TECHNICAL RESEARCH REPORT

Power Balancing in Multiple Spot-Beam Satellite Systems for Multicast Support

by Gun Akkor, John S. Baras, Michael Hadjitheodosiou

CSHCN TR 2004-15
(ISR TR 2004-28)

The Center for Satellite and Hybrid Communication Networks is a NASA-sponsored Commercial Space Center also supported by the Department of Defense (DOD), industry, the State of Maryland, the University of Maryland and the Institute for Systems Research. This document is a technical report in the CSHCN series originating at the University of Maryland.

Web site http://www.isr.umd.edu/CSHCN/

Power Balancing in Multiple Spot-Beam Satellite Systems for Multicast Support

Gun Akkor, John S. Baras, and Michael Hadjitheodosiou
Electrical and Computer Engineering Department, and
Center for Satellite and Hybrid Communication Networks,
University of Maryland, College Park, MD 20742, USA.
e-mail: {akkor, baras, michalis}@isr.umd.edu

***Abstract*— We address the problem of optimizing resource sharing and flow control in a multiple spot-beam broadband satellite system that supports both unicast and multicast flows. Satellite communication systems, with their wide-area coverage and ubiquitous access to large number of users, clearly have an inherent advantage in supporting distributed applications that require concurrent transmission of content to multiple users. In order to remain competitive against other broadband technologies, 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 a high load variation across the spot-beam queues may significantly under-utilize the system and be perceived unsatisfactory by potential users when both unicast and multicast flows are active in the system. We propose an optimization based-approach to balance the load in the system and conclude that it is possible to increase the average session rates of all active flows by up to 30% after this optimization is applied.**

***Index Terms*— System design, multicast delivery, satellite networks, power allocation, mathematical optimization.**

I. Introduction

The role of satellite systems 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 and on-board packet processing are currently under development [1]. These new systems will offer higher data rates and will enable the use of small, low-power, and low-cost user terminals. Therefore, they are likely to become more competitive against other broadband communication solutions 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.

Despite 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 (as perceived speed or performance) is required whether the provider uses unicast or multicast to deliver content. In order to make multicast delivery rewarding to both parties, next generation satellite systems should take into account that both unicast and multicast flows will co-exist in the system, and make sure that system resources are shared optimally between these flows. The latter issue is particularly important, since satellite bandwidth is scarce and satellite systems have to make the most out of the available resources to remain competitive against other broadband technologies.

In this paper, we address this problem from the perspective of resource sharing and flow control in a

This material is based upon work supported by NASA under award 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.

multiple spot-beam satellite system that supports both unicast and multicast flows. We show that a high load variation across the spot-beam queues may significantly under-utilize the system and decrease user satisfaction, when both unicast and multicast flows are active in the system. We propose an optimization based-approach to balance the load in the system, and in doing so, take into account that both multicast and unicast flows will co-exist and compete for the system resources.

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 III, we formulate our problem in an optimization framework. Section IV provides the solution, and Section V discusses the analysis framework we have developed for testing the performance of our approach. In Section VI, we present numerical performance results. Last section concludes the paper and draws attention to future work on this subject.

II. Motivation

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

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. This choice will later affect our channel model. The use of multiple narrow spot-beams allows satellite power to be concentrated into densely populated areas, and enables the use of low-power, low-cost user terminals that offer two-way direct communication. It also provides efficient utilization of the available satellite bandwidth by high frequency reuse across spot-beam locations. An on-board processor and switch forward packets to one or more spot-beam queues.

In this multiple spot-beam system, packets of several active flows are queued at the NOC. The NOC forwards the packets to the satellite at a rate limited by the uplink capacity of the system. The on-board processor

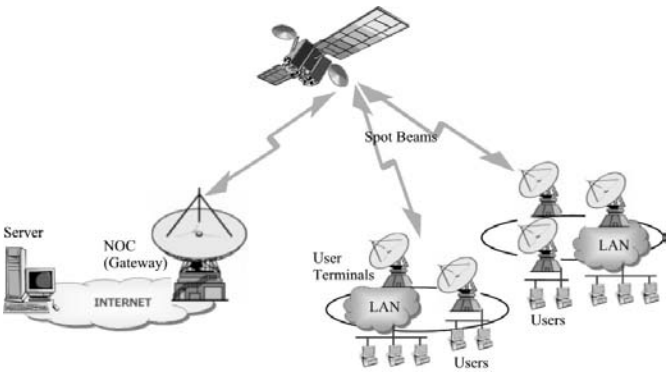

Fig. 1. Satellite communication system architecture. The satellite provides broadband access to users across multiple spot-beam locations.

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, receivers 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. Therefore, while the packets of a unicast flow affect the load on only one spot-beam queue, in case of a multicast flow, a single session may affect the load on several spot-beam queues. This may have direct implications on the rate each flow is served, as well as the user satisfaction.

At every queue, multiple flows (unicast and multicast) share the total service rate of the queue. The rate-share of a flow belonging to a particular queue depends on the number of flows currently active in the queue, the type of the flows, and the rate allocation policy between different type of flows, — i.e. unicast and multicast. In order to avoid over-flowing of any of the on-board queues, the input rate of a flow at the NOC queue have to be determined by the rate the flow can be served at the spot-beam queues. For a unicast flow, the *maximum sustainable input rate* at the NOC queue is equal to the rate-share of the flow at the spot-beam queue that it has been forwarded to. However, for a multicast flow, the maximum sustainable rate is equal to the *minimum* of the *supportable* rate-shares the flow gets across multiple spot-beam queues. This requirement would cause all receivers of a multicast session to adjust their rates to this minimum, and would negatively effect user satisfaction if there is a high variation among the supportable session

rates.

In this system, a high variation may be the result of several factors, such as the distribution of users across geographical spot-beam locations, uneven effective downlink channel rates due to climatic variations, and time of the day. In this paper, we propose an optimization-based approach for load balancing across spot-beam queues in order to minimize the rate variance multicast flows experience across multiple spot-beam queues. We show that this type of load balancing 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.

III. 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, corresponding to each on-board antenna, and $L = M/K$ columns, where we assume, without loss of generality, that L is an integer. We denote by $\mathcal{A}_l,\ l = 1, 2, \ldots, L$, 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 $b_j,\ \ j = 1, 2, \ldots, M$, 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 (BER) as a function of allocated power. The power levels for all beams satisfy:

$$0 \le p_j \le P_{tot}, \quad j = 1, 2, \ldots, M, \quad \text{and} \tag{1}$$

$$\sum_{j \in \mathcal{A}_l} p_j = P_{tot}, \quad l = 1, 2, \ldots, L, \tag{2}$$

where P_{tot} is the total available system power.

A flow f_i, for $i = 1, 2, \ldots, N$, 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

$$w_{ij} = 0 \qquad \text{if } i \notin \mathcal{B}_j, \tag{3}$$

$$0 < w_{ij} \le 1 \qquad \text{if } i \in \mathcal{B}_j, \tag{4}$$

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

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 *maximum supportable rate* of

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

at the spot-beam queue b_j. However, the *maximum sustainable 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.

$$\lambda_i = \min_{j:\ i \in \mathcal{B}_j} \{\lambda_{ij}\}, \tag{7}$$

in order to avoid overflowing of the spot-beam queues.

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

$$\mathbf{p}^* = \arg\min_{\mathbf{p}} \sum_{i=1}^{N} \sigma_i^2, \tag{8}$$

subject to constraints

$$0 \le p_j \le P_{tot} \quad j = 1, 2, \ldots M, \tag{9}$$

$$\sum_{j \in \mathcal{A}_l} p_j = P_{tot} \quad l = 1, 2, \ldots, L, \tag{10}$$

$$\text{given} \quad \mathbf{s} = [s_1 \ldots s_M], \tag{11}$$

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 service shares of every flow, — i.e. $\{w_{ij}\}$. In the next section, we will provide the solution to (8). In the remainder of this paper, the rate-power curve is of assumed to be of the form $r_j = \beta(s_j)\cdot p_j\ \forall j$. This assumption is later validated in Section V.

IV. Solution

When no load balancing is considered, the simplest assignment would be to set power levels to

$$p_j = \frac{P_{tot}}{K},\ \forall j. \quad (16)$$

We call this assignment, *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 rates for each spot-beam queue and consequently every session rate. In the remainder of this paper, EAS policy is used as the base case for comparison. The solution to (8) is referred to as the *balanced-antenna-share* (BAS) policy and is denoted by the solution vector $\mathbf{p}^{\text{BAS}}$.

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) $\mathcal{U}^c$, 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

$$\mathbf{p}^{\text{BAS}} = [\mathbf{p}_{\mathcal{E}}^{\text{BAS}}|\mathbf{p}_{\mathcal{U}}^{\text{BAS}}|\mathbf{p}_{\mathcal{U}^c}^{\text{BAS}}]^{\text{T}}. \quad (17)$$

Under the BAS policy, empty spot-beam queues are removed from the calculation by setting $p_j = 0,\ \forall j\in\mathcal{E}$. All the remaining queues have active flows, therefore, we would like to have strictly positive power assignments for all p_j such that $j\notin\mathcal{E}$. The queues with only unicast flows have to be 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 keep the EAS policy assignments for such queues, and set $p_j = p_j^{\text{EAS}},\ \forall j\in\mathcal{U}$. Having determined the power levels for the first two components of the solution vector, where $\mathbf{p}_{\mathcal{E}}^{\text{BAS}} = \mathbf{0}$ and $\mathbf{p}_{\mathcal{U}}^{\text{BAS}} = \mathbf{p}_{\mathcal{U}}^{\text{EAS}}$, the values for the power vector $\mathbf{p}_{\mathcal{U}^c}^{\text{BAS}}$, of cardinality $M_{\mathcal{U}^c}$ can be calculated as

$$\mathbf{p}_{\mathcal{U}^c}^{\text{BAS}} = \mathbf{X}^{\text{-1}}\cdot\mathbf{B}^{\text{T}}\cdot(\mathbf{B}\cdot\mathbf{X}^{\text{-1}}\cdot\mathbf{B}^{\text{T}})^{-1}\cdot\mathbf{d}, \quad (18)$$

where, $\mathbf{X}$ is a $M_{\mathcal{U}^c}$x$M_{\mathcal{U}^c}$ matrix, $\mathbf{B}$ is a Lx$M_{\mathcal{U}^c}$ matrix, and $\mathbf{d}$ is a Lx1 vector.

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

$$a_{jj} = \sum_{i=1}^{N}\frac{2}{N_i}\cdot\beta(s_j)^2\cdot w_{ij}^2, \quad (19)$$

and $\mathbf{V}$ is a Nx$M_{\mathcal{U}^c}$ matrix with entries,

$$v_{ij} = \frac{1}{N_i}\cdot\beta(s_j)\cdot w_{ij}. \quad (20)$$

The entries of the matrix $\mathbf{B}$ represents the mapping of spot-beam queues to antenna groups and given by

$$b_{lj} = \begin{cases} 1, & \text{if } j\in\mathcal{A}_l \\ 0, & \text{if } j\notin\mathcal{A}_l \end{cases} \quad (21)$$

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

$$d_l = P_{tot} - \sum_{j\in(\mathcal{U}\bigcap\mathcal{A}_l)} p_j^{\text{EAS}}, \quad (22)$$

for $l = 1, 2, \ldots, L$. The service rate vector $\mathbf{r}^{\text{BAS}}$ is determined by $\mathbf{p}^{\text{BAS}}$ and the channel state vector as

$$\mathbf{r}^{\text{BAS}} = \min(\mu(\mathbf{p}^{\text{BAS}}, \mathbf{s}), r_{\max}\cdot\mathbf{1}), \quad (23)$$

where $r_{\max}$ is the maximum system downlink rate determined by the set of applicable modulation and coding methods.

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

V. 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 bit-error-rate, 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. And, 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.

A. The rate-power curve

The rate-power curve is based on the following link power-budget calculation adapted from a commercial satellite application [1]–[3]. 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

$$\text{EIRP} = P_t + G_t - L_t, \tag{24}$$

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

$$L_o = L_p + L_r, \tag{25}$$

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

$$C/N_o = \text{EIRP} - L_o + G/T - k, \tag{26}$$

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

$$E_b/N_o = C/N_o - R_b \text{ in dB}. \tag{27}$$

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 (27), to determine the rate that achieves the target BER for a given power and rain attenuation level:

$$R_b = P_t + \beta(L_r), \tag{28}$$

where $\beta(L_r) = G_t - L_t - L_p + G/T - k - E_b/N_o - L_r$. It is possible to express (28) in linear terms:

$$R_b = \beta(L_r) \cdot P_t \text{ in bps}. \tag{29}$$

We will use (29) in calculating the rate-power relationship per rain attenuation level of the channel. Fig. 2 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, not all rates are achievable depending on the set of modulation and coding schemes available for implementation. The numerical values for link-budget parameters are given in Table I.

G_t	L_t	L_p
46.50	0.50	210.75
G/T	k	E_s/N_o
16.37	-228.60	3.56

TABLE I
NUMERICAL VALUES FOR LINK-BUDGET PARAMETERS TAKEN FROM [2]

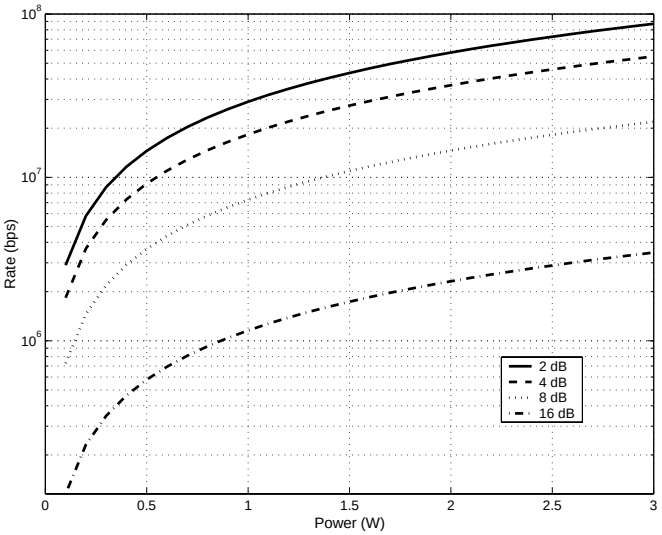

Fig. 2. Rate-power curves for different rain attenuation levels.

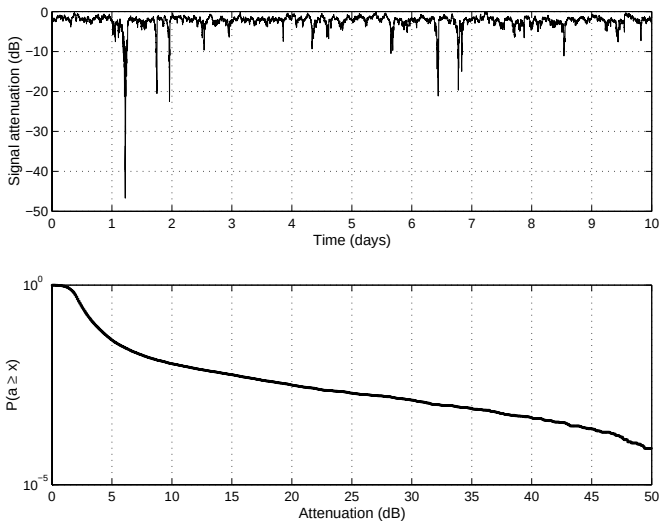

Fig. 3. A sample attenuation time series and the cumulative distribution function of rain attenuation.

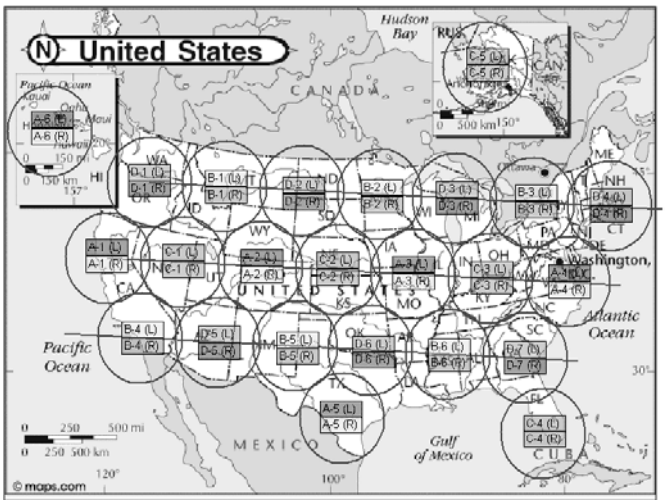

Fig. 4. Locations of the 48 beams in two polarizations over the United States for the satellite system

B. Channel model

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 [4]. 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. The simulator generates a time-series with 68 seconds resolution. Each attenuation level sample in decibels is input to (28), which through the link-budget calculation gives the downlink rate as a function of allocated antenna power. Fig. 3 shows a sample realization of the rain attenuation time series and the corresponding cumulative distribution function for the channel model simulator.

C. 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 the unicast and multicast flows forwarded to each spot-beam location and the distribution of 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 the commercial satellite system described in [2], [3]. Fig. 4 shows the approximate locations of the $M = 48$ spot-beams in two polarizations over the United States 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.

This 48 spot-beams share the access to $K = 4$ on-board antennas. The antenna assignments are as shown in Table II. The alpha-numeric identifiers used in Table II denotes the spot-beam locations and polarizations on Fig. 4. 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 [5]. Assuming that a flow f_i is more likely to be forwarded to spot-beam queue b_j if the beam illuminates a larger fraction of the total population, we calculated the probability distribution plotted in Fig. 5. 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.

D. 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 [6], authors propose a policy that allocates resources as a logarithmic function of the number of users downstream of the

ANT1	ANT2	ANT3	ANT4
D1-L	D1-R	B1-L	B1-R
D2-L	D2-R	B2-L	B2-R
D3-L	D3-R	B3-L	B3-R
D4-L	D4-R	B4-L	B4-R
D5-L	D5-R	B5-L	B5-R
D6-L	D6-R	B6-L	B6-R
D7-L	D7-R	A1-R	C1-L
C1-R	A1-L	A2-R	C2-L
C2-R	A2-L	A3-R	C3-L
C3-R	A3-L	A4-R	C4-L
C4-R	A4-L	A5-R	C5-L
C5-R	A5-L	A6-R	A6-L

TABLE II

SPOT-BEAM VS ANTENNA ASSIGNMENTS

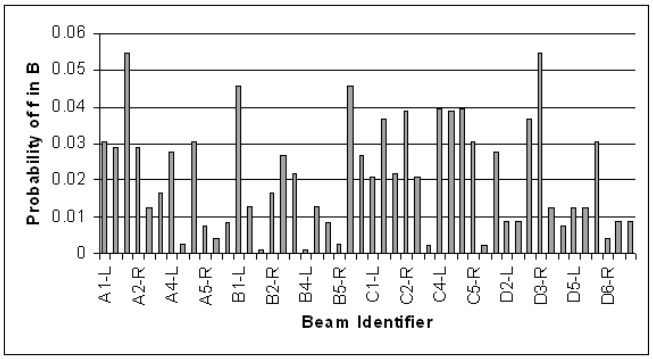

Fig. 5. Connection probability distribution

bottleneck, and show that it achieves the best tradeoff between user satisfaction and fairness among unicast and multicast flows. 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:

$$w_{ij} = \begin{cases} 0, & \text{if } n_{ij} = 0 \\ \frac{1+\log(n_{ij})}{\sum_{i \in \mathcal{B}_j} 1+\log(n_{ij})}, & \text{if } n_{ij} \neq 0 \end{cases} \quad (30)$$

In the next section, we calculate the optimal 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.

VI. 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 according to the distribution function given in Fig. 5 between the NOC and spot-beam locations. The multicast group size is distributed log-normally with mean group size $G = \log(25)$ users and standard deviation 0.5. The maximum downlink rate is $r_{\max} = 92$Mbps. In the numerical results, the number of unicast connections are kept at a fixed $L_u = 250$, while the number of multicast connections are varied. The number of unicast connections are chosen such that, in the absence of multicast connections and under perfect channel conditions, the average service rate of a unicast connection is on the order of 10Mbps/beam/connection.

The first set of results look at the performance of the algorithm for a fix number of unicast and multicast connections while channel conditions change over time. At every unit time, the channel states for all 48 spot-beams are sampled and power distribution levels are re-calculated. Table III lists the configuration parameters for this setup, and the statistical information on the channel attenuation level (A) averaged over 10000 time units. The top portion of Fig. 6 shows the average rate increase experienced by the sustainable rates of all active flows during time interval 2000 - 2120 (solid line) and the time average of average rate increase (dashed line) over the whole test duration of 10000 time units. We observe that, during this interval, the flows experience an average rate increase of 15-40% under the BAS policy compared to their rates under the EAS policy. Over the whole test duration, the time average of the rate improvement shows a 23.6% increase over all flows.

The bottom portion of Fig. 6 shows the percentage of active flows that experience a rate increase over the same interval. Note that, because of the way our cost function is constructed, at the end of the optimization, some queues might get power levels that are lower than their power levels under the EAS policy. Consequently, such queues will be served at a lower rate, causing some flows to get service rates that are lower than their rates under the EAS policy. However, as shown in Fig. 6, around 55-80% of all flows (solid line) experience a rate increase, with a time average of 70% (dashed line) over the course of the analysis.

The second set of results look at the effectiveness of the algorithm under changing multicast group dynamics. In this test, we sample the channel states of all 48 spot-beams, and then vary the number of multicast groups that are active in the system and the configuration of the flows. For each value of L_m, 1000 different flow configurations are generated and results are averaged.

L_u	L_m	G	$P_{tot}(W)$
250	25	34.16	15
$\max(A)$ (dB)	$\min(A)$ (dB)	avg(A) (dB)	std(A) (dB)
12.07	0.42	2.35	2.27

TABLE III

CONFIGURATION PARAMETERS FOR TIME ANALYSIS

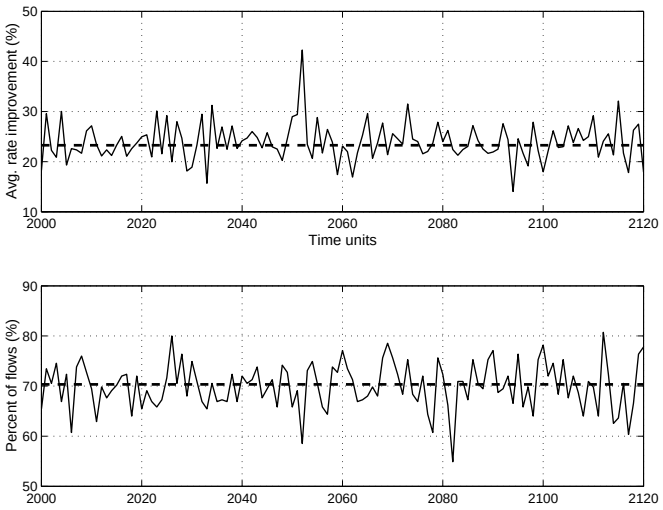

Fig. 6. (a) Average rate increase experienced by the sustainable rates of all active flows at the NOC under BAS policy compared to the sustainable rates under EAS policy, (b) Percentage of flows experiencing a rate increase over time.

The configuration parameters for this case are given in Table IV, together with statistical information on the channel attenuation.

In Fig. 7, we plot the change in the sustainable rates of all active flows as the number of multicast groups is varied. We observe that, on the average, active flows experience a rate increase of 25-30%. However, we see that there is a high variation in the individual experiences of flows. While some flows experience a rate increase of up to 125%, there are also flows that observe a decrease in their sustainable rates up to 75%. Therefore, it is important to look at what percentage of the total flows experience a rate increase.

L_u	L_m	G	$P_{tot}(W)$
250	$10-50$	28.7	15
$\max(A)$ (dB)	$\min(A)$ (dB)	avg(A) (dB)	std(A) (dB)
6.8	0.4	2.33	1.32

TABLE IV

CONFIGURATION PARAMETERS FOR GROUP DYNAMICS ANALYSIS

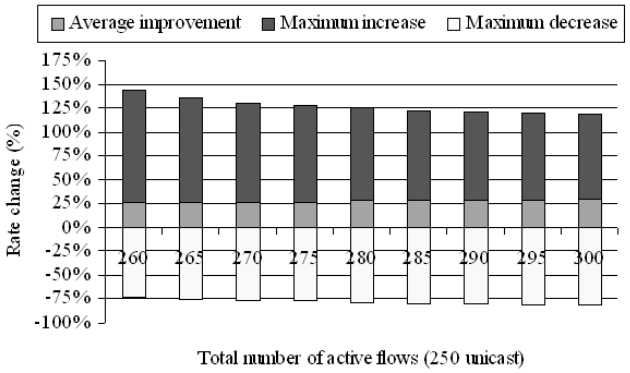

Fig. 7. Maximum rate increase, decrease, and the average rate improvement experienced by the sustainable rates of active flows at the NOC under BAS policy compared to the sustainable rates under EAS policy.

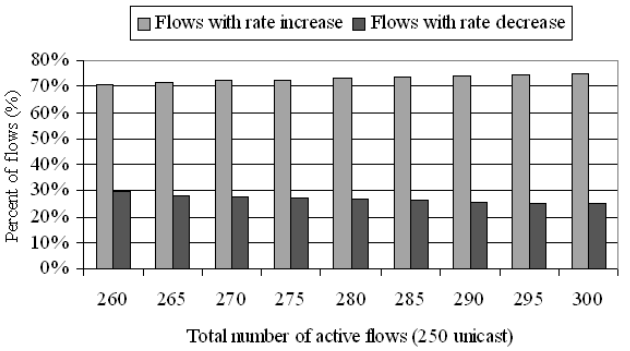

Fig. 8. Percentage of flows experiencing a rate increase or a decrease in their sustainable rates at the NOC under BAS policy compared to their sustainable rates under EAS policy.

In Fig. 8, we look at how this percentage varies with changing group dynamics. In all cases, approximately 70% of all active flows experience a rate increase, with this number increasing slightly as the number of multicast groups increase. This is because our algorithm tries to minimize the rate variance experienced by multicast flows only, while keeping the EAS shares of queues with no multicast flows. This observation is also evident from Fig. 9, which plots the number of empty, number of unicast only, and number of mixed queues over the range of test cases. Fig. 9 shows that more beams become empty or unicast only when number of multicast groups active in the system is small.

Finally, we look at the success of our algorithm in reducing the average rate variance of multicast flows. Fig. 10 plots the decrease in the overall variance for

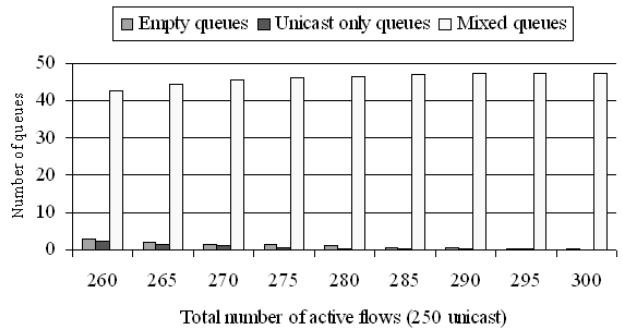

Fig. 9. Number of empty, unicast only, and mixed queues as a function of varying group dynamics.

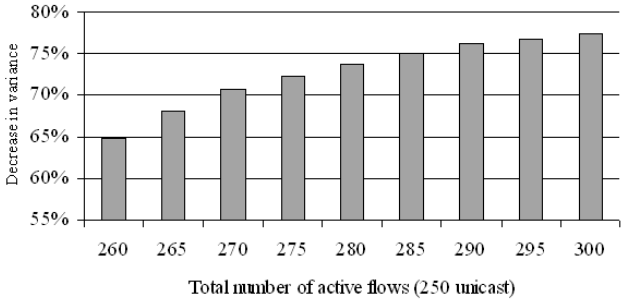

Fig. 10. Percent decrease in the average value of rate variance among all active multicast groups under BAS policy compared to the starting rate variance under EAS policy.

multicast flows. It is possible to reduce the rate variance by 65-78%. This number increases with increasing number of multicast groups, since more queues are empty or unicast only when there are only few multicast groups.

VII. 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 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 algorithm increases the sustainable session rates by up to 30% when averaged over all active flows. However, the algorithm's fairness might be an issue. While a majority of the active flows experience an increase in their sustainable session rates, some flow rates are reduced up to 75%. The algorithm does not have a control on which flow rates are reduced. Because the algorithm's main objective is to minimize the rate variances of multicast flows, most of the flows with a rate decrease are unicast. Therefore, an extension of this algorithm is under study to provide lower bounds on the power levels such that no flow gets a lower sustainable rate than the one it would get under the EAS policy. However, it is possible that this type of a lower bound may prove to be too restrictive. An alternative would be to attach priority levels to every flow to determine which flows could be forced to have a reduced rate. Finally, a more extensive study is under way to see how the rate restrictions imposed by this method would interact with the flow control mechanisms of individual transport protocols.

References

[1] E. Lutz, M. Werner, and A. Jahn, *Satellite systems for personal and broadband communications*. Springer-Verlag, 2000.

[2] 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.

[3] E. J. Fitzpatrick, "SPACEWAY system summary," *Space Communications*, vol. 13, no. 1, pp. 7–23, 1995.

[4] U.-C. Fiebig, "A time-series generator modeling rain fading and its seasonal and diurnal variations," in *1st International Workshop of COST-Action 280*, Malvern, UK, 2002.

[5] 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://eire.census.gov/popest/data/states/tables/NST-EST2003-01.php

[6] A. Legout, J. Nonnenmacher, and E. W. Biersack, "Bandwidth allocation policies for unicast and multicast flows," *IEEE/ACM Trans. Networking*, vol. 9, no. 4, pp. 464–478, August 2001.