

## Multicast-aware power allocation in multiple spot-beam satellite communication systems

Gun Akkor<sup>\*,†</sup>, Michael Hadjitheodosiou<sup>‡</sup> and John S. Baras<sup>§</sup>

*Institute for Systems Research, and Center for Satellite and Hybrid Communication Networks,  
University of Maryland, College Park, MD 20742, U.S.A.*

### SUMMARY

In this paper, we address the problem of user heterogeneity in satellite multicast from the perspective of resource allocation in a multiple spot-beam satellite system that supports both unicast and multicast flows. Satellite communication systems, with their wide-area coverage and direct access to large number of users, clearly have an inherent advantage in supporting multicast applications. In order to remain competitive against other broadband technologies, however, next generation satellite systems will be required to support both unicast and multicast flows and offer optimal sharing of system resources between these flows. We show that user heterogeneity across spot-beam queues may result in lower allocated session rates for active flows, and be perceived as unsatisfactory by potential users when both unicast and multicast flows are active in the system. We propose an optimization-based approach that allocates resources with the goal of smoothing user heterogeneity, and show that resulting session rates are higher on the average for both unicast and multicast flows. This is achieved through the re-distribution of system power among spot-beam queues, by taking into account the load on the queues and the channel states. We conclude that it is possible to increase the average session rates of multicast flows by 25–100%, and the rates of unicast flows by 15–40% compared to the pre-optimization levels. Copyright © 2006 John Wiley & Sons, Ltd.

KEY WORDS: system design; multicast delivery; satellite networks; power allocation

### 1. INTRODUCTION

The role satellite systems play in today's communication infrastructure is changing rapidly. This change is fueled by two main ingredients. The first one is the technological advances in the design of new satellite systems. Next generation satellite communication systems that utilize higher frequency bands, such as the Ka-band, and support spot-beam technology, on-board packet processing and switching are currently under development [1–3]. These new technologies

\*Correspondence to: Gun Akkor, Patton Electronics, Co., 7622 Rickenbacker Drive, Gaithersburg, MD 20879, U.S.A.

<sup>†</sup>E-mail: gakkor@patton.com

<sup>‡</sup>E-mail: michalis@isr.umd.edu

<sup>§</sup>E-mail: baras@isr.umd.edu

Contract/grant sponsor: NASA; contract/grant number: NCC8235

allow higher data rates and enable the use of small, low-power, and low-cost user terminals, making satellite communication systems more competitive against other broadband communication solutions (cable, DSL) in providing integrated voice, data, and multimedia communications. The second component is the set of new applications, such as on-demand multimedia content delivery, distance learning, and distributed software updates, that have recently emerged in the Internet. These applications are distributed in nature and require concurrent transmission of the same content to multiple users. Satellite communication systems, with their wide-area coverage, direct and ubiquitous access to large number of users, clearly have an inherent advantage in supporting such services [4].

Despite the promising technologies and the potential for multicast content delivery over satellite networks, however, such services remain largely unavailable due to the lack of an incentive to deploy them. From the network service providers' point of view, there will be an incentive to use multicast delivery only if it results in considerable bandwidth savings and allows deployment of new applications. The problem of providing users with an incentive to use multicast delivery is more difficult. From a user's point of view, a high service satisfaction, such as high transmission speed or low delay, is required whether the provider uses unicast or multicast to deliver the content. In order to make multicast delivery rewarding to both parties, the requirements of multicast communication and the constraints of the underlying system, such as the network topology and the physical transmission medium, should be tuned to each other [5].

One key performance issue that has to be addressed in multicast deployment is the *user (receiver) heterogeneity*. Multicast communication is significantly complicated by the widely varying capabilities of the user group in terms of available network bandwidth, supportable transmission rates, and traffic load across different network domains. This heterogeneity forces the multicast session to adjust to the capabilities of the user(s) with the worst performance. Therefore, the group experiences a performance that is below the network's capability.

In this paper, we address the problem of user heterogeneity from the perspective of resource allocation in a multiple spot-beam satellite system that supports both unicast and multicast flows. We show that user heterogeneity across spot-beam queues may result in lower allocated session rates for active flows, and be perceived as unsatisfactory by potential users when both unicast and multicast flows are active in the system. We propose an optimization-based approach that allocates resources with the goal of smoothing user heterogeneity, and show that resulting session rates are higher on the average for both unicast and multicast flows.

The rest of the paper is organized as follows. In the next section, we outline the problem in the context of our target satellite system architecture, and identify the key issues. In Section 3, we formulate our problem in an optimization framework. Section 4 provides the solution, and Section 5 discusses the analysis framework we have developed for testing the performance of our approach. In Section 6, we present numerical performance results. Last section concludes the paper and draws attention to future work on this subject.

## 2. MOTIVATION

User heterogeneity in multicast communication is a well-known problem and has been investigated extensively in the literature [6–12]. However, bulk of this research targets primarily wireline networks, where the design and solution spaces for the problem are significantly




Figure 1. Satellite communication system architecture. The satellite provides broadband access to users across multiple spot-beam locations.

different than that of the satellite networks. In this section, we identify the sources of user heterogeneity in the context of satellite networks.

In this paper, we consider a star topology satellite network, where a Ka-band, geosynchronous satellite provides broadband services to a large number of users located inside its footprint. The satellite supports multiple spot-beams and on-board packet switching technologies that allow transmission of data to multiple users in multiple spot-beam locations (Figure 1). Users that are equipped with two-way direct communication terminals access the backbone terrestrial network through a gateway node referred to as the network operations centre (NOC).

In this multiple spot-beam satellite system, we consider the delivery of content from NOC to users in response to requests. We assume that each delivery session is identified by a unique flow, and packets of several active flows are queued at the NOC, which forwards them to the satellite at a rate limited by the uplink capacity of the system. An on-board processor and switch forward the packets to one or multiple spot-beam queues, duplicating the packets in the latter case. A packet belonging to a unicast flow is forwarded to a single spot-beam queue, corresponding to the spot-beam location, in which the end user resides. In case of a multicast flow, however, members of the multicast session may reside in multiple spot-beam coverage areas, and therefore, packets need to be duplicated and forwarded to multiple spot-beam queues on-board the satellite (Figure 2).

At every spot-beam queue, several flows (unicast and multicast) *share* the total service rate of the queue. The rate-share of a flow that belongs to a particular queue depends on several factors, such as the number of flows currently active in that queue, the type of the flows, and the rate allocation policy between different type of flows, i.e. unicast and multicast. We call this rate-share as the *supportable session rate* of the flow at that particular queue. A unicast flow would have a single supportable session rate determined by the queue corresponding to the spot-beam location, in which the end user resides. However, a multicast flow could have a set of supportable session rates if the user group were to span more than one spot-beam location. In order to avoid over-flowing of any of the on-board queues, the input rate of a flow at the NOC queue has to be determined by the minimum rate that the flow can be served at the spot-beam queues. We refer to this rate as the *maximum sustainable session rate* of the flow. For a unicast flow, the maximum sustainable session rate at the NOC queue equals to its single supportable session rate. The sustainable session rate of a multicast flow, on the other hand, is determined by

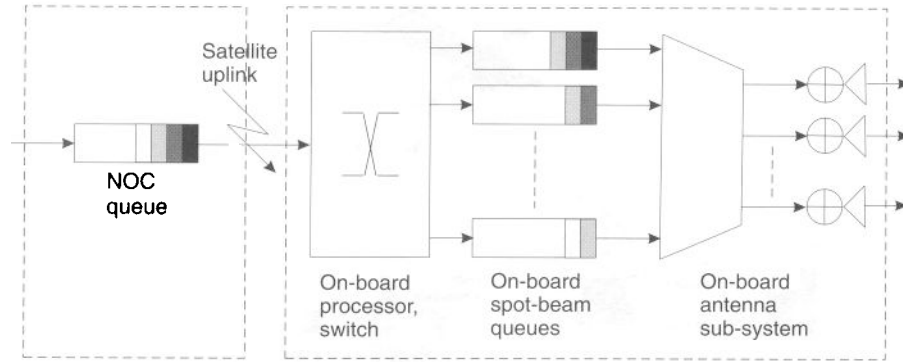


Figure 2. Conceptual view of the on-board satellite system and queues.

the *minimum* of its supportable session rates across multiple spot-beam queues. Therefore, heterogeneity in the supportable session rates of a multicast flow would limit its maximum sustainable rate, and may be deemed unsatisfactory by some users.

In this system, heterogeneity in the supportable session rates of a multicast flow would occur depending on several factors:

- The service rate of each spot-beam queue varies as a function of the allocated power and the channel state, for a given modulation scheme and bit error rate (BER) target. Therefore, not all queues can be served at the same effective rate.
- The packets of a unicast flow affect the load on only a single spot-beam queue, while, in case of a multicast flow, a single session may affect the load on several spot-beam queues. Factors, such as the distribution of users across geographical spot-beam locations and the type of the flows, affect the individual rate-shares of the flows.
- The rate allocation policy on how unicast flows are treated in comparison to multicast flows sharing the same queue, directly affects how the service rate of the queue is shared between active flows. Therefore, supportable session rate of a multicast flow is determined by the number of active unicast flows in the same queue, causing a variation across different spot-beam queues.

In this paper, the goal is to minimize the heterogeneity multicast flows experience by smoothing the variation in the supportable session rates of all flows. We show that this type of smoothing could result in higher rate allocations for most active flows, improving the total utilization of the system. In the following section, we describe this approach in an optimization framework and specify the parameters of interest.

### 3. PROBLEM FORMULATION

In this system,  $M$  on-board spot-beam queues are served by  $K$  on-board antennas in a time-divided manner. The downlink transmission is organized into bursts, each of which occupies a fixed time interval. During a burst, an antenna serves only one spot-beam queue. We define the

time it takes to serve each spot-beam queue only once with no antenna idling as a *transmission round*. A transmission round can be viewed as a frame of  $K$  rows, each corresponding to an on-board antenna, and  $L = M/K$  columns, where we assume, without loss of generality, that  $L$  is an integer. We denote by  $\mathcal{A}_l$ ,  $l = 1, 2, \dots, L$ , as the set of spot-beam queues that are served simultaneously (corresponding to a column of the frame).

The transmission rate  $r_j$  of spot-beam  $b_j$ ,  $j = 1, 2, \dots, M$ , at the time of its burst interval, depends on the allocated power  $p_j$ , and the current channel state  $s_j$ , according to a general concave rate-power curve  $\mu_j(p_j, s_j)$ . For any state  $s_j$  of the downlink channel, rate-power curve represents the rate, under a specific set of coding schemes, that achieves a target bit error rate as a function of the allocated power. The power levels of all beams satisfy:

$$0 \leq l_j \leq p_j \leq P_{\text{tot}}, \quad j = 1, 2, \dots, M \quad (1)$$

and

$$\sum_{j \in \mathcal{A}_l} p_j = P_{\text{tot}}, \quad l = 1, 2, \dots, L \quad (2)$$

where  $P_{\text{tot}}$  is the total available system power and  $\{l_j\}_{j=1}^M$  is a set of lower bounds on the power levels of the queues.

A flow  $f_i$ , for  $i = 1, 2, \dots, N$ , which is forwarded to spot-beam queue  $b_j$  is assigned a rate-share  $w_{ij}$  of the service rate of the queue, depending on the load of the queue, and the type of the flows forwarded to it, such that

$$w_{ij} = 0 \quad \text{if } i \notin \mathcal{B}_j \quad (3)$$

$$0 < w_{ij} \leq 1 \quad \text{if } i \in \mathcal{B}_j \quad (4)$$

$$\sum_{i \in \mathcal{B}_j} w_{ij} = 1, \quad j = 1, 2, \dots, M \quad (5)$$

where  $\mathcal{B}_j$  is the set of all flows that are forwarded to the spot-beam queue  $b_j$ . Therefore, the packets of flow  $f_i$  could be served at a *supportable session rate* of

$$\lambda_{ij} = w_{ij} \cdot r_j = w_{ij} \cdot \mu_j(p_j, s_j) \quad (6)$$

at the spot-beam queue  $b_j$ . However, the *maximum sustainable session rate* of the flow at the NOC queue is limited to the minimum rate that the flow could be served across spot-beam queues, i.e.

$$\lambda_i = \min_{j: i \in \mathcal{B}_j} \{\lambda_{ij}\} \quad (7)$$

in order to avoid overflowing of the spot-beam queues (Figure 3).

For unicast flows, there exists a single spot-beam queue index  $j$  for which  $i \in \mathcal{B}_j$ , corresponding to the beam where the destination user resides. However, for multicast flows, there are several indices for which this may be true. The variation in  $\{\lambda_{ij}\}_{j=1}^M$  can be minimized by adjusting the service rates of spot-beam queues, i.e.  $\{r_j\}_{j=1}^M$ . The service rates, in turn, depend on the allocated power levels and the channel states. Therefore, our goal is to minimize this

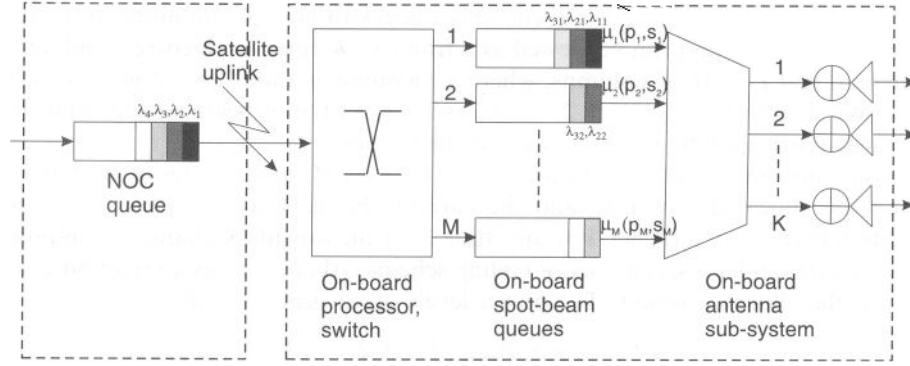


Figure 3. Representation of supportable session rates and maximum sustainable session rates for four active flows.

variation by arranging the power levels allocated to each queue, subject to a total power constraint, and a set of given channel states. In other words, we would like to find the optimal vector of power levels  $\mathbf{p}^* = [p_1^* \dots p_M^*]$  that would minimize the sum of the rate variances of all multicast flows across spot-beam queues:

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \sum_{i=1}^N \sigma_i^2 \quad (8)$$

subject to constraints

$$0 \leq l_j \leq p_j \leq P_{\text{tot}}, \quad j = 1, 2, \dots, M \quad (9)$$

$$\sum_{j \in \mathcal{A}_l} p_j = P_{\text{tot}}, \quad l = 1, 2, \dots, L \quad (10)$$

$$\text{given } \mathbf{s} = [s_1 \dots s_M] \quad (11)$$

where

$$\sigma_i^2 = \frac{1}{N_i} \sum_{j=1}^M x_{ij} \cdot (\lambda_{ij} - m_i)^2 \quad (12)$$

$$m_i = \frac{1}{N_i} \sum_{j=1}^M x_{ij} \cdot \lambda_{ij} \quad (13)$$

$$x_{ij} = \begin{cases} 1 & \text{if } i \in \mathcal{B}_j \\ 0 & \text{if } i \notin \mathcal{B}_j \end{cases} \quad (14)$$

$$N_i = \sum_{j=1}^M x_{ij} \quad (15)$$

Note that for unicast flows,  $N_i = 1$ , and  $\sigma_i^2 = 0$ . Therefore, unicast flows do not contribute to the cost function, but they affect the solution since they change the total load on the system and consequently the rate-shares of every flow, i.e.  $\{w_{ij}\}$ . In the remainder of this paper, the rate-power curve is assumed to be of the form  $r_j = \beta(s_j) \cdot p_j$ ,  $\forall j$ . This assumption is later validated in Section 5. In the next section, we provide a solution to (8).

#### 4. SOLUTION

When no distinction is made among the spot-beam queues, the simplest assignment would be to set equal power levels for all, such that

$$p_j = \frac{P_{\text{tot}}}{K}, \quad j = 1, 2, \dots, M \quad (16)$$

We call this assignment, the *equal-antenna-share* (EAS) policy and denote it by the vector  $\mathbf{p}^{\text{EAS}}$ . Given the channel state vector  $\mathbf{s}$ , the power vector  $\mathbf{p}^{\text{EAS}}$  completely determines the service rate of each spot-beam queue and the sustainable session rate of each active flow:

$$r_j^{\text{EAS}} = \min(\beta(s_j) \cdot p_j^{\text{EAS}}, r_{\text{max}}), \quad j = 1, 2, \dots, M \quad (17)$$

$$\lambda_i^{\text{EAS}} = \min_{j:i \in \mathcal{J}_j} \{w_{ij} \cdot r_j^{\text{EAS}}\}, \quad i = 1, 2, \dots, N \quad (18)$$

where  $r_{\text{max}}$  is the maximum system downlink rate determined by the set of available modulation and coding methods. In the remainder of this paper, we use the EAS( $\mathbf{p}^{\text{EAS}}, \lambda^{\text{EAS}}$ ) policy as the basis for comparison.

Equation (17) states that, the EAS policy power assignments may be in excess of the power levels required to achieve the maximum system downlink rate for a given channel state. From (18), we can conclude that the supportable session rates of a multicast flow, which are higher than the minimum in (18) could be reduced without effecting the maximum sustainable session rate of that flow. Consequently, it may be possible to maintain the same session rate at a lower queue service rate, resulting in a lower power level requirement for a given channel state. Combining these two observations, it is possible to calculate the set of *minimum power levels* that will maintain the EAS session rates:

$$p_j^{\text{min}} = \max_{i:i \in \mathcal{J}_j} \left\{ \frac{\lambda_i^{\text{EAS}}}{\beta(s_j) \cdot w_{ij}} \right\}, \quad j = 1, 2, \dots, M \quad (19)$$

Observe that  $p_j^{\text{min}}$  is always less than or equal to  $p_j^{\text{EAS}}$  for all spot-beam queues. Therefore, the power difference between the two levels can be re-distributed to other spot-beam queues to possibly improve the session rates of active flows. The power vector with power levels given as in (19) is denoted by  $\mathbf{p}^{\text{MIN}}$ .

We refer to the solution to (8) as the *balanced-antenna-share* (BAS) policy and denote it by the vector  $\mathbf{p}^{\text{BAS}}$ . We consider two different solutions to (8). In the first case, which we will refer to as the BAS-I policy, the lower bounds on the power levels are set to zero for all beams, i.e.  $l_j = 0$ ,  $\forall j$ . The total available system power is re-distributed among the spot-beam queues. This setting allows some queues to get power level assignments that are lower than the minimum power levels given in (19). Consequently, some flows may be served at rates lower than their

rates under the EAS policy. In this case, the fairness of the algorithm becomes an issue, since it does not have control over which flow rates are reduced as a result of the optimization performed. In the second case, which we will refer to as the BAS-II policy, we set lower bounds on the power levels, such that no flow gets a lower sustainable session rate than their rates under the EAS policy, i.e. set  $l_j = p_j^{\min} \forall j$ . This choice, however, restricts the optimization space since only the excess power could be distributed, but guarantees that the optimization will return rates that are no worse than the EAS policy rates for all active flows. We will elaborate more on the fairness and the effects of having lower bounds in Section 6. In the next two sections, we present the solution under both policies.

#### 4.1. Solution under BAS-I policy

Before proceeding with the solution, we classify spot-beam queues into three sets: (i)  $\mathcal{E}$ , the set of empty queues for which  $\mathcal{B}_j = \emptyset$ , (ii)  $\mathcal{U}$ , the set of spot-beam queues with only unicast flows, and (iii)  $\mathcal{W}^c$ , the set of beam queues with both unicast and multicast flows. Based on this classification, the solution power vector can be re-arranged, without loss of generality, as

$$\mathbf{p}^{\text{BAS-I}} = [\mathbf{p}_{\mathcal{E}}^{\text{BAS-I}} | \mathbf{p}_{\mathcal{U}}^{\text{BAS-I}} | \mathbf{p}_{\mathcal{W}^c}^{\text{BAS-I}}]^T \quad (20)$$

Under BAS-I policy, empty spot-beam queues are removed from the calculation by setting  $p_j = 0, \forall j \in \mathcal{E}$ . The queues with only unicast flows are excluded from the calculations as well, because, independent of their service rates, the unicast flows that are forwarded to such queues will have zero rate variance. Therefore, we assign minimum power levels for such queues in order to guarantee that the session rates of the unicast flows are no worse than the EAS session rates in unicast only queues, and set  $p_j = p_j^{\min}, \forall j \in \mathcal{U}$ .

Having determined the power levels for the first two components of the solution vector, where  $\mathbf{p}_{\mathcal{E}}^{\text{BAS-I}} = \mathbf{0}$  and  $\mathbf{p}_{\mathcal{U}}^{\text{BAS-I}} = \mathbf{p}_{\mathcal{U}}^{\text{MIN}}$ , the values for the power vector  $\mathbf{p}_{\mathcal{W}^c}^{\text{BAS-I}}$ , of cardinality  $|\mathcal{W}^c|$  can be calculated using a Lagrangian formulation as

$$\mathbf{p}_{\mathcal{W}^c}^{\text{BAS-I}} = \mathbf{G} \cdot \boldsymbol{\alpha} + \mathbf{Z} \cdot \mathbf{d} \quad (21)$$

where

$$\mathbf{G} = \mathbf{X}^{-1} - \mathbf{X}^{-1} \cdot \mathbf{B}^T \cdot (\mathbf{B} \cdot \mathbf{X}^{-1} \cdot \mathbf{B}^T)^{-1} \cdot \mathbf{B} \cdot \mathbf{X}^{-1} \quad (22)$$

$$\mathbf{Z} = \mathbf{X}^{-1} \cdot \mathbf{B}^T \cdot (\mathbf{B} \cdot \mathbf{X}^{-1} \cdot \mathbf{B}^T)^{-1} \quad (23)$$

and  $\mathbf{X}$  is a  $|\mathcal{W}^c| \times |\mathcal{W}^c|$  matrix,  $\mathbf{B}$  is a  $L \times |\mathcal{W}^c|$  matrix,  $\mathbf{d}$  is a  $L \times 1$  vector, and  $\boldsymbol{\alpha}$  is a  $|\mathcal{W}^c| \times 1$  vector of Lagrangian multipliers.

The matrix  $\mathbf{X}$  is given by  $(\mathbf{A} - 2 \cdot \mathbf{V}^T \cdot \mathbf{V})$ , where  $\mathbf{A}$  is a  $|\mathcal{W}^c| \times |\mathcal{W}^c|$  diagonal matrix with entries

$$a_{jj} = \sum_{i=1}^N \frac{2}{N_i} \cdot \beta(s_j)^2 \cdot w_{ij}^2 \quad \forall j \in \mathcal{W}^c \quad (24)$$

and  $\mathbf{V}$  is a  $N \times |\mathcal{W}^c|$  matrix with entries

$$v_{ij} = \frac{1}{N_i} \cdot \beta(s_j) \cdot w_{ij} \quad \forall j \in \mathcal{W}^c \quad (25)$$



and  $i = 1, 2, \dots, N$ . The entries of the matrix  $\mathbf{B}$  represents the mapping of spot-beam queues to antenna groups and given by

$$b_{ij} = \begin{cases} 1 & \text{if } j \in \mathcal{A}_l \\ 0 & \text{if } j \notin \mathcal{A}_l \end{cases} \quad \forall j \in \mathcal{U}^c \quad (26)$$

and  $l = 1, 2, \dots, L$ . The vector  $\mathbf{d}$  represents the remaining power available for distribution to the spot-beam queues in set  $\mathcal{U}^c$  following the power assignments to queues in set  $\mathcal{U}$ , and given by

$$d_l = P_{\text{tot}} - \sum_{j \in (\mathcal{U} \cap \mathcal{A}_l)} p_j^{\text{min}} \quad (27)$$

for  $l = 1, 2, \dots, L$ . In the solution of (21), a non-zero Lagrangian multiplier implies that the corresponding power level must be zero, and we have  $p_j$  strictly greater than zero when the multiplier  $\alpha_j$  vanishes in (21).

The service rate vector  $\mathbf{r}^{\text{BAS-I}}$  is determined by  $\mathbf{p}^{\text{BAS-I}}$  and the channel state vector as

$$r_j^{\text{BAS-I}} = \min(\beta(s_j) \cdot p_j^{\text{BAS-I}}, r_{\text{max}}) \quad (28)$$

for  $j = 1, 2, \dots, M$ .

#### 4.2. Solution under BAS-II policy

Following a classification similar to that of the previous case, we assign  $\mathbf{p}_\delta^{\text{BAS-II}} = \mathbf{0}$ , and  $\mathbf{p}_\psi^{\text{BAS-II}} = \mathbf{p}_\psi^{\text{MIN}}$ . Under BAS-II policy, all queues that are not empty are guaranteed the minimum power level assignment given in (19). Therefore, we start the solution with a base power vector,  $\mathbf{p}^{\text{BASE}}$ , such that

$$\mathbf{p}^{\text{BASE}} = [\mathbf{0} | \mathbf{p}_\psi^{\text{MIN}} | \mathbf{p}_\psi^{\text{MIN}}]^T \quad (29)$$

Total power used by each subgroup of spot-beam queues under this assignment is given by

$$d_l^{\text{BASE}} = \sum_{j \in \mathcal{A}_l} p_j^{\text{BASE}} \quad (30)$$

and the remaining power for redistribution is equal to

$$d_l = P_{\text{tot}} - d_l^{\text{BASE}} \quad (31)$$

for  $l = 1, 2, \dots, L$ . Therefore, if  $\exists l$  such that  $d_l = 0$ , there remains no additional power to distribute to spot-beam queues in  $\mathcal{A}_l$  while satisfying the minimum requirement, and those queues will have to remain at their minimum power level assignments. Based on this observation, we can further classify the queues in  $\mathcal{U}^c$  as those in set

$$\mathcal{H} \triangleq \bigcup_l \{j : j \in (\mathcal{U}^c \cap \mathcal{A}_l) \text{ and } d_l > 0\} \quad (32)$$

and its complement,  $\mathcal{H}^c$  such that  $\mathcal{U}^c = \mathcal{H} \cup \mathcal{H}^c$ . Also, let  $\mathcal{L}$  be defined as

$$\mathcal{L} \triangleq \{l : d_l > 0\} \quad (33)$$

Then, the power level assignments for  $\mathbf{p}_{\mathcal{H}}^{\text{BAS-II}} = \mathbf{p}_{\mathcal{H}}^{\text{MIN}}$ , and the power level assignments for  $\mathbf{p}_{\mathcal{H}}^{\text{BAS-II}}$  of cardinality  $|\mathcal{H}|$  is given by

$$\mathbf{p}_{\mathcal{H}}^{\text{BAS-II}} = \mathbf{p}_{\mathcal{H}}^{\text{MIN}} + \mathbf{G} \cdot \boldsymbol{\alpha} + \mathbf{Z} \cdot \mathbf{d} \quad (34)$$

where  $\mathbf{G}$  and  $\mathbf{Z}$  are as defined in (22) and (23), respectively.  $\mathbf{X}$  is a  $|\mathcal{H}| \times |\mathcal{H}|$  matrix,  $\mathbf{B}$  is a  $|\mathcal{L}| \times |\mathcal{H}|$  matrix,  $\mathbf{d}$  is a  $|\mathcal{L}| \times 1$  vector, and  $\boldsymbol{\alpha}$  is a  $|\mathcal{H}| \times 1$  vector of Lagrangian multipliers. Equation (34) represents the power that can be distributed in addition to minimum power assignments.

The matrix  $\mathbf{X}$  is given by  $(\mathbf{A} - 2 \cdot \mathbf{V}^T \cdot \mathbf{V})$ , where  $\mathbf{A}$  is a  $|\mathcal{H}| \times |\mathcal{H}|$  diagonal matrix with entries,

$$a_{ij} = \sum_{i=1}^N \frac{2}{N_i} \cdot \beta(s_j)^2 \cdot w_{ij}^2 \quad \forall j \in \mathcal{H} \quad (35)$$

and  $\mathbf{V}$  is a  $N \times |\mathcal{H}|$  matrix with entries,

$$v_{ij} = \frac{1}{N_i} \cdot \beta(s_j) \cdot w_{ij} \quad \forall j \in \mathcal{H} \quad (36)$$

and  $i = 1, 2, \dots, N$ . The entries of the matrix  $\mathbf{B}$  represents the mapping of spot-beam queues to antenna groups and are given by

$$b_{ij} = \begin{cases} 1 & \text{if } j \in \mathcal{A}_l \\ 0 & \text{if } j \notin \mathcal{A}_l \end{cases} \quad \forall l \in \mathcal{L}, \forall j \in \mathcal{H} \quad (37)$$

The vector  $\mathbf{d}$  represents the remaining power available for distribution to the spot-beam queues in set  $\mathcal{H}$  following the minimum power assignments to the queues and is given by (31) for all  $l \in \mathcal{L}$ . In the solution of (34), a non-zero Lagrangian multiplier implies that the corresponding power level must remain at the minimum power level, and we have  $p_i > p_i^{\text{min}}$  when the multiplier  $\alpha_j$  vanishes in (34). The final solution vector is represented as

$$\mathbf{p}^{\text{BAS-II}} = [\mathbf{0} | \mathbf{p}_{\mathcal{H}}^{\text{MIN}} | \mathbf{p}_{\mathcal{H}}^{\text{BAS-II}} | \mathbf{p}_{\mathcal{H}'}^{\text{BAS-II}}]^T \quad (38)$$

The service rate vector  $\mathbf{r}^{\text{BAS-II}}$  is determined by  $\mathbf{p}^{\text{BAS-II}}$  and the channel state vector as

$$r_j^{\text{BAS-II}} = \min(\beta(s_j) \cdot p_j^{\text{BAS-II}}, r_{\text{max}}) \quad (39)$$

for  $j = 1, 2, \dots, M$ .

In the next section, we describe our analysis framework for evaluating the effectiveness of this approach.

## 5. EVALUATION

In order to evaluate the effectiveness of our approach, we first have to define several components that directly affect its performance. The first component is the *rate-power curve* that determines the rate that achieves a target BER, given the allocated power level and the channel state. The next component is the *channel model* that the channel states are based up on. In order to realistically reflect the distribution of flows across spot-beam queues and to determine the queue-antenna mappings, we have to describe the *spot-beam configuration* of our architecture. Lastly, we have to determine the *rate-allocation policy* between the unicast and multicast

flows that share the same spot-beam queue. The following sections describe these components in detail.

### 5.1. The rate-power curve

The rate-power curve is based on the following link power-budget calculation adapted from an application for a commercial satellite system [2, 13, 14]. For a given transmit power  $P_t$  in decibel Watts (dBW), the equivalent isotropically radiated power (EIRP) for the antenna system in dBW is given by

$$\text{EIRP} = P_t + G_t - L_t \quad (40)$$

where  $G_t$  and  $L_t$  are the antenna gain, and the losses in the transmitting equipment in decibels (dB), respectively. The losses due to signal propagation through the atmosphere and rain attenuation are calculated as

$$L_o = L_p + L_r \quad (41)$$

where  $L_p$  and  $L_r$  are the losses due to propagation, and rain attenuation, respectively, both in dB. Then, the ratio of signal power to noise power spectral density in decibel Hertz (dBHz) follows as

$$C/N_o = \text{EIRP} - L_o + G/T - k \quad (42)$$

where  $G/T$  in decibels per Kelvin (dB/K) is called the *figure of merit* of the receiver determined by the antenna gain  $G$  (dB) and its overall noise temperature  $T$  in Kelvin (K), and  $k$  is the Boltzmann constant in dBW/K/Hz. For a bit rate of  $R_b$  in dBHz, the ratio of bit energy to noise power density becomes

$$E_b/N_o = C/N_o - R_b \quad \text{in dB} \quad (43)$$

The rain attenuation becomes substantial at Ka-band frequencies, and is the most important factor. Therefore, we assume that—all other effects remaining constant—we can express the rate as a function of the transmit power  $P_t$  and the rain attenuation level  $L_r$  for a given  $E_b/N_o$  value that guarantees a target BER for a given coding and modulation scheme. Consequently, one can rewrite (43), to determine the rate that achieves the target BER for a given power and rain attenuation level:

$$R_b = P_t + \beta(L_r) \quad (44)$$

where  $\beta(L_r) = G_t - L_t - L_p + G/T - k - E_b/N_o - L_r$ . It is possible to express (44) in linear terms:

$$R_b = \beta(L_r) \cdot P_t \quad \text{in bps} \quad (45)$$

We will use (45) in calculating the rate-power relationship for a given rain attenuation level of the channel. Figure 4 shows rate-power relationship for different levels of rain attenuation. In this paper, we assume that rate is a continuous function of power, even though, in real systems, all rates may not be achievable depending on the set of modulation and coding schemes available for implementation. The numerical values for the link-budget parameters are given in Table I.




Figure 4. Rate-power curves for different rain attenuation levels.

Table I. Numerical values for link-budget parameters as taken from Reference [14].

$G_t$ (dB)	$L_t$ (dB)	$L_p$ (dB)
46.50	0.50	210.75
$G/T$ (dB/K)	$k$ (dB/K/Hz)	$E_s/N_0$ (dB)
16.37	-228.60	3.56

### 5.2. Channel model

The choice of the frequency band is not restrictive for our problem setting, but we believe that, next generation systems are moving in the direction of using higher frequency bands, because higher bands offer wider bandwidth segments that are not available at more crowded lower frequency bands. Therefore, we use a Ka-band channel model in our evaluations.

In order to determine the rain attenuation levels for the Ka-band channel, we use a model that is based on the simulator developed at DLR (German Aerospace Center), Institute for Communications and Navigation [15,16]. The model is based on specific channel model parameters from the DLR measurement campaign carried out at Oberpfaffenhofen near Munich, Germany, in the years 1994 till 1997 with the 40 GHz beacon of the Italian satellite ITALSAT. The channel simulator generates a time-series of attenuation, and calculates the cumulative distribution of attenuation. It is also possible to extract the probability of being in a fade exceeding a given duration and exceeding a fading depth given as parameter.




Figure 5. A sample attenuation time series and the cumulative distribution function of rain attenuation.

The simulator generates a time-series with 64 s resolution. Each attenuation level sample in decibels is input to (44), which through the link-budget calculation gives the downlink rate as a function of allocated antenna power. Figure 5 shows a sample realization of the rain attenuation time series and the corresponding cumulative distribution function for the channel model simulator.

### 5.3. Beam and antenna configuration

In order to evaluate the performance of our approach, we need to create unicast and multicast flows between the NOC and the spot-beam locations. However, the number of unicast and multicast flows forwarded to each spot-beam location and the distribution of the multicast users across these locations should reflect the possible load imbalance in a real multiple spot-beam satellite system. Therefore, we first consider the beam locations and the antenna assignments of a geostationary satellite proposed for a commercial satellite system [13, 14]. Figure 6 shows the approximate locations of the  $M = 48$  spot-beams in two polarizations over the U.S. for this system as indicated by 24 circles. In each circle, the upper and lower identifiers denote the left- and right-polarized spot-beam signals, respectively.

These 48 spot-beams share the access to  $K = 4$  on-board antennas. The antenna groups are as shown in Figure 7. Spot-beams operating at the same frequency band are served by the same antenna. Next, based on the approximate geographical area covered by each spot-beam, we have calculated the approximate population illuminated by each spot-beam, using the most recent U.S. Census Data [17]. Assuming that a flow  $f_i$  is more likely to be forwarded to

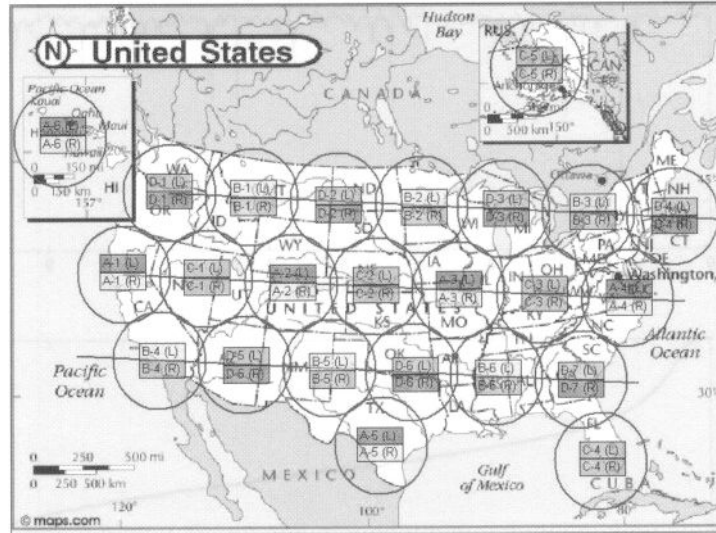


Figure 6. Locations of the 48 beams in two polarizations over the U.S. for the example satellite system.

Uplink	Low (GHz)	29.500	29.620	29.740	29.860
Frequencies	High (GHz)	29.620	29.740	29.860	29.980
Downlink	Low (GHz)	19.700	19.820	19.940	20.060
Frequencies	High (GHz)	19.820	19.940	20.060	20.180
		A 1-R	C 1-L	C 1-R	A 1-L
		A 2-R	C 2-L	C 2-R	A 2-L
		A 3-R	C 3-L	C 3-R	A 3-L
		A 4-R	C 4-L	C 4-R	A 4-L
Beam Family and		A 5-R	C 5-L	C 5-R	A 5-L
Polarization of Regional		A 6-R	D 1-R	D 1-L	A 6-L
Spot-Beams		B 1-L	D 2-R	D 2-L	B 1-R
		B 2-L	D 3-R	D 3-L	B 2-R
		B 3-L	D 4-R	D 4-L	B 3-R
		B 4-L	D 5-R	D 5-L	B 4-R
		B 5-L	D 6-R	D 6-L	B 5-R
		B 6-L	D 7-R	D 7-L	B 6-R

Figure 7. Beam family and polarization of regional spot-beams. Spot-beams operating at the same frequency band are served by the same antenna.

spot-beam queue  $b_j$  if the spot-beam illuminates a larger fraction of the total population, we calculated the probability distribution plotted in Figure 8. This distribution gives the probability of a flow being forwarded to a spot-beam for all 48 spot-beams and is used to create flows between the NOC and the spot-beam locations.

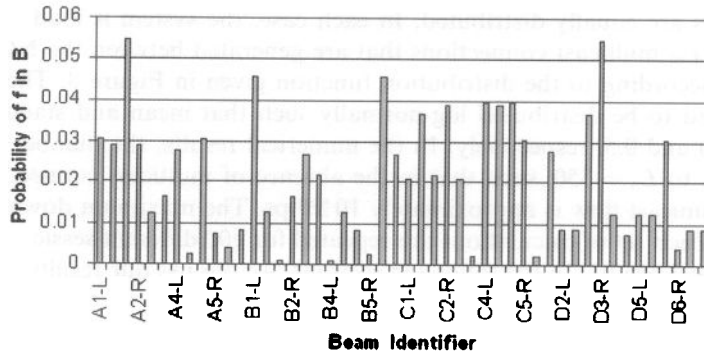


Figure 8. Connection probability distribution. This probability distribution is used to create flows between the NOC and the spot-beam locations.

#### 5.4. Rate allocation policy

Finally, we have to determine how the service rate of each spot-beam queue is shared among the unicast and multicast flows forwarded to the beam. The policy determines how multicast flows are treated compared to unicast flows sharing the same bottleneck, in this particular case, the same spot-beam queue. In order to give some incentive for multicast delivery over unicast, rate allocation proportional to the number of users downstream of the bottleneck has been suggested in the literature. In Reference [18], authors compare three policies that allocate resources: (i) as a logarithmic function of the number of users downstream of the bottleneck, (ii) as a linear function of the number of users, and (iii) as equally between unicast and multicast flows independent of the number of users. Their analysis show that logarithmic allocation achieves the best tradeoff between user satisfaction and fairness among unicast and multicast flows compared to other alternatives. In this paper, we adopt the same policy.

The rate-share  $w_{ij}$  of a flow  $f_i$  in spot-beam queue  $b_j$  is determined by  $n_{ij}$ , which is the number of receivers of the flow that resides in the area illuminated by the beam:

$$w_{ij} = \begin{cases} 0 & \text{if } n_{ij} = 0 \\ \frac{1 + \log(n_{ij})}{\sum_{i \in \mathcal{I}_j} 1 + \log(n_{ij})} & \text{if } n_{ij} \neq 0 \end{cases} \quad (46)$$

for  $i = 1, 2, \dots, N$  and  $j = 1, 2, \dots, M$ .

There are possibly other classes of strategies for allocating resources between unicast and multicast flows, but it is out of the scope of this paper. In the next section, we calculate the power levels of all spot-beam queues and the maximum sustainable rates of every flow under BAS policy and compare our results to the values under EAS policy.

## 6. RESULTS

In this section, we will present numerical results on the performance of our approach. The results on BAS policy are given in comparison to the performance under the EAS policy, i.e.

when power levels are equally distributed. In each case, the system is loaded with  $L_u$  unicast connections, and  $L_m$  multicast connections that are generated between the NOC and the spot-beam locations, according to the distribution function given in Figure 8. The multicast group size  $G_m$  is assumed to be distributed log-normally such that mean and standard deviation of  $\log(G_m)$  is  $\log(25)$  and 0.5, respectively. In the numerical results, the number of active unicast connections is set to  $L_u = 350$ , such that in the absence of multicast connections, the average session rate of a unicast flow is approximately 10 Mbps. The maximum downlink rate is set to  $r_{\max} = 92$  Mbps. Each set of calculations are repeated for 100 different session configurations to obtain the statistical results. In the following sections, we discuss our results under both BAS-I and BAS-II policies.

### 6.1. Results under BAS-I policy

In this section, we start looking at the results under BAS-I policy. The first set of results looks at the performance of the algorithm for a fix number of unicast and multicast connections while channel conditions change over time. In this scenario, there are  $L_u = 350$  active unicast connections, and  $L_m = 25$  active multicast connections. In our analysis we assume that the optimization is performed every 64 s, which is the step size of our channel simulator for generating an attenuation level sample. An independent instance of the channel simulator is run for each spot-beam. At every step, the channel states for all 48 spot-beams are sampled and power distribution levels under the BAS-I policy are re-calculated. Using the rate-power curve, the service rates of all spot-beam queues, and the sustainable session rates of all active flows are calculated for each time step. The test is performed for a simulated time of 720 min corresponding to  $T = 675$  samples. In Figure 9, we plot the per cent change in the sustainable session rates of all active flows averaged over the total number of samples, given by

$$\bar{\eta}_i^{\text{BAS-I}} = 100 \cdot \frac{1}{T} \sum_{t=1}^T \left( \frac{\lambda_i^{\text{BAS-I}}[t] - \lambda_i^{\text{EAS}}[t]}{\lambda_i^{\text{EAS}}[t]} \right) \quad (47)$$

for  $i = 1, 2, \dots, L_u + L_m$ .




Figure 9. Per cent rate change in the sustainable session rates of all active flows averaged over the test duration under BAS-I policy, compared to the session rates under EAS policy.



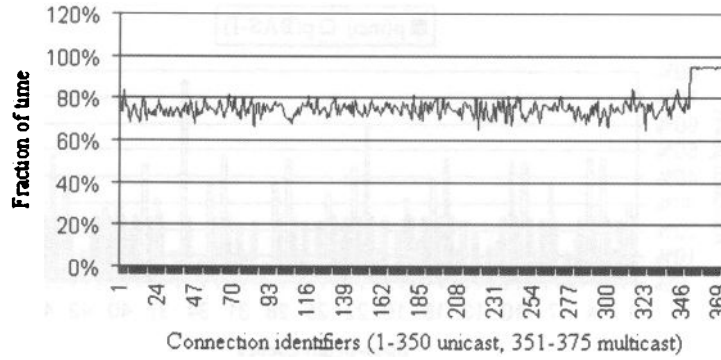


Figure 10. Per cent of total time BAS-I session rates are equal to or better than EAS session rates for all active flows.

We observe that all active multicast and unicast sessions are served at higher average session rates compared to the EAS policy case. The multicast flows experience an average increase of 95% in their sustainable session rates, while unicast flows experience average gains of 37%. Unicast flows experience a more moderate improvement compared to multicast flows. There are two factors behind this behaviour. First, the optimization policy tries to minimize the rate variance experienced by all multicast flows, without taking into account the rates of the unicast flows sharing the same queues as the multicast flows. As a result, multicast flows benefit the most from the re-arrangement of power levels across spot-beam queues. Therefore, the fairness of the BAS-I policy at a per-flow level is an issue, even though, the net system throughput is increased.

Secondly, BAS-I policy allows power levels to go to zero, therefore, the service rates of some spot-beam queues drop down to levels that are lower than their EAS rates at the end of the optimization. Consequently, the flows incident to them have lower session rates. As a result, the *instantaneous rate* of an active flow may drop down to a level lower than the EAS rate, even though the average rate of the flow remains higher than the EAS rate. Therefore, it is important to look at the percentage of total time, the sustainable session rates of all active flows remain at a level equal to or higher than their EAS session rates. In Figure 10, we look at this metric given by

$$\bar{\tau}_i^{\text{BAS-I}} = 100 \cdot \frac{1}{T} \sum_{t=1}^T \mathbf{1}(\lambda_i^{\text{BAS-I}}[t] \geq \lambda_i^{\text{EAS}}[t]) \quad (48)$$

for  $i = 1, 2, \dots, (L_u + L_m)$ , where  $\mathbf{1}(\cdot)$  is the indicator function. We observe that for unicast connections, the flow rates are below the EAS rates approximately 25% of the time, while the number is approximately 5% for multicast flows over the same duration. Therefore, for unicast flows, the instantaneous flow rates drop below EAS rates for a significant fraction of the time. From a user point of view, this fluctuation in the session rate of an active flow may not be desirable for some applications, even though the session rate is higher on the average.

In Figure 11, we plot the percentage of total power assigned to each spot-beam queue, and in Figure 12, the corresponding service rates, given by

$$\bar{p}_j^{\text{BAS-I}} = 100 \cdot \frac{1}{T} \sum_{t=1}^T p_j^{\text{BAS-I}}[t] \quad (49)$$

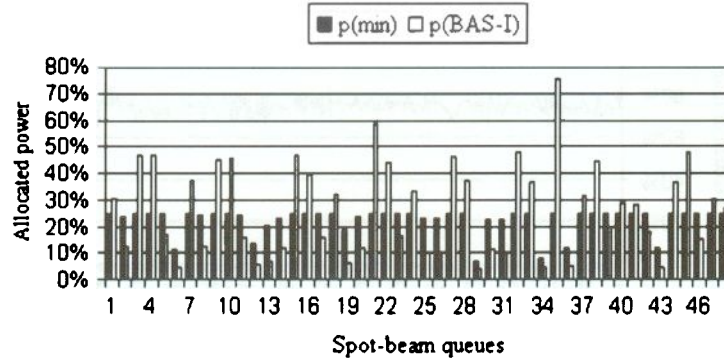


Figure 11. Average power assigned to each spot-beam queue as percentage of total system power over the test duration under BAS-I policy.

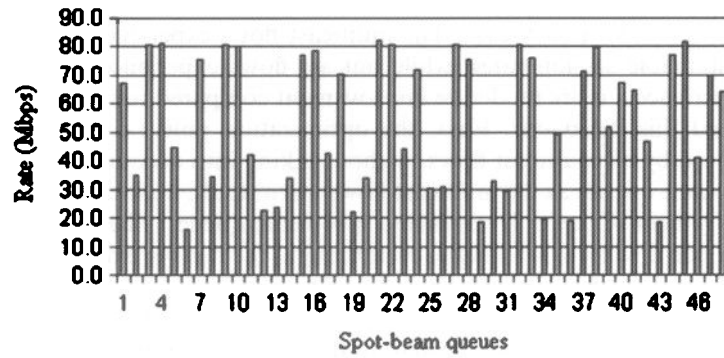


Figure 12. Average service rate of the spot beam queues over the test duration under BAS-I policy.

and

$$\bar{r}_j^{\text{BAS-I}} = \frac{1}{T} \sum_{t=1}^T r_j^{\text{BAS-I}}[t] \quad (50)$$

respectively, for  $j = 1, 2, \dots, M$ . Observe that several spot-beam queues have average power levels that are below the minimum power levels that would maintain the EAS session rates. Consequently, all unicast flows forwarded to such queues have lower session rates than their EAS rates giving rise to the behaviour we observe in Figure 10.

The second set of results look at the average performance of the BAS-I policy under changing group dynamics. In these experiments, there are  $L_u = 350$  unicast flows, while the number of active multicast flows is varied between  $L_m = 5$  and 30. In Figure 13, we plot the change in the sustainable session rates averaged over all active unicast and multicast flows, given by

$$\bar{\eta}_u^{\text{BAS-I}} = \frac{1}{L_u} \sum_{i=1}^{L_u} \bar{\eta}_i^{\text{BAS-I}} \quad (51)$$

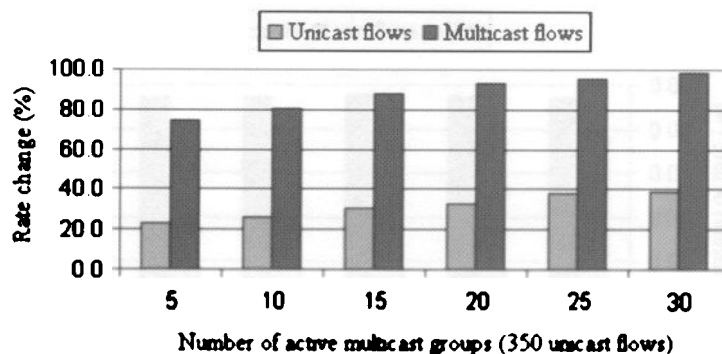


Figure 13. Average rate change in the sustainable session rates of all unicast flows averaged over all active unicast flows under the BAS-I policy and average rate change in the sustainable session rates of all multicast flows averaged over all active multicast flows under the BAS-I policy.

and

$$\bar{\eta}_m^{\text{BAS-I}} = \frac{1}{L_m} \sum_{i=L_u+1}^{(L_u+L_m)} \bar{\eta}_i^{\text{BAS-I}} \tag{52}$$

respectively, as the number of active multicast connections is varied.

We observe that unicast flows experience an increase of 20–40% in their average session rates while the number of active multicast connections is varied. The increase in the average session rate becomes more significant as the system is loaded with more active multicast flows. This is because, when there are very few multicast flows, most of the queues have only unicast connections and our optimization framework excludes such queues, i.e. such queues remain with their minimum power assignments. As the unicast and multicast flow mix increases, more queues are included in the re-distribution of power. Multicast flows experience 70–100% increase in their average sustainable session rates. Finally, in Figure 14, we plot the percentage of time flows having higher session rates than their EAS rates as the number of active multicast connections is varied.

From the results of this section, we conclude that BAS-I policy increases the average session rates for both types of flows, however, (i) the policy is unfair against unicast flows, and some unicast flows may see a decrease in their rates, and (ii) although the session rates are higher compared to EAS session rates on the average, instantaneous values may drop below the EAS values. BAS-II policy elevates these issues by imposing bounds on the power levels to guarantee that session rates remain at or above EAS policy values *at all times* for both types of flows. In the next section, we look at the performance of the BAS-II policy under similar test settings.

### 6.2. Results under BAS-II policy

In this section, we look at the same set of metrics under the BAS-II policy. In Figure 15, we plot the per cent change in the sustainable session rates of all active flows averaged over a similar test duration of 720 min corresponding to  $T = 675$  samples. Under BAS-II policy, unicast flows experience an average increase of around 17% in their sustainable session rates, while multicast flows have an increase of around 27%. We observe that, compared to BAS-I policy, the increase

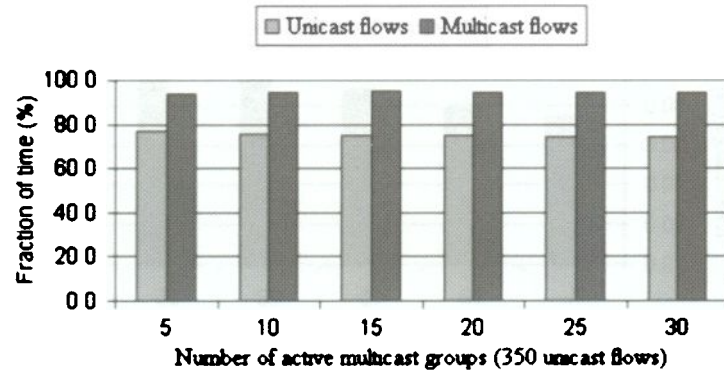


Figure 14. Per cent of total test time BAS-I session rates of unicast flows are equal to or better than EAS session rates averaged over all active unicast flows and per cent of total test time BAS-I session rates of multicast flows are equal to or better than EAS session rates averaged over all active multicast flows.

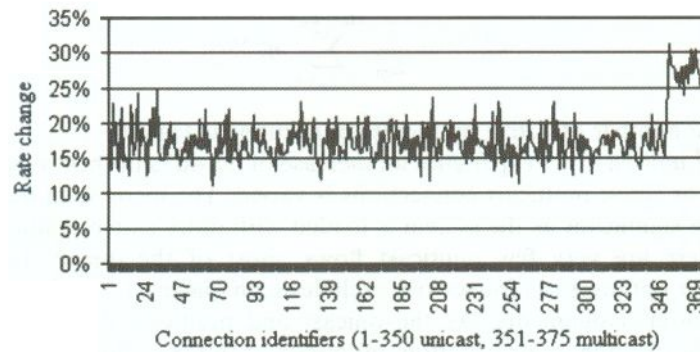


Figure 15. Per cent rate change in the sustainable rates of all active flows averaged over the test duration under BAS-II policy, compared to the session rates under EAS policy.

in the session rates of both types of flows is down, however, instantaneous session rates of all unicast and multicast flows remain above the EAS session rates at all times for all flows, i.e.  $\tau_i^{\text{BAS-II}} = 100\%$ ,  $\forall i$ .

In Figure 16, we plot the percentage of total power assigned to each spot-beam queue, and in Figure 17, the corresponding service rates. Note that under this policy, all power levels are at or over the minimum power levels required to maintain EAS policy session rates for all active flows. We observe that, the average power levels and service rates across spot-beam queues are now more uniformly distributed compared to the levels under the BAS-I policy.

The second set of results look at the average performance of the BAS-II policy under changing group dynamics. In these experiments, there are  $L_u = 350$  unicast flows, while the number of active multicast flows are varied between  $L_m = 5$  and 30. In Figure 18, we plot the change in the sustainable session rates averaged over all flows as the number of active multicast connections is varied. Under BAS-II policy, both types of flows benefit less from the optimization, however, their average rates remain 15–30% above their EAS session rates. One




Figure 16. Average power assigned to each spot-beam queue as percentage of total system power over the test duration under BAS-II policy.

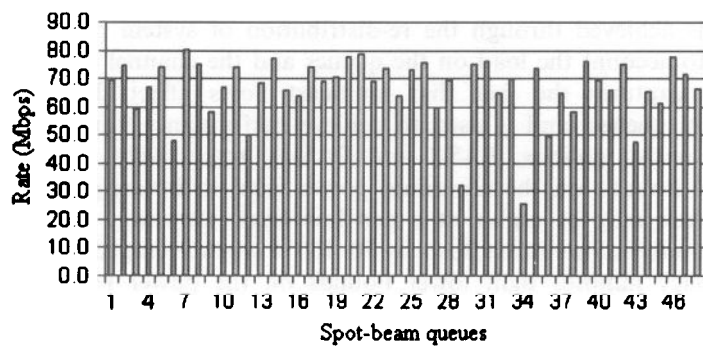


Figure 17. Average service rate of the spot beam queue over the test duration under BAS-II policy.




Figure 18. Average rate change in the sustainable session rates of all unicast flows averaged over all active unicast flows under the BAS-II policy and average rate change in the sustainable session rates of all multicast flows averaged over all active multicast flows under the BAS-II policy.

important observation to make is that this time mixing of unicast and multicast flows negatively affect the improvement of multicast flows (compare to Figure 13). This follows from the fact that when more active flows are present under BAS-II policy, the lower bound constraints restrict the optimization space, and the system has less flexibility in distributing the power. Therefore, we would expect under BAS-II policy that a very heavy loaded system would cause all flows to remain at their EAS policy session rates.

We can conclude that BAS-II policy still attains desirable performance improvements in terms of the average sustained session rates, while solving the fairness related issues of the BAS-I policy. In the next section, we sum up our observations and provide future directions on this work.

## 7. CONCLUSION

In this paper, we have introduced an optimization framework for balancing the spot-beam queue service rates such that the sum of the rate variances of all active multicast flows is minimized. This is achieved through the re-distribution of system power among spot-beam queues, taking into account the load on the queues and the channel states. The rate variance metric effectively captures the fact that multicast flows affect the load distribution of multiple spot-beam queues, and is used to achieve performance improvements from it. We provide two alternative policies, BAS-I and BAS-II, respectively. BAS-I policy does not impose any lower bounds on the minimum power level to be assigned to each spot-beam queue, and therefore may be unfair at a per-flow level. However, the policy increases the sustainable session rates of multicast flows by 70–100% when averaged over all active multicast flows. BAS-II policy imposes tight lower bounds on the power levels and, therefore, the multicast flows experience an increase of up to 30%. However, the policy also guarantees that the unicast flows do not see a performance degradation in their rates at the end of the optimization.

Depending on the application, the type of the flows, and the service rate guarantees provided by the service provider, it may not be necessary to require a strict minimum rate for all active flows. Therefore, an extension to the current policies is under study to provide a quality of service (QoS) or priority based minimum rate requirement to all active flows. In this alternative policy, a QoS level is attached to all flows to determine which flows are allowed a reduced rate, and what are the minimum rate requirements. This information is used to determine the minimum power levels for each spot-beam queue. On the overall, we conclude that it is possible to improve the performance of the system by careful tuning of system parameters to the requirements of the flows supported by it, and provide an example to the fact that future systems must be designed to support multiple types of flows.

## ACKNOWLEDGEMENTS

This material is based upon work supported by NASA under grant number NCC8235. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Aeronautics and Space Administration.

## REFERENCES

1. Gargione F, Iida T, Valdoni F, Vatalaro F. Services, technologies, and systems at Ka band and beyond: a survey. *IEEE Journal on Selected Areas in Communications* 1999; **17**(2):133–144.
2. Lutz K, Werner M, Jahn A. *Satellite Systems for Personal and Broadband Communications*. Springer: Berlin, 2000.
3. Evans JV. Communication satellite systems for high speed Internet access. *IEEE Antennas and Propagation Magazine* 2001; **43**(5):11–22.
4. Chitre P, Gupta R. Satellite systems for multimedia and Internet traffic. *IEEE Microwave Symposium Digest* 2001; **2**:20–25.
5. Akkor G, Hadjitheodosiou M, Baras JS. Transport protocols in multicast via satellite. *International Journal of Satellite Communications* 2004; **22**(6):611–627.
6. Li X, Paul S, Ammar M. Layered video multicast with retransmissions: evaluation of hierarchical rate control. *Proceedings of IEEE INFOCOM*, San Francisco, CA, U.S.A., vol. 3, March 1998; 1062–1072.
7. Vicisano L, Crowcroft J, Rizzo L. TCP-like congestion control for layered multicast data transfer. *Proceedings of IEEE INFOCOM*, San Francisco, CA, U.S.A., vol. 3, March 1998; 117–130.
8. Bhattacharyya S, Kurose JF, Towsley D, Nagarajan R. Efficient rate-controlled bulk data transfer using multiple multicast groups. *Proceedings of IEEE INFOCOM*, San Francisco, CA, U.S.A., March 1998.
9. Roca V. Packet scheduling for heterogeneous multicast transmissions. *Proceedings of P/HSN*, Salem, MA, U.S.A., August 1999; 101–116.
10. Donahoo MJ, Ammar MH, Zegura EW. Multiple channel multicast scheduling for scalable bulk-data transport. *Proceedings of IEEE INFOCOM*, New York, NY, U.S.A., vol. 2, 1999; 847–855.
11. Tunpan A, Corson MS. On satellite multicast to heterogeneous receivers. *Proceedings of IEEE ICC*, Helsinki, Finland, vol. 10, June 2001; 3150–3154.
12. Gau R-H, Haas ZJ, Krishnamachari B. On multicast flow control for heterogeneous receivers. *IEEE/ACM Transactions on Networking* 2002; **10**(1):86–101.
13. Hughes Communications Galaxy Inc. Application of Hughes Communication Galaxy, Inc. before the Federal Communications Commission for two Ka-band domestic fixed communication satellites, December 1993.
14. Fitzpatrick EJ. SPACEWAY system summary. *Space Communications* 1995; **13**:7–23.
15. Fiebig U-C. A time-series generator modelling rain fading and its seasonal and diurnal variations. *Proceedings of the 1st International Workshop of COST-Action 280*, Malvern, U.K., 2002.
16. Fiebig U-C. A time-series generator modelling rain fading. *Proceedings of the Open Symposium on Propagation and Remote Sensing*, Garmisch-Partenkirchen: URSI Commission F, 2002.
17. Population Division U.S. Census Bureau. Table NST-EST2003-01. Annual estimates of the population for the United States, and for Puerto Rico: April 1, 2000 to July 1, 2003. Release date: December 18, 2003 [Online]. Available: <http://c2k0.census.gov/popest/data/states/tables/NST-EST2003-01.php>
18. Legout A, Nonnenmacher J, Biersack EW. Bandwidth allocation policies for unicast and multicast flows. *IEEE/ACM Transactions on Networking* 2001; **9**(4):464–478.