University of Würzburg
Institute of Computer Science
Research Report Series

Performance Comparison of Different Class-and-Drop Treatment of Data and Acknowledgements in DiffServ IP Networks

Stefan Köhler and Uwe Schäfer

Report No. 237 14th August 1999

University of Würzburg
Department of Computer Science
Lehrstuhl für Informatik III
Am Hubland, D-97074 Würzburg, Germany
koehler@informatik.uni-wuerzburg.de

Performance Comparison of Different Class-and-Drop Treatment of Data and Acknowledgements in DiffServ IP Networks

Stefan Köhler and Uwe Schäfer
University of Würzburg
Department of Computer Science
Lehrstuhl für Informatik III
Am Hubland, D-97074 Würzburg, Germany
koehler@informatik.uni-wuerzburg.de

Abstract

In differentiated services IP networks the sender of a TCP connection determines the class of data packets he emits. The receiver chooses the class of the acknowledgements sent back, independently of the received class. In this work, we examine the impact of different drop precedence in a class and the assignment of different classes for data and acknowledgements. The results show that the throughput of a TCP connection depends not only on the data class, but also on the right choice for the acknowledgements. Some combinations of classes for data and ACKs could even lead to an "unfair" use of bandwidth. On the other hand, for a high throughput the selection of the drop precedence is in most cases only important for data packets.

1. Introduction

Internet Service Providers are looking for strategies to offer "differentiated services" to satisfy customer demand for quality of service (QoS) and of its potential to increase revenues. How to support these differentiated services is still a subject of research. One of the models, the DS service architecture [6], tries to implement different levels of services for individual or aggregated flows. The current framework allows the sender to mark its packets by setting the appropriate bits in an IP header field [2]. The network provider then checks the packets at the network border for conformance to service contracts. If the packet conforms to the contract it is marked as IN (in profile) otherwise the packet is marked as OUT (out of profile). The treatment of the IN and OUT packets in the core network depends on the per-hop behavior of the traffic class. Per-hop behavior (PHB) is defined as the externally observable forwarding behavior applied at a DS-compliant node to a DS behavior aggregate [5]. The DS framework is independent of the routing decision and thus does not define any end-to-end service. It achieves scalability by implementing complex classification and conditioning functions only at network boundary nodes. The service is defined through the different treatment (PHB) of the marked packets in the routers.

There are several ongoing discussions in the Internet Engineering Task Force (IETF) DiffServ working group about the interaction between UDP and TCP traffic. In contrast to UDP, TCP uses congestion control based on a window mechanism to reduce its transmission rate. This leads in situations with congestion to unfairness between UDP and TCP. There are already several papers [12, 13] discussing the effects between UDP and TCP and trying to clarify whether it is necessary to use three instead of two levels of drop precedence. As far as we know, no investigation on the effect of different classes for TCP data and acknowledgments in a DS network has been done. Thus, we concentrate on the behavior of TCP connections in

- different proposed traffic classes and
- a single traffic class with different drop precedences.

The aim of this paper is to evaluate the behavior of TCP under these conditions and to clarify the impact on the goodput of a TCP connection.

The paper is organized as follows. Section 2 describes different PHBs, the traffic conditioners which are located at the boundary nodes to mark the packets appropriately and the active queue-management to implement the different PHBs. The topology, configuration and parameters used in the simulations are presented in Section 3. Section 4 discusses the results. Finally, we give a conclusion and an outlook in Section 5.

2. Components of a Differentiated Service Network

We consider in this paper the following defined [3, 4] forwarding per-hop behaviors:

2.1. Expedited Forwarding PHB

The intention of the expedited forwarding (EF)-PHB is to build a low loss, low latency, low jitter, assured bandwidth, end-to-end service through DS domains. Such a service appears to the endpoints like a point-to-point connection or a "virtual leased line" [4] and has also been described as Premium service [11]. In RFC 2598 no explicit treatment of the marked packets is defined, but packets marked for EF-PHB may be remarked at a DS domain boundary only to other codepoints that satisfy the EF-PHB. Packets marked for EF-PHB should not be denoted or promoted to another PHB by a DS domain. Consequently, packets which exceed the agreed rate are dropped.

2.2. Assured Forwarding PHB

The Assured Forwarding (AF)-PHB [3] specifies four traffic classes with three drop precedence levels (colors) in order to provide differentiated services to the customers in IP networks. The level of forwarding assurance of an IP packet in the AF class depends on

1. the forwarding resources that have been allocated to the AF class,
2. the current load within the AF class and,
3. in case of congestion within the class, the drop precedence of the packet.

A DS node does not reorder IP packets of the same microflow, no matter if they are in or out of the profile, as long as they belong to the same AF class. This is motivated by the fact that reordering of TCP packets might cause severe performance problems for the TCP connections. A DS node should implement at least one of the four AF classes, but it is not required to implement all of them. More details on the behavior of the AF-PHB can be found in [3]. Unmarked Traffic is treated as Best Effort.

2.3. Traffic Conditioners

Traffic conditioners are used to shape respectively meter the traffic entering or leaving a DS domain. Beside the "normal" Token Bucket algorithm which guarantees that the burstiness of a flow is bounded in such a way that the flow never exceeds the rate $b + r \cdot t$, where b is the bucket size, t the time and r the token arrival rate, we implemented the following proposal by Heinanen and Guerin [1].

Two Rate Three Color Marking (trTCM) with Three Drop Precedence

The Three Color Marker (TCM) uses two token buckets (P and C) to meter an IP packet stream and marks its packets either green (DP0), yellow (DP1), or red (DP2). The two token buckets have different depths called Peak Burst Size (PBS) and Committed Burst Size (CBS). The first bucket is filled with the Peak Information Rate (PIR) and the second with the Committed Information

Rate (CIR). The conditioner operates in one of two modes, Color Aware and Color Blind mode, where the colors stand for particular codepoints in the IP header. In case of the AF-PHB, the color can be coded as the drop precedence of the packet. In both modes the token buckets P and C are initialized to PBS respectively CBS at time 0. Thereafter, the token count T_p is incremented PIR times per second by one to an upper bound of PBS and the token count T_c is incremented CIR times per second by one up to CBS. If a packet of size B bytes arrives at time t in Color Blind mode, the following happens:

- the packet is red, if $T_p(t) - B < 0$,
- otherwise the packet is yellow and T_p is decremented by B, if $T_c(t) - B < 0$,
- otherwise the packet is green and both T_p and T_c are decremented by B.

If the traffic conditioner is in the Color Aware mode, it reacts as follows:

- the packet is red, if it has been precolored as red or if $T_p(t) - B < 0$,
- otherwise the packet is yellow and T_p is decremented by B, if the packet has been precolored as yellow or if $T_c(t) - B < 0$,
- otherwise the packet is green and both T_p and T_c are decremented by B.

2.4. Active Queue-Management

Beside traffic conditioners which are necessary to control the incoming traffic of a DS domain, a further mechanism is needed to determine the PHBs of the nodes.

Random Early Drop (RED) Queue

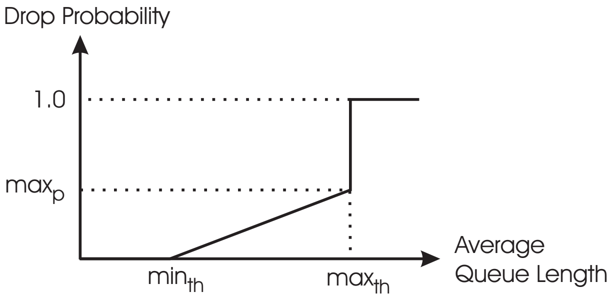

Figure 1: Increasing drop probability in RED.

In RED Queues four parameters need to be configured. Two parameters define the thresholds, min_{th} and max_{th}, where random packet drops occur with respect to the average queue length, see Figure 1. The value max_p determines the drop probability at max_{th} and a weight w_q is used to calculate the average queue length (AQL). An arriving packet at time t is accepted with probability $1 - P(AQL(t))$, where $P(AQL(t))$ is the drop probability. The average queue size is calculated using a low-pass filter with an exponential weighted moving average [9]. The intention behind RED is to control the average queue length through max_{th}, min_{th}, and max_p, as well as the degree of burstiness reflected through w_q.

n-RED

An extension of RED is n-RED in order to support the drop precedences of several classes, where the parameter values differ between the classes.

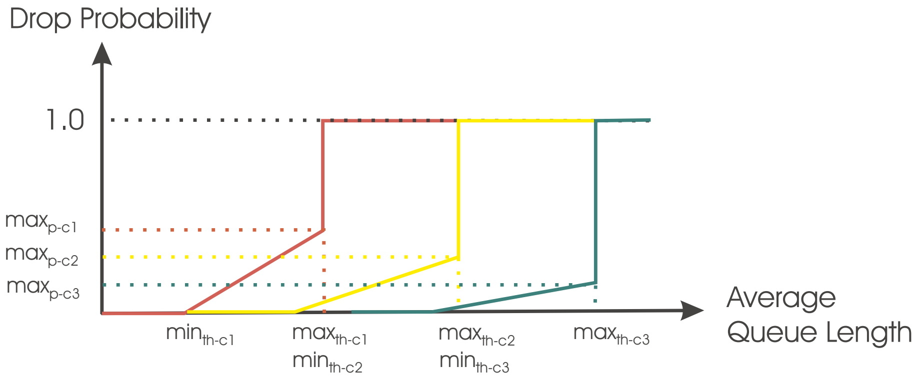

Figure 2: Drop probability for different traffic classes in n-RED ($n = 3$).

We use the overlap n-RED model as shown in Figure 2. The parameter n defines the level of drop precedence. The average queue length for packets at level n is the sum of packets in class 0 to n. The average queue length determines the drop probability as explained in the previous section. If the different drop precedences are adopted to a color as shown in Figure 2, the n-RED model can be used to define an AF-PHB. There are further possibilities to extend the RED model, for instance the RIO (RED with IN and Out) mechanism calculates an average queue length for IN packets (AQL_{IN}), and another for all packets in the queue (AQL_{ALL}). Both packet classes use the same threshold parameters, whereas the drop probability for IN packets depends on AQL_{IN} and the drop probability for OUT packets on AQL_{ALL}.

The mechanisms we use to build EF- and AF-PHB are only proposals and not defined in an RFC. There is ongoing discussion in the DS working group about optimal mechanisms to obtain service discrimination.

3. Example Network and Parameters

For our simulations we used the network simulator **ns** version 2.1b5 [8] developed at UC Berkeley, LBL, USC/ISI and Xerox PARC. The code has been modified to implement the traffic conditioners, multi-color RED and RIO queues.

3.1. Network Topology

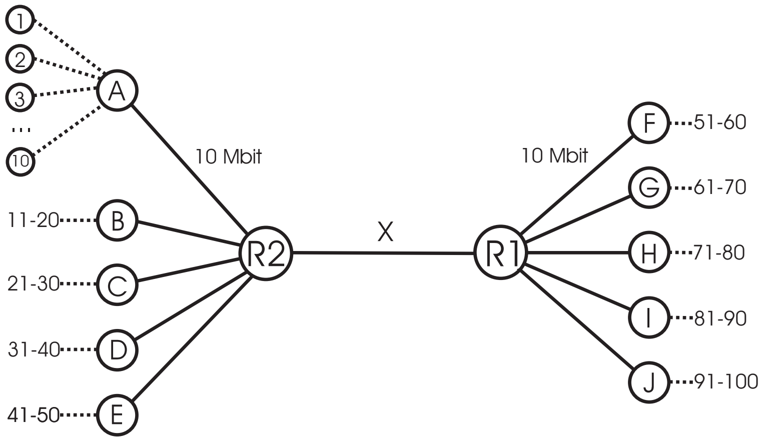

Figure 3: Network topology used for simulation.

For simplicity and comparability with other simulations we chose a similar topology to [7] or [12]. As shown in Figure 3, sources and destinations are connected via a bottleneck link between router R1 and router R2. Five nodes with traffic conditioners to meter the incoming traffic are placed on both sides. Each node is connected over a 10 Mbps link to the router. Each node contains ten individual TCP sources. All sources are ftp-sources and use NewReno to transfer ftp-packets. Two sources send according to the simulation scenario Premium respectively Green traffic (20%), three send Assured respectively Yellow (30%) and five sources use Best Effort respectively Red (50%) to transfer the packets. Thus, there are 50 sources on each side connected to a sink on the opposite side and vice versa. The link between router R1 and router R2 has a bandwidth of $X = 50$ Mbps[1] at 100% load. To simulate a bottleneck we vary the bandwidth X from 100 Mbps to 35.7 Mbps (50% to 140% load on the link). We observe the ftp connection between source 1 and source 51. The remaining 99 traffic sources create background traffic during the whole simulation. The background sources begin their transmission randomly in the first 500ms. After 10 seconds the background traffics get stabilized and the examined source starts to transmit. The observed source is interrupted after 60 seconds of transmission for 20 seconds. This procedure is repeated a total of 10 times. The background TCP connections send ACKs of the same class (drop precedence) as the data. Only the examined connection uses different classes respectively drop precedences for data and ACKs.

Two scenarios are investigated:

- **Two-Bit Differentiated Services Architecture**
  The "classical" model from [11] is implemented. It contains three classes: Premium (EF-PHB), Assured (AF-PHB), and Best Effort. The traffic within a class has identical drop precedence. This scenario shows the impact of different classes for data and ACKs in a TCP connection. Two Token Buckets are placed at each node but not in the routers. The first Token Bucket controls the Premium Rate. If the Premium traffic exceeds its contracted rate the packets get lost. Premium traffic which follows the contract is placed in a preferentially treated FIFO queue. The second Token Bucket observes the Assured Rate. If Assured traffic exceeds its rate, the packets are only remapped to Best Effort. Beside these traffic conditioners, RIO queues are placed in each node and in the routers to treat Assured and Best Effort traffic.

[1] This corresponds to 5 nodes connected with 10Mbps.

- **Two Rate Three Color Marking with Three Drop Precedence**
  To concentrate on the effects of different drop precedence in an AF class, the traffic in the simulations belongs to the same class. The only difference between the TCP connections is the different drop precedence. A trTCM traffic conditioner is placed at each node to control the Color and a n-RED queue. The routers use only n-RED queues without traffic conditioners.

3.2. Parameters for Simulation Study

For all simulations we used a certain set of reference parameters:

Delay

To cover a wide range of different scenarios, the RTT is varied from 20ms to 200ms (see Table 1). For every RTT value a scenario is simulated where all the connections have an equal RTT and a scenario where different RTTs are mixed to avoid possible oscillations of the TCP connections. A mixed scenario consists of traffic sources with 3 different RTTs. As shown in row two in Table 1 the link of the first traffic source to the router has a delay of 3ms, the second a delay of 23ms, the third a delay of 3ms and so on. Both scenarios show nearly equivalent results (see Appendix).

node to router [ms]	**R1 to R2 [ms]**	**router to destination [ms]**	**RTT [ms]**
3	4	3	20
3-23-3-23-48	4	3-23-3-23-48	20
10	10	10	60
15	20	15	100
23-48-3-23-48	4	23–48-3-23-48	100
30	40	30	200
48-23-3-23-48	4	48-23-3-23-48	200

Table 1: Simulated delays, RTT is the round trip time of the observed connection.

Load

It is difficult to define load for a TCP connection, because TCP adapts its window size to the available resources. In the simulation scenario each node is connected over a 10 Mbps link to the router. Five nodes are connected to each router. Thus, if the link from R1 to R2 has a bandwidth of 50 Mbps in every direction, the system is well defined and has a load of 100%. To create a bottleneck, the bandwidth of the router link is varied in the following way:

load [%]	50	100	110	120	130	140
bandwidth [Mbps]	100	50	45.5	42	38.5	36

Table 2: Simulated loads of the bottleneck link

Traffic Conditioners and Queue Management

The values for the traffic conditioners are motivated by the fact that 20% of the traffic should be Premium respectively AF11 and 30% of the traffic should be Assured respectively AF12.

Premium CIR	250 kByte/s = 2 Mbps
Premium CBS	50 kByte
Assured CIR	375 kByte/s = 3 Mbps
Assured CBS	70 kByte

Table 3: Two-Bit DS token bucket parameters

CIR	250 kByte/s = 2 Mbps
CBS	50 kByte
PIR	375 kByte/s = 3 Mbps
PBR	70 kByte

Table 4: trTCM token bucket parameters

Theoretical studies still have to be done to formulate general rules which determine the right parameters for traffic conditioners and queues in a DS network. To be comparable, we oriented our choice on previous simulations [7, 10, 12, 13] and performed some studies to find appropriate parameters.

$\min_{\text{out}}$	35
$\max_{\text{out}}$	50
p_{out}	0.1
$\min_{\text{In}}$	55
$\max_{\text{In}}$	65
p_{In}	0.05
w_q	0.002 (=500 packets)
queue limit	90

Table 5: RIO queue parameters

$\min_{\text{red}}$	35
$\max_{\text{red}}$	50
p_{red}	0.3
$\min_{\text{yellow}}$	45
$\max_{\text{yellow}}$	60
p_{yellow}	0.2
$\min_{\text{green}}$	60
$\max_{\text{green}}$	70
q_{green}	0.1
w_q	0.002 (=500 packets)
queue limit	90

Table 6: trTCM parameters

The maximum transfer unit (MTU) of the ftp connections is 1000 Bytes for all sources. In the following section, we concentrate on the most significant results. More results can be found in the Appendix.

4. Simulation Results

In this section we present numerical results for the TCP goodput based on the parameters given in the previous section. Each point in the figures represents the average amount of data that was transferred within 60s. The errorbars are the 95% confidence intervals. We group the points according to their data class and connect the points with different acknowledgement classes.

In the rest of the paper we use yx as abbreviation for a connection with class y for the data and class x for its ACKs (e.g. PB for Premium data and Best-Effort ACKs). The notation Premium, Assured, Best Effort express Px, Ax, and Bx.

4.1. Two-Bit Differentiated Services

We implement the Two-Bit DS model to estimate the impact of different classes for data and ACKs in a TCP connection. In the following two sections we present the results for different delays and loads.

Influence of Different Delays

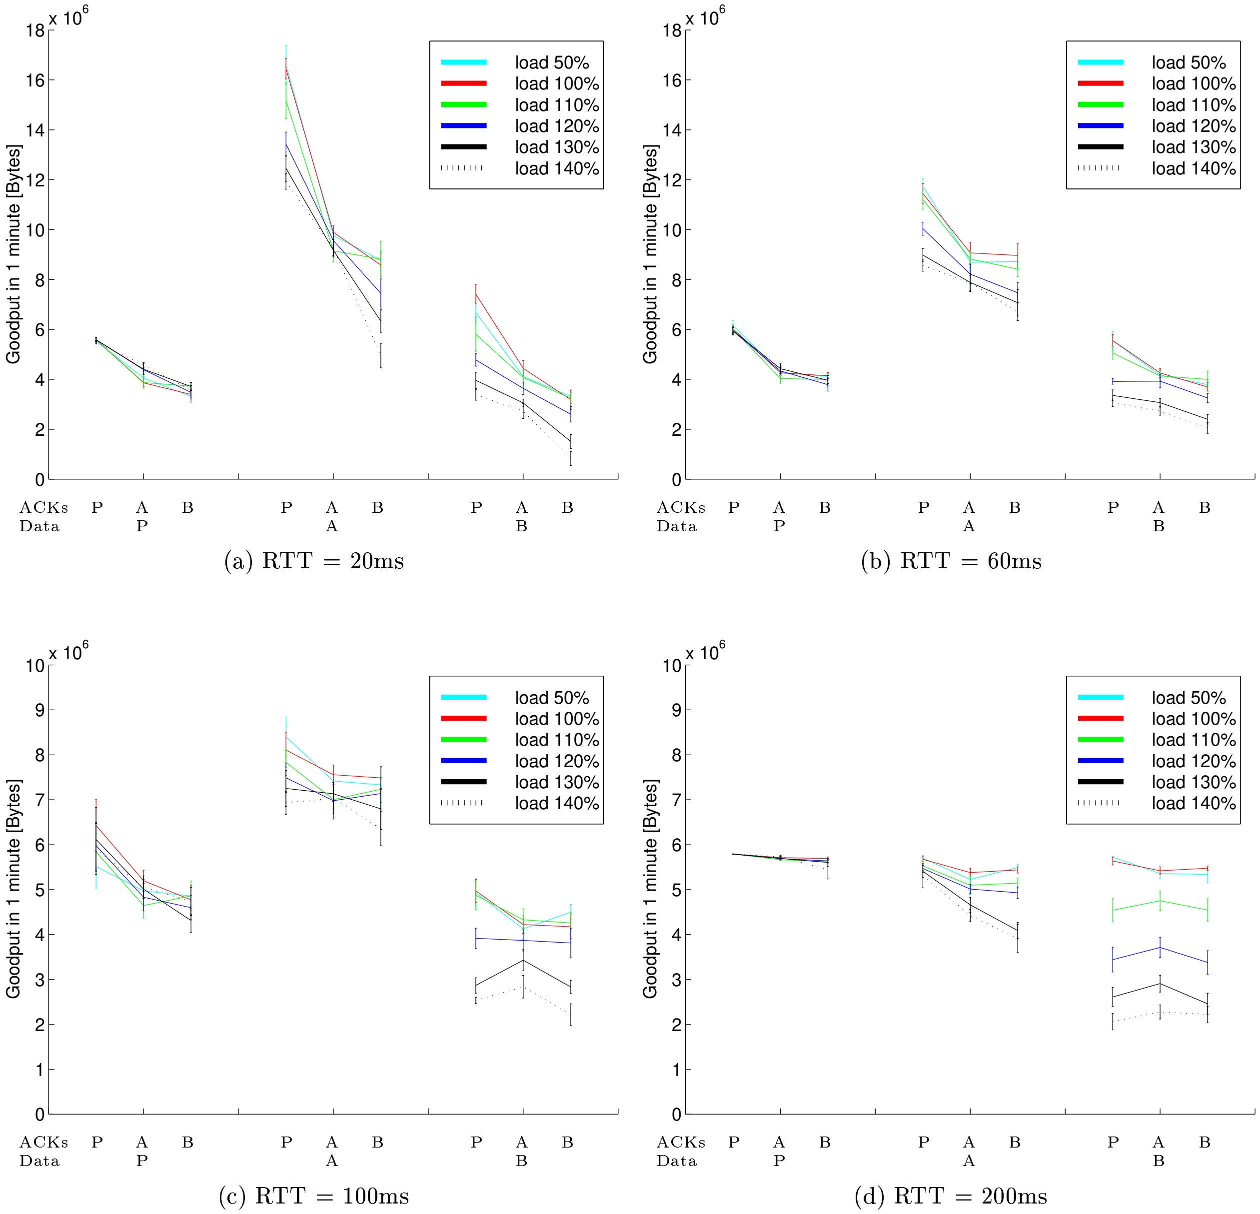

Figure 4: Comparison of goodput for different class combinations with varying load

Figure 4(a) shows the results for different loads at a fixed round trip time of 20ms. In this scenario the influence of the ACK class is very clear. As expected the rate for PP connections is independent of the load and – compared to Assured data connections – very low since in Two-Bit DS the dropper discards all Premium packets that are out of profile. The boundary for the throughput depends of course on the chosen parameters for the Premium class. PA and PB connections have an even lower goodput since beside the control of the data rate the ACKs have a higher delay and a probability to get lost in the RIO queue. Due to their small size, acknowledgements are not dropped or remapped at the traffic conditioner. The goodput of PB connections is about 60% in comparison to PP connections. Thus, even for connections in the Premium class with small RTTs the appropriate choice of ACKs is important.

The effect that the goodput does not only depend on the class with which the data is sent, but also on the class of the ACKs, is even more evident for connections that use Assured data. The throughput of AP connections in Figure 4(a) is up to 140% higher than that of AB connections. However, for Assured and for Best Effort connections the variation of the goodput at different loads is significantly higher than for Premium connections. This results from the absence of droppers and the use of RIO queues. These mechanisms cause also that the Assured connections always have a higher goodput than Best Effort connections.

Figures 4(a) to 4(d) show the evolution of the goodput, when the RTTs of the connections increase. With rising round trip time the effect of the Premium dropper lessens as can be see in Figure 4(c). This results in a higher variation of the PP goodput. In general, for a higher round trip time the difference between the classes decreases.

Due to a higher round trip time, the token buckets influence the throughput not any longer and the Premium connections get a higher goodput than the other classes, see Figure 4(d). This is based on the preferential treatment of Premium packets.

In the Figures 4(a) to 4(d) all connections in the respective simulations have an identical round trip time. The results are equivalent for an environment in which the connections have varying round trip times (see Appendix).

Influence of Different Loads

The figures are based on the same simulation presented in the previous section. The goodput is depicted with varying RTTs (including the mixed scenarios) over the different class combinations for a specified load.

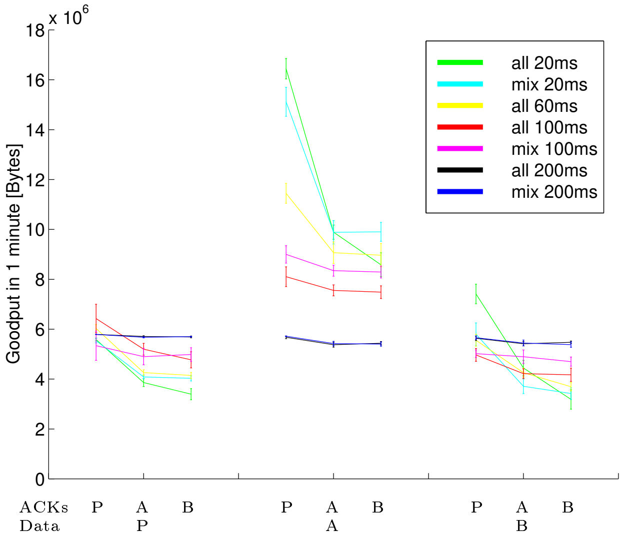

(a) load 100%

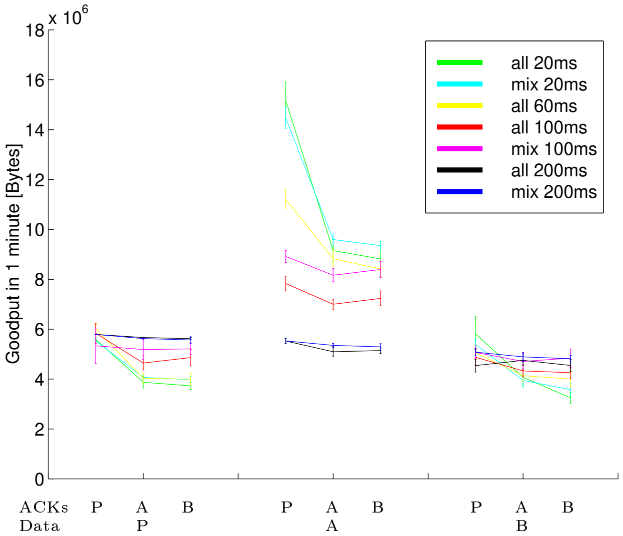

(b) load 110%

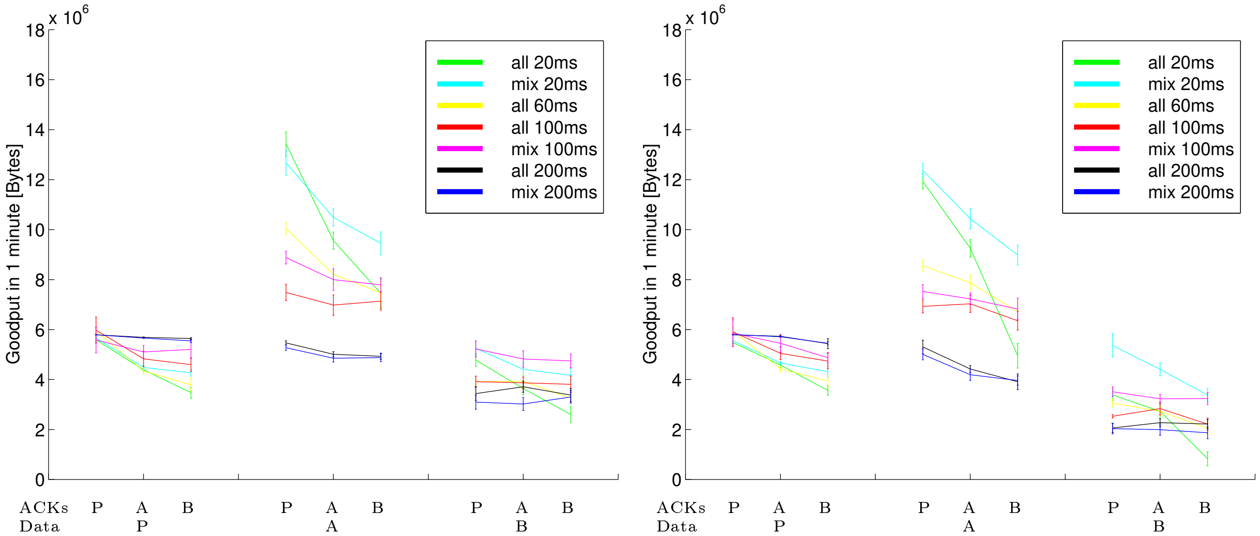

(c) load 120%

(d) load 140%

Figure 5: Comparison of goodput for different class combinations with varying RTTs

Figure 5(a) shows the goodput in dependence of the round trip times and a load of 100%. The figure indicates that with rising RTT the goodput of the different classes approaches an equal level. For low round trip times the Assured connections raise their throughput at the expense of the Best Effort connections. As seen in the figures the goodput is nearly doubled for Assured than for Best Effort connections. Further investigations are needed to clarify the amount of unfairness in these situations.

Whereas at low load the class of the ACKs is the main factor for goodput, at higher loads the ACKs get less important for the goodput. This evolution is presented in the Figures 5(a) to 5(d). In Figure 5(d) more Best-Effort data packets are thrown away due to the high level of congestion and therefore these connections get a lower throughput. This effect could also be observed for low RTT connections. In Figure 5(a) and 5(d) the goodput of an AB connection with a RTT of 20ms is lower than the goodput for the same connection and a RTT of 60ms, whereas the AP connections take profit out of the situation.

All simulations show that Premium and Assured connections are extensively protected by the used mechanisms and that sending Assured data with Premium acknowledgements results in the highest throughput, especially at low round trip times. This behavior should be taken into consideration for pricing models in DS networks. With Premium data packets the throughput can only reach the contracted rate.

4.2. Three Color Marking with Three Drop Precedence

To investigate the influence of different drop precedence in an AF class the whole traffic in the simulations belongs to the same class. Only three different levels of drop precedence are distinguished, labeled AF_{11}, AF_{12}, and AF_{13}, where the first number in the index describes the class and the second the drop precedence. Thus, AF_{11} has the lowest drop probability, followed by AF_{12}.

Influence of Different Delays

The goodput is illustrated with varying load over the different drop precedence combination for a specified RTT in Fig. 6.

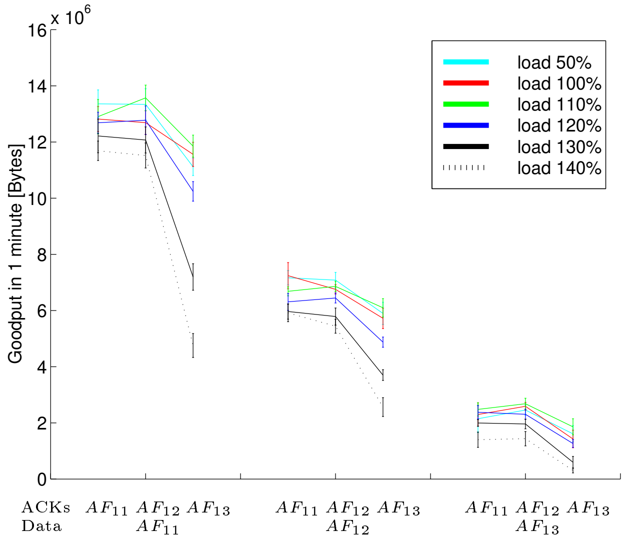

(a) RTT = 20ms

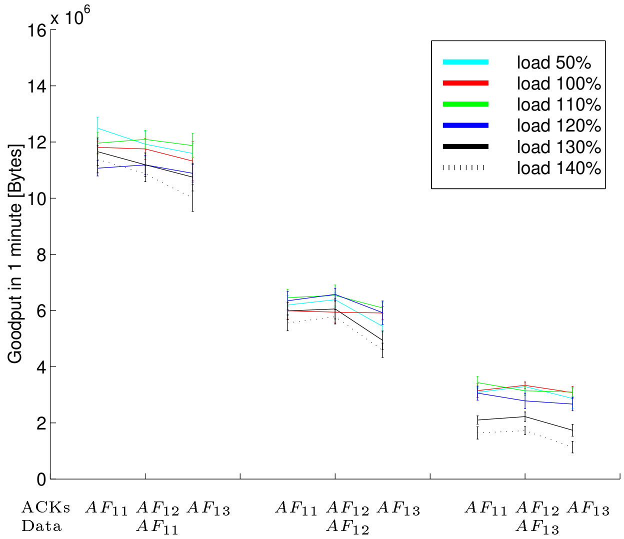

(b) RTT = 60ms

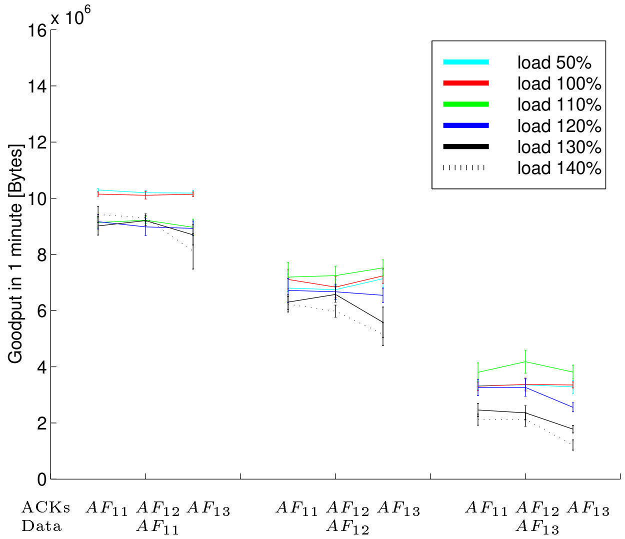

(c) RTT = 100ms

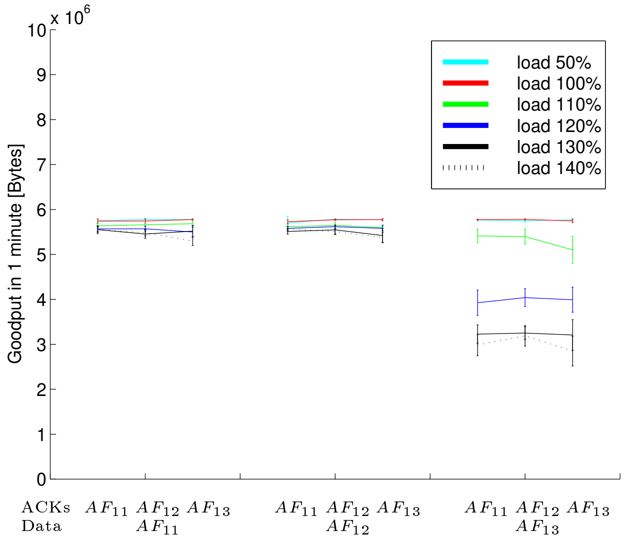

(d) RTT = 200ms

Figure 6: Comparison of goodput for different combinations of drop precedences with varying load

In Figure 6(a) the goodput for each data drop precedence remains within a relatively small range of values. In comparison to the Two Bit model, the throughput depends only on the drop precedence for the data except for very high loads and not on the drop precedence for ACKs. In this situation the n-RED queue protects AF_{11} and AF_{12} at the expense of AF_{13} packets.

As seen in the Figures 6(a) to 6(d), the difference in the goodput between the single drop precedences decreases with increasing round trip times. It is the same behavior we have seen in the Two-Bit DS model. In Figure 6(d) the goodput is independent of the different drop probabilities for low loads. In case of an increasing load the classes with lower drop probability are protected at the expense of the classes with higher drop probability.

Influence of Different Loads

The following figures are based on the same simulations presented in the previous section. The goodput is depicted in Figure 7 with different RTTs (including mixed scenarios) over the combination of drop precedence for a specified load.

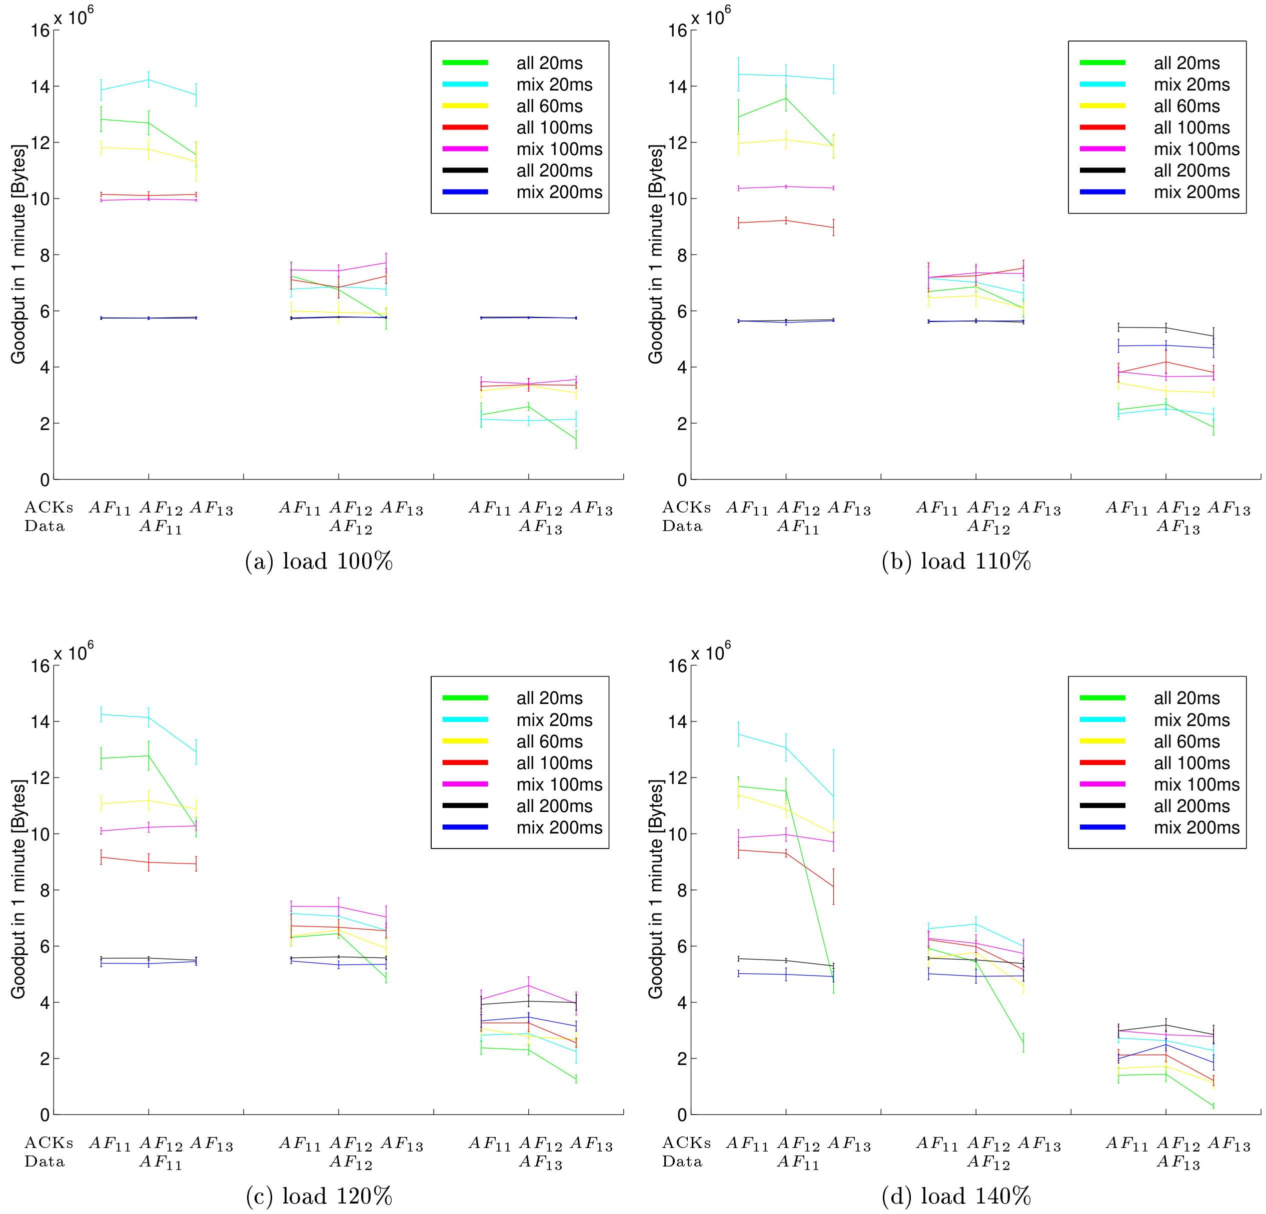

(a) load 100%

(b) load 110%

(c) load 120%

(d) load 140%

Figure 7: Comparison of goodput for different combinations of drop precedence with varying RTTs

In comparison to Two-Bit DS, the confidence intervals of all measured samples are smaller. This is caused by the use of the same queue for all packets. In Two-Bit DS two queues are used, one priority queue for Premium packets and one for Assured and Best Effort packets.

In contrast to Two-Bit DS the acknowledgement class has not a huge influence. The performance of a TCP connection mainly depends on the data class. This results from the fact that the loss of an ACK can be compensated by the TCP protocol, whereas this is not the case when a data packet is lost. Only in regions where the RTT is small and the network is overloaded the class of the ACKs influences the performance, cf. Fig. 7(a). Under these special conditions the goodput could be lowered more than 50%, see Figure 7(d), but as previously mentioned, in most scenarios only the drop precedence of the data is the important parameter for the goodput.

5. Conclusions and Outlook

In this paper we investigated the impact of different class and drop treatment for data and ACKs in a DS environment. The results of our simulations lead to the following conclusions:

- The throughput of a TCP connection in a DS network does not only depend on the sender but also on the receiver. The appropriate choice of the ACK class has a big influence on the throughput.
- In most cases, the use of different drop precedences within a class for ACKs has no influence on the performance worth mentioning. In general, the drop probability of the data determines the goodput. The drop precedence of the acknowledgements is only important for high loads and low round trip times.
- The simulations show that it is possible to get a better throughput with Premium respectively EF-PHB ACKs. In some cases an even unfair gain is made.
- The performance gain for some combinations of data and ACKs should be considered in pricing models.
- Our studies facilitate the use of the EF class for real time traffic because the user gets a fixed, load-independent share of the bandwidth.

Some questions arise from our investigation. They are left to further studies and discussions:

- It should be specified which traffic should be transported in the EF class. In our opinion TCP can not take advantage of this traffic class except when the network is overloaded.
- EF-PHB traffic can not exceed a specified rate. To get a higher throughput the TCP connection may be be split into an EF and AF class part. In this case there is the possibility of reordering TCP packets which affects the throughput. Further studies are needed to estimate if it is worth to use tow data classes for TCP or not.
- A more complex topology and realistic traffic types, especially adaptive, non-adaptive and short-lived (e.g. WWW) traffic flows should be used to investigate the influence of different classes and drop precedences.
- The influence of the routing decision has to be considered (e.g. shortest path for premium traffic).

6. Acknowledgements

The authors would like to thank Norbert Vicari and Kenji Leibnitz for the productive discussions of the presented results. The financial support of the Deutsche Telekom AG (Technologiezentrum Darmstadt) is appreciated.

References

[1] J. Heinanen, R. Guerin, A Two Rate Three Color Marker, Internet Draft <heinanen-diffserv-trtcm-01.txt>, May 1999

[2] K. Nichols, S. Blake, F. Baker, D. Black, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, RFC 2474, December 1998

[3] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, Assured Forwarding PHB Group, RFC 2597, June 1999

[4] V. Jacobson, K. Nichols, K. Poduri, An Expedited Forwarding PHB, RFC 2598, June 1999

[5] M. Carson, W. Weiss, S. Blake, Z. Wang, D. Black, E. Davies, An Architecture for Differentiated Services, RFC 2475, December 1998

[6] Davies, E., Keshav, S., Verma, D., Carlson, M., Ohlman, B., Blake, S., Bernet, Y., Binder, J., Wang, Z., Weiss, W., A Framework for Differentiated Services, Internet Draft <draft-ietf-diffserv-framework-02.txt>, February 1999

[7] J. Ibanez, K. Nichols, Preliminary Simulation Evaluation of an Assured Service. Internet Draft <draft-ibanez-diffserv-assured-eval-00.txt>, August, 1998.

[8] NS simulator, Version 2.1b5 Available from http://www-mash.cs.berkeley.edu/ns

[9] S. Floyd, V. Jacobson, Random Early Detection gateways for Congestion Avoidance. IEEE/ACM Transactions on Networking, Vol 1, Num 4, August 1993, pp. 397-413.

[10] D. Clark, W. Fang, Explicit Allocation of Best Effort Packet Delivery Service. IEEE/ACM Transactions on Networking, Vol 6, Num 4, August 1998, pp. 362-373.

[11] K. Nichols, V. Jacobson, L. Zhang, A Two-bit Differentiated Services Architecture for the Internet, Internet Draft <draft-nichols-diff-svc-arch-00.txt>, November 1997

[12] M. Goyal, P. Misra, R. Jain, Effect of Number of Drop Precedences in Assured Forwarding Internet Draft <draft-goyal-dpstdy-diffserv-02.txt>, June 1999

[13] N. Seddigh, B. Nandy, P. Pieda, Study of TCP and UDP Interaction of the AF PHB Internet Draft <draft-nsbnpp-diffserv-tcpudpaf-00.txt>, June 1999

A. Additional Simulations

In this section we present additional results of simulations where the background traffic has different RTTs. The parameters are described in Section 3. As previously mentioned, the results are similar to the results in Section 4 for a non mixed environment.

A.1. Two Bit DS – Influence of different delays

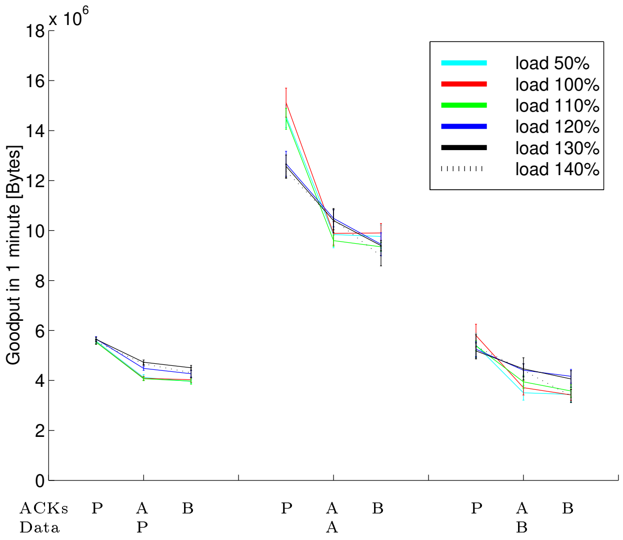

(a) RTT=20ms (connections with different RTTs)

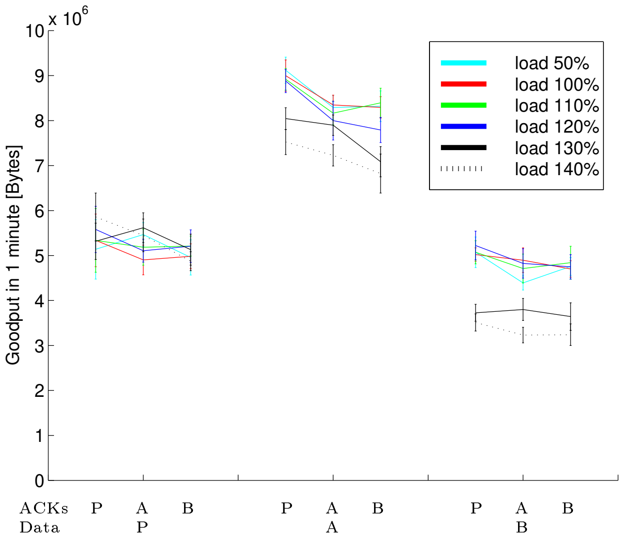

(b) RTT = 100ms (connections with different RTTs)

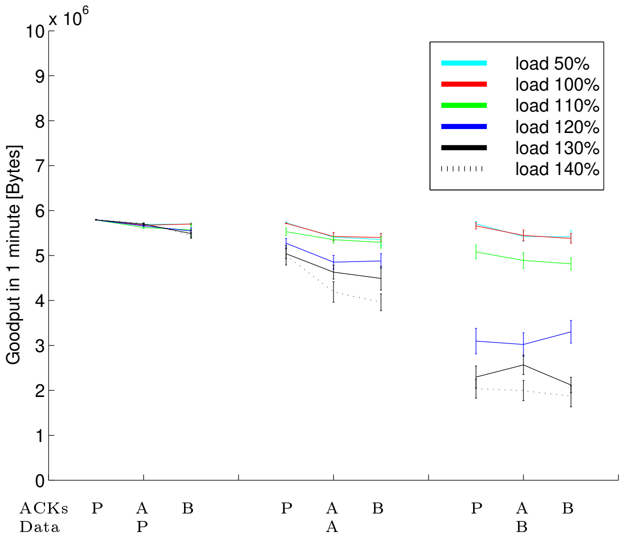

(c) RTT = 200ms (connections with different RTTs)

Figure 8: Comparison of goodput for different class combinations with varying load

A.2. Three Color Marking with Three Drop Precedence – Influence of Different Delays

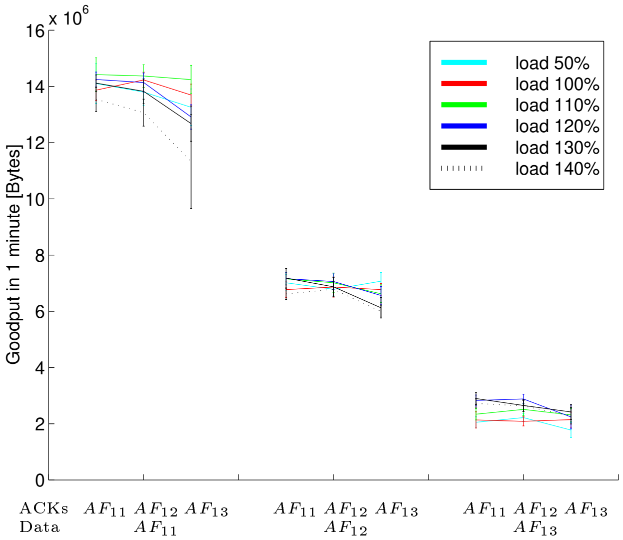

(a) RTT = 20ms (connections with different RTTs)

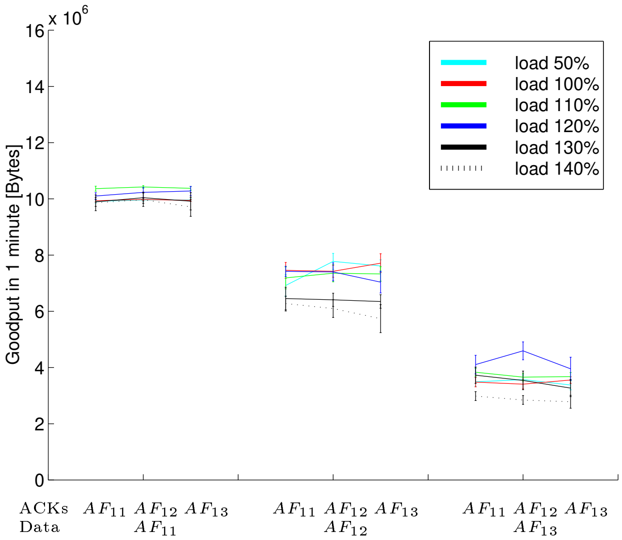

(b) RTT = 100ms (connections with different RTTs)

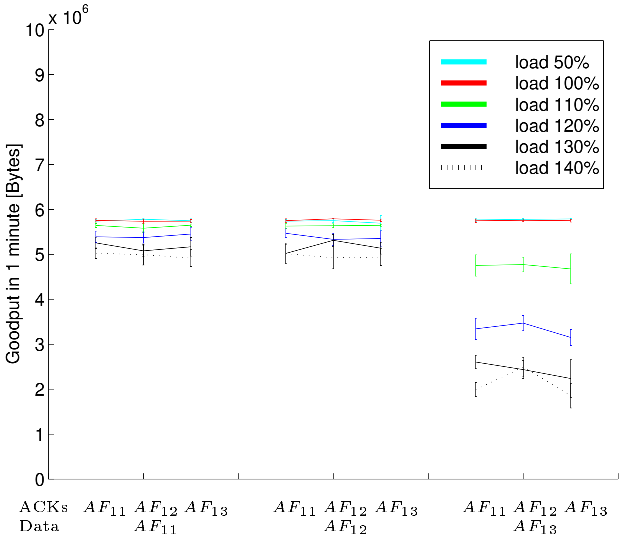

(c) RTT = 200ms (connections with different RTTs)

Figure 9: Comparison of goodput for different combinations of drop precedences with varying load