static configuration of service guarantees

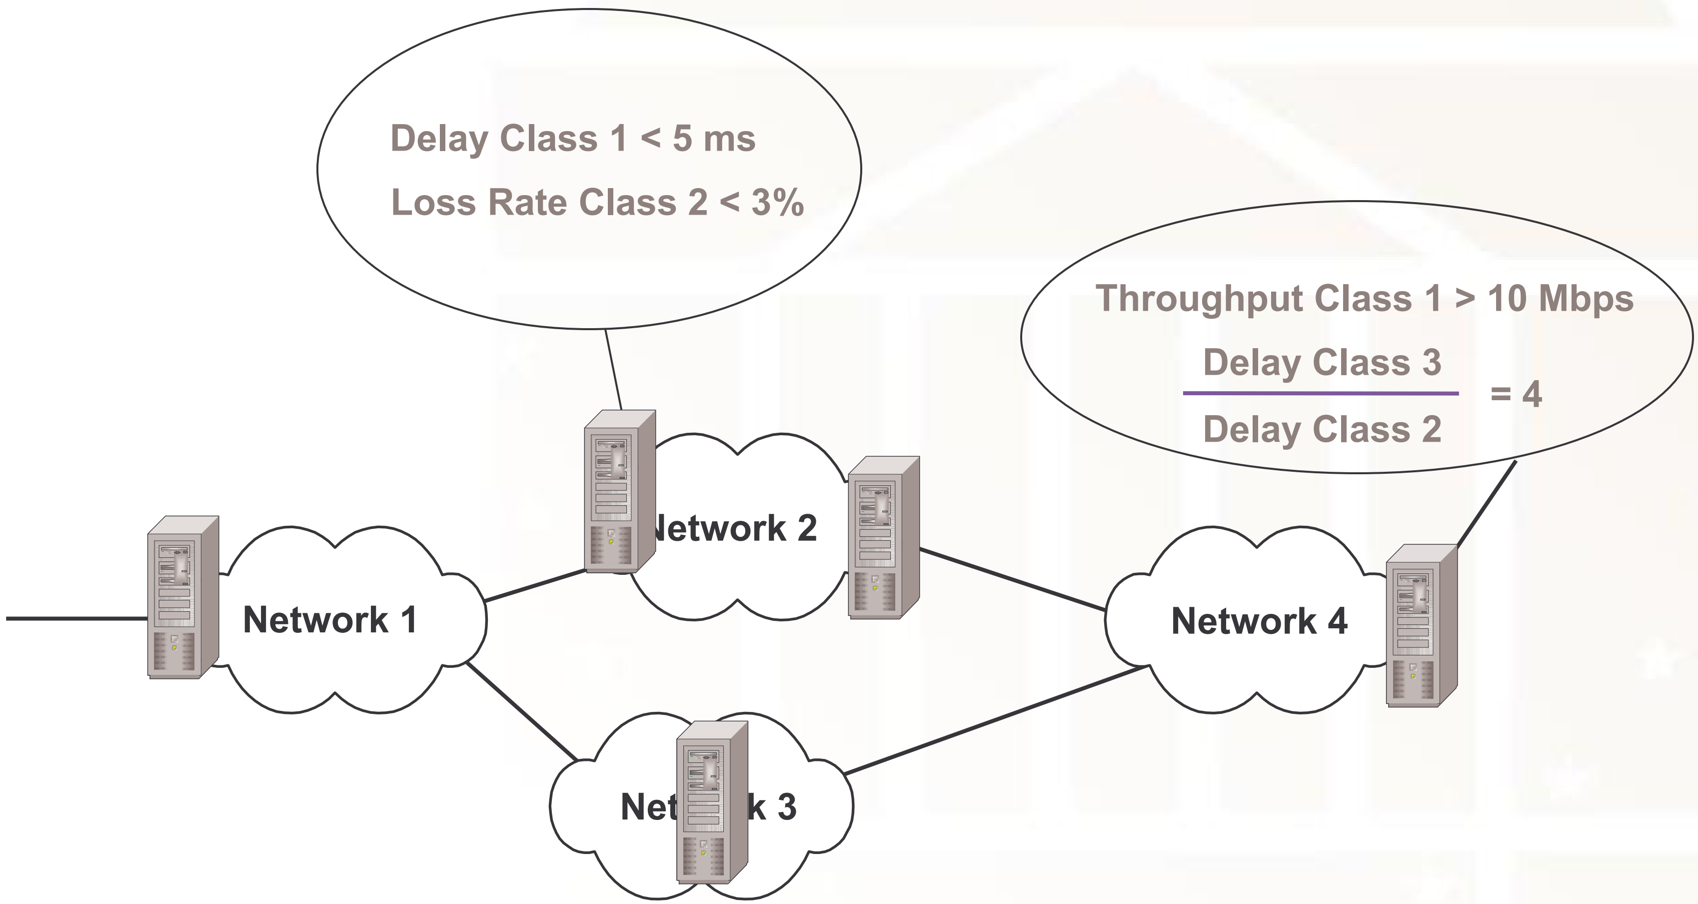

JoBS – Joint Buffer Management and Scheduling

- JoBS (Joint Buffer Management and Scheduling):
  - Defines any algorithm that tracks traffic arrivals to allocate service rates and to drop traffic in a single step
- Operations:
  - Service rate allocation to traffic classes
  - Service rate allocation is periodically adjusted
  - If no feasible rate allocation exists for meeting all service guarantees, drop traffic
  - If necessary, relax service guarantees according to a pre-specified order
- JoBS can realize the Quantitative Assured Forwarding service

Arrivals, Departures, Losses at an Output Queue

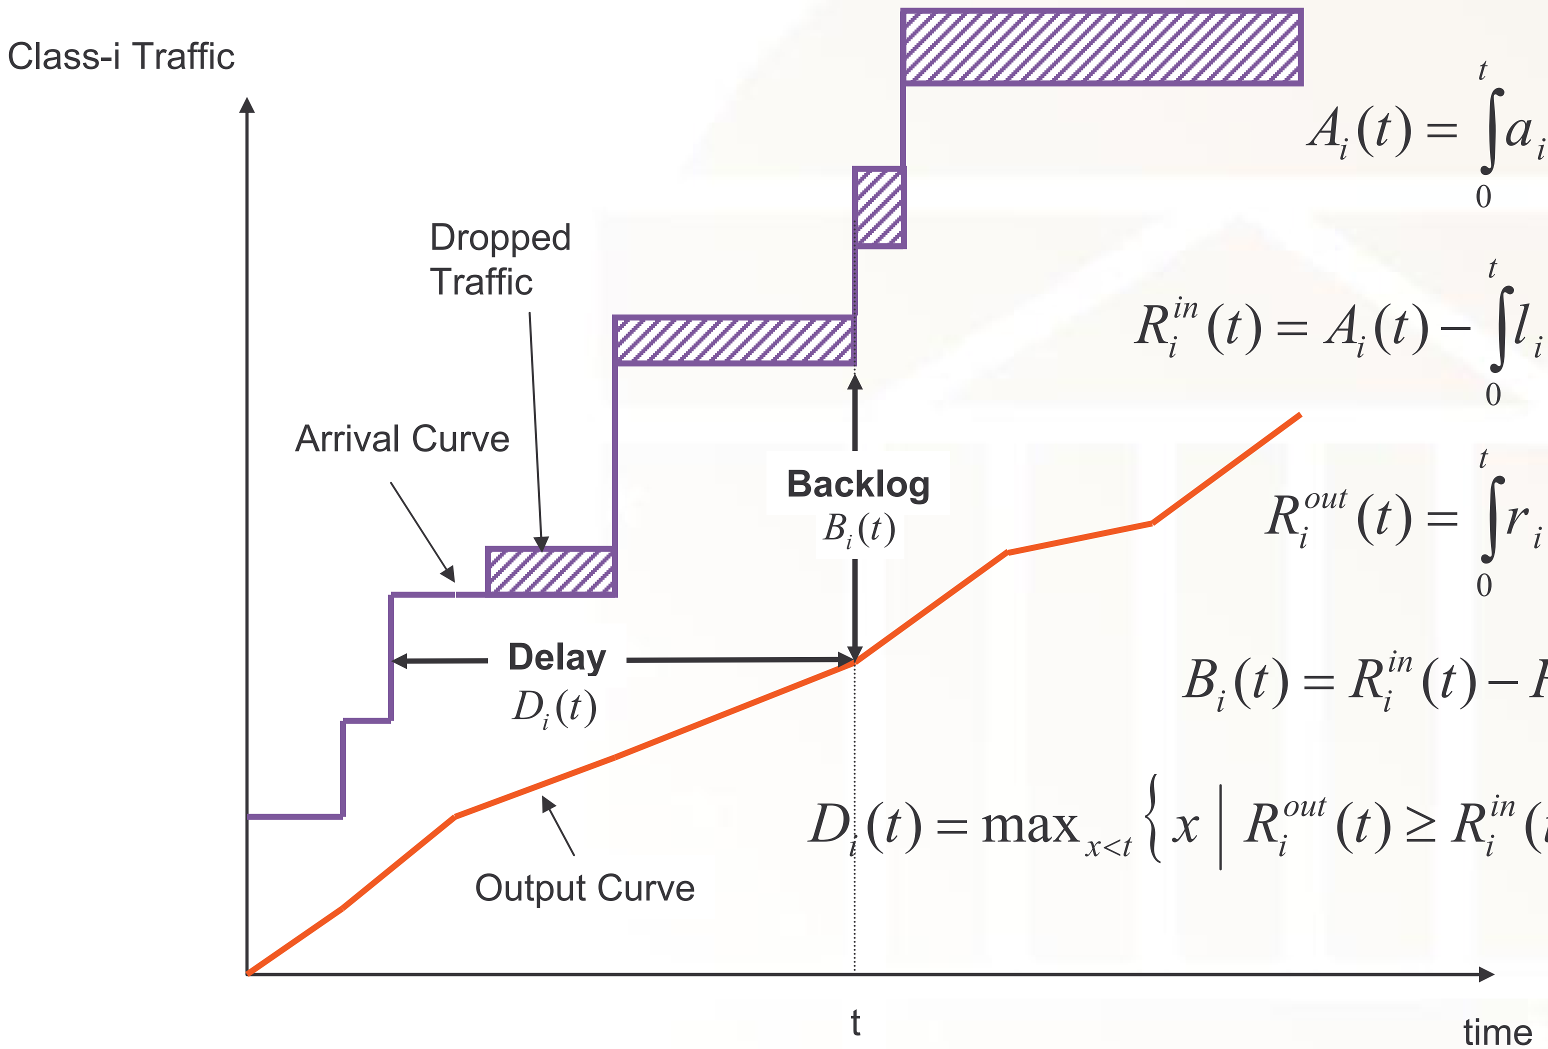

$$A_i(t) = \int_0^t a_i(x)dx$$

$$R_i^{in}(t) = A_i(t) - \int_0^t l_i(x)dx$$

$$R_i^{out}(t) = \int_0^t r_i(x)dx$$

$$B_i(t) = R_i^{in}(t) - R_i^{out}(t)$$

$$D_i(t) = \max_{x<t}\left\{x \;\middle|\; R_i^{out}(t) \geq R_i^{in}(t-x)\right\}$$

Rate Projections

Assumptions for projections at times s for delays at t>s:

- Current rate allocation does not change
- No new arrivals

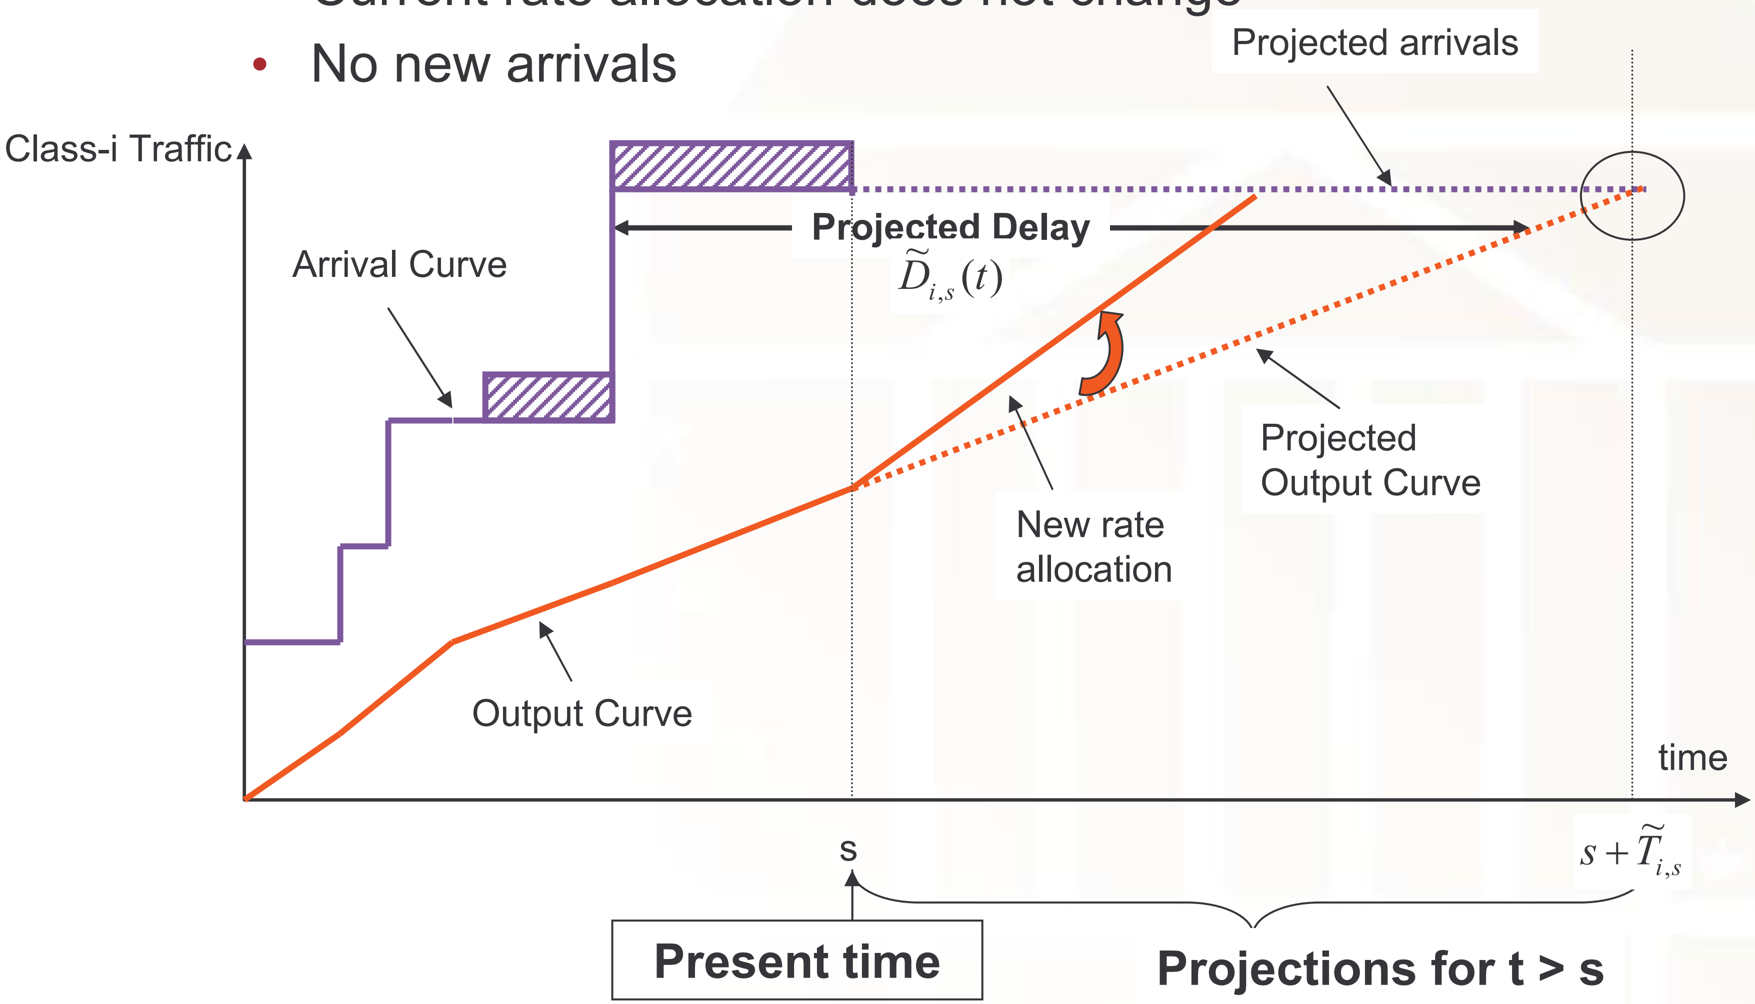

Optimization-Based Algorithm

- New rate allocations and drop decisions are obtained from an optimization

Minimize:	losses and changes to the rate allocation
Subject to:	- absolute bounds on loss, and delay
	- proportional service differentiation
	- system constraints (e.g., buffer size)

- If constraint system becomes infeasible, relax constraints in a specified order

Formulation of the Optimization

Objective function

$$\textbf{min} \quad C^2 \sum_i l_i(s) \; + \; \sum_i \left(r_i(s) - r_i(s^-) \right)^2$$

Physical Constraints

$$\sum_i r_i(t) = C, \quad \sum_i B_i(t) \leq B$$

Service Guarantees

$$\begin{cases} p_{i,s} \leq L_i \\ \max_{s<t<s+\tilde{T}_{i,s}} \tilde{D}_{i,s}(t) \leq d_i \\ \dfrac{\overline{D}_{i+1,s}}{\overline{D}_{i,s}} \approx k \quad \text{where } \overline{D}_{i,s} = \dfrac{1}{\tilde{T}_{i,s}} \displaystyle\int_s^{s+\tilde{T}_{i,s}} \tilde{D}_{i,s}(x)dx \\ \dfrac{p_{i+1,s}}{p_{i,s}} \approx k' \quad \text{where } p_{i,s} = \displaystyle\int_{t_0}^{s} l_i(x)\,dx \Big/ \int_{t_0}^{s} a_i(x)\,dx \end{cases}$$

Evaluation by Simulation

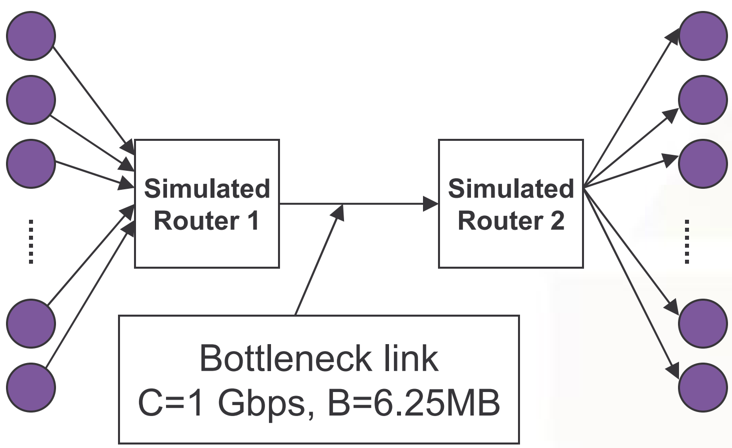

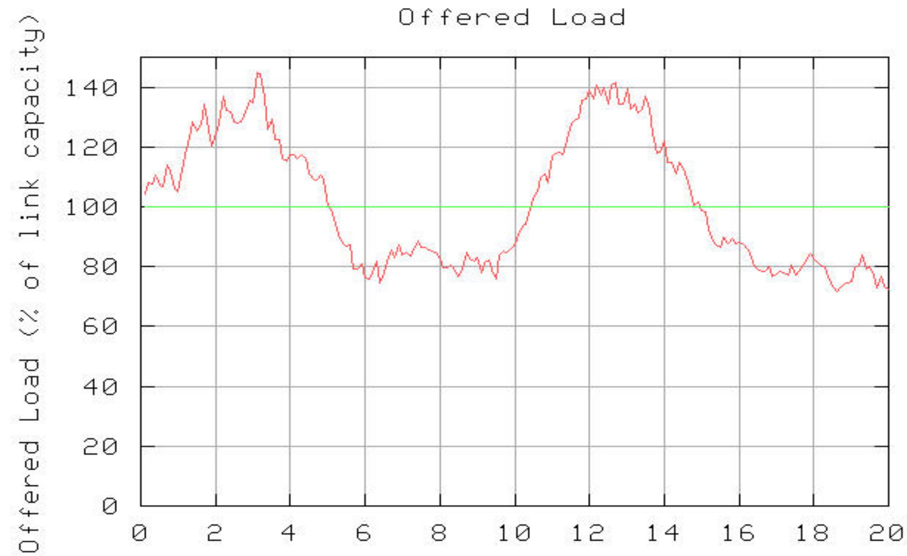

Class	Abs. Delay Guarantee	Proportional Delay Diff. Factor	Proportional Loss Diff. Factor
1	1 ms	-	
			2
2	-		
		4	2
3	-		
		4	2
4	-		

Simulation Results: Delay

$$\frac{\text{Class-4 delay}}{\text{Class-3 delay}} \approx 4$$

$$\frac{\text{Class-3 delay}}{\text{Class-2 delay}} \approx 4$$

Class-1 delay ≤ **1 ms**

$$\frac{\text{Class-(i+1) loss}}{\text{Class-i loss}} \approx 2$$

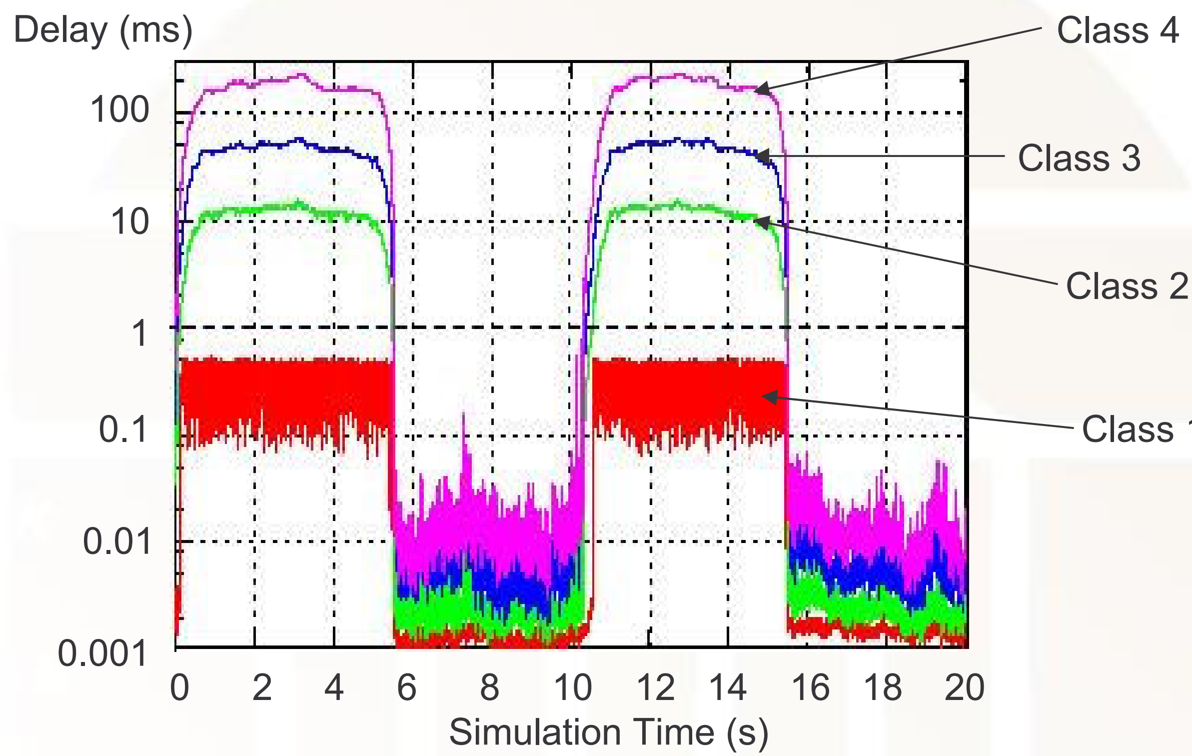

Simulation Results: Loss Rate Ratios

$$\frac{\text{Class-4 delay}}{\text{Class-3 delay}} \approx 4$$

$$\frac{\text{Class-3 delay}}{\text{Class-2 delay}} \approx 4$$

Class-1 delay ≤ **1 ms**

$$\frac{\text{Class-(i+1) loss}}{\text{Class-i loss}} \approx 2$$

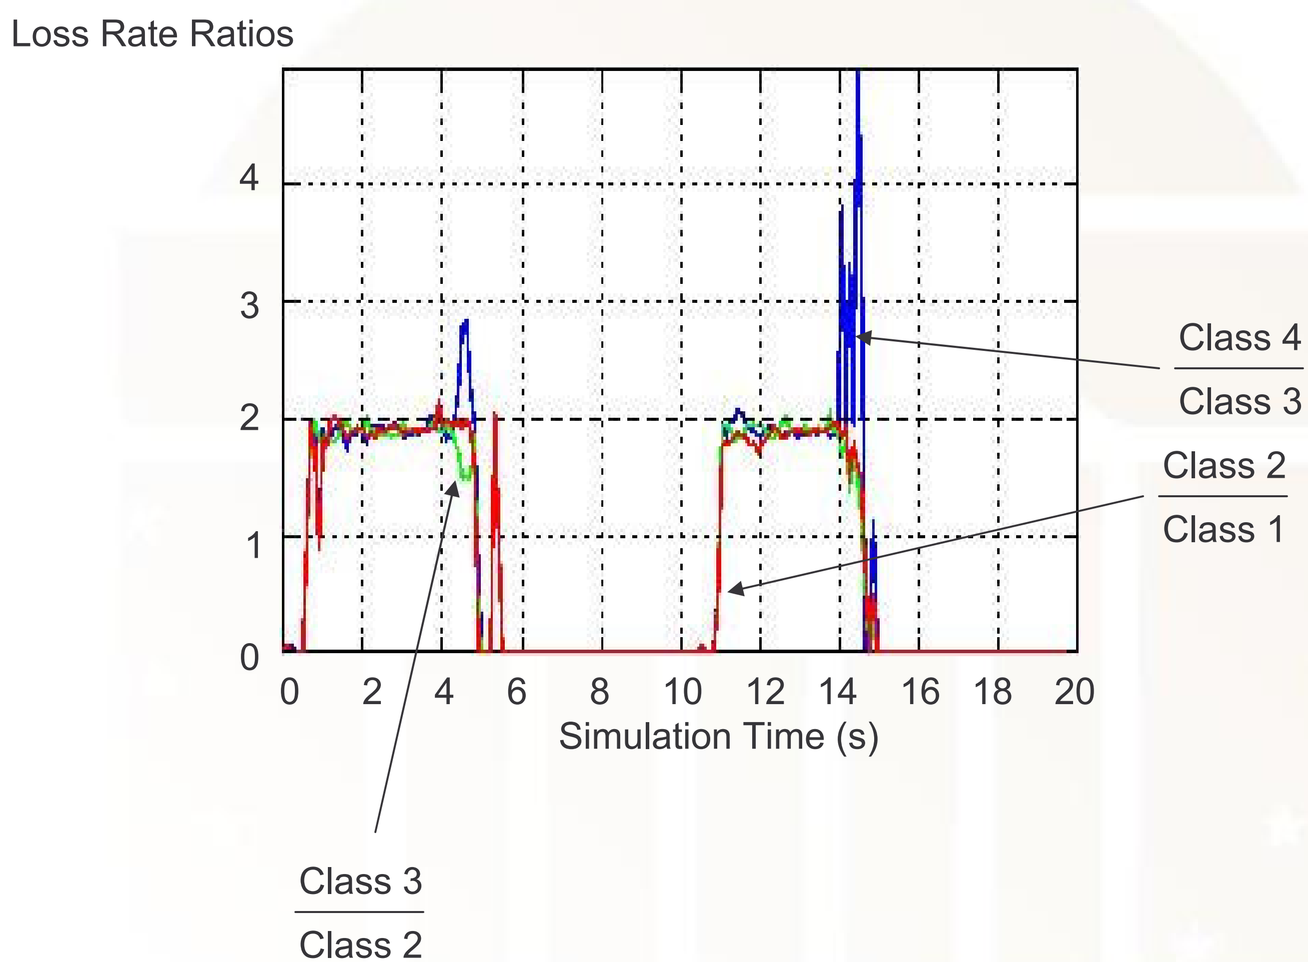

Feedback Control Solution to JoBS

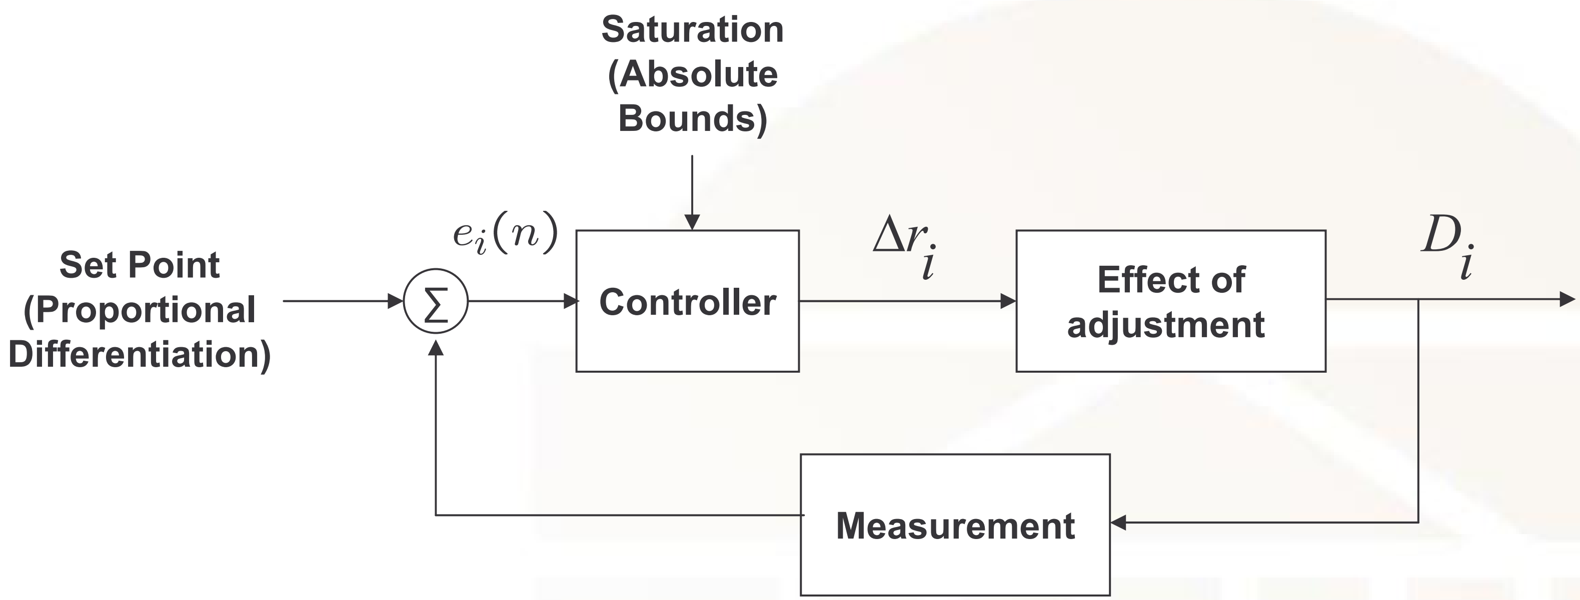

- Proportional controller: $\Delta r_i(n) = K(n) \cdot e_i(n)$
  - $e_i(n)$ is the deviation of the class-i delay from the desired proportional differentiation
  - $K(n)$ is a proportional coefficient
- Linearization of the non-linear system around an operating point.
  - Allows to use linear control theory tools (e.g., derivation of a stability condition)
- Losses are handled by a similar feedback mechanism

Feedback Control Solution to JoBS

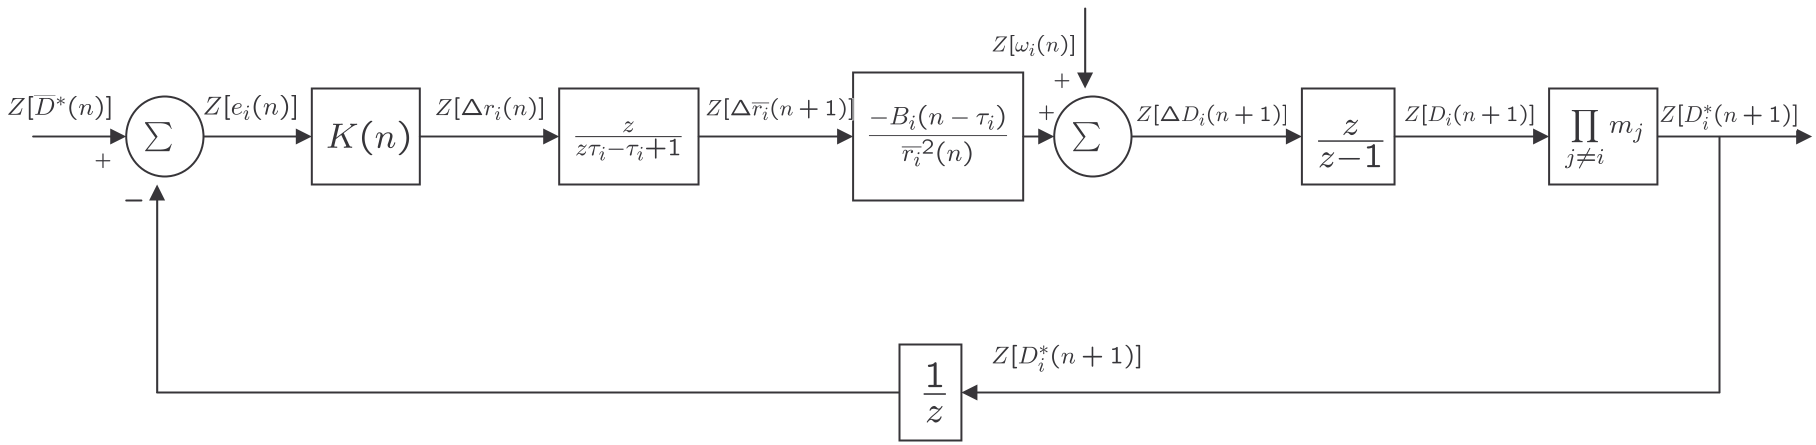

- Proportional controller: $\Delta r_i(n) = K(n) \cdot e_i(n)$
  - $e_i(n)$ is the deviation of the class-i delay from the desired proportional differentiation
  - $K(n)$ is a proportional coefficient
- Linearization of the non-linear system around an operating point.
  - Allows to use linear control theory tools (e.g., derivation of a stability condition)
- Losses are handled by a similar feedback mechanism

Conditions on the Delay Controllers

- Stability condition (proportional differentiation):

$$-2 \cdot \min_i \left\{ \frac{B_i(n)}{\prod_{j \neq i} m_j \cdot D_i^2(n)} \right\} \leq K(n) \leq 0$$

- Saturation effects (absolute delay/throughput guarantees):

$$K(n) \geq \max_i \left(\frac{r_{i,\min}(n) - r_i(n-1)}{e_i(n)} \right)$$

with

$$r_{i,\min}(n) = \max \left\{ \frac{B_i(n)}{d_i - D_i(n)}, \mu_i \cdot \chi_{B_i(n) \geq 0} \right\}$$

Evaluation in Testbed Network

- Implementation in FreeBSD kernel
  - Testbed of 6 Pentium IIIs 1Ghz with multiple Ethernet interfaces
  - Routers connected by 100 Mbps point-to-point Ethernet links
- Software is now part of ALTQ 3.1

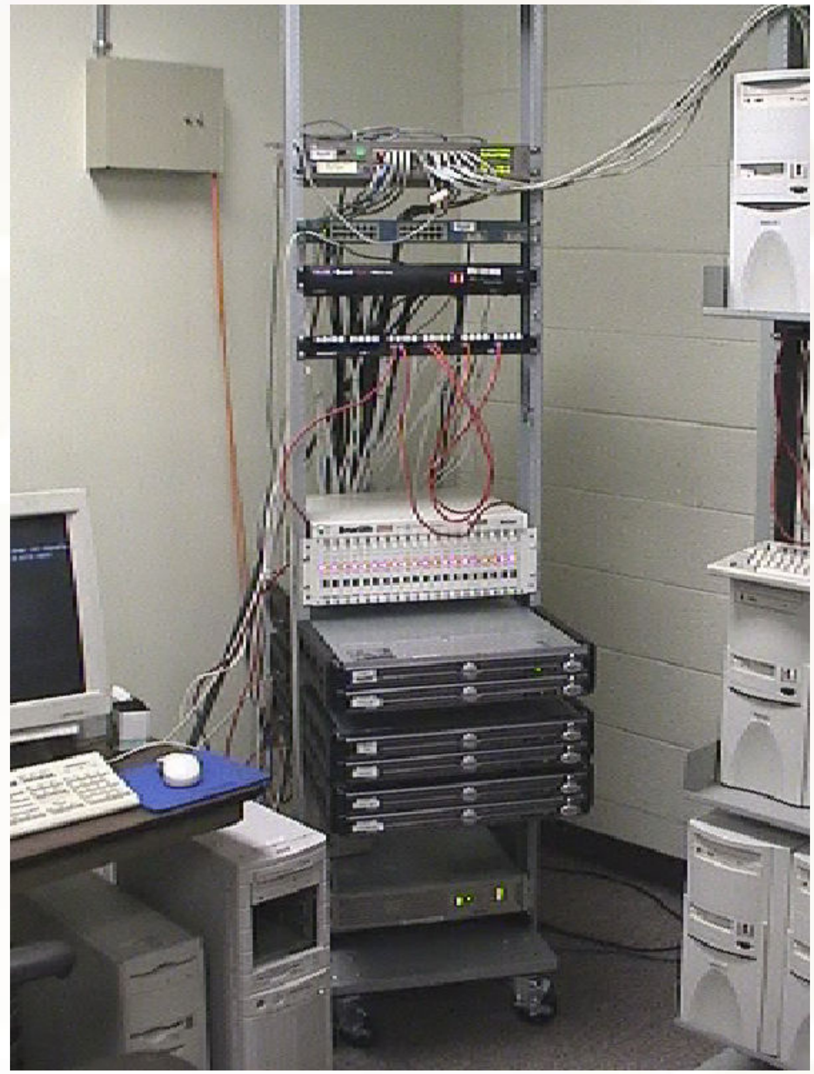

Experimental Setup

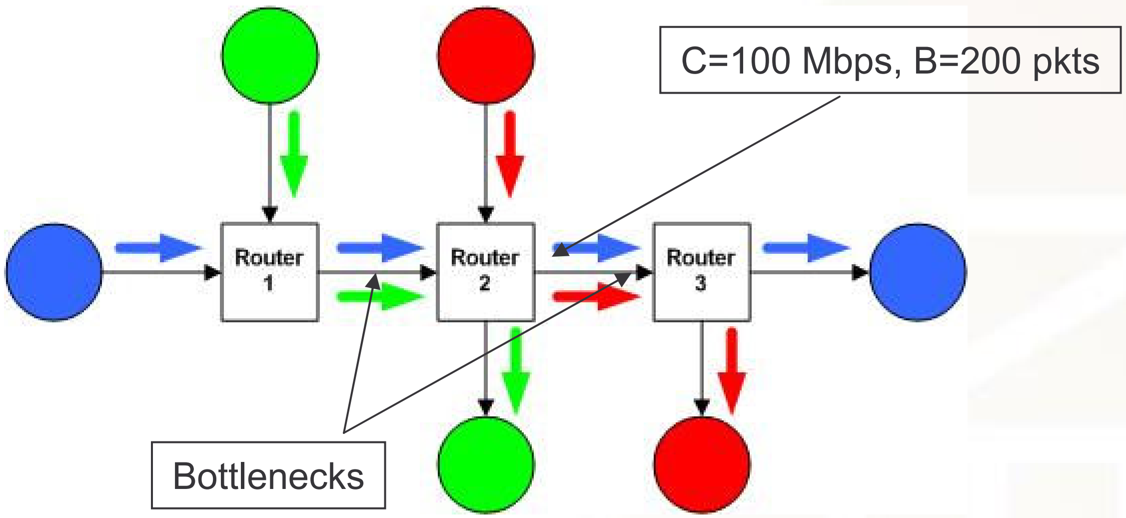

Class	No. of Flows	Proto.	Traffic
1	6	UDP	On-off
2	6	TCP	Greedy
3	6	TCP	Greedy
4	6	TCP	Greedy

Class	Abs. Delay Guarantee	Abs. Loss Guarantee	Abs. Throughput Guarantee	Proportional Delay Diff. Factor	Proportional Loss Diff. Factor
1	8 ms	1 %	-	-	-
2	-	-	35 Mbps		
3	-	-	-	2	2
4	-	-	-	2	2

Delay Differentiation (at Router 1)

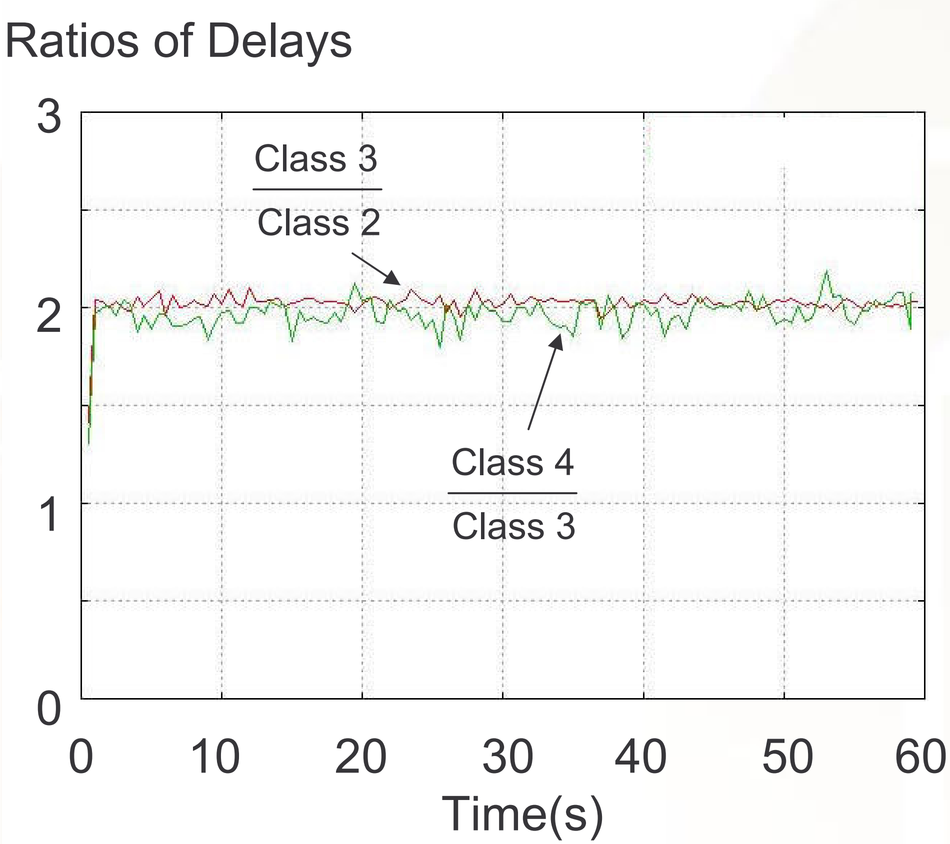

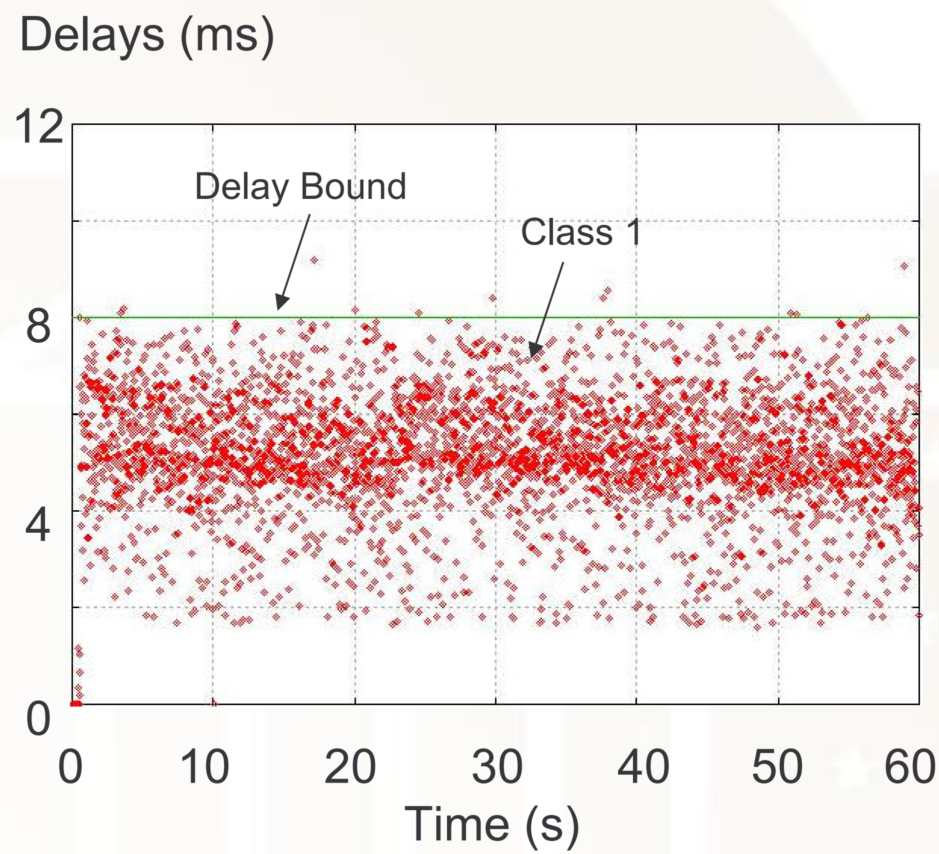

→ Similar results can be observed at Router 2

Loss Differentiation (at Router 1)

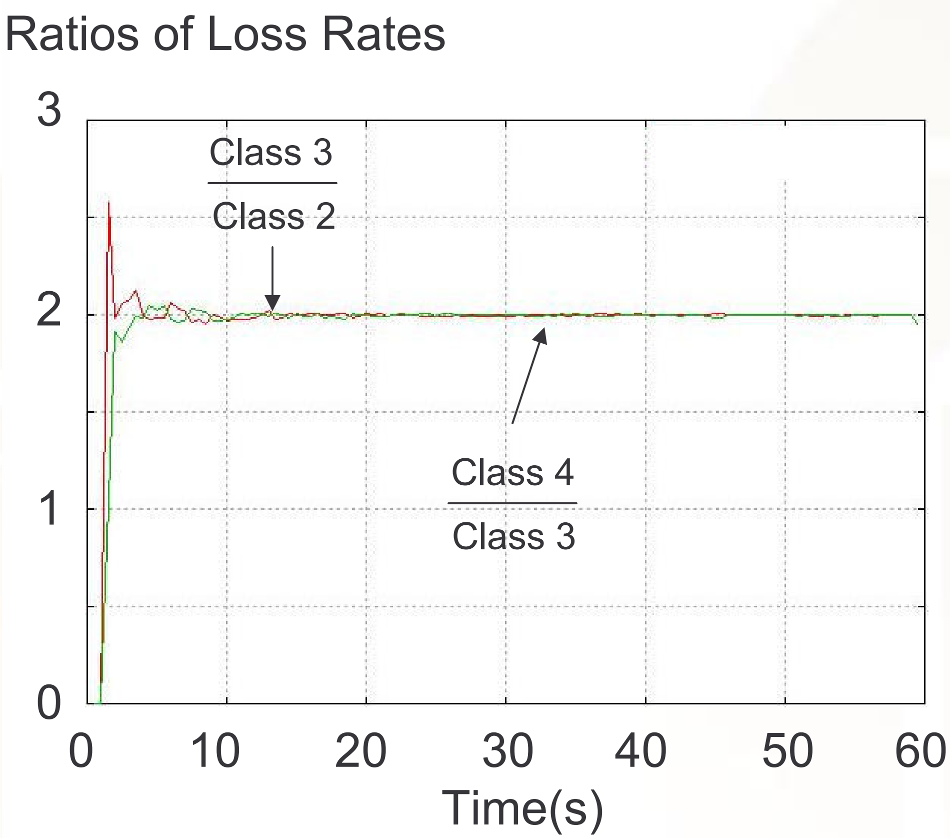

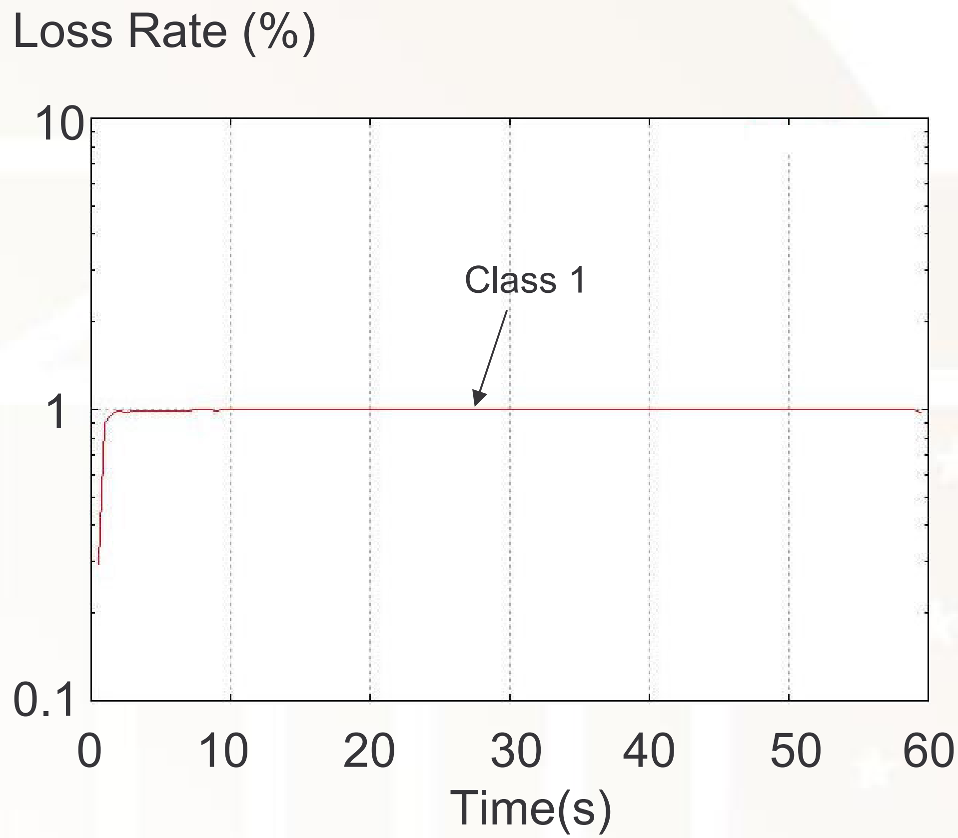

Throughput Differentiation (at Router 1)

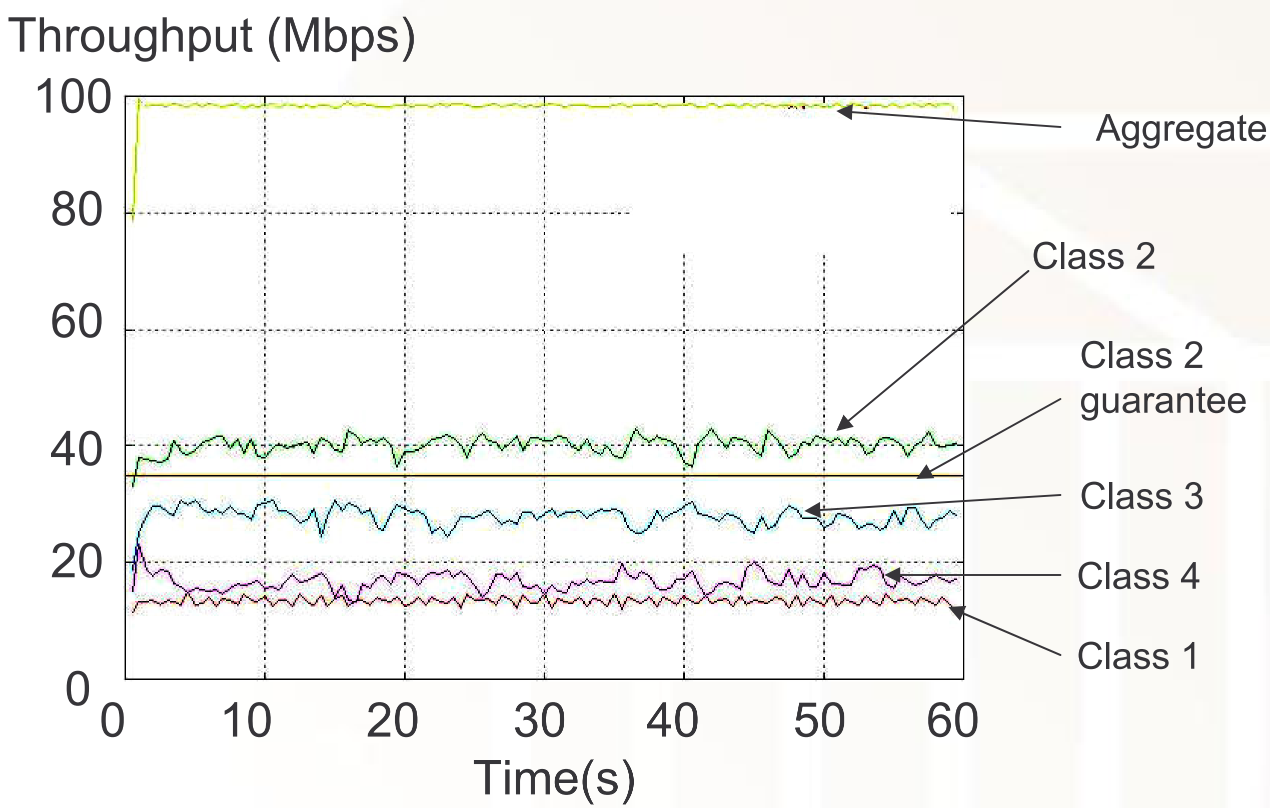

Extending JoBS to TCP traffic

- JoBS drops traffic if necessary
  - TCP traffic is sensitive to losses
- JoBS is a hop-by-hop scheme
  - TCP is an end-to-end protocol
  - Somehow must reduce traffic input to the network to avoid losses and service violations

- TCP congestion control:
  - Classical TCP congestion control: reduce traffic when there is a loss
  - Explicit Congestion Notification (ECN): routers can signal congestion to TCP sources when congestion is impending

- **Key Idea:**
  - Exploit congestion control mechanisms of TCP to reduce traffic
  - Combine ECN with JoBS scheme to control rate of TCP traffic (UDP traffic is just dropped at routers)

Controlling End-to-End TCP Traffic

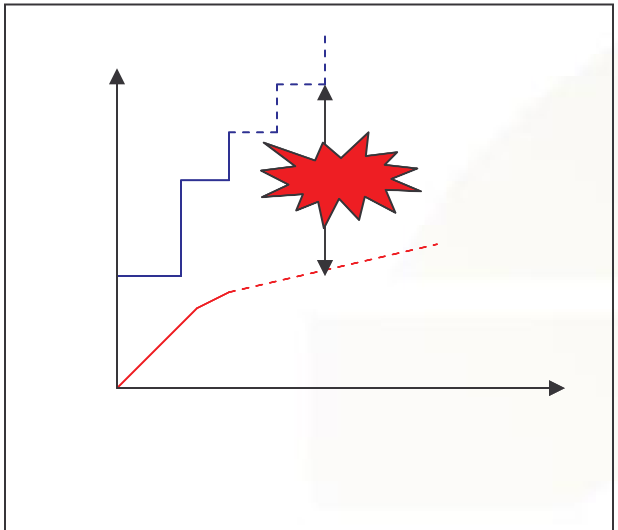

- Traffic must be reduced
- Mark packet(s) with congestion signal (ECN bit)

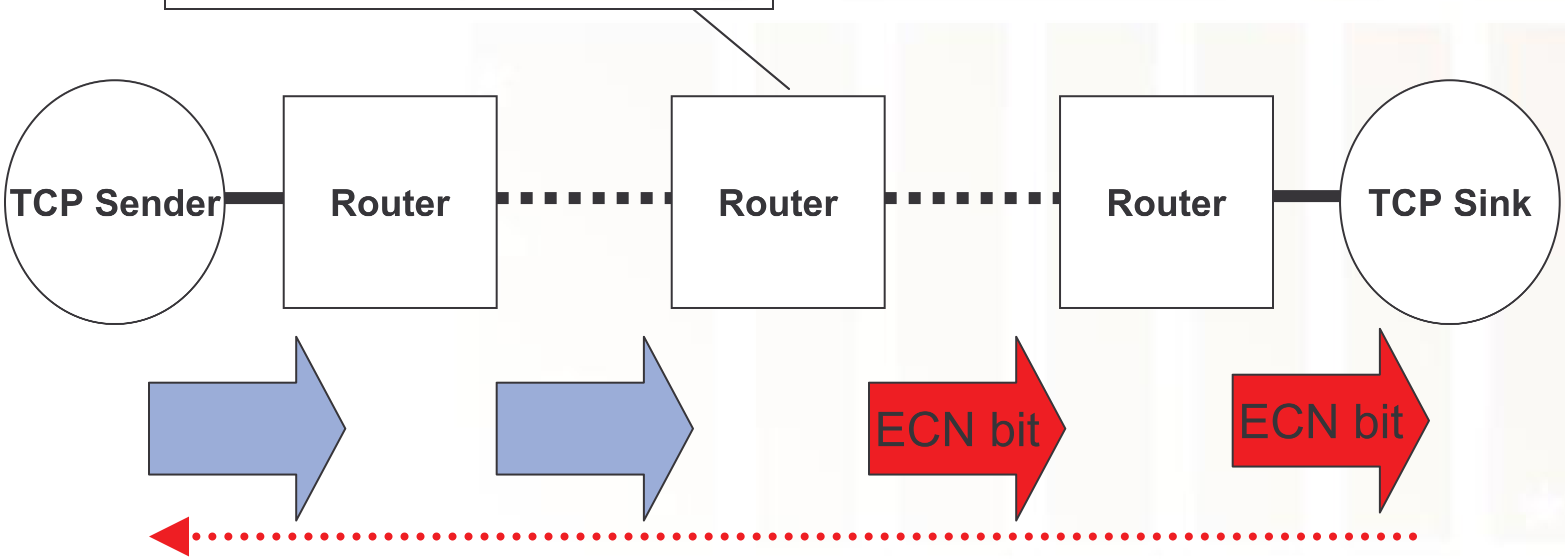

Avoiding Per-Flow Operations

- TCP congestion control operates on a per-flow basis
  - Violation of our design constraints
- Key observation:
  - Majority of TCP traffic is generated by a limited number of flows (“heavy-hitters”)
- Mechanisms:
  - Identify heavy-hitters via flow filtering
  - Control traffic from heavy-hitters via ECN marking
- Properties:
  - Can be used for loss avoidance and traffic regulation for service guarantees
  - Does not require any changes to TCP

ns-2 Simulation Setup

Round-trip times between 44 and 80 ms

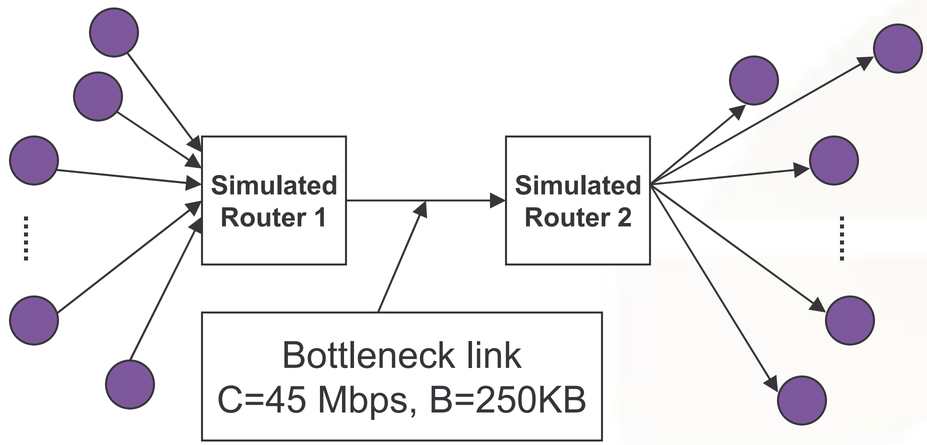

Class	No. of On-Off TCP Flows	No. of Greedy TCP Flows
1	5	3
2	10	3
3	15	3
4	20	3

Class	Abs. Delay Guarantee	Abs. Loss Guarantee	Abs. Throughput Guarantee	Proportional Delay Diff. Factor	Proportional Loss Diff. Factor
1	10 ms	1 %	5 Mbps	-	-
2	-	-	-		
3	-	-	-	4	2
4	-	-	-	4	2

Class-1 Delay Guarantees

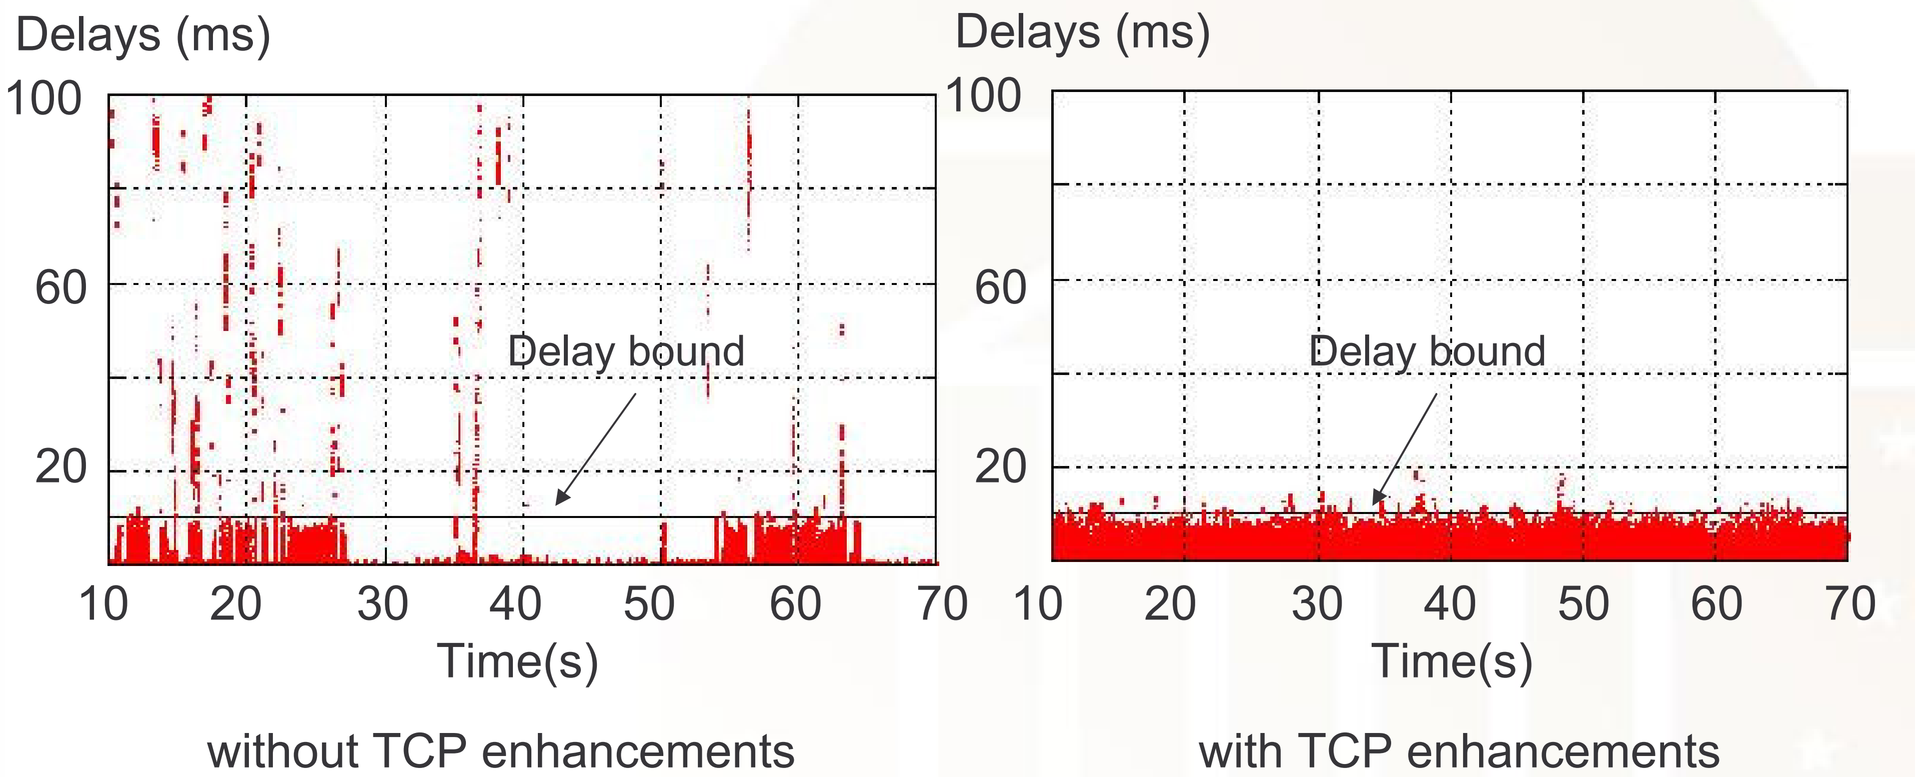

Loss Differentiation

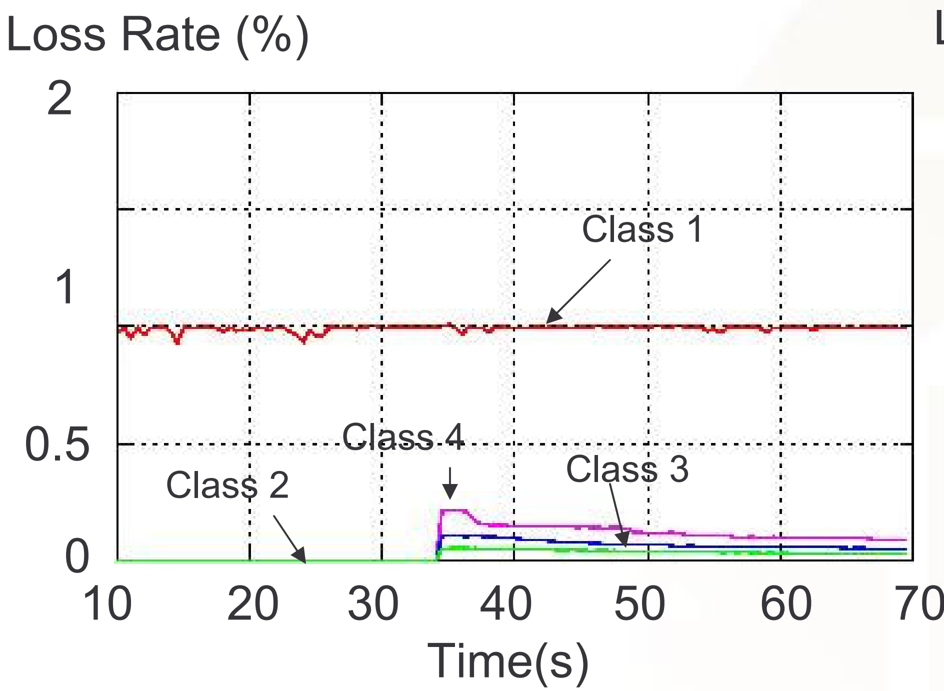

without TCP enhancements

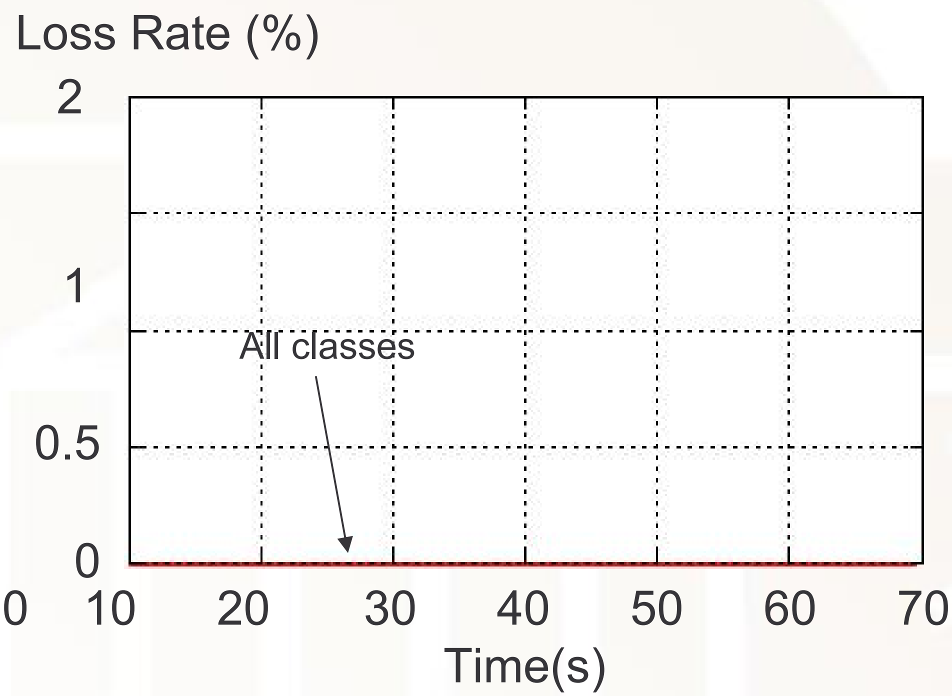

with TCP enhancements

Throughput Differentiation

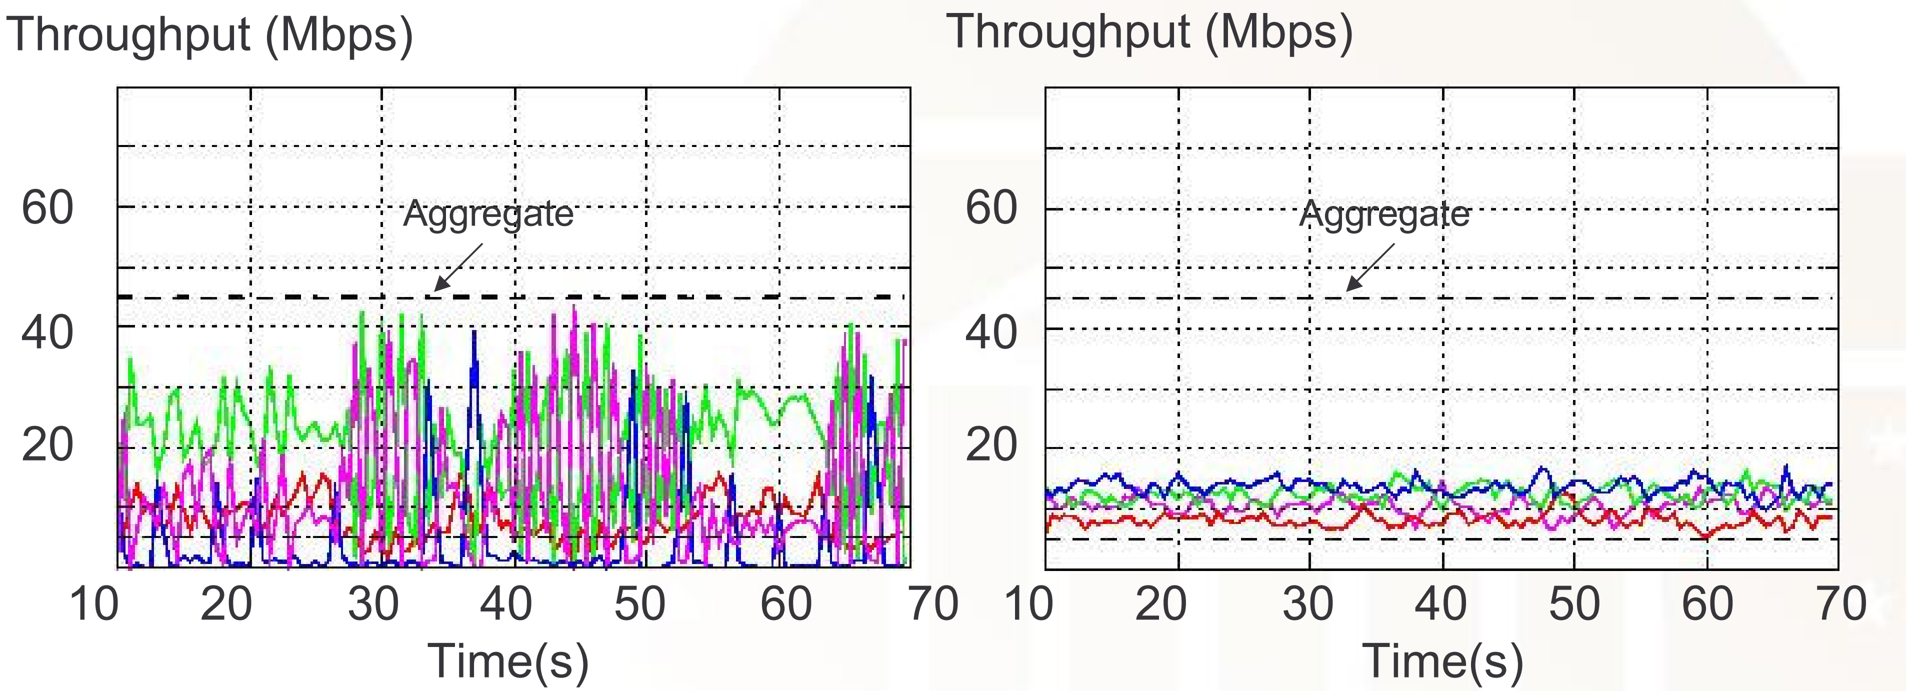

without TCP enhancements

with TCP enhancements

(Class 1 is in RED, Class 2 is in GREEN, Class 3 is in BLUE and Class 4 is in PURPLE)

Summary

- Quantitative Assured Forwarding
  - Subsumes all class-based architectures
- Key scheme: Joint Buffer Management and Scheduling
  - Reference algorithm (optimization)
  - Feedback-control algorithm
  - Implementation
  - Demonstrated efficiency by analysis, simulation, and measurements
- Reconcile end-to-end service with per-hop differentiation
  - ECN marking allows for avoiding losses in TCP
  - Fairness at the flow level without per-flow reservations
  - Can be a viable alternative to admission control and traffic policing

Future Work

- Class selection policies
  - Class selection by the network
    - Raise inter-domain cooperation issues
  - Class selection by applications
    - Need to enforce collaboration
- Extending the architecture to provide market incentives
  - “One of the impediments to the deployment of new services on the Internet is the lack of market incentives to improve network services and applications and to use them efficiently.” (NSF workshop report, 2003)

Questions?