Performance Analysis of the Counters-Based Modified Traffic Conditioner in a DiffServ Network

Maria-Dolores Cano, Fernando Cerdan, Joan Garcia-Haro, Josemaria Malgosa-Sanahuja
Department of Information Technologies and Communications
Polytechnic University of Cartagena
Campus Muralla del Mar s/n (Ed. Hospital de Marina) 30202 Cartagena, Spain
Ph. +34 968 325953 Fax +34 968 325338
{mdolores.cano, fernando.cerdan, joang.haro, josem.malgosa}@upct.es

Abstract

The Counters-Based Modified (CBM) traffic conditioner was presented in a previous work as a feasible election for the Assured Service implementation with DiffServ. In this work, we evaluate the end-to-end performance of TCP Reno sources with the CBM traffic conditioner in topologies with one and three network nodes under miscellaneous characteristics, i. e. variable contracted target rates and round trip times. In addition and for the same scenarios, the robustness of the couple CBM-RIO is tested when Assured Service and best-effort sources coexist and compete for network resources. Strong assurance of achieving target bandwidths and even allocation of spare bandwidth across different flows can be provided as shown in simulation results.

1. Introduction

Differentiated Services (DiffServ) paradigm [1] has been standardized as one of the promising solutions for providing (Quality of Service) QoS in IP networks. The DiffServ architecture tries to create a simple scheme that provides a range of QoS levels by moving complexity toward the edge of the network. Packets are classified and marked to receive a particular treatment on the nodes along their path. This treatment is known as per-hop behavior (PHB). Complex classification and conditioning functions (metering, marking, shaping) need only to be implemented at boundary nodes. Meanwhile, interior nodes perform a set of forwarding PHBs to aggregates of traffic that have been appropriately marked.

One of the PHBs standardized by the IETF is the Assured Forwarding per-hop behavior (AF PHB) [2]. The idea behind AF PHB is to assure a minimum throughput (the target rate or contracted rate) to a connection, while enabling consuming more bandwidth if the network load is low. This distribution of the excess bandwidth should be done in a fair way, where two different concepts can be understood as fairness. The first considers fairness as the even distribution of excess bandwidth among all connections that compose the aggregate. The second defines fairness as a proportional distribution of the outbound bandwidth with respect to the contracted rate. In this paper we adopt the first definition.To achieve the AF goals, packets of individual flows are marked belonging to one of the four independently forwarded AF classes. As detailed in [2], within each AF class an IP packet can be assigned one of three different levels of drop precedence. In case of congestion, the drop precedence of a packet determines the relative importance of the packet within the AF class.

A significant step in providing the Internet Assured Service was the introduction of RIO (RED (Random Early Detection) In and Out) in [3]. Once a packet is marked as in-of-profile (IN) or out-of-profile (OUT), the aggregate reaches the router device, where RIO is applied. RIO is the combination of two RED [4] algorithms with different drop probability curves, so that OUT packets are more likely to be discarded. RIO uses a single FIFO queue to service both IN and OUT packets. The probability of dropping an OUT packet depends on the total number of packets arriving to the node, while the probability of dropping an IN packet depends exclusively on the buffer occupancy of IN packets.

During last years, different traffic conditioners have been proposed in related literature [5][6][7][8][9]. Nevertheless, it has been demonstrated the difficulty in finding a traffic conditioner able to achieve the AF goals despite the interaction with the AF PHB implementation mechanisms. Some of the proposals cannot guarantee strictly contracted rates because of the dependence on network parameters such as Round Trip Time (RTT). Others, even under agreeable conditions where there is no diversity in network parameters, present a complex configuration; thus, contracted target rates are not ensured if any slight modification in configuration parameters is done. Additionally, there are traffic conditioners that guarantee target rates but, when distributing the excess bandwidth, they do not perform a fair share (in none of the two contemplated definitions). Last trends in the development of traffic conditioners either require the use of signaling or need per-flow-monitoring in the router with the consequent scalability problems. Moreover, even in these last cases there is a lack of fairness in the outbound bandwidth distribution among the TCP sources of the aggregate.

The Counters-Based Modified (CBM) traffic conditioner introduced in [10] is presented as an alternative approach to achieve a fair distribution of excess bandwidth among the different TCP sources that compose the aggregate in the Assured Service. CBM applies the Counters-Based (CB) mechanism [11] for packet marking. As shown in [11], CB guarantees the in-profile bandwidth allocation in scenarios with variable RTT and target rates. Once contracted rates are ensured, CBM provides a fair share of the spare bandwidth by means of probabilistically dropping OUT packets at the traffic conditioner. With CBM, complexity remains at the Assured Service capable host before the RIO buffer management scheme. CBM assumes some signaling to determine the dropping probability of an OUT packet, like the amount of excess bandwidth or the average RTT of all connections, which makes it more feasible than other proposed traffic conditioners such as [5][6][8][9].

In this paper, we study the performance throughout simulations of the couple CBM-RIO first introduced in [10]. While in [10] it is explained how the CBM mechanism works showing its behavior, here we focus on studying end-to-end TCP performance when the CBM is employed as traffic conditioner in heterogeneous topologies. The analysis in terms of assuring contracted target rates and distribution of excess bandwidth to TCP Reno sources, is firstly done in a single node topology with miscellaneous characteristics (different target rates, round trip times and share of resources with best-effort connections). CBM is also tested in heterogeneous three-node topologies to investigate its robustness in a more complex network framework, closer to real situations. As we show later in simulation results, it is possible to afford fairness in the sharing of excess bandwidth by using the CBM traffic conditioner without losing accuracy in assuring contracted rates in different DiffServ domains.

The rest of this paper is organized as follows. Section 2 reviews the characteristics of the CBM implementation. In Section 3, we present the scenario and assumptions for carrying out simulations. In Section 4, simulation results are presented and discussed. The paper concludes in Section 5 summarizing the most important facts.

2. CBM Traffic Conditioner

The Counters-Based Modified (CBM) was introduced in [10] based on the idea that if all sources introduce the same number of out-of-profile packets into the network (assuming all packets have a similar size), then each source can get the same portion of excess bandwidth. This ideal behavior is affected by the odd characteristics of each TCP connection. To confront these influences, it was suggested that connections that are sending OUT packets beyond their ideal fair quota should be penalized. This penalty was based on probabilistically dropping OUT packets in the traffic conditioner.

In [10] it was shown that connections with small target rates and low round trip delays generate more out-of-profile packets between consecutive in-of-profile packets than other connections, thus getting more network resources. From these observations, the CBM was developed to work as follows. Placed next to the TCP source (out of the reach of the final user), it has a variable that counts the number of packets that have been marked as OUT between two consecutive IN packets. Every time a packet is marked as OUT, the CBM traffic conditioner checks this variable. If the variable does not exceed a minimum value *min*, then the OUT packet is injected into the network. If it exceeds a maximum value *max*, then the OUT packet is dropped. Finally, if the variable remains between *min* and *max*, the OUT packet is dropped with probability *p*.

Accordingly, to employ CBM it is necessary to configure the *max* and *min* thresholds as well as to calculate the dropping probability *p*. As explained in [10], to tune the *max* and *min* parameters we follow equations (1) and (2), where MSS stands for Maximum Segment Size. The excess bandwidth could be seen as another TCP source whose maximum TCP window size is determined by the product BW_{excess} times $RTT_{average}$. Therefore, we set the *max* limit to this value. It is well known that in TCP/IP, a simple additive increase and multiplicative decrease algorithm satisfies the sufficient conditions for convergence to an efficient state of the network, and it is used to implement congestion avoidance schemes. For this reason, a practical *min* value is half the *max* value.

$$max = \left\lceil \frac{Bandwidth_{excess} \cdot RTT_{average}}{MSS} \right\rceil \qquad (1)$$

$$min = \left\lceil \frac{max}{2} \right\rceil \qquad (2)$$

The estimation of RTT can be obtained by periodically signaling from the router device. The TCP protocol implements an algorithm that estimates the RTT of the current connection. This estimation is periodically sent to the router device, which calculates the average RTT. This value is then returned to the traffic conditioner, where packets are marked and/or dropped. Notice that per-flow state monitoring in the router is not required, in the sense that the router does not contain information on each individual active packet flow unlike traffic conditioner implementations from [5] [6] [8] or [9].

$$p = 2 \cdot \frac{target\_rate/link\_rate}{1 + target\_rate/link\_rate} \qquad (3)$$

On the other hand, the dropping probability p is shown in equation (3). Each source has a different value of p, between 0 and 1, based on its contracted rate. The equation for the dropping probability gives a slightly more preference to connections with small targets. A deeper explanation can be found in [10]. Notice that equation (3) is only applied when the number of OUT packets is in the interval (*min*, *max*). The simplified pseudo-code of the entire CBM algorithm is written in Fig. 1.

```
 Initially:
    Counter1=1
    Counter2=link_rate/target_rate
    Counter3=0
Calculate the values for the probability p and the limits max
 and min
 For each unit of time:
    Counter2--
    If counter2 <= 0
        Counter1++
        Counter2=link_rate/target_rate
    if there is a packet arrival
        if counter1>0
            the packet is marked as IN
            counter1--
            counter3=0
        else
            the packet is marked as OUT
            counter3++
            if time>start_dropping_time
                if counter3>max
                    the OUT packet is dropped
                else if counter3>min
                    the OUT packet is dropped with probability p
                otherwise the OUT packet is accepted
```

Figure 1. Simplified pseudo-code of the CBM traffic conditioner algorithm.

3. Scenarios for Simulation

The CBM traffic conditioner is evaluated throughout simulations in a single node topology depicted in Fig. 2 and a three-node topology shown in Fig. 3. TCP traffic is generated by eight TCP Reno sources transmitting at link rate, which has been set to 33 Mbps.

The simulation tool used in this work for the sliding window protocol of TCP Reno sources was developed in [12]. It was extensively utilized in [13], and was applied to validate the analytical study carried out in [14]. Some characteristics of this simulation tool are: TCP sources have been selected as greedy for a worst case to achieve a relative high network congestion state, destinations only send acknowledgements, which are never lost or delayed, and the maximum window size equals the product bandwidth delay as usual for WAN environments.

We employ a large packet size of 9,188 bytes, which corresponds to classical IP over ATM (Asynchronous Transfer Mode) and could represent Differentiated Services over MPLS (Multi Protocol Label Switching), where the use of the ATM technology is inherent. Routers located inside the network, buffer and forward the aggregated traffic. The queue management employs RIO (i.e., twin RED algorithms to preferentially drop OUT packets). The RIO parameters are [40/70/0.02] for IN packets and [10/40/0.2][1] for OUT packets. Weight_in and Weight_out RED parameters used to calculate the average queue size have been chosen equal to 0.002 as recommended in [4].

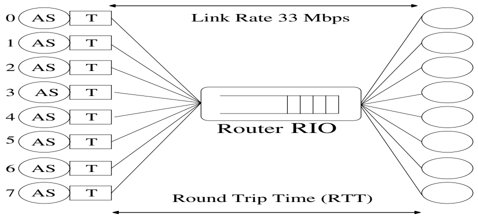

Figure 2. General single node topology used in simulations where the bottleneck is the router device (T≡Traffic Conditioner).

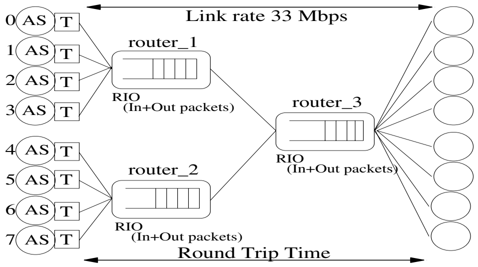

Figure 3. General three-node topology used in simulations.

We consider five different scenarios in an under-subscribed situation (traffic load ≤ 60%). The oversubscribed scenario (traffic load > 60%) is less interesting in this study since the excess bandwidth represents a very small portion of the total available bandwidth. To verify the impact of target rates and the influence of RTTs in both topologies, different values are used along simulations. Characteristics of each scenario are summarized in Table 1. Simulation results have a confidence interval of 95% that has been calculated with a normal distribution function using 30 samples, with an approximate value of ±0.002 for all fairness calculations, and ±0.01 for the achieved target rates.

4. Simulation Results

In this section, we present and discuss simulation results carried out in the scenarios described in Table 1. Firstly, it is shown how CBM leads to a fair share of the excess bandwidth in a simple topology composed of one node (one bottleneck). We also present results of the interaction of Assured Service connections with best-effort connections competing for the outbound bandwidth. Finally, we evaluate the performance of TCP crossing

---

[1] [minth, maxth, maxp]

three routers following the same steps (assuring contracts, excess bandwidth fairness and robustness when sharing resources with best-effort connections).

Table 1. Scenarios for simulations (eight TCP Reno sources).

	Target rates (Mbps)	RTT (ms)
Scenario A	2.5	50
Scenario B	1-1-2-2-3-3-4-4-	50
Scenario C	2.5	10 to 80 at intervals of 10
Scenario D	1-1-2-2-3-3-4-4	10 to 80 at intervals of 10
Scenario E	4-4-3-3-2-2-1-1	10 to 80 at intervals of 10

4.1. Single node topology with Assured Services sources

To evaluate fairness we use the fairness index *f* denoted by expression (4), where x_i is the excess throughput of source *i*, and *n* is the number of sources that compose the aggregate [15]. The closer to 1 in the *f* value, the more the fairness obtained. We use the term throughput meaning *goodput* in calculations of the fairness index. The *max* and *min* thresholds in each scenario are included in Table 2.

$$f = \frac{\left(\sum_{i=1}^{n} x_i\right)^2}{n \cdot \sum_{i=1}^{n} x_i^2}\ ;\ f \le 1 \qquad (4)$$

Table 2. *max* and *min* values used in each scenario with CBM in a single node topology (from eq. 1 and 2).

Scenario	A	B	C	D	E
Link rate (Mbps)	33	33	33	33	33
Σ target rates (Mbps)	20	20	20	20	20
BW_{excess} (Mbps)	13	13	13	13	13
$RTT_{average}$ (ms)	50	50	45	45	45
max (#OUT packets)	9	9	8	8	8
min (#OUT packets)	5	5	4	4	4

Fig. 4 depicts the different *f* values obtained from simulations in topology from Fig. 2, and compares them in the same scenarios to other traffic conditioners that do not implement probabilistic OUT packets dropping (CB, TSW and LB). Simulations for the TSW and LB traffic conditioners have been carried out taking into consideration the performance evaluation study from [11].

Fairness indexes included in Fig. 4 reveal that it is possible to assure fairness in the excess bandwidth sharing with the CBM traffic conditioner, achieving an *f* value close to 0.95. Although the LB and TSW algorithms attain a high *f* value in scenarios A and B respectively, it should be noted that using these mechanisms inbound bandwidths are not guaranteed [11]. Therefore, the underlying idea of keeping all connections sending a similar number of OUT packets is presented as a comparatively improvement in the development of traffic conditioners for the Internet Assured Service.

In addition, Fig. 5 is included to confirm that the dropping of OUT packets in the CBM traffic conditioner along with its interaction with the RIO buffer management scheme do not affect the inbound bandwidth assurance.

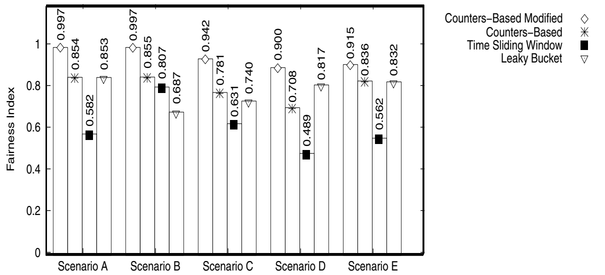

Figure 4. Fairness indexes for Counters-Based Modified (CBM), Counters-Based (CB), Time Sliding Window (TSW) and Leaky Bucket (LB) in scenarios A to E in a single node topology

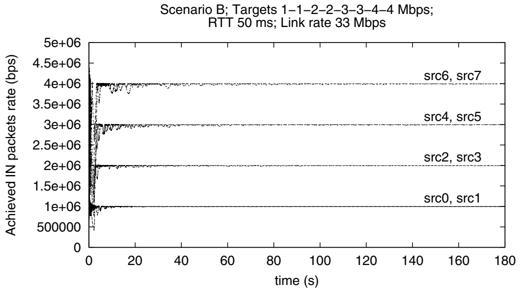

Figure 5. Achieved rates for IN packets in scenario B in a single node topology.

4.2. Single node topology with Assured Service Sources and Best-Effort Sources

In this subsection, best-effort (BE) sources compete with Assured Service (AS) sources for the available excess bandwidth. The topology used for this set of simulations corresponds to Fig. 2, eight TCP Reno connections, where the first four connections have an Assured Service but the last four connections belong to the best-effort class. The fact of being best-effort implies that all packets generated by these sources are considered as out-of-profile, and they do not have contracted target rates. We have conducted simulations for the five scenarios explained in Table 1 with slight modifications commented below. Link rate is kept at 33 Mbps.

In scenario A, the AS sources have a target rate of 5 Mbps, and all sources (included the BE ones) have a RTT of 50 ms. The limits *max* and *min* are 9 and 5 respectively. From simulation results, best-effort sources and Assured Service sources get nearly the same amount of spare bandwidth and a fairness index of 0.937 is reached. Scenario B is equal to scenario A, but the four AS connections have contracted rates of 4-5-6 and 7 Mbps each. In this case, where thresholds *max* and *min* are 8 and 4 packets, the *f* value is 0.864. In this situation, best-effort sources do not cause starvation either, and let the assured connections to get a portion of the excess bandwidth close to the ideal value.

In scenario C, the AS sources have a target rate of 5 Mbps. The RTT value ranges from 10 ms to 40 ms at increments of 10 ms (AS connections), and 50 ms to 80 ms at intervals of 10 ms (BE connections). The limits *max* and *min* take a value of 8 and 4 packets respectively. From simulation results, BE sources achieve a *goodput* close to the ideal value with a difference of 0.5 Mbps between the maximum and minimum reached *goodputs*. The effect of having different values of RTT is hardly noticeable in the distribution of the outbound bandwidth, which is reflected in a fairness index of 0.847.

Finally, the most complex scenarios D and E also present an *f* value over 0.8. In scenario D, the four AS sources have contracted rates of 4-5-6 and 7 Mbps, and a RTT that goes from 10 ms to 40 ms at intervals of 10 ms. The RTT for the BE sources ranges from 50 ms to 80 ms in increments of 10 ms. Scenario E only differs from D in the target rates of the AS connections, being in this case 7-6-5 and 4 Mbps. The limits *max* and *min* take a value of 7 and 4 packets in both scenarios. Fig. 6 shows the *goodput* of BE sources in Scenario D. In this case, the difference between the maximum and minimum reached *goodput* is less than 0.5 Mbps.

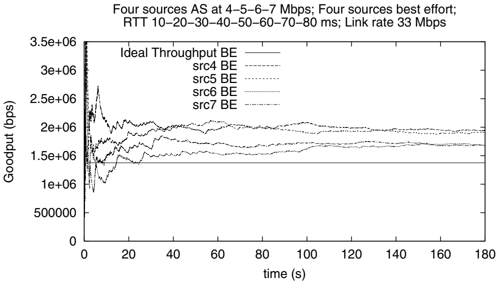

Figure 6. *Goodput* (bps) for BE sources in Scenario D.

Fairness indexes obtained in each scenario illustrate the good performance accomplished with the CBM traffic conditioner in the distribution of the excess bandwidth when AS connections and best-effort connections coexist. The best-effort users only generate out of control OUT packets that difficult the provisioning of a consistent network service. We show that the robustness of the couple CBM-RIO makes the entire service structure resistant to these users who try to maximize the bandwidth they attain from the network, since all AS sources get their target rates and also benefit from the excess bandwidth.

4.3. Three-node topologies

In this section, we evaluate the performance of TCP crossing three routers; a notably more complex and heterogeneous topology than the usual one employed in the related literature. We study three different cases from general topology of Fig 3, where node tagged as *router_1* implements the RIO mechanism with parameters [40/70/0.02] for IN packets and [10/40/0.2] for OUT packets. The node tagged as *router_2* executes RIO with the same parameters as *router_1* or the RED algorithm with parameters [10/40/0.2]. *Router_3* receives traffic from *router_1* and *router_2*, and also implements RIO with the same parameters. As indicated in Section 3, we have used recommendations from [4].

CBM traffic conditioners are placed beside the TCP sources when an Assured Service is to be contracted, otherwise sources belong to the best-effort class and their packets are treated as out-of-profile without performing traffic conditioning. The *max* and *min* values (eq. 1, 2 and 3) employed in the CBM mechanism are shown in Table 3. We calculate these values for *router_1* and *router_2*, assuming that *router_1* does not know about the existence of *router_2* and vice versa. Simulation results have a confidence interval of 95% that has been calculated as indicated in Section 3. Next, we describe the three cases under examination.

Table 3. *max* and *min* limits used in each scenario with CBM in a three-node topology (from eq. 1 and 2).

Scenario	A	B	C	D	E
Link rate (Mbps)	33	33	33	33	33
Values for sources#0 to 3					
Σ target rates (Mbps)	10	6	10	6	14
BW_{excess} (Mbps)	23	27	23	27	19
$RTT_{average}$ (ms)	50	50	25	25	25
max router_1 (#OUT packets)	16	19	8	10	7
min router_1 (#OUT packets)	8	9	4	5	4
Values for sources#4 to 7					
Σ target rates (Mbps)	10	14	10	14	6
BW_{excess} (Mbps)	23	19	23	19	27
$RTT_{average}$ (ms)	50	50	65	65	65
max router_2 (#OUT packets)	16	13	21	17	24
min router_2 (#OUT packets)	8	7	12	9	12

Case 1. This first case is composed of three routers RIO and 8 TCP Reno sources with an Assured Service contracted as illustrated in Fig. 3. Sources generate traffic at link rate, set to 33 Mbps. The different characteristics of scenarios A, B, C, D and E were described in Table 1.

As observed in simulation results (see Table 4), CBM and RIO allow users to obtain their contracted rates despite of the miscellaneous situations. The fact of dropping out-of-profile packets before entering the aggregate makes TCP sources to adapt to the network characteristics regarding the available bandwidth and their respective contracted rates. Once the target rates are achieved, each connection gets a similar portion of excess bandwidth as indicated in Fig. 9, where the index fairness is over 0.8 for all scenarios except scenario D. The odd distribution of excess bandwidth in scenario D can be explained as follows. *Router_1* manages connections with small targets and low RTT, while *router_2* has to deal with large target rates and high RTT. In a single node topology, this does not represent a problem due to the good interaction between CBM and RIO. However, in this case, the task of distributing excess bandwidth is done mostly by *router_3*. This router uniquely makes use of RIO, and consequently hardly manages to provide a fair excess bandwidth sharing (*f=0.623*).

Table 4. Achieved rate in Mbps for IN packets in case 1

Source\Scenario	A	B	C	D	E
0	2.50	0.99	2.50	0.99	3.85
1	2.49	1.00	2.50	0.99	3.99
2	2.49	1.99	2.49	2.00	3.00
3	2.49	1.99	2.49	2.00	2.99
4	2.49	2.99	2.50	2.99	1.99
5	2.49	2.99	2.50	2.99	1.99
6	2.49	3.99	2.49	3.95	1.00
7	2.50	3.99	2.48	3.70	1.00

Case 2. In this situation, eight TCP Reno sources have an Assured Service as depicted in Fig. 3 but one of the nodes, *router_2*, does not implement RIO but RED. We have conducted simulations with the same different scenarios of case 1. This example is interesting for an ISP (Internet Service Provider) because it can guarantee an Assured Service with a simpler implementation (i.e. RED, that is basically a FIFO scheme avoiding global synchronization), or even it can reconfigure more easily its network resources.

Results exhibit that the inbound bandwidth is still guaranteed after a transient interval. As an example see Fig. 7, where we show the first 180 seconds of simulation time for the most complex scenario (D). The fairness index is over 0.8 except for scenario D (see Fig. 9). Again, this is the worst situation, since *router_2* endures more assured traffic (IN packets) with larger RTT in its connections, and it does not implement packet differentiation (RED mechanism). The dropping of OUT packets in the CBM traffic conditioners make sources to balance the distribution of the outbound bandwidth without interfering in the assured contracted rates that are guaranteed by *router_3*.

Case 3. Topology of case 2 allows for a flexible network reconfiguration. Now, we are interested in studying the effect of best-effort traffic in order to foresee the robustness of CBM-RIO in this topology. This type of traffic could represent malicious users that get into the assured network or transient ISP resource allocation. Usually, DiffServ implementations do not mix best-effort and assured traffics in the same queue, but locating their corresponding packets in different queues that belong to different Assured Forwarding classes. Accordingly, the coexistence of best-effort sources with Assured Service ones would represent transitory situations in which the ISP cannot reallocate those connections or configure a more specific architecture. Therefore, obtaining good results under these circumstances (i.e. ensuring contracted target rates and fair distribution of outbound bandwidth) would imply that the ISP is able to face up special short-live cases like these ones.

In simulations, traffic is generated at link rate by twelve TCP Reno sources. Sources 0 to 3 and 6 to 9 have an Assured Service, while sources 4 and 5 from *router_1* and sources 10 and 11 from *router_2* are best-effort (see Fig. 8). We also use miscellaneous attributes as depicted in Table 5. The best-effort connection packets are treated as out-of-profile and these sources do not have contracted target rates, thereby trying to get as much excess bandwidth as possible.

In spite of not having packet differentiation in *router_2* (RED), results illustrate that contracted rates for Assured Service sources are guaranteed (see Table 6). Again, we experience some problems in scenario D, since sources 8 and 9 remain below their target rates. These two connections in scenario D have two of the largest RTT (90 ms and 100 ms) and the greatest contracted rates (4 Mbps), and both enter the RED router. Therefore, *router_3* has to deal with most of the inbound bandwidth, which together with the presence of best-effort traffic makes sources 8 and 9 to not completely fulfill their targets.

Due to the substantial differences in delays and targets among connections, it is not possible to strictly guarantee contracted rates. Additional improvements should be necessary to achieve more robustness if the ISP would consider this topology as quite likely. Furthermore, the Internet Assured Service does not provide a hard guarantee. The fact of ensuring targets with variations that in the worst case achieve 70 % of target rates could be thought of as an advance in providing service differentiation with the assured service approach.

Table 6. Achieved rate in Mbps for IN packets in case 3.

Source\Scenario	A	B	C	D	E
0	2.50	1.00	2.49	1.00	3.30
1	2.49	0.99	2.50	1.00	3.99
2	2.49	1.99	2.49	1.99	2.99
3	2.49	1.99	2.50	2.00	2.99
6	2.49	2.99	2.49	2.90	1.99
7	2.49	2.99	2.49	2.70	1.99
8	2.50	3.97	2.50	2.70	1.00
9	2.49	3.99	2.49	2.60	1.00

Regarding the excess bandwidth, the fact of having mixed CBM AS traffic and best-effort traffic in the same router (*router_1* and *router_2*) favors the generation of less OUT packets from best-effort connections, therefore the fairness index is kept above 0.75 excluding scenario D (see Fig. 9).

Table 5. Target rates, RTT and *max-min* limits, if applicable, for TCP Reno sources in case 3 (from source#0 to source#11 respectively).

	Contracted target rates (Mbps)	RTT (ms)	*Sources# 0-3*		*Sources# 6-9*	
			Max	*Min*	*Max*	*Min*
Scenario A	2.5-2.5-2.5-2.5-0-0-2.5-2.5-2.5-2.5-0-0	50	16	8	16	8
Scenario B	1-1-2-2-0-0-3-3-4-4-0-0	50	19	9	13	7
Scenario C	2.5-2.5-2.5-2.5-0-0-2.5-2.5-2.5-2.5-0-0	10 to120 at intervals of 10	11	6	30	15
Scenario D	1-1-2-2-0-0-3-3-4-4-0-0	10 to120 at intervals of 10	13	7	25	13
Scenario E	4-4-3-3-0-0-2-2-1-1-0-0	10 to 120 at intervals of 10	10	5	35	18

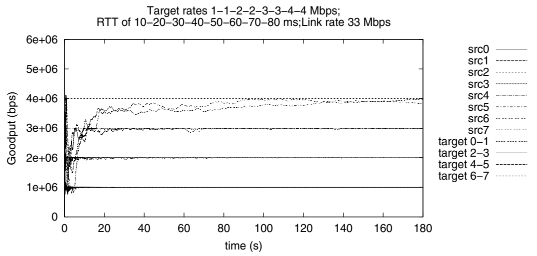

Figure 7. Guaranteed contracted target rates of all sources with CBM in case 2 and scenario D.

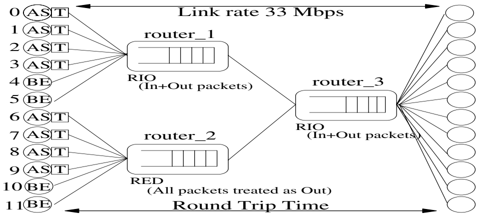

Figure 8. Three-node topology for case 3.

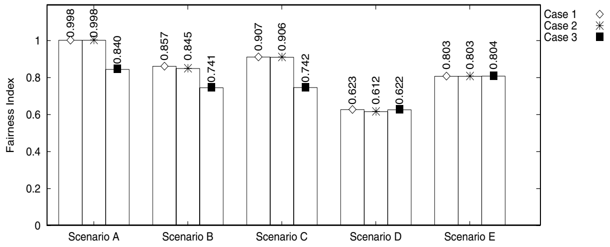

Figure 9. Fairness index with CBM in three-node topologies.

5. Conclusions and further work

In this paper, we have carried out a performance analysis of the Counters-Based Modified (CBM) traffic conditioner. The CBM, used together with the RIO buffer management scheme, guarantees target rates and distributes excess bandwidth equitable among competing sources. The CBM traffic conditioner reaches this objective discarding out-of-profile packets before joining the aggregate, with a probability that depends on the target rate, the excess bandwidth, and an estimation of the average RTT of all connections.

We present simulation results in miscellaneous TCP environments (different target rates, different round trip times, and share of resources with best-effort connections) with a single node and a three node topology closer to a real framework. Simulation results reveal that in most cases contracted target rates are guaranteed regardless the particularities of the simulation scenario. Excess bandwidth in these environments is distributed among sources in a way that the fairness index remains above 0.8 in all toplogies and scenarios, which can be considered as reasonable for an Assured Service.

Acknowledgements

This work was supported by the Spanish Research Council under project FAR-IP TIC2000-1734-C03-03.

References

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

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

[3] D. Clark and W. Fang, “Explicit Allocation of Best-Effort Packet Delivery Service”, IEEE/ACM Transactions on Networking, Vol. 6 No. 4, pp. 362-373, August 1998.

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

[5] W. Lin, R. Zheng, J. Hou, “How to make assured service more assured”, Proceedings of the 7th International Conference on Network Protocols (ICNP'99), pp. 182-191, Toronto, Canada, October 1999.

[6] B. Nandy, N. Seddigh, P. Pieda, J. Ethridge, “Intelligent Traffic Conditioners for Assured Forwarding Based Differentiated Services Networks”, Proceedings of Networking 2000, LNCS 1815, Paris, France, pp.540-554, May 2000.

[7] M. Goyal, A. Durresi, P. Misra, C. Liu, R. Jain, “Effect of number of drop precedences in assured forwarding”, Proceedings of Globecom 1999, Rio de Janeiro, Brazil, Vol. 1(A), pp. 188-193, December 1999.

[8] I. Alves, J. De Rezende, L. De Moraes, “Evaluating Fairness in Aggregated Traffic Marking”, Proceedings of IEEE Globecom’2000, San Francisco, USA, pp. 445-449, November 2000.

[9] I. Andrikopoulos, L. Wood, G. Pavlou, “A fair traffic conditioner for the assured service in a differentiated services internet”, Proceedings of IEEE International Conference on Communications ICC2000, New Orleans, LA, Vol. 2, pp. 806-810, June 2000.

[10] Maria-Dolores Cano, Fernando Cerdan, Joan Garcia-Haro, Josemaria Malgosa-Sanahuja, “Counters-Based Modified Traffic Conditioner”, Lecture Notes in Computer Science (QofIS 2002), Vol. 2511, pp. 57-67, Springer-Verlag, 2002.

[11] Maria-Dolores Cano, Fernando Cerdan, Joan Garcia-Haro, Josemaria Malgosa-Sanahuja, “Performance Evaluation of Traffic conditioner Mechanisms for the Internet Assured Service”, in Quality of Service over Next-Generation Data Networks, Proceedings of SPIE Vol. 4524, pp. 182-193, 2001.

[12] F. Cerdan, O. Casals, “Performance of Different TCP Implementations over the GFR Service Category”, ICON Journal, Special Issue on QoS Management in Wired &Wireless Multimedia Communications Network, Vol.2, pp.273-286, Baltzer Science, January 2000.

[13] V. Bonin, F. Cerdan, O. Casals, “A simulation study of Differential Buffer Allocation”, Proceedings of 3rd International Conference on ATM, ICATM’2000, pp. 365-372, Germany, June 2000.

[14] V. Bonin, O. Casals, B. Van Houdt, C. Blondia, “Performance Modeling of Differentiated Fair Buffer Allocation”, Proceedings of Conference on Telecommunications Systems, Dallas, USA, 2001.

[15] R. Jain, “The Art of Computer Systems Performance Analysis”, John Wiley and Sons Inc., 1991.