JoBS
Joint Buffer Management and Scheduling for Differentiated Services

Jörg Liebeherr and Nicolas Christin

University of Virginia
Computer Science Department
Charlottesville, VA 22904

[jorg|nicolas]@cs.virginia.edu

Outline

- **The Context**
  - DiffServ with AF guarantees
- **The Problem**
  - Differentiated Services, Differentiated Services Enhancements
- **Our Approach**
  - Joint Buffer Management and Scheduling (JoBS)
- **Conclusions**

Integrated Services (IntServ)

- ***Provide absolute per-flow guarantees:***
  - upper bound on delay
  - upper bound on loss rate

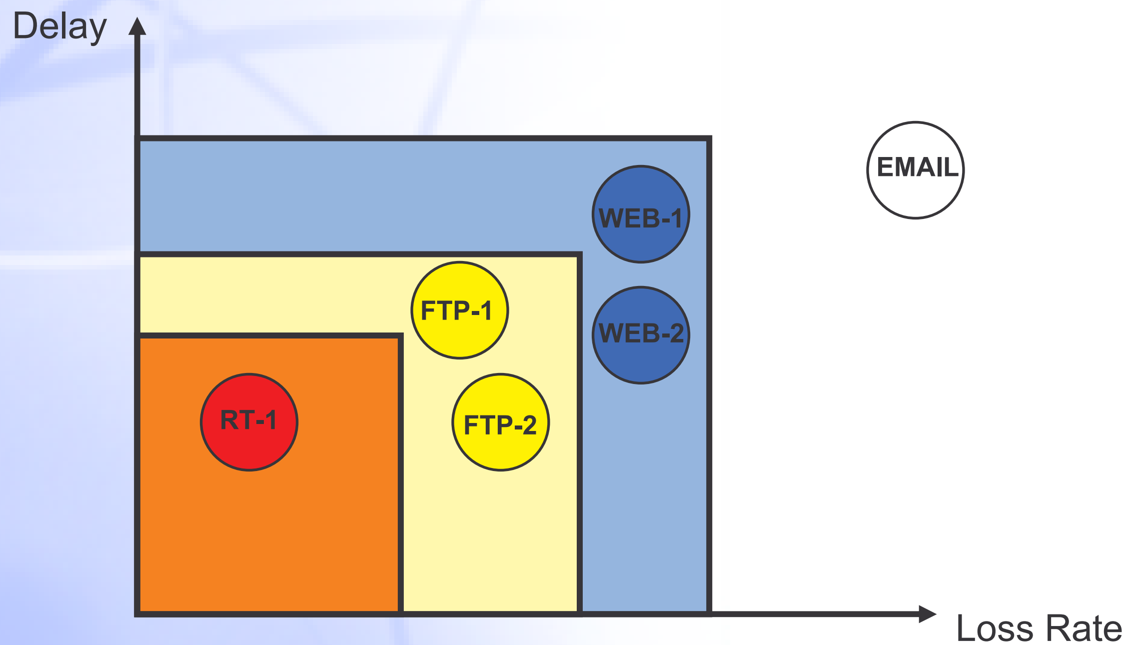

Differentiated Services (AF/DiffServ)

- ***Provide service differentiation between traffic classes***
  - Service differentiation is only qualitative

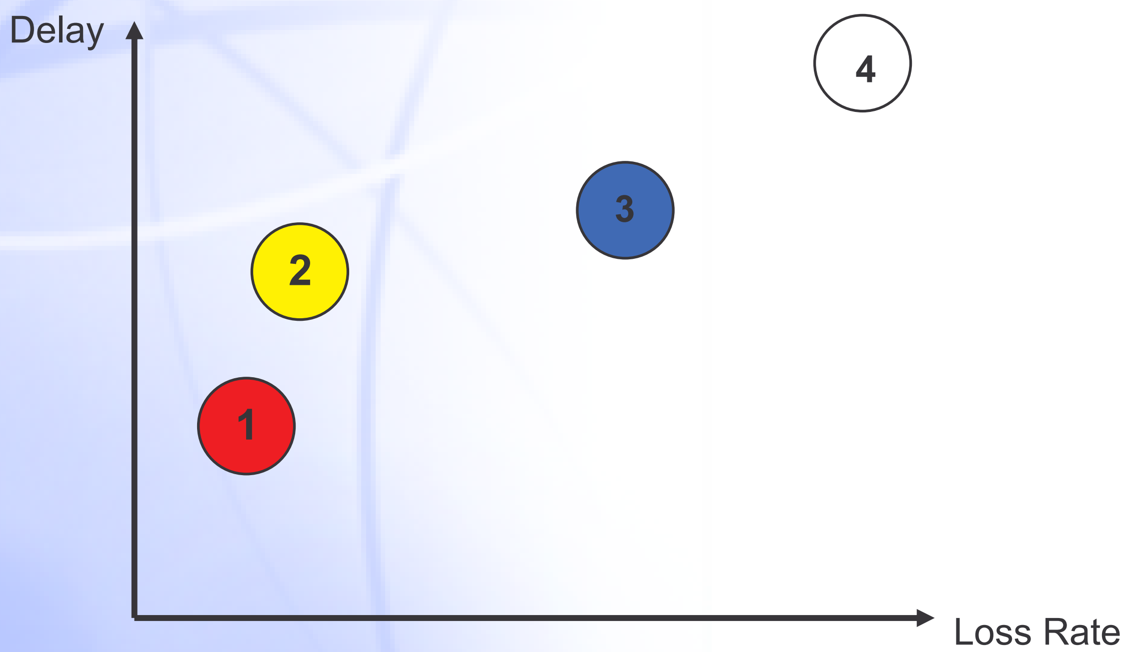

Differentiated Services Enhancements

- ***Provide quantifiable guarantees within an AF/ DiffServ context***

  e.g., Proportional Delay and Loss Differentiation

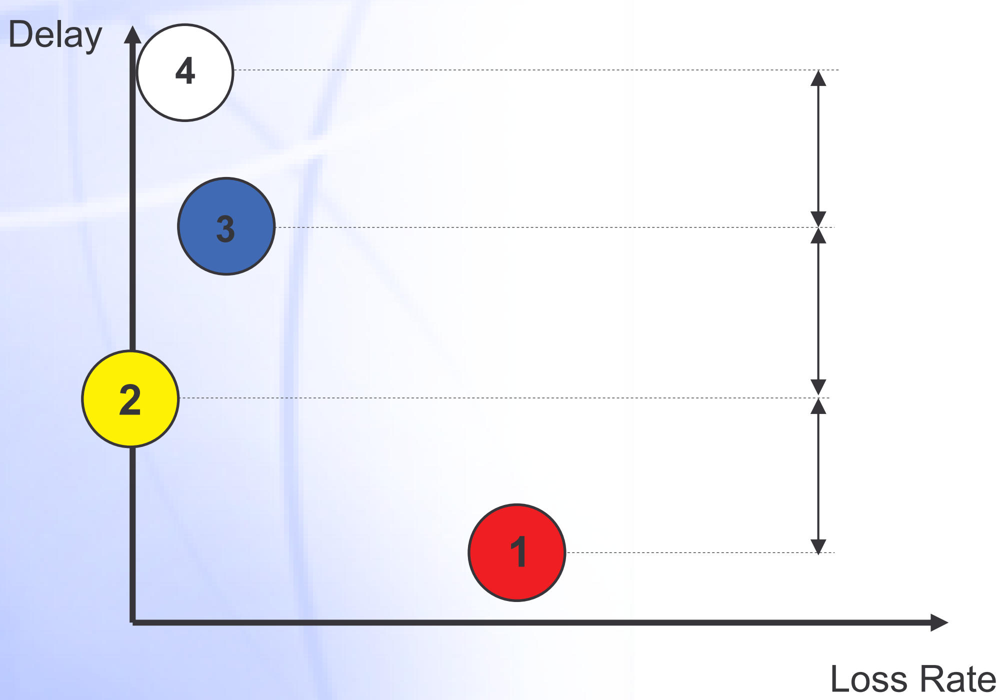

JoBS: Joint Buffer Management and Scheduling for Differentiated Services

JoBS: Joint Buffer Management and Scheduling

- Offer **proportional and absolute** guarantees for **both loss and delay**

$$\frac{\text{Class-2 delay}}{\text{Class-1 delay}} \cong 4 \quad \text{or:} \quad \frac{\text{Class-2 loss rate}}{\text{Class-1 loss rate}} \cong 2$$

$$\text{Class-2 delay} \leq 5 \qquad \text{Class-2 loss rate} \leq 10^{-9}$$

- If necessary, relax guarantees in a given preference order

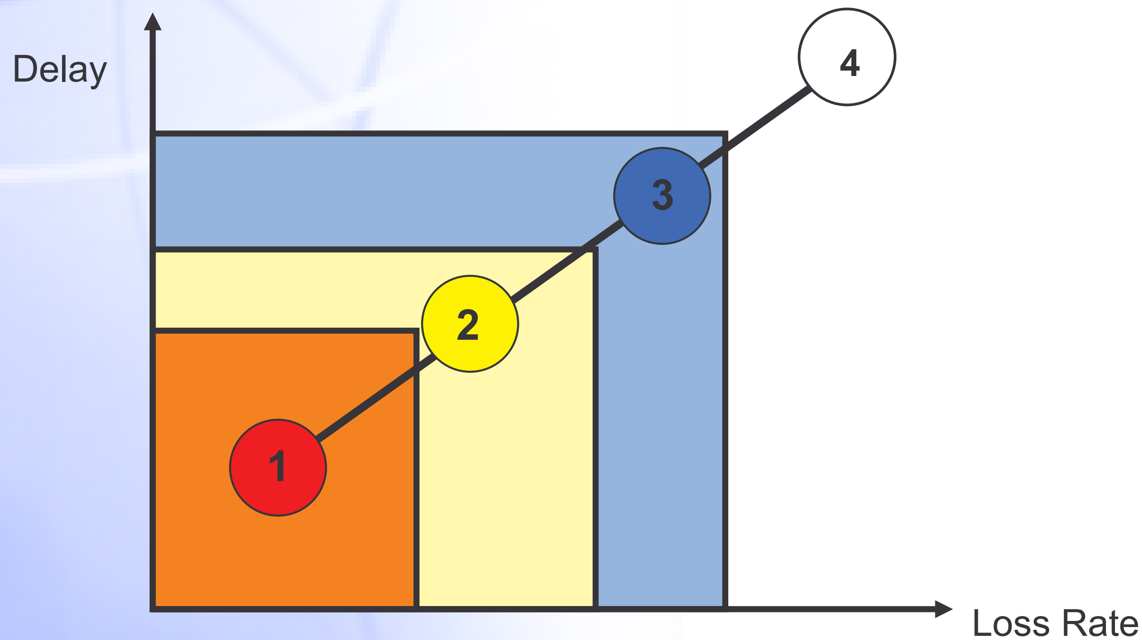

JoBS: Joint Buffer Management and Scheduling for Differentiated Services

Related Work (on enhanced AF Service)

- Proportional Delay and Loss Differentiation (Dovrolis et. al.)
- Mean-Delay Proportional Scheduler (Barghavan et. al.)
- ABE Service (Hurley et. al.)
- Several papers at this workshop

JoBS – Joint Scheduling and Buffer Management

- ***JoBS operates at the output port of a router***
- ***JoBS mechanisms:***
  - Service rate allocation to traffic classes
  - Service rate allocation is periodically adjusted
  - Rate allocation is based on projections of delays and loss rate
  - If no feasible rate allocation exists, drop traffic

- ***Contributions***
  - Progress on the question: How strong can we make AF service?
  - Propose a formal framework to view both loss and delay differentiation in an DiffServ/AF context

Arrivals, Departures, and Losses at a node

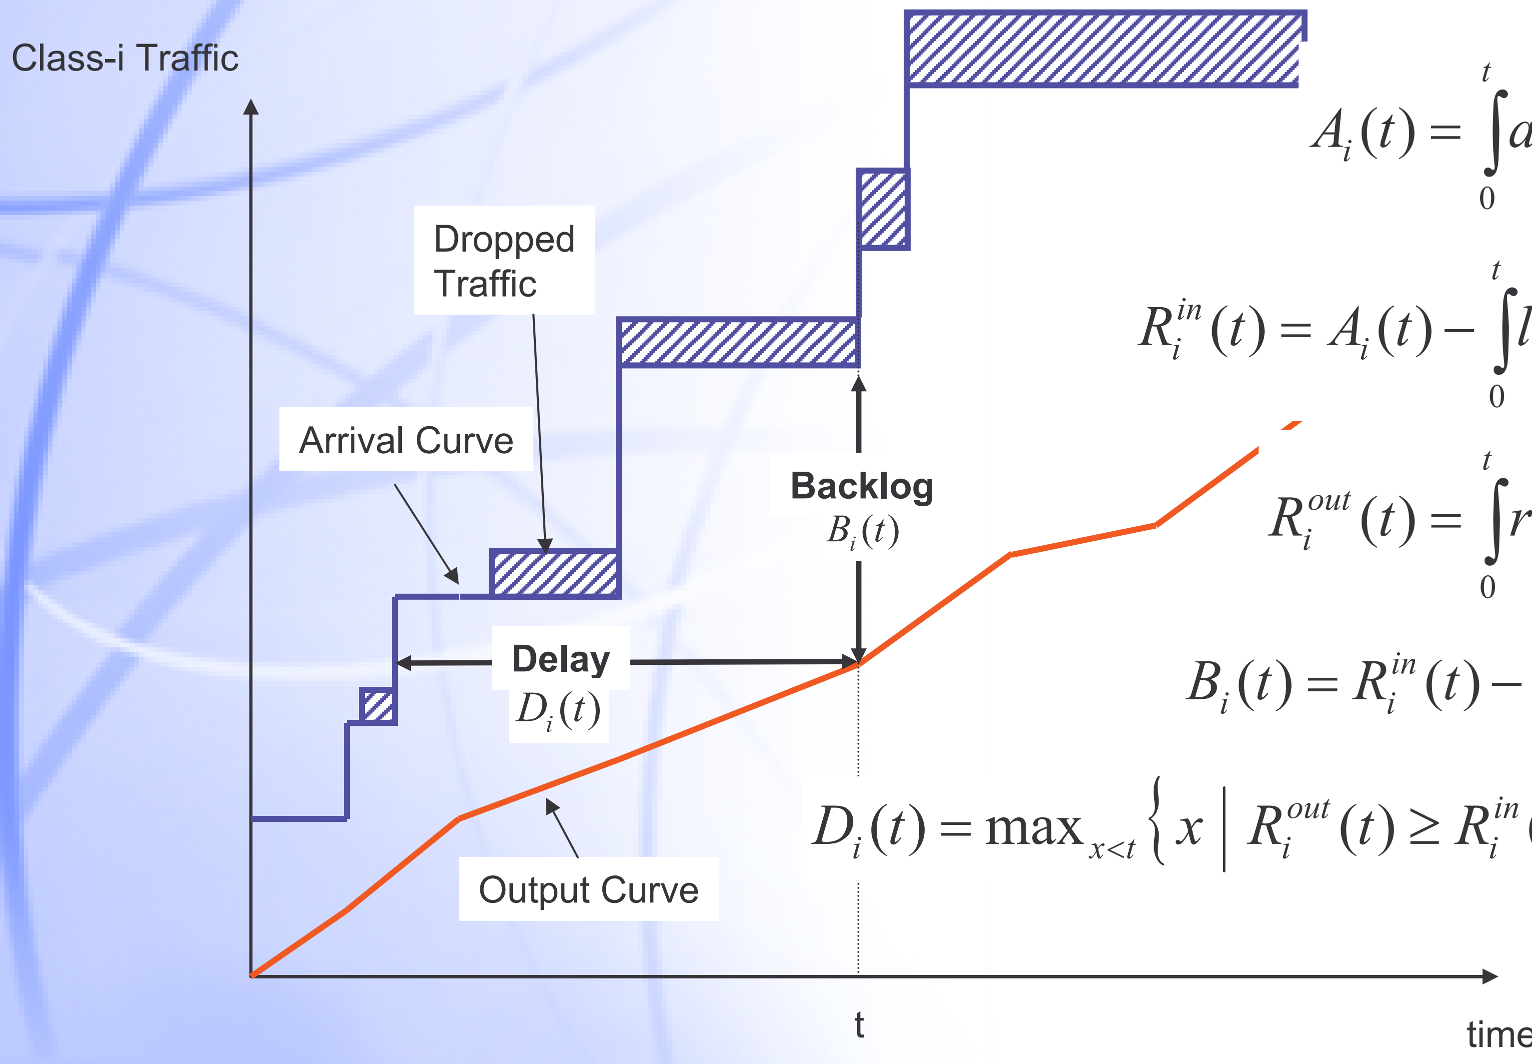

$$A_i(t) = \int_0^t a_i(x)dx$$

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

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

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

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

Rate Projections

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

- Current rate allocation does not change
- No new arrivals

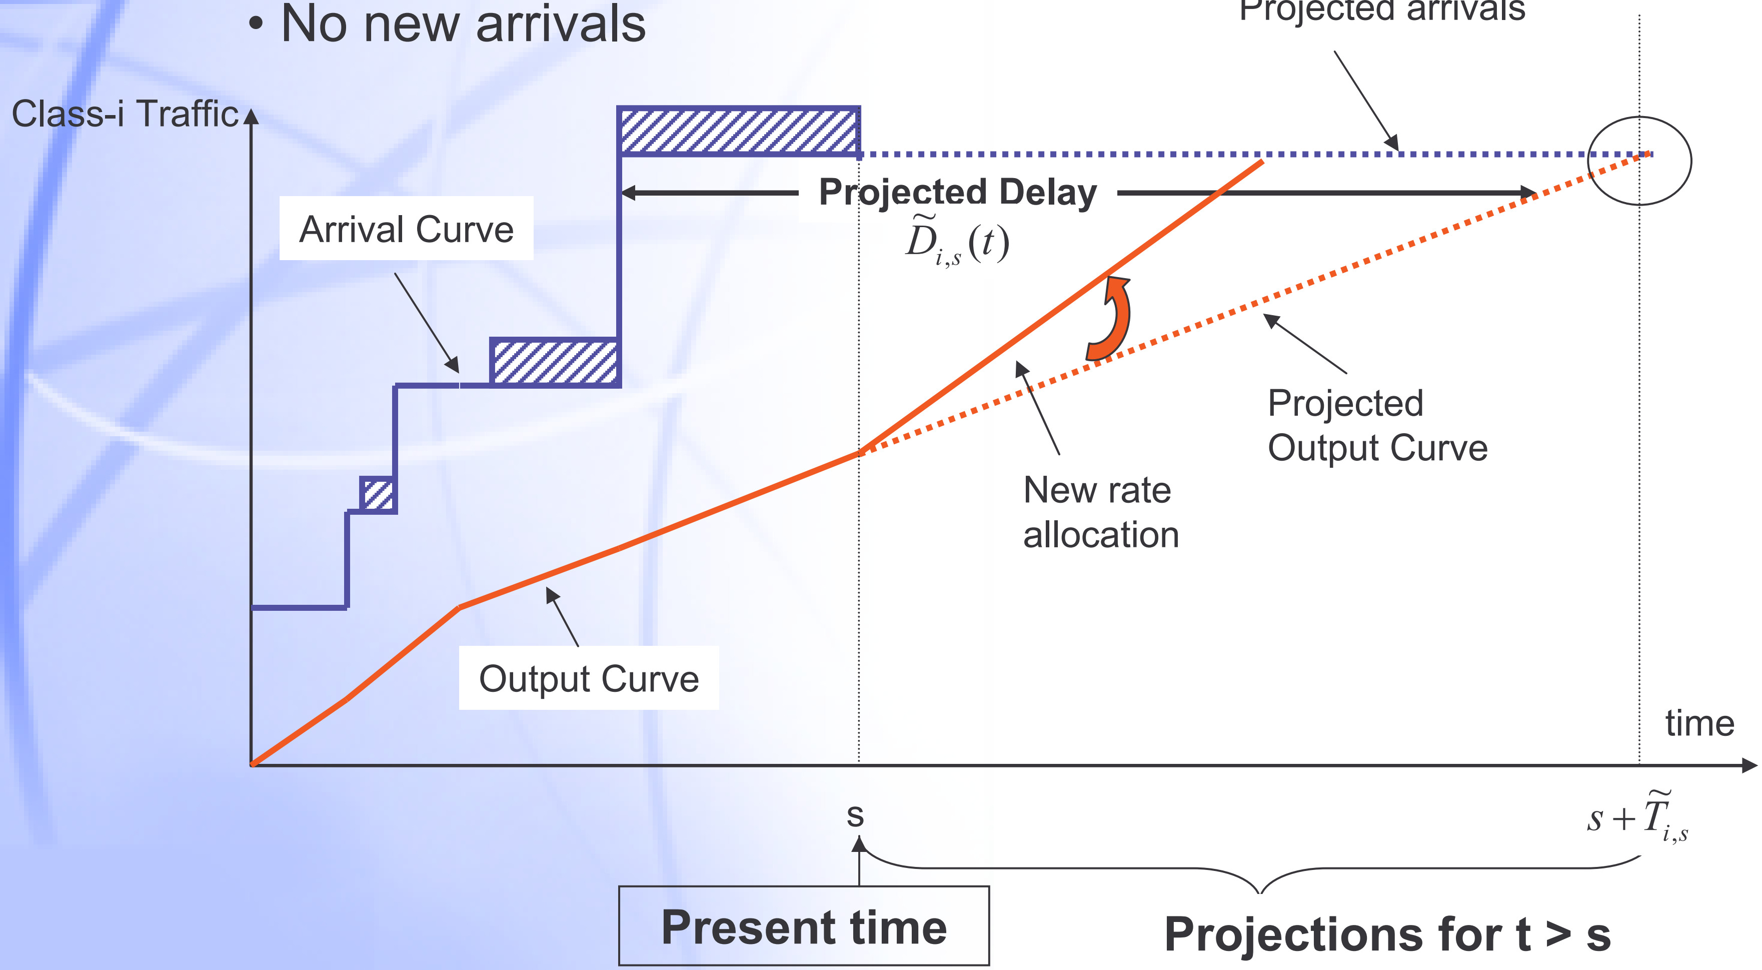

Projections

For $s < t \leq s + \widetilde{T}_{i,s}$:

$$\widetilde{R}_{i,s}^{in}(t) = R_i^{in}(s)$$

$$\widetilde{R}_{i,s}^{out}(t) = R_i^{out}(s) + (t-s)r_i(s)$$

$$\widetilde{B}_{i,s}(t) = \widetilde{R}_{i,s}^{in}(t) - \widetilde{R}_{i,s}^{out}(t)$$

$$\widetilde{D}_{i,s}(t) = \max_{t-s<x<t}\left\{x \;\middle|\; \widetilde{R}_{i,s}^{out}(t) \geq \widetilde{R}_{i,s}^{in}(t-x)\right\}$$

JoBS

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

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

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

System Constraints

- Scheduler is work-conserving

$$\sum_i r_i(t) = C$$

- Finite buffer size

$$\sum_i B_i(t) \le B$$

Delay Constraints

- **Absolute delay constraints**

$$\max_{s<t<s+\tilde{T}_{i,s}} \tilde{D}_{i,s}(t) \le d_i$$

- **Relative delay constraints**

$$\frac{\overline{D}_{i+1,s}}{\overline{D}_{i,s}} \approx k$$

where

$$\overline{D}_{i,s} = \frac{1}{\tilde{T}_{i,s}} \int_{s}^{s+\tilde{T}_{i,s}} \tilde{D}_{i,s}(x)\,dx$$

Loss Constraints

- Loss is defined as the dropped traffic in the current busy period

$$p_{i,s} = \frac{\int_{t_0}^{s} l_i(x)\, dx}{\int_{t_0}^{s} a_i(x)\, dx}$$

- **Absolute loss constraints**: $p_{i,s} \le L_i$

- **Relative loss constraints**: $\frac{p_{i+1,s}}{p_{i,s}} \approx k'$

Objective function

- **First Goal:** Avoid losses if possible
- **Second Goal:** Hang on to current rate allocation

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

- This is a non-linear optimization problem

  → specify heuristic algorithm which approximates optimal solution

JoBS (heuristic)

- Decompose optimization into a number of small computations
- Use virtual-clock type algorithm to implement rate allocation
- Run rate allocation only once for every N packets, or if buffer overflows

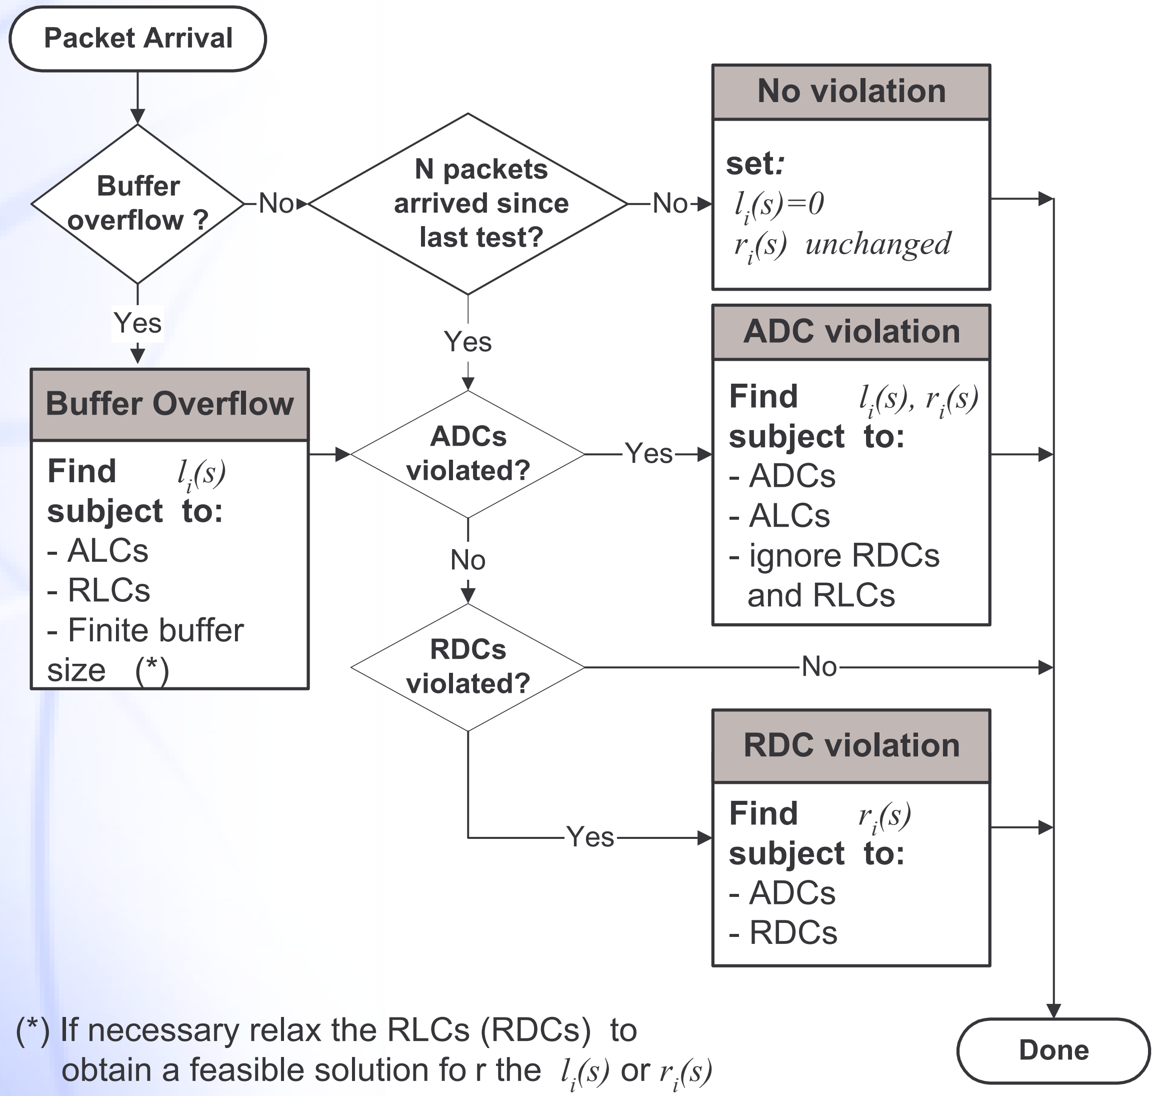

(*) If necessary relax the RLCs (RDCs) to obtain a feasible solution fo r the $l_i(s)$ or $r_i(s)$

Experimental Setup: Single Node

- Output link capacity = 1 Gbps,
- Buffer size = 6.25MB,
- Bursty arrival pattern: superposition of 200-550 Pareto sources (α=1.2).
- The offered load curve varies between 70% and 150% of the link capacity,
- 4 traffic classes,
- Each class contributes 25% of the total traffic.

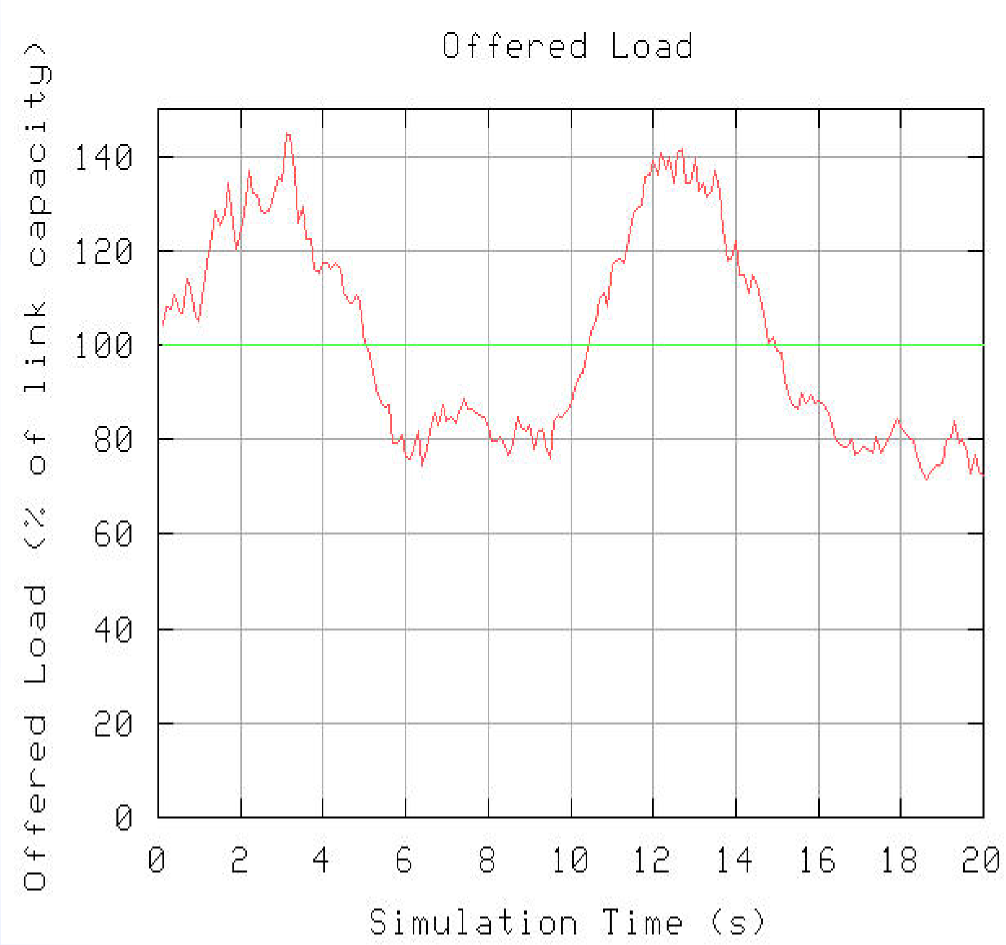

Simulation Results: Delay

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

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

$$\frac{\text{Class-2 delay}}{\text{Class-1 delay}} \approx 4$$

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

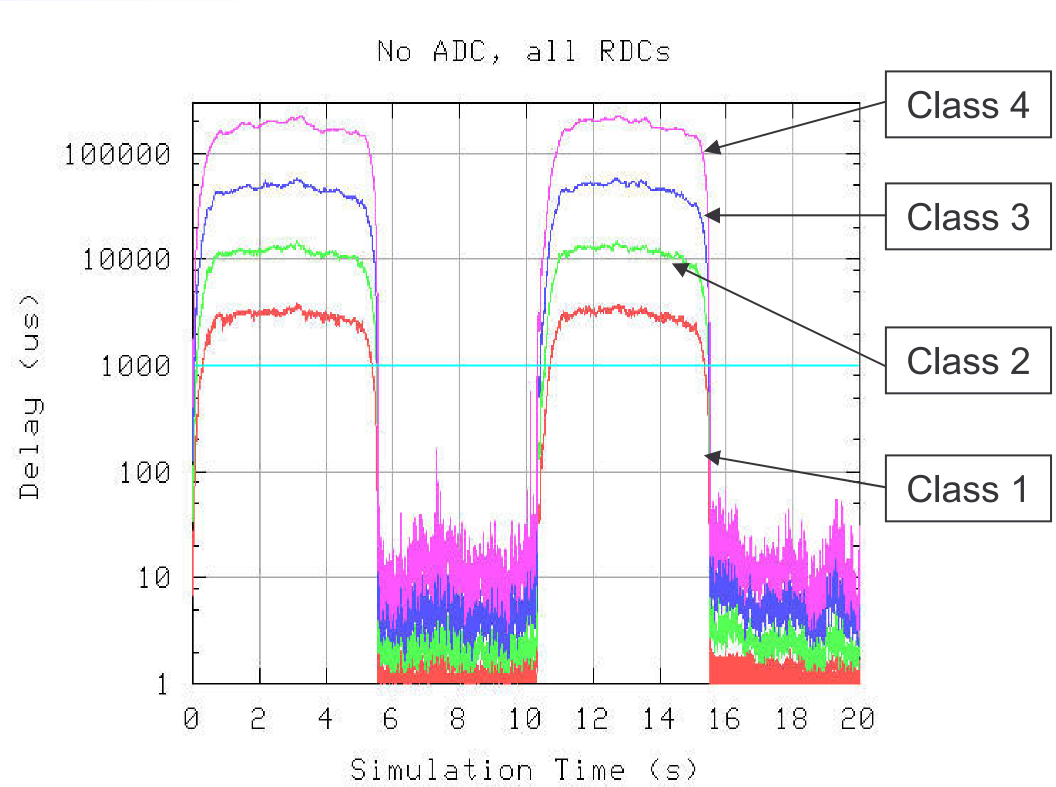

- ***JoBS (heuristic)***

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**

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

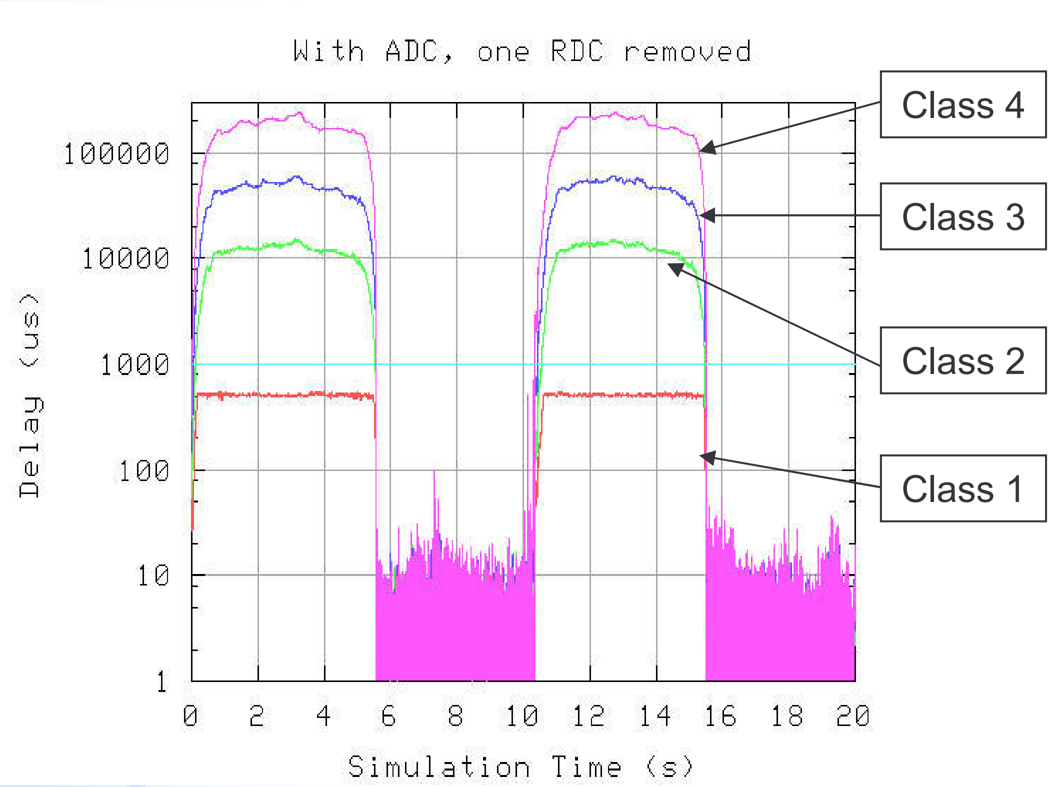

- ***JoBS (heuristic)***

Simulation Results: Loss

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

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

$$\frac{\text{Class-2 delay}}{\text{Class-1 delay}} \approx 4$$

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

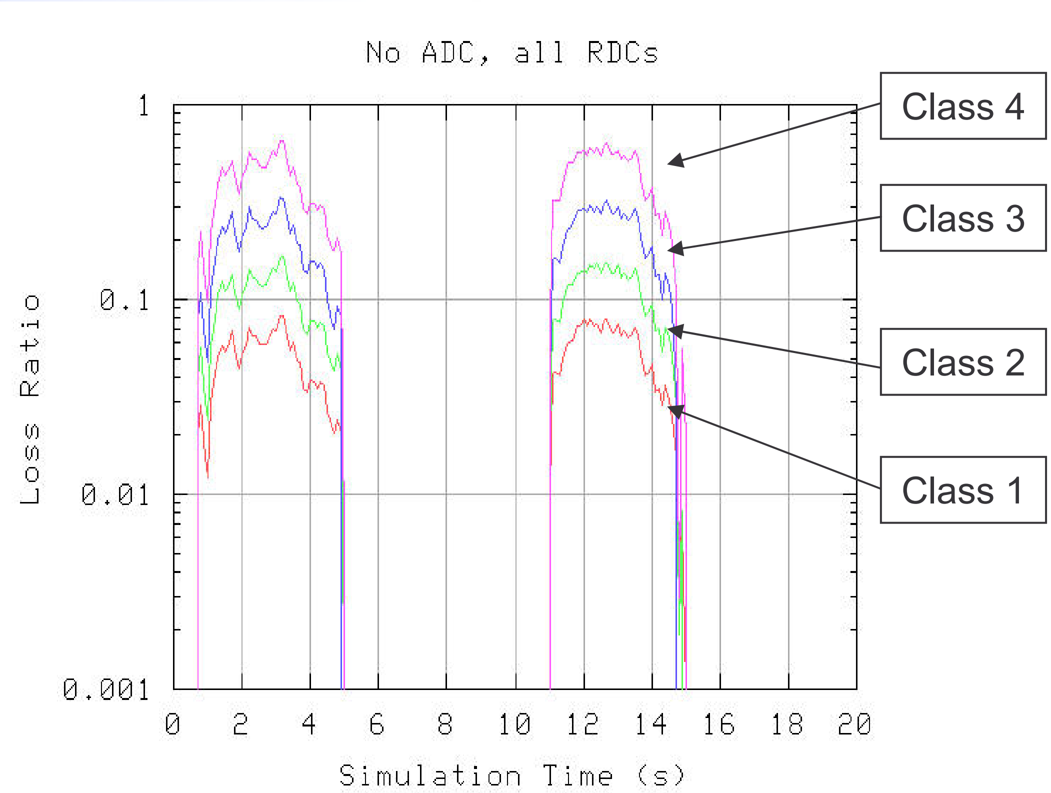

- ***JoBS (heuristic)***

Comparison with Literature: Delay

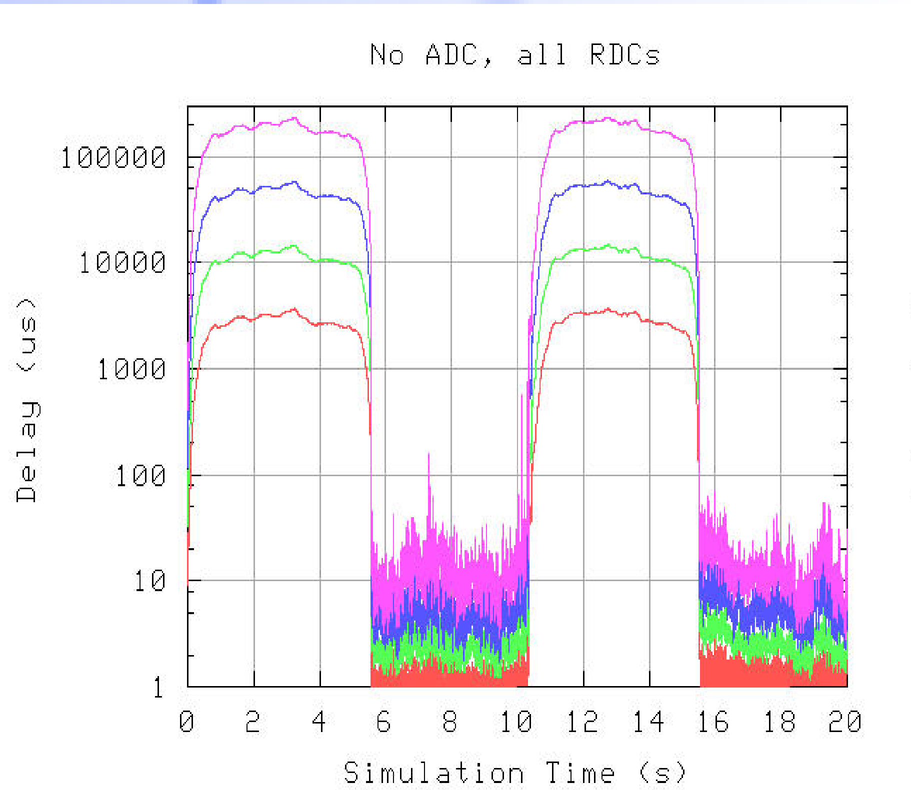

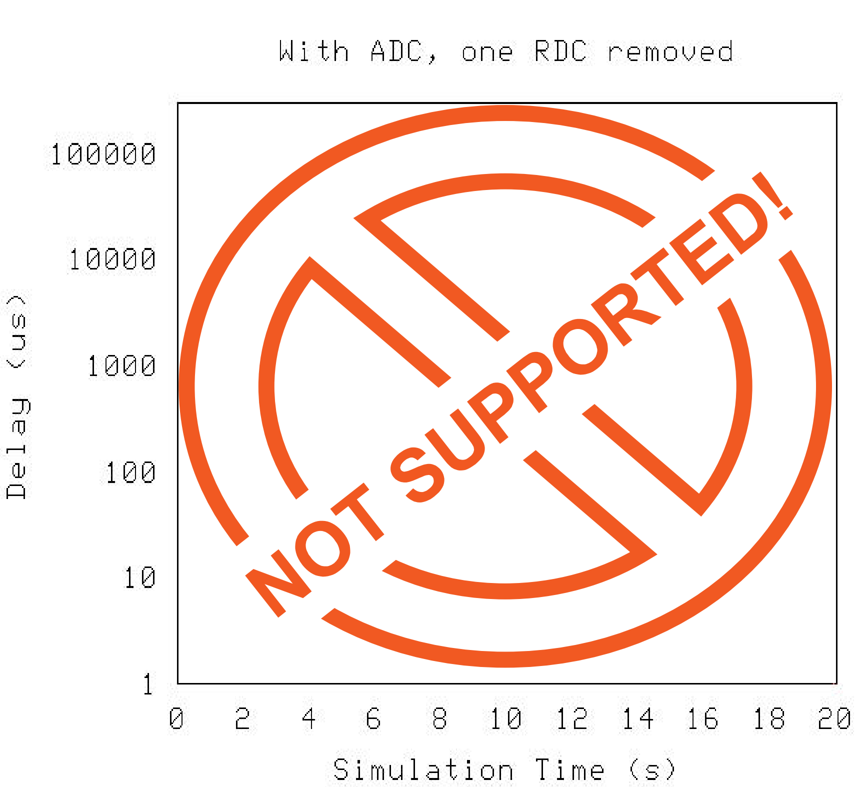

- ***Waiting Time Priority+Proportional Loss Rate Dropper by Dovrolis et.al.***

Comparison with Literature: Loss

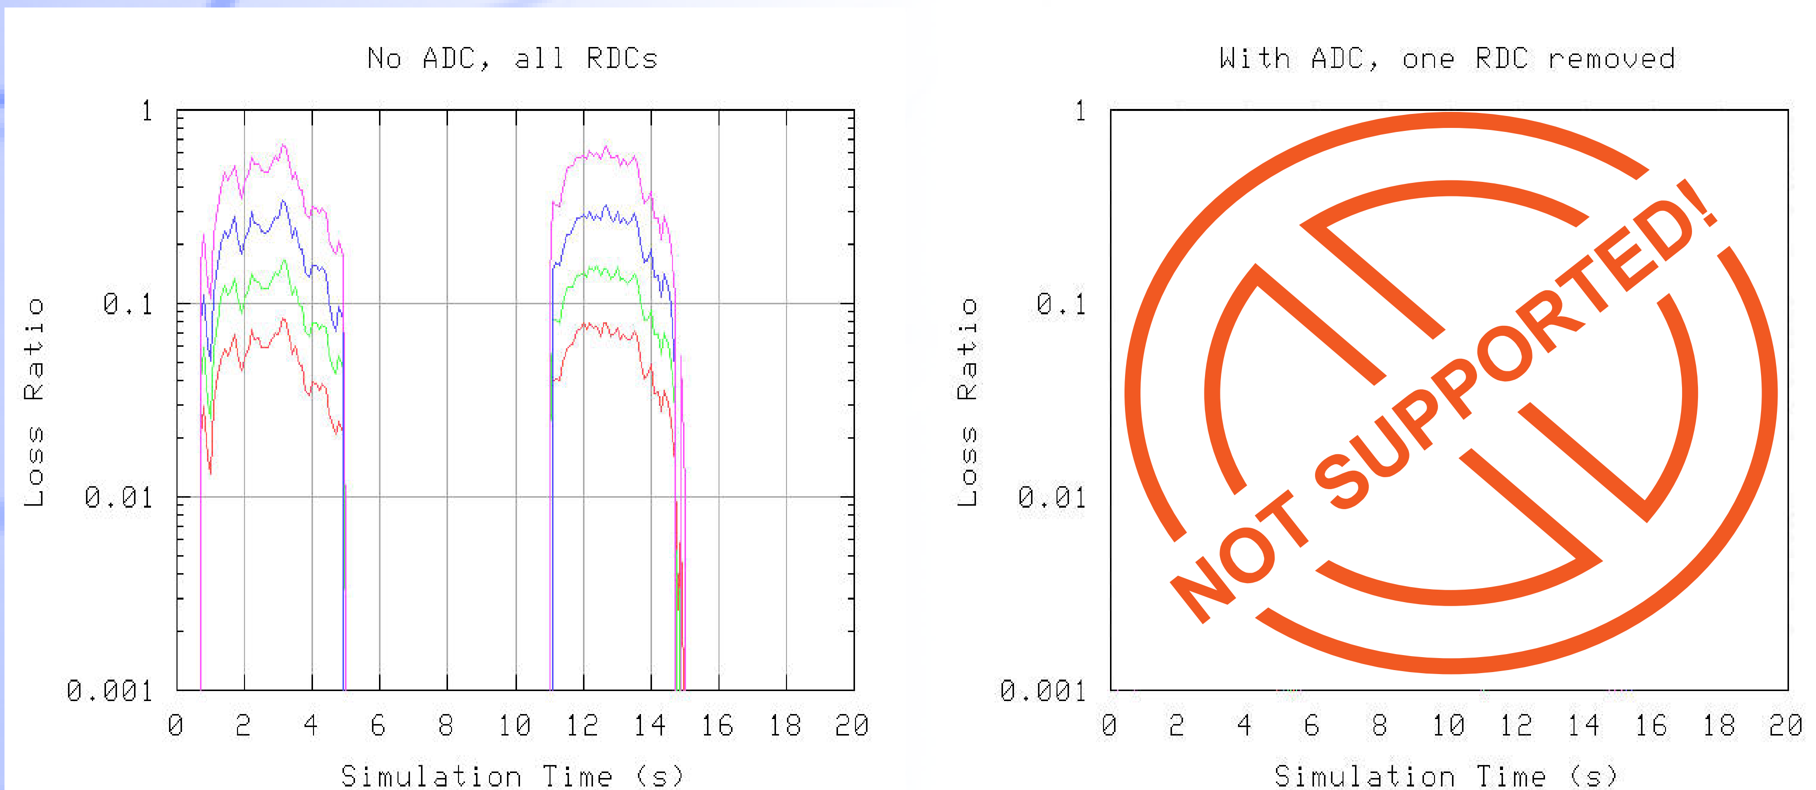

- ***Waiting Time Priority+Proportional Loss Rate Dropper by Dovrolis et.al.***

Conclusions

- Formal approach to enhanced differentiated services
- Tackle both loss and delay differentiation
- Provide absolute as well as relative guarantees.

- ***Current issues:***
  - Integration of TCP congestion control mechanisms
  - Implementation in ALTQ (100 Mbps) and Intel IXP (1 Gbps)

- Additional information (including applet demonstration) available at

  `mng.cs.virginia.edu`