

Routing of Real-time Traffic in a Transformational Communications Architecture¹

A. Durresi^{a*}, D. Dash^b, B. L. Anderson^b, R. Kannan^a, S. Kota^c, R. Jain^b

^aLouisiana State University, 298 Coates Hall, Baton Rouge, LA 70803

^bThe Ohio State University, Columbus, Ohio, 43210

^cHarris Corporation West Coast Operation

*durresi@csc.lsu.edu, Telephone: 225-578-3902

Abstract—In this paper we study a constraint-based routing strategy using label switching in a multi-layered hierarchical satellite constellation such as Transformational Communications Architecture. Both quality (Bit Error Rate on free space optical and radio inter-satellite links) and bandwidth availability on a satellite link are taken into account when setting up routes for high priority real-time traffic such as VoIP, which is sensitive to delay and jitter. To protect the real-time traffic from being swamped by bursty best-effort traffic we propose to have a separate queue for high priority traffic. Packets from several real-time flows are aggregated onto the same LSP (Label Switched Path) based on destination and priority, and are shielded from each other by a proportional dropping policy, where packets from flows exceeding their Committed Information Rate (CIR) have high dropping probability during congestion. The performance of the prioritized load balancing routing algorithm on a multi-layered satellite network is simulated and analyzed.

TABLE OF CONTENTS

1. INTRODUCTION	1
2. CONSTRAINED MULTI-PATH ROUTING	3
3. BACKGROUND AND RELATED WORK.....	5
4. A TCA ARCHITECTURE	6
5. HIERARCHIAL ROUTING.....	8
6. TRAFFIC AGGREGATION	10
7. SIMULATION RESULTS.....	11
7. CONCLUSIONS	16
REFERENCES.....	17

1. INTRODUCTION

Satellite networks with their potential for global coverage and high bandwidth availability are an attractive option for establishing an “internet in the sky” [10, 30]. In the absence

of a wired infrastructure, any individual host or an ad hoc network can access the rest of the wired network through satellites. Satellite networking has evolved from simple bent-pipe routing for geo-stationary orbit (GEO) satellite networks to on-board switching capabilities in low earth orbit (LEO) broadband satellite constellations [6] and High Altitude Platforms (HAPs) [31].

The next generation of networking will involve integration of terrestrial and space networks with High Altitude Platforms (HAPs) providing last mile connectivity to certain sensitive areas (e.g. disaster relief, battlefields) where high bandwidth and accessibility are necessary. With the advance of free space optical technology it is possible to have a network of these HAPs, which could be unmanned aerial vehicles, communicating with each other as well as with satellites through Inter Orbit Links (IOL). With radio access they make it possible for small hand-held terminals on the ground to access satellites with high speed connections [11, 12].

Several US government agencies announced in 2002 an initiative to build a new space communication infrastructure, which was named “Transformational Communications Architecture” (TCA) [15]. TCA will be a joint effort with full involvement by the Assistant Secretary of Defense for Command, Control, Communications and Intelligence and the Defense Information Systems Agency. The Transformational Communications Office has support from across the military services, the Defense Information Systems Agency, the intelligence community and National Aeronautical and Space Agency (NASA).

The Transformational Communications Architecture is shown in Figure 1 [15]. This new space infrastructure will be similar to the Internet architecture and will be secure. Its primary use is military but later on public services would be

¹ 0-7803-8155-6/04/\$17.00© 2004 IEEE

² IEEEAC paper #1215, Version 3, Updated December 10, 2003

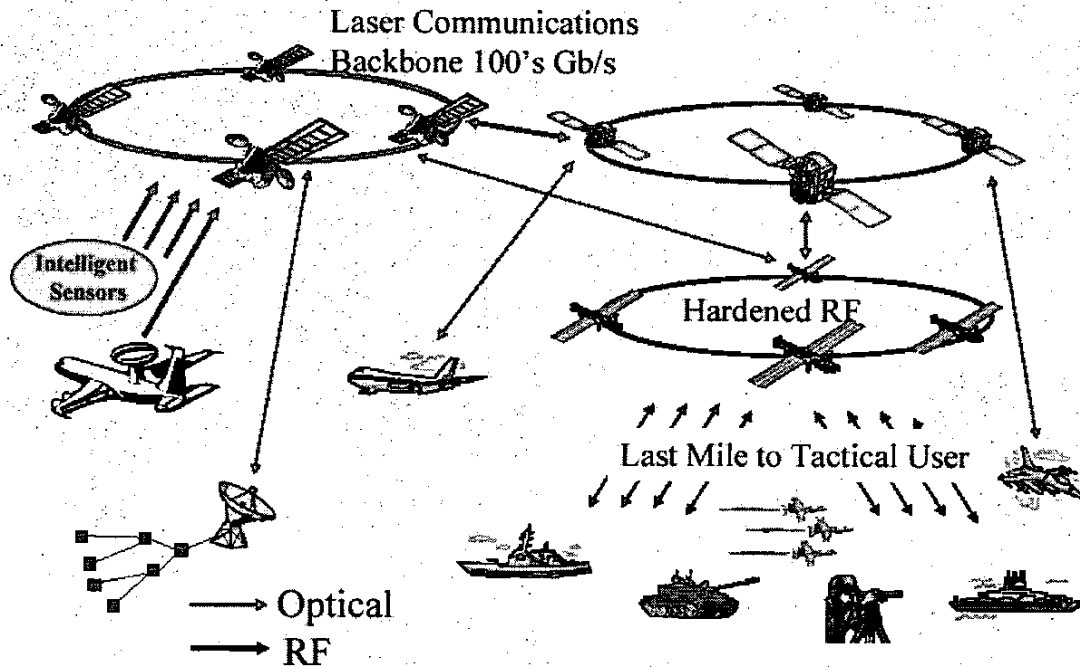


Figure 1. Transformational Communications Architecture

carried out. In such complex architecture it is important to design a simple routing and congestion control scheme in order to use efficiently the network resources. We consider a three-layer architecture of geo-stationary (GEO) satellites, low earth orbit (LEO) satellites and high altitude platforms (HAPs) to simulate an architecture similar to the TCA.

GEOs orbiting in high altitude geo-stationary orbits (36000 km) individually are unattractive for delay sensitive traffic because of large propagation delay, consequently making them unsuitable for real-time applications.

All these issues motivated the deployment of low-earth orbit (LEO) satellites that orbit the Earth at a height of just 500 to 1,000 miles, which in turn necessitates the use of multiple satellites that constantly orbit around the earth in fixed planes, to provide constant service to any area. The low altitude orbit makes LEOs capable of providing smaller, more energy-efficient spot beams, and delivers latency potentially equal to (or better than) transcontinental fiber optic cable. Frequency reuse is also an important advantage considering the limited and costly frequency spectrum while increasing the system capacity. With the advent of multiple spot beams, inter-satellite links (ISLs) between satellites and on board switching and processing capabilities, these constellation of low-earth orbit (LEO) satellites along with their terrestrial gateway servers form Autonomous systems (AS). One of the distinct advantages of LEO satellite networks over GEO networks is the reduction in

propagation delay, making them an attractive option for routing real time traffic.

HAPs can be deployed in areas with heavy and sensitive traffic (e.g. battlefields, disaster relief) and provide access to the high-speed satellite network for terrestrial users with mobile and hand-held terminals. HAPs can easily form a high-speed network among themselves and satellites in the higher layers [13] with optical links.

We propose and study a multi-layered satellite network with GEOs acting as the backbone routers, LEOs as the second layer and HAPs deployed in specific local regions. With the term HAP we include also other types of flying vehicles. This three-layered constellation provides high bandwidth access to all types of users and low latency to delay sensitive applications. Such architecture is similar to the Transformational Communications architecture proposed in [15]. Figure 2 shows the three layer satellite network architecture of GEOs, LEOs and HAPs where we apply a prioritized load balancing routing algorithm. We identify traffic as either best-effort or high priority, with QoS guarantees for high priority traffic.




Figure 2. Satellite multi-layer architecture

The rest of the paper is organized as follows; Section 2 discusses the motivation for constrained multi-path routing and load balancing in layered satellite networks; Section 3 reviews the background research work, Section 4 discusses the selected architecture, Section 5 presents hierarchical routing, Section 6 discusses traffic aggregation, in Section 7 are presented our simulation results, and Section 8 concludes.

2. CONSTRAINED MULTI-PATH ROUTING

Given the nature of satellite constellations, static routing protocols based on topological properties of the network like minimum hop path or minimum delay will give rise to congestion at some points in the network (e.g. satellites that are visible over major cities) and a lot of unused bandwidth at other points in the network, leading to under-utilization of network resources and degradation of service offered. The number of active sessions in such networks is uncontrolled and no kind of priority policy is implemented. Consequently, high priority calls are routed over (and share) paths heavily used by low-priority calls.

The need for QoS in satellite networks is motivated by several reasons. With an explosion of network traffic in terms of users and applications, ISPs want to offer different levels of service based on business priorities of the users or applications. With applications varying from real time interactive traffic (e.g. VoIP), real time non-interactive traffic (e.g. streaming video) to non-real time traffic (e.g.

web traffic) it is necessary to differentiate between the levels of service provided. High-speed networks should be able to support different degrees of Quality of Service (QoS) to different applications. For example, real-time traffic generated by multimedia applications has radically different requirements than best-effort traffic. So real-time applications require tight bounds on transfer delay (in the order of hundreds of milliseconds).

Real-time applications such as Voice over IP (VoIP) and streaming video are susceptible to changes in the transmission characteristics of data networks. Real-time traffics are also susceptible to network behaviors referred to as delay, jitter and packet loss, which can degrade the voice application to the point of being unacceptable to the average user.

It becomes essential therefore, to separate real-time traffic from non-real time, and route it over explicit paths that meet the desired QoS requirements. Satellite networks employing static routing based on topological properties of the network lead to congestion in some parts of the network, which is heavily used by all classes of traffic even when there is a lot of leftover bandwidth in the network. By employing a load-balancing algorithm with delay and quality of the link as constraints we manage to protect real-time flows from each other as well as from bursty best-effort traffic.

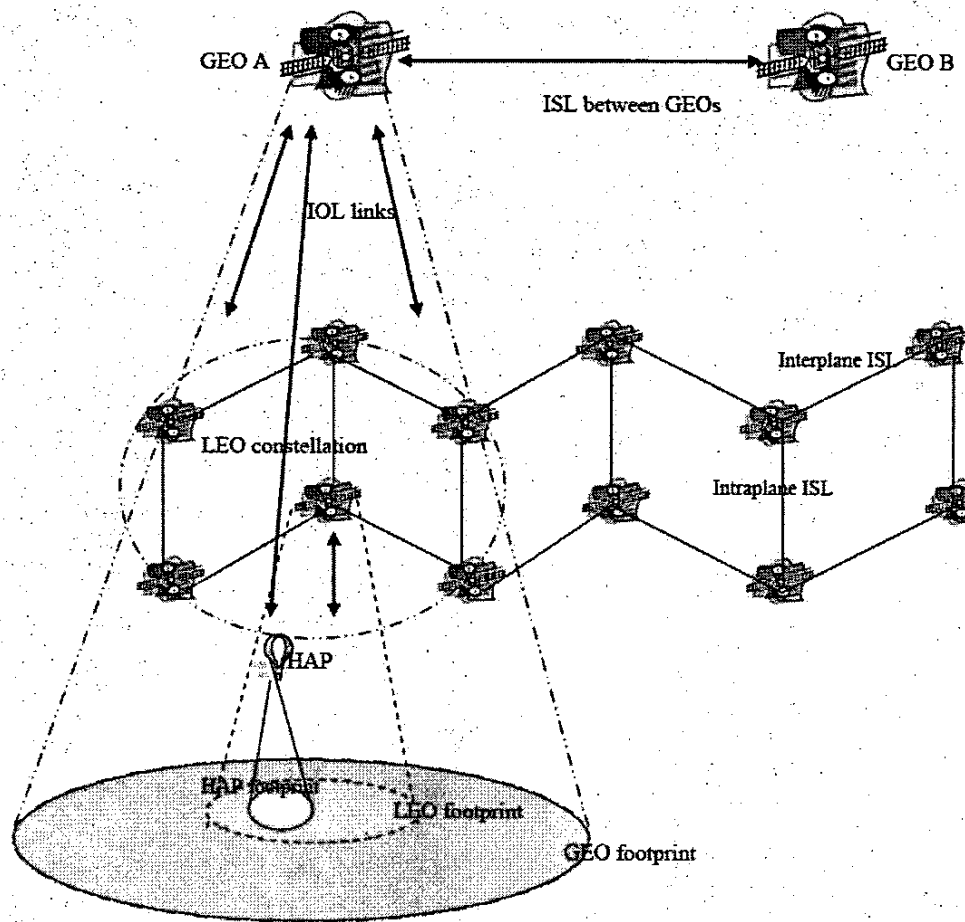


Figure 3. Proposed three layer architecture

2.1 Quality Constraint

Free-space optical communication between satellites in a distributed network can permit high data rates of communication between different places on Earth. To establish optical communication between any two satellites requires that the line of sight of their optics be aligned during the entire communication time. Because of the large distance between the satellites and the alignment accuracy required, the pointing from one satellite to another is complicated because of vibrations of the pointing system caused by two fundamental stochastic mechanisms: tracking noise created by the electro-optic tracker and vibrations derived from mechanical components [32]. Vibration of the transmitter beam in the receiver plane causes noise in the received optical power. Vibrations of the receiver telescope relative to the received beam decrease the heterodyne mixing efficiency. These two factors increase the bit-error rate (BER) of inter-satellite links. We introduce quality of

the link as an additional constraint while setting up path for loss sensitive traffic (e.g. VoIP calls).

2.2 QoS proportional to CIR

Committed Information Rate (CIR) of a flow is a service guarantee by the network provider of a certain minimum end-to-end bandwidth for the flow. A flow with higher CIR usually pays more than a user of a flow with lower CIR.

Currently both terrestrial and satellite networks employ the max-min fairness rule while dropping packets during congestion. According to the max-min fairness rule [33], each flow is allocated bandwidth as follows:

$$\text{Alloc}(i) = \text{Min} \{ \text{send}(i), rr \} \quad (1)$$

$$\text{Alloc}(i) < \text{Available Bandwidth} \quad (2)$$

Here $send(i)$ is the data rate of the i th flow and rr is the maximum rate that satisfies the above inequality. Any flow sending more than rr will have its throughput reduced to rr . In this scheme, the CIR of the flow is not considered in bandwidth allocation. Consequently a flow with higher CIR experiences packet dropping without even exceeding the CIR of the flow.

In the Proportional Allocation of Bandwidth (PAB) scheme [18], bandwidth allocated to a flow is commensurate with the CIR of the flow, i.e., QoS provided to a flow depends in the priority of the flow; higher the CIR better the QoS. Bandwidth is allocated such that all flows have identical flow rate to CIR ratio. However this requirement must be satisfied with full network utilization. Therefore in PAB the allocation of bandwidth is given by:

$$\text{Alloc}(i) = \text{Min}\{ \text{send}(i), \text{frac} * \text{CIR}(i) \} \quad (3)$$

$$\text{Alloc}(i) < \text{Available Bandwidth} \quad (4)$$

Here, $\text{CIR}(i)$ gives the CIR of the i th flow and frac is the maximum fractional multiplier (between 0 and 1) that satisfies the above inequality. The frac determines the maximum data rate of a flow as a fraction of its CIR. If the data rate of a flow is below its allowed throughput $\text{frac} * \text{CIR}$ then the flow does not suffer any packet loss. Further if a flow has a data rate less than its allowed fraction of CIR, then the remaining excess bandwidth is also shared among other flows in proportion to their CIR. No flow is allowed to send more than its CIR during congestion. The throughput of any flow sending more than the allowed fraction of CIR is reduced to its maximum allowed data rate. Thus PAB differentiates between flows and allocates bandwidth in proportion to the CIR of the flows.

2.3 Multi-protocol Label Switching (MPLS)

Due to the high-speed mobility of the nodes and ISL handoffs between satellites in the LEO layer there are several issues regarding routing of IP traffic over satellite networks [7]. MPLS has evolved as an IP-based QoS architecture, though originally developed for IP over ATM integration [2, 4]. It combines the traditional datagram service with the virtual circuit approach. Basically, a LSP is a virtual circuit that allows the setup of explicit routes for packets of a class, a critical capability for traffic engineering. Similar to the DiffServ [1] approach the ingress LSP at the edge of the network classifies a packet into classes and sets the initial label. MPLS also supports bandwidth reservation for classes, again enforced by a packet scheduler. MPLS along with constraint routing provides us with the option of routing over non-shortest paths subject to constraints of bandwidth and delay [5].

Since MPLS operates independently of layer 3 and will use IP routing methods, standard IP QoS can be enforced during the LSP setup process. LSPs with specific bandwidth requirements delay bounds can be set up using constraint-

based routing and have labels associated with them. Consequently appropriate traffic can be routed along their desired QoS path.

3. BACKGROUND AND RELATED WORK

In this Section we survey work related to our own, both to point out the many contributions of previous researchers and to place our contributions in the proper context. One of the challenges of satellite networks is the development of specialized routing algorithms. The network routing problem in LEO systems encompasses the overall service strategy (connection-oriented or connectionless), the routing strategy (centralized or distributed), the actual protocols or algorithms used to manage the dynamic nature of the network, and the satellite handover strategy used by terminals and satellites. The routing algorithms for satellite networks should compute paths with low communication and computational overhead, and adapt the routing decisions to the dynamic satellite network topology in real time.

3.1 Connection-oriented routing

Connection-oriented routing has been the focus of the research for the low-Earth orbit (LEO) satellite networks in recent years. Connection-oriented routing proposals assume ATM-like switches in the satellites [19, 20, 21, 23]. A number of papers have examined issues related to virtual connection routing in connection-oriented LEO networks. One difficulty with connection-oriented routing arises when the communications session outlasts the visibility period of the initial and final satellites of the end-to-end path the connection must necessarily be handed over to successor satellites. The connections must be established and maintained in the satellite network, which is a very dynamic environment. The probabilistic routing protocol (PRP) introduced in [23] considers a technique that can be used to select initial routes through the satellite network that have a low probability of requiring a connection reroute. There are different approaches to solve the connection-oriented routing problem. In [20], Werner proposed subdividing the time-varying LEO topology into intervals (states) of static topology, enumerating all of the possible virtual circuit combinations, and then picking a path that minimizes delay jitter by selecting a path across a series of states according to some optimization technique. Similar results by the same author are also reported in [19]. The algorithm presented in [21] uses snapshots of the constellation to optimize the paths.

3.2 Connectionless routing

With the explosive growth of the Internet, connectionless routing is being pushed to satellite networks. To realize this, satellites carry IP switches that forward packets independently. These IP switches are connected to each other as well as to ground stations. There are several

proposals regarding the IP-based routing in satellite networks [22, 24, 25, 26, 27]. Mauger and Rosenberg introduce the concept of defining a logical, virtual topology of cells on the ground, and performing routing of the packets with reference to the fixed virtual model [22]. Satellites that move above a given region become the embodiment of the virtual node. By providing a fixed virtual topology and by using virtual connections obtained through a restricted set of routing plans, the satellite network can provide quality-of-service guarantees. The authors recognize that there may be discrepancies between the virtual model and the actual interconnection of terminals to satellite links (since terminal handover may be performed independently of reassignment of satellites to virtual nodes), and compensate for this by proposing that ownership of cells is broadcast to all adjacent nodes so that routing to the final satellite in the path can be accomplished. The so-called *Darting* algorithm delays the exchange of topology update information until it is necessary to send data packets [24]. However, it is shown in [25] that the *Darting* algorithm does not reduce the protocol overhead. The datagram routing algorithm [26] aims to route the packets on minimum propagation delay paths using a distributed routing protocol. The routing protocol presented in [27] uses a hybrid approach that uses geographic-based routing and shortest path routing with limited scope.

3.3 Multi-layered routing

The paper by Shacham is one of the earliest works on multi-satellite networks to discuss many of the issues; namely, distributed routing protocols and addressing, as well as topology control and transport protocols [28]. Shacham advocates link-state routing that utilizes the predictability of topology changes and computation of multiple paths between nodes, as well as quality-of-service routing. The paper also discusses addressing, and is the first to propose basing addresses on the geographical locations of the terminals. In [9], a two-layered satellite network architecture consisting of LEO and medium-Earth orbit (MEO) satellite networks and a routing algorithm are proposed. MLSR [8] proposes a three-layer satellite architecture that calculates shortest delay paths between the satellites in the satellite network and the gateways on earth.

In our approach we use a three-layer architecture with aggregated flow routing. Simple IP switching of VoIP calls could lead to packets arriving out of order at the receiver and increased jitter. So we set up virtual circuits that can be dilated (expanding resources reserved along the path) to accommodate packets within the same Forwarding Equivalence Class (FEC) providing high bandwidth access to all types of users and low latency to delay sensitive applications.

4. A TCA ARCHITECTURE

In this Section we present a new three-layered architecture, similar to the transformational communications architecture, that we use for our simulations.

4.1 Interconnections and Coverage

The satellites maintain three types of links:

1. *User Data Links (UDLs)* between terrestrial users and the satellite network. A user can have UDLs to multiple satellites in each of the layers and vice-versa.
2. *Inter orbit links (IOLs)* exist between the layers of the satellite hierarchy. Each satellite is linked by IOLs to the satellites under its coverage as well as to the satellites in the upper layer that cover it.
3. *Inter Satellite Links (ISLs)* exist in the LEO layer which can be interplane ISLs, intraplane ISLs or cross-seam ISLs. The HAPs also communicate with each other through line of sight optical links.

Three satellites in the GEO orbit are sufficient to cover the entire earth surface (between 81°N and 81°S). Three GEOs divide the earth surface into three fixed domains, the size of which correspond to the elevation angle of the GEO satellites since the GEOs will always be over their respective positions with respect to the earth surface.

The second layer consists of satellites in the LEO orbit. Satellites in the LEO orbit revolve around the earth at high speeds, which introduces handoffs, since a satellite now serving a region will fall below the elevation mask as it moves away and has to handover traffic to the incoming satellite. Handoffs can be performed asynchronously or synchronously. In asynchronous handoffs the footprint of the satellite moves along with the satellite (satellite-fixed cells) and calls are handed off to the satellite following it in a successive fashion. Whereas, in synchronous handoffs the footprints are fixed (the earth's surface is divided into fixed cells) and the satellite currently over the region hands over the whole footprint when it moves out to the successor satellite. In an earth-fixed system, the satellites continuously train their antennas onto a fixed footprint for a period of time during which they are visible over the region [34, 35]. For example Teledesic is a proposed commercial LEO satellite constellation with a earth-fixed footprint system [36]. In this paper we assume earth-fixed footprints. Figures 4(a) and 4(b) show the satellite-fixed and earth-fixed mechanisms respectively.

Figure 5 shows the footprints resulting from a two layered architecture with satellites both in the GEO and LEO orbits, superimposing the footprints of the satellites in the GEO layer over the earth-fixed footprints for the LEO satellites. Assuming A, B & C are the three satellites in the GEO layer

we have three domains where each domain is further subdivided into cells/sub-domains, the size of which corresponds to the footprint of the LEO satellites, assuming earth-fixed footprints for LEO satellites. Depending on




Figure 4. Mobile and Fixed Footprints




Figure 5. Domains and Cells

whether HAPs are deployed, the cell/sub-domain could be further subdivided into smaller cells (see Figure 3).

4.2 Addressing

In Figure 5 DOMAIN A represents the footprint from GEO A, sub-divided into cells by several LEO footprints and the footprint of LEO A.4 is further subdivided by the presence of HAPs under A.4 into A.4.1 and A.4.2.

Addressing for the purpose of deciding the egress node is also derived from the hierarchical structure. In the above figure a user in cell A.4.2, where $A.4.2 \subset A.4 \subset A$, can be reached by all three: HAP (A.4.2), LEO (A.4) and GEO (A). Consequently, all users in cell A.4.2 can have an address prefix (e.g.A.4.2) that identifies the egress nodes through which the user can be reached.

We assume earth-fixed footprints for the LEO satellites, so that handovers between LEO satellites of ground to satellite links (GSLs) are always synchronous and periodic. In Figure 5 when the LEO I over cell A.1 moves over to cell A.2, and LEO II over cell A.2 is moving out, LEO II forwards its routing table with the active connections and their next hops to LEO I. An active connection is always routed over the same logical topology, even though the LEO satellites keep changing. We reduce the dynamic LEO constellation into a logically fixed topology [8].

5. HIERARCHIAL ROUTING

GEOs being at the highest altitude act as the “eyes” for the satellite constellation, and keep track of the LEO and HAP movement and interconnections under its coverage. Furthermore most of the interconnections (e.g. the ISLs and cross-seam links between LEO satellites) are periodic and can be predicted.

5.1 Static Connection Matrix

Given the regular topology of the satellite networks, our logically fixed topology for the LEOs and the assumption that HAPs can be considered stationary with respect to LEOs and GEOs, a static connection matrix can be easily maintained and updated at regular intervals or manually (in case HAPs are deployed or satellite/link failures).

The static connection matrix gives us all the interconnections between all the nodes in the constellation and the propagation delays between them. Since propagation delay varies (but is periodic and can be predetermined) for inter-plane ISLs in the LEO layer, the matrix is periodically changing. Figure 6 shows an interconnected topology. The shaded portion refers to the cross-seam plane across which the ISLs will be switched off rapidly because of counter-rotating orbits. The GEOs are connected to each other through line-of-sight links and to the satellites in the lower layers through IOLs. Satellites in adjacent orbits/planes in the LEO layer are interconnected by interplane ISLs at all times except at cross-seams where links are handed over very fast, and at poles where they are switched off. The link propagation delay on interplane ISLs changes with latitude (link propagation delay reduces as the latitude increases, since orbits converge at the poles), but since this behavior is also periodic the link delays between neighbors can be predicted. Adjacent satellites in the same orbit/plane are interconnected by intra-plane ISLs, which

are never switched off and have constant delay. Apart from being periodically uploaded it also reacts to sudden topology changes e.g. induction of new nodes to the constellation (whenever HAPs are deployed over a region, their IOL delays to both the LEO and GEO satellites are uploaded to the matrix).




Figure 6. Example Connection Topology

In essence the connection matrix gives us the next-hop neighbors for any node with the propagation delay of the link between them. Table 1 is a fragment of the connection matrix for the topology in Figure 6 (\checkmark indicates presence of a direct link and \times indicates they are not next hop neighbors). The static connection matrix is uploaded (via the satellite in the GEO layer or terrestrial gateways) to every node on the network periodically. The static connection matrix can be uploaded by the respective GEO satellites in every domain since each GEO can communicate with all nodes (LEO and HAP) under its coverage; thus when a HAP is deployed within a GEO’s domain, the static matrix is uploaded by the GEO upon recognizing the HAP within its domain.

Table 1. Static Connection Matrix

	A	B	B.1	B.2	B.3	A.7	A.8	A.9
A	N/A	\checkmark	\times	\times	\times	\checkmark	\checkmark	\checkmark
A.7	\checkmark	\times	\checkmark	\times	\times	N/A	\checkmark	\times
A.8	\checkmark	\times	\times	\checkmark	\times	\checkmark	N/A	\checkmark
A.9	\checkmark	\times	\times	\times	\checkmark	\times	\checkmark	N/A

5.2 Optimal Path Selection

In this section we outline the steps taken to select an optimal path for a flow with two objectives: one is to ensure Quality of Service (QoS) for priority traffic within the delay

constraints for the flow and the other is to ensure that the load on the network is balanced.

In any form of QoS routing with some constraints the following have to be considered:

1. Ability to compute a path at the source, and to compute the path in such a way that the computation can take into account not just some scalar metric but also a set of constraints that should not be violated.
2. Ability to distribute the information about network topology and attributes associated with links throughout the network.
3. Ability to route explicitly.
4. Ability to reserve resources and subsequent modification and distribution of the link attributes.

We will be dealing with all the above requirements in the rest of the paper.

Given the costly nature of satellite resources over-provisioning is not an option in satellite networks, so we resort to a form of class-based queuing. The overall bandwidth for each link is split, so that we have two queues, a fraction μ ($0 < \mu < 1$) of the link capacity for high priority traffic and the rest $(1-\mu)$ for best-effort traffic. Before reserving resources for a high priority call/flow on a link, we have to see whether bandwidth needed for the flow is available on the link (from now on whenever we refer to setting up paths it implies high priority traffic; no resources are reserved for best effort traffic). The CIR (see Section 1) for the flow is the bandwidth that has to be reserved on the link. Hence the summation of the CIR's of the set of flows (f) has to be less than or equal to the fraction (μB) of the link bandwidth (B) reserved for high priority traffic.

$$\mu B \geq \sum_f \text{CIR}(f) \quad (5)$$

No admission control is implemented for best-effort calls; consequently if bandwidth normally reserved for high priority traffic is available, best-effort calls will be entertained on the link, but may be switched or pre-empted if a request for bandwidth comes from a high priority call. By dividing the queue with the fraction reserved for high priority traffic, we essentially separate the real-time VoIP flows from TCP flows.

5.3 Constraint Routing

Apart from bandwidth availability, we also consider delay and quality of signal (BER, see Section 1) of a path as additional constraints during the path selection process.

Let Ω be the generic cost function associated with a path p , related to the network topology and quality of the links on the path, where Ω is a function of the total propagation

delay (D_p) on the path, and the quality degradation (bit error rate) of the path (Q_p).

$$D_p = \sum_l d(l) \quad (6a)$$

$$Q_p = \sum_l q(l) \quad (6b)$$

where l is the set of links on the path, d is the link delay and q reflects the quality degradation (bit error rate) of the link. The routing function first determines the set of paths satisfying the generic cost requirements, i.e. the maximum delay bound (D) and minimum quality degradation (Q) acceptable for the specific traffic. Paths with Q_p and D_p beyond the acceptable limits for the flow are rejected.

$$\Omega(p) = \infty \quad ; \quad \{D_p > D, Q_p > Q\} \quad (7a)$$

$$\Omega(p) = D_p \quad ; \quad \{D_p \leq D, Q_p \leq Q\} \quad (7b)$$

Paths $p1$ and $p2$ are considered equivalent, provided they are between the same source and destination pair, when $\Omega(p1) = \Omega(p2)$.

For each set of equivalent paths, we then associate a dynamic cost based on the residual bandwidth of the path. For each link (l) along the path in which we have an available bandwidth B_l , the residual bandwidth of the path is:

$$R(p) = \min(B_l) \quad (8)$$

where link l belongs to path p .

For high priority calls the residual bandwidth is in terms of B_p (bandwidth assigned for high priority traffic) assigned for that link, whereas for best-effort traffic it is in terms of the available bandwidth on the link.

The overall cost Ψ to the path p is then assigned:

$$\Psi(p) = \Omega(p) / R(p) \quad (9)$$

The paths are then sorted according to Ψ , the first path in the set being the primary path and the other being the secondary paths.

5.4 Dynamic Traffic Load

From the static connection matrix and link propagation delays we have a static snapshot of the network with no knowledge about the traffic load.

Intra-domain link state exchange

Within each domain all the nodes (satellites in LEO, GEO and HAPs) measure their output buffers on all outgoing links (GSL, ISL and IOL) to determine the residual bandwidth on the links and flood it within the domain. Each node within the domain now has information about the bandwidth availability on all the links within the domain and also to the border nodes of adjacent domains (e.g. see

Figure 6 link state of the link between A.7 and B.1 is also flooded within the domain). This link state exchange along with the static connection matrix is used to build the intra-domain routing table for each node and can be used to find paths within the domain as well to the border nodes of adjacent domains.

Information Exchange across Domains (Inter-domain)

One of the distinct advantages of dividing the network into domains is reducing the overhead of information exchange (remote nodes do not have to exchange information) and also keeping information exchange local prevents the nodes from using outdated information.

An aggregated routing table for a domain (e.g. Domain B) includes the maximum residual bandwidth paths from each border node at one end of the domain (e.g. B, B.1, B.2 & B.3) to every border node at the other end of the domain (e.g. B, B.13, B.14 & B.15). Table 2 represents the information in the aggregated routing table for domain B under the satellite GEO B (see Figure 6) with one entry shown i.e. from border node B.1 to B.13 with the available bandwidth and quality of the path.

Table 2. Aggregated Routing Table for Domain B

	B	B.13	B.14	B.15
B	N/A			
B.1		B.1-B.4-B.7-B.10-B.13 {10Mb, Q factor}		
B.2				
B.3				

The aggregated routing table reduces a domain into a single logical node with several ingress links and egress links with all combinations of switching the flow from any ingress point to an egress point along with the corresponding available bandwidth.

5.5 Routing Strategy

We outline the steps to setting up the LSP, as the example below elucidates (refer to Figure 7). *S* and *D* are the source and destination nodes respectively.

- *S* looks up the static connection matrix to determine the domain and location of *D* in the network topology and also the minimal delay paths.
- Since Domain B is the transit domain between the source destination pair, the aggregated routing table for Domain B and the intra-domain routing table for Domain A are used by *S* to determine the available paths to reach Domain C.
- The paths are ranked according to the steps outlined in Section 5.2.

- Routing from the border node in C to the egress node *D* is done intra-domain in domain C during the LSP setup process.




Figure 7. LSP setup for S-D pair

6. TRAFFIC AGGREGATION

In this Section we outline the traffic aggregation and packet dropping policy implemented.

6.1 Real-Time Traffic Aggregation

The topology and link attribute distribution have already been outlined in Section 5, so mechanisms to handle requirements (1) and (2) mentioned in Section 5.2 are already discussed.

Here we outline the procedure to establish explicit routes (requirement (3) in Section 5.2). In the previous Section we outlined the procedure of an LSP setup. Each LSP is associated with a Forwarding Equivalence Class (FEC) so all packets with the same FEC are switched on that LSP. Once an LSP is setup, all nodes, including the source on that LSP, can send packets that have the same destination and class; accordingly when a new call has to be accommodated on the same LSP, the bandwidth of the LSP is dilated (provided bandwidth is available) according to the requirements (Committed Information Rate) of the new call. By implementing packet forwarding instead of connection-oriented flow switching, we intend to improve the scalability of the scheme. The example below explains how the forwarding component works. We have already explained how routing between a source destination pair belonging to two different domains, is accomplished. Here for the sake of simplicity both the source and destination belong to the same domain.

In Figure 8, A is the source node and B is the destination node. We have four calls originating from A and ending at B. Say path I has 10Mb for the path between A to B and path II has 10 Mb, but path I has less end-to-end delay and

so is the best path. Paths I and II are not necessarily completely disjoint.




Figure 8. QoS Routing

When flow 1 (CIR=4Mb/s) comes in, LSP-I is setup along path I and resources along it are reserved according to the CIR of the flow. Flow 2 (CIR=6Mb/s) comes in next. Since it belongs to the same FEC as call 1, provided path I still has 6Mb bandwidth, LSP-I is dilated from 4Mb/s to 10Mb/s and call 2 is forwarded on LSP-I. Flow 3(CIR=5 Mb/s) comes in next, but since there is no more available bandwidth on Path I, source A goes for Path II and LSP-II is setup with an bandwidth of 5 Mb/s and so on.

Absolute IP packet forwarding would result in packets from the same call being routed over different paths, which would end up increasing jitter. By aggregating packets and switching them as bundled flows we minimize overhead as well as ensure that packets belonging to a particular call are always switched through the same path from source to destination.

Since several calls with the same FEC will be sharing the same LSP, there is always the possibility of a flow exceeding its CIR (Committed Information Rate), i.e., Flow 1 increases its rate to 6Mb/s. This could lead to packet dropping for both flows. Also since the CIR of flow 2 is 6Mb/s, it is obvious that the sender of flow 2 is paying more and it should be protected.

6.2 Packet Dropping Policy

Our Packet Dropping Policy guarantees that during congestion flows get a share of IP available bandwidth, which is in proportion to their CIR. This is achieved by implementing a form of dropping policy [18], based on the principles of Differentiated Services [16, 17].

At the edge of the network packets are marked with labels, which encode the ratio of flow rate to CIR for the flow. These labels are used to differentiate between packets from different flows by the core routers. The core routers perform multi-level threshold based dropping. Packets from flows that exceed their CIR will be marked with lower priority labels. The excess traffic will have higher probability to be dropped during congestion, but can use the network if there are available resources. Our technique to implement Proportional dropping [18] involves two main components:

the labeling of the packets at the edge of the network, and dropping of packets at the core router.

7. SIMULATION RESULTS

In this Section we outline the simulation details; tools, models and parameters used for quantitative evaluation of our proposed scheme.

We have used the Network Simulator [3] as our simulation tool. NS-2 is a discrete event simulator targeted at networking research. It is an open-source simulator available for free public use. NS-2 is available on several platforms such as FreeBSD, Linux, SunOS and Solaris. NS-2 also builds and runs under Windows. In our simulations NS-2 is built over a Linux operating system. NS-2 provides substantial support for simulation of transport protocols (e.g. TCP), routing, and multicast protocols over wired and wireless (local and satellite) networks. NS-2 is an object-oriented simulator written in C++, with an OTcl interpreter as the front-end. As a user one can define network objects such as nodes, the interconnections between them and events (sending packets, link outages, stopping packet transmission) to simulate a particular scenario. The trace files generated are analyzed to determine the quantitative performance of the scheme being simulated.

7.1 Simulation Configuration

Figure 3 is a snapshot of our simulation model with all three layers and their interconnections. In all simulations we use a polar-orbiting configuration with 66 satellites as the LEO layer. Furthermore, we assume three equally spaced GEOs on the equatorial orbit. VoIP can adequately represent high priority traffic, as it is interactive and has strict delay and jitter requirements. A high priority call is interpreted as an aggregate of several VoIP calls of the order of Megabits/s(Mb/s) where each VoIP call is an exponential on-off source at 8 Kb/s according to the G.729 standard [14]. Background web traffic was simulated as TCP flows with infinite ftp sources. In our study we analyze three real-time VoIP calls (RT1, RT2, RT3) originating from the terrestrial ground station S and terminating at terrestrial ground station D. The queue was split with $\mu = 0.8$ (fraction of queue reserved for VoIP traffic).

Figure 9 shows the simulation model pictorially, with three GEO satellites in an equatorial orbit and the LEO satellites in a polar orbiting constellation. Table 3 lists the parameters for the satellites in the LEO and GEO layers. These are the default parameters for all our simulations unless otherwise specified. The HAPs are only local to the LEO satellites under whose coverage they are deployed and move within a 20 km radius.




Figure 9. Simulation Model

7.2 Delay Packet Loss and Jitter Analysis

In this section we compare the performance of the multi-path algorithm on the multi-layered architecture with that of Bellman's shortest-path algorithm [37]. We use the network topology and parameters outlined in Section 7.1 for simulations in this section. Both the multi-path algorithm proposed in this paper and Bellman's shortest path algorithm are used to route traffic through the network in separate instances and the output data analyzed. Three VoIP (real-time traffic) generators are attached to the source node, all three flows have the same destination node. Background traffic in the form of TCP flows is generated between the source-destination pair.

The parameters measured are delay, jitter and packet loss. We define the parameters as pertaining to our simulations: *Delay* is measured as the time interval from when a packet is queued at the output queue of the source node to when it is received at the destination node. Delay can be measured as either one-way or round-trip delay. VoIP typically tolerates delays up to 150 ms beyond which the quality of the call becomes unacceptable. In our simulations we measure one-way mean delay.

Jitter is the *variation* in delay over time, between arrivals of consecutive packets. High values of jitter degrade the quality of a VoIP call. We measure jitter as the standard deviation of one-way delay.

Packet loss is expressed as the percentage of packets lost in transit between the source and the destination. The optimum network requirements for acceptable VoIP traffic are:

System Parameters	LEO Layer	GEO Layer
Altitude (km)	780	35786
Orbit Type	Polar	Equatorial
Number of Planes	6	1
Satellites in a Plane	11	3
Inter-ISLs	2	-
Intra-ISLs	2	2
ISL Bandwidth (Mb/s)	5.0	20.0

• One way delay

- $\mu < 150\text{ms}$
- Jitter $< 30\text{ ms}$
- Packet Loss $< 3\%$

Scenario 1: Shortest Path Routing. In this case shortest path routing is implemented to route the traffic between the source-destination pair. Bellman's algorithm calculates the shortest path between two nodes, without taking into account residual bandwidth on the path, all the flows RT1, RT2, RT3 and the best effort TCP flows are routed on the same shortest path. Table 4 shows the packet loss, delay and jitter values for the three VoIP flows (RT1, RT2, RT3) with shortest path routing. In the absence of any form of load balancing all the traffic flows, VoIP and best-effort traffic are routed along the shortest path between the nodes.

Table 3. Simulation Parameters

Table 4. Shortest Path routing

Traffic Source	Packet Loss(%)	Mean Delay(ms)	Jitter(ms)
RT 1	11.89	172.477	54.0916
RT 2	22.46	187.028	39.526
RT 3	36.34	201.222	22.362

Packets from all three VoIP flows and the best effort TCP flows share the same output queue even though alternate paths between the source destination pair are underutilized. Without any form of reservation of resources or priority policy the router does not discriminate between packets from VoIP flows and best-effort flows and packets are dropped indiscriminately during congestion leading to high values of packet loss for the three VoIP flows. Sharing the same queue with best-effort traffic introduces increased and variable delay for VoIP packets increasing the jitter and mean delay values.

Scenario 2: Load Balancing. Now we implement the multi-path algorithm on the network topology (Section 7.2). The output queues are split with a fraction ($\mu=0.8$) reserved for the VoIP flows. Path selection for the VoIP calls follows the steps outlined in Chapter 4. Best-effort traffic (TCP) flows are routed along the shortest path; no load balancing is implemented for the TCP flows. Table 5 shows the packet loss, delay and jitter values for the three VoIP flows with the load balancing multi-path algorithm in effect. MPLS constraint routing is used to set up Label Switched Paths for the three VoIP flows.

Table 5. With Load balancing

Traffic source	Packet Loss(%)	Mean Delay (ms)	Jitter(ms)
RT 1	0.67	65.0703	4.697
RT 2	0.71	60.4791	1.486
RT 3	1.0	80.9157	1.58593


The average constraint-based LSP setup time was 0.06 seconds. The VoIP flows are now protected from the best-effort packets with 0.8 of the link capacity reserved for them. Also with load balancing the VoIP flows are routed over alternate paths, within the delay bounds, increasing network utilization while reducing the jitter and delay values. RT1 experiences some jitter since all of the best effort traffic flows along the path used by RT1. By reserving resources for the VoIP calls we reduce the packet loss. Setting up Label Switched Path for a VoIP flow ensures packets belonging to the flow follow the same path always and do not need re-ordering at the receiver and reduce jitter.

Comparison. We compare the performance of the load-balancing algorithm with Bellman's shortest path algorithm as the background traffic increases. Average values of delay, jitter and packet loss for the three flows (RT1, RT2 and RT3) are plotted against increasing best-effort traffic.

Figure 10 is a plot of the average mean delay of the three flows against increasing background traffic. As the background best effort traffic is increased the values for average delay for the three VoIP flows with shortest path routing keep increasing as the output queue gets increasingly congested. Whereas, increasing background has little or no effect on the VoIP flows within the load-balancing scheme, since they are protected from the best effort traffic as well as routed on paths with available bandwidth and not necessarily on the shortest-path only.

Figure 11 is a plot of the average jitter of the three flows against increasing background traffic. As the background best effort traffic is increased the values for average jitter for the shortest path routing keep increasing, whereas with the load-balancing scheme the jitter values increase very gradually and are within acceptable limits for a VoIP call.

Figure 12 is a plot of the average packet loss of the three flows against increasing background traffic. As with delay and jitter VoIP calls within the shortest path scheme suffer increasing packet loss with increasing background traffic. Resource reservation and class based queuing (split queue) implemented for the load-balancing scheme ensures packets from the VoIP flows do not compete with best-effort traffic and hence do not suffer even though the best effort traffic load is increased.

**Figure 10.** Average mean delay




Figure 11. Average jitter




Figure 12. Average Packet Loss

7.3 Goodput and Utilization Ratio

We define Goodput as the effective data rate perceived at the receiver's end. In this section we measure the effectiveness of the QoS mechanism, i.e., of isolating the VoIP traffic from the best-effort traffic. Simulations are carried out with a bottleneck link in the path set up for the VoIP flow. The bandwidth of the bottleneck link is 2.3 Mb/s. In one instance of the simulation the VoIP flow shares the link with the best effort traffic without any QoS mechanism in place, then we measure the goodput of the VoIP flow with QoS routing where the VoIP flow is isolated from the best effort traffic by splitting the queue and reserving resources for the VoIP flow. Figure 13 is a plot between increasing background traffic against VoIP goodput. As we can infer from the plot without QoS routing the VoIP goodput reduces with increasing background traffic, i.e., with the best effort traffic increasing the goodput possible for the VoIP flow keeps decreasing in order to have acceptable delay and jitter values for the VoIP flow. Whereas, with the QoS mechanism in effect the VoIP packets are isolated from the best-effort packets keeping the VoIP goodput constant even when background traffic is increased.




Figure 13. Goodput vs. Background Traffic

with the QoS scheme the isolated VoIP flow does not have to compete with best-effort traffic and VoIP, goodput is maintained at the receiver node without having to dilate the bandwidth of the bottleneck link and the utilization ratio remains independent of the best effort traffic.

Next we analyze the utilization of the bottleneck link in terms of the real-time goodput flowing through it. We define a term *utilization ratio* where:

Utilization ratio = real-time goodput / Link bandwidth
 The bandwidth of the VoIP call on the link is kept constant. The goodput of the VoIP call is measured at the receiving node and is kept constant at the cost of increasing the bandwidth of the bottleneck link. Figure 14 is a plot of the utilization ratio against increasing background traffic. For the simulation instance without the QoS scheme the link bandwidth of the bottleneck has to be increased with increasing background traffic to maintain the VoIP goodput, which translates to a decrease in the *utilization ratio*. While

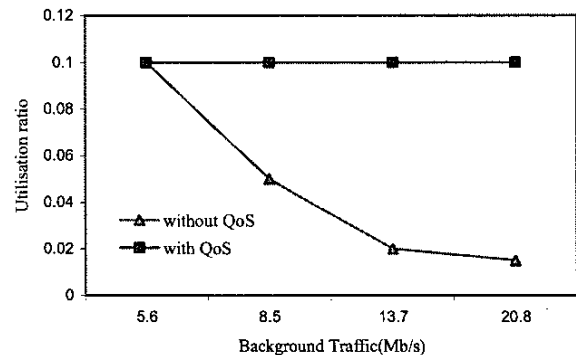


Figure 14. Utilization ratio vs. Background traffic




Figure 15. Multiple Congested links

7.3 Congestion on Multiple Links and Proportional Dropping

In this section we analyze the performance of the proportional dropping policy. Aggregating packets from several VoIP flows onto the same LSP could result in unfair dropping of packets during congestion (see Sections 6.1 and 6.2).

We simulate the configuration shown in Figure 15. Flows travel different distances in the network. There are $N+1$ terrestrial routers. A terrestrial router is connected to the next terrestrial router through the satellite router. At the routers R_0 to R_{N-1} , flows enter the network and at the router R_N , all the flows leave the network. At router R_0 , flow S_0 enters the network. At router R_i flows $S(i*2)+1$ to $S(i+1)*2$ enter the network. In each experiment set, the number of congested links varied from 4 to 6.

The total CIR of all the flows is less than the link bandwidth, and flows that exceed their CIR experience packet dropping during congestion.

Performance measure

To calculate the effectiveness of the dropping policy we define an Allocation Ratio $[AR(i)]$ for source i , which is calculated as;

If $\{Flow\ rate\ (i)\ \leq\ CIR(i)\}$

If $\{Flow\ rate\ (i)\ \leq\ CIR(i)\}$

$$AR(i) = Throughput(i) / Flow\ rate(i)$$

Else $\{Flow\ rate\ (i)\ > CIR(i)\}$

$$AR(i) = Throughput(i) / CIR\ (i)$$

Ideally for a VoIP flow the $AR(i) \cong 1$ i.e. the maximum bandwidth the flow can have during congestion is not more than its CIR.

Figure 16 is a graph of the allocation ratio for Source (0) against increasing traffic (Mb). Each increase in link number by one increases the number of competing sources by two, increasing the traffic load. With proportional dropping implemented, the flow from Source $S(0)$ receives its share of the bandwidth, whereas without the policy the throughput for the flow keeps dropping as the load offered to the network is increased.




Figure 16. Allocation ratio For UDP flow $S(0)$

7.4 Quality Constrain

The free space optical links are prone to several optical impairments. An optical correlator that can measure online the quality degradation of optical links is presented in [29]. We use the information about the optical signal quality in our routing scheme.

In this section we compare our protocol with quality degradation of a link as constraint to a shortest-path first protocol. Goodput of the network is defined as the number of calls of acceptable quality established across the same. We define quality degradation of a link as Q_i which is the BER of the link i . For a path p :

$$Q_p = \sum_i Q_i \quad (10)$$

If Q_p is within the acceptable quality desired for the call, the path is within the set of paths along which the call could be routed (see Eq 7(a)).




Figure 17. Comparison of Network Goodput

Figure 17 compares the goodput of a shortest path protocol to our quality-constrained protocol. Network goodput (calls of acceptable quality established) for the routing protocol with quality gives a higher number of acceptable connections because it can adapt based on the quality of links in the network. As the number of connection requests increase, more calls become unacceptable in the shortest path first routing protocol.

7.5 Communication Overhead Comparison

In this section we analytically compare the communication overhead of routing table calculations in our domain based routing protocol with that of a flat distributed routing protocol.

In Figure 18 is illustrated the comparison between the two protocols. The X-axis represents the number of satellites in a single layer topology. The Y-axis represents the total

communication overhead in terms of transmission units, where a transmission unit represents an intra-domain link state packet or an inter-domain aggregated link state packet. The distributed protocol with no domains or hierarchical structure and with distributed routing table calculation experiences a sharp increase in communication overhead as the number of satellites in the topology increases. Every satellite in the topology broadcasts link state information to every other satellite in the network. Once a satellite receives all the link state broadcasts from every other node in the network it constructs the routing table.

The domain routing protocol, however, limits link-state broadcasts to between satellites within the same domain only, and an aggregated routing table for the neighboring domains is broadcast within the domain by the GEO satellite of the domain. For the domain protocol, the X-axis in Figure 18 represents the satellites in the LEO layer only, other than the three GEO satellites, which divide the LEO layer in three domains. The increase in communication overhead in the domain protocol is a lot less as compared to the flat distributed protocol and consequently more scalable as the size of the topology increases.




Figure 18. Communication Overhead Comparison

8. CONCLUSIONS

We studied a TCA-like three-layer architecture consisting of satellites in the GEO layer, LEO layer and HAPs, which allows terrestrial users with hand-held terminals to connect to the high-speed satellite network. Ease of deploying HAPs and integrating it with the rest of the network make it an attractive option for global access. Dividing the satellite network into domains based on the GEO satellite footprints reduces the communication overhead for the network. Given the geography sensitive nature of satellite networks where a satellite over a major city might experience heavy traffic, whereas a neighboring satellite over an ocean is under-utilized, the load-balancing algorithm ensures that traffic from neighboring satellites does not further clog the already loaded satellite and also improves the network utilization. Given the sensitive nature of VoIP calls to

congestion, dividing the queue and isolating them from low priority best effort traffic ensures that optimal delay, jitter and packet loss requirements for the VoIP flows are maintained across the network. Flow aggregation between VoIP flows with the same FEC increases the granularity of switching reducing overhead at the core routers. The proportional dropping policy protects VoIP flows from each other, ensuring that during congestion only packets from a VoIP flow, which has exceeded its CIR, are dropped. The Multi-level dropping policy ensures flows receive QoS commensurate with their CIRs. High BER on inter-satellite links due to mechanical vibration and antennae tracking issues motivated the quality constraint, which improves the number of acceptable calls established across the network.

REFERENCES

- [1] Li, T., Rechter, Y. "A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)", RFC 2430, October 1998.
- [2] Awduche, D.O., Malcolm, J., O'Dell, M., McManus, J. "Requirements for Traffic Engineering Over MPLS", draft-ietf-mpls-traffic-eng-00.txt, October 1998
- [3] Network Simulator-ns2. <http://www.isi.edu/nsnam/ns>
- [4] Rosen, E.C., Vishwanathan, A., Callon, R. "Multiprotocol Label Switching Architecture", draft-ietf-mpls-arch-04.txt, Feb. 1999
- [5] Praveen Bhaniramka, Wei Sun, Raj Jain "Quality of Service using Traffic Engineering over MPLS: An Analysis," Proc. 25th Annual IEEE Conference on Local Computer Networks (LCN 2000), Tampa, Florida USA, November 8 - 10, 2000, pp. 238-241.
- [6] T. R. Henderson and R. H. Katz, Network Simulation for LEO Satellite Networks, 18th AIAA International Communications Satellite Systems Conference (ICSSC), Oakland, California, April 2000.
- [7] L. Wood, A. Clerget, I. Andrikopoulos, G. Pavlou, and W. Dabbous, "IP Routing Issues in Satellite Constellation Networks," International Journal of Satellite Communications, vol. 19, no. 1, pp. 69-92, January/February 2001.
- [8] Ian F. Akyildig, Eylem Ekici, Michael D Bender. "MLSR: A Novel Routing Algorithm for Multilayered Satellite IP Networks" IEEE/ACM Transactions on Networking, Vol. 10, No. 3, June 2002.
- [9] J. Lee and S. Kang. "Satellite over Satellite (SOS) network: A novel architecture for satellite network." In Proc. IEEE INFOCOM, vol. 1, March 2000, pp. 315-321.
- [10] T.R. Henderson "Networking over next generation Satellite Systems", PhD thesis, University of California at Berkeley.
- [11] <http://www.bakom.ch/imperia/md/content/english/funk/forschungundentwicklung/studien/HAPS.pdf>
- [12] <http://www.elec.york.ac.uk/comms/haps-main.html>
- [13] M. Antonini, E. Cianca, A. De Luise, M. Pratesi, M. Ruiggieri, "Stratospheric Relays: Potentialities of New Satellite-High Altitude Platforms Integrated scenarios", IEEE Aerospace Conference 2003, Big Sky, MT, March 8-15, 2003.
- [14] International Telecommunication Union. <http://www.itu.int/rec/recommendation.asp?type=folders&lang=e&parent=T-REC-G.729>
- [15] Christine M. Anderson, Transformational Communications, presented at GSAW 2002, available online at: <http://sunset.usc.edu/GSAW/gsaw2002/s8/canderson.pdf>
- [16] S. Blake et al. "An Architecture for Differentiated Services", RFC 2475, December 1998.
- [17] Z. Cao, Z. Wang, E. Zegura, "Rainbow Fair Queuing: Fair Bandwidth Sharing Without Per-Flow State", *Proceedings IEEE INFOCOM 2000*, Conference on Computer Communications, p.922-31 vol.2.
- [18] A. Durrezi, P.K. Jagannathan, R. Jain "Scalable Proportional Allocation of Bandwidth in Satellite IP Networks" IEEE Aerospace Conference 2003, Big Sky, MT, March 8-15, 2003.
- [19] M. Werner, C. Delucchi, H. Vogel, G. Maral, and J. De Ridder, "ATMbased routing in LEO/MEO satellite networks with intersatellite links," *IEEE J. Select. Areas Commun.*, vol. 15, pp. 69-82, Jan. 1997.
- [20] G. Berndl, M. Werner, and B. Edmaier, "Performance of optimized routing in LEO intersatellite link networks," in *Proc. IEEE 47th Vehicular Technology Conf.*, vol. 1, May 1997, pp. 246-25.
- [21] H. S. Chang, B.W. Kim, C. G. Lee, Y. Choi, S. L. Min, H. S. Yang, and C. S. Kim, "FSA-based link assignment and routing in low-earth orbit satellite networks," *IEEE Trans. Veh. Technol.*, vol. 47, pp. 1037-1048, Aug. 1998.
- [22] R. Mauger and C. Rosenberg, "QoS guarantees for multimedia services on a TDMA-based satellite network," *IEEE Communication Magazine*, vol. 35, pp. 56-65, July 1997.

- [23] H. Uzunalioglu, I. F. Akyildiz, and M. D. Bender, "A routing algorithm for LEO satellite networks with dynamic connectivity," *J. Wireless Networks (WINET)*, vol. 6, no. 3, pp. 181–190, June 2000.
- [24] K. Tsai and R. Ma, "Darting: A cost effective routing alternative for large space-based dynamic topology networks," in *Proc. IEEE MILCOM*, 1995, pp. 682–687.
- [25] R. A. Raines, R. F. Janoso, D. M. Gallagher, and D. L. Coulliette, "Simulation of two routing protocols operating in a low earth orbit satellite network environment," in *Proc. IEEE MILCOM*, vol. 1, Nov. 1997, pp. 429–433.
- [26] E. Ekici, I. F. Akyildiz, and M. D. Bender, "A distributed routing algorithm for datagram traffic in LEO satellite networks," *IEEE/ACM Trans. Networking*, vol. 9, pp. 137–147, Apr. 2001.
- [27] T. R. Henderson and R. H. Katz, "On distributed, geographic-based packet routing for LEO satellite networks," in *Proc. GLOBECOM*, vol. 2, Dec. 2000, pp. 1119–1123.
- [28] N. Shacham. Protocols for Multi-Satellite Networks. *Proceedings of IEEE MILCOM '88*, pages 501–505, 1988.
- [29] Betty Lise Anderson, Feras Abou-Galala, Arjan Durresi. High-resolution optical correlator for monitoring health of an optical link,. Accepted for publication in "Applied Optics", 2003.
- [30] Clarke, Arthur C. "Extra-Terrestrial Relays: Can Rocket Stations Give World-wide Radio Coverage?" *Wireless World*, October 1945, p. 306.
- [31] T. C. Tozer and D. Grace, "High Altitude Platforms for Wireless Communications" *Electronics and Communication Engineering Journal*, June 2001.
- [32] Shlomi Arnon, Stanley R. Rotman, and Norman S. Kopeika, "Performance limitations of a free-space optical communication satellite network owing to vibrations: heterodyne detection" *Applied Optics*, Vol. 37, No. 27, 20 September 1998.
- [33] I. Stoica, S. Shenker, H. Zhang, "Core-Stateless Fair Queuing: Achieving Approximately Fair Bandwidth Allocations in High Speed Networks", *Computer Communication Review*. vol.28, no.4, p. 118-30, October 1998.
- [34] D. Patterson and M. Sturza. U.S. Patent No. 5408237: Earth-Fixed Cell Beam Management for Satellite Communication System, 1995.
- [35] M. Rahnema. U.S. Patent No. 5596722: Packet Routing System and Method for Achieving Uniform Link Usage and Minimizing Link Load, 1997.
- [36] D. Patterson and M. Sturza U.S. Patent No. 5736959: Earth-fixed cell beam management for satellite communication system using dielectric lens-focused scanning beam antennas. http://www.delphion.com/cgi-bin/viewpat.cmd/US05736959__?OUT_FORMAT=pdf&MODE=fstv.
- [37] R. Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16:87-90, 1958.



Arjan Durresi, assistant professor in the Computer Science Department at Louisiana State University conducts research in QoS over IP networks, wireless and satellite networks, network security, congestion control, traffic management and optical networking. He has 15 years of experience in industry and academic research. He has a BSEE, MS and Ph.D. from Polytechnic University of Tirana and a Superior Specialization certificate in telecommunications from La Sapienza University in Rome, Italy and Italian Telecommunications Institute.

Durga Dash has a B.E. Electronics and Communications Engineering and a M.S. from the Ohio State University

Betty Lise Anderson, associate professor the Department of Electrical Engineering at The Ohio State University, conducts research in photonics and optical engineering. Her particular areas of interest are optical communications, microwave photonics, coherence, and optoelectronic device physics. She obtained her BSEE from Syracuse University in 1978, and her MS/PhD in Materials Science/Electrical Engineering from the University of Vermont in 1988 and 1990, respectively, after working in industry for 9 years.



Rajgopal Kannan obtained his B.Tech in Computer Science and Engineering from IIT-Bombay in 1991 and the Ph.D in Computer Science from the University of Denver in 1996. He is currently an Assistant Professor in the Computer Science department at Louisiana State University. His areas of interest are in network security, wireless sensor networks, routing and multicasting protocols, distributed systems and game-theory. He has served on the program committee of ICCCN, MobiDe etc. and is a reviewer for NSF and several Journals and Conferences

Sastri Kota is a technical consultant with Harris Corporation in areas of network architectures and protocols. Mr. Kota has several years of experience in military and commercial network systems design, analysis and held technical and management positions at Lockheed Martin, SRI International, Ford Aerospace, XEROX, The MITRE and Computer Sciences Corp. Currently, he is the US Chair for the ITU-R Working Party 4B and Rapporteur for the ITU-R WP 4B FSS Ka-band performance. He was the Chair for the Wireless ATM working group of the ATM Forum.



Raj Jain is an adjunct professor at the Computer and Information Science Department of the Ohio State University. He is a fellow of IEEE and ACM and serves on the editorial boards of Computer Networks, Computer Communications and the Journal of High Speed Networks. He is the author of two popular books: FDDI Handbook: High Speed Networking using Fiber and Other Media" published by Addison-Wesley and "The Art of Computer Systems Performance Analysis" published by Wiley.