A Lightweight, Link-layer, Source-based Routing Protocol for LEO Satellite Networks

H. Y. Chen , Jim Hutchins, and Bruce Mah
Sandia National Laboratories, MS 9011, P.O. Box 969, Livermore, CA 94550
Hycsw@ca.sandia.gov

Abstract:

This paper addresses the networking issues of an ongoing project that evaluates the feasibility of a system consisting of a large constellation of low-earth-orbit (LEO) Laser Communication Satellites (LCS). Due to the strict payload constraints imposed by the LCS technology, we propose a lightweight, link-layer, source-based routing protocol that will extend the LEO infrastructure to the fast growing terrestrial IP and ATM networks. This protocol relies on the ground-stations to calculate source routes based on a position-dependent addressing scheme and a fixed constellation topology, thereby allowing satellites to simply forward fixed-size packets. We design this protocol to be topology agile, fault tolerant, and network protocol independent. In addition, it will deliver packets in order and provide load balancing.

1. Introduction

The Low Earth Orbit (LEO) satellite technology has gained much interest in recent years because of its ability to provide lower latency and cheaper networks than the Geo-synchronous alternative. Several LEO constellations [1] are currently being deployed, and will be in full production by year 2003. These networks employ large LEO satellites weighing between 450 and 2,207 kg, and most will only provide low- to mid-bandwidth voice services. In contrast, this project will evaluate the feasibility of a network consisting of a large constellation of small (10 to 100 kg) Laser Communication Satellites (LCS) [2] to deliver broadband services at multi-gigabit speeds. We propose to implement the LEO infrastructure as a ”wire in space” using a stateless, lightweight protocol, so that the satellite’s communications subsystem can meet the strict LCS payload constraint. This protocol is topology-agile, provides load balancing, and supports both the Internet Protocol (IP) [3] and Asynchronous Transfer Mode (ATM) [4] networks. As a link layer protocol, it can leverage on the standardization efforts at the Internet Engineering Task Force (IETF) and the ATM Forum to provide broadband services such as voice, video, and data.

We organize this paper as follows. Section 2 provides the background information for our analysis and design. Section 3 presents an architectural overview of our link-layer protocol and its functional components. Section 4 demonstrates applications of this protocol to extend the terrestrial IP and ATM networks. Section 5 summarizes the current work and discusses our future plans.

2. Background Information

Over the past decade, the Internet has burgeoned into a worldwide information highway consisting of approximately 5 million hosts on over 45,000 interconnected networks. This unprecedented growth, together with the introduction of multimedia workstations, has spurred the development of innovative applications that require high speed, low latency, and real-time transport. Today's Internet can neither scale in its bandwidth nor guarantee the Quality of Services (QoS) necessary to meet these performance requirements. The current trend is to use the ATM technology as the underlying infrastructure for the next generation of enterprise and global IP networks. Aside from providing a transparent interface to the best-effort service, these IP-over-ATM proposals [5] also offer extensions to include future Differentiated Internet Services [6], using real-time protocols [7, 8] that are being defined at the IETF. These new protocols will enhance the current Internet by allowing network resources to be allocated and guaranteed to real-time IP flows. Following this trend, we will evaluate an ATM solution for the LCS LEO constellation to complement the terrestrial infrastructure. In addition to IP, it will also provide native-mode ATM services that are defined by the ATM forum. Since the LCS power subsystem will consume a large percentage of its total weight budget, the communication subsystem's power requirement will be an important evaluation criterion beside performance considerations.

2.1. The ATM based LEO Constellation

In order to support ATM onboard a LCS satellite, the communication subsystem must implement ATM's User-to-Network Interface (UNI) [9] and Private Network-to-Network Interface (PNNI) protocol [10], where UNI is a signaling protocol for connecting end-users, and PNNI is a routing as well as a signaling protocol for connecting ATM switches. The PNNI routing protocol extends IP's link state protocol [11] to add resource information such as the maximum and available link-bandwidth, the switch buffer resources, guaranteed delay and delay variations, etc. The PNNI signaling protocol completes a connection request by generating a source-route using the PNNI routing and resource information, if it is the originating switch, or by executing a Call Admission Control (CAC) algorithm to allocate or deny the requested resources, if it is a transit switch. Connections have hard states, which require memory for storage. Because these connection states must be maintained until explicitly released by end users, the procedure for handoff between satellites must provide reliable transfer in order to prevent disruptions to connectivity. A reliable handoff procedure requires yet another protocol, which in turn will demand additional compute-cycles, especially since satellite's transitions occur fairly frequently. Based on the above analysis, we predict that an ATM solution would demand a significant amount of compute and memory resources, thereby requiring a large power and weight budget.

LCS's have a projected average lifetime of 5 years, which represents a mean-time-between-failure of only a few weeks in large constellations. A failed satellite will create a communication hole in the topology, thereby denying its ground-stations access to the infrastructure. Furthermore, because the state of an ATM connection is distributed among all switches in its path, a failed satellite can extend its denial of service to all connections having this satellite in their path, regardless of where these connections were originated. This situation is further exacerbated because this hole will orbit around its plane, disrupting yet another set of connections every transition period, which can be as often as 10 to 30 minutes. This rippling effect can be detrimental to the network operation and will continue until the faulty satellite is replaced. Although this phenomenon can be prevented by over-provisioning of overhead satellites coupled with a complex fault-handling procedure, this solution will be costly, and the additional processing overhead will further increase the satellite's total weight consumption.

2.2. IP Based LEO Constellation

As demonstrated in the previous section, mobility in the underlying infrastructure can complicate solutions that are connection-oriented. Therefore, we choose to evaluate an alternative approach that employs connectionless IP for the LEO infrastructure. If satellites in the LEO constellation were to support IP routing explicitly, then their communication subsystem would periodically execute IP's complicated routing algorithms, as well as maintain the resulting routing databases. Routers calculate IP routes using routing protocols such as Routing Information Protocol (RIP) [12] and Open Shortest Path First (OSPF) [11]. RIP is a distance vector protocol that uses the Bellman-Ford algorithm to calculate the best paths to its destinations. OSPF, on the other hand, uses a link state protocol to achieve faster network convergence and to reduce bandwidth overhead, but is more compute and memory intensive. It calculates its shortest path using the Dijkstra algorithm. In addition to unicast routes, routers must also run multicast routing protocols to support the emerging collaborative, multimedia applications. Multicast routers typically support the Internet Group Management Protocol (IGMP) [13], and one of the following multicast routing protocols: the Distance Vector Multicast Routing Protocol (DVMRP) [14], the Protocol Independent Multicast - Sparse Mode (PIM-SM) [15] protocol, or the Core Based Tree (CBT) [16] protocol. As expected, these unicast and multicast routing protocols will produce routing information that will require a proper handoff process in order to minimize any interruptions during satellite transition. Thus, we anticipate that an IP solution, though simpler than ATM, will also demand a fair amount of compute and memory resources, and therefore, would likely exceed the small weight budget allowed by the LCS technology as well.

3. The ‘‘ Wire in Space’’ Approach

Previous analyses demonstrate that onboard IP and ATM support can be compute and memory intensive. Because increased processing and memory elements demand additional power, larger batteries will be necessary, which, in turn, will increase the satellite's total weight budget. In order to build smaller, cheaper, and faster communication satellites, we propose to implement the LEO infrastructure as a "wire in space," such that its sole function is to forward packets. As such, the communication satellites do not have to compute complicated routing and signaling protocols, which will also eliminate the need to keep state information such as routing entries and connection tables. Furthermore, we chose to forward fixed-size, 64-byte, packets, simplifying both the VLSI implementation and the memory management of the onboard packet switching equipment. Moreover, a 64-byte packet-size can simultaneously accommodate IP and ATM packets, by encapsulating 53 bytes of payload within an 11-byte protocol header. In the IP case, segmentation and reassembly (SAR) will be performed at the ground-station to satellite interface, using technology similar to ATM to accommodate variable length packets that are typically larger then 53 bytes.

The following sections present the functional components of the link-layer protocol. They are designed with the following assumptions. 1) The constellation consists of near polar circular orbits. 2) The LCS payload constraints will impose an upper limit of 4 high-speed, inter-satellite, laser cross-links. 3) Each satellite can multiplex communications to multiple ground units using either RF or atmospheric laser technologies. At the end of this section, we hope to have demonstrated that, using this approach, the satellite's packet forwarding operations are few and simple, and can easily be implemented in hardware.

3.1. The Position Dependent Addressing Scheme

A LEO satellite completes its orbit in approximately 100 to 300 minutes depending on its altitude; therefore, a ground-station will have a different satellite in view periodically (e.g., approximately every 10 to 30 minutes if there were 10 satellites per orbit). If hardware addresses were assigned to moving satellites, all ground-stations would have to periodically update their hardware-to-network address translation table and, in the case of ATM, the connection database, which can be very costly in large networks. Therefore, we propose to build a logical topology that assigns unique addresses to regions of given longitude and latitude. We demonstrate this addressing scheme in Figure 1 using a 4 x 8

LEO constellation. Using this configuration, the ascending polar orbits, east of the Prime Meridian, become the descending orbits after crossing the North Pole, thereby providing global coverage with only half of the orbits required normally. However, this configuration creates a seam at the Prime Meridian; the two orbits on either side of the Prime Meridian circulate in opposite directions, and therefore, preclude the implementation of inter-orbit links between them. We assume that LEO satellites are traveling with a constant speed in stable orbits, and consequently, will be able to track their own positions and assume a new address at appropriate times (see Figure 2). Similar to the technique used by the Global Positioning System, our satellites should be able to resynchronize their clock with a control ground-station to correct inevitable drifts during the course of their operation.

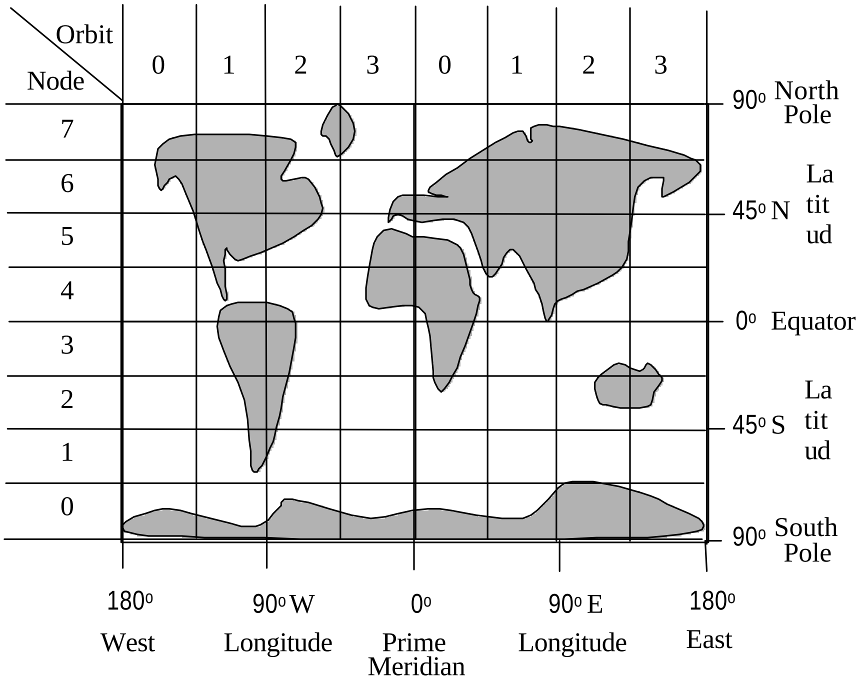

Figure 1. Address assignments via terrestrial coordinates

Each address consists of three components: an orbit, a node, and a channel identifier (ID). The orbit-ID identifies the longitude region of a given orbit. Within an orbit, node-IDs reflect the current latitude positions of the orbiting satellites. We propose to assign the first four channel-ID's (channels 0, 1, 2, and 3) to the inter-satellite laser links, and the next two (channels 4 and 5) to trusted ground-stations that perform the network control functions. Channel IDs from ‘6’ and up are used to multiplex the satellite accesses among its ground units. Typically, a ground unit will adopt the orbit-ID and node-ID assigned to its location. It will acquire a unique channel-ID from the overhead satellite through an initialization procedure at system startup time. Ground units will keep their addresses for the duration of their uptime.

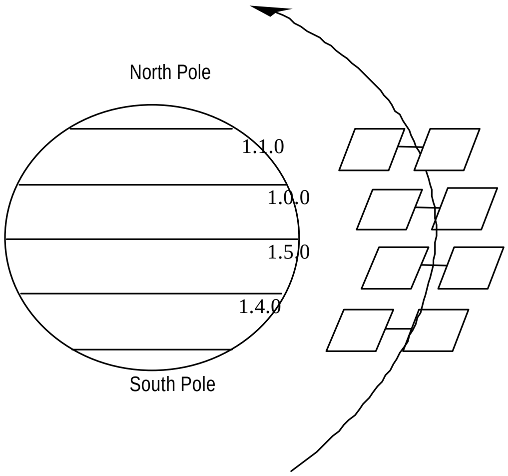

Figure 2. A satellite assumes a new address as it enters the next region.

3.2. The Protocol Header Format

We describe, in this section, the header fields that are used by the protocol components in order to carry out their functions. We will present their detailed applications in a later section when appropriate.

1. Epoch – contains the sequence number representing the periods between satellite transitions.
2. Type – differentiates the unicast, multicast, broadcast, and the various control packets.
3. Priority - prioritizes different traffic classes for preferential treatments.
4. Source - identifies the hardware address of the source ground-station.

5. Destination -
    - is the hardware address of the destination in unicast packet.
    - is a node-mask in broadcast or multicast packet, where each bit-position in the mask represents a corresponding satellite in the orbit.
    - is a wildcard representing the first reachable ground-station.
6. Link mask - reports the inter-satellite link-status using the corresponding bit values.
7. Flow-control - starts or stops to a traffic flow.
8. Index – represents hop-counts in a source-route initially, will also serve as pointers to an active entry as the packet proceeds in its path.
9. Source-route – consists of an 'index' number of 2-bit forwarding directions for corresponding satellites in the path.
10. HEC – contains header error checksum to ensure accurate deliveries of packets.

Because we assume a 64-byte packet size, we have only 11 bytes (64 – 53 = 11) to specify the protocol header, which will impose limitations on the size of node-address as well as source-route, and hence the size of network diameter.

3.3. The Source-based Routing Algorithm

This protocol relies on designated ground-stations, also referred to as control stations, to maintain the topology database and to respond to topology queries from other ground-stations. Typically, control stations will be pre-configured with the topology information, which consists of the number and inclination of circular orbits, the number of satellites per orbit, the number of inter-satellite cross-links per satellite, and the link characteristics such as half or full duplex, the direction of flow, and speed, etc. Control stations keep the topology up to date by incorporating the most recent status reports from all satellites. In order to acquire the topology database, all ground-stations are configured with the hardware and protocol address of one or more control stations, as well as corresponding source-routes to reach them. Based on the acquired LEO topology, ground-stations can proceed to calculate, using the Dijkstra or a similar algorithm, the shortest paths to all other ground-stations.

Recalculation of the shortest paths occurs only when the link-state changes. We propose that every satellite announces its changes immediately, thereby allowing all ground-stations, control and generic, to update the shortest paths in a timely fashion to bypass invalid links. We will cover this fault recovery mechanism in more detail in a later section. Because ground-stations send all packets with a common destination over the same source-based route, they will deliver these packets in sequence, under normal conditions. The ability to delivery packets in sequence is pertinent to the ATM technology, as ATM cells that are delivered out-of-order will preclude their reassembly at the destination.

3.4. The Unicast Operation

To forward a unicast packet, a source station constructs a source-route based on a shortest-path to the packet's destination. A source-route is formatted according to the street sign algorithm [17], which allows the satellites at each hop to steer this packet to travel left, right, north or south accordingly. A source-route consists of an index field followed by a series of 2-bit forwarding directions, where each direction holds the channel-ID of the output cross-link. The index field points to the active entry for the current satellite and will be decremented at each hop as the packet propagates along its source-route. A ‘‘null’’ index indicates that the packet has reached its last-hop satellite, and that the satellite should now deliver the packet to the target ground-station identified by the channel-ID. In order to prevent miss-delivery of packets, the destination satellite should always verify that the packet header's orbit- and node-ID are identical to its own.

3.5. The Broadcast and Multicast Operation

Because our satellites are completely stateless, the protocol header must contain both routing as well as state information so that broadcast-packets can be forwarded to all destinations without looping. Without proper treatment, looped broadcast packets will be forwarded indefinitely in this topology, and therefore, could consume increasing amounts of bandwidth as more and more broadcast packets are introduced into the network. Because, we have very limited space in the protocol header to store this information, we are forced to implement broadcasting by sending one copy of this packet per orbit via unicast. Hardware-based branching occurs only after the packet has reached its target orbit. Under this scheme, a ground-station will construct one source-route to reach a satellite in every orbit and mark all the bits in the node mask to indicate that all satellites in the orbit are to receive this packet. Thus, the node mask is used to provide the routing information to reach target satellites, as well as the state information to prevent looping in the orbit. A satellite first examines its bit position in the node mask to determine if it is the targeted receiver. If so, it will clear this bit and replicate the packet to all of its ground units. It then decides whether this packet should be forwarded further by examining the node mask. If there are still marked bits in the mask, the satellite will forward the packet in a pre-selected direction around the orbit; otherwise, it will drop the packet to prevent looping. When there are failed links in an orbit, a ground-station may have to send multiple copies of a broadcast packet to this orbit, one per fragment. As shown in Figure 3, a ground-station must first locate all satellites that have broken links, opposite to the pre-selected direction. It then sends one copy of this packet to each of these satellites, marking the bit positions of reachable satellites within that fragment, thereby allowing the satellites to proceed with the branching as described earlier.

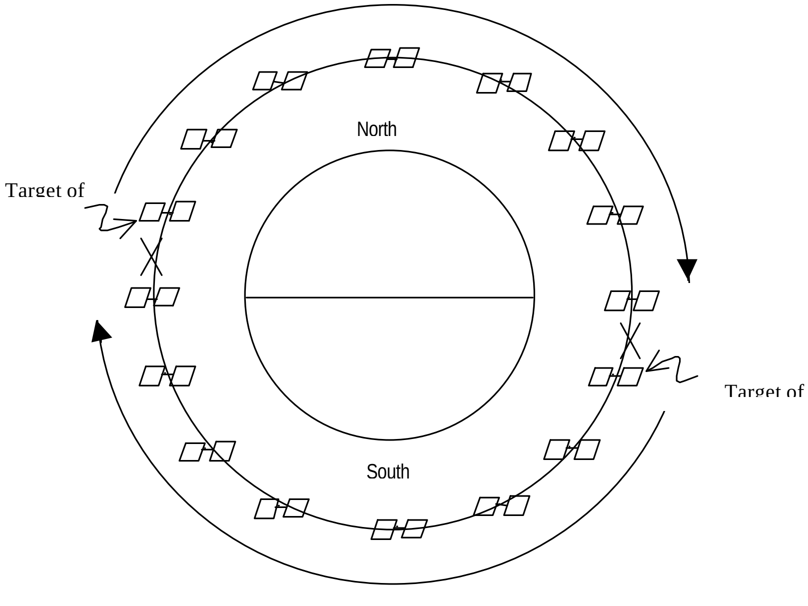

Figure 3. One copy per segment for orbit with broken cross-links

With this approach, multicast is just a subset of the broadcast operation. Multicast packets will only be sent, via unicast, to those orbits that have satellites with group members; likewise, the packet replication will only occur at those satellites that have members reflected in the node mask.

3.6. The Dynamic Fault Detection and Automatic Fail Over

Circular polar orbits are popular because satellites in its orbit can maintain a constant angle with their neighbors. This characteristic allows satellites to use fixed optics in their laser terminals without the need for heavier tracking devices. Without tracking, however, the laser cross-links are more prone to losing the line-of-sight of neighboring satellites and therefore, experiencing higher bit-error rates. Laser Communication Satellites typically have a five-year lifetime. Therefore, a constellation consisting of 120 such satellites will probably experience a 15-day (5 * 365 / 120) mean time between independent failures. In order to deal with the dynamics of the satellite and link status, this protocol requires that each satellite send unsolicited status reports as soon as changes are detected. A failed satellite will be detected when all of its neighbors report a link failure in its direction.

The status reports are intended for all ground-stations and, therefore, should be delivered via a broadcast mechanism, and which, unfortunately, is not directly available to our satellites because of their limited intelligence. To overcome this obstacle, a satellite wraps its status report in a control message and forwards it, using a predefined algorithm, to any ground-station that it can reach. This ground-station is responsible to forward this report to all other ground units as soon as possible, using the broadcast mechanism described earlier. Note that the orbit- and node-ID in the control packet's header identifies the owner of the link report.

3.7. The Simple Handoff Procedure

Since our packet delivery architecture does not maintain routing or connection information in the satellites, they need only to handoff packets that are queued at the instance of transition. At this time, the satellite will hand shake with its successor to enter a transition mode to handoff these packets. Once in this mode, the receiver of the handoff will re-process and -queue these packets for proper deliveries. Meanwhile, it will hold newly arrived packets in a temporary queue until the transition is completed. Unfortunately, this simple handoff procedure will not catch packets that are still in-flight at the moment of transition. Note that we have reserved a 3-bit epoch field in the header to record the epoch number between satellite transitions; the epoch number is maintained by each ground-station and will be incremented after each transition. A receiving ground-station detects an out-of-order delivery by examining the arriving packet's epoch value from a source, if it is less than that of the previous packet, the packet must have been delivered out-of-order and therefore, should be dropped accordingly.

The limited header space limits the size of the epoch field. However, we increment this field rather slowly; for example, a 3-bit epoch value wraps around every 80 to 240 minutes (assuming a 10-30 minute transition period). Therefore, the small size of the epoch field will not be of concern. In fact, we can also use this mechanism to detect out-of-order deliveries caused by dynamically rerouted traffic flows due to changes in link state. In this case, we will also increment the epoch number when a ground station recalculates its shortest paths. This measure can prevent out-of-order packets from consuming precious downstream bandwidth and processing overhead, thereby improving the overall network performance.

3.8. The Flow Control Algorithm

Under heavy load, switches on the satellites need buffer space to prevent packet losses from buffer overflow. Because this protocol delivers 64-byte fragments from larger, original packets, it is likely that fragments from many packets will be interleaved at arrival; therefore congestion losses can effect a larger number of packets, and which will exacerbate the degradation in network performance [18]. Unfortunately, the LCS weight limitation discourages the use of large onboard buffers, we, therefore, propose a simple, hop-by-hop, stop-and-go flow control mechanism (Figure 4) to reduce its memory requirements. We will use the flow-bit in the packet header to flag a stop command when the downstream node detects congestion and a go-command after the condition is relieved. Note that a flow control command can be piggybacked on a data packet, in the reverse direction, if one is pending; otherwise, the flow-control command should be dispatched as soon as possible using a null packet. Because it will take the upstream node at least one round-trip time (RTT) to respond, in order to accommodate in-flight packets, the downstream node needs to keep a minimum of one (bandwidth x RTT) memory beyond its congestion threshold to prevent any packet loss. Simulation studies will be conducted to evaluate the effects of buffer size, scheduling algorithm, and congestion control schemes on network performance.

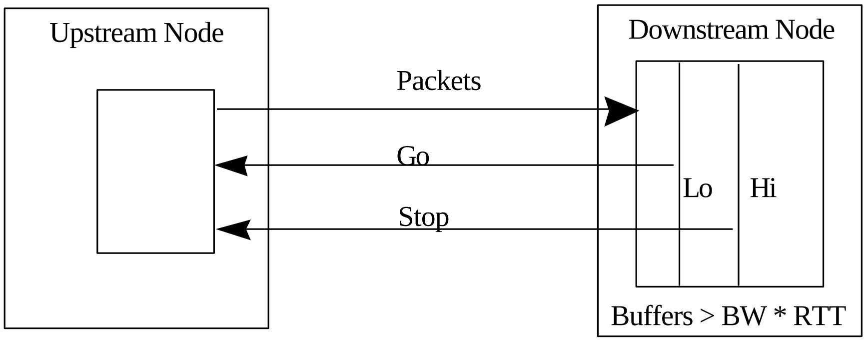

Figure 4. Hop-by-hop, stop-and-go flow control.

4. Example Applications

This section illustrates the applications of our link layer protocol to support IP and ATM. We will draw upon the background information and the protocol specification from earlier sections to support our arguments.

4.1. The Extension to ATM

In order to extend the terrestrial ATM network, the ground-stations of our LEO network will function as ATM switches, supporting the UNI and PNNI protocols. Because our link-layer protocol provides accesses to multiple ground-stations, it needs to provide the unicast, broadcast, and multicast capability to achieve this goal. As described in section 3, our multi-access link-layer protocol has these capabilities. However, the ATM technology only works on point-to-point links. ATM switches on either end of a point-to-point link exchange their identities during the initial handshake, and therefore, can proceed to exchange routing and signaling messages, using reserved ATM virtual circuits, without ambiguity regarding the source. Unfortunately, this condition no longer holds true when the link is connecting multiple ATM switches. We will have to rely on the packet's source hardware address, in addition to the reserved ATM circuit, to differentiate the source as well as the type of this incoming message. As such, the UNI and PNNI protocols can function properly, thereby allowing terrestrial ATM devices to extend their connections over this infrastructure.

LEO infrastructures are often highly meshed, and therefore will provide multiple shortest paths to a common destination. We propose to assign ATM virtual connections to shortest paths of equal weight, in a round-robin fashion, to balance the network load.

4.2. The Extensions to IP

In order to provide interoperability to the existing terrestrial IP network, the LEO infrastructure must have the ability to disseminate routing information generated by RIP or OSPF. As described in the previous section, our protocol has the basic components to meet this demand. However, it forwards fixed-size, 64-byte packets, which is contrary to IP's packet delivery mechanism. IP delivers variable size packets that can range from 64 bytes to a maximum transfer unit (MTU) of the packet's originating link (e.g., 1500 bytes for Ethernet and 4352 bytes for FDDI). Therefore, ground-stations must provide a hardware segmentation and reassembly mechanism similar to that of the ATM technology. We propose to adopt the ATM hardware and its Adaptation Layer-5 (AAL5) protocol [19] for this purpose, but without ATM's complex signaling protocol. We will use the source's hardware address to identify the reassembly queue at the destination ground-station for incoming IP traffic. Similar to the ATM case, we can provide load leveling by assigning IP flows to equally weighted shortest paths in a round-robin fashion. Under this scheme, the receiving ground-stations will need the packet's source-route in addition to the source ground-station's hardware address as the identifier for the reassembly queue.

5. Summary and Future Work

This work defines a packet-delivery architecture for a constellation network consisting of a large number of low-earth-orbit (LEO) Laser Communication Satellites (LCS). Due to the strict payload constraints imposed by the LCS technology, we propose a stateless, lightweight, link-layer, source-based routing protocol to extend the mobile LEO satellite infrastructure to the fast growing terrestrial IP and ATM networks. The architecture will allow smaller, cheaper, and faster communication satellites by limiting the satellite's function to only packet forwarding, and by placing all compute-intensive responsibilities in the ground-stations. Therefore, our constellation acts as a " wire in space" and has the following innovative features. It has a multi-access ATM link, as opposed to the conventional point-to-point paradigm. It provides in-sequence packet delivery, a hardware-assisted broadcast/multicast operation, multi-protocol support (i.e., ATM and IP), and multimedia services (i.e., voice, video, data). In addition, its routing scheme is topology agile, and has built-in fault detection and fail over capabilities. Because the constellation acts as a "wire in space," satellites will be completely stateless and will require the 11-byte protocol header to carry routing as well as state information, thereby limiting the size of the address field and, consequently, the network's diameter.

For the coming year, we plan to build an experimental testbed as a proof of concept. This testbed will consist of prototype satellite communication subsystems, and hardware devices to emulate the mobility of the satellite infrastructure, as well as the bit-error-pattern of the satellite laser and/or RF links. In addition, we will use simulation techniques to investigate potential system limitations in the presence of dynamic network traffic. We hope to fine-tune our protocol design, based on our simulation results, in order to alleviate any performance bottlenecks, especially in the area of congestion and flow control. Thus far, this protocol has addressed the support for fixed ground-stations only. In the coming year, we plan to evaluate the feasibility of adopting a proxy agent to offload the complexity from mobile units, thereby providing mobile users with acceptable performance at a lower per-unit cost.

Acknowledgment

The authors appreciate the technical and editorial input from Phil Bryson and Rich Palmer. Special thanks are due to Phil for coining the “wire in space” phrase, which has proven to be a powerful way to communicate the architectural concept. We are also grateful to Jeffrey Kern, Hilary Jones, Byron Dean, Lyndon Pierson and Lee Schoeneman for the information related to the satellite communication technologies.

References

[1] L. Wood, “Lloyd’s Satellite Constellations,” http://www.ee.surrey.ac.uk/Personal/constellations/, September 1998.

[2] Sandia National Laboratories’ internal communication.

[3] ISI, USC, “DOD Standard Internet Protocol,” RFC 760, January 1980.

[4] M. de Prycker, “Asynchronous Transfer Mode-solution for broadband ISDN,” Ellis Horwood Limited, ISBN 0-13-053513-3, 1991.

[5] H. Chen, R. Tsang, J. Brandt, and J. Hutchins, “A Survey of IP over ATM Architectures,” ftp://ca.sandia.gov/pub/papers/IP-over-ATM.ps, January 1997

[6] M. Carlson, W. Weiss, S. Blake, Z. Wang, and E. Davies, “ An Architecture for Differentiated Services,” http://ietf.org/internet-drafts/draft-ietf-diffserv-arch-01.txt, August 1998.

[7] M. Schulzrinne, S. Casner, R. Frederick, V. Jacobson, “RTP, A Transport Protocol for Real-Time Applications,” January 1996.

[8] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, “Resource ReSerVation Protocol (RSVP), Version 1 Functional Specification,” Internet Draft, November 1996.

[9] The ATM Technical Committee, “ATM User Network Interface (UNI) Specification Version 3.1,” Prentice Hall, 1995, pp151-294.

[10] The ATM Technical Committee, “Private Network-Network Interface Specification Version 1.0,” af-pnni-0055.000, March 1996.

[11] J. Moy, “The OSPF Specification,” RFC 1131, January 1993.

[12] C. Hedrick, “Routing Information Protocol,” RFC 1058, June 1988.

[13] W. Fenner, “Internet Group Management Protocol, Version 2”, Internet Draft, October 1996.

[14] T. Pusateri, "Distance Vector Multicast Routing Protocol," Internet Draft, September 1996.

[15] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification," Internet Draft, September 1996.

[16] Street sign A. Ballardie, S. Reeve, N. Jain, "Core Based Trees (CBT) Multicast– Protocol Specification," Internet Draft, September 1996.

[17] N. F. Maxemchuck, "Routing in the Manhattan street network," IEEE Transactions on Communications, vol. 35 no. 5, May 1987, pp. 503-512

[18] C. Fang, H. Chen, and J. Hutchins, "A Simulation Study of TCP Performance in ATM Networks," Proceedings of IEEE GLOBECOM'94, November 1994, pp.1217-1223, ftp://ca.sandia.gov/pub/papers/Globecom94.ps.

[19] G. C. Sackett, C. Y. Metz, "ATM and Multiprotocol Networking," McGraw-Hill, 1996, pp. 175-177.