Statistical Bounds on the Drop Probability of Assured Forwarding Services in DiffServ Interior Nodes under the Processor Sharing Scheduling Discipline

Brahim Bensaou
Computer Science Department
Hong Kong University of
Science and Technology
Clear Water Bay, Kowloon, Hong Kong
brahim@cs.ust.hk

Shixin Zhuang* and Xiren Cao
E & E Engineering Department
Hong Kong University of
Science and Technology
Clear Water Bay, Kowloon, Hong Kong
sxzhuang@bu.edu & eecao@ee.ust.hk

Abstract

This paper addresses the problem of modelling and analysing an interior node in IETF's DiffServ services model. Specifically it is concerned with the performance of the DiffServ assured forwarding service category in presence of a premium service class. In this model, the network node shares its outgoing link capacity between a Premium service representing the Expedited Forwarding (EF) per-hop behavior, and two classes of Assured service, that represent two classes of the Assured forwarding (AF) per-hop behavior. It this paper, the traffic is modelled as Markov modulated fluid sources, and we focus on a system where out of profile traffic is dropped at the edge of the network thus both AF queues support only one drop precedence. Using a decomposition approach, approximations, and spectral analysis, we are able to derive upper and lower bounds on the tail of the distribution of the buffer content for both AF classes given a generalized processor sharing scheduling is used to differentiate the two classes. Such approximate analysis of the interaction between traffic classes can help to achieve a better understanding of this type of networks; enables the provision of throughput differentiation as defined by the AF PHB through the GPS scheduler while quantifying delay; and finally helps simplify greatly the design of bandwidth brokers that do not rely on long term bandwidth (over) provisioning.

1 Introduction

In the current Internet best-effort service model, the quality of the communication depends on the temporary status of incoming traffic and availability of network resources, which are impossible to predict beforehand. As a consequence, applications do not expect any guarantees and the network does not promise to deliver packets according to application requirements', if ever. Emergence of new types of applications that require Quality of Service (QoS), led the Internet Engineering Task Force (IETF) to considering a number of extensions to the current Internet service model. The Integrated Service (IntServ) architecture set the target of supporting end-to-end per flow guarantees, however it faces scalability challenges that are until today unsolved. A compromise approach, adopted by the IETF, is the Differentiated Service (DiffServ) architecture, whose aim is to deliver QoS to classes of service by deploying an appropriate arsenal of techniques, such as traffic shaping/policing, packet marking, bandwidth provisioning and so on. DiffServ provides however weaker forms of assurance than IntServ in order to achieve scalability.

The general idea of the DiffServ architecture is to classify and tag packets into a small number of classes at the edge of the network and to deploy mechanisms inside the network to treat various classes of packets differently. Figure 1 depicts a DiffServ-capable domain. IETF has standardized two types of forwarding behaviors or Per-Hop Behaviors (PHB) in the standard terminology: Expedited Forwarding (EF) [1], and Assured Forwarding (AF) [2], which are built to provide a premium service and an assured service, respectively [4]. Edge nodes of the domain, classify incoming traffic according to the service profiles in the Service Level Agreement (SLA), and tag packets with the appropriate differentiated service code point (DSCP). For each service class, the meter measures the flows and mark them as in-profile or out-profile packets. The traffic conditioner shapes the incoming flows by either dropping the out-profile packets, or forwarding them with certain probability. Core nodes only check the DSCP and forward packets according to the corresponding PHBs. By performing the many-to-one mapping from individual flows to aggregations, scalability is achieved.

*Brenda Zhuang is now with Boston University

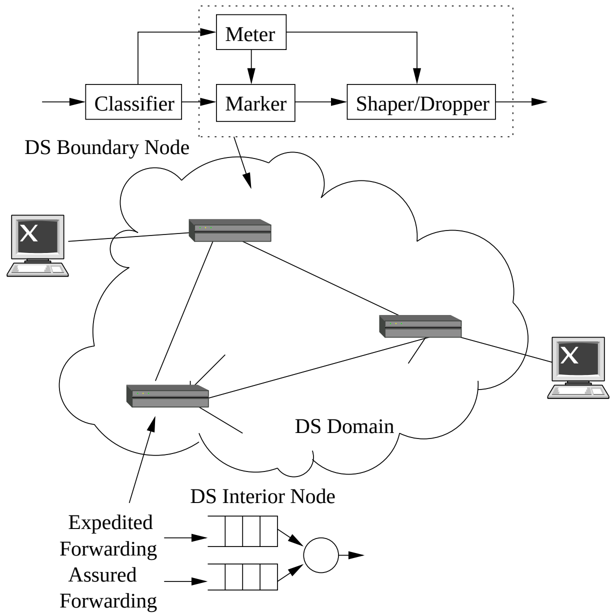

Figure 1: Differentiated Service Domain.

The performance of DiffServ architecture has been studied in various models. In [5], an analytical model of both EF and AF, assuming Poisson traffic, is given. Expressions of the performance measures that characterize the services, such as loss probability, delay and queue length distribution are then obtained. The Premium service is simply modelled as an $M/M/1/K$ queueing system. The RED In/Out (RIO) scheme is considered for AF with a Poisson arrival process. Validation of this simple model is carried out through simulation to argue in favor of the simple Poisson traffic assumptions. In [6], loss and delay behaviors of EF/AF are studied using a combination of router mechanisms, such as threshold dropping and priority scheduling, and packet marking schemes: for example, edge-discarding and edge-marking. The performance of the system is considered under both Poisson arrivals and Markov modulated bursty traffic. With a constant server capacity, it is shown that both implementations result in equivalent performance of packet loss probability. In [9], the model aims to improve TCP throughput with Assured service, and provide understanding of TCP end-to-end behavior in DiffServ network.

While most of the work contributes to the better understanding of traffic performance and behavior in DiffServ networks, we notice that the models investigate different schemes proposed to support PHBs by studying isolated Premium service or Assured service. Interaction between the two classes is seldom taken into account realistically. Our work intends to build on and add to previous contributions to help understand the interactions between classes of service. This is of paramount importance when designing bandwidth brokers that use the bandwidth efficiently. In this paper, we address the analysis of an interior node that supports Premium service and multiple classes of Assured service. The traffic is modelled as Markov Modulated Fluid Sources. Due to the complexity of the general multiple-queue scheduling case, we focus on a model where one EF queue and two-classes AF queues share the link capacity. In this model we assume that out of profile traffic is dropped at the edge of the network thus all AF traffic is in profile. Note that IETF recommends 4 AF classes with up to 3 dropping precedences each, however, most implementations consider one to two AF classes with two drop precedence. We adopt a decomposition technique as used in [10] and the approximation using spectral analysis developed in [11] and [12] in order to study the tail distribution of the AF classes queues under a combined priority and Generalized Processor Sharing (GPS) service hierarchy. We are able to derive an upper and a lower bound on the tail distribution of individual queues in the system. Note that, while AF traffic requires throughput differentiation, which is achieved by the GPS scheduler, the tail distribution helps quantify the delay and/or loss experienced by a given class and therefore enables the design of intelligent bandwidth brokers without relying on over-provisionning.

The rest of this paper is organized as follows. In Section 2, we propose and describe the system model including some elements of proposed mechanisms to support differentiated services. Analysis of the bounds and discussion of important observations is carried out in Section 3. In Section 4, some simple numerical investigations are carried out both analytically and by simulation and results are presented to show the tightness of the bounds. Finally, concluding remarks and discussion of future directions are given in Section 5.

2 System Model

The system model contains two key components of the DiffServ architecture, boundary node and core node. Since the architecture aims to push the complexity to the boundary nodes, and forward packets according to simple rules at the core nodes inside the domain, the nodes have different functions and structure. At the boundary node, the incoming flows are admitted according to pre-established agreements. For different traffic classes, the conditioner measures the statistical characteristics of the flows and mark them according to the SLA profile. The traffic is then shaped with the leaky bucket. In the core node, all packets would be forwarded according to the appropriate PHBs. The core node can be viewed as a queueing system with multiple service disciplines. A high priority queue served with strict priority drains the premium traffic at a guaranteed rate R_{Hmax}, which is designed not to exceed the link capacity. The lower priority queues, designed to provide AF services, share the remaining bandwidth according to the GPS scheduling policy to implement proportional differentiation as shown in the Figure 2

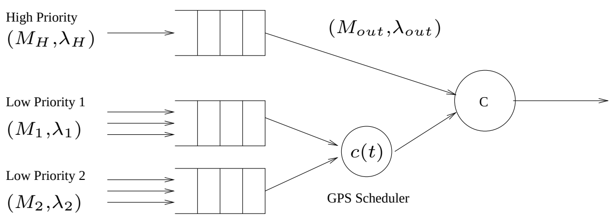

Figure 2: System Model Abstraction

The arrival process, such as the Internet traffic, at boundary nodes is known to be very bursty. After applying the leaky bucket algorithm as a shaper, the output traffic from the boundary node is of phase type [7]. For mathematical tractability, the input traffic to the core node is approximated by a superposition of Markov modulated fluid sources, instead of a superposition of output processes from the leaky bucket. At the core node, in order to provide low-loss, low-delay and low-jitter services as specified by IETF, EF traffic must comply with a very strict traffic profile so that the EF queue is always empty. The total link capacity is assumed to be constant C.

In this first model, the input traffic is assumed to be a superposition of Markov modulated fluid sources, each of which is characterized by its irreducible generator M_i, $i = H, 1, 2$ on state space $S_i = \{1, 2, \cdots, N_i\}$. The input rate vector is denoted by $\lambda_i = \{\lambda_i^1, \cdots, \lambda_i^{N_i}\}$. Without loss of generality, we assume that the AF queues share the available bandwidth with a rate proportional GPS fair share assignment of ϕ_1 and ϕ_2, which satisfy $\phi_1 + \phi_2 = 1$.

Let $\overline{\lambda}_i$ denote the average input rate of session i sources. Since the servers are work-conserving, the stability condition for this system is ensured by $\overline{\lambda}_1 + \overline{\lambda}_2 + \overline{\lambda}_H < C$.

For $i = H, 1, 2$, let $r_i(t) \in \{\lambda_i^1, \cdots, \lambda_i^{N_i}\}$ denote the arrival rate of session i at time t. For any time interval $[\tau, t]$, the arrival curve is denoted by $A_i(\tau, t) = \int_\tau^t r_i(u)du$. Similarly, the service rate at time t is $s_i(t)$, and denote session i's service curve by $S_i(\tau, t) = \int_\tau^t s_i(u)du$. The instantaneous backlog of session i at time t is expressed by

$$Q_i(t) = \sup_{\tau \le t}\{A_i(\tau, t) - S_i(\tau, t)\} \tag{1}$$

Since the high priority traffic always abides by the traffic contract profile set up in advance, the arrival rate will never exceed R_{Hmax}. As a consequence the input process and the output process are the same i.e., (M_H, λ_H) and (M_{out}, λ_{out}) are the same. For the low priority queues, the service rate at time t is influenced by the instantaneous service rate of the high priority traffic, thus, $c(t) = C - r_H(t)$. According to GPS scheduling, each session i is guaranteed an instantaneous service rate $g_i(t) \ge \phi_i c(t)$, for $i = 1, 2$. The average guaranteed service rate is denoted by $\overline{g}_i$.

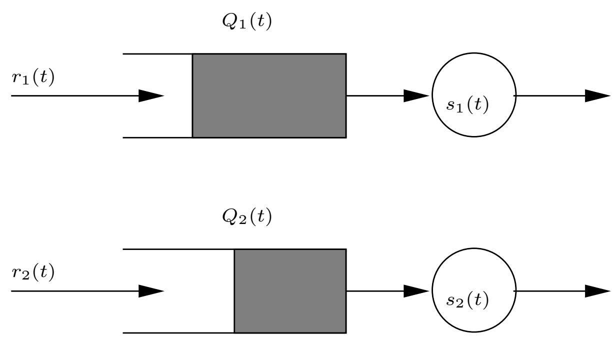

Figure 3: Decomposition of Two Queue GPS System

In order to bound the tail of the two queues in the GPS system, we adopt a decomposition approach. The GPS system can be decomposed into two queues as in Figure 3. The service rates of the decomposed queues, $s_1(t)$ and $s_2(t)$, depend on each other through the following relation:

$$s_1(t) = \begin{cases} c(t) - s_2(t) & if \quad Q_1(t) > 0 \\ r_1(t) & if \quad Q_1(t) = 0 \end{cases} \tag{2}$$

$$s_2(t) = \begin{cases} c(t) - s_1(t) & if \quad Q_2(t) > 0 \\ r_2(t) & if \quad Q_2(t) = 0 \end{cases} \tag{3}$$

For session 1, the sample path equation of the queueing process can be written

$$\frac{dQ_1}{dt} = r_1(t) - [c(t) - s_2(t)], if \;\; Q_1(t) > 0 \tag{4}$$

Due to the symmetry of the system, a similar result can be obtained for queue 2 by reversing the roles of queue 1 and queue 2. We will thus focus only on session 1 in the following analysis.

In GPS, when both queues are non empty, the service rate of session 1 is $s_1(t) = g_1(t) = \phi_1 c(t)$. When there is one queue that is empty, the other queue is served with the total available capacity. Thus, for $Q_1(t) > 0$ we have

$$\frac{dQ_1}{dt} = r_1(t) - [\phi_1 c(t) + (\phi_2 c(t) - r_2(t))\mathbf{1}_{\{Q_2(t)=0\}}] \tag{5}$$

where $\mathbf{1}_{\{.\}}$ is the indicator function.

In this paper, without loss of generality, we model the high priority traffic as a single Markov modulated on/off source. It is easy to extend the model to the case where the process is a superposition of a set of such on/off sources. It is characterised by the generator M_H and rate vector λ_H,

$$M_H = \begin{pmatrix} -\alpha_H & \alpha_H \\ \beta_H & -\beta_H \end{pmatrix}; \quad \lambda_H = (0, r_{H\max}),$$

where $r_{H\max} \leq R_{H\max}$. The available bandwidth for the low priority queues $c(t) = C - r_H(t)$ is thus also a Markov modulated process. The traffic of session i ($i = 1, 2$) is modelled as the aggregation of K_i on/off sources. Multiple sources can be lumped into one single aggregated Markov modulated source as studied in [11].

3 Analysis

We can observe from the model that the low priority queues depend on each other. Thus, the analysis of their exact behaviour is complex. However, we can study the decomposed system to derive statistical bounds on the queues distributions.

3.1 Lower Bound

To obtain a lower bound on the queue backlog, for example, of session 1, we need to find the instantaneous backlog $\widetilde{Q}_1(t)$ so that $\widetilde{Q}_1(t) \leq Q_1(t)$ is satisfied for all t. Since the decomposed queue and original queue have the same arrival process, a service rate process $\widetilde{s}_1(t)$ that is independent of session 2's queue should be chosen to give $\widetilde{Q}_1(t)$ as given in (4).

Let $[t_1, t_2]$ be the busy period of session 1, the service curve during the interval is denoted by $\widetilde{S}_1(t_1, t_2) = \int_{t_1}^{t_2} \widetilde{s}_1(v)dv = \int_{t_1}^{t_2}(c(u) - \widetilde{r}_2(u))du$, where $c(u)$ is the instantaneous available service rate, and $\widetilde{r}_2(t)$ is the chosen arrival rate of session 2 that is compatible with the choice of $\widetilde{s}_1(t)$, i.e., to achieve the lower bound of session 1. Since the modulating process from the higher priority queue always imposes the same effect on both queues, the lower priority queues are served with a bandwidth that is smaller than the link total capacity. The lower bound scenario and its equivalent system are shown in Figure 4.

We consider a session 2 without buffer and all excess traffic is dropped. By choosing $\widetilde{s}_1(t) = \phi_1 c(t) + (\phi_2 c(t) - r_2(t))^+ = c(t) - \widetilde{r}_2(t)$, where $\widetilde{r}_2(t) = \min\{r_2(t), \phi_2 c(t)\}$, we have the one-queue system for session 1 with a modulated service process $\widetilde{s}_1(t)$, resulting in $\widetilde{Q}_1(t) \leq Q_1(t)$.Thus we have $s_2(t) \leq \widetilde{r}_2(t)$.

Let $[t_1, t_2]$ a time interval contained in the session's busy period in the decomposed system. Then, $\widetilde{S}_1(t_1, t_2) \leq$

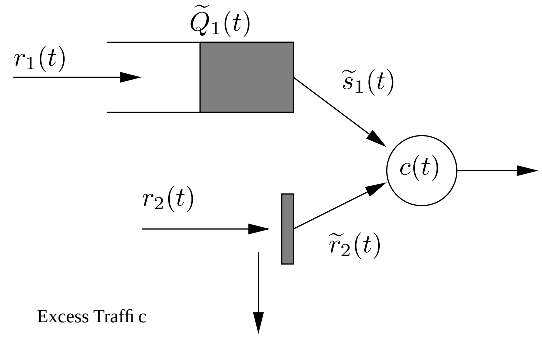

(a) Within the AF sub-system

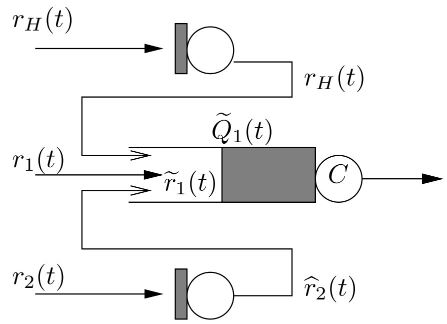

(b) Within the full system

Figure 4: Low Bound of Session 1 $\widetilde{Q}_1(t)$

$S_1(t_1, t_2)$ as shown in the following.

$$\begin{aligned} \widetilde{S}_1(t_1, t_2) &= c(t_1, t_2) - \int_{t_1}^{t_2} \widetilde{r}_2(u)du \\ &\leq c(t_1, t_2) - \int_{t_1}^{t_2} s_2(u)du \\ &\leq c(t_1, t_2) - S_2(t_1, t_2) \\ &\leq S_1(t_1, t_2) \end{aligned} \tag{6}$$

According to (4), $\widetilde{Q}_1(t) \leq Q_1(t)$ holds with $\widetilde{r}_2(t) = \min\{r_2(t), \phi_2 c(t)\}$.

The dynamics of the queueing process is thus described as follows.

$$\begin{aligned} \frac{d\widetilde{Q}_1(t)}{dt} &= r_1(t) - [c(t) - \widetilde{r}_2(t)], \ if \ \widetilde{Q}_1(t) > 0 \\ &= [r_1(t) + \widetilde{r}_2(t) + r_H(t)] - C \end{aligned} \tag{7}$$

Since $\widetilde{r}_2(t) = \min\{r_2(t), \phi_2 c(t)\}$, session 2 service rate $\widetilde{r}_2(t)$ and the high priority arrival process $r_H(t)$ are correlated. Thus, their superposition cannot be expressed simply. However, we can construct another process $\widehat{r}_2(t)$, which is Markovian, and is independent of $r_H(t)$ and satisfies $\widehat{Q}_1(t) \leq \widetilde{Q}_1(t) \leq Q_1(t)$ for all t. Let $\widehat{r}_2(t) = \min\{r_2(t), \phi_2 c_{\min}\}$, and $c_{\min} = \min_t\{c(t)\}$. It is obvious that $\widetilde{r}_2(t) \leq \widehat{r}_2(t)$. In this case, $c_{\min} = C - r_{Hmax}$.

The rate process $\widehat{r}_2(t)$ is characterised as the output process of session 2 queue. It is approximated by a Markov modulated fluid source, $(\widehat{M}_2, \widehat{\lambda}_2)$, including underload states and one busy state, $S = S_U \cup s_b$. The construction of this approximation process can be found in [12] in details.

The equivalent process is Markovian and the superposition of three sources is mathematically tractable. Accordingly, when $r_2(t) > \phi_2 c_{min}$, session 2 queue is in an overload state where the sojourn time is approximately Markovian. Since the buffer is zero, the excess traffic is ignored. When $r_2(t) \leq \phi_2 c_{min}$, the queue is in an underload state and the rate dynamics are finely modelled. Thus, we can construct the generator and rate vector as in [12].

The equivalent input process is characterised by the Kronecker product of the individual sources $r_1(t)$, $\widehat{r}_2(t)$ and $r_H(t)$. Though $\widehat{r}_2(t)$ is considered only to simplify the numerical investigation and is expected to result in a looser bound than $\widetilde{r}_2(t)$, numerical results show that the lower bounds are still relatively tight. In this equivalent system, the queue tail distribution of session 1 can be solved using the spectral method as described in [13].

3.2 Upper Bound

To understand the upper bound of the GPS scheduling, we consider the system in two possible cases.

1. Unbiased system:

$$\overline{\lambda}_1 < \overline{g}_1 \; and \; \overline{\lambda}_2 < \overline{g}_2 \tag{8}$$

2. Biased system:

$$\overline{\lambda}_1 < \overline{g}_1 \; and \; \overline{\lambda}_2 > \overline{g}_2 \; or \; \overline{\lambda}_1 > \overline{g}_1 \; and \; \overline{\lambda}_2 < \overline{g}_2 \tag{9}$$

First inspecting the unbiased case, we consider the upper bound of session 1. Since the average input rate of session 2 is less than its average guaranteed rate, the upper bound will be achieved by assuming session 1 obtains the "worst possible service". That is, $\widetilde{s}_1(t)$ is chosen to make $\widetilde{Q}_1(t) \geq Q_1(t)$ for all t. In this case, we study the situation when session 1 activity always encounters the busy period of session 2, and thus it only obtains the strict guaranteed minimum service. At time t, the service rate of session 1, $\widetilde{s}_1(t)$ shares the available capacity $c(t)$ with the service process of session 2, $d_2(t)$. The equivalent system is shown in Figure 5.

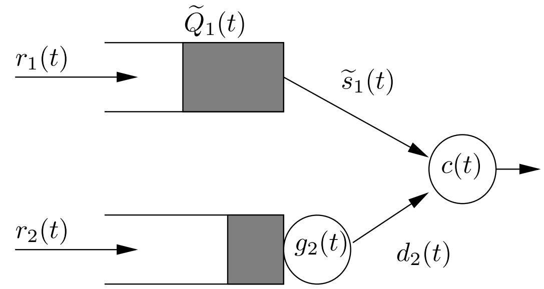

(a) Within the AF sub-system

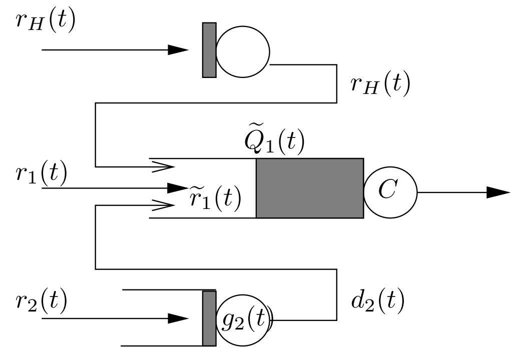

(b) Within the full system

Figure 5: Upper Bound of Session 1 $\widetilde{Q}_1(t)$ in Unbiased System

We start by characterising the departure process of session 2 in this situation. The arrival process of session 2 is a Markov modulated process, and the service rate is another Markov modulated process. At time t, the arrival rate $r_2(t)$ and service rate $g_2(t)$ act as a producer and consumer system coupled by a queue. From observation of the departure sample path, the departure process is characterised by both producer and consumer processes. However, since now session 2 has a buffer, the output process $d_2(t)$ is no longer Markovian. Thus, for mathematical tractability, following the approximations in [12], we lump all the activity of session 2 queue when the buffer is non-empty (i.e., the queue's backlog time) into one single state, while when the queue is empty, the system transitions are accurately modelled. The departure process is thus approximated by a Markov process according to the method developed in [12].

In order to upper bound $Q_1(t)$, the equivalent arrival process to session 1 $\widetilde{r}_1(t)$ is the sum of three processes, $r_1(t)$, $d_2(t)$ and $r_H(t)$. The dynamics of session 1 queue bound $\widetilde{Q}_1(t)$ is given by

$$\frac{d\widetilde{Q}_1(t)}{dt} = r_1(t) - [c(t) - d_2(t)],\ if\ \widetilde{Q}_1(t) > 0$$
$$= [r_1(t) + d_2(t) + r_H(t)] - C \tag{10}$$

In the biased system, there is always one traffic source with smaller average rate than the average guaranteed backlog clearing rate, while the other one is larger than the assigned average service rate. Without loss of generality, we take the case for $\overline{\lambda}_1 > \overline{g}_1$ and $\overline{\lambda}_2 < \overline{g}_2$.

The upper bound of session 1 would be the same as the previous unbiased case, because the average rate of session 2 is smaller than the average guaranteed rate. At time t, when the arrival rate $r_2(t)$ is higher than the guaranteed backlog clearing rate, the departure rate is limited to the available capacity, $g_2(t)$. At the same time, the buffer will build up. The server busy period prolongs for some time until the buffer is emptied with the guaranteed rate.

On the other hand, the upper bound of session 2 is straightforward. The average rate of session 1 is larger than the average assigned rate. In the "worst possible service" case of session 2, session 1 would always achieve its assigned service rate in such case. Thus, we consider the system, where the service rate is $\widetilde{s}_2(t) = c(t) - \phi_1 c(t) = \phi_2 c(t)$. The dynamics of the queue evolution is expressed as

$$\frac{d\widetilde{Q}_1(t)}{dt} = r_1(t) - \phi_2[C - r_H(t)],\ if\ \widetilde{Q}_1(t) > 0$$
$$= [r_1(t) + \phi_2 r_H(t)] - \phi_2 C \tag{11}$$

It is equivalent to the system where the arrival process is the superposition of $r_1(t)$ and $\phi_2 r_H(t)$, and the service rate is constant, $\phi_2 C$. In all cases, the queue tail distribution can therefore be solved using the techniques in [11, 13].

4 Numerical example and Validation

In this section, we present some numerical results obtained from the previous analysis. Simulations are also carried out using Network Simulator, version 2 (NS-2), to validate the analysis.

4.1 Settings

The traffic parameters of the first case study are given in Table 1. The high priority (EF) traffic is generated by a single on/off source, while the two-queue GPS input traffics are modelled each as a superposition of K_i on/off sources. For $i = 1, 2$, the source sends fluid at rate λ_i when it is on. The average on duration is $1/\alpha_i$ and the average off period is $1/\beta_i$. The channel capacity is normalised as in [11] and assumed to be $C = 2.37$.

In NS-2 we implemented in the DiffServ module the composite scheduler that combines the strict priority for EF traffic and we inherited the Weighted Interleaved Round Robin (WIRR) implementation that exists in NS-2 as a packet by packet approximation of the the GPS scheduler for the AF traffic. In the simulation, a policy is established between a source and a destination nodes. All flows matching that source-destination pair are treated as a single traffic aggregate. Each traffic aggregate has an associated policer, meter type and initial code point.

The topology of the simulation is shown in Figure 6. The constant capacity of the core node is set in different scenarios such that buffer overflow is observable at the queues.

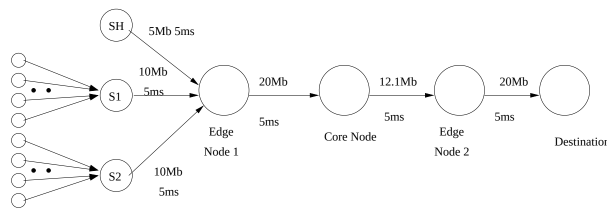

Figure 6: Simulation Topology in NS-2

As pointed out by many researchers, solving the models above exactly proves to lead sometimes to some singularity problems due to the roundoff and underflow errors. As such, in the numerical solution, we only approximate the tail distribution by the largest eigenvalue approximation $Q_i(x) \sim A_0 e^{z_0 x}$ where A_0 is the overload probability in a bufferless system. This has been repeatedly shown to constitute a very reasonable approximation.

4.2 Results

Figure 7 and 8 show the bounds for a GPS assignment of $(\phi_1, \phi_2) = (0.5, 0.5)$ respectively and the simulation scenario. The fair share assignments represent the unbiased case when considering the upper bounds.

Table 1: Traffic Characteristics in Simulation Case 1.

	$\frac{1}{\alpha_i}$	$\frac{1}{\beta_i}$	λ_i	K_i
Higher Priority	$10ms$	$90ms$	$1Mb$	1
Session 1	$10ms$	$90ms$	$1Mb$	9
Session 2	$10ms$	$90ms$	$1Mb$	9

The bounds are relatively tight in Case 1 as they both fit within the same order of magnitude. Simulations show that the bounds can provide a good approximation for the actual dropping rate at the network node.

In the study of core node performance on different assignment of buffer size, we have focused on how service differentiation is supported in the proposed scheme. Figure

9 and Figure 10 show the bounds when the GPS assignment is $(\phi_1, \phi_2) = (0.3, 0.7)$, and $(0.2, 0.8)$ and where the traffic is described in Table 2.

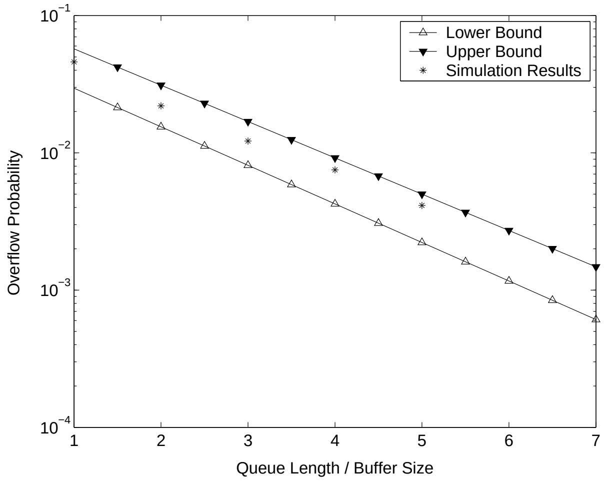

Figure 7: Bounds on Queue Length Distribution of $(\phi_1, \phi_2) = (0.5, 0.5)$ with Simulation Results

Table 2: Parameters for the System Case 2.

	α_i	β_i	λ_i	K_i
Higher Priority	0.5	1	1	1
Lower Priority Session 1	0.4	1	1	10
Lower Priority Session 2	0.4	1	1	10

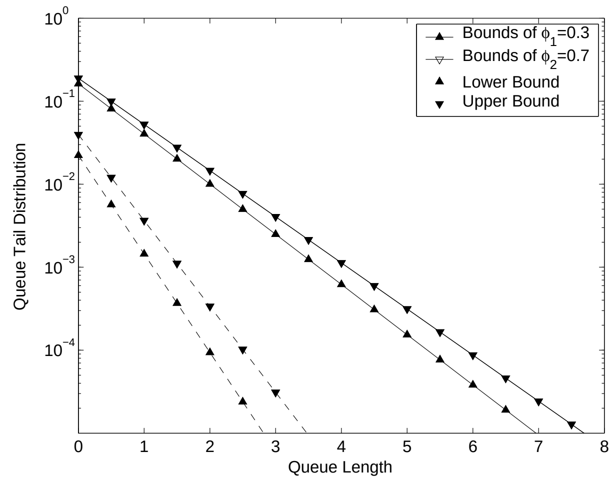

Figure 9: Bounds on Queue Length Distribution of $(\phi_1, \phi_2) = (0.3, 0.7)$

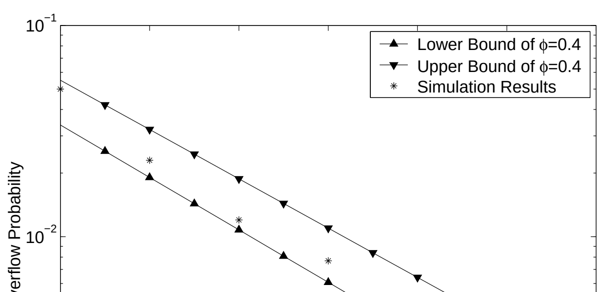

Figure 8: Bounds on Queue Length Distribution of $(\phi_1, \phi_2) = (0.4, 0.6)$ with Simulation Results

In addition to the throughput differentiation guaranteed by the GPS scheduler, we observe that with different assignment of fair share to AF classes, we can actually achieve a quite clear differentiation between classes in terms of overflow (delay). However, such delay differentiation depends on the available link capacity (out of a total of $10.1Mbps$ is this scenario) for the low priority classes.

5 Conclusion

In this paper, a model of a DiffServ interior node providing different PHBs is given. The node serves EF class with a strict priority queueing and the AF traffic shares the rest of the link bandwidth according to a generalised processor sharing scheduling scheme. For mathematical tractability we assumed Markov modulated fluid arrival traffic. From this model, we were able to obtain upper and lower bounds on the tail distributions of the AF queues based on spectral analysis techniques and the decomposition approach. Numerical results show the tightness of

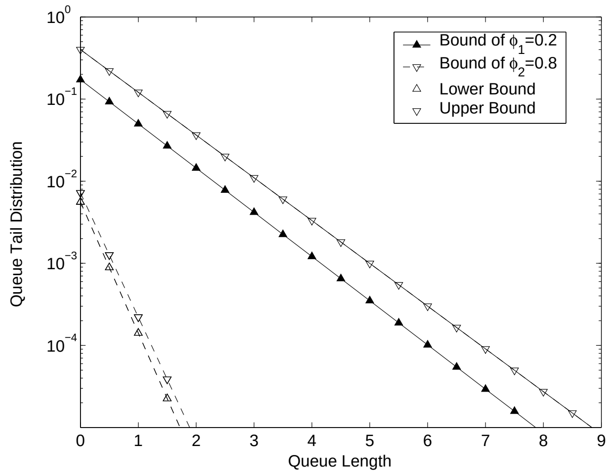

Figure 10: Bounds on Queue Length Distribution of $(\phi_1, \phi_2) = (0.2, 0.8)$

the bounds. These results can be very useful in many practical design issues of DiffServ. To cite only one areas, the design of efficient bandwidth brokers that perform dynamic admission of traffic between adjacent DiffServ domains would benefit greatly from such model, especially if closed form approximations are derived from this model.

References

[1] V. Jacobson, K. Nichols, and K. Poduri, *"An expedited forwarding PHB"*, RFC 2598, 1999.

[2] J. Heinanen, F.Baker, W, Weiss, and J. Wroclawski, *"Assured forwarding PHB group"*, RFC 2597, 1999.

[3] S. Blake, D. Black, M. Carlson, E. Davis, Z. Wang and W. Weiss, *"An architecture for differentiated services"*, RFC 2475, 1998.

[4] K. Nicholes, V. Jacobson, and L. Zhang, *"A Two-Bit Differentiated Services Architecture for the Internet"*, RFC 2638, 1999.

[5] M. May, J. C. Bolot, A. Jean-Marie, and C. Diot, *"Simple Performance Models of Differentiated Services Schemes for the Internet"*, IEEE Infocom, 1999.

[6] S. Sahu, D. Towsley, and J. Kurose, *"A Quantitative Study of Differentiated Services for the Internet"*, IEEE Global Internet, 1999.

[7] M. Xie *"Stochastic Performance Estimation in A GPS Server"*, Thesis, 2001.

[8] A. Elwalid, and D. Mitra *"Fluid Models for the Analysis and Design of Statistical Multiplexing with Loss Priorities on Multiple Classes Of Bursty Traffics"*, IEEE Infocom, 1992.

[9] I. Yeom, and A. L. N. Reddy, *"Modeling TCP Behavior in a Differentiated Service Network"*, IEEE/ACM transaction on networking,2001 vol. 9, pp. 31-46

[10] F. Lo-Presti, Z. L. Zhang, and D. Towsley, *"Bounds, Approximations and Applications for a Two-Queue GPS System"*, Technical Report 95-109,Department of Computer Science, University of Massachusetts December 1995

[11] D. Anick, D. Mitra, and M. M. Sondhi, *"Stochastic Theory of a Data-Handling System with Multiple Sources"*, The bell system technical journal 61 pp. 1871-1894 (1982)

[12] A. Elwalid, and D. Mitra, *"Analysis, Approximations and Admission Control of a Multi-Service Multiplexing System with Priorities"*, IEEE Infocom, pp. 463-472 (1995).

[13] D. Mitra, *"Stochastic Theory of a Fluid Model of Producers and Consumers Coupled by a Buffer"*, Adv. Appl. Prob., vol. 20, pp. 646-676 (1988).

[14] A. K. Parekh, and R. G. Gallager *"A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single Node Case"*, IEEE/ACM transaction on networking,1993 vol. 1, pp. 344-357