ABSTRACT

Title of dissertation: MULTICAST COMMUNICATION SUPPORT
                           OVER SATELLITE NETWORKS

                           Gün Akkor, Doctor of Philosophy, 2005

Dissertation directed by: Professor John S. Baras
                           Department of Electrical and Computer Engineering

In this dissertation, we focus on providing multicast communication support over satellite networks. We investigate the possible performance enhancements in terms of the throughput, capacity, and scalability of a Ka-band, multiple spot-beam satellite communication system that supports unicast and multicast services. The role satellite systems play in today's communication infrastructure is changing rapidly, fueled by the technological advance in the design of new satellite systems, and by the new multimedia service applications, such as on-demand multimedia content delivery, distance learning, and distributed software updates that would benefit from the wide-area coverage, direct and ubiquitous access capability of the satellite systems.

These applications require concurrent transmission of the same content to multiple users. In order for multicasting-based services to grow over satellite networks, there must be an incentive to deploy them. We address the problem of user heterogeneity that

occurs when multicast users that are located across several different spot-beam locations experience different channel conditions. We propose a novel power allocation scheme for smoothing out the heterogeneity experienced by the multicast groups, while making sure that unicast users get a fair share of system resources as well.

Our power allocation scheme would benefit from user feedback in determining the channel conditions. However, collecting feedback from a large set of users is a challenging task in satellite systems, since access to the uplink bandwidth is to be shared between several users, and the resources are usually limited. We introduce a novel algorithm that reduces the volume of feedback information that is to be transmitted over the satellite segment of the network, while maintaining that the relevant information is collected in a timely manner.

Finally, we focus our attention to the potential benefits of integrating packet level forward error correction coding to packet delivery for reliable multicast services over satellite networks. Forward error protection helps recover corrupted data, and minimizes the need for retransmissions over the satellite channel. We investigate the use of a special form of forward error correcting (FEC) code and couple it with an adaptive control mechanism to dynamically adjust the number of encoding packets forwarded to the users.

MULTICAST COMMUNICATION SUPPORT OVER SATELLITE NETWORKS

by

Gün Akkor

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2005

Advisory Committee:

Professor John S. Baras, Chair
Professor Ashok K. Agrawala
Professor Anthony Ephremides
Professor Şennur Ulukuş
Professor Min Wu

To my parents...

ACKNOWLEDGMENTS

Manousakis, and Georgios Theodorakopoulos. Their support and collaboration have made my experience in this program a memorable one. I would like to thank everyone who in one way or the other helped me to complete my degree.

Finally, my gratitude is far less than they deserve for my wonderful parents. Without their love, endless support, and understanding, this would not be possible.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

xiv

Chapter 1

Introduction

The role satellite systems play in today's communication infrastructure is changing rapidly. Two main ingredients fuel this change. The first one is the technological advance 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 allow higher data rates, and enable the use of small, low-power, and low-cost user terminals, thereby making satellite communication systems more competitive in price and performance against other broadband communication solutions (cable, digital subscriber line (DSL)). Therefore, satellite communication systems are emerging as a solution for providing integrated voice, data, and multimedia communications to a broader range of users.

The second component is the new multimedia service applications for home and business users, such as on-demand multimedia (voice, data, video) content delivery, distance learning, distributed software updates, telemedicine, and e-commerce, 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 and providing connectivity to geographically dispersed regions. According to a report cited by Chitre and Gupta [4], communication satellites carried an estimated 10.5 Gigabit/second of Internet traffic in the year 2000. Therefore, despite recent concerns on the future of financial markets, the share of satellite communication systems in the global communication infrastructure is likely to increase as these new applications become more mature.

In order for multicasting-based services to grow over satellite networks, there must be an incentive to deploy them. From 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 service delay, is required whether the provider uses unicast or multicast to deliver the content. Therefore, it is imperative that the next generation satellite systems provide support for these services, and make sure that they perform favorably compared to existing services and applications.

In this dissertation, we investigate the possible performance enhancements in terms of the throughput, capacity and scalability of a multiple-spot beam satellite communication system that supports both unicast and multicast services. The favorable characteristics of a satellite system, such as the global reach, flexible bandwidth-on-demand, and broadcast capability, as well as the shortcomings, such as the relatively long propagation delays and the complications of a wireless channel, constitute a unique design

space for supporting multicast services. In this dissertation, we combine this design space with the advances in system design, such as on-board switching and spot-beam technology, to achieve performance improvements.

In the remainder of this chapter, we provide an overview of the relevant satellite network architectures and current system proposals. Then, we outline the organization of the remaining chapters and highlight our contributions to the literature.

1.1 Satellite systems for multimedia and Internet traffic

In the recent years, the amount of data traffic carried over the Internet has increased tremendously as the number of Internet users grew from an estimated 241.5 million in the year 2000, to a projected number of 775.7 million in the year 2005. Although, the capacity and the speed of the links, routers, and switches in the terrestrial backbone network have been keeping up with this surge, as service providers increase their capacity, access to the backbone network remains poor for most of the users. While much of the broadband access is via cable (cable modem) or over telephone lines (DSL), there is still a significant segment of users that does not have these types of access, especially in parts of the world where the terrestrial infrastructure is less developed. Therefore, there is a market opportunity for satellite service providers in expanding their offerings to support multimedia and Internet traffic.

The first company to provide Internet access via satellite was Hughes Network Systems (HNS) with their *DirecPC*® service, which since then became the *DirecWay*®
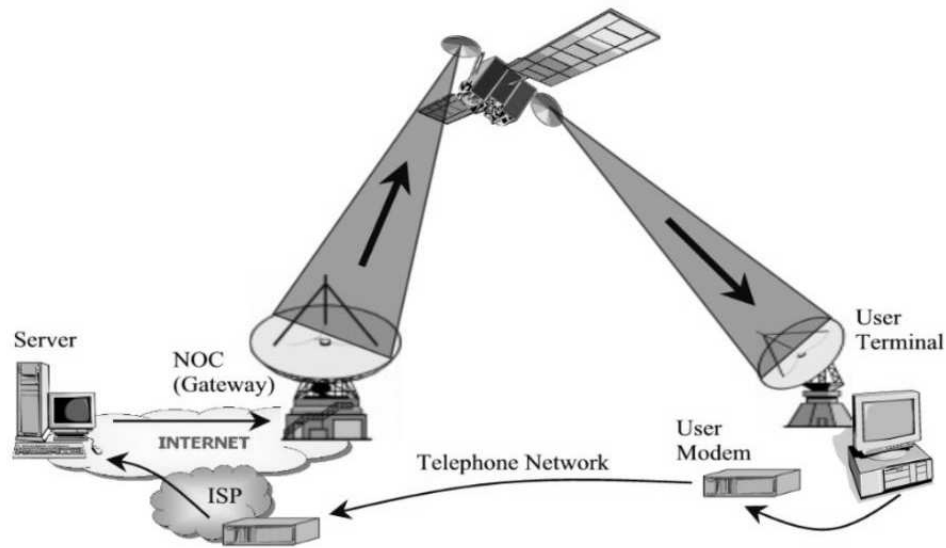
Figure 1.1: Hybrid satellite system approach of DirecPC.

system. The DirecPC service employed a hybrid satellite system, where requests for content (file download, WWW access) were made over the telephone network using a modem connection, while the responses were delivered by the DirecPC Ku-band satellite. DirecPC system required a subscription to an Internet Service Provider (ISP) for modem connection, as well as to the HNS system. A hybrid user terminal sent the requests to the Internet server over the telephone network by encapsulating them in a header that routed the response to the HNS satellite gateway rather than back to the user. The gateway was responsible for decapsulating and reformatting the traffic for satellite transmission. The traffic was sent over the satellite network directly to the user terminal (Fig. 1.1).

DirecWay system is the next-generation system that supports two-way connectivity over the Ku-band satellite. DirecWay user terminals receive and transmit data over
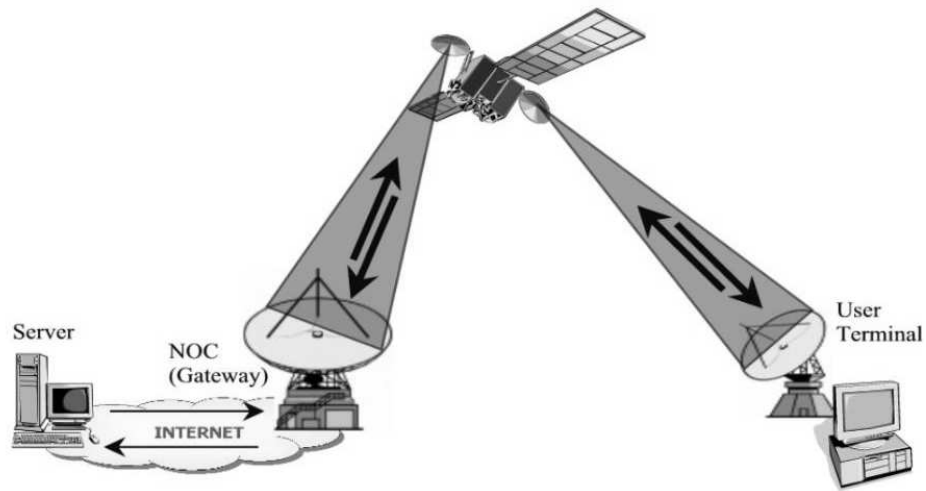
Figure 1.2: Two-way satellite system approach of DirecWay.

the satellite network and through the DIRECWAY Network Operation Center, thereby avoid the requirement for the additional connectivity through the telephone network. DirecWay system supports downlink (satellite-to-user terminal) transmission rates of up to 48 Mbps, and uplink (user terminal-to-satellite) transmission rates of 64, 128, and 256 Kbps (Fig. 1.2). DirecWay system is one of the major competitors in the market for providing Internet services over a satellite only network.

In addition, there have been several other proposals for launching new systems operating primarily at the Ku and Ka frequency bands. Because of the spectrum congestion of the C and Ku frequency bands at the geostationary orbital arc, there is no available orbital location to place a new satellite without causing interference to any of the existing systems. This have lead to proposals for low-earth-orbit (LEO) or medium-earth-orbit (MEO) satellite constellations at the Ku-band for providing global or near-global cover-

age. A list of these systems could be found in [3]. However, due to weak market conditions most of these systems are either abandoned or postponed. Besides marketing and financial concerns, the proposed LEO and MEO systems that operate at the Ku-band were also challenged by the satellite operators of geosynchronous-earth-orbit (GEO) systems that operate at the same band due to the possibility of interference to their systems. We believe that Skybridge system by Alcatel Telecommunications is the only one among the listed systems that is still being pursued. The original Skybridge system was made up of a constellation of 80 LEO satellites that provided a near-global coverage. However, currently, Skybridge services are supported by existing GEO satellites that operate at the C- and Ku-bands, and the LEO constellation is planned to be deployed subsequently if the demand for the services increases in time. Inside the footprint of the satellites, Skybridge gateways aggregate traffic to and from a large number of user terminals, and connect them to corporate networks and Internet service providers. Skybridge system supports downlink transmission rates of up to 45 Mbps, and uplink transmit rates of 64 K to 1 Msymbols/second depending on the terminal transmit power [5].

The availability of geosynchronous orbital locations and the wider frequency band segments at the Ka-band have lead to proposals for GEO Ka-band satellite systems, despite the possibility of severe signal attenuation due to rain. A list of these systems could be found in [3]. These systems have also suffered from weak market conditions and the unpredictability of the customer base, and therefore, have either gone through extensive revisions or been abandoned. Among these systems are the Spaceway system

proposed by Hughes Network Systems, and the Euroskyway system proposed by the Italian Alenia Aerospazio. Spaceway's original design comprises of a system of Ka-band satellites that provide coverage using 48 narrow spot-beams ($1^o - 3^o$) per satellite and 125 MHz of transmission bandwidth per spot-beam. Spaceway satellites support downlink transmission rates of up to 92 Mbps, and uplink transmission rates of 384 Kbps to 1.5 Mbps depending on the terminal. Moreover, they include on-board processing and buffering capabilities that allow packet switching on-board the satellite from one spot-beam to another [6]. Euroskyway system, on the other hand, is made up of a system of Ka-band satellites that support on-board switching and provide downlink transmission rates of up to 32 Mbps, and uplink transmit rates of 160 Kbps to 2 Mbps depending on the terminal type [7]. Additional information on these and similar systems may be found in [2], even though some of the technical specifications do not agree with other sources due to frequent technical revisions on the design of these systems.

Even though the availability and the commercial success of these systems are not known, there are some common traits among them that provide a template for a next-generation system. First of all, utilizing the Ka-band frequency is very desirable, because of the availability of orbital locations, and wider bandwidth segments that would enable the support for multimedia applications. Use of multiple-spot beam coverage enables satellite power to be focused on a smaller area, thereby allowing smaller, low-cost, and low-power user terminals. In order to provide high connectivity, and to support multicast applications between users that are located in different spot-beam coverage ar-

eas, however, new systems must also utilize on-board switching and buffering functions. Finally, to make the most out of the broadcast capability and the wide-area coverage of satellites, and to remain competitive against other broadband solutions, these systems must provide the support for emerging multicast services.

In this dissertation, we focus on a satellite system *template* with the characteristics outlined as above, and investigate the requirements for providing multicast support. We leave the discussion of the properties of our system to the subsequent chapters, and choose to introduce them in parallel with the relevant problems we address. In the next section, we give an overview of these problem and our contributions to the literature.

1.2 Multicast communication support over satellite networks

1.2.1 User heterogeneity in multicast communications

Utilizing the Ka-band and having multiple spot-beams bring up additional challenges in the design of systems. At Ka-band frequencies, transmitted signals are highly susceptible to degradation due to rain and other atmospheric effects. For high availability, these systems must have means for mitigating the effects of rain attenuation. Moreover, since rain events happen in areas of smaller dimension compared to the satellite footprint, users across different spot-beam locations are likely to have very different channel characteristics. These effects become more pronounced in case of multicast communication, since the performance of the communication session depends, collectively, on the performance of a group of users that are possibly located across several different

spot-beam locations and experiencing different channel conditions. Therefore, there is a need for handling this performance heterogeneity among the users of a multicast session in order to have a quality of service that is acceptable to all of them. Chapter 2 investigates this problem, and proposes a novel power allocation policy for smoothing out the heterogeneity experienced by the multicast group members, while making sure that other users (such as unicast) get a fair share of the system resources as well. Part of our results on this chapter has been presented previously. Earlier work on this subject has been presented in the AIAA International Communications Satellite Systems Conference (ICSSC 2004) [8]. Some of our recent results are included in the paper submitted to the International Journal of Satellite Communication and Network for consideration for publication.

1.2.2 Efficient feedback collection

Any system that adapts to changing propagation conditions would benefit from user feedback. However, collecting feedback from a large set of users is a challenging task in satellite communication systems, since access to the uplink bandwidth is to be shared between several users, and the resources are usually limited. Therefore, allocating a large portion of the available uplink resources to feedback collection is not a viable option. In Chapter 3, we introduce a novel algorithm that reduces the volume of feedback information that is to be transmitted over the satellite segment of the network, while maintaining that the relevant information is collected in a timely manner. We achieve

this goal by exploiting the nature of the feedback to be collected and by using the fact that most of the time, only the behavior of a subset of users needs to be tracked. The power allocation policy of Chapter 2 relies on the knowledge of channel conditions across spot-beams in determining the power levels. Therefore, it requires the collection of channel state information from the users. In this chapter, we also integrate our feedback collection algorithm into the power allocation policy of the previous chapter, and discuss their performance when the two cooperates. Part of the results of this chapter has been presented in the IEEE GLOBECOM Conference (2003) [9] and in the Conference on Information Science and Systems (CISS 2004) [10]. More recent results have been submitted to upcoming IEEE GLOBECOM Conference (2005).

1.2.3   Using FEC for reliable multicast delivery

For reliable multicast services over satellite networks, retransmissions of missing data segments are undesirable due long propagation delays and time-varying channel conditions. Therefore, primary focus on this topic has been on providing some form of forward error protection by transmitting redundant information along with the forwarded data. Forward error protection helps recover missing data segments, and thus, minimizes the need for retransmissions over the satellite network. In Chapter 4, we investigate the potential benefits of using packet level forward error correcting (FEC) codes as a part of a reliable multicast transport protocol. We discuss the limitations of fixed-block length FEC codes, and test the use of a special form of packet level FEC codes,

namely the Luby's Transform (LT) codes, which are capable of generating a large number of encoding packets for the same group of input packets. This flexibility in the number of encoding packets improves protocol's robustness against bursty periods of unfavorable channel conditions (e.g. when the transmitted signal experiences deep fading). However, use of forward error protection also reduces the bandwidth efficiency of the transmission. Therefore, we propose a scheme for adaptively adjusting the number of transmitted encoding packets to the changing propagation conditions. Part of the results of this chapter has been presented in the Conference on Information Science and Systems (CISS 2004) [10]. A more general discussion on reliable multicast transport protocols was presented in the International Journal on Satellite Communication and Networking [11].

Chapter 2

Multicast-aware Power Allocation in Multiple Spot-beam Satellite

Systems

2.1 Introduction

Unlike unicast communication, where the attributes of the communication session are adjusted according to the capabilities of a *single* user (receiver), multicast communication is significantly complicated by the widely varying capabilities of the members of the *multicast group*, in terms of available network bandwidth, supportable transmission rates, and traffic load. Heterogeneity in the performance of the receiver group becomes more pronounced in wide-area multicast, where the group members are likely to span several different network domains. This heterogeneity forces the multicast session to adjust to the capabilities of the user or the subset of users with the worst performance. Therefore, the group experiences a performance level that is below the network's capability.

There are several papers in the literature that address the problem of receiver heterogeneity in terrestrial multicast communication, based on what is identified as the source of heterogeneity, and the type of the multicast application. In terrestrial networks, receiver heterogeneity is generally attributed to the varying congestion levels

across different network domains, and to the network bandwidth available to the users behind different bottleneck links. In this case, the multicast session rate has to be set to the level supportable by the most congested link, penalizing the users with higher available bandwidth. To overcome this problem, several authors have suggested layered organization of the transmitted content [12, 13]. In layered multicast transmission, sender distributes the content using layers of increasing bandwidth, while receivers add or drop layers based on their perception of available bandwidth on the path to the sender. This allows each receiver to match its desired rate. However, layered distribution of content is effective, only if the users behind the same bottleneck link act in a coordinated manner, because the action of a receiver dropping a layer (for example, due to perceived congestion) would not be effective unless all the receivers sharing the same bottleneck link drop the same layer. Moreover, if a receiver causes congestion on the bottleneck link by adding a new layer, another might interpret the resulting losses as a consequence of its current level and drop a layer. Therefore, fairness of the protocol requires that all of them have a coordinated behavior [13]. As another point, a layered organization is only possible if the content to be delivered supports it. Layered organization is more effective for video and audio-conferencing applications, where content can be effectively organized into layers of increasing quality [14]. However, for bulk data transfer, this is not always possible. By using a proper arrangement, a similar approach could be used [15, 16], if the data can be processed beforehand. For data generated on the fly, it becomes more complex to find a suitable data organization without wasting network

capacity.

In multiple spot-beam satellite networks, receiver heterogeneity is also the result of temporal and spatial variation of the physical communication channel. Transmitted signals over the wireless channel are susceptible to fading due to atmospheric and rain attenuation. Rain fading becomes significant for systems operating at Ka-band frequencies and beyond. Spatial and temporal analysis of rainfall rates in different parts of the world show that the size of intense rain cells range from hundreds of meters to a few kilometers [17], and are of small dimension compared to the footprint of a spot-beam (650km in diameter for a 1^o narrow spot-beam of a geosynchronous earth orbit satellite). Moreover, instantaneous rainfall rates of two locations that are more than 100km apart can be considered to be independent [18]. Therefore, users residing in different geographical locations are likely to experience highly variable channel conditions.

Mitigating the effects of rain attenuation has been an area of focus in the research community. The measures against signal degradation can be grouped into two as diversity techniques and compensation techniques. Diversity techniques, such as site- and frequency-diversity, avoid signal degradation by switching between the signals obtained at different receiver sites or between different frequency bands. Site diversity takes advantage of the fact that probability of high attenuation occurring simultaneously at two or more receiver sites is significantly lower than any single site. Therefore, signals obtained at multiple sites can be combined (at some central location) to improve the signal-to-noise ratio. Frequency diversity makes use of the fact that signals suffer

more from atmospheric attenuation as the frequency of operation increases. Hence, a Ka-band system may switch to lower frequency bands (C or Ku) when the attenuation due to rain exceeds a certain threshold. Other diversity techniques include time- and orbital-diversity [19].

Compensation techniques, such as adaptive coding and modulation, transmission rate reduction, and power control, avoid signal degradation by restoring the signal quality to the initial level. Adaptive coding involves changing the amount of redundancy introduced to the transmitted information as the quality of the link changes. The probability of successful transmission increases as the redundancy level is increased, however, transmission (bandwidth) efficiency decreases at the same time. Adaptive modulation schemes decrease the required signal-to-noise ratio for achieving a target bit error rate (BER) by reducing the spectral efficiency (in bps/Hz) of the transmitted signal when fading occurs. A satellite system may switch between higher-order modulation schemes, such as 16-PSK (phase shift keying), 64-PSK, or 254-QAM (quadrature amplitude modulation), under clear sky conditions, and lower-order modulation schemes like BPSK (binary PSK) and QPSK (quadrature PSK) under deep fading. Data rate reduction, on the other hand, achieves an extra margin over the required signal-to-noise ratio by decreasing the information data rate whenever the system experiences deep fading. Finally, in adaptive power control schemes, transmitted power is adjusted dynamically based on the attenuation levels. Power control can be applied on the satellite downlink by changing the satellite effective isotropic radiated power (EIRP), or on the uplink,

by controlling the power of earth stations. Adaptive power control (APC) is more efficient than providing fixed power margins, since severe attenuation occurs typically in short durations. APC can also be easily combined with other compensation techniques if larger fade margins are required [19, 20].

While diversity and compensation techniques improve transmission quality of the users located inside the footprint of the satellite spot-beam, they do not take into account the interaction between the users, such as those belonging to the same multicast group, and the type of the communication (unicast or multicast). The techniques dealing with user heterogeneity in terrestrial multicast communication, on the other hand, focus on the requirements of the multicast communication, however, do not address the heterogeneity arising from the impairment of the physical channel. This chapter bridges the gap, and investigates the theoretical gains of a multiple spot-beam satellite system that supports both unicast and multicast flows in terms of its utilization and capacity, when adaptive power control and transmission rate reduction are integrated as tools to deal with user heterogeneity in satellite multicast communication.

We show that user heterogeneity across spot-beam queues arises not only as a result of the load variation and the types of communication sessions, but also because of the variation of the physical channel conditions. This heterogeneity results in lower allocated session rates for active flows, and may 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 controls system power with the goal of smoothing

16
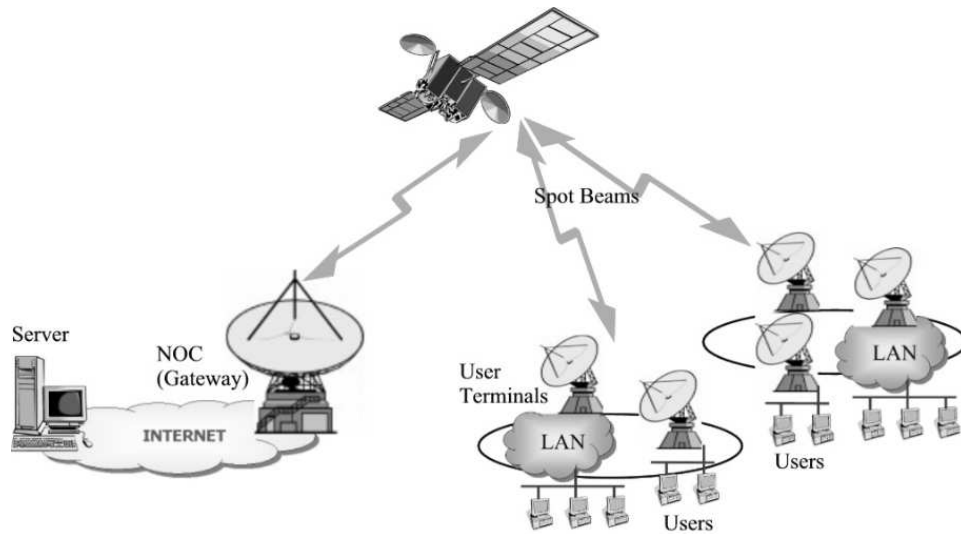
Figure 2.1: Satellite communication system architecture. The satellite provides two-way broadband access to users across multiple spot-beam locations.

user heterogeneity, and show that resulting session rates are higher on the average for both unicast and multicast flows.

2.2 Overview

In this chapter, 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 processing and switching technologies that allow transmission of data to multiple users in multiple spot-beam locations (Fig. 2.1). 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 center (NOC).

17
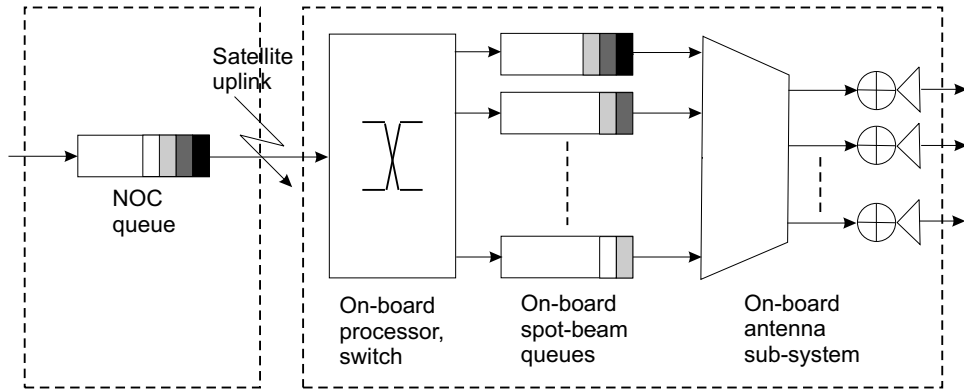
Figure 2.2: Conceptual view of on-board satellite system and queues.

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 (Fig. 2.2).

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 several supportable session rates if the user group was 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 have to be determined by the minimum rate 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 the *minimum* of its supportable session rates across (potentially) 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 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 one 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 every flow.

- 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 chapter, the goal is to minimize the heterogeneity that 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.

2.3   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
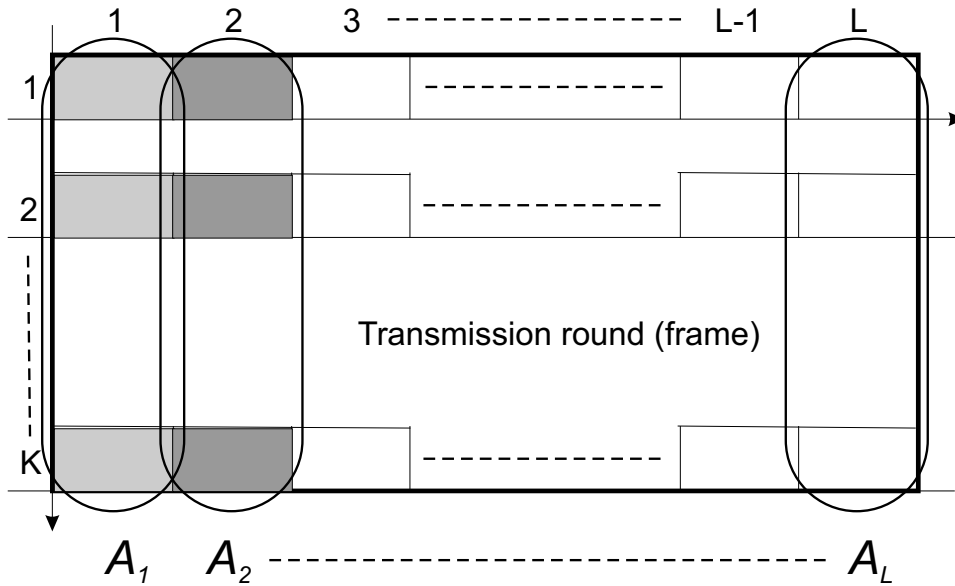
Figure 2.3: Transmission round concept.

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 $\mathcal{A}_l,\ l = 1, 2, \ldots, L$, as the set of spot-beam queues that are served simultaneously, corresponding to a column of the frame (Fig. 2.3).

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 as a function of the allocated power. The power levels of

all beams satisfy:

$$0 \leq l_j \leq p_j \leq P_{tot}, \quad j = 1, 2, \ldots, M, \quad \text{and} \qquad (2.1)$$

$$\sum_{j \in \mathcal{A}_l} p_j = P_{tot}, \quad l = 1, 2, \ldots, L, \qquad (2.2)$$

where P_{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 $i = 1, 2, \ldots, N$, 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

$$w_{ij} = 0 \qquad \text{if } i \notin \mathcal{B}_j, \qquad (2.3)$$

$$0 < w_{ij} \leq 1 \qquad \text{if } i \in \mathcal{B}_j, \qquad (2.4)$$

$$\sum_{i \in \mathcal{B}_j} w_{ij} = 1 \quad j = 1, 2, \ldots, M, \qquad (2.5)$$

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

$$\lambda_{ij} = w_{ij} \cdot r_j = w_{ij} \cdot \mu_j(p_j, s_j), \qquad (2.6)$$

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.

$$\lambda_i = \min_{j:\, i \in \mathcal{B}_j} \{\lambda_{ij}\}, \qquad (2.7)$$

in order to avoid overflowing of the spot-beam queues (Fig. 2.4).

22
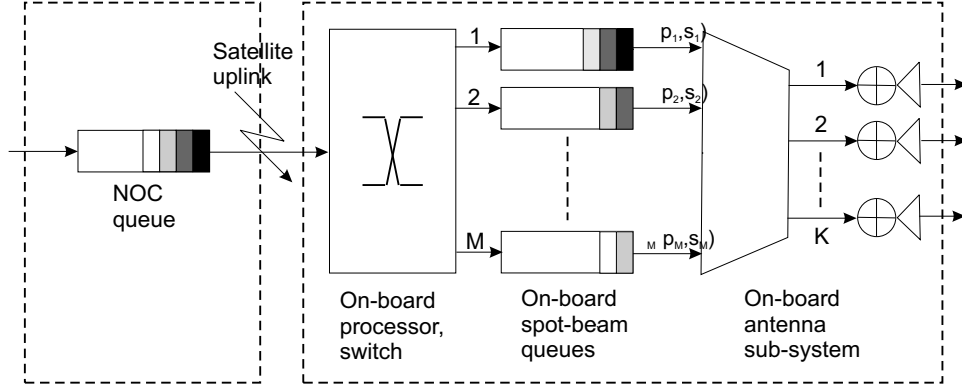
Figure 2.4: Representation of supportable session rates and maximum sustainable session rates for four active flows.

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 allocated to each queue, subject to a total power constraint, and a given set of 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 flows across spot-beam queues:

$$\mathbf{p}^* = \arg\min_{\mathbf{p}} \sum_{i=1}^{N} \sigma_i^2, \qquad (2.8)$$

subject to constraints

$$0 \leq l_j \leq p_j \leq P_{tot} \quad j = 1, 2, \ldots M, \tag{2.9}$$

$$\sum_{j \in \mathcal{A}_l} p_j = P_{tot} \quad l = 1, 2, \ldots, L, \tag{2.10}$$

$$\text{given} \quad \mathbf{s} = [s_1 \ldots s_M], \tag{2.11}$$

where,

$$\sigma_i^2 = \frac{1}{N_i} \sum_{j=1}^{M} x_{ij} \cdot (\lambda_{ij} - m_i)^2, \tag{2.12}$$

$$m_i = \frac{1}{N_i} \sum_{j=1}^{M} x_{ij} \cdot \lambda_{ij}, \tag{2.13}$$

$$x_{ij} = \begin{cases} 1, & \text{if } i \in \mathcal{B}_j \\ 0, & \text{if } i \notin \mathcal{B}_j \end{cases}, \tag{2.14}$$

$$N_i = \sum_{j=1}^{M} x_{ij}. \tag{2.15}$$

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 chapter, 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 2.5. In the next section, we provide a solution to (2.8).

2.4 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

$$p_j = \frac{P_{tot}}{K}, \ j = 1, 2, \ldots, M. \tag{2.16}$$

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:

$$r_j^{\text{EAS}} = \min\left(\beta(s_j) \cdot p_j^{\text{EAS}}, r_{max}\right), \quad j = 1, 2, \ldots, M, \tag{2.17}$$

$$\lambda_i^{\text{EAS}} = \min_{j:\, i \in \mathcal{B}_j} \left\{ w_{ij} \cdot r_j^{\text{EAS}} \right\}, \quad i = 1, 2, \ldots, N, \tag{2.18}$$

where $r_{\max}$ is the maximum system downlink rate determined by the set of available modulation and coding methods. In the remainder of this chapter, we use the EAS policy as the basis for comparison.

Equation (2.17) 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 (2.18), we can conclude that the supportable session rates of a multicast flow, which are higher than the minimum in (2.18) 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 observa-

tions, it is possible to calculate the set of *minimum power levels* that will maintain the EAS session rates:

$$p_j^{\min} = \max_{i:\, i \in \mathcal{B}_j} \left\{ \frac{\lambda_i^{\text{EAS}}}{\beta(s_j) \cdot w_{ij}} \right\}, \; j = 1, 2, \ldots, M. \tag{2.19}$$

Observe that $p_j^{\min}$ is always less than or equal to p_j^{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 (2.19) is denoted by $\mathbf{p}^{\text{MIN}}$.

We refer to the solution to (2.8) as the *balanced-antenna-share* (BAS) policy and denote it by the vector $\mathbf{p}^{\text{BAS}}$. We consider two different solutions to (2.8). 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 (2.19). 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 its rate 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

26

active flows. We will elaborate more on the fairness and the effects of having lower

bounds in Section 2.6. In the next two sections, we present the solution under both

policies.

2.4.1   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) $\mathcal{U}^c$, 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

$$\mathbf{p}^{\text{BAS-I}} = [\mathbf{p}_{\mathcal{E}}^{\text{BAS-I}} | \mathbf{p}_{\mathcal{U}}^{\text{BAS-I}} | \mathbf{p}_{\mathcal{U}^c}^{\text{BAS-I}}]^{\text{T}}. \tag{2.20}$$

Under BAS-I policy, empty spot-beam queues are removed from the calculation by set-

ting $p_j = 0, \ \forall j \in \mathcal{E}$. The queues with only unicast flows are excluded from the cal-

culations 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, 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 $\mathbf{p}_{\mathcal{U}^c}^{\text{BAS-I}}$, of

cardinality $|\mathcal{U}^c|$ can be calculated using a Lagrangian formulation as

$$\mathbf{p}_{\mathcal{U}^c}^{\text{BAS-I}} = \mathbf{G} \cdot \alpha + \mathbf{Z} \cdot \mathbf{d}, \tag{2.21}$$

where

$$\mathbf{G} \;\; = \;\; \mathbf{X}^{-1} - \mathbf{X}^{-1} \cdot \mathbf{B}^{\mathrm{T}} \cdot (\mathbf{B} \cdot \mathbf{X}^{-1} \cdot \mathbf{B}^{\mathrm{T}})^{-1} \cdot \mathbf{B} \cdot \mathbf{X}^{-1}, \qquad (2.22)$$

$$\mathbf{Z} \;\; = \;\; \mathbf{X}^{-1} \cdot \mathbf{B}^{\mathrm{T}} \cdot (\mathbf{B} \cdot \mathbf{X}^{-1} \cdot \mathbf{B}^{\mathrm{T}})^{-1}, \qquad (2.23)$$

and $\mathbf{X}$ is a $|\mathcal{U}^c|$x$|\mathcal{U}^c|$ matrix, $\mathbf{B}$ is a Lx$|\mathcal{U}^c|$ matrix, $\mathbf{d}$ is a Lx1 vector, and α is a $|\mathcal{U}^c|$x1

vector of Lagrangian multipliers.

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

with entries

$$a_{jj} = \sum_{i=1}^{N} \frac{2}{N_i} \cdot \beta(s_j)^2 \cdot w_{ij}^2, \;\; \forall j \in \mathcal{U}^c, \qquad (2.24)$$

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

$$v_{ij} = \frac{1}{N_i} \cdot \beta(s_j) \cdot w_{ij}, \;\; \forall j \in \mathcal{U}^c, \qquad (2.25)$$

and $i = 1, 2, \ldots, N$. The entries of the matrix $\mathbf{B}$ represents the mapping of spot-beam

queues to antenna groups and given by

$$b_{lj} \;\; = \;\; \begin{cases} 1, & \text{if } j \in \mathcal{A}_l \\ 0, & \text{if } j \notin \mathcal{A}_l \end{cases} , \;\; \forall j \in \mathcal{U}^c, \qquad (2.26)$$

and $l = 1, 2, \ldots, L$. The vector $\mathbf{d}$ represents the remaining power available for distribu-

tion to the spot-beam queues in set $\mathcal{U}^c$ following the power assignments to queues in set

$\mathcal{U}$, and given by

$$d_l = P_{tot} - \sum_{j \in (\mathcal{U} \cap \mathcal{A}_l)} p_j^{\text{MIN}}, \qquad (2.27)$$

28

for $l = 1, 2, \ldots, L$. In the solution of (2.21), 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 α_j vanishes in (2.21).

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

$$r_j^{\text{BAS-I}} = \min(\beta(s_j) \cdot p_j^{\text{BAS-I}}, r_{\max}), \tag{2.28}$$

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

2.4.2 Solution under BAS-II policy

Following a classification similar to that of the previous case, we assign $\mathbf{p}_{\mathcal{E}}^{\text{BAS-II}} = \mathbf{0}$, and $\mathbf{p}_{\mathcal{U}}^{\text{BAS-II}} = \mathbf{p}_{\mathcal{U}}^{\text{MIN}}$. Under BAS-II policy, all queues that are not empty are guaranteed the minimum power level assignment given in (2.19). Therefore, we start the solution with a base power vector, $\mathbf{p}^{\text{BASE}}$, such that

$$\mathbf{p}^{\text{BASE}} = [\mathbf{0}|\mathbf{p}_{\mathcal{U}}^{\text{MIN}}|\mathbf{p}_{\mathcal{U}^c}^{\text{MIN}}]^{\text{T}}. \tag{2.29}$$

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

$$d_l^{\text{BASE}} = \sum_{j \in \mathcal{A}_l} p_j^{\text{BASE}}, \tag{2.30}$$

and the remaining power for redistribution is equal to

$$d_l = P_{tot} - d_l^{\text{BASE}}, \tag{2.31}$$

for $l = 1, 2, \ldots, L$. 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

$$\mathcal{H} \triangleq \bigcup_l \{j : j \in (\mathcal{U}^c \cap \mathcal{A}_l) \text{ and } d_l > 0\} \qquad (2.32)$$

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

$$\mathcal{L} \triangleq \{l : d_l > 0\} . \qquad (2.33)$$

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

$$\mathbf{p}_{\mathcal{H}}^{\text{BAS-II}} = \mathbf{p}_{\mathcal{H}}^{\text{MIN}} + \mathbf{G} \cdot \alpha + \mathbf{Z} \cdot \mathbf{d}, \qquad (2.34)$$

where G and Z are as defined in (2.22) and (2.23), respectively. $\mathbf{X}$ is a $|\mathcal{H}|\text{x}|\mathcal{H}|$ matrix, $\mathbf{B}$ is a $|\mathcal{L}|\text{x}|\mathcal{H}|$ matrix, $\mathbf{d}$ is a $|\mathcal{L}|\text{x}1$ vector, and α is a $|\mathcal{H}|\text{x}1$ vector of Lagrangian multipliers. Equation (2.34) represents the power that can be distributed in addition to minimum power assignments.

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

$$a_{jj} = \sum_{i=1}^{N} \frac{2}{N_i} \cdot \beta(s_j)^2 \cdot w_{ij}^2, \ \forall j \in \mathcal{H}, \qquad (2.35)$$

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

$$v_{ij} = \frac{1}{N_i} \cdot \beta(s_j) \cdot w_{ij}, \ \forall j \in \mathcal{H}, \qquad (2.36)$$

30

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

$$b_{lj} = \begin{cases} 1, & \text{if } j \in \mathcal{A}_l \\ 0, & \text{if } j \notin \mathcal{A}_l \end{cases}, \quad \forall l \in \mathcal{L}, \ \forall j \in \mathcal{H}. \tag{2.37}$$

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 (2.31) for all $l \in \mathcal{L}$. In the solution of (2.34), a non-zero Lagrangian multiplier implies that the corresponding power level must remain at the minimum power level, and we have $p_j > p_j^{\min}$ when the multiplier α_j vanishes in (2.34). The final solution vector is represented as

$$\mathbf{p}^{\text{BAS-II}} = [\mathbf{0} | \mathbf{p}_{\mathcal{U}}^{\text{MIN}} | \mathbf{p}_{\mathcal{H}}^{\text{BAS-II}} | \mathbf{p}_{\mathcal{H}^c}^{\text{BAS-II}}]^{\text{T}}. \tag{2.38}$$

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

$$r_j^{\text{BAS-II}} = \min(\beta(s_j) \cdot p_j^{\text{BAS-II}}, r_{\max}), \tag{2.39}$$

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

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

2.5 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.

2.5.1 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, 6, 21]. 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

$$\text{EIRP} = P_t + G_t - L_t, \tag{2.40}$$

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

$$L_o = L_p + L_r, \tag{2.41}$$

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

$$C/N_o = \text{EIRP} - L_o + G/T - k, \tag{2.42}$$

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

$$E_b/N_o = C/N_o - R_b \text{ in dB.} \tag{2.43}$$

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

$$R_b = P_t + \beta(L_r), \tag{2.44}$$

where $\beta(L_r) = G_t - L_t - L_p + G/T - k - E_b/N_o - L_r$. It is possible to express (2.44) in linear terms:

$$R_b = \beta(L_r) \cdot P_t \text{ in bps.} \tag{2.45}$$

We will use (2.45) in calculating the rate-power relationship for a given rain attenuation level of the channel. Figure 2.5 shows rate-power relationship for different levels of

G_t (dB)	L_t (dB)	L_p (dB)
46.50	0.50	210.75
G/T (dB/K)	k (dB/K/Hz)	E_b/N_o (dB)
16.37	-228.60	3.56

Table 2.1: Numerical values for link-budget parameters.

rain attenuation. Throughout this analysis, we assume that rate is a continuous function of power, even though, in real systems, not all rates may 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 2.1.

2.5.2 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 [22, 23]. The model is based on specific channel model parameters from the DLR measurement campaign carried out at Oberp-
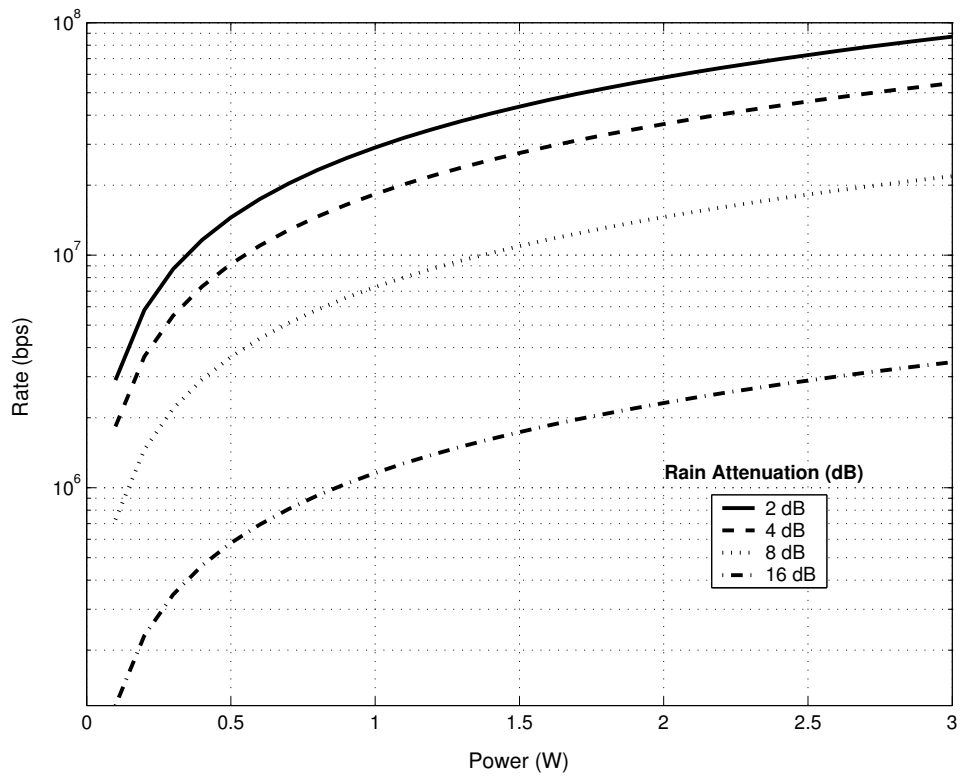
Figure 2.5: Rate-power curves for different rain attenuation levels.

faffenhofen 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 64 seconds resolution. Each attenuation level sample in decibels is input to (2.44), which through the link-budget calculation gives the downlink rate as a function of allocated antenna power. Figure 2.6 shows a sample realization of the rain attenuation time series and the corresponding cumulative distribution function for the channel model simulator.

2.5.3 Spot-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 corresponding antenna assignments of a geostationary satellite proposed for a commercial satellite system in [6, 21]. Figure 2.7 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-
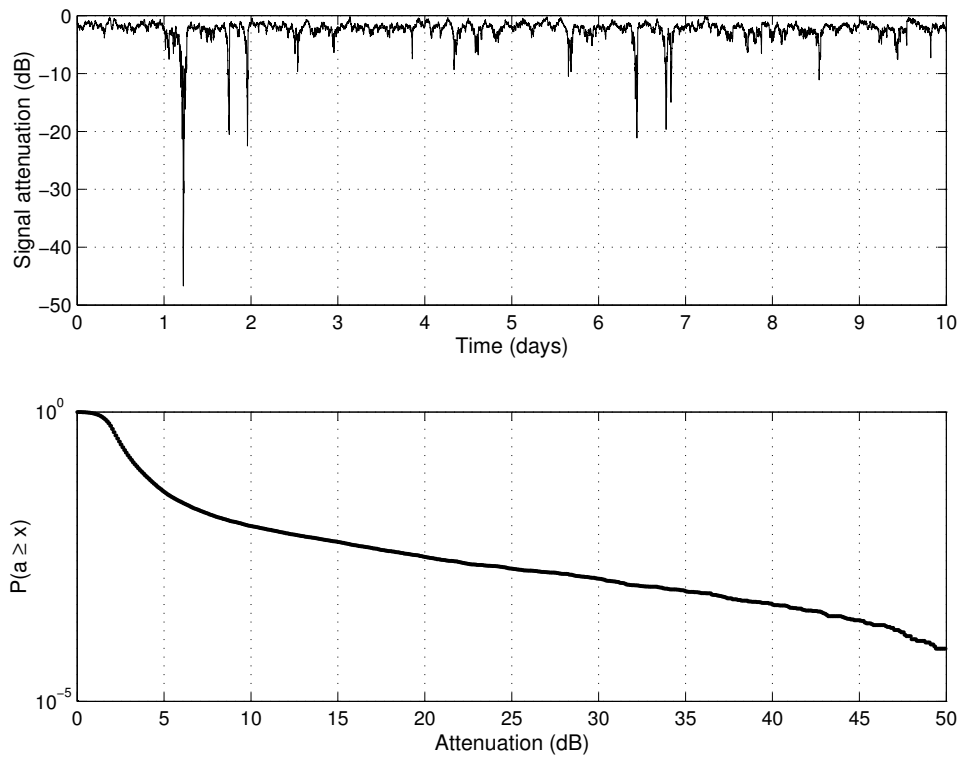
Figure 2.6: A sample attenuation time series and the cumulative distribution function of rain attenuation.

Figure 2.7: Locations of the 48 beams in two polarizations over the United States for the example satellite system.

and right-polarized spot-beam signals, respectively.

The frequency assignments for each beam family are as shown in Fig. 2.8. 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 [24]. Assuming that a flow f_i is more likely to be forwarded to spot-beam queue b_j if the spot-beam illuminates a larger fraction of the total population, we calculated the probability distribution plotted in Fig. 2.9. This distribution gives the probability of a flow being forwarded to a spot-beam for all 48 spot-beams and is used to initiate sessions between the NOC and the spot-beam locations.

Uplink Frequencies	Low (GHz)	29.500	29.620	29.740	29.860
	High (GHz)	29.620	29.740	29.860	29.980
Downlink Frequencies	Low (GHz)	19.700	19.820	19.940	20.060
	High (GHz)	19.820	19.940	20.060	20.180
		A1-R	C1-L	D1-L	B1-R
		A2-R	C2-L	D2-L	B2-R
		A3-R	C3-L	D3-L	B3-R
		A4-R	C4-L	D4-L	B4-R
Beam Family and		A5-R	C5-L	D5-L	B5-R
Polarization of Regional		A6-R	D1-R	D6-L	B6-R
Spot-Beams		B1-L	D2-R	D7-L	A1-L
		B2-L	D3-R	C1-R	A2-L
		B3-L	D4-R	C2-R	A3-L
		B4-L	D5-R	C3-R	A4-L
		B5-L	D6-R	C4-R	A5-L
		B6-L	D7-R	C5-R	A6-L

Figure 2.8: Beam family and polarization of regional spot-beams. Spot-beams operating at the same frequency band are served by the same antenna.
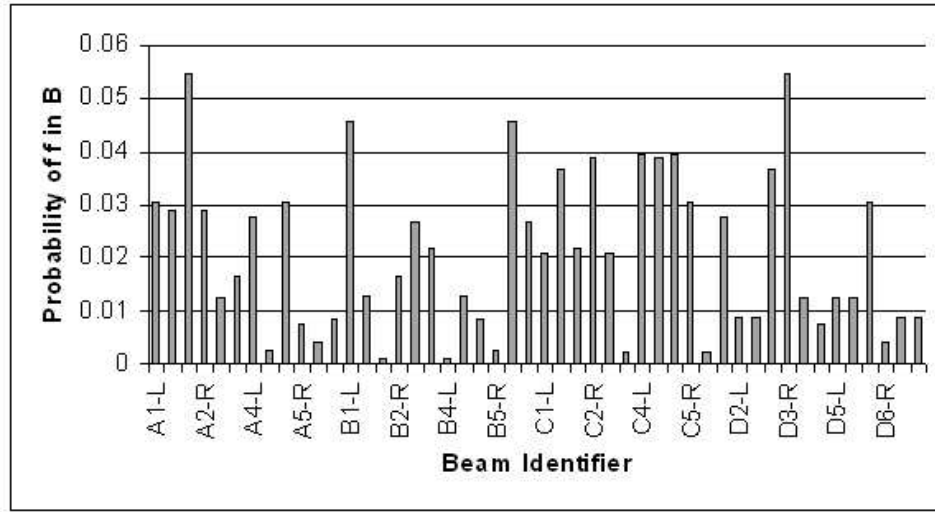
Figure 2.9: Connection probability distribution. This probability distribution is used to create connections between the NOC and the spot-beam locations.

The spot-beams share the access to $K = 4$ on-board antennas in TDM fashion. Therefore, we have to determine the sets, $\mathcal{A}_l$, $l = 1, 2, \ldots, L = M/K = 12$, of spot-beams that correspond to the allocation of TDM slots. The spot-beam queues that belong to the same set will access the antennas simultaneously. Therefore, two spot-beams operating at the same frequency band should not belong to the same set. In addition, to limit interference, the left- and right-polarized signals serving the same spot-beam location should operate in different sets. The problem of allocating spot-beams to TDM slots can be put into an optimization framework by itself. However, it is out of the scope of our problem setting. In this chapter, we assume that the system would make an attempt to distribute the load evenly among the antennas, and therefore, divide the spot-beam queues into 4 groups based on their connection probability values given in

	A_1											A_{12}
1	A1-R	A2-R	A3-R	A4-R	A5-R	A6-R	B1-L	B2-L	B3-L	B4-L	B5-L	B6-L
2	B3-R	B6-R	B5-R	B2-R	B1-R	B4-R	A5-L	A4-L	A1-L	A2-L	A3-L	A6-L
3	C5-L	C1-L	C2-L	C4-L	C3-L	D6-R	D4-R	D2-R	D1-R	D3-R	D7-R	D5-R
K=4	D1-L	D4-L	D7-L	D2-L	D3-L	D5-L	D6-L	C4-R	C5-R	C1-R	C2-R	C3-R
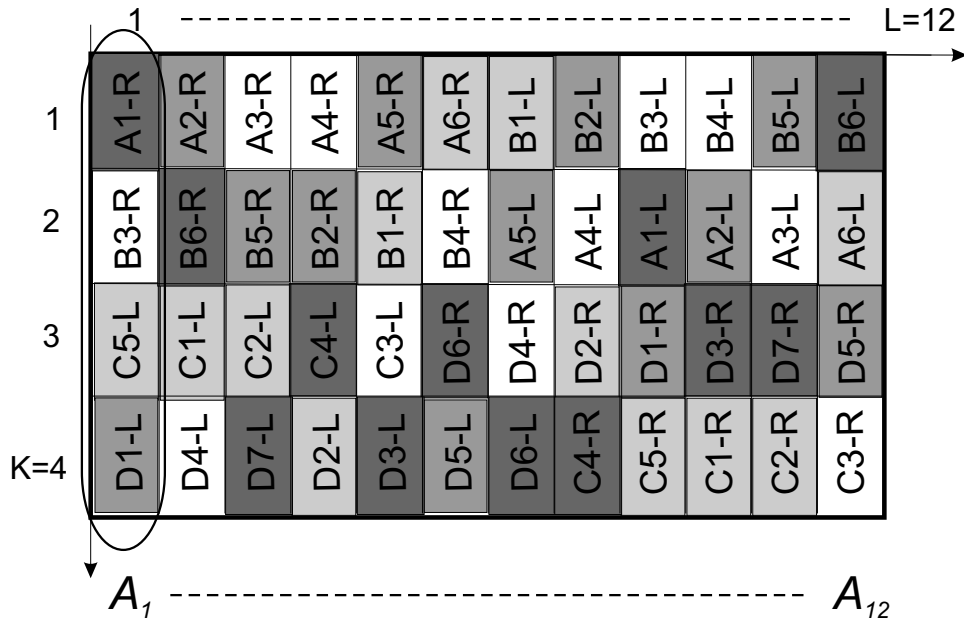
Figure 2.10: Allocation of spot-beams queues to transmission frame slots.

Fig. 2.9. We, then, pick one spot-beam queue from each group when assigning them to the sets subject to the constraints given as above. The resulting sets are shown in Fig. 2.10. Though this assignment is possibly not unique, any assignment that adheres to the constrains that we have imposed is sufficient for our purposes.

2.5.4   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 down-

stream of the bottleneck has been suggested in the literature. In [25], 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 dissertation, we test our algorithm with linear and logarithmic weight allocation policies. Independent allocation of weights to unicast and multicast flows is of no interest to us, since we believe that, for effective deployment of multicast delivery, there is a need to allocate more resources to multicast sessions over unicast ones.

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. In linear allocation, the weights are given by

$$
w_{ij} \;=\; \begin{cases} 0, & \text{if } n_{ij} = 0 \\[2ex] \dfrac{n_{ij}}{\sum\limits_{k \in \mathcal{B}_j} n_{kj}}, & \text{if } n_{ij} \neq 0 \end{cases} \tag{2.46}
$$

for $i = 1, 2, \ldots, N$ and $j = 1, 2, \ldots, M$. This scheme allocates n times more resources to a multicast flow with n downstream users than a unicast flow with one downstream user. However, this policy may be too aggressive in favoring multicast flows. The alternative is to use a logarithmic distribution. For logarithmic allocation, the weights

are distributed as

$$
w_{ij} = \begin{cases} 0, & \text{if } n_{ij} = 0 \\[2ex] \frac{1+\log(n_{ij})}{\sum\limits_{k \in \mathcal{B}_j} 1+\log(n_{kj})}, & \text{if } n_{ij} \neq 0 \end{cases} \tag{2.47}
$$

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

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.

2.6   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 Fig. 2.9. 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.

2.6.1 Results under BAS-I policy with logarithmic weights

In this section, we start looking at the results under BAS-I policy when using logarithmic weight allocation. 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 seconds, 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 minutes corresponding to $T = 675$ samples. In Fig. 2.11, we plot the percent change in the sustainable session rates of all active flows averaged over the total number of samples, given by

$$\overline{\eta}_i^{\text{BAS-I}} = 100 \cdot \frac{1}{T} \sum_{t=1}^{T} \left(\frac{\lambda_i^{\text{BAS-I}}[t] - \lambda_i^{\text{EAS}}[t]}{\lambda_i^{\text{EAS}}[t]} \right). \tag{2.48}$$

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

We observe that all of the active multicast and unicast sessions are served at higher average session rates compared to the EAS policy case. The multicast flows experience

44
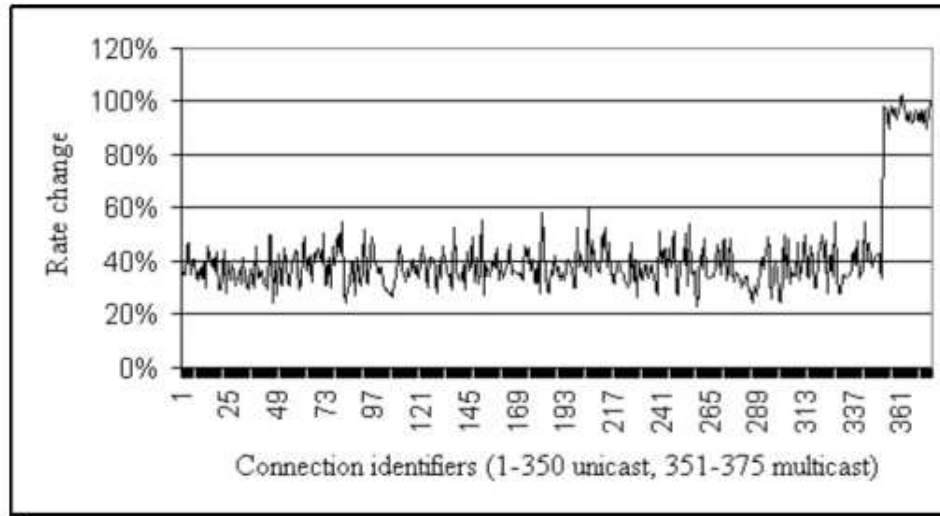
Figure 2.11: Percent rate change in the sustainable session rates of all active flows averaged over the test duration under BAS-I policy, compared to the session rates under EAS policy when using logarithmic weight allocation.

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 behavior. 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

45

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 Fig. 2.12, we look at this metric given by

$$\overline{\tau}_i^{\text{BAS-I}} = 100 \cdot \frac{1}{T} \sum_{t=1}^{T} \mathbf{1} \left(\lambda_i^{\text{BAS-I}}[t] \geq \lambda_i^{\text{EAS}}[t] \right), \tag{2.49}$$

for $i = 1, 2, \ldots, (L_u + L_m)$, 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 this 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 thought the session rate is higher on the average.

In Fig. 2.13, we plot the percentage of total power assigned to each spot-beam queue, and in Fig. 2.14 the corresponding service rates, given by

$$\overline{p}_j^{\text{BAS-I}} = 100 \cdot \frac{1}{T} \sum_{t=1}^{T} p_j^{\text{BAS-I}}[t], \text{ and} \tag{2.50}$$

$$\overline{r}_j^{\text{BAS-I}} = \frac{1}{T} \sum_{t=1}^{T} r_j^{\text{BAS-I}}[t], \tag{2.51}$$

respectively, for $j = 1, 2, \ldots, M$. Observe that several spot-beam queues have average power levels that are below the minimum power levels that would maintain the EAS
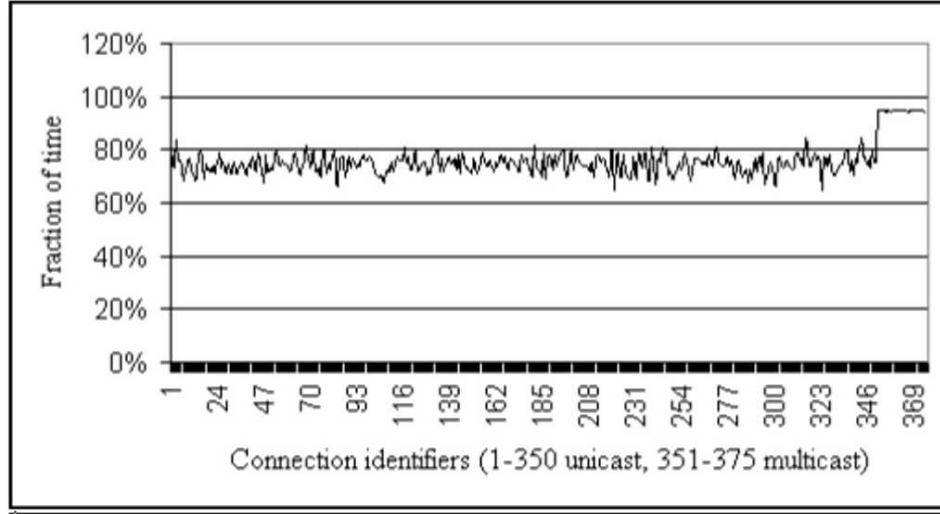
Figure 2.12: Percent of total test time BAS-I session rates are equal to or better than EAS session rates for all active flows when using logarithmic weight allocation.

session rates. Consequently, all unicast flows forwarded to such queues have lower session rates than their EAS rates giving rise to the behavior we observe in Fig. 2.12.

The second set of results looks 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 groups is varied between $L_m = 5$ to 30. In Fig. 2.15, we plot the change in the sustainable session rates averaged over all active unicast and multicast flows, — i.e.

$$\overline{\eta}_u^{\text{BAS-I}} = \frac{1}{L_u} \sum_{i=1}^{L_u} \overline{\eta}_i^{\text{BAS-I}}, \tag{2.52}$$

and

$$\overline{\eta}_m^{\text{BAS-I}} = \frac{1}{L_m} \sum_{i=L_u+1}^{(L_u+L_m)} \overline{\eta}_i^{\text{BAS-I}}, \tag{2.53}$$

as the number of active multicast groups is varied.

47
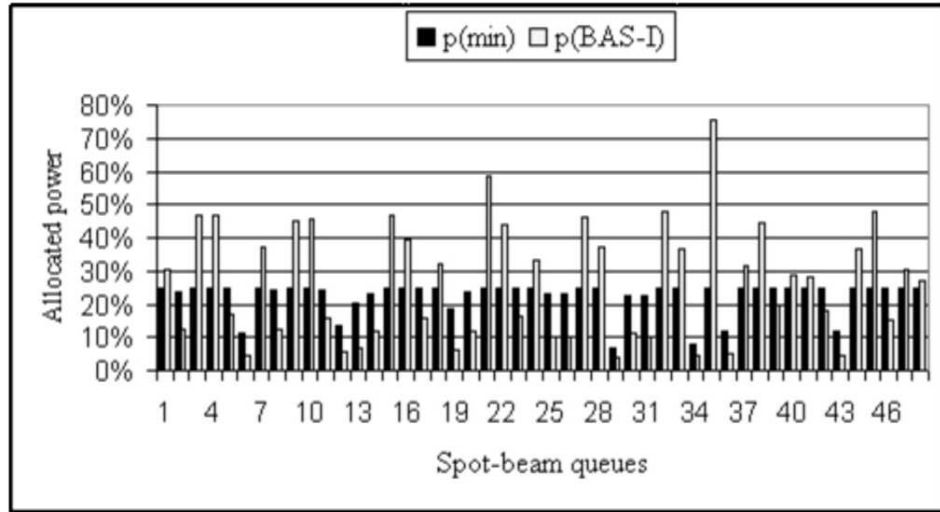
Figure 2.13: Average power assigned to each spot-beam queue as percentage of total system power over the test duration under BAS-I policy when using logarithmic weight allocation.
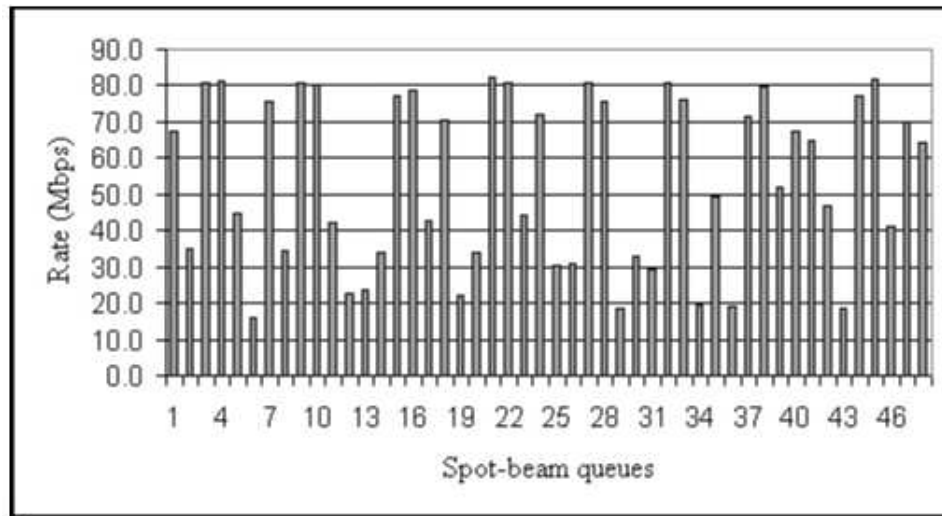
Figure 2.14: Average service rates of the spot beam queues over the test duration under BAS-I policy when using logarithmic weight allocation.

Figure 2.15: Average rate change in the sustainable session rates of unicast flows averaged over all active unicast flows and average rate change in the sustainable session rates of multicast flows averaged over all active multicast flows under the BAS-I policy when using logarithmic weight allocation.

49

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 Fig. 2.16, we plot the percentage of time flows have 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.

Figure 2.16: Percent 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 percent 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 when using logarithmic weight allocation.
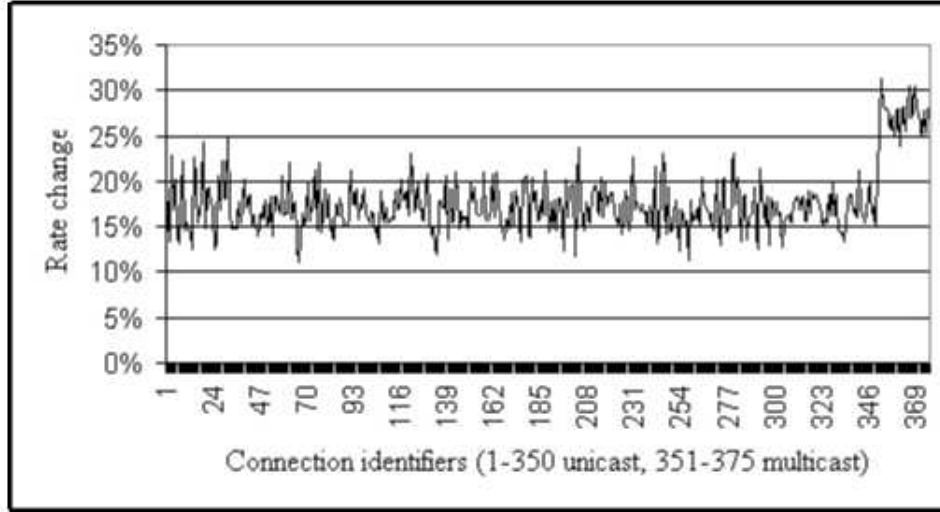
Figure 2.17: Percent rate change in the sustainable rates of all active flows averaged over the test duration under BAS-II policy, compared to the session rates under EAS policy when using logarithmic weight allocation.

2.6.2 Results under BAS-II policy with logarithmic weights

In this section, we look at the same set of metrics under the BAS-II policy when using logarithmic weight allocation. In Fig. 2.17, we plot the percent change in the sustainable session rates of all active flows averaged over a similar test duration of 720 minutes 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 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. $\overline{\tau}_i^{\text{BAS-II}} = 100\%$, $\forall i$.

52
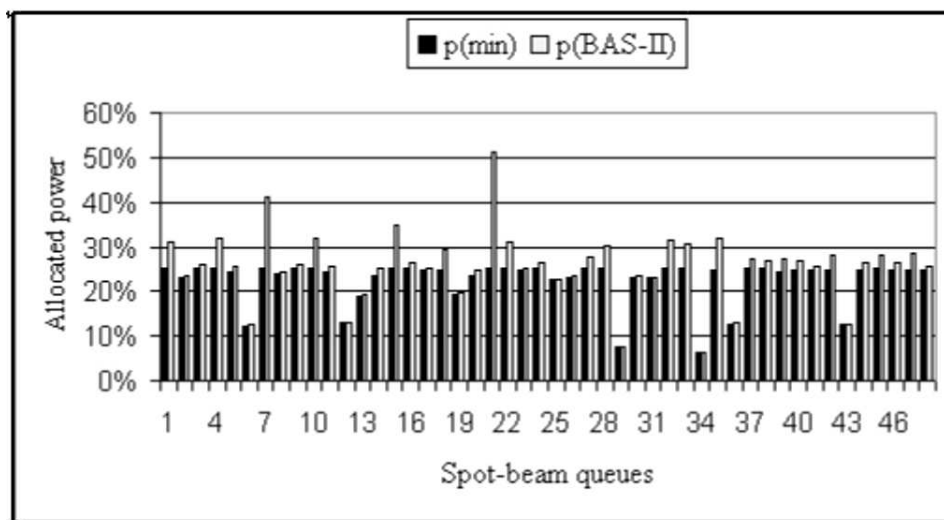
Figure 2.18: Average power assigned to each spot-beam queue as percentage of total system power over the test duration under BAS-II policy when using logarithmic weight allocation.

In Fig. 2.18, we plot the percentage of total power assigned to each spot-beam queue, and in Fig. 2.19, 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 is varied between $L_m = 5$ to 30. In Fig. 2.20, we plot the change in the sustainable session rates averaged over all flows as
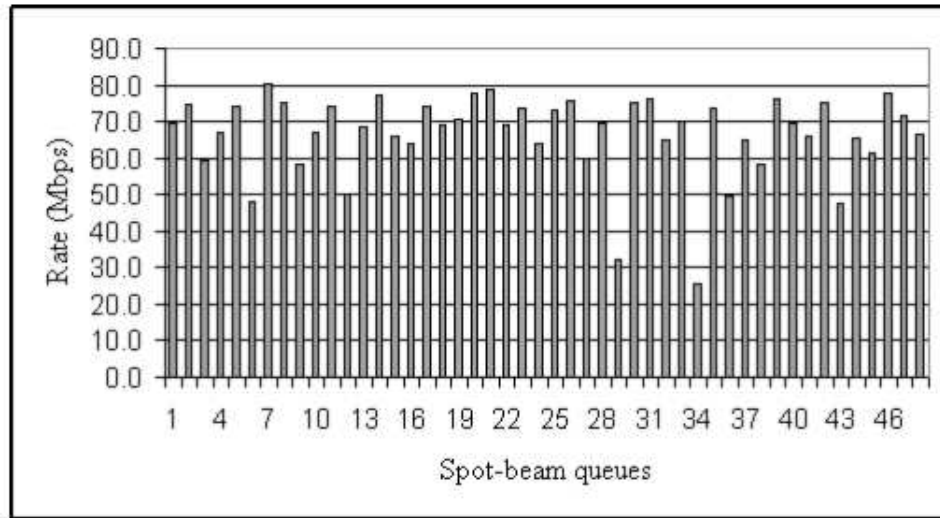
Figure 2.19: Average service rates of the spot beam queues over the test duration under BAS-II policy when using logarithmic weight allocation.

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 important observation to make is that this time mixing of unicast and multicast flows negatively affect the improvement of multicast flows (compare to Fig. 2.15). 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

54
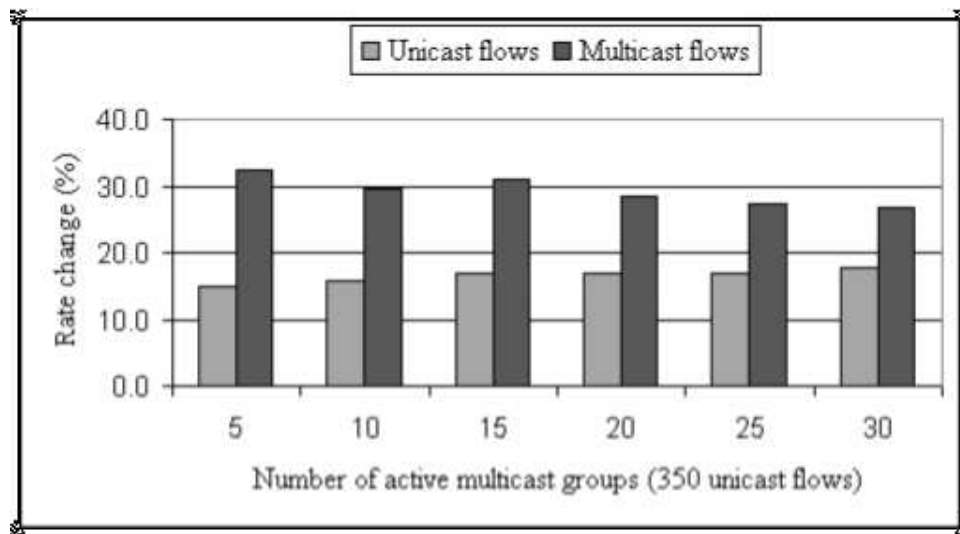
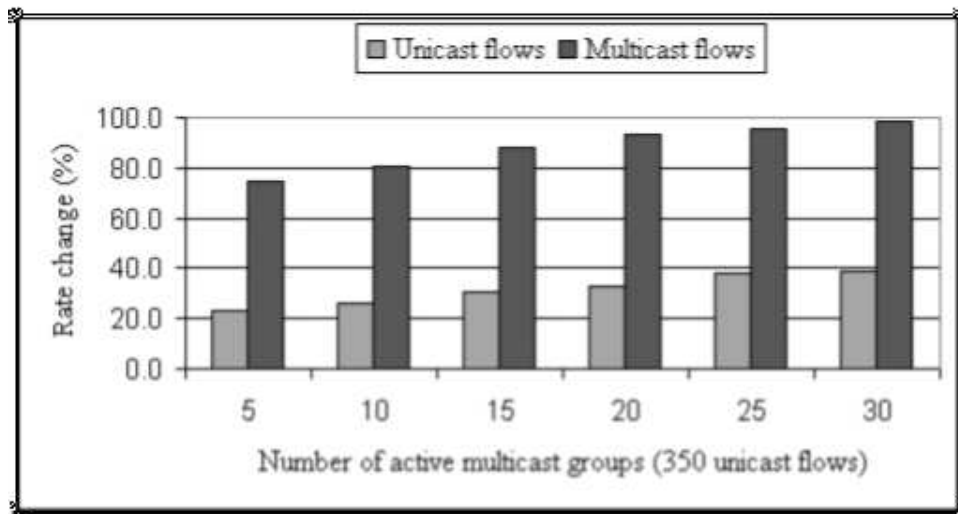Figure 2.20: Average rate change in the sustainable session rates of unicast flows averaged over all active unicast flows and average rate change in the sustainable session rates of multicast flows averaged over all active multicast flows under the BAS-II policy when using logarithmic weight allocation.

issues of the BAS-I policy. In the following sections, we look at the effect of having linear weight allocation on the performance of our allocation policies.

2.6.3 Results under BAS-I policy with linear weights

In this section, we make a head-to-head comparison to the results of the previous sections. We first look at the performance 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 as in the previous cases, but the weights are allocated linearly in the system. In Fig. 2.21, we plot the percent change in the sustainable session rates of all active flows averaged over the total number of samples, and in Fig. 2.22, the corresponding figure for the total fraction of time session rates are above the EAS rates. In both cases, we observe a very similar performance trend in comparison to the cases when using logarithmic weights. In the next set of figures, we compare how having linear weights fares when the number of active multicast groups is varied. Figures 2.23 and 2.24 compare the average rate change experienced by unicast and multicast flows, respectively under the two different weight allocation schemes. We observe that with both logarithmic and linear weights, the performance improvement over the EAS policy case remains in the same range. Therefore, we conclude that, in terms of the relative improvement over the EAS policy, the results are insensitive to the choice of the weights.

Figures 2.25 and 2.26 compare the fraction of total test time BAS-I session rates
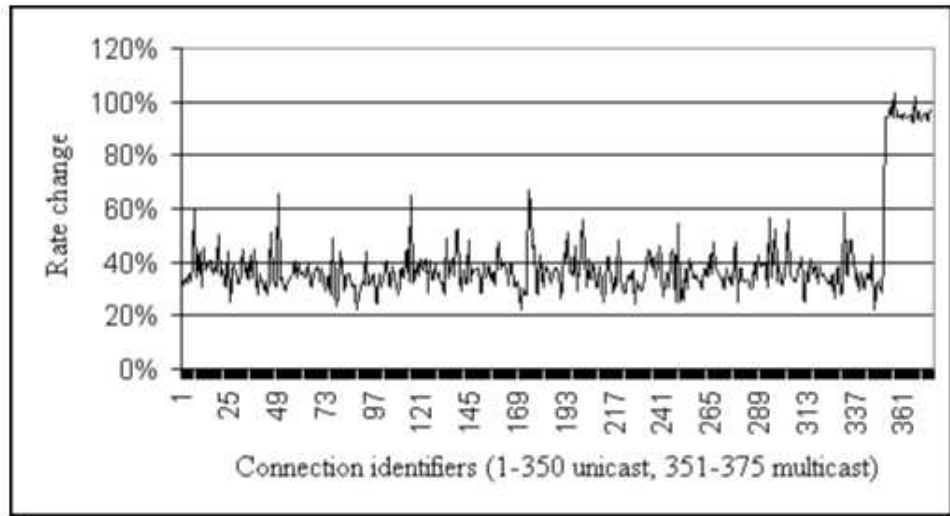
Figure 2.21: Percent rate change in the sustainable session rates of all active flows averaged over the test duration under BAS-I policy, compared to the session rates under EAS policy when using linear weight allocation.
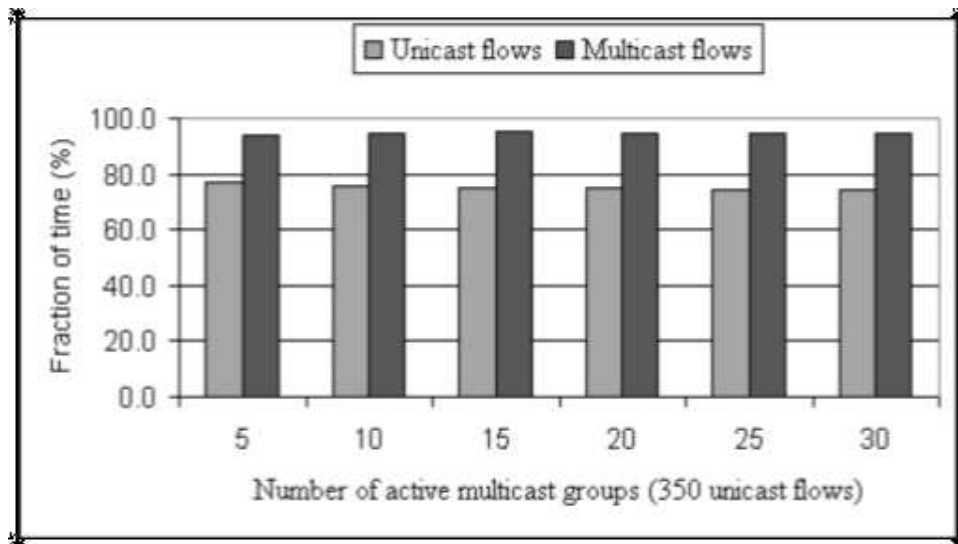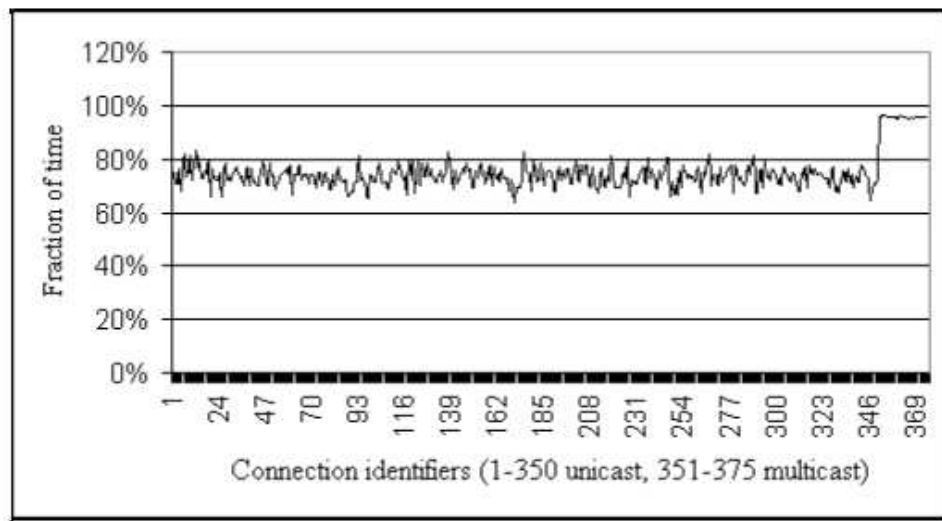

Figure 2.22: Percent of total test time BAS-I session rates are equal to or better than EAS session rates for all active flows when using linear weight allocation.
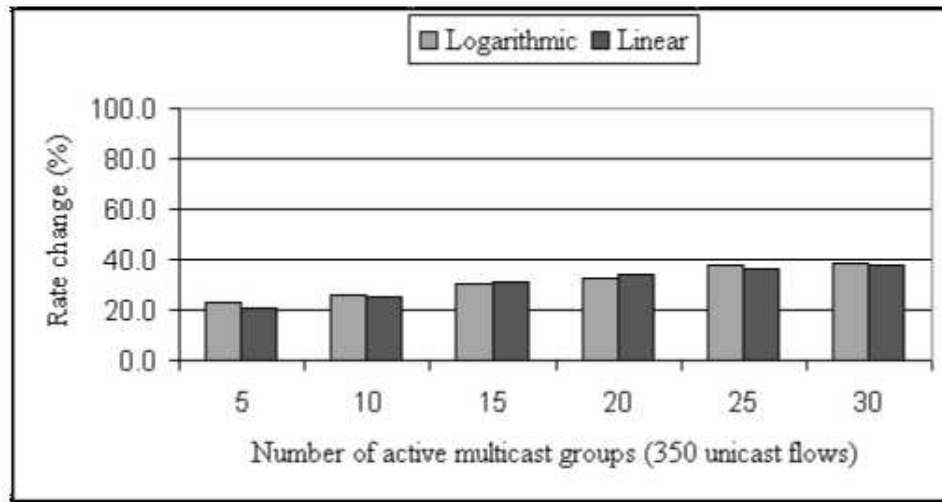
Figure 2.23: Head-to-head comparison of percent rate change in the sustainable session rates of unicast flows averaged over all active unicast flows under the BAS-I policy for logarithmic and linear weight allocations.
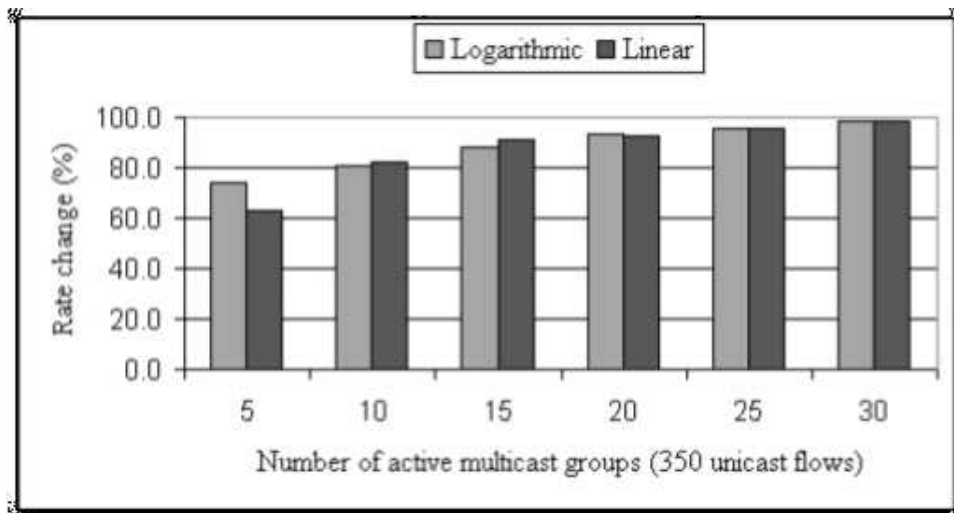
Figure 2.24: Head-to-head comparison of percent rate change in the sustainable session rates of multicast flows averaged over all active multicast flows under the BAS-I policy for logarithmic and linear weight allocations.
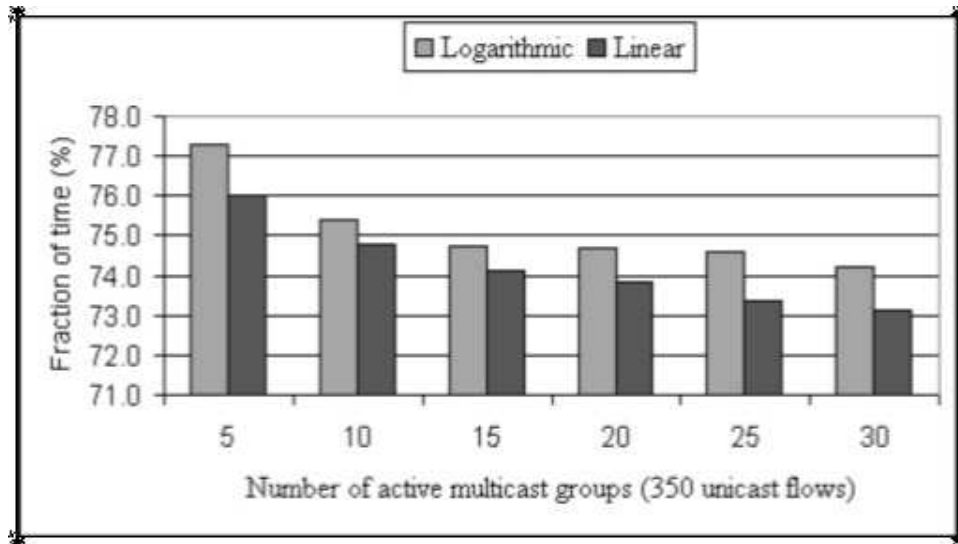
Figure 2.25: Head-to-head comparison of fraction of total test time BAS-I session rates of unicast flows are equal to or better than EAS session rates when averaged over all active unicast flows for logarithmic and linear weight allocations.

of unicast and multicast flows are equal to or better than EAS session rates, respectively, when using logarithmically and linearly allocated weights. We observe that having linear weight allocation mostly favors multicast flows in terms of the fraction of total time these flows are allocated a rate above their EAS rates. This result is attributed to the fact that under linear weight allocation, consistently more bandwidth is allocated to multicast flows than unicast flows. However, on the overall, our observations suggest that weight allocation scheme does not have a pronounce effect on the performance of our algorithm.
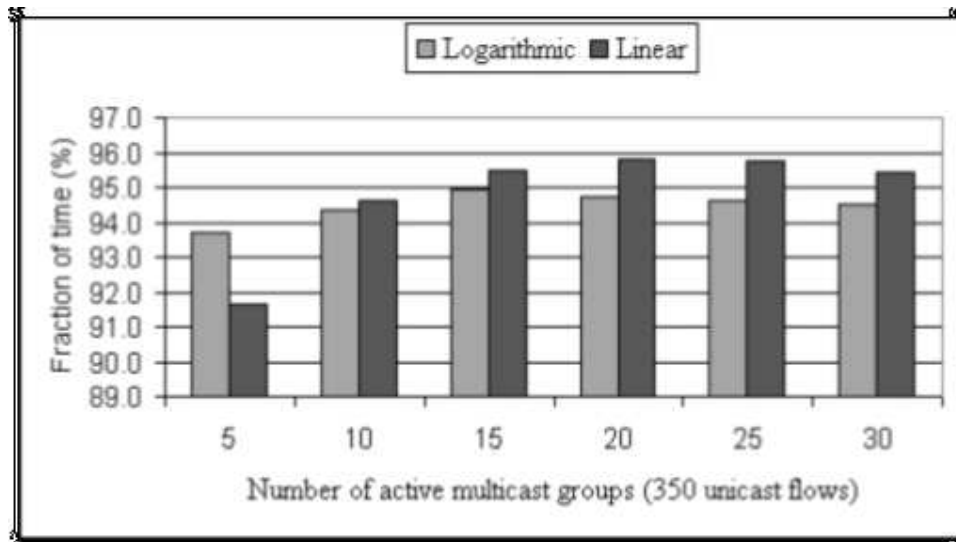
Figure 2.26: Head-to-head comparison of fraction of total test time BAS-I session rates of multicast flows are equal to or better than EAS session rates when averaged over all active multicast flows for logarithmic and linear weight allocations.

2.6.4 Results under BAS-II policy with linear weights

This section presents the corresponding performance graphs when using linear weights under BAS-II policy. Figure 2.27 plots the average rate change in the sustainable session rates of all active flows over the test duration when using linear weights. Figures 2.28 and 2.29 provide head-to-head comparison of logarithmic and linear weights in terms of the percent rate change experienced by unicast and multicast flows, respectively, under BAS-II policy. These results strengthen the conclusion that the choice of weight allocation scheme does not have a pronounce effect on the performance of both policies.

60
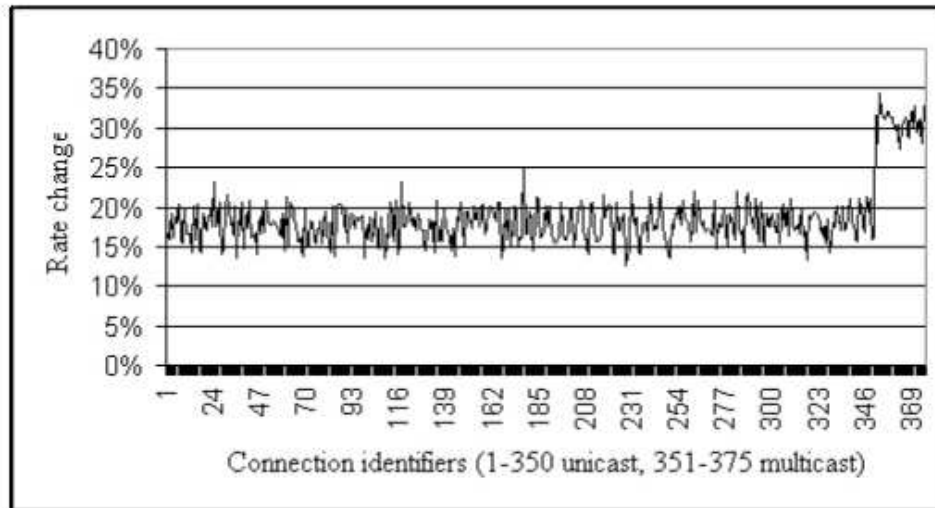
Figure 2.27: Percent rate change in the sustainable session rates of all active flows averaged over the test duration under BAS-II policy, compared to the session rates under EAS policy when using linear weight allocation.
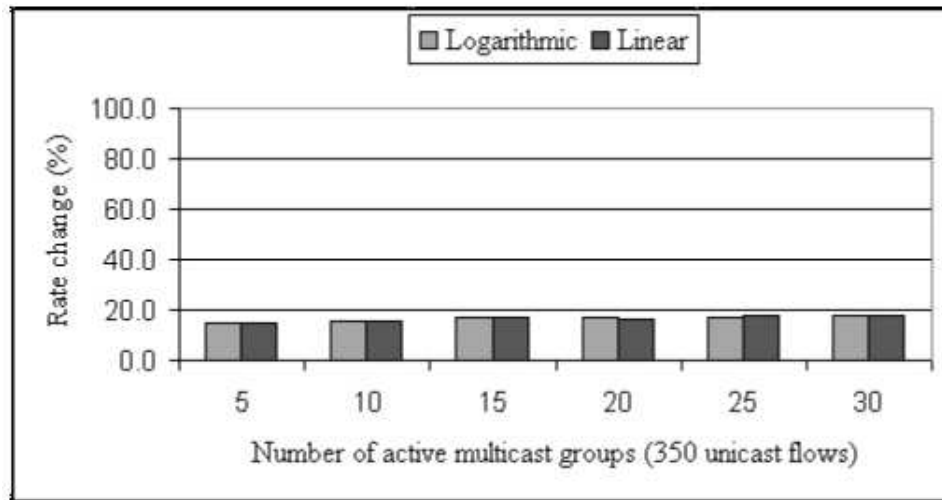
Figure 2.28: Head-to-head comparison of percent rate change in the sustainable session rates of unicast flows averaged over all active unicast flows under the BAS-II policy for logarithmic and linear weight allocations.
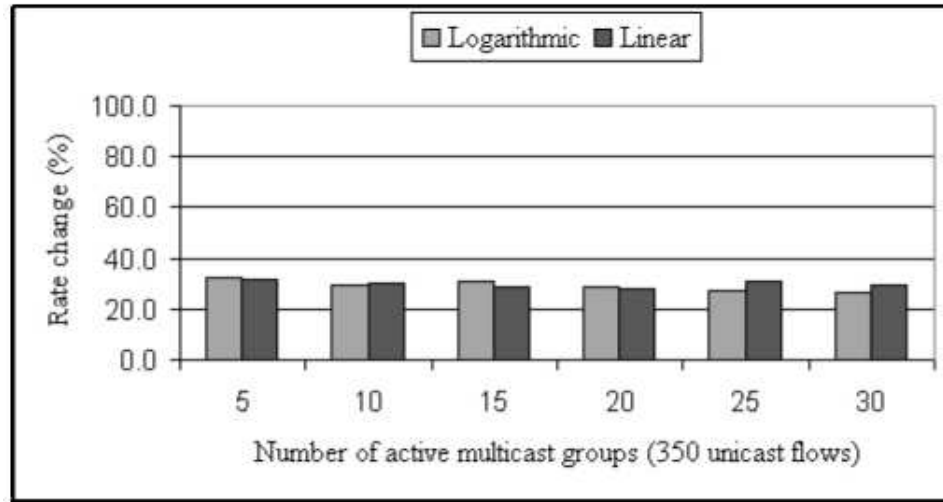
Figure 2.29: Head-to-head comparison of percent rate change in the sustainable session rates of multicast flows averaged over all active multicast flows under the BAS-II policy for logarithmic and linear weight allocations.

2.7 Summary

In this chapter, 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 provided 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.

Chapter 3

A Multiple Subset Sum Formulation for Efficient Feedback Collection

3.1 Introduction

The previous chapter overviewed various adaptive schemes for mitigating the effects of atmospheric and rain fading in satellite communication systems, and focused on an adaptive power allocation and rate reduction policy for overcoming user heterogeneity in multicast communication over satellite networks. Although the policy explicitly relies on the knowledge of channel states for calculating the effective downlink rates of each spot-beam queue, we have simply assumed that this information was readily available to the system, and have not discussed how it was obtained. In general, however, any scheme that adapts system parameters to the changing propagation conditions would typically have to perform a series of common functions that involves the monitoring, prediction, and estimation of the channel conditions.

User feedback would be a valuable input for a number of such components, however, collecting periodic feedback from a large number of users would result in the well-known *feedback implosion* problem. Feedback implosion is identified as a major problem when a large number of users try to transmit their feedback messages to a central node, causing high traffic concentration and backlog in the network. Although not

unique to satellite based networks, feedback implosion phenomenon has additional side effects in this context, since a significant portion of the uplink resources may be hold up to transmit these feedback messages instead of useful data traffic, and the access to the shared medium may be clogged [11, 26]. The fact that the satellite return channel (uplink) is a shared medium, and that the satellite spectrum is limited and expensive, makes it necessary to minimize the amount of bandwidth required for transmission of user feedback.

Feedback implosion problem has been studied in the literature as part of several terrestrial multicast transport protocols, which rely on user feedback for providing reliable transmission. In this context, feedback implosion arises when individual users confirm reception of a data packet by sending positive acknowledgments, or request retransmission of a data packet by sending negative acknowledgments to the source. Therefore, several existing terrestrial multicast protocols adopt feedback suppression and feedback aggregation schemes to control the number, and the flow of feedback packets. Feedback aggregation is applicable in networks where a logical or physical hierarchy is possible. Feedback aggregation relies on intermediate entities in the network, such as routers, to collect, filter-out and combine feedback information as it traverses the network towards the source. For example, an intermediate router could collect positive acknowledgments sent by several users for the same data packet, and forward only one (aggregated) acknowledgment towards the source. This aggregation towards the root of the hierarchy cuts back on the number of feedback reports and prevents receivers from contacting the

source directly, enabling the protocol to scale over a larger set of receivers.

Feedback aggregation schemes require additional functionality at the intermediate routers of the network [27]. This raises questions about whether such router extensions can be standardized, deployed, and automatically configured in the backbone network. In order to circumvent this problem, feedback suppression algorithms control the number of feedback packets created and forwarded by every user, instead of controlling the volume of feedback information by aggregation. In terrestrial networks, a typical implementation of feedback suppression would require users to listen to the feedback messages that were multicasted or broadcasted to the network by other users, and to suppress their own feedback if the message is redundant or is covered by some other message. For example, a user could suppress its own negative acknowledgment for a missing packet, if some other user has already sent a negative acknowledgment for the same packet, since the source would be multicasting the packet to all users as soon as the first acknowledgment is received. Coupled with random timers and statistical backoff, this type of feedback suppression would, in an ideal setting, let only one feedback report to be returned to the sender for every missing packet [28–31]. However, multicast (broadcast) forwarding of feedback reports must be scoped by use of *time-to-live* (TTL) fields to minimize the flooding of the network by feedback packets.

Revisiting the typical satellite deployment scenario from the previous chapter (Fig. 3.1), feedback aggregation would not be a viable solution in this architecture, since the star topology of the network does not lend itself into a hierarchy, and users

66
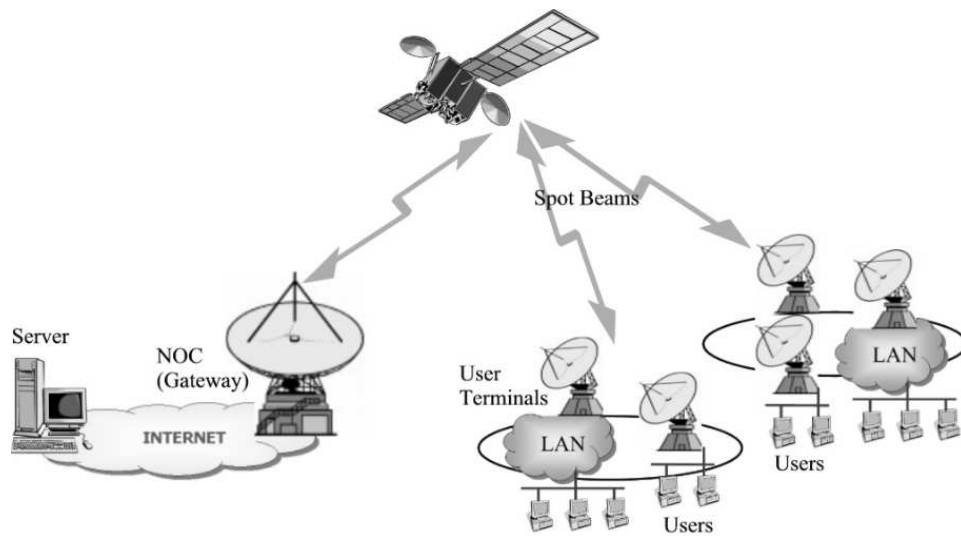
Figure 3.1: Satellite communication system architecture. The satellite provides broad-band access to users across multiple spot-beam locations.

have to share the uplink access in each spot beam location. Even if feedback aggrega-tion techniques are used in the local domains, such as the local area networks, feedback implosion would remain to be a challenge due to the potentially large number of termi-nals. A feedback suppression implementation similar to that of the terrestrial networks is also not possible, since in general, there exists no secondary network or communi-cation among the users to distribute the feedback messages. Even though the satellite can replay back the feedback packet it received from a user by broadcasting it to all the receivers, delay incurred during the coordination would be unacceptable in most cases, because of the long propagation delays. For example, a user would have to wait on the order of seconds to listen to the feedback messages of others before deciding on whether to transmit its own message.

However, other possible implementations of feedback suppression is possible, if, for example, a representative (or a set of representatives) among the users could be selected to transmit feedback messages for the group. Several proposals look at efficient transmission of feedback messages over the satellite uplink without relying on terrestrial return paths. In [32], the proposed algorithm uses a round-robin time division multiple access (TDMA) scheme on the uplink channel. The disadvantage is that, it does not scale well with increasing number of receivers, since more uplink bandwidth has to be allocated for feedback transmission. In [33], the authors group receivers into clusters, where only the cluster heads transmit feedback over the satellite link. Although this scoping reduces the number of users that have to access the uplink channel, the approach assumes that receivers inside a cluster are able to communicate with each other through a secondary network. Some of the more recent work tries to avoid any communication and collaboration between the users. To our knowledge, the approach described in [34] is one of the first proposals to solve the feedback implosion problem over satellite networks without any secondary communication between the users. The disadvantage of this approach is that the feedback resolution requires multiple round trip times over the high latency satellite channel. In [35], the authors propose forming receiver clusters based on their spatial correlations, e.g. rain fall precipitation data, which is a prime factor that degrades service quality. Receivers in the same cluster use different return channels, while receivers in different clusters may use the same channel since there is a low probability of receivers located away from each other sending feedback simultane-

ously.

We believe that the requirement for terrestrial return links, as well as the requirement for secondary communication between the users, are restrictive, and in general contradicts with the reasons for deployment of satellite networks. The feedback implosion and the uplink bandwidth sharing problem should be studied without these assumptions. In this chapter, we present a *feedback implosion suppression algorithm* for efficient feedback collection from a large set of users without relying on any collaboration among them, and on any secondary network infrastructure. We show that, in some of the cases, it is possible to reduce the volume of user feedback by exploiting the nature of the information and by considering that (i) the feedback may contain redundant information (due to correlations among the users), and (ii) the protocol may need to track the behavior of only a subset of users, e.g. the group of users with worst case channel conditions. In the following sections, we first look at the problem as a part of a protocol that collects periodic channel state information from users, in order to complement the policy behavior of the previous chapter. We later integrate the algorithm with the power allocation policy of the previous chapter and look at the performance results when the two cooperates.

3.2 CSI collection protocol

In this chapter, we refer to the network architecture given in Fig. 3.1, and assume that a protocol performs between the satellite gateway (NOC) and a group of direct users that

are located inside the footprint of the satellite, in an attempt to collect periodic channel state information (CSI). Users submit their channel state information by accessing the uplink channel and sending a feedback message. In each spot-beam, the users access the uplink in multiple-frequency time-division multiple access (MF-TDMA) mode, where multiple frequency channels are allocated for uplink access, and the TDMA scheme is used within each frequency channel. The users acquire access by requesting a number of time-slots from the NOC, which invokes a resource allocation algorithm to share the uplink capacity (available time-slots) of the spot-beam among all active users. The result of the resource allocation is broadcast to all active users. We assume that a feedback message can be transmitted in a single uplink time-slot.

First, we present the CSI collection protocol behavior that will be relevant to the operation of our feedback implosion suppression algorithm. The CSI of interest, in this case, is the signal attenuation level due to atmospheric and rain fading as measured at the input of the user receiver equipment. The goal of the protocol is to calculate the *maximum* signal attenuation level of each spot-beam location by collecting periodic reports on the attenuation levels of all active users inside coverage area. The maximum signal attenuation level is used as input to various fade mitigation algorithms in order to compensate for the user with worst channel conditions.

It is assumed that the set of all active users are readily available at the NOC at all times, since this information is also required for login, bandwidth allocation, and possibly billing purposes. Therefore, when a user terminal u_i becomes active, the protocol

initializes and keeps a state variable $\hat{s}_i$ at the NOC for recording the attenuation level of that user. At this point the protocol has no information on the attenuation level of the user, therefore, the initial value of the variable is set to the expected attenuation level $\overline{A}$. The expected attenuation level may be calculated through empirical data (off-line), or by modeling. The protocol updates the state variables for all active users, and calculates a new maximum for every spot-beam b_j at every collection period by following a two-step process:

1. Let $s_i(kT)$ denote the attenuation level sample measured by user u_i at the start of collection period k at time $t = kT$, where T is the collection period. Active users send their measurements to the NOC, which collects and updates the corresponding state variable values:

$$\hat{s}_i[k] \leftarrow s_i(kT) \ \forall i \in \mathcal{U}_k, \tag{3.1}$$

   where $\mathcal{U}_k$ is the set of active users at the time of collection period k.

2. Let $\mathcal{U}_k^j$ denote the set of active users inside the coverage area of spot-beam b_j. At the NOC the maximum attenuation level for the spot-beam is calculated as

$$s_{\max}^j[k] \overset{a}{=} \max_{i \in \mathcal{U}_k^j}\{\hat{s}_i[k]\} \overset{b}{=} \max_{i \in \mathcal{U}_k^j}\{s_i(kT)\}, \tag{3.2}$$

   for $j = 1, 2, \ldots, M$, where M is the number of spot-beams, and $\mathcal{U}_k = \bigcup_{\forall j} \mathcal{U}_k^j$.

The update operation in (3.1) requires all active users to acquire access to up-link time-slots to transmit their feedback messages. However, observe that, for every

spot-beam, feedback volume would be minimized if only the user with the maximum attenuation level responded at every collection period. We can consider two extreme scenarios for illustrative purposes: (i) all users communicate among each other through a secondary network (possibly a terrestrial connection) before deciding on whether to transmit a feedback message, and suppress the feedback messages of all but the user with the maximum; (ii) every user is assigned a separate uplink time-slot, over which the feedback message is transmitted to the satellite. The former scenario requires additional infrastructure and collaboration among users, which we believe is an overly restrictive requirement and is usually contradictory to the reasons for deployment of a satellite network in the first place. The latter situation gives rise to the feedback implosion problem as well as the waste of uplink resources.

In order to reduce the volume of feedback information that is transmitted through the network, our FIS algorithm modifies the behavior of the CIS collection protocol such that, with FIS in place, only a subset $\hat{\mathcal{U}}_k \subseteq \mathcal{U}_k$ of active users report their measurements using at most C_k uplink time-slots for the collection period. Therefore, not all state variables are updated as in (3.1) at the end of collection period k:

$$\hat{s}_i[k] \leftarrow s_i(kT) \ \forall u_i \in \hat{\mathcal{U}}_k, \tag{3.3a}$$

$$\hat{s}_i[k] \leftarrow \hat{s}_i[k-1] \ \forall u_i \notin \hat{\mathcal{U}}_k. \tag{3.3b}$$

Consequently, the maximum attenuation levels calculated for each spot-beam using the state variables may not be equal to the maximum of the attenuation samples of active users measured at time $t = kT$, — i.e equality (b) in (3.2) may not hold in general.

72

Therefore, we denote the maximum attenuation levels calculated when FIS algorithm is in place by $\hat{s}^j_{\max}[k]$ for $j = 1, 2, \ldots, M$, where:

$$\hat{s}^j_{\max}[k] = \max_{i \in \mathcal{U}^j_k}\{\hat{s}_i[k]\}, \tag{3.4}$$

and $\mathcal{U}_k = \bigcup_{\forall j} \mathcal{U}^j_k$.

The goal of the FIS algorithm is, therefore, to limit the error between the actual maximum in (3.2) and the partial maximum in (3.4) by determining the users in set $\hat{\mathcal{U}}_k$ and their assignments to a given number of uplink time-slots at every collection period. In the next section, we describe the FIS algorithm for determining set $\hat{\mathcal{U}}_k$.

3.3 FIS algorithm

3.3.1 Algorithm formulation

In this section, we will describe the algorithm behavior for a single spot-beam. We assume that identical copies of the algorithm run for each spot-beam in order to suppress the feedback information from users covered by them. Let C^j_k be the maximum number of time-slots that can be allocated to transmission of feedback messages for collection period k in spot-beam b_j. The system may reserve a fixed number of time-slots for feedback transmission, or the value of C^j_k can be determined by the remaining time-slots after the user requests for data transmissions are accommodated. The FIS algorithm selects users and assigns them to one of the C^j_k available uplink time-slots by solving a *multiple subset sum* problem (MSSP) [36]. Multiple subset sum problem is a variant

of the bin packing problem in which the number of bins is given. Mathematically, the MSSP can be described as follows:

Definition 1 *Given a set $\{u_i\}_{i=1}^{N}$ of* items, *each item u_i having a positive weight w_i, and a set $\{h_l\}_{l=1}^{C}$ of identical* bins *with positive capacity H, what is the assignment that maximizes:*

$$\max \sum_{l=1}^{C}\sum_{i=1}^{N} w_i x_{il} \tag{3.5a}$$

$$\textit{subject to } \sum_{i=1}^{N} w_i x_{il} \leq H, \ \forall \, l, \tag{3.5b}$$

$$\sum_{l=1}^{C} x_{il} \leq 1, \ \forall \, i, \tag{3.5c}$$

$$x_{il} \in \{0,1\}, \ \forall \, i, l. \tag{3.5d}$$

Without loss of generality, we assume that $N \geq C$, otherwise the problem is trivially solved.

At the start of the CSI collection period k, using the set of active users in $\mathcal{U}_k^j$, and the state variables $\{\hat{s}_i[k-1]\}_{i \in \mathcal{U}_k^j}$, the FIS protocol solves an instance of the MSSP where,

$$N = |\mathcal{U}_k^j|, \ \text{cardinality of the set } \mathcal{U}_k^j, \tag{3.6a}$$

$$C = C_k^j, \ \text{number of available time-slots}, \tag{3.6b}$$

$$w_i = \min\{\hat{s}_i[k-1], A_{\max}\}, \ \forall i \in \mathcal{U}_k^j, \tag{3.6c}$$

$$H = A_{\max}, \tag{3.6d}$$

where $A_{\max}$ is the maximum attenuation level that can be compensated by the fade countermeasures implemented by the system. Equation (3.6c) allows all users with a reported attenuation level larger than $A_{\max}$ be treated as equals since the system does not have the means to compensate for their attenuation level above this limit.

Let $\mathbf{x}^k$ be the solution vector that represents the assignment of items to bins. The assignment of items to bins represents the assignment of users to uplink time-slots, such that $x_{il}^k = 1$ if user u_i is *selected* and assigned to time-slot l for $\forall i \in \mathcal{U}_k^j$, and $l = 1, 2, \ldots, C_k^j$, else $x_{il}^k = 0$.

The NOC broadcasts the solution vector over the network as a part of the existing on-demand bandwidth allocation procedure and reserves the necessary uplink time-slots for the collection period. The MSSP formulation provides a procedure for selecting and assigning users to available uplink time-slots, however, it may assign more than one user to a particular time-slot. Therefore, the FIS algorithm must have a *collision avoidance* strategy to limit the number of collisions. In order to avoid collisions, a user u_i chooses to transmit its feedback message during its assigned time-slot with probability $p_i(k)$. The procedure for choosing collision avoidance probability values is discussed in the subsequent sections. The set $\hat{\mathcal{U}}_k^j$ is constructed as:

$$\hat{\mathcal{U}}_k^j = \{u_i : x_{il}^k = 1 \cap \mathbf{1}(\mathcal{E}_i)\}, \tag{3.7}$$

where $\mathbf{1}(.)$ is the indicator function, $\mathcal{E}_i$ is the event that u_i transmits in period k, such that $E(\mathbf{1}(\mathcal{E}_i)) = p_i(k) \; \forall i$.

In the following section, we discuss the important properties of this algorithm formulation that enable efficient feedback suppression in the context of CSI collection.

3.3.2 Algorithm properties

MSSP is a NP-hard problem, and we assume that *best-fit-decreasing* (BFD) approximation is implemented for the solution [36], not only because it is one of the fastest heuristics, but also it has desirable properties in achieving the goals of the FIS algorithm.

Property 1 *BFD algorithm starts the selection and allocation of the items first from the ones with the largest weights. The algorithm always selects the first C_k^j largest items. The remaining items are selected only if there is room.*

Property 1 states that the first C_k^j users with the highest reported attenuation levels are always selected and assigned to a time-slot in the current collection period. This approach gives priority to users that have reported a high attenuation level in the past collection periods, over the users that reported a lower attenuation level. The remaining users face the possibility of suppression depending on their attenuation levels and the number of available time-slots. Assuming channel state varies slowly, this property effectively reduces the size of the user set that will be selected in the next period.

Property 2 *Items with larger weights, if selected, are more likely to occupy a bin that has few other items since they fill most of the capacity of a bin.*

Property 2 decreases the probability that a user with a high reported attenuation level will be involved in a collision in the current collection period, since it is more likely to be assigned to a time-slot with few other users. Property 1 and Property 2 together achieves the goal of reducing the feedback volume by restricting the transmissions to a group of users that are expected to attain the maximum attenuation level.

Property 3 *Items with smaller weights are only selected if there is available bin space after placement of items with larger weights.*

Finally, Property 3 states that users with lower reported attenuation levels are *suppressed* by the *selection* process. The suppression becomes more severe as the number of available uplink time-slots decreases, since the MSSP algorithm runs with fewer bins. It may be argued that selecting the largest C_k^j items and allocating them into bins one-by-one for the collection period k would have been a better strategy, since that would have avoided collisions between users. However, note that the allocation decision at the NOC is based solely on the *last successfully reported* attenuation level of each user, rather than the current value measured by the user. Such a strategy would cause a user with a previously low attenuation level to remain suppressed for a long time, even if it currently measures a high attenuation level, and would consequently reduce the robustness of FIS algorithm against the changes in the channel conditions. This observation also is the basis for our collision avoidance strategy which we describe in the next section.

3.3.3 Collision avoidance strategy

At every collection period, an active user only knows its channel assignment for that period and the value of its current attenuation level measurement. We assume that users can also listen to their own transmissions and detect collisions. Therefore, all users keep a record of the value of the last attenuation sample they successfully transmitted to the NOC. This value is equal to the value of the state variable kept at the NOC for the user, which the NOC always uses in calculating the maximum as well as in determining the channel assignment vector for the FIS algorithm. The transmission probability $p_i(k)$ of the user u_i should be proportional to the discrepancy between its current measured level $s_i(kT)$ and the last reported level, $\hat{s}_i[k-1]$, since the error would not only affect the selection and suppression criterion in the subsequent collection periods, but more importantly would increase the error between the true maximum and the calculated maximum. Therefore, a user assigned to an uplink time-slot chooses to transmit it feedback with probability proportional to the square of the difference between the current attenuation level and the last successfully reported value. The user behaves more aggressively in transmitting its feedback message, if there is a significant change in its condition. However, in order to put a cap on the aggressiveness of a user, the functional relationship between the square of the difference and probability of transmission should saturate as the difference goes to infinity. We choose the family of curves of the form:

$$f(y, z) = \frac{2}{(1 + e^{-y^2/z})} - 1.0, \tag{3.8}$$

78
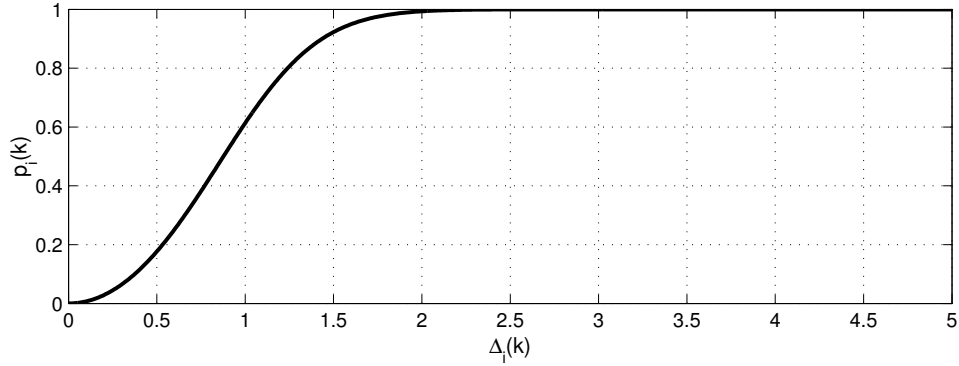
Figure 3.2: Probability of transmission as a function of Δ.

where, z determines how sharp the curve saturates. Let $\Delta_i(k)$ be defined as

$$\Delta_i(k) = |\hat{s}_i[k-1] - s_i(kT)| \tag{3.9}$$

for user u_i at collection period k. Then, the probability of transmission for the user in this interval is given by

$$p_i(k) = \frac{2}{\left(1 + e^{-\Delta_i(k)^2/z}\right)} - 1.0, \tag{3.10}$$

for a given z. In Fig. 3.2, we plot this curve for $z = 0.7$.

3.4 Performance analysis

3.4.1 Analysis results

In this section, we first look at the performance of the FIS algorithm for a fixed set of $U = 200$ active users inside one spot-beam coverage area over a simulated time of 1280 minutes. Therefore, in the remainder of this section, we drop the index for spot-beams. This test duration corresponds to a total of $K = 1200$ CSI collection periods

79

with an interval of $T = 64$ seconds, which is the granuality of our channel simulator (Section 2.5.2). The test is repeated 100 times, and average results are reported.

During the simulation, a separate instance of the channel simulator is run for each user, modeling its attenuation level. The attenuation time series generated for a user is independent from other users. In general, this assumption may not always be valid, since attenuation due to rain typically occurs in cells of certain diameter that may include more than one user. In this case, the users that are located close to each other geographically may have correlated attenuation levels. However, a study by Fukuchi [18] found that this correlation decreases rapidly with increasing distance, and that instantaneous rainfall rates at two locations more than about 100 km apart can be regarded as independent. Also, we may argue that such a correlation would only improve the performance of our algorithm by decreasing the number of users with different attenuation measurements.

In Fig. 3.3, we plot a realization of the time series of maximum channel attenuation calculated from the user samples (dashed-line),

$$s_{\max}[k] = \max_{i \in U}\{s_i(kT)\}, \; k = 1, 2, \ldots, K, \tag{3.11}$$

versus the maximum channel attenuation level calculated from state variables at the NOC (solid-line),

$$\hat{s}_{\max}[k] = \max_{i \in U}\{\hat{s}_i[k]\}, \; k = 1, 2, \ldots, K, \tag{3.12}$$

for $C = 2, 10, 20$ uplink time-slots. We observe that the time-series are in close agreement for $C = 20$ and $C = 10$, but the discrepancy becomes more significant for $C = 2$. We calculate the average absolute error between the two curves as a function of number
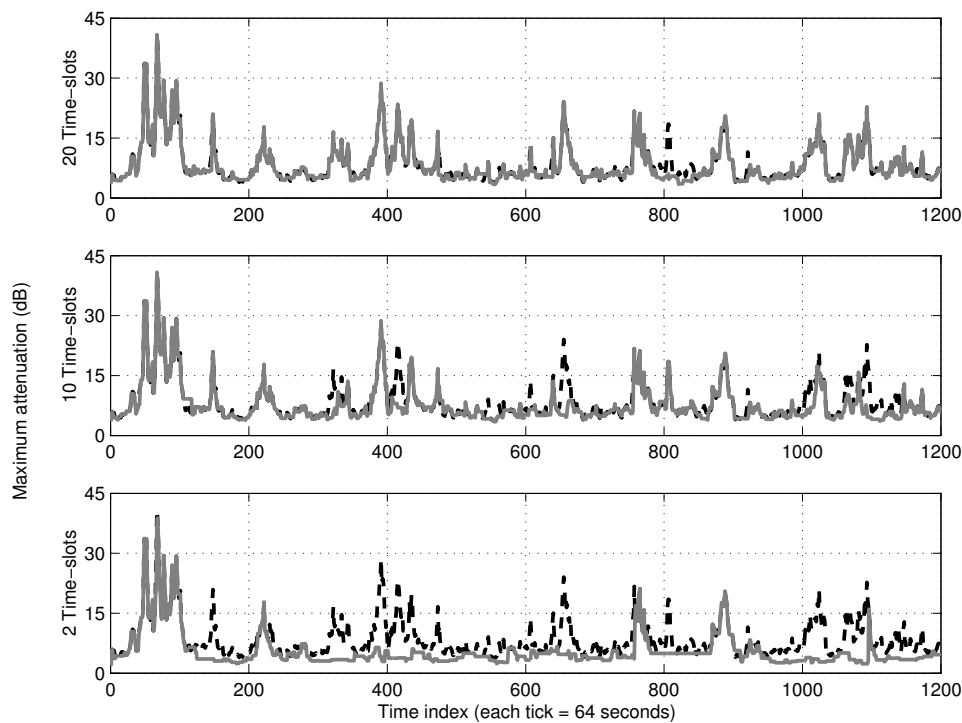
Figure 3.3: Time-series plot of actual channel attenuation level (dashed-line) versus the level calculated at the NOC (solid-line) for a set of 200 users.

of available uplink time-slots:

$$error(C) = \frac{1}{1200} \sum_{k=1}^{1200} |s_{\max}^C[k] - \hat{s}_{\max}^C[k]|, \tag{3.13}$$

which we tabulate in Table 3.1.

C	30	20	10	2	1
Error (dB)	0.3973	0.7117	1.4855	3.7733	4.5522

Table 3.1: Average absolute error.

In Fig. 3.4(a), we plot the fraction of the total users suppressed by the FIS algorithm as a function of the number of available uplink time-slots. A user's feedback may

81

be suppressed in one of the two ways: (i) it is not selected by the MSSP solution, or (ii) its feedback experiences a collision and is not reported successfully to the NOC. Figure 3.4(a) shows the contribution of the two cases in each instance. We observe that, $95 - 99\%$ of users are suppressed on the average by the FIS algorithm. Considering that the algorithm successfully keeps track of the maximum level for up to 10 time-slots (corresponds to 5% of total user population), this is a promising result showing that the algorithm correctly identifies the group of users that attain the maximum attenuation level. The type of suppression changes as the number of available uplink time-slots decreases. More users are suppressed through the selection process than channel collisions. When there are available time-slots, MSSP formulation selects and assigns more users. This leads to more collisions on the average per user in the assigned time-slots, however, also gives the opportunity to sample a larger user space if some feedback successfully reaches to the NOC. When there are very few time-slots, MSSP formulation suppresses most of the users. This leads to fewer collisions per user, however, increases the error if the wrong subset of users are assigned to the time-slots (compare to Table 3.1).

In Fig. 3.4(b), we look at the time-slot usage profile as a function of available uplink time-slots. We see that a time-slot is utilized successfully $30 - 35\%$ of the time. This percentage remains relatively constant independent of the ratio of the number of available time-slots to the number of users.

In order to assess the performance of our algorithm under different user group sizes, we plot, in Fig. 3.5, the average error in the calculated attenuation level versus the
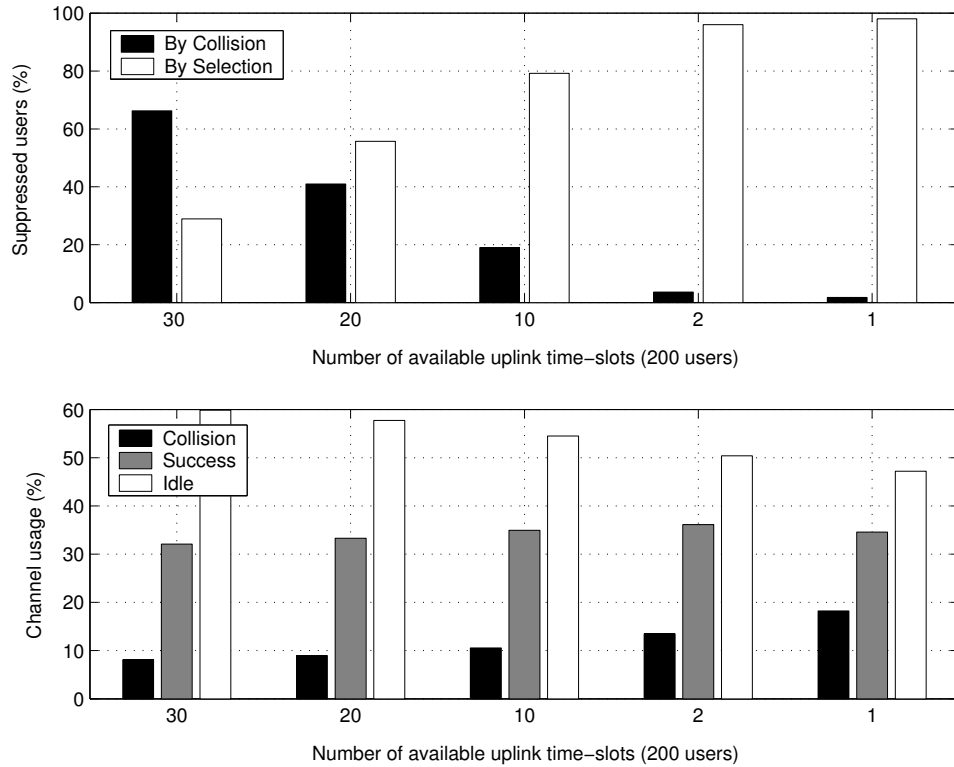
Figure 3.4: (a) Fraction of the users suppressed by the FIS algorithm as a function of the number of available uplink time-slots (200 users). (b) Channel usage profile as a function of the number of available uplink time-slots (200 users).
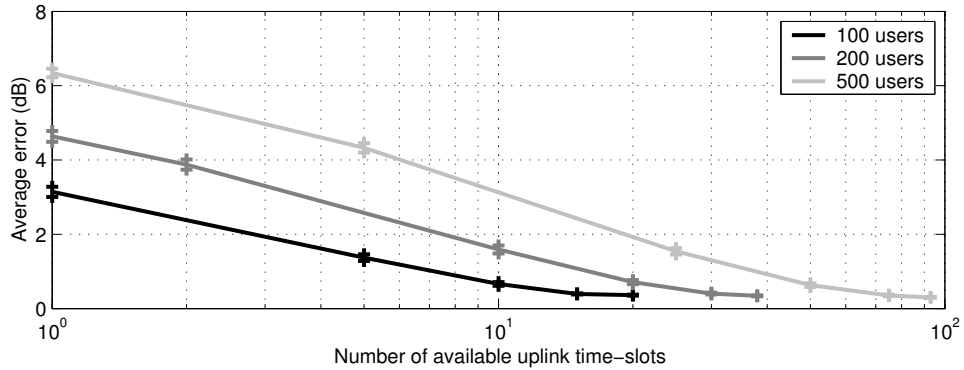
Figure 3.5: Average error versus number of available uplink time-slots as a function of the number of active users.

number of available uplink time-slots for a range of user groups. We observe that for a given number of time-slots, the average error of FIS algorithm increases as the number of users increase. This result is expected since the algorithm is forced to select and allocate users from a larger set to a fixed number of time-slots. Therefore, more users are suppressed by the algorithm, resulting in less accurate calculation of the maximum attenuation level. However, more importantly, if the ratio of available time-slots to the number of users is kept constant, then we observe that the algorithm performs equally well for all group sizes.

Although average error in the calculated attenuation level is an indicator of the performance of the FIS algorithm, from a system provider's or user's perspective, a more important performance metric would be the effect of this error as an input to a fade compensation algorithm. In order to assess this performance, we look at the use of the attenuation level as an input to a simple downlink power control scheme. We
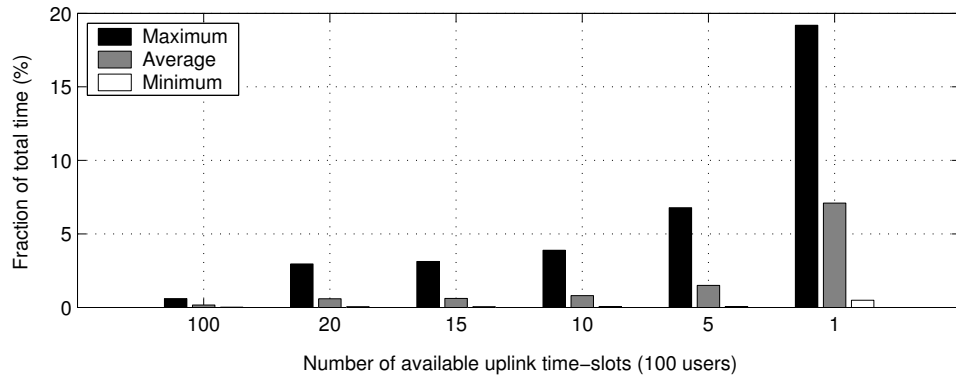
Figure 3.6: Percentage of total time BER target is not met by a user as a function of number of available uplink time-slots (100 users).

assume that the system has a 10dB link-budget power margin that is used to compensate for the signal attenuation due to atmospheric and rain fading. The system allocates this power margin based on the calculated maximum attenuation in order to meet a target bit-error-rate (BER). With no FIS in place, the system will not be able to meet this BER target, only when the maximum attenuation level is more than the 10dB power margin. However, with FIS in place, the error in calculating the maximum attenuation level may cause the system to underestimate the actual attenuation level. This would result in additional time instances when the BER target is not met by some users. On the other hand, over-compensation would result in signal enhancement, leading to an increased risk of co- and cross-channel interference. In Figures 3.6 to 3.8, we look at the maximum, average and minimum percentage of time a user will not meet the BER target as a function of the number of available uplink time-slots for user groups of 100, 200 and 500.
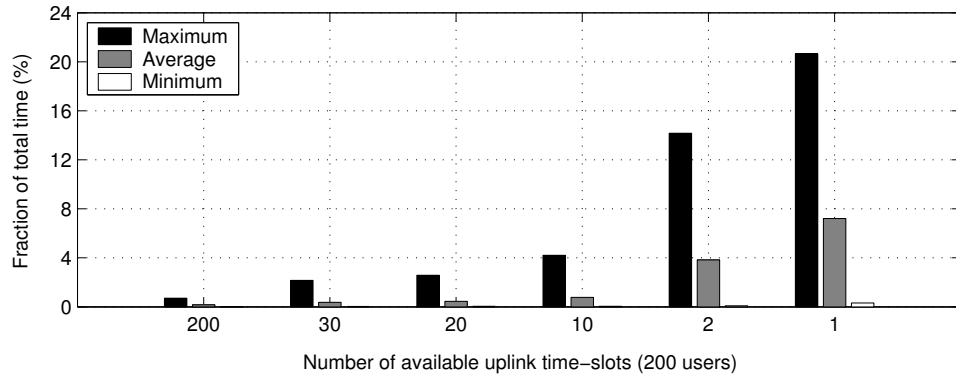
Figure 3.7: Percentage of total time BER target is not met by a user as a function of number of available uplink time-slots (200 users).
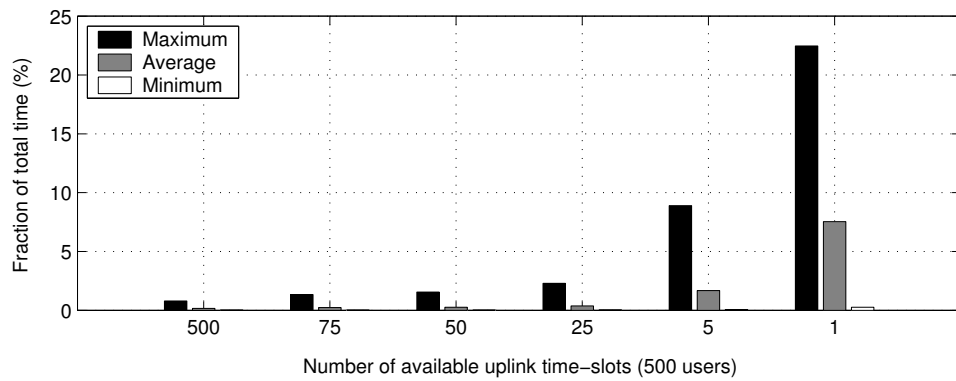
Figure 3.8: Percentage of total time BER target is not met by a user as a function of number of available uplink time-slots (500 users)

We observe that, average performance degradation experienced by a user remains acceptable even if the number of available time-slots is as low as 5% of the total user population.

3.5   Multicast-aware power allocation with FIS

In this section, we re-visit the multicast-aware power allocation policy that we have presented in the previous chapter. In the previous chapter, we assumed that the algorithm had perfect knowledge of the channel states for every spot-beam at the time of the optimization. In this section, we assume that CSI is collected from the active users inside the coverage area of every spot-beam, and the channel state of the spot-beam is determined by the user with the worst channel condition. In order to minimize the total number of uplink channels required for information collection in every spot-beam, we use the FIS algorithm and compare the results to the case where we had the perfect knowledge of the channel states.

We keep the same test setup as in the previous chapter. We create L_u unicast and L_m multicast connections between the NOC and spot-beam locations using the connection probability function given in Fig. 2.9. Once this process is completed, total number of unicast and multicast users in each spot-beam location is known. Instead of running a separate instance of our channel simulator per spot-beam, this time we run a separate instance of the channel simulator for every user. Since the time resolution of the channel simulator is $T = 64$ seconds, NOC updates the channel state of each spot-beam by

collecting CSI from all the active users inside that spot-beam location every T seconds. The test is performed for a simulated time of 720 minutes, where the channel states are updated for a total of $K = 675$ times. Each test is repeated for 100 different session configurations to achieve the statistical results. We use logarithmic weight allocation when running the power-allocation policy.

We look at three different scenarios for CSI collection:

- **Perfect CSI case:** Information is collected from every active user inside the spot-beam coverage area using separate return channels for every user. The channel state for spot-beam b_j at the end of collection period k, where there are $|\mathcal{U}_k^j|$ active users, is given by

$$s_{\max}^j[k] = \max_{i \in \mathcal{U}_k^j}\{s_i(kT)\}, \tag{3.14}$$

for $j = 1, 2, \ldots, M$, where M is the number of spot-beams. In this scenario, the number of time-slots used for each round of information collection is equal to the number of active users, — i.e. $C_k^j = |\mathcal{U}_k^j|, \ \forall j$.

- **Binpacking case:** Binpacking formulation is an extension of MSSP formulation, where the number of bins is not given, but determined by the minimum number required to fit all items. In our problem setting, this translates into having a variable number of allocated time-slots in each collection period and for every spot-beam, since the number of users and the channel state samples of these users vary over time. In binpacking case, at each collection time, all users are assigned to a time-slot, and potentially the number of allocated slots may be equal to the number of

active users. However, due to the possibility of collisions, the set of users that successfully report their CSI is a subset of the active user set. The channel state for spot-beam b_j at the end of collection period k is given by

$$\hat{s}^j_{\max}[k] = \max_{i \in \mathcal{U}^j_k}\{\hat{s}_i[k]\}, \tag{3.15}$$

for $j = 1, 2, \ldots, M$. The number of time-slots used in this case is always less than or equal to the number of active users, — i.e. $C^j_k \leq |\mathcal{U}^j_k|, \ \forall j$.

- **Subset Sum case:** This is the final variation, where we use the MSSP formulation, and the number of channels allocated is given and fixed through out the tests. In this scenario, at each collection time, not all users are allocated a channel. Moreover, of the users that are allocated a channel, not all successfully transmit their CSI information due to collisions. Hence, the channel state for spot-beam b_j at the end of collection period k is given by

$$\hat{s}^j_{\max}[k] = \max_{i \in \mathcal{U}^j_k}\{\hat{s}_i[k]\}, \tag{3.16}$$

for $j = 1, 2, \ldots, M$. The number of time-slots used in this case is fixed, — i.e. $C^j_k = C, \ \forall j$.

In the next two sections, we look at the performance of BAS-I and BAS-II policies for each of the above scenarios.

3.5.1 Performance of BAS-I policy with FIS

In this section, we compare the performance of BAS-I policy under different feedback implosion suppression scenarios. In order to assess the changes in performance of our policy, we first have to compare the extend of suppression introduced by each FIS scenario, and the corresponding time-slot usage profile. In Fig. 3.9, we look at the time-slot usage profile under binpacking scenario, when there are $L_u = 350$ active unicast connections and $L_m = 5$ to 30 active multicast groups. Under binpacking scenario, the number of slots used per collection period and per spot-beam varies over time. In Fig. 3.9 we plot three metrics: (i) the maximum of the average time-slot usage over all spot-beams, given by

$$C_{\text{max}} = \max_j \frac{1}{K} \sum_{k=1}^{K} C_k^j, \tag{3.17}$$

(ii) time-slot usage averaged over all spot-beams, given by

$$C_{\text{avg}} = \frac{1}{M} \sum_{j=1}^{M} \frac{1}{K} \sum_{k=1}^{K} C_k^j, \tag{3.18}$$

and (iii) the average suppression efficiency, which is the ratio of average number of time-slots used to average number of active users. We observe that, the FIS algorithm achieves a suppression ratio of $70\% - 74\%$, but maximum number of time-slots used by a spot-beam may go up to 3 times higher than the average time-slot usage. In Fig. 3.10, we look at the time-slot usage profile under the subset sum scenario. For subset sum scenario, we have to specify the maximum number of bins available to the FIS algorithm. In order to test the performance to the limit, under this scenario, we set the maximum
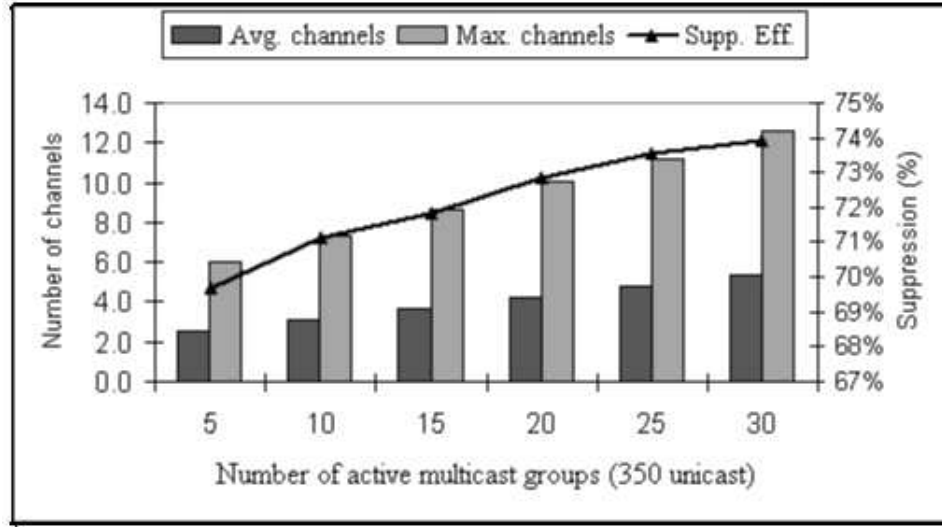
Figure 3.9: Time-slot usage profile and suppression efficiency of spot-beams as a function of active multicast groups when using binpacking formulation.

number of time-slots available to every spot-beam to be equal to the average number of channels used for the binpacking case, rounded off to the nearest integer, — i.e. set $C_k^j = C_{\mathrm{avg}} \ \forall j, k$. We observe that, average number of channels used by a spot-beam is still lower than the maximum we set, since for some very lightly loaded spot-beams not all bins are required in the solution of the subset sum problem. The suppression ratio is forced to be higher, since maximum number of channels have a tight cap, going up to a maximum of 78%.

Under both binpacking and subset sum scenarios, the FIS algorithm achieves very good suppression ratios. However, it is important to test how the performance of the power allocation policy, which relies on the channel state information obtained by the FIS algorithm, is affected compared to the perfect CSI case. The change in performance

Figure 3.10: Time-slot usage profile and suppression efficiency of spot-beams as a function of active multicast groups when using subset sum formulation.

is a consequence of the fact that with partial CSI, the power levels calculated by the policy may over- or under-estimate the actual power requirements of the spot-beams. As a result of this, the rates allocated to particular flows may not reflect the actual supportable rates.

In Fig. 3.11, we look at the average rate change in the sustainable session rates of unicast flows for each of the feedback collection scenarios. We observe that the performance under binpacking scenario is on a par with the perfect CSI scenario. However, the performance under subset sum scenario remains sub par for all cases. Similar results are observed for multicast flows as well. These results are attributed to the fact that under binpacking scenario, the FIS algorithm has access to as many time-slots as needed under time-varying conditions. Therefore, binpacking scenario closely follows the per-

Figure 3.11: Average rate change in the sustainable session rates of unicast flows averaged over all active unicast flows under the BAS-I policy for each of the feedback collection scenarios.

Figure 3.12: Average rate change in the sustainable session rates of multicast flows averaged over all active multicast flows under the BAS-I policy for each of the feedback collection scenarios.

Figure 3.13: Percent 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 for each of the feedback collection scenarios.

formance of perfect CSI scenario. Under subset sum scenario there is a tight cap on the maximum number of available time-slots. This inflexibility under changing propagation conditions causes degradation in performance. However, improvement in the average session rates of unicast and multicast flows remains higher than 15% and 50%, respectively even though the average number of allocated time-slots is in the range of 22% to 28% of the average number of users, which shows that FIS algorithm works well in conjunction with our power-allocation policy. A similar performance trend is observed when we look at the fraction of time session rates of unicast and multicast flows remain over the EAS session rates under the BAS-I policy (Figures 3.13 and 3.14).

Finally, we look at the change in the value of cost function of the optimization for

94

Figure 3.14: Percent 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 for each of the feedback collection scenarios.

each of the feedback collection scenarios. We observe that the optimization performs equally well in minimizing the cost function, reducing its value to $90 - 95\%$ of the starting value for all scenarios. Therefore, the performance degradation we see is not a result of the performance of the optimization, but because of the discrepancy between the attenuation level values the algorithm uses versus the ones that sessions actually experience.

In the next section, we look at the performance of BAS-II policy under different feedback collection scenarios.

Figure 3.15: Percent change in the value of the cost function after optimization with respect to the starting value under BAS-I policy for each of the feedback collection scenarios.

3.5.2 Performance of BAS-II policy with FIS

Figures 3.16 to 3.19 show similar performance trends in terms of time-slot usage, suppression efficiency, and average rate change in the sustainable session rates of unicast and multicast flows. Binpacking case performs on a par level with perfect CSI case, while subset sum case causes performance degradation in the average rate increase that the flows experience as a result of the optimization. A more important performance change in the results of BAS-II policy comes in the fraction of time flows attain rates that are higher than EAS rates after the optimization. BAS-II policy guarantees that the session rates of all flows remain at a rate higher than the EAS session rates at all time under perfect CSI. However, in Figures 3.20 and 3.21, we observe that BAS-II policy no
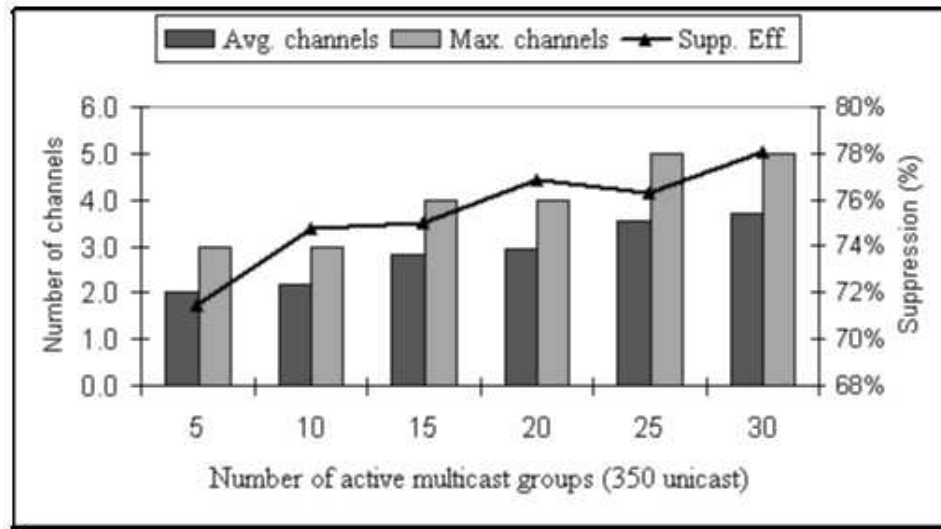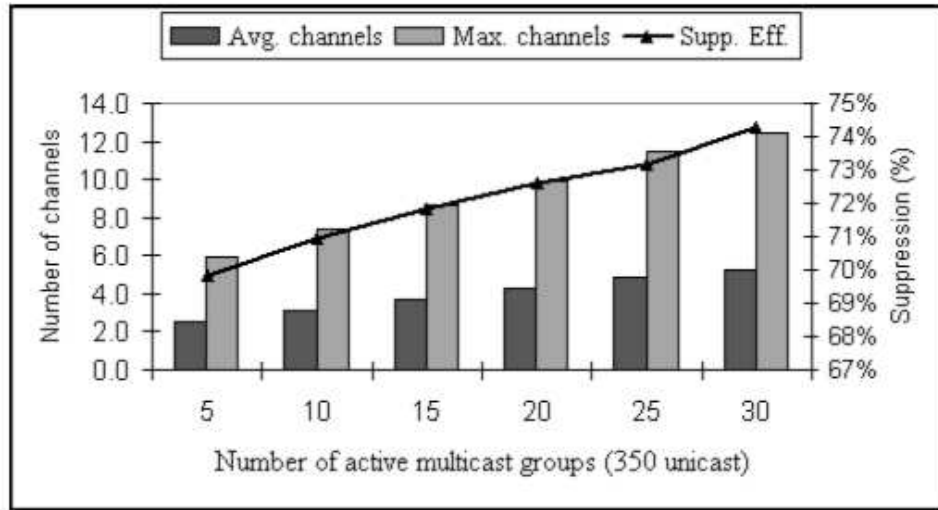
Figure 3.16: Time-slot usage profile and suppression efficiency of spot-beams as a function of active multicast groups when using binpacking formulation.

longer has this guarantee due to the discrepancy between the attenuation level values.

Finally, we look at the change in the value of cost function of the optimization for each of the feedback collection scenarios. We observe that under BAS-II policy, optimization reduces the value of the cost function to $40 - 65\%$ of the start value for all scenarios. We observe that this is considerably less than the values attained under BAS-I policy. This result, however, is a direct consequence of the fact that BAS-II policy works with a much smaller search space due to the restrictions imposed by the minimum power requirements, as discussed earlier in the previous chapter.
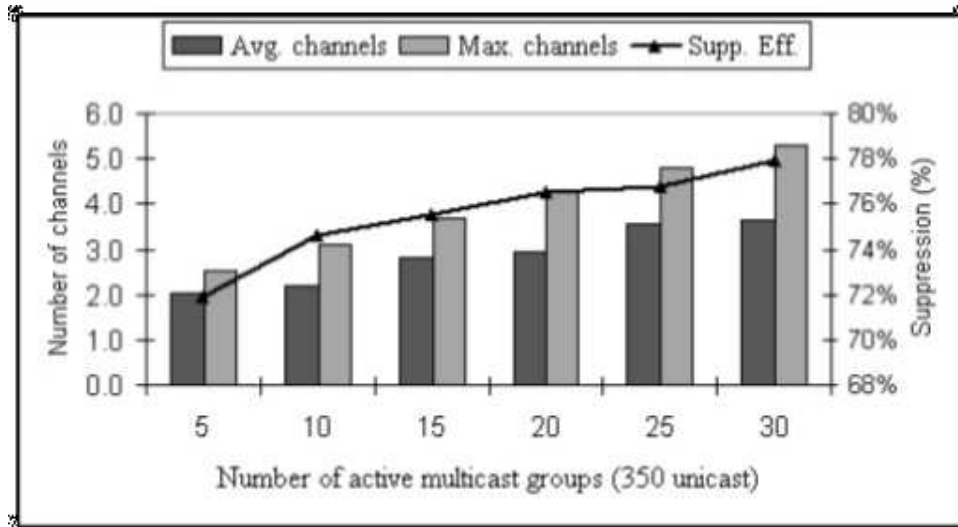
Figure 3.17: Time-slot usage profile and suppression efficiency of spot-beams as a function of active multicast groups when using subset sum formulation.
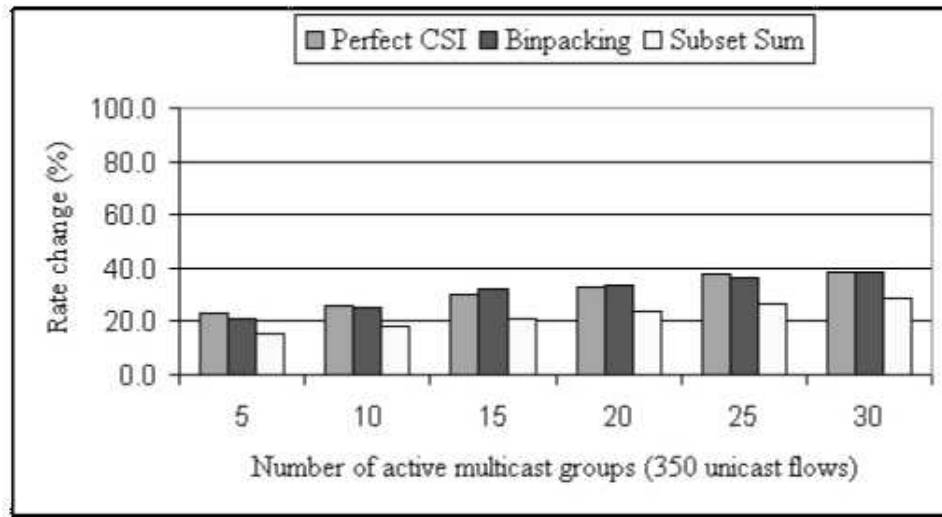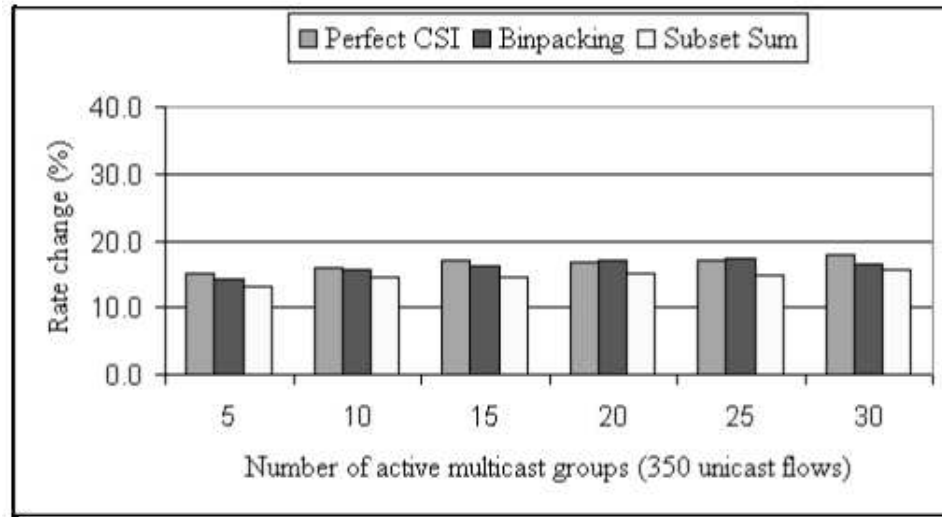

Figure 3.18: Average rate change in the sustainable session rates of unicast flows averaged over all active unicast flows under the BAS-II policy for each of the feedback collection scenarios.
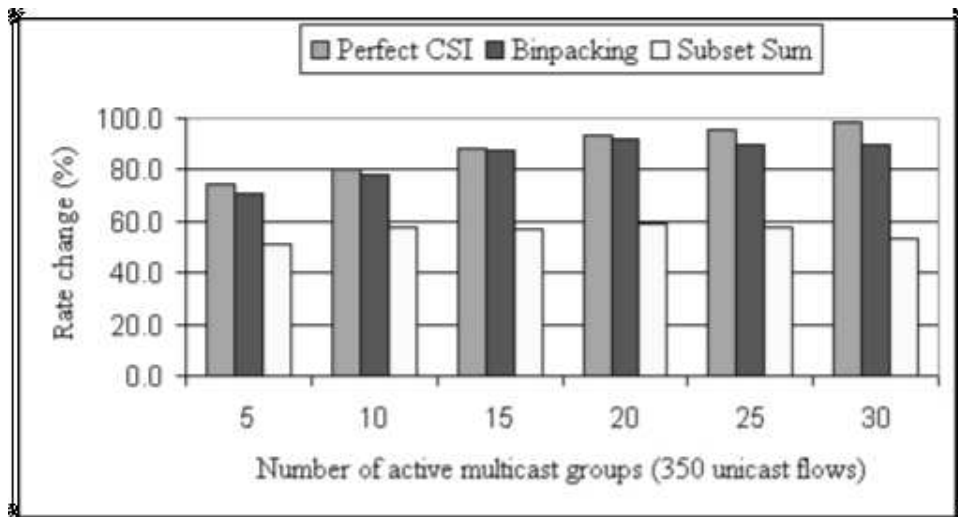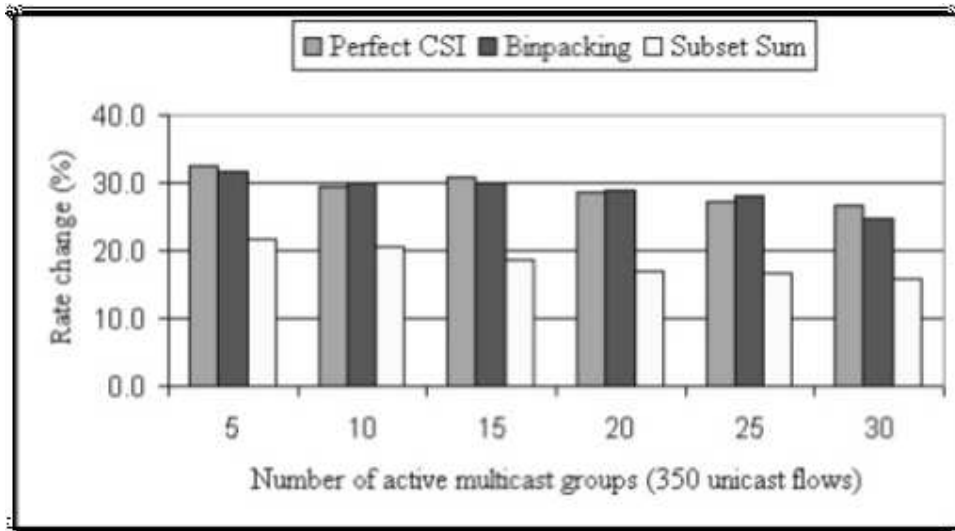
Figure 3.19: Average rate change in the sustainable session rates of multicast flows averaged over all active multicast flows under the BAS-II policy for each of the feedback collection scenarios.
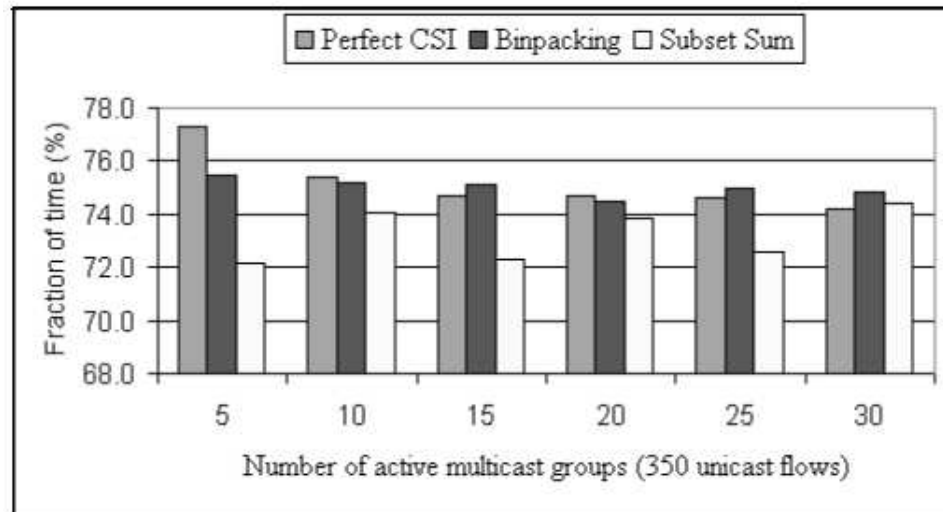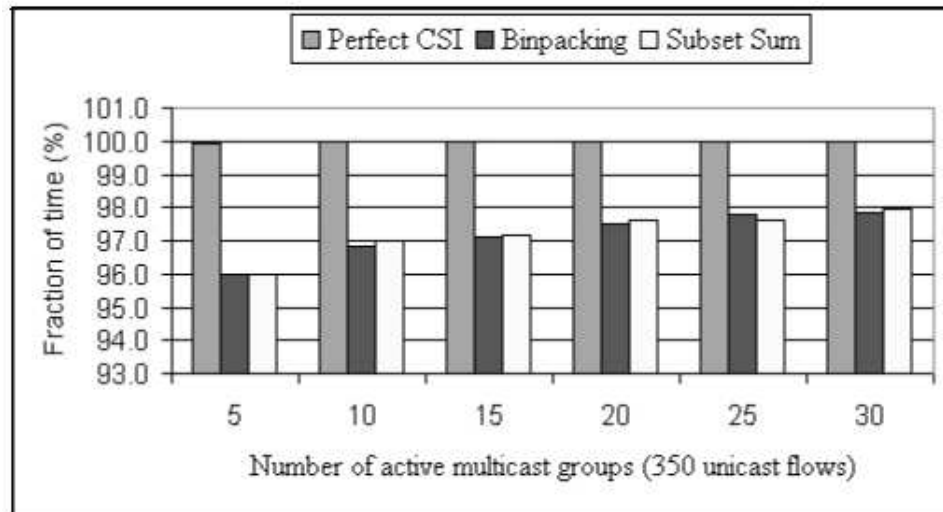

Figure 3.20: Percent of total test time BAS-II session rates of unicast flows are equal to or better than EAS session rates averaged over all active unicast flows for each of the feedback collection scenarios.
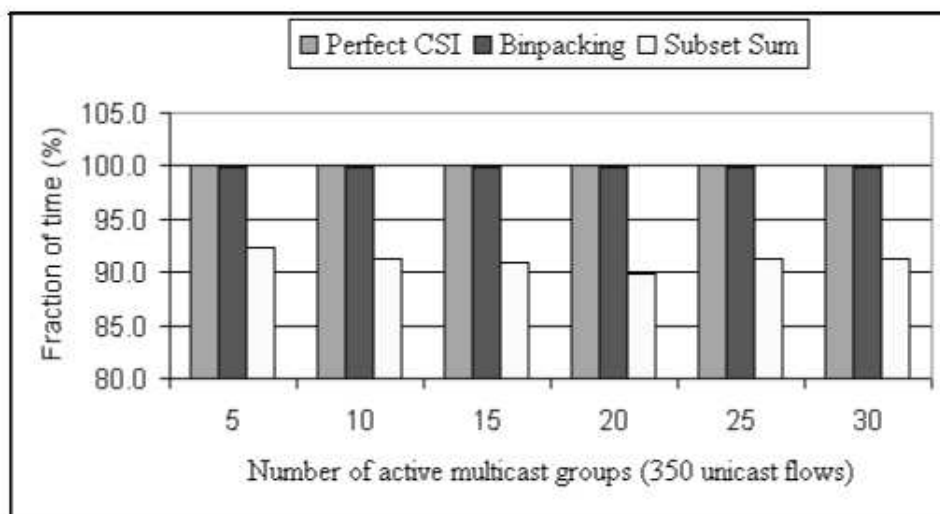
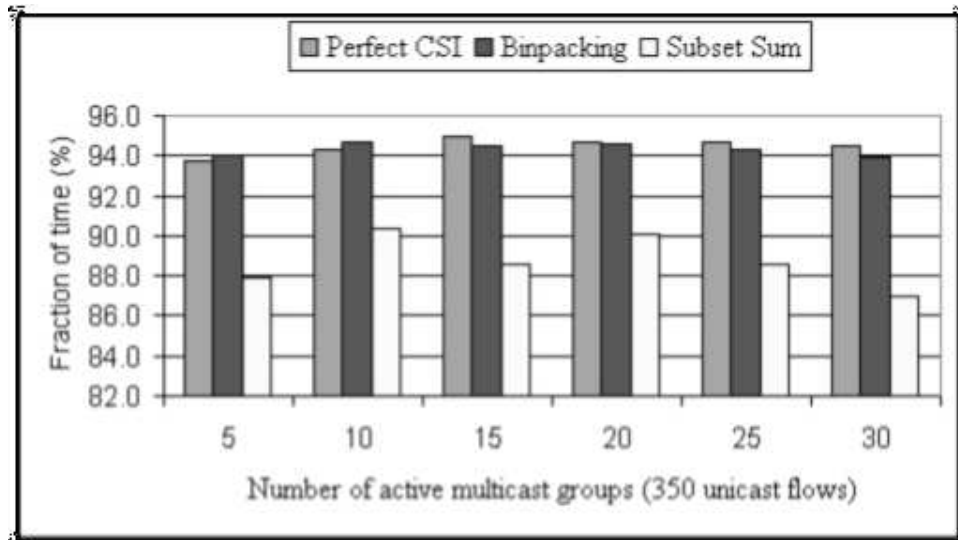Figure 3.21: Percent of total test time BAS-II session rates of multicast flows are equal to or better than EAS session rates averaged over all active multicast flows for each of the feedback collection scenarios.
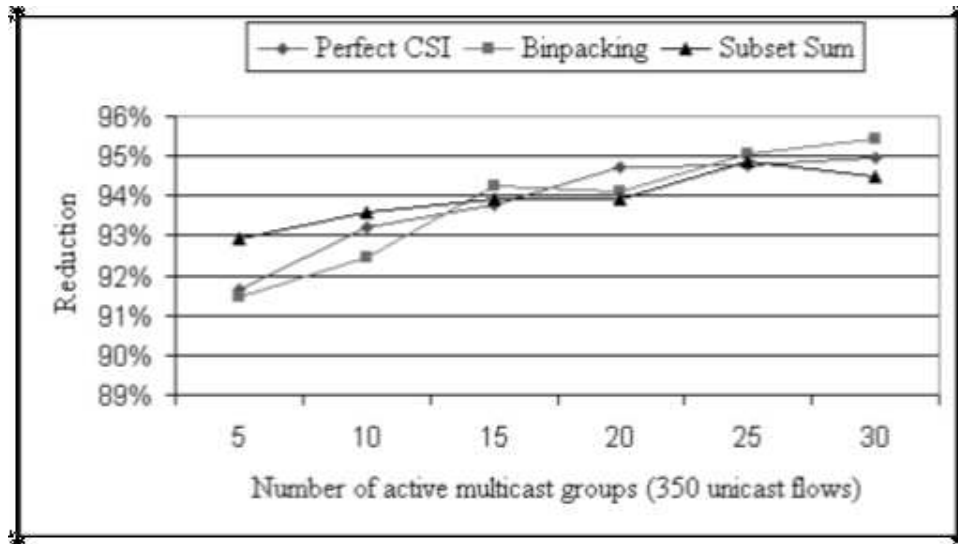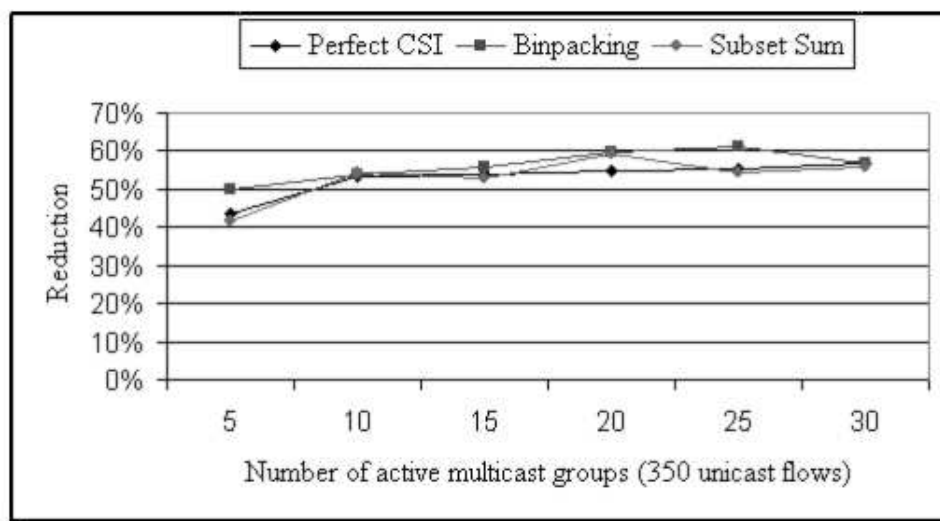

Figure 3.22: Percent change in the value of the cost function after optimization with respect to the starting value under BAS-II policy for each of the feedback collection scenarios.

100

3.6 Summary

In this chapter, we presented a feedback implosion suppression algorithm that uses a multiple subset sum (MSS) formulation in order to select and assign users to a limited number available uplink channels (time-slots). MSS formulation is based on the observation that for most of the multicast scenarios where user feedback is used to probe the condition of the members of the multicast group, it is sufficient to only track the behavior of only a certain subset of users (e.g. subset of users with the worst channel conditions). In this case, it is possible to prioritize the users based on the *importance* of their feedback to the task at hand, and allocate them accordingly. However, in most of the cases, the priority level of a user is not known to or can not be communicated to other users to help in the allocation process. Moreover, the decision mechanism may have partial or outdated information on the group. Therefore, determining the right subset of users to collect feedback is a challenging task.

Our formulation successfully achieves this task, and successfully tracks the user(s) with the highest channel attenuation, when the number of available uplink time-slots is larger than 5% of the multicast group size. In the MSS formulation, the number of available uplink slots is given and fixed. However, it may not be possible to know, in advance, the required number of uplink time-slots for acceptable performance. Therefore, in this chapter, we also investigated the extension of the algorithm to binpacking formulation, where the number of uplink time-slots is not fixed. The algorithm uses as many uplink time-slots as needed in order to allocate all active users according to the

selection mechanism. The binpacking formulation outperforms the MSS formulation due to this flexibility. However, the number of uplink time-slots used by the binpacking formulation can potentially be as large as the number of active users, in which case the suppression ratio goes to zero. On the other hand, binpacking formulation can be used as a guideline to set the maximum number of available time-slots in MSS formulation. Finally, in this chapter, we integrated our feedback implosion suppression algorithm to the power allocation policy of the previous chapter. The channel state information required by the power allocation policy is collected from the users through the feedback implosion algorithm. We showed that it is possible to reach suppression ratios of $70 - 78\%$ while maintaining the initial performance of the power allocation policy.

Chapter 4

Integrated FEC Coding for Reliable Multicast Transport

4.1 Introduction

In the previous chapters, we have discussed the advantages of supporting multicast services over satellite networks, and presented two features that enable multicast-aware power allocation and efficient feedback collection. These features provide support for distribution of multicast content to multiple users. In addition, some of the multicast applications in the Internet, such as distributed computing, software updates, distance-learning, and telemedicine, may also require or benefit from *reliable* multicast communication. Therefore, in this chapter, we focus on some of the issues related to the reliable delivery of multicast content over satellite networks.

Reliable delivery requires that all data segments transmitted over the network are received correctly by all receivers. In order to achieve this goal, missing or corrupted data segments must be detected and recovered. In traditional schemes, such as the automatic-repeat-request (ARQ), receivers detect missing data segments and automatically request for retransmissions by transmitting a feedback message of one of the following forms to the sender:

- **Positive acknowledgments (ACK):** Receivers return ACK packets to the sender,

indicating the data segments that have been received correctly.

- **Negative acknowledgments (NACK):** Receivers return NACK packets to the sender, indicating only the data segments that are missing.

- **A hybrid approach:** Receivers use both ACK(s) and NACK(s).

In these schemes, a missing ACK or a received NACK for a particular data segment acts as a trigger, in response to which the sender initiates retransmission of the missing segment. However, a recovery mechanism that is based solely on ARQ does not perform well in case of multicast communication. ACK-based loss notification causes *feedback implosion*, since every member of the multicast group has to separately acknowledge the successful receipt of every data segment. Additionally, the sender is required to keep track of the size, and the current state of the members of the multicast group, in order to identify the last data segment correctly received by all. Both problems negatively affect the efficiency and scalability of the protocol. NACK-based loss notification alleviates some of these problems. The performance comparison study presented in [37] confirms that NACK-based multicast transport protocols deliver better performance than their ACK-based counterparts in wireline terrestrial networks. In NACK-based feedback, receivers detect missing packets by checking for gaps in the packet sequence numbers and report missing segments to the sender. The sender neither needs to know the size of the receiver group, nor keeps track of the current state of every receiver in its group. Also, the number of NACK packets is expected to be less than that of ACK packets at low error rates. A disadvantage, however, is that it is more difficult for the sender

to know when it can free the transmission buffers, since not receiving a NACK may either indicate successful transmission of the corresponding data segment, or loss of the notification packet itself.

These performance related issues are only aggravated over satellite networks. For a GEO satellite, the propagation delay between the user terminals and the satellite is about 125ms, which adds up to a round-trip-time (RTT) of about 500ms for a two-way system over the satellite segment. Therefore, ARQ-based recovery of individual data segments increases transmission delay significantly. In addition, retransmission of individual data segments, in response to individual receiver requests, does not make good use of the broadcast capabilities of the satellite channel, because retransmitted segments only benefit one receiver while discarded by many others. Moreover, because of the large bandwidth-delay product typical to such networks (e.g. 20Mbps·500ms for a GEO satellite operating at K/Ka-band frequency), the transport protocol would put a large amount of data onto the network before getting responses for old segments. Therefore, the number of outstanding data segments that have not yet been acknowledged would be large, forcing the protocol to keep a large amount of data in the transmission buffers. Finally, deep channel fading due to atmospheric effects may lead to bursty periods where several data segments are lost back to back. Recovery of these data segments over the high latency network through the repeat-request process would be inefficient, and may trigger feedback implosion even in the NACK-based case.

Although it is not possible to completely abandon an ARQ-based scheme, and still

provide full reliability, we conclude that in order to support reliable multicast services over a satellite network, the transport protocol should act proactively to protect transmitted data against losses, and consequently, to minimize the need for retransmissions. Forward error correction (FEC) coding is a well-known technique for protecting data against losses and corruption [38–40]. FEC coding can be applied against bit-level corruption, in which case the FEC encoder generates $(n - k)$ parity symbols (e.g. bytes) for a k-symbol packet. Such codes are capable of recovering the k-symbol packet if the number of corrupted symbols are less than or equal to $t = (n - k)/2$. This, in effect, decreases the effective loss rate of the channel and causes the upper layers to see a more reliable network [41]. However, unrecoverable packets are still retransmitted from the source, and different packets are retransmitted in response to requests of different receivers. An alternative application of FEC codes to multicast protocols is as an erasure code at the packet level. [40, 42–44]. In this case, lower network layers detect and discard corrupted packets (possibly by application of a separate FEC coding), and the erasure code works at the higher network layers (e.g. transport layer) to recover missing packets. The FEC encoder generates n encoding packets from k original data packets, and the original k data packets can be recovered as long as k out of the transmitted n encoding packets are received correctly. Use of FEC coding at the packet level offers some solution to a number of problems.

One immediate benefit of packet level FEC coding is reducing the number of missing data packets. This minimizes the number of feedback messages (e.g. NACK),

as well as the need for retransmissions, and thereby improves efficiency and scalability. FEC coding has an additional benefit in case of satellite multicast services. Unlike unicast communication, where the retransmission of a data packet exclusively benefits the receiver that has made the request, in multicast, the protocol has the option to either unicast or multicast the retransmitted packet. Multicasting of retransmitted packets would benefit multiple receivers in case of correlated loss. However, it may as well degrade channel utilization in case of uncorrelated loss. Because of the broadcast nature of the satellite channel, however, all packets can be received by all members of the multicast group. Therefore, transmission of encoding packets significantly improves the channel utilization, because FEC encoding packets can be useful to all receivers in the recovery process of an encoding block, even when different encoding packets are received at different receivers. Also, the protocol becomes more scalable in the number of receivers. The protocol does not need to know which packets have been lost by the receivers of the session per encoding block, but only the maximum number of packets lost by any receiver of the session. Therefore, the feedback from a single receiver has been reduced [42, 45–47].

Many of the recent reliable multicast transport protocols integrate some form of packet level FEC coding [48–50] in their design. However, implementation of packet level FEC coding is difficult for typical data packet sizes. Also, number of encoding packets that can be transmitted against packet losses is limited to the block size of the code. Therefore, there has been an effort in the research community that is initiated

107

primarily by the IETF working groups, to define schemes that could work with large packet sizes and over large block sizes [43, 44]. As an example, in [51] authors propose a scheme using Tornado codes, which are capable of working over large block sizes.

In this chapter, we investigate the potential benefits of using packet level FEC coding for reliable multicast delivery over a satellite-based communication system. We consider application of a special form of packet level FEC codes, namely the LT codes [52, 53], which are capable of generating a large number of encoding packets for the same group of input packets. We discuss the possibility of adjusting the number of encoding packets adaptively in response to changing channel conditions and receiver feedback. In the next section, we start by describing our target network scenario, and discuss where the application of packet level FEC would fit within the network layers.

4.2 Network scenario

We assume that the multicast content (software, webpage, file) is located at a server on an IP based terrestrial network. This content is to be distributed reliably to multiple user locations, which are within the coverage area of the satellite network. In this scenario, users that are equipped with two-way direct communication terminals access the terrestrial network through a gateway node we refer to as the Network Operations Center (NOC). Therefore, the NOC is responsible for fetching the content from the server, and then distributing it over the satellite network (Fig. 4.1).

We assume that the session is established using the idea of connection split-

Figure 4.1: Satellite communication system architecture.

ting [54, 55] that has recently been of much interest to the research community. In this scenario, the session is split into three components: (a) a TCP/IP connection is established between the server and the satellite gateway, (b) over the satellite portion of the network, a satellite reliable multicast transport protocol is run, and (c) between the user terminal and the user machines another TCP/IP connection is established. We assume that data reliability is ensured natively by the TCP protocol in sub-connections (a) and (c). Therefore, we are primarily interested in the behavior and performance of the transport protocol over the satellite portion of the network. Our decision to go with the idea of connection splitting in this type network scenario is motivated by several observations. The first one stems from the particular structure of the network topology where the connections from the server to the gateway, and from the user terminal to the user machines are essentially point-to-point. Inside the satellite core, on the other hand,

the star topology allows satellite to push content in single hop to all users. Secondly, the transmission characteristics, such as the transmission delay and error patterns, over the satellite portion are widely different than those that are over the terrestrial network. Finally, satellite service providers are more likely to favor this setup compared to an end-to-end solution between the server and user machines, since the former allows them to have total control over the traffic flow inside their network for activities ranging from resource allocation, and security to billing.

Under this setup, we assume that Digital Video Broadcasting (DVB) [56] standard for data broadcasting is implemented for delivering data packets inside the satellite core network, although, the discussion in the remainder of this chapter does not depend on the behavior of a particular packet delivery scheme. DVB standard for data broadcasting allows encapsulation of IP datagrams into MPEG-2 Transport Stream Packets using the Multiple Protocol Encapsulation (MPE) standard [56]. The incoming stream of IP datagrams are segmented into 188-Bytes (4 Bytes packet header and 184 Bytes payload) transport packets at the satellite gateway. The DVB channel encoder encodes each packet, first using a shortened (204,188) RS code (adds 16 parity bytes to each transport packet), and then, passing the packet bit stream through an interleaver. The interleaved symbols are encoded by a rate punctured convolutional encoder capable of supporting variable rates of 1/2, 2/3, 3/4, 5/6, and 7/8, and transmitted through the physical channel.

In the rest of this chapter, we assume that DVB channel coding is capable of detecting all corrupted packets and discarding them. Therefore, only correctly received
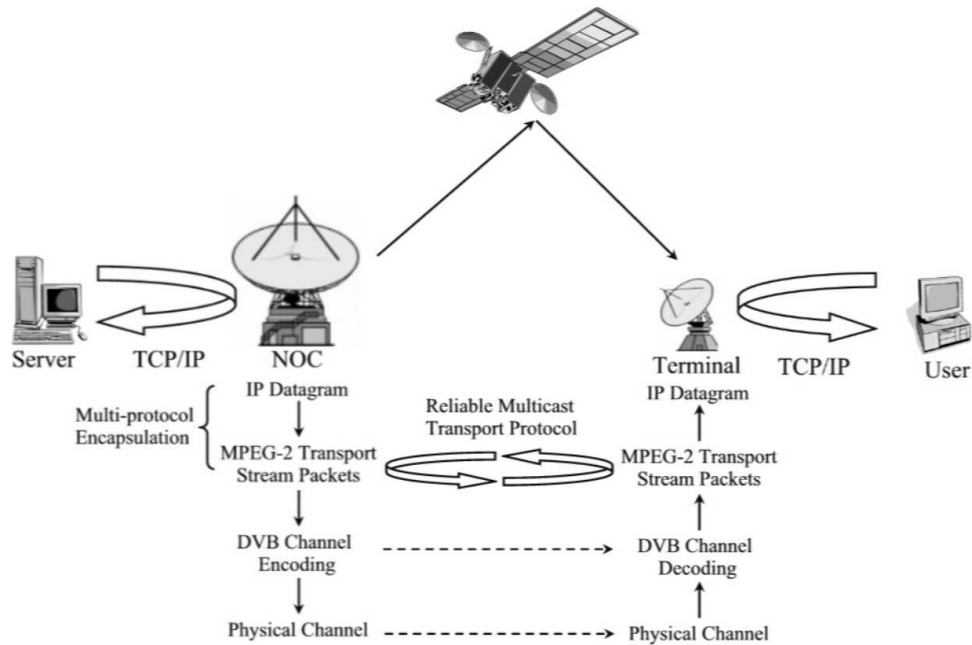
Figure 4.2: Proposed breakup of the communication session.

packets are forwarded to the MPEG transport layer, and missing packets have to be
recovered at this point. Reliability at the MPEG transport layer is assisted by the appli-
cation of the packet level FEC coding as an erasure code (Fig. 4.2). In the next section,
we give an overview of the structure of the class of packet level codes we consider for
application under this network scenario, and define how coding is integrated into the
transmission of data packets.

4.3 Packet level FEC coding for reliable multicast communication

Application of packet level FEC coding using some of the traditional block codes, such
as Reed-Solomon codes, are limited because of the difficulty in implementing them
for large symbol sizes, and the restrictions on the encoding block sizes. In order to

111

be applicable to our scenario, the FEC coding should be able to accept packet sizes of several hundred bytes as input symbols. Additionally, the code should be able to generate a large number of encoding packets, preferably, on demand, to overcome bursty packet losses.

In order to achieve some of these goals, we make use of a special class of FEC codes, referred to as the LT codes [52, 53] in the literature. Unlike block codes, there is no predetermined number of encoding symbols that can be generated for LT codes. Instead, the LT encoder can generate as few or as many unique encoding symbols as required on demand. Moreover the input symbol sizes for the LT encoder can be arbitrary. The construction of the LT codes allows end-users to recover an original batch of k transport packets, if they manage to accumulate a slightly larger set of $K = (1 + \epsilon) \cdot k$ of the encoding packets that are generated from the same input packet batch. Additionally, the number of encoding packets that can be generated from the same input packet batch is typically much larger than the size of the batch. The benefit of this construction is two-fold. First, it allows the satellite gateway to transmit additional encoding packets as long as there are users that have not yet accumulated enough packets from the batch. Secondly, it simplifies the user requests for additional packets, since now only information that has to be forwarded to the satellite gateway is comprised of the number of additional encoding packets (from the same batch) that is required by each receiver until all receivers complete the reception. This construction also makes very good use of the broadcast nature of the satellite channel since every transmitted encoding packet bene-

fits all the users equally. In the next section, we provide an overview of the construction

of LT codes.

4.3.1 Construction of LT codes

The process of generating an encoding packet from a batch of input packets is concep-

tually easy to describe [53]. We assume that data to be distributed can be organized into

B batches of k input packets. Let l be the size of an input packet. For an input batch, an

encoding packet of size l is generated by the following procedure (Fig. 4.3):

1. From a degree distribution (a good degree distribution is the key to the operation

   of LT codes and it will be discussed in more detail), randomly choose the degree

   d of the encoding packet.

2. Choose, uniformly at random d distinct input packets as neighbors of the encoding

   packet.

3. The value of the encoding packet is the bitwise exclusive-or (XOR) of the d neigh-

   bors (d input packets).

When using the encoding packets to recover the original input packets, the decoder

needs to know the degree and the set of the neighbors of each encoding packet. This in-

formation can be communicated to the decoder in many different ways but one efficient

example is to use a key to seed the random number generators that generates the degree

value and the neighbor list. A good candidate for the key is the sequence number in the

header of each encoding packet. Then, the sequence number of each encoding packet can be used to obtain its degree and neighbor information.

The decoder uses the following rule to recover input packets. If there is at least one encoding packet that has exactly one neighbor (degree $d = 1$), then the neighbor (corresponding input packet) can be recovered immediately since it is the copy of the encoding packet. The recovered input packet is XOR'ed into any remaining encoding packets that also have that input packet as a neighbor. Recovered input packet is removed from the neighbors of these encoding packets and the degree of each such encoding packet is decreased by one to reflect this removal. The decoding process continues at each step by processing encoding packets with degree one, and completes successfully if all input packets are recovered in this manner. Therefore, the initial degree distribution has to be such that there exists a non-empty set of encoding packets with degree one at the end of each step. If the set of encoding packets with degree one becomes empty before all input packets are recovered, then the decoding fails.

The LT process is a generalization of the classical process of throwing balls randomly into bins. A classical analysis shows that $K = \ln(k/\delta) \cdot k$ balls are necessary on the average to ensure that each of the k bins is covered by at least one ball with probability of at least $1 - \delta$. In the analysis of the LT process, encoding packets are analogous to balls and input packets are analogous to bins and the process succeeds if at the end all input packets are covered. The analysis in [53] shows a construction of a degree distribution such that the input batch of k packets can be recovered using

114
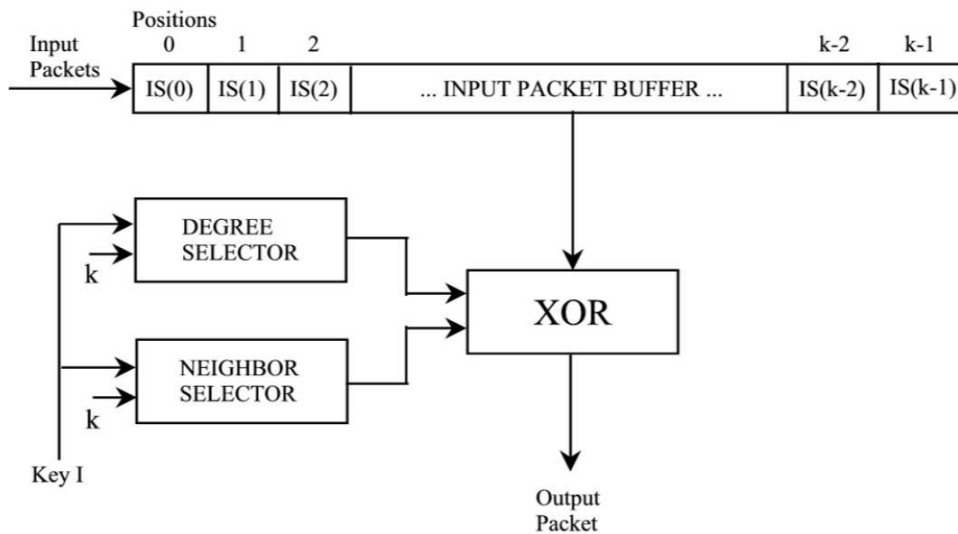
Figure 4.3: Block diagram for LT encoding.

$K = k + O(ln^2(k/\delta) \cdot \sqrt{k})$ encoding packets. In the next section, we discuss how LT codes can be integrated into the transmission of MPEG-2 transport stream packets.

4.3.2 Integration of LT codes to packet transmission

In this section, we will describe how a generic transport protocol could incorporate LT codes into the transmission of packets in order to provide reliability over the satellite segment of the network. We assume that data can be organized into B batches of k MPEG-2 transport stream packets at the NOC following the multi-protocol encapsulation. Let K be the number of encoding packets, such that the probability of a receiver recovering the original batch of k input packets after accumulating K encoding packets is greater than $1 - \delta$, for some $\delta > 0$. Each batch is encoded using LT coding, and encoding packets are forwarded to the DVB channel encoder for transmission.

Transmission of batches proceeds in transmission rounds. In transmission round, r, where $1 \leq r \leq B$, the protocol attempts transmission of batch r for the first time. Because of the construction of LT codes, users have to accumulate at least K encoding packets in order to successfully recover any input batch. Therefore, for batch r, the protocol has to transmit at least K encoding packets in this first attempt. The number of encoding packets transmitted in the first attempt of a batch is controlled by a transmission window, which we refer to as W_1. In round r, $W_1(r)$ encoding packets are generated for batch r and placed in an outgoing queue Q_1. Only after the transmission of $W_1(r)$ encoding packets for the batch, the users can start the recovery process. Therefore, the value of $W_1(r)$ is at least equals to $W_1^{\min} = K$, for $1 \leq r \leq B$. However, protocol may choose to transmit more than K encoding packets for a batch in its first attempt in order to proactively account for the possible packet losses during the transmission.

A user may fail to complete the recovery process of a batch following the initial attempt, for two reasons: (i) packet loss over the channel is too severe, and not enough packets have been accumulated, or (ii) K encoding packets are not sufficient for decoding of this particular instance of the LT encoding which may happen with probability δ, for some $\delta > 0$. In these two cases, the recovery of the batch is not completed after the first attempt, and additional encoding packets are transmitted in the subsequent attempts. A second window, W_2, controls the transmission of additional packets for all failed batches. In transmission round r, the protocol transmits additional encoding

packets for batches $i = 1, \ldots, r - 1$, such that

$$W_2(r) = \sum_{i=1}^{r-1} A_i(r), \tag{4.1}$$

where $A_i(r)$ is the additional number of encoding packets required at round r for a failed batch i with initially $A_1(1) = 0$ and $W_2(1) = 0$. Additional encoding packets for failed batches are stored in an outgoing queue Q_2. Q_2 always holds additional encoding packets for previously failed batches. If all previous batches are recovered before the current round, then $W_2(r) = W_2^{\min} = 0$.

The transmission of packets in queue Q_1 and queue Q_2 follow a time division, where transmission of packets in Q_1 is followed by the transmission of the packets of in Q_2. This constitutes the completion of transmission round r. After completion of the transmission round, the NOC awaits feedback reports from the users. When the protocol receives a feedback report from a user u, it gets information on the decoding status of the last transmitted and all previously failed batches. This feedback report consists of a set of indices of the batches that failed at user u, $I_u \subseteq \{1, 2, \ldots, r\}$, and the number of additional encoding packets required in the subsequent round to complete the decoding in the next round, — i.e. $\{A_i^u(r+1)\}_{i \in I_u}$. After collecting feedback reports, the protocol calculates the maximum number of additional encoding packets required for all batches, over all users:

$$A_i(r + 1) = \max_{u:i \in I_u} \{A_i^u(r + 1)\}, \ i = 1, \ldots, r. \tag{4.2}$$

The protocol then proceeds with the next transmission round until all batches are completed successfully by all users. Ideally, the goal of the protocol would be to com-

plete the transmission in B rounds. However, because of possible packet losses, it may be required to transmit additional encoding packets for several batches after round B. In this case, $W_1(r) = 0$, for $r > B$, since initial transmission attempts for all batches have been completed.

4.3.3 Adaptive transmission of encoding packets

From user perspective, the protocol must make sure that every batch can be recovered at the end of the first attempt of transmission, because additional rounds add to the overall delay. At the same time, protocol should transmit as few encoding packets as possible (ideally completing the recovery after transmitting K encoding packets) for bandwidth efficiency. These two requirements constitute a trade-off. In one extreme scenario, the protocol may set $W_1 >> K$, in which case, all users complete decoding of a batch after initial attempt with very high probability. On the other hand, if packet loss rate is low, this approach would result in transmission of many unnecessary encoding packets. In the other extreme, the protocol may set $W_1 = K$, and transmit additional encoding packets reactively in response to user requests. This guarantees that minimum number of encoding packets are transmitted for the batch. However, because, the propagation delay for a two-way satellite system is much larger than the transmission delay at high data rates, this approach would result in increased delays.

In this section, we look at the possibility of using channel state condition for determining the number of encoding packets transmitted in each round. At the end of trans-

mission round r, each user transmit a set of indices $I_u \subseteq \{1, 2, \ldots, r\}$ corresponding to the batches that are not completed at the end of the round, and the number of additional encoding packets requested in the next round, — i.e. $\{A_i^u(r+1)\}_{i \in I_u} \ \forall u \in \mathcal{U}$. This information can be used to calculate the average loss rate of the user in the previous round by,

$$P_u(r+1) = \begin{cases} \frac{1}{|I_u|} \sum\limits_{i \in I_u - \{r\}} \frac{A_i^u(r+1)}{A_i(r)} + \frac{A_r(r+1)}{W_1(r)}, & r \in I_u \\ \frac{1}{|I_u|} \sum\limits_{i \in I_u} \frac{A_i^u(r+1)}{A_i(r)}, & r \notin I_u \end{cases} \tag{4.3}$$

In the next transmission round, these loss probability estimates are used to scale the number of encoding packets transmitted to the users according to the following rules:

$$P(r+1) = \lceil \max_{u \in \mathcal{U}}\{P_u(r+1) : P_u(r+1) \leq \gamma\} \rceil \tag{4.4}$$

$$W_1(r+1) = P(r+1) \cdot K, \tag{4.5}$$

$$W_2(r+1) = \sum_{i=1}^{r}(P(r+1) \cdot A_i(r+1)), \tag{4.6}$$

where, $P(r+1)$ is the maximum packet loss rate based on the feedback information collected at the end of the previous round, and γ is a threshold value. The threshold value is required to limit the number of additional encoding packets transmitted for the current round when the maximum packet loss rate is very high. We are going to use γ as a tunable parameter. Intuitively, if some user is experiencing a very deep fade, resulting in a very high packet loss rate, even the large number of additional encoding packets transmitted would not be enough for it to accumulate the necessary packets for the purpose of decoding. Therefore, the benefit to the worst case user would be minimal while the bandwidth efficiency for users in better channel conditions drops drastically.

In the next section, we analyze the potential benefits of this approach for an example scenario, in an attempt to gain better understanding of the uses of packet level FEC coding as an aid in reliable multicast communication.

4.4 Performance analysis

4.4.1 Simulation setup

In this section, we look at the distribution of a 184Mb file to $U = 20$ users over the satellite network. The file is organized as $B = 100$ batches of $k = 10000$ MPEG-2 transport stream packets (184 Bytes payload and 4 Bytes header). The analysis in [53] shows that the decoding overhead for batches of this size over repeated experiments never exceeds 10% for LT codes. Therefore, as a worse case analysis, we assume that a user successfully recovers an input batch if it accumulates at least $K = 11000$ encoding packets.

In order to study the performance of packet distribution with integrated LT coding, we need to establish connection between the channel states of each user and the corresponding packet loss rates. The channel state for each user is simulated independently according to the channel model previously discussed in Section 2.5.2. This model generates a time series of signal attenuation over a satellite channel operating at Ka-band frequencies. For every user, signal attenuation level sample generated by the simulator is input to the downlink link power-budget calculation for the satellite system (see Section 2.5.1), for a transmission power level of $P = 5.0$ Watts and downlink data rate of

92 Mbps. The link power-budget calculation determines the ratio of bit energy to noise power spectral density (E_b/N_o) for every user.

Assuming that the DVB channel encoder uses the standard concatenated coding with (188,204) RS outer code to encode each 188-Bytes LT encoding packet, ideally interleave the packet symbols, and then pass it through a rate $1/2$ convolutional coder, the packet error probability at the output of the DVD channel decoder for a user is given by the following well-known approximations:

1. For the convolutional inner code, the probability that the output bit-stream will differ from the original bit-stream in exactly d positions is given by

$$P_d = \frac{1}{2}\text{erfc}(\sqrt{\frac{E_b}{N_o} \cdot d}), \tag{4.7}$$

   for $BPSK$ modulation.

2. The probability of bit error at the output of the convolutional decoder can be approximated by

$$P_{cb} \approx \sum_{d=d_f}^{d_{\max}} c_d \cdot P_d, \tag{4.8}$$

   where d_f is the free distance of the convolution code, and c_d is the number of possible output bit-streams that differ from the original one by d positions. The first eighteen values of c_d for the $d_f = 10$ convolutional code used by the DVB encoder is given by

$$c_d = [36\ 0\ 211\ 0\ 1404\ 0\ 11633\ 0\ 77433\ 0\ 5026900$$
$$3322763\ 0\ 21292910\ 0\ 134365911\ 0]. \tag{4.9}$$

3. The symbol error probability can be upper-bounded by

$$P_s \leq \min(8 \cdot P_{cb}, 1.0). \tag{4.10}$$

During the simulations, symbol error probability of each user is calculated using this analysis and the corresponding propagation conditions. Based on the assumption that the symbol errors within a packet occur independently of each other due to the presence of symbol interleaver, for each user packet erasure events are determined. Packet erasure events take into account the error detecting and correction capabilities of the outer RS code. In the next section, we provide our simulation results under this test setup.

4.4.2 Results

The behavior of the algorithm depends heavily on the packet loss patterns of the users, which depend on the channel propagation conditions. Therefore, we look at 100 different instantiations of the file transfer described in the previous section over a period of 24 hours simulated time. There are several metrics of interest in this scenario:

1. **Average number of encoding packets per batch:** This is the average number of encoding packets transmitted per batch over the lifetime of the file transfer (over 100 batches) including the packets transmitted during the initial attempt and any subsequent attempts. The file transfer consists of 100 batches of $K = 11000$ encoding packets. Therefore, this average is always equal to or greater than K.

2. **Average number of transmission rounds per batch:** This is the average number of attempts made by the NOC in order to complete the transmission of a batch over

the lifetime of the file transfer. At least one attempt is made in order to transmit a batch, hence, this average is always equal to or greater than one.

3. **Session time:** This is the time it takes for all users to complete the reception of the file transfer.

4. **Average size of W_1:** This is the average size of the transmission window W_1 over the lifetime of the file transfer.

5. **Average size of W_2:** This is the average size of the transmission window W_2 over the lifetime of the file transfer.

In Figures 4.4 to 4.7, we look at these five metrics as averaged over all 100 session instantiations. The first observation we make is that the channel conditions are favorable for most of the 24 hours period. The average channel attenuation during this period remains is in the range of $2 - 3$dB, and the 5dB. transmit power, coupled with the strong DVD channel encoding, is capable of keeping the symbol error rates at low figures even when the channel attenuation is as high as $5 - 6$dB. Therefore, most of the sessions complete batch transmissions in one round, with transmitting only few more additional encoding packets per batch. In Fig. 4.4, we plot the average number of encoding packets transmitted per batch in order to complete decoding by all users as a function of the threshold γ. We observe that batches are completed on average only with $75 - 80$ additional encoding packets. Moreover, as we increase the threshold value, the algorithm transmits more packets proactively. Because the algorithm decides on the number of

123

Figure 4.4: Average number of encoding packets per batch transmitted in order to complete the file transfer as averaged over all 100 session instantiations.

packets to be transmitted proactively based on an estimate of the actual packet loss rate, and because the channel conditions may vary by the time the feedback reaches NOC, the increased value of the threshold results in the transmission of unnecessary encoding packets (observed as the increased number of packet per batch completion).

However, proactive transmission of encoding packets also helps decrease the average number of rounds required to complete a batch by all active users. This results can be observed in two different ways; (i) directly from Fig. 4.5, where average number of transmission rounds per batch is plotted as a function of the packet loss threshold γ, or (ii) from the fact that the average size of W_1 increases (i.e. more packets are transmitted with the initial attempt), while the average size of W_2 decreases (i.e. fewer packets are transmitted in the subsequent rounds). These results are plotted in Figures 4.6 and 4.7.
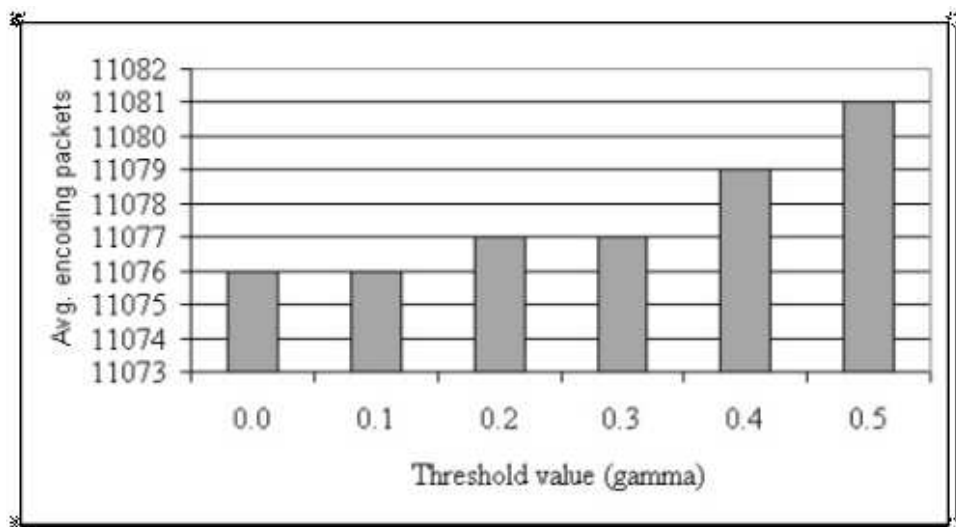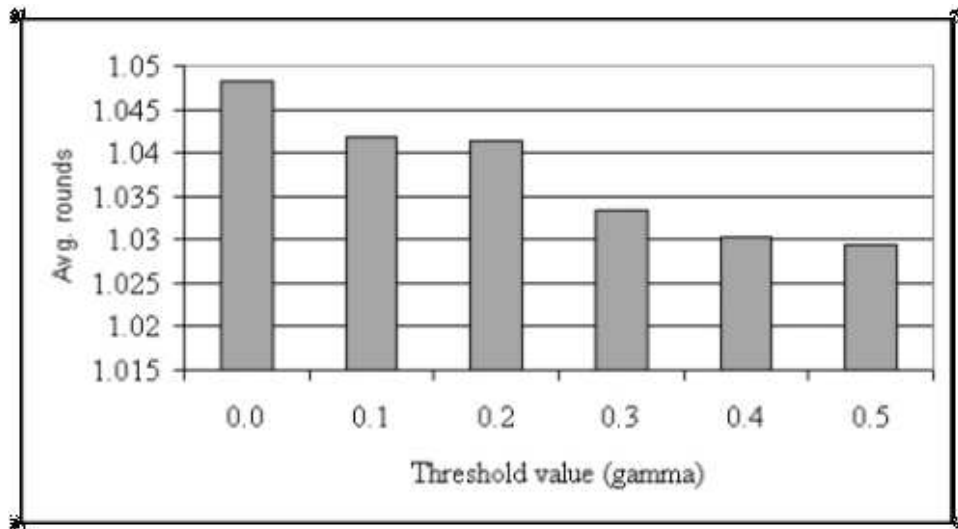
Figure 4.5: Average number of transmission rounds per batch in order to complete the file transfer as averaged over all 100 session instantiations.
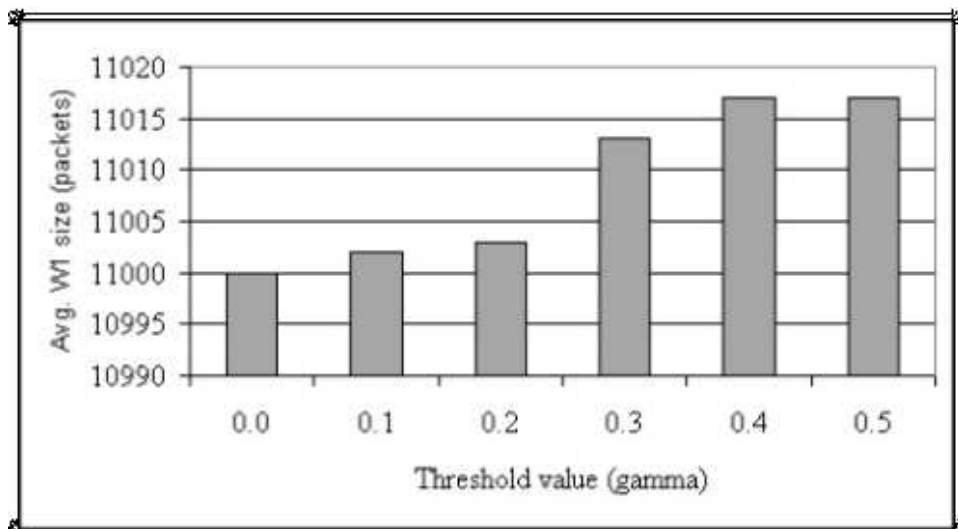
Figure 4.6: Average size of the transmission window W_1 as averaged over all 100 session instantiations.
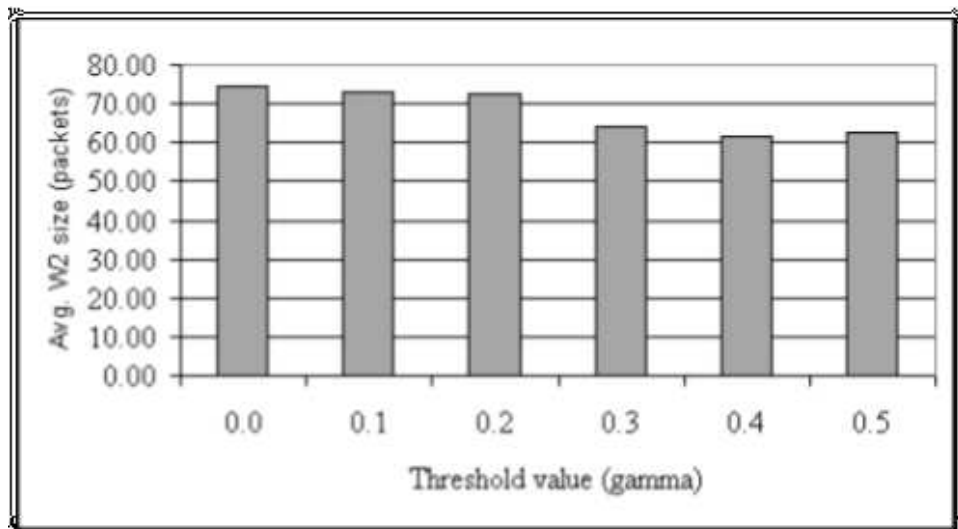
Figure 4.7: Average size of the transmission window W_2 as averaged over all 100 session instantiations.

Although the average number of transmission rounds per batch decreases as more packets are transmitted in advance, the session time to complete the transmission may actually increase as a result (Fig. 4.8). This counter-intuitive observation results from the fact that even though the number of rounds per batch decreases, the time it takes to complete a single round increases because of the additional encoding packets.

When averaged over all 100 instantiations of the file transfer, the results are not able to reflect clearly the benefits of using an adaptive window size for proactively transmitting additional encoding packets. This is because most of the sessions that are initialized during the 24hr period actually encounter favorable channel conditions and therefore complete the file transfer without the need for retransmissions or additional encoding packets. As a result of this, when we look at the average results over all sessions, the

126
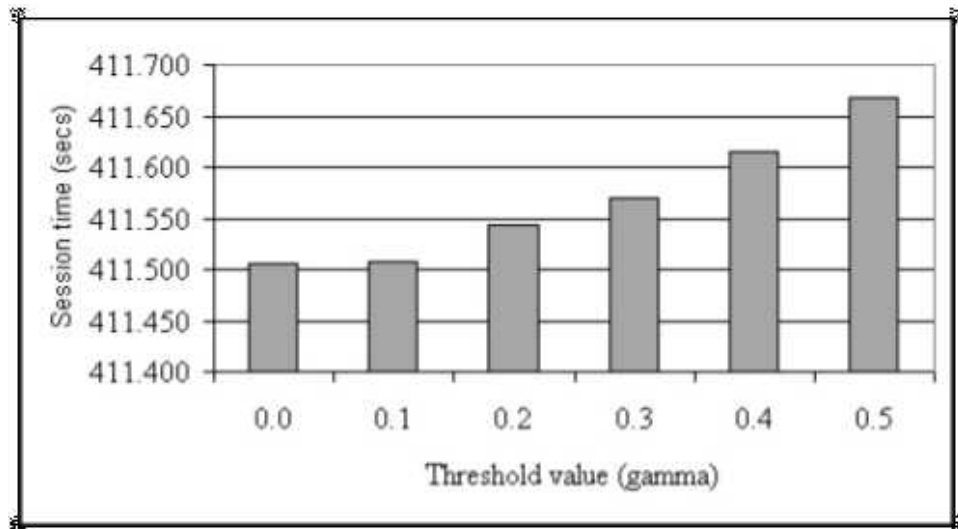
Figure 4.8: Average time it takes to complete the file transfer at all users, as averaged over all 100 session instantiations.

additional benefits of having adaptive mechanisms appear minimal. Therefore, in the next set of results, we look at the same performance metrics over only the sessions that required on the average more than $K = 11000$ encoding packets per batch to complete the file transfer.

In Fig. 4.9, we plot the average number of encoding packets required per batch for these sessions. We observe that for these sessions the number of additional encoding packets required to complete a batch can be considerably larger than the minimum number K. The flexibility in the number of encoding packets that can be generated for a batch by the LT coding structure proves very valuable particularly in these cases. However, the effect of changing the loss rate threshold follows a similar trend for these instances as well. As more encoding packets are transmitted proactively, the algorithm

Figure 4.9: Average number of encoding packets per batch transmitted in order to complete the file transfer as averaged over only the sessions that required more than K encoding packets per batch.

faces the risk of transmitting more than necessary since the calculations are based on estimates. The number of transmission rounds required per batch, on the other hand, decreases with the proactive approach (Fig. 4.10).

Similar trends are observed for the sizes of transmission windows W_1 and W_2, even though the dynamic range is much larger for this subset of sessions (Figures 4.11 and 4.12). Clearly, the results establish the conclusion that as more encoding packets are transmitted proactively, the requirement for additional encoding packets in the subsequent rounds decreases. However, the results also show that the algorithm faces the risk of transmitting to many encoding packets in advance, while trying to predict future packet losses, thereby increasing the time to complete the session for all users
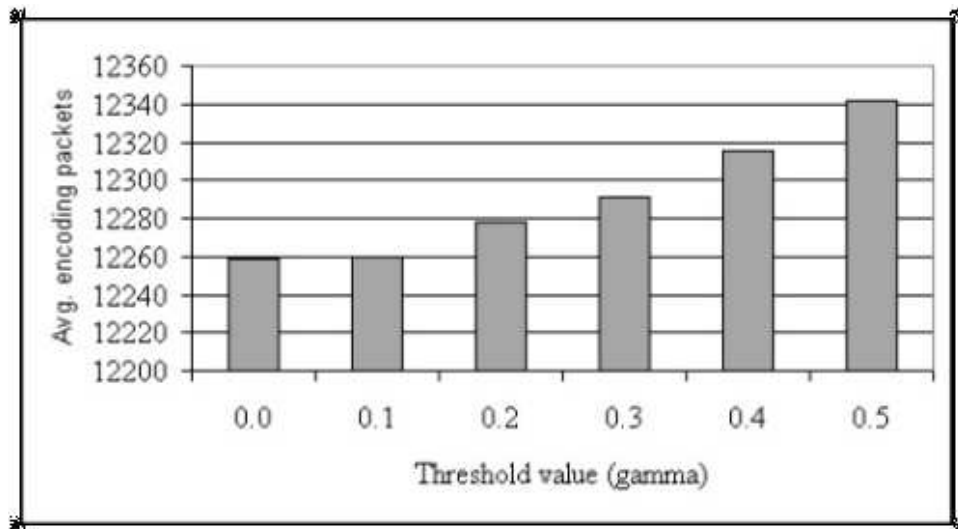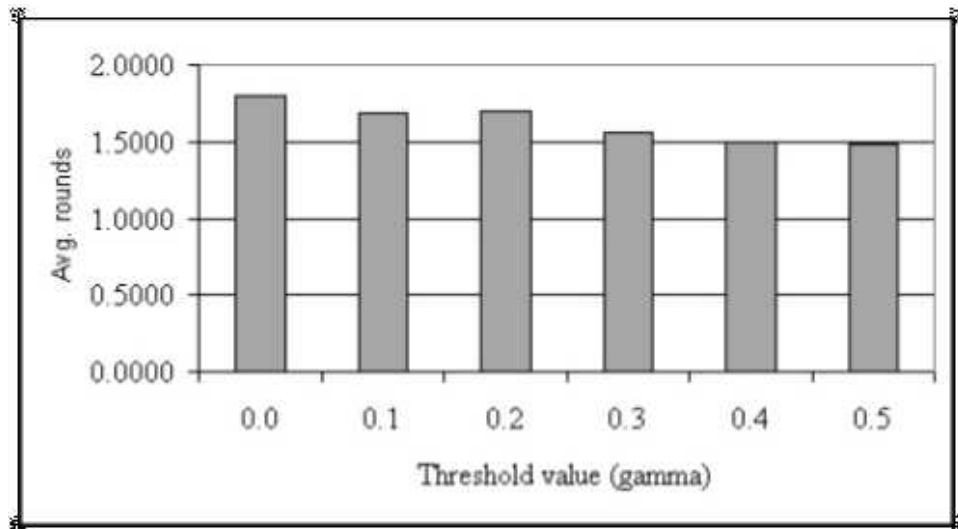
Figure 4.10: Average number of transmission rounds per batch in order to complete the file transfer as averaged over only the sessions that required more than K encoding packets per batch.

(Fig. 4.13).

In the next section, we sum up our observations from this chapter and comment on the insights we gained on the use of packet level FEC coding.

4.5   Summary

In this chapter, we have discussed the potential benefits of using packet level FEC coding as an aid in reliable multicast communication. Integration of packet level FEC coding is useful in reliable multicast communication because it helps in recovering from packet losses without relying on retransmission of specific packets. This is the most important feature of the packet level FEC codes in satellite multicast transmission. A transmitted
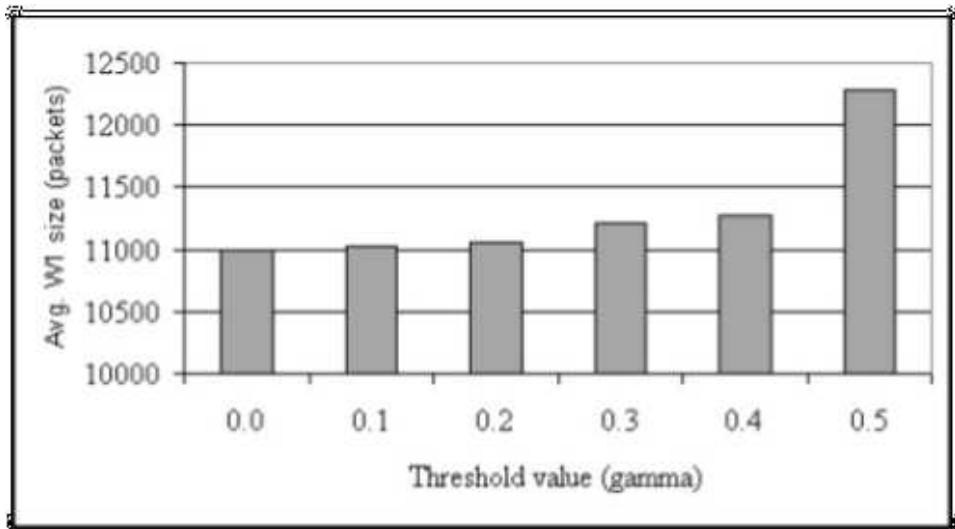
Figure 4.11: Average size of the transmission window W_1 as averaged over only the sessions that required more than K encoding packets per batch.
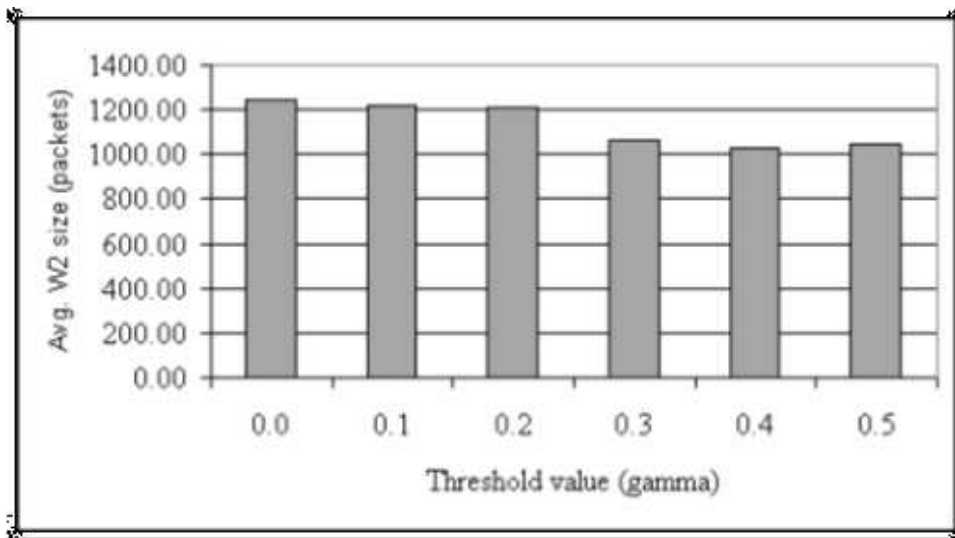
Figure 4.12: Average size of the transmission window W_2 as averaged over only the sessions that required more than K encoding packets per batch.
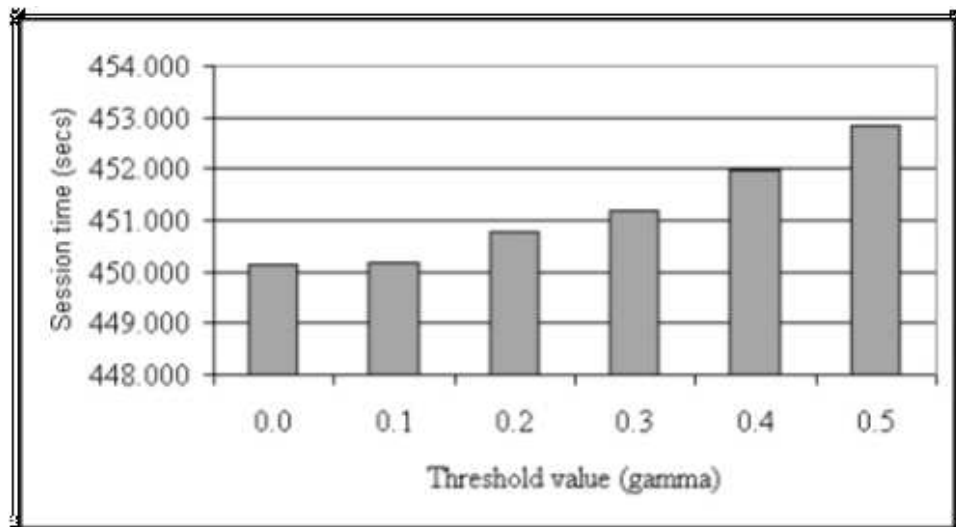
Figure 4.13: Average time it takes to complete the file transfer at all users, as averaged over only the sessions that required more than K encoding packets per batch.

encoding packet recovers losses of different packets in different users, and the broadcast nature of the satellite channel enables easy dissemination of encoding packets to several users in single hop. However, packet losses in satellite networks occur in bursts that happen in time-scales that are much larger than the duration of a single packet transmission. Therefore, in order to be useful, packet level FEC coding should be able to generate a fairly large number of encoding packets per input batch. Otherwise, the system would run out of encoding packets to transmit before all users accumulate enough packets to continue with the decoding of a batch. The LT encoding structure we have discussed in this chapter has been the choice primarily because of its ability to generate a very large number of encoding packets per input batch.

We also investigated the possibility of adjusting the number of encoding packets

131

transmitted in response to the estimated packet losses. The benefit of this approach appears to be limited for several reasons. First of all, the packet loss rate remains low for most of the operational time due to the existence of link power-margins and strong channel coding implemented for typical systems. Therefore, for most of the sessions, transmission of additional encoding packets in response to user requests prove to be efficient enough. In cases where channel fading is deep enough to cause severe packet loss patterns, it is difficult to estimate the packet loss rates in advance to act proactively, especially for short session durations. Therefore, the algorithm risks transmitting more encoding packets than needed. In conclusion, we expect adaptive transmission of encoding packets to be more application specific and be applied to dissemination of large data, rather than being a standard part of a multicast transmission protocol.

Chapter 5

Conclusions

In this dissertation, we focused on investigating the possible performance enhancements in terms of the throughput, capacity and scalability of a multiple-spot beam satellite communication system that supports both unicast and multicast services. Towards this goal, Chapter 2 and Chapter 3 were dedicated to providing schemes for more efficient allocation and use of system resources when multicast applications share the system resources with unicast applications. On the other hand, Chapter 4 investigated the fundamentals of reliable multicast communication over satellite networks from the transport protocol perspective.

In this work, we targeted a star topology satellite network, where a Ka-band, geosynchronous satellite provided broadband services to a large number of users located inside its footprint. The multiple spot-beam, on-board packet processing and switching technologies allowed transmission of data to multiple users in multiple spot-beam locations, while users that are equipped with two-way direct communication terminals accessed the backbone terrestrial network through the network operations center. In this network setup, we highlighted the performance related issues that arise when multicast sessions are introduced to the system, and when they compete with existing unicast sessions. This network template, and its discussion throughout the dissertation, provides

a guideline for future work on this topic by clearly differentiating the design and solution spaces for multicast communication over satellite system, from that of the wireline systems.

More specifically, in Chapter 2, we looked at the problem of user heterogeneity in multicast communication. In a multiple spot-beam satellite networks, we showed that user heterogeneity occurs a result of the temporal and spatial variations of the physical communication channel. We argued that this variations become significant particularly for systems operating at the Ka-band frequencies and beyond. These observations have lead to the conclusion that users residing in different geographical locations are likely to experience highly variable channel conditions, and in order for multicast services to perform favorably, the system has to compensate for these variations, not in an individual basis, but as a group. Towards this end, we showed that schemes primarily targeted for the mitigation of the effects of signal fading, such as adaptive power allocation, could be made aware of the number, distribution and the type of the sessions operating at the higher network layers. Our contribution allows adaptive power allocation to multiple spot-beams with the goal of smoothing user heterogeneity, and result in average service rates that are higher for both unicast and multicast sessions. This work also pushes the envelope for the possibility of more cross-layer interaction between functionalities that have been considered to operate separately.

In Chapter 3, we provided methods for efficient feedback collection from a large group of users. We identified efficient feedback collection as a key issue, not only be-

cause it complements the operation of our multicast-aware power allocation scheme, but also because it is a key part in providing additional multicast services, such as reliable multicast communication. Our discussion shows that providing efficient feedback collection and at the same time controlling the volume of information flowing through the network is a more complicated task in our target scenario than similar approaches targeted primarily for wireline networks where users share or receive information among them more easily. However, our framework proved that it is possible, in most cases, to reduce the volume of user feedback by exploiting the nature of the information to be transmitted and by considering that only a subset of users need to be tracked. We first formulated our solution as a stand alone application to the problem of collecting channel state information from a set of users that can access to only a limited number of channels for communicating their feedback information. The solution, which is based on the idea that the users with the *breaking news* should be given immediate priority in communicating their information, proved successful in limiting the volume of feedback information generated, while maintaining the collection of relevant information. Following our stand-alone evaluation, we integrated this algorithm to our multicast-aware power allocation scheme to complement its operations. Our contributions in this chapter showed that, it is possible to rely on user feedback for various multicast services, even in the case of large number of users, without allocating a large portion of system resources.

The last part of the thesis investigated multicast communication support from a different angle that looked at the problem from the point of view of individual multicast

sessions rather than the global optimization of system resources among all active unicast and multicast sessions. In Chapter 4, we evaluated the use of packet level FEC coding as a part of a generic transport protocol that provides reliable multicast communication. Our experimental setup was more geared towards understanding the benefits and limits of having packet level FEC coding in this context than providing a full protocol description. Our experiments confirmed that packet level FEC coding is efficient in recovering from losses in the context of satellite networks because of the possibility of recovering different packets losses in different users by dissemination of encoding packets. The main contribution of this chapter is in the use of a more complicated and more realistic channel model in experiments. Our setup included a time-series generator for simulating the channel attenuation levels, the output of which is fed into a real-life link-budget calculation in order to estimate the packet loss rates. This setup was more complicated, and presumably more realistic, than previous cases in the literature, which set the packet loss rates as an external parameter, rather than going through the complete stack we have used. This setup was crucial in providing the insights we have obtained on the adaptive application of packet level FEC coding.

Although this thesis covered a significant number of topics related to multicast communication support over satellite communication networks, there remain equally important open problems to dwell up on. The next section, highlight some of these problems for future work on this topic.

5.1 Suggestions for future study

The following points remain open to further study:

- In Chapter 2, we tested our power allocation policy for different number of sessions and different session configurations, however, they were set at the beginning of the analysis and remained fixed in time during the course of the simulations. A more realistic scenario would have to investigate dynamic session configurations that take into account the session connection and termination rates and session durations into effect.

- In Chapter 2, we assumed that the power-rate curve was a continuous function. However, in a real system, the set of power levels that can be supported and the corresponding transmission rates for a given set of modulation schemes are likely to be finite. Therefore, it is desirable test this approach under more stringent system constraints.

- It may be an interesting direction to link the minimum power levels associated with each active session in Chapter 2 to some form of a quality of service (QoS) metric. Minimum power levels guarantee the support of a minimum rate that can be negotiated before the start of the session as a QoS level.

- In Chapter 3, we chose a simple functional relationship for determining the transmission probability of individual users. Our analytical results were limited in favor of using a more complicated channel model, which did not lend itself into a

easy mathematical description. Therefore, we relied on simulations and intuitive analysis. It would be interesting to see if this can be taken a step further to get better understanding of the form of the transmission probability curve.

- In Chapter 4, we relied on simple averages to estimate the packet loss rate for future frames. However, it might be possible to extract more information on the channel behavior if other parameters such as fade durations, and the conditional probability distributions for fade levels could be obtained from the underlying channel model. Therefore, an interesting direction, which was outside the scope of this dissertation, would be extract statistical information on the channel behavior and to incorporate it to the estimation of future packet loss rates for adaptive packet level FEC encoding.

BIBLIOGRAPHY

[1] F. Gargione, T. Iida, F. Valdoni, and F. Vatalaro, "Services, technologies, and systems at Ka band and beyond: A survey," *IEEE J. Select. Areas Commun.*, vol. 17, no. 2, pp. 133–144, February 1999.

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

[3] J. V. Evans, "Communication satellite systems for high speed Internet access," *IEEE Antennas and Propagation Magazine*, vol. 43, no. 5, October 2001.

[4] P. Chitre and R. Gupta, "Satellite systems for multimedia and Internet traffic," in *Proc. of IEEE Microwave Symposium Digest*, vol. 2, May 2001, pp. 20–25.

[5] A. Bertout, G. Christiansen, J. Couetet, and P. Spencer, "Skybridge services," Alcatel Telecommunications Review, Tech. Rep., 4th Quarter 2001, http://www.alcatel.com/atr/.

[6] E. J. Fitzpatrick, "SPACEWAY system summary," *Space Communications*, vol. 13, pp. 7–23, 1995.

[7] G. Losquadro and V. Marziale, "The Euroskyway network with on-board processing to support fast Internet and interactive TV services for residential users," in *Proc. of IEEE-IBC Colloquium*, London, October 2000, www.euroskyway.it.

[8] G. Akkor, J. S. Baras, and M. Hadjitheodosiou, "An optimization-based approach to flow control and resource allocation problem in satellite communication systems," in *Proc. of AIAA ICSSC 2004*, Monterey, CA, May 2004.

[9] ——, "A feedback implosion suppression algorithm for satellite reliable multicast," in *Proc. of IEEE GLOBECOM*, vol. 7, San Francisco, CA, December 2003, pp. 3765–3769.

[10] G. Akkor and J. S. Baras, "A reliable multicast transport protocol with feedback implosion suppression for satellite communication systems," in *Proc. of CISS*, Princeton, NJ, March 2004.

[11] G. Akkor, M. Hadjitheodosiou, and J. S. Baras, "Transport protocols in multicast via satellite," *Int. J. Sat. Comm.*, vol. 22, no. 6, pp. 611–627, November/December 2004.

[12] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-driven layered multicast," in *Proc. of ACM SIGCOMM*, 1996, pp. 117–130.

[13] L. Vicisano, J. Crowcroft, and L. Rizzo, "TCP-like congestion control for layered multicast data transfer," in *Proc. of IEEE INFOCOM*, vol. 3, March 1998, pp. 117–130.

[14] X. Li, S. Paul, and M. Ammar, "Layered video multicast with retransmissions: Evaluation of hierarchical rate control," in *Proc. of IEEE INFOCOM*, vol. 3, March 1998, pp. 1062–1072.

[15] V. Roca, "Packet scheduling for heterogeneous multicast transmissions," in *Proc. of PfHSN*, August 1999, pp. 101–116.

[16] M. J. Donahoo, M. H. Ammar, and E. W. Zegura, "Multiple channel multicast scheduling for scalable bulk-data transport," in *Proc. of IEEE INFOCOM*, vol. 2, 1999, pp. 847–855.

[17] C. Enjamio, E. Vilar, and F. P. Fontan, "Spatial distribution of rainfall rate: benefits of the site diversity as a dynamic fade mitigation technique," in *Proc. of COST Action 280, First International Workshop*, July 2002.

[18] H. Fukuchi, "Quantitative analysis of the effect of adaptive satellite power control as a rain attenuation countermeasure," in *Proc. of IEEE International Symposium on Antennas and Propagation*, vol. 2, 1994, pp. 1332–1335.

[19] A. D. Panagopoulos, P.-D. M. Arapoglou, and P. G. Cottis, "Satellite communications at Ku, Ka, and V bands: propagation impairments and mitigation techniques," *IEEE Communication Surveys and Tutorials*, vol. 6, no. 3, Third Quarter 2004, www.comsoc.org/pubs/surveys.

[20] M. M. J. L. van de Kamp, "Climatic radiowave propagation models for the design of satellite communication systems," Ph.D. dissertation, Eindhoven Technical University, Eindhoven, The Netherlands, November 1999.

[21] 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.

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

[23] ——, "A time-series generator modeling rain fading," in *Proc. of Open Symposium on Propagation and Remote Sensing*. Garmisch-Partenkirchen: URSI Commission F, 2002.

[24] 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

[25] 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.

[26] M. W. Koyabe and G. Fairhurst, "Reliable multicast via satellite: A comparison survey and taxonomy," *Int. J. Sat. Comm.*, vol. 19, no. 1, pp. 3–23, 2001.

[27] P. Radoslavov, C. Papadopoulos, R. Govindan, and D. Estrin, "A comparison of application-level and router-assisted hierarchical schemes for reliable multicast," in *Proc. of IEEE INFOCOM*, 2001, pp. 229–238.

[28] J. P. Macker and R. B. Adamson, "The multicast dissemination protocol toolkit," in *Proc. of IEEE MILCOM 1999*, vol. 1, November 1999, pp. 626–630.

[29] S. Floyd, V. Jacobson, C. Liu, S. McCanne, and L. Zhang, "A reliable multicast framework for light-weight sessions and application level framing," *IEEE J. Select. Areas Commun.*, vol. 15, pp. 407–421, 1997.

[30] T. Speakman and et al., "Pragmatic general multicast (PGM) reliable transport protocol specification," Internet Society Request for Comments, December 2001, RFC3208.

[31] S. K. Kumar, S. Bhattacharyya, M. Keaton, D. Kiwior, J. Kurose, D. Towsley, and S. Zabele, "Scalable fair reliable multicast using active services," *IEEE Network*, January 2000.

[32] C. Wang and V. C. M. Leung, "Performance evaluations of SRMTP for reliable multicasting over satellite networks," in *Proc. of WCNC*, vol. 3, March 2003, pp. 1813–1818.

[33] G. Cao and Y. Wu, "Reliable multicast via satellite," in *Proc. of IIT: Coding and Computing*, April 2001, pp. 183–188.

[34] S. Cho and I. F. Akyildiz, "A poker-game-based feedback suppression algorithm for satellite reliable multicast," in *Proc. of IEEE GLOBECOM 2002*, November 2002.

[35] E. Ichihara, K. Kikuchi, T. Tsuchida, K. Kawazoe, and H. Kazama, "Reliable IP-multicast protocol for bi-directional satellite communication systems," in *Proc. of AIAA ICSSC*, April 2003, Online Technical Information.

[36] M. R. Garey and D. S. Johnson, *Computers and intractability: a guide to the theory of NP completeness*. San Francisco: W. H. Freeman and Company, 1979.

[37] D. Towsley, J. Kurose, and S. Pingali, "A comparison of sender-initiated and receiver-initiated reliable multicast protocols," *IEEE J. Select. Areas Commun.*, vol. 15, no. 3, pp. 398–406, April 1997.

[38] R. E. Blahut, *Theory and Practice of Error Control Codes*. MA: Addison Wesley, 1984.

[39] A. J. MacAuley, "Reliable broadband communication using a burst erasure correcting code," in *Proc. of ACM SIGCOMM*, Philadelphia, PA, September 1990.

[40] L. Rizzo, "Effective erasure codes for reliable computer communication protocols," *ACM Computer Communication Review*, vol. 27, no. 2, pp. 24–36, April 1997.

[41] S. Cho, A. Goulart, I. F. Akyildiz, and N. Jayant, "An adaptive FEC with QoS provisioning for real-time traffic in LEO satellite networks," in *Proc. of IEEE ICC*, June 2001, pp. 2938–2942.

[42] J. Nonnenmacher, E. W. Biersack, and D. Towsley, "Parity-based loss recovery for reliable multicast transmission," *IEEE/ACM Trans. Networking*, vol. 6, no. 4, pp. 349–361, August 1998.

[43] M. Luby, L. Vicisano, J. Gemmell, L. Rizzo, M. Handley, and J. Crowcroft, "Forward error correction building blocks," Internet Society Request for Comments, December 2002, RFC3452.

[44] ——, "The use of forward error correction in reliable multicast," Internet Society Request for Comments, December 2002, RFC3453.

[45] L. Rizzo and L. Vicisano, "A reliable multicast data distribution protocol based on software FEC techniques," in *Proc. of HPCS*, 1997.

[46] D. Rubenstein, S. Kasera, D. Towsley, and J. Kurose, "Improving reliable multicast using active parity encoding services," in *Proc. of IEEE INFOCOM*, 1999.

[47] I. Rhee, S. R. Joshi, M. Lee, S. Muthukrishnan, and V. Ozdemir, "Layered multicast recovery," in *Proc. of IEEE INFOCOM*, vol. 2, 2000, pp. 805–813.

[48] A. Donner, S. Bovelli, and S. Shabdanov, "Reliable multicast based on DVB-RCS," in *Proc. of AIAA ICSSC*, May 2002, Online Technical Information.

[49] H. Ernst, S. Shabdanov, A. Donner, and S. Scalise, "Reliable multicast for land mobile satellite channels," in *Proc. of AIAA ICSSC*, April 2003, Online Technical Information.

[50] F. Tommasi, S. Molendini, and A. Tricco, "Design of the satellite reliable distribution protocol (SRDP)," in *Proc. of IEEE GLOBECOM*, vol. 6, December 2003, pp. 3284–3289.

[51] J. W. Byers, M. Luby, M. Mitzenmacher, and A. Rege, "A digital fountain approach to reliable distribution of bulk data," in *Proc. of ACM SIGCOMM 1998*.

[52] M. Luby, "Information additive code generator and decoder for communication systems," U. S. Patent No. 6,307,487, October 2001.

[53] ——, "LT codes," in *Proc. of Symposium on Foundation of Computer Science*, 2002.

[54] X. Zhou and J. S. Baras, "TCP over GEO satellite hybrid networks," in *Proc. of MILCOM 2002*, vol. 1, October 2002, pp. 29–34.

[55] Y. Shang and M. Hadjitheodosiou, "TCP splitting protocol for broadband aeronautical satellite network," in *Proc. of DASC 2004*, vol. 2, 2004, pp. 1–9.

[56] *Digital Video Broadcasting (DVB): DVB specification for data broadcasting*, ETSI EN 301 192 v.1.3.1 ed., European Telecommunication Standards Institute, 2003.