A New Proposal for Assuring Services in Internet

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 32 5953
Fax +34 +968 32 5338
{mdolores.cano, fernando.cerdan, joang.haro, josem.malgosa}@upct.es

Abstract. *In this paper we present a new mechanism to provide an assured service in terms of target rate and fair excess bandwidth, like the Internet Assured Service. Research in Internet Assured Service faced up both questions in separate ways proposing different traffic conditioners to work with the RIO buffer management, and proposing different modifications to this buffer management, among others. In this work, we suggest using a buffer management scheme different from RIO that also treats in-of-profile and out-of-profile packets differently but avoiding interference between them. This scheme is used together with the Counters Based traffic conditioner because of its high accuracy in guaranteeing target rates. We evaluate and compare by simulation the performance of our proposal using TCP RENO sources. One important issue to be considered is that the proposal is a feasible alternative to the standard architecture for Differentiated Services in Internet.*

Keywords: *traffic conditioner, buffer management, RIO, DiffServ, QoS.*

1 Introduction

Differentiated Services (DiffServ) approach [1] appeared as a need to provide QoS solutions 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. Thus, 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, whereas interior nodes perform a set of forwarding PHBs to aggregates of traffic that have been appropriately marked.

Currently, the IETF has two PHBs with the status of proposed standards, the Expedited Forwarding per-hop behavior (EF PHB) and the Assured Forwarding per-hop behavior (AF PHB). The idea behind AF PHB [2] is to ensure a minimum throughput (target rate or contracted rate) to a connection, while enabling consuming more bandwidth if the network load is low. To achieve this goal, packets of individual flows are marked belonging to one of the four independently forwarded AF classes. 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 congested DiffServ node tries to protect packets with a lower drop precedence value from being lost by preferably discarding packets with a higher drop precedence value. Note that minimum throughput is also called in-profile bandwidth or inbound bandwidth, and excess bandwidth can be also referred as outbound bandwidth along this study.

Despite of the abundant literature written about the AF-PHB (e.g. [3][4][5][6]), any solution has been found to face up its two goals, assuring the inbound bandwidth and offering a fair distribution of the excess bandwidth if available. The Counters-Based (CB) traffic conditioner presented in [7] has been demonstrated to perform comparatively better than other traffic conditioners. This mechanism based on counters guarantees the in-profile bandwidth allocation in topologies with variable round trip times and different target rates. Its easy configuration and high accuracy make it suitable for general use. Only two counters are needed to implement this algorithm, and any configuration parameter is required. It also includes a simple mechanism to avoid accumulation of “credits” when a source stops transmitting data. Nevertheless, it might be not necessary depending on the type of TCP traffic injected to the network, the TCP implementation (RENO, SACK, etc.) and the use of a suitable buffer management scheme able to avoid any type of connection misbehaviors.

From the comparative simulation study carried out in [7], this traffic conditioner together with RIO (RED [8] (Random Early Detection) In and Out) [9] over performs the TSW (Time Sliding Window)-RIO and Leaky Bucket-RIO mechanisms in terms of guaranteeing inbound bandwidth, with fluctuations in the achieved rate that do not exceed 1% of the connection target rate. Both TSW [9] and Leaky bucket have such amount of parameters to tune, additionally to the TCP source and RIO parameters that makes very difficult its configuration unlike the couple CB-RIO. Nevertheless, CB-RIO fails in providing a fair excess bandwidth sharing among sources. The simplicity of the CB algorithm motivated us to propose a different buffer scheme (Dual Queuing) that together with the traffic conditioner allows an easier parameter configuration in order to provide both, a good performance in most of the cases and the possibility of trying an analytical study.

In this paper, we introduce a new approach for achieving fairness in the excess bandwidth distribution and assuring inbound bandwidth for competing TCP connections. The idea is based on using a buffer management in the router different from the traditional RIO, and applied it with the Counters-Based algorithm [7] configuring the Counters-Based Dual Queuing (CBDQ) mechanism. The basic insight behind this scheme is to eliminate interference between IN and OUT packets by placing them in different queues in the router and using a suitable scheduling algorithm. As we show in simulation results, it is possible to provide fairness in the excess bandwidth sharing among the flows of the aggregate, and to strictly guarantee target rates to individual TCP sources in terms of *goodput* by using this approach. It is interesting to remark that despite of DiffServ mechanisms are not built to provide and end to end service, the use of TCP sources makes sense to study the performance of these TCP connections in terms of throughput without taking into account retransmitted packets, which is usually called *goodput*.

The rest of this paper is organized as follows. Section 2 details the characteristics of our proposal, the Counters-Based Dual Queuing implementation. In Section 3, we present the scenario and assumptions to carry out simulations. In Section 4, simulation results are shown. The paper concludes in Section 5 summarizing the most relevant facts.

2 Out Profile Bandwidth Fairness

As explained earlier, the Counters-Based traffic conditioner worked out the inbound bandwidth assurance problem [7]. Consequently, the remaining challenge to be figure out is the fair share of excess bandwidth among connections.

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 evaluate fairness we use equation (1), where x_i is the excess throughput of source *i*, and *n* is the number of sources that compose the aggregate [10]. The closer to 1 in the *f* value, the more the fairness obtained. As indicated in the introduction section, we employ the term throughput meaning *goodput* for the index fairness calculations. With this goal, next section discusses our proposal, the Counters-Based Dual Queuing scheme.

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

2.1 Dual Queuing Buffer Management

Dual Queuing Buffer Management comes out with the conception of applying a different buffer management scheme in the interior node rather than the traditional RIO implementation. Like in RIO, the underlying notion is to give a different treatment to IN and OUT packets in the router. Nevertheless, when IN and OUT packets share the same buffer there is some kind of interference between both types of traffic, which makes difficult to handle them in order to provide a fair excess bandwidth sharing. This effect has been widely shown in ATM networks with the GFR service [11]. GFR is an ATM service that provides similar goals than the Internet Assured Service usually tested with TCP sources, since most applications rely on TCP/IP protocols.

Dual Queuing Buffer Management consists of buffering and forwarding IN and OUT packets in the router from two separate FIFO queues. A suitable scheduling algorithm is also used to decide which queue is served first, the queue buffering the IN packets or the queue buffering the OUT packets. See Figure 1.

To reduce complexity we employ a scheduling algorithm that serves both queues in a weighted-round robin fashion. The scheduler serves the queue buffering IN packets with a probability that matches the total contracted traffic load ρ, while the queue buffering OUT packets is served with probability 1-ρ. If the scheduler visits an empty queue it switches to the other queue in order to have a work conserving system.

$$\rho = \frac{\sum_i target\_rate(i)}{link\_Rate} \quad (2)$$

In addition, both queues use a threshold to limit the maximum number of packets that can be stored. These thresholds are HBO_{in} and HBO_{out} for the IN and OUT queue respectively, where HBO stands for High Buffer Occupancy.

This buffer management scheme can be used with any of the traffic conditioner mechanisms mentioned in the introduction. Nevertheless, we have studied the performance of this scheme when we use a Counter Based traffic conditioner given its good behavior providing the inbound profile. For this reason we have called this scheme CBDQ (Counters-Based Dual Queuing). It should be noted that this is an alternative to the AF PHB proposed standard to provide an assured service. The key difference lies on using a second queue to OUT packets that may result in a non-in sequence packet service. In our opinion, the standard tries to force a minimal difficult implementation to interior nodes using a FIFO queue. Nevertheless this FIFO queue needs a rather complex buffer management like RIO with a bad traffic conditioner interaction, failing in providing a fair excess bandwidth sharing among TCP sources. It is clear that TCP overcomes the possible lost of sequence of out-of-profile packets, and today the use of two different queues along with a scheduling mechanism is not an important implementation drawback with current technology, even in extremely high speed networks.

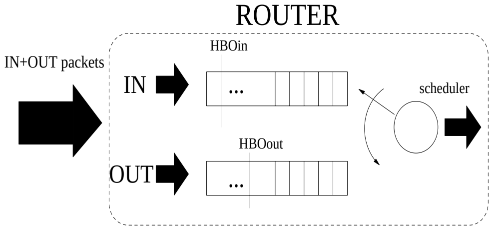

Fig. 1. Counters-Based Dual Queuing scheme.

In this paper we perform a simulation study in order to foresee a general good behavior of the couple Counters-Based Dual Queuing as the basis to try an analytical study that helps to tune the buffer thresholds that let us a suitable network dimensioning. This analytical study is out of the scope of this paper.

3 Scenarios for simulations

Topology selected for our simulations is illustrated in Figure 2. TCP traffic is generated by 8 TCP Reno sources that transmit at the link rate, which has been set to 33 Mbps. To verify the impact of using different target rates for each source, different values are used along the simulations. We also measure the influence of different RTTs. In the TCP homogeneous scenario (same RTT for all connections), round trip delay between sources and destinations is set to 50 ms. In the TCP heterogeneous scenario, this value varies from 10 ms to 80 ms in increments of 10 ms. In our simulations, HBO_{in} and HBO_{out} have been set to 20 packets and 10 packets respectively.

As a first insight, we have used a large packet size of 9,188 bytes that corresponds to the classical IP over ATM, there is an increasing interest in supporting differentiated services in technologies like MPLS, where the use of ATM technology seems inherent. We investigate if the different systems involved (TCP sources, traffic conditioner and the buffer management scheme) are able to meet the expected service guarantees. For this reason, we have conducted some of our simulations (see next section) with other packet sizes along with TCP RENO greedy sources for a relatively high network load. The simulation tool used in this work for the sliding window protocol of TCP Reno sources was developed in [11] and has been also extensively used in [12] and [13]. In addition, it was applied to validate the analytical study carried out in [14].

For the set of simulations, we consider five different scenarios with a contracted traffic load around 60% of the link rate. In these situations, we have a 40% of the link rate as excess bandwidth, so we can measure how fair the share is among the different sources.

Scenario A. All connections have same RTT and same contracted rates. Simulations have been run with target rates of 2.4 Mbps and RTT of 50 ms for all connections.

Scenario B. All connections have same RTT and different contracted rates. Simulations have been conducted with target rates of 1, 1, 2, 2, 3, 3, 4 and 4 Mbps respectively and a RTT of 50 ms for all connections.

Scenario C. All connections have different RTT and same contracted rates. Simulations have been performed with target rates of 2.4 Mbps for all connections and RTT from 10 to 80 ms at increments of 10 ms.

Scenario D. All connections have different RTT and different contracted rates (sources with small targets have small RTT). Simulations have been carried out with target rates of 1, 1, 2, 2, 3, 3, 4 and 4 Mbps and RTT from 10 to 80 ms at increments of 10 ms.

Scenario E. All connections have different RTT and different contracted rates (sources with small targets have large RTT). Simulations have been run with target rates of 4, 4, 3, 3, 2, 2, 1 and 1 Mbps and RTT of 10 to 80 ms at increments of 10 ms.

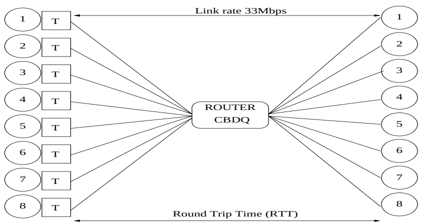

Fig. 2. Topology for simulations (T≡Traffic Conditioner).

4 Simulation results

In next sections, we present simulation results carried out in the scenarios described above. The fairness index offers an idea of the bias in the excess bandwidth distribution. Moreover, we demonstrate that using the Dual Queuing buffer management scheme does not affect the Counters-Based performance regarding the in-profile bandwidth assurance. Finally, some simulations have been done to observe the effect of having different packet sizes.

4.1 Fairness index

Simulations have been run for the five cases described in Section 3. By using this implementation in the router, results in terms of excess bandwidth sharing clearly improve respect to the traffic conditioner (TSW, LB, Counters-Based)-RIO implementation as shown in Table 1.

From results presented in this table, we can see that the improvement in the fairness index is small (around 0.1) in cases A and B. This is because of the unusual characteristic of these two situations, i.e., same RTT for all connections. Under this circumstance, many pairs traffic conditioner-RIO perform well in terms of achieving a fair excess bandwidth share, included the Time Sliding Window, the Leaky Bucket and the Counters-Based.

In cases C, D an E, the effect of having a different RTT for each source determines the improvement in the fairness index when using the CBDQ approach. Clearly, the interaction traffic conditioner (TSW, LB or CB)-RIO in a heterogeneous scenario, different RTT and different target rates, originates an unfair distribution of the outbound bandwidth. The use of a buffer management scheme that treats IN and OUT packets in a different way, but without interfering each other confers much better results.

Table 1. Fairness index (TC≡Traffic Conditioner; BMS≡Buffer Management Scheme; S≡Scenario).

TC-BMS	**S. A**	**S. B**	**S. C**	**S. D**	**S. E**
CBDQ	0.99	0.97	0.97	0.88	0.95
CB–RIO	0.85	0.85	0.78	0.71	0.84
TSW–RIO	0.58	0.81	0.63	0.49	0.56
LB–RIO	0.85	0.68	0.74	0.81	0.83

We have also performed some simulations using a shared memory for the OUT queue, i.e. the OUT queue is separated in several virtual queues in order to see the effect in the fairness index. This fact could benefit the fairness index taking into account that OUT packets are buffered randomly among the virtual queues, as well as they are served randomly too when the scheduler visits the shared OUT queue (with visit probability 1-ρ). Nevertheless, obtained simulation results (not shown) reveal that the fairness index gain is negligible. This is due to the interaction between the Counters-Based traffic conditioner and the dual queuing scheme that makes all TCP connections to send approximately the same amount of OUT packets, the key issue for a fair bandwidth sharing, independently of its contracted target rate and round trip time. In addition, it is important to remark that the amount of parameters used in this scheme are lesser and simpler to configure than parameters involved in RIO and the TSW or Leaky Bucket components. Furthermore, as we have commented earlier, this simplicity makes us to start a possible analytical approach in order to provide a proper network dimensioning.

Another important concern is that the interaction between the Counters-Based traffic conditioner and the RIO buffer management scheme was studied in [7]. In this work, it was shown that despite of the effect that parameters such as RTT or target rates produce in the overall performance, the inbound bandwidth was guaranteed. We present simulation results to confirm that the interaction between the Counters-Based traffic conditioner and the CBDQ buffer management scheme keeps on assuring contracted rates.

Figures 3 to 5 depict the achieved rate for IN packets in scenarios B through D. It is evident from the figures that after a short transient, the in-profile bandwidth is hard guaranteed. Note that, this IN throughput (*goodput*) is part of the global *goodput* performance of the TCP source, where each source gets a fair share of the bandwidth according to the definition given in Section 2. The use of a different buffer scheme, rather than the RIO, does not influence the completion of the traffic conditioner.

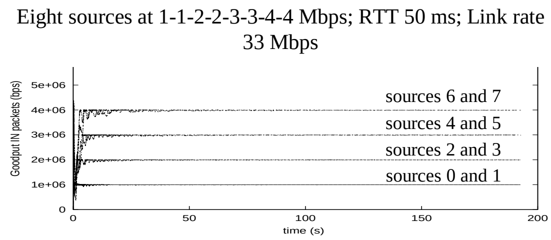

Fig. 3. Achieved rate for IN packets in scenario B.

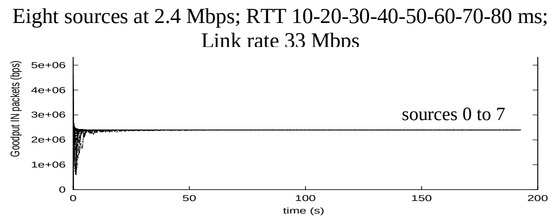

Fig. 4. Achieved rate for IN packets in scenario C.

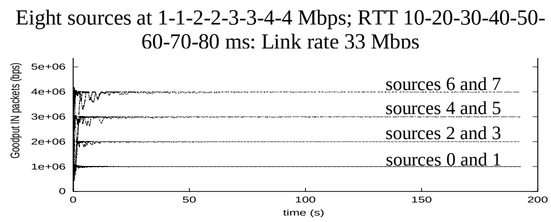

Fig. 5. Achieved rate for IN packets in scenario D.

4.2 Packet size dependency

Given that nodes may receive packets of any length, we have considered of interest to observe how variations in the packet size can affect the performance in terms of biasing in the excess bandwidth distribution. In this case we have carried out a set of simulations for scenarios A through E, and packet sizes of 1,500, 5,300 and 9,188 bytes. Although 5,300 bytes packet size does not match any known service, it has been included to study the trend of the fairness index as the packet size decreases. Therefore, in this section we perform new simulations to calculate the fairness index with small and medium packet sizes. As shown in Table 2, the index fairness value remains over 0.8 for nearly all cases with the Counters-Based Dual Queuing mechanism.

Table 2. Fairness index in CBDQ with packet size variations.

Packet size	**S. A**	**S. B**	**S. C**	**S. D**	**S. E**
1,500 bytes	0.99	0.90	0.64	0.40	0.86
5,300 bytes	0.99	0.94	0.90	0.76	0.98
9,188 bytes	0.99	0.97	0.97	0.88	0.95

Table 3 includes the fairness index obtained in simulations to study the interaction of different packet sizes travelling simultaneously in the network. We have employed three different packet sizes, 9,188 bytes (sources 0, 3 and 6), 5,300 bytes (sources 1, 4 and 7) and 1,500 bytes (sources 2 and 5). As it is observed, we get worst results for the index fairness when there are packets of distinct sizes going through the network. Despite of some results may favor TSW or Leaky Bucket, we should take into account that both mechanisms lack of enough accuracy providing the inbound rates. In addition, given their difficult parameter configurations, simulation results in Table 3 cannot be considered as a general conclusion. It should be clear that the simulation results presented in Table 3 are taken in the worst possible conditions, so we must keep in mind that CBDQ has the lesser and easiest parameter configuration that may help to obtain in general better results. We have left this issue for further study.

Table 3. Fairness index in CBDQ with simultaneous different packet sizes.

TC-BMS	**S. A**	**S. B**	**S. C**	**S. D**	**S E**
CBDQ	0.70	0.56	0.63	0.45	0.61
CB-RIO	0.64	0.38	0.68	0.76	0.66
TSW-RIO	0.56	0.51	0.57	0.49	0.63
LB-RIO	0.75	0.63	0.73	0.62	0.70

5 Conclusions and further work

Despite of extensive research in providing both a target rate and a fair excess bandwidth sharing in a DiffServ domain, none of the solution presented faces up the problem in a suitable way in terms of feasibility and scalability as we do in this work. We have proposed the use of a different buffer management scheme in the router rather than the traditional RIO algorithm. This buffer management breaks a basic rule specified in the DiffServ architecture defined in

RFC 2474, since OUT packets may be not served in the arrival sequence.

The scheme gives a different treatment to IN and OUT packets in the router using two separate FIFO queues to avoid interaction between both types of traffic (Dual Queuing buffer management). A scheduling algorithm is employed to decide which queue is first served. The scheduling serves the queue with IN packets with a probability equal to the contracted load i.e., the sum of the target rates divided by the link rate, while the queue with OUT packets is visited according to the complementary probability. The combination of the Counters-Based traffic conditioner and this dual queuing buffer management gives rise to the Counters-Based Dual Queuing (CBDQ) implementation.

Simulation results show that, with this implementation at router nodes is possible to provide a target rate and a fair share of the excess bandwidth in terms of *goodput* to each TCP RENO source that composes the aggregated that arrives to the router. Therefore, we can provide an assured service like the Internet Assured Service.

We have compared CBDQ results with the two most important Assured Service schemes (TSW-RIO and Leaky Bucket-RIO). CBDQ ensures the inbound rate of each individual TCP source, while providing and end to end *goodput* performance. Moreover, CBDQ allows each source to benefit from excess bandwidth with a fairness index close to 0.8 in all cases, unlike the other two schemes.

Summarizing, the combination Counters-Based Dual Queuing lacks of complexity parameter configuration unlike the TSW or Leaky Bucket – RIO combinations that together with the TCP characteristics make these Assured Service Implementations almost impossible to control, to provide the expected target rate and a fair excess bandwidth share depending on network resources. The simulation study presented in this work should be considered as a first step in providing an analytical approach to the behavior of the CBDQ with TCP sources that gives rise to a suitable network dimensioning.

Acknowledgements

This work was supported by the Spanish Research Council under grant 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] B. Nandy, N. Seddigh, P. Pieda, J. Ethridge, "Intelligent Traffic Conditioners for Assured Forwarding Based Differentiated Services Networks", Proceedings of Networking 2000, Paris, France, pp. 540-554, May 2000.

[4] 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.

[5] 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.

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

[7] Maria-Dolores Cano, Fernando Cerdan, Joan Garcia-Haro, Josemaria Malgosa-Sanahuja, "Performance Evaluation of Traffic conditioner Mechanisms for the Internet Assured Service", Proceedings of ITCOM'01, Denver, USA, Vol. 4524, pp. 182-193, August 2001.

[8] 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.

[9] 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.

[10] R. Jain, "The Art of Computer Systems Performance Analysis", John Wiley and Sons Inc., 1991.

[11] 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.

[12] F. Cerdan, O.Casals, "Mapping an Internet Assured Service on the GFR ATM Service", Lecture Notes in Computer Science 1815 (Networking 2000), pp. 398-409, Springer-Verlag.

[13] V. Bonin, F. Cerdan, O.Casals, "A simulation study of Differential Fair Buffer Allocation", Proceedings of 3rd International Conference on ATM, ICATM'2000, Germany, June 2000.

[14] V. Bonin, O.Casals, B. Van Houdt, C. Blondi, "Performance Modeling of Differentiated Fair Buffer Allocation", Proceedings of the 9th International Conference on Telecommunications Systems, Dallas, USA.