ASSURED SERVICE EVALUATION

José F. de Rezende
rezende@gta.ufrj.br
Grupo de Teleinformática e Automação (GTA)
Universidade Federal do Rio de Janeiro
COPPE/EE - Programa de Engenharia Elétrica
Caixa Postal 68504 - 21945-970 - Rio de Janeiro - RJ - Brasil

Abstract

This article focuses on the differentiated services model as a basis for the provision of QoS in the Internet. The objective is to evaluate the performance of TCP traffic flows crossing domains where nodes implement the Assured Forwarding PHB. The results obtained show the adequacy of this service in the supply of QoS guarantees and disclose to the problems of fairness in the bandwidth sharing.

1 Introduction

With the increased processing capability of personal computers and the appearance of new distributed multimedia applications, the Internet tends to receive a larger amount of multimedia traffic. The best-effort model adopted by the Internet for the transport of information doesn't support the requirements of these new applications. Proposals are being studied to find a way to define different service levels, i.e. services on the Internet with a specific QoS. The proposals typically use two approaches: resource reservation and exploitation of the ToS field (Type of Service) from the IP header. These approaches are known as Integrated Services [1] and Differentiated Services [2], respectively. The Integrated Services approach provides service discrimination through the explicit allocation and reservation of resources. This model demands a large amount of signaling between nodes and maintaining state information at each node for each flow. The Differentiated Services approach intends to define a simple group of mechanisms to treat packets with different priorities according to the information carried in the ToS field of the IP packet header. This approach requires neither state information per flow nor signaling to each node, which increases its potential for scalability.

This article focuses on the Differentiated Services (DS) model as a basis for providing QoS in the Internet. In the proposal developed by the Differentiated Services (*DiffServ*) working group, the idea is to rename the ToS field of the IP header, and call it the DS field. When entering a DS domain, packets from a traffic flow have their DS field marked in agreement with the contract established between the user that generated the flow and those responsible for the domain. From the analysis of this field, routers inside the DS domain can provide different treatment for the packets of aggregated flows [3]. This treatment, known as Per Hop Behavior (PHB), is defined in each node of the DS domain. Two PHBs are currently in wide discussion; the Expedited Forwarding PHB (EF-PHB) [4] and the Assured Forwarding PHB (AF-PHB) [5]. These treatments per node when coupled with traffic conditioning mechanisms at the borders provide a service level within the domain. A detailed and exhaustive evaluation of the service level offered by these mechanisms needs to be accomplished prior to its deployment. The objective of this article is to evaluate the performance of TCP flows crossing DS domains that implement the Assured Forwarding PHB.

This article is organized into several sections. Section 2 describes the Assured Service, detailing the characteristics of the AF-PHB. In section 3, the mechanism of the active queue management is described. In section 4, the simulation results obtained, and the analysis of the results is presented. Finally, section 5 draws conclusions and discusses the future of this work.

2 Assured Service

The Assured Service is a service supplied by a DS domain that implements the AF-PHB at each node. The Assured Service model is not based on a strict service guarantee, but on a service expectation to be obtained, even when there are periods of congestion. The admission control plays an important role in this model, because it is entrusted with guaranteeing the existence of the necessary resources along the entire path between the origin and the destination. An Assured Service user establishes a contract with the provider that defines a profile for the service expected to be obtained.

Traffic conditioning mechanisms at the borders control the Assured traffic that enters or leaves the DS domain. The meter is used to compare the traffic with the contracted service profile and its state relative to a packet can be used to affect the actions of the marker, policer and shaper. In the Assured Service model, a warranty exists, such that marked packets are delivered with a high probability, as long as the aggregated traffic doesn't exceed the contracted rate defined by the service profile. However, a user is allowed to exceed the contracted profile rates, with the understanding that the excess traffic won't be delivered with the same high probability. As packets enter a DS domain, they are treated at each node according to the PHB identified by the DS field present in each packet. In the Assured Service, each packet has a *codepoint* encoded in the DS field which identifies the AF-PHB.

The AF-PHB [5] allows a DS domain provider to offer different levels of assurances for forwarding IP packets for a user of the domain. AF-PHB provides the delivery of IP packets in four independent delivery classes (AF classes), where each class is allocated a certain amount of resources, such as buffers and bandwidth, in each DS node. Within each AF class, an IP packet can be marked, either by the user or by the DS domain, with three levels of loss precedence. In case of congestion within the AF class, the loss precedence of each packet determines the relative importance of that packet. A congested DS node preferentially discards packets with a higher loss precedence value, while avoiding to discard packets with a smaller precedence value. The implementation of an AF-PHB tries to minimize long-term congestion, but allowing short-term congestion, such as the result of traffic bursts. This requires an algorithm for active queue management, which is discussed in the next section.

3 Active Queue Management

In order to support the requirements of the Assured Service, an active queue management mechanism should contain a function that monitors the instantaneous congestion level and computes a smoothed congestion level. The algorithm for packet discarding uses this smoothed congestion level to determine when packets should be discarded. This can be obtained by using a random function to discard packets. An example of such algorithm is RED (Random Early Detection) [6].

RED is a mechanism implemented in routers designed to be used with the TCP protocol. In this mechanism, the router monitors the size of its queue, and when it detects an imminent period of congestion, the router notifies the source implicitly by discarding one of its packets. This way, the source is notified based on the timeout corresponding to the discarded packet. This determines that the source should adjust its congestion window. The RED routers discard packets earlier than typical routers so that the source decreases its congestion window earlier.

As previously stated, for each AF class, routers at the border of the DS domain monitor and mark packets from individual (or aggregated) flows in agreement with the different loss precedence levels supported by the domain. If only two loss precedence levels are used, packets from a flow that obey the contracted service profile are marked as *IN* (in-profile) and packets that are outside the service profile are marked as OUT (out-of-profile). In this case, the queue management by the internal routers of the DS domain can be accomplished by the adoption of two RED algorithms. One for packets that are in conformity with the profile, i.e. packets marked with *IN*, and another for the OUT packets. This mechanism is known by RIO (RED with *IN* and OUT) [7]. The second RED algorithm is configured more aggressively than the first algorithm, since there is a priority on discarding OUT packets.

The RIO calculates an average queue size for *IN* packets, and another for all packets in the queue (*IN*+OUT). When a packet arrives, RIO determines if the packet is *IN* or OUT. If it is an *IN* packet, the router only calculates the average size of the queue for the *IN* packets (avg_{IN}). Otherwise, the algorithm calculates the average size of the queue taking into account all packets (avg_{TOTAL}). The probability of discarding an *IN* packet depends on avg_{IN}, and the probability of discarding OUT packets depends on avg_{TOTAL}. Three parameters exist for each RED algorithm. The three parameters, min_{IN}, max_{IN} and P_{maxIN}, define the normal operation phase $[0, min_{IN})$, the congestion avoidance phase $[min_{IN}, max_{IN})$, and the congestion control phase $[max_{IN}, \infty)$ for *IN* packets. In the same form, min_{OUT}, max_{OUT} and P_{maxOUT} defines the corresponding phases for the OUT packets.

4 Simulation results

The simulation of the Assured Service was done using the ns-2 simulator [8]. For this, it was necessary to implement the mechanisms that compose this service, i.e. traffic conditioning mechanisms, such as policers and packet markers, and the mechanism of active queue management that implements the AF-PHB itself. The mechanism used to check the traffic conformity at the border routers, and so to mark packets in agreement with service profiles, is the leaky-bucket mechanism. Specifically, packets that arrive in the leaky-bucket and find a credit are marked as *IN* and immediately sent. Those packets that don't find a credit are marked as *OUT*. The fill rate of the bucket is the expected rate, i.e. the service profile for which the Assured Service should supply warranties. The bucket size, i.e. the maximum allowed burst size, is an evaluated parameter to determine the influence on the performance.

The topology used in the simulations is shown in figure 1. The choice of the topology was based on the following criterion in an attempt in obtaining results that were as realistic as possible. RTT values were chosen so that TCP flows reach a considerable throughput (delay-bandwidth product). Increasing RTT values for evaluated TCP flows. Creation of the aggregation effect, i.e. the access links have a larger capacity than that of the bottleneck link, so that bursts are formed in this last link. In the topology used in the simulation, the bottleneck link is between nodes $r1$ and $r2$. It exists flow merging and splitting points so that bursts of packets arrive at the same time in the queues.

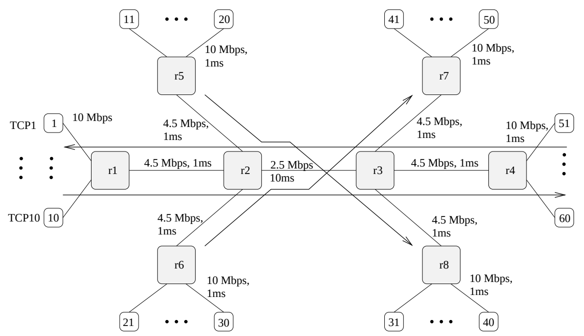

Figure 1: network topology.

4.1 Adaptive traffic

In a first simulation the fairness of the bandwidth sharing for the Assured Service was evaluated. For this, 10 TCP connections (N_{AF}), TCP1 to TCP10, were connected from node 1 to node 10. Each connection i has a service profile equal to p_i, where $p_i = i * p_1$ with $i = 1..10$. In these simulations, the best-effort traffic that competes with the Assured Service traffic are long-lived adaptive flows, i.e. flows that use end-to-end congestion control. This traffic is generated by 20 FTP flows (N_{BE}) crossing the DS domain originating from node 11 to 30, as indicated by the arrows from figure 1. To check the conformity of the traffic generated by each Assured flow, a leaky-bucket was used with a bucket size of 20 packets (or TCP segments), where each segment has a size of 1000 bytes. The following parameters for the queues *IN* and *IN+OUT* were used, respectively: [15,30,0.002,0.02] and [7,15,0.002,0.1][1].

The curve of figure 2 shows the goodput achieved by the connections with the smallest and the largest service profile, TCP1 and TCP10 respectively, as a function of the reserved bandwidth. The total reserved bandwidth is shown as a percentage of the total capacity of the bottleneck link. These curves also show the fair bandwidth (f_i) that should be allocated for each of these connections. The calculation of these values is made through the equation 1. It is considered fair bandwidth for each Assured flow, the reserved rate plus a weighted share of the excess bandwidth.

$$f_i = p_i + (C_b - \sum_{i=1}^{N_{AF}} p_i) * \frac{N_{AF}}{N_{AF} + N_{BE}} * \frac{p_i}{\sum_{i=1}^{N_{AF}} p_i} \quad (1)$$

The results show that connections with lower profile rates always obtain a larger bandwidth than considered fair. As the profile rate increases it becomes harder to these connections to obtain this fair bandwidth. In situations of an over-provisioned network (less than 20%) and under-provisioned network (more than 60%), connections with larger profile rates do not achieve their fair bandwidth. The same behavior was observed for all other connections of this simulation. This was already remarked in a previous work [9].

In a second simulation, the monitored connections (TCP1 to TCP10) have different RTTs, with the purpose of considering the effect of this parameter in the bandwidth obtained by a TCP flow. The RTT of each connection is varied from 36ms (TCP1) to 126ms (TCP10) through the modification of the propagation delay from the sources to node $r1$ (between 5ms and 50ms). The curve of figure 3 shows the goodput achieved by the connections TCP1 and TCP10 as a function of the reserved bandwidth. These curves also

[1] $[min\_thresh, max\_thresh, q_{weight}, P_{max}]$

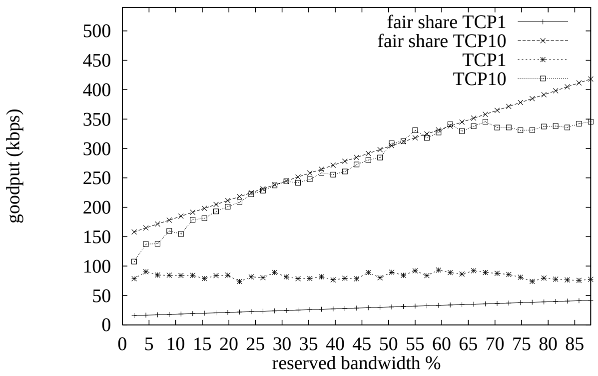

Figure 2: TCP goodput and fairness.

show the fair share for each of these connections. The curves show that when the RTT increases it becomes more difficult to the connections with larger profiles to achieve their fair rates. For instance, the connections TCP7, TCP8, TCP9 and TCP10 do not realize their fair rate. The cause is that connections with a small RTT send more oportunistic traffic (OUT packets) than the connections with a large RTT. Despite that to lead to a higher probability of losing packets, these connections grab more bandwidth than connections with large RTTs.

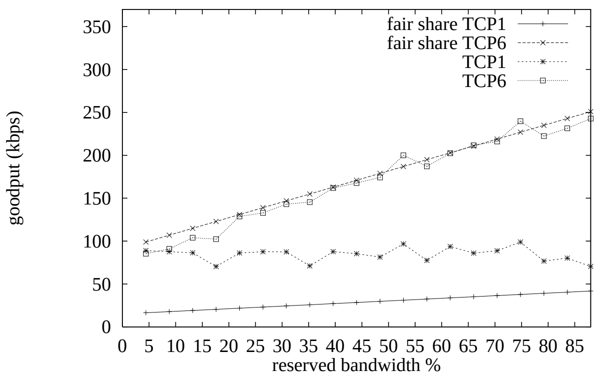

Figure 3: TCP goodput and RTT.

In a third simulation, the bucket size of the leaky-bucket that checks the conformity of the traffic was varied. Figure 4 shows the goodput achieved by the connections from TCP6 through TCP10 as a function of the bucket size in number of packets. The service profile of each connection is given by $p_1 = 18kbps$, which gives approximately 40% of the total capacity of the bottleneck link.

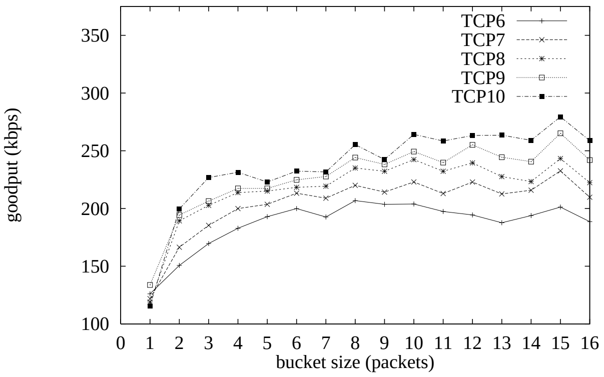

Figure 4: TCP goodput as function of the bucket size.

The results show that the throughput of the connections with a larger service profile rate is more sensitive to the bucket size than those with a smaller service profile rate. The connections reached a stable value only for bucket sizes of tewlve packets or higher. The smaller the bucket size, the larger the probability of losing credits. These losses reduced the effective marking rate of a traffic flow to a value below the contracted profile. In the connections with a larger service profile, the bucket was filled at a higher rate. This increased the rate of credit losses even more. When the bucket size was increased, those connections regained a marking rate corresponding to the contracted profile.

4.2 Non-Adaptive traffic

In addition to the adaptive traffic, the simulations included best-effort traffic that did not use congestion control. This was represented by UDP sources that generated either constant rate or bursty traffic. These sources, known as non-adaptives, do not decrease their sending rate during periods of congestion, which can worsen the state of the network even further.

Figures 5(a) and 5(b) show the goodput achieved by connections TCP1 and TCP10 when competing with constant bit-rate (CBR) and bursty (on-off) best-effort traffic, respectively. These simulations used 20 best-effort sources (N_{BE}) whose traffic followed the same path as the adaptive best-effort connections, i.e. from node 11 to node 30. The sending rates of the best-effort traffic in each case totaled the overall non reserved capacity of the bottleneck link. The CBR sources use packet sizes of 1500 bytes. The on-off sources are characterized by inactivity and activity times of 0.6 and 0.4

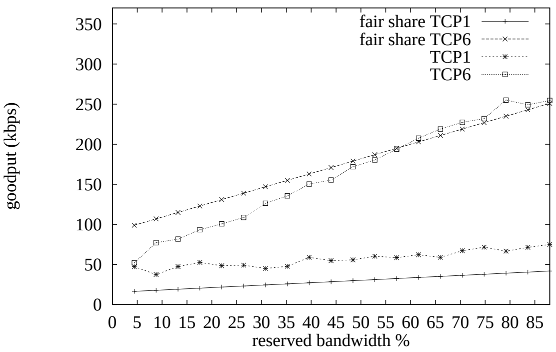

(a) *CBR* traffic

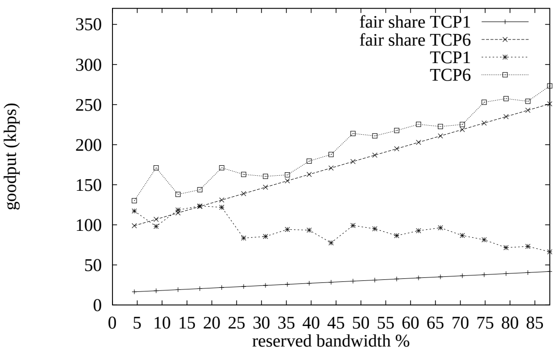

(b) *ON-OFF* traffic

Figure 5: TCP goodput with non-adaptive sources.

seconds, respectively.

The results show that when competing with non-adaptive traffic, the RIO mechanism protects the *IN* packets of the flows with reserved rates from the best-effort packets, guaranteeing the contracted service profile. Nevertheless, the assured flows with higher profile rates and RTT do not get the share of bandwidth that is considered fair. For instance, the connections TCP7, TCP8, TCP9 and TCP10 do not realize their fair rate. The connection TCP6 only get their fair rate for an allocated bandwidth larger than 60%.

5 Conclusions

This article presents a performance evaluation of the Assured Service for the Differentiation of Services in the Internet. This evaluation is done through simulation using the ns-2 simulator. The results obtained show that this service is capable to offer QoS guarantees in terms of bandwidth as soon as the network is well provisioned. In the accomplished simulations, the Assured Service traffic continues receiving the contracted service rate even in presence of non-adaptive flows, i.e. those that don't perform an end-to-end congestion control. However, TCP flows that have higher RTT values and higher contracted service rates don't achieve the expected rate. The perspectives of this work are the evaluation of a mechanism of active queue management with more than two loss precedence values.

References

[1] R. Braden and D. Clark, "Integrated Services in the Internet Architecture: an Overview," *Internet RFC*, Jun. 1994. `rfc1633`.

[2] S. Blake, D. Black, M. Carlson, E Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," *Internet RFC*, Dec. 1998. `rfc2475`.

[3] K. Nichols, S. Blake, F. Baker, and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," *Internet RFC*, Dec. 1998. `rfc2474`.

[4] V. Jacobson, K. Nichols, and K. Poduri, "An Expedited Forwarding PHB," *Internet RFC*, June 1999. `rfc2598`.

[5] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured Forwarding PHB Group," *Internet RFC*, Jun. 1999. `rfc2597`.

[6] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, Aug. 1993.

[7] 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, Aug. 1998.

[8] Network Simulator – NS (version 2), `http://www-mash.cs.berkeley.edu/ns/`

[9] J. Ibanez and K. Nichols, "Preliminary Simulation Evaluation of an Assured Service," *Internet Draft*, Aug. 1998. `draft-ibanez-diffserv-assured-eval-00`.