              A Rate Adaptive Shaper for Differentiated Services

Status of this Memo

Copyright Notice

Abstract

   This memo describes several Rate Adaptive Shapers (RAS) that can be
   used in combination with the single rate Three Color Markers (srTCM)
   and the two rate Three Color Marker (trTCM) described in RFC2697 and
   RFC2698, respectively. These RAS improve the performance of TCP when
   a TCM is used at the ingress of a diffserv network by reducing the
   burstiness of the traffic. With TCP traffic, this reduction of the
   burstiness is accompanied by a reduction of the number of marked
   packets and by an improved TCP goodput. The proposed RAS can be used
   at the ingress of Diffserv networks providing the Assured Forwarding
   Per Hop Behavior (AF PHB). They are especially useful when a TCM is
   used to mark traffic composed of a small number of TCP connections.

1. Introduction

   In DiffServ networks [RFC2475], the incoming data traffic, with the
   AF PHB in particular, could be subject to marking where the purpose
   of this marking is to provide a low drop probability to a minimum
   part of the traffic whereas the excess will have a larger drop
   probability. Such markers are mainly token bucket based such as the
   single rate Three Color Marker (srTCM) and two rate Three Color
   Marker (trTCM) described in [RFC2697] and [RFC2698], respectively.

   Similar markers were proposed for ATM networks and simulations have
   shown that their performance with TCP traffic was not always
   satisfactory and several researchers have shown that these
   performance problems could be solved in two ways:

1. increasing the burst size, i.e. increasing the Committed Burst
   Size (CBS) and the Peak Burst Size (PBS) in case of the trTCM, or

2. shaping the traffic such that a part of the burstiness is removed.

The first solution has as major disadvantage that the traffic sent to
the network can be very bursty and thus engineering the network to
provide a low packet loss ratio can become difficult. To efficiently
support bursty traffic, additional resources such as buffer space are
needed. Conversely, the major disadvantage of shaping is that the
traffic encounters additional delay in the shaper's buffer.

In this document, we propose two shapers that can reduce the
burstiness of the traffic upstream of a TCM. By reducing the
burstiness of the traffic, the adaptive shapers increase the
percentage of packets marked as green by the TCM and thus the overall
goodput of the users attached to such a shaper.

Such rate adaptive shapers will probably be useful at the edge of the
network (i.e. inside access routers or even network adapters). The
simulation results in [Cnodder] show that these shapers are
particularly useful when a small number of TCP connections are
processed by a TCM.

The structure of this document follows the structure proposed in
[Nichols]. We first describe two types of rate adaptive shapers in
section two. These shapers correspond to respectively the srTCM and
the trTCM. In section 3, we describe an extension to the simple
shapers that can provide a better performance. We briefly discuss
simulation results in the appendix.

2. Description of the rate adaptive shapers

2.1. Rate adaptive shaper

The rate adaptive shaper is based on a similar shaper proposed in
[Bonaventure] to improve the performance of TCP with the Guaranteed
Frame Rate [TM41] service category in ATM networks. Another type of
rate adaptive shaper suitable for differentiated services was briefly
discussed in [Azeem]. A RAS will typically be used as shown in
figure 1 where the meter and the marker are the TCMs proposed in
[RFC2697] and [RFC2698].

```
                                 Result
                              +----------+
                              |          |
                              |          |
                              |          V
           +--------+     +-------+     +--------+
 Incoming  |        |     |       |     |        |  Outgoing
 Packet ==>|  RAS   |==>  | Meter |==>  | Marker |==>Packet
 Stream    |        |     |       |     |        |  Stream
           +--------+     +-------+     +--------+
```
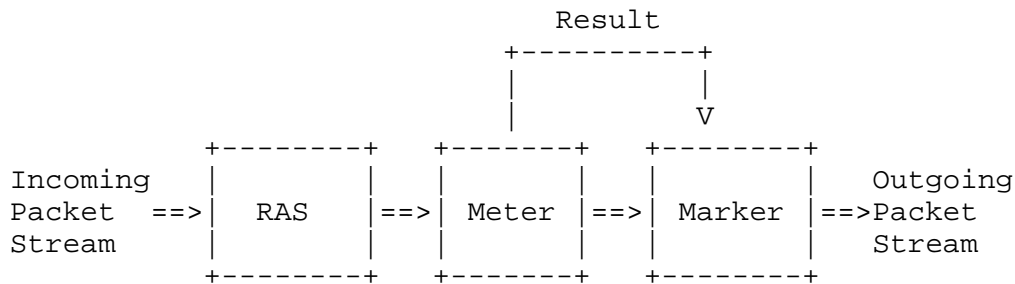
Figure 1. Rate adaptive shaper

The presentation of the rate adaptive shapers in Figure 1 is somewhat
different as described in [RFC2475] where the shaper is placed after
the meter. The main objective of the shaper is to produce at its
output a traffic that is less bursty than the input traffic, but the
shaper avoids to discard packets in contrast with classical token
bucket based shapers. The shaper itself consists of a tail-drop FIFO
queue which is emptied at a variable rate. The shaping rate, i.e.
the rate at which the queue is emptied, is a function of the
occupancy of the FIFO queue. If the queue occupancy increases, the
shaping rate will also increase in order to prevent loss and too
large delays through the shaper. The shaping rate is also a function
of the average rate of the incoming traffic. The shaper was designed
to be used in conjunction with meters such as the TCMs proposed in
[RFC2697] and [RFC2698].

There are two types of rate adaptive shapers. The single rate rate
adaptive shaper (srRAS) will typically be used upstream of a srTCM
while the two rates rate adaptive shaper (trRAS) will usually be used
upstream of a trTCM.

2.2. Configuration of the srRAS

The srRAS is configured by specifying four parameters: the Committed
Information Rate (CIR), the Maximum Information Rate (MIR) and two
buffer thresholds: CIR_th (Committed Information Rate threshold) and
MIR_th (Maximum Information Rate threshold). The CIR shall be
specified in bytes per second and MUST be configurable. The MIR
shall be specified in the same unit as the CIR and SHOULD be
configurable. To achieve a good performance, the CIR of a srRAS will
usually be set to the same value as the CIR of the downstream srTCM.
A typical value for the MIR would be the line rate of the output link
of the shaper. When the CIR and optionally the MIR are configured,
the srRAS MUST ensure that the following relation is verified:

        CIR <= MIR <= line rate

   The two buffer thresholds, CIR_th and MIR_th shall be specified in
   bytes and SHOULD be configurable. If these thresholds are
   configured, then the srRAS MUST ensure that the following relation
   holds:

        CIR_th <= MIR_th <= buffer size of the shaper

   The chosen values for CIR_th and MIR_th will usually depend on the
   values chosen for CBS and PBS in the downstream srTCM. However, this
   dependency does not need to be standardized.

2.3. Behavior of the srRAS

   The output rate of the shaper is based on two factors. The first one
   is the (long term) average rate of the incoming traffic. This
   average rate can be computed by several means. For example, the
   function proposed in [Stoica] can be used (i.e. EARnew = [(1-exp(-
   T/K))*L/T] + exp(-T/K)*EARold where EARold is the previous value of
   the Estimated Average Rate, EARnew is the updated value, K a
   constant, L the size of the arriving packet and T the amount of time
   since the arrival of the previous packet). Other averaging functions
   can be used as well.

   The second factor is the instantaneous occupancy of the FIFO buffer
   of the shaper. When the buffer occupancy is below CIR_th, the output
   rate of the shaper is set to the maximum of the estimated average
   rate (EAR(t)) and the CIR. This ensures that the shaper buffer will
   be emptied at least at a rate equal to CIR. When the buffer
   occupancy increases above CIR_th, the output rate of the shaper is
   computed as the maximum of the EAR(t) and a linear function F of the
   buffer occupancy for which F(CIR_th)=CIR and F(MIR_th)=MIR. When the
   buffer occupancy reaches the MIR_th threshold, the output rate of the
   shaper is set to the maximum information rate. The computation of
   the shaping rate is illustrated in figure 2. We expect that real
   implementations will only use an approximate function to compute the
   shaping rate.

```
                     ^
    Shaping rate     |
                     |
                     |
           MIR       |                         =========
                     |                      //
                     |                    //
         EAR(t)      |----------------//
                     |              //
                     |            //
           CIR       |===========
                     |
                     |
                     |
                     |-----------+---------+---------------------->
                            CIR_th       MIR_th Buffer occupancy
```
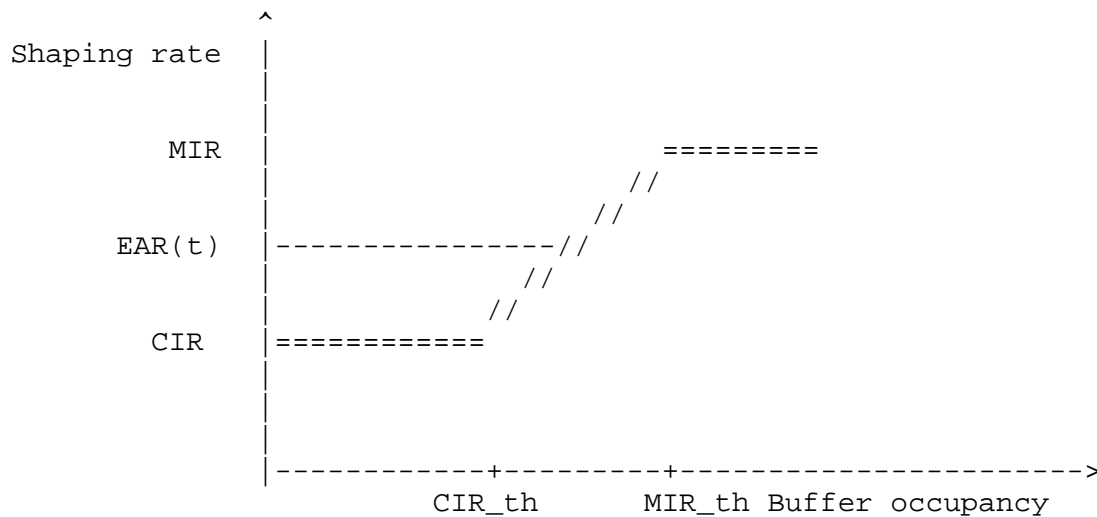
Figure 2. Computation of shaping rate for srRAS

2.4. Configuration of the trRAS

   The trRAS is configured by specifying six parameters: the Committed
   Information Rate (CIR), the Peak Information Rate (PIR), the Maximum
   Information Rate (MIR) and three buffer thresholds: CIR_th, PIR_th
   and MIR_th. The CIR shall be specified in bytes per second and MUST
   be configurable. To achieve a good performance, the CIR of a trRAS
   will usually be set at the same value as the CIR of the downstream
   trTCM. The PIR shall be specified in the same unit as the CIR and
   MUST be configurable. To achieve a good performance, the PIR of a
   trRAS will usually be set at the same value as the PIR of the
   downstream trRAS. The MIR SHOULD be configurable and shall be
   specified in the same unit as the CIR. A typical value for the MIR
   will be the line rate of the output link of the shaper. When the
   values for CIR, PIR and optionally MIR are configured, the trRAS MUST
   ensure that the following relation is verified:

           CIR <= PIR <= MIR <= line rate

   The three buffer thresholds, CIR_th, PIR_th and MIR_th shall be
   specified in bytes and SHOULD be configurable. If these thresholds
   are configured, then the trRAS MUST ensure that the following
   relation is verified:

           CIR_th <= PIR_th <= MIR_th <= buffer size of the shaper

   The CIR_th, PIR_th and MIR_th will usually depend on the values
   chosen for the CBS and the PBS in the downstream trTCM. However,
   this dependency does not need to be standardized.

2.5. Behavior of the trRAS

   The output rate of the trRAS is based on two factors. The first is
   the (long term) average rate of the incoming traffic. This average
   rate can be computed as for the srRAS.

   The second factor is the instantaneous occupancy of the FIFO buffer
   of the shaper. When the buffer occupancy is below CIR_th, the output
   rate of the shaper is set to the maximum of the estimated average
   rate (EAR(t)) and the CIR. This ensures that the shaper will always
   send traffic at least at the CIR. When the buffer occupancy
   increases above CIR_th, the output rate of the shaper is computed as
   the maximum of the EAR(t) and a piecewise linear function F of the
   buffer occupancy. This piecewise function can be defined as follows.
   The first piece is between zero and CIR_th where F is equal to CIR.
   This means that when the buffer occupancy is below a certain
   threshold CIR_th, the shaping rate is at least CIR. The second piece
   is between CIR_th and PIR_th where F increases linearly from CIR to
   PIR. The third part is from PIR_th to MIR_th where F increases
   linearly from PIR to the MIR and finally when the buffer occupancy is
   above MIR_th, the shaping rate remains constant at the MIR. The
   computation of the shaping rate is illustrated in figure 3. We
   expect that real implementations will use an approximation of the
   function shown in this figure to compute the shaping rate.

```
                       ^
     Shaping rate      |
                       |
             MIR       |                              ======
                       |                          ///
                       |                       ///
             PIR       |                    ///
                       |                  //
                       |                //
           EAR(t)      |----------------//
                       |            //
                       |          //
             CIR       |===========
                       |
                       |
                       |
                       |------------+---------+--------+-------------------->
                            CIR_th      PIR_th   MIR_th  Buffer occupancy
```
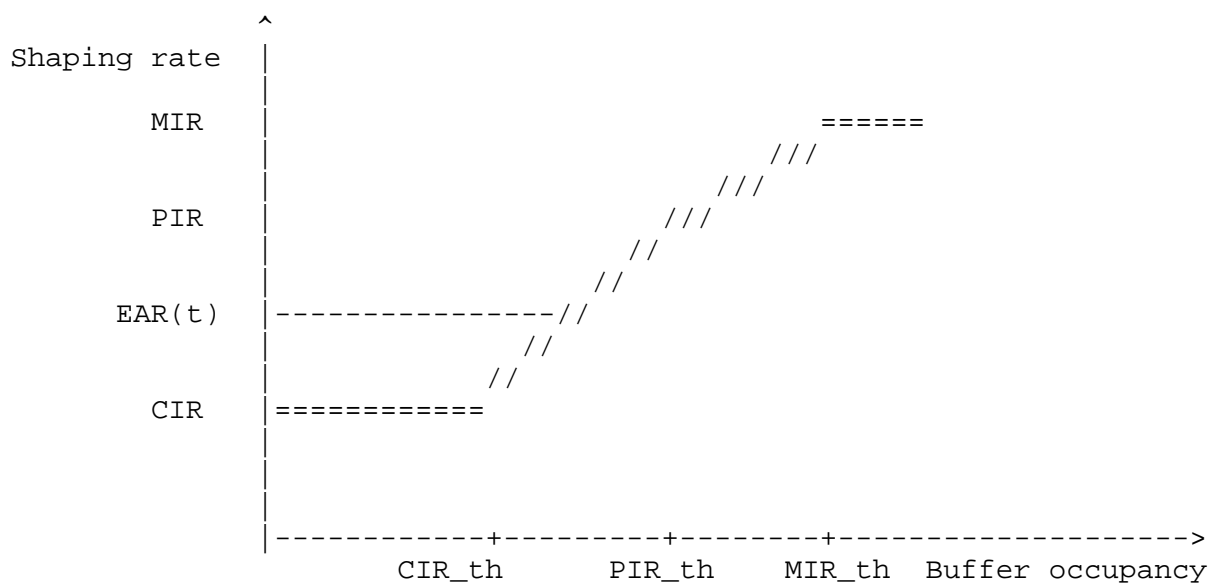
                Figure 3. Computation of shaping rate for trRAS

3. Description of the green RAS.

3.1. The green rate adaptive shapers

   The srRAS and the trRAS described in the previous section are not
   aware of the status of the meter. This entails that a RAS could
   unnecessarily delay a packet although there are sufficient tokens
   available to color the packet green. This delay could mean that TCP
   takes more time to increase its congestion window and this may lower
   the performance with TCP traffic. The green RAS shown in figure 4
   solves this problem by coupling the shaper with the meter.

```
                      Status          Result
                  +----------+ +----------+
                  |          | | |        |
                  |          | | |        |
                  V          | |          V
             +--------+   +-------+   +--------+
  Incoming   | green  |   |       |   |        |   Outgoing
  Packet ==>|  RAS   |==>| Meter |==>| Marker |==>Packet
  Stream     |        |   |       |   |        |   Stream
             +--------+   +-------+   +--------+
```
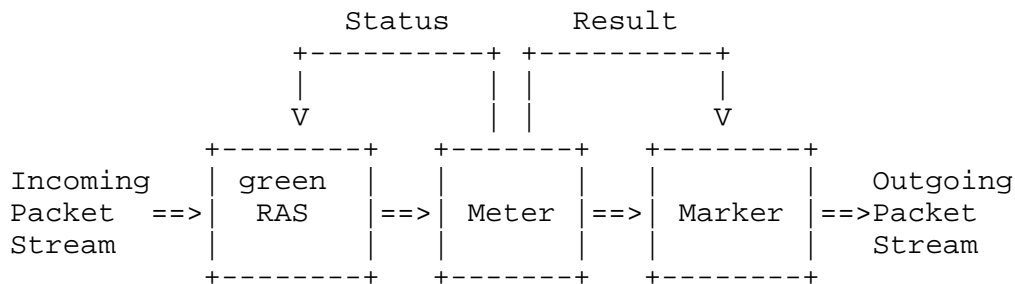
                   Figure 4. green RAS

   The two rate adaptive shapers described in section 2 calculate a
   shaping rate, which is defined as the maximum of the estimated
   average incoming data rate and some function of the buffer occupancy.
   Using this shaping rate, the RAS computes the time schedule at which
   the packet at the head of the queue of the shaper is to be released.
   The main idea of the green RAS is to couple the shaper with the
   downstream meter so that the green RAS knows at what time the packet
   at the head of its queue would be accepted as green by the meter. If
   this time instant is earlier than the release time computed from the
   current shaping rate, then the packet can be released at this time
   instant. Otherwise, the packet at the head of the queue of the green
   RAS will be released at the time instant calculated from the current
   shaping rate.

3.2. Configuration of the Green single rate Rate Adaptive Shaper
     (GsrRAS)

   The G-srRAS must be configured in the same way as the srRAS (see
   section 2.2).

3.3. Behavior of the G-srRAS

   First of all, the shaping rate of the G-srRAS is calculated in the
   same way as for the srRAS. With the srRAS, this shaping rate
   determines a time schedule, T1, at which the packet at the head of
   the queue is to be released from the shaper.

   A second time schedule, T2, is calculated as the earliest time
   instant at which the packet at the head of the shaper's queue would
   be colored as green by the downstream srTCM. Suppose that a packet
   of size B bytes is at the head of the shaper and that CIR is the
   Committed Information Rate of the srTCM in bytes per second. If we
   denote the current time by t and by $Tc(t)$ the amount of green tokens
   in the token bucket of the srTCM at time t, then T2 is equal to
   $\max(t, t+(B-Tc(t))/CIR)$. If B is larger than CBS, the Committed
   Burst Size of the srTCM, then T2 is set to infinity.

   When a packet arrives at the head of the queue of the shaper, it will
   leave this queue not sooner than $\min(T1, T2)$ from the shaper.

3.4 Configuration of the Green two rates Rate Adaptive Shaper (G-trRAS)

   The G-trRAS must be configured in the same way as the trRAS (see
   section 2.4).

3.5. Behavior of the G-trRAS

   First of all, the shaping rate of the G-trRAS is calculated in the
   same way as for the trRAS. With the trRAS, this shaping rate
   determines a time schedule, T1, at which the packet at the head of
   the queue is to be released from the shaper.

   A second time schedule, T2, is calculated as the earliest time
   instant at which the packet at the head of the shaper's queue would
   be colored as green by the downstream trTCM. Suppose that a packet
   of size B bytes is at the head of the shaper and that CIR is the
   Committed Information Rate of the srTCM in bytes per second. If we
   denote the current time by t and by $Tc(t)$ (resp. $Tp(t)$) the amount of
   green (resp. yellow) tokens in the token bucket of the trTCM at time
   t, then T2 is equal to $\max(t, t+(B-Tc(t))/CIR, t+(B-Tp(t))/PIR)$. If B
   is larger than CBS, the committed burst size, or PBS, the peak burst
   size, of the srTCM, then T2 is set to infinity.

   When a packet arrives at the head of the queue of the shaper, it will
   leave this queue not sooner than $\min(T1, T2)$ from the shaper.

4. Assumption

   The shapers discussed in this document assume that the Internet
   traffic is dominated by protocols such as TCP that react
   appropriately to congestion by decreasing their transmission rate.

   The proposed shapers do not provide a performance gain if the traffic
   is composed of protocols that do not react to congestion by
   decreasing their transmission rate.

5. Example services

   The shapers discussed in this document can be used where the TCMs
   proposed in [RFC2697] and [RFC2698] are used. In fact, simulations
   briefly discussed in Appendix A show that the performance of TCP can
   be improved when a rate adaptive shaper is used upstream of a TCM.
   We expect such rate adaptive shapers to be particularly useful at the
   edge of the network, for example inside (small) access routers or
   even network adapters.

6. The rate adaptive shaper combined with other markers

   This document explains how the idea of a rate adaptive shaper can be
   combined with the srTCM and the trTCM. This resulted in the srRAS
   and the G-srRAS for the srTCM and in the trRAS and the G-trRAS for
   the trTCM. Similar adaptive shapers could be developed to support
   other traffic markers such as the Time Sliding Window Three Color
   Marker (TSWTCM) [Fang]. However, the exact definition of such new
   adaptive shapers and their performance is outside the scope of this
   document.

7. Security Considerations

   The shapers described in this document have no known security
   concerns.

8. Intellectual Property Rights

   The IETF has been notified of intellectual property rights claimed in
   regard to some or all of the specification contained in this
   document. For more information consult the online list of claimed
   rights.

9. Acknowledgement

   We would like to thank Emmanuel Desmet for his comments.

10. References

    [Azeem]   Azeem, F., Rao, A., Lu, X. and S. Kalyanaraman, "TCP-
                    Friendly Traffic Conditioners for Differentiated
                    Services", Work in Progress.

    [RFC2475]   Blake S., Black, D., Carlson, M., Davies, E., Wang, Z.
                    and W. Weiss, "An Architecture for Differentiated
                    Services", RFC 2475, December 1998.

    [Bonaventure]   Bonaventure O., "Integration of ATM under TCP/IP to
                    provide services with a minimum guaranteed bandwidth",
                    Ph. D. thesis, University of Liege, Belgium, September
                    1998.

    [Clark]   Clark D. and Fang, W., "Explicit Allocation of Best-
                    Effort Packet Delivery Service", IEEE/ACM Trans. on
                    Networking, Vol. 6, No. 4, August 1998.

    [Cnodder]   De Cnodder S., "Rate Adaptive Shapers for Data Traffic
                    in DiffServ Networks", NetWorld+Interop 2000 Engineers
                    Conference, Las Vegas, Nevada, USA, May 10-11, 2000.

    [Fang]   Fang W., Seddigh N. and B. Nandy, "A Time Sliding
                    Window Three Colour Marker (TSWTCM)", RFC 2859, June
                    2000.

    [Floyd]   Floyd S. and V. Jacobson, "Random Early Detection
                    Gateways for Congestion Avoidance", IEEE/ACM
                    Transactions on Networking, August 1993.

    [RFC2697]   Heinanen J. and R. Guerin, "A Single Rate Three Color
                    Marker", RFC 2697, September 1999.

    [RFC2698]   Heinanen J. and R. Guerin, "A Two Rate Three Color
                    Marker", RFC 2698, September 1999.

    [RFC2597]   Heinanen J., Baker F., Weiss W. and J. Wroclawski,
                    "Assured Forwarding PHB Group", RFC 2597, June 1999.

    [Nichols]   Nichols K. and B. Carpenter, "Format for Diffserv
                    Working Group Traffic Conditioner Drafts", Work in
                    Progress.

   [Stoica]   Stoica I., Shenker S. and H. Zhang, "Core-stateless
                   fair queueuing: achieving approximately fair bandwidth
                   allocations in high speed networks", ACM SIGCOMM98, pp.
                   118-130, Sept. 1998

   [TM41]   ATM Forum, Traffic Management Specification, verion
                   4.1, 1999

Appendix

A. Simulation results

   We briefly discuss simulations showing the benefits of the proposed
   shapers in simple network environments. Additional simulation results
   may be found in [Cnodder].

A.1 description of the model

   To evaluate the rate adaptive shaper through simulations, we use the
   simple network model depicted in Figure A.1. In this network, we
   consider that a backbone network is used to provide a LAN
   Interconnection service to ten pairs of LANs. Each LAN corresponds
   to an uncongested switched 10 Mbps LAN with ten workstations attached
   to a customer router (C1-C10 in figure A.1). The delay on the LAN
   links is set to 1 msec. The MSS size of the workstations is set to
   1460 bytes. The workstations on the left hand side of the figure
   send traffic to companion workstations located on the right hand side
   of the figure. All traffic from the LAN attached to customer router
   C1 is sent to the LAN attached to customer router C1'. There are ten
   workstations on each LAN and each workstation implements SACK-TCP
   with a maximum window size of 64 KBytes.

```
       2.5 msec, 34 Mbps                        2.5 msec, 34 Mbps
        <-------------->                         <-------------->
   \+---+                                                    +---+/
   -| C1|--------------+                        +------------|C1'|-
   /+---+              |                        |            +---+\
   \+---+              |                        |            +---+/
   -| C2|-----------+  |                        |  +---------|C2'|-
   /+---+           |  |                        |  |         +---+\
   \+---+           |  |                        |  |         +---+/
   -| C3|---------+ |  |                        |  | +-------|C3'|-
   /+---+         | |  |                        |  | |       +---+\
   \+---+         | |  |                        |  | |       +---+/
   -| C4|-------+ +-+--------+      +---------+-+ +--------|C4'|-
   /+---+       |   |        |      |         | |         +---+\
   \+---+       +---|        |      |         | |---+      +---+/
   -| C5|-----------|  ER1   |-----|   ER2   |-----------|C5'|-
   /+---+       +---|        |      |         | |---+      +---+\
   \+---+       |   |        |      |         | |         +---+/
   -| C6|-------+   +--------+      +---------+   +--------|C6'|-
   /+---+       ||||                          ||||        +---+\
   \+---+       ||||    <------->             ||||        +---+/
   -| C7|------------+|||    70 Mbps          |||+--------|C7'|-
   /+---+          |||    10 msec             |||         +---+\
   \+---+          |||                        |||         +---+/
   -| C8|-------------+|                      ||+---------|C8'|-
   /+---+           ||                        ||          +---+\
   \+---+           ||                        ||          +---+/
   -| C9|--------------+                      |+----------|C9'|-
   /+---+            |                        |           +---+\
   \+---+            |                        |           +----+/
   -|C10|---------------+                     +-----------|C10'|-
   /+---+                                                 +----+\
```
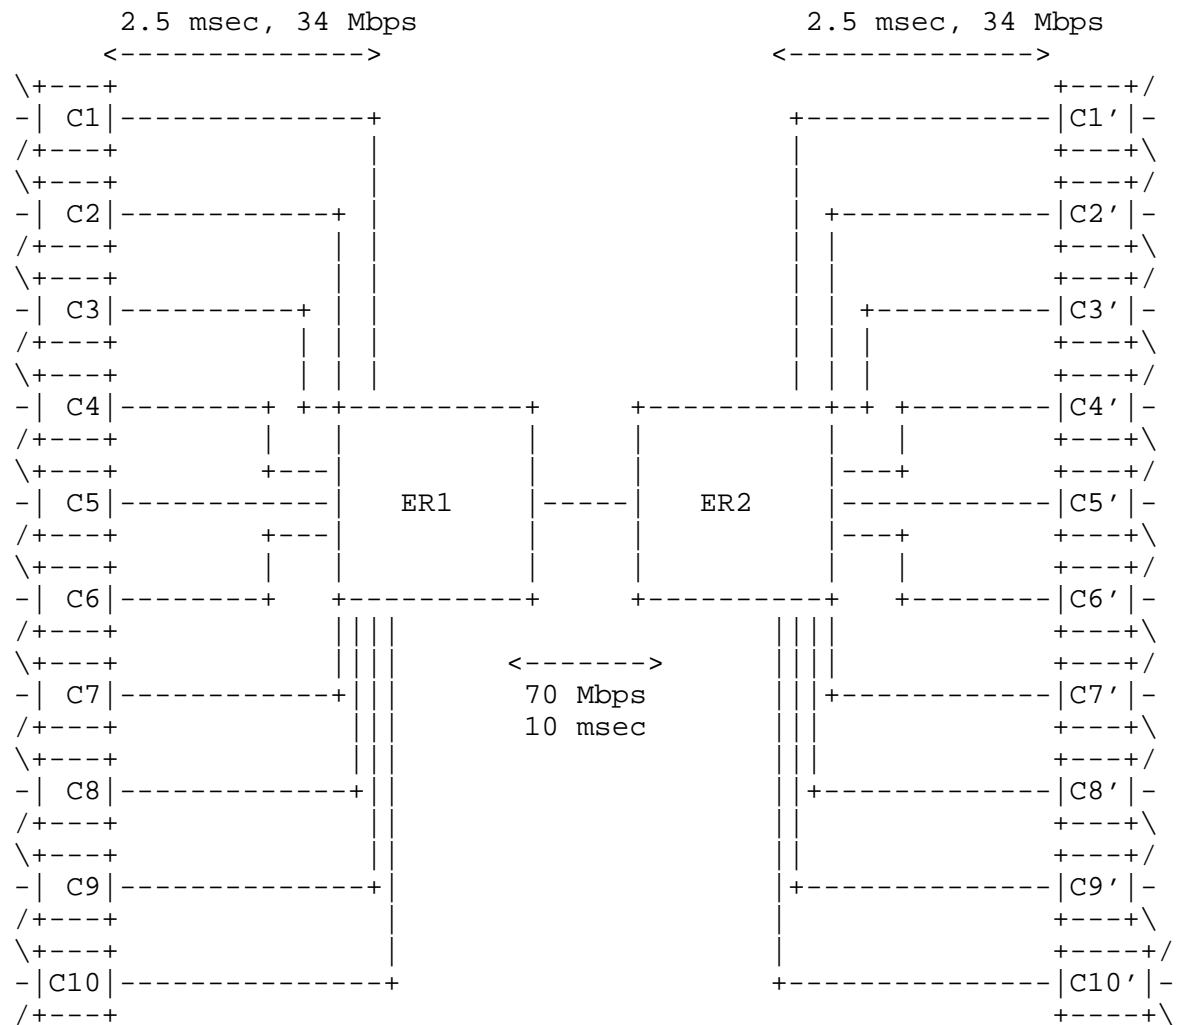
<div align="center">Figure A.1. the simulation model.</div>

The customer routers are connected with 34 Mbps links to the backbone
network which is, in our case, composed of a single bottleneck 70
Mbps link between the edge routers ER1 and ER2. The delay on all the
customer-edge 34 Mbps links has been set to 2.5 msec to model a MAN
or small WAN environment. These links and the customer routers are
not a bottleneck in our environment and no losses occurs inside the
edge routers. The customer routers are equipped with a trTCM
[Heinanen2] and mark the incoming traffic. The parameters of the
trTCM are shown in table A.1.

Table A.1: configurations of the trTCMs

Router	CIR	PIR	Line Rate
C1	2 Mbps	4 Mbps	34 Mbps
C2	4 Mbps	8 Mbps	34 Mbps
C3	6 Mbps	12 Mbps	34 Mbps
C4	8 Mbps	16 Mbps	34 Mbps
C5	10 Mbps	20 Mbps	34 Mbps
C6	2 Mbps	4 Mbps	34 Mbps
C7	4 Mbps	8 Mbps	34 Mbps
C8	6 Mbps	12 Mbps	34 Mbps
C9	8 Mbps	16 Mbps	34 Mbps
C10	10 Mbps	20 Mbps	34 Mbps

All customer routers are equipped with a trTCM where the CIR are 2
Mbps for router C1 and C6, 4 Mbps for C2 and C7, 6 Mbps for C3 and
C8, 8 Mbps for C4 and C9 and 10 Mbps for C5 and C10. Routers C6-C10
also contain a trRAS in addition to the trTCM while routers C1-C5
only contain a trTCM. In all simulations, the PIR is always twice as
large as the CIR. Also the PBS is the double of the CBS. The CBS
will be varied in the different simulation runs.

The edge routers, ER1 and ER2, are connected with a 70 Mbps link
which is the bottleneck link in our environment. These two routers
implement the RIO algorithm [Clark] that we have extended to support
three drop priorities instead of two. The thresholds of the
parameters are 100 and 200 packets (minimum and maximum threshold,
respectively) for the red packets, 200 and 400 packets for the yellow
packets and 400 and 800 for the green packets. These thresholds are
reasonable since there are 100 TCP connections crossing each edge
router. The parameter maxp of RIO for green, yellow and red are
respectively set to 0.02, 0.05, and 0.1. The weight to calculate the
average queue length which is used by RED or RIO is set to 0.002
[Floyd].

The simulated time is set to 102 seconds where the first two seconds
are not used to gather TCP statistics (the so-called warm-up time)
such as goodput.

A.2 Simulation results for the trRAS

For our first simulations, we consider that routers C1-C5 only
utilize a trTCM while routers C6-C10 utilize a rate adaptive shaper
in conjunction with a trTCM. All routers use a CBS of 3 KBytes. In
table A.2, we show the total throughput achieved by the workstations
attached to each LAN as well as the total throughput for the green
and the yellow packets as a function of the CIR of the trTCM used on
the customer router attached to this LAN. The throughput of the red

packets is equal to the difference between the total traffic and the
green and the yellow traffic. In table A.3, we show the total
throughput achieved by the workstations attached to customer routers
with a rate adaptive shaper.

Table A.2: throughput in Mbps for the unshaped traffic.

	green	yellow	total
2Mbps [C1]	1.10	0.93	2.25
4Mbps [C2]	2.57	1.80	4.55
6Mbps [C3]	4.10	2.12	6.39
8Mbps [C4]	5.88	2.32	8.33
10Mbps [C5]	7.57	2.37	10.0

Table A.3: throughput in Mbps for the adaptively shaped
traffic.

	green	yellow	total
2Mbps [C6]	2.00	1.69	3.71
4Mbps [C7]	3.97	2.34	6.33
6Mbps [C8]	5.93	2.23	8.17
8Mbps [C9]	7.84	2.28	10.1
10Mbps [C10]	9.77	2.14	11.9

This first simulation shows clearly that the workstations attached to
an edge router with a rate adaptive shaper have a clear advantage,
from a performance point of view, with respect to workstations
attached to an edge router with only a trTCM. The performance
improvement is the result of the higher proportion of packets marked
as green by the edge routers when the rate adaptive shaper is used.

To evaluate the impact of the CBS on the TCP goodput, we did
additional simulations were we varied the CBS of all customer
routers.

Table A.4 shows the total goodput for workstations attached to,
respectively, routers C1 (trTCM with 2 Mbps CIR, no adaptive
shaping), C6 (trRAS with 2 Mbps CIR and adaptive shaping), C3 (trTCM
with 6 Mbps CIR, no adaptive shaping), and C8 (trRAS with 6 Mbps CIR
and adaptive shaping) for various values of the CBS. From this
table, it is clear that the rate adaptive shapers provide a
performance benefit when the CBS is small. With a very large CBS,
the performance decreases when the shaper is in use. However, a CBS
of a few hundred KBytes is probably too large in many environments.

Table A.4: goodput in Mbps (link rate is 70 Mbps) versus CBS
in KBytes.

CBS	2_Mbps_unsh	2_Mbps_sh	6_Mbps_unsh	6_Mbps_sh
3	1.88	3.49	5.91	7.77
10	2.97	2.91	6.76	7.08
25	3.14	2.78	7.07	6.73
50	3.12	2.67	7.20	6.64
75	3.18	2.56	7.08	6.58
100	3.20	2.64	7.00	6.62
150	3.21	2.54	7.11	6.52
200	3.26	2.57	7.07	6.53
300	3.19	2.53	7.13	6.49
400	3.13	2.48	7.18	6.43

A.3 Simulation results for the Green trRAS

We use the same scenario as in A.2 but now we use the Green trRAS
(G-trRAS).

Table A.5 and Table A.6 show the results of the same scenario as for
Table A.2 and Table A.3 but the shaper is now the G-trRAS. We see
that the shaped traffic performs again much better, also compared to
the previous case (i.e. where the trRAS was used). This is because
the amount of yellow traffic increases with the expense of a slight
decrease in the amount of green traffic. This can be explained by
the fact that the G-trRAS introduces some burstiness.

Table A.5: throughput in Mbps for the unshaped traffic.

	green	yellow	total
2Mbps [C1]	1.10	0.95	2.26
4Mbps [C2]	2.41	1.66	4.24
6Mbps [C3]	3.94	1.97	6.07
8Mbps [C4]	5.72	2.13	7.96
10Mbps [C5]	7.25	2.29	9.64

Table A.6: throughput in Mbps for the adaptively shaped
traffic.

	green	yellow	total
2Mbps [C6]	1.92	1.75	3.77
4Mbps [C7]	3.79	3.24	7.05
6Mbps [C8]	5.35	3.62	8.97
8Mbps [C9]	6.96	3.48	10.4
10Mbps [C10]	8.69	3.06	11.7

The impact of the CBS is shown in Table A.7 which is the same
scenario as Table A.4 with the only difference that the shaper is now
the G-trRAS. We see that the shaped traffic performs much better
than the unshaped traffic when the CBS is small. When the CBS is

large, the shaped and unshaped traffic performs more or less the
same. This is in contrast with the trRAS, where the performance of
the shaped traffic was slightly worse in case of a large CBS.

Table A.7: goodput in Mbps (link rate is 70 Mbps) versus CBS
in KBytes.

CBS	2_Mbps_unsh	2_Mbps_sh	6_Mbps_unsh	6_Mbps_sh
3	1.90	3.44	5.62	8.44
10	2.95	3.30	6.70	7.20
25	2.98	3.01	7.03	6.93
50	3.06	2.85	6.81	6.84
75	3.08	2.80	6.87	6.96
100	2.99	2.78	6.85	6.88
150	2.98	2.70	6.80	6.81
200	2.96	2.70	6.82	6.97
300	2.94	2.70	6.83	6.86
400	2.86	2.62	6.83	6.84

A.4 Conclusion simulations

From these simulations, we see that the shaped traffic has much
higher throughput compared to the unshaped traffic when the CBS was
small. When the CBS is large, the shaped traffic performs slightly
less than the unshaped traffic due to the delay in the shaper. The
G-trRAS solves this problem. Additional simulation results may be
found in [Cnodder]

Authors' Addresses

    Olivier Bonaventure
    Infonet research group
    Institut d'Informatique (CS Dept)
    Facultes Universitaires Notre-Dame de la Paix
    Rue Grandgagnage 21, B-5000 Namur, Belgium.

    EMail: Olivier.Bonaventure@info.fundp.ac.be
    URL: http://www.infonet.fundp.ac.be

    Stefaan De Cnodder
    Alcatel Network Strategy Group
    Fr. Wellesplein 1, B-2018 Antwerpen, Belgium.

    Phone: 32-3-240-8515
    Fax: 32-3-240-9932
    EMail: stefaan.de_cnodder@alcatel.be

Full Copyright Statement

Acknowledgement