

## WIRELESS MULTIMEDIA AND INTERNET VIA SATELLITE

**Mark A. Sturza**

**Farzad Ghazvinian**

**Teledesic Corporation**

2300 Carillon Point

Kirkland, Washington 98033 USA

+1 818 907 1302 voice

+1 818 907 1357 fax

mark@teledesic.com

+1 206 803 1409 voice

+1 206 803 1404 fax

farzad@teledesic.com

### **Abstract**

There is a significant worldwide demand for broadband communications capacity. The cost to bring modern communications to poor and remote areas is so high that many of the world's people cannot participate in our global community. The solution is a satellite-based broadband network whose service cost in rural, remote areas is comparable to that of wireline networks in advanced urban areas. Such a network can provide a variety of services including Internet access, multimedia conferencing, video conferencing, videotelephony, distance learning, and voice.

### **The Demand for Switched, Broadband Services**

*"The Internet is the most important single development in the world of computing since the IBM PC was introduced in 1981."* [1]. The emergence of the World Wide Web and network-centric computing provide a compelling model for switched, broadband services. Peer-to-peer networking, based on the ubiquity and exponential performance improvements of personal computing, is transforming the way individuals live and businesses create value. Switched broadband connections provide a variety of services including Internet access, multimedia conferencing, video conferencing, videotelephony, distance learning, and voice.

The Internet is still at a primitive stage of development, comparable to the first personal computers in the late 1970's. At that time, it was difficult to imagine the pervasiveness and range of applications of personal computing today. By contrast, the World Wide Web already provides a revealing glimpse of the promise of the Internet, with tens of thousands of companies and millions of individuals exploring, publishing and developing on this new medium. Information can be transmitted anywhere.

Well, not quite anywhere. The promise of the information age is constrained by the lack of access to switched, broadband services in much of the developed, and virtually all of the developing, world. The Teledesic Network will provide these switched, broadband connections on demand, and will extend their reach to anywhere on Earth.

### **The Demand for Broadband Access**

Most of the world lacks the infrastructure for even the most basic telecommunications services:

- Over half the world's population lives more than two hours from the nearest telephone.
- Manhattan has more telephones than all of Africa, where approximately 121,000 of the 151,000 villages have no telephones at all.
- 75% of the world's telephones are located in just 9 countries.
- In India, there are 7 million telephones, clustered in a few large cities, to serve 860 million people.
- The world's telephone waiting list, see Figure 1, is large and growing, and it does not include the huge number of people who have not joined waiting lists because they have no hope of ever getting service.

Even those areas with basic voice service get access through 100-year-old technology – analog, copper networks – most of which will never be upgraded to support digital, broadband services. The developing world, which constitutes 85% of the world's population, has less than 10% of the computer networks.




Figure 1. Global Waiting List for Basic Telephone Service [2]

While access bandwidth is growing in tiny steps, personal computer processing power, which is fueling the demand for bandwidth, continues the exponential growth dictated by Moore's Law, resulting in a "bandwidth bottleneck", as shown in Figure 2.




Figure 2. Bandwidth Bottleneck [3]

While fiber will continue to be the technology of choice for connecting countries and trunking between central offices, it will not be economical in the local loop outside of the most advanced urban areas of the developed world. Widespread broadband deployment will be expensive. PacTel has estimated the cost to upgrade California's infrastructure to support broadband capabilities at \$15 billion, while NTT has projected the cost of fibering Japan at \$900 billion. In Australia, which has strong public policy supporting universal access, 60% of the \$40 billion cost of broadband services – or about 2½ times the cost of the Teledesic Network – would be necessary to provide service to just 30% of the population. The deployment of fiber requires significant infrastructure and backhoes do not obey Moore's Law.

Because demand for broadband services will be spotty and diffuse in the early years, the economics of wireline access – which are dependent on distance and density – will not allow broadband build-out for much of the world's geography and the vast majority of its population. Simply building out broadband in urban areas will generate significant revenues over the next decade. However, without location-insensitive access, broadband will never fulfill its potential. The broadband communications revenue forecast is shown in Figure 3.




Figure 3. Broadband Communications Revenue Forecast [4]

### Implementing Broadband Wireless Access

The deployment of ubiquitous, switched, broadband services will be built on the foundation of wireless economics. Wireless systems require radio spectrum, and broadband service requires a lot of spectrum.

As a practical matter, not enough bandwidth can be aggregated at lower frequencies to provide a new wireless broadband service on a large scale. The short wavelengths at higher frequencies have the advantage of allowing compact terminals and antennas to support high bandwidth applications. The disadvantages of these short wavelengths are heightened terrain blockage and rain attenuation. These disadvantages can be mitigated by sending signals vertically rather than horizontally, so satellite-based systems are attractive.

The lowest frequency band internationally allocated for fixed satellite service, with sufficient bandwidth for broadband applications is Ka-band. Ka-band satellite systems are the solution to the wireless broadband challenge.

Geostationary satellites can never provide fiber-like delays to support services that are seamlessly compatible with existing terrestrial networks. By contrast, low-Earth-orbit satellite systems allow delays comparable to fiber. Low altitude also reduces signal loss and terminal power requirements and allows the use of more compact electronics and antennas. This is the solution Teledesic has adopted.

### Distributed vs. Centralized Architecture

Just as networks on the ground have evolved from centralized systems built around a single mainframe computer to distributed networks of interconnected PCs, space-based networks are evolving from centralized networks relying on a single geostationary satellite to distributed networks of interconnected low-Earth-orbit satellites. In geostationary networks, loss of a single satellite is catastrophic. To reduce this possibility to an acceptable level, reliability must be engineered to the point of diminishing returns where further gains in reliability are only achieved at high cost.

With a distributed network, like Teledesic's, reliability is engineered into the network rather than the individual satellites, reducing the complexity and cost of the individual satellites and enabling more streamlined, automated manufacturing processes. With its distributed architecture, dynamic routing, and robust scalability, the Teledesic Network emulates the most famous distributed network of all, the Internet, while adding the benefits of low-latency and location-insensitive access.

### The Latency Factor

Without knowing for certain what applications and data protocols a broadband network will be called upon to accommodate in the 21<sup>st</sup> Century, it is reasonable to assume that those applications will be developed for the terrestrial fiber networks in the advanced urban areas of the developed world. To ensure seamless compatibility with those networks, a satellite system must be designed with the same essential characteristics as fiber networks – broadband channels, low error rates, and low delays.

Traditional geostationary satellites (GEOs), orbiting at 36,000 km, can never support fiber-like delays. Even at the speed of light, round-trip communications through a geostationary satellite entails a minimum transmission latency – end-to-end delay – of approximately one-half second. This latency causes the annoying delay in many intercontinental telephone calls, impeding understanding and distorting the personal nuances of speech. What can be an inconvenience on analog voice transmissions, however, can be untenable for videoconferencing and for many data applications. By contrast, Teledesic satellites will orbit 50 times closer to Earth than traditional geostationary satellites, allowing delays comparable to fiber.

Applications are developed for terrestrial networks, not for special networks with non-standard characteristics. Companies that build networks that are not compatible with the predominant data protocols and applications are taking a big business risk that their systems will be usable only for specialized, proprietary applications. History has not looked favorably upon companies that have made big bets on inferior products. And since telecommunications customers make purchasing decisions based on their most demanding – not their average – application, geostationary satellite systems may not be a feasible choice if even a relative minority of services are latency-sensitive. In fact, the great majority of data applications are adversely affected by high latency.

### Latency in TCP/IP

Excessive latency causes otherwise high-bandwidth connections to communicate at a fraction of their capacity. Additionally these effects arise not with obscure data protocols or obsolete hardware, but with almost all implementations of the world's most successful data protocol, TCP/IP, which connects the global Internet and is the standard for corporate networking.

Table 1 shows some of the connections a data packet traverses on its way across the Internet from Teledesic's offices in Kirkland, Washington, USA to the US Antarctic Research Base in McMurdo Sound, Antarctica. Because the Internet is designed as a best-effort delivery mechanism, there is no way to guarantee that a packet's trip will be completed in a given time. In fact, congestion has caused the 29 ms delay between Access One's host (Hop #3) and Sprint's network (Hop #5). However, this delay is insignificant compared to the connection over a geostationary satellite between Hops 9 & 11. Although geostationary links can provide up to 250 ms one-way access, this particular link has caused a 574 ms delay in transmission. This 250 ms minimum delay is not a matter of immature technology or implementation issues, but the result of traveling over 72,000 km up to and back from a geostationary satellite.

Table 1. Kirkland, WA, USA to McMurdo, Antarctica TCP/IP Example

HOP#	TIME	INTERNET HOST
1	2 ms	gateway.teledesic.com [206.213.86.1]
3	14 ms	internet1.accessone.com [198.68.191.1]
5	43 ms	sl-ana-2-F0/0.sprintlink.net [144.228.70.2]
8	54 ms	icm-fix-w-H2/0-T3.icp.net [144.228.10.22]
9	56 ms	ARC1.NSN.NASA.GOV [192.203.230.5]
11	630 ms	MCMURDO.NSN.NASA.GOV [128.161.115.2]
14	636 ms	mcmurdo.gov [157.132.103.51]

If the data packet is just part of a file of images to be transferred, why does it matter if it arrives half a second after it is transmitted? It turns out that for TCP/IP, and in fact for all lossless-protocols that guarantee the integrity of the data transmission, latency is a constraining factor on the usable average data rate.

Since the packet may be lost in transmission, a copy of it must be kept in a buffer on the computer in Kirkland until an acknowledgment that the packet arrived successfully is received from the Antarctica computer. (All common data protocols, including Novell's SPX/IPX and OSI, operate on this principle.) The data packet's trip over the geostationary connection takes 250 ms at best, and the acknowledgment packet takes at least another 250 ms to get back, so the copy of the data packet cannot be removed from the buffer in Kirkland for at least 500 ms. Since packets cannot be transmitted unless they are stored in the buffer, and the buffer can only hold a limited number of packets, no new packets can be transmitted until old ones are removed when their acknowledgments are received.

Specifically, the default buffer size in the reference implementation of TCP/IP is 4 Kbytes (32 kilobits). This means that at any given moment, only 32 kilobits can be in transit and awaiting acknowledgment. No matter how many bits the channel can theoretically transmit, it still takes at least half a second for any 32 bits to be acknowledged. So, the maximum data throughput rate is 32 kilobits per ½ second, or 64 Kbps.

Let's put this in perspective. If you take off-the-shelf hardware and software, hook up a broadband geostationary link, and order a E-1 line (2.048 Mbps), you expect to be able to transmit about 2 Mbps worth of data. In fact, any connection via a geostationary satellite is constrained to only 64 Kbps, which is 3% of the purchased capacity. "A very high data rate channel with latency is effectively a low throughput channel." [5].

Changing protocols is not a feasible solution to this situation. The trend in data networking is toward a single "pipe" carrying many types of data (including voice and other real-time data). It is therefore neither useful nor economical to transmit specific kinds of data using custom, proprietary protocols. Can the implementations of standard protocols, such as TCP/IP, be modified to support higher buffer sizes? Yes, in theory, although these modifications are rarely simple or convenient, as computers on both sides of any connection need to be upgraded. Moreover, the maximum buffer size possible in TCP/IP is 64 KB, which still only provides 1.024 Mbps, or 50% of the E-1 rate over a geostationary link.

Even worse, if the geostationary link is not at one of the endpoints of the data transmission but is instead an intermediate connection (as in the Antarctica example above), there is no method to notify the transmitting computer to use a larger buffer size. Thus, while data packets can seamlessly traverse multiple fiber and fiber-like networks (such as Teledesic), geostationary links are unsuitable for seamless intermediate connections.

The interplay of latency and buffer sizes does not effect all data transmissions, only lossless ones. For real-time data, such as voice and video, where it is not essential that all data be transmitted, "lossy" protocols can transmit higher data rates with less overhead. Unfortunately, real-time applications, such as voice telephony and videoconferencing, are precisely the applications most susceptible to degraded quality as a result of high latency.

When customers evaluate GEO versus LEO broadband satellite links, they will need to decide whether they are willing to make do with bandwidth constraints on lossless connections and "choppy" real-time applications, or whether they want connections with the same essential characteristics of fiber. Instead of attempting to modify the entire installed base of network equipment with which one might want to communicate, obtaining seamless compatibility with existing terrestrial networks is incredibly attractive. As both bandwidth requirements and the use of real-time data accelerate, the benefits of the fiber-like service that Teledesic offers become increasingly compelling.

### **The Teledesic Network**

Teledesic plans to begin service by the year 2002. Teledesic does not intend to market services directly to end-users. Rather, it will provide an open network for the delivery of such services by others. The Teledesic Network will enable local telephone companies and governments in host countries to extend their networks, both in terms of geographic scope and in the kinds of services they offer. Ground-based gateways will enable service providers to offer seamless links to other wireline and wireless networks.

