The Impacts of Aggregation on the Performance of TCP Flows in DS Networks

Virpi Laatu, Jarmo Harju and Pekka Loula
Tampere University of Technology, Pohjoisranta 11 P.O Box 300, FIN-28101 Pori, Finland
virpi.laatu@tut.fi, jarmo.harju@tut.fi, pekka.loula@tut.fi

Abstract

This paper examines the fairness issues among individual TCP flows in a Differentiated Services (DS) network. In a DS network, service agreements for customer traffic are typically specified at aggregate levels, rather than on the level of individual flows. When we deal with aggregated sources in such networks, we need to consider not only the fairness issues among aggregates but also the fairness issues among individual flows within an aggregate. The aim of the measurements was to evaluate how aggregate marking affects the performance of individual TCP flows when competing with each other within the same Assured Forwarding (AF) class. Since aggregation is a key factor in any DS capable network, it is essential that an improvement of QoS (Quality of Service) should be seen also by individual flows.

1 Introduction

The Differentiated Services (DS) architecture [1] is a respectable way to prove preferential treatment to certain packets inside the IP network, but it also brings new problems, especially to TCP users. There are many problems when running TCP over DS networks. One of the main problem is that TCP is unaware of target rates, which means that TCP adjusts its flow rate to fit available bandwidth. Another problem is the interaction between different flows in the same bottleneck link. Many researches have concluded that the interaction between TCP and UDP traffic may cause UDP traffic to impact TCP traffic in adverse manner. The problems are mainly caused by TCP's congestion control. In the case of congestion, the TCP flow slows down its transmission rate, whereas UDP flow will not back off. There are also problems between TCP flows with different characteristics, such as TCP flows with different round trip times (RTTs), Web connections and FTP flows.

Things will become more complex, because in a DS network, due to scalability traffic marking at edge nodes is mainly performed on the aggregated level. This kind of aggregation marking mechanisms can further deteriorate the performance of TCP flows. Since aggregation is a key factor in DS networks, it is important to study the effects of aggregation, especially from end-user's point of view. The use of aggregation raises a number of concerns, such as fairness concern. The fairness between the users is an important issue, since we expect that the service discrimination will be based on some kind of pricing models. In this study, we evaluate how per-aggregation marking mechanism affects fairness within an aggregate and between aggregates that are entering an Internet service provider's (ISP's) network.

Our study focuses on services built on top of the Assured Forwarding (AF) PHB [2]. In order to support the requirements of the AF service, the routers of a DS network need to implement more sophisticated scheduling and buffer management policies that discriminate packets according to the service classes they belong to. However, these traffic markers and scheduling mechanisms are designed mainly to support service differentiation for aggregated flows and not tailored for TCP traffic, which constitutes the majority of the traffic on the Internet today.

The fair sharing of the available bandwidth is a major issue, especially in the context of the AF service. Many studies have tackled fairness problems in the AF service based networks. However, these research efforts [e.g., 3, 4, 5, 6, 7] have almost exclusively looked at fairness issues on a per-aggregation basis, not within an aggregate, as addressed in this study. Some studies [8, 9, 10, 11, 12, 13] have dealt with fairness between individual flows within an aggregate, but they applied to an aggregate consisting of TCP and UDP traffic or the traffic marking is performed on a per-flow basis, or their proposed improvements are too complex or not scale well. Medina et al. in [14] evaluated how configuration parameters for WRED (Weighted RED) affect the bandwidth sharing between individual AF flows within an aggregate. However, WRED parameters were only configured for UDP traffic. In fact, WRED parameters have to be optimised for TCP, because they are usually dropped in the case of congestion.

We are not aware of any other previous study that evaluated the problems of fair bandwidth allocation among individual TCP flows of an aggregate when using an aggregate marking scheme. Thus, the focus of our experiments was to determine how significantly the aggregate marking mechanism affects the bandwidth distribution among the TCP flows within a single AF class from an end-user's perspective.

The rest of this paper is organized as follows: Section 2 describes briefly the Assured Service scheme. In Section 3, we describe the measurement scenario used in the study, i.e., what things were measured, how, and in what kind of conditions. In Section 4, we present the results of the measurements. Section 5 summarizes our findings and concludes the paper.

2 Assured Service scheme

In this section, we describe the components that compose the Assured Service scheme used in this study, i.e., traffic conditioning mechanisms, such as a policer, a traffic meter and a packet marker, and an active queue management technique that implements the AF PHB group itself.

The AF PHB group as specified in [2] provides the delivery of IP packets in four independent traffic classes (AF classes), each with three level of drop precedence (low/green, medium/yellow and high/red). The Assured Service relies on packet monitoring and marking mechanisms, performed by the traffic conditioner at the edge node of a DS network, and an active queue management (AQM) mechanism at the core nodes. ISP ensures that the aggregate traffic generated complies with the traffic profiles specified in the SLAs between the users and the network. In this study, we considered the SLAs that are made on a per-customer (e.g., a small company) basis rather than on a per-connection basis. We used a dual token-bucket based mechanism called Two Rate Three Color Marker (trTCM) [15] to check the traffic conformity at the edge routers, and to mark packets in agreement with service profile. A trTCM measures incoming traffic from the customer and marks the packets based on two rates, Committed Information Rate (CIR) and Peak Information Rate (PIR), and their associated burst sizes, Committed Burst Size (CBS) and Peak Burst Size (PBS), respectively. CBR is used as the green token bucket size, whereas PBS is used as the yellow token bucket size. When the customer's measured traffic is within the contracted average sending rate (CIR), the packets are marked as green. The packets that exceed the contracted rate are not discarded immediately, but marked as yellow instead when the rate falls below the PIR. All packets that exceed the PIR are marked as red.

In case of congestion within the AF class, the packets with high drop precedence are discarded first. There are many alternative AQM policies to be used at the core routers, such as Random Early Detection (RED) [16], RIO (RED with IN and OUT) [17], and Core-stateless Fair Queuing (CSFQ) [18]. Following the AF specification [2], we chose RIO-like AQM policy, which implements the AF PHB using the three-priority (color) version of the RED. This mechanism is known as GRED (Generalized RED). GRED allows multiple drop precedence levels within an AF class. In our experiments, we used three sets of RED parameters, one for each color. The parameters setting for GRED are listed in Table 1.

Table 1. GRED parameters used for experiments

Parameters	Green	Yellow	Red
	Queue length L=120 kB		
Min_{th}	0.375L	0.1875L	0.0875L
Max_{th}	0.75L	0.375L	0.1875L
Max_p	0.02	0.05	0.1
W_q	0.002	0.002	0.002

For our traffic markers, we used a CIR of 800 kbps with a bucket size of 15 packets (CBS). PIR was set 1.4 Mbps, while PBS was set to 18 packets. The rate for AF packets is limited to 3 Mbps. Link sharing between the AF class and best effort (BE) traffic was implemented using a CBQ (Class Based Queuing) mechanism.

For simplicity, we used one AF class for all traffic generated during the experiments. However, more than one class can be considered, allowing more flexibility in assigning applications to the different AF classes.

3 Measurements

In order to illustrate how equal service subscribers (equal service profiles), emitting the different number of flows, compete for the available bandwidth, we conducted the testbed shown in Fig. 1. Our DS testbed was set up using the Linux QoS support. We conducted experiments where two traffic aggregates (called Aggr1 and Aggr2), with identical target rates (800 kbps) but different numbers of flows compete for the available bandwidth. Aggr1 consists of eight individual TCP sources, while Aggr2 consists of two individual TCP sources.

The scenarios were composed of ten measured TCP flows with a mixture of TCP and UDP flows as background traffic. The background traffic was used to create bottleneck for the AF traffic on the link between Core1 and Core2 nodes The total sum of green colored background TCP packets was limited to 800 kbps. PIR for aggregated background TCP flows was set to 1.4 Mbps. In addition, we added unbounded best effort UDP traffic to fill up the rest of the available link bandwidth on the 10 Mbps bottleneck link. During the measurements, only one AF class was used. The bandwidth for AF packets was limited to 3 Mbps.

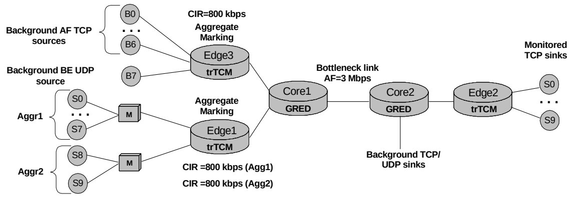

Fig. 1. The test network topology

When we deal with aggregated flows in DS networks, we need to consider not only the aggregate throughput, but also goodput or throughput achieved by individual flows. The behavior was mainly evaluated by observing the distribution of bandwidth among individual TCP flows generated by an FTP application. In these experiments, we consider the following performance metrics: (1) utilization of the committed rate (CIR), (2) the average goodput and throughput obtained by an FTP flow, (3) the fairness achieved in the allocation of committed, excess (yellow and red packet rate) and total bandwidth among different flows, (4) the throughput achieved by the aggregated source. For each flow, the number of green colored packets delivered to the corresponding destination is calculated (the green throughput). Utilization of the committed rate by a flow is measured as the ratio of the green throughput of a flow and the expected fair share of committed rate

(CIR of aggregate rate shared equally by each flow within an aggregate). To quantify the level of the fairness, we used the fairness index [19] and the standard deviation.

Each test was carried out five times with the same parameters to gain confidence in our results. We assumed that after ten seconds the background traffic gets stabilized and the monitored sources could start their transmission. The monitored TCP flows were started simultaneously. All tests lasted for 120 seconds.

4 Experimental results

Fig. 2 shows the total bandwidth achieved by each of the individual TCP flows within the two monitored traffic aggregates. The results show that both aggregates fail to achieve their aggregated target rates. Also, the variance of the throughput seen by individual flows is relatively large (bigger unfairness) among the flows within an aggregate Aggr1. As seen from this example run, some flows within an aggregate Aggr1 back off, while the other flows obtain more than their fair share of the available bandwidth.

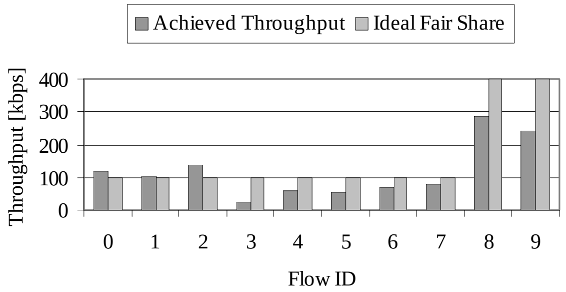

Fig. 2. Measured per-flow throughputs

Table 2 summarizes the statistical results for the experiment. The results show that the flow aggregate with eight flows (Aggr1) always gets better share of the available bandwidth. The results show that the aggregate containing two flows (Aggr2) achieves only on average 510 kbps, whereas the aggregate Aggr1 gets around 630 kbps. It appears also from the results that standard deviation is higher within the aggregate Aggr1 than it is in the case of the aggregate Aggr2.

Table 2. Measurement results for aggregates

	Aggr1 (8 flows)		Aggr2 (2 flows)	
	Mean	Sdev	Mean	Sdev
Utilization of CIR	0.80	0.11	0.63	0.02
Per-flow goodput [kbps]	75.3	28.8	243.3	16.1
Per-flow thr.put [kbps]	78.4	29.9	257.6	16.3
Aggregate thr.put [kbps]	629.9	21.4	515.1	14.5

The fairness indices between individual flows and flow aggregates were also calculated. Table 3 shows the obtained fairness indices between all the ten monitored TCP flows, whereas Table 4 tells the obtained fairness indices between the two flow aggregates. The results show that the overall fairness (total) is quite poor among individual flows. However, the overall fairness between the two aggregates stays on a satisfactory level.

Table 3. Fairness indices between individual flows

Committed		Excess		Total	
Mean	Sdev	Mean	Sdev	Mean	Sdev
0.881	0.022	0.838	0.064	0.890	0.031

Table 4. Fairness indices between flow aggregates

Committed		Excess		Total	
Mean	Sdev	Mean	Sdev	Mean	Sdev
0.984	0.017	0.791	0.026	0.991	0.001

5 Conclusions

A fair distribution of bandwidth is one of the main concerns in all cases where equal treatment of the users' flows is assumed. Fairness is an important factor to guarantee the level of agreed service and satisfy the end-users' expectations.

In this study, we examined the effects that the Assured Services mechanism has on the behavior of TCP flows. We studied how fairness is affected by changing the number of flows injected into the traffic aggregate. The results showed that there is significant variation in the performance seen by individual end-users when using the aggregate marking scheme. Moreover, our measurements showed that the aggregate containing many flows outperforms the aggregates containing fewer flows in terms of achieved aggregate throughput. On the other hand, we found that the unfairness seems to be more evident for the aggregate with bigger number of flows.

Since aggregation is a key factor in any DS network, an improvement of QoS should be seen by individual flows. Also a fair distribution of committed bandwidth is one of the main concerns. Without addressing the fairness problems in the context of AF service, it is unlikely that any kind of quantifiable guarantees can be provided by this service.

References

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, An Architecture for Differentiated Services, *RFC2475*, December 1998.

[2] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, Assured Forwarding PHB Group, *RFC2597,* June 1999.

[3] B. Nandy, J. Ethridge, A. Lakas, A. Chapman, Aggregate Flow Control: Improving Assurances for Differentiated Services Network, *Proceedings of the IEEE Conference on Computer Communications (Infocom'01)*, Alaska, USA, April 2001.

[4] H. Su, M. Atiquzzaman, ItswTCM: A New Aggregate Marker to Improve Fairness in DiffServ, *Computer Communications journal*, Vol. 26, 2003, pp. 1018-1027.

[5] B. Nandy, N. Seddigh, P. Pieda, J. Ethridge, Intelligent Traffic Conditioners for Assured Forwarding Based Differentiated Services Networks, *Proceedings of the IFIP High Performance Networking*, Paris, June 2000.

[6] Y. Chait, C.V. Hollot, V. Misra, D. Towsley, H. Zhang, J.C.S. Lui, Providing Throughput Differentiation for TCP Flows using Adaptive Two-Color Marking and Two-Level AQM, *Proceedings of the IEEE Conference on Computer Communications (Infocom'02),* New York, June 2002.

[7] A. Habib, S. Fahmy, B. Bhargava, Design and Evaluation of an Adaptive Traffic Conditioner for Differentiated Services Networks, *Proceedings of the IEEE International Conference on Computer Communication and Networks (IC3N'01),* Arizona, October 2001, pp. 90-95.

[8] I.B.H. de Alves, J.F. de Rezende, L.M. de Moraes, Evaluating Fairness in Aggregated Traffic Marking, *Proceedings of the IEEE Global Communications Conference (Globecom'00),* San Francisco, November 2000.

[9] S. Yi, X. Deng, G. Kesidas, C.R. Das, Providing Fairness in DiffServ Architecture, *Proceedings of the IEEE Global Communications Conference (Globecom'02),* Taipei, Taiwan, November 2002.

[10] I. Yeom, A.L.N. Reddy, Adaptive Marking for Aggregated Flows, *Proceedings of the IEEE Global Communications Conference (Globecom'01),* San Antonio, Texas, November 2001.

[11] A. Das, D. Dutta, A. Helmy, Fair Stateless Aggregate Traffic Marking using Active Queue Management Techniques, *Proceedings of the IEEE International Conference on Management of Multimedia Networks and Services (MMNS'02),* Santa Barbara, USA, October 2002, pp. 211-223.

[12] I. Yeom, A.L.N. Reddy, Impact of Marking Strategy on aggregated flows in a Differentiated Services Network, *Proceedings of the IEEE International Workshop on Quality of Service (IWQoS'99),* London, June 1999, pp. 156-158.

[13] T. Hatano, H. Shigeno, K. Okada, Y. Matsushita, Fair Rate TCP for Fair Share of Bandwidth over DiffServ Networks, *Proceedings of the IEEE Pacific Rim Conference on Communications, Computers and Signal Processing (PACRIM'03),* Canada, August 2003.

[14] O. Medina, J. Orozco, D. Ross, Bandwidth Sharing under the Assured Forwarding PHB, *Minutes of the 5th TF-NGN meeting,* irisa.fr/bibli/public/pi/2001/1478/1478.html, October 2001.

[15] J. Heinanen, R. Guerin, A Two Rate Three Color Marker, *RFC2698,* September 1999.

[16] S. Floyd, V. Jacobson, Random Early Detection Gateway for Congestion Avoidance, *IEEE/ACM Transactions on Networking,* Vol. 1, No. 4, August 1993, pp. 397-413.

[17] D. David, C. Fang, W. Fang, Explicit Allocation of Best Effort Packet Delivery Service. *IEEE/ACM Transactions on Networking,* Vol. 6, No. 4, August 1998, pp. 362-373.

[18] A. Das, D. Dutta, A. Helmy, Fair Stateless Aggregate Traffic Marking Using Active Queue Management Techniques. *Proceedings of the IEEE International Conference on Management of Multimedia Networks and Services (MMNS'02)*, Santa Barbara, USA, October 2002, pp. 211-223.

[19] R. Jain, The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation and Modelling, *John Wiley and Sons, Inc.*, 1991, pp. 36-37.