Impact of marking strategy on aggregated flows in a differentiated services network *

Ikjun Yeom
A. L. Narasimha Reddy
Dept. of Elec. Engg.
Texas A & M University
College Station, TX 77843-3128
{ikjun,reddy}@ee.tamu.edu

Abstract

Diff-serv framework is receiving wide attention as an architecture for implementing service guarantees in the Internet. Current framework allows customers to mark their packets and the network provider to check them for conformance to service contracts. This paper looks at the impact of marking strategies employed by aggregated sources (customers) on the provided service in a diff-serv network. The paper shows that aggregation reduces the impact of differences in round-trip-times on throughput guarantees. The paper also shows that proportional marking of packets can lead to large differences in delivered throughputs to individual flows within the aggregation. The paper proposes two new marking algorithms that improve the fairness among the individual flows within an aggregation. The paper also shows that a customer can achieve higher throughput by marking a disproportionate number of packets IN along the paths that are experiencing congestion while remaining within the service contract. The impact of such a strategy on the network provider and the individual flows is studied.

Keywords : Differentiated service, Aggregation, Quality of service.

1 Introduction

Current Internet provides best-effort service to end-users. Recently, many multimedia applications have been developed. Multimedia applications are sensitive to available bandwidth and delay experienced in the network. To satisfy these requirements, differentiated service (diff-serv) framwork for the Internet has been proposed.

Diff-serv framework is a proposal to provide service guarantees over networks by providing different drop preferences [1, 2, 3]. In this framework, the routers at the edge of the network monitor and mark packets of flows (individual or aggregated). The packets of a flow that obey the service profile are marked IN (in profile) and the packets that are beyond the service profile are marked OUT (out-of-profile). The network gives preference to IN packets while dropping OUT packets disproportionately at the time of congestion. The router doesn't distinguish between packets of individual flows and can use FIFO style scheduling mechanisms. This preferential drop mechanism is expected to provide better throughput for IN packets than OUT packets. With appropriate network provisioning, it is expected that this could result in bandwidth guarantees. Diff-serv framework allows aggregated sources as well as individual sources. This paper looks at the QoS and bandwidth management issues at an aggregated source.

When several individual sources are aggregated, output traffic of the aggregation is not like traffic of one big single source in the following respects: (a) Each source within the aggregation responds to congestion individually. (b) The aggregation has mutiple destinations, and each source within the aggregation experiences different delays and congestion. Therefore, when we deal with aggregated sources in diff-serv networks, we need to consider not only the throughput achieved by aggregated sources but also the throughput achieved by individual flows.

We will assume that a customer with an aggregated source employs his/her own marker to manage individual flows within the aggregation. The network provider may monitor and remark packets to ensure compliance of the contract. In this paper, we study how a customer can mark packets within an aggregation to achieve specific performance goals while staying within the contract-profile. We will also discuss

*This work was supported in part by a Texas ATP grant and by an NSF Career Award

the impact of such strategies on the network provider.

The rest of the paper is organized as follows. Section 2 presents simple simulations with aggregation to motivate the rest of the paper. In Section 3, we propose new marking algorithms to meet performance goals of individual flows. Section 4 presents simulations with the proposed algorithms. In Section 5, we discuss the simulation results and present related work. Section 6 concludes the paper and points to future work.

2 Effect of aggregated sources

In this section, we discuss effect of aggregated sources on other flows in diff-serv networks. To study the effect of aggregation, we conducted simple simulations using ns-2 [6].

Figure 1 shows the simulated network topology. There are five aggregated sources and five single sources, and each aggregated source consists of three individual sources. We used a TCP/Reno agent for each individual source. The marker used time sliding marking algorithm proposed in [3]. The router uses RED parameters 20/40/0.5 for the OUT packets and 50/100/0.02 for the IN packets. The bottleneck bandwidth is set to 9 Mbps. Four aggregated sources and four single sources reserve 2 Mbps, 1 Mbps, 0.5 Mbps, and 0.1 Mbps. One aggregated source and one single source are best-effort. Since a total of 7.2 Mbps is reserved, the subscription level is 80%. We ran two simulation experiments using this configuration. In the first experiment, we used the same RTT, 60 ms for all the 20 flows. In the second experiment, we used RTTs of 40 ms, 60 ms, and 80 ms for the three sources of an aggregated pool, and an RTT of 60 ms for single sources. Flows within an aggregated source may have different RTTs as they may be talking to different hosts after passing through the same bottleneck link.

Figure 2 shows the results of the simulations. From Figure 2(a), we can observe that aggregated sources obtain more bandwidth than single source with the same reserved bandwidth. The three flows within an aggregated source claim three times as much of the shared excess bandwidth than a single source and hence the difference between the throughputs of aggregated sources and the individual sources. Throughput of aggregated sources with the same RTTs is not much different from the throughput of aggregated sources with different RTTs. This observation may indicate that aggregation may blunt the effect of RTT differences. We will show more results to this effect later in the paper.

Figure 2(b) shows the rates achieved by individual flows within an aggregation at different reservations.

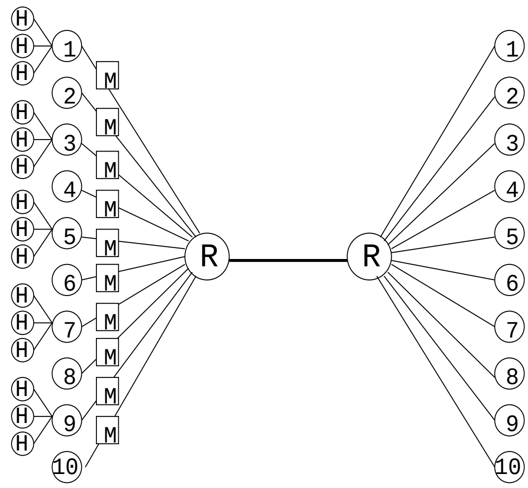

Figure 1: Simple topology for simulation with aggregated sources

As RTT is increased, throughput is decreased. This shows that different RTTs impact the throughputs of individual flows even though the impact on aggregation is reduced.

3 New marking strategy for aggregated sources

In Section 2, we have shown that the impact of different RTTs on achieving reserved aggregated bandwidth are reduced by aggregated sources. The total throughput of aggregated sources are maintained even if RTTs of individual sources are different. However, there is still unfairness in assigning the reserved bandwidth to individual users. Is it possible to fairly share the aggregated bandwidth independent of the RTTs of individual flows or the behavior of individual flows in the aggregation?

We consider maintaining state for each flow within an aggregation at the boundary router. Average sending rate of a flow is maintained as state information for each flow at the marker of the aggregated source. This information is used in balancing resources across the different flows within the aggregation.

If we apply the marking strategy proposed by [3] for aggregated sources, average sending rate of n aggregated sources at time t, $r_{agg}(t)$, is

$$r_{agg}(t) = \sum_{i=1}^{n} r_i(t) \qquad (1)$$

When a packet arrives and $r_{agg}(t)$ exceeds reserved rate, r_{resv}, at that instance, the packet is marked as OUT with probability, P_{OUT},

$$P_{OUT}(t) = (r_{agg}(t) - r_{resv})/r_{agg}(t) \qquad (2)$$

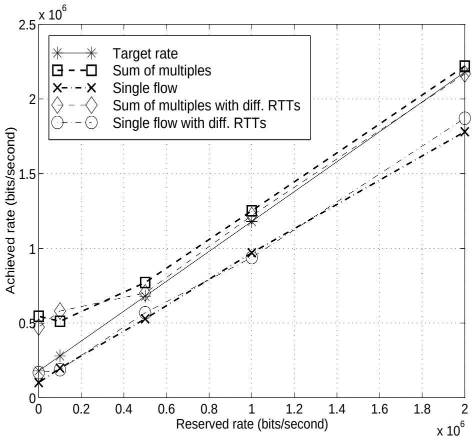

(a) Throughputs of aggregations and single sources

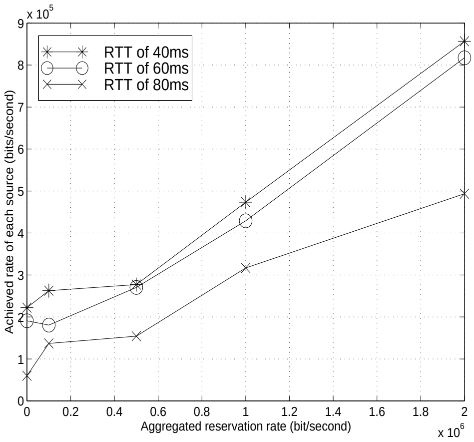

(b) Impact of different RTTs within aggregation

Figure 2: Effect of aggregated sources

Therefore, the number of OUT packets of each source, N_i, in interval, τ, is

$$N_i = \int_0^\tau \frac{\max(r_{agg}(t) - r_{resv}, 0)}{r_{agg}(t)} \times r_i(t)\ dt \quad (3)$$

Then, the number of OUT packets of aggregated sources, N_{agg} is

$$N_{agg} = \sum_{i=1}^{n} N_i \quad (4)$$

$$= \int_0^\tau \max(r_{agg}(t) - r_{resv}, 0)\ dt \quad (5)$$

From (3), a packet is determined to be marked OUT or IN not by the sending rate of its source but by the aggregated sending rate. Even if the sending rate of a source is less than the individual target rate, $r_{each} = r_{resv}/n$ (since RTT is longer than others or it experienced congestion), the packets from that flow are marked OUT with the same probability as others. We term this marking strategy *Proportional marking* in this paper. We will use fair sharing of aggregated bandwidth as a target of this study. It is easy to generalize this to allow different bandwidth allocations if necessary.

To remove this unfairness within aggregated sources, we propose a new marking algorithm. In designing the new marking algorithm, there are two conditions we need to consider: (1) The total number of OUT packets should be the same as in the original algorithm. (2) The packets from a flow with sending rate less than r_{each} should be marked OUT with a relatively smaller probability (to improve fairness). A third condition requires that any new marking policy do not introduce IN-marked packets that may later be marked OUT by network provider due to voilation of service contract. We assume that a sliding window marking algorithm [3] is used by the network for testing compliance.

Figure 3 shows the proposed algorithm. For each packet arrival, the algorithm calculates average rates of individual sources as well as the average rate of aggregation. If the average rate of aggregation is less than the reserved rate, it marks the packet as IN. Otherwise, it compares the average rate of the individual flow with the individual target rate. We study the following two choices for the individual target (IT),

$$IT = \begin{cases} r_{resv}/\text{No. of active flows} \\ \qquad \text{for equal sharing of IN packets,} \\ r_{agg}(t)/\text{No. of active flows} \\ \qquad \text{for achieving equal throughput.} \end{cases} \quad (6)$$

By setting the individual target as the first choice, every individual flow is expected to get a fair-share of the aggregated reserved bandwidth, and the total throughput (achieved by IN and OUT packets) of an individual source may vary due to variations in available excess bandwidth along the different paths of the flows. The second choice is for achieving equal throughput by assigning more IN packets to the flows going through congested links. If the average rate of an individual source is less than the individual target, the packets of that flow are marked IN at a higher rate.

From Figure 3, it is clear that the proposed algorithm satisfies the second condition of marking a higher fraction of packets IN when the source rate is less than r_{each}. From the algorithm, P_{OUT} is given by

$$P_{OUT}(t) = \begin{cases} (r_{agg}(t) - r_{resv})/r_{exceed}(t) & \text{if } r_i(t) > r_{each} \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

Here, r_{exceed} is the sum of r_i that is greater than r_{each}. The number of OUT packets of each source, N_i is

$$N_i \;=\; \int_0^{\tau} Pout(t) \times r_i(t)\; dt \quad (8)$$

Since P_{OUT} is zero for sources with r_i less than r_{each}, the number of OUT packets of aggregated sources, N_{agg} is

$$N_{agg} \;=\; \sum_{exceed} N_i \quad (9)$$

$$\;=\; \int_0^{\tau} \max(r_{agg}(t) \;-\; r_{resv}, 0)\; dt \quad (10)$$

From (5) and (10), it is seen that the proposed algorithm satisfies the first condition.

4 Simulation results

In this section, we present simulation results and show how the new algorithm achieves fair sharing of bandwidth within an aggregation. We will show that this bandwidth management results in improved throughput realization. We modified ns-2 [6] to implement the new marking algorithm. In all simulations, we used a TCP-Reno agent in ns-2 as a source and FTP application as a traffic generator.

4.1 Dealing with different RTTs

This section describes the impact of different RTTs within aggregated sources and how the new algorithm deals with this impact. Figure 4 shows a simple network topology used in this simulation. There are five markers and aggregated sources, each aggregated source consists of ten individual sources. Bandwidth of every link except the link between the two routers is 10 Mbps, and bandwidth of the link between two routers is limited to 6 Mbps. Each aggregated source reserves 1 Mbps. We assign RTT_j^i, RTT of i^{th} individual source in j^{th} aggregated source as;

$$RTT_j^i \;=\; 130 \;+\; 4 \times (i-1) \times (j-5.5)\; (\text{ms}) \quad (11)$$

This results in five aggregated sources with varying differences in RTTs. For example, aggregated source 1 has a (min RTT, max RTT) = (130 ms, 130 ms) compared to that of aggregated source 5 with (58 ms, 202 ms).

The router uses RED parameters 20/40/0.5 for OUT and 40/100/0.02 for IN packets. With this simulation setup, we conducted two simulation experiments. In the first experiment, we use the proportional marking algorithm, and in the second experiment, we use the new marking algorithm.

Figure 5 shows the achieved rates of individual sources. The horizontal line in the figure shows the individual target rate. In Figure 5(a), it is clear that there exists unfair bandwidth sharing within aggregated sources. This unfairness increases as the differences in RTTs are increased. For example, a source achieves 4.8 times the bandwidth of another source in the fifth aggregation. With the new marking algorithm, most sources reach their target rates as shown

```
Algorithm:

for each packet arrival
  calculate the new avg_agg_rate and avg_rate_i
  if avg_agg_rate <= target
    mark this packet IN
  else
    if avg_rate_i > individual_target
      calculate sum_exceeds
      With probability Pout=(avg_agg_rate-target)
                                   /sum_exceeds
        mark this packet OUT
      else
        mark this packet IN
    else
      mark this packet IN
end

Fixed parameter:
  target: reserved rate for aggregated sources

Saved variables:
  avg_agg_rate: average rate of aggregated sources
  avg_rate_i: average rate of ith individual source
  individual_target: target/the number of active flows
  sum_exceeds: the sum of avg_rate_i which is greater
            than individual_target

Other:
  Pout: current packet OUT-marking probability
```

Figure 3: New marking algorithm for aggregated sources

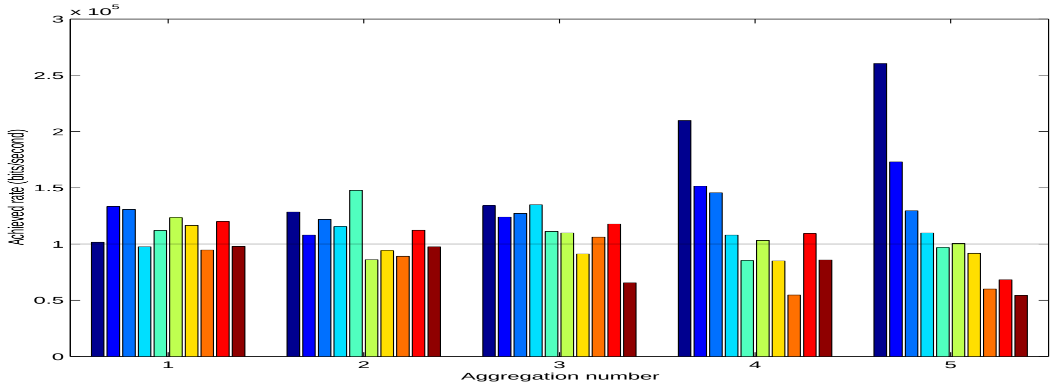

(a) Proportional marking

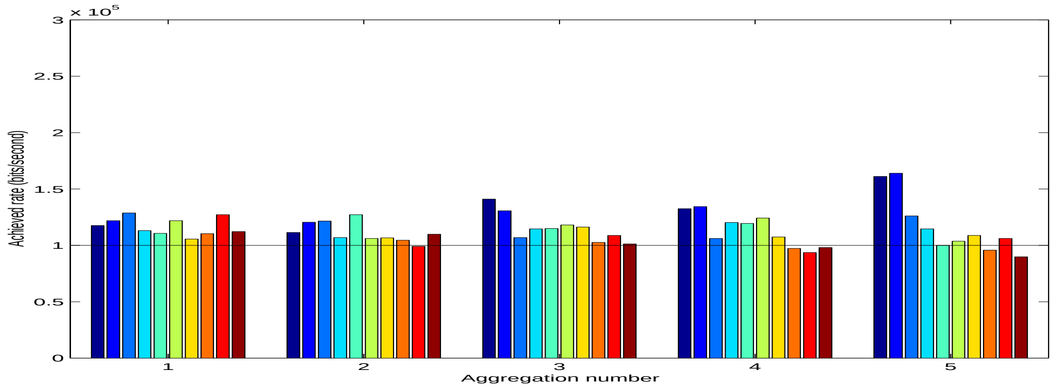

(b) New marking

Figure 5: Dealing with different RTTs within aggregated sources

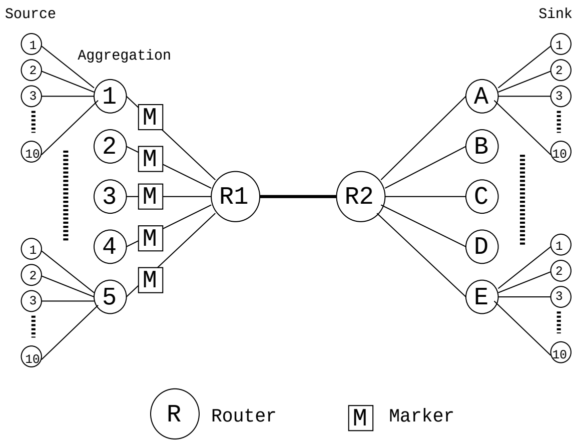

Figure 4: Aggregated network topology with different RTTs

in Figure 5(b). The bandwidth is shared more fairly among the different sources.

To compare the results quantitatively, we present Table 1. The average throughput of each aggregation is not much different from each other in both schemes even though RTTs of individual sources are different. This again shows that the impact of different RTTs can be reduced by aggregation. The row STD shows the standard deviation among the individual rates within an aggregation. It is observed that STD increases significantly with increased RTT differences within an aggregation. The new marking algorithm achieves significantly smaller variation compared to proportional marking. The row Max/Min compares the maximum and minimum rates realized within an aggregation. Again, it is observed that fairness is considerably improved with the new marking algorithm.

Table 1: Quantitative comparisons

Algorithm	Proportional marking					New marking				
Aggregation	1	2	3	4	5	1	2	3	4	5
RTT Max/Min	1	1.32	1.77	2.42	3.48	1	1.32	1.77	2.42	3.48
Average(Kbps)	112.7	110.0	112.2	113.8	114.5	116.9	111.4	115.5	113.3	117.0
STD (Kbps)	14.3	19.3	21.2	44.4	62.1	7.7	8.8	12.3	14.8	25.3
Max/Min	1.41	1.72	2.06	3.83	4.80	1.22	1.28	1.39	1.43	1.82

4.2 Aggressive bandwidth management in congested networks

Results from the earlier section showed that a customer can mark packets effectively to improve fairness while staying within the contract-profile. Can we extend this idea further to allocate different amounts of IN-profile bandwidth based on the congestion currently being experienced by the flows? This, in effect, is the same problem as before, but in a different context: Can we effectively manage IN-profile bandwidth to improve performance of specific flows (with longer RTTs or those experiencing congestion) within an aggregation? However, as we will point out later, reallocating IN-profile bandwidth based on congestion may have important consequences on the network providers.

Figure 6 shows the network topology used for simulations. There are two aggregated sources, and each aggregated source consists of ten individual sources. Each aggregated source reserves 1 Mbps. The network consists of a 1 Mbps link between the router and node 'A' and 2 Mbps link between the router and node 'B'. Individual sources 1∼8 of each aggregated source are connected to node 'A' and individual sources 9 and 10 are connected to node 'B'. In this topology, therefore, the link between the router and node 'A' is 160% subscribed, and the link between the router and node 'B' is 20% subscribed if we assume that each individual source expects to get 0.1 Mbps ($=r_{resv}$/number of individual sources). The network as a total has enough capacity (3 Mbps) to support the two aggregated sources (total reservation of 2 Mbps). Due to the dynamic nature of flows, one of the links may be oversubscribed as in this example. Each individual source has same RTT as 40 msec. With this simulation topology, we conducted two simulation experiments. In the first simulation, we applied proportional marking to the marker for aggregation 1 and the new algorithm 1 (for IN packet equal sharing) for aggregation 2. In the second simulation, we applied proportional marking for aggregation 1 and the new algorithm 2 (for achieving equal throughput) for aggregation 2.

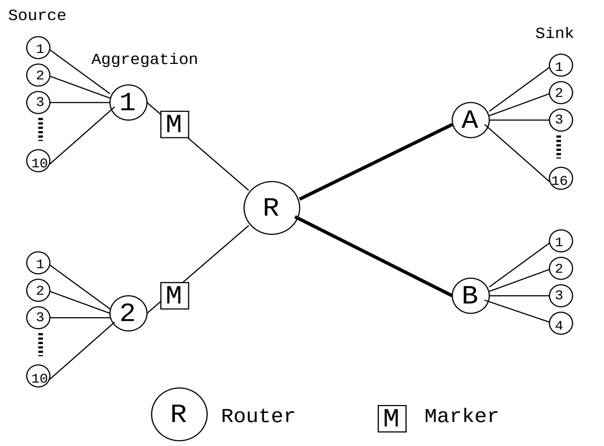

Figure 6: Aggregated network topology with different congestion levels

Figure 7 shows the results of Simulation 1. Each bar shows the achieved throughput of individual source, and the dark portion in each bar indicates the throughput achieved by IN packets. In Figure 7(a), it is observed that each packet is marked OUT with the same probability even if its source cannot reach its reserved rate, and IN packets are unfairly distributed to flows achieving higher rates. In Figure 7(b), however, each individual source achieves a fair-share of the reserved rate and shares IN packet throughput equally even with different levels of congestion. Clearly, fair sharing of IN-profile bandwidth improved the performance of the flows through the congested link. Both the aggregated sources stay within the contract-profile, but the second source achieved a higher bandwidth through the congested link than the first source. This is a result of managing IN-profile bandwidth effectively by distributing it fairly among the individual sources. When source 1 and source 2 compete for bandwidth on the congested link, source 2 achieves higher share due to marking higher number of packets IN.

Figure 8 shows the results of Simulation 2. The algorithm 2 is more aggressive by the fact that it sends

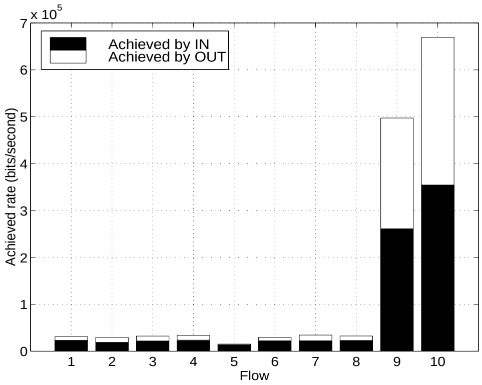

(a) Proportional marking

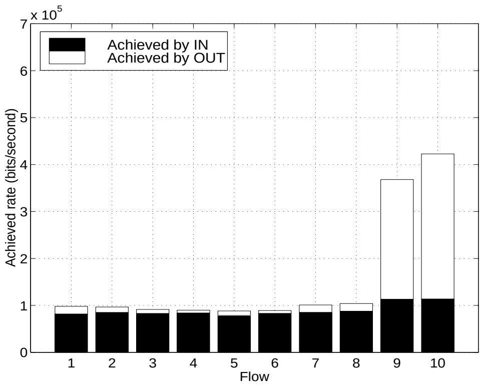

(b) Algorithm 1

Figure 7: Throughput comparisons with proportional marking and the algorithm 1

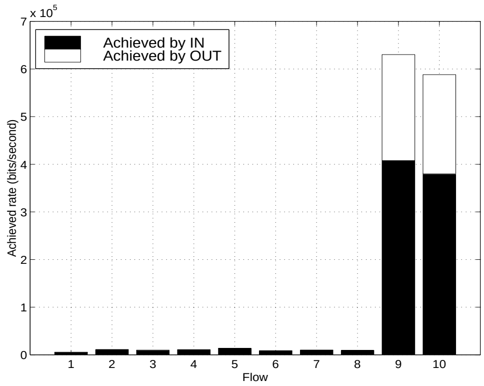

(a) Proportional marking

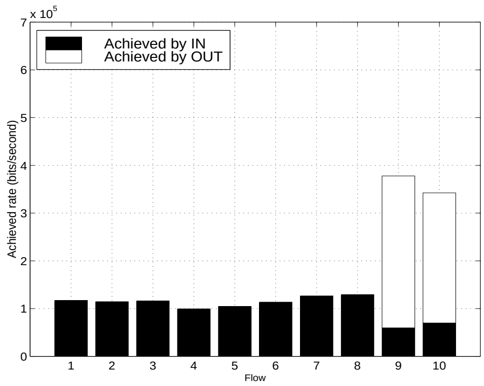

(b) Algorithm 2

Figure 8: Throughput comparison of proportional marking and the algorithm 2

more IN packets on congested links than on uncongested links so as to get more bandwidth in congested links. Again, aggregated source 1 used proportional marking and aggregated source 2 used the algorithm 2. In Figure 8(a), the flows through congested link loose more bandwidth than the flows in Figure 7(a). The flows using the algorithm 2 in Figure 8(b) get more bandwidth than the flows in Figure 7(b).

The algorithm 2 is more aggressive than the algorithm 1 in trying to meet the performance goals. The goal here is to achieve 0.1 Mbps for each individual source while staying within the contract-profile. As can be seen from Figure 8(b), the new marking algorithm allocates more IN-profile bandwidth to sources observing congestion than the ones that are not experiencing congestion. As a result, these sources claim a larger share of the congested link bandwidth, exceeding the individual targets of 0.1 Mbps. The flows within aggregated source 1 achieve significantly less bandwidth due to proportional marking. These two experiments show that individual marking strategies employed by customers can impact each other even when every source stays within the contract-profile.

4.3 Scalability

From the previous experimental results, it is clear that the proposed marking algorithms reduce the impact of differences in RTTs among the flows within an aggregation and achieve better throughput with the same reserved rate in congested links. However, this improvement in throughput is caused not by improvement of network performance nor increase of resources but by aggressive behavior of sending more IN packets through the congested links. How well do these algorithms and allocated benefits scale if every source employes the proposed algorithm?

To observe the scalability of these strategies, we conducted two sets of simulation experiments. In the first set, we used the same simulation topology as in Figure 6. In the second set, we added one best-effort aggregated source of ten individual sources. Each set consists of three experiments. For both the aggregated sources, we applied proportional marking in the first experiment, the algorithm 1 in the second experiment, and the algorithm 2 in the third experiment in each set.

The results are presented in Figure 9. In Figure 9(a), the total throughput of each experiment is similiar to the throughputs of the others since network resources are limited even though the proposed algorithms send more IN packets. Only difference between proportional marking and the proposed algorithms is that throughput of proportional marking strategy in the congested link is achieved by both IN and OUT packets while throughputs of the proposed algorithms in the congested link are achieved only by IN packets. This results in different performance in the experiments with best-effort flows. In Figure 9(b), IN packet throughput of each flow is protected even in the presence of best-effort flows, but the OUT packet throughput is reduced by best-effort flows. As a result, the new marking algorithms achieve higher throughputs through the congested link.

4.4 Impact on Network

Figure 10 shows the instantaneous queue length of the router on the congested link. The graph marked 'Exp.1' shows the queue length when both the aggregated sources use proportional marking scheme, and the graph marked 'Exp.2' shows the queue length when both the aggregated sources use Algorithm 1. The graph marked 'Combined' shows queue length when one aggregated source uses proportional marking and the other uses Algorithm 1. The graph, 'Exp.1' shows that queue length is maintained between 20 and 40 packets after slow-start since maximum and minimum thresholds for OUT packets are 20 and 40 packets. The queue length in Exp.2 is much higher than in Exp.1 because most packets sent to congested link are marked IN. Since IN packets are configured with the RED (min, max) thresholds of (40, 100), the packets observe longer queue lengths (and delays) at the router. In the combined experiment, queue lengths again approach Exp.1 after a burst of losses around 12 seconds. Since the drop thresholds are configured differently for IN and OUT packets, the new marking algorithm resulted in larger queuing delays through the congested path in the network. It is noted that instantaneous queue lengths may exceed the RED thresholds since RED uses average queue lengths.

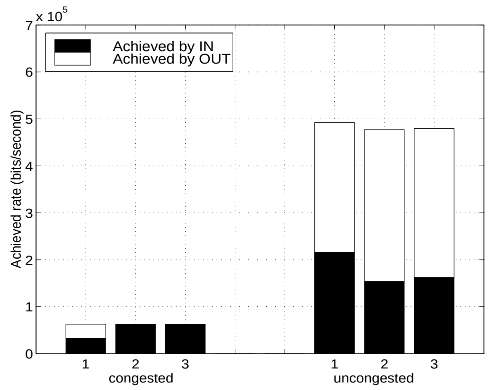

(a) Without best-effort flows

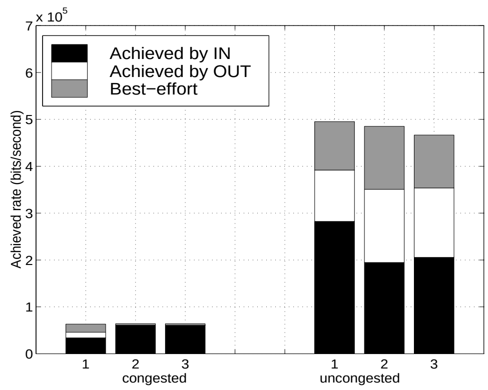

(b) With best-effort flows

Figure 9: Impact of marking algorithms on throughput

4.5 Impact on other flows at transient time

In this section, we consider the situation when users switch their packet marking schemes for achieving desired throughput and discuss the impact of marking schemes on the individual flows at transient time. We did a simple simulation with topology described in Figure 6. Aggregated source 1 uses proportional marking scheme during the whole simulated period, while aggregated source 2 switches from proportional marking to Algorithm 1 at the simulated time of 30 seconds.

Traces of congestion windows (cwnds) of two individual sources sending packets on the congested link within each aggregated source are presented in Figure 11. During the interval from 0 to 30 seconds, cwnds of two sources are not much different from each other. During 30-42 seconds, source 2 continues increasing its window even though the congestion level

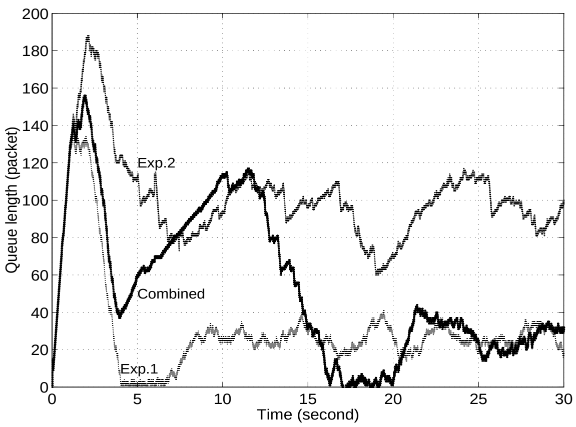

Figure 10: Queue lengths of different marking algorithms

on this link hasn't changed. The individual sources are still TCP sources. Due to the higher service level of these packets (all packets marked IN), the packets are not dropped. During this transition (until the queue lengths build up to a level where IN packets have to be dropped), this source continues sending packets at a higher rate even though the link is congested. As a result, cwnd of source 1 is reduced to one.

5 Discussion and Related work

Our simulation experiments show that: (a) Even though individual TCP flows may respond to congestion, this congestion avoidance backoff can be muted by moving to next higher service level (as seen in Fig. 11). (b) During such transitions, the congestion may actually increase (as observed by increased queue lengths) even when individual flows respond to congestion (Fig. 10. (c) When employed universally, these strategies can result in a bidding war for resources as each aggregated source shifts its resources to congested links (by moving up the service levels) and (d) the congested links eventually settle down to serving packets of highest service level while shutting down service for other packets (best-effort flows in our simulations, as seen in Fig. **??**).

Pricing [19, 20] will have an important effect on a number of the above observations, specifically on the nature of moving up the service levels. It has been suggested [21] that resources should be priced based on the level of congestion to balance load evenly across the network links. Network providers may employ fair-sharing techniques [11, 13, 12, 15, 14, 16] to balance resource utilization among competing aggregated sources at the time of congestion. In such a case, shifting resources to congested links will likely have less impact than observed in this study.

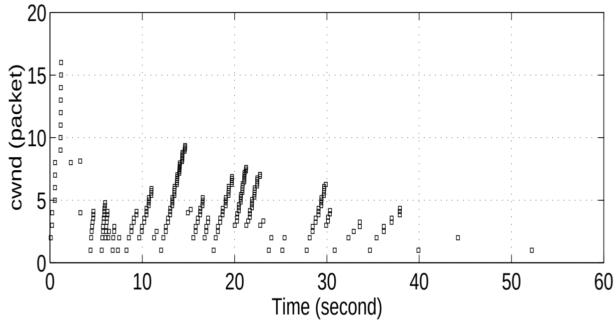

(a) Source 1

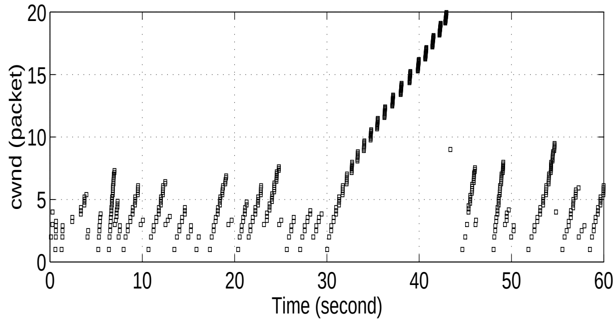

(b) Source 2

Figure 11: Congestion windows of different marking algorithms

Aggregated sources also pose interesting new questions. As observed in this study and by several others, aggregated sources do not behave like TCP sources even when all the constituent flows within the aggregation are TCP flows. With suggested modifications to TCP [10], individual TCP flows will appear less responsive to congestion. If the aggregated source readjusts its resources among the individual flows, even when an individual flow backs off, it is likely that another flow within the aggregation may send more packets through the congested link. All of these issues point to the need for studying network bandwidth management and network dynamics further.

Recent work on diff-serv networks mostly dealt with individual sources [1, 2, 3, 4, 5]. Adaptive marking to achieve throughput targets for single sources is studied in [5]. Aggregation of individual traffic sources and the resulting traffic distributions have been studied [17, 18]. Aggregation in the context of RSVP has been studied in [7, 8, 9]. Our work looked at the problem of sharing available resources (IN-profile bandwidth) among the individual flows of an aggregation to achieve specific performance goals.

6 Conclusions

In this paper, we have studied the impact of aggregated sources on the guarantees/performance provided by a differentiated-services network. We have shown that differences in RTTs of individual flows have less of an impact on the realized throughputs of aggregated sources than on the throughputs of individual sources. We have shown that proportional marking of packets within an aggregation can lead to significant unfairness among the flows within the aggregation. We have proposed new marking algorithms for aggregated sources that improve fairness among the individual flows within the aggregation. The proposed algorithms maintain state for individual flows at the boundary marker to achieve specific performance goals. We have presented simulation results to show the impact of the proposed algorithms on realized throughputs, network congestion and the scalability of the proposed approaches.

References

[1] J. Ibanez and K. Nichols, "Preliminary Simulation Evaluation of an Assured Service," draft-ibanez-diffserv-assured-eval-00.txt, INTERNET DRAFT, August, 1998.

[2] K. Nichols, V. Jacobson, and L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet," draft-nichols-diff-svc-arch-00.txt, INTERNET DRAFT, December, 1997.

[3] David D. Clark and Wenjia Fang, "Explicit Allocation of Best-Effort Packet Delivery Service," November, 1997. http://diffserv.lcs.mit.edu/exp-alloc-ddc-wf.pdf

[4] Anindya Basu and Zheng Wang, "A Comparative Study of Schemes for Differentiated Services," *Bell labs Technical Report*, Aug. 1998.

[5] W. Feng, Dilip D. Kandlur, D. Saha, and Kang G. Shin, "Adaptive Packet Marking for Providing Differentiated Services in the Internet." *Proc. of Int. Conf. on Network Protocols*, Oct. 1998.

[6] Network simulator (Ns), University of California at Berkeley, CA, 1997. Available via http://www-nrg.ee.lbl.gov/ns/.

[7] Andreas Terzis, Lixia Zhang, and Ellen L. Hahne, "Reservations for Aggregate Traffic: Experiences from an RSVP Tunnels Implementation," *IEEE/IFIP IWQoS '98*, May, 1998.

[8] Steven Berson and Subramaniam Vincent, "Aggregation of Internet Service States," *IEEE/IFIP IWQoS '98*, May, 1998.

[9] Olov Schelen and Stephen Pink, "Aggregating Resource Reservations over Multiple Routing Domains," *IEEE/IFIP IWQoS '98*, May, 1998.

[10] W. Feng, Dilip D. Kandlur, D. Saha, and Kang G. Shin, "Understanding TCP Dynamics in an Integrated Services Internet," *Proceedings of the International Workshop on NOSSDAV*, May, 1997.

[11] A. Demers, S. Keshav and S. Shenker, "Analysis and simulations of a fair queuing algorithm," *Proc. of ACM SIGCOMM*, 1989.

[12] S. Golestani, "Duration limited statistical multiplexing of delay-sensitive traffic," *Proc. of INFOCOMM*, 1991.

[13] A. Parekh and R. Gallager, "A generalized processor sharing approach to flow control in integrated service networks: The multiple node case," *Proc. of IEEE INFOCOMM'93*, Mar. 1993.

[14] B. Suter, T. V. Lakshman, D. Stiliadis and A. Choudhury, "Design considerations for supporting TCP with per-flow queuing," *Proc. of INFOCOMM '98*, Apr. 1998.

[15] I. Stoica, S. Shenker and H. Zhang, "Core-Stateless Fair queueing: Achieving approximately fair bandwidth allocations in high speed networks," *Proc. of ACM SIGCOMM*, Oct. 1998.

[16] R. Guerin, S. Kamat, V. Peris and R. Rajan, "Scalable QOS provision through buffer management," *Proc. of ACM SIGCOMM '98*, Oct. 1998.

[17] W. Willinger, M. Taqqu, R. Sherman, and D. Wilson, "Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level," *Proc. ACM SIGCOMM '95*, 1995.

[18] M. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", *Proc. ACM SIGMETRICS '96*, May 1996.

[19] R. Cocchi, D. Estrin, S. Shenker, and L. Zhang. Pricing in Computer Networks: Motivation, Formulation and Examples. *ftp://parc.ftp.xerox.com/pub/net-research/policy.ps.Z.*

[20] D. Clark. A model for cost allocation and pricing in the internet. *MIT workshop on Internet Economics*, 1995.

[21] I. Stoica and H. Zhang. LIRA: An approach for service differentiation in the internet. *Proc. of NOSSDAV '98*, June 1998.