

Proportional Differentiated Services: Delay Differentiation and Packet Scheduling

Constantinos Dovrolis
University of Delaware
dovrolis@cis.udel.edu

Dimitrios Stiliadis
Bell Laboratories
stiliadi@dnrc.bell-labs.com

Parameswaran Ramanathan
University of Wisconsin
parmesh@ece.wisc.edu

ABSTRACT

The proportional differentiation model provides the network operator with the ‘tuning knobs’ for adjusting the per-hop QoS ratios between classes, independent of the class loads. This paper applies the proportional model in the differentiation of queueing delays, and investigates appropriate packet scheduling mechanisms. The second part of this work [2] considers the corresponding loss rate differentiation and buffer management problem.

Starting from the Proportional Delay Differentiation (PDD) model, we derive the average queueing delay in each class, show the dynamics of the class delays under the PDD constraints, and state the conditions in which the PDD model is feasible. The feasibility model of the model can be determined from the average delays that result with the Strict Priorities scheduler. We then focus on scheduling mechanisms that can implement the PDD model, when it is feasible to do so. The Proportional Average Delay (PAD) scheduler meets the PDD constraints, when they are feasible, but it exhibits a pathological behavior in short timescales. The Waiting Time Priority (WTP) scheduler [3], on the other hand, approximates the PDD model closely, even in the short timescales of a few packet departures, but only in heavy load conditions. PAD and WTP serve as motivation for the third scheduler, called Hybrid Proportional Delay (HPD). HPD approximates the PDD model closely, when the model is feasible, independent of the class load distribution. Also, HPD provides predictable delay differentiation even in short timescales.

Parts of this paper were first published at ACM Sigcomm 1999 [1]. The work of C.Dovrolis and P.Ramanathan reported in this paper was supported in part by the National Science Foundation under Grant No. MIP-9526761, and by a research grant from Cisco Systems Inc.

1 Introduction

Providing some kind of service differentiation in the Internet has been a problem of enormous commercial and research importance in at least the last ten years. The Differentiated Services (DiffServ) architecture [4] was proposed as a more scalable solution to this problem, compared to previous approaches such as the Integrated Services (IntServ) architecture [5]. The DiffServ vision is that stateless priority mechanisms at the network core can be combined with stateful mechanisms at the network edges, in order to construct versatile end-to-end services. Among the first DiffServ proposals was the Virtual Leased Line (VLL) service [6] and the Assured service [7]. Recent findings, however, have raised concerns about the effect of aggregation on the VLL service [8], and about the provisioning complexity of the Assured service [9, 10]. Several other models for scalable differentiated services have been proposed. To name a few, the *SCORE* architecture can provide per-flow service guarantees without per-flow state in the core routers [11, 12]. [13] proposes a core-stateless scheme in which resource management and admission control are performed only at egress routers. The Asymmetric Best Effort [14] provides two service classes: one for delay-sensitive applications and another for throughput-sensitive applications.

A fundamentally different approach in the DS framework is the *relative differentiated services*. In this approach, the network traffic is grouped into N classes of service which are ordered, such that *Class i is better (or at least no worse) than Class $(i - 1)$ for $1 < i \leq N$, in terms of local (per-hop) metrics for the queueing delays and packet losses* [15]. The eight *Class Selector PHBs*, standardized by the IETF [16], follow the relative service differentiation model. In this context, applications and users do not get an absolute service level assurance, such as an end-to-end delay bound or throughput, since there is no admission control and resource reservations. Instead, the network assures that higher classes will offer better QoS than lower classes, and so it is up to the applications and users to select the class that best meets their requirements, cost, and policy constraints [17]. From this point of view, the relative service differentiation model follows the architectural principle of *end-system adaptation*, since it provides the class selection as an additional dimension in the end-system adaptation space.

There are several ways in which a network can provide relative differentiated services. For example, the differentiation can be based on appropriate pricing (i.e., higher classes are more expensive) [18], or on capacity provisioning (i.e., higher classes get more bandwidth relative to their expected

loads) [19]. Such mechanisms, however, cannot always provide consistent class differentiation, because the relative QoS differentiation between classes varies with the class loads. Other mechanisms, such as strict class prioritization, provide consistent differentiation that does not depend on the load variations, but they do not allow the network operator to adjust the relative QoS between classes. Such tuning knobs are necessary in a practical setting, since the network operators must be able to adjust the QoS difference between classes based on pricing and policy objectives. Our basic premise, starting from these limitations of other relative differentiation models, is that a relative differentiated services architecture has to be:

- *Predictable*, in the sense that the differentiation should be consistent (i.e., higher classes are better, or at least no worse) independent of the class load variations, and
- *Controllable*, meaning that the network operators should be able to adjust the relative QoS difference between classes based on their criteria.

As a target for predictable and controllable relative differentiation, we consider the *proportional differentiation model* [15]. According to this model, the basic performance measures for packet forwarding locally at each hop are ratioed proportionally to certain *class differentiation parameters* that the network operator chooses. Even though there is no wide consensus on the most appropriate performance measures for packet forwarding, it is generally agreed that a better network service means lower queueing delays and lower likelihood of packet losses. In this paper, we apply the proportional model in the case of queueing delay differentiation. The sequel paper [2] focuses on loss differentiation and appropriate buffer management, considering also the case of coupled delay and loss differentiation.

In the first part of this paper, we propose and study the Proportional Delay Differentiation (PDD) model. The PDD model controls the ratios of the average queueing delays between classes, based on the specified Delay Differentiation Parameters (DDPs). Starting from the PDD model, we derive the average queueing delay in each class given a class load distribution. We then show certain dynamic properties for the class delays in PDD, and state the conditions under which the PDD model is feasible. In the second part of the paper, we study three schedulers for the PDD model. The Proportional Average Delays (PAD) scheduler appears to always meet the PDD model, when it is feasible to do so. PAD, however, exhibits unpredictable behavior in short timescales. The Waiting

Time Priorities (WTP) scheduler, on the other hand, approximates the PDD model closely even in short timescales, but only in heavy load conditions. A third scheduler, called Hybrid Proportional Delays (HPD), combines the operation of PAD and WTP. Our simulation study shows that, in heavy load conditions, above 90% or so, both PAD and WTP meet the PDD model, and so does HPD. HPD approximates the PDD model closely, when the model is feasible, independent of the class load distribution. Also, HPD provides predictable delay differentiation even in short timescales.

The structure of the paper is as follows. In Section 2, we study the PDD model and present its dynamic properties and feasibility conditions. Sections 3, 4, and 5 focus on the PAD, WTP, and HPD schedulers, respectively. In Section 6, we review other scheduling disciplines in the area of relative and proportional delay differentiation. We conclude in Section 7.

2 Proportional Delay Differentiation (PDD) model

We consider the following queueing model of a packet multiplexer (or ‘link’). A *lossless, work-conserving*, and *non-preemptive* link with capacity C (bytes per second) services N First-Come First-Served (FCFS) queues, one for each traffic class¹. The order in which packets are selected for service is determined by the *packet scheduler* \mathcal{S} . The lossless property requires that the utilization of \mathcal{S} is less than 100%, and that there are enough buffers for packets that need to be queued. The work-conserving property means that the scheduler does not idle when there are waiting packets. The non-preemptive property means that a packet transmission is always carried out to completion.

The interarrivals between packets of class i are IID random variables, generated by a renewal arrival process for that class. The packet sizes in class i are also IID random variables, generated by a renewal packet size process for that class. The arrival and packet size processes of different classes are independent, and, in general, different. Some additional assumptions are made later, where needed. The *aggregate traffic stream* that arrives in the queueing system is determined by the superposition of the N packet streams that arrive in each class. Let λ_i be the average input rate in class i (packets per second), and $\lambda = \sum_{i=1}^N \lambda_i$ the average rate of the aggregate traffic stream. Let \bar{L}_i be the *average size* (bytes) of class i packets, and $\bar{L} = \sum_{i=1}^N (\lambda_i \bar{L}_i) / \lambda$ the average size among all packets. The link *utilization* is $u = \bar{L}\lambda / C < 1$.

¹We use the terms ‘class’ and ‘queue’ interchangeably.

The *Proportional Delay Differentiation (PDD)* model aims to control the *ratios* of the average class queueing delays based on the *Delay Differentiation Parameters (DDPs)* $\{\delta_i, i = 1, \dots, N\}$. Specifically, let \bar{d}_i be the *average queueing delay*, or simply average delay, of the class i packets. The PDD model requires that the ratio of average delays between two classes i and j is fixed to the ratio of the corresponding DDPs

$$\frac{\bar{d}_i}{\bar{d}_j} = \frac{\delta_i}{\delta_j} \quad 1 \leq i, j \leq N \quad (1)$$

Our convention is that *higher classes provide better service, i.e., lower queueing delays*, and so $\delta_1 > \delta_2 > \dots > \delta_N > 0$. In the following, we choose Class-1 as the ‘reference class’ and set $\delta_1=1$. Then, the PDD model requires that the average delay of each class i is a certain fraction δ_i of the average delay of Class-1,

$$\bar{d}_i = \delta_i \bar{d}_1 \quad i = 2, \dots, N \quad (2)$$

independent of the aggregate load, or the class load distribution.

2.1 Per-class average delays in the PDD model

In given load conditions, the $N-1$ ratios of the PDD model specify uniquely the average delays of the N classes. The key additional relation in the mapping from delay ratios to class delays is the *conservation law* [3, 20], which constrains the average class delays in any work-conserving scheduler. Specifically, the conservation law requires that for any work-conserving scheduler \mathcal{S} that causes an average delay \bar{d}_i in class i , we must have that

$$\sum_{i=1}^N \lambda_i \bar{L}_i \bar{d}_i = \lambda \bar{L} \bar{d}_{ag} = \bar{q}_{ag} \quad (3)$$

where \bar{d}_{ag} and \bar{q}_{ag} is the average delay (in seconds) and the average backlog (in bytes), respectively, in a FCFS scheduler that has the same capacity and services the same aggregate traffic stream as \mathcal{S} . The conservation law (in the previous form) holds under arbitrary distributions for the packet interarrivals and packet sizes, as long as the first moment of these distributions (λ_i and \bar{L}_i) exist. The conservation law implies that even though a scheduler \mathcal{S} can affect the relative magnitude of the class delays, making the delay of one class lower than the delay of another class, this is a ‘zero-sum’ game because the weighted sum of (3) has to be equal to the average backlog \bar{q}_{ag} of the aggregate traffic stream. The average backlog \bar{q}_{ag} is a significant *invariant* in this balance, as it does not depend

on the scheduling discipline \mathcal{S} , but only on the aggregate traffic stream and on the capacity C .

Suppose now that the scheduler \mathcal{S} satisfies the PDD model of (1). With the additional constraint (3) that the conservation law imposes, we can show that the average delay in class i is

$$\bar{d}_i = \frac{\delta_i \bar{q}_{ag}}{\sum_{n=1}^N \delta_n \lambda_n \bar{L}_n} \quad i = 1, \dots, N \quad (4)$$

Consequently, even though the PDD model consists of $N-1$ *relative* constraints, the average delay in each class is determined in an absolute sense when the PDD model is applied to a certain traffic stream with class loads $\{\lambda_i\}$, average packet sizes $\{\bar{L}_i\}$, and average aggregate backlog \bar{q}_{ag} .

2.2 Delay dynamics in the PDD model

In this section, we further assume that the N classes have the same packet size distribution. We can then set $\bar{L}_i = \bar{L} = 1$, and normalize the backlog measures to *average packet* units. The average delay in each class is then simplified to

$$\bar{d}_i = \frac{\delta_i \bar{q}_{ag}}{\sum_{n=1}^N \delta_n \lambda_n} \quad i = 1, \dots, N \quad (5)$$

Additionally, we assume that the N classes have the same interarrival distribution, but with a possibly different average interarrival $1/\lambda_i$.

Based on Equation (5), we can now investigate the variations in the average class delays under the constraints of the PDD model, as the aggregate load, the class load distribution, or the DDPs vary. We refer to the following properties as the *delay dynamics in the PDD model*. The proofs of these properties are given in Appendix I.

Property 1: Increasing the input rate of a class, increases (in the wide sense²) the average delay of all classes.

This property shows that there is no isolation between classes in the PDD model, i.e., when the delay of a class increases, due to additional load in that class, the delays of all classes will also increase.

²'Increasing a function $f(x)$ in the wide sense' means that the corresponding function is non-decreasing, i.e., $df(x)/dx \geq 0$.

Property 2: Increasing the rate of a higher class causes a larger increase in the average class delays than increasing the rate of a lower class.

In the case of two classes, for instance, suppose that the class rates become either $\lambda'_1 = \lambda_1 + \epsilon$ and $\lambda'_2 = \lambda_2$, or $\lambda''_1 = \lambda_1$ and $\lambda''_2 = \lambda_2 + \epsilon$ ($\epsilon > 0$). Even though the conservation law requires that the weighted average of the class delays is the same in both cases ($\lambda'_1 \bar{d}'_1 + \lambda'_2 \bar{d}'_2 = \lambda''_1 \bar{d}''_1 + \lambda''_2 \bar{d}''_2$), the class average delays in the second case are larger, i.e., $\bar{d}''_1 > \bar{d}'_1$ and $\bar{d}''_2 > \bar{d}'_2$, because of the PDD constraints. This property shows that higher classes cost more, in terms of queueing delay, than lower classes.

Property 3: Decreasing the delay differentiation parameter of a class increases (in the wide sense) the average delay of all other classes, and decreases (in the wide sense) the average delay of that class.

This property implies that if the delay of a class is reduced, by lowering its DDP, then the delay of all other classes will increase.

The following two properties are important in the Dynamic Class Selection (DCS) framework [17], in which users with a certain maximum end-to-end delay constraint search for the minimum acceptable class. The first property states that when one or more users move to a higher class, the delay of all classes increases. The class delays decrease, on the other hand, when one or more users move to a lower class. The second property shows that when a user switches from one class to another, the user still observes a consistent class ordering, i.e., the higher class provides a lower delay.

Suppose that the class load distribution changes from $\{\lambda_n\}$ to $\{\lambda'_n\}$, with $\lambda'_i = \lambda_i - \epsilon$, $\lambda'_j = \lambda_j + \epsilon$, and $\lambda'_k = \lambda_k$ for all $k \neq i, j$ ($\epsilon > 0$). Let \bar{d}'_n be the average delay in class n when the class load distribution is $\{\lambda'_n\}$.

Property 4: If $i > j$ then $\bar{d}'_i \leq \bar{d}'_j$ for all $n = 1 \dots N$. Similarly, if $i < j$ then $\bar{d}'_i \geq \bar{d}'_j$.

Property 5: If $i > j$ then $\bar{d}'_j \geq \bar{d}'_i$. Similarly, if $i < j$ then $\bar{d}'_j \leq \bar{d}'_i$.

2.3 Feasibility of the PDD model

We have assumed so far that the PDD model is *feasible*, i.e., that there exists a work-conserving scheduler that can meet the constraints of (2). However, this is not always possible. Given the load distribution and the average backlog of the aggregate traffic, the PDD model not only specifies the $N - 1$ delay ratios, but also the N average delays of Equation (4). There may not exist a work-conserving scheduler that can set the average delay of each class as in Equation (4) however. Intuitively, the reason is that the average delay of a class (or of any set of classes) has a minimum value due to the inherent load in that class (or set of classes). This minimum average delay would result if that class (or set of classes) was given strict priority over the rest of the traffic.

Formally, given the input rates $\{\lambda_i\}$, the average class packet sizes $\{\bar{L}_i\}$, and the average backlog \bar{q}_{ag} of the aggregate traffic stream, we say that *a set of DDPs $\{\delta_i, i = 2, \dots, N\}$ is feasible when there exists a work-conserving scheduler that can set the average delay of each class as in (4)*. So, the set of DDPs $\{\delta_i, i = 2, \dots, N\}$ is feasible when the set of class delays $\{\bar{d}_i = \delta_i \bar{q}_{ag} / (\sum_{n=1}^N \lambda_n \delta_n \bar{L}_n), i = 1, \dots, N\}$ is feasible.

The necessary and sufficient conditions for the feasibility of a set of N average class delays, given the N class loads, were derived by Coffman and Mitrani in [21]. Under general assumptions³, a set of N average delays $\{\bar{d}_i, i = 1, \dots, N\}$ is feasible if and only if the following $2^N - 2$ inequalities hold,

$$\sum_{i \in \phi} \lambda_i \bar{d}_i \bar{L}_i \geq \bar{d}_\phi^{SP} \sum_{i \in \phi} \lambda_i \bar{L}_i = \bar{q}_\phi^{SP} \quad \text{for all } \phi \in \Phi \quad (6)$$

where Φ is the set of $2^N - 2$ nonempty proper subsets of the set $\{1, 2, \dots, N\}$. \bar{q}_ϕ^{SP} and \bar{d}_ϕ^{SP} are the average backlog and the average delay, respectively, of the traffic in the set ϕ , *if ϕ was given strict priority over all other classes*. The Coffman-Mitrani inequalities state that any set of classes ϕ has a lower bound on its average backlog and this minimum value results when the traffic in ϕ is serviced with the highest priority in a Strict Priority (SP) scheduler.

In the remainder of this section, we make the additional assumption that all classes have the same packet size distribution, and we set the average packet size to $\bar{L} = \bar{L}_i = 1$. With this assumption, Regnier [22] extended the Coffman-Mitrani results showing that the necessary and sufficient feasibility

³Most of [21] assumes Poisson arrivals. The particular result that we include here holds, as [21] also states, for general interarrival and packet size distributions that are different between classes.

conditions of (6) can be reduced from 2^N-2 to the following $N-1$ inequalities

$$\sum_{i=k}^N \lambda_i \bar{d}_i \geq \bar{d}_{k,N}^{SP} \sum_{i=k}^N \lambda_i = \bar{q}_{k,N}^{SP} \quad k = 2, \dots, N \quad (7)$$

where $\bar{q}_{k,N}^{SP}$ and $\bar{d}_{k,N}^{SP}$ are the average backlog \bar{q}_{ϕ}^{SP} and average delay \bar{d}_{ϕ}^{SP} , respectively, for $\phi = \{k, \dots, N\}$. Regnier's inequalities impose a lower bound on the average backlog of the $N-1$ subsets of the k highest classes ($k = 1, \dots, N-1$).

Let \bar{d}_i^{SP} be the average delay of Class- i in the SP scheduler which services class m with higher priority than class n for $m > n$ (i.e., Class- N is given highest priority and Class-1 is given lowest priority). With this notation, we have that $\bar{d}_{k,N}^{SP} \sum_{i=k}^N \lambda_i = \sum_{i=k}^N \lambda_i \bar{d}_i^{SP}$, and so the Regnier inequalities become

$$\sum_{i=k}^N \lambda_i \bar{d}_i \geq \sum_{i=k}^N \lambda_i \bar{d}_i^{SP} \quad k = 2, \dots, N \quad (8)$$

From (8) and (4), we can now show that the necessary and sufficient condition for the feasibility of the PDD model when $N = 2$ is

$$\bar{d}_2 = \frac{\delta_2 \bar{q}_{ag}}{\lambda_1 + \delta_2 \lambda_2} \geq \bar{d}_2^{SP} \quad (9)$$

Since $\bar{q}_{ag} = \lambda_1 \bar{d}_1^{SP} + \lambda_2 \bar{d}_2^{SP}$, it is easy to show that the feasibility condition becomes

$$\delta_2 \geq \frac{\bar{d}_2^{SP}}{\bar{d}_1^{SP}} \quad \text{or} \quad \frac{1}{\delta_2} \leq \frac{\bar{d}_1^{SP}}{\bar{d}_2^{SP}} \quad (10)$$

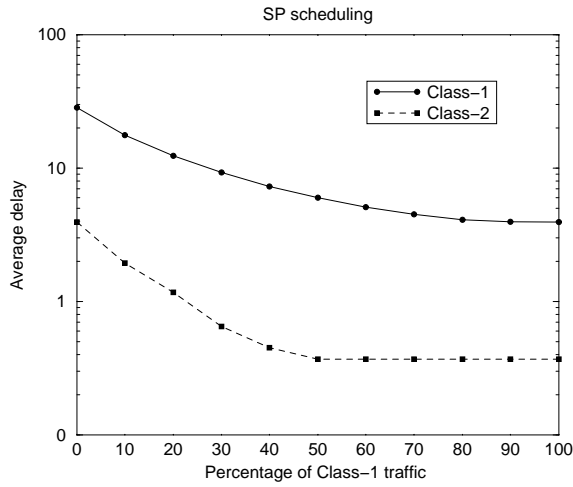
So, a given delay ratio $\delta_2 < 1$ between Class-2 and Class-1 is feasible if and only if the corresponding average delay ratio in the SP scheduler is not larger than δ_2 . Note that the average delays of the two classes in SP are a function of the input rate λ , of the load distribution (λ_1, λ_2) , and of the aggregate average backlog \bar{q}_{ag} , and so the feasibility of the PDD model depends on the same factors.

When $N > 2$, the feasibility conditions for the PDD model are the following $N-1$ inequalities

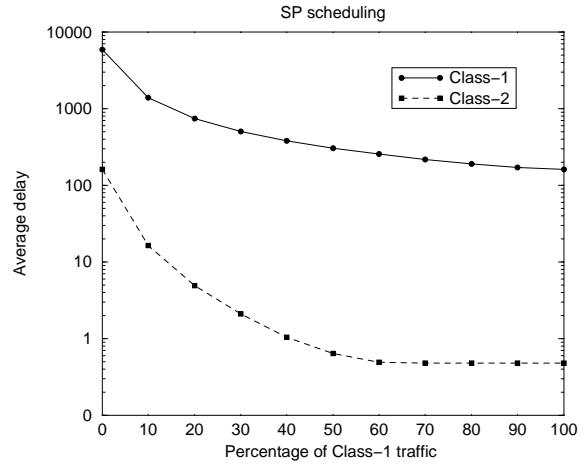
$$\sum_{i=k}^N \lambda_i \delta_i \geq \frac{S}{\bar{q}_{ag}} \sum_{i=k}^N \lambda_i \bar{d}_i^{SP} \quad k = 2, \dots, N \quad (11)$$

where $S = \sum_{i=1}^N \lambda_i \delta_i$ and $\bar{q}_{ag} = \sum_{i=1}^N \lambda_i \bar{d}_i^{SP}$.

To illustrate the feasibility conditions graphically, Figure 1 shows the per-class average delays with SP for the case of two classes in two utilization points ($u=75\%$ and $u=95\%$). Also, Figure 2 shows the

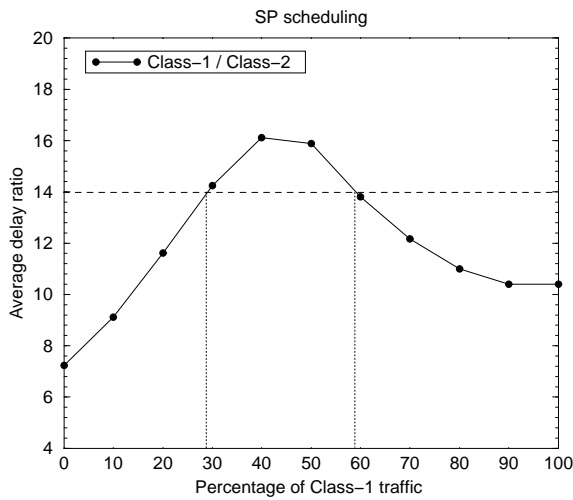


(a) Average delays, $u=75\%$

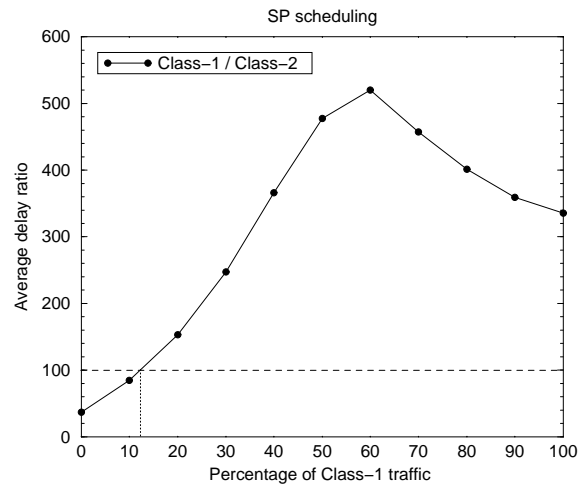


(b) Average delays, $u=95\%$

Figure 1: Average delays with SP as a function of the load distribution.



(a) Ratios of average delays, $u=75\%$



(b) Ratios of average delays, $u=95\%$

Figure 2: Average delay ratios with SP as a function of the load distribution.

average delay ratios $\bar{d}_1^{SP}/\bar{d}_2^{SP}$ for the same load conditions. These graphs resulted from a simulation of the SP scheduler⁴. When $u=75\%$ (Figure 2-a) and the load distribution is $(\lambda_1, \lambda_2)=(70,30)$, the DDP $1/\delta_2=8$ is feasible, while the DDP $1/\delta_2=16$ is infeasible. On the other hand, if $1/\delta_2=14$, the load distribution must be such that $0.3 < \lambda_1/\lambda < 0.6$ in order for the PDD model to be feasible. When $u=95\%$ (Figure 2-b), the DDP $1/\delta_2=32$ is always feasible, while the DDP $1/\delta_2=100$ is only feasible when at least 12% of the traffic belongs to Class-1.

For some queuing models, the average delays with the SP scheduler can be analytically computed. For the M/G/1 queueing model with $N = 2$ classes, for instance, we have that $\bar{d}_1^{SP}/\bar{d}_2^{SP} = 1/(1 - \lambda)$ [20]. This result, which is based on Poisson arrivals, shows that the feasibility range increases as the aggregate load increases, but it hides the dependency of the feasibility range on the class load distribution.

If the number of classes is $N > 2$, such a graphical interpretation of the feasibility conditions is not straightforward. One can still use (11) to examine numerically if the $N-1$ conditions are satisfied for the given DDPs, load distribution, and aggregate backlog. In order to examine the feasibility of the PDD model, we need to know the average delays \bar{d}_k^{SP} that would result with a SP scheduler. If the traffic does not conform well to a mathematically tractable queueing model, one needs to estimate or measure them experimentally. Most existing routers provide an SP scheduler, and so a network provider can measure the class average delays \bar{d}_k^{SP} deploying the SP scheduler for short time intervals in the actual workload. Alternatively, if the SP deployment is considered harmful due to its starvation effects, the delays \bar{d}_k^{SP} can be estimated emulating an SP scheduler with the actual workload of the link.

3 Proportional Average Delay (PAD) scheduling

Another way to interpret the PDD model is that the *normalized average delays*, defined as $\tilde{d}_i = \bar{d}_i/\delta_i$, must be equal in all classes, i.e.,

$$\tilde{d}_i = \frac{\bar{d}_i}{\delta_i} = \frac{\bar{d}_j}{\delta_j} = \tilde{d}_j \quad 1 \leq i, j \leq N \quad (12)$$

⁴The simulation model and parameters that we used in this paper is described in Appendix II.

A scheduler that aims to equalize the normalized average delays among all classes is described next. We refer to this algorithm as *Proportional Average Delay (PAD)* scheduling.

Let $B(t)$ be the set of backlogged classes at time t , $D_i(t)$ be the sequence of Class- i packets that departed before time t (in order of their departure), and d_i^m be the delay of the m 'th packet in $D_i(t)$. Assuming that there was at least one departure from class i before t , the normalized average delay of class i at t is

$$\tilde{d}_i(t) = \frac{1}{\delta_i} \frac{\sum_{m=1}^{|D_i(t)|} d_i^m}{|D_i(t)|} = \frac{1}{\delta_i} \frac{S_i}{P_i} \quad (13)$$

where $|D_i(t)|$ is the cardinality of $D(t)$, S_i is the sum of queueing delays of all packets in $D_i(t)$, and $P_i = |D_i(t)|$ is the number of packets in $D_i(t)$.

Suppose that a packet has to be selected for transmission at time t . *PAD chooses the backlogged class j with the maximum normalized average delay at t ,*

$$j = \arg \max_{i \in B(t)} \tilde{d}_i(t) \quad (14)$$

The packet at the head of queue j is transmitted, after a queueing delay $d_j^{P_j+1}$. The variables S_j and P_j are then updated as $S_j = S_j + d_j^{P_j+1}$ and $P_j = P_j + 1$, and the new normalized average delay $\tilde{d}_j(t)$ is computed from (13).

The selection of the maximum $\tilde{d}_i(t)$ term, as in (14), requires at most $N-1$ comparisons, which is a minor overhead for the small number of classes that we consider here. The main computational overhead of PAD is a division, for the calculation of S_j/P_j , after each packet departure⁵. With current commodity microprocessors ($\approx 500\text{MHz}$), a floating-point division takes less than 100ns (≈ 50 cycles). This delay would not be an issue for network interfaces of up to 1Gbps, which transmit a (minimum-size) 40-byte packet in about 320ns. For even faster network interfaces, one can avoid the division operation and compute $\tilde{d}_i(t)$ using an exponential running-average of the form $\tilde{d}_i(t) = (1 - 2^{-w})\tilde{d}_i(t) + 2^{-w}d_i^{P_i+1}/\delta_i$, where w is a positive integer.

The basic idea in PAD is that if a packet is serviced from class j with the maximum normalized average delay, the delay of that packet will not increase any more (since it is serviced), and thus the increase of S_j due to that packet will be minimized. So, servicing a packet from class j tends to reduce the difference of $\tilde{d}_j(t)$ from the normalized average delays of the other classes. In the

⁵Obviously, the factor $1/\delta_i$ can be precomputed and does not require a division. If it is a power of two, it does not require a multiplication either.

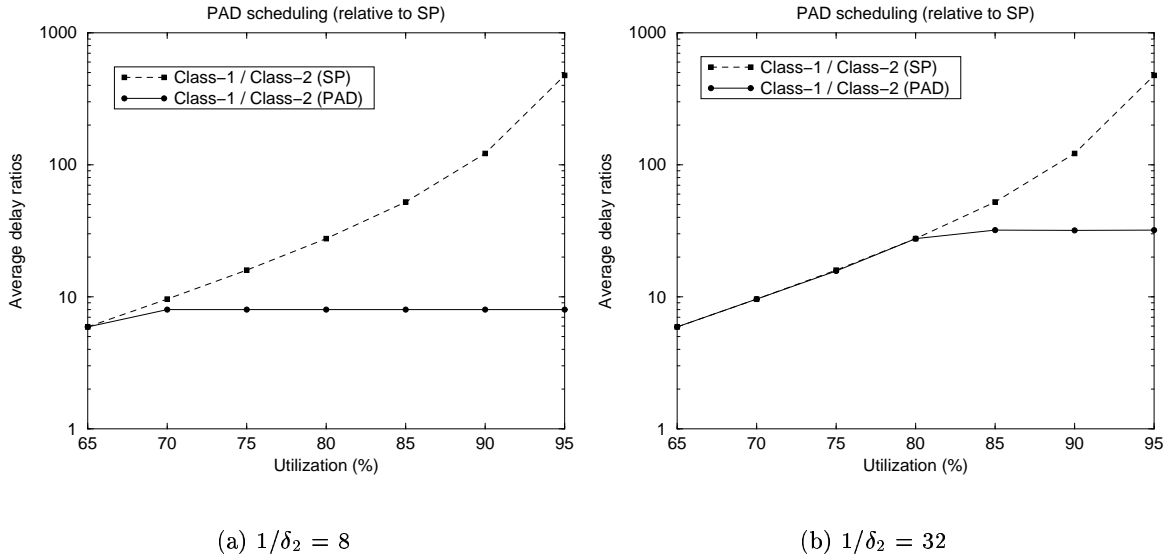
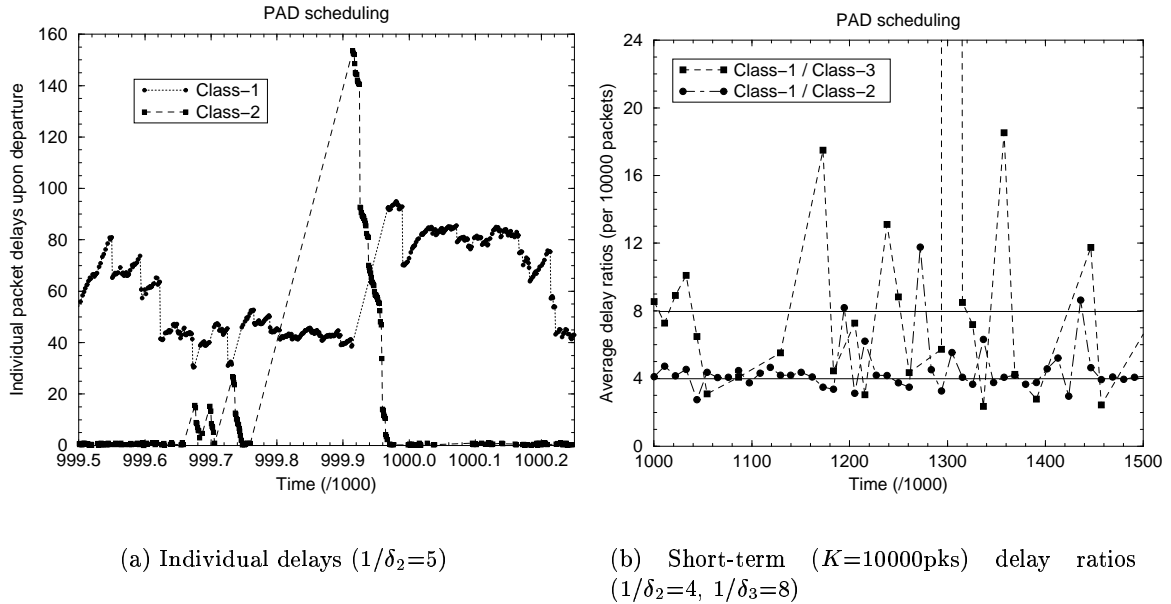


Figure 3: Average delay ratios with PAD and SP as a function of the utilization.

long-run, if the scheduler always minimizes the differences between the normalized average delays in this manner, we expect that the normalized average delays will be about the same.

Our simulation experiments have shown that PAD can achieve the PDD model constraints, when the given set of DDPs is feasible. To illustrate this point, Figure 3 shows typical simulation results for $N=2$ classes. The two graphs show the ratio \bar{d}_1/\bar{d}_2 of the average delays with PAD and SP for two DDP selections and for a uniform load distribution ($\lambda_1=\lambda_2$), as a function of the utilization u . Recall from Equation (10) that in the case of two classes, a given DDP $1/\delta_2$ is feasible if and only if it is lower than the ratio \bar{d}_1/\bar{d}_2 in SP. As shown in Figure 3-a, PAD achieves the specified delay ratio $1/\delta_2=8$ in all utilization points that this DDP is feasible. The resulting delay ratio is less than eight only when $u < 70\%$, but SP does not lead to a larger delay ratio either, meaning that this DDP is infeasible in that utilization range. Similarly, in Figure 3-b, PAD achieves the specified delay ratio $1/\delta_2=32$ when the utilization is higher than about 80%, which is also when this DDP is feasible.

PAD, however, exhibits a pathological behavior in short timescales. To illustrate this behavior, Figure 4-a shows the queuing delays of individual packets at their departure instant. The utilization in this simulation is $u=90\%$ and the load distribution to $(\lambda_1, \lambda_2) = (70, 30)$. Note that during most of the time, the Class-2 packets depart with minor queuing delays. Occasionally though, as it occurs shortly after time 999.900, some Class-2 packets depart with significantly larger queuing delays,



(a) Individual delays ($1/\delta_2=5$)

(b) Short-term ($K=10000$ pks) delay ratios ($1/\delta_2=4$, $1/\delta_3=8$)

Figure 4: Individual packet delays and short-term delay ratios with PAD.

even larger than the Class-1 delays in the corresponding time frame. Let us analyze in detail what happens in the time period 999.760-999.960. Just before 999.760, the Class-2 queue is almost empty and its queueing delays are close to zero. From 999.760 to 999.920 (phase-1), \tilde{d}_2 is less than \tilde{d}_1 , and there are no departures from Class-2. The arriving Class-2 packets during phase-1 accumulate large waiting times, up to the duration of this phase (160 time units), but the normalized average delay \tilde{d}_2 does not change since there are no departures from that class. At about 999.920, \tilde{d}_1 becomes less than \tilde{d}_2 , after servicing exclusively packets of that class during phase-1. From that point and until 999.960 (phase-2), all the backlogged Class-2 packets from phase-1 depart back-to-back without interventions from Class-1 packets. At about 999.960, the two normalized average delays become comparable, and the scheduler starts servicing packets again in a more uniform manner. This example shows that *because PAD is unaware of the waiting times of backlogged packets, it occasionally allows higher classes to experience much larger delays than their long-term average delays, or the queueing delays of lower classes.*

It is also instructive to examine the ratios of short-term average delays between classes with PAD, in relatively short timescales. Figure 4-b shows the ratios of average delays, measured in every $K=10000$ successive packet departures, for $N=3$ classes. The utilization is $u=90\%$, and the

load distribution is $(\lambda_1, \lambda_2, \lambda_3) = (50, 30, 20)$. Note that PAD does a poor job in meeting the PDD constraints in this timescale. This should be expected, since PAD attempts to equalize the *long-term* normalized average delays and not the normalized average delays in the *last K departures*. Another important point in this graph is that there are several time periods in which the ratio \bar{d}_1/\bar{d}_2 is larger than the ratio \bar{d}_1/\bar{d}_3 , meaning that Class-3 experiences higher delays than Class-2. These are incidents of unpredictable delay differentiation. *A good scheduler for the PDD model should not only achieve or closely approximate the constraints of (2), but it should also provide predictable delay differentiation in short timescales.* A scheduler that emphasizes on this predictability issue is studied next.

4 Waiting Time Priority (WTP) scheduling

The Waiting Time Priority scheduling algorithm was first studied by L.Kleinrock in 1964 [3] under the name *Time-Dependent Priorities*. A packet is assigned a priority that increases proportionally to the packet's waiting time. Higher classes have larger priority-increase factors. The packet with the highest priority is serviced first (in non-preemptive order). In the following, we describe WTP in a different way to highlight its similarities and differences with PAD.

Suppose that class i is backlogged at time t , and that $w_i(t)$ is the *head waiting time* of class i at t , i.e., the waiting time of the packet at the head of class i at t . We define the *normalized head waiting time* of class i at t as

$$\tilde{w}_i(t) = w_i(t)/\delta_i \tag{15}$$

Every time a packet is to be transmitted, the WTP scheduler selects the backlogged class j with the *maximum normalized head waiting time*,

$$j = \arg \max_{i \in B(t)} \tilde{w}_i(t) \tag{16}$$

The similarities with PAD are now obvious. In the same way that PAD chooses for service the class with the maximum normalized average delay, WTP chooses for service the class with the maximum normalized head waiting time. PAD attempts to minimize in this manner the differences of the class normalized average delays, while *WTP attempts to minimize the differences between the normalized waiting times of successively departing packets.* Suppose, hypothetically, that WTP manages to

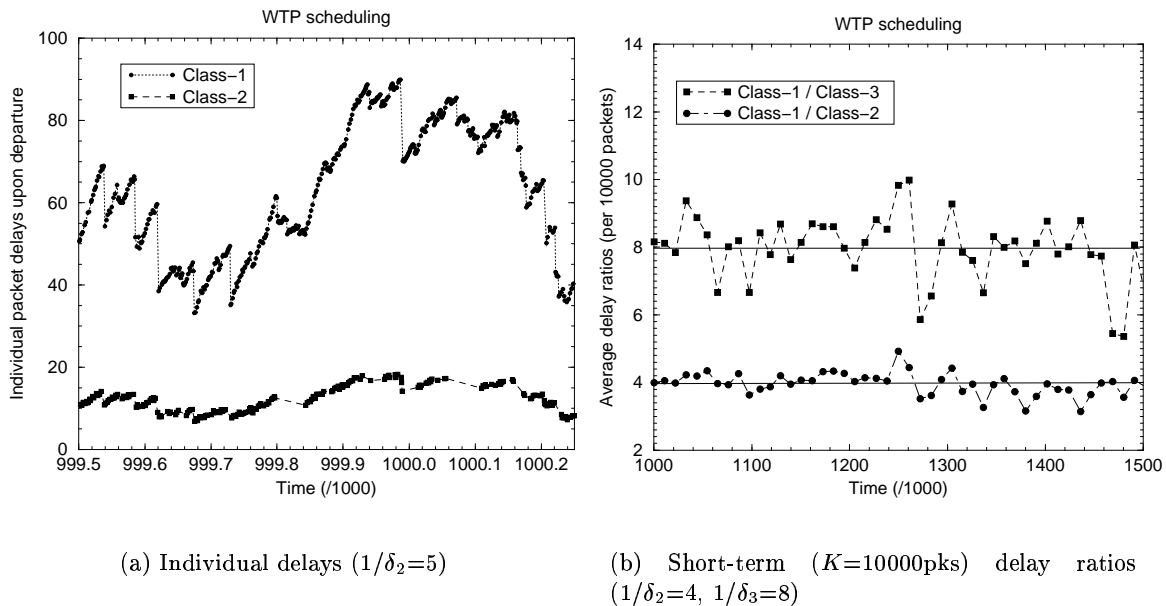


Figure 5: Individual packet delays and short-term delay ratios with WTP.

always reduce these differences to zero. The queueing delays of successively departing packets will be then proportional to the given DDPs,

$$\frac{d^m}{d^{m+1}} = \frac{\delta_{c(m)}}{\delta_{c(m+1)}} \quad (17)$$

where d^m and $c(m)$ are the queueing delay and the class, respectively, of the m -th departing packet. Of course, this goal is not always feasible. Equation (17) shows, though, that *WTP attempts to achieve a proportional differentiation between the delays of successive packet departures.*

In terms of scalability and performance, WTP requires at most $N-1$ comparisons for each packet transmission, which is a minor overhead for the eight classes or so that we consider. To avoid the multiplication of the waiting time $w_i(t)$ with the factor $1/\delta_i$, the waiting time can be increased by $1/\delta_i$ after each time unit (instead of incrementing it). An implementation requirement, which also applies to PAD, is that packets have to be timestamped upon arrival so that their delays can be measured; we do not expect this requirement to be an important difficulty in practice.

WTP is an excellent scheduler in terms of providing higher classes with lower delays in short timescales. To illustrate the behavior of WTP in short timescales, the two graphs (a) and (b) of

Figure 5 show the individual delays of successive packet departures, and the ratios of short-term class delays averaged in every $K=10000$ packets, respectively. To compare WTP and PAD, these graphs refer to the same time interval and the same traffic stream as the graphs in Figure 4. First, note that in Figure 5-a WTP services the Class-2 packets with lower queueing delays, even in the shortest timescales of successive packet departures. A more careful examination of that graph also shows that the individual packet delays in Class-2 are approximately proportional to the corresponding Class-1 packet delays, and that the proportionality factor is about five, which is the specified DDP ($1/\delta_2=5$). In other words, as stated in Equation 17, *WTP attempts to provide proportional delay differentiation between successive packet departures, and this objective is closely approximated when the delays are sufficiently large.*

Second, note that in Figure 5-b the correct class ordering is also maintained when the delays are measured over short timescales, in this graph every $K=10000$ packets. Figures 9 and 10 in the next section, show that this is also true in a wide span of timescales, ranging from a few packets to many thousands of packets. So, *with WTP the average queueing delays in higher classes are smaller than the average delays in lower classes, independent of the location or the duration of the averaging time window.* Also note that the measured delay ratios in Figure 5-b vary around the specified DDPs, $1/\delta_2=4$ and $1/\delta_3=8$, even though the deviations are sometimes significant. In summary, WTP provides predictable delay differentiation, and it approximates the PDD model in short timescales. Does WTP, however, meet the PDD constraints of (1)?

In the special case of Poisson arrivals, we can show analytically that WTP converges to the PDD model as the utilization approaches 100%. The following result is proved in Appendix III.

Proposition 1: If the packet arrivals in each class are generated from a Poisson distribution, the WTP scheduler converges to the PDD model of (1) as the utilization u tends to 100%,

$$\frac{\bar{d}_i^{WTP}}{\bar{d}_j^{WTP}} \rightarrow \frac{\delta_i}{\delta_j} \quad \text{as } u \rightarrow 1 \quad (18)$$

Since the Poisson assumption has been shown to be often invalid in packet networks [23, 24], we have also examined the validity of Equation (18) with simulations, using the infinite-variance Pareto interarrival distribution. The general observation from our empirical study is that *WTP meets the PDD model as the aggregate backlog in the link \bar{q}_{ag} tends to infinity.*

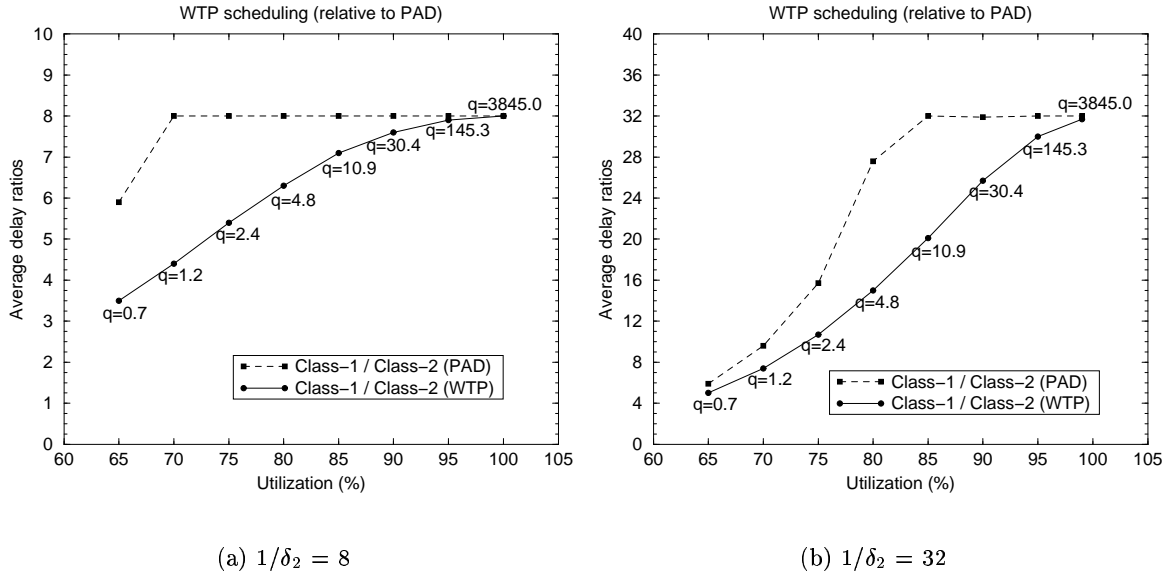


Figure 6: Average delay ratios with WTP and PAD as a function of the utilization.

How close to ‘infinity’ does the aggregate backlog have to be, though, in order for WTP to closely approximate the PDD model? Figure 6 shows the achieved delay ratios for the case of two classes and for a uniform load distribution, as a function of the utilization. The aggregate backlog for each utilization point is also shown (denoted by q in the graphs). Note that the aggregate backlog does not depend on the load distribution when all classes have the same statistical characteristics, or on the specified DDPs. The corresponding curves for PAD are also shown, for comparison purposes.

Figure 6 shows that WTP tends to the PDD model as the utilization approaches 100% and the aggregate backlog increases. For a moderate delay ratio, such as $1/\delta_2=8$, the required average backlog for a close approximation of the specified DDP is a few tens of packets. This is not an unusual average backlog for heavily utilized, high-speed links [25]. In order to closely approximate a more drastic delay differentiation with WTP however, such as $1/\delta_2=32$, an average backlog of hundreds of packets is needed. This may be impractical, especially when there are not enough buffers, or when the network operators prefer to drop packets (and so limit their maximum delay) instead of queueing them for hundreds of packet transmission times. When the utilization is less than about 85%, WTP shows a large deviation from the PDD model, either because the specified DDP is infeasible (as in Figure 6-b for $u < 85\%$), or simply because WTP is not a particularly good scheduler in terms of meeting the long-term average delay ratios of the PDD model.

5 Hybrid Proportional Delay (HPD) scheduling

As shown in §3, PAD attempts to minimize the differences between normalized average class delays. This objective makes PAD capable of meeting the PDD model when the given DDPs are feasible. WTP, on the other hand, attempts to minimize the differences between normalized head waiting times. As a consequence of this behavior, WTP provides higher classes with lower delays even in the shortest timescales. Can we combine the operation of these two scheduling algorithms in order to create a scheduler that is both close to the PDD model, when the DDPs are feasible, and also that provides a predictable delay differentiation in short timescales? This is the objective of the *Hybrid Proportional Delay (HPD)* scheduling discipline, described next.

HPD, just like PAD and WTP, maintains a delay metric for each class i , normalized by the corresponding DDP δ_i . Specifically, if $\tilde{d}_i(t)$ is the normalized average delay in class i at t , as defined in (13), and $\tilde{w}_i(t)$ is the normalized head waiting time in class i at t , as defined in (15), the *normalized hybrid delay* in HPD is

$$\tilde{h}_i(t) = g\tilde{d}_i(t) + (1 - g)\tilde{w}_i(t) \quad (19)$$

where g is the *HPD parameter* ($0 \leq g \leq 1$). When a packet is to be transmitted, HPD chooses the backlogged class j with the maximum normalized hybrid delay,

$$j = \arg \max_{i \in B(t)} \tilde{h}_i(t) \quad (20)$$

When g is zero HPD becomes equivalent to WTP, while when g is one HPD becomes equivalent to PAD. For other values of g , HPD combines the operation of PAD and WTP resulting in a hybrid behavior.

An ‘optimal’ selection of g would attempt to maximize a given performance metric that evaluates HPD’s ability to approximate the PDD model, and at the same time to provide predictable delay differentiation in short timescales. We do not pursue such an optimal selection of g here, however. Instead, through an empirical study, we found that $g=0.875$ is a ‘good’ value in the trade-off between the behaviors of PAD and WTP. Simulation results supporting this choice are given at the end of this section. Note that $0.875=1-2^{-3}$, which also means that the multiplication with g and $1 - g$ can be replaced with shift and subtract operations.

In the following, we examine different aspects of the HPD behavior and compare it with PAD

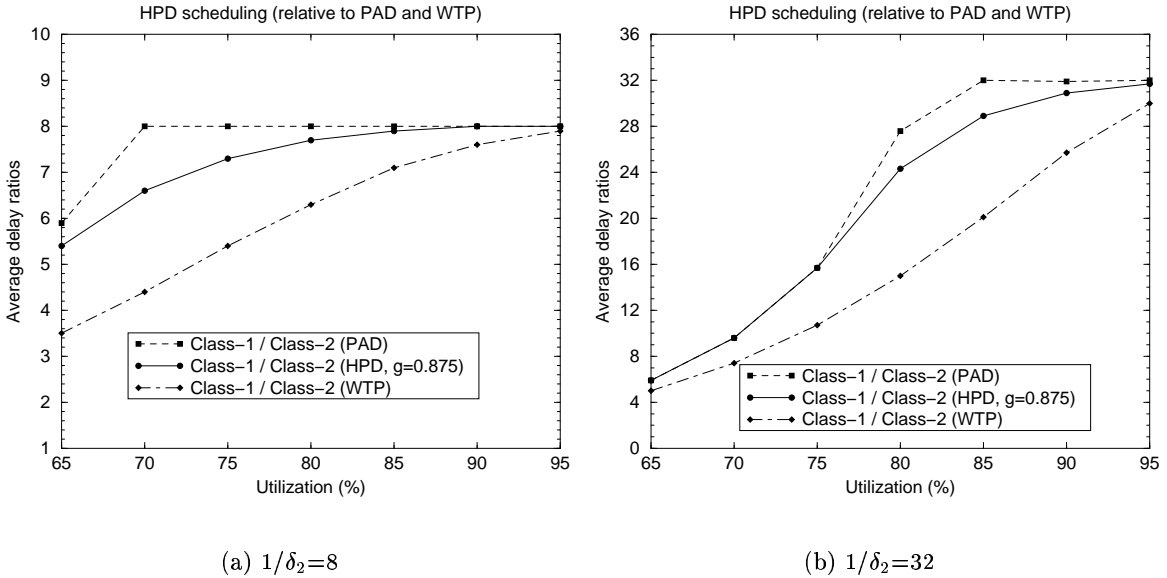


Figure 7: Average delay ratios with HPD as a function of the utilization.

and WTP using simulation results.

Effect of the aggregate load on the average delay ratios: Figure 7 shows the effect of the utilization u on the HPD average delay ratios. Results are shown for two classes, when the specified delay ratio is either $1/\delta_2=8$ or $1/\delta_2=32$. The load distribution between the two classes is uniform. For comparison purposes we also show the PAD and WTP curves.

As the utilization increases, all three schedulers tend to the specified delay ratios. The differences between HPD and PAD in heavy load conditions, above 90% or so, are minor. When the utilization is between 70% to 90%, HPD is significantly closer to the PDD constraints than WTP. The largest relative error between HPD and PAD in this operating region is about 20% in these graphs. This error can be further reduced with a higher value of the HPD parameter g , if the network operator values the importance of meeting the specified DDPs more than providing a consistent delay ordering between classes in short timescales. In light load conditions, less than about 70% or 75%, the PDD model becomes infeasible, and so the PAD scheduler cannot achieve the specified delay ratios either.

Effect of the load distribution on the average delay ratios: Figure 8 shows the effect of the class load distribution on the HPD average delay ratios. Results are shown for four classes and seven class load distributions, when the utilization is 75% and 95%. The DDPs are: $\delta_1/\delta_2 = \delta_2/\delta_3 =$

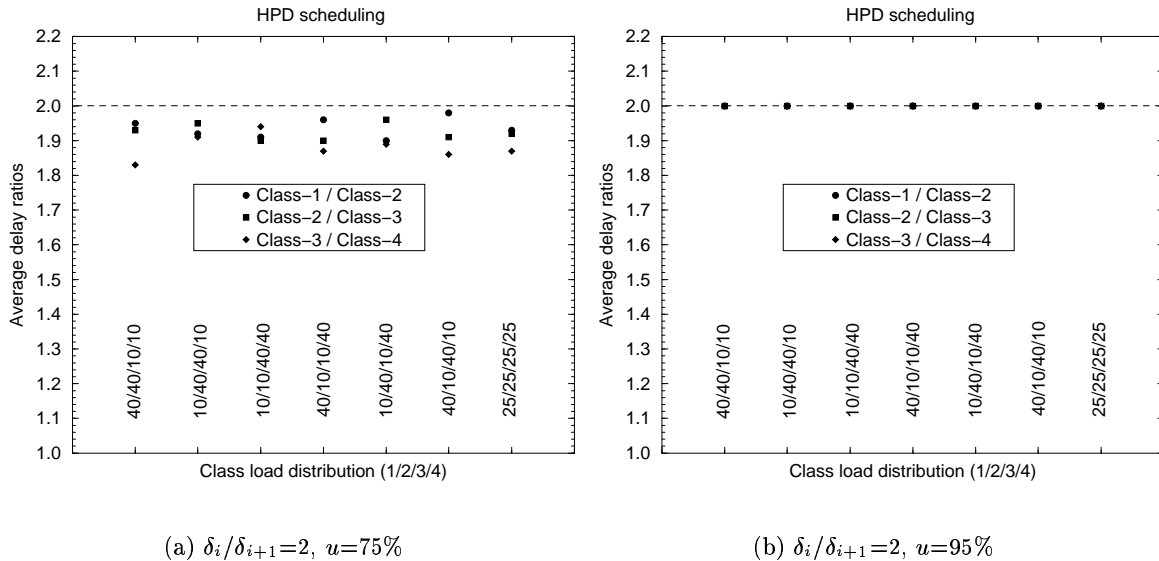


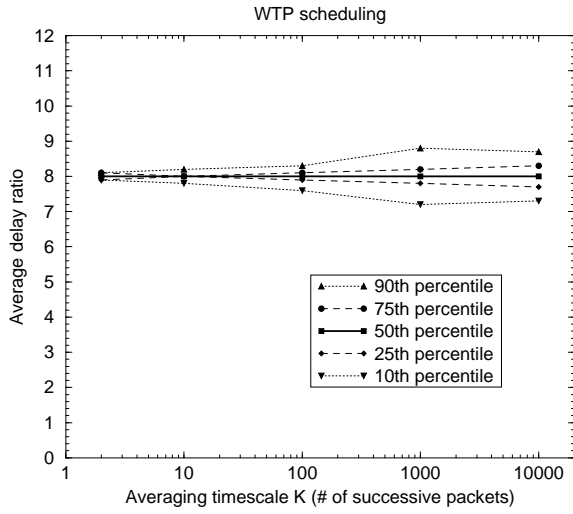
Figure 8: Average delay ratios with HPD as a function of the load distribution.

$\delta_3/\delta_4=2$, i.e., each class should have a twice as large average delay than the next higher class.


Ideally, HPD should meet the PDD model independent of the class load distribution. The set of DDPs is feasible in both load conditions. When the utilization is 75% (Figure 8-a), HPD closely approximates the specified delay ratio between successive classes, and the deviations from the PDD model are less than 10%. In heavy load conditions (Figure 8-b), HPD accurately meets the specified delay ratios in all load distributions. In general, HPD manages to closely approximate the PDD model, when it is feasible, *independent of the class load distribution*. The approximation errors increase as the utilization decreases, and as the required delay differentiation between classes becomes more extreme.

Delay differentiation with HPD and WTP in short timescales: In evaluating the predictability of a scheduler, our primary focus is to examine whether higher classes encounter lower delays in all timescales. It has already been illustrated in the previous graphs that HPD provides higher classes with lower delays, but only in terms of long-term averages. Is the delay ordering between classes, however, also met in short timescales?

Figures 9 and 10 investigate the delay differentiation that HPD and WTP produce in short timescales of different length. The graphs in these two figures show *five percentiles of the distribution of delay ratios between two classes, when the delay ratios are measured in successive time windows of*

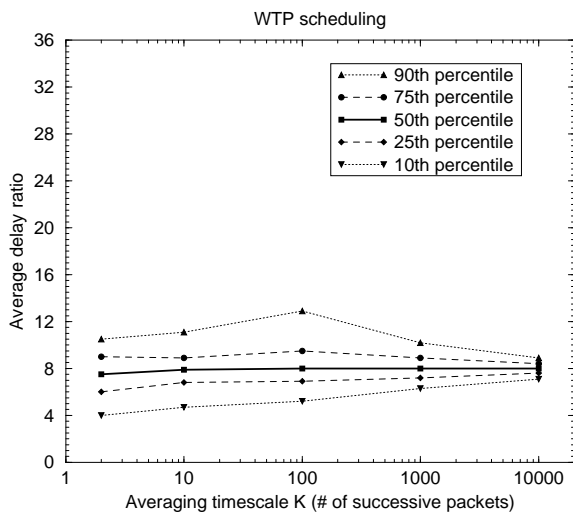


(a) WTP, $1/\delta_2=8$

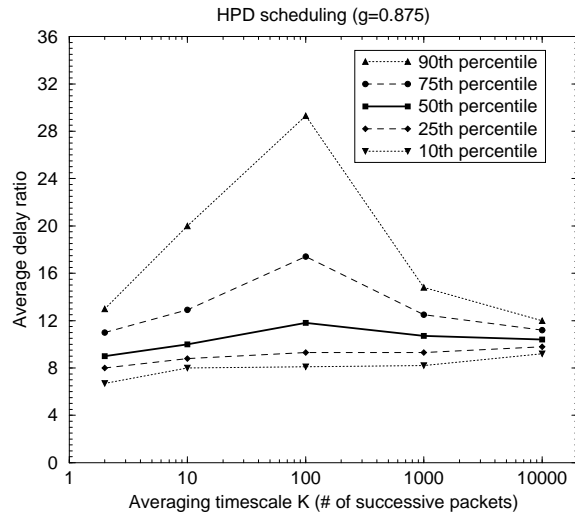


(b) HPD, $1/\delta_2=8$

Figure 9: Percentiles of the average delay ratios with WTP and HPD as a function of the averaging timescale K ($u=95\%$).



(a) WTP, $1/\delta_2=8$



(b) HPD, $1/\delta_2=8$

Figure 10: Percentiles of the average delay ratios with WTP and HPD as a function of the averaging timescale K ($u=80\%$).

K packet departures. *K* determines the timescale in which we measure the class delays and compute their ratio. In these experiments, *K* ranges from successive packet departures ($K=2$) to several thousands of packet departures ($K=10000$). The five percentiles shown give a comprehensive view of the distribution of the resulting delay ratios, providing the median delay ratio (50th percentile), the 25th and 75th percentiles, as well as two tail delay ratios (10th and 90th percentiles). Note that when there are no packet departures from both classes in a time period of *K* packet departures, a delay ratio is not measured. Also, the duration of the simulation runs is adjusted so that we get approximately the same number of delay ratios for all values of *K*. In both figures, the specified DDP is $\delta_2=1/8$ and the class load distribution is uniform.

Figure 9 shows the results of these experiments in heavy load conditions ($u=95\%$). As shown, the resulting distributions of delay ratios are quite narrow, centered around the specified DDPs $\delta_1/\delta_2=8$, in both HPD and WTP. This means that in heavy load conditions, HPD inherits the excellent predictability of WTP, i.e., it provides lower delays to higher classes, and it approximates closely the PDD model. This behavior is consistent in a range of timescales, from successive departures to several thousands of packet departures. In Figure 10, on the other hand, the utilization is $u=80\%$. In this case, the HPD delay ratio distributions are significantly wider than the WTP distributions. Additionally, the HPD distributions are not centered exactly around the specified $\delta_1/\delta_2=8$, but they indicate slightly higher delay ratios. Even if HPD does not approximate the PDD model as closely as WTP does in moderate or low load conditions, however, it still provides lower delays to the higher class, and so it produces predictable delay differentiation.

Choosing a ‘good’ HPD parameter value: Figure 11 shows the effect of the HPD parameter g on the ratio of average delays for two classes. In heavy load conditions ($u=95\%$), the selection of g does not matter much, and the PDD model is closely approximated for practically any g . The reason is that in heavy load conditions both PAD and WTP meet the PDD model, and so does HPD. In lower load conditions ($u=75\%$) though, g needs to be close to one in order for HPD to approximate the PDD model well. Note that the dependency of the delay ratio on g is convex, meaning that the approximation error decreases faster as g approaches one.

A good value of g has to also take into account the predictability of HPD. Figure 12 shows five percentiles for the average delay ratios with HPD for $g=0.8$ and $g=0.95$, as it was done in Figure 10-b for $g=0.875$. The class load distribution is uniform, and the utilization is $u=80\%$. Note that as g decreases, the width of the resulting delay ratio distributions decreases as well. This

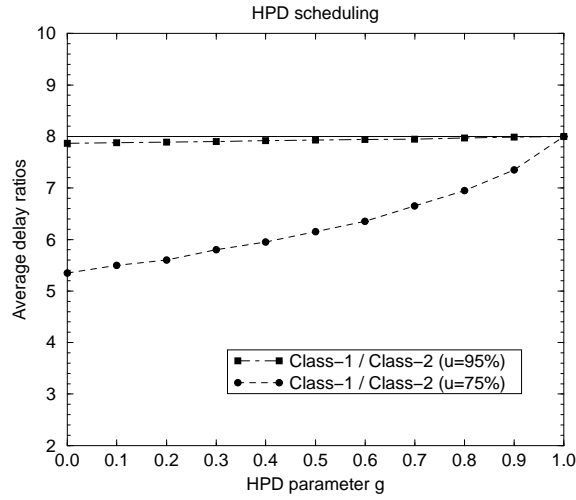
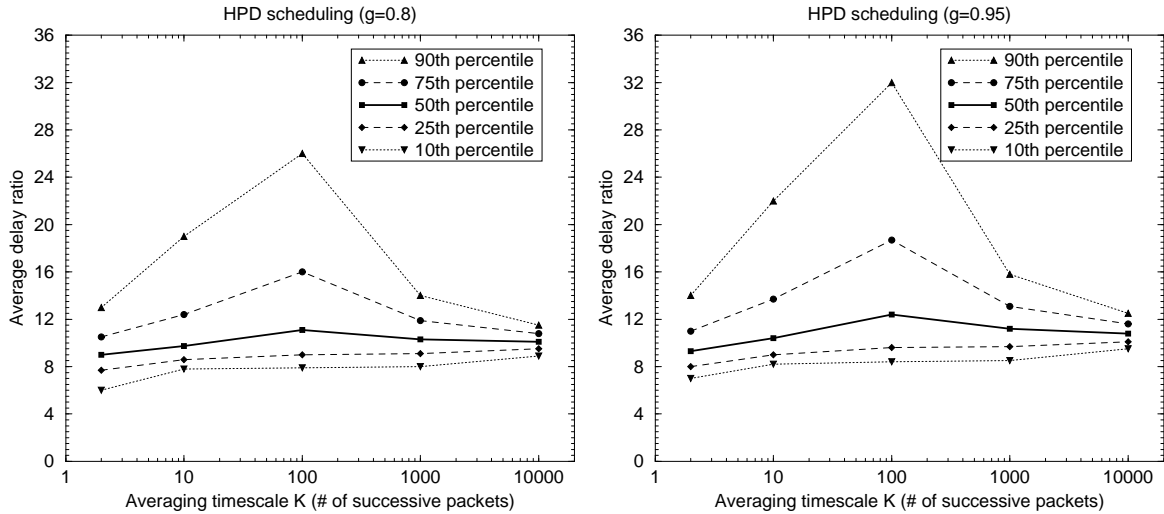


Figure 11: Effect of the HPD parameter g on the average delay ratios ($1/\delta_2=8$).



(a) HPD, $g=0.8$

(b) HPD, $g=0.95$

Figure 12: Percentiles of the average delay ratios with HPD for two values of g , as a function of the averaging timescale K ($u=80\%$).

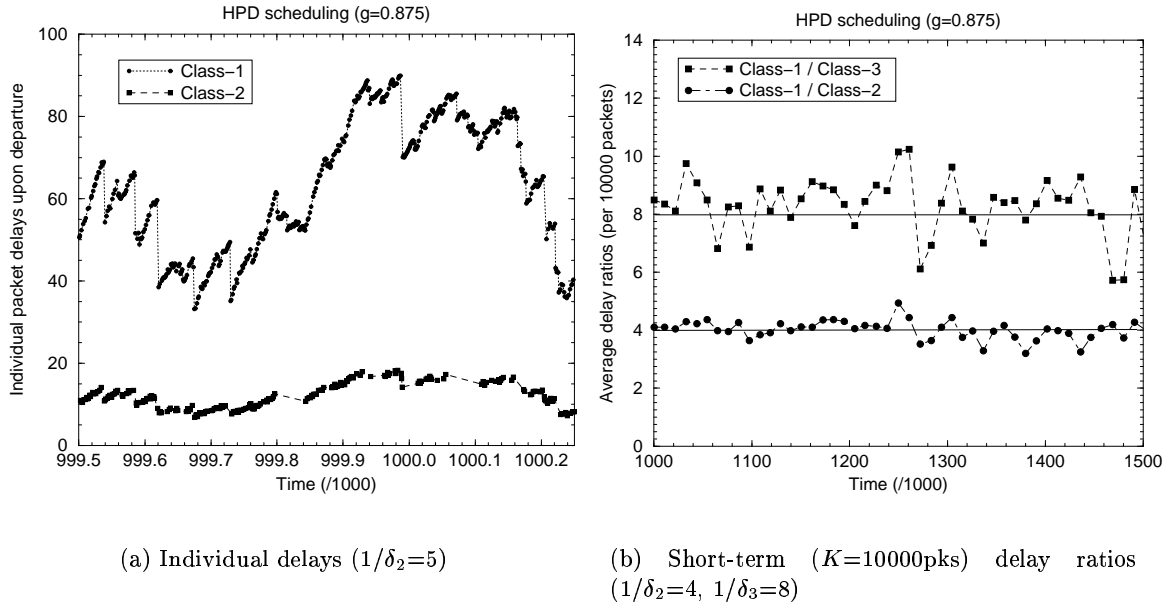


Figure 13: Individual packet delays and short-term delay ratios with HPD.

is expected, since as g decreases, HPD becomes closer to WTP. As a good and practical trade-off between approximating the PDD model and achieving consistent delay differentiation in short timescales, we choose $g=0.875$.

To further illustrate that HPD provides consistent delay differentiation in short timescales when $g=0.875$, Figure 13 shows the same sample paths of individual delays and short-term delay ratios for HPD, as Figure 5 shows for WTP, and Figure 4 shows for PAD. The HPD parameter is set to $g=0.875$. Comparing carefully the two graphs in Figures 13-a and 5-a shows that, with $g=0.875$, HPD provides almost indistinguishable results with WTP.

A different traffic model: heterogeneous ON-OFF sources: The simulation experiments presented thus far used the same arrival and packet size distributions for all classes. Additionally, the simulated traffic model was Pareto interarrivals with fixed packet sizes. In this paragraph, we simulate a radically different traffic model. Specifically, the traffic in each class is generated by an ON-OFF source. The ON phase is exponentially distributed with an average duration of 30, 20, 10, and 1 time units for Classes 1, 2, 3, and 4, respectively. The OFF phase is Pareto distributed with a scale parameter α of 1.2, 1.5, 1.7, and 1.9 for Classes 1, 2, 3, and 4, respectively.

Figure 14 shows the average delay ratios that result with HPD for seven class load distributions

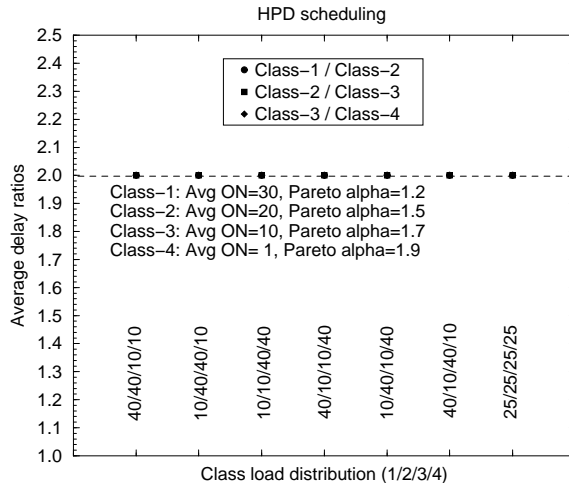


Figure 14: Average delay ratios with HPD in different load distributions and heterogeneous ON-OFF traffic ($u=95\%$).

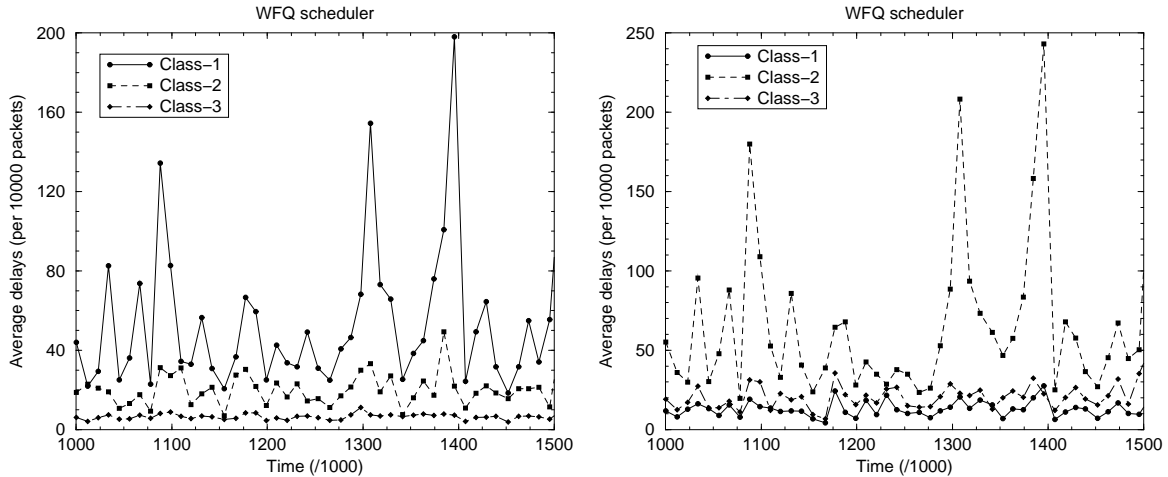
in heavy load conditions ($u=95\%$). The load distributions are in bytes. The DDPs are: $\delta_1/\delta_2 = \delta_2/\delta_3 = \delta_3/\delta_4=2$, i.e., each class should have a twice as large average delay than the next higher class. The important point in this graph is that HPD tends to the PDD model in heavy load conditions, even with this radically different traffic model. Similar results hold for PAD and WTP.

6 Related work on delay differentiation

In this section, we review some other scheduling algorithms for relative or proportional delay differentiation. We start with the family of *link sharing schedulers*, showing that they can provide controllable delay differentiation, but they are too sensitive to class load distribution changes. The next scheduler (BPR) is a dynamic variation of GPS, and it can also approximate the PDD model. Finally, we present recent contributions made by other researchers on the problem of proportional delay differentiation.

6.1 Link sharing schedulers

Several packet schedulers aim to *provide each class with a minimum bandwidth share of the link's capacity*. Examples of such schedulers are the packet-based approximations of GPS [26] (such as WFQ [27]), as well as the Class Based Queueing (CBQ) [28] and Hierarchical Packet Fair Queueing



(a) $(\lambda_1, \lambda_2, \lambda_3) = (50, 30, 20)$

(b) $(\lambda_1, \lambda_2, \lambda_3) = (40, 35, 25)$

Figure 15: Average delays with a three-class WFQ scheduler in two slightly different load distributions.

(H-PFQ) [29] link sharing schedulers. These mechanisms were developed in the link sharing context, where different organizations or users are guaranteed a certain fraction of a link’s capacity, sharing any available excess bandwidth. The link sharing schedulers have been also proposed for providing relative delay differentiation. For example, [19] proposes the use of such a scheduler to implement the *Olympic service model*, which consists of the ‘Gold’, ‘Silver’, and ‘Bronze’ classes.

A drawback of the link sharing schedulers is that slight changes in the class load distribution affect dramatically the resulting delay differentiation. Figure 15 shows short-term average delays in a three-class WFQ scheduler, measured in every $K=10000$ packet departures. The WFQ weights are selected as $(w_1, w_2, w_3) = (0.5, 0.35, 0.3)$ so that, when the load distribution is $(\lambda_1, \lambda_2, \lambda_3) = (50, 30, 20)$, the average delay in Class-1 is about 50, in Class-2 about 20, and in Class-3 about 7 time units. In Figure 15-a, notice that the selected WFQ weights create a predictable delay differentiation, providing the maximum delays to Class-1 and the lowest delays to Class-3. If, however, only 10% of the Class-1 traffic moves to Class-2 and Class-3, leading to the load distribution $(\lambda_1, \lambda_2, \lambda_3) = (40, 35, 25)$, the class ordering is violated, causing the maximum delays for Class-2 and the minimum delays for Class-1 (see Figure 15-b).

Obviously, such a ‘priority inversion’ between the offered classes would be unacceptable for the

higher class users. The previous example illustrates that link sharing schedulers are quite sensitive to changes in the class load distribution. Such ‘priority inversions’ can occur when the load distribution is not as it was expected when the weights were selected. This can occur when the load distribution is nonstationary, or when the load distribution varies in short timescales due to traffic burstiness. One way to deal with this problem is to dynamically adjust the link share weights based on the current class loads; this is the basic idea behind the following scheduler.

6.2 Backlog Proportional Rate (BPR) scheduling

BPR is a dynamic version of GPS, in which *the class weights are dynamically adjusted based on the instantaneous class loads*. Specifically, let $r_i(t)$ be the service rate that is assigned to queue i at time t . For two backlogged queues i and j , the service rate allocation in BPR follows the proportional constraint:

$$\frac{r_i(t)}{r_j(t)} = \frac{\delta_j q_i(t)}{\delta_i q_j(t)} \quad (21)$$

where $q_i(t)$ is the backlog of queue i at time t . If the queue i is empty at time t , $r_i(t) = 0$. The sum of the assigned service rates $\sum_{i=1}^N r_i(t)$ must be equal to the link capacity C when the scheduler is busy.

The main finding of [1] is that BPR approximates the PDD model of Eq 2 under heavy load conditions. This asymptotic result can be explained informally as follows. In heavy load conditions, all BPR queues are almost always backlogged, due to the ‘simultaneous queue clearing’ property stated and proved in [1]. So, if we write (21) in terms of averages and use Little’s law $\bar{q} = \lambda \bar{d}$, we have that:

$$\frac{r_i}{r_j} = \frac{\delta_j \bar{q}_i}{\delta_i \bar{q}_j} = \frac{\delta_j \bar{d}_i \lambda_i}{\delta_i \bar{d}_j \lambda_j} \quad (22)$$

where r_i is the average output (service) rate in class i . In steady-state and lossless conditions, however, the output rates r_i are equal to the input rates λ_i ($r_i = \lambda_i$), and so BPR tends to the PDD constraints,

$$\frac{\bar{d}_i^{BPR}}{\bar{d}_j^{BPR}} = \frac{\bar{d}_i}{\bar{d}_j} = \frac{\delta_i}{\delta_j} \quad (23)$$

Even though BPR tends to the PDD model, [1] showed that it does not do so as well as WTP. Specifically, the deviations from the PDD model for the same load conditions are higher, and the delay ratios in short timescales have a wider distribution. In addition, BPR suffers from the *simultaneous*

queue clearing property, according to which all queues in a BPR busy period become empty at the same time. This effect causes occasional delay bursts in the high class packets.

6.3 Recent contributions on proportional delay differentiation

The WTP scheduler received further attention in [30]. The authors, starting from Kleinrock's analysis of WTP [3] and assuming Poisson arrivals, derived the feasible load distribution range for a given set of DDPs. They also proposed a numerical method for calculating the scheduling weights (when $N > 2$), so that WTP can achieve a feasible set of DDPs for any given load conditions.

Another scheduler for proportional delay differentiation, called *Mean-Delay Proportional (MDP)*, was proposed in [31]. MDP is similar to HPD, in the sense that it also chooses for service at time t the class with the maximum normalized 'delay' metric $\tilde{D}_i(t)$. The average delay metric in MDP takes into account both the previously departed packets, and also *all the backlogged packets*. [32] and [33] proposed two proportional delay schedulers (called *Proportional Queue Control Mechanism (PQCM)* and *Dynamic Weighted Fair Queueing (D-WFQ)*, respectively) which are based on the GPS rate allocation mechanism. Both schedulers adjust the GPS weights periodically, say every Δ seconds, based on the measured class backlogs and input rates, attempting to achieve proportional average delay differentiation in the next window of Δ seconds.

Other researchers proposed different relative differentiation models than the PDD model of §2, and studied scheduling mechanisms that can implement them. For instance, [34] proposed the *Local Optimal Proportional Differentiation (LOPD)* scheduler, which attempts to provide proportional average delay differentiation *in each busy period*, assuming that there will be no further arrivals after each packet departure. LOPD is optimal for this differentiation objective of [34]. [35] extended the PDD model in the direction of *deadline violation probabilities*, proposing a scheduler that aims for proportional delay violation probabilities among classes. Finally, [36] proposed a *Joint Buffer management and Scheduling (JoBS)* mechanism which determines both the scheduling order and the packet drop decisions. The basic idea is that JoBS makes a prediction after every packet arrival for the delays of the backlogged traffic, and modifies the service rates so that all QoS constraints are met. If this is not possible, some packets are dropped.

7 Conclusions

This paper applied the general model of proportional differentiation in the context of queueing delays. The dynamics and feasibility of the PDD model were derived, and three schedulers that approximate the PDD model in heavy load conditions were designed. The three schedulers differ in the trade-off between approximating the PDD model closely, and providing consistent delay differentiation in short timescales.

The proportional differentiation architecture has three additional components. First, it can be applied in the context of coupled delay and loss differentiation, considering a lossy packet multiplexer. The Proportional Loss Differentiation (PLD) model, as well as the design of appropriate packet droppers was studied in [2]. Second, the relative differentiation that the proportional model offers can be combined with dynamic class selection at the end-points, in order to meet absolute and end-to-end QoS objectives. A first investigation of this technique has appeared in [17]. Third, a network operator can use the PDD model and a scheduler such as HPD to achieve a certain average delay in each class, as described in [37]. This provisioning process determines the link capacity and the scheduling DDPs, given the desired average delay and the target traffic load in each class.

Future work in the context of proportional delay differentiation can pursue a deeper mathematical treatment of the proposed schedulers. The following issues deserve further investigation. First, the feasibility conditions for the PDD model can be analyzed in greater depth, assuming a simpler traffic model. Second, it is likely that PAD can be proven to be optimal in meeting the PDD model within a certain family of scheduling algorithms and for certain traffic conditions. Third, it is likely that WTP can be proven to converge to the PDD model in heavy load conditions, even without assuming Poisson arrivals, as long as the aggregate average backlog is sufficiently large.

References

- [1] C. Dovrolis, D.Stiliadis, and P.Ramanathan, “Proportional Differentiated Services: Delay Differentiation and Packet Scheduling,” in *Proceedings of ACM SIGCOMM*, Sept. 1999.
- [2] C. Dovrolis and P.Ramanathan, “Proportional Differentiated Services, Part II: Loss Rate Differentiation and Packet Dropping,” in *IEEE/IFIP International Workshop on Quality of Service (IWQoS)*, June 2000.

- [3] L. Kleinrock, *Queueing Systems, Volume II*. John Wiley and Sons, 1976.
- [4] S. Blake, D.Black, M.Carlson, E.Davies, Z.Wang, and W.Weiss, *An Architecture for Differentiated Services*, Dec. 1998. IETF RFC 2475.
- [5] P. P. White, “RSVP and Integrated Services in the Internet: A Tutorial,” *IEEE Communications Magazine*, pp. 100–106, May 1997.
- [6] V. Jacobson, K. Nichols, and K.Poduri, *An Expedited Forwarding PHB*, June 1999. RFC 2598.
- [7] D. D. Clark and W. Fang, “Explicit Allocation of Best Effort Packet Delivery Service,” *IEEE/ACM Transactions on Networking*, vol. 6, pp. 362–373, Aug. 1998.
- [8] A. Charny and J. L. Boudec, “Delay Bounds in a Network with Aggregate Scheduling,” in *Proceedings QOFIS*, Oct. 2000.
- [9] I. Stoika and H.Zhang, “LIRA: An Approach for Service Differentiation in the Internet,” in *Proceedings NOSSDAV*, 1998.
- [10] S. Sahu, P.Nain, D.Towsley, C.Diot, and V.Firoiu, “On Achievable Service Differentiation with Token Bucket Marking for TCP,” in *Proceedings of ACM SIGMETRICS*, June 2000.
- [11] I. Stoika and H.Zhang, “Providing Guaranteed Services Without Per Flow Management,” in *Proceedings of ACM SIGCOMM*, Sept. 1999.
- [12] I. Stoika, S.Shenker, and H.Zhang, “Core-Stateless Fair Queueing: Achieving Approximately Fair Bandwidth Allocations in High Speed Networks,” in *Proceedings of ACM SIGCOMM*, Sept. 1998.
- [13] C. Cetinkaya and E.W.Knightly, “Egress Admission Control,” in *Proceedings of IEEE INFOCOM*, Mar. 2000.
- [14] J. L. Boudec, M. Hamdi, L. Blazevic, and P.Thiran, “Asymmetric Best Effort Service for Packet Networks,” in *Proceedings Global Internet Symposium*, Dec. 1999.
- [15] C. Dovrolis and P.Ramanathan, “A Case for Relative Differentiated Services and the Proportional Differentiation Model,” *IEEE Network*, Oct. 1999.
- [16] K. Nichols, S.Blake, F. Baker, and D.L.Black, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*, Dec. 1998. IETF RFC 2474.

- [17] C. Dovrolis and P.Ramanathan, “Dynamic Class Selection: From Relative Differentiation to Absolute QoS,” in *Proceedings of IEEE International Conference on Network Protocols (ICNP)*, Nov. 2001.
- [18] A. M. Odlyzko, “Paris Metro Pricing: The Minimalist Differentiated Services Solution,” in *Proceedings IEEE/IFIP International Workshop on Quality of Service*, June 1999.
- [19] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, *Assured Forwarding PHB Group*, June 1999. RFC 2597.
- [20] G. Bolch, S.Greiner, H.Meer, and K.S.Trivedi, *Queueing Networks and Markov Chains*. John Wiley and Sons, 1999.
- [21] E. Coffman and I.Mitrani, “A Characterization of Waiting Time Performance Realizable by Single-Server Queues,” *Operations Research*, vol. 28, pp. 810–821, May 1980.
- [22] J. Regnier, “Priority Assignment in Integrated Services Networks,” Tech. Rep. LIDS-TH-1565, LIDS-MIT, Dec. 1986.
- [23] V. Paxson and S.Floyd, “Wide Area Traffic: The Failure of Poisson Modeling,” *IEEE/ACM Transactions on Networking*, vol. 3, pp. 226–244, June 1995.
- [24] W. Willinger, M.S.Taqqu, R.Sherman, and D.V.Wilson, “Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level,” in *Proceedings of ACM SIGCOMM*, pp. 100–113, Sept. 1995.
- [25] B. Reynolds, “RED Analysis for Congested Network Core and Customer Egress,” tech. rep., QualNet, Jan. 1999. North American Network Operators’ Group (NANOG) meeting.
- [26] H. Zhang, “Service disciplines for guaranteed performance service in packet-switching networks,” *IEEE Proceedings*, vol. 83, pp. 1374–1399, Oct. 1995.
- [27] A. Demers, S.Keshav, and S.Shenker, “Analysis and Simulation of a Fair Queueing Algorithm,” in *Internetworking: Research and Experience*, pp. 3–26, 1990.
- [28] S. Floyd and V. Jacobson, “Link-Sharing and Resource Management Models for Packet Networks,” *IEEE/ACM Transactions on Networking*, vol. 3, pp. 365–386, Aug. 1995.

- [29] J. Bennett and H.Zhang, “Hierarchical Packet Fair Queueing Algorithms,” *IEEE/ACM Transactions on Networking*, vol. 5, pp. 675–689, Oct. 1997.
- [30] M. K. H. Leung, J. Lui, and D. Yau, “Characterization and Performance Evaluation for Proportional Delay Differentiated Services,” in *Proceedings International Conference on Network Protocols (ICNP)*, Oct. 2000.
- [31] T. Nandagopal, N.Venkitaraman, R.Sivakumar, and V.Bharghavan, “Delay Differentiation and Adaptation in Core Stateless Networks,” in *Proceedings of IEEE INFOCOM*, Mar. 2000.
- [32] Y. Moret and S.Fdida, “A Proportional Queue Control Mechanism to Provide Differentiated Services,” in *International Symposium on Computer and Information Systems (ISCIS)*, Oct. 1998.
- [33] C. C. Li, S.-L. Tsao, M. C. Chen, Y. Sun, and Y.-M. Huang, “Proportional Delay Differentiation Service Based on Weighted Fair Queueing,” in *Proceedings IEEE International Conference on Computer Communications and Networks (ICCCN)*, Oct. 2000.
- [34] H. Saito, C. Lukovszki, and I. Moldovan, “Local Optimal Proportional Differentiated Services Scheduler for Relative Differentiated Services,” in *Proceedings IEEE International Conference on Computer Communications and Networks (ICCCN)*, Nov. 2000.
- [35] S. Bodamer, “A Scheduling Algorithm for Relative Delay Differentiation,” in *Proceedings of the IEEE Conference on High Performance Switching and Routing (ATM 2000)*, June 2000.
- [36] J. Liebeherr and N.Christin, “Buffer Management and Scheduling for Enhanced Differentiated Services,” Tech. Rep. CS-2000-24, University of Virginia, Aug. 2000.
- [37] C. Dovrolis and P.Ramanathan, “Class Provisioning in Proportional Differentiated Services Networks,” in *Proceedings of Scalability and Traffic Control in IP Networks (SPIE ITCOM304 conference)*, Aug. 2001.

Appendix I: Delay dynamics in PDD

In the following proofs, $S = \sum_{n=1}^N \lambda_n \delta_n$.

Proof of Property (1).

Since the DDPs are positive, an increase in the average delay of one class causes an increase in the average delays of all classes. Since an increase in the average rate of the aggregate traffic stream cannot cause a decrease in all class delays, we have that

$$\frac{\partial \bar{d}_i}{\partial \lambda_j} \geq 0 \quad 1 \leq i, j \leq N \quad (24)$$

◇

Proof of Property (2).

From Equation (5) we have that

$$\frac{\partial \bar{d}_i}{\partial \lambda_j} - \frac{\partial \bar{d}_i}{\partial \lambda_k} = \frac{\delta_i}{S^2} \left[S \left(\frac{\partial \bar{q}_{ag}}{\partial \lambda_j} - \frac{\partial \bar{q}_{ag}}{\partial \lambda_k} \right) + \bar{q}_{ag} (\delta_k - \delta_j) \right] \quad (25)$$

Since the interarrivals and packet sizes follow the same distributions in all classes, the average backlog \bar{q}_{ag} does not depend on the class that the traffic belongs to, and so $\partial \bar{q}_{ag} / \partial \lambda_j = \partial \bar{q}_{ag} / \partial \lambda_k$. If $k < j$, $\delta_k > \delta_j$, and so

$$\frac{\partial \bar{d}_i}{\partial \lambda_j} > \frac{\partial \bar{d}_i}{\partial \lambda_k} \quad \text{when } k < j \quad (26)$$

◇

Proof of Property (3).

From Equation (5) we have that

$$\frac{\partial \bar{d}_i}{\partial \delta_j} = \frac{-\bar{q}_{ag} \delta_i \lambda_j}{S^2} \leq 0 \quad i \neq j \quad (27)$$

and

$$\frac{\partial \bar{d}_i}{\partial \delta_i} = \frac{\bar{q}_{ag} S - \bar{q}_{ag} \delta_i \lambda_i}{S^2} \geq 0 \quad (28)$$

The equalities are necessary for the case $\lambda_j = 0$ and $\lambda_i = \lambda$, respectively.

◇

Proof of Property (4).

Let λ'_k be the rate and \bar{d}'_k be the average delay of a class k after the load transition from class i to class j . We have that $\lambda'_i = \lambda_i - \epsilon$ ($0 < \epsilon \leq \lambda_i$), $\lambda'_j = \lambda_j + \epsilon$, and $\lambda'_k = \lambda_k$ for all $k \neq i, j$. So,

$$S' = \sum_{n=1}^N \delta_n \lambda'_n = S + \epsilon(\delta_j - \delta_i) \quad (29)$$

When $i > j$, $\delta_j > \delta_i$, and so $S' \geq S$. Note that $S' = S$ when $\lambda_i = \lambda_j = \lambda$. Thus, from Equation (5), $\bar{d}'_k \leq \bar{d}_k$ for all $k = 1 \dots N$. Similarly, when $i < j$, it follows that $\bar{d}'_k \geq \bar{d}_k$. Note that both load distributions have the same average aggregate backlog \bar{q}_{ag} , because the interarrivals and packet sizes follow the same distributions in all classes.

◇

Proof of Property (5).

As in the proof of the previous property, $S' = S + \epsilon(\delta_j - \delta_i)$. Suppose that $i > j$ and thus $\delta_j > \delta_i$. It is easy to see that $\delta_j S \geq \delta_i S'$, because $S \geq \epsilon \delta_i$. Consequently,

$$\bar{d}'_j = \delta_j \bar{q}_{ag} / S' \geq \delta_i \bar{q}_{ag} / S = \bar{d}_i \quad (30)$$

Similarly, when $i < j$, it follows that $\bar{d}'_j \leq \bar{d}_i$.

Appendix II: Simulation model

A random source generates the traffic for each source. The interarrivals between packets of the same class follow a Pareto distribution with shape parameter $\alpha = 1.5$. The distribution has infinite variance when $\alpha < 2$, leading to bursty interarrivals. Pareto interarrivals with infinite variance is probably an appropriate model for highly aggregated network traffic. We have experimented with several distributions for the packet sizes, including the exponential distribution, multimodal distributions around certain common values (e.g., 40, 550, and 1500 bytes), and also with fixed-size packets. In this paper, we only include simulation results for fixed-size packets; we did not observe any qualitative differences in the results when simulating variable-size packets. In all simulation results presented in this paper, the interarrival and packet size distributions are the same for all classes. An exception is the results shown in Figure 14.

The capacity of the scheduler is normalized to $C=1$ average-size packet per time unit. The utilization of the scheduler is $u = \lambda/C$, where λ is the average rate of packet arrivals. The class load distribution is specified in terms of packets. For example, when we say that Class-1 is 70% of the aggregate traffic, we mean that 70% of the packets belong to Class-1. The time unit in all graphs is the time it takes to transmit an average-size packet. The reported delays in all graphs are in these time units. In the case of fixed-size packets, when we say that the queuing delay of a packet is 10 units, we mean that the packet waited in the queue for the transmission of ten other packets.

Appendix III: WTP behavior in heavy load conditions

Kleinrock derived the average queuing delays with WTP [3] in the special case of Poisson arrivals. Assuming that all classes have the same average packet size, and normalizing the service rate to $C=1$, Kleinrock's result states that:

$$\bar{d}_i^{WTP} = \bar{d}_i = \frac{\bar{d}_0/(1-\lambda) - \sum_{k=1}^{i-1} \lambda_k \delta_k (1 - \delta_i/\delta_k)}{1 - \sum_{k=i+1}^N (1 - \delta_k/\delta_i)} \quad i = 1 \dots N \quad (31)$$

where \bar{d}_0 is the average remaining service time for the packet that is being transmitted upon the arrival of a new packet. Recall that $\delta_1 = 1 > \delta_2 > \dots > \delta_N > 0$. We prove Equation (18) using induction in the following manner: after showing that $\bar{d}_2/\bar{d}_1 = \delta_2$, we will assume that $\bar{d}_k/\bar{d}_1 = \delta_k$ for all $k = 2 \dots m < N$ with $m \in \{2, \dots, N-1\}$, and then show that $\bar{d}_{m+1}/\bar{d}_1 = \delta_{m+1}$.

For the initial induction step, it is easy to show from (31) that

$$\frac{\bar{d}_2}{\bar{d}_1} = \frac{1 - \sum_{k=2}^N \lambda_k (1 - \delta_k) - \lambda_1 (1 - \delta_2)}{1 - \sum_{k=3}^N \lambda_k (1 - \delta_k/\delta_2)} \quad (32)$$

and so,

$$\frac{\bar{d}_2}{\bar{d}_1} = \frac{(1-\lambda) + \sum_{k=1}^N \lambda_k \delta_k + \lambda_1 (\delta_2 - 1)}{(1-\lambda) + \sum_{k=1}^N \lambda_k \delta_k / \delta_2 + \lambda_1 (1 - 1/\delta_2)} \quad (33)$$

As $\lambda \rightarrow 1$, and thus, $u \rightarrow 100\%$,

$$\frac{\bar{d}_2}{\bar{d}_1} = \frac{\tilde{S} + \lambda_1 (\delta_2 - 1)}{\tilde{S}/\delta_2 + \lambda_1 (1 - 1/\delta_2)} = \delta_2 \quad (34)$$

where $\tilde{S} = \lim_{\lambda \rightarrow 1} \sum_{k=1}^N \lambda_k \delta_k$. This completes the proof for $N=2$.

For $N > 2$, the inductive assumption is that, as $\lambda \rightarrow 1$,

$$\frac{\bar{d}_k}{\bar{d}_1} \rightarrow \delta_k \quad \text{for all } k = 2 \dots m < N \quad (35)$$

with $m \in \{2, \dots, N-1\}$. Then, the average delay of the $(m+1)$ 'th class is

$$\bar{d}_{m+1} = \frac{\bar{d}_0/(1-\lambda) - \sum_{k=1}^m \lambda_k \delta_k \bar{d}_1 (1 - \delta_{m+1}/\delta_k)}{1 - \sum_{k=m+2}^N \lambda_k (1 - \delta_k/\delta_{m+1})} \quad (36)$$

and the ratio of \bar{d}_{m+1} and \bar{d}_1 becomes

$$\frac{\bar{d}_{m+1}}{\bar{d}_1} = \frac{1 - \sum_{k=1}^N \lambda_k (1 - \delta_k) - \sum_{k=1}^m \lambda_k \delta_k (1 - \delta_{m+1}/\delta_k)}{1 - \sum_{k=1}^N \lambda_k (1 - \delta_k/\delta_{m+1}) + \sum_{k=1}^m \lambda_k (1 - \delta_k/\delta_{m+1})} \quad (37)$$

As $\lambda \rightarrow 1$, let $\tilde{S} = \lim_{\lambda \rightarrow 1} \sum_{k=1}^N \lambda_k \delta_k$, $\tilde{S}_m = \lim_{\lambda \rightarrow 1} \sum_{k=1}^m \lambda_k \delta_k$, and $\tilde{\lambda}_m = \lim_{\lambda \rightarrow 1} \sum_{k=1}^m \lambda_k$. Using this notation, we have that as $\lambda \rightarrow 1$,

$$\frac{\bar{d}_{m+1}}{\bar{d}_1} \rightarrow \frac{\tilde{S} - \tilde{S}_m + \delta_{m+1} \tilde{\lambda}_m}{\tilde{S}/\delta_{m+1} - \tilde{S}_m/\delta_{m+1} + \tilde{\lambda}_m} = \delta_{m+1} \quad (38)$$

which completes the proof.