AIAA 2000–1237

NETWORK SIMULATION FOR LEO SATELLITE NETWORKS

Thomas R. Henderson and Randy H. Katz *
Department of Electrical Engineering and Computer Science
University of California, Berkeley
Berkeley, CA 94720

Abstract

Advances in satellite technology are enabling the deployment of large constellations of Low-Earth-Orbiting (LEO) satellites. In this paper, we describe the design and capabilities of a simulation tool that we hope will enable further research on networking over LEO systems. The tool, which has been integrated as part of the freely distributed `ns` simulator, enables detailed packet-level simulations of multiple access protocols, unicast and multicast routing, transport protocols such as TCP, and application performance over future broadband LEO networks (and also GEO networks as a special case). We describe the design of this tool, and discuss how we modeled certain critical parameters and system attributes, including link handoffs between ground terminals and satellites and between satellites themselves. We next illustrate some fundamental delay performance results of LEO systems, obtained by studying simulation models of both a broadband version of the Iridium network and a system based on the proposed Teledesic constellation. We obtain lower bounds on the delay performance that can be achieved in these constellations. Finally, we quantify the performance degradation incurred by using hop counts rather than propagation delays as the cost metrics used by shortest-path routing algorithms. We find that, on average, there is little difference in the routes obtained by the two approaches, although in the higher latitudes the differences can occasionally be substantial.

Introduction

Satellite constellations composed of low-earth-orbiting (LEO) satellites have recently become a reality. Although there are clearly some significant economic hurdles along the road to deployment, there is little doubt that in the near future some communications satellites systems offering broadband data services will be deployed at orbital altitudes lower than that of a geostationary (GEO) orbit. At least two such systems (Skybridge, Teledesic) have already received significant levels of funding and are under development.

The analytical study of the performance of data communications protocols and applications in a LEO satellite network is often made very difficult by the complexity and time-varying nature of the network topology. Therefore, it is helpful to study simulation models of such systems. In this paper, we describe the design of a packet-level simulation tool useful for studying both LEO and GEO networks. The strength of this simulation tool is that it is integrated with a freely-available networking research tool, known as `ns`, already in wide use in the research community. After first providing a high-level overview of those characteristics of LEO networks that are relevant to network simulation, we describe the simulation environment that we constructed and justify our choices for the various simulation parameters we needed to configure. Our simulation models revealed some interesting fundamental delay performance characteristics of LEO networks, which we subsequently illustrate. We conclude by demonstrating how the simulator can be used to study packet routing in LEO networks by exploring the performance degradation incurred by using hop counts rather than propagation delays as shortest-path routing metrics.

LEO Network Characteristics

Constellation Design

Most commercially-proposed LEO constellation designs place the satellites in a number of near-circular orbital planes, in which the satellites are uniformly distributed around the orbit. *Polar orbit* constellations, such as the Iridium and proposed Teledesic systems, are those in which the orbital planes have inclinations close to 90^o, while *inclined orbit* constellations (otherwise known as "Walker" constellations), such as used by the Globalstar system, use an orbital inclination angle less than 90^o so as to obtain a more uniform global coverage.[1] In this paper, we focus on

---

*{*tomh,randy*}*@cs.berkeley.edu.*

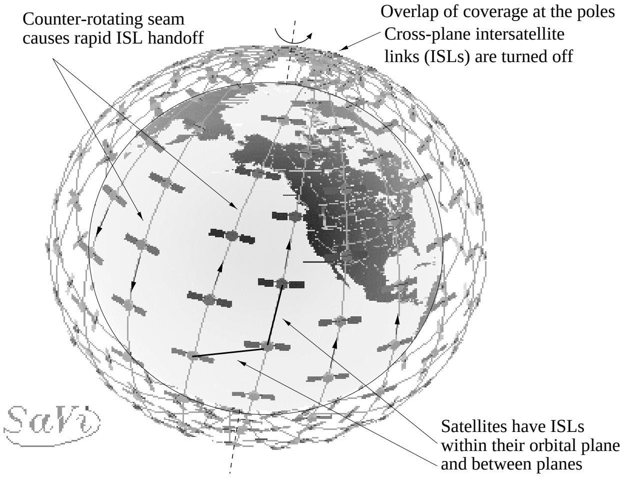

Fig. 1 Example of a polar-orbiting satellite constellation. The figure was generated using the SaVi software developed by the Geometry Center at the University of Minnesota.

polar orbit constellations, such as is illustrated in Figure 1, although other constellation patterns could be designed in our simulator by altering the orbits accordingly. The closer the orbital inclination tends towards a purely polar orbit, the more difficult it is to launch the satellites. Although this constellation design has a concentration of coverage at the poles, it has been found to be a superior design for total Earth coverage with a modest number of satellites,[2] and it has the advantage that most of the intersatellite communications links are not rapidly time-varying. In general, a user anywhere on the Earth's surface should be able to view (above a certain elevation mask*) at least one satellite at any time. We assume that, in general, more than one satellite may be above the elevation mask. Each satellite is equipped with an antenna system capable of directed coverage of portions of the Earth's surface. To obtain higher system capacity, the antenna system incorporates frequency reuse via decomposition of the coverage area into a number of smaller spot beams (i.e., cells). At an altitude on the order of one thousand kilometers, the satellites orbit the Earth roughly every two hours, so that continuous coverage requires link handoff between terminals and satellites. The footprint track of the satellites also has an east-west component as well as the north-south component, since as the satellites orbit in their fixed plane, the Earth rotates beneath them.

The fact that there are multiple satellites above a given terminal's elevation mask does not necessarily imply that the terminal can communicate with more than one satellite. To communicate with a satellite, a terminal must also lie within the radiation pattern of

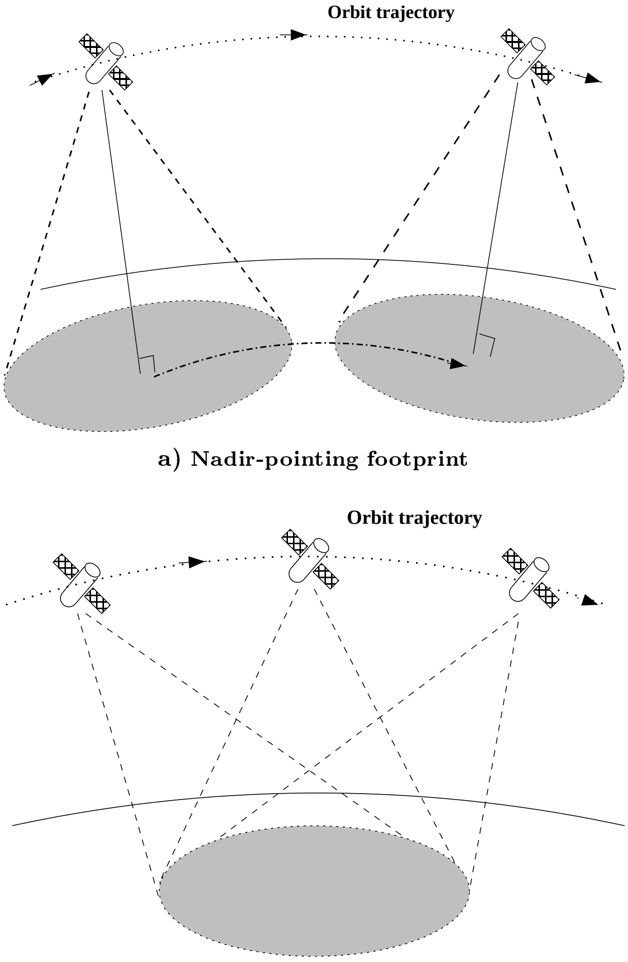

a) Nadir-pointing footprint

b) Earth-fixed footprint

Fig. 2 Satellite-fixed (nadir-pointing) vs. Earth-fixed footprints.

that satellite's directional antenna. Since power management is a concern in LEO systems, especially on the dark side of the Earth, satellite systems such as Iridium deactivate redundant antenna beams to reduce power.[3,4] However, by providing coverage to an area from more than one satellite, the system availability can be increased in several ways. First, the system can compensate for shadowing by terrain and buildings by offering alternative satellites. Second, during daylight hours, if the satellite is located along the same line of sight as the sun, communication will be impossible for a period of time even if the satellite is high in the sky. This is known as a *sun outage* and occurs also in C-band GEO systems, although only for a few minutes each day around the spring and fall equinoxes when the sun crosses the Earth's equatorial plane. Third, increased bandwidth can be provided to a particular geographic area (for example, an area with a lot of users) by using spot beams from neighboring satellites.

---

*The *elevation mask* is the minimum elevation angle of the satellite (with respect to the tangent to the Earth's surface at the terminal's location) above which communications are considered to be possible.

Intersatellite Links

In LEO systems that provide complete global coverage, the satellites are connected via a network of intersatellite links (ISLs). Typically, a given satellite will have ISLs to between four and eight of its nearest neighbors– payload constraints will likely prohibit the use of more than eight ISLs. ISLs are projected to be high-capacity HF or optical links– therefore, in this type of system, the bottleneck links will be the ground-to-satellite links (GSLs), due to the limited RF spectrum available for such links. There are three types of ISLs. *Intraplane* ISLs, which connect a satellite to two or four of its nearest neigbors within the same plane, can be treated as fixed links in the topology. *Interplane* ISLs, which connect a satellite to its nearest neighbors in adjacent, co-rotating planes, are time-varying links for a number of reasons. First, the distance between satellite planes changes as a function of latitude. Second, phasing may not be maintained between the planes, causing the satellites of different planes to slowly drift with respect to one another. Third, the interplane ISLs are switched off in the vicinity of the poles because the antenna pointing mechanism cannot track the rapidly changing angle between the satellites fast enough.[4,5] Finally, note that in a polar constellation (Figure 1), there are two regions in which the planes are counter-rotating, thereby forming a "seam" in the topology. *Cross-seam* ISLs are a special case of interplane ISLs. Cross-seam ISLs, if they exist, are rapidly handed off to the next satellite. If cross-seam ISLs do not exist, communication between two locations on opposite sides of the seam must be routed over a pole. The Iridium system does not support cross-seam ISLs, while Teledesic plans to support them. Keller and Salzwedel have analyzed the problem of cross-seam ISLs in the Iridium system and have concluded that no special antenna steering requirements are necessary,[6] although link acquisition may be challenging. Teledesic plans to use two interplane ISLs per satellite, but at the seam, only one ISL will be active; the other will be used to acquire the next (cross-seam) satellite.[7]

Handoffs

In LEO systems, each satellite covers a portion of the earth's surface with a radiation pattern, or footprint. Each satellite's footprint is typically divided into a number of equal-sized cells, and a phased array antenna on-board the satellite periodically illuminates each cell and then "hops" to another one, creating a "hopping beam" (or scanning beam) schedule over time. The purpose of small cells and electronically-steered hopping beams is threefold. First, if the radiated power is concentrated on a small area, the link budget improves and terminals can use smaller antennas or power. Second, as in cellular systems, system capacity can be increased via frequency reuse. Finally, by varying the hopping dwell intervals, varying amounts of capacity can be allocated to the different cells.

Since the satellites move with respect to the Earth's surface, connections between a terminal and a satellite must be handed over to another satellite when the current satellite drops too low above the horizon. For example, the view time for an Iridium satellite is roughly ten minutes. Therefore, each system must have a technique for controlling handoffs of active communications sessions. There are two general techniques available, depending on the capabilities of the satellite antenna system. The first technique, *asynchronous* handoff, is appropriate for satellite antenna systems that have a nadir[†] pointing footprint. As the satellite moves across the sky, its footprint sweeps across the surface with a constant velocity (on the order of 5-10 km/s), as shown in Figure 2a. When a terminal reaches the edge of the current (leading) footprint, it is handed off to a new satellite whose (trailing) footprint is entering the area. This is the technique used in the Iridium system, and the handoff process is managed by a central control station that monitors each terminal to detect when it nears a coverage boundary.[3] At any point in time, some fraction of the terminals will be near a coverage boundary, so the system must be continually involved in handing off terminals.

An alternative handoff approach has been proposed by Restrepo and Maral.[8] If the satellite system is capable of electronically steering its beam so that it compensates for its motion, the satellite footprint can be fixed for a small interval. As shown in Figure 2b, this leads to "Earth-fixed cells" on the ground. After some time, all of the satellites will be moving away from their respective footprints; the system can then periodically reassign each satellite to a new fixed footprint on the ground. With this approach, the handoffs are *synchronous* since all handoffs occur when the network reorganizes, and the topology will remain static for on the order of tens of seconds to a few minutes. Note that if the system period, which is defined as the least common multiple of the satellite orbital period and the Earth's rotation period, is small, the constellation configuration can be thought of as evolving through a small set of discrete states. Although some authors[9,10] emphasize the importance of a small system period, we assume that the network connectivity between the ground terminals and the satellite network will never be cyclic, so its influence on routing is less significant.[‡]

[†] The nadir point is the point on the Earth's surface that is intersected by a line between the satellite and the center of the Earth. It is also sometimes referred to as the subsatellite point.

[‡] The Teledesic and Iridium systems do not have a small system period due to their choice of altitude.

	Iridium	**Teledesic**
Altitude	780 km	1375 km
Planes	6	12
Satellites per plane	11	24
Orbit inclination (deg)	86.4	84.7
Interplane separation (deg)	31.6	15
Seam separation (deg)	22	15
Elevation mask (deg)	8.2	40
Max. ISLs per satellite	4	8
Cross-seam ISLs	no	yes

Table 1 Parameters for the Iridium and (proposed 288 satellite) Teledesic systems. Both systems are examples of polar orbiting constellations.

Summary of Constellation Parameters

Table 1 summarizes key properties of the (proposed) Teledesic[11,12] and Iridium[1] constellations. Of these values, we will show later in the paper that the seam separation, the elevation mask, and the presence of cross-seam ISLs have important implications on the delay performance of the system.

The `ns` Simulator and Our Extensions

We constructed our satellite simulation tool by building on the UCB/LBNL network simulator known as `ns`, now under development as part of the VINT project.[13] `ns` is a event-driven simulator originally derived from the REAL network simulator.[14] The simulator has an object-oriented architecture, and simulation objects are typically implemented as *split objects*: partly in C++, and partly in an object-oriented extension of the Tcl scripting language known as OTcl.[15] Simulation objects exist simultaneously in both language realms, and functionality can typically be added in either language (generally, functionality that requires per-packet processing is best implemented in C++, while more infrequently processed code is more flexibly implemented in OTcl). The state between the split implementation is made consistent through the use of *bound* instance variables, in which any changes to such variables in one language domain are immediately visible in the other. `ns` is a particularly strong choice for TCP research, since many TCP variants (Tahoe, Reno, NewReno, Vegas, etc.) are standard parts of the simulator. `ns` has also been used extensively for multicast routing and transport protocol research. Until recently, `ns` did not focus on providing detailed simulations of the link and physical layers, but UCB's BARWAN and CMU's Monarch research groups have contributed support for Local Area Networks (LANs), wireless channel error models, and wireless ad-hoc routing protocols.[16,17] `ns` is freely available at `http://www-mash.cs.berkeley.edu/ns/`. However, `ns` was not initially designed to support terminal mobility or dynamic topologies. Consequently, we were forced to introduce several new components into the simulator to more faithfully model LEO networks. Along the way, we attempted to make sure that our extensions were compatible with other aspects of the simulator. Finally, we must emphasize that our simulator is most suitable for simulations of protocols above the physical layer in the ISO model. Orbital motion modeling is basic (e.g., no nodal precession) and the physical channel does not model detailed channel conditions. We chose to abstract out the precise details of physical motion and channel modeling since we were most interested in routing and transport protocol research.

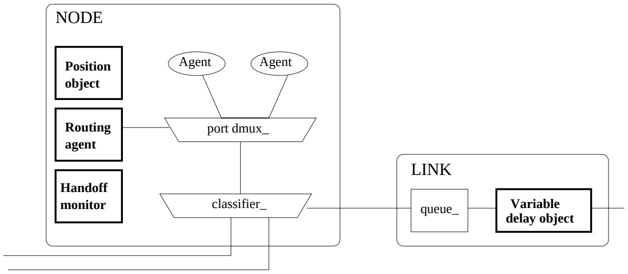

Fig. 3 Extensions to the `ns` simulator (new elements listed in bold type).

`ns` Extensions

Figure 3 illustrates the major additions to the simulator.[§] We first introduced a spherical coordinate system, and added a *position object* to each network node. This position object can be given an initial coordinate and an equation which describes its trajectory through the coordinate system as a function of time. We centered the spherical coordinate system at the Earth's center, with the z-axis aligned with the Earth's rotation axis. This alignment simplified the description of polar orbits and trajectories for Earth terminals. The *link delay object*, which previously returned a fixed propagation delay, was changed to return a value based on the instantaneous positions of the two nodes at the end of the link.

We next built more sophisticated link models than are generally used in `ns`, based on the wireless networking implementations contributed by the BARWAN and Monarch research groups.[16,17] Figure 4 illustrates the main components of a network interface stack. Link layer and medium-access (MAC) protocols can be implemented, and the size and type of the interface queue can be configured. Outgoing packets are copied to and incoming packets are copied from different *channel* objects, thereby enabling the modeling of separate uplinks and downlinks. In addition, received packets can be passed to a propagation model object, where they can be errored according to random probability

---

[§] Since `ns` evolution is on-going, the exact structure of these enhancements embedded in the publicly available simulator is subject to change.

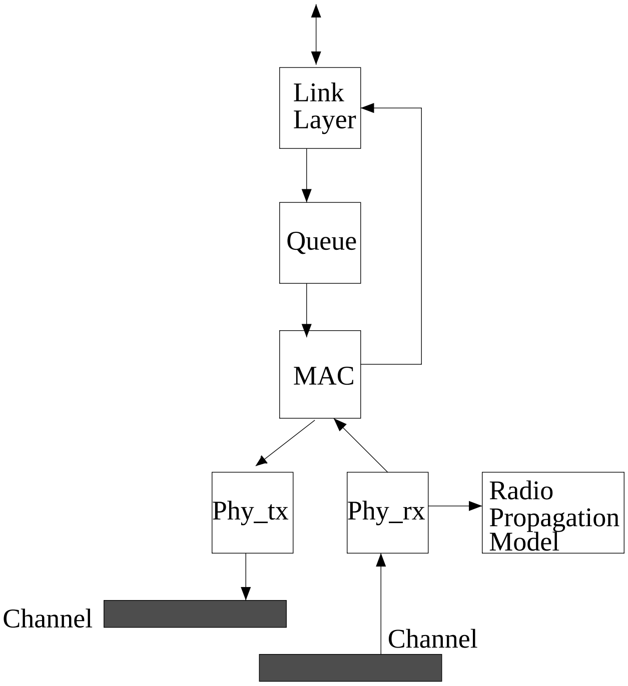

Fig. 4 Elements of the satellite network interface.

distributions (which may be a function of the distance between the sender and the receiver).

Our use of the simulator focused on LEO packet routing, and in the remainder we describe in more detail implementation aspects related to routing. The largest piece of coding involved link handoffs, because `ns` previously did not permit links to be dynamically detached and reattached to different nodes. Furthermore, we needed to introduce *handoff agents* to govern the handoffs. These agents are responsible for monitoring for opportunities to take down, bring up, or handoff links. Various policies for performing the handoffs can be implemented with this structure. Finally, we enabled the construction of dynamic, distributed routing agents for experimentation with distributed routing.

The default routing code in `ns` uses an all-pairs shortest path algorithm to compute new routes for each node in the simulator whenever the topology changes. This algorithm is useful to populate routing tables initially for static topologies, but is very computationally expensive when applied to dynamically changing topologies because it has complexity of roughly $O(n^3)$. To speed up our simulations, we implemented single-source shortest path algorithms and configured the simulator to optionally compute routes on demand (whenever a packet needed to be sent), which yielded a run-time performance improvement of up to two orders of magnitude.

In summary, by composing simulations with existing networking modules and protocols (such as TCP or UDP), `ns` enables packet-level simulations of protocols at the MAC, link, network, transport, and application levels of the OSI reference model. Interested readers can find more information on our implementation in the `ns` documentation.[¶] In the remaining sections, we illustrate some fundamental delay performance results obtained by using the simulator to study the Iridium and proposed Teledesic constellations.

Basic Delay Performance Results

Our primary research focus has been on LEO packet routing performance for polar orbiting constellations. In this section, we illustrate some of the capabilities of our simulation tool by describing some fundamental delay (latency) performance results inherent in the topology of such constellations.

In the remainder of this paper, we make frequent reference to the Iridium[18] and Teledesic[19] satellite constellations; in particular, to a broadband version of the Iridium system that performs packet, not circuit, switching, and the 288-satellite polar orbiting (proposed) Teledesic constellation. Rather than explore the entire design space of possible satellite constellations, we have chosen to focus on the Iridium and Teledesic constellation topologies as examples of feasible LEO systems because they represent two designs that have been considered commercially viable from a frequency management (interference), orbital deployment, and economic perspective. As of this writing, the Teledesic constellation design is subject to change.

Parameter Selection

Beyond the key simulation attributes described above, the following additional details help to more fully describe our models. With respect to the constellation configuration, we made the following two minor simplifications. First, we did not model the minimal orbital eccentricity found in the topologies; our orbits were purely circular. Second, we did not model any drifts in nominal satellite position with respect to the original constellation design, assuming instead that, where possible, the placement of satellites in adjacent orbits will be staggered so as to maximize ground coverage (i.e., in Teledesic, where satellites are nominally spaced at intervals of 15 degrees in each orbit, we offset the position of satellites in adjacent planes by 7.5 degrees). While such a staggering is optimal, it is unclear whether satellite operators will expend the fuel necessary to maintain this phasing (both Iridium and Teledesic plan to hold constant the relative positions of satellites within a particular orbit, but in the Teledesic system there are no guarantees of maintaining any phasing between satellites in different planes).

Iridium satellites are connected to their four nearest neighbors: two satellites in the same orbital plane, and one each in the adjacent planes. Satellites along the counter-rotating seam only have three active ISLs

---

[¶] Available at `http://www-mash.cs.berkeley.edu/ns/`

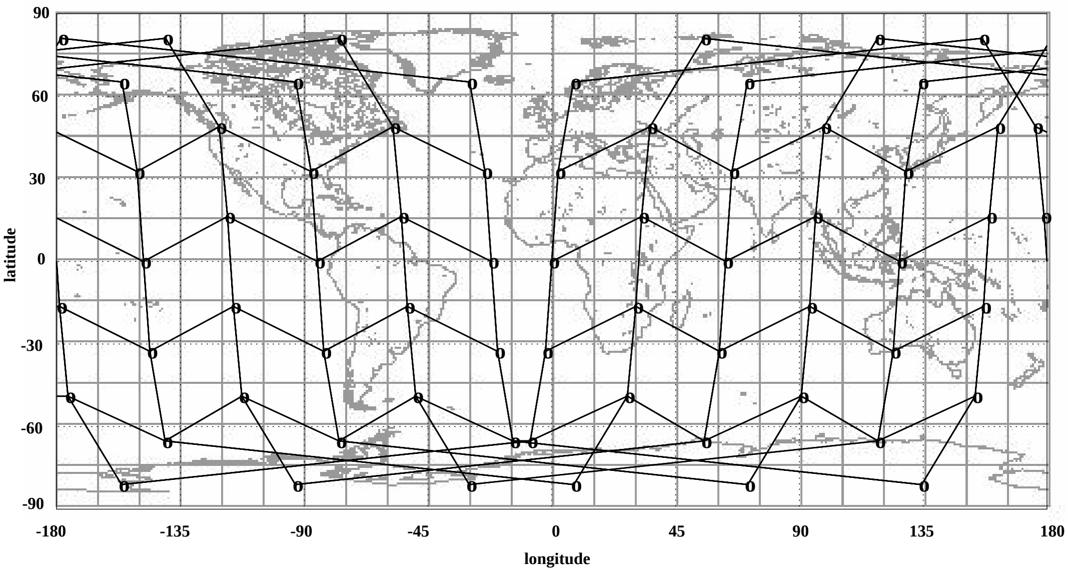

Fig. 5 Snapshot of the Iridium constellation, illustrating active ISLs.

if cross-seam ISLs are turned off (in our simulator, we could also selectively enable cross-seam ISLs for the Iridium topology but generally experimented without them). It is only the cross-seam ISLs that require satellite handoffs, since the intraplane ISLs are static links, and the interplane ISLs only need to be deactivated and reactivated near the poles. The Teledesic system connects to eight nearest neighbors: the four closest satellites within the same plane, one satellite each in the two adjacent planes, and one satellite each two planes away. At the counter-rotating seam, only one ISL is active across the seam and the other is used to acquire the next satellite to be handed off to. We configured the GSLs to be full duplex links at 1.5 Mb/s (i.e., we considered a broadband version of the Iridium system), and the ISLs to be 155 Mb/s for Teledesic and 25 Mb/s for Iridium. The particular values of these bandwidths were not important since we were only considering a minimal amount of traffic. Figures 5 and 6 illustrate snapshots of satellite positions and active intersatellite links for Iridium and Teledesic, respectively. The plots were generated by outputting satellite and link position information and then superimposing the data on a rectangular map projection obtained from the Xerox PARC Map Viewer. Note in the Iridium topology the lack of cross-seam ISLs and the absence of interplane ISLs in the high latitudes.

Various policies for performing handoffs between network nodes are possible– the exact choice of handoff mechanism is sensitive to the satellite hardware capabilities, and Iridium and Teledesic have not publicly revealed their techniques. We implemented both asynchronous and synchronous handoffs as described above. Asynchronous handoffs between ground terminals and satellites work as follows. Each terminal periodically checks whether the satellite that is serving it has dropped below the elevation mask for the terminal. In our simulations, we performed this check every ten seconds, on average (we added a random dither to the timeout interval so that it would vary between five and fifteen seconds); we did not regard the exact value of this timeout parameter as being critical, although too small of a choice leads to slower simulations. Upon checking, if the terminal discovers that the current satellite has dropped below the elevation mask, the terminal searches for another satellite that is above the mask and connects to the first such one found. The technique of synchronous handoffs assumes that topology changes occur only at certain times– our simulator can also be configured such that all nodes perform a topology check synchronously.

We next describe two simulation parameters that are highly dependent on the antenna steering capabilities. Interplane ISLs are deactivated whenever one or both satellites are above a given latitude threshold. We typically set this threshold to 70 degrees, since analysis by Werner indicates that Iridium should be able to maintain ISLs between 60 degrees north and south latitude,[5] and a Motorola patent by Rahnema claims that an Iridium-like constellation is able to keep these links active up to 68 degrees latitude.[20] Although we conjecture that the denser Teledesic constellation may be able to steer these beams beyond a 70 degree latitude, we have no evidence to support this. Handoff agents on board the satellites monitor for this occurrence as well (again, we check every ten seconds on average). Finally, cross-seam ISLs cannot be maintained near the points where the counter-rotating planes intersect; in our simulations, we deactivated these ISLs whenever the satellites were within eight degrees of longitude of one another. More information in the public domain about the antenna steering ca-

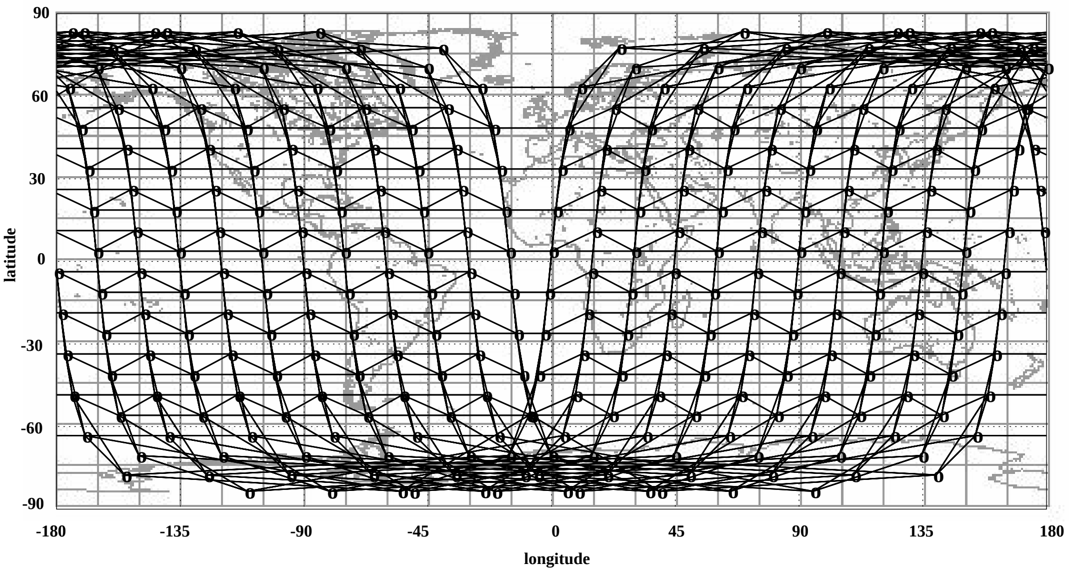

Fig. 6 Snapshot of the Teledesic constellation, illustrating active ISLs.

pabilities of ISLs is needed to make these parameter guesses more accurate. In general, we found that our results were not highly sensitive to these two parameters.

Since our studies were focused on fundamental routing and propagation delay performance measurements, we simplified our simulations (and dramatically improved simulation runtime) by not modeling additional delays due to multiple access contention, framing, and link layer protocols, nor did we consider queueing delays in the network due to heavily loaded links. We also did not model or experiment with link outages or errors due to terrain or sun outages, propagation impairments, or thermal noise. Our rationale for these simplifications was that, while investigating the potential for network load balancing through routing is a good candidate for future research, our simulations on the fundamental routing properties of LEO networks did not require the level of detail that would have resulted from modeling all of the parameters listed above. Nevertheless, the simulator allows for such additional models to be inserted for future research.

Simulation Results

The basic packet delay performance of LEO satellite constellations has never been thoroughly described in the literature. In this section, we quantify typical delay profiles that might be seen by users of future LEO networks.

One of the advantages of LEO systems over GEO satellites is the reduction in propagation delay between the Earth and satellite. Although the end-to-end latency can often be reduced from a quarter of a second to tens of milliseconds by using a LEO system, the delay in a LEO system is inherently variable. Our first experiments with our LEO network simulator were designed to study this delay variability.

Figure 7 is a scatter plot of the end-to-end delay experienced by 10,000 different single packet exchanges ("pings") using the Teledesic system. This simulation was designed to illustrate the range of end-to-end delays that users of these systems might experience. In the simulation, which ran for 20,000 seconds of simulation time, we repeated the following steps every two seconds. We first selected two points at random on the Earth's surface, and instantiated a link between each terminal and the first eligible satellite found (a satellite was considered "eligible" if it was above the terminal's elevation mask). We then configured one of the terminals to send a packet to the other, and measured the one-way delay. The LEO system used a centralized shortest-path routing algorithm based on minimization of the current propagation delay of each link– the routes were centrally computed and instantaneously loaded into each node in the simulator. Although this method of routing violates the speed of light limitation, it represents an upper bound on the achievable performance of a routing algorithm designed to obtain shortest paths. The distance plotted is the great circle distance between the two terminals. The figure illustrates that the end-to-end propagation delay in the Teledesic system is usually below 100 ms if shortest path routes can be found (fewer than 1% of our data points exceeded 100 ms). Also, independent of the distance between the two terminals, a user may encounter an end-to-end delay that can differ by roughly 30 ms, depending on the particular configuration of the satellite constellation. This performance represents a lower bound on the achievable delay and delay variability that can be provided by a LEO satellite network.

Figure 8 illustrates a similar delay scatter plot for the Iridium constellation, were it to include cross-seam

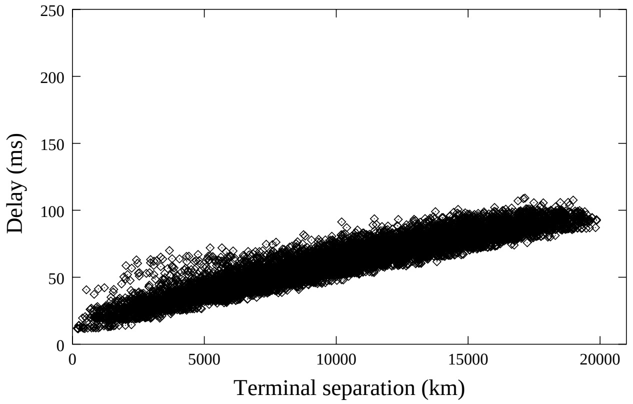

Fig. 7 Scatter plot of the one-way delay experienced by 10,000 different pings between random locations on the Earth's surface, when global min-delay shortest path routing is used (Teledesic constellation).

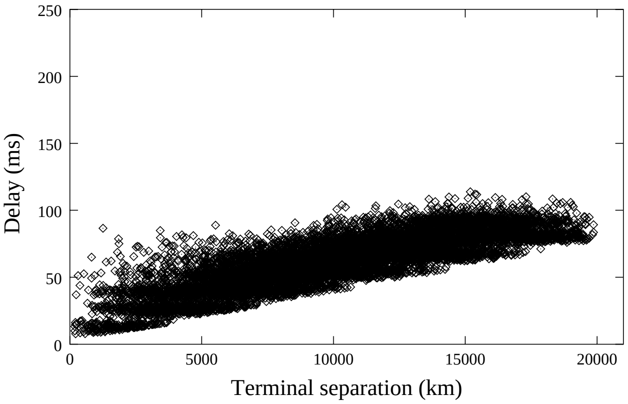

Fig. 8 Scatter plot of the delay experienced by 10,000 different pings, when global min-delay shortest path routing is used (Iridium constellation).

ISLs (we have included them here for comparison purposes). The delay performance of this constellation is similar to that of Teledesic (Figure 7) in terms of the lower bound, but the Iridium constellation exhibits higher variability due to the sparser satellite coverage.

Another way to observe the delay variability is to examine plots of a single session over a long period of time. Figure 9 plots end-to-end delay performance between a terminal located in New York and one in San Francisco over the course of one day. The data points are the delay experienced by a packet sent every 60 seconds. The end-to-end delay varies over a range of roughly 23-60 ms. Over an 11,000 second timespan beginning at time 57,600, the delay is noticeably increased. Even though the Teledesic constellation that we have considered uses cross-seam ISLs, there are certain instances in the mid-latitudes where they cannot be easily maintained due to counter-rotating planes intersecting at mid latitudes. At a smaller timescale (Figure 10), it can be seen that the delay changes slowly as the satellites move with respect to one another, while handoffs somewhere along the route cause a step change in the delay of up to 8 ms. Such changes may cause packet reordering within the network. Although we did not experiment with the performance of TCP connections over such paths, first-order calculations suggest that the amount of packet reordering due to these delay changes should not trigger false fast TCP retransmissions for low to modest transmission rates.‖

A similar plot between the same two terminals for the Iridium constellation is more interesting (Figure 11). Since Iridium does not employ cross-seam ISLs, whenever the seam lies between the two endpoints (which happens twice daily), the packets must be routed over the poles, causing a large increase in delay. Moreover, if a session is active across this seam at the critical handoff, the step increase or decrease in latency will be around 60 ms, which can cause a large amount of packet reordering. This kind of delay variability is inherent in a constellation that does not use cross-seam ISLs, and in the case of Iridium, the step change can be as large as 90 ms. Even without the increased delay at the counter-rotating seam (presuming that cross-seam ISLs are possible in such a constellation), Iridium exhibits much more delay variability than Teledesic, with the delay varying from 20 to 75 ms in much larger discrete steps. This is a direct consequence of having fewer satellites in the constellation, since every routing change that results in a different number of satellite hops also changes the path length by a more significant amount. For example, in Figure 11, the cluster of points around 20 ms is due to the path only traversing two satellite hops, while the cluster of points around 80 ms results from certain instances in time when five satellite hops are required.

Routing Cost Metrics

In global shortest path computations, delay and hop count are two commonly minimized metrics. In the satellite mesh, a minimization of hop count, while potentially easier to compute, is suboptimal because the links have different propagation delays (shorter near the poles). To study the degradation incurred by using hop counts instead of delay as the cost metric, we ran simulations for two identical sets of source-destination pairs (again, 10,000 pings with the endpoints selected at random), with the simulator configured to compute global shortest paths based on link propagation delays (min-delay) on one hand, and hop counts (min-hop) on the other. We then calculated the difference in delay experienced for each ping. Figure 12 plots, as a function of the number of satellite hops, the average and maximum delay degradation from using min-hop instead of min-delay routing for a Teledesic-like con-

---

‖For a 1 Mb/s session with 500 byte packets, a 8 ms delay decrease could cause at most 2 packets to be reordered.

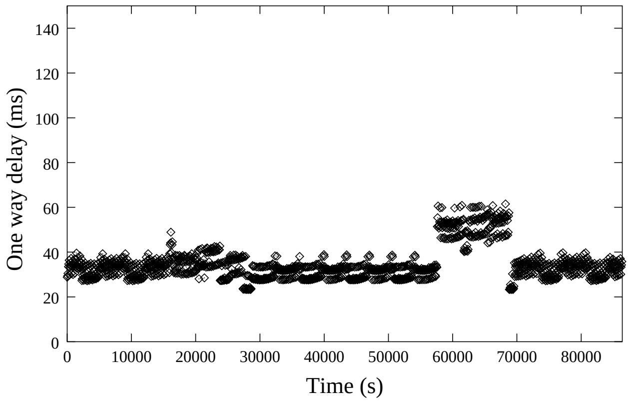

Fig. 9 Delay variation between New York and San Francisco over the course of one day, for the Teledesic constellation.

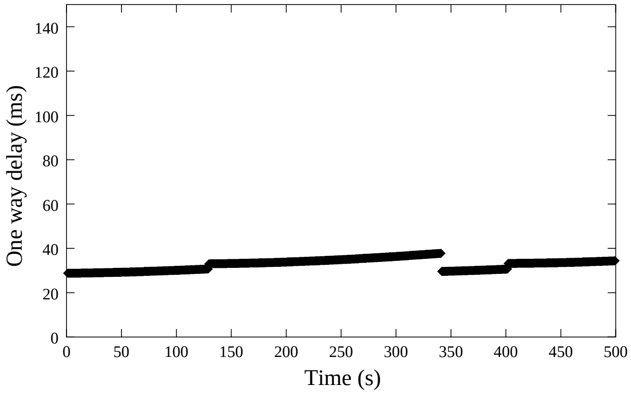

Fig. 10 A view of the previous plot at a smaller timescale. End-to-end delays are characterized by slow variations over tens of seconds punctuated by step increases and decreases in the delay of generally no more than 8 ms.

stellation. The error bars around the average values represent one standard deviation.

Although on average the penalty for using hop count as the routing metric is generally below 10 ms, the maximum difference can be quite high. These outliers were due to particular configurations in the constellation where there were a multiplicity of minimum hop paths through the mesh, some of which used more (short delay) links in the low latitudes, and some of which used more (longer delay) links in the high latitudes. In these outlier cases, the minimum hop path that was found first was one that included a lot of low latitude satellites. Figure 13 illustrates an example (using the Teledesic constellation) of how two routes with the same number of hops can have very different end-to-end delays. Another interesting feature of the data is that the maximum and average delay difference decreases for the very largest distances. Furthermore, the difference between min-hop and min-delay paths in the Iridium system are not as large, because there are fewer satellites and hence, fewer candidate paths to choose from.

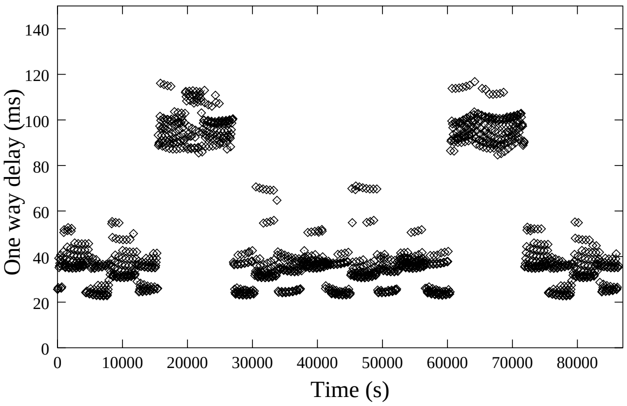

Fig. 11 Delay variation between New York and San Francisco over the course of one day, for the Iridium constellation without cross-seam ISLs.

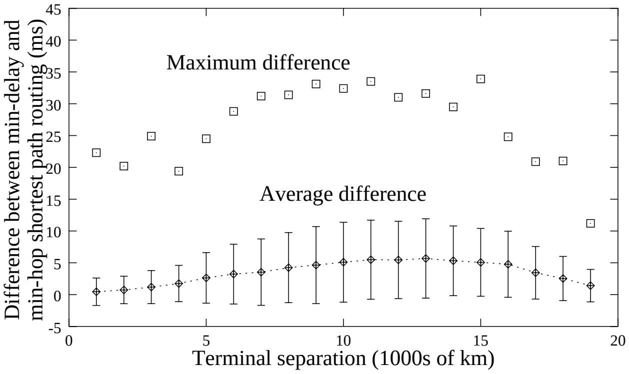

Fig. 12 Average and maximum delay difference between min-hop and min-delay shortest path routing, as a function of the great-circle distance between terminals (Teledesic constellation).

Related Work

In this section, we identify other works that are relevant to the modeling and simulation of LEO satellite systems. Although several other researchers have used simulation in their study of LEO systems, we are aware of only one other publicly available simulation tool for research on LEO systems: the LEONART (Low Earth Orbit Analysis and Research Tools), developed as part of the European ACTS project.[21]

Overviews of LEO satellite networks in general can be found in two recently published books and a tutorial paper on LEO communication systems. Pattan describes orbital mechanics, constellation design, multiple access, frequency issues, and antenna subsystems in.[1] Jamalipour focuses on two key issues: the implications of the projected non-uniform traffic density around the globe, and an analysis of spread spectrum multiple access.[22] Finally, Maral's tutorial paper on LEO satellite systems is a very good overview of the state of the art circa 1990.[23] The paper discusses orbital configurations, basic routing issues, multiple access, and link analyses.

Regarding the modeling of LEO systems, Wood's

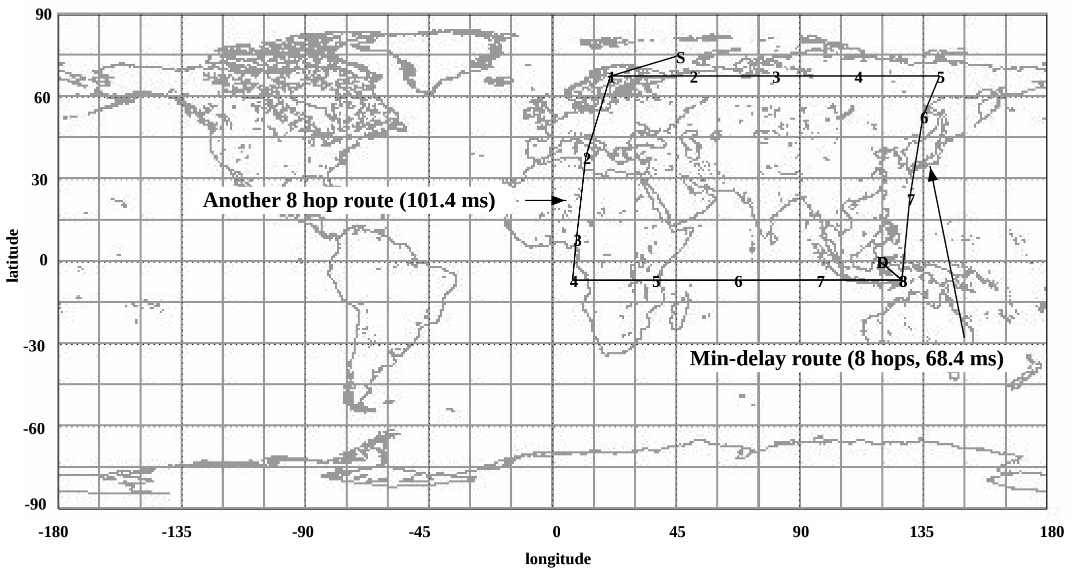

Fig. 13 An example of how min-hop routing can occasionally select a route with a much larger delay than that of min-delay routing. Both routes contain eight satellite hops, but one route uses hops in the higher latitudes which incur much less propagation delay.

thesis is one of the most comprehensive discussions of LEO graph topology issues and tradeoffs in constellation design.[24] Gavish and Kalvenes have studied the relationship between satellite altitude and LEO delay performance, system capacity, and power system design.[4] They find that the altitude of satellites can be a critical design parameter depending on the various constraints of the systems. Finally, Werner describes LEO topology design issues, constraints on intersatellite links, and capacity and traffic engineering aspects, and presents a formal model for the analysis of network connectivity requirements.[5]

Aside from the references listed immediately above, a number of papers focus primarily on constellation design. The work by Adams and Rider is often credited as the basis for the Iridium constellation.[25] The paper by Beste analyzes the design of satellite constellations to provide different levels of continuous, redundant coverage.[2] One of the earliest simulation models of LEO networks, used to evaluate different constellation designs, is described in the paper by Clare, Wang, and Atkinson.[26]

Detailed research literature on proposed commercial LEO systems is difficult to find. There have been a number of high-level papers on the Iridium system– the papers by Grubb and Brunt are probably the most accessible and representative of the group.[27,28] Hubbel contrasts the Iridium and (cellular) AMPS systems with respect to signaling and handoffs, providing useful details about how Iridium handoffs work.[3] Fossa has studied the performance of Iridium in the event of the loss of several satellites.[29] The Teledesic system is an outgrowth of an original system proposal called Calling,[30] which was developed independently and concurrently with Iridium in the late 1980s. Although Calling was oriented towards telephony services, Teledesic has evolved into a broadband system based on packet switching. There is very little publicly available literature on Teledesic's current designs. The paper by Sturza describes various system design issues, while the presentation of Braun describes Teledesic in the context of extending the reach of the Internet through the system.[11,19]

The first author's dissertation[31] provides further references and additional results on LEO satellite simulations.

Summary

We have described the design and construction of a satellite network simulator, based on the freely available `ns` simulator, suitable for packet-level simulations of the MAC, link, network, transport, and application layers. This simulator revealed some interesting fundamental delay performance properties of LEO networks, especially pertaining to the effects of whether or not cross-seam ISLs are present in polar-orbiting constellations. We obtained lower bounds of typical delay profiles that users of future broadband LEO systems may encounter. We also quantified the performance degradation incurred by using hop counts instead of propagation delays as the metrics in shortest-path routing algorithms. Our `ns` extensions for simulating satellite networks have been incorporated into the main `ns` distribution and are freely available to the research community.

Acknowledgments

We thank the many contributors to the VINT project for developing the simulation platform on which we could build our simulator. Among the VINT participants, we thank in particular Giao Nguyen, who assisted us in many aspects of developing the simulation code and who contributed the initial wireless modules to `ns`. Marie-Jose Montpetit and Don Hoffman of Teledesic both were very helpful in sharing proposed Teledesic constellation parameters with us. Finally, we thank Lloyd Wood for providing valuable feedback on various aspects of the simulator. This work was supported by DARPA contract DAAB07-95-C-D154, by the State of California under the MICRO program, and by Hughes Aircraft Corporation, Metricom, Fuji Xerox, Daimler-Benz, and IBM.

References

[1] Pattan, B., *Satellite-Based Cellular Communications*, McGraw Hill, 1998.

[2] Beste, D., "Design of Satellite Constellations for Optimal Continuous Coverage," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 14, No. 3, 1978, pp. 466–473.

[3] Hubbel, Y., "A Comparison of the IRIDIUM and AMPS Systems," *IEEE Network*, Vol. 11, No. 2, March 1997, pp. 52–59.

[4] Gavish, B. and Kalvenes, J., "The Impact of Satellite Altitude on the Performance of LEOS-based Communication Systems," *Wireless Networks*, Vol. 4, No. 2, 1998, pp. 199–212.

[5] Werner, M., Jahn, A., Lutz, E., and Bottcher, A., "Analysis of System Parameters for LEO/ICO-Satellite Communication Networks," *IEEE Journal on Selected Areas in Communications*, Vol. 13, No. 2, 1995, pp. 371–81.

[6] Keller, H. and Salzwedel, H., "Link Strategy for the Mobile Satellite System Iridium," *Proceedings of the 1996 IEEE Vehicular Technology Conference*, Vol. 2, 1996, pp. 1220–1224.

[7] Hoffman, D., "personal communication," 1998.

[8] Restrepo, J. and Maral, G., "Cellular Geometry for World-Wide Coverage by non-GEO Satellites using 'Earth-Fixed Cell' Technique," *Space Communications*, Vol. 14, 1996, pp. 179–189.

[9] Pratt, S. et al., "An Operational and Performance Overview of the IRIDIUM Low Earth Orbit Satellite System," *IEEE Communications Surveys*, Vol. 1, No. 3, 1999, pp. 2–10.

[10] Chang, H., Kim, B., Lee, C., Choi, Y., Min, S., Yang, H., and Kim, C., "Topological Design and Routing for LEO Satellite Networks," *Proceedings of IEEE Globecom '95*, 1995, pp. 529–35.

[11] Braun, H.-W., "Architecture and Performance of Large Internets, Based on Terrestrial and Satellite Based Infrastructure," *Presentation at 1997 ACM Sigmetrics Conference*, 1997.

[12] Montpetit, M.-J., "personal communication," 1998.

[13] Bajaj, S. et al., "Improving Simuation for Network Research," *Technical Report, University of Southern California*, 1999.

[14] Keshav, S., "REAL: A Network Simulator," *Technical Report 88/472, University of California, Berkeley*, 1988.

[15] Wetherall, D. and Lindblad, C., "Extending Tcl for Dynamic Object-Oriented Programming," *Proceedings of the Tcl/Tk Workshop*, July 1995.

[16] Brewer, E. et al., "A Network Architecture for Heterogeneous Mobile Computing," *IEEE Personal Communications*, Vol. 5, No. 5, Oct. 1998, pp. 8–24.

[17] Johnson, D. and Maltz, D., "Protocols for Adaptive Wireless and Mobile Networking," *IEEE Personal Communications*, Vol. 3, No. 1, Feb. 1996, pp. 34–42.

[18] Leopold, R. and Miller, A., "The IRIDIUM Communications System," *IEEE Potentials*, Vol. 12, No. 2, April 1993, pp. 6–9.

[19] Sturza, M., "The Teledesic Satellite System," *Proceedings of 1994 IEEE National Telesystems Conference*, 1994, pp. 123–126.

[20] Rahnema, M., "U.S. Patent No. 5430729: Method and Apparatus for Adaptive Directed Route Randomization and Distribution in a Richly Connected Communication Network," 1995.

[21] DeRidder, J.-J., "LEONART: A Low Earth Orbit Satellite Communication Systems Analysis Framework," *Proceedings of the 4th International Conference on Telecommunications Systems Modeling and Analysis, Nashville, TN*, March 1996.

[22] Jamalipour, A., *Low Earth Orbital Satellites for Personal Communication Networks*, Artech House, 1998.

[23] Maral, G., de Ridder, J., Evans, B., and Richharia, M., "Low Earth Orbit Satellite Systems for Communications," *International Journal of Satellite Communications*, 1991, pp. 209–25.

[24] Wood, L., *Network Performance of Non-Geostationary Constellations Equipped with Intersatellite Links*, Masters thesis, University of Surrey, 1995.

[25] Adams, W. and Rider, L., "Circular Polar Constellations Providing Continuous Single and Multiple Coverage Above a Specified Latitude," *Journal of Astronautical Sciences*, Vol. 35, No. 2, April 1987, pp. 155–192.

[26] Clare, L., Wang, C., and Atkinson, M., "Multiple Satellite Networks: Performance Evaluation via Simulation," *Proceedings of MILCOM '87*, 1987, pp. 404–10.

[27] Grubb, J., "IRIDIUM Overview," *IEEE Communications Magazine*, Vol. 29, No. 11, Nov. 1991.

[28] Brunt, P., "IRIDIUM– Overview and Status," *Space Communications*, Vol. 14, No. 2, 1996, pp. 61–68.

[29] Fossa, C., "A Performance Analysis of the IRIDIUM Low Earth Orbit Satellite System with a Degraded Satellite Constellation," *Mobile Computing and Communications Review*, Vol. 2, No. 4, 1997, pp. 54–61.

[30] Tuck, E. et al., "The Calling Network," *International Journal of Satellite Communications*, Vol. 12, No. 1, 1994, pp. 45–61.

[31] Henderson, T., *Networking over Next-Generation Satellite Systems*, Ph.D. Dissertation, University of California, Berkeley, 1999.