



Stateless RD Network Services

Maxim Podlesny, Sergey Gorinsky

► **To cite this version:**

Maxim Podlesny, Sergey Gorinsky. Stateless RD Network Services. Mark Crovella; Laura Marie Feeney; Dan Rubenstein; S. V. Raghavan. 9th International IFIP TC 6 Networking Conference (NETWORKING), May 2010, Chennai, India. Springer, Lecture Notes in Computer Science, LNCS-6091, pp.135-147, 2010, NETWORKING 2010. .

HAL Id: hal-01056317

<https://hal.inria.fr/hal-01056317>

Submitted on 18 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Stateless RD Network Services

Maxim Podlesny¹ and Sergey Gorinsky²

¹ University of Calgary, Calgary, Canada

² Madrid Institute for Advanced Studies in Networks (IMDEA Networks),
Madrid, Spain

mpodlesn@ucalgary.ca, sergey.gorinsky@imdea.org

Abstract. Rate-Delay (RD) Network Services constitute a promising differentiated-services architecture for multi-provider networks, by offering users a choice between high throughput or low queuing delay at bottleneck links. An RD router provides service differentiation via transmission scheduling and by managing two FIFO queues. To ensure strict delay bounds, an RD router tracks arrival times of packets in the D service queue, and discards late packets at the queue head. However, maintaining the per-packet state is undesirable for complexity and cost reasons. In this paper, we present a Stateless RD (S-RD) router design that provides low queuing delay to the D service exclusively via buffer dimensioning, without requiring any per-packet state. After proving analytically that the S-RD design meets the delay guarantees, we use simulation to evaluate the performance of the stateless design, confirming that S-RD routers preserve the delay bounds of RD network services. As case studies, we consider Voice-over-IP (VoIP) and Web browsing as particular examples of Internet applications. The extensive simulation results demonstrate that S-RD Network Services significantly improve VoIP quality and increase the goodput of short-lived Web flows, without degrading the throughput of long-lived flows.

Keywords: Quality of service, resource allocation, Internet applications

1 Introduction

The RD (Rate-Delay) Network Services [1] constitute a promising recent design for service differentiation in multi-provider network environments. In particular, the D (Delay) service assures low queuing delay at bottleneck links; it is suitable for Internet telephony and other applications that require low end-to-end packet delays. While delay is the most important consideration for delay-sensitive applications, the loss rate tends to be higher for the D service. In turn, the higher loss rate increases the occurrence of TCP (Transmission Control Protocol) [2] retransmission timeouts, which can negatively affect the performance of web and other TCP-based applications.

In this paper, we propose a modification to the RD router design. The original design involves recording packet arrival times in order to schedule transmissions and guarantee low queuing delay. This state information requires additional

memory and processing, which increases the router cost (e.g., expensive Static Random Access Memory (SRAM) [3]), makes the router more complex, and can affect router performance. Our modification avoids the per-packet state information required for arrival-time tracking. Our theoretical analysis and simulations confirm that the proposed S-RD (Stateless RD) Network Services still assure bounded queuing delays.

To understand the effects on application-perceived performance, we perform two case studies, in which we assess the performance of typical Internet applications (VoIP and web browsing) with S-RD Network Services. Our extensive simulations examine a wide variety of network topologies and traffic scenarios. The simulation results show that S-RD Network Services significantly improve the quality of Internet telephony. More surprisingly, web browsing also benefits from the S-RD Network Services.

The rest of the paper is organized as follows. Section 2 describes the S-RD design. In Section 3, we present the theoretical analysis of this stateless version. Section 4 evaluates the S-RD Network Services using ns-2 simulations. Sections 5 and 6 report our assessment of application-perceived performance for VoIP and web browsing, respectively. Finally, Section 7 concludes the paper with a summary of its contributions.

2 RD Network Services

2.1 Overview

The key idea in RD Network Services is to separate delay-oriented traffic and throughput-oriented traffic into two classes and serve them using separate queues. Class D provides low queueing delay, while class R provides high(er) per-flow throughput. The two main parameters of the RD Network Services are the delay constraint d , which is the maximum queueing delay for flows from class D, and the desired ratio k between per-flow rates for classes D and R.

The queues are scheduled such that the ratio of traffic volume serviced from the D and R queues is maintained close to $\alpha = \frac{n_D}{kn_R}$, where n_D and n_R represent the number of class D and class R flows, respectively. For a link of capacity C , the effective service rates for the D and R queues are:

$$R_D = \frac{n_D C}{n_D + kn_R}, \quad R_R = \frac{kn_R C}{n_D + kn_R} \quad (1)$$

We will refer to the traffic sent from the R and D queues during the recalculation period as L_D and L_R , respectively. The size of buffer for the D queue is configured so that the draining time of the D queue is close to d . Thus, the control rules employed to allocate the buffers for each class are the following:

$$B_D = \frac{n_D C d}{n_D + kn_R}, \quad B_R = \min \left\{ B_{max}; B - \frac{n_D C d}{n_D + kn_R} \right\} \quad (2)$$

where B is the total buffer dedicated to a link, B_{max} is the size of a buffer required for effective support of throughput-greedy TCP traffic [4].

Routers periodically reallocate the buffers to queues for the D and R queues according to Equation (2). The numbers of flows in the classes are estimated through the time-stamp vector algorithm [5]. Packets from each queue are served in a FIFO (First-In First-Out) manner for scheduling the departures of packets. The selection of a queue to be served is based on the values of L_D and L_R since the last reset of L_D , L_R . If L_D is no more than αL_R , then the D queue is chosen for transmission. Otherwise, the next transmission is from the R queue. To enforce the delay constraint, the router tracks the arrival time of each packet in the D queue, and drops a packet from the head of queue D if the queueing delay of the packet exceeds d . In the event of queue overflow, the DropTail policy is used for dropping packets from either class.

2.2 Stateless Design

In the stateless version of the RD Network Services, we remove the requirement to explicitly track packet arrival times. Instead, we carefully limit the buffer size for class D traffic to ensure that queueing delay does not exceed d . In particular, we specify the size of the D buffer as follows:

$$B_D = \lfloor (d - w)R_D \rfloor^+, \quad w = \frac{2}{C} \left(\frac{S_D^{max}}{\alpha} + S_R^{max} \right) \quad (3)$$

where S_D^{max} and S_R^{max} are the maximum sizes of packets from classes D and R, respectively, and w is a delay adjustment. Through the subsequent analysis, we derive the foregoing expression for w , and prove that configuring the buffer size of queue D of the S-RD link via Equation (3) provides strict support for the delay constraint.

3 Analysis

In this section, we formally analyze the worst-case delay bounds for a backlogged S-RD router. There are in fact two different cases: one queue is backlogged, or both queues are backlogged. For space reasons, we ignore the simple case³ of a single backlogged queue, and examine the more interesting case when both the D and R queues are backlogged.

The dual-queue backlog scenario also has two cases, based on the S-RD buffer configuration. The first (trivial) one corresponds to a buffer of size zero, which causes no queuing delay. The second case involves a non-zero buffering delay, i.e., $d - w > 0$, where d is the delay constraint.

Our goal is to consider the latter case, and derive the minimum value of w required for supporting queuing constraint if $B_D = \lfloor (d - w)R_D \rfloor^+$. Let us consider an arbitrary packet p from the D queue. We assume that p arrives at

³ For this case, it suffices to add one more counter, beyond L_D and L_R , to indicate the traffic that must depart from the D queue in order to avoid exceeding the delay constraint.

the D queue at time t_a and departs from the D queue at time t_d . Suppose that at times t_a, t_d :

$$\frac{L_D(t_a)}{L_R(t_a)} = \alpha + \delta(t_a), \quad \frac{L_D(t_d)}{L_R(t_d)} = \alpha + \delta(t_d) \quad (4)$$

where $L_R(t_a) > 0$ and $L_R(t_d) > 0$. Thus, if $\delta(t) < 0$ then the packet sent at time t is from the D queue, otherwise - from the R queue. We consider the scenario where both the D and R queues are backlogged during time period $[t_a; t_d]$. We will refer to the traffic sent from the D and R queues during time period $[t_a; t_d]$ as ΔL_D and ΔL_R , $\Delta L_R > 0$, respectively. We prove the following theorem:

Theorem 1. *For any packet p in any traffic pattern, the maximum queuing delay is $d-w$ provided that:*

$$\frac{\Delta L_D}{\Delta L_R} \geq \alpha \quad (5)$$

Proof. Indeed, if inequality (5) holds, then $\frac{R'_D}{R'_R} \geq \alpha$, where R'_D, R'_R are actual serving rates for the D and R queues during $[t_a; t_d]$, respectively, and $R'_D \geq R_D$. Since $B_D = (d-w)R_D$, the maximum packet delay does not exceed $d-w$. ■

Since $\Delta L_D = L_D(t_d) - L_D(t_a)$, $\Delta L_R = L_R(t_d) - L_R(t_a)$, we can rewrite inequality (5) as follows:

$$\frac{L_D(t_d) - L_D(t_a)}{L_R(t_d) - L_R(t_a)} \geq \alpha \quad (6)$$

Let us denote the left side of inequality (6) as γ . Then, using expressions in Equation (4) and performing a simple transformation, we establish that:

$$\gamma = \alpha + \left(\delta(t_d) + \frac{L_R(t_a)}{\Delta L_R} (\delta(t_d) - \delta(t_a)) \right) \quad (7)$$

Therefore, inequality (5) holds if and only if:

$$\delta(t_d) + \frac{L_R(t_a)}{\Delta L_R} (\delta(t_d) - \delta(t_a)) \geq 0 \quad (8)$$

Let us now prove the following:

Theorem 2. *For any packet p in any traffic pattern, $\frac{\Delta L_D}{\Delta L_R} \geq \alpha$ if and only if:*

$$\delta(t_a) \leq 0, \quad \delta(t_d) \geq 0 \quad (9)$$

Proof. First, let us prove this is a sufficient condition. Indeed, if $\delta(t_a) \leq 0$, $\delta(t_d) \geq 0$, then inequality (8) holds for any values of $L_R(t_a)$ and ΔL_R , i.e., for any traffic pattern and any packet p . Second, let us prove it is a required condition. Suppose that it is not true. We need to consider three possible cases:

Case 1: $\delta(t_a) \leq 0$, $\delta(t_d) < 0$. Then from inequality (8) we have that:

$$\frac{\Delta L_R}{L_R(t_a)} + 1 \leq \frac{\delta(t_a)}{\delta(t_d)} \quad (10)$$




Fig. 1. Schedule of packet departures when $\delta(t_a) > 0$, $\delta(t_d) \geq 0$.

Since the left side of inequality (10) exceeds 1, and there exist traffic patterns and packets p such that the right side is less than 1, we have a contradiction.

Case 2: $\delta(t_a) > 0$, $\delta(t_d) \geq 0$. Then from inequality (8) we derive that:

$$\frac{\Delta L_R}{L_R(t_a)} + 1 \geq \frac{\delta(t_a)}{\delta(t_d)} \quad (11)$$

Since there exist traffic patterns and packets p such that the left side of inequality (11) is smaller than 2, whereas the right side of inequality (11) exceeds 2, we have a contradiction.

Case 3: $\delta(t_a) > 0$, $\delta(t_d) < 0$. Inequality (8) leads us to:

$$\frac{\Delta L_R}{L_R(t_a)} + 1 \leq \frac{\delta(t_a)}{\delta(t_d)} \quad (12)$$

Since the left part of inequality (12) is positive, and its right side is negative, we have a contradiction. Thus, we have shown that our assumption cannot be true, which means that (9) is a required condition. ■

From Theorem 1 and Theorem 2 we conclude that (9) expresses a sufficient condition for supporting queueing delay of at most $d - w$ for any packet with an arbitrary traffic pattern at the S-RD link.

Theorem 3. *For any packet p in any traffic pattern, the maximum queueing delay is $d - w$ if and only if $\frac{\Delta L_D}{\Delta L_R} \geq \alpha$.*

Proof. The sufficiency of this condition follows from Theorem 1. Let us now prove the necessity. Let us consider an arriving packet p that completely fills the buffer of the D queue, i.e., the enqueueing of that packet causes $q_D = B_D$, where q_D is the size of queue D. Indeed, if $\frac{\Delta L_D}{\Delta L_R} < \alpha$, then $\frac{R'_D}{R'_R} < \alpha$, and $R'_D < R_D$. From the fact that $B_D = (d - w)R_D$ we conclude that the maximum packet delay of a packet p exceeds $d - w$. ■

Let us now consider all possible cases when the packet delay exceeds $d - w$ for a packet p .

Case 1: $\delta(t_a) > 0$, $\delta(t_d) \geq 0$. In Figure 1, we show the schedule of packet departures in the considered case. According to Theorem 2 and Theorem 3, if packet p arrived at time t_1 and departed at t_d , then its queueing delay would not exceed $d - w$, since $\delta(t_1) < 0$, $\delta(t_d) \geq 0$. In the interval $[t_a; t_1]$, there is no potential arrival time t'_a of packet p at which $\delta(t'_a) < 0$, so the queueing delay of




Fig. 2. Schedule of packet departures when $\delta(t_a) \leq 0$, $\delta(t_d) < 0$.

packet p can exceed the delay constraint by the length of the interval $[t_a; t_1]$. We will refer to the length of that interval as w' , and to the amount of R traffic sent during this interval as X . Suppose that a D packet departing at time t_0 has size S_D , then the following inequalities must hold:

$$L_D(t_0) \leq \alpha L_R(t_0), \quad L_D(t_0) + S_D > \alpha L_R(t_0), \quad L_D(t_0) + S_D \leq \alpha L_R(t_0) + X\alpha \quad (13)$$

Lemma 1. *The value of X in the worst case scenario that satisfies the inequalities in (13) is:*

$$X = \frac{S_D^{max}}{\alpha} + S_R^{max}. \quad (14)$$

Proof. By inspection, it is clear that the X defined by equation (14) satisfies 13. Next, we need to show that there is no smaller solution. Let us suppose that there exists a smaller solution X' :

$$X' = X - \Delta X \quad (15)$$

where X is defined by equation (14), $\Delta X > 0$, $\Delta X < X$, ΔX is an integer, i.e., there exists ΔX such that X' satisfies inequalities in (13). Let us assume that the traffic scenario is such that the first inequality in (13) achieves equality. Then from the third inequality in (13) and $L_D(t_0) = \alpha L_R(t_0)$ we derive that:

$$S_D \leq S_D^{max} + \alpha S_R^{max} - \alpha \Delta X \quad (16)$$

Assuming that $S_D = S_D^{max}$ and $S_R^{max} < \Delta X$, we have that inequality (16) is not valid. Since there exists a traffic scenario such that the third inequality in (13) is not valid, we have a contradiction. Finally, we mention that S_R^{max} in (14) reflects that traffic is in packets, i.e., not fluid. ■

From Lemma 1 we conclude that the maximum queuing delay in excess of $d - w$ in the considered case is as follows:

$$w' = \frac{1}{C} \left(\frac{S_D^{max}}{\alpha} + S_R^{max} \right) \quad (17)$$

Case 2: $\delta(t_a) \leq 0$, $\delta(t_d) < 0$. In Figure 2 we demonstrate how the packets are scheduled for this case. If during time interval $[t_1; t_2]$, the link continued to serve the D queue up to packet p instead of packets from class R, then, according to Theorems 2 and 3, queuing delay of packet p would not exceed $d - w$, since

$\delta(t_a) < 0$, $\delta(t'_d) \geq 0$, where t'_d would be its departure time. Therefore, queuing delay of packet p can exceed the delay constraint by the length of the interval $[t_1; t_2]$. We refer to the length of that interval as w'' , and to the amount of D traffic sent during this time interval as Y . As in Case 1, Y is defined by the right side of equation (14). Therefore, w'' is the same as w' defined by equation (17).

Case 3: $\delta(t_a) > 0$, $\delta(t_d) < 0$. Since this case is a combination of the two previous ones, the maximum queuing delay in excess of $d - w$ is the sum of w' and w'' :

$$w = \frac{2}{C} \left(\frac{S_D^{max}}{\alpha} + S_R^{max} \right) \quad (18)$$

This expression completes the derivation of 3. Since we did not use the information that p fills the buffer of queue D while considering the three possible cases of exceeding the $d - w$ delay, we have in fact proved the following:

Theorem 4. *Sizing the D buffer according to Equation (3) ensures that the S-RD router algorithm supports maximum queuing delay d for class D.*

4 Simulation Evaluation of S-RD Network Services

In this section, we evaluate our modification to the original design of the RD Network Services through simulations using version 2.29 of ns-2 [6]. We run the experiments in a dumbbell topology where the bottleneck and access links have capacities 100 Mbps and 200 Mbps, respectively. The bottleneck link carries 100 long-lived D flows and 100 long-lived R flows in both directions and has propagation delay 50 ms. We choose propagation delays for the access links so that the propagation RTT (Round-Trip Time) for the flows is uniformly distributed between 104 ms and 300 ms. In addition, there is one web server and one web traffic receiver connected to the bottleneck link. The long-lived flows from the both classes join the network during the initial 1 s of an experiment. The size of web flows is described by the Pareto distribution with the average of 30 packets and shape index 1.3, and the web flows arrive at the network according to a Poisson process. All flows employ TCP NewReno [7] and data packets of size 1 KB. We configure link buffer to $B = B_{max} = C \cdot 250$ ms, where C is the capacity of the link. Every experiment lasts 60 s, and we repeat it five times for each of the considered parameter settings. We apply the same settings for the design and the measurements as in [1]. We average the utilization and loss rate over the whole experiment with exclusion of its first five seconds. We report the maximum packet delay observed over all runs of each experiment. All our simulation results show that the S-RD Network Services support the required throughput differentiation between the classes.

Capacity scalability. To explore the impact of the link capacity, we vary the bottleneck link speed from 10 Mbps to 1 Gbps. The arrival rate of the web-like flows is 50 flows per second (fps). Figure 3 shows that there is no violation of the delay constraint across the range of link speeds considered. In addition, the loss rate of class D shows only a marginal increase compared to the original RD design. The worst-case delay bound is tight for high link speeds, but not as tight

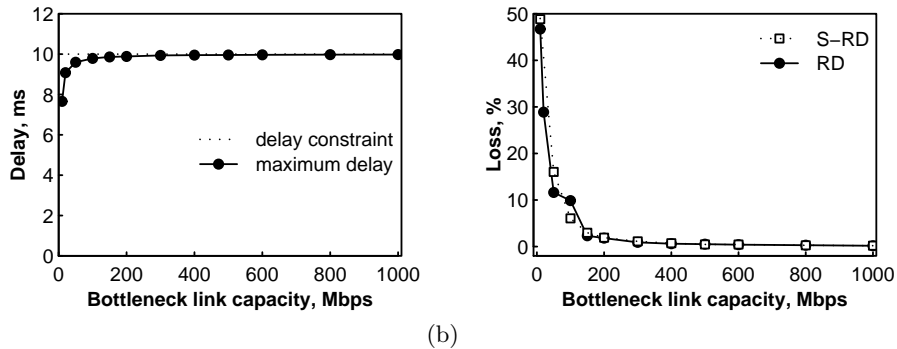


Fig. 3. Capacity scalability of the S-RD algorithm: (a) maximum queuing delay of class D; (b) comparison of the loss rate for class D to original version of the RD link.

Table 1. Categories of voice transmission quality


R-factor range	MOS	Quality category	User satisfaction
90 - 100	4.34 - 4.50	Best	Very satisfied
80 - 90	4.03 - 4.34	High	Satisfied
70 - 80	3.60 - 4.03	Medium	Some users dissatisfied
60 - 70	3.10 - 3.60	Low	Many users dissatisfied
50 - 60	2.58 - 3.10	Poor	Nearly all users dissatisfied

for lower link speeds. The maximum queuing delay observed for low bottleneck link capacities close to 10 Mbps is due to the increased significance of the term w in Equation (3) specifying the size of the D buffer.

We also studied the impact of the intensity of web-like traffic, the value of delay constraint, and the maximum packet sizes for classes D and R by varying the appropriate parameter. The results indicate that the delay constraint is strictly supported over the whole range of the varied parameter. For space reasons, we do not provide these results here.

5 VoIP Traffic Results

To evaluate the quality of the delivered service for VoIP, we use Mean Opinion Score (MOS) [8], a subjective score for voice quality ranging from 1 (Unacceptable) to 5 (Excellent). To estimate a MOS score through network characteristics, we employ the E-Model [9], which assesses VoIP quality by accounting for network characteristics like loss and delay. The E-Model uses the R-factor, which is computed as a function of all of the impairments occurring with the voice signal. The R-factor ranges from 0 to 100, with 100 being the best, which is reflected on Table 1.



(a) (b)
Fig. 4. Influence of the intensity of the web-like flows: (a) Average MOS; (b) average utilization of class R.

5.1 Evaluation Methodology

To generate VoIP traffic and perform measurements of voice quality, we use the tool developed in [10], an additional module of the network simulator ns-2. We use the same network topology as in Section 4 with the same traffic from R class and web-like traffic from both classes, but the bottleneck link delay is 10 ms. Instead of long-lived D flows, there are 100 VoIP flows with the same propagation RTTs of 150 ms. The value of d is 50 ms. Web flows arrive with the intensity of 50 fps. We perform five experiments for each settings, and each experiment lasts for 70 sec. To encode the speech, we employ AMR (Adaptive Multi-Rate) Audio Codec [11] operating at audio bit rate of 12.2 kbps. The parameters we measure are average MOS and the average utilization of class R. While measuring MOS, first ten seconds of the experiment are neglected. All flows join the network during the initial 1 s. We compare the performance of the S-RD Network Services with the performance of DropTail.

5.2 Simulation Results

Influence of the web-like traffic. To study the influence of the web-like flows, we change the intensity of the web-like flows from 1 fps to 150 fps. In Figure 4, we observe that the S-RD Network Services demonstrate better performance for VoIP over almost the whole range of the varied parameter, while the R traffic achieves the same bottleneck link utilization as the DropTail link.

Transient behavior. In this experiment, VoIP flows join the network during the whole experiment lasting for 600 s. There are 500 VoIP flows that start arriving from the beginning. The arrival process is Poisson with an average rate of 1 fps. Whereas the average MOS with the DropTail link is 2.97, MOS with the S-RD Network Services is 4.16. The utilization of class D is 84.45% and 83.85% with the S-RD Network Services and DropTail link, respectively. Thus, the S-RD Network Services deliver better service for VoIP in the considered dynamic scenario.

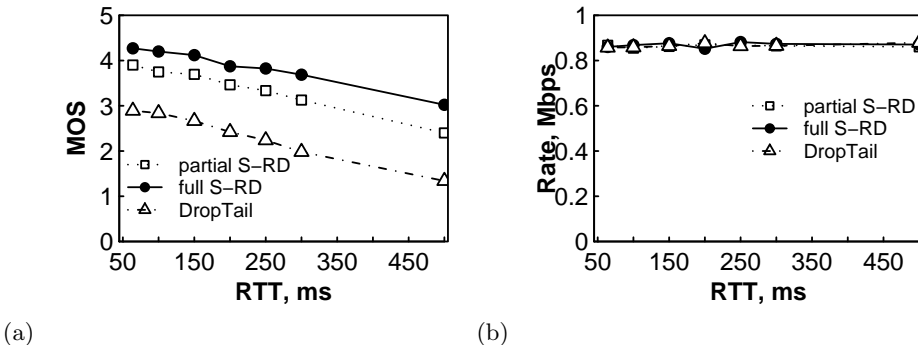



Fig. 5. Performance under the partial deployment for different propagation RTTs: (a) Average MOS; (b) average per-flow throughput of class R.

Impact of VoIP population size. To examine the scalability of the design concerning the population of VoIP flows, we vary the number of them from 100 to 500. The results show that the number of VoIP flows does not affect the quality of VoIP. In particular, MOS with the S-RD Network Services is between 4.14 and 4.16 whereas MOS with the DropTail link is between 3.01 and 3.13. The consistent performance of VoIP over the whole range of the varied parameter is because a VoIP flow requires a relatively small connection throughput. The R flows achieve 82-87% utilization of the bottleneck link.

Partial deployment. In this experiment, we explore the situation when VoIP flows use the service provided by two different ISPs. There is one bottleneck within each ISP, 50 VoIP flows, 50 R flows going through both the ISPs, and two groups of 50 R flows each traversing a single ISP. In particular, we can consider two deployment scenarios, which reflect the deployment of the S-RD Network Services by only one ISP and by both the ISPs. The propagation RTT of VoIP flows is varied between 64 ms and 500 ms. In Figure 5, we observe that even under the partial deployment of the S-RD Network Services, which is labeled as "partial" in the graph, VoIP flows get better service. Moreover, the full deployment of the design (labeled as "full") further improves the VoIP quality. More importantly, the improvements of VoIP quality do not affect the service delivered to the R class concerning the flow rates.

6 Web Traffic Results


To evaluate the performance of a web application, which generates flows with different sizes, we calculate the average goodput of the web-like flows as the average of the goodput of each web-like flow. The goodput of a web-like flow is the ratio between the flow size and its FCT (Flow Completion Time).



(a)

(b)

Fig. 6. Influence of the intensity of the web-like flows: (a) average utilization of the long-lived flows; (b) average goodput of the web-like flows.



(a)

(b)

Fig. 7. Influence of the number of the long-lived flows: (a) average utilization of the long-lived flows; (b) average goodput of the web-like flows.

6.1 Evaluation Methodology

In the experiments, we employ a dumbbell topology with the same experimental settings as in Section 5. To compare the S-RD Network Services design, we also run the experiments under the same settings for the DropTail link. There are 100 long-lived flows in the forward and reverse directions that are served as class R. The R flows join the network during the first second of an experiment. The value of d is 50 ms. In addition, there is one web server and one web traffic receiver connected to the bottleneck link. The web server generates flows with the same parameters as in Section 4, which are served as class D.

6.2 Simulation Results

Influence of the web-like traffic. We study the influence of the intensity of the web-like flows by varying their arrival rate in the interval between 5 fps and




Fig. 8. Influence of propagation RTT of the web-like flows: (a) average utilization of the long-lived flows; (b) average goodput of the web-like flows.

400 fps. In Figure 6, we see that S-RD improves the performance of the web-like flows over the whole range of the varied parameter. On the other hand, the performance of the long-lived flows deteriorates when the intensity of the web-like flows exceeds 100 flows per second. The decreased performance of web-like flows beyond intensities of 250 fps is attributed to the increased packet loss rate.

Influence of the long-lived flows. To explore the population scalability, we vary the number of long-lived flows between 50 and 600. The intensity of the web-like flows is 100 fps. Figure 7 shows that the S-RD Network Services provide consistently better performance for the web-like flows over the whole range of the varied parameter, whereas the long-lived flows have the same goodput with the S-RD Network Services and DropTail link. In particular, the former improves the goodput of the web-like flows by 50%-200%.

Influence of propagation RTT of the web-like traffic. In this experiments, we vary the propagation RTT of the web-like flows in the range between 30 ms and 500 ms. The number of long-lived flows is 100. In Figure 8, we observe that the S-RD Network Services significantly improve the goodput of the web-like flows for small RTTs. Besides, the performance of the S-RD Network Services and DropTail link for the long-lived flows is similar except for RTTs less than 50 ms, for which the S-RD Network Services show slightly lower bottleneck link utilization than the DropTail scheme.

7 Conclusion

In this paper, we proposed a modification to the original router design of the RD Network Services. The new simpler design does not track packet arrival times. Our theoretical analysis and simulations showed that the proposed S-RD Network Services guarantee the low bounded queuing delay of the original design. As case studies, we explored the application-perceived performance of VoIP and web browsing with the S-RD Network Services. Through the extensive simulations, we determined that the architecture improves the application performance.

Acknowledgement

Financial support for the work was provided in part by iCORE (Informatics Circle of Research Excellence) in the Province of Alberta. The authors are thankful to Carey Williamson for his invaluable comments about this paper.

References

1. Podlesny, M., Gorinsky, S.: RD Network Services: Differentiation through Performance Incentives. In: Proceedings ACM SIGCOMM 2008. (August 2008)
2. Postel, J.: Transmission Control Protocol - DARPA Internet Program Protocol Specification. RFC 793 (September 1981)
3. Iyer, S., Kompella, R., McKeown, N.: Designing Packet Buffers for Router Line Cards. Technical Report, TR-02-HPNG-031001 (October 2002)
4. Villamizar, C., Song, C.: High performance TCP in ANSNET. ACM SIGCOMM Computer Communication Review **24**(5) (October 1994) 45–60
5. Kim, H., O'Hallaron, D.: Counting Network Flows in Real Time. In: Proceedings IEEE GLOBECOM 2003. (December 2003)
6. McCanne, S., Floyd, S.: ns Network Simulator. <http://www.isi.edu/nsnam/ns/>.
7. Floyd, S., Henderson, T.: The NewReno Modification to TCP's Fast Recovery Algorithm. RFC 2582 (April 1999)
8. ITU-T: Methods for subjective determination of transmission quality. Recommendation P.800 (August 1996)
9. Bergstra, J., Middelburg, C.: The E-model, a Computational Model for Use in Transmission Planning. ITU-T Recommendation G.107 (June 2006)
10. Bacioccola, A., Cicconetti, C., Stea, G.: User-level Performance Evaluation of VoIP Using ns-2. In: Proceedings NSTools 2007. (October 2007)
11. Sjoberg, J., Westerlund, M., Lakaniemi, A., Xie, Q.: Real-Time Transport Protocol (RTP) Payload Format and File Storage Format for the Adaptive Multi-Rate (AMR) and Adaptive Multi-Rate Wideband (AMR-WB) Audio Codecs. IETF RFC 3267 (June 2002)