

Broadband Access Networks and the Emergence of Voice over IP (VoIP): an Economic Analysis for Cable and ADSL

Daniel Fryxell¹, Steven Lanning², Marvin Sirbu³
djf@andrew.cmu.edu, slanning@lucent.com, sirbu@cmu.edu

ABSTRACT

This paper analyzes the cost structure of broadband wireline access networks based on the Internet Protocol (IP) offering integrated Internet access and voice services for households and small businesses. The first network follows a telco-based architecture. The architecture for the cable plant is similar to recent designs for Digital Loop Carrier (DLC) systems with fiber to the serving area and copper pairs to the customer premise. However, the PSTN infrastructure is completely replaced by an IP network that is extended to the home using ADSL technology. The second network follows a cable-based architecture. The cable plant follows recent trends in the design of Hybrid Fiber Coaxial (HFC) networks. The IP network is extended to the home using a cable modem platform that follows the DOCSIS reference model from Cable Labs. The IP architecture, which supports Internet access and voice, is basically the same for both networks. Voice service is based on VoIP with IP soft switches at the central office/headend and Multi Terminal Adapters (MTA) at customer premises.

Our model calculates forward looking (TELRIC) capital costs for constructing a broadband access network and provides a systematic way to compare different broadband architectures and technologies. Drawing on a rich data set containing detailed demographic, geographic and geologic information on telecommunications serving areas throughout the U.S., it allows detailed comparisons by region and density zone.

The results show that a broadband network based on a Hybrid Fiber/Coaxial architecture has lower initial first capital costs than a comparable network based on ADSL over copper pairs over all density levels. The difference is due to the higher cost of DSL access multiplexers as compared to HFC optical nodes and Cable Modem Termination Systems. The cost of the HFC architecture is, however, more sensitive to voice and data traffic assumptions than is the ADSL architecture. Finally, the HFC architecture cost advantage is even larger, if part of the plant cost can be attributed to video delivery.

Keywords: ADSL, HFC, Cable Modem, VoIP, Broadband, Access Network.

1 Department of Engineering & Public Policy, Carnegie Mellon University

2 Bell Laboratories, Lucent Technologies

3 Department of Engineering & Public Policy and Information Networking Institute, Carnegie Mellon University

1 Introduction

One of the most critical aspects of the communications networks that support the Internet is the access infrastructure connecting businesses and households to regional and national backbones. Although large corporations can afford sophisticated high-capacity access links, the existing access solutions for residential customers and small businesses rely mostly on the public switched telephone network (PSTN). This network, which was built to carry voice, is not adequate for data communications.

Several always-on broadband access solutions have been recently developed, including Digital Subscriber Line (xDSL), cable modems, fixed wireless and satellite. Cable modem technology was the first broadband access solution deployed in large scale for the residential market and it gained an early lead that has persisted until today. In an attempt to catch up with cable operators in the broadband market, telephone companies started to aggressively roll out ADSL technology, which offers broadband data rates over the copper pairs of legacy telephone networks. Despite the promise of future competitive wireless solutions, currently these two wireline technologies enjoy a considerable advance in terms of deployment and are the only alternatives widely available in the market.

Another technological trend that has been spurred by the Internet is the convergence of voice and data. This idea started to take shape from the development of large data networks and the increasing weight of data traffic as a fraction of total traffic. Network operators soon realized that, instead of keeping two separate networks in parallel, they could obtain considerable savings in equipment and operations costs by integrating voice as just one more type of data traffic. However, carrying voice over a data network requires, not only the development of directory services and servers to manage calls, but also strict Quality of Service (QoS) in terms of packet throughput and delay. That can not be assured by a best-effort IP network, as today's Internet. For this reason, this convergence trend had been limited until very recently to transport in backbones of service providers and private networks of large firms. The investment required to go the next step, and provide the full functionality for voice calls, can only be justified by a large market for VoIP, where calls are initiated as IP at the customer premise, instead of having central office gateways making the analog to IP conversion. The critical mass that is required for this conversion may come from the opportunities opened by the extension of data networks to households and small businesses using always-on broadband platforms. Although the long-term convergence of data and voice seems certain, the ability and economic viability of offering voice over IP applications over the emerging broadband access networks may be critical to shape the pace of this convergence.

In the balance of this paper, we present an economic analysis of two architectures for an IP-based local access network, providing both voice and data services. We compare the cost of a network based on ADSL technology with the cost of an alternative architecture based on hybrid fiber/coax and cable modem technology. Section 2 describes the architectures of the proposed networks. The cost model is presented in section 3 and its results are presented and discussed in section 4.

2 Description of the Networks

2.1 ADSL Access Network

2.1.1 ADSL technology

Asymmetric Digital Subscriber Line (ADSL) is a digital technology for the local loop (the link between the telephone central office and the end customer that is also known as “the last mile”) that delivers broadband data rates over the legacy copper infrastructure used for POTS (Plain Old Telephone Service). This technology extends the usable spectrum in copper loops from the 4KHz used by analog telephone service up to about 1.1 MHz. By taking advantage of this large increase in usable bandwidth, ADSL technology allows the expansion of the existing access capacity by two orders of magnitude and permits telephone operators to offer advanced data services without requiring the expensive installation of new cables.

ADSL is an asymmetric transmission technology that offers a larger transmission capacity in the downstream direction (from the central office to the customer premise) than the upstream direction. This characteristic is a perfect fit for typical Internet access activity, where the download of files and other information constitutes the bulk of traffic. Average upstream data rates are typically no more than 10% of the downstream data rates.⁴ In contrast to Internet access over the telephone system, ADSL is an always-on technology that allows a connection to be permanently active.

As specified in the ADSL PHY recommendations (ANSI T1.413 issue 2 and ITU G.922.1), ADSL delivers downstream data rates up to 6.144 Mbps and upstream data rates up to 640Kbps⁵. However, the maximum data rate that can be achieved over a copper pair depends on the length of the line, the type of wire, the existence of bridge taps, and the level of electromagnetic noise. The highest data rates can only be achieved for 24-gauge loops up to 9,000 feet. The downstream speed is reduced to 1.5 Mbps for loops up to 18,000 feet and longer loops are generally not eligible for ADSL service, though other lower speed xDSL services may be feasible.

As described in the ADSL Forum Reference Model (ADSL Forum, 1997), an ADSL line is terminated at a Remote ADSL Transceiver Unit (ATU-R) on the customer premise side and at a Central Office Transceiver Unit (ATU-Cs) on the network side. In the most popular configuration, these two devices create a bearer channel for transport of ATM cells. On the customer premise, the ATU-R is integrated in a network termination unit, usually known as an ADSL modem, which may be an internal device directly connected to the bus of a PC or a stand-alone external unit with one or more interfaces for customer premise equipment. ADSL modems may also include higher level functions such as switching, bridging, routing and working as a gateway between networks using different protocols. On the central office side, the ATU-C is housed in an access node, commonly known as a Digital Subscriber Line Access Multiplexer (DSLAM). DSLAMs, which may house hundreds of ATU-Cs, statistically concentrate the typically bursty individual streams into high-bandwidth upstream links.

⁴ The rise of peer-to-peer applications, such as Napster may significantly change this pattern. Preliminary data from one DSL provider found a 2.5:1 ratio of downstream to upstream traffic.

⁵ Higher data rates are optional. Most ADSL products currently in the market allow downstream data rates up to 8Mbps and upstream data rates up to 1Mbps.

Current standards for ADSL access architectures are based on an ATM layer running over the ADSL physical layer. In this model, there are ATM virtual circuits connecting the ATU-R's to an aggregation point, which may be located in the access network or at the point of presence of an ISP. The former has been the preferred solution for LECs, either using RFC-1483 bridging over ATM or IP tunneling based on PPP/L2TP, because it eliminates the complexity and cost of having to manage a large number of virtual circuits starting at the premises and going all the way across the access network to terminate at the ISP's routers.

2.1.2 Architecture of the access network up to the local office

The proposed architecture for the ADSL-based access network envisions end-to-end IP technology with ADSL in the local loop to provide complete integration of voice and Internet access.

Outside plant

The outside plant for the ADSL-based network follows current architectures for telephone networks, with copper pairs connecting customer premises to the Central Office (CO) or to Remote Terminals (RT) such as Digital Loop Carrier (DLC) systems. RTs are used in neighborhoods where distance limitations or economic considerations do not make it appropriate to have copper lines all the way from the customer premises to the CO. In this mixed copper/fiber configuration, RTs are the interface between the fiber and copper segments of the local loop. The copper sub-loops are terminated at line cards in the remote terminals, which are then connected to the central office by fiber optics cables. ADSL technology imposes stricter distance limitations than the PSTN, with the higher-speed ADSL configurations requiring shorter copper loops. In this context, the use of remote terminals, which in an ADSL architecture house the transceivers (ATU-Cs) and DSLAMs, acquires a special relevance because it is the only way to serve customer locations that can not be reached using all-copper local loops. DLC equipment capable of supporting ATU-C line cards are referred to as Integrated DLC (IDLC) and are currently being deployed by SBC and other ILECs⁶.




Figure 1. Architecture for the ADSL-based access network.

⁶ See http://www.sbc.com/Technology/data_strategy/project_pronto/

The ADSL architecture is based on the system reference model proposed by the ADSL Forum, as described in (ADSL Forum, 1997). Yet, in contrast to this reference model, where the Public Switched Telephone Network (PSTN) infrastructure for voice is kept in place, here the PSTN is completely replaced by an IP network that offers integrated voice and data services. For this reason, no circuit-switching equipment is used in our architecture. Moreover, there is no need for splitters in the local loop because ADSL does not share lines with POTS.

Customer premise

Customer premises are equipped with a Residential Gateway (RG) that provides the interface between the customer premise equipment (CPE) and the access network. In terms of basic network connectivity, the RG is simply a standalone ADSL modem with bridging functionality that includes an interface to the customer premise network. This interface may be ATM, Ethernet, HomePNA or a combination of these alternatives. Households are more likely to rely on a HomePNA solution because it uses the existent telephone cables and does not require the installation of new wires⁷. Small businesses may prefer a higher-capacity solution as Ethernet, despite the fact that it requires wiring the location with Category 5 cable. In order to offer voice service, the residential gateway includes also a Multimedia Terminal Adapter (MTA) equipped with RJ11 ports for analog telephone service.

To ensure lifeline telephone service during a power outage, reliable power must be provided to the RG which in turn powers conventional handsets. ADSL modems installed today do not need power backup because they are used for Internet access only and are not powered by the carrier. We assume that RGs are locally powered at the premise as is common practice today for ADSL modems. Since the power is local, the power backup is provided by rechargeable batteries. We think that a solution with network powered RGs would have advantages in terms of economies of scale and reliability of a power backup system and labor cost for the installation and maintenance of the RG. Since the current generation of ADSL modems consumes about 40 times the power of an analog voice line, extensive changes to conventional LEC power systems would be required to support network powering of RGs. A network-powered architecture for ADSL may be developed in the future. Yet, due to the difficulty in estimating its cost, we assumed locally powered RGs in the cost model.

The residential gateway is placed outside the home at the demarcation point between the network and the customer premise, where today we can find the Network Interface Device (NIDs) for the PSTN or the Network Interface Unit (NIUs) for cable telephony. Currently, ADSL modems are typically installed inside the premise close to a PC. There are two reasons for placing the RG at the NID location outside the home. First, in order to achieve the reliability required for lifeline telephone service, the RG should be positioned where it can be easily maintained by the network provider, and out of reach of the customer. Moreover, for a universal service--not an Internet service for computer-literate early adopters--it is not realistic to assume that customers are able to install and maintain RGs by themselves. Furthermore, carrier installation can insure the reliability of all power and wiring connections. Second, both the line entering the premise on the network side and the telephone inside wiring are easily accessible from there. It is important to have easy access to the inside wiring because it is used for telephone service and eventually for a HomePNA network. A RG placed elsewhere would require a new cable to be extended from the entry point into the premise to the RG, bypassing the telephone inside wiring and bringing the

⁷ See <http://www.homepna.org/>

ADSL line directly to the RG. Moreover, the MTA and the HomePNA port would have to be connected back to the telephone lines. In contrast, if the RG is placed at the demarcation point, no reconfiguration of customer premise wiring is needed.

The main disadvantage of having the RG at the demarcation point outside the home is that it requires an environmentally hardened case and electronics capable of withstanding extreme temperatures. Also, when the RG is locally powered, its installation may require a power cable to be extended from an interior power outlet, through the wall to the RG.

Note that we are assuming that the RG is owned by the network provider. This follows naturally from the decision to place the RG outside the home under the control of the network provider. We think that a business model where the customer only pays the monthly bill makes more sense for a network that offers lifeline telephone service. Finally, placing the RG outside the home may be more consistent with the FCC's Part 68 rules governing the demarcation between the network and Customer Premises Equipment, which, in the U.S., are detariffed.

Central office

At the interface between the cable plant and the in-office equipment, there is a Main Distribution Frame (MDF), which terminates and provides electrical protection for copper loops, and a Fiber Distribution Frame (FDF), which terminates fiber cables.

Copper loops are connected to ADSL transceivers (ATU-Cs) in DSLAM banks, which are then connected to access servers. These devices aggregate the network links from both field and in-office DSLAMs and provide the ATM functionality to terminate the virtual circuits initiated at RGs. The connections to ISPs are managed by the access server using Layer 2 Tunneling Protocol (L2TP) or an equivalent protocol. Point to Point Protocol (PPP) connections initiated at the customer premise are partially authenticated by the server and directed to the desired ISP using an IP tunnel. This architecture gives the ISP the opportunity to authenticate the user and to assign it an IP address in its address space. In this architecture for the access network, the access server allows the LEC to keep control over the end customers while still providing them with the ability of accessing several ISPs. Moreover, the access server allows the LEC to offer its own level 3 services, such as voice over IP.

The interoffice optical rings are based on packet over Synchronous Optical Network (SONET) technology and use OC-48 circuits. There is one Add Drop Multiplexer (ADM) per each OC-48 circuit to which the office is connected. This device extracts/inserts OC-1 channels from/into the OC-48 circuit.

2.2 Cable Access Network

2.2.1 Cable modem technology

There are currently two communications networks that offer connectivity to homes: the telephone network and the cable television network. While ADSL offers broadband access over telephone lines, cable modems, as the name says, offer broadband access over cable networks. Internet access over cable requires the use of cable modems at the customer premises and a Cable Modem Termination System (CMTS) at the headend. These devices appropriate a portion of the usable RF spectrum in the cable plant to create a pair of digital channels for data transmission.

Currently, cable modems in the US follow the DOCSIS 1.1 recommendations from CableLabs for two-way broadband access over cable (CableLabs, 2000). In DOCSIS 1.1, the downstream channel uses

one of the 6 MHz blocks originally created for TV channels to offer data rates of 30 Mbps using QAM-64 modulation (in low-noise systems, where QAM-256 modulation can be used, this value goes up to 43 Mbps). The upstream channel allocates a narrower block in the low-frequency part of the spectrum that is reserved for upstream communications (between 5 and 42 MHz). Upstream data rates are a function of the upstream channel bandwidth (from 200 KHz to 3.2 MHz) and the modulation scheme used (QPSK or QAM-16). The range of possible data rates for upstream channels goes from 320Kbps to 10 Mbps, but typical values currently being used are in the 1.3Mbps to 5.1 Mbps range.

In contrast to ADSL, where every user has a dedicated access channel, data over cable relies on data channels shared by multiple users using statistical multiplexing. Channel sharing allows an efficient use of the available capacity by multiplexing the bursty data traffic of several cable modems over a high-capacity channel. This implementation extends the savings associated with packet-switching networks to the access part of the network. Depending on the number of subscribers and their traffic pattern, a channel can cover the entire area served by a headend or, at the opposite extreme, there may be one or even multiple data channels per optical node.

The data over cable platform offers an always-on link to the Internet. Given the packet-switching nature of the system, only the actual transmission of data uses network capacity.

2.2.2 Architecture of the access network up to the local office

The cable-based broadband access network uses architectures and technologies developed for the cable television market. Modern cable systems follow Hybrid Fiber Coaxial (HFC) architectures for the cable plant offering a substantially larger bandwidth than the old architectures based on tree-and-branch coaxial systems. The HFC physical infrastructure provides support for an IP-based broadband access network over a cable modem platform for two-way communications over cable.

HFC physical infrastructure

As can be seen in Figure 2, an HFC network has three main sections: fiber backbone, distribution and drops. Fiber optic cables connect a distribution hub to optical nodes located in neighborhoods serving 150-1000 homes. Optical nodes are the interface between the fiber backbone and the distribution. On the distribution side, they have several independent output legs where coaxial cables are connected. Nodes are distribution centers from where systems of coaxial cable extend the network to the customer premises. This coaxial part of the network can further be divided into distribution and drops. The distribution takes the network to the streets just outside customer locations, while drops are cables connecting the distribution to individual customers. The coaxial systems that make up the distribution network have amplifier cascades of typically 3 or 4 units. Amplifier cascades in these coaxial systems are short enough to be entirely operated at high signal levels, similar to the distribution part of old tree and branch architectures, while keeping the signal distortion low (see Ciciora, 1999). Splitters and directional couplers are passive elements that split coaxial cables input into two or more outputs and divide the power of the input signal among the outputs according to the desired balance. Taps are passive elements that connect customer drops to distribution cables. Finally, the distribution plant also requires power supplies because amplifiers and nodes are active elements and need power to operate. We assume a distributed power architecture, with power supplies placed across the distribution coaxial systems. In this architecture, there are no power cables connecting the distribution coaxial systems served by different

nodes and therefore the reach of each power supply is limited to the coaxial system to which it is electrically connected.




Figure 2. Architecture for the cable-based access network.

Customer premise

The architecture for the customer premise network in the cable-based architecture is similar to the one described for ADSL. The RG offers the same functionality, but the ATU-R function is replaced by a cable modem. Just as for ADSL and for the same reasons, the RG is placed at the demarcation point between the network and the customer premise and is owned and maintained by the cable company.

In terms of power, we also assume that the RG is powered from the premise and that power backup is provided by batteries in the RG. In fact, cable modems are currently powered by the premise and this approach is also used in some networks supporting cable-telephony NIUs. As discussed before, we think that a solution with RGs powered by the network would have advantages in terms of economies of scale and reliability of the power backup system and an easier installation of the RG. Some cable companies are already adopting this approach to power NIUs for cable telephony, but NIUs have a power consumption as low as 2W compared with the typical 10W of cable modems; and current take rates for cable telephone service are low. The large amount of power that the network would have to provide for RGs assuming high take rates would require larger power units than the ones used today. This fact makes it difficult to obtain accurate estimates for the cost of network power systems. For this reason, and because we want to have similar solutions for ADSL and cable, we opted for customer premise power plus battery backup in the cost model.

Distribution hub

At the distribution hub, transmission equipment for the cable plant includes optical panels, optical splitters, laser transmitters, optical receivers and electrical combiners/splitters. In the downstream

direction, the signals for the several services delivered over the cable plant are combined in a way to produce an electrical signal for each laser with the RF channels that the nodes served by that laser should receive. In order to allow one laser to serve multiple optical nodes, the optical signal coming out of the laser is directed into an optical splitter where it is divided into several identical outputs signals that are then send to downstream fiber strands. This design is acceptable to the extent that one downstream data channel has enough capacity to serve all the locations on multiple optical nodes; as penetration or traffic increases, the cable operator must add additional CMTSs and reduce the number of fibers fed with a common signal. In the upstream direction, optical receivers –for each fiber-- convert the signals generated by the return lasers into the electrical signals that are directed to CMTSs and other equipment supporting eventual additional services.

The data network functionality is provided by CMTSs--following the DOCSIS 1.1 reference model for two-way communications over cable--a router and ADMs. The router connects the several CMTS devices in the hub to the interoffice network. This router offers a path for traffic flowing between CMTSs or between a CMTS and the interoffice network. Although the CMTSs and the router are different network elements, the equipment currently in the market typically integrates these two functions in the same physical device. The most common configuration is to have CMTSs in the form of cards in a router chassis. Finally, ADMs provide an interface to insert and extract channels into/from the interoffice SONET ring.

2.3 Interoffice Network and Main Office (Tandem / Headend)

The networks studied here include only local access up to the main office. Higher level networks are outside the scope of this work. The network can be divided into local offices (distribution hubs for cable and central offices for ADSL), a main office (headend for cable and tandem office for ADSL) and interoffice facilities. The main office is the central access point of the network and includes gateways to external networks that may be either national backbones or other access networks. Here, the connections to Internet Service Providers, which are gateways for the Internet or managed IP backbones, are of special importance.

The interoffice transport facilities connect the main office to every local office that is part of the access network. This network consists of several interconnected fiber optics rings using OC-48 SONET technology. The OC-48 circuits in the interoffice rings are further divided into OC-1 channels that can be allocated individually to establish transport links between pairs of local offices or between a local office and the main office. There are channels connecting every local office to the main office. Direct channels between pairs of local offices are established only if there is enough traffic to justify it. Otherwise, traffic follows an indirect path via the core router in the main office.




Figure 3. General configuration for the interoffice network and the main office.

The basic elements in an IP main office are ADMs, which extract/insert basic OC-1 channels from/into the SONET rings, and a core router. This device is at the core of the access network and provides the logical switching functionality to handle non-direct traffic between local offices and traffic to ISPs. The main office is also the location for electronic equipment that does not need to be close to the local cable plant, such as network management, servers for IP services, caches and firewalls.

2.4 Voice over IP (VoIP)

In a few short years, Voice over IP has progressed from simple software programs allowing multimedia PC users to chat using microphones and headsets, to a sophisticated and complete architecture for delivering traditional telephony over packet as opposed to circuit-switched networks.

The VoIP architecture used in this paper follows the basic principles of the PacketCable 1.0 reference model from Cabelabs (CableLabs, 1999a). Although PacketCable was originally developed for cable networks, its basic design principles and functionality are valid for any architecture offering VoIP with guaranteed quality of service (QoS). Thus, the approach to VoIP is the same for both HFC and ADSL networks.

At the customer premise, in addition to the cable or ADSL modem, there is an MTA, equipped with RJ11 ports for analog telephone handsets, performing the VoIP gateway function that converts between analog voice signals and VoIP packets. We assume the configuration recommended in PacketCable 1.0 with the MTA integrated in the RG, but it might eventually be a standalone unit connected to the premise network. As we see it, an architecture based on MTAs with RJ11 ports for analog handsets may be just a transition solution to be followed by IP handsets that directly send packets over a premise data network.

At the headend, an IP “soft switch” provides the extra processing functionality required for processing voice calls. This device performs the functions of a Call Management Server, Signaling Gateway and Media Gateway Controller in the PacketCable 1.0 architecture.

3 Description of the Cost Model

This section describes a cost model for broadband access networks providing voice and Internet access as described above. The cost model can be divided into two sub-models: the ADSL sub-model covering networks of telephone operators and using ADSL technology and the cable sub-model covering networks of cable television operators and using cable modem technology. Since these two sub-models are integrated over the same platform using the same geographic and traffic assumptions, our broadband cost model provides a systematic way to compare and contrast the costs structures of ADSL and cable. Here, we provide only an overview of the model and a short description of some of its engineering rules, assumptions and parameters. For a detailed description of the model see (Fryxell, 2000).

There are a few main points about the cost model that are important to explain up front, as they are critical to understand the structure of the model, how it was built and the results that it generates.

1. The scope of the model is limited to the access network.

The cost model covers local access networks and includes customer premises, outside plant, several local offices (central office or distribution hub), a main office (tandem office or headend) and interoffice facilities connecting the offices included in the model. These networks have a scope roughly equivalent to U.S. Local Access and Transport Area (LATAs) in the telephone world. This includes only the network up to the main office and the links to external networks. National or regional backbones are outside the scope of the model.

2. Costs are calculated as though all equipment is newly installed ("Green field" model)

The cost is determined from a forward-looking perspective, meaning the cost of building the network today with current technology. For this reason, the architectures that we use do not necessarily represent the preferred architectures for companies upgrading legacy networks. Moreover, the fact that there are sunk costs in legacy LEC and cable networks is not taken into consideration.

It is important to understand that a forward-looking cost is not necessarily the most useful information for all situations, for example, an incumbent network operator facing an upgrade decision. However, this approach is important to analyze the business opportunity for new entrants contemplating an overlay network and it is the method adopted by regulators to determine the Total Element Long Run Incremental Costs (TELRIC) that are used to set cost-based prices for unbundled network elements, universal service subsidies and interconnection charges⁸.

3. The model calculates only initial first capital costs for the network infrastructure.

Expenses for maintenance, operations and support, as well as marketing, G&A and other costs of running a company are not included.

4. The cost of network elements dedicated to video is not included in the cable model.

It is very common in telecommunications that new services share network elements with older services being offered over the same network. In the case of the cable network, the IP-based access network shares part of the cable plant and buildings with video broadcast. Since the ADSL network is exclusively dedicated to IP-based access, there is an asymmetry that makes it difficult to compare the two architectures because they do not offer the same services.

In order to compare the two architectures we have to determine the cost of the cable network in a way that isolates the IP-based access network and excludes video broadcast. The first step is to exclude completely all the elements that are used only for video delivery. However, this may not be enough because some network elements are shared by both services. The allocation of the cost of shared network for pricing purposes is a business decision to be taken by the company offering the services and is always somewhat arbitrary. This common cost is an overhead that may be allocated based on the elasticity of demand for each service, proportionally to the remaining costs that can be directly allocated to each service, following a usage-based scheme, or just based on the order by which services are added. Under this latter scheme, the business opportunity for a new service is simply based on any incremental cost (including revenue loss) that it requires. Our default assumption is to allocate all the costs to the IP access network as if it were the only service being offered over the network. A sensitivity analysis is then performed to assess the impact of allocating a fraction of the cost of the common elements to video broadcast.

5. The cost model is based on the HAI Model.

Our broadband cost model is built over the basic structure of the HAI Model 5.0a⁹, developed by HAI Consulting Inc., of Boulder, Colorado. The HAI model is a cost model for the PSTN that was developed to submit to the FCC for the purpose of estimating the forward-looking economic costs of basic telephone service, unbundled network elements, local network interconnection and access. This model is probably the main tool in use today to estimate forward-looking costs for the PSTN.

We use the HAI model only as the starting point. We kept its structure and data sets, but the engineering rules and parameters for the outside cable plant were extensively modified to reflect the characteristics of wireline broadband networks instead of the PSTN. These modifications were especially large for the cable network, where the outside plant follows an HFC architecture as opposed to the traditional telephone networks based on copper pairs. The switching and interoffice network of the HAI Model, which is circuit-based, were also replaced by an integrated IP network providing voice and Internet access.

The main advantage of using the framework of the HAI Model is the possibility of accessing its rich data set containing detailed demographic, geographic and geological information of the areas served by each telephone company in the US. This feature allows the model to provide network construction cost estimates area by area based on real data for the entire US. Other data inherited from the HAI model include the location of offices, estimates of voice traffic, and cost and engineering parameters for placing cable, cable infrastructure (poles, conduits, trenching, etc.) and buildings.

3.1 Structure of the Model

The database that we use is organized by local and main offices, which, due to the PSTN origin of the database, are real central offices and tandem offices used by telephone companies. Each local office is associated with an area of coverage, a main office and a telephone company that owns it.

⁸ A recent decision by the U.S. Eighth Circuit Court of Appeals has rejected the FCC's use of TELRIC principles for setting UNE prices and remanded the issue back to the Commission. The FCC has not yet indicated if it will appeal the decision to the U.S. Supreme Court.

⁹ The first releases of this model were named Hatfield Model.

The model provides input data by state and by LEC. Each run of the model calculates the cost of the access network for every LATA served by the LEC in that state. Given the name of the state and the company, a query to the database returns the local offices, main offices and their respective locations. Since some areas include more than one main office and we defined access network as having only one main office, the total access network for a given area may include several sub-networks, each one including one main office and multiple local offices. Once the local offices that are part of the network have been identified, the model runs a query to the database for each one of these offices that returns the clusters of customer locations that are served by that office. Then, the calculations for network equipment and the associated investment are performed in the reverse direction.




Figure 4. Flow of operations in the cost model.

The first step of the cost computation is performed for every cluster individually and includes the distribution and part of the feeder within the limits of the clusters. Then, the investment for the feeders connecting the clusters to their respective local offices¹⁰ is calculated. The calculations for the local offices, main office and interoffice network come next. Finally, the investments are grouped and the annual capital carrying costs are determined.

3.2 Geography of the Network

The HAI Model database that we use provides the locations of real central offices and tandem offices in the PSTN. We assume that for a pure ADSL architecture telephone companies would still keep their offices in the same locations. For the HFC network design, we also assume these office locations for distribution hubs and headends. From a cost modeling perspective, the most important issue regarding the placement of distribution hubs is the number of customer locations that they serve. We compared the numbers obtained from the database with the ones commonly referenced for hubs in the trade literature and we concluded that this assumption is reasonable for the purpose of estimating costs.

The area served by a local office is divided into clusters of customer locations. There is a record in the database for each cluster describing its demographic, geographic and geological characteristics. Clusters are the smallest units for which we use real data. Further sub-divisions of a cluster and the specific location of houses within a cluster are exclusively based on assumptions. The database from the HAI Model has 331,488 clusters for the entire US. The average area per cluster is 3.986 square miles and the average number of locations (households + businesses) is 343.

¹⁰ In Figure 5, this part of the feeder is the main feeder plus the sub-feeder.




Figure 5. Feeder and clusters.




Figure 6. Cluster grid.

As showed in Figure 5, the model connects the local office to every cluster using a feeder cable following a star architecture. For the ADSL network, the feeder may be fiber or copper, but for the cable network it is always fiber. The feeder follows a right-angle path, with four perpendicular main feeder routes leaving the local office and sub-feeders connecting the main feeder to the center of each cluster. Main feeder and sub-feeder distances are computed from the radial distance and the angle alpha that can be found on the record describing the cluster. A distribution grid is applied to each cluster following the methodology used in the HAI Model. As showed in Figure 6, the cluster is divided into equal-size lots with a 2x1 aspect ratio. The number of lots, and therefore the size of the grid, is a function of the number and type of customer locations. The “grid” topology has branches extending East and West where the distribution cables pass to serve the individual subscriber locations. For a complete description of how the grid is designed see (HAI, 1998b).

3.3 Locations, Voice Lines and Traffic

Locations and voice lines

The networks analyzed here serve households and small businesses. Medium and large corporations have access to more expensive higher-capacity networking solutions. We adopt the simplifying assumption that all the customers are either households or small businesses and therefore there are no economies of scale associated with the fact that the same network provider would probably also offer access solutions to large businesses.

The database records describing each cluster contain the number of locations and PSTN voice lines. We use these values to define the number of locations and voice lines passed by the network. The numbers of households and residential lines are the same as the numbers read from the records describing the clusters. Business users are divided into two categories: large and small. We assume that large businesses will not be served by either ADSL or HFC networks, but by direct on-net fiber connections. We filter out the demand from large businesses using estimates of the percentage of business that are small businesses and the percentage of business lines that are small-business lines. After the initial adjustment, the number of business lines is further corrected such that the number of business lines contained within each cluster is not larger than the number of businesses passed times a parameter that

defines the maximum number of lines per business. These numbers are adjusted network-by-network such that the resulting number of lines per business passed, after the two corrections, is consistent with typical values for small businesses.

Also important to determine the number of locations and voice lines to be used in the cost computations are the take rates defining the penetration of each service in the market and the percentage of locations passed that are actually connected to the network. Since locations have to be connected by a customer drop before they can subscribe to any service, the fraction of on-network locations have to be at least as high as the highest take rate for any service.

Voice traffic

Voice traffic is determined based on the concept of voice line. This concept is derived from the PSTN, where a voice line translates into a physical copper loop and a port at the local switch. For an integrated IP architecture, where multiple voice conversations share the same access link, the concept of voice line has simply the meaning of a logical voice line, which is used to calculate voice traffic demand.

The average numbers of voice lines per household and small business are determined at the cluster level based on the number of residential lines, business lines, households and businesses passed for each cluster. The busy hour traffic per residential line and business line are derived from voice traffic currently carried by the PSTN as reported by LECs in the ARMIS report for the respective area. The traffic is divided into intra-office (intra distribution hub for the cable architecture), local interoffice (to somewhere else in the access network) and inter-network (to a point external to the access network).

The default bit rate for voice is assumed 96 Kbps. This parameter assumes use of a G.711 codec with packet size of 10 ms without silence suppression at a 64 Kbps bit rate. The headers include 12 bytes of RTP, 8 bytes of UDP and 20 bytes of IP for a total overhead of 32Kbps. The transmission of voice packets in the HFC plant assumes header suppression as described in (CableLabs, 2000). See section 4 for a sensitivity analysis on voice codecs.

Internet access traffic

Traffic for Internet access is calculated on a per location basis. The business traffic is not based on real business locations, but on the concept of small business equivalent. Since the traffic for businesses is likely to be subject to large variations according to the type and size of business, it is not appropriate to treat all locations as equal. We make the assumption that the Internet traffic for each business location is proportional to its number of voice lines. We estimate the traffic for a small business equivalent with a fixed number of lines, which by default is four. For each actual business in the database, we determine how many small business equivalents it represents by looking at the ratio of voice lines. The Internet traffic per business in a cluster is then given by

$$\frac{\text{business lines}}{\text{lines per } SB_{\text{Equivalent}}} * \text{traffic per } SB_{\text{equivalent}}$$

Internet access traffic for the busy hour is calculated using the parameters for the average data rate for an active location and an estimate of the fraction of active locations during the busy hour. By computing the traffic per location as the product of these parameters, the model accounts for the fact that at any given moment, not all households and businesses are using their Internet access.

There are few useful studies for estimating likely packet traffic by residences and businesses who have broadband access. Moreover, new applications and technologies are constantly emerging --such as

Napster-- and it is impossible to predict how future traffic patterns will evolve. For households, according to research done at AT&T Laboratories, a typical broadband Internet surfer looking at text and still pictures can read data at only 40 Kbps, averaged over a 3-minute interval, and needs an average upstream link speed of just 10% that value (Dutta-Roy, 1999). According to the same source, the busy hour activity ratio, measured as the percentage of subscribers active during the busy hour, is between 30% and 50%. We think that 40Kbps is on the high side because users will not be reading web pages at full speed all the time. Early adopters are likely to have a higher bandwidth usage, but the average per user would certainly go down for higher take rates. Using this information from AT&T and the experience with ADSL service of the institution where two of the authors of this paper work, we decided to adopt 20Kbps as the default average residential downstream data rate for Internet access. We are using the lowest end of the interval for activity rate. Even so, assuming 1 in 3 during the busy hour may still be a little high. For example, ISPs size their PSTN modem banks to a peak activity rate around 10%. Moreover, according to (Morgan, 1998), if we look at the early busy hour around 4PM, when businesses are active, the residential activity rate is about 80% to 90% of its busiest hour of the day. For a small business equivalent we assume 80Kbps downstream, 25% of that for upstream and an activity rate of 90% during the busy hour. We should note that there is a large uncertainty in these numbers and that they should be regarded only as a starting point to conduct sensitivity analysis on bandwidth usage.

3.4 Customer Premise

The prices for ADSL and cable modems obtained from cost models, industry journal, list prices and resellers vary considerably. Retail prices for external ADSL Modems with basic transmission and bridging functionality, PPP support and an Ethernet interface have come down from \$400/\$500 two years ago to around \$250/\$280 today. Cable modems have historically been a little cheaper due to an initial higher scale of deployment, but today they are roughly in the same price range as ADSL modems. The wholesale price at which a network provider buys this type of equipment is likely to be substantially lower. Just by applying a 30% discount, we get a cost per modem between \$175 and \$200. Since we want a forward-looking price for a network where a large number of these devices would be deployed, the cost would most likely be closer to the lower end of this range. However, we must not forget that in the architecture used here, the RG is placed outside the customer premise and therefore has to be environmentally hardened. The extra cost for an environmentally hardened case of the type used for NIDs and NIUs is close to \$20 and electronic components that withstand extreme temperatures are also more expensive. With these additional costs, we estimate a price of \$225 for the part of the RG that is the broadband modem (ADSL or cable).

With respect to the MTA functionality for voice, it was more difficult to get price estimates. We only found one reference from an MSO pointing to an incremental cost of around \$60 to \$70 to integrate the functionality of a VoIP terminal with two RJ11 ports in a RG. We use \$60 as the default value for a 2-port residential MTA and \$120 for a 6-port business MTA.

Currently, the power consumption for cable and ADSL modems is in the range of 7W to 10W. For an RG, we still have to add the power for the MTA functionality. Providing power backup for this kind of device during the 8 hours recommended by Bellcore would require a quite expensive and bulky battery. However, we believe that power consumption for RGs may be reduced considerably with the integration of all the operations into just a few chips. Moreover, the RG consumption during a power

outage can be reduced even further by moving the RG to a “lifeline state.” During these periods, the PC interface functionality would be turned off, the speed of transmission would be reduced from full-rate to a slower sub-rate and the device could go into a sleep mode when not in use. We believe that these measures may lead to devices that won’t require much more than the 2 to 5 Watts¹¹ that NIUs currently consume. In order to assure 8 hours of power backup for telephone service, a battery will have to provide from 16Wh to 40Wh. By comparison with the price for similar backup batteries in other electronics products, we believe that the battery and charger should add from \$50 to \$70 to the cost of an RG. We adopt \$60 as default value in the model.

We use an average installation time of 80 minutes per RG. This includes a short survey (10 minutes), installation of the case with connection of drop cable and telephone wires (30 minutes), installation of power line (35 minutes), and testing (5 minutes). At a loaded labor rate of \$50 per hour, the installation adds \$67 to the cost of the RG. These estimates reflect the opinions of experts in the installation of NIDs for telephone networks and NIUs for cable telephony.

3.5 ADSL Specific Elements

Customer Drops

The drop investment is determined following the methodology used in the HAI Model, where the drop distance and the fraction of buried versus aerial drops is a function of the density of lines¹² in the cluster (HAI, 1998a and HAI, 1998b). The investment required for customer drops was divided into cable and labor. Since, in contrast to the PSTN, the ADSL architecture has only one loop per location, the cable part was adjusted to account for this reduction in the number of loops. Installation labor was not adjusted for fewer drops per location, as the labor for installing drop cable is essentially insensitive to the number of pairs in the cable.

Distribution

The model subdivides the cluster into rectangular serving areas such that the loop distance from the center of each serving area, where a remote terminal is located, to the furthest customer location -- following a right-angle path-- is smaller than the maximum permissible length for the copper portion of the loop. This limit is a function of the desired ADSL speed.

The distribution is divided into aerial, buried and underground fractions¹³. The network equipment includes cables with copper pairs, terminals where customer drops are connected, a distribution frame that interfaces the distribution copper pairs to the feeder, conduits and poles. The investment required for this part of the network includes also all the costs involved in placing the cable and installing the rest of the equipment. Since the architecture for the distribution plant follows the copper-pair model used in the PSTN, we adopted the engineering rules and cost information used in the distribution module of the HAI Model. For a detailed description of these calculations see (HAI, 1998a and HAI, 1998b).

¹¹ This range of power consumption is not very far from what is possible using existent chips for ADSL modems. In fact, the power consumption for chips performing all the ADSL functionality ranges between 1.4W and 3.6W (see the chipset table on www.xDSL.com).

¹² The density of lines is used as a proxy for the density of buildings in the area. Low-density areas are likely to be rural areas where drops are longer, mostly aerial, and cheaper to install on a per foot basis.

¹³ The percentage of each type is a function of the density of lines in the cluster.

Remote Terminals (RTs)

Remote terminals are needed when the total distance of the main feeder and sub-feeder plus the distance from the center of the serving area to the furthest customer is larger than the maximum permissible length for the copper part of the loop. Otherwise, there is no RT in the cluster and the feeder is copper all the way to the central office.

RTs include power, housing and DSLAMs with enough ADSL line cards for all the loops terminating at RTs plus a percentage of spare ports. The model chooses between small-scale and large-scale DSLAM versions, with 48 and 240 ADSL line ports per unit respectively. The cost for housing and power is determined based on the same costs for PSTN RTs in the HAI Model, adjusted for the space and power requirements of ADSL line cards. DSLAMs are configured with redundant power supplies, switching processors and network cards for high-reliability operation. The cost for DSLAMs was obtained by applying a 30% discount to list prices of equipment currently in the market. Since we were not able to find prices for environmentally hardened versions for outside plant installations, we assume a 20% increase over prices for central office equipment.

Feeder

The feeder may be fiber or copper depending on whether the cluster contains RTs or not. As shown in Figure 5, it is divided into main feeder, sub-feeder and eventually, when there are several RTs in the cluster, in-cluster feeder connecting RTs. This last section of the feeder shares part of the cable structure (poles, trenches, conduits, etc.) with the distribution and is computed at the same time as the rest of the cluster. The main feeder and sub-feeder follow the same configuration as the feeder in the HAI Model. We use the methodology followed in the HAI Model, with four main feeder routes and sub-feeders connecting these routes to the center of each cluster. For a detailed description see (HAI, 1998b).

Central Offices

The investment associated with the main distribution frame is computed based on a cost per line as suggested in the HAI Model for the PSTN. The fiber distribution frame is composed of 24-fiber optical panels that we assume equivalent to those used to connect interoffice fiber rings in the HAI Model. Central Office DSLAMs are similar to those used in remote terminals, but do not need to be environmentally hardened. Only large-scale DSLAMs with 240 ADSL line ports are assumed used in the central office. Several DSLAM shelves are integrated in the same access system and share a single OC-3 network port to the access server. The access server investment reflects the cost for a fully redundant server (redundant network ports, redundant routing engine and redundant switching engine) based on list prices with a discount of 30%. The price for ADMs is the same as used by the HAI Model for similar equipment.

The computation of the investment for the building uses parameters from the HAI Model that are a function of the number of PSTN lines served by the office. These parameters include the area of the building and surrounding lot, the cost for power systems and costs per square foot for land and construction. The number of lines is a proxy for the different cost areas in terms of land and construction. In terms of area, the model determines an equivalent number of PSTN lines based on the number of ADSL loops that terminate as copper and require ADSL ports at the central office. Since ADSL port termination technology requires more space than POTS, a conversion factor, which by default is set to 3, is used. The investment required for power systems is also determined based on the number of copper loops, but assumes ADSL equipment requires 6 times the power per line of a POTS network.

3.6 Cable Specific Elements

Customer drops

We assume that in terms of labor involved in installing drops, coaxial cables are roughly equivalent to copper pairs and therefore the labor cost to install the drops is exactly the same as in the ADSL architecture. With respect to cable, we account one drop per location connected to the network and assume the same drop distance as for ADSL.

Distribution

We assume a cluster consists of a grid of streets with houses along the horizontal and these linked by vertical avenues. Tapped branch cables are used along the horizontal streets of the distribution grid that pass by each customer premise; these branches are then connected along the vertical to a node or to a network amplifier. In the most basic configuration, we have two connecting cables (one North and one South) connecting the node to several branches. As showed in Figure 8, these cables do not have taps, but have splitters and directional couplers connecting the cable to distribution branches. The model places taps along branch cables to serve the locations passed by each branch. Typical taps may be 2-way, 4-way or 8-way¹⁴ and only a fraction of the ports is actually connected to drops. The percentage of each type of tap and the tap fill rate are a function of the cluster density characteristics. The model uses a table with these values that was compiled based on the analysis of diagrams of coaxial systems in real cable networks.




Figure 7. Branch of distribution.




Figure 8. Vertical connecting cable.

The number of line extenders is determined such that there is enough signal amplification to meet attenuation and power loss for taps along branches. The attenuation along the cable is simply a function of its length and diameter. The power derived for each tap depends on the number of ports and the minimum required power at each port, which is given by the minimum power at the customer premise plus the attenuation of the drop. A precise determination of the power derived for each tap would require the knowledge of its exact location and number of ports. Since the available information about clusters is not detailed enough to obtain these values, we make the assumption that all the taps are equal and have a number of ports given by the average number of ports per tap in the cluster. Furthermore, we assume a constant distance between taps within each cluster. Amplifiers are then placed such that the minimum allowable signal level at the input to an amplifier is not violated. The model follows the typical principle in the design of cable networks of having the same operating level¹⁵ for all amplifiers and adjusting the

¹⁴ An n-way tap can feed a maximum of n drops.

¹⁵ Signal level at the output of the amplifier.

gain¹⁶ for each amplifier according to the signal level at its input. While signal to noise ratio constraints give the minimum allowable signal level at the input to amplifiers, distortion limitations define the operating level.

In order to extend the area covered by a node without violating the maximum amplifier cascade design rule, the model uses Network Amplifiers (NAs) connected by express cables. These cables can reach further than what would be possible with ordinary distribution cables because they do not have any power loss for taps, splitters or directional couplers. Their maximum length is limited by attenuation only. NAs are a kind of second level “distribution center,” similar to nodes, with several output ports providing connection to branch cables and vertical cables. Typically, NAs are prepared for a future conversion to optical nodes and they are equipped with monitoring transceivers that allow a centralized management system to detect eventual problems with the amplifier or any of its output legs. It should be noted that the practice followed here of associating NAs with express cables is not a design rule followed in all HFC implementations. It all depends on the particular geographic and demographic configuration of each area. However, we believe that the generic architecture followed by the model provides reasonable estimates for the number of NAs and line extenders required for the network.

The power required for each node plus the coaxial system around it is calculated assuming a consumption of 1 Amp for the optical node plus 0.7 Amp per NA and 0.46 Amp per line extender. These currents were derived from (Donaldson and Midkiff, 1999) and assume some voltage loss over the distribution plant. Each distribution serving area is then equipped with large (15 Amp) or small (6 Amp) 90 VAC power supplies equipped for standby operation. The costs used for power supplies include line conditioner, DC-AC inverter, battery charger, cabinet, two 36-volt strings of 12-volt batteries, network monitoring transponder and installation.

Division of cluster into distribution serving areas

The area covered by an optical node may be either limited by the maximum cascade of amplifiers or by the maximum number of customer locations that can be served by a node. The maximum allowable number of households per node is a parameter of the model. There is not such a parameter for businesses. Instead, business locations are converted to an equivalent number of households based on the relation between the traffic for businesses and households. When an optical node is not sufficient to cover the entire cluster, the cluster is divided into distribution serving areas with a node for each one of them. All the serving areas in a cluster have the same size and are rectangular in shape. Given the limitation to the maximum cascade of amplifiers, there is a trade off between how many amplifiers should be used in the vertical versus the horizontal directions. Solutions with more amplifiers extended vertically have “taller” serving areas and therefore need fewer serving areas stacked vertically. The inverse happens when the priority is given to horizontal amplifier cascades.

The algorithm that divides the clusters into serving areas first defines the feasible space for this partition based on the particular geography and demography of the cluster and then searches for the best configuration along the boundary of the feasible space. The best configuration for the distribution plant is the one that minimizes the number of nodes per cluster given the constraints for the maximum number of amplifiers in cascade (network amplifiers plus line extenders) and maximum number of customer

¹⁶ In reality, the gain of the amplifier is also constant. The adjustment is given by the insertion of pads that create a loss before the amplifier.

locations served by a node. When there are multiple solutions with equal number of nodes, the one with fewer NAs is preferred. This assures that NAs and express cables are used only when they contribute to reduce the number of nodes. Express cables represent an additional cable cost, though, by limiting attenuation, they may allow a single node to serve a larger geographical area. For this reason, when the use of express cables does not contribute to decreasing the number of serving areas per cluster, they are not used at all.

Optical nodes

The model assumes one bi-directional optical node per serving area. The cost used for optical nodes is based on a configuration with an optical receiver, an upstream data-grade 1310 nm Fabry-Perot laser transmitter, net monitoring transponder and four coaxial legs equipped with bridger switches, which provide the capability to isolate RF ingress in the upstream spectrum. Each fiber node terminates four fiber strands: one upstream, one downstream and two spares.

Feeder (fiber backbone)

The calculations for the feeder in the cable model are basically the same as in the ADSL model. There is a main feeder, a sub-feeder and, when there are multiple nodes per cluster, in-cluster feeder. The only difference is that the feeder is always fiber because there is at least one optical node per cluster.

Distribution Hubs

The model assumes distribution hubs are equipped with optical patch panels, optical transmission equipment, integrated CMTS/router and ADMs. By default, a laser transmitter followed by a four-way optical splitter serves a group of four nodes. In this architecture, all the nodes in a group receive the same signal. When downstream data channels serve fewer than 4 nodes, the size of these broadcast groups is reduced accordingly. In the upstream direction, there is one optical receiver per node. The CMTSs are cards in a router chassis. CMTS cards have one downstream channel and may have 1, 2, 4 or 6 upstream channels. The model chooses the configuration that minimizes cost. As default, downstream channels are QPSK modulated with a data rate of 30 Mbps and upstream channels are 16-QAM modulated with a data rate of 5 Mbps. The price of this element is derived from list prices with a discount of 30%. The ADMs that provide the interface to the SONET optical rings are essentially the same used in the ADSL network.

The computations for the floor area and power systems are also similar to the ADSL case, but the equivalent number of POTS lines is now based on the number of optical nodes with a default factor of 10 lines per node. This factor leads to an area of 500 square feet for the smallest distribution hubs and 1000 square feet for the largest ones.

3.7 Interoffice Transport and Main Office (Tandem Office / Headend)

Our model uses a module from the HAI Model that computes the configuration of the interoffice optical rings and provides the distances involved assuming that cables follow a right-angle route between offices. The main office is connected to five different ISP POPs using links that are half a mile long each. The interoffice rings use OC-48 circuits, while the cable links between the tandem and the IXC/ISP's POPs use OC-3 circuits. It is not cost-effective to multiplex OC-3 channels into a higher capacity OC-48 circuit because these links are very short. Savings in fiber would be outweighed by the extra cost in ADMs. Each OC-48 circuit in a ring requires two strands and each point-to-point circuit, such as the ones

between the main office and ISPs, requires four strands for redundancy. We assumed no use of Dense Wavelength Division Multiplexing (DWDM). The costs of fiber cable and cable infrastructure (protection, poles, conduits, and installation) are the same as in the HAI Model.

As shown in Figure 3, the VoIP architecture requires an IP soft switch at the main office. This switch's assumptions are based on equipment designed for network providers offering signaling and call management functions for VoIP. These products are still in the very early stages of market deployment and do not always provide the full functionality and the scale that would be appropriate for large networks. We based our cost estimates on list prices with a 30% discount, but we regard them as rough estimates for what the cost of a fully functional soft switch may be in the future.

The calculations for the building and power systems are based on the same tables used for the local offices. We use the land and construction costs per square foot that are associated with the highest density zone. The building is assumed to have the same 10,000 square feet that the HAI Model provides for PSTN tandem offices. In terms of power systems, the model assumes that a main office is equivalent to a 10,000 line central office in the PSTN world. Finally, as suggested in the HAI Model, we use a 40% reduction factor to account for savings related to the large percentage of main offices that are co-located with local offices.

3.8 Capital Carrying Cost

The initial investment for each network element is converted into a flow of monthly payments for the economic life of the network element. The presentation of cost under the form of a monthly cash flow allows the comparison of investments for different elements with different economic lives on an equal basis. Moreover, these monthly cash flows provide the monthly revenue that the network operator has to receive from its customers in order to amortize the initial investment.

For each network element, or class of elements, the value that we want to determine, P , is the periodic flow of payments that should be received from customers in order to pay for the capital carrying cost associated with the initial investment in that element. The cost of capital for the network operator is the weighted average of the cost of equity and the cost of debt given by

$$r = r_e w_e + r_d w_d$$

where w_e and w_d are the percentages of equity and debt. The operational results are subject to a corporate tax at the rate t . Taxes are assumed due on all revenue less depreciation and interest on debt. Assuming the same monthly period for revenues and taxes and linear depreciation for tax purposes, the monthly payment P for a network element with an economic life of N monthly periods is determined from the expression

$$Initial\ Investment = \sum_{i=1}^N \frac{P^*}{(1+r^*)^i}$$

where the return on investment and payment adjusted for taxes are given by

$$r^* = r - r_d w_d t$$

$$P^* = (1-t)P + \frac{Initial\ Investment}{N} t$$

The parameters required in order to calculate the capital carrying costs are the cost of capital and the economic lives of network elements. We use the economic lives for outside plant and buildings

adjusted for net salvage values proposed in (HAI, 1999a). These values are based on the average projection lives as determined by the three way meetings (FCC, State Commission, and ILEC) for the RBOCSs and SNET. The same study provides economic lives of 16.50 and 10.07 years for circuit switching and digital transmission equipment, respectively. We think that in integrated data architectures there is no differentiation between switching and transmission equipment. All the network electronics is data communications equipment and should have a single economic life. We selected the value proposed for digital transmission equipment for this economic life. For the sake of consistency, we use the same value for CPE, though it can be argued that the economic life of CPE will be shorter as newer models offer new capabilities.

It is important to note that our economic life for electronic equipment is considerably larger than the values typically found in cost models for data networks, which in some cases are as low as 3 to 4 years. The reason for our relatively long economic life is that, in contrast to most of these models that are created for risky small-scale investments in ISP networks, our model addresses more stable investments in large networks of local access providers, either Local Exchange Carriers (LEC) or a cable companies. Carriers are more likely to base their access network architecture on a stable technology platform for compatibility reasons.

4 Results and Analysis

We have used our model to calculate illustrative results for the entire state of Delaware. We assume that the network serves all the households and small businesses in the state. This means that there is a take rate of 100% of potential customers. The area covered by the network has 1 headend and 33 distribution hubs.

The results for the ADSL architecture were obtained for a maximum copper length of 18k feet. The results for the cable architecture were obtained for a maximum allowable cascade of 4 amplifiers and a maximum node size given a limit of one downstream and one upstream data channel per node. This number is 355 households per node for the default configuration going down to as low as 276 for the bandwidth multiplier of 5 (see sensitivity to data rate for Internet access) and as high as 1,500 for the lowest bit-rate codec (see sensitivity to voice bit rate).

Table 1 shows the characteristics of this area broken into density zones. The density zones are defined by the number of voice lines per square mile. We selected this measure because it is also used in the HAI Model and it is employed therein to determine other parameters that are a function of the density zone (percentage of aerial versus buried or underground structure, costs of trenching, etc.). Density also determines the extent to which the fixed costs of cable placement are shared among multiple customers. In the PSTN, the number of voice lines determines the size of the cable whose placement costs are to be shared. For a broadband network, the relevant measure is the number of locations, rather than the number of voice lines, as multiple "voice lines" may be carried over a single broadband drop. However, the density of lines is also a good proxy for density of locations. In terms of households, locations or lines convey roughly the same information. The same is not necessarily true for businesses, where there is less uniformity, but, as can be seen in the fourth column of Table 1, density zones for voice lines or for locations are roughly the same. Table 1 also shows that the distance from the cluster to the local office is another measure that is highly correlated with the density zones. On average, the distance from the cluster to the distribution hub decreases for higher density zones. Not surprisingly, distribution hubs are

typically located in the highest density areas. Finally, most of the locations passed by the network are households. Businesses are mostly concentrated in the top half of the density zones.

Density zone (lines/sqr mile)	Households	Businesses	Locations per sqr mile	Residential voice lines	Business voice lines	Clusters	Average distance from cluster to local office (feet)
5	501	18	14	544	65	10	43,706
100	49,294	2,289	42	53,863	6,951	169	30,064
200	22,704	1,036	100	24,689	3,045	48	21,797
650	41,617	3,660	258	46,062	16,764	80	18,020
850	8,646	948	464	9,448	5,159	17	17,956
2,550	82,260	8,364	940	91,189	35,972	156	13,596
5,000	37,636	3,886	1,922	41,348	15,173	68	12,500
10,000	33,046	4,017	2,615	36,398	19,328	62	11,741
10,000+	19,196	3,598	5,133	20,748	20,307	48	4,663
Total	294,900	27,815	175	324,290	122,763	658	18,592

Table 1. Characterization of the area served by the network.

As shown in Figures 9 and 10, for both broadband architectures CPE accounts for the largest fraction of the total network cost. For the cable architecture, CPE is 43% of total cost; for the ADSL architecture, even though the dollar value is similar, CPE accounts for only 28% of the larger total cost. Remote and in-office DSLAMs make up another third of ADSL architecture costs. CPE and loop plant up to the local office dominate costs in both architectures: 87% for cable and 85% for ADSL. Offices and office equipment account for 11% for cable and 14% for ADSL. Finally, the interoffice facilities are irrelevant in terms of total cost, accounting for no more than 2% to 3%.




Figure 9. Cable architecture cost split.




Figure 10. ADSL architecture cost split.

Overall, the HFC architecture costs roughly 35% less than the ADSL architecture. This difference is mainly due to the large cost difference between optical nodes and remote DSLAMs. Despite the fact that only 68% of clusters are served by remote DSLAMs while all clusters are served by optical nodes, the ADSL architecture shows a cost for remote DSLAMs that is roughly 17 times larger than the cost for optical nodes in the cable architecture. The relatively higher cost of a DSLAM when compared to a CMTS is also the reason why local office costs are higher for the ADSL network. Furthermore, the ADSL architecture requires a unique twisted copper pair for each location, while coaxial cable is shared among many subscribers. Thus, distribution plant costs are also higher for ADSL than for the cable architecture. In general there is a much greater degree of equipment and plant sharing in the cable architecture as compared to the ADSL architecture: coaxial cable, optical nodes and the CMTS are all shared. By contrast, in the ADSL architecture, copper loops and ATU-C interfaces are required for each

location. Only in the feeder plant does ADSL have an advantage; serving areas for ADSL remote terminals are typically larger than for optical nodes, resulting in fewer fibers in the ADSL architecture compared with the cable architecture.




Figure 11. Cost per location for ADSL and cable architectures.




Figure 12. Cost per household vs. small business.

Figure 12 shows how the cost per household compares with the cost per small business in both architectures. The cost of common elements that are sensitive to traffic has been allocated on a per Kbps basis. The cost for non-traffic sensitive (NTS) elements – e.g. CPE, distribution, copper feeder and DSLAMs was allocated on a per location basis. We should note that this is not the only way to allocate costs. Indeed, the proper allocation of shared costs in the telecommunications industry has been debated for decades.

The difference in cost per location between small businesses and households is due in part to CPE cost (business MTA, serving more "voice lines," is more expensive than a residential MTA). Moreover, a larger share of feeder and office equipment is being allocated to business locations due to their higher traffic. The cable architecture has fewer NTS costs and more shared plant; thus, allocating common costs by traffic results in a greater difference between households and business in the cable architecture. This higher sensitivity to traffic in cable architectures leads to very similar costs for business location in both networks.

5	100	200	650	850	2550	5000	10000	10000+	Total
100%	90%	85%	73%	72%	65%	70%	59%	15%	68%

Table 2. Percentage of clusters with ADSL remotes by density zone.

Figures 13 and 14 show that cable's cost advantage over ADSL extends across all density zones. The cable advantage is relatively small for the low-density zones, where the dominant term is the feeder and the potential for sharing in the cable architecture is not fully realized. This difference grows to its maximum value in the intermediate-density zones and then decreases slightly for the high-density zones. This last effect is actually related to the distance between the cluster and the local office and only indirectly related to density zones. As can be seen in Table 2, the highest density zones are associated with shorter distances to the local office and therefore most of the loops are entirely copper. The concentration of DSLAMs in the local office leads to savings associated with cheaper equipment, larger scale and higher fill rates.




Figure 13. Cable architecture cost by density zone.




Figure 14. ADSL architecture cost by density zone.

Figures 15 and 16 show the sensitivity of ADSL and cable networks costs to Internet traffic assumption. The four columns provide the results when the default data rates (20K/2K for households and 80K/20K for small businesses) are multiplied by 0.5, 1, 2 and 5 respectively. Results show that the ADSL network is basically insensitive to any changes in data rates. In fact, as long as the bandwidth is within the limits of ADSL technology, there is basically no effect on the part of the network between CPE and DSLAMs. Only access servers, routers and SONET equipment are affected by this change. These terms increase by an average of 53% when the data rate multiplier goes from 0.5 to 5, despite the fact that the total cost of the network increases by little more than 2%. The cable network, where the access is shared, is considerably more sensitive to bandwidth usage. While the total cost increases by 14%, nodes increase by 32% and the total cost of CMTSS, routers and SONET equipment increases by an average of 176%.




Figure 15. Sensitivity to Internet access bandwidth for the ADSL architecture.




Figure 16. Sensitivity to Internet access bandwidth for the cable architecture.

Figures 17 and 18 show the sensitivity of cable and ADSL architectures to the voice codec data rate. The codecs analyzed here are G.711 with packet size of 10ms (64Kbps plus 32Kbps overhead), G.711 with packet size of 30ms (64Kbps plus 10.7Kbps overhead), G.728 with packet size of 10ms (16Kbps plus 32Kbps overhead) and G.729A with packet size of 30ms (8Kbps plus 10.7Kbps overhead). For a discussion of codecs for VoIP see (CableLabs, 1999b).

The results are similar to those observed for Internet access. As expected, the ADSL network is insensitive to bandwidth usage. The total cost increases less than 1% from the least bandwidth intensive to the most bandwidth intensive codec. The cost of access servers, routers and SONET equipment increases by an average of 4%, which is less than what we observed for Internet access. The cable network is considerably more sensitive. The cost for CMTSs, routers and SONET equipment increases by an average of 95% and the total cost increases by 6%. It is interesting that, with an increase of 32%, the optical nodes are more sensitive to variations in voice than Internet access. This is due to voice having a larger upstream data rate and the lack of capacity of the upstream channel being the most important factor for how bandwidth affects the number of nodes.




Figure 17. Sensitivity to voice data rate for the ADSL architecture.




Figure 18. Sensitivity to voice data rate for the cable architecture.

As explained in the description of the cost model for the cable architecture, the data access network shares part of the cable plant and buildings with video broadcast. Our default assumption was to allocate all the costs to data as if it were the only service being offered over the network. Figure 19 shows the sensitivity of the results to the fraction of the common elements that is allocated to data versus video. Common elements include downstream cable plant, buildings and interoffice facilities. Results show that the ability to share part of the network with video delivery leads to an even greater cost advantage of cable over ADSL.

Figure 20 shows what happens if we assume that broadband access equipment costs come down substantially in the future due to improved electronics. The elements included in this sensitivity analysis are CPE (not including installation, backup batteries and case), DSLAMs, access servers and CMTSs (including the router sharing the same chassis). The ADSL model is very sensitive to DSLAM costs, and thus the difference between the two architectures is reduced. Cable's advantage over ADSL would be reduced from 35% to 22% if electronics costs were to drop by half.




Figure 19. Sensitivity to fraction of common elements allocated to data vs. video for the cable architecture.




Figure 20. Sensitivity to broadband electronics cost for the ADSL and cable architectures.

5 Conclusions

This TELRIC cost analysis shows that an HFC architecture has a significantly lower cost per subscriber than ADSL for the provision of an integrated broadband access network providing both Internet access and Voice over IP. HFC networks allow sharing of distribution plant and CMTS costs across multiple subscribers, while twisted pair loops and ADSL DSLAM ports must be dedicated to each location. DSLAM port costs are subject to significant future cost reductions as integrated circuit densities increase; by contrast, a higher percentage of HFC costs are in outside feeder and distribution plant, which is less likely to benefit from Moore's law.

A TELRIC analysis is most relevant when considering the costs likely to be incurred by a new entrant overbuilding a network in competition with incumbent cable or telephone service providers. Looked at from the perspective of upgrading existing networks, it appears that the cost of upgrading existing cable plant to HFC and installing CMTSs is less than the comparable cost of adding to existing telephone plant new integrated digital loop carrier equipment or central office DSLAMs to support ADSL to all households. Moreover, we have assumed that the entire cost of the HFC network is allocated to voice and Internet access services. Considering that for a modest incremental investment in transmission equipment, the HFC plant can also provide delivery of entertainment video, thus garnering significantly increased revenue, the HFC architecture looks even more attractive.

The model allows us to look at the cost of providing integrated broadband services to areas of different population densities. For these two wireline architectures, there is nearly a four to one difference in cost per location served between the dense urban core and rural areas with less than 5 subscribers per square mile for both architectures. The HFC architecture is less costly than the ADSL architecture for all density levels.

The greater degree of equipment sharing in the HFC architecture means that more incremental investment is required as subscriber data or voice traffic increases. That is, the cost of providing service using an HFC architecture increases with traffic growth more than does the cost of an ADSL network. The cost of the HFC network is particularly sensitive to the amount of traffic in the upstream direction.

Thus, growth in applications such as videoconferencing or peer-to-peer networking (*e.g.* Napster, interactive games) would have a relatively more significant impact on the HFC architecture.

Over the last 18 months, many observers have expressed fear that monopoly local operators will dominate broadband access in the future, giving them an inordinate degree of control over both the Internet and content delivery. To the extent that ADSL provides an economically viable alternative, this is less likely to be a concern. Our results suggest that HFC networks have a significant advantage in capital costs. However, when marketing, billing, operations support, and other costs are taken into account, the difference may not appear so striking in percentage terms. In any case, recent announcements by major cable MSOs, AT&T and Times Warner, coupled with recent court rulings suggest that cable operators will eventually open up their broadband access networks to multiple ISPs.

Finally, the annual cost for an integrated broadband architecture providing both voice and data service, is not much different than the cost of a network providing voice service only (see Fryxell, Sirbu and Wanichkorn, 1998). Broadband networks, however, are far more flexible, easily allowing multiple "voice lines" to share a single access link. Moreover, when the local access network handles voice calls as IP, there will be no barrier to carrying all Interexchange traffic over IP networks at significant savings.

6 Acknowledgements

This research was supported in part by grants from Carnegie Mellon University's Information Networking Institute, the Internet and Telecommunications Convergence Consortium based at M.I.T. and by the Portuguese Foundation for Science and Technology. The opinions and conclusions contained in this paper are solely those of the authors and do not necessarily represent the views of the research sponsors.

7 References

- ADSL Forum, "ADSL Forum System Reference Model," TR-001, 1997.
- ANSI (American National Standard Institute), "Network and Customer Installation Interfaces – Asymmetric Digital Subscriber Line (ADSL) Metallic Interface," ANSI T1.413-1998, November 1998.
- Cable Television Laboratories Inc., "Data-Over-Cable Service Interface Specifications: Radio Frequency Interface Specification," SP-RFIV1.1-I04-000407, Fourth Interim Release, April 7, 2000.
- Cable Television Laboratories Inc., "PacketCable 1.0 Architecture Framework Technical Report," PKT-TR-ARCH-V01-991201, December, 1999.
- Cable Television Laboratories Inc., "PacketCable Audio/Video Codecs Specification," PKT-SP-CODEC-I01-991201, First Interim Release, December, 1999.
- CED Magazine, "Modernizing Network Power: Antiquated Approaches Set MSOs Up for Failure," January 2000.
- Ciciora, W., Farmer, J. and Large, D., "Modern Cable Television Technology: Video, Voice, and Data Communications," Morgan Kaufmann Publishers, San Francisco, 1999.
- Donaldson, G. and Midkiff, R., "Harnessing Network Power," *CED Magazine*, February 1999.
- Dutta-Roy, A., "A Second Wind for Wiring," *IEEE Spectrum*, September 1999.

- Fryxell, D., Sirbu, M. and Wanichkorn, K. "An IP-Based Local Access Network: Economic and Public Policy Analysis," in Sharon Gillett and Ingo Vogelsang (eds.), *Competition, Regulation, and Convergence: Current Trends in Telecommunications Policy Research. Selected Papers from the 26th Telecommunications Policy Research Conference*. Lawrence Elbaum Associates, Mahwah, New Jersey, USA, 1999.
- Fryxell, D., "Broadband Local Access Networks: An Engineering, Economic and Public Policy Analysis of Cable Modems and ADSL," Carnegie Mellon University, Fall 2000.
- HAI Consulting Inc., "HAI Model Release 5.0a: Inputs Portfolio," January, 1998.
- HAI Consulting Inc., "HAI Model Release 5.0a: Model Description," February, 1998.
- Morgan, S., "The Internet and the Local Telephone Network: Conflicts and Opportunities," *IEEE Communications Magazine*, January 1998.