Internet Routing over LEO Satellite Constellations

Paolo Narvaez
pnarvaez@mit.edu

Antoine Clerget
Antoine.Clerget@sophia.inria.fr

Walid Dabbous
Walid.Dabbous@sophia.inria.fr

October 9, 1998

Abstract

Satellite constellations offer interesting perspectives for Internet traffic routing, since they benefit from a worldwide coverage and do not have the main drawback of geosynchronous satellites : propagation delay. However, because of their low altitude, LEO satellites are in constant motion and exhibit specific issues for routing. This article presents a general architecture applicable to the integration of satellite constellations in the Internet.

1 Introduction

Satellite links have long been used to transmit Internet traffic. Today, the conventional communications satellites that convey this traffic are located in geostationary orbits (GEO), maintaining a constant position relative to the earth. These GEO satellites are used as bent pipes (i.e. act as mirrors) and establish a fixed communication link between two ground stations. These links consist of only one satellite hop and are completely static and transparent. No routing takes place at the satellite level.

The main drawback of using GEO satellites for communications is the long propagation delay due their high altitude orbit. This leads both to low interactivity, which is a problem for real-time applications such as telephony, and bad performance for some transport protocols such as TCP, which needs a few round trip times to reach the optimal transmission rate ([2], [3], [4]). Moreover, a ground transmitter has to have a high gain antenna to reach the GEO satellite. In order to solve these problems, there is a strong interest in developing constellations of low earth orbit (LEO) satellites.

Because of their low altitude, LEO satellites are in constant motion relative to the earth. To cover all day specific areas on the earth, it is necessary to have a satellite constellation. A LEO satellite constellation consists of multiple rings of satellites spread as evenly as possible over the earth's surface. Because of their mobility, routes between two earth points need to be continuously modified to maintain connectivity. Furthermore, because of the proximity of these satellites to the earth, many connections will have to traverse more than one satellite in the constellation. This raises new problems relative to the establishment, routing and maintenance of connections between two points on the earth surface, as well as their integration with ground networks.

In this paper, we present a general architecture to integrate a LEO constellation with the Internet. We consider how to transpose the various aspects of IP routing to the satellite environment. We discuss how to propagate IP routes through the constellation as well as how to optimize the IP forwarding decisions. Our proposed architecture allows the constellation to interoperate transparently with the ground IP networks, while using internal routing mechanisms that exploit the advantages of satellite routing.

Our approach falls within the IP routing context and does not assume any specific underlying technology for the satellite network. We only take into consideration IP traffic and look at what mechanisms are needed to route this type of traffic over the special characteristics of LEO constellations.

In section 2, we start by describing what aspects of satellite routing are different from terrestrial routing. Section 3 describes some work related to satellite routing. Section 4 proposes a general IP routing architecture for the satellite constellations. Section 5 introduces the various metrics used to perform adequate routing. Section 6 and 7 describe internal and external routing mechanisms. Section 8 concludes the paper.

2 Satellite Routing

2.1 Mobility

Most terrestrial routing protocols, such as OSPF and RIP, rely heavily on exchanging information between routers in order to distributively gather information on the topology of the network. But due to the high mobility of low earth orbit satellites, and therefore to the constantly changing topology of the constellation, these informations become very quickly irrelevant. Therefore, individual satellites cannot act as conventionnal routers : to integrate a satellite constellation into the Internet, one cannot just use the IP routing protocols on the satellite segments. Discovering and establishing routes over a satellite constellation is inherently a different problem than

routing over terrestrial links.

Eventhough the topology of the constellation is highly dynamic, it is predictable. The orbit trajectories can be calculated offline and can be used throughout the constellation period. Moreover, unlike terrestrial links, where new links or destinations can be added on a day to day basis, a satellite constellation is fixed in size. There are only a few unexpected changes (mainly failures) in the topology of the constellation. Therefore, unlike terrestrial networks, there is not much need to exchange topology information.

Since the constellation topology is highly predictable and topology feedback is not necessary, the satellite network can dedicate more resources exchanging information on the traffic load. Furthermore, because the network is highly symmetric and uniform, it is easier for a satellite to perform routing functions. Routing algorithms can exploit the regularity of the topology to optimize forwarding decisions.

2.2 Worldwide coverage

For routing purposes, today's Internet is divided into many different autonomous systems (AS). Each AS contains a set of routers that can exchange packets amongst themselves using some consistent routing protocol. When packets travel between autonomous systems, they have to cross a pair of border gateways.

In a typical AS, there is only one border gateway for any given destination : all traffic crossing AS X to reach network Y will be routed through the same gateway G regardless of changing costs that may take place within AS X. This is because even if more than one exterior gateway is advertised, internal routing protocols such as OSPF will consider the outside metric more important than the inner one.

In terrestrial networks, the external cost metric is always considered of a higher value than the internal metric because it represents the cost of sending a packet through another AS over a large geographic distance, whereas the internal metric measures the cost of the transmission over a small local network.

In satellite constellations, the length of the internal path can easily be as large as that of the external one. Therefore, the routing protocol should consider the internal cost metric to be as important as the external one.

As a result, in the satellite context, it might be reasonable to consider more than one gateway for a given destination. In other words, depending on the satellite that has to forward the packet, the border gateway may be different. Using more than one exterior gateway increases the choice of routes and therefore improves the optimality of the overall routing.

3 Previous Work

The directionality of the satellite beams as well as the multiple hops that are required to connect two ground stations require the LEO constellation to perform extensive routing functions within the satellite constellation. Recent work [1, 5, 6, 7] in the literature has concentrated on how to establish an ATM connection from one ground point to another through LEO constellations.

The work in [7] proposes establishing a separate virtual path connection (VPC) for every pair of source-destination satellites. Before entering the space segment, all the ATM virtual connections (VCs) sharing the same pair of satellite entry and exit points are aggregated into the same VPC. Within the satellite constellation, all switching is performed only according to the VP label (not VC label). As the topology of the constellation varies because of the satellite motions, the VPs are reconfigured internally. Therefore, all the internal routing done at the satellite level is transparent to the different end-to-end VCs.

The work in [6] introduces some new ideas to ensure quality of service (QoS) guarantees within the satellite constellation. The main idea is to define a network of virtual nodes that is independent of the actual physical network. The virtual nodes are geographically static and are embodied by whichever satellite at the time is nearest to the appropriate geographical position. At the ATM layer all routing decisions and bandwidth allocations are performed using the fixed topology of the virtual nodes. The actual path taken by the ATM cells is transparent to the ATM layer. Therefore, QoS parameters can be enforced through the virtual node topology, eventhough the underlying satellite layer is connectionless and time-varying.

However, none of this work addresses the question of how to integrate the constellation or its ATM network into the Internet. This paper discusses on a possible architecture for the integration of a satellite constellation network in the Internet, whatever the underlying technology used for internal routing.

4 Architecture

4.1 System Integration

In section 2, it has been shown that satellite constellations have different routing requirements than terrestrial links. Both for practical and financial reasons, and because of the satellite network's homogeneity and dynamicity, doing IP routing in every satellite is probably not the best solution. The routing integration we propose here should be independent from the underlying technology of the satellite constellation. Internally, the satellite constellation should use its own routing protocols. However, externally, the constellation needs to interact transparently with terrestrial networks and the rest of the Internet as if it were just another IP network.

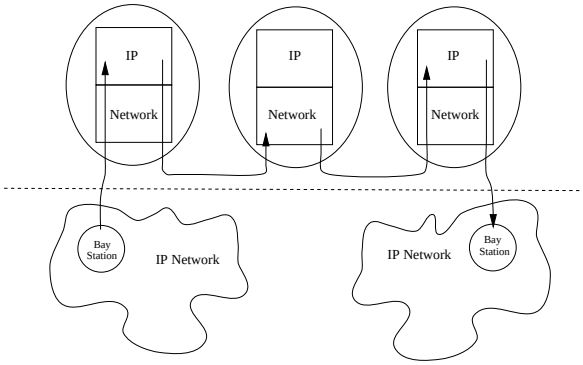

Figure 1: The IP layer is only visited at when entering and exiting the constellation. Within the constellation, the network layer routes the packet directly to the exit point.

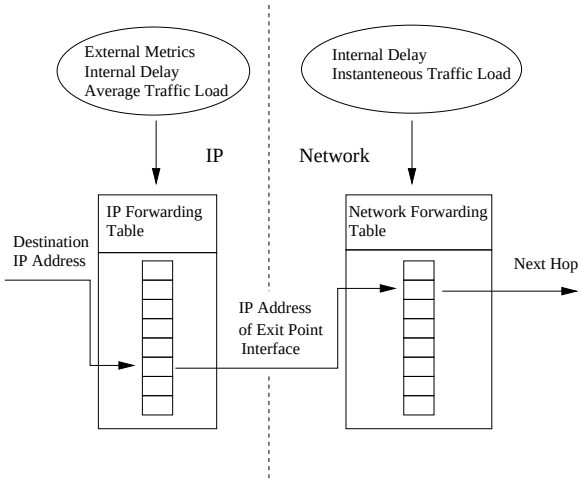

Figure 2: Two Level Routing Table

We propose to configure the entire satellite constellation into a single autonomous system (AS). Any ground station that is connected to the constellation will see the same autonomous system. It will always connect to the same constellation IP address, regardless of what satellite is actually above. Likewise, the same IP address can be used at any time to reach a given ground station. In other words, each IP address identifies an entry/exit interface to/from the constellation. The border IP-gateways are the exit and entry points of the constellation with a given IP address, and are therefore virtual entities at fixed geographical positions.

When an IP packet enters the constellation, the border gateway determines the packet's exit point and uses a proprietary method to route that packet to its exit point within the constellation. Routing is therefore performed in two different layers : The IP layer and the network layer. The IP layer determines the IP exit interface from the constellation, which is the next hop on the IP routing path. The network layer decides how the packet arrives at that interface.

Figure 1 illustrates how a packet travels through the satellite constellation. Figure 2 shows the two level routing table architecture.

4.2 Network address

When a packet arrives in the constellation, the border IP-router looks in its IP routing table to determine the IP address of the exit interface. It is then converted to a network address, thanks to a Constellation Address Resolution Protocol, and the proprietary internal routing protocol is then in charge of carrying this packet through the constellation. To have IP connectivity with the constellation, the ground terminal must register to its border gateway, which updates a network address / IP address conversion table. This can be done using mechanisms similar to those in ATMARP [8]. The entry is removed from the table when the terminal disconnects.

The network address can be a :

1. Terminal or Connection ID

   This will in particular be the case for connection oriented internal routing protocols. If the constellation dynamically allocates IP addresses, and if it is possible to map these on network addresses, the Address Resolution Protocol becomes very simple, and we can avoid using conversion tables.

2. Geographical position and Medium Address (Code, Time Slot, and/or Frequency)

   The geographical position can be represented by the latitude and longitude of the ground terminal (determined for example with a Global Positionning System - GPS). We can also use fixed cell numbers. What should the granularity of the geographic addressing be, i.e. how should we choose the size of the cells ? At all times, there has to be a satellite that can forward our packet to the ground and that therefore covers the entire cell. When a cell is leaving the coverage of a satellite A, another satellite B must be here to cover the entire cell. This implies, as shown on figure 3 that the cell must be smaller that the overlap of two neighbor satellites' footprints.

3. Satellite and Beam ID or Moving Cell ID, Medium Address (Code, Time Slot, and/or Frequency)

   For a given fixed host on the ground, these addresses will change over time. To avoid constantly updating our IP/network address conversion tables, it would be preferable to use the knowledge the gateway has of the satellite constellation's topology at that time to convert temporary addresses to fixed addresses. We can then store these fixed addresses in our conversion table, and have this type of addressing behave as the previous one.

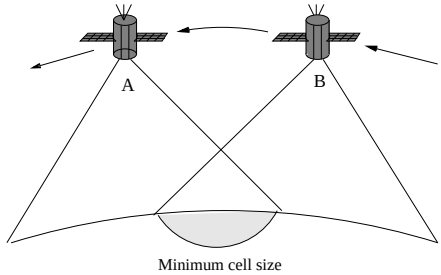

Figure 3: Minimum cell size

4.3 Border gateways

The integration of the satellite constellation is performed transparently for the ground at the IP layer. The virtual IP border-gateways and the ground hosts run some compatible exterior routing protocol such as EGP or BGP. Through these connections, the ground and the satellites both import and export inter-AS routes. The gateways establish internal BGP connections to share these inter-AS routes. They should also exchange their network address.

In section 4.1, we explained that virtual entities at fixed geographical positions should play the role of the border gateways. These border gateways are responsible for forwarding the IP packets originating from a region. Here are three ways of implementing this :

- The border gateway, (i.e. the instance of the BGP/EGP protocol with the IP routing table and the IP / network address conversion table) is implemented in a satellite covering the region.
- All the border gateway's tables are stored and updated in every satellite. The gateway process is activated in the satellite when it covers the region.
- The border gateway is implemented in a dedicated ground station.

With the first implementation, the IP routing tables in the satellites have to be transfered from satellite to satellite as they move along their orbit. This method requires a large amount of computation performed periodically. The second method involves a lot of memory in the satellites. With the third method, putting the border gateways on the ground enables us to construct a system with minimal memory and computationnal requirements since we can move the IP tables away from the satellite. Moreover, this method has the advantage of allowing terrestrial exchanges between the border routers on the ground. In other words, the EGP/BGP/... connections can be set on terrestrial links. However we need to put in place terrestrial infrastructures, and packets must be bounced back to earth for IP routing, but this may be necessary if we want to have light systems in the satellite.

Under this third framework, we have to place stations on the ground. For all ground hosts to be able to access their border gateway directly with a unique satellite, we need to have a ground facility per cell (section 4.2). But installing a ground facility per cell may not be always feasible. For example, for the Teledesic constellation, this would mean installing more than a thousand ground stations[1]. However we can lower the number of ground facilities if we accept additionnal traffic on the constellation. The only important thing is for every host to know the network address of a border gateway - if possible the closest. This third implementation is therefore the most realistic, and will be the one used in this paper.

4.4 Path of an IP packet

Once the ground host has IP-connectivity with the LEO-constellation, a BGP (or some other exterior routing protocol) connection is established and the accessible networks are advertised as for any other AS.

To route an IP packet over the constellation, it is first sent to the border gateway as a local network packet[2]. The border gateway decapsulates the IP packet, determines the next hop (IP router) which is the constellation exit gateway using the external routing protocol[3], and forwards the packet with the internal routing protocol. This packet reaches the exit border gateway, which determines the IP address of the ground AS border gateway. It is internally routed, and broadcasted over the exit cell by the LEO satellite. Only the ground receiver whose MAC address matches will forward the packet on the rest of the Internet.

The problem with this mechanism is that a given packet must be bounced back twice to earth before exiting the constellation, which is costly in terms of bandwidth - since this happens for every packet - and delay (eventhough this stays quite reasonnable for Low Earth Orbit satellites).

Some of these bounces can however be avoided :

1. Large Users

   It may be interesting to integrate an important ground router connected to the LEO constellation, i.e. the access router of a large ISP, in the constellation's autonomous system. It then establishes internal BGP connections with the other border gateways and share the informations concerning external route advertisments. It acts as a border gateway, exchanging as well its network address. By choosing to integrate these *"Large Users" (LU)* in the constellation's autonomous system, we accept to carry all external

[1] With 24 sats/plane at $H = 1400$ km, min. elevation angle $\alpha = 40°$, the footprint radius is $R.(\pi/2 - \alpha - asin(\frac{R}{R+H}.cos(\alpha)) = 1233\ km$, the cell radius $C < C_{max} = F/2 - \pi.R/24 = 400\ km$. Min number of cells $N > 4.R^2/C_{max}^2 = 1000$

[2] Using the internal routing if not directly reachable by the satellite

[3] section 7

routes to and from the Large User through the constellation, generating extra traffic. However, the LU will be able, as a border gateway, to reach directly the constellation exit point, avoiding the first bounce. If the exit point is also a LU, we can avoid both bounces.

2. Medium Users

   The ground routers connected to the constellation, especially those who are geographically close to a border gateway (i.e. in a same satellite beam), that have enough computing power, may want to listen to the BGP internal traffic and network address exchanges. They will then be able to forward directly IP packets to Large Users, and avoid the first bounce.

3. IP / network address conversion

   If the network address can easily be computed from the IP address (4.2), the border gateways can specify the 'next hop' when advertising external routes on the internal BGP/EGP connections. Doing so, the entry border gateways will be able to avoid bouncing the IP packet to the exit border gateway.

4.5 Ground terminal's mobility

The ground terminal's mobility is a problem for internal routing. The network layer must therefore integrate mechanisms to deal with mobility. In that case, if the network address is a terminal or connection ID, mobility will be completely transparent to the IP layer. Otherwise, the mobile terminal will be assigned a new network address when it leaves its cell, which, except for very high speed moving terminals, doesn't happen very often. A simple way to deal with this mobility at the IP level is then to disconnect and reconnect to the constellation, to update the IP/network address conversion tables. Eventhough it may be suboptimal, continuing to access the constellation through the same border gateway is still possible, and avoids updating all the other border gateways routing tables.

5 Routing Criteria

In order to propose routing systems for satellite constellations, we need to consider different measures for selecting routes. We here consider the following metrics : external metric, internal delay, expected traffic load, and instantaneous traffic load. Each of these measures is acquired in a different way. In this section, we will briefly discuss each of these metrics separately.

5.1 External Metric

The external metric is a cost associated with every IP address (or aggregation of addresses) outside of the satellite constellation. The external cost can be obtained through the border routers running some exterior routing protocol (i.e., BGP). These routers periodically advertise their connectivity to groups of IP addresses together with their cost.

The external metric can change every time the exterior routing protocol sends routing updates. However, in practice, external metric changes occur much more infrequently than changes in local routing information. Nevertheless, this information is dynamic, and internal feedback within the constellation is needed to update this cost.

The external cost is needed to calculate the best exit point for the satellite constellation. It essentially measures how desirable a border router is for the purpose of forwarding a packet to a given destination.

5.2 Internal Delay

The internal delay is the propagation delay of the path between the entry and exit points of the satellite constellation. It is proportional to the physical distance between these two points.

Because of the satellite movements, this delay is highly time-varying. However, because of the periodic motion, all the delays of the constellation are periodic. They can be calculated offline and then used when the satellite is in the appropriate position. Therefore, even though the delays are constantly changing, no calculation needs to be performed on-line since everything is known ahead of time.

The satellite constellation can use this metric to minimize the distance and propagation time for every packet traversing the satellite. It measures how easy (in distance) it is to send a packet to some particular border router. Along with the external metric, the internal delay is used to find out which border router will lead to the shortest path to a given destination.

5.3 Expected Traffic Load

The expected traffic load is an a priori estimate of the amount of traffic that is traversing some satellite path within the constellation. The two major factors that can influence the expected load are the geographic position and the time/date.

Again, because of the constellation's periodicity, the expected load can be calculated offline and used throughout the period. The expected load can also be modified adaptively to capture variations in traffic patterns over long time periods.

The expected traffic load can be used to find the routes within the constellation that are on average least loaded. This metric is necessary to predict what the traffic conditions will be on some other part of the constellation before packets are forwarded there.

5.4 Instantaneous Traffic Load

The instantaneous traffic is the amount of data that is traversing a satellite link at some particular time. Such data might

be obtained by measuring the queues of the various satellites or the utilization of the various links.

Since the instantaneous load is constantly changing, continuous feedback is needed to update this information. However, because of the large propagation delay and the fast rate of change of the traffic load, it is impractical to propagate this information beyond a local neighborhood. Therefore, the instantaneous traffic load can only be used to optimize local routing. It can be used by neighboring routers to divert local traffic going into temporarily loaded links.

6 Internal Routing

Given a network destination address, the network layer in a satellite must decide how to forward the packet. For this purpose, it can use the periodic information about the constellation topology as well as the dynamic feedback on the instantaneous traffic load.

Because of the highly symmetric and uniform structure of the satellite constellation, routing between the entry and exit point of the satellite constellation can probably be done efficiently. Since the total number of exit points (cells) is limited, it is unnecessary to establish virtual circuits for intra-constellation routing ; next hop routing seems to be simple enough.

Since each satellite has perfect knowledge of the constellation's topology, the network layer can decide which next hop is best for every destination. Each satellite independently decides which is the best next-hop (next satellite) to forward the packet. The network layer can also use local information on instantaneous traffic load to modify the next hop if it is too congested.

An easy routing scheme that could be implemented is some modified version of the hot potato scheme. Under such scheme, a packet is forwarded to the neighboring satellite that is closest to the exit point. If the link to that satellite is congested beyond a threshold, the packet is forwarded to the next closest satellite, and so on. Once again, because of periodicity, the list of best next hops (for the distance metric) can be pre-computed and stored in a table. An example of a network routing algorithm could then simply be to choose the first next hop closer to our destination that is sufficiently un-congested.

6.1 Quality of Service Routing

Traditionally the Internet has been a network supporting best effort traffic. Currently, most IP networks cannot offer any QoS guarantees to IP traffic. However, it is expected that schemes enabling QoS routing will become widely available in the future Internet.

In order to accommodate QoS traffic within the constellation, the network layer needs to be modified. The network layer routing must be made more rigid and the paths more static.

The satellites in the path between two constellation IP addresses are constantly changing, Therefore, the flow state stored on a particular satellite is of little use some time later when the satellite is in a new position, serving other paths. In order to be able to store useful flow information in the constellation, we use the concept of virtual node introduced in [6].

In this model, there are an equal number of virtual nodes as satellites in the constellation. However, the virtual nodes are static and cover a fixed geographic position. A particular virtual node is embodied by whichever physical satellite is in its corresponding area.

The virtual node contains all the flow state information necessary for performing QoS routing. When a physical satellite leaves a virtual node to another satellite, it transmits to the new satellite all the relevant flow information.

At the network level, routing is performed through the virtual nodes. The choice of paths is restricted so that a flow between two constellation IP addresses always follows the same virtual nodes (or a limited number virtual node paths). In this manner, the virtual nodes can collect flow information for the packets traversing the constellation. The virtual nodes, similarly to ground routers, can then use whatever scheduling priority scheme is desired to ensure the appropriate QoS for some particular flows.

7 External Routing

The cost metric sent with the route advertisement from the constellation to the earth is a function of the original external metric, the internal delay to the exit IP address, and perhaps the average traffic load. In this manner, the ground network receives an adequate metric of how costly it is to send a packet to destination through the satellite constellation starting at that geographic point.

In order to construct its routing table, the IP layer must optimize over several metrics: the external metric of the destination, the delay through the constellation, and the expected traffic. While the last two parameters are predictable and periodic, the external metric is dynamic and can change in an unpredictable manner. Therefore, the IP routing table needs to be recomputed on-line whenever the parameters change.

The optimization can be achieved by running a shortest path tree algorithm with a maximum depth of two arcs. The first set of arcs represents the path to each of the exit points. The weight of each arc is the internal delay and expected traffic to the exit point. This weight can also be modified with the expected congestion to that exit point. The second set of arcs represents the path from the exit point to the outside destinations and are weighted by the external metric. Finding the shortest path tree will determine which exit point must be used for each destination.

Figure 4 illustrates such a tree. Nodes 1, 2, 3, 4, and 5 are

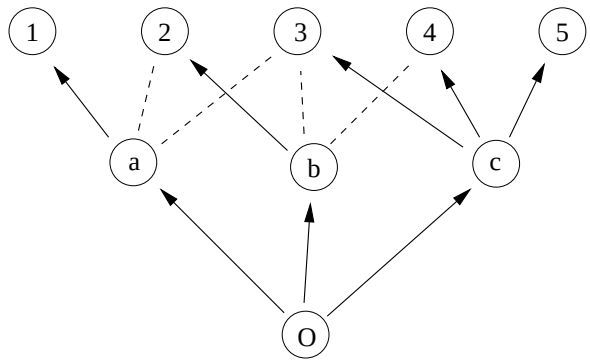

Figure 4: Shortest Path Tree Calculation

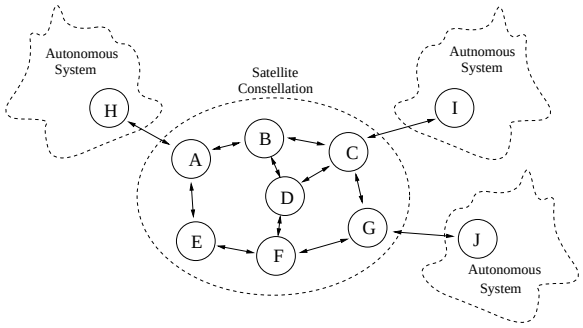

Figure 5: Autonomous Systems

IP destination networks. Nodes a, b, and c are IP addresses of interfaces between three satellites and the earth. Node O is the entry border IP-router which is performing the optimization. Each of the destination network is reachable by more than one interface (dashed or solid lines). However, after the optimization is performed, routing will only take place along the shortest path tree. Therefore, the exit point for network 3 will be interface c, even though it is also reachable by interface b.

The optimal metric calculated for every destination can also be used by the exterior routing protocol. This metric is the one that is exported to the terrestrial routers to advertise route availability.

As discussed in section 2, this routing scheme also allows different satellites in the constellation to route packets to the same destination through different exit points. As an example, see Figure 5. There are three autonomous systems connected through the border routers H,I, and J to the satellite constellation. There is a network N that is reachable through both I and J. They both advertise the same external metric to C and G respectively. Under this routing scheme, it is possible that A, B, C, and D will send packets destined to N through I, while E, F, and G will send them to J.

Satellite A will also send a route advertisement to H notifying its autonomous system the existence of a path to N through the constellation. The advertised path is the best of the two paths for A (through I rather than J).

8 Conclusion

Because of its high mobility and worldwide coverage, integrating a LEO satellite constellation in the Internet raises number of issues that don't appear when integrating a terrestrial network. Transforming each satellite into an IP router doesn't seem to be a reasonnable solution : it requires a lot of computations to be performed periodically ; it is probably a very expensive solution and isn't necessarily adapted to the coexistent non-IP applications. We therefore propose here a layered architecture, where internal routing is carried out by a proprietary protocol, possibly even connection oriented, and external routing is done in a seamless way with the rest of the Internet. The set of metrics we considered relevant for route selection are the external metric, the internal delay, the expected traffic load and the instantaneous traffic load. Using these metrics we proposed examples of internal routing, based on next hop routing (hot potato scheme) and external routing, where we find a shortest path tree to determine the optimal exit gateway.

References

[1] I. Akylidiz and S.H. Jeong, "Satellite ATM Networks: A Survey," IEEE Communications Magazine, July 1997.

[2] Mark Allman, Chris Hayes, Hans Kruse, and Shawn Ostermann, "TCP Performance Over Satellite Links," in Proceedings of the 5th International Conference on Telecommunication Systems, March 1997

[3] Mark Allman, "Improving TCP Performance Over Satellite Channels," Master's thesis, Ohio University, June 1997.

[4] Mark Allman, Chris Hayes, and Shawn Ostermann, "An Evaluation of TCP Slow Start Modifications," Submitted to Computer Communication Review.

[5] J. Gilderson and J. Cherkaoul, "Onboard Switching for ATM via Satellite," IEEE Communications Magazine, July 1997.

[6] R. Mauger and C. Rosenberg, "QoS Guarantees for Multimedia Services on a TDMA-Based Satellite Network," IEEE Communications Magazine, July 1997.

[7] M. Werner et. al., "ATM-Based Routing in LEO/MEO Satellite Networks with Intersatellite Links," IEEE Journal on Selected Areas in Communications, Vol. 15, No. 1, January 1997.

[8] M. Laubach and J. Halpern, "Classical IP and ARP over ATM", RFC 2225, April 1998.