Simple Performance Models of Differentiated Services Schemes for the Internet

Martin May, Jean Bolot, and Alain Jean-Marie

INRIA Sophia-Antipolis
2004 route des Lucioles, BP 93
06902 Sophia-Antipolis Cedex

E-mail: `mmay,bolot,ajm@sophia.inria.fr`

HTTP://www.inria.fr/rodeo/personnel/mmay/diff-serv/

June 12 , 1998

Overview

- **Augment Best-Effort of the current Internet**
  - No guarantees
  - No isolation (all packets receive the same service)
- **IntServ**
  - Hard guarantees
  - Flow isolation
  - Protocol to propagate the reservation information in the network (RSVP, YESSIR)
  - Define Services (Guaranteed Service, Controlled Load, ...)
  - ISSLL (Integrated Services over Specific Lower Layers)

The “new” Solution: Diff-serv

- **The concept of differentiated services is as old as the Internet**
  – IP TOS bits defined but never used
  – The DE bit in Frame Relay
  – CLP bit in ATM
- **Diff-Serv Architecture**
  – Offer services for aggregated traffic
  – Instead of maintaining state of each and every flow, discriminate the packets according to their TOS field
  – Forcing as much complexity out of the core of the network into edge
  – All flows entering the network must be policed
  – Offer ether throughput or delay improvements

The Differentiated Services Architecture

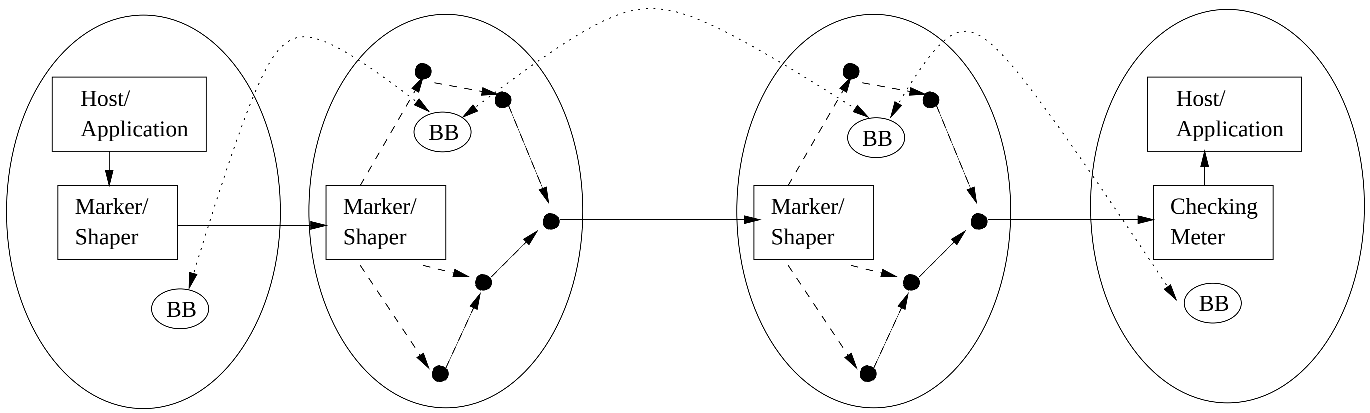

Figure 1: General Diff-Serv Architecture

The main trends

- **Actual Architectures**
  - SAPS (Service Allocation Profile Scheme)
  - IOS (Cisco's Internetwork Operating System)
  - TDSA (Two-bit Differentiated Services Architecture)
  - SRP (Scalable Resource Reservation Protocol)
  - USB (User-Share Differentiation)
- **Proposed Services**
  - Premium Service (Expedited Forwarding)
  - Assured Service (Drop Preference)

Motivation for our Work

- **Quantify Service for High and Low Priority**
  - Efficient Throughput (Assured Service)
  - Delay (Premium Service)
  - Denial of Service
- **Useful for:**
  - Tariffing
  - Dimensioning
  - Applications (Layered Coding)
- **Approach:** Analysis, verified by Simulation and Experimentation

A Model for the Assured Service

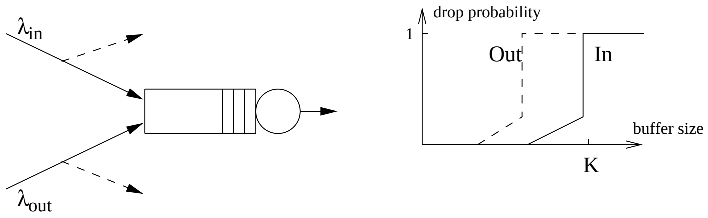

- Profile Meters tag Packets
- We first drop Low Priority Packets in case of congestion
- Arbitrary Drop Probability $\alpha^H(n)$: FIFO, RED, etc.
- We assume aggregated traffic $=>$ Poisson Process
- Simulation of LRD Traffic
- Three Variants of RIO

Comparison of Taildrop, RIO, and Threshold Dropping

- **TAIL** No specific buffer management scheme: $\alpha^H(n) = \alpha^L(n) = 1$, $n = 0..K-1$.
- **RIO** Accept everybody until the buffer content is equal to $K/2$. Then drop with a probability that increases linearly to 10% for tagged packets, and to 90% for non-tagged packets. Specifically,

$$\begin{cases} \alpha^T(n) = 1 \\ \alpha^{NT}(n) = 1 \end{cases} \quad n \leq K/2, \quad \begin{cases} \alpha^T(n) = 1 - 0.1 \times (2n-K)/K \\ \alpha^{NT}(n) = 1 - 0.9 \times (2n-K)/K \end{cases} \quad K/2 \leq n \leq K-1 \tag{1}$$

- **THRESH** Accept everybody until the buffer contents reaches $K/2$, then drop all non-tagged packets but accept all tagged packets. Formally,

$$\alpha^T(n) = 1, \quad 0 \leq n \leq K-1, \quad \begin{cases} \alpha^{NT}(n) = 1 & n \leq K/2 \\ \alpha^{NT}(n) = 0 & n > K/2 \end{cases} \tag{2}$$

A Model for the Assured Service (cont)

- **Throughput (tagged/non-tagged packets)**

$$\lambda^H_{\text{eff}} = \lambda p \sum_{n=0}^{K-1} \alpha^H(n)\pi(n), \quad \lambda^L_{\text{eff}} = \lambda \sum_{n=0}^{K-1} \alpha^L(n)\pi(n). \tag{3}$$

- **Probability of Loss**

$$\pi^H_{\text{rej}} = 1 - \sum_{n=0}^{K} \alpha^H(n)\pi(n), \quad \pi^L_{\text{rej}} = 1 - \sum_{n=0}^{K} \alpha^L(n)\pi(n). \tag{4}$$

The Impact of RIO

- **Efficient Throughput**

$$\lambda^T_{\text{eff}} = \phi\,\frac{1}{\mu} + \frac{1}{\rho}\frac{p}{\mu\alpha(K-1)}\left(\frac{\alpha^T(K-2)}{\alpha(K-2)} - \frac{\alpha^T(K-1)}{\alpha(K-1)}\right) + O(\frac{1}{\rho^2}) \quad (5)$$

$$\lambda^{NT}_{\text{eff}} = (1-\phi)\,\frac{1}{\mu} + \frac{1}{\rho}\frac{1-p}{\mu\alpha(K-1)}\left(\frac{\alpha^{NT}(K-2)}{\alpha(K-2)} - \frac{\alpha^{NT}(K-1)}{\alpha(K-1)}\right) + O(\frac{1}{\rho^2}) \quad (6)$$

$$\lambda_{\text{eff}} = \frac{1}{\mu} + O(\frac{1}{\rho^K}) \quad (7)$$

with

$$\phi = \frac{p\alpha^T(K-1)}{p\alpha^T(K-1) + \overline{p}\alpha^{NT}(K-1)}. \quad (8)$$

The Impact of RIO cont.

■ Observations

- Equation (8) shows that ϕ depends only on the probability of being accepted in the *last* buffer position, but *not* on the general shape of the drop function α
- The effective throughput for tagged and non-tagged packets converges slowly as ρ increases to the steady state values ϕ/μ and $(1-\phi)/\mu$
- To achieve a $(\phi, 1-\phi)$ load-independent sharing of the bandwidth chose:

$$\frac{p\alpha^T(n)}{p\alpha^T(n) + \overline{p}\alpha^{NT}(n)} = \phi \, , \quad i.e. \quad \alpha^{NT}(n) = \frac{1-\phi}{\phi} \frac{p}{1-p} \alpha^T(n) \quad (9)$$

Comparison of Tail, RIO, Thresh (Assured Service)

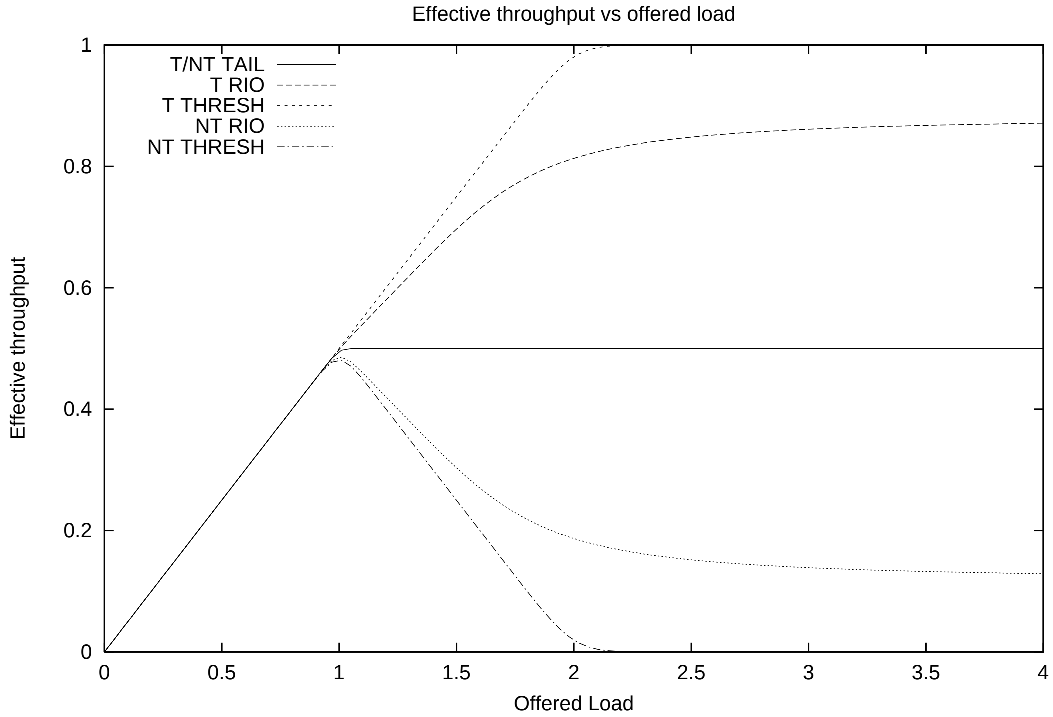

Figure 2: Effective throughput vs offered load, $K = 100$, $p = 0.5$ (T stands for tagged packets, NT for non-tagged packets)

Comparison of Varius Types of Traffic

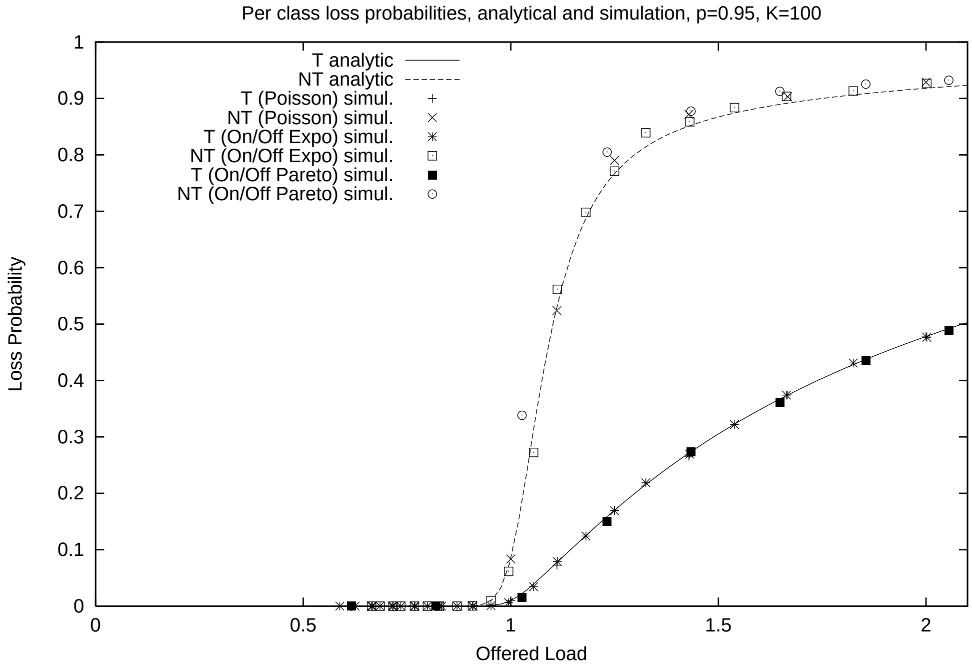

Figure 3: Comparing analysis and simulation: loss probability vs. load, buffer size $K = 100$, fraction of tagged packets $p = 0.95$ (T stands for tagged packets, NT for non-tagged packets)

Comparison between Discrimination on total and per-class Populations (Asured Service)

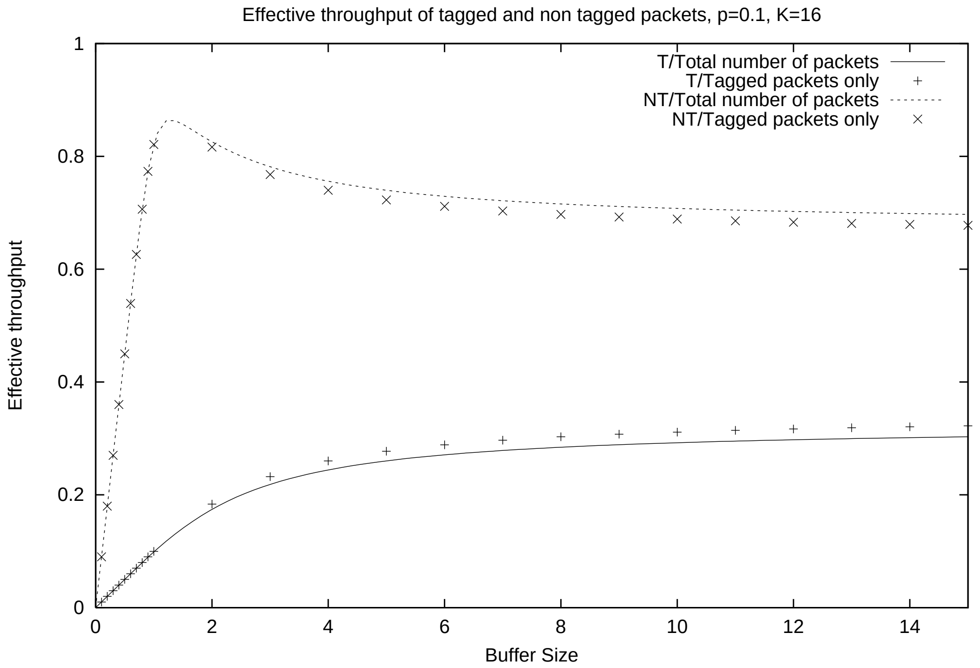

Figure 4: Effective throughput for tagged packets with RIO and drop probability depending on total number in queue (lines) and number of tagged packets in queue (points)(T stands for tagged packets, NT for non-tagged packets)

Simulation Results - Assured Service

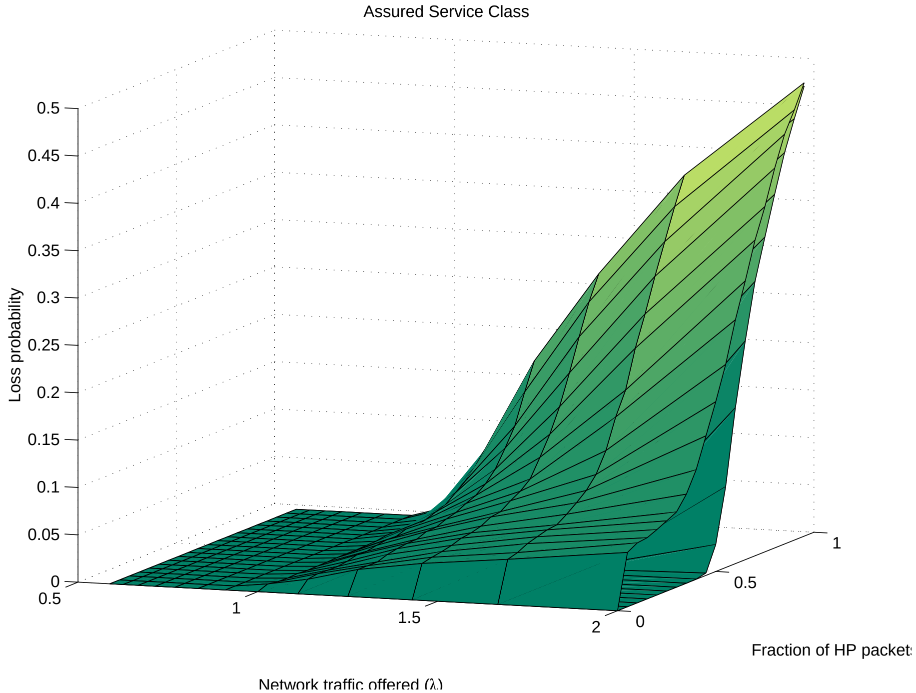

Figure 5: Loss probability with different parameter sets: tagged packets. The lower surface corresponds to the parameter set proposed by Clark, the upper surface corresponds to the other parameter set we used.

Simulation Results - Best Effort Service

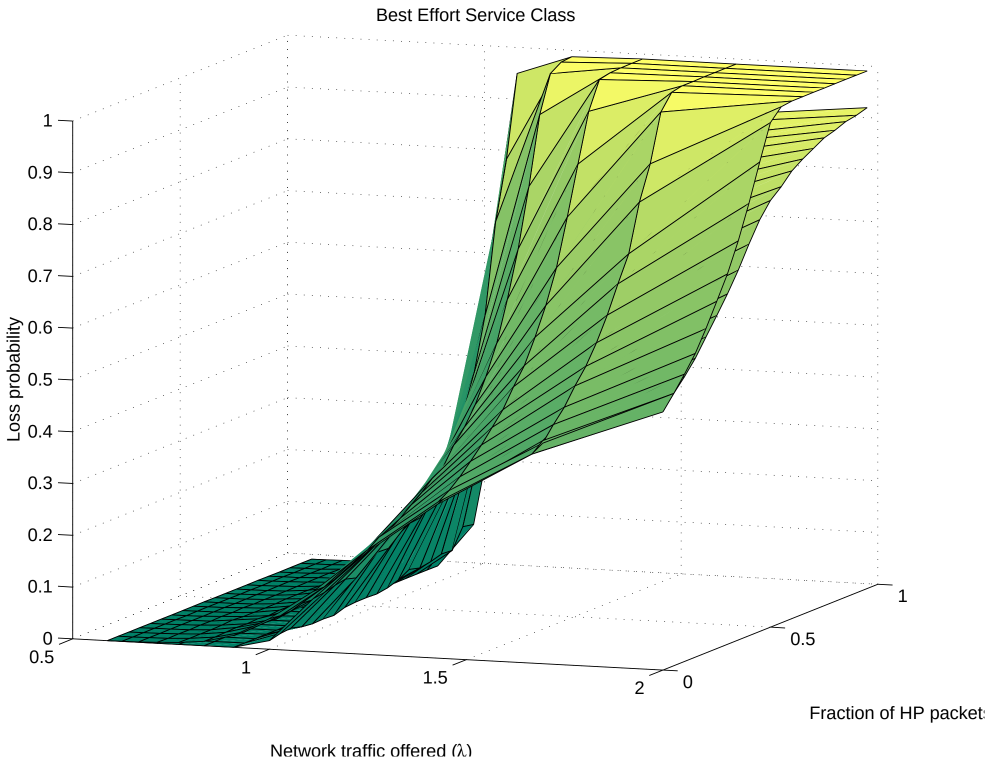

A Model for the Premium Service

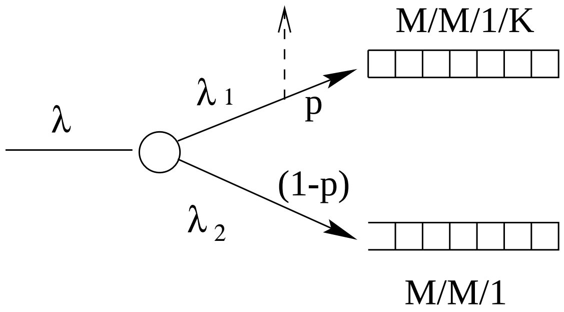

- **Offers a Premium Service in Terms of Delay**
- **With Admission Control it offers a CBR like Service**
- **Our model offers tight upper bounds for the sufferd Queueing Delay**

The Premium Service (cont.)

- **Low Priority Packets are transmitted only when the High Priority Queue is empty**
- **For the High Priority Queue we get the "classic" Mean Response Time:**

$$ER_2 \leq ER_2^{\infty} := \frac{1}{\mu} \frac{1}{1-\rho_1} \frac{1}{1-\rho} . \tag{10}$$

- **We get an Upper Bound for the Mean Waiting Time in the Low Priority Queue with the Vacation Time of the High Priority Queue**

Comparison of Simulation and Analytical Results (Premium Service)

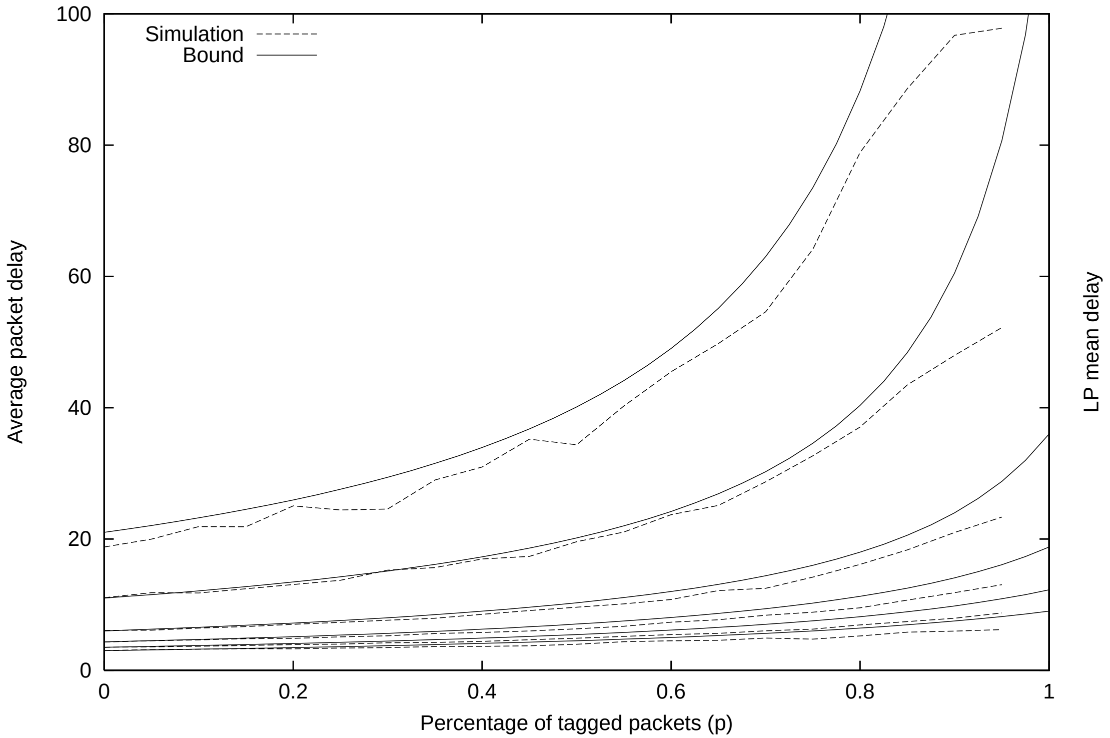

Simulation Results - Delay with Premium Service

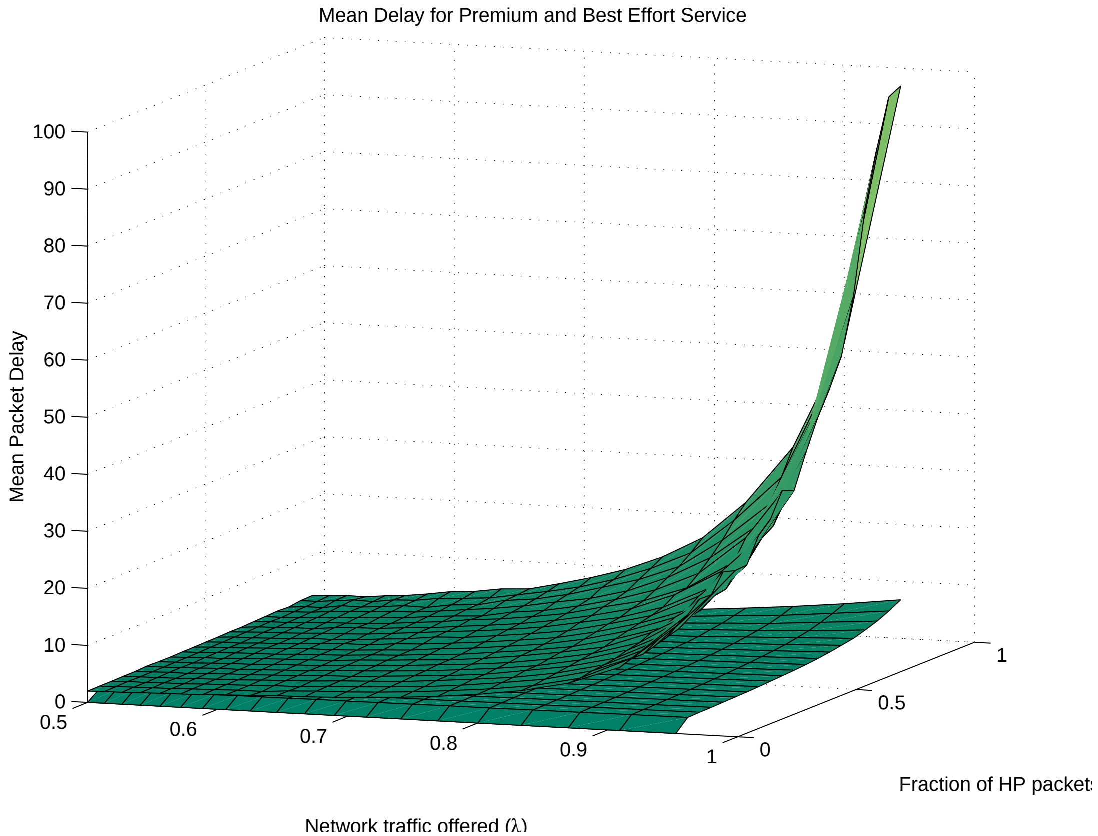

Results with different Traffic Types (Premium Service)

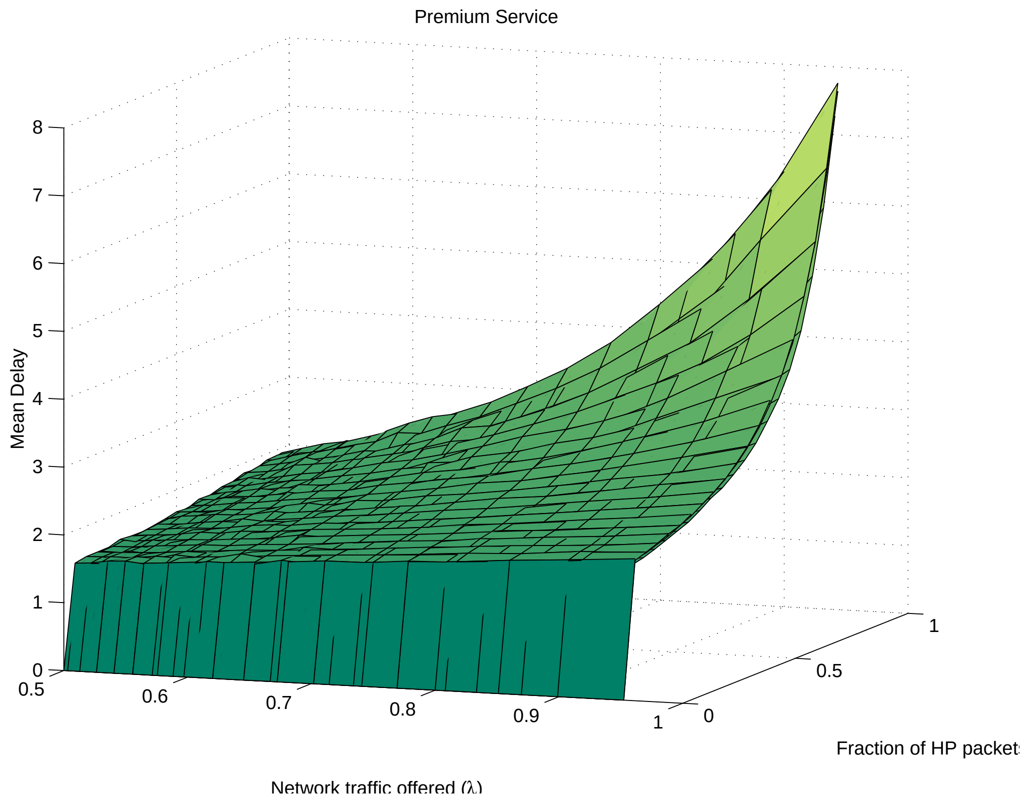

Results with different Traffic Types (Best Effort Service)

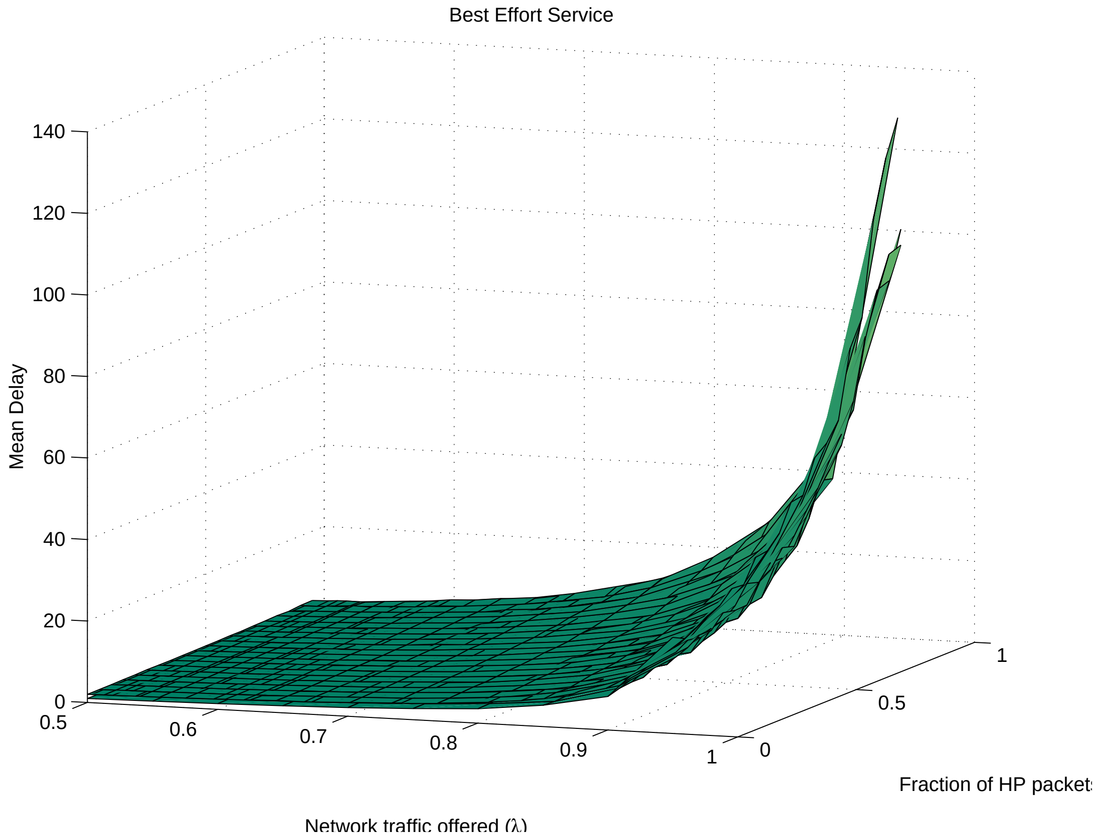

Results

- **Dimensioning:** Tradeoff between Delay and Loss Probability
- **Delay Difference between two Service Classes** Problems with Layered Applications
- **Tariffing:** Jitter between two Service Classes (Premium Service)

Difference in Loss Probability between two Service Classes (Assured Service)

- **RIO:** The effective throughput does *not* depend on the general shape of the RIO drop function, but on the probability of being accepted in the last buffer position.

Future Work

- **Develop a model in which we could quantify the impact of service discrimination on a** *per connection* **basis, as opposed to a** *per aggregate-of-connections* **basis as is done here**
- **Extend the models for a multi-class Premium and Assured Service**
- **Develop a pricing scheme where we use the models described in this paper to calculate the difference of service we get in each service class**

Experimental Results - Jitter

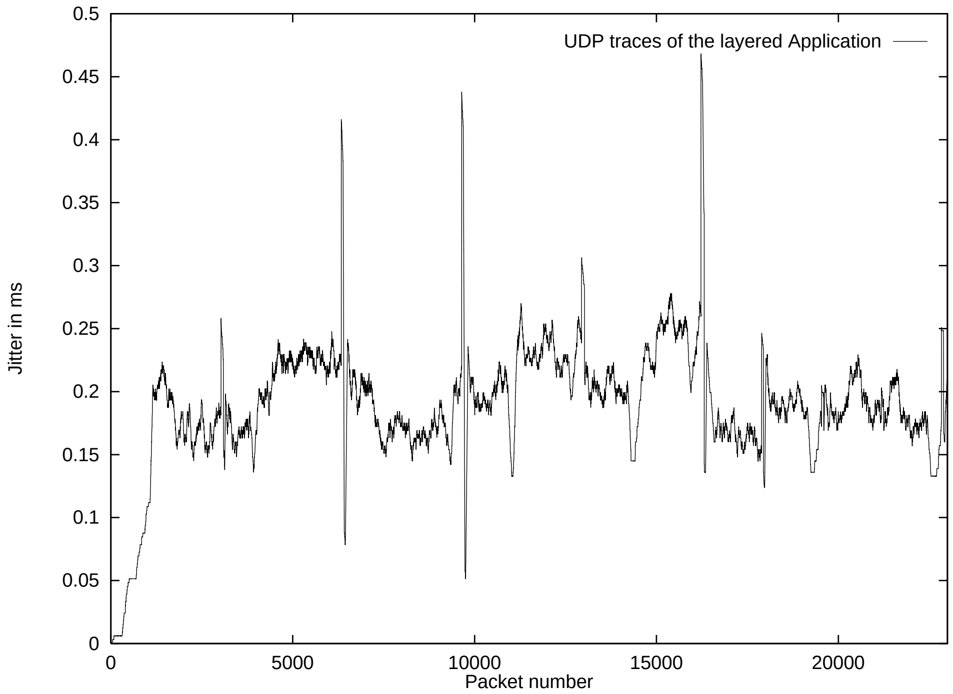

Figure 6: Evolutions of the delay difference of two simultaneously send packets. One using the Premium Service, the other best effort

Discrimination for the Premium Service

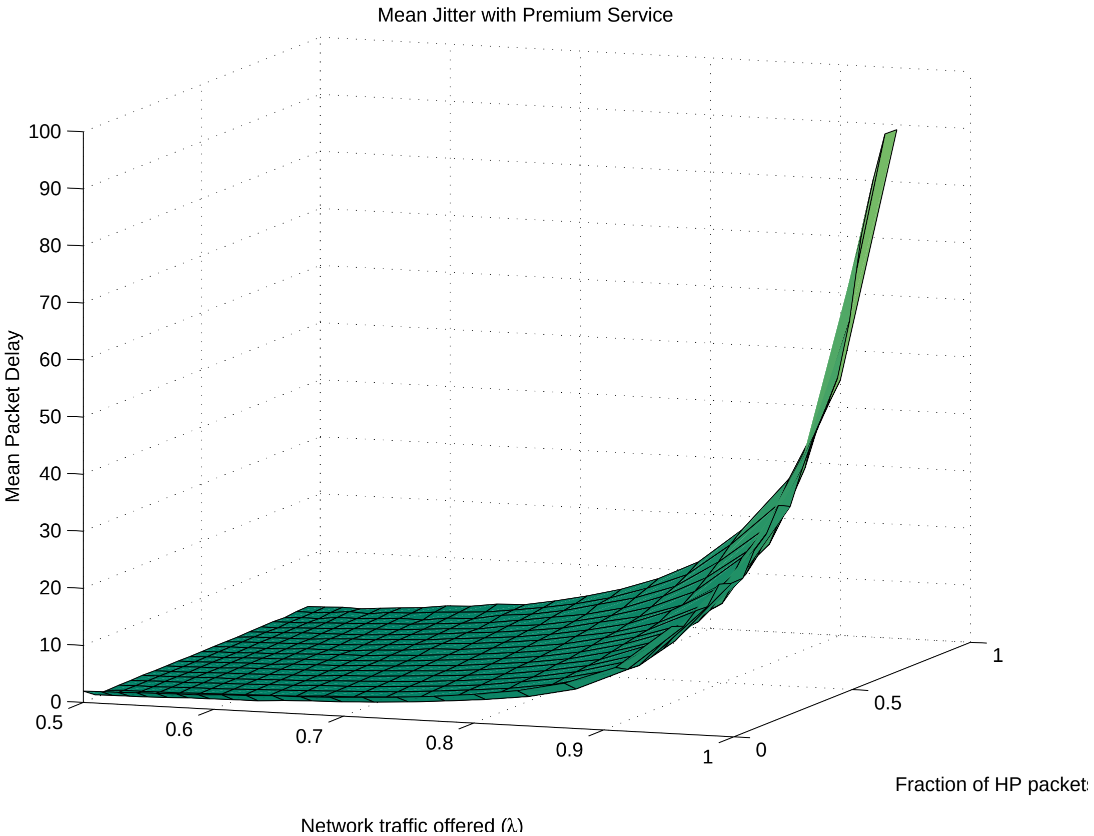

Discrimination for the Assured Service

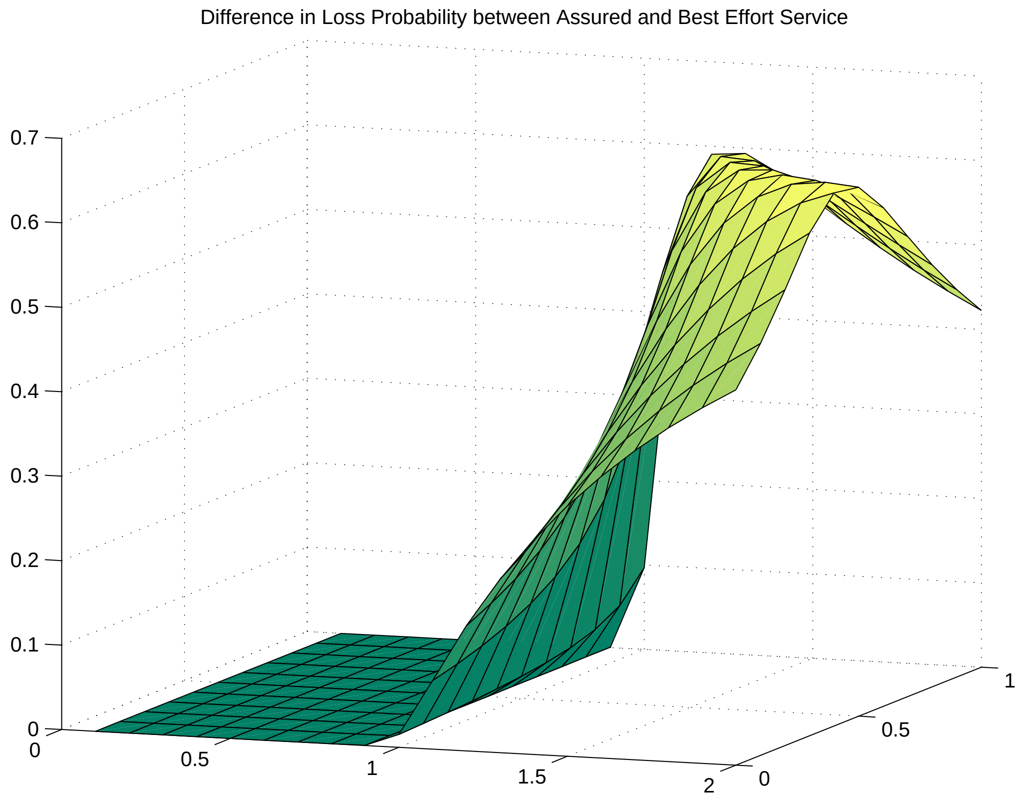

Tradeoff between delay and loss Probability for $\rho = 0.7$

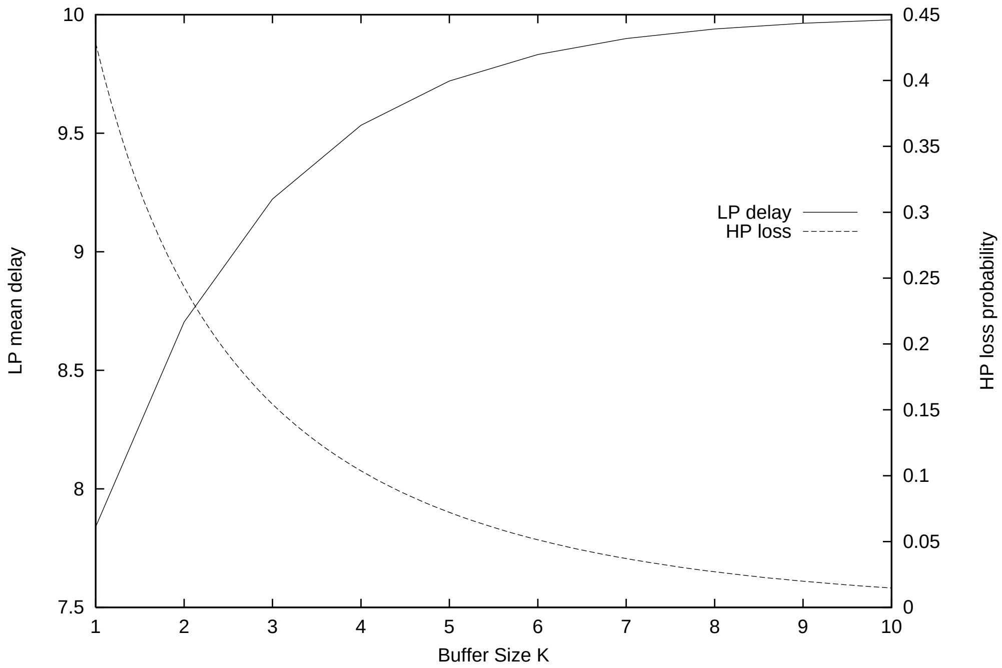

Figure 7: Tradeoff between best-effort delay and Premium Service loss probability as a function of the Premium Service buffer size. For the case of $\rho = 0.7$