ABE: Providing a Low-Delay Service within best-effort

Paul Hurley
Mourad Kara
Jean-Yves Le Boudec
Patrick Thiran

March 27, 2001

Abstract

We propose Alternative Best-Effort (ABE), a novel service for IP networks, which relies on the idea of providing low-delay at the expense of maybe less throughput. The objective is to retain the simplicity of the original Internet single class, best effort service, while providing low delay to interactive, adaptive applications. With ABE, every best effort packet is marked as either *green* or *blue*. Green packets are guaranteed a low bounded delay in every router. In exchange, green packets are more likely to be dropped (or marked using congestion notification) during periods of congestion than blue packets. For every packet, the choice of one or other colour is made by the application, based on the nature of its traffic and on global traffic conditions. Typically, an interactive application with real-time deadlines, such as audio, will mark most of its packets as green, as long as the network conditions offer a large enough throughput. In contrast, an application that transfers binary data such as bulk data transfer, will seek to minimise overall transfer time and send blue traffic. We propose router requirements that aim at enforcing benefit for all types of traffic, namely, that green traffic achieves a low delay and blue traffic receives at least as much throughput as it would in a flat (legacy) best-effort network. ABE is different from differentiated or integrated services in that neither packet colour can be said to receive better treatment, thus flat rate pricing may be maintained, and there is no need for reservations or profiles. In this paper, we define the ABE service, its requirements, properties and usage. We discuss the implications of *replacing* the existing IP best effort service by the ABE service. We propose and analyse an implementation based on a new scheduling method called Duplicate Scheduling with Deadlines (DSD). It supports any mixture of TCP, TCP Friendly and non TCP Friendly traffic.

Keywords Alternative Best-Effort, ABE, Traffic Control, best-effort, low delay service, real-time traffic.

1 Introduction

We present an enhancement to the IP best-effort service, Alternative Best-Effort (ABE), which relies on the idea of providing low-delay at the expense of maybe less throughput. The motivation for such a service is twofold. Firstly, there now exists interactive multimedia applications which perform well across a wide range of loss and throughput conditions [1, 2], but for which delay often remains the major impediment [3]. Secondly, unlike differentiated services (see Section 5), we would like to design a service where it is not required to police how much traffic uses the low delay capability, in order to retain the operational simplicity of a single class network.

The paper is organised as follows. In Section 2.1 we define the ABE service and analyse its implications. In particular, we identify and discuss in Section 2.2 the central issue called *green does not*

hurt blue. We also discuss migration issues from the traditional IP service (flat best-effort) to ABE. As a proof of concept, a router implementation is described in Section 3; it is based on the combination of a scheduler called Duplicate Scheduling with Deadlines (DSD) and a traditional control loop (Section 3.3). Implementations were done in the Linux kernel, in the Dummynet network emulator, and in the ns-2 simulator [4]. Simulation results of the ns-2 implementation are shown in Section 4. Section 5 reviews related work and positions ABE with respect to other proposals in differentiated services.

2 The ABE service

2.1 Definition of the service

ABE is defined as follows.

1. ABE packets are marked either green or blue[1].

2. Green packets receive a low, bounded delay at every hop. Realistic values of the per-hop delay bound are discussed later in this section.

3. *Green does not hurt blue:* If some source decides to mark some of its packets green rather than blue, then the quality of the service(delay and throughput) received by sources that mark all their packets blue remains the same or becomes better. This definition is made more specific in Section 2.2.

4. All ABE packets belong to one single best effort class. If the total load is high, then every source may receive little throughput. However, entirely blue sources would experience more throughput than entirely green sources sharing the same network resources.

A consequence of these requirements is that green packets are more likely to be dropped during bouts of congestion than blue packets, or, if Explicit Congestion Notification (ECN) [5] is used, to be marked with the congestion bit. ECN provides congestion feedback to the source by marking a bit in the packet header, enabling it to adjust to feedback without necessarily dropping its packets. For simplicity, in the rest of the paper we consider only non ECN-capable systems.

In essence, ABE can be thought of as allowing an application to trade delay for loss or less throughput, by marking some packets green. The third requirement, green does not hurt blue, derives from the objective that the colour chosen by an application need not be policed. Indeed, if the third requirement is enforced, then an application which decides to mark some packets green must do so because it values the low delay more than a potential increase in loss (or decrease in throughput); otherwise, it would mark its packets blue. In all cases, there is no penalty for other applications, which might choose to mark all their packets blue. This requirement also plays a role in interworking and migration (Section 2.4). Note that ABE supports traffic that may be solely TCP Friendly traffic or non TCP Friendly, or a mixture of the two.

In Figure 1 we illustrate how a TCP friendly source would use the ABE service, by showing a simple simulation where an interactive, adaptive audio source competes with n background sources for one bottleneck. The source has the choice of marking packets blue or green. Assume the source has a required minimum rate R_0 in order to function properly, for a given loss pattern in the network. The rate R_0 is shown by the horizontal dashed line. Also assume that the source is able to forward-correct packet losses, as long as the minimum rate is achieved (see [1] for such an application example; note

---

[1]The choice of the terms blue and green, two neutral colours, is to indicate that none of the two has priority over the other, while green, the colour of the traffic light signal for *go*, indicates low queueing delay.
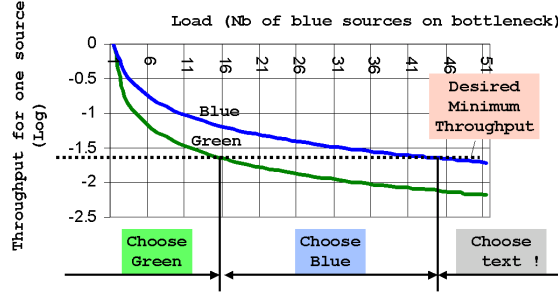
2

Figure 1: A possible strategy for a multimedia source using the ABE service.

that this would not be needed if ECN was used). The choice between green or blue is left to the audio application. It depends on its utility function $u(R, D)$, for a given throughput R and end-to-end network delay D. On this simplified example, we assume that the utility function for our source satisfies (1) $u(R, D) = 0$ for $R < R_0$ and (2) $u(R, D)$ is a decreasing function of D only for $R \geq R_0$. In other words, once a minimum rate R_0 is achieved which provides enough intelligibility, delay becomes the major impediment. For this source, the optimal strategy is to be green in the low load region, blue in the moderate load region, and to disconnect when the load is too high. Note that this example is oversimplified. In general we expect more complex utility functions to be used; see [7] for an actual audio source using such a utility function.

An ABE aware source would probably use a colour mixing strategy, where they would send some green packets and some blue. This would for example be used by a colour adaptation algorithm which sends probe packets of either colour, in order to determine which region the source is currently operating in. This is perfectly permissible and considered normal practice. In fact, apart from possibly policing TCP Friendliness if so required [8, 9], a network supporting ABE does not need to analyse individual flows. Unlike the multimedia source above, a source using TCP is probably more interested in its throughput and should thus mark all its packets blue. Intuitively, it is because ARQ protocols such as TCP are more sensitive to packet loss than to queueing delay, though queueing delay does have an impact (see Section 4 for an illustration).

The IETF mandates that non-TCP sources be TCP Friendly [10], namely, the source should not receive more throughput than a TCP flow would. This is for reasons of fairness and for avoiding congestion collapses. However, it is still the case that many multimedia flows are *not* TCP Friendly. Thus, the requirement that green does not hurt blue applies even if green traffic originates from non TCP Friendly sources. Note that non TCP friendly sources may, in some cases, severely hurt other, TCP Friendly sources, and this is true with or without ABE. The requirement simply means that giving low delay to such sources does not make things worse. Section 4 illustrates by simulation how an implementation based on Duplicate Scheduling with Deadlines supports non TCP Friendly sources.

At very high bit rates, queueing delays are in general expected to be low and high speed backbones probably will not need any delay differentiation. Hence, we currently expect ABE routers to be implemented at network peripherals, where bit rates are of the order of a few Mb/s (or even less for cellular radio systems). The value of the delay bound offered to the green service depends on how many hops are used by one flow. A multimedia flow probably uses a small number (2 to 6) of low speed hops. An interactive audio application has a delay budget of 100-150 msec, out of which 50 msec may be allocated to network delay. As a result, we expect the green per-hop delay bound to be set to a value in the range of 5 to 20 msec.

2.2 Green does not hurt Blue

In this section we define more accurately the service requirement that green does not hurt blue, as introduced in Section 2.1. We subdivide the requirement into two parts. The first part addresses the case of non-TCP friendly sources:

Definition 2.1 *(Local Transparency to Blue) Consider the scenario, flat best-effort, in which a node would forget the colour and thus treat all ABE packets as one single best effort class. The node satisfies local transparency to blue if, for each packet that is blue in the original (ABE) scenario:*

1. *the delay is not larger in the real, ABE scenario than in the flat best-effort scenario*

2. *if the blue packet is not dropped (or marked with a congestion notification) in the flat best-effort scenario, then it is not either in the real, ABE scenario.*

In Section 3, we describe a scheduler (DSD) which provides local transparency to blue.

Now if some sources sending green traffic are TCP Friendly, and greedy, local transparency to blue may not be sufficient. Indeed, a common interpretation of TCP Friendliness is that the source data rate should not exceed θ given by

$$\theta = \frac{s}{R\sqrt{\frac{2p}{3}} + 3t_1\sqrt{\frac{3p}{8}}p(1 + 32p^2)} \tag{1}$$

where R is the round trip time, p the rate of loss events, t_1 the TCP retransmit time (roughly speaking, proportional to the round trip time), and s is the packet size [11]. Thus, it is quite possible that, by becoming green, a TCP Friendly source would be allowed a higher data rate, due to the reduction in round trip time. Such a source would generate more packets than if it was blue, and there is the risk that, in some cases, it would hurt blue packets. This leads to the second part of the requirement:

Definition 2.2 *(Throughput Transparency to Blue) Assume that sources employ a rate adaptation algorithm which conforms to a loss-throughput formula such as Equation (1). To provide blue with throughput transparency, the ABE node should ensure that an entirely green flow gets a lesser or equal throughput than if it were blue.*

Unlike local transparency, throughput transparency seems impossible to implement exactly: On the one hand, it requires knowing the round trip time for every flow, which is not feasible in practice. On the other hand, the rate adaptation algorithm implemented by a source may significantly deviate from a straight application of Equation (1). Indeed, the dependency of rate on round trip time in Equation (1) is not necessarily a desirable feature of a rate adaptation algorithm. It should not be confused with the fact that a source using many hops should receive less throughput. This latter fact is desirable, but it is implemented by having a higher loss ratio. Further discussion on this is provided in [12]. Fixes have been suggested to rate adaptation algorithms, that would remove the dependency of rate on round-trip time [13]. If such fixes where to become widespread, then throughput transparency to blue would be an automatic consequence of local transparency to blue (which can be exactly implemented, for example with DSD). Thus, we consider the requirement for throughput transparency to be loose.

Our approach to providing throughput transparency to blue uses a controller that acts upon a parameter g of the DSD scheduler, which controls the service received by green packets. g is a factor that determines, in case of a tie between green and blue, the probability of forwarding a green packet. The delay and loss ratio are monitored, and using Equation (1), the controller adjusts g to make sure throughput transparency is maintained. The controller solves the issue of evaluating the round trip

time by observing that an under-evaluation of green round trip times is conservative for blues. Thus, the controller assumes that all flows are greedy and have a total round trip time equal to the queueing time at this node plus a fixed, virtual base value (20 msec in the implementation). This value is taken to be small so that it is unlikely that the real value could be below it. A potential problem could be that the drop probability becomes higher than necessary for either green flows with higher round trip times, or non-greedy green flows, or non-adaptive green flows. However, our simulations indicate that the combination of the DSD scheduler and the controller does not seem to produce that problem.

An alternative coarse implementation may consist of avoiding any single green flow from getting too large a fraction of throughput, by examining the drop history for greens, as proposed in [9]. Researching such an alternative is the subject of future work.

2.3 Router Requirements

Following from the discussion in the previous section, a router implementing ABE must:

1. Provide low, bounded delay to green packets; the delay bound is fixed by network management, probably in the 5-20 msec range.

2. Provide local transparency to blue (Definition 2.1).

3. Provide throughput transparency to blue (Definition 2.2).

4. Preserve packet sequence within blue and within green.

5. Keep green packet loss as low as possible, while adhering to the above requirements.

The first three requirements directly derive from the previous discussion. The fourth requirement is because an implementation should try to make the use of green in the service as attractive as possible.

In today's Internet, it is considered desirable to preserve packet ordering, though this is not always enforced. Similarly, an ABE node is expected to preserve packet order as much as possible; however, the delay preference given to green may result in a green packet overtaking a blue one. It is desirable that any ABE implementation is also work-conserving, though this is not strictly a requirement.

2.4 Inter-working and Migration

ABE could be used by an operator in two distinct ways; either as a separate service, or as a replacement to the flat (existing) best effort IP service. In this paper we focus on the latter. Replacing flat best effort by ABE requires a rule for assigning a colour to packets that do not have one (such packets come from a non-ABE source or network). By default, the rule is to assign a blue colour to packets. Indeed, because of the characteristic that green does not hurt blue, ABE unaware sources receive the same service as they would if the network would be flat best effort. An operator might thus introduce ABE and let customers and other carriers gradually move to ABE, without any specific change to charging or control policies.

Conversely, consider an ABE aware source which uses a concatenation of networks, some ABE, some flat best effort. We have mentioned earlier that an ABE aware source probably has to implement a colour adaptation algorithm. Now, depending on traffic conditions, the ABE source might see small or large delays even for green traffic. This implies that the colour adaptation algorithm should not make any quantitative assumption about the value of end-to-end delay guaranteed for green traffic. [14] discusses how the ABE colour can be encoded in the IP packet header.

3   Implementation

In this Section we present a router implementation model to support the ABE service. This implementation assumes that the router has only output port queueing. It is based on a new scheduling concept, Duplicate Scheduling with Deadlines (DSD). We have also undertaken other implementations, based on different scheduling concepts, which include: a differential dropper based implementation in the ns2 simulator [15], (and first outlined in [16]) and in the Linux kernel [17], and a dummy packet based implementation in the Dummynet emulator [17]. We describe DSD, discuss its compliance to the ABE router requirements, as per Section 2.3, and show some experiment results from simulations.

Before delving into the details of the scheme's description, its motivation is first explained. One of the first schemes to implement ABE that might spring to mind is a FCFS scheduling discipline with a threshold drop policy to filter green packets. In such a scheme, blue packets would be accepted when the buffer is not full, while green packets would only be accepted if they can be served with no greater a delay than some maximum d.

Most of the time, though, there would be little or no incentive in being green. What is desired is to provide green with the best service possible while still ensuring that green does not hurt blue. Any significant extra gain by blue packets is at the expense of green ones. The gain blue packets would enjoy under ABE should be kept to a minimum such that there is still an incentive to use green packets whenever appropriate. This can be formalised by the following optimisation problem: Minimise the number of green losses subject to the following constraints:

- Green packets receive a no larger queueing delay than d.

- Local transparency to blue (Definition 2.1) holds.

- The scheduling is work conserving.

- No reordering: Blue (respectively green) packets are served in the order of arrival.

A solution to this problem is the DSD, a new scheduling algorithm based on the concept of *duplicates*. Deadlines are assigned to packets as they arrive, green and blue packets are queued separately, and the deadlines of the packets at the head of blue and green queues are used to determine which one is to be served next.

As previously discussed, throughput transparency as well as local transparency is required for ABE to ensure rate-adaptive green flows do not hurt blue. This is facilitated by the use of a parameter g, which is used in deciding which queue should be served in the event that the deadlines of the packets at the head of each queue can both be met if the other queue was served beforehand. The value of g used at any given time is determined by a control loop as described in Section 3.3. We can now describe DSD in detail.

3.1   Duplicate Scheduling with Deadlines (DSD)

An example of how DSD works is given in Figure 2. Duplicates of all incoming packets are sent to a virtual queue with a buffer size *Buff*. A duplicate is admitted if the virtual buffer is not full. Packets in the virtual queue are served according to FCFS at rate c, as they would be in a flat best-effort. The times at which duplicates will be served are used to assign blue packets *deadlines* at which they would have (approximately) been served in flat best-effort. The original arriving packets are fed according to their colour into a green and a blue queue. Blue packets are always served at the latest their deadline permits subject to work conservation. Green packets are served in the meantime if they have been in the queue for less than d seconds, and are dropped otherwise.
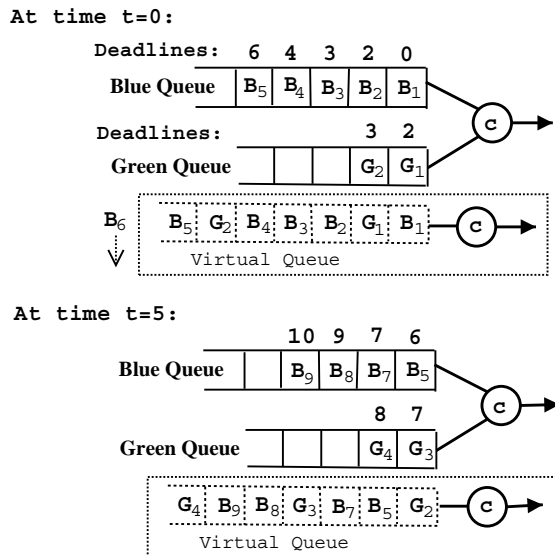
Figure 2: Two snapshots as an example of DSD, at time $t = 0$ (top) and $t = 5$ (bottom). For this example, all packets are the same size and "packet" time is used. To facilitate understanding, we consider first the case where green packets do not undergo the green acceptance test and where $g = 1$. The maximal buffer size is $Buff = 7$ packets. The maximum green queue wait is $d = 3$ packets. B and G denote blue and green packets respectively. In the first snapshot, B_1 is served at time $t = 0$ in order to meet its deadline, then G_1, B_2, B_3, B_4. G_2 has to be dropped from the green queue because it has to wait for more than $d = 3$, whereas B_6 had to be dropped because the virtual queue length was $Buff$ when it arrived. At time $t = 5$, we reach the situation of the second snapshot. As no blue packet has reached its deadline yet, G_3 can be served, followed by B_5, B_7, G_4, B_8, and B_9.

A blue packet is dropped if its duplicate was not accepted in the virtual queue. Otherwise, it is tagged with a deadline, given by the time at which its duplicate will be served in the virtual queue, and placed at the back of the blue queue.

A green packet is accepted if it passes what is called the *green acceptance test* and dropped otherwise. A green packet arriving at time t fails the test if the sum of the length of the green queue at time t (including this packet), and of the length of the first part of the blue queue that contains packets tagged with a deadline less than or equal to $t + d + pg_{new}$, where pg_{new} is the transmission delay for the incoming green packet, is more than $c(d + pg_{new})$, and passes otherwise. The use of the test ensures the total buffer occupancy, namely the sum of the green and blue queue lengths, does not exceed $Buff$, which is discussed in Section 3.2.

Consider again the example in Figure 2, except green packets are now enqueued only if they pass the green acceptance test. This amounts here to accepting a green packet at time t if the number of green packets in the queue at time t, augmented by the number of blue packets in the queue with a deadline between $[t, t+4]$ is no more than 4. The only difference from Figure 2 is that G_2 is no longer enqueued. Indeed, when it arrived, the green queue already contained packet G_1, and the blue queue contained packets B_1, B_2 and B_3. The total queue length at time was 5 packets (including G_2), and so G_2 fails the test.

An accepted green packet is then assigned a deadline which is the sum of its arrival time plus its maximum waiting time d, and placed at the back of the green queue. At each service time, a decision is made as to which queue to serve. The serving mechanism's primary function is to ensure that blue

7

packets are always served no later than their deadlines. The best performance green could receive would be to then serve the green queue as much as possible, subject to this restriction. However, as previously discussed, in addition to local transparency, throughput transparency is needed to ensure green adaptive applications do not benefit too much from lower delay.

It can happen at service time that both blue and green packets at the head of their respective queues are able to wait, as letting the other packet go first would still allow it to be served within its deadline. For the purpose of supporting throughput transparency, when this situation arises, the packet serving algorithm uses the current value of the *green bias* g, a value in the range $[0, 1]$, to determine the extent to which green is favoured over blue. More precisely, when both blue and green packets can wait, g is the probability that the green packet is served first. The value $g = 1$ corresponds to the case where green is always favoured. Conversely, the value $g = 0$ corresponds to the systematic favouring of blue packets. In the example in Figure 2, the packets served would have thus been, successively, B_1, B_2, G_1, B_3, B_4, B_5, B_7, G_3, G_4, B_8 and B_9.

A value of g less than 1 causes the delay for green traffic to be increased. This increase in delay for green TCP Friendly traffic reduces their throughput, thereby enabling blue traffic to increase its throughput. Increasing the delay of non TCP Friendly traffic may not reduce their throughput, but blue flows are, in the worst case, as equally protected from this type of traffic as they would have been in a flat best-effort service. The value of g chosen is made according to a control loop which is described in Section 3.3.

All green packets who miss their deadline, by waiting for more than d seconds (these packets are said to have become *stale*, are removed from the green queue. At service time, the possible events that arise and packet served by DSD are as follows:

Event	What is Served?
Both queues empty	Nothing
Green queue empty, blue queue not empty	Head of blue queue
Blue queue empty, green queue not empty	Head of green queue
Head of blue queue cannot wait	Head of blue queue
Head of blue queue can wait, head of green queue cannot	Head of green queue
Head of green queue and of blue queue can wait	With probability g, head of green queue else head of blue queue

Pseudocode for DSD is given below. Let *now* be the current time, p.deadline denote the latest time a packet p can remain in the queue (whose value is tagged onto packet p), and p.transmissionDelay denote its transmission delay.

Packet enqueueing Algorithm

```
packet p arrives at the output port
dup = p
Add dup to the virtual queue

if p is blue
  if dup was dropped from virtual queue
    drop p
  else
    vd = queueing delay received by dup
         in virtual queue
    p.deadline = now + vd
    add p to blue queue
else // p is green
  if p fails "green acceptance test"
```

```
            drop p
        else
            p.deadline = now + d
            add p to green queue
```

Packet Serving Algorithm

```
    drop stale green packets, those packets from green
    queue who cannot be served within their deadline

    headGreen = packet at head of green queue
    headBlue = packet at head of blue queue

    if headGreen = 0 // no green to serve
        if headBlue !=0
            serve headBlue
    else if headBlue = 0 // no blue to serve
        serve headGreen
    else // both queues contain packets
        p_g = headGreen.transmissionDelay
        dead_g = headGreen.deadline
        p_b = headBlue.transmissionDelay
        dead_b = headBlue.deadline

        if now > dead_b − p_g
            serve headBlue // because it cannot wait
        else if now > dead_g − p_b
            serve headGreen // because it cannot wait
        else with probability g
            serve headGreen
        else
            serve headBlue
```

"green acceptance test"
```
    pg_new = transmission delay for p
    l_g = length of green queue
    l_b = length of packets in blue queue with
            deadlines <= now + d + pg_new
    if l_g + l_b > c * (d + pg_new)
        return "p fails test"
    else
        return "p passes test"
```

It is not mandated that the virtual queue employ drop-tail queueing although in the simulation results shown it is. An Active Queue Management scheme such as RED [18] can be supported for blue traffic by applying it to the virtual queue, and using those results in assigning losses and deadlines.

Removing stale green packets, those packets from the green queue whose deadline cannot be met, involves a search of this queue up to the first alive green packet. In practice, these stale green packets can be cleaned up can between service times, as was done in our dummynet implementation, and it has proven sufficiently fast. However, for really high speed networks this search may prove expensive. As such, further optimisations of this algorithm may be needed, and are the subject of on-going work.

Some of the building blocks in DSD are similar to those in other scheduling techniques. The calculation and tagging of deadlines to each arriving packet is also performed by Earliest Deadline First (EDF)[19] schedulers and its variants. However, EDF sorts packets according to deadlines, whereas DSD remains FIFO within each of its two queues, and the deadlines are used at service to determine whether the head of the green or the head of the blue queue should be served. The use of a virtual queue has been used many times, for example in an admission control context [20].

3.2   Properties of DSD

Let us list now some of the most important properties of DSD:

1. Buffer space constraint: the total buffer occupancy for real packets (green and blue counted together) is always less than $Buff$, the size of the virtual queue used for duplicates.

2. All accepted blue packets will be served by their deadlines. Accepted blues are thus served at the same time as, or earlier than, they would have been in flat best-effort.

3. All green packets are served before d, or are otherwise dropped. Low bounded (per hop) delay for the green packets is enforced by dropping a green packet that waits or would have to wait d seconds in the queue.

4. The green acceptance test does not unnecessarily drop green packets, in the following sense. If all enqueued green packets are to be served, then it is impossible to serve, within d seconds, an incoming green packet that arrived at time t and would violate the green admission test. In addition, if $g = 1$, the green admission test is optimal in the sense that it accepts exactly the green packets that will be served within d seconds; otherwise. Note that if $g < 1$, some green packets may become stale and be dropped by the Packet Serving Algorithm.

Points 2 and 3 follow immediately from Section 3.1 and the pseudo-code. Points 1 and 3 are proven in [6].

3.3   Control loop for DSD

For the reasons described in Section 2.2, unlike local transparency, maintaining throughput transparency is by its nature approximate. g is used as a control parameter to balance the throughputs of green and blue. These are estimated from Equation (1), using a fixed value to represent the non-queueing delay portion of the round trip-time of a flow. This value is chosen to be small, since this favours blue traffic. For the purposes of the control, flows are assumed to be greedy, since this also increases the protection to blue flows.

Estimates for the delay and loss ratio for both green and blue traffic are monitored. Let $\theta_b(t)$ and $\theta_g(t)$ be these estimates for the blue and green throughput respectively at time t. The value of g is chosen so that their ratio is close to a desired value γ, which is slightly larger than 1, to provide blues with a small advantage in throughput, and to offer a safety margin for protection from errors in throughput estimation.

g is updated every T seconds according to the control law,

$$g(t + T) = (1 - \alpha)g(t) + \frac{\alpha}{1 + (\gamma\theta_g(t)/\theta_b(t))^K},$$ (2)

where $\alpha \in (0, 1)$ and $K > 0$ are two control parameters. T is a chosen parameter of the system which determines the rate of update of g. The initial value of g upon commencement of control can be chosen to be 1, namely, $g(0) = 1$.
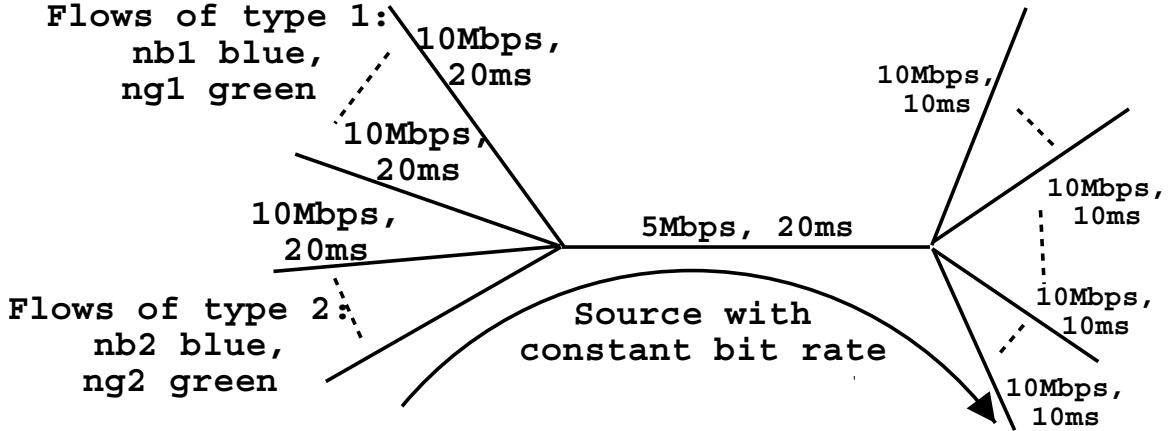
Figure 3: Simulation topology.

Let us briefly explain the rationale behind this choice of control law, which we do not claim to be optimal. In the ideal case where $\theta_b = \gamma\theta_g$, there should not (a priori) be any bias against blue nor green, and the value of g should be $1/2$. If θ_b is larger than $\gamma\theta_g$, then g must be increased, and vice versa if θ_b is smaller than $\gamma\theta_g$. We wish to maintain symmetry in the amount by which we increase or reduce g: the amount by which g is increased if $\theta_b/\gamma\theta_g$ is multiplied by some factor A should be the same amount by which g is decreased if $\theta_b/\gamma\theta_g$ is divided by the same factor A. Denoting by $\xi = \ln(\theta_b/\gamma\theta_g)$, the targeted g should therefore be an increasing function F of ξ with central symmetry around 0, and such that $F(0) = 1/2$, $F(\xi) = 0$ for $\xi \to -\infty$ and $F(\xi) = 1$ for $\xi \to +\infty$. Such a function is the sigmoid function

$$F(\xi) = \frac{1}{1 + \exp(-K\xi)}$$

where K is the slope of the function at the origin. The larger K, the closer the sigmoid function to the step (Heavyside) function

$$\overline{F}(\xi) = \begin{cases} 1 & \text{if} \quad \xi > 0 \\ 1/2 & \text{if} \quad \xi = 0 \\ 0 & \text{if} \quad \xi < 0. \end{cases}$$

The control law

$$g(t + T) = g(t) + \alpha(F(\xi) - g(t))$$

where α is the adaptation gain, will therefore bring g to the targeted value. If $0 \leq \alpha \leq 1$, this control law keeps $g(t)$ between 0 and 1 at all times t. Replacing ξ by $\ln(\theta_b/\gamma\theta_g)$ in this equation, we get the control loop equation for the green bias as given in Equation (2).

The use and control of the green bias g is only one possible scheme, and its performance can be improved, for example, by taking into account the deadlines of all the packets in the queues, and not just those at the head. Such extensions are the subject of further investigation.

4 Simulation Results

In this section, we show simulations, using ns-2, of DSD run on the topology shown in Figure 3. There are $n_{b,1}$ blue sources and $n_{g,1}$ green sources with an outgoing link propagation delay of 20ms (sources of type 1), and $n_{b,2}$ blue and $n_{g,2}$ green sources with an outgoing 10Mbps link of propagation delay 50ms (sources of type 2). All sources pass the 5Mbps link L of propagation delay 20ms, and
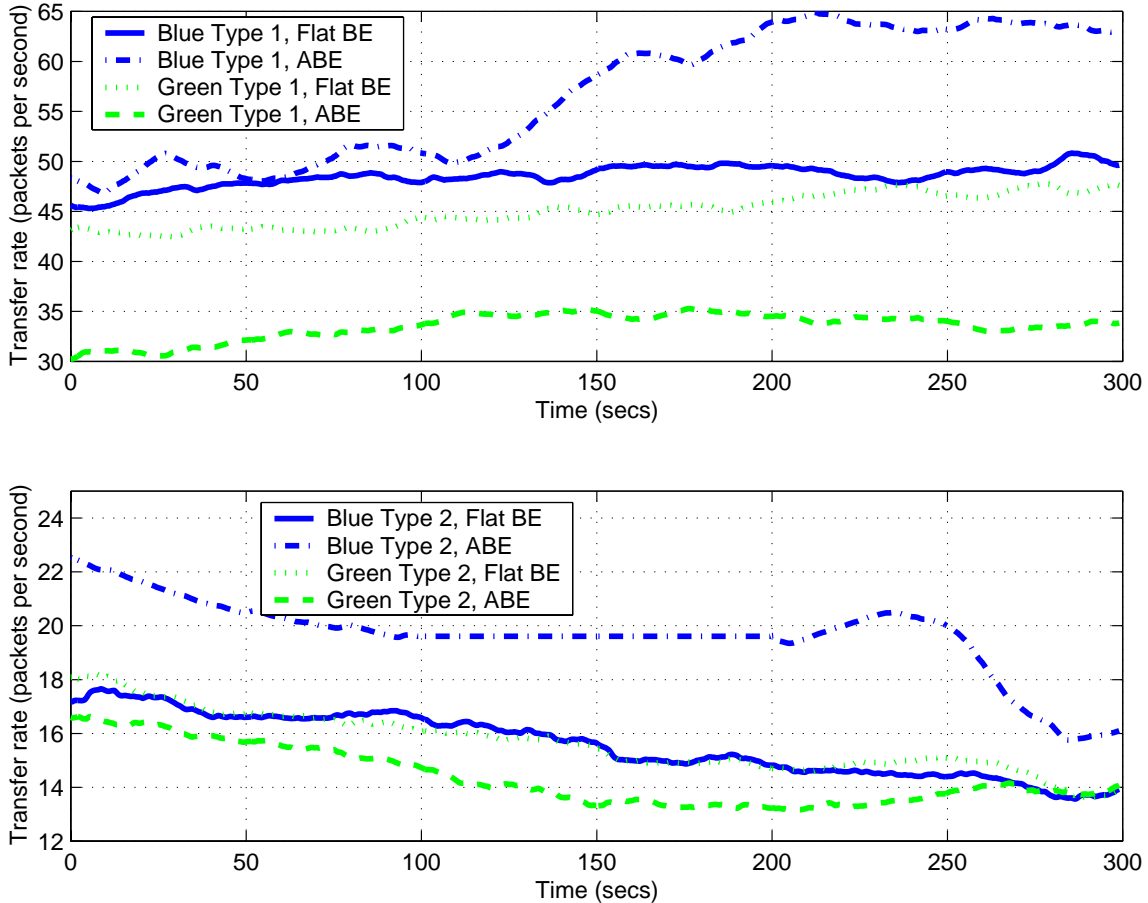
Figure 4: Average packet transfer rate for green and blue connections, as a function of time t, when the router implemented DSD and when it implemented flat best-effort. The results are obtained by simulating the network described on Figure 3, with 5 blue flows and 5 green flows of each type, namely $n_{b,1} = n_{b,2} = n_{g,1} = n_{g,2} = 5$, and no CBR traffic.

terminate via a 10Mbps link of propagation delay 10ms. These blue sources are TCP Reno, and the green sources are the TCP Friendly algorithm as described in [7]. There is also green traffic which sends a constant rate r (CBR) and passes through the link L.

The router buffer size was 60 packets (i.e. $Buff = 60$) and the maximum delay green can queue for, d, was $0.04s$. For simplicity, the size of all packets is fixed at 1000 bytes. The control loop updates its value of g every $0.5s$ (i.e. $T = 0.5$), the gain parameter α was 1.1, and the conservative value of $20ms$ was taken to be the round-trip time used for estimating throughput. The router distinguishes green and blue by a bit in the packet header. Each simulation ran for 300 seconds of simulated time.

The first goal of this simulation study was to show that green does not hurt blue, under a variety of conditions; namely when there are flows of various round-trip times, where green flows may be either TCP Friendly or non TCP Friendly, may be greedy or not and where flows may send a mixture of green and blue packets. In addition to this, we illustrate that green flows benefit from low delay (at the expense of less throughput), and show the effect DSD has on the loss rates of each traffic type. The reference simulation is what happens in the flat best-effort scenario, in the absence of ABE, where all packets are treated equally at the router.

Density Plots for Green Traffic Queueing Delay with and without ABE
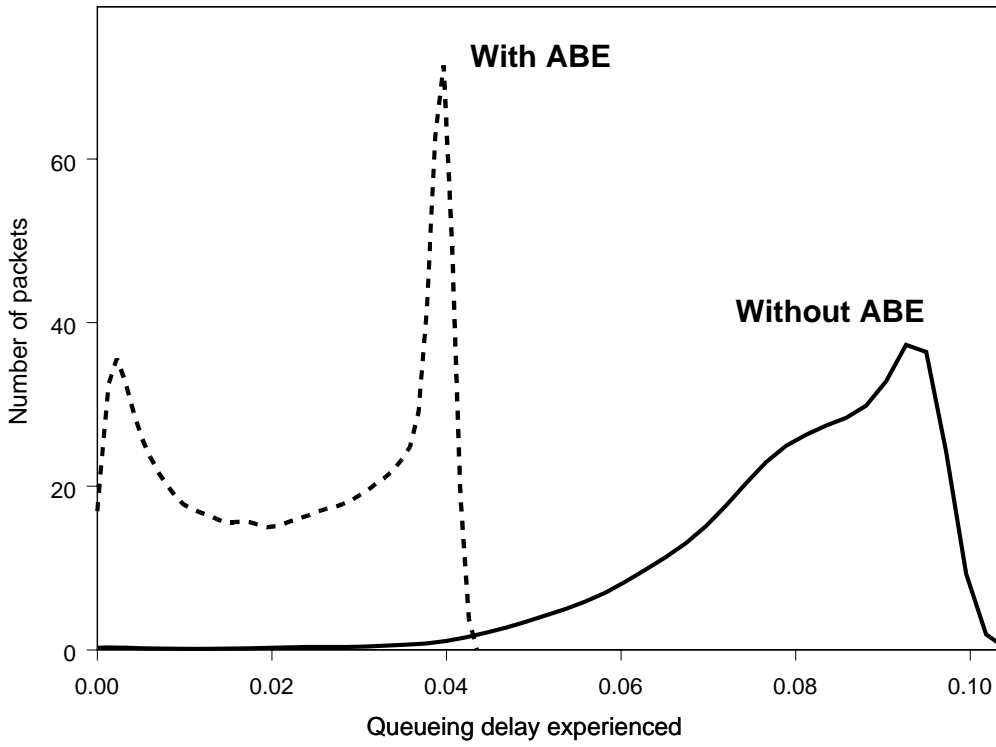

Figure 5: Density plot of queueing delay received by green packets under ABE/DSD and flat best-effort. 5 blue TCP and 5 green TCP Friendly flows of each type ($n_{b,1} = n_{b,2} = n_{g,1} = n_{g,2} = 5$)

We first examine some scenarios when there are only TCP and TCP Friendly flows. For the case where there are 5 blue TCP and 5 green TCP Friendly flows of each type ($n_{b,1} = n_{b,2} = n_{g,1} = n_{g,2} = 5$), Figure 4 shows the average transfer rate for each blue and green connection, of both types at each time t. Figure 5 shows the end-to-end delay distributions received for green packets under ABE and flat best-effort. Blue flows of each type receive more throughput with ABE than the did in flat best-effort, thus benefiting from the use of ABE. Green flows receive less, and in exchange, the green queueing delay is small and bounded by $d = 0.04s$. The green loss ratio was 4.97% when using ABE, and 3.3% in the flat best-effort, while the blue loss ratio decreased from 3.2% to 2.5% when moving to ABE. The extra throughput that blue flows of type 1 receive over type 2 flows follows from the lower round trip-time they experience.

The same number of blue and green sources does not occur in general, and ABE is designed to work independently of asymmetry in the amount of green and blue traffic. For the case where there are 5 blue TCP and 3 green TCP Friendly flows of type 1 ($n_{b,1} = 5$, $n_{g,1} = 3$) and 3 blue TCP and 5 green TCP Friendly flows of type 2 ($n_{b,2} = 3$, $n_{g,2} = 5$), Figure 6 shows that again green does not hurt blue.

The situation where blue traffic is TCP, and green traffic is no longer TCP Friendly, but a constant bit-rate source is now examined. Here there are 5 blue TCP flows of each type ($n_{b,1} = n_{b,2} = 5$) and CBR green traffic which sends at 1Mbps. The number of packets received for each blue traffic type and for the CBR source is shown as a function of time in Figure 7. What we see is that the blue traffic
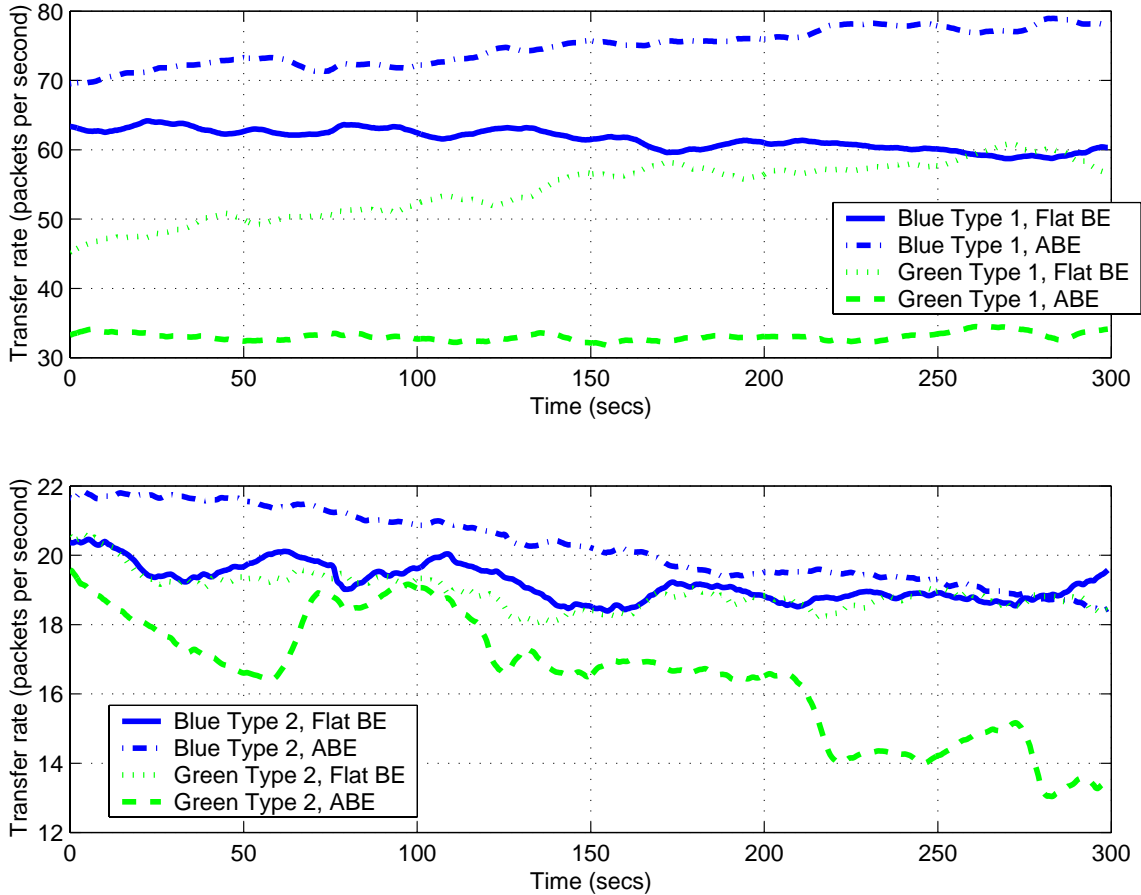
Figure 6: Average packet transfer rate per green and blue connection, as a function of time t, when the router implemented ABE/DSD and when it implemented flat best-effort. The results are obtained by simulating the network described on Figure 3, with $n_{b,1} = 5, n_{g,1} = 3, n_{b,2} = 3, n_{g,2} = 5$

receives more slightly throughput with DSD than with flat best-effort, due to the local transparency property, and the non TCP Friendly CBR traffic receives less.

We now look at the scenario where there is blue TCP traffic ($n_{b,1} = n_{b,2} = 5$), and green traffic is composed both of TCP Friendly sources ($n_{g,1} = n_{g,2} = 5$) and CBR traffic of rate 1Mbps. The average packet transfer rate for the blue and green of type 1, and for the CBR source as a function of time is shown in Figure 8. The results for type 2 traffic is omitted for easy of reading.

Finally, we look at the case where TCP friendly flows mix their traffic. In this instance, the flows do not logically decide, based on network conditions, how much traffic to send as blue and how much to send as green, and this is the subject of further work. Rather, a simple method is used here, where, by randomisation, the TCP Friendly sources of type 1 send approximately 20% blue packets and 80% green. Again $n_{b,1} = n_{b,2} = n_{g,1} = n_{g,2} = 5$ and there is CBR green traffic of 1Mbps. Figure 9 shows that again, green does not hurt blue. The mixed colour sources get a throughput that lies between what they would have got had they been 100% green and what they would have got if that had been 100% blue.
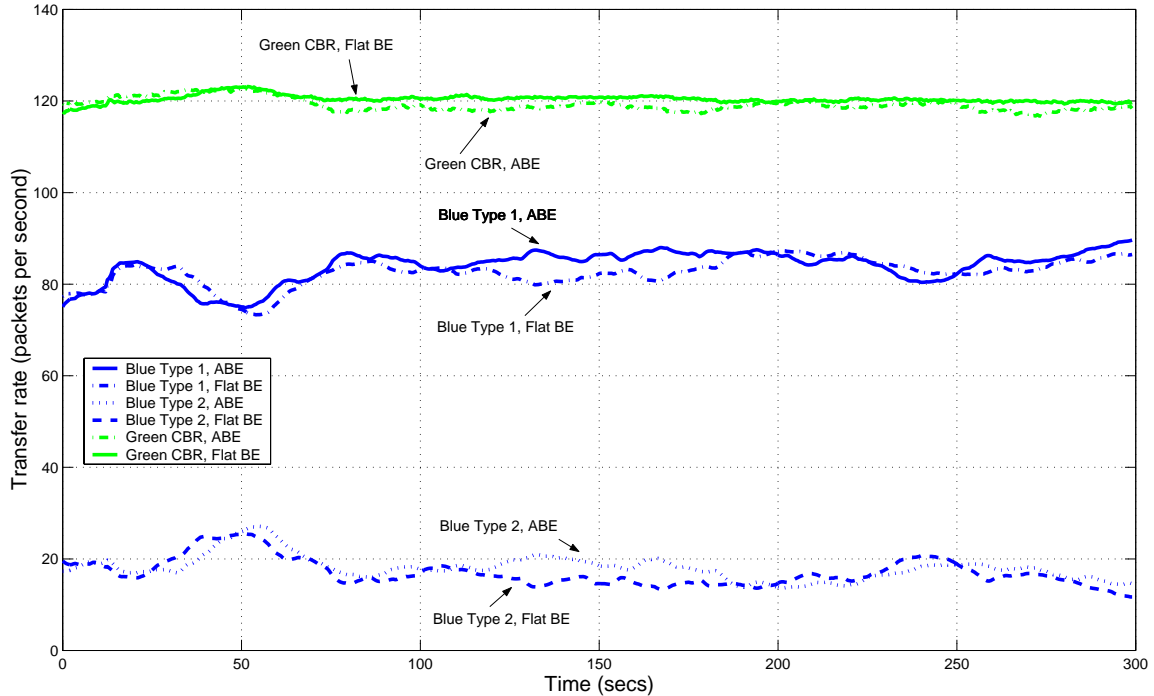
Figure 7: Average packet transfer rate per green and blue connection, as a function of time t, when the router implemented ABE/DSD and when it implemented flat best-effort. There are 5 blue flows of each type and a CBR flow of 1Mbps which is green.

5   Related Work

As the Internet evolves towards a global communication infrastructure, a number of propositions exist for providing QoS architectures. These aim to support more sophisticated services than those provided by flat best-effort services. Currently there are two broad families for QoS provision and both are based on some form of priority and service differentiation.

The first family of solutions, Integrated Services, (IntServ) uses reservations (admission control) and requires routers to manage per flow states and perform per flow operations. It also requires per flow accounting and charging. The second family of solutions, Differentiated services (DiffServ), is based on a coarser notion of QoS, focussing on aggregates of flows in the core routers and intending to differentiate between service classes rather than provide absolute per flow QoS measures.

Integrated services have been shown to exhibit much higher flexibility and assurance level than those provided by Differentiated services. However the main disadvantages of these services are that they are less scalable and robust than differentiated services. Hence, these latter services have been the focus of attention lately mainly because they move the complexity of QoS provision from the core to the edges of the network where it may be feasible to maintain a restricted amount of per-flow state. Often Integrated services are identified as being supported by statefull network architectures (because of the per-flow management), whilst differentiated services as underpinned by stateless network architectures.

An example of such an architecture is SCORE (Scalable Core) proposed by Stoica and Zhang [21] with the aim of providing guaranteed services without per-flow state management. They proposed the Dynamic packet State (DPS) technique to estimate the aggregate reservation rate and use that estimate to perform admission control. To achieve this, they perform Core-Stateless fair queueing

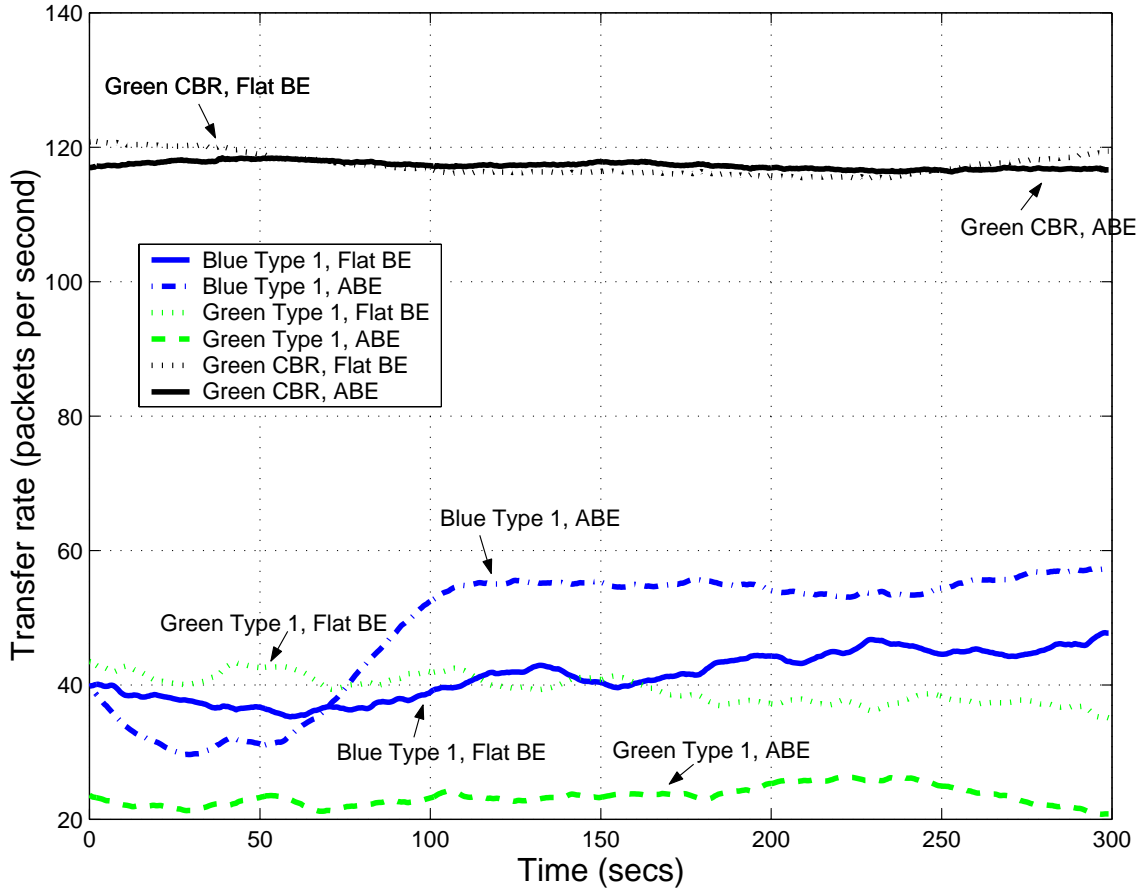Figure 8: Average packet transfer rate per green and blue connection of Type 1, and for the CBR source as a function of time t, when the router implemented ABE/DSD and when it implemented flat best-effort. The CBR source sends at 1Mbps and there are 5 of each other type of source running.

(CSFQ) using DPS to encode dynamic per flow state in the context of approximating the Fair queueing algorithm. Nandagopal et al. [22] also proposed a core stateless QoS architecture (called Corelite) which offers per-hop per-class relative average delay differentiation and end-to-end delay adaptation.

There are several proposals for supporting QoS through differentiated services. Crowcroft [23] proposed a low delay service, analysed by May et al [24], coded with a single bit. Turning on this bit ensures that the packet receives serving priority while constrained to a smaller buffer size. Depending on the input traffic and the buffer sizes of both types of traffic, this typically would result in the low delay traffic also having more throughput. Similarly, Expedited Forwarding [25] (EF) aims to provide extremely low loss and low queueing delay guarantees. SIMA [26] offers applications the choice of a level (0-7) of how "real-time" its traffic is, with each level having relatively lower delay and loss ratio than the previous one.

Dovrolis et al [27] described a proportional differentiation model where the quality between classes of traffic is proportional and thus can be performed independently of the load within each class. Central to their work was the utilisation of two packet schedulers BPR (Backlog Proportional Rate) and WTP (Waiting-Time Priority) to approximate the behaviour of the proportional differentiation model. Moret and Fdida [28, 29] also described a two-class proportional differentiation model called Proportional Queue Control Mechanism (PQCM). Both studies propose controlling the relative queueing
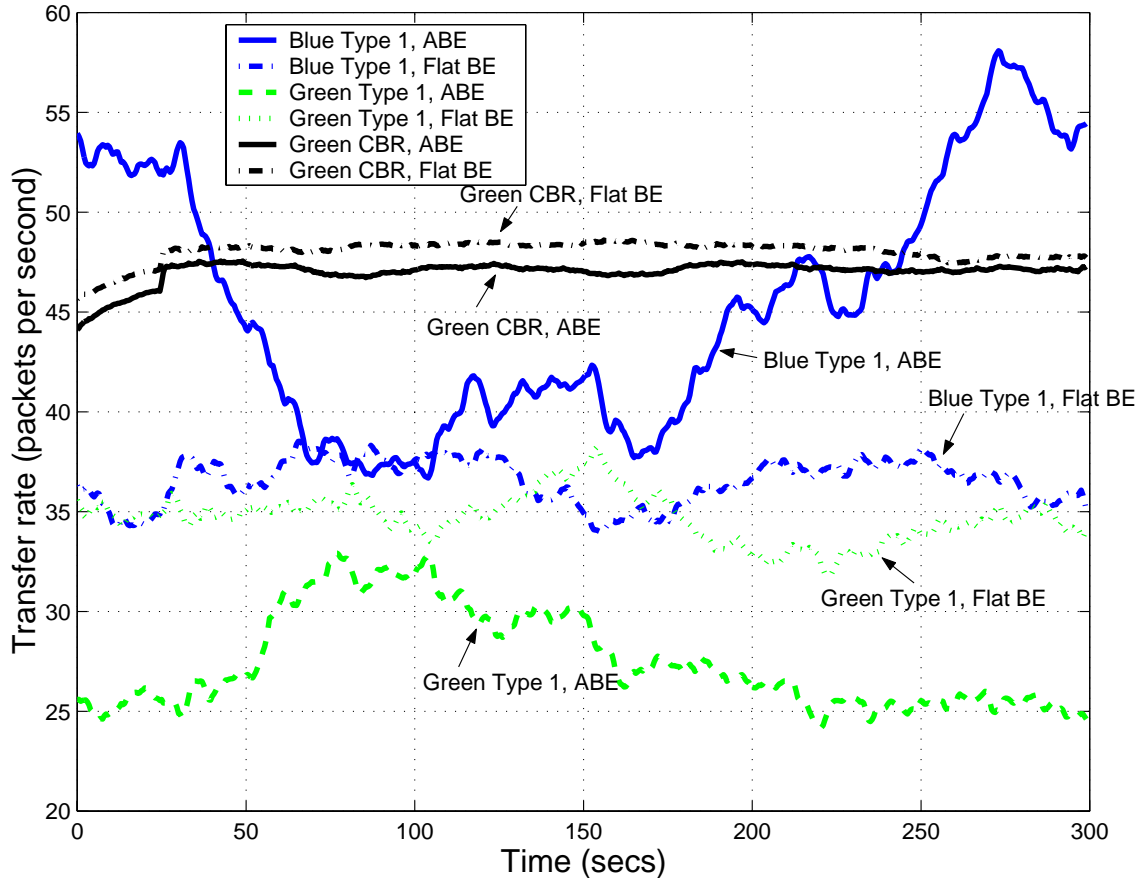
Figure 9: Average packet transfer rate per green and blue connection, as a function of time t, when the router implemented ABE/DSD and when it implemented flat best-effort. The results are obtained by simulating the network described on Figure 3, with $n_{b,1} = n_{g,1} = 3 = n_{b,2} = 3 = n_{g,2} = 5$, and the TCP Friendly sources of type 1 sent 20% blue and 80% green packets.

delays between classes.

All of these proposals couple low delay with improved throughput, and are some form of priority. They can be used to support adaptive and non-adaptive interactive applications, provided that some form of admission control is performed. They can provide a premium service, at a price that has to be higher than the best effort service (otherwise all traffic would use the better service). In contrast, ABE green packets cannot be said to receive a better treatment than blue ones and ABE may be introduced as a replacement for the existing best effort service. On the other hand, ABE is not suited to support multimedia applications which require hard guarantees and cannot adapt.

Assured Forwarding (AF) [30] is also a differentiated service. It divides AF traffic into classes within each there are distinct levels of drop precedence. It offers an assurance that IP packets are forwarded with high probability as long as the aggregate input traffic within a class does not exceed an agreed profile. The authors also suggest that an AF class could be used to implement a low delay service where low loss is not an objective, by allocating an AF class with a low buffer space (call it the low delay AF class). Such a service is in principle different from ABE, which views all blue and green packets as one class; the service received by green packets is dependent on the amount of green *and* blue traffic. In contrast, the performance of an AF low delay class is not expected to be affected by

the amount of best effort traffic. In that sense, the low delay AF class is a differentiated service which requires differentiated charging, contrary to ABE. Hence, ABE can be viewed as being positioned between the flat best-effort service and AF.

Lastly, destination drop might appear as an alternative to ABE that would require no support from the network. This alternative would consist in having the destination drop all packets that arrive too late, say after a transit deadline. However, it wastes network resources, since packets are dropped after being carried by the network, and the overall performance of such a scheme can become very poor [15].

6 Conclusions

We have described ABE, a new service which enables best-effort traffic to experience a low delay, at the expense of possibly more throughput. ABE is targeted at providing low delay, with no concept of reservation or signalling and while retaining the spirit of a flat rate network. The service choice of green or blue is self-policing since the user/application will be coaxed into choosing one or the other or indeed a mixture of both, based on its traffic profile objectives. ABE allows a collection of rate-adaptive multimedia applications to drive the network into a region of moderately high load and low delay. It also allows such an application to trade reduced throughput for low delay, thus in some cases increasing its utility. The design of a multimedia adaptive application that would exploit the new degree of freedom offered by ABE can be found in [7]. Note however that ABE also brings benefits if there are non-adaptive UDP applications.

It should be stressed that ABE is a new service in its own right and not a substitute for reservation or priority services. In contrast, with ABE, both delay sensitive (green) and throughput sensitive (blue) traffic share the same resources, and high load in any of the two pools affects the other. We proposed to introduce ABE as a replacement for the existing best effort Internet service.

We have defined the ABE service, its requirements and properties. We also addressed deployment issues. In addition, we have presented a router implementations based on a new scheduling scheme (DSD), and discussed its compliance. Our simulation results show the benefits of the new degree of freedom offered by ABE in the best-effort services. We have found that under ABE, blue packets received more throughput than under a flat best-effort network while giving a low bounded delay to green packets.

References

[1] J. Bolot, S. Fosse-Parisis, D. Towsley. Adaptive FEC-Based Error Control for Interactive Audio in the Internet. *Proceedings of IEEE Infocom 99*.

[2] C. Diot, C. Huitema, T. Turletti. Multimedia Application should be Adaptive. *HPCS*, Aug. 1995.

[3] C. Huitema. Quality today in the internet. ftp.telecordia.com/pub/huitema/stats/quality.today.html

[4] ns v2 simulator. http://www.isi.edu/nsnam/ns

[5] S. Floyd. TCP and Explicit Congestion Notification. *ACM Computer Communication Review*. V. 24 N. 5, October 1994, p. 10-23.

[6] P. Hurley, M. Kara, J.Y. Le Boudec, P. Thiran. A Novel Scheduler For a Low Delay Service Within Best-Effort. Technical Report, DSC/2001/014, EPFL-DSC, February 2001.

[7] C. Boutremans, J.Y. Le Boudec. Adaptive delay aware error control for internet telephony. Technical Report DSC/2000/031, EPFL-DSC, http://dscwww.epfl.ch, 2000.

[8] B. Suter, T.V. Lakshman, D. Stiliadis, A. Choudhury. Design Considerations for Supporting TCP with Per-flow Queueing. *Proceedings of IEEE INFOCOM 98*.

[9] S. Floyd, K. Fall. Promoting the Use of End-to-End Congestion Control in the Internet. IEEE/ACM Transactions on Networking, August 1999.

[10] TCP Friendly web site. http://www.psc.edu/networking/tcp_friendly.html

[11] J. Padhye, V. Firoiu, D. Towsley, J. Kurose. Modeling TCP Throughput: A Simple Model and its Empirical Validation. *Proceedings of SIGCOMM'98.*

[12] M. Vojnovic, J.-Y. Le Boudec, C. Boutremans. Global Fairness of additive-increase and multiplicative-decrease with heterogeneous round-trip times. Proceedings of IEEE INFOCOMM'2000, Tel Aviv, Israel, March 2000.

[13] T. Henderson, E. Sahouria, S. McCanne, R. Katz. Improving Fairness of TCP Congestion Avoidance. *Proceedings of IEEE Globecom '98, Sydney, Australia, Nov. 1998*

[14] The Alternative Best-Effort Service. Internet Draft, Work in Progress. draft-hurley-alternative-best-effort-01.txt

[15] P. Hurley, J.Y. Le Boudec, P. Thiran. The Alternative Best-Effort Service. Technical Report Research Report DSC1999/036, EPFL-DSC, http://dscwww.epfl.ch, 1999.

[16] P. Hurley, J.Y. Le Boudec. A Proposal for an Asymmetric Best-Effort Service. Proceedings of IEEE/IFIP IWQoS '99, London, England., May 1999

[17] ABE Project Web Page. http://www.abeservice.org

[18] S. Floyd, V. Jacobson. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking*, V.1 N.4, August 1993, p.397-413.

[19] H. Zhang. Service Disciplines for Guaranteed Performance Service in Packet-Switching Networks. Proceedings of the IEEE, Vol. 83 No. 10, October 1995.

[20] T. Ferrari, W. Almesberger, J.Y. Le Boudec. SRP: a Scalable Resource Reservation Protocol for the Internet. Computer Communications, September 1998, Vol. Vol 21 Number 14 No. Special issue on 'Multimedia networking', 1200-1211.

[21] I. Stoica, H. Zhang. Providing Guaranteed Services without Per Flow Management. 81-94, Proceedings of the ACM Sigcomm 99.

[22] T. Nandagopal, N. Venkitaraman, R. Sivakumar, V. Bharghavan. Relative Delay Differentation and Delay Class Adaptation in Core-Stateless Networks. IEEE Infocom 2000, Tel Aviv, Israel.

[23] J. Crowcroft. All you need is just 1 bit. Keynote Presentation. *IFIP Conf. on Protocols for High Speed Networks*, Oct. 1996.

[24] M. May, J. Bolot, C. Diot, A. Jean-Marie. 1-bit Schemes for Service Discrimination in the Internet: Analysis and Evaluation. *Technical Report no 3238*, INRIA.

[25] V. Jacobson, K. Nichols, K. Poduri. An Expedited Forwarding PHB. RFC2598.

[26] K. Kilkki, J. Ruutu. Simple Integrated Media Access - an Internet Service Based on Priorities 6th International Conference on Telecommunication Systems, 1998.

[27] C. Dovrolis, D. Stiliadia, P. Ramanathan. Proportional Differentiated Services: Delay Differentiation and Packet Scheduling. Proceedings of ACM SIGCOMM'99.

[28] Y. Moret, S. Fdida. A Proportional Queue Control Mechanism to Provide Differentiated Services. *International Symposium on Computer Systems*, Belek, Turkey, October 1998.

[29] Y. Moret, S. Fdida. A proportional Queue Control Mechanism to Provide Differentiated Services. International Symposium on Computer System, Belek, Turkey, October 1998.

[30] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski. Assured Forwarding PHB Group. RFC2597.