Our baseline also assumes 2x20 MHz of spectrum availability. Exhibit 4-AA shows the economic impact of spectrum availability assumptions. Note that the lack of spectrum increases the cost of the buildout in unserved areas by nearly 5%. The cost impact is relatively small because 2x10 MHz of spectrum is sufficient for 82% of the cell sites (see Exhibit 4-S). The cost impact in areas with negative NPV is even smaller (less than 3%). This is because the cell sites in these areas are typically smaller, so that they also have fewer HUs in them (see Exhibit 4-X for the impact of cell radius on the Investment Gap), which reduces the spectrum needs for the cell sites. Consequently, the impact on the Investment Gap in these areas is also small.

We have not yet addressed the fact that no U.S. service provider currently has more than 2x10MHz of contiguous spectrum in the 700MHz band. But both Verizon Wireless and AT&T Wireless do have noncontiguous spectrum holdings of over 2x20MHz of spectrum across different bands. However, these bands will not all have similar propagation characteristics.

A common deployment strategy used in such situations is to use the lower frequency bands with superior propagation characteristics to serve households further away from the cell site. The higher frequency bands, which can have superior capacity through the use of MIMO techniques, are then reserved for serving those closer to the cell site. This ensures that each available spectrum band is efficiently used.

Cost per cell site

Exhibit 4-AB shows a cost breakdown of a wireless network for all unserved areas. Note that the cost of the network is dominated by last-mile and second-mile costs, which we shall refer

Exhibit 4-Z: Sensitivity of Costs and Investment Gap to Subscriber Capacity Assumptions—Change in Costs and Investment Gap Under Different Downlink Capacity Assumptions

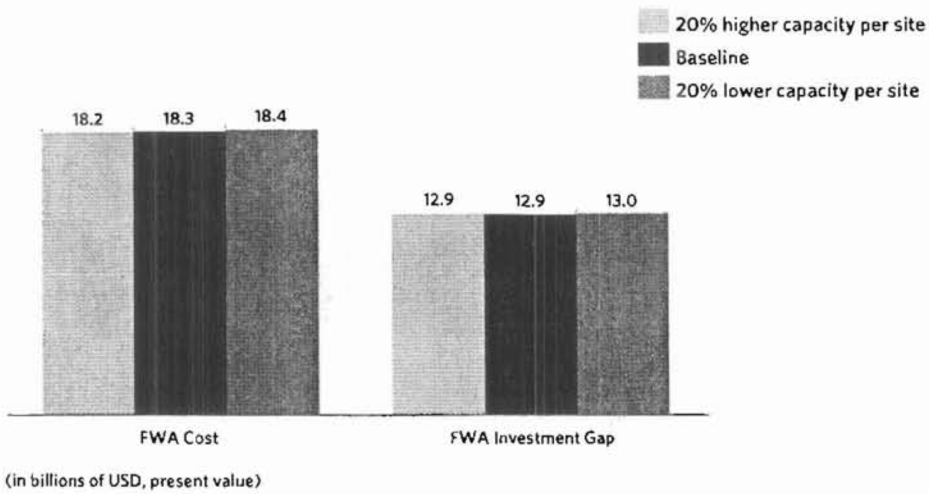

Exhibit 4-AA: Impact of Spectrum Availability on FWA Economics—Change in FWA Costs and Investment Gap Under Different Spectrum Availability Assumptions

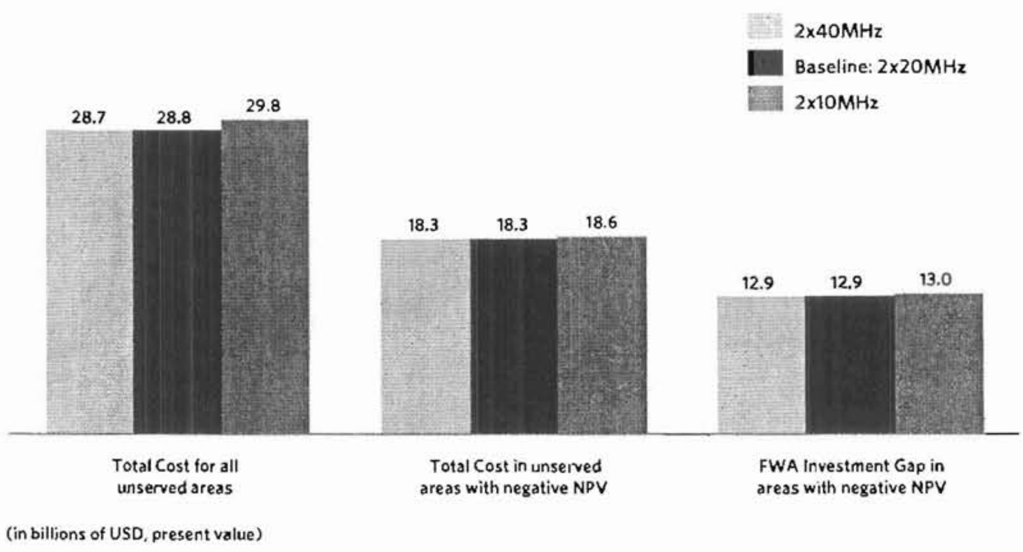

to as simply *site costs*; these account for more than 67% of the total costs. Exhibit 4-AC shows that tower construction/lease and second-mile backhaul costs constitute 68% of the cost of deploying, operating and maintaining a cell site.

Tower construction/lease costs comprise 34% of site costs. To model site costs appropriately, we create one set of hexagonal cells that cover the entire country for each analyzed cell-size (2, 3, 5 and 8 miles). These hexagonal cells represent the wireless cells. Each cell needs to contain at least one tower. To account for the fact that existing services imply existing towers, we turn to several data sources. First, we used the Tower Maps data set of tower locations.[67] For cells that do not include a tower site in that data set, we used 2G and 3G coverage as a likely indicator of cell site availability. Specifically, we assumed that the likelihood of a tower's presence is half the 2G/3G coverage in the hexagonal cell area. For example, a cell that is fully covered by 2G/3G service has only a 50% chance of having a tower site. In areas without a tower, we assume that a new tower needs to be constructed 52.5% of the time;[68] the remainder of the time we assume a cell site can be located on an existing structure (e.g., a grain silo or a church steeple).

In practice, the cost of deploying a wireless network in an area without any wireless coverage today should be higher because of the likely absence of any existing wireless network infrastructure that the provider can leverage. And, with our assumptions above, we capture that effect.

Our cost assumptions in the model indicate that the total 20-year cost of constructing and maintaining a tower is $350K to $450K. By comparison, the total cost of co-locating on an existing structure is only $165K to $250K. Further, our model

Exhibit 4-AB: Cost Breakdown of Wireless Network Over 20 Years[69]

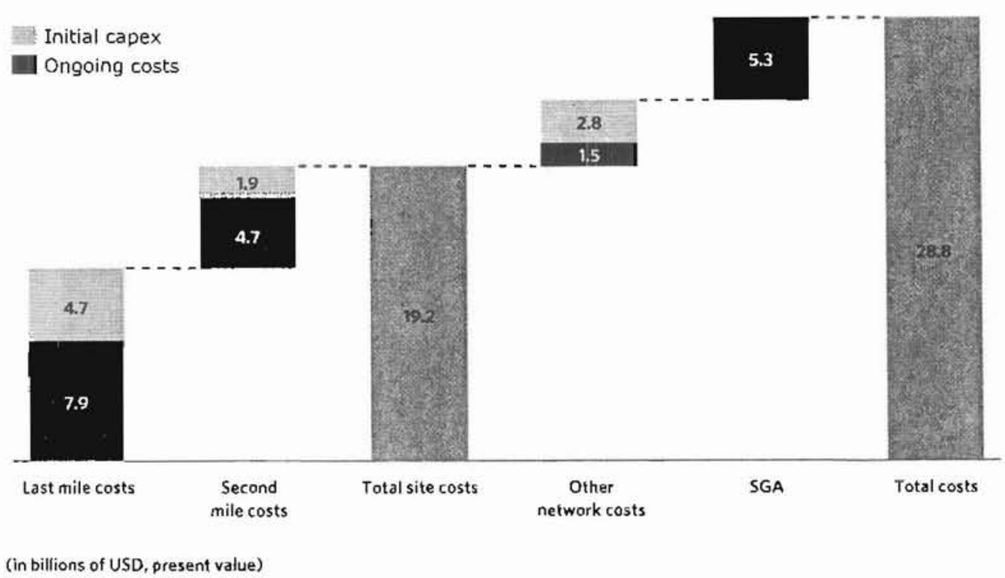

Exhibit 4-AC: Breakdown of Total Site Costs for Wireless Network in Unserved Areas

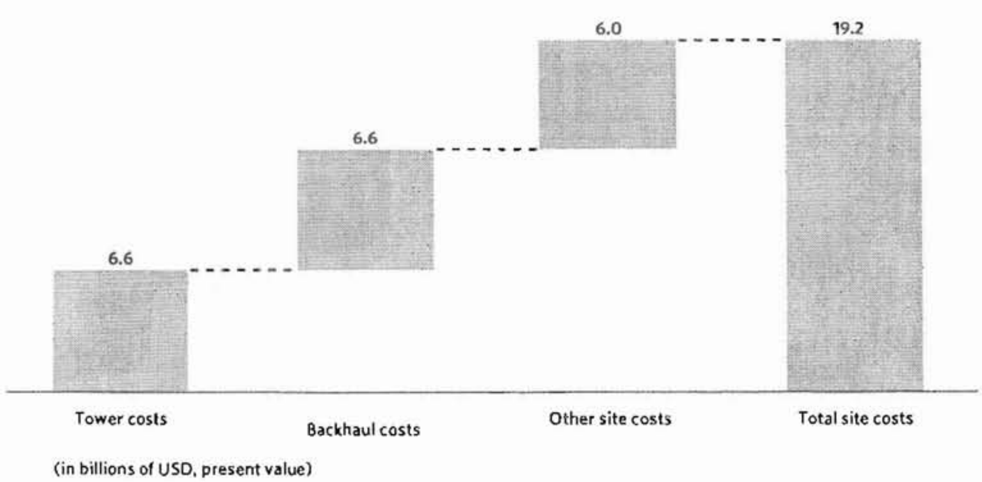

shows that new tower construction is necessary around 15% of the time.

Second-mile backhaul

Our baseline model for the FWA network uses a Hybrid Fiber Microwave (HFM) backhaul architecture with limited microwave penetration. Specifically, we allow a maximum of four hops. Recall that a network architecture that allows a deeper microwave penetration will reduce network costs at the expense of a possible reduction in reliability. Recognizing this trade-off between reliability and cost, we analyze how a restriction on the number of hops affects the cost of the FW buildout and the investment gap. Specifically, we analyze two HFM architectures and compare them with a fiber-only network: (1) Very limited microwave penetration: an HFM network where we allow a maximum of four hops; and (2) Moderate microwave penetration: an HFM network where we allow a maximum of four hops.

In each scenario, we constrained the capacity of the microwave link to 300 Mbps. That limits our ability to daisy-chain microwave links, because the cumulative backhaul needs of all cell sites upstream of a link in the chain cannot exceed the capacity of that link. For example, returning to Exhibit 4-U, the capacity of the link between Cell sites 2 and 3 must be greater than the cumulative backhaul needs of Cell sites 1 and 2; otherwise, one of Cell sites 1 or 2 will require a fiber connection.

Exhibit 4-AD compares the initial investment for the three scenarios. We note that the cost of limiting the number of hops is small—less than 5% when we limit it to two instead of four. This is because most of the unserved regions do not constitute large contiguous areas and can, therefore, be served using a small cluster of cell sites. As a result, the limitation does not severely impact cost. In fact, in the scenario where we allow deep microwave penetration, more than 85% of the cell sites using microwave backhaul connect to a fiber-fed cell site in two or fewer hops.

Exhibit 4-AD: Cost of an HFM Second-Mile Backhaul Architecture—Initial Investment with Different Second-Mile Backhaul Network Architectures

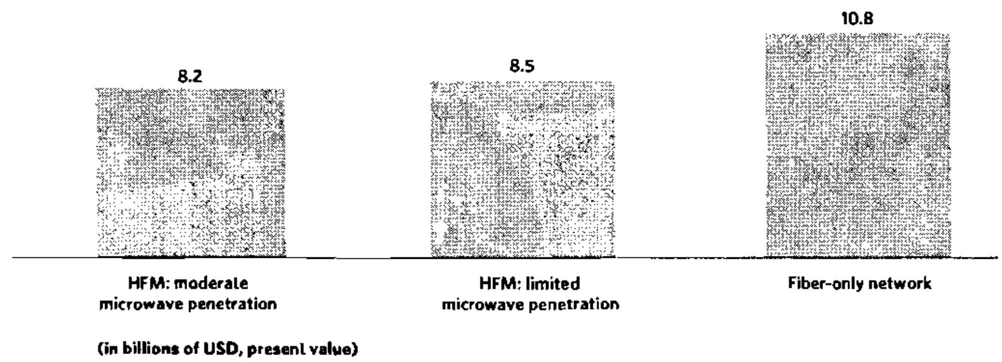

Exhibit 4-AE: Cost Assumptions and Data Sources for Wireless Modeling

Parameter	Source and comments
Tower construction	Mobile Satellite Ventures filing under Protective Order
BTS	Mobile Satellite Ventures filing under Protective Order
Ancillary Radio Access Network	Mobile Satellite Ventures filing under Protective Order
Core network equipment	Mobile Satellite Ventures filing under Protective Order
Site operations	Mobile Satellite Ventures filing under Protective Order
Land Cover	http://www.landcover.org/data/landcover/ (last accessed Feb. 2010) Summary File 1, US Census 2000
Elevation	NOAA GLOBE system http://www.ngdc.noaa.gov/mgg/topo/gltiles.html (last accessed Feb. 2010)
Microwave radio	Dragonwave
Microwave operations	Level-(3) filing under Protective Order
Fiber installation, equipment, operations and maintenance	See cost assumptions for FTTP
Wireless CPE	Based on online price information available for different manufacturers

Conclusions

In order to engineer a wireless network to provide a service consistent with the National Broadband Availability Target, we use the uplink speed target and supplement it with terrain data to compute a maximum cell radius for four different terrain types. In the downlink, we calculate a maximum subscriber capacity per cell site.

A significant driver of variation in per site costs is tower availability and backhaul costs. For backhaul, a Hybrid Fiber Microwave (HFM) architecture results in a lower cost; but a fiber-only network does have the benefit of deeper fiber penetration.

Next, we conduct a sensitivity analysis of our model parameters and assumptions. Not surprisingly, spectrum availability and spectrum bands can have a significant impact on the cost the FWA network as well as the investment gap.

12,000-foot-loop DSL (Digital Subscriber Line)

Telephone networks have traditionally been two-way (or duplex) networks, arranged in a hub-and-spoke architecture and designed to let users make and receive telephone calls. Telephone networks are ubiquitous in rural areas, in part because local carriers have had the obligation to serve all households in their geographic area; this is known as the carrier-of-last-resort obligation. In addition, some telephone companies have historically relied upon implicit subsidies at both the federal and state levels to provide phone service. More recently, they have received explicit financial support through the federal Universal Service Fund (USF). The USF was designed to ensure that all households have access to telephone service at rates that are reasonably comparable to urban rates.

Thousands of independent telephone companies provided service in local markets. But when the telephone network was originally constructed, a single operator, AT&T, dominated it. In 1984, AT&T divested its access network into seven Regional Bell Operating Companies (RBOCs). Over time, the original seven RBOCs have consolidated into three: AT&T (formerly Southwestern Bell, Pacific Telesis, Ameritech, BellSouth and non-RBOC SNET), Verizon (formerly NYNEX, Bell Atlantic and non-RBOC GTE) and Qwest (formerly US WEST).

Exhibit 4-AF:
Breakout of Voice Line Ownership – Telco Consumer Telephone Access Lines Market Share (3Q 2009)[70]

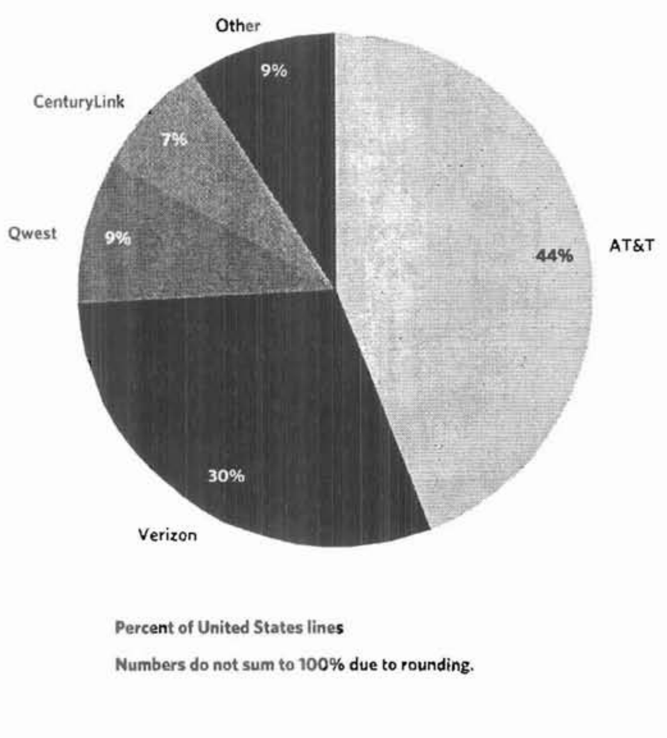

Percent of United States lines

Numbers do not sum to 100% due to rounding.

Consolidation has occurred among smaller Incumbent Local Exchange Carriers (ILECs) as well, with many of them consolidating into CenturyLink, Windstream, Frontier and Fairpoint. Yet well over a thousand small ILECs remain. Today, there are more than 1,311 Telco operators,[71] but the three RBOCs own 83% of voice lines.[72] See Exhibit 4-AF.

The evolution of modern telephone company networks has required significant investments in network capabilities in order to offer broadband access. In the late 19th and early 20th centuries, these networks were built for plain old telephone service (POTS), which provides basic voice service between users over twisted-pair copper wires. These wires, or "loops," were installed between the home and the telephone exchange office via an underground conduit or telephone poles. The basic telephone network architecture and service, originally designed for two-way, low frequency (~4 kilohertz, or kHz), all-analog transmissions with just enough capacity to carry a single voice conversation, are still used today by most homes and businesses. In fact, this network is the basis for the high-speed broadband service known as Digital Subscriber Line (DSL) offered by telecommunications companies.

With the advent of the modem, telephone networks were the first networks to provide Internet access. After all, millions of homes were already "wired" with twisted-pair copper lines that provided POTS. Initially, dial-up Internet used the same analog network designed for voice to deliver Internet access at speeds of up to 56 kilobits-per-second (kbps). To offer high-speed access, the network needed to be reengineered to handle digital communications signals and upgraded to handle the tremendous capacity needed for broadband data and broadcast transmissions. Although twisted-pair copper cables are capable of carrying high-capacity digital signals, the network was not optimized to do so. The large distance between a typical home and telephone exchange offices, as well as the lack of high-speed digital electronics, stood in the way of broadband deployments.

Steps to upgrade telephone networks for broadband:

- Invest in fiber optic cable and optic/electronics to replace and upgrade large portions of the copper facilities for capacity purposes
- Replace and redesign copper distribution architecture within communities to "shorten" the copper loops between homes and telephone exchanges
- Deploy new equipment in the exchanges as well as the homes (DSL equipment) to support the high capacity demands of DSL and broadband
- Develop the technology and equipment necessary for sophisticated network management and control systems

➤ Implement back-office, billing and customer service platforms necessary to provide the services common among telephone operators today

DSL provided over loops of 12,000 feet (12 kft) is a cost-effective solution for providing broadband services in low-density areas. In fact, it is the lowest cost solution for 10% of the unserved housing units. DSL over 12 kft loops meets the broadband target of a minimum speed threshold of 4 Mbps downstream and 1 Mbps upstream, and the backhaul can easily be dimensioned to meet the BHOL per user of 160 kbps.[73] Since DSL is deployed over the same existing twisted-pair copper network used to deliver telephone service, it benefits from sunk costs incurred when first deploying the telephone network.

Capabilities

DSL over loops of 12,000 feet typically uses ADSL2/ADSL2+ technology, which was first standardized in 2005 and which uses frequencies up to 2.2 MHz. As ADSL2+ over 24AWG gauge wire provides rates of 6 Mbps downstream and 1 Mbps upstream, the technology meets the speed requirements for broadband service of 4 Mbps down and 1 Mbps up. Figure 4-AH illustrates how loop length affects speed for ADSL2+.

The technology can perform 1 Mbps upstream on 12 kft of 24 AWG twisted-pair copper loops.[74] In this case, 24 AWG wire is assumed with no bridged taps. Performance with 22 AWG wire, which is often used in rural areas, would yield higher bitrates, while use of 26 AWG wire would yield lower rates.

In order to provide faster speeds than those listed above, DSL operators can bond loops and continue to shorten loop lengths. The bonding of loops can be used to multiply the speeds by the number of loops to deliver rates over 30 Mbps if sufficient numbers of copper loops are available.[75] The performance improvements that can be achieved by shortening loops from 12 kft to 5,000 feet or 3,000 feet and replacing existing technology with VDSL2 are discussed in the DSL 3-5 kft section below. Shortening loops requires driving fiber closer to the end-user; while costly, it could provide much faster speeds that could serve as an interim step for future fiber-to-the-premises (FTTP) deployments. Investment in 12 kft DSL, therefore, provides a path to future upgrades, whether the upgrade is to 5 kft or 3 kft loops or FTTP.

For the small-to-medium enterprise business community, copper remains a critical component in the delivery of broadband. Ethernet over Copper (EoC), often based on the G.SHDL standard, is a technology that makes use of existing copper facilities by bonding multiple copper pairs electronically. EoC can provide speeds between 5.7 Mbps on a single copper pair

Exhibit 4-AG: Telco-Plant Upgrades to Support Broadband

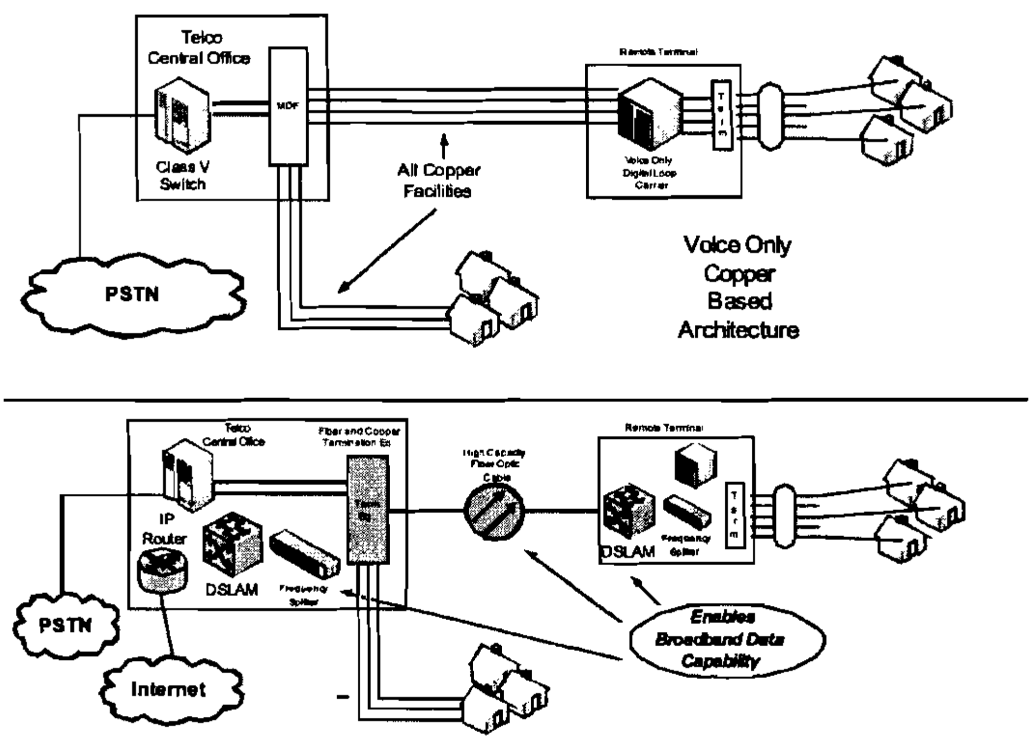

and scale up to 45 Mbps, or potentially higher, by bonding multiple copper pairs. Though middle and second mile connectivity of 100 Mbps is likely necessary, bonded EoC technology can serve as a useful and cost-effective bridge in many areas. Moreover, the embedded base of copper plant is vast—one market study shows that more than 86% of businesses today are still served by copper.[76] Although service providers may prefer to deploy fiber for new builds, existing copper likely will be part of the overall broadband solution, particularly for last- and second-mile applications, for the next several years.

In addition to bonding and loop shortening, marginal speed improvements and increased stability of service levels with ADSL2+ can be achieved through the use of Level 1 dynamic spectrum management (DSM-1).[77] DSM-1 is physical layer network management software that enables reliable fault diagnosis on DSL service. This advancement is available today and may increase bit-rates by up to 10% on ADSL2+.[78] Additionally, DSM-1 helps to ensure stability and consistency of service such that carriers can reach the theoretical 4 Mbps even at high take rates within a copper-wire binder.

We model a 12 kft DSL network that meets the speed and capacity requirements defined in the discussion of 4Mbps downstream requirement in Chapter 3. As outlined in the network design considerations below, we note network sharing in DSL networks does not start until the second mile. The modeled ADSL2+ technology exceeds the speed requirement and includes costs associated with loop conditioning when appropriate. In addition, the modeled build ensures that second and middle-mile aggregation points are connected to the Internet backbone with fiber that can support capacity requirements.

A fundamental operational principle for DSL is that all of the bandwidth provisioned on the last-mile connection for a given end-user is dedicated to that end-user. Unlike HFC, Fixed Wireless, and PON, where the RF spectrum is shared among multiple users of that spectrum and thus subject to contention among them, the last-mile DSL frequency modulated onto the dedicated copper loop and associated bandwidth are dedicated. Sharing or contention with other users on the network does not occur until closer toward the core of the network, in the second and middle mile, where traffic is aggregated (see Exhibit 4-AI). This second- and middle-mile network sharing still occurs in all other access network technologies as well. The "sharing" concept is introduced in detail in the capacity planning discussion in the Network Dimensioning section below.

The ADSL 2+ standard is widely deployed today in telco DSL networks and is assumed to be the minimum required to achieve 4 Mbps downstream and 1 Mbps upstream. The last mile access network ADSL2+ is defined in ITU-T Recommendation G.992.5[11]. The technology provides rates of 6 Mbps downstream and 1 Mbps upstream on the longest loops of a Carrier Serving Area (CSA) (3.7 km or 12 kft of 24 AWG twisted-pair copper loop), with much higher rates attainable on shorter loops.[79]

We perform our analysis and cost calculations based upon a maximum 12 kft properly conditioned copper loop. Loop conditioning costs are applied to those loops that have never been conditioned to offer DSL. For example, if the statistical model showed any DSL speeds for a given census block, we do not apply the loop-conditioning cost since we assume it had already occurred. We believe that only about 1 million homes nationwide have DSL available at a speed below the 4 Mbps

Exhibit 4-AJ: Downstream Speed of a Single ADSL2+ Line as a Function of Loop Length[80] (24 AWG)

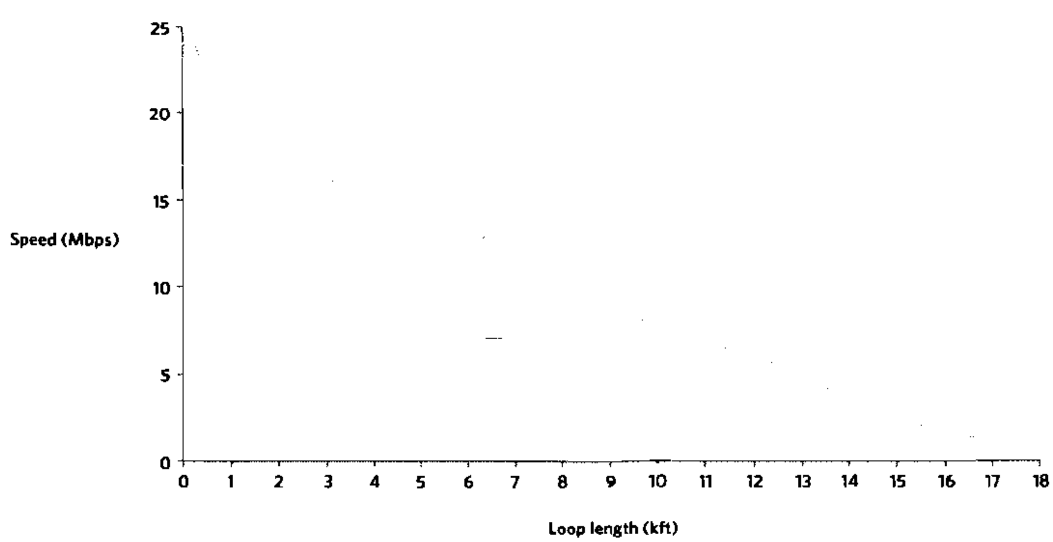

Exhibit 4-AI: DSL Network Diagram

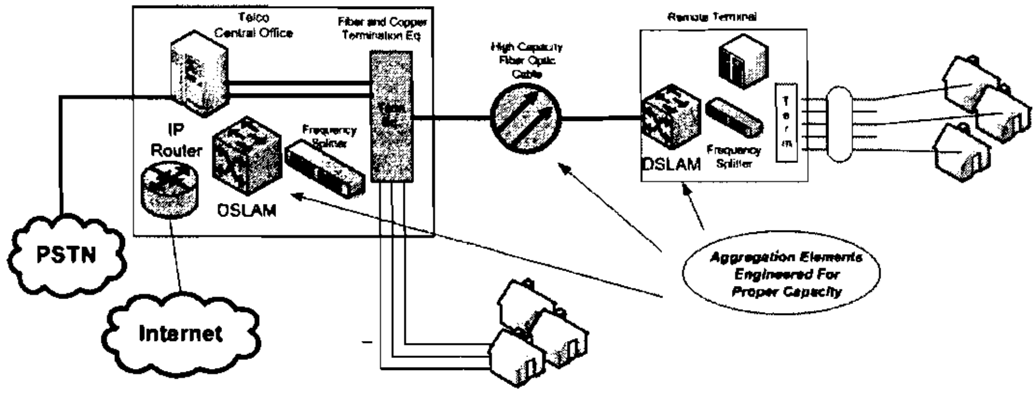

Exhibit 4-AJ: Capacity of a DSL Network–Simultaneous Streams of Video in a DSL Network[81,82]

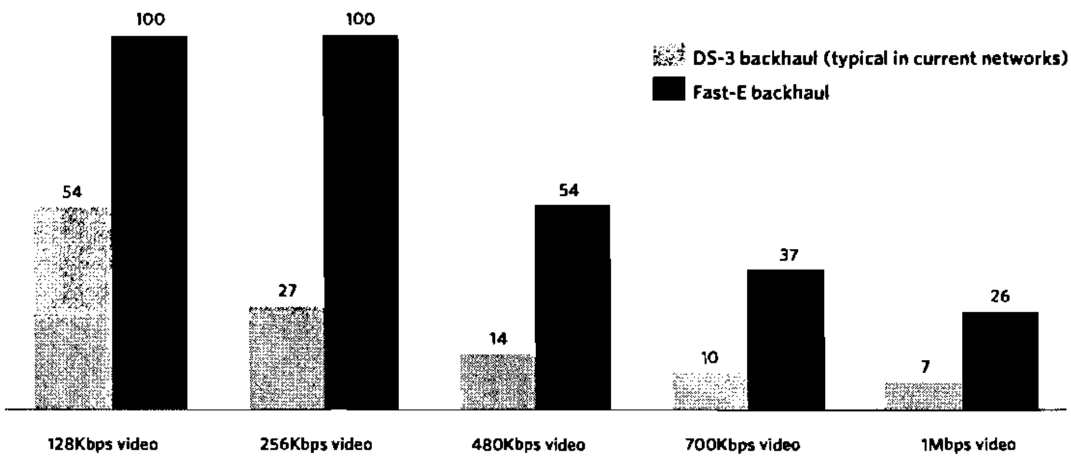

Exhibit 4-AK: Economic Breakdown of 12,000-foot DSL

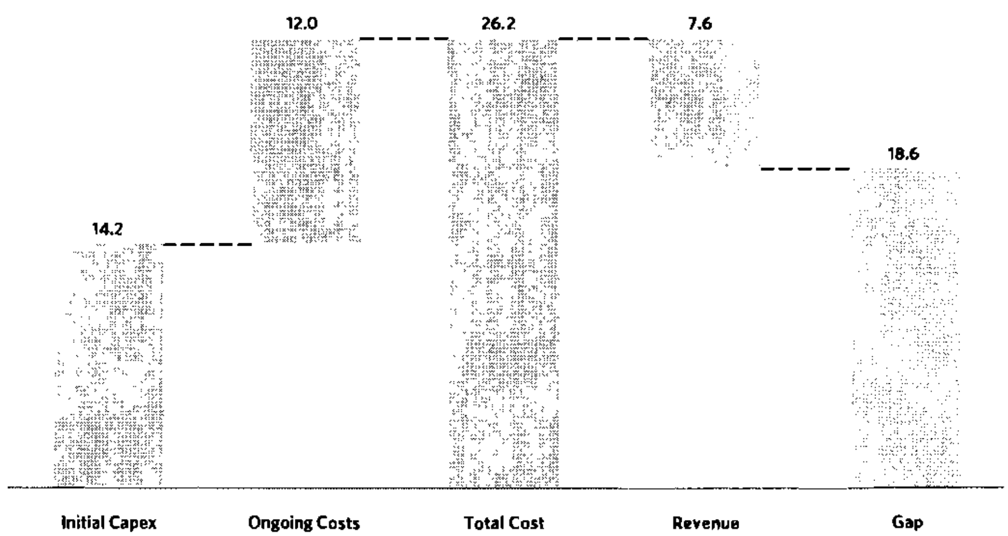

target speed. In the remaining areas, comprising about 6 million housing units, the model includes loop-conditioning costs.

We model the ADSL2+ access network such that DSLAMs are connected to the central office and other middle- and second-mile aggregation points using fiber optic-based Ethernet technology that provides backhaul capacities more than sufficient to meet a 4 Mbps down and 1 Mbps up end-user requirement. Moreover, we calculate the estimated average BHOL per user to be 160 kbps. A typical DSLAM serves between 24-384 subscribers. Since Ethernet-based backhaul provides a minimum of 100 Mbps (a.k.a. Fast-E) bandwidth, scaling to as much as 1 Gbps (a.k.a. Gig-E), the middle- or second-mile aggregation point has sufficient backhaul capacity required to support 4 Mbps down and 1 Mbps up. The resulting capacity of such a DSL network dimensioned with a Fast-E backhaul is shown in Exhibit 4-AJ.

In a DSL network with fewer subscribers, as will be the case in rural areas with low population density, the fraction of users who could simultaneously enjoy video streams of a given data rate would go up proportionately. The dimensioning discussed above is in contrast to the capacity of the network with conventional backhaul provisioning of ~1 Mbps in the shared portions of the network for every 14.5 users.[83]

Economics

The economics of the DSL network depend on revenues, operating costs and capital expenditures. Using granular cost data from DSL operators and vendors, the model calculates the gap to deploy 12 kft DSL to unserved markets as $18.6 billion. Exhibit 4-AK shows the breakout among initial capital expenditure, ongoing costs and revenue.

Initial Capex

Initial capital expenditures include material and installation costs for the following: telco modem, NID, protection, aerial or buried copper drop, DSLAM, cabinet, ADSL2+ line card,

Exhibit 4-AL: Data Sources for DSL Modeling

Material Costs	Source
Telco Modem	Windstream filing under Protective Order
For port sizes of 24 - 1,008:	
DSLAM Unit	Windstream filing under Protective Order
Cabinet	Windstream filing under Protective Order
Allocated Aggregation Cost (CO Ear)	Windstream filing under Protective Order
ADSL2+ line cards	Windstream filing under Protective Order
Fiber optic cabling	FTTH Council
Aerial Drop	Windstream filing under Protective Order
Buried Drop	Windstream filing under Protective Order
NID	Windstream filing under Protective Order
Protection	Windstream filing under Protective Order
Copper cable (24 and 22 AWG)	Windstream filing under Protective Order
Drop terminal/ building terminal (DTBT)	Windstream filing under Protective Order
Feeder distribution interface (FDI)	Windstream filing under Protective Order
Material Labor Costs	
FDI Splicing and Placing labor cost	Windstream filing under Protective Order
DTBT Splicing and Placing labor cost	Windstream filing under Protective Order
Telco Drop and NID labor cost	Windstream filing under Protective Order
Structure Labor Costs	
Duct, Innerduct and Manhole labor cost	Windstream filing under Protective Order
Loop Conditioning cost	Windstream filing under Protective Order
Poles. Anchor and Guy labor cost	Windstream filing under Protective Order
Buried Excavation labor cost under various types of terrain- normal, hardrock and softrock	Windstream filing under Protective Order

allocated aggregation cost, fiber cable up to 12 kft from the end-user, feeder distribution interface and drop terminal/building terminal, as well as the engineering costs for planning the network and the conditioning required on loops (i.e., the removal of load coils[84] and bridged taps[85]). For a detailed list of inputs into our model and the source for each, please refer to Exhibit 4-AL.

Ongoing Costs

Ongoing costs include: replacement capital expenditures required to replace network components at the end of their useful lives, network administration, network operations center support, service provisioning, field support, marketing and SG&A.

Revenues

Revenues are calculated by taking the Average Revenue Per User (ARPU)—which varies according to the level of broadband service/speed provided as well as whether the bundle of services provided includes voice, data and video—and multiplying it by the average number of users. For 12 kft DSL, only data ARPUs are used as incremental to voice, which is assumed present due to the fact that DSL technology utilizes twisted-pair copper wires originally installed and used for POTS.

Satellite

Broadband-over-satellite is a cost-effective solution for providing broadband services in low-density areas. In fact, it could reduce by $14 billion the gap to deploy to the unserved if the 250,000 most-expensive-to-reach housing units were served by satellite broadband. Satellite broadband, as provided by next generation satellites that will be launched as early as 2011, meets our Broadband Availability Target requirements by offering a minimum speed threshold of 4 Mbps downstream and 1 Mbps upstream and BHOL per user of 160 kbps.

Capabilities

Satellite operators are in the midst of building high capacity satellites that will dramatically augment the capacity available for subscribers in the next two years. ViaSat and Hughes, for example, plan to launch high-throughput satellites in 2011 and 2012, and offer 2-10 Mbps and 5-25 Mbps download-speed services, respectively. Upload speeds will likely be greater than the 256 kbps offered today, but no specific upload speeds have been announced. Since satellites are technically constrained by the total capacity of the satellite (>100Gbps), operators could change plans to offer customers at least 1 Mbps upstream even if it is not currently planned. Since the next-generation satellites will be able to offer 4 Mbps downstream and 1 Mbps upstream, satellite broadband meets the technological requirements for inclusion in the National Broadband Plan.

Technical limitations

Over the last decade, satellite technology has advanced to overcome some of the common drawbacks previously associated with it. Due to the properties of the spectrum band used for this service (Ku band downlink 11.7-12.7 GHz, uplink 14-14.5 GHz; Ka band downlink 18.3- 20.2 GHz; uplink 27.5-31 GHz), inclement weather can have an effect on service. However, the ability to dynamically adjust signal power, modulation techniques and forward error correction have all reduced degradation of service except in the most severe of weather conditions.

Since the satellites are in geosynchronous orbit nearly 22,300 miles above the earth, there is a round-trip propagation delay of 560 milliseconds associated with a typical PING (user to ISP and back to user). Recently, integrated application acceleration techniques, including TCP acceleration, fast-start and pre-fetch, have helped mitigate satellite latency for some Web-browsing experiences.[86]

Despite these technological advancements to improve the Web-browsing experience, the latency associated with satellite would affect the perceived performance of applications requiring real-time user input, such as VoIP and interactive gaming. Not only does this delay have a potentially noticeable effect on applications like VoIP, but it would also be doubled in cases where both users were using satellite broadband (e.g., if two neighbors, both served by satellite VOIP, talked on the telephone). Given that most voice calls are local, this could become a significant issue for rural areas if all calls must be completed over satellite broadband.

Spot beams

Broadband satellites use multiple spot beams to provide nationwide coverage. Spot beams use the same spectrum over and over in different geographies, providing more total throughput for a given amount of spectrum. The multiple re-use of frequencies across the coverage area for a satellite provider is similar to a cellular system that reuses frequencies in a "cell." Furthermore, because a spot beam focuses all its energy on a very specific area, it makes more efficient use of the available satellite power.

Nevertheless, a satellite's bandwidth to an end user is provided by and limited to the bandwidth of the spot beam covering that geographic area as well as the total satellite capacity. Therefore, potential network chokepoints for a satellite broadband network include total satellite capacity and spot beam bandwidth.[87] Each spot beam is designated over a section of the United States; once a spot beam is assigned to a certain geographic area, it generally cannot be re-allocated, shifted or moved to cover another area.

With its first leased satellite in 2005 and again with its own satellite in 2007, WildBlue found itself running out of capacity in high-demand regions.[88] In fact, ViaSat plans to aim bandwidth at exactly the same regions where WildBlue's capacity has run out.[89] Many unserved do not live in high-demand areas. These are among the factors that play a role in the capacity assumed available for broadband as discussed below.

Capacity

Providing sufficient capacity for a large number of broadband subscribers, e.g. all of the unserved, may prove challenging with satellite broadband. ViaSat and Hughes believe these next generation satellites have the capacity to serve as many as 2 million homes each;[90] ViaSat has stated on the record that its ViaSat-1 satellite will be capable of providing approximately 1 million households with Internet access service at download speeds of 4 Mbps and upload speeds of 1 Mbps.[91]

Treating satellite as a substitute for terrestrial service, however, requires that satellite be able to deliver service comparable to terrestrial options. Practically speaking, that means that satellite needs to support an equivalent BHOL per user.[92] We believe that the satellite industry could support more than 1.4 million subscribers in 2011 (note that this combines existing capacity with what is planned on being launched) and a total of more than 2.0 million subscribers in 2012 (after the launch of Hughes's next generation satellite, Jupiter). The picture becomes less clear, however, as we look to 2015, when the number of subscribers that current and planned satellites can support would decrease as demand per user grows. End-user demand has been growing at rates as high as 30% annually.[93]

We make certain assumptions in quantifying the number of subscribers that the entire U.S. satellite broadband industry could support with the launch of ViaSat-1 in 2011 and Jupiter in 2012. As there have been no commitments to launch new broadband satellites after 2012, we create a five-year outlook on satellite broadband capacity based on the following assumptions (see Exhibit 4-AM):

- ViaSat will launch a 130 Gbps satellite in early 2011.[94] A comparable satellite, Jupiter, will be launched by Hughes in 2012.[95]
- "Total Downstream Capacity" is 60% of "Total Capacity."
- "Total Usable Downstream Capacity" factors in 10% loss, which includes factors such as utilization and a potential loss of capacity from geographic clustering in which a non-uniform distribution of subscribers would engender certain spot beams to not be fully utilized.

Busy hour offered load (BHOL) assumption

Busy hour offered load, or BHOL, is the average demand for network capacity across all subscribers on the network during the busiest hours of the network. Understanding BHOL is critical for dimensioning the network to reduce network congestion. A more detailed discussion on BHOL can be found later in the Network Requirements section, but the basis for our assumption in satellite is explained here.

Suppose we want to dimension a network that will continue to deliver 4 Mbps. In order to estimate the BHOL for such a network in the future, we first note that average monthly usage is doubling roughly every three years, based on historical growth.[96] There is a significant difference between average usage and the typical user's usage with average usage heavily influenced by extremely high bandwidth users. Next, it becomes crucial to pick the right starting point (i.e., today's BHOL). As the mean user on terrestrial based services is downloading roughly 10 GB of data per month, busy hour loads per user for terrestrial networks translate to 111 kbps busy hour load, assuming that 15% of traffic is downloaded during the busy hour. Terrestrial-based services like cable and DSL experiencing busy hour loads of close to 111 kbps today form the "high usage" case in Exhibit 4-AN.

If we exclude the extremely high-bandwidth users, the average user downloads about 3.5 GB/month, which under the same assumptions for the busy hour would translate to 39 kbps busy hour load. The "medium usage" case in Exhibit 4-AN takes the 39 kbps as a starting point and grows to 160 kbps in 2015; it is this case that we use for our analysis of satellite as well as other networks. The "low usage" case assumes a user downloads 1 GB/month, which translates to 11 kbps; that is roughly what level of service satellite providers offer today of 5-10 kbps.[97] Using 11 kbps as a starting point, the "low usage" case applies the same growth rate as the medium and high usage cases. Exhibit 4-AN summarizes the three usage cases.

Exhibit 4-AM: Available Satellite Capacity Through 2015

Year	2009	2010	2011	2012	2013	2014	2015
Total Capacity (Gbps)	35	35	165	295	295	295	295
Total Downstream Capacity (Gbps)	21	21	99	177	177	177	177
Total Usable Downstream Capacity (Gbps)	19	19	89	159	159	159	159

One reason why the BHOL-per-user might be lower for satellite: satellite operators' fair access policies, which are essentially usage caps, and a degree of self-selection in those who choose satellite-based broadband. However, in a world where users do not self-select into satellite, it is far from certain the extent to which these reasons will still be valid.

Using the above-mentioned assumptions under the "medium usage" case, the satellite industry could support nearly 1 million subscribers by 2015 (see Exhibit 4-AO). Note that each successive year, the satellites can support fewer subscribers due to the doubling of the BHOL every few years noted above. Each next-generation satellite can support approximately 440,000 subscribers using the usage forecast for 2015. Given that the satellite industry in the United States currently supports roughly 900,000 subscribers, this presents a potential difficulty in meeting the needs of the industry's current subscriber base, plus new net additions. If satellite broadband is offered at a level of service comparable to that of terrestrial broadband under the "medium usage" case and BHOL growth continues, satellite providers will need to devote significant incremental capacity to their existing customer base.

Since satellite providers today offer BHOL of between 5 kbps and 10 kbps,[98] our terrestrial-based BHOL assumptions would represent a marked increase in the service level of satellite providers. ViaSat has said on the record that its ViaSat-1 will support a "provisioned bandwidth" (a concept very similar to busy hour load) of 30-50 kbps.[99] However, satellite operators may not be planning for yearly growth comparable to historical terrestrial rates. Thus, despite the growth in satellite capacity between 2010 and 2012, the number of subscribers capable

Exhibit 4-AN: Satellite Usage Scenarios[100]

Year	2009	2010	2011	2012	2013	2014	2015
Busy Hour Load (Kbps) @ 27% growth y-o-y							
Low usage	11	14	18	22	28	36	46
Medium usage	39	49	62	79	100	126	160
High usage	111	141	178	225	285	360	455

Exhibit 4-AO: Satellite Capacity Based on Low, Medium and High Usage Scenarios

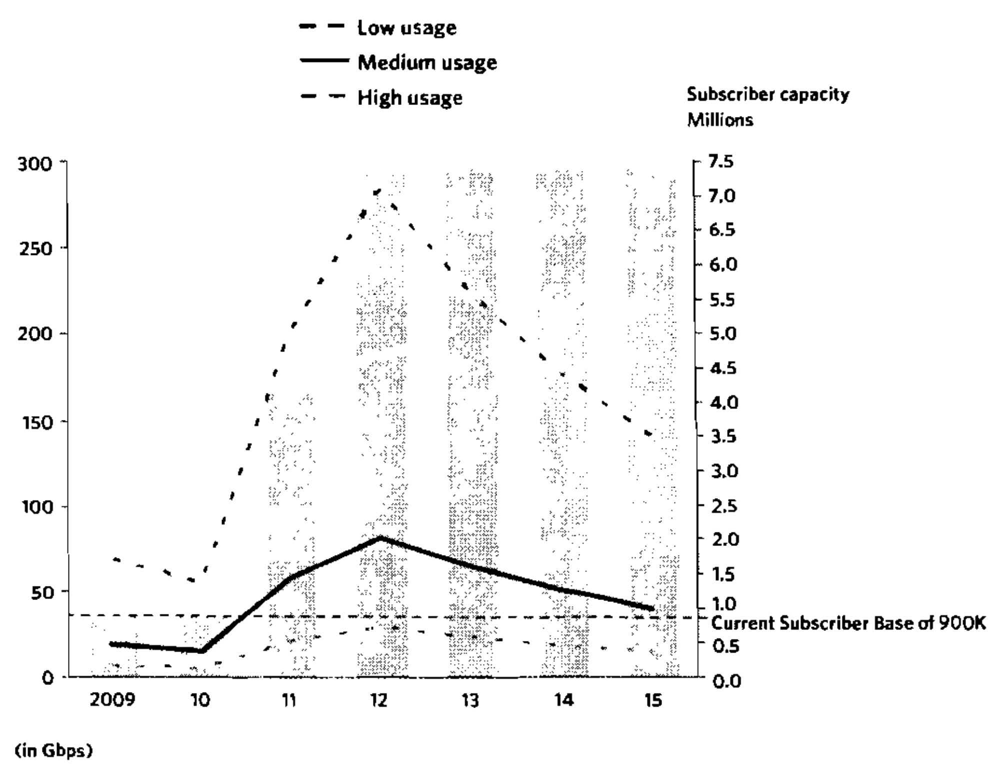

of being supported with our assumptions starts to fall quickly after 2012, absent additional satellite launches. Due to the limited capacity, we do not assume satellite in the calculation of the gap figure of $23.5 billion, but we have contemplated a case in which 250,000 of today's unserved subscribe to broadband over satellite.[101]

If satellite is used to serve the most expensive 250,000 of the unserved housing units, it will reduce the gap. Some 250,000 housing units represent 3.5% of all unserved, <0.2% of all U.S. households, and account for 57%, or $13.4 billion, of the total gap. Exhibit 4-AP shows the remaining gap if satellite is used to serve the most expensive census blocks containing a total of 250,000 subscribers.

The map in Exhibit 4-AQ identifies the location of the highest gap census blocks with a total of 250,000 housing units that we assume are served by satellite in Exhibit 4-AP.

Economics

Nearly all of the costs for satellite broadband are fixed and upfront with the development, construction and launch of the satellite. Each next-generation satellite costs approximately $400 million, which includes satellite construction, launch insurance and related gateway infrastructure.[102] Operating costs for a satellite broadband operator are typically lower than for a wired network provider. Because a single satellite can provide coverage for the entire country with the exception of homes on the north face of mountains or with dense tree cover, the cost of satellite broadband remains constant regardless of household density, which makes it a great option for remote areas.

However, due to the capacity constraints of each satellite, and the growth in use discussed above, satellite operators likely need to continue adding new satellites over time. Estimates of the initial capital expenditure to provide all 7 million of the unserved housing units using satellite broadband service are near $10 billion, including the cost of up to 16 next-generation satellites as well as the CPE and installation for each end-user, assuming the "medium usage" scenario. Timing may be an issue if satellite broadband were deployed as the only means of reaching the unserved, as a next-generation satellite takes approximately three years to build.[103]

Additionally, with each satellite capable of supporting roughly 440,000 subscribers using our assumptions, satellite operators could be forced to potentially more than double their current monthly subscriber fees, which today range from $60-80 per month, in order to maintain the same return on investment as today.

The cost-per-subscriber is driven by the high up-front costs associated with building and launching a satellite. As capacity required per-subscriber increases, the number of subscribers that each satellite can support drops. That drop, in turn, means that there are fewer subscribers over whom to amortize high fixed costs. Thus the average cost-per-subscriber increases, creating less favorable economics over time or requiring higher monthly fees to be charged to the end-user as described above.

Even with greater efficiency of planned satellites like ViaSat-1 or Jupiter, which provide more capacity per launch, the average capex-per-subscriber will only grow with the increase in effective load-per-user. See Exhibit 4-AR, which shows the average capex per subscriber at various levels of monthly usage. The levels of usage correspond to the low, medium and high usage cases described above.

In Exhibit 4-AR, the capex of a satellite (including build, launch and insurance), the associated gateway infrastructure and the CPE is divided by the number of subscribers, depending on the usage characteristics. Note that the average cost calculation may in fact overstate the true cost of a given subscriber over the lifetime of the satellite.

Exhibit 4-AP: Economics of Terrestrially Served if Most Expensive Housing Units are Served with Satellite[104]

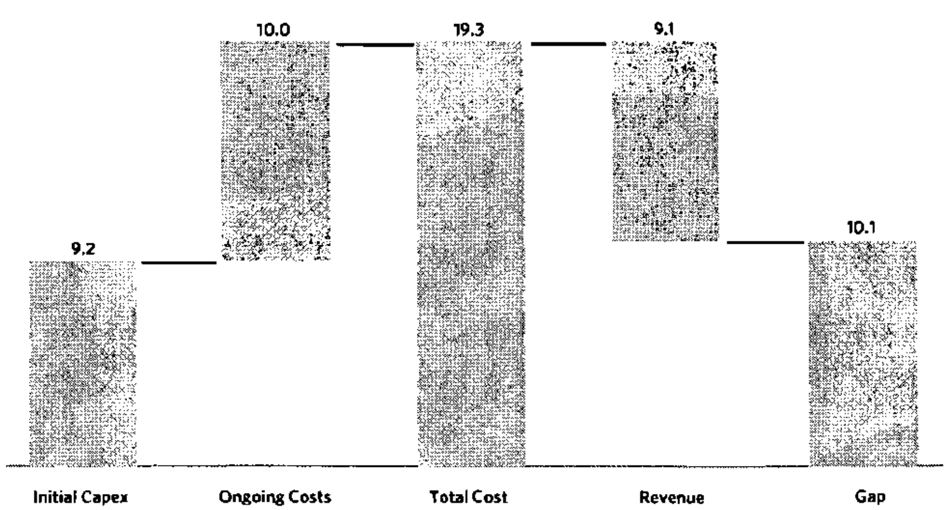

Buy down

Due to the relatively high price of satellite broadband service, there may be a need for a subsidy of the monthly ARPU for those served by satellite broadband. Current ARPU for satellite broadband is generally $60-80 per month depending on speed tier, service provider and choice of whether to purchase CPE upfront or pay a monthly fee for it.[105] For illustrative purposes, assuming a starting point of $70 per month, end-user support to reduce the price to $35 monthly would cost $105 million annually (250,000 people x $35 difference in ARPU x 12 months).

Exhibit 4-AQ:
Location of Highest-Gap Housing Units

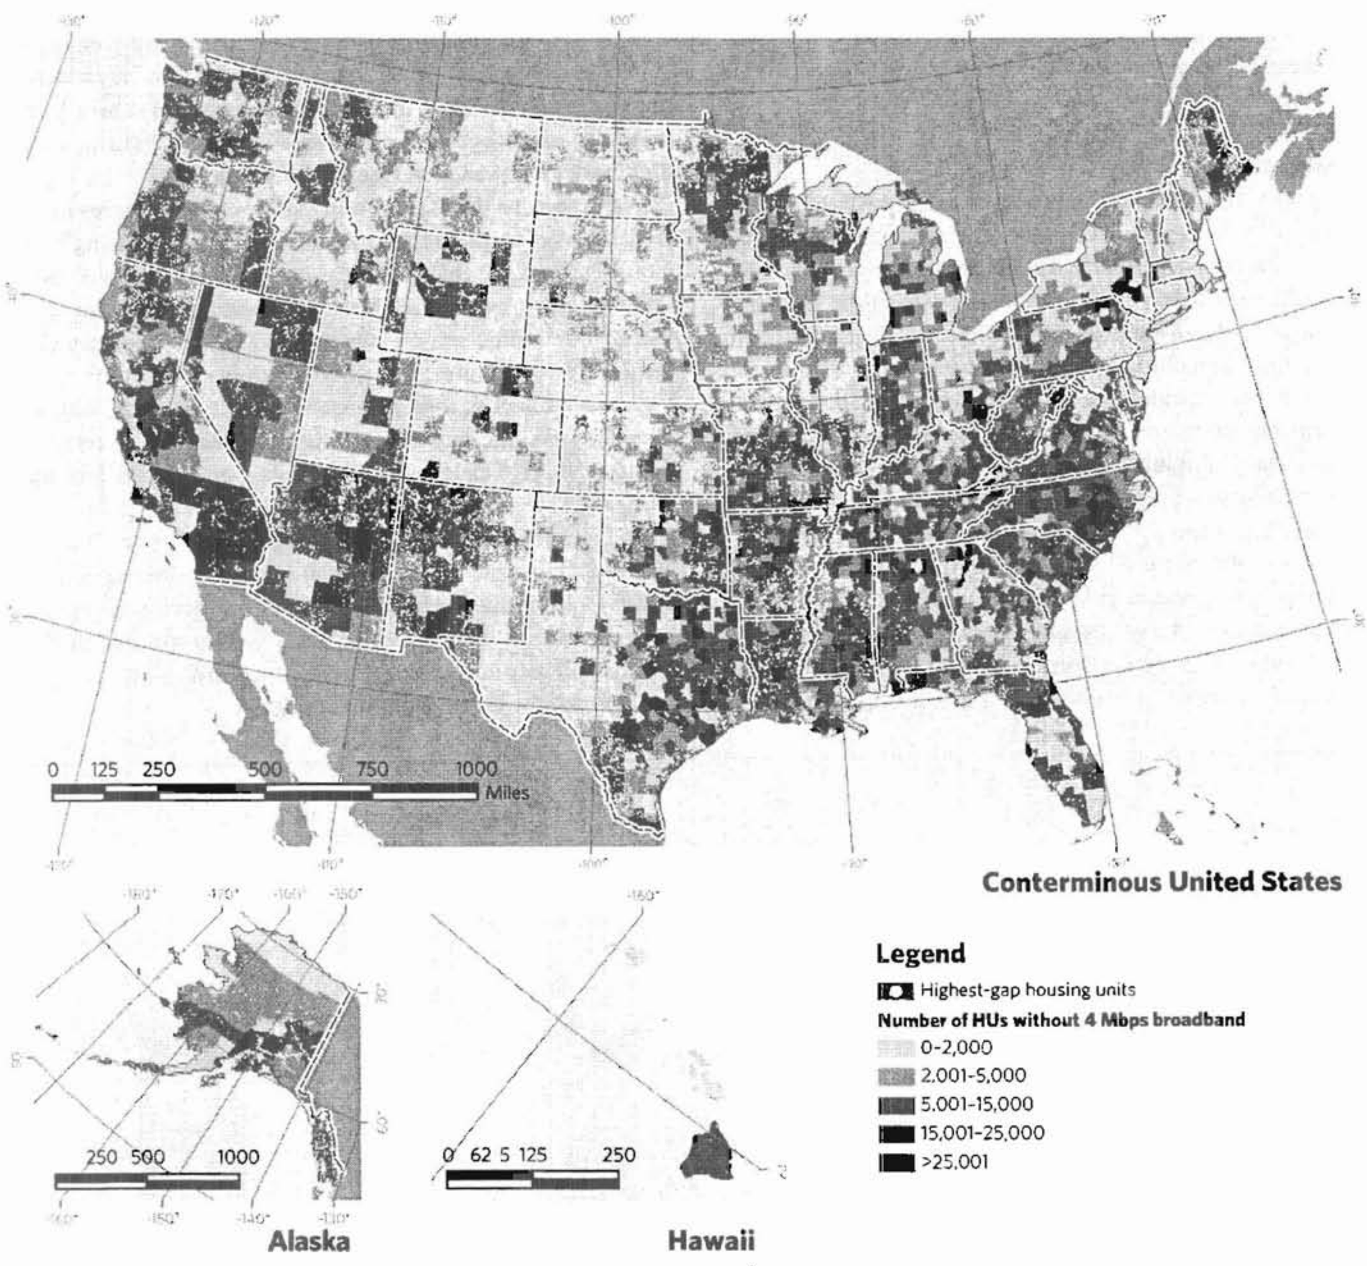

Over 20 years, discounting at 11.25%, the present value of this annual amount is over $800 million.

As discussed above, if satellite operators were to assume a higher use case to provide a level of service comparable to terrestrial providers and to double their price to ensure consistent return on investment (note that the ability to generate enough cash flow affects their ability to finance future satellites), the required subsidy would grow proportionately. Assuming a contemplated starting price of $120, the subsidy required would be $255 million annually (250,000 people x $85 difference in ARPU x 12 months) to yield an end-user price of $35. Over 20 years, the present value of this annual expenditure is roughly $2 billion.

Despite these challenges, we believe that satellite can still provide an economically attractive service for some, and that satellite providers can be an alternative to terrestrial providers, both wired and wireless. However, as we explain further in Chapter 3, uncertainty—principally about the optimal role satellite might play in the disbursement process—has led us to not explicitly include satellite in the base-case calculation.

Exhibit 4-AR:
Satellite Capex per Subscriber— Average cost/POP at Scale

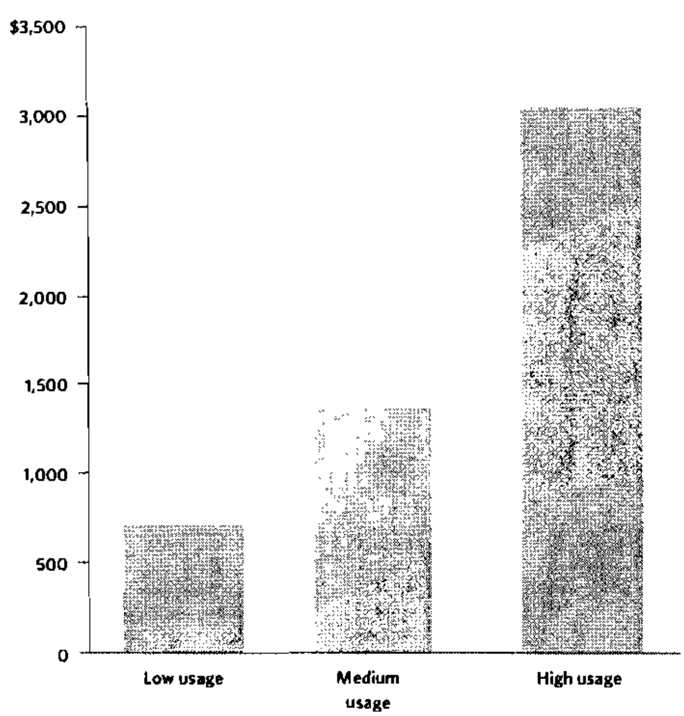

TECHNOLOGIES NOT INCLUDED IN THE BASE CASE

Fiber-to-the-premises (FTTP)

Fiber-to-the-premises (FTTP) offers the greatest potential capacity of any of the technologies considered, making it the most future-proof alternative. The tradeoff for this is the additional construction cost incurred to extend fiber all the way to the premises, making FTTP the most capital-intensive solution considered. On the operational side, the extension of fiber enables the removal of all active components in the outside plant, providing FTTP with a substantial operational savings over competing technologies with active electronics in the outside plant.[106] However, in unserved areas in particular, these savings are insufficient to overcome the initial capital expenditure burden, making FTTP the solution with the highest lifetime cost and the highest investment gap.

Capabilities

There are three basic types of FTTP deployments: point-to-point (P2P) networks, active Ethernet networks and passive optical networks (PON). PON makes up more than 94% of the current residential FTTP deployments in the United States.[107] PON has the advantage of offering lower initial capital expenditure requirements and lower operating expenditures relative to P2P and Active Ethernet deployments, respectively. As such, our analysis utilized PON as the modeled FTTP network.

Exhibit 4-AS shows the capabilities of the varieties of PON currently in use in the United States.[108]

While the majority of homes currently passed by FTTP deployments in the United States are passed by BPON networks, more new deployments are utilizing GPON.[109] PON is a shared medium, meaning that a portion of the access network running between the headend and the passive optical splitter is shared among multiple end-users.

Typical PON deployments share a single fiber in the feeder portion of the access network among 32 end-users. See Exhibit 4-AT. For BPON, this yields a fully distributed downstream capacity of 19.4 Mbps and upstream capacity of 4.8 Mbps per end-user. For GPON, these capacities increase to 78 Mbps downstream and 39 Mbps upstream. As these speeds do not factor in any oversubscription, with a reasonable oversubscription of 15:1,[110] an operator with either a BPON or GPON deployment could easily offer its customers a product with download speeds exceeding 100 Mbps, far exceeding what we anticipate being required in the foreseeable future.[111] As such, FTTP clearly is a candidate from a capability standpoint for delivering broadband to the unserved.

Future PON architectures

PON architectures continue to evolve. The full standard for the next evolution of GPON is expected to be completed in June

2010, with deployments starting in 2012. It will offer download speeds of 10 Gbps and upload speeds of 2.5 Gbps and 10 Gbps, and it will be able to coexist on the same fiber as GPON. Deployments of the next evolution of EPON could even predate those of GPON, offering download speeds of 10 Gbps and upload speeds of 1 Gbps and 10 Gbps.[112] See Exhibit 4-AU.

Beyond these near-term standards, numerous long-term ideas are being presented. For example, Wave Division Multiplexing PON would replace the splitter with an arrayed wave guide and utilize a different wavelength for each end-user. This would effectively eliminate the sharing of the fiber in the second mile that takes place with existing PON varieties, enabling dedicated end-user capacities of 10 Gbps or more.

Exhibit 4-AS: Capabilities of Passive Optical Networks (PON)

	BPON	EPON	GPON
Standard	ITU-T G.983	IEEE 802.3ah	ITU-T G.984
Bandwidth	Downstream up to 622 Mbps	Downstream up to 1.25 Gbps	Downstream up to 2.5 Gbps
	Upstream up to 155 Mbps	Upstream up to 1.25 Gbps	Upstream up to 1.25 Gbps
Downstream wavelength(s)	1490 and 1550 nm	1550 nm	1490 and 1550 nm
Upstream wavelength	1310 nm	1310 nm	1310 nm
Transmission	ATM	Ethernet	Ethernet, ATM, TDM

Exhibit 4-AT: Passive Optical Network (PON) FTTP Deployment

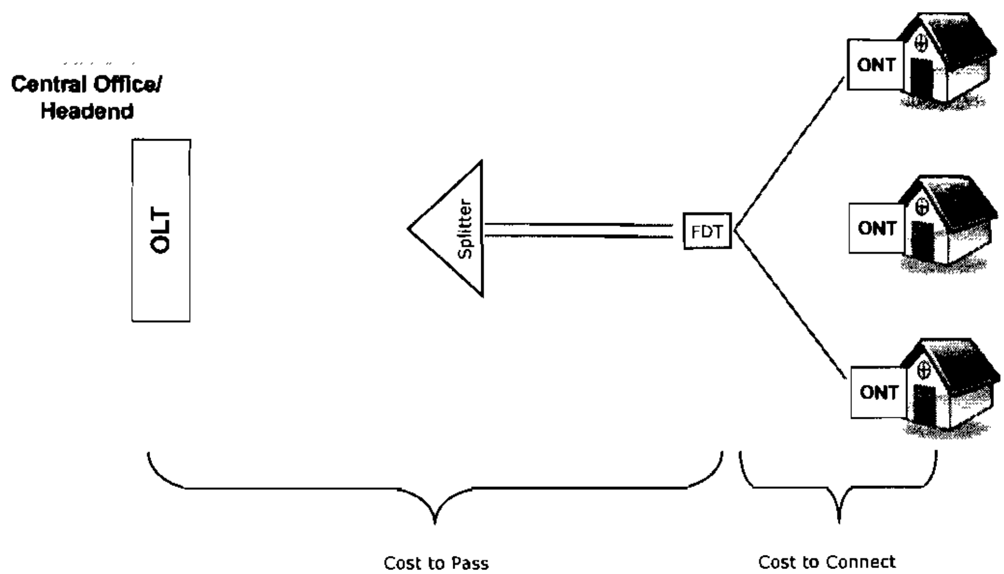

Exhibit 4-AU: Future PON Architectures

	10G GPON	10G EPON
Bandwidth (upstream/downstream)	10/2.5 Gbps or 10/10 Gbps shared	10/1 Gbps or 10/10 Gbps shared
Positives	Compatible with existing GPON	First completed
Key challenges	10 Gbps upstream not viable for single-family units	10 Gbps upstream not viable for single-family homes; 1 Gbps upstream too little bandwidth

FTTP economics

To build FTTP to deliver broadband to the 7 million housing units that are classified as unserved (at a broadband definition of 4 Mbps download and 1 Mbps upload) would lead to an investment gap of $62.1 billion.

The initial capital expenditure averages out to be slightly more than $5,000 per premises. This initial capex value comprises two pieces: the cost to pass a premises and the cost to connect a premises. (These costs are detailed in Exhibit 4-AV.)

The cost to connect a premises is the smaller of the two charges, typically averaging about $650-$750/premises.[113] The cost to connect is entirely success-driven and consists of the installation of the fiber drop and equipment at the customer premises. Making up the bulk of the $5,000 initial capex cost of a FTTP deployment is the cost to pass a premises; this is the cost to build the fiber network distributed over the premises capable of being serviced by the network. Cost-to-pass is typically spoken of in terms of all premises passed by a FTTP deployment, but the more meaningful number is cost-to-pass per subscriber, which takes into account penetration rate. With fiber installation costs ranging between $10,000 and $150,000 per mile, depending on a variety of factors including deployment methodology, terrain and labor factors,[114] the cost to pass is highly sensitive to penetration rate and household density.

Using several data points provided by existing FTTP providers, we are able to establish the following empirical relationship between the cost-to-pass for a FTTP deployment and household density, using standard curve-fitting techniques[115] (see Exhibit 4-AW):

Cost per home passed = $\$701.59 * e^{(8.19/\text{Household density})}$

where Household density is in homes per square mile.

As one can see, the unserved segment starts to intersect the cost-to-pass curve just as the curve starts to steepen significantly. At about 10 households per square mile, the cost-per-premises passed is slightly less than $1,600. Halving the density to five housing units per square mile more than doubles the cost-to-pass, to more than $3,600. At this level, factoring in average broadband penetration of roughly 65% and including the cost to connect each premises yields a cost-per-subscriber in excess of $6,000. Due to the low densities of the unserved segment and given the current expectation of bandwidth demand over the coming years, even with an optimistic scenario for increasing broadband adoption, FTTP may be prohibitively expensive when alternative technologies can also meet bandwidth demands.

The final category of costs is one where FTTP holds a significant advantage: the cost-to-serve. By extending fiber all the way from the serving office or headend to the customer premises, an FTTP network eliminates the need for any active components in the outside plant. This can reduce ongoing maintenance and support expenditures by as much as 80% relative to an HFC plant.[116] However, on a monthly basis for a typical scale network deployment, this savings amounts to just a few dollars per subscriber, and as such is generally insufficient to offset the initial capital expenditure burden.

Exhibit 4-AV:
Breakout of FTTP Gap

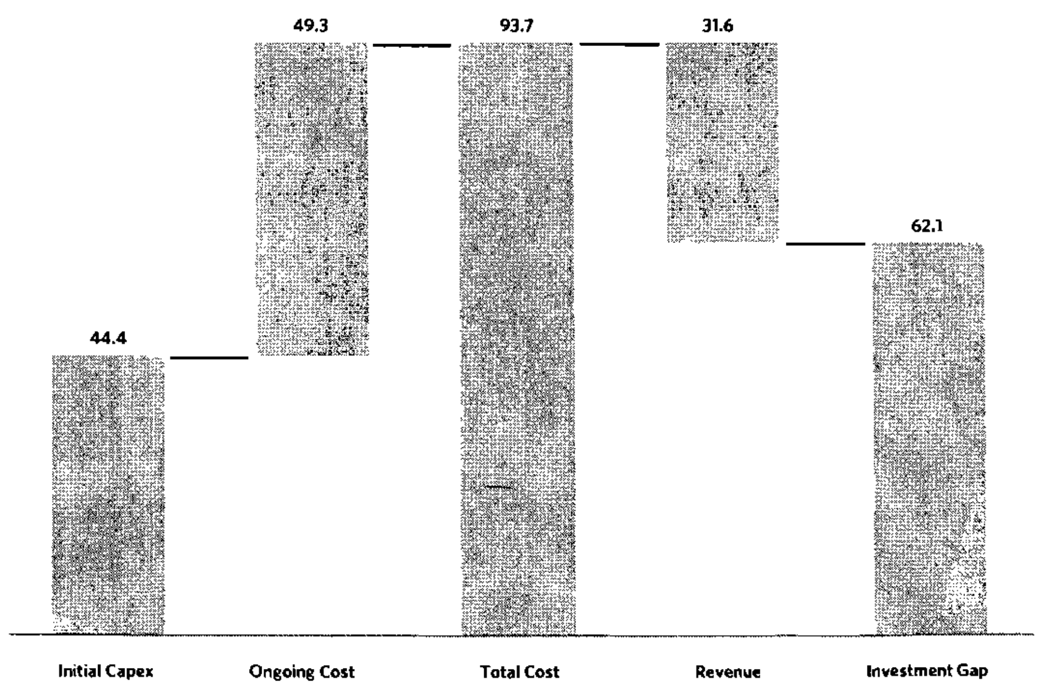

FTTP Deployment

The cost information above can be displayed in a simple financial model that can be used to easily estimate the viability of a FTTP deployment in addition to the model that calculates the cost of the investment gap across the country. See Exhibit 4-AX.

First, consider cost per home passed. In this example, we use $850, a value that would cover roughly 80% of the United States. Factoring in a 40% penetration rate, a value taken from the high end of Verizon's publicly stated 2010 target rate for its competitive deployments,[117] we get a $2,125 cost-to-pass per subscriber. Adding in the cost-to-connect, inflated to account for churn and equipment replacement over the life of the network, we get a rough estimate of $3,225 total investment per subscriber. At this level, an operator could succeed with a monthly EBITDA of

Exhibit 4-AW: Cost to Pass with FTTP by Density of Homes[118]

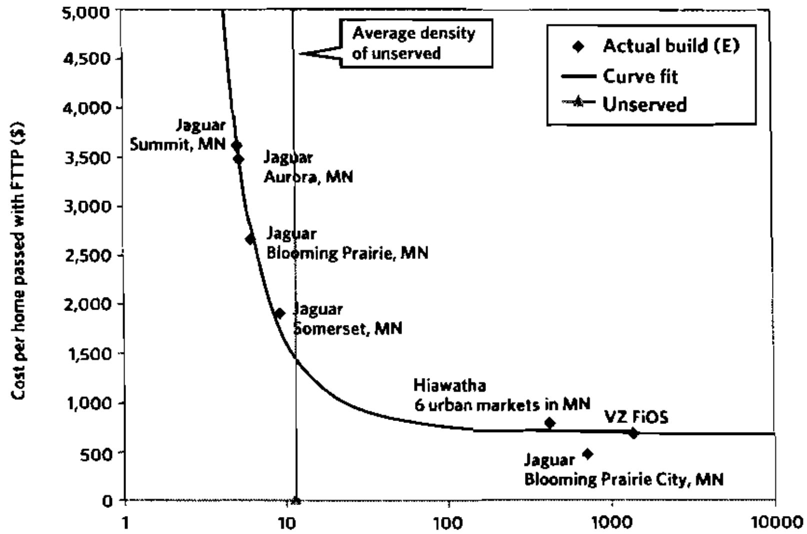

(USD, present value)

Household density (HH per mi²)

Curve fit to provide cost for any density

$R^2 = 0.992$

R^2 adjusted = 0.990

Exhibit 4-AX: Simple Financial Model to Calculate Breakeven EBITDA for FTTP

Cost per home passed	$850	
Take rate	40%	
Plant cost per sub	$2,125	
Cost to connect/maintain	$1100	$650 initial cost to connect; $450 CPE/churn replacement
Total capex per sub	$3,225	
Profit for NPV = 0 (over 20 years)	$3,225	Low profit target
Taxes (@ 20% effective rate)	$806	
PV of EBITDA required	$4,031	Ignores potentially significant interest payments and impact on required EBITDA
EBITDA per month (@ 10% WACC over 20 years)	$42.50	Assumes customer lifetime is 5 years

Key questions:
- How will EBITDA required for breakeven change as density-driven costs change?
- Is that EBITDA target reasonable?

$42.50/subscriber, a value that is roughly in line with estimates of margins for some of the largest providers in the country.

Next, we calculate the cost to deploy FTTP in each county in the country using the curve fit calculated in Exhibit 4-AW. Applying that cost to the financial model laid out in Exhibit 4-AX, one can calculate the EBITDA required for FTTP to break even in each county; the results are shown in Exhibit 4-AY. Note that a successful FTTP entrant would need to have roughly $38 in monthly EBITDA from each customer at the assumed 40% take rate to provide returns to capital in the denser half of the country.

It is important to note that for an incumbent, much of the revenue associated with a FTTP deployment cannibalizes its existing revenue. As such, an incumbent telco would only want to factor in the incremental revenue offered by a FTTP deployment, namely additional data revenue and video revenue. This has the effect of significantly reducing the viability of FTTP deployments currently for many incumbent providers.

Due largely to this cost structure, there have been few large incumbent providers overbuilding their existing footprints with FTTP. To date, the bulk of FTTP deployments have been driven by a single RBOC, Verizon, which has deployed FTTP in the denser, suburban and urban areas in its footprint, and by Tier 3 ILECs, CLECs, municipalities and other small providers. These providers have deployed FTTP in areas that are less densely populated than those of Verizon, but they have been able to largely replicate the RBOCs' cost structure by achieving an average penetration rate that is nearly double that of the RBOC (54% vs. 30 %).[119]

3,000 - 5,000 foot DSL

Despite providing faster broadband speeds than 12 kft DSL and being capable of delivering video services, DSL over loops of 3,000 (3 kft) feet or 5,000 (5 kft) feet has a higher investment gap when providing broadband services in low-density unserved areas. DSL over 3-5 kft loops delivers broadband speeds well in

Exhibit 4-AY: Esitmated Monthly EBITDA Required to Break Even on an FTTP Build Across the Country[120]

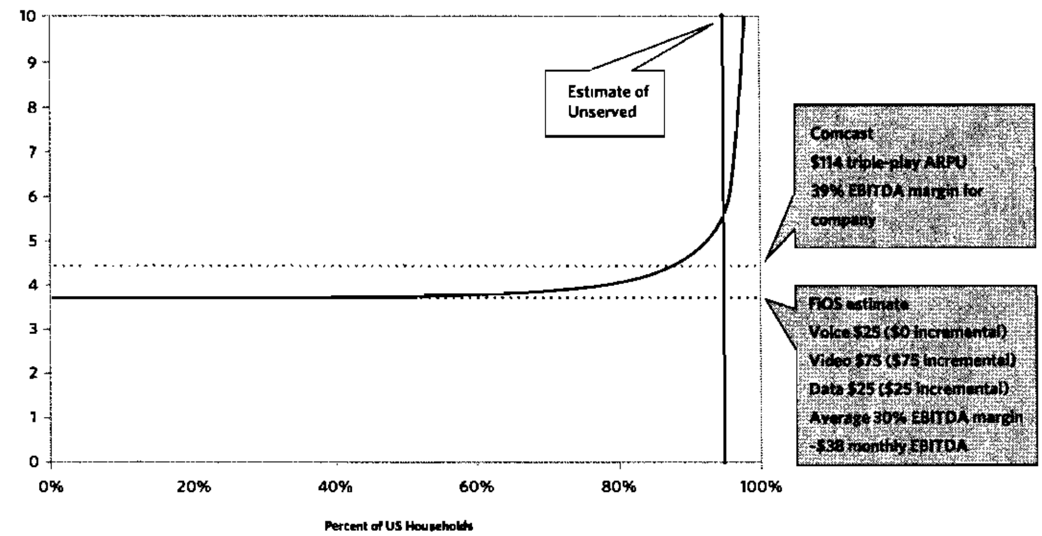

Exhibit 4-AZ: Data Sources for FTTP Modeling

Item	Source
Optical light terminal (OLT)	Calix protective order filing
Fiber distribution hub (FDH)	FTTH Council
optical splitter	FTTH Council
Fiber drop terminal (FDT)	FTTH Council
Optical network terminal (ONT)	FTTH Council, Calix protective order filing
fiber optic cabling	FTTH Council
aerial placement	FTTH Council
buried placement	FTTH Council
operating/maintenance expenses	Hiawatha Broadband protective order

excess of the 4 Mbps downstream and 1 Mbps upstream target. However, due to the cost of driving fiber an additional 7,000 to 9,000 feet closer to the end user, 3 kft DSL and 5 kft DSL are more costly solutions than 12 kft DSL and, thus, have higher investment gaps than 12 kft DSL in all unserved markets.

Capabilities

DSL over loops of 3 kft or 5 kft typically uses VDSL2 technology, which was first standardized in 2006 and uses frequencies up to 30 MHz. While there may be some VDSL technology still being used today, many operators are replacing it with VDSL2. Therefore, we will examine the capabilities of VDSL2 technology at 3 kft and 5 kft.

VDSL2 can provide 35 Mbps downstream and 6 Mbps upstream over 3 kft loops, and it can provide 20 Mbps downstream and 2 Mbps upstream over 5 kft loops. As VDSL2 over 24 AWG wire provides rates well above 4 Mbps downstream and 1 Mbps upstream, the technology meets the speed requirements for broadband service. Exhibits 4-BA and 4-BB illustrate how loop length affects speed for VDSL2. Of course, speeds realized in the field are heavily dependent on plant quality, so

Exhibit 4-BA: Downstream Speed of a Single VDSL2 Line at Various Loop Lengths[121]

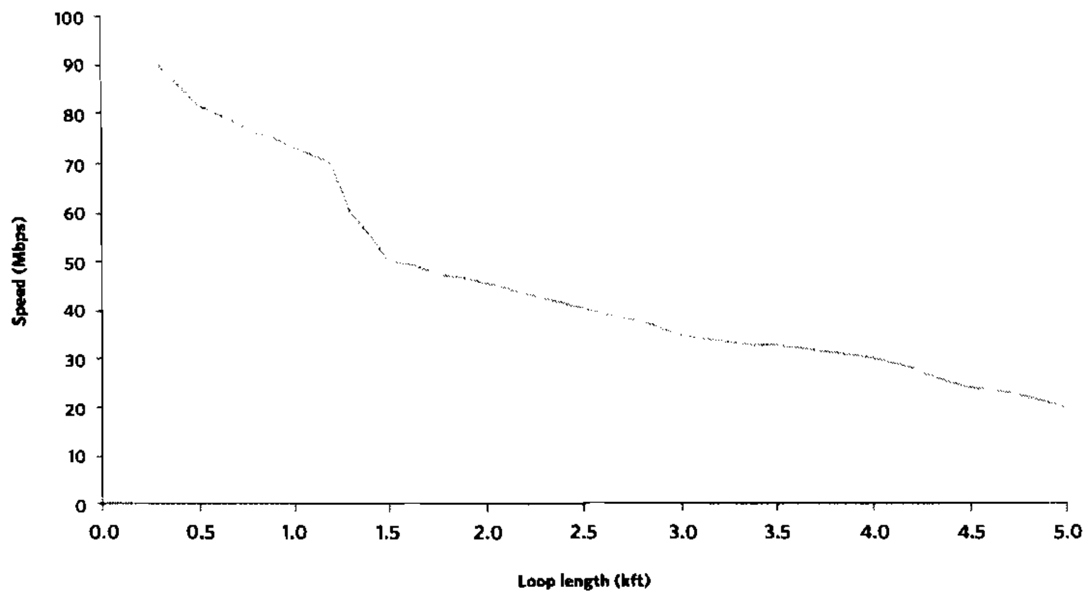

Exhibit 4-BB: Upstream Speed of a Single VDSL2 Line at Various Loop Lengths[122]

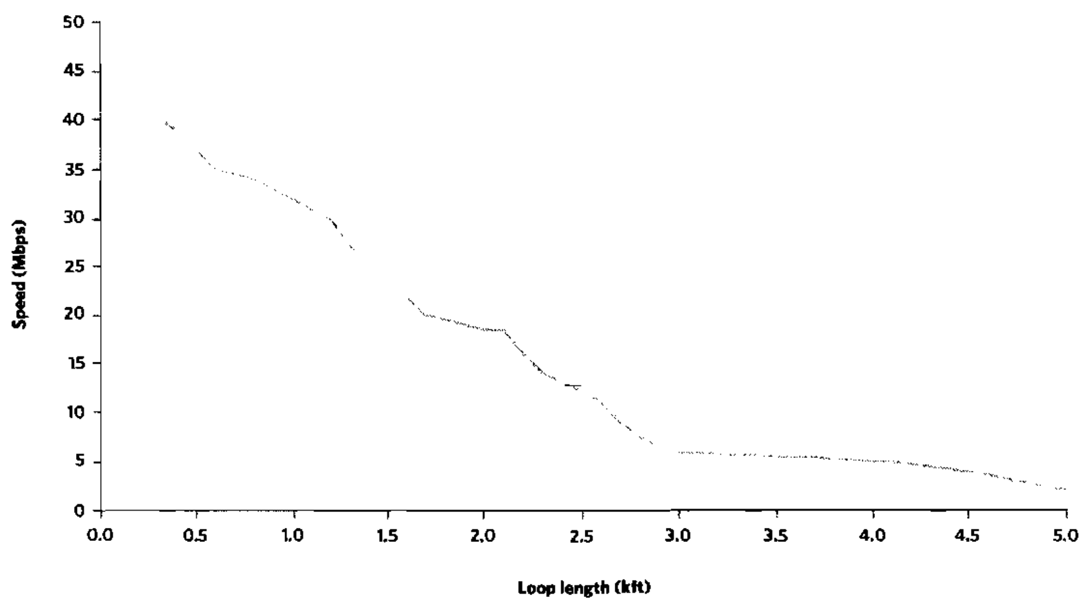

any degradation in the copper plant will lead to lower speeds for a given loop length.

In this case, 24 AWG wire is assumed with no bridged taps. Performance with 22 AWG wire, which is often used in rural areas, would yield higher bitrates, while use of 26 AWG wire would yield lower rates.

For VDSL2, performance can be improved through vectoring, bonding or a combination of the two. Vectoring, or Dynamic Spectrum Management level 3 (DSM-3), has shown improved performance in lab tests by canceling most of the crosstalk between VDSL2 lines sharing the same binder and is currently being tested in the field. The bonding of loops, assuming there are two copper pairs available, would enable the doubling of the speed achieved to the end-user. A combination of vectoring and bonding could produce downstream speeds over 300 Mbps if lab and field tests prove successful. Exhibits 4-BC and 4-BD illustrate the performance of bonded and vectored VDSL2.

Operators who have shortened loops from 12 kft to 3-5 kft and currently use VDSL2 technology have seen DSL technology offer faster speeds in the past decade.[123] Current and future

Exhibit 4-BC: Downstream Speed of VDSL2 Variants[124]

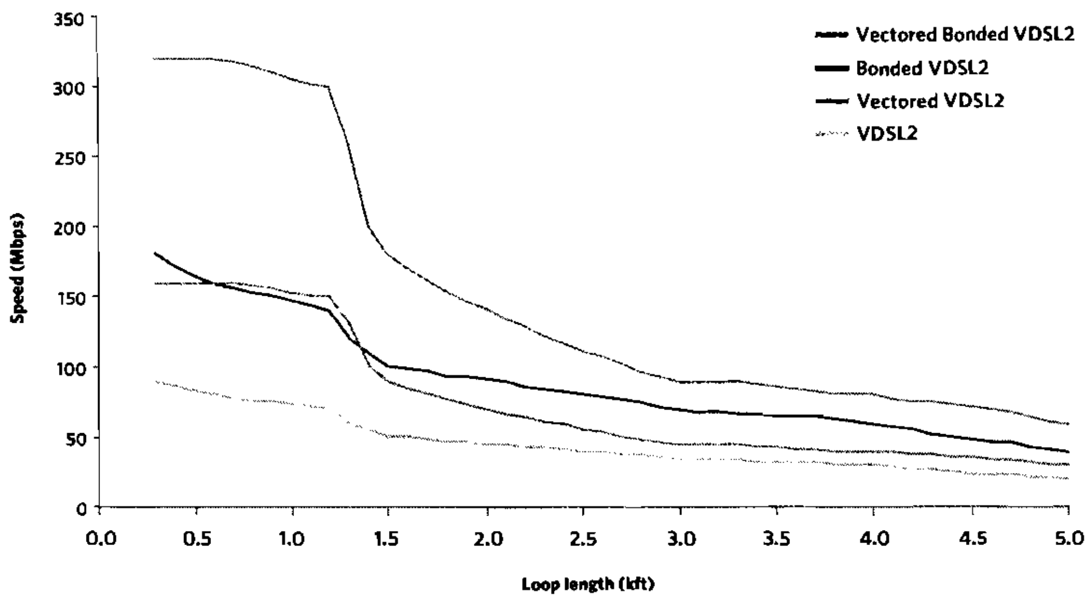

Exhibit 5-BD: Upstream Speed of VDSL2 Variants[125]

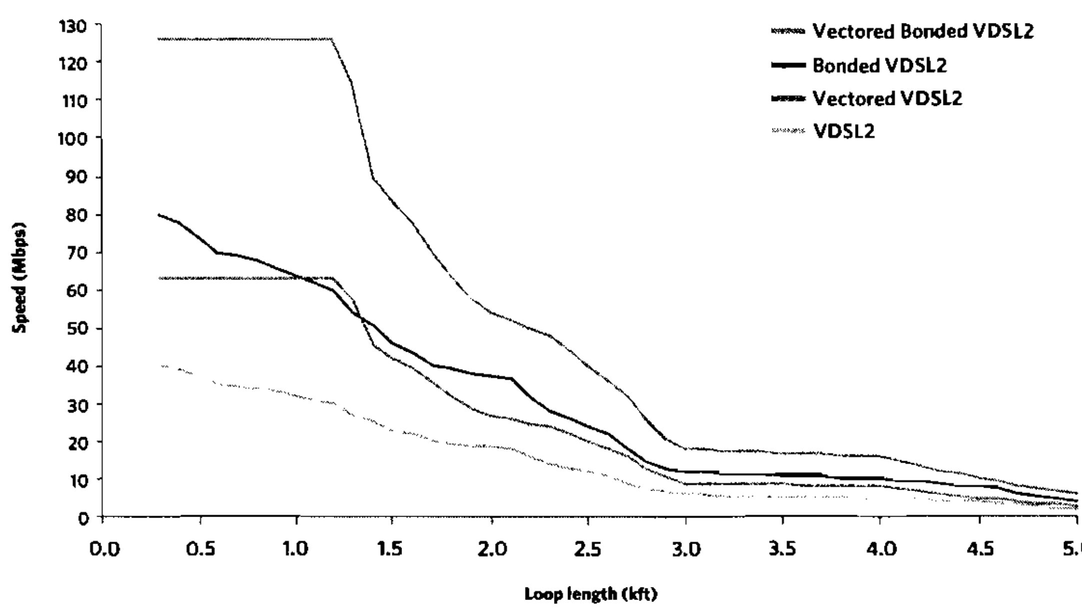

technology improvements, such as the three levels of DSM, are likely to continue to improve speeds as well as the stability of the service provided. Further development of and investment in these improvements, along with bonding, are likely due to DSL's prevalence worldwide.

We model the VDSL2 access network in a similar fashion to the ADSL2+ network described (see above for details). In essence, we assume VDSL2 DSLAMs are connected to central office and other middle- and second-mile aggregation points with fiber-optic-based Ethernet technology providing backhaul capacities that are more than sufficient to meet the end-user requirement. Costs associated with loop conditioning are included when appropriate.

Economics

Like those of the 12 kft DSL network, the economics of the 3 kft DSL and 5 kft DSL networks depend on revenues, operating costs and capital expenditure. Using granular cost data from DSL operators, the model calculates the investment gap to deploy 3 kft DSL to unserved markets as $52.7 billion and the investment gap to deploy 5 kft DSL to unserved markets as $39.2 billion. The total gaps for 3 kft and 5 kft DSL are more than twice as costly as the respective number to deploy 12 kft DSL to the unserved, despite 3-5 kft DSL earning nearly 3x the revenue of 12 kft DSL because their ARPUs include video as well as data. The cost differential is mainly driven by the high cost of driving fiber closer to the end user, less so by the higher cost of VDSL2 technology versus ADSL2+ technology. The following waterfall charts show the breakout among initial capital expenditure, ongoing costs and revenue. See Exhibits 4-BE and 4-BF.

Initial Capex

Initial capital expenditures include material costs and installation for the following: telco modem, NID, protection, aerial or buried copper drop, DSLAM, cabinet, VDSL2 line card, allocated aggregation cost, fiber cable up to 3 kft or 5 kft from the end-user (respectively), feeder distribution interface and drop terminal/building terminal, as well as the engineering costs for planning the network and the conditioning required on loops (i.e., the removal of load coils and bridged taps).

Ongoing Costs

Ongoing costs include replacement capital expenditure required to replace network components at the end of their useful lives, network administration, network operations center support, service provisioning, field support, marketing and SG&A.

Revenues

Revenues are calculated by taking the ARPU—which varies according to the level of broadband service/speed provided as well as whether the bundle of services provided includes voice, data and video—and multiplying it by the average number of users. For 3 kft and 5 kft DSL, data and video ARPUs are used as the incremental services to voice, which is assumed present due to the fact that DSL technology utilizes the twisted pair of copper wires originally installed and used for POTS. VDSL2's higher speeds at 3 kft and 5 kft could support both video and data, although not all real-world operators of VDSL2 choose to offer both services today. The addition of video revenue is not enough to compensate for the incremental investment required to drive fiber within 3 kft and 5 kft of the end user for the unserved.

Exhibit 4-BE: Breakout of 3,000-Foot DSL Gap

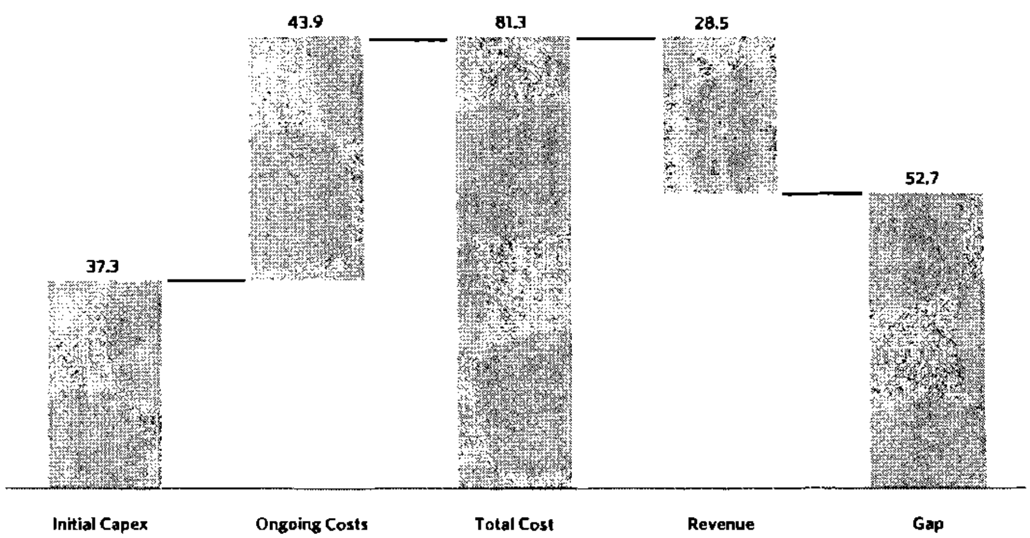

Material and labor costs for 3 kft and 5 kft DSL are the same as for 12 kft DSL except for VDSL2 line cards, which are sourced from a Qwest filing under Protective Order.

15,000 foot DSL

DSL over loops of 15,000 feet (15 kft) is a very cost-effective solution for providing Internet access in low-density areas but fails to meet the Broadband Availability Target.

Capabilities

DSL over 15 kft loops typically uses ADSL2/ADSL2+ technology. ADSL2+ over 24 AWG wire provides rates of 2.5 Mbps downstream and 600 kbps upstream; therefore, the technology does not meet the speed requirements for broadband service under the Broadband Availability Target. Refer to Exhibit 4-AH in the 12 kft DSL section for a further understanding of how downstream speed varies with loop-length distance.

Hybrid Fiber-Coax Networks

The focus in this section will be on high-speed data connectivity provided by hybrid-fiber-coax (HFC), or cable, networks. We'll look first at the capabilities of HFC networks, then at the economics of these services.

Our analysis indicates that the capabilities of HFC networks far exceed end-user speed and network capacity requirements, as shown above and in the National Broadband Plan. Therefore, by

Exhibit 4-BF: Breakout of 5,000-Foot DSL Gap

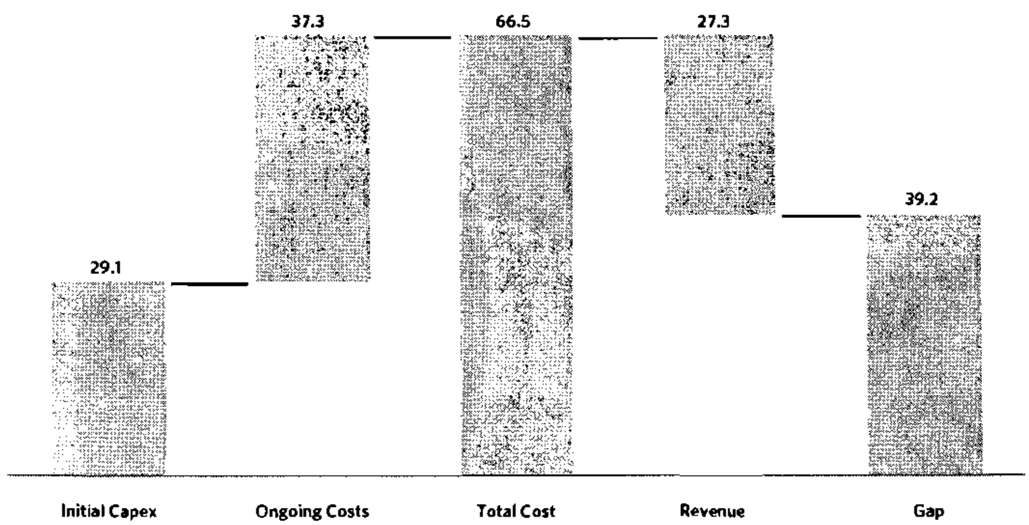

Exhibit 4-BG: Breakout of 15,000-Foot DSL Gap

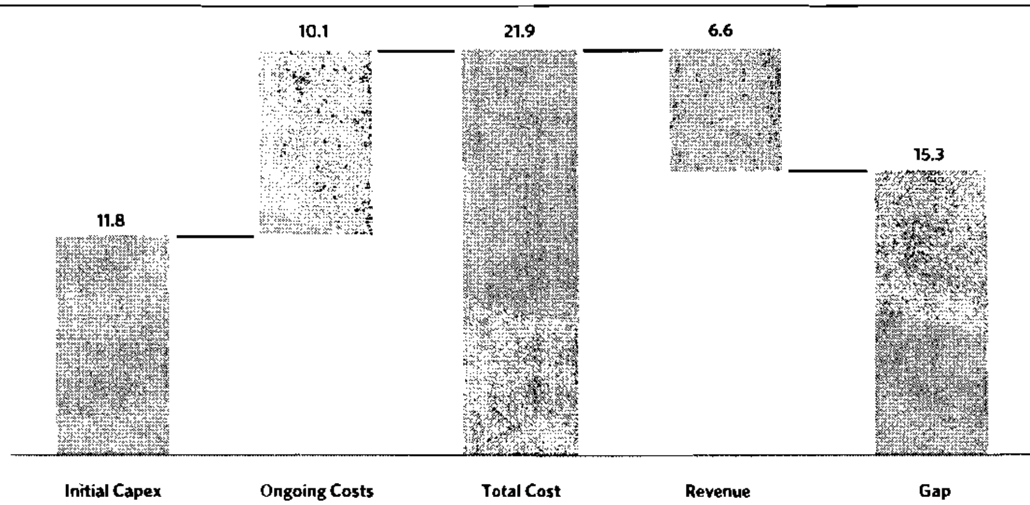

definition, homes within the HFC footprint are considered served. However, the investment gap to deploy HFC networks in unserved areas is larger than that of DSL or fixed wireless as noted above.

The near-ubiquity of HFC networks that can provide high-speed broadband access is a tremendous asset that puts the United States in a unique position among other countries. HFC networks were initially designed to deliver one-way video, but have evolved over time to allow two-way transmission of data and voice in addition to video. Today, cable systems pass roughly 90% of U.S. households with high-speed data services; in addition, more than 90% of homes are passed by cable plant, with 50% of those homes taking at least basic cable video service, thereby amounting to 63 million subscribers.[126] Some 52% of broadband subscribers in the United States subscribe to cable-based service, the second highest rate among OECD countries.[127]

History

When cable systems were initially constructed, the industry was highly fragmented, with many small firms operating networks in local markets. Today, there is very little overlap in cable networks because, in most markets, cable operators received exclusive rights to operate in their geography in the form of a franchise agreement granted by local franchising authorities. It is important to note that cable companies have not been subjected to the same network-sharing or carrier-of-last-resort obligations as the telephone companies; however, cable companies do not receive Universal Service Fund (USF) monies to offset the costs of constructing and maintaining their networks. Maintaining one network per geographic area greatly reduced the network cost-per- subscriber, which, along with having monopoly or near-monopoly control over the video market, has allowed these networks to be successful in the face of large up-front capex requirements.

Exhibit 4-BH:
Breakout of Cable Coverage– Share of Homes Passed by Cable Companies

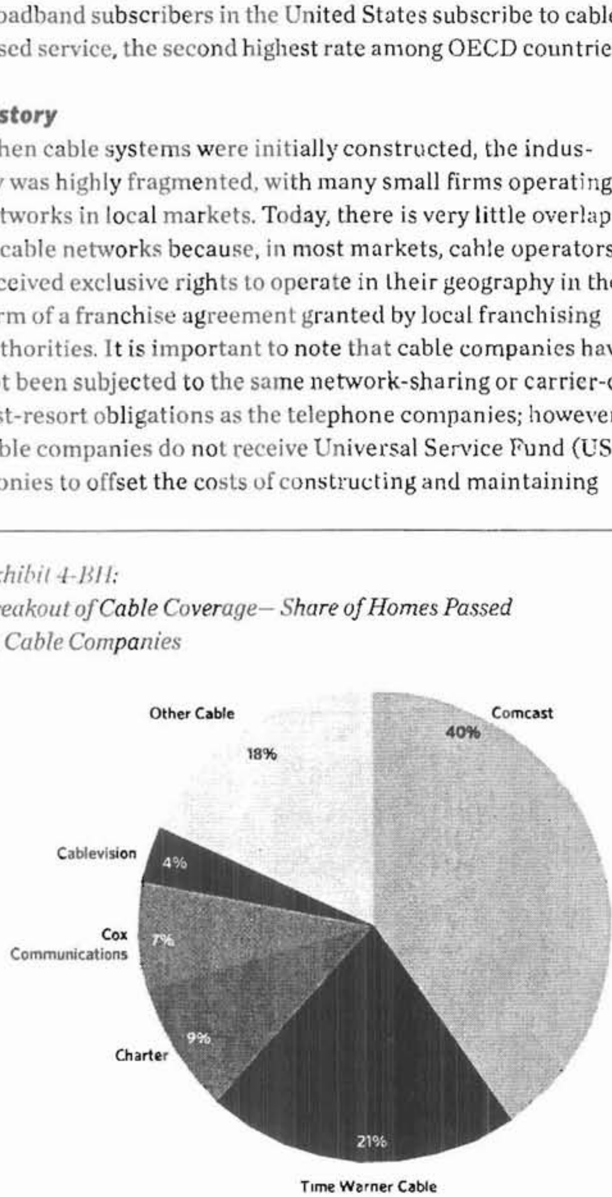

Numbers do not sum to 100% due to rounding.

Due to the complementary nature of footprints and scale advantages in content acquisition, the cable industry has experienced significant consolidation over the years. Today, there are almost 1,200 cable system operators but, as shown in Exhibit 4-BH, the top five companies pass 82% of homes passed by cable video service.[128]

Cable MSOs have spent $161 billon from 1996-2009 on capital expenditures; in part, this was used to enable broadband capabilities.[129] Cable systems were originally constructed to provide one-way video signals, so customers initially could not send information back through the network. In the early deployment of cable (1950s-1970s), the networks were known as CATV (Community Antenna Television) and were built to provide TV and radio services. The network was designed to support all-analog, one-way transmissions from the community satellite antennas (cable headends) to end-user televisions over coaxial cable.

In the 1990s with the advent of the Internet and passage of the 1996 Telecommunications Act, cable companies began upgrading their networks to provide the two-way transmission capabilities required for Internet data traffic and telephony in addition to TV/radio signals. The network needed to be reengineered to handle two-way transmissions of digital communication signals and upgraded to handle higher capacity demands. The original "tree and branch" architecture of cable systems was ideal for transmitting TV signals from the head-end to the home television. However, video transmission over coaxial cable was still susceptible to noise and interference and required amplifiers, line extenders and other active electronics to ensure that the signal would reach end-user TV sets with acceptable quality. Unfortunately, these active electronics a) were not capable of passing signals in the upstream direction and b) were often not spaced properly within the cable plant for upstream transmission. As a result cable companies invested in HFC upgrades throughout the 1990s to overcome these problems. Such upgrades were seen as attractive since millions of homes were already "wired" with high capacity coaxial cable and the revenue potential of triple play services created a compelling business case. Exhibit 4-BI illustrates some examples of the infrastructure upgrades required for HFC networks.

Steps to upgrade cable networks for broadband:

- Invest in fiber optic cable and optic/electronics to replace and upgrade coaxial cable for capacity purposes

- Replace and redesign headend equipment, line transmission equipment, set top boxes to allow for two-way data transmission, and add DOCSIS modems
- Deploy telephone switching equipment and interconnection facilities to provide VoIP services
- Develop the technology and equipment necessary for more sophisticated network management and control systems
- Implement the back-office, billing and customer service platforms necessary to provide the standard triple play services common among cable operators today

Capabilities

Cable companies coupled their investments in two-way upgrades with a standardization effort. Cable-based broadband relies on Data Over Cable Service Interface Specification (DOSCIS). The first release of DOCSIS was in 1997, with DOCSIS 2.0 released in 2001 and the third-generation standard (DOCSIS 3.0) now being deployed widely. DOCSIS 2.0, currently the most widely deployed, provides up to 36 Mbps of downstream bandwidth and up to 20 Mbps upstream, while DOCSIS 3.0 provides up to 152 Mbps of downstream bandwidth and up to 108 Mbps of upstream (with four bonded channels).[130]

As noted above, cable systems provide shared bandwidth in the last mile, with multiple homes sharing a fixed amount of bandwidth at a single node. Ultimately, bandwidth-per-customer is driven both by the number of customers (and their usage) per node and the total bandwidth available per node. Given typical busy-hour usage rates (see Network Dimensioning section), users on a DOCSIS 2.0 system can receive up to 10 Mbps;[131] under DOCSIS 3.0, that number will increase substantially, to 50 Mbps.[132] Actual figures, however, depend on a large number of variables, including not only the DOCSIS specification, but also spectrum allocation and use and the number of homes per node.

Impact of cable-system spectrum

Spectrum in cable plants, as in over-the-air broadcasting, is a measure of how much "real estate" is devoted to transmitting signals. Most two-way cable plants use 450 MHz or more of spectrum, with many having been upgraded to provide 750 MHz or more. Each analog television channel requires 6 MHz of spectrum. Exhibit 4-BJ shows the spectrum allocation for a typical 750 MHz, DOCSIS 2.0 deployment.

Note that all upstream communications take place in low-frequency spectrum, below 52 MHz. FCC rules requiring that broadcast Channel 2 be carried on Channel 2 of the analog spectrum (54 - 60 MHz) established the low end of downstream spectrum.[133] Cable companies' outside plant equipment is tuned for this: band-pass filters allow upstream traffic only below 52 MHz. In addition, band-pass filters in consumer electronics are tuned to block potentially large amplitude upstream signals only below 52 MHz.

The 52-MHz upper bound on upstream spectrum places limits on upstream bandwidth. First, because it would require

Exhibit 4-BI: Upgrades to Enable Broadband Services

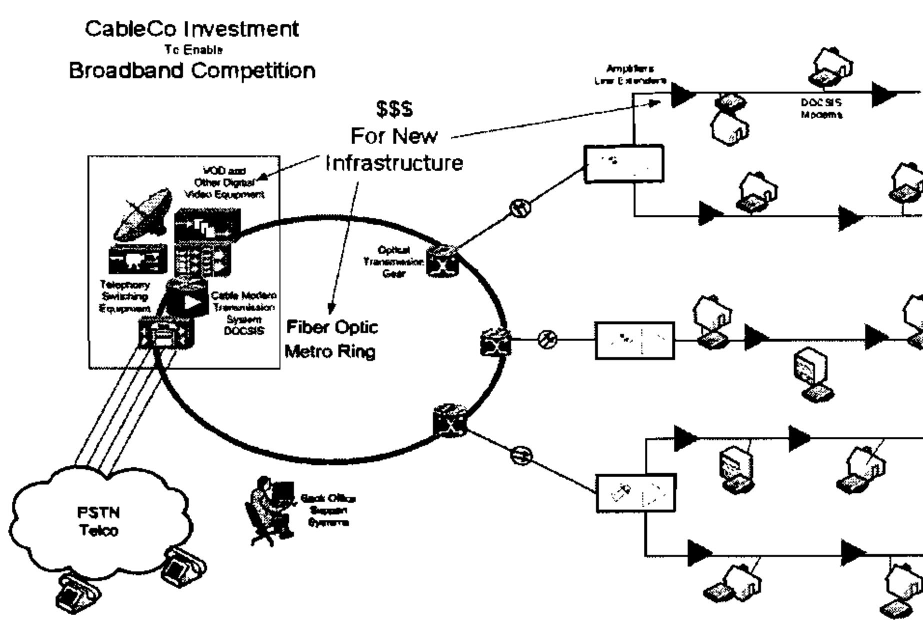

changes to cable plant and consumer electronics, adding spectrum for upstream use above the 52 MHz would be difficult and costly. In addition, interference at low frequencies (e.g., from motor noise, ham and CB radio, walkie-talkies) could reduce usable upstream spectrum significantly.[134] While DOCSIS 3.0 allows for the bonding of multiple channels to increase upstream capacity, these other spectrum issues will likely provide real-world limits to upstream capacity.

Downstream bandwidth faces fewer constraints; cable companies can devote higher-frequency 6 MHz channels to downstream capacity. In addition, DOCSIS 3.0 allows carriers to devote four or even eight channels to downstream data communications.

Cable companies use Quadrature Amplitude Modulation ("QAM") to increase the bandwidth transmitted over a given amount of spectrum (the Mbps-per-MHz), with typical deployments featuring 16, 64 or 256 QAM. In typical DOCSIS 2.0 deployments, the downstream direction is 64 or 256 QAM and the upstream is 16 QAM. As an example, consider a typical DOCSIS 2.0 deployment with one 6 MHz downstream channel at 64 QAM which delivers approximately 36 Mbps.

Cable companies can create additional capacity for downstream bandwidth (or for additional broadcast video channels, or other services like video-on-demand) through a number of means. The most obvious may be to increase the frequency of the cable plant, but this requires extensive upgrades in outside plant and is often very expensive.

There are a number of less expensive options available. As discussed above, going from DOCSIS 2.0 to DOCSIS 3.0 allows the cable system to devote more frequency, assuming it can be made available, to data while keeping the plant total unchanged. Cablevision estimated the cost of its DOCSIS 3.0 rollout at about $70 per home passed (there may be additional success-based expense, e.g., CPE). Scale economies may bring that number 10-20% lower for larger MSOs.[135]

Another option is Switched Digital Video (SDV). In the current HFC architecture, all video channels are sent to all subscribers with filtering of channels for different subscription services made by the set-top box. SDV transmits only those channels to a given node when those channels are in use by a subscriber. This means that the majority of channels are not transmitted most of the time, thereby using fewer channels in aggregate. SDV is therefore a relatively inexpensive technique to reclaim on the HFC network bandwidth to be used for other purposes. Cisco Systems estimates the cost of SDV at $12-$16 per home passed.[136] A number of MSOs are moving forward with SDV,[137] although concerns exist for third party providers of DVRs like TiVo.[138]

Another approach is analog reclamation. In analog reclamation, often termed "going all digital," cable companies move away from transmitting analog signals entirely. A single analog channel takes up 6 MHz (the equivalent of more than 30 Mbps as noted above); the same spectrum (or bandwidth) can carry 10 digital standard-definition channels or three high-definition channels. Analog reclamation can therefore "add" a substantial number of channels to a typical system. For example, by

Exhibit 4-BJ: Spectrum Allocation in Cable Plant