Performance Study of End-to-End Resource Management in ATM Geostationary Satellite Networks with On-Board Processing

Güray Açar [1] and **Catherine Rosenberg**[2*]

[1]Imperial College,
Dept of Electronics & Electrical Eng.
London, UK

[2]Purdue University,
School of Electrical & Computer Eng.
West-Lafayette, USA

Abstract

Because of their large geographic coverage, inherent broadcast capabilities, and fast deployment features, network operators plan to use satellite-based networks to supplement existing wire-line and legacy networks to bring broadband and multimedia services to end-users. Satellites are multiple access systems with very limited transmission capacity compared to terrestrial network nodes. Therefore, end-to-end resource management for such systems is key to deliver acceptable Quality of Service (QoS) to services while providing adequate efficiency.

We study in some details, using simulations, a method integrating Bandwidth on Demand with Call Admission Control in geostationary satellite networks with an on board switch. In particular, we describe a set of algorithms and study the impact of different types of bandwidth reservation on the QoS received by the connections and the network efficiency.

---

[*]Corresponding author: School of ECE, West Lafayette, IN 47907-1285, USA. Email: cath@ecn.purdue.edu. Phone: +1-765-494-0034, Fax: +1-765-494-0880.

1. Introduction

Already very successful in broadcasting entertainment services, digital satellite systems are viewed as viable service vehicles. The demand for Internet and multimedia services and the subsequent need for higher bandwidth drive network operators to seek mechanisms to cost-effectively provide broadband access. As a result, network operators plan to use satellite-based networks to supplement existing wire-line and legacy networks to bring broadband and multimedia services to end-users ([2], [4], [5], and [6]).

Satellites have been part of the Internet since its early days, providing Internet backbone connectivity to Internet Service Providers (ISPs), intranet solutions for geographically dispersed corporations and, more recently, direct-to-user Internet access.

Many proposals have been made to national and international regulatory agencies for allocation of spectrum for broadband applications using low earth orbit (LEO) satellites, medium earth orbit (MEO) satellites, and geo-synchronous (GEO) satellites. In this paper we concentrate on GEO based satellite networks. GEO based systems deliver continuous services to a specific region with a single satellite. They play an ever-increasing role in the public and private internets, due mostly to their large geographic coverage, inherent broadcast capabilities and fast deployment. They are attractive to support data, audio and video streaming; bulk data transfer such as software update or dissemination of Web caches; and applications involving limited interactivity such as distance learning. They are also attractive to provide broadband access to users who are either beyond the reach of the terrestrial network, or have particular needs for broadcast/multicast applications or fast deployment.

Several journals have devoted special issues on broadband satellite networks. See [4] and [5] for good surveys on the systems, issues and solutions.

Different commercial systems using Bent Pipe (BP) satellites currently exist. They are best suited to one of the following services:

- Backbone connectivity or trunking (gateway to gateway),
- Meshed VSAT (Very Small Aperture Terminal) networks for LAN interconnection (terminal to terminal),
- Internet access (terminal to gateway).

We define a Next Generation Satellite Network (NGSN) as a system based on a GEO multi-beam satellite that operates typically in the Ka-Band and comprises an on-board processor (OBP). More precisely a NGSN (see Figure 1) will be characterized by the support of:

- **Two-way capability**: user terminals are able to send, at rates ranging from a few hundred of kbits/s up to a few Mbits/s and to receive at rates of several tens of Mbits/s.
- **Spot beams technology**: instead of covering the whole footprint of a satellite by a global beam, this footprint is divided into a number of spot beams. The benefits of spot beams are twofold: a) the power requirements of user terminals are reduced. This permits the use of smaller antennas in the ground segment and hence reduces the terminal cost. b) frequency can be re-used between beams, which increases the capacity of the space segment.

- **On-Board Processing** (OBP) capabilities, whereby the satellite does not act as a repeater, but includes switching functions that increases the connectivity in the sky by directing incoming data onto the appropriate beam(s).

A typical NGSN comprises several tens of spot beams and the total capacity of the system is in the range of several Gigabits/s. On-board processing satellites with high-gain multiple spot-beams and on-board switching (OBS) capabilities have been considered as key elements of the next generation of satellite systems. The OBP comprises at least the traffic plane of a switch while the control plane can be partly in the satellite and partly ground based. The users will be offered high-speed bi-directional digital links through low cost and small size terminals. Direct connections to other users of the NGSN are provided as well as enhanced access to terrestrial networks through access nodes named Gateways. The access is enhanced because thanks to the connectivity in the sky, a user can directly reach the gateway closer to its destination, hence bypassing a large portion of a potentially clogged network. This system will provide bandwidth-on-demand to users. Although employing an OBS results in more complexity on board the satellite and hence more risks, the following are the advantages of on-board switches:

- Improved interconnectivity.
- Mesh capability (i.e., one hop terminal to terminal communications, irrespective of the terminals' beams) with simple terminals.
- Bandwidth on demand with half the delay (see later).
- Efficient multicasting through the use of a multicast OBS.

So far 3 requirements have driven the choice of the transport technology in the space segment: 1) the need to inter-operate seamlessly with IP services, 2) the fact that the multiple access is MF-TDMA which calls for a fixed size packet (i.e., a cell based technique), and 3) the need to provide different levels of Quality of Service (QoS) to the users. These requirements have yielded most of the proposed systems to be ATM (Asynchronous Transfer Mode) based or at least ATM-like (a vague term saying that the designer is keeping only what he likes from the original ATM). ATM has been used extensively in the core network due to the availability of high speed, low cost switches with QoS capabilities and interfaces for routers.

In this paper, we consider a NGSN with an OBP ATM switch. In this system, any user of any beam can connect directly with any other user or gateway of any beam. Two types of access units are defined in the system: Satellite Access Units (SAUs) connect individual users or LANs to the network, and gateways (GWs) connect the satellite network to other networks. SAUs are connected to the satellite through MF-TDMA (Multi-Frequency Time Division Multiple Access) links on the uplink and TDM on the downlink. Gateways can either be connected to the satellite using very high-speed point-to-point bi-directional TDM link or MF-TDMA links depending on the traffic being envisaged. The Network Control Center (NCC) is in charge of most of the signaling and management functions in the satellite network.

Satellite networks are multiple access systems with limited transmission capacity compared to terrestrial networks. Therefore, end-to-end resource management for such systems is key to deliver acceptable Quality of Service (QoS) to users while providing adequate efficiency (i.e. a level of Grade of Service (GoS) that entails the use of such systems).

Bandwidth on Demand (BoD) is central to end-to-end resource management. It is defined as a set of MAC (Medium Access Control) protocols and algorithms that allow connections to request resources on a demand basis, while the connection is already in progress, in an environment where many bursty connections share a common medium access link. We designed a method integrating a BoD process with a Call Admission Control (CAC) scheme for an ATM GEO satellite network in [1], [8], and [9]. In this paper, we present further studies and results on the integrated method introduced in these papers. More precisely, we discuss performances in terms of delay and efficiency using a large-scale simulation program. We show how reservation and CAC affect the NGSN performances.

In a bent pipe system, terminals in a beam use BoD to request bandwidth on the uplink, which is the scarce resource to be managed efficiently and fairly by the BoD controller. In an OBP system, the problem is complicated by the fact that there may be more than one type of scarce resource to protect. Indeed, while the uplink resource is still scarce, the buffering in the satellite can also be very limited and hence there is a need to co-ordinate the sharing of the uplink resources in the different beams so that the traffic for a given output port in the OBS is not overflowing the corresponding buffer. Indeed, it would be inefficient to use a scarce uplink resource to send packets that will be discarded in the satellite because of lack of buffering. While some discarding can be acceptable, it is important to keep it under control. Obviously if the capacity of the downlink can be dimensioned slightly over the capacity of the uplink and if the buffering is of a size similar to what it would be in a terrestrial ATM switch, then there may be no need for co-ordination. Otherwise, the BoD controller could have to co-ordinate the sharing on all the uplinks and output ports, which would make the process much more complicated. Usually ATM switches in terrestrial networks have a significant amount of buffering to handle burst level congestion, to allow best-effort traffic to have a reasonable chance of going through and to offer an acceptable Grade of Service (i.e., call blocking) to VBR (Variable Bit Rate) connections. If the amount of buffer space in the OBS is large enough we can assume that the BoD process has only to perform the fair and efficient sharing on each uplink separately. The large buffers on board the satellite will allow statistical multiplexing and will absorb occasional burst level congestions on some overloaded downlink beams. In this paper, we assume enough buffer space on board the satellite and downlink capacities slightly larger than uplink capacities. Hence, the task of the BoD process is to share the available capacity on each uplink among the active SAUs of the beam as in a bent pipe system.

The BoD process consists of the following phases:

1) A first phase during which each SAU computes the resource requirements for individual ATM VCs (Virtual Connection) or for groups of ATM VCs (BoD Entities). In [1] we presented a set of **RRE (Resource Requirement Estimation)** algorithms to perform this first phase. In [3] we discussed how aggregation should be performed to maintain efficiency and fairness.

2) The second phase consists in signaling the resource requirements in the form of **<u>Resource Requests</u>** (RR) from the SAUs to the BoD controller. Since the satellite comprises an OBP, the controller could be satellite based.

3) The third phase is crucial. The BoD controller has to compute, for each beam, the fair and efficient allocation of time-slots (TS) to the VCs (or BoD entities). This results in the

creation of the **Burst Time Plan** (BTP). An algorithm to perform the fair and efficient sharing of resources is presented in [3], [7], [8], and [9].

4) The fourth phase consists in signaling the response from the BoD controller to the SAUs (broadcast of the BTP).

5) The last phase is performed by each SAU that has to allocate the TS received to its different connections. This is an internal scheduling phase.

Section-2 will briefly describe the BoD and CAC integrated process. Emphasis will be given to the resource allocation phase at the BoD controller, the resource requirement estimation phase at the SAUs, and the method introduced in [8] and [9] integrating the BoD process with the CAC. In section-3, we present a short description of our simulation program followed by the results of our simulation analysis. In our simulations, we analyzed the impacts of the BoD and CAC parameters on queuing delay, efficiency, and service segregation among different service classes. Finally, in section-4, we present our conclusions. Note that, although we will present the results of our studies in an ATM framework, they can easily be extended to an IP QoS framework.

2. BoD and CAC Integrated Process

In this section, we briefly describe the resource allocation scheme at the BoD controller, and the Resource Requirement Estimation (RRE) algorithms at the SAUs. We then elaborate on the integration of the BoD process and the CAC in a QoS framework.

Note that not all connections will use BoD. For QoS reasons, some connections will only rely on statically allocated resources obtained at call set-up. This is the case for example for real-time connections. Others will use BoD and will negotiate at call set-up, depending on their type, two parameters SR and BR that will be introduced later.

2.1 RRE Algorithms and BoD Controller

SAUs send RRs (Resource Request) periodically to the BoD controller. We assume in our studies a RR period that is equal to one MF-TDMA frame. Hence, SAUs can send RRs every MF-TDMA frame. The reply of the BoD controller to the signaled RRs arrives at the SAUs in the form of a BTP (Burst Time Plan). The time interval between RR signaling and the reception of the corresponding BTP is the response time. If the BoD controller is on board the geostationary satellite (OBP case), the response time will be at least 250 msec, because of 125-msec propagation delay between the SAU and the GEO satellite.

Each BoD entity (an individual ATM VC or a group of ATM VCs using BoD) is assigned a Static Resource (SR) and a Booked Resource (BR) by the satellite network at call set-up. These parameters can be zero depending on the type of the entity (i.e., VBR or UBR (Unspecified Bit Rate)). In the following, we will assume that the BoD process deals with individual VCs. The other case was partly dealt with in [8] and requires further study. SR is the amount of resource (a number of time slots per frame) that is statically allocated to the connection. On top of SR, each connection using the BoD will be allocated, in a frame k, an amount of resource corresponding to x_k TSs, based on its current request RR_k (generated by the SAU few frames ago, see later), on its BR, on the requests of the other connections, and the available capacity. BR is the amount of resource that is booked (i.e., reserved) for the connection at call set-up. The connection is guaranteed an amount of resource of at least BR on top of SR in each frame, if it requests for it. Equation 2-1 illustrates the relation among RR_j, BR_j, and x_j for an ATM VC_j, where the subscript k is omitted for simplicity sake.

$$\text{If } \begin{Bmatrix} RR_j \le BR_j \Rightarrow x_j = RR_j \\ RR_j > BR_j \Rightarrow x_j = BR_j + BE_j^* \end{Bmatrix} \qquad \text{Equation 2-1}$$

where BE_j^* is the best-effort share of the available capacity for the VC_j in the current frame k. The computation of BE_j^* is based on Game Theory [3], [4], and [8]. Very briefly, assuming that there are N BoD connections and M non-BoD connections in the current frame, the BE_j^*'s are chosen such that $\prod_{j=1}^{N} BE_j^*$ is maximized, subject to the following constraints being satisfied:

1. $BE_j^* \le \max\left(0, \left(RR_j - BR_j\right)\right) \qquad \forall j$

2. $\sum_{j=1}^{N} BE_j^* \le C_A$ where C_A, is the available best-effort capacity.

Note that the available best-effort capacity, C_A, is computed as $C_A = C - \sum_{j=1}^{(N+M)} SR_j - \sum_{j=1}^{N} \min\left(RR_j, BR_j\right)$, where C is the total uplink transmission capacity. Hence we have chosen to share the available capacity C_A fairly among <u>all</u> the BoD connections. We will discuss this in more details later.

Recall that we have assumed sufficient buffer space in the on-board switch and hence each beam is processed separately. Also, note that the BoD controller has to make sure that the total transmission rate allocated to connections within a SAU does not exceed the transmission rate of the terminal equipment.

The resource requirement estimation (RRE) phase, during which the resource requirement of a connection is computed, is an essential component to the BoD process. It is explained in the following.

Let i be the number of frames during one response time. Assuming the BoD controller is on board the satellite, and assuming a frame size of 64 msec, i will be 4 or 5 depending on the amount of processing time required during the resource allocation phase at the controller. If the frame size is 32 msec, i will be 8 or 9. If the RRE algorithm computes a resource requirement for a given BoD connection in frame k, then the corresponding BTP would arrive and be effective after a complete response time, i.e., in frame i+k. This is the target frame. In other words, the objective of the RRE algorithm is to estimate the amount of resource that the connection will require **during the target frame**, which is one response time (i.e. i *MF-TDMA frames)* after the moment the resource requirement is computed.

The RRE algorithm should neither overestimate nor underestimate the resource requirement of the connection. Note that it is impossible to know whether the RRs that are still in the fly (requests made during the frames k-1, k-2, …, k-i+1) will be fully granted by the BoD controller.

While designing our RRE algorithms, we have chosen to be conservative in that our algorithms aim to preclude any TS (Time Slot) waste. Therefore, they aim to compute the number of ATM cells that will **certainly** be ready to be transmitted in the *target frame*. Accordingly, there are two basic **assumptions** behind our RRE algorithms:

i. there will be no cell arrivals from the moment RRE is invoked till the end of the target frame,

ii. all past (i-1) RRs will be fully granted by the BoD controller.

The RRE algorithm must be such that the SAU and the BoD controller are somehow synchronized in terms of the request-reservation process. There is a need for a memory element that will remember those RRs that are not fully granted by the BoD controller and those resources that are not fully used by the SAU (possibly because of traffic policing at the SAU). We have developed two approaches to deal with this problem.

The first approach, which is called the BoD controller without memory (RRE–1), places all memory at the SAUs. In this approach, the BoD controller has no memory of the past RRs (i.e., it does not remember that it has not granted completely a past request to a connection). The SAUs remember those RRs that are only partially accepted by the BoD controller. The RRE algorithm measures the current queue length, subtracts from it the total number of ATM cells that are expected to depart from the SAU buffers until the end of the target frame. The result (if positive) constitutes the number of TSs that will be requested for the target frame. The algorithm can be expressed by the Equation 2-2.

$$RR(k) = \left[q(k) - N(k) - (i+1) \cdot SR - \sum_{r=1}^{i-1} RR(k-i+r) \right]^{+} \qquad \text{Equation 2-2}$$

where q(k) is the number of ATM cells of the connection waiting in the SAU buffer at the beginning of the k^{th} frame, N(k) is the number of TSs allocated by the BoD controller for the

k^{th} frame, i is the number of frames in one response time, and RR(k) is the resource request that was made in the k^{th} frame. Note that $[x]^+ = x$ if $x \geq 0$, and 0 otherwise.

If the BoD controller cannot fully grant a RR, the remaining part of the RR (that is not accepted) will remain in the buffer after the corresponding BTP has arrived. Hence, the queue length is the memory element that we required; the SAU will automatically re-request the part of the RR that was not granted by the BoD controller after reception of that BTP. Note that the SAU will have to wait for (at least) one more response time before the remaining part of the RR is granted.

The advantages of this approach are its scalability and reliability. The network can be expanded easily without the need to worry about the memory space at the BoD controller. If a request is lost, the SAU will know it after one response time and will be able to compensate for it. The disadvantage of this approach is that it does not react fast to a mismatch between what the connection has asked for and what the controller has allocated.

The second approach, which is called BoD controller with memory (RRE-2) places some part of the memory at the BoD controller. If the BoD controller cannot fully grant a RR_k, which was sent in the k^{th} frame, to a connection, it keeps in memory what it was not able to allocate, say t_k, and add it to the next request coming from this connection (i.e., RR_{k+1}=RR_{k+1}+t_k). If there are no requests in frame k+1, then it creates one of value t_k. In other words, if the BoD controller can only partially accept a RR, it remembers that it owes the remaining part and will try to allocate it in the next frame. The SAU, on the other hand, remembers also how much resource the BoD controller owes, and avoids re-requesting this amount of resource. In order to implement this approach we need two more variables per connection. The first one represents the memory of the BoD controller and is held at the BoD controller, while the second one is held at the SAU and is the vision that the SAU has of the BoD controller memory. These variables are necessary to deal with RR losses during transmission. They are:

$$t(k) = [t(k-1) + RR(k) - N(k+i)]^{+} \qquad k \geq 0 \qquad \text{where } t(-1) \equiv 0$$

There will be one t(k) variable at the BoD controller for each VC. It is the memory of the BoD controller that keeps account of those TSs that were requested but could not be allocated to the connection. Note that RR(k) is the kth RR arriving at the BoD controller, and N(k) is the number of TSs allocated for the kth frame. The response to RR(k) will be received by the SAU in the (k+i)th RR period, corresponding to N(k+i). Similarly we define a variable p(k) to be used at the SAU for each VC, which is defined below.

$$p(k) = [p(k-1) + RR(k-i) - N(k)]^{+} \quad k>i \text{ where } p(k)\equiv 0 \text{ for all } k\leq i.$$

Similarly p(k) is to keep track of those TSs that the BoD controller <u>owes</u> to the connection. As long as there is no RR losses the equation below must hold.

p(k)=t(k-i) k>i

Our second RRE algorithm is represented by the following equation.

$$RR(k) = \left[q(k) - N(k) - (i+1) \cdot SR - \sum_{r=1}^{i-1} RR(k-i+r) - p(k) \right]^{+} \qquad \text{Equation 2-3}$$

Note that RRE-2 does not re-request for partially accepted past RRs. It knows that the BoD controller will remember. Note that if a RR is not fully granted, it will take at least one frame time (as opposed to one response time, i.e., <u>i</u> frames for RRE-1) for the remaining part of the

RR to be allocated. Also, note that there is a need for a mechanism to re-establish the request-reservation synchronization between the SAU and the BoD controller when a RR is lost.

The second approach is advantageous because it increases the responsiveness of the BoD process but it puts more burden on the BoD controller, which could be a problem if the controller is on board the satellite, and makes the process more sensitive to RR losses.

2.2 BoD & CAC Integration and QoS

In ATM networks, Connection Admission Control (CAC) is a network process that receives as an input, a connection request that specifies the traffic descriptor and QoS requirements of the connection and returns a response granting or denying the admission request. The objective of the CAC is to limit the number of connections within the network so that each connection receives sufficient amount of network resources to meet its guaranteed QoS requirements. Roughly speaking, on a link with a homogeneous user profile (i.e., connections with identical traffic characteristics and QoS requirements) the CAC would limit the number of users to a number N, so that, $N \cdot E_r \leq r \cdot C$, where E_r is the effective rate allocated to each connection, ρ is the desired load, and C is the total link capacity.

In a real network, users have different traffic characteristics and different QoS requirements. Managing such networks and offering differentiated QoS to different traffic classes require some segregation among service classes and hence scheduling schemes more sophisticated than FIFO (First In, First Out) in the network nodes. What we call segregation is the ability for a scheduler to protect the QoS of each class from the behavior of the others. In general, we expect two things from a switch: a good segregation between service classes (e.g., between

CBR (Constant Bit Rate), VBR (Variable Bit Rate) and UBR (Unspecified Bit Rate) traffic classes)[1] and a acceptable segregation within a class between different connections. Different schedulers are being used. Some switches and routers use Head of the Line (HoL) priority. For instance, in a scenario where only VBR and UBR connections are supported, the scheduler could assign full priority to VBR connections in order to provide guaranteed QoS to these connections. In that case the VBR class is well protected. Another very popular scheduler is Weighted Fair Queueing (WFQ) in which the server is shared among the different classes (or connections) based on allocated weights.

Just like a scheduler in a terrestrial switch, the BoD process is the resource manager in the satellite network. It provides regulated access to satellite resources for a high number of very bursty and uncoordinated connections spread over a large area. Static Resource (SR) and Booked Resource (BR) are crucial, because they are the means by which the users of the network can be guaranteed QoS. Note that SR and BR are assigned per connection. We could also design a scheme where the SR and BR are negotiated for a group of connections within an SAU. This is not what we will assume in the following. Using SR and BR, the BoD process allows us to segregate not only among service classes but also among the connections

---

[1] There are several traffic classes defined in ATM. The ones that we considered are CBR, real-time VBR, nrt VBR, GFR (Guaranteed Frame Rate) and UBR. CBR and rt-VBR have stringent delay constraints and hence do not use BoD (there are only allocated SR at call set-up) while UBR is best-effort and is guaranteed nothing. GFR is best-effort with a minimum guaranteed rate and is not mentioned here to keep the discussion manageable. Finally, nrt-VBR is guaranteed a very low cell loss ratio and a bounded delay. Hence nrt-VBR, GFR and UBR would use BoD.

within the same service class. For instance, in a network with non real time (nrt) VBR and UBR BoD connections, we may allocate a non zero SR and BR for VBR connections while UBR connections would not be allocated any static or reserved resources. SR and BR will not only favor VBR connections against UBR connections by making sure that some of the uplink capacity is dedicated to the VBR connections, but also guarantee various levels of QoS to different applications within the VBR service class by reserving different amount of resources for each connection. Note that if a connection does not need its SR for a period of time, only other connections within the same SAU can use the corresponding TSs while for BR, any other connections in the beam can use the corresponding TSs.

The method for CAC and BoD integration, which is explained in [8] and [9] in detail, must ensure that a VC_j will be accepted into the network if the sum of what is to be statically allocated to *j* (i.e., SR_j) and what is to be booked for *j* (i.e., BR_j) is less than the total capacity of the multiple access link, minus the sum of the already allocated capacity and minus the sum of the booked capacity for all ongoing calls *k* on the uplink. This is expressed in the Equation 2-4 below.

$$\mathrm{SR}_j + BR_j + \sum_k SR_k + \sum_k BR_k \leq C, \qquad \text{Equation 2-4}$$

where C is the total capacity of the satellite uplink. For reasons to be explained in section-3.3, Equation 2-4 will be modified later. Hence, if a UBR connection is not allocated any SR and BR at call set-up, it is not subject to CAC.

In the case of a bent pipe system, if the uplink CAC can accept the new connection, the satellite system can accept it and the terrestrial CAC process (if any) proceeds. In the OBP case, the CAC in the satellite has to proceed. As long as the buffers in the satellite have a size similar to those in terrestrial ATM switches, a CAC similar to a terrestrial CAC could be implemented. However, the SAUs should implement policing to make sure that what goes on the uplink (which has been in some sort transformed by BoD) is compliant with the CAC in the satellite so as to avoid excessive losses on board the satellite.

It is clearly seen in the discussion so far that SR and BR are crucial parameters of the BoD process. SR and BR assigned to connections do not only determine the QoS guaranteed to the connections, but also, together with the total capacity of the system, determine the maximum number of nrt-VBR connections[2] that can be supported by the network. In other words, SR and BR assignment to the BoD connections is an issue that has great impact on both the QoS guaranteed to the connections and the GoS (Grade of Service) of the network. In this respect, the network operator aims to assign each connection an appropriate [SR, BR] pair so that it receives the required QoS, while the GoS of the network is optimized.

The choice of the right values for SR and BR for a connection is a complicated issue. The connections can request (and very probably will get), using BoD, much more than its BR thanks to statistical resource sharing among all the connections. On the other hand, SR is a fixed amount of resource that is allocated to the connection every frame. Therefore the same amount of resource, say x, may provide a different QoS to a connection depending on whether it is allocated in a static-only fashion (i.e. SR=x, BR=0) or a booked-only fashion (i.e. SR=0,

---

[2] In a system supporting GFR, GFR connections would be allocated SR=0 and BR≡MCR (minimum cell rate) and hence would be subject to CAC.

BR=x). SR and BR could also have different impacts on the GoS of the network, depending on how the BoD process is designed. Indeed, we have described in section-2 a BoD that treats all the VCs the same irrespective of their types when sharing the capacity available for best-effort, i.e, C_A. However, we could design other BoD processes that share C_A differently. Recall that C_A is the whole capacity available after distributing all the SRs and the portion of the BRs that was requested. We could, for example, design a scheme in which a portion α of C_A is allocated to the nrt-VBR traffic first and then the remaining of the available capacity should be offered to all the connections on a fair and efficient way basis. We could also try to emulate a WFQ scheduling by allocating a minimum "best-effort" bandwidth to the UBR connections and a minimum "best-effort" bandwidth to the nrt-VBR. In order to take full advantage of the statistical multiplexing effect of the BR to enable a better GoS, we could also imagine that the booked capacity that has not been requested in frame k, i.e., $B_A = \sum_{i=1}^{N} BR_i - \sum_{i=1}^{N} \min(RR_j, BR_j)$ is shared first among the nrt-VBR connections that requested more than their BR. We would then be able to take advantage of some statistical multiplexing effect, that would allow us to reserve less for a connection using BR only than for the same connection using SR only.

SR and BR can be envisaged as the tools provided to the network operator in order to fine-tune the trade-off between the QoS guaranteed to connections and the GoS of the network. In the next section, we will present the results of our simulations. We will present the impact of SR and BR on the service segregation in the satellite network, the delay characteristics experienced by the connections, and the amount of resources (i.e., Time Slots) wasted by the connections.

3. Simulation and Performance Analysis

3.1 Simulator Program

The simulation program that we developed is a time-driven simulation where each iteration corresponds to one millisecond of simulation time. It is written in C++ though we have avoided using an OOP (Object Oriented Programming) approach in some modules to improve the execution speed.

Only one uplink beam of the satellite is simulated. We assumed only one connection per SAU, because we aim to study the queuing delay and resource waste as a result of BoD operations only. In a SAU with more than one connections, queuing delay and resource waste will not only be influenced by the BoD process but also by the statistical multiplexing within the SAU buffer. Each connection is an ATM cell stream that is generated by an ON-OFF source with geometrically distributed ON and OFF periods, where cell emission rate during the ON periods is the PCR (Peak Cell Rate) of the connection, and the source is silent during OFF periods (i.e., 2-state Markov Modulated Deterministic Process). Note that the parameters of this cell arrival process are chosen so that each connection has a PCR=192 kbps and a SCR (Sustainable Cell Rate) that is set according to the purposes of the simulation. The mean burst length (i.e., mean ON period) is equal to 200 msec where the mean inter-burst time (i.e., mean OFF period) is varied to attain different SCRs. Also note that the Maximum Burst Size (MBS) of each connection is kept equal to 512 cells by means of a leaky bucket.

In our simulations, we assumed that we had 2 MF-TDMA carriers and a frame of 64 TS (Time Slots). Each TS is one msec long, and can carry one ATM cell. Accordingly the maximum useful transmission rate is 384 kbps, and one TS per frame corresponds to a useful transmission rate of 6kbps.

The BoD controller is assumed to be on board the satellite. One-way propagation delay between the SAUs and the satellite is assumed to be 125 msec. In our simulations we assumed that the processing time at the BoD controller, the transmission time for signaling the RRs and broadcasting the BTP will add up to a total of 6 msec. Under these assumptions the response time is 256 msec. Hence there are 4 frames within one response time. Note that the processing time at the BoD controller, and the time required to signal RRs from SAUs to the BoD controller are most likely to be longer than 6msec in a real satellite network, in fact probably of the order of a frame.

Recall that SR and BR are expressed in terms of number of TSs per frame. The total of SR and BR for a connection will also be expressed in terms of TSs per frame, and denoted as TTS (Total Time Slots). SCR and PCR of a connection will be expressed in terms of number of cells per frame (which is equivalent under our assumptions to a number of TSs per frame). The PCR of each connection in our simulations is 32 cells/frame (i.e. 192 kbps).

3.2 Service Segregation

The objective of this part of our study is to illustrate the influence of SR and BR on service segregation within the satellite network. We define two types of connections, type-1 and type-2. These two types are identical in the traffic characteristics of the cell streams that they offer into the network. Both types of connections offer a mean cell arrival rate of 4 cells/frame. The only difference is that type-1 connections are not assigned any SR or BR (i.e., corresponding to UBR connections) while type-2 connections are assigned BR only (i.e., corresponding to nrt-VBR or GFR connections). Note that SR is kept zero in this part of the study.

We assume that there are N1 type-1 and N2 type-2 connections in the network. The total number of connections in the network, N=N1+N2, is kept constant at a value that ensures a uplink load (i.e., total mean cell arrival rate/total capacity) of 0.9. Hence there are 29 active connections in the network (i.e., N=29). Note that we have kept the load under control to be able to compare results. In reality, our CAC scheme cannot limit the number of type-1 connections as they are not assigned any SR or BR and the CAC is only based on these parameters. Therefore there is no way to control the network load. As mentioned earlier, it could be a good idea to have B_A, the resources booked but not requested during this frame, to be shared among the type-2 connections first. However, in our tests we have assumed that the whole available capacity C_A is shared among all connections on equal grounds (see section-2.2).

The first test we conducted was to measure the mean queuing delay in the system when we had only connections of type-1. The result of this simulation can be seen in black dashed line in Figure 2 for both RRE-1 and RRE-2 algorithms respectively. Then we introduced type-2 connections with various TTS values (TTS=BR, because SR=0), and measured the mean queuing delay for both type-1 and type-2 connections. The number of type-2 connections in the network was increased from 1 to the maximum possible number, which is computed using the CAC scheme expressed by Equation 2-4 in section-2.2. The results for both RRE-1 and

RRE-2 algorithms are presented in Figure 2. Note that the TTS values are changed from 4 TS/frame to 32 TS/frame, and that solid lines show the mean queuing delay for type-2 connections, while dotted lines of the same color show the mean queuing delay for type-1 connections.

A first observation is that Figure 2 illustrates the superiority of RRE-2 algorithm over RRE-1 algorithm in terms of mean queuing delay experienced by each connection. Note that we are studying the queuing delay in this paper. SAU-to-SAU or SAU-to-gateway cell transfer delay will be approximately 250 msec (i.e., one round trip propagation delay between earth and the GEO satellite) longer than the queuing delay. In this case the minimum mean SAU-to-SAU or SAU-to-gateway transmission delay will be on the order of 550 msec. Queuing delays longer than one response time (i.e., 256 msec in our studies) are caused by RRs that are not fully granted by the BoD controller. While the RRE-1 algorithm requires at least one more response time to re-request the remaining part of a partially accepted RR, the RRE-2 algorithm can re-request it after one frame time (64 msec in our studies). For small values of TTS having RRs not fully granted by the BoD controller is a likely event at the load envisaged in the simulation. Therefore, the difference between the mean queuing delay values for RRE-1 and RRE-2 algorithms is more visible for small values of TTS. For TTS=4 TS/frame, the difference between the mean queuing delay values for both algorithms may go higher than 50 msec.

Figure 2 illustrates the segregation, in terms of mean queuing delay, between type-1 and type-2 connections provided by non-zero TTS. The more we increased TTS for the type-2 connections (i.e., the more resources we were able to book for type-2 connections) the shorter their mean queuing delay became. In return the mean queuing delay experienced by type-1 connections increased.

An interesting observation is that increasing TTS beyond a certain point has almost no effect on the mean queuing delay while it has a large impact on the maximum number of type-2 connections that can be admitted and hence on the network revenue.

As it is seen in Figure 2, it is impossible to reduce the mean queuing delay lower than one response time in the network (i.e. 256 msec if the BoD controller is on board the satellite) with any of the 2 RRE algorithms. For those applications that require a shorter mean queuing delay than one response time we need to assign a non-zero SR. In the next section, we present our study on the influence of SR and BR on queuing delay and resource waste.

3.3 Mean Queuing Delay and Resource Waste

In this part of our study, the user profile is essentially the same as the one described in the preceding section. Type-2 connections are assigned SR and BR such that (SR+BR)>0 and SR is not necessarily equal to zero anymore. The total number of connections N=N1+N2 is kept constant so that the network load is kept constant. Note that we will vary the network load, mean cell arrival rate (i.e., the SCR), and the [SR, BR] pair corresponding to type-2 connections according to the purposes of our tests.

In our first test, each connection (both type-1 and type-2) has the same mean cell arrival rate, SCR=4 cells/frame. Type-2 connections are assigned BR only (i.e., SR=0). The number of type-2 connections is always at the maximum possible value, depending on the BR assigned to each of them, and the CAC scheme expressed by Equation 2-4 in section-2.2. Under these conditions, we measured the mean queuing delay for type-2 connections for various values of BR and the network load. Figure 3 shows the results for both RRE algorithms. It is seen in this figure that as the network load increases from 0.8 to 0.95 the mean queuing delay for a connection with BR=4 TS/frame and SR=0 increases from 375 msec to 600 msec, if RRE-1 algorithms is implemented. The increase is less drastic when RRE-2 algorithm is used. In other words, the mean queuing delay experienced by a type-2 connection is very much dependent on the current network load for low values of BR. As we increase the BR for each type-2 connection, the mean queuing delay and its variation with respect to the network load will reduce. However, as we mentioned before, increasing BR means reducing the number of type-2 connections in the network. Here we observe again the trade-off between the number of type-2 connections (i.e., the GoS) and the mean queuing delay experienced by each type-2 connection.

As the BR is further increased the reduction in the mean queuing delay and its variation with respect to the current network load becomes insignificant for both RRE algorithms. This is obvious since for large BR the connections do not really need to compete for best-effort traffic. Note that the mean queuing delay cannot be less than one response time no matter what BR value is assigned as long as SR=0.

In the second test, the network load is kept constant at 0.9. Each type-2 connection is assigned both SR and BR, where SR+BR=TTS. We have assigned different values to the TTS of type-2 connections. SR is varied from 0 to BR provided that (SR+BR) is kept constant at the assigned value of TTS. The mean queuing delay is measured for type-2 connections, and presented in Figure 4 for SCR=4, 8 and 16 cells/frame. We have chosen to show results for

RRE-2 algorithm only for lack of space. In Figure 4 a somehow surprising result is presented. Indeed we see that increasing SR while keeping (SR+BR)=TTS, **increases** the mean queuing delay experienced by type-2 connections for TTS values that are smaller than a critical value. We define here the Critical TTS as the lowest TTS value for which increasing SR always causes reduction in the mean queuing delay for type-2 connections. If TTS is less than the Critical TTS, increasing SR may increase the mean queuing delay for type-2 connections rather than decreasing it. On the contrary, if TTS is greater than the Critical TTS, increasing SR will always reduce the mean queuing delay for type-2 connections.

Remember that SR can only be used by the type-2 connection to which it is assigned since we have assumed a single connection per SAU. As we have explained in section-2.2, the more resources are assigned statically, the less statistical resource sharing takes place among the connections. For a type-2 connection increasing SR will have two opposing side-effects:

1. A reduction in the mean queuing delay, because every time a burst arrives at the buffer a number of cells will leave quickly using the statically allocated TSs.

2. An increase in the mean queuing delay, because the reduction in the statistical resource sharing will increase the probability that the BoD controller will not fully grant some RRs.

Depending on the traffic characteristics of the connections, the [SR, BR] pair assigned to the type-2 connections, and the current network load, one of the side-effects listed above will prevail, and the mean queuing delay will either decrease or increase.

In Figure 5 we present the ratio Critical TTS/SCR as a function of the burstiness of the connections (i.e., PCR/SCR). It is clearly seen in this figure that as the burstiness of the connections decreases, the critical TTS approaches the mean cell arrival rate of the connection. As the burstiness of the connections increase the Critical TTS/SCR ratio increases almost linearly. Note that to obtain Figure 5, we have varied the SCR of a connection by varying the mean OFF period of the source while keeping the mean ON period constant. We could observe a different relation between the Critical TTS/SCR ratio and the burstiness of the source if we had adopted a different method to vary the SCR of the connections. Observing Figure 5 reveals that for a connection with a burstiness factor of 8, the Critical TTS is 6 times the SCR of the connection. Hence the connection must be assigned a TTS greater than or equal to 6 times its SCR if we want to be sure that the mean queuing delay experienced by the connection will decrease by increasing the SR/BR ratio. This has obviously a very bad impact on the GoS. Furthermore, there will be a lot of resource waste at the SAUs because SR will not be used during the time the source is idle. Figure 6 presents the percentage of TSs that are wasted at the SAU with respect to the SR assigned to the connection for various TTS values, for SCR=4 cells/frame, 8 cells/frame, and 16 cells/frame respectively. As it is seen in this figure, TS waste is especially high for bursty connections. Finally, Figure 7 shows the variation of the mean queuing delay for type-2 connections with respect to the percentage of TS wasted at the SAUs because of non-zero SR for various TTS values and for SCR= 4, 8, and 16 cells/frame. For SCR=4 cells/frame, which represents the most bursty connection, the steady reduction in mean queuing delay occurs for only those TTS values that are higher than or equal to 20 TS/frame. More than 80% of the total TSs allocated to a type-2 connection are wasted in order to reduce the mean queuing delay from 300 msec to 100 msec. As the burstiness of the connections decreases the TS waste percentage corresponding to a mean queuing delay value decreases as well. Note that, even for SCR=16 cells/frame, which represents the less bursty connection, almost 35 % of all TSs allocated to a type-2 connection are wasted for a mean queuing delay value of 100 msec for type-2 connections.

The second side-effect of increasing SR, namely the increase in the mean queuing delay because the statistical resource sharing is reduced, can be mitigated by limiting the amount of resources that can be statically allocated to connections. This implies a modification of our CAC scheme, which was expressed by Equation 2.4 in section-2.2. The modified CAC scheme is shown in Equation 3-1 below.

$$\left\{ \begin{array}{c} SR_j + \sum_k SR_k \le e \cdot C \\ SR_j + BR_j + \sum_k SR_k + \sum_k BR_k \le C \end{array} \right\} \text{ where } 0 < e < 1 \qquad \text{Equation 3-1}$$

This corresponds to a 2-stage CAC scheme. The first stage checks if the total amount of static resources to be assigned to all connections exceeds the upper limit on the amount of resource that can be statically allocated to connections, which is denoted as $e \cdot C$. The second stage of CAC, which is identical to the CAC proposed in section-2.2, checks if the total SR and BR to be assigned to connections exceeds the capacity of the network. The right value to be chosen for e is a further research topic.

3.4 Queuing Delay Histograms

In this section, we provide a more detailed view of the queuing delay characteristics. We ran tests that built histograms for the queuing delay experienced by type-2 connections for both RRE-1 and RRE-2 algorithms. The network load is kept constant at 0.9.

Figure 9 shows the queuing delay histograms for RRE-1 algorithm. In Figure 9, SR+BR=TTS is kept constant at 12 TS/frame. The first part of the plot shows the histogram for SR=0 and BR=TTS. In the second part, TTS is divided equally between SR and BR, and in the third part TTS = SR. It is clearly seen in the first part of Figure 9 that most of the cells experience a queuing delay of about 320 msec, or of about 520 msec. The first peak represents those cells that could successfully depart from the SAU buffer after one response time, and the second peak represents those cells that had to wait for another response time before the necessary resource is allocated. In case of heavy congestion at the BoD controller, we could have observed more than two peaks, each at a multiple of the response time. As we increase the SR/BR ratio, the histograms become wider. In this case, we also observe small periodical peaks superimposed on the histograms. This is because the TSs that are allocated statically to a connection are placed back-to-back within the MF-TDMA frame at fixed locations, which causes periodical back-to-back departures from the SAU buffer. As SR rises so does the height of these periodical peaks.

Figure 10 shows the histograms for RRE-2 algorithm. When compared to the first part of Figure 9, the first part of Figure 10 presents only one peak at 320 msec, the histogram decreasing then almost exponentially to zero. This is because of the memory at the BoD controller. Because of this memory element, in case of congestion, those cells wait at the SAU buffer for a number of frames as opposed to a number of response times.

Perhaps the most important conclusion we should draw from all these histograms is the sudden increase in the variance of the queuing delay when SR>0. A closer observation of the histograms reveals that assigning a connection a non-zero SR can triple the range of the queuing delay.

High variance in the queuing delay can become a serious problem in case for example where TCP/IP is running on top of ATM. Latency is central to TCP operations, and high delay variations can trigger many throughput-related problems. In the case of important variations in the delay, the most frequent problem that we can expect to encounter will be excessive and unnecessary retransmissions from acknowledgment timers expiring whenever the round-trip time becomes significantly larger than expected. A related effect, which may also be encountered, is if the end-systems interpret the increased latency as congestion/loss related, and start modifying their congestion windows. Hence, very severe degradations of performance can be encountered. Therefore, allocating SR, i.e., allocating statically TSs to connections can have adverse effects on their performance if not done carefully. In fact, there is a need to conduct a much larger scale study on the interactions between TCP/IP and the BoD process.

4. Conclusions

Satellites are multiple access systems with very long propagation delay and scarce transmission capacity compared to terrestrial networks. End-to-end resource management is necessary for the efficient sharing of satellite resources, and for QoS support. The BoD process is central to end-to-end resource management in satellite networks, and SR (Static Resource) and BR (Booked Resource) are the means with which the BoD process guarantees various levels of QoS to connections and the CAC guarantees GoS to some of the connections.

Our simulations illustrated that SR and BR can be successfully used to segregate between different service classes. First we studied the case when SR=0. For values of BR that are close to the SCR of the connection, the mean queuing delay experienced by the connection heavily depends on the current network load. As the network load increased from 0.8 to 0.95, the mean queuing delay increased from 375 msec to 600 msec. The connections that use BoD with non-zero SR+BR, such as nrt-VBR or GFR, are the ones that bring the most revenue to the network. It is possible to reduce the mean queuing delay experienced by these connections by increasing the BR that is assigned to them. We have observed that, for SCR=4 cells/frame, a BR value of 10 cells/frame yields a mean queuing delay between 320 msec and 400 msec depending on the current network load. However, we should be careful since increasing BR decreases the maximum number of this type of connections that can be accepted into the network. Note that we have kept the network load constant in our studies. In other words, we have limited the number of type-1 connections in the network. Type-1 connections represented UBR connections as they are not assigned any SR or BR. In a real satellite network, because CAC cannot be implemented on UBR connections, the only way to ensure that the network load will be within reasonable limits is to well dimension the network capacity. Alternatively, as we have explained in section-2.2, we may choose a different allocation philosophy for the best effort capacity at the BoD controller.

Further reduction in mean queuing delay could be obtained by the introduction of SR. In our, studies we have observed that if SR+BR (i.e., TTS) is not high enough, increasing SR may

increase the mean queuing delay experienced by the type-2 connections. This is because increasing the SR decreases the total amount of resources that can be statistically shared among the connections in the network. We defined the Critical TTS value as the minimal value of SR+BR such that the mean queuing delay experienced by the type-2 connections decrease monotonically when increasing SR. We found a linear relation between the ratio Critical TTS/SCR and the burstiness (i.e., PCR/SCR) of type-2 connections. For highly bursty connections we need to assign a TTS much higher than the connections' SCR in order to reduce the mean queuing delay by increasing SR. This not only decreases the number of type-2 connections that can be accepted into the network (hence reducing the GoS) but also increases the TS waste at the SAUs. Our simulation analysis illustrated that, for type-2 connections, a reduction of the mean queuing delay from 300 msec down to 100 msec is only possible with TS waste percentage between 35% and 80% depending of the burstiness of the connections.

We have proposed to limit the amount of resources that can be statically allocated to type-2 connections in order to attain shorter mean queuing delay for type-2 connections with smaller TS waste percentages.

Another side effect of SR is the sudden increase in the variance of the queuing delay. This may yield very low end-to-end throughput if end-to-end flow control protocols that are based on sliding windows such as TCP are running above ATM specific layers. The impact of SR on TCP should be further studied.

5. Illustrations and Figures

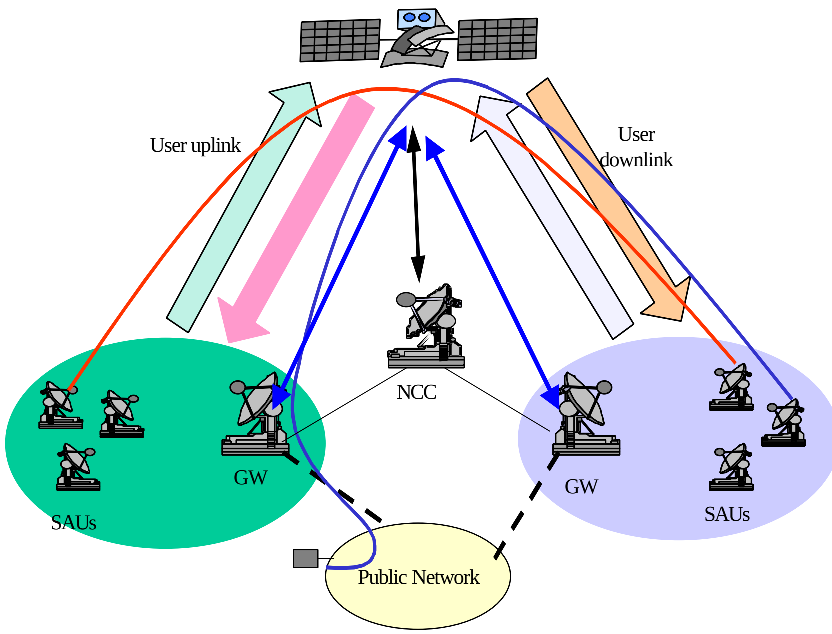

Figure 1: Next Generation Satellite Network

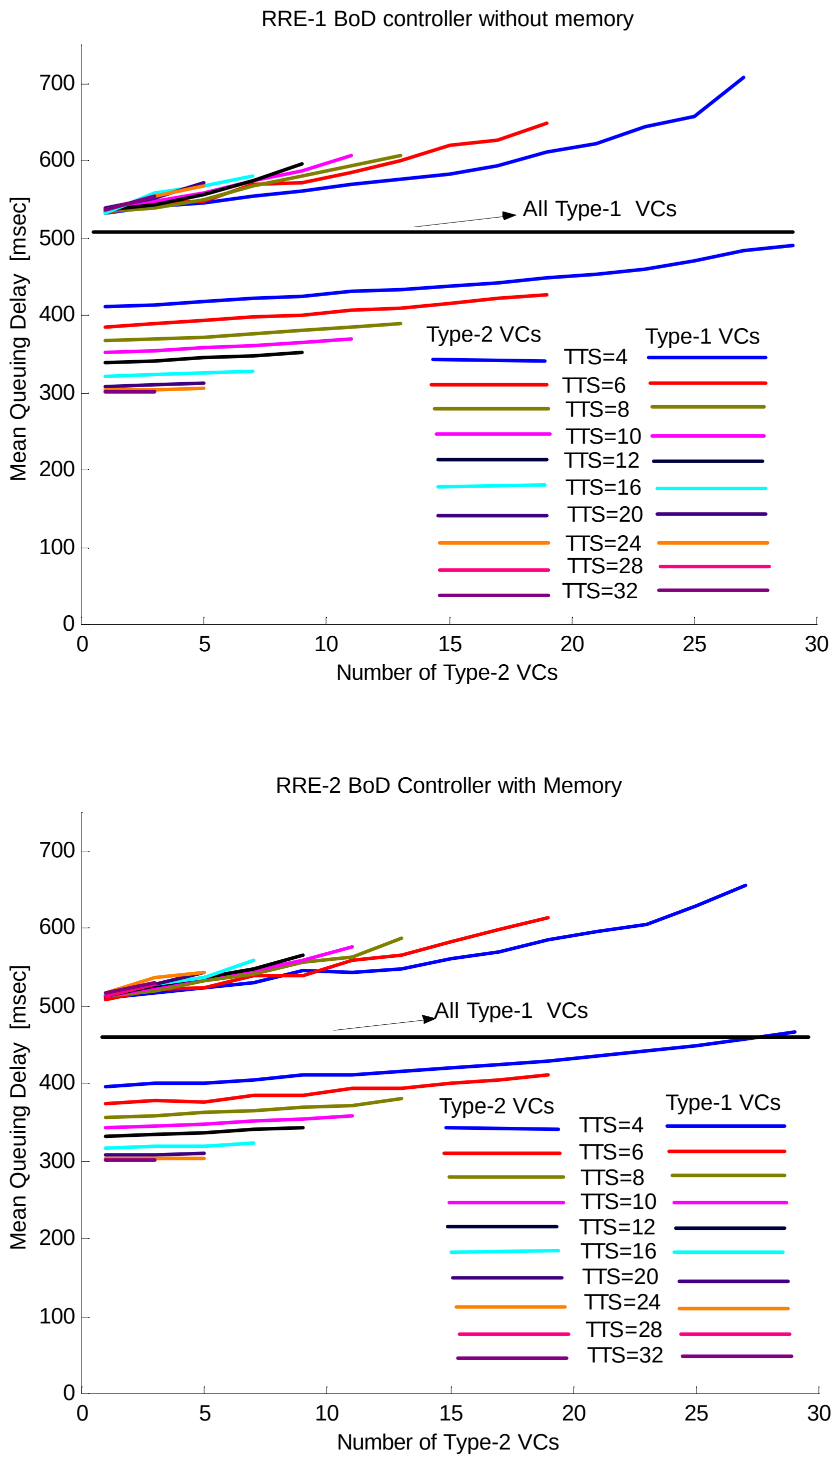

Figure 2: Mean Queuing Delay vs Number of Type-2 Connections

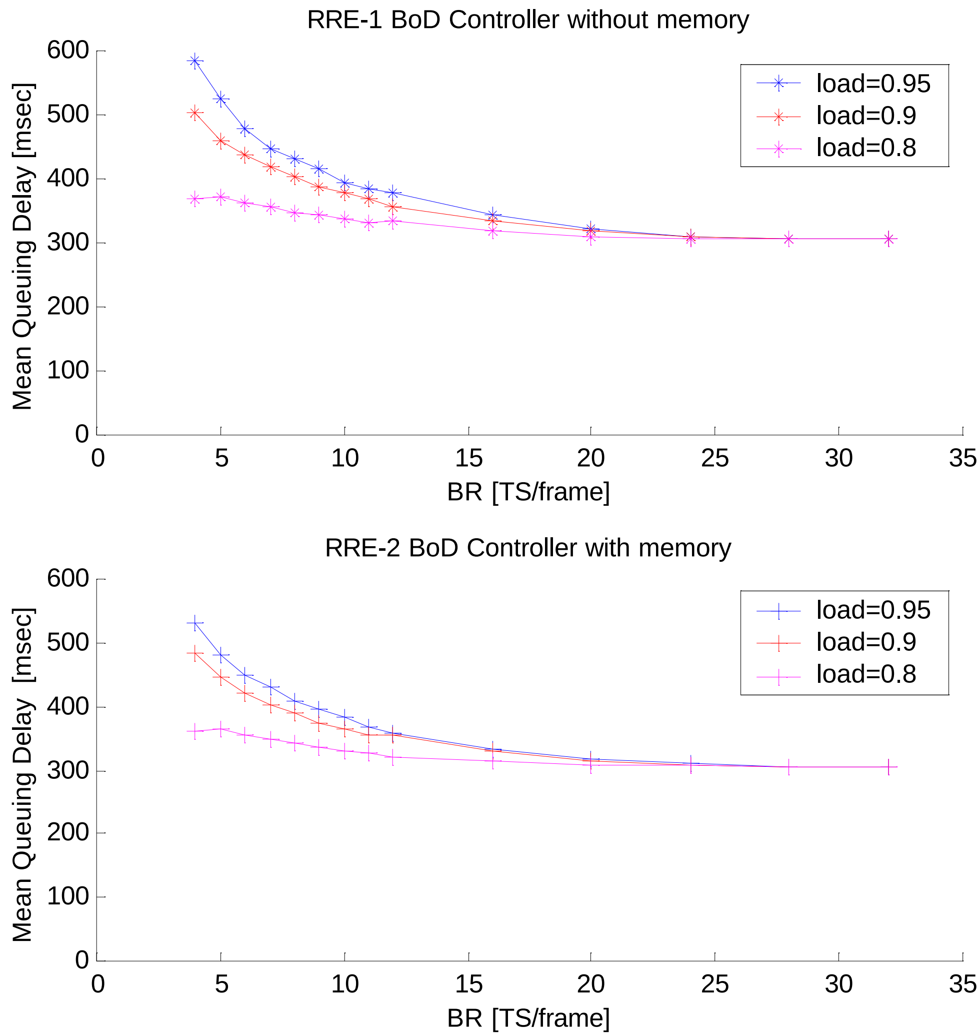

Figure 3: Mean Queuing Delay vs Number of Booked TSs (SR = 0)

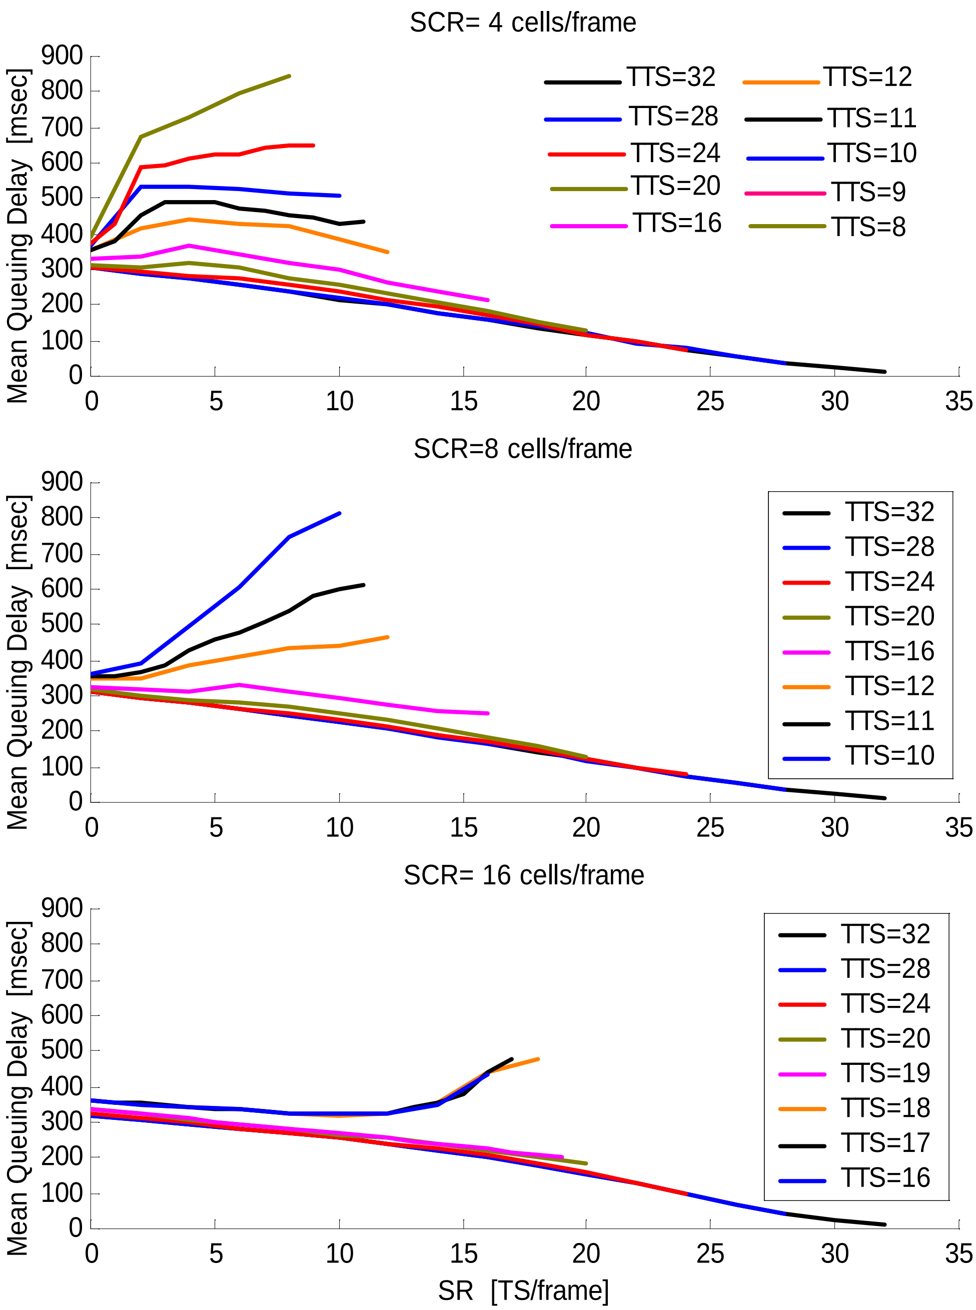

Figure 4: Mean Queuing Delay vs SR

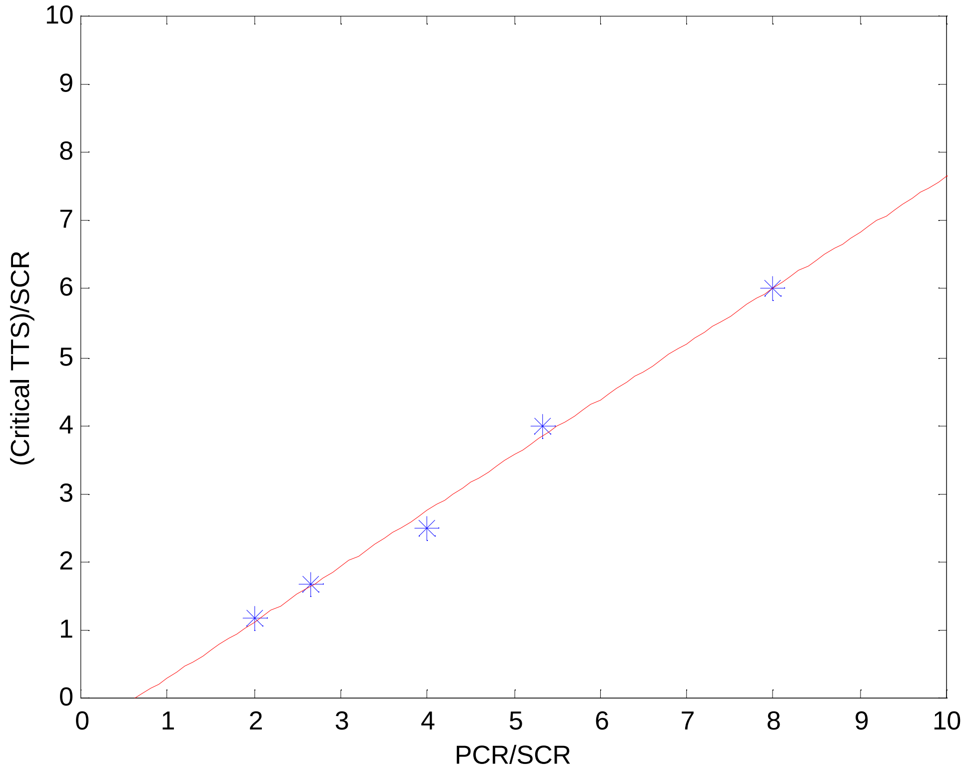

Figure 5: Critical TTS/SCR vs Burstiness

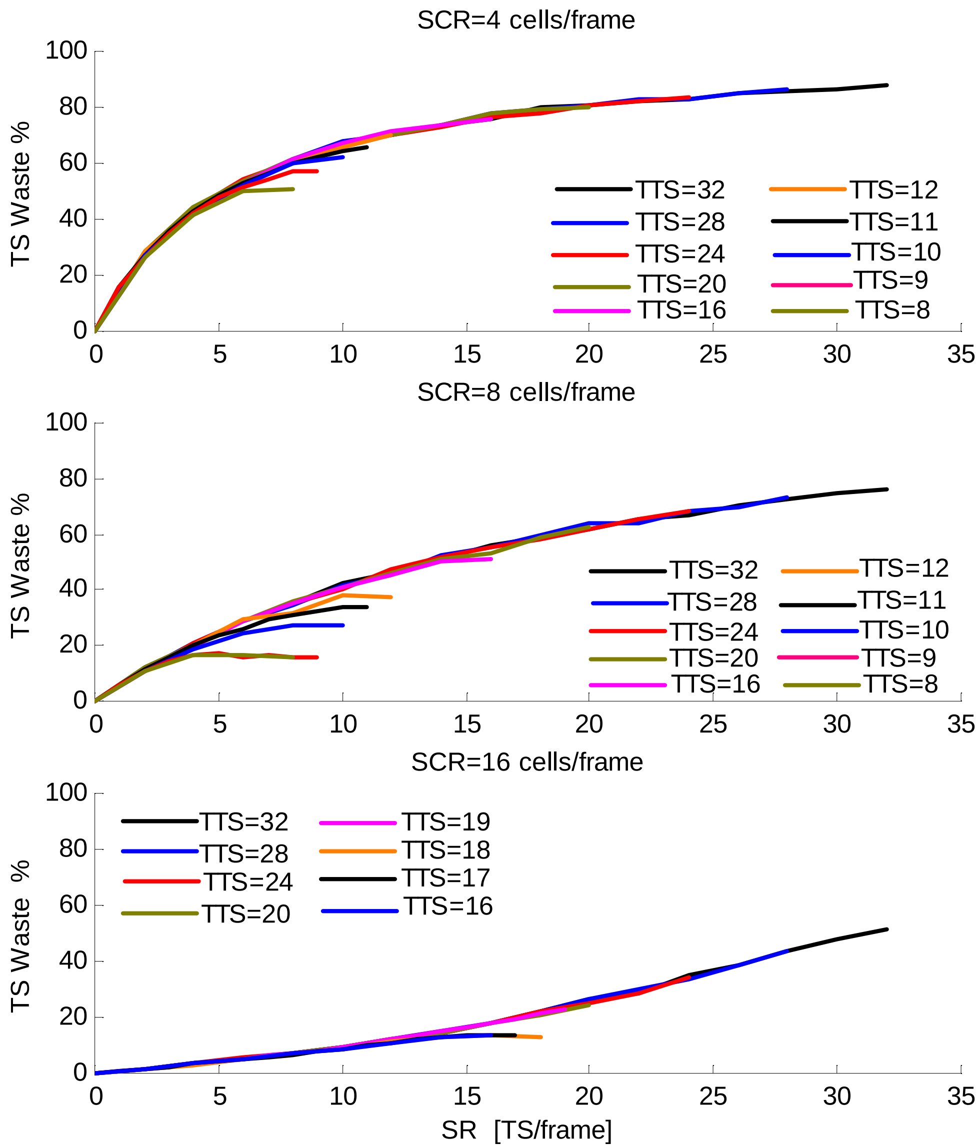

Figure 6: TS Waste Ratio vs SR for various TTS

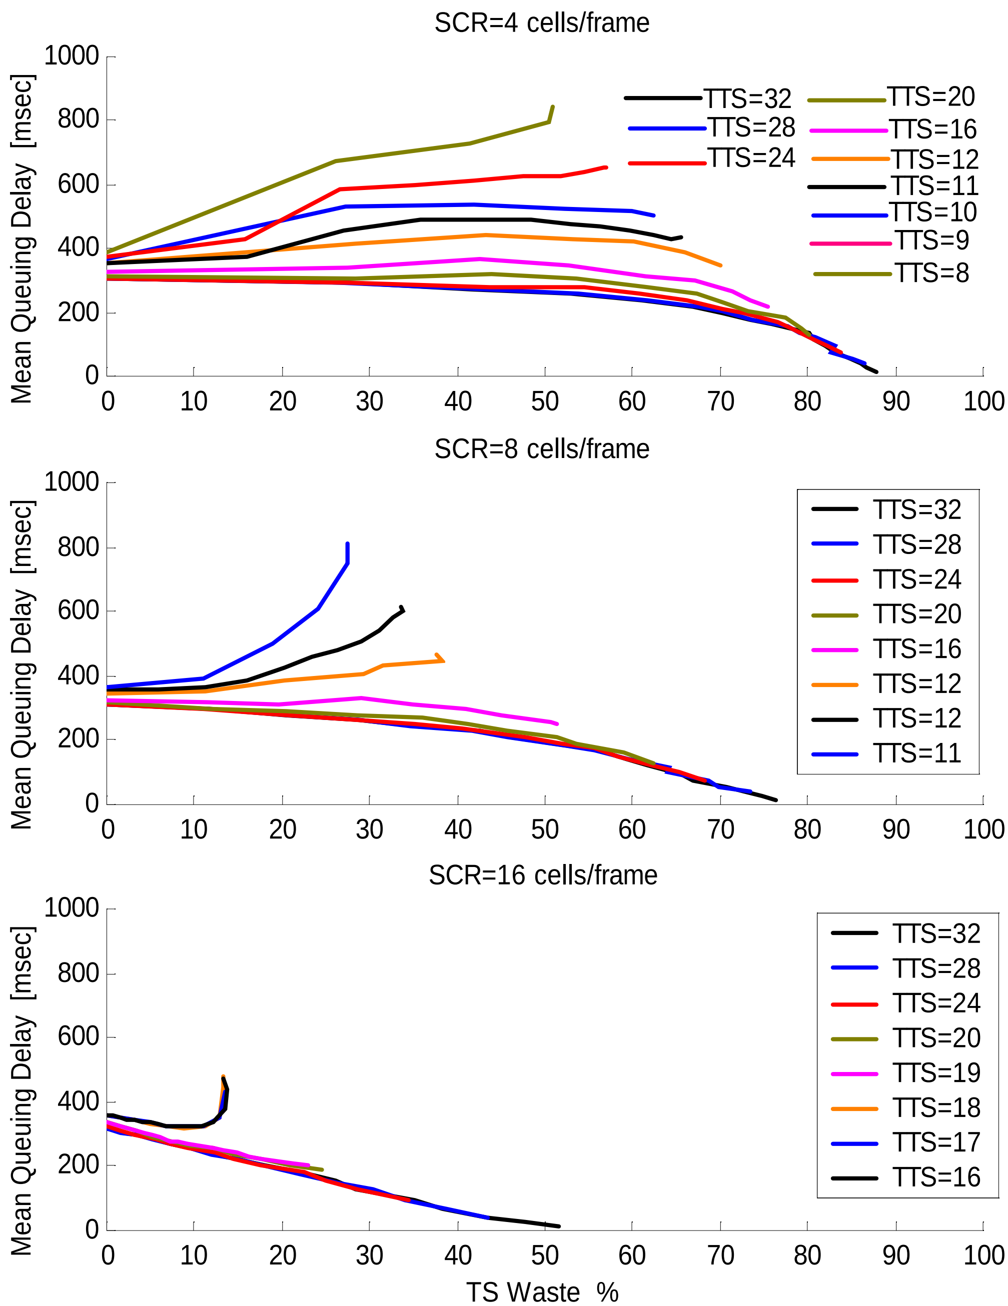

Figure 7: Mean Queuing Delay vs TS Waste

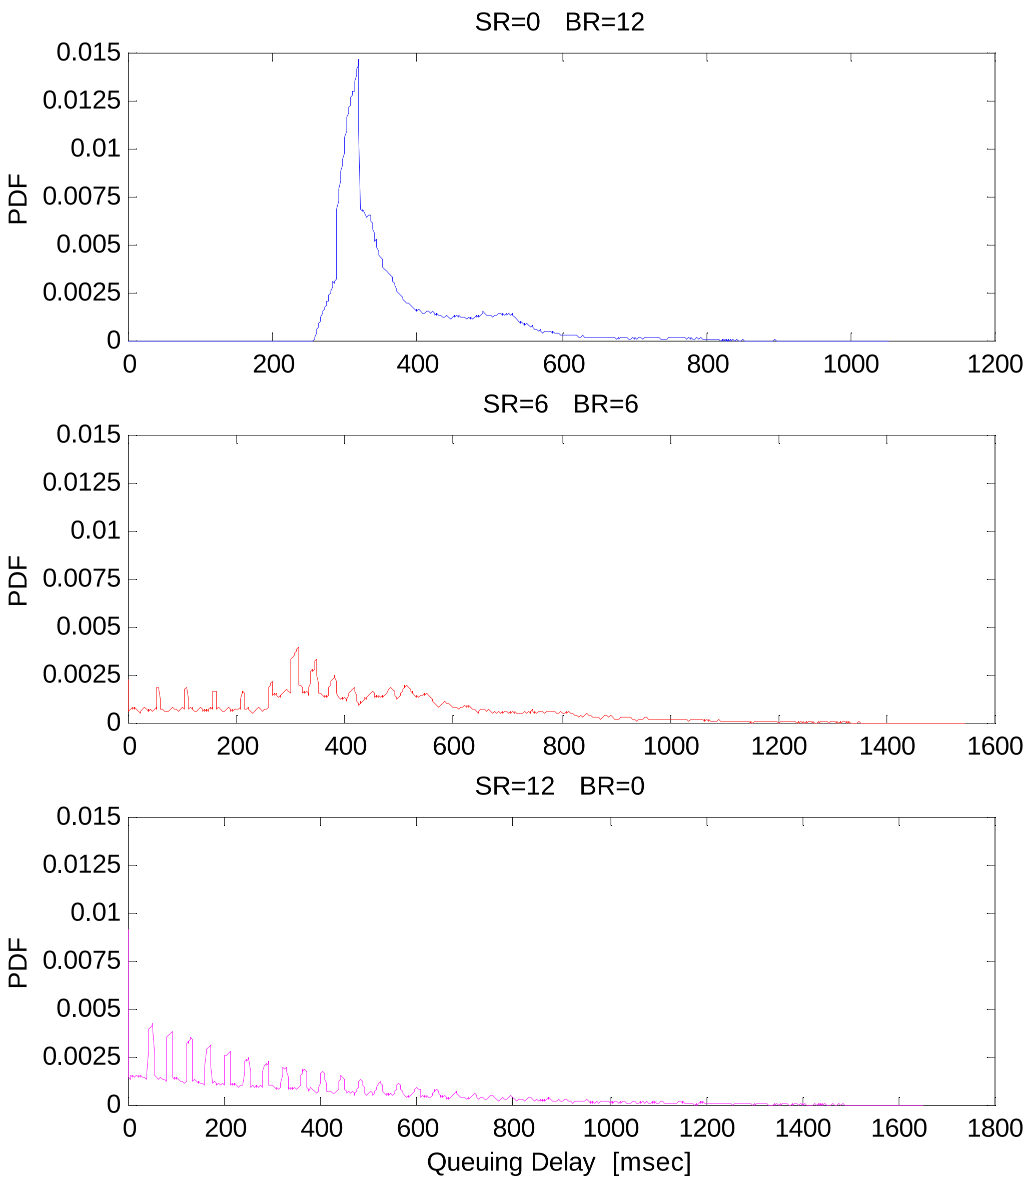

Figure 9: Queuing Delay Histograms for RRE-1 and TTS = 12 TS/frame

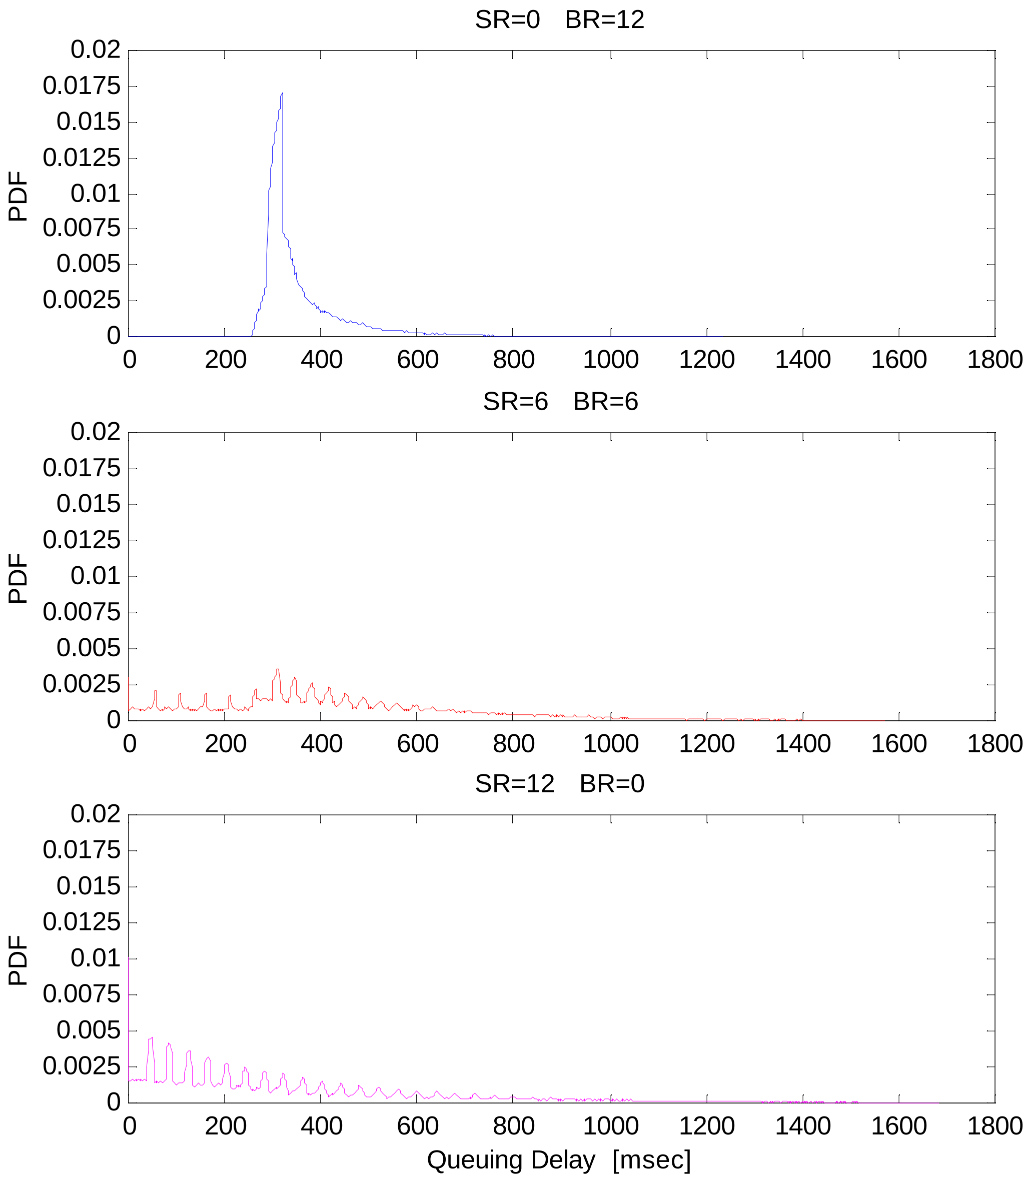

Figure 10: Queuing Delay Histograms for RRE-2 and TTS = 12 TS/frame

6. References

[1] G. Açar and C. Rosenberg, "Algorithms to compute for Bandwidth on Demand Requests in a Satellite Access Unit," *Proceedings of 5th Ka Band Utilization Conference*, Taormina, Italy, October 1999.

[2] Ian F. Akyildiz and S. Jeong, "Satellite ATM Networks: A Survey," *IEEE Communications Magazine, Vol. 35, No. 7, pp. 30-44, July 1997.*

[3] A. Girard, C. Rosenberg and M. Khemiri, "Fairness and Aggregation: A Primal Decomposition Study," *IFIP Networking 2000*, Paris, France, May 2000.

[4] *IEEE Communications Magazine*, Vol. 35, No. 7, pp. 6-65, July 1997.

[5] *IEEE Communications Magazine*, Vol. 37, No. 3, pp. 8-72, March 1999.

[6] *IEEE Journal on Selected Areas in Communications*, Vol. 17, No. 2, February 1999.

[7] H. Yaiche, R. Mazumdar and C. Rosenberg, "A Game Theoretic Framework for Rate Allocation and Charging of Available Bit Rate (ABR) Connections in ATM Networks," *IFIP Broadband Communications'98*, Stuttgart, March 1998.

[8] C. Rosenberg, "End-to-end Resource Management for ATM On Board Processor Geostationary Satellite Systems," *Proceedings of 4th Ka Band Utilization Conference, Venice, Italy, pp. 481-488, November 1998.*

[9] C. Rosenberg, Y. Qian, R. Hu, and H. Abu-Amara, “Integrated Connection Admission Control and Bandwidth on Demand Algorithm for a Multimedia Satellite Network,” submitted to IEEE Trans. On Networking, Feb. 20000.