

Connecting Homes to the Internet: An Engineering Cost Model of Cable vs. ISDN

Sharon Eisner Gillett

Telemedia, Networks and Systems Group,
Laboratory for Computer Science; and
Research Program on Communications Policy,
Center for Technology, Policy and Industrial Development
Massachusetts Institute of Technology

Using the World Wide Web at 28.8 Kbps (or less) can be a frustrating experience: a multimedia page that takes a fraction of a second to download at Ethernet speeds takes many seconds at modem rates. Two enhancements to existing infrastructure have the potential to deliver more satisfactory residential Internet access: ISDN telephone service, and upgraded cable TV networks.

While ISDN dedicates bandwidth to each user, cable networks support a shared bandwidth approach similar to that used in computer Local Area Networks (LANs). This report describes the technologies and evaluates qualitative differences between the two approaches. It presents quantitative results of capital cost models based on case studies of the PSICable deployment in Cambridge, MA, and the Internet over ISDN service offered by Internex, Inc. in the San Francisco, CA area.

The report finds that cable's shared-bandwidth approach has superior economic characteristics. For example, 500 Kbps Internet access over cable can provide the same average bandwidth and four times the peak bandwidth of ISDN access for less than half the capital cost per subscriber. The economy of the shared bandwidth approach is most evident when comparing the per-subscriber cost per bit of peak bandwidth: \$0.60 for the 500 Kbps cable service versus close to \$16 for ISDN. Cable-based access also has better service characteristics: it can support both full-time Internet connections and higher peak bandwidths, such as a 4 Mbps cable service that provides thirty-two times the peak bandwidth of ISDN.

The report concludes with an analysis of the barriers to diffusion of cable and ISDN Internet access, including business and policy factors. It finds that the closed market structure for cable subscriber equipment has not been as effective as the open market for ISDN equipment at fostering the development of needed technology. Furthermore, monopoly control of residential communications infrastructure—whether manifest as high ISDN tariffs or simple lack of interest from cable operators—limits business opportunities for Internet service providers.

Acknowledgments

This thesis would not have been possible without the support of my husband, Walter, and the cooperation of my son, Max.

I thank my advisors, David Tennenhouse and Lee McKnight. David encouraged me from the beginning to pursue my interests, suggested the topic, and challenged me to produce a work of this magnitude. Lee provided much-appreciated encouragement and kept me sane throughout the process.

I would also like to thank Butler Lampson for showing me by example how to construct a logical argument and express it in writing.

A number of industry experts contributed their knowledge to my data collection efforts, and I thank them for their time and willingness to assist in this research. Particularly helpful were Will Biedron of Continental Cablevision; Robert Berger of Internex, Inc.; Jim Naro of BBN (NEARNet); Walter Mansell of MIT's Telecommunications Services; Ron Hoffman of MIT's Distributed Computing and Network Services; and Dr. David Clark of MIT's Laboratory for Computer Science. Of course, I remain responsible for any errors of fact or interpretation.

My fellow graduate research assistants made it fun. Special thanks for help along the way go to Joe Bailey, David Gingold, Tom Lee, Jae Roh, and Russell Rothstein of the Research Program for Communications Policy; and Chris Lindblad, Bill Stasior, and David Wetherall of the Telemedia, Networks and Systems Group of MIT's Laboratory for Computer Science.

Finally, I am most grateful for the financial support of my sponsors: ARPA's "Interface Technologies for Networked Computing" project, number N00174-93-C-0036; Siemens Corporation; and ARPA's "Design and Development of Distributed Video Systems" project, number #F30602-92-C-0019.

to Walter and Max

Table of Contents

Table of Contents	5
Chapter One: Introduction	7
1.1.....Comparing ISDN vs. cable	9
1.2.....Why focus on Internet access?	11
1.3.....Perspective	11
1.4.....This report	13
Chapter Two: The Internet	15
2.1.....Overview	15
2.2.....Functions required by Internet technology	23
2.3.....Limitations of the access line model	28
2.4.....Incremental improvements to existing infrastructure	30
Chapter Three: Cable technology	35
3.1.....Cable networks	35
3.2.....Internet access over the cable plant	42
3.3.....Evaluation of cable LAN as Internet access network	50
Chapter Four: ISDN technology	54
4.1.....Local telephone networks	54
4.2.....Internet access over the local telephone network	67
4.3.....Evaluation of ISDN as Internet access network	71
Chapter Five: Cost model of Internet over cable	78
5.1.....Overview	79
5.2.....Technology reference model	81
5.3.....Subscriber demand model	95
5.4.....Traffic model	97
5.5.....Summary	101
Chapter Six: Cost model of Internet over ISDN	102
6.1.....Overview	102
6.2.....Technology reference model	103
6.3.....Subscriber demand model	108
6.4.....Summary	109
Chapter Seven: Results	111
7.1.....Result variables	111

7.2.....Input parameters	113
7.3.....Average capital cost	115
7.4.....Penetration limitations	117
7.5.....Components of capital cost	119
7.6.....Sensitivity analysis	128
7.7.....Entry cost	150
7.8.....Summary	154
Chapter Eight: Conclusions	155
8.1.....Summary comparison of services	155
8.2.....Technology, business, and policy barriers	162
8.3.....Policy recommendations	169
8.4.....Suggestions for further research	173
8.5.....Summary and conclusion	176
Bibliography	179

Chapter One

Introduction

The Internet is a network of networks that interconnects computers around the world, supporting both business and residential users. In 1994, a multimedia Internet application known as the World Wide Web became popular. The higher bandwidth needs of this application have highlighted the limited Internet access speeds available to residential users. Even at 28.8 Kilobits per second (Kbps)—the fastest residential access commonly available at the time of this writing—the transfer of graphical images can be frustratingly slow.

This report examines two enhancements to existing residential communications infrastructure: digital local telephone service (ISDN), and cable television networks upgraded to pass bi-directional digital traffic. It analyzes the potential of each enhancement to deliver Internet access to residential users. It validates the hypothesis that upgraded cable networks can deliver residential Internet access more cost-effectively, while offering a broader range of services.

The research for this report consisted of case studies of two commercial deployments of residential Internet access, each introduced in the spring of 1994:

- Continental Cablevision and Performance Systems International (PSI) jointly developed PSICable, an Internet access service deployed over upgraded cable plant in Cambridge, Massachusetts;
- Internex, Inc. began selling Internet access over ISDN telephone circuits available from Pacific Bell. Internex's customers are residences and small businesses in the "Silicon Valley" area south of San Francisco, California.

To develop conclusions from these case studies, the data is analyzed using the engineering cost model methodology described by Reed.¹ First, reference models are developed for the incremental technology needed to deliver residential Internet access over each type of infrastructure. These models, developed using spreadsheet-style software, consist of a series of equations expressing the relationships between the different technological components.² The equations are parameterized by variables such as residential housing density and the average bandwidth needs of customers. Next, component costs are collected and incorporated into the models' equations in order to provide estimates of the capital cost of each type of access. (Operating costs are outside the scope of this report.) Finally, the sensitivity of each model's results to variations in the input parameter values is tested.

The results of this analysis support the hypothesis articulated above. As will be shown, cable-based Internet access can provide the same average bandwidth and higher peak bandwidth more economically than ISDN. For example, 500 Kbps Internet access over cable can provide the same average bandwidth and four times the peak bandwidth of ISDN access for less than half the cost per subscriber. In the technology reference model of the case study, the 4 Mbps cable service is targeted at organizations. With this model, the 4 Mbps cable service can provide the same average bandwidth and thirty-two times the peak bandwidth of ISDN for only 20% more cost per subscriber. When this reference model is altered to target 4 Mbps service to individuals instead of organizations, 4 Mbps cable access costs 40% **less** per subscriber than ISDN. (In the rest of this report, the case study's service is referred to as "Organizational 4 Mbps" and the re-targeted service as "Individual 4 Mbps.") The economy of the cable-based approach is most evident when comparing the per-subscriber cost per bit of peak bandwidth: \$0.30 for Individual 4 Mbps, \$0.60 for Organizational 4 Mbps, and \$2 for the 500 Kbps cable services—versus close to \$16 for ISDN. However, the potential penetration of cable-based access is constrained in many cases (especially for the 500 Kbps service) by limited upstream channel bandwidth. While the penetration limits are quite sensitive to several of the input parameter assumptions, the cost per subscriber is surprisingly less so.

Because the models break down the costs of each approach into their separate components, they also provide insight into the match between what follows naturally from the technology and how existing business entities are organized. For example, the models show that subscriber equipment is the

¹See (Reed, 1992), pp. 8-9, and (Reed, 1993), p. 1. See also (Johnson and Reed, 1990) and (Sirbu, Reed and Ferrante, 1989).

²The models were developed using the program Demos ("DEcision MOdeling Software") from Lumina Decision Systems, Inc. Demos provides an object-oriented graphical user interface to spreadsheet-style data manipulation, providing a straightforward way to express in equations the relationships among different model components.

most significant component of average cost. When subscribers are willing to pay for their own equipment, the access provider's capital costs are low. This business model has been successfully adopted by Internex, but it is foreign to the cable industry. As the concluding chapter discusses, the resulting closed market structure for cable subscriber equipment has not been as effective as the open market for ISDN equipment at fostering the development of needed technology. In addition, commercial development of both cable and ISDN Internet access has been hindered by monopoly control of the needed infrastructure—whether manifest as high ISDN tariffs or simple lack of interest from cable operators.

The rest of this chapter discusses the hypothesis stated above in more detail. After providing a framework for the rest of the document, it concludes with a perspective on the contributions of this work in relation to previous research.

1.1. Comparing ISDN vs. cable

ISDN and upgraded cable networks will each provide different functionality (e.g. type and speed of access) and cost profiles for Internet connections. It might seem simple enough to figure out which option can provide the needed level of service for the least cost, and declare that option “better.” A key problem with this approach is that **it is difficult to define exactly the needed level of service for an Internet connection.** The requirements depend on the applications being run over the connection, but these applications are constantly changing. As a result, so are the costs of meeting the applications' requirements.

Until about twenty years ago, human conversation was by far the dominant application running on the telephone network. The network was consequently optimized to provide the type and quality of service needed for conversation. Telephone traffic engineers measured aggregate statistical conversational patterns and sized telephone networks accordingly. Telephony's well-defined and stable service requirements are reflected in the “3-3-3” rule of thumb relied on by traffic engineers: the average voice call lasts three minutes, the user makes an average of three call attempts during the peak busy hour, and the call travels over a bidirectional 3 KHz channel.³

In contrast, data communications are far more difficult to characterize. Data transmissions are generated by computer applications. Not only do existing applications change frequently (e.g. because of software upgrades), but entirely new categories—such as Web browsers—come into being quickly, adding different levels and patterns of load to existing networks. Researchers can barely measure these patterns as quickly as they are generated, let alone plan future network capacity based on them.

³See (White, 1993), p. 13. The channel is actually 4 KHz wide, but only about 3 KHz are used to carry the signal, with the rest reserved as spacing between channels.

The one generalization that does emerge from studies of both local and wide-area data traffic over the years is that computer traffic is bursty.⁴ It does not flow in constant streams; rather, the level of traffic varies widely over almost any measurement time scale (Fowler and Leland, 1991). Dynamic bandwidth allocations are therefore preferred for data traffic, since static allocations waste unused resources and limit the flexibility to absorb bursts of traffic.

This requirement addresses traffic patterns, but it says nothing about the absolute level of load. How can we evaluate a system when we never know how much capacity is enough? In the personal computing industry, this problem is solved by defining “enough” to be “however much I can afford today,” and relying on continuous price-performance improvements in digital technology to increase that level in the near future. Since both of the infrastructure upgrade options rely heavily on digital technology, another criteria for evaluation is the extent to which rapidly advancing technology can be immediately reflected in improved service offerings.

I hypothesize that cable networks satisfy these evaluation criteria more effectively than telephone networks because:

- Coaxial cable is a higher quality transmission medium than twisted copper wire pairs of the same length. Therefore, fewer wires, and consequently fewer pieces of associated equipment, need to be installed and maintained to provide the same level of aggregate bandwidth to a neighborhood. The result should be cost savings and easier upgrades.⁵
- Cable's shared bandwidth approach is more flexible at allocating any particular level of bandwidth among a group of subscribers. Since it does not need to rely as much on forecasts of which subscribers will sign up for the service, the cable architecture can adapt more readily to the actual demand that materializes.
- Telephony's dedication of bandwidth to individual customers limits the peak (i.e. burst) data rate that can be provided cost-effectively. In contrast, the dynamic sharing enabled by cable's bus architecture can, if the statistical aggregation properties of neighborhood traffic cooperate, give a customer access to a faster peak data rate than the expected average data rate.

⁴For examples of this conclusion, see (Shoch and Hupp, 1980), (Fowler and Leland, 1991), and (Claffy, Polyzos and Braun, 1992).

⁵Such savings should be particularly evident when comparing operational expenses, which are outside the scope of this report.

1.2. Why focus on Internet access?

Internet access has several desirable properties as an application to consider for exercising residential infrastructure. Internet technology is based on a peer-to-peer model of communications. Internet usage encompasses a wide mix of applications, including low- and high-bandwidth as well as asynchronous and real-time communications. Different Internet applications may create varying degrees of symmetrical (both to and from the home) and asymmetrical traffic flows. Supporting all of these properties poses a challenge for existing residential communications infrastructures.

Internet access differs from the future services modeled by other studies described below in that it is a real application today, with growing demand.⁶ Aside from creating pragmatic interest in the topic, this factor also makes it possible to perform case studies of real deployments.

Finally, the Internet's organization as an "Open Data Network" (in the language of (Computer Science and Telecommunications Board of the National Research Council, 1994)) makes it a service worthy of study from a policy perspective.⁷ The Internet culture's expectation of interconnection and cooperation among competing organizations may clash with the monopoly-oriented cultures of traditional infrastructure organizations, exposing policy issues. In addition, the Internet's status as a public data network may make Internet access a service worth encouraging for the public good. Therefore, analysis of costs to provide this service may provide useful input to future policy debates.

1.3. Perspective

This report explores the cost of connecting homes to the Internet through cable and ISDN infrastructures and discusses the technology, policy and business issues involved. It applies the engineering cost model methodology used in (Reed, 1993), (Reed, 1992), (Johnson and Reed, 1990), and (Sirbu, Reed and Ferrante, 1989) to a new communications service: Internet access.

The works referenced above focus on future infrastructure capital costs for telephone and cable networks to deliver telephony, distributive video (e.g. broadcast television), and switched video services (e.g. video on demand).

⁶See (Gillett, 1995) for a discussion of the Internet access market, including residential access. (Flynn, 1995) estimates this market at \$135 million, while The Maloff Company, a Michigan consultancy focusing on business and the Internet, estimates that the Internet access market grew from \$119 to \$521 million from March 1994 to January 1995. This data is derived from summaries of Maloff's "Internet Service Provider Marketplace Analysis" (1993-4 and 1994-5) which are available on the Internet at <http://www.trinet.com/maloff>.

⁷As a practical matter, the non-proprietary nature of Internet technology also makes Internet service easier to research than commercial on-line services, for which the local access technologies described in this report can also be used. In some cases, commercial and Internet services may share use of the same infrastructure.

They focus on the question of integration: are there economies of scope or scale for either cable or telephone companies to provide multiple services over a single "Integrated Broadband Network?" This report answers a different question: do the architectural differences between existing cable and telephone networks make either one a more cost-effective Internet access medium? In so doing, it continues the exploration of architectural differences in (Gillett, Lampson and Tennenhouse, 1994) and (Tennenhouse, Lampson, Gillett et al., 1995), but with a narrower focus on cable and ISDN as the infrastructures, and Internet access as the application. It also extends the qualitative discussion of architectural differences between bus and star networks found in (Calhoun, 1992).

In exploring the question of architectural differences stated above, this report takes existing infrastructure costs as given, except for capital investments specifically needed to support the Internet access application. It contributes technology reference models for those investments, identifying the assumptions and parameters that can cause cost estimates to vary. In the sense that it isolates the incremental investments needed to support Internet access, it builds on the work described in (Mitchell, 1990) for telephony.

In developing the technology reference models, this report brings together two technical literatures that are often separate: infrastructure technology (cable and telephone networks) and Internet technology. (Berger, 1994) offers an exception to this separation, providing a brief user-level overview of Internet connections delivered over a wide variety of access media. Building on that article, this report delves into more technical detail for two particular residential access media, namely cable and ISDN. This report also builds on (Estrin, 1982), which integrates data communications with cable and telephone network technology but was written before the rapid diffusion of Internet-style networking. Much of Estrin's description of data-over-cable technology is still remarkably accurate. New since 1982, however, are the diffusion of home computers and the development of real computer networks, such as the Internet and commercial on-line services, to which residential subscribers can actually connect.

The technology reference models developed in this report are exercised using 1994 cost data gathered from two case studies of specific implementations of residential Internet access. The importance of this data, which is of necessity limited in its generality and lifetime, lies in its ability to support quantitative comparison of costs. Not absolute but **relative** costs are central to the hypothesis and its validation. Other researchers developing related models have stated caveats that apply equally well to this work:

It should be stressed...that there is still a great deal of uncertainty in the industry about costs. The technology is moving rapidly and costs are falling. At the same time there is a great deal of

hype, particularly in the media, about the capabilities of the technology and the speed at which new technology will be introduced. While every effort has been made to establish 'reality' nevertheless it should be recognised that estimating costs, particularly beyond the next few years is a difficult exercise. Caution should be exercised in using the estimates presented. (Communications Futures Project, 1994), p. 3.

The value of this work lies less in its cost data than in the analytical framework it develops through its models.

This report's focus on cable and ISDN infrastructures and the Internet access application differentiates it from the modeling performed by the Communications Futures Project of the Australian Bureau of Transport and Communications Economics, described in (Communications Futures Project, 1994), (Communications Futures Project, 1994), (Communications Futures Project, 1994), and (Communications Futures Project, 1994). That project has developed cost models for a variety of different infrastructure technologies, including Hybrid Fiber Coax (HFC), Fiber to the Home (FTTH), Asymmetrical Digital Subscriber Line (ADSL), and a variety of wireless broadcast technologies.⁸ The models consider a wide variety of general applications, including telephony, video, and so-called interactive services. Incorporating a service diffusion model of future subscriber demand, the models forecast Net Present Value results for a variety of evolutionary scenarios. The incorporation of demand modeling distinguishes the Communications Futures Project from other infrastructure modeling efforts, including this one. Their series of reports presents a much more extensive range of results than this report, which instead provides significantly more cost and technology detail for its fewer areas of focus.

This report considers providers' capital costs only; neither operational costs nor pricing are discussed. As such, this report may provide input to, but is not directly part of, the emerging "Internet economics" literature dealing with flat vs. usage-based Internet pricing policies. (Bailey and McKnight, 1995) reviews this literature and discusses future directions in this field.

1.4. This report

This report investigates the hypothesis that the shared-bandwidth architecture of cable networks enables them to provide more satisfactory and cost-effective Internet access than the dedicated-bandwidth architecture of ISDN. Before this hypothesis can be explored, it must be fully understood. Accordingly, Chapters 2 through 4 provides background on the multiple technologies involved in providing access to the Internet over cable and ISDN. Chapter 2 discusses the Internet. Chapter 3 describes cable

⁸These acronyms are explained in more detail in Chapters 3 and 4.

infrastructure and how Internet access can be provided over it, while Chapter 4 serves the same function for ISDN.

The methodology employed for this research begins with the construction of engineering cost models for cable and ISDN Internet access. These spreadsheet-style models are based on capital cost data collected from case studies of two commercial deployments. Chapters 5 and 6 describe the details of the cost models of cable and ISDN Internet access, respectively. Where Chapters 3 and 4 discuss technology issues at a general level, Chapters 5 and 6 describe the specific implementations used in the two cases studied. Given this structure, the reader who is less interested in the technological “whys” behind these implementations can skip Chapters 3 and 4 and still understand the cost models.

The cost models produce two kinds of results: quantitative cost comparisons and their implications for business strategy and public policy. Chapter 7 reports the quantitative results. It compares the cost of Internet access over cable and ISDN under an original set of parameter values, then examines the sensitivity of these results to changes in these values. Chapter 8 considers the implications of these results, including a discussion of the technological, business, and policy barriers to rapid diffusion of either type of Internet access. It concludes with a set of policy recommendations and suggestions for further research.

Chapter Two

The Internet

When a home is connected to the Internet, residential communications infrastructure serves as the “last mile” of the connection between the home computer and the rest of the computers on the Internet. This chapter describes the Internet technology involved in that connection, while Chapters 3 and 4 discuss the technologies involved in the two different residential communications infrastructures, cable and ISDN, considered in this report. This chapter does not discuss other aspects of Internet technology in detail; that is well done elsewhere.⁹ Rather, it focuses on the services that need to be provided for home computer users to connect to the Internet. As background, it begins with an overview of the Internet, including a definition of Internet access (also referred to as a connection to the Internet).

2.1. Overview

The Internet is a collection of computer networks that interconnect by using a common addressing scheme and suite of communication protocols. The protocol suite is known by the name of its transport and network layer components: the *Transmission Control Protocol and Internet Protocol (TCP/IP)*.¹⁰

An important—and often misunderstood—aspect of the definition of the Internet is that the physical connections between networks are unspecified. In fact, they are quite heterogeneous, ranging from 1200 bits per second (bps) dialup telephone links to dedicated leased 45 Mbps or faster circuits. It is therefore misleading to speak of “the speed of the Internet.” The Internet is a service overlay on top of a heterogeneous, international physical

⁹ See, for example, (Comer, 1991). Internet technology specifications are published as “Request for Comments” documents (RFCs) by the Internet Engineering Task Force. RFCs may be found in a number of locations on the Internet, including <http://www.internic.net>. These documents update some of the protocol descriptions in Comer’s book.

¹⁰This thesis uses the convention of italicizing terms when they are first defined. See (Comer, 1991) for a thorough description of TCP/IP and related Internet protocols.

communications infrastructure. The speed with which data can traverse the network is largely determined by the speed of the underlying communication links.

It is also important to distinguish internets from “The Internet.”¹¹ An internet is a generic term referring to any interconnected set of networks. The Internet refers to a particular connected set of networks, running TCP/IP, with a single, globally administered address space. An enterprise, such as a corporation or university, may create an internal internet using any protocols or addressing scheme it wishes. If this internet is to be physically connected to the global Internet, however, at least some of its computers must be able to handle data formatted according to the Internet Protocol (IP). These computers must also be assigned IP addresses that do not conflict with those of other computers on the Internet.

There is no single, accepted criteria for deciding whether a particular network or computer is part of the Internet. Instead, different types of access fall along a continuum depending upon the level of service they provide. As Crocker observes:

The Internet is many things to many people. It began as a technology and has grown into a global service. With the growth has come increased complexity in details of the technology and service, resulting in confusion when trying to determine whether a given user is “on” the Internet. Who is on the Internet? What capabilities do they have?¹²

He then proposes four categories of end-user access, each offering a different level of Internet service:

1. *Full access*: the attached computer runs TCP/IP and other key Internet protocols, and is always attached to the Internet.
2. *Client access*: the user’s computer runs Internet applications, but need not always be connected or able to support all Internet protocols.
3. *Mediated access*: the user’s computer runs no Internet applications. Instead, it connects to a server that runs those applications on its behalf and relays the results.

¹¹Definitions of different kinds of internets (e.g. enterprise internets) vs. The Internet are found on slides produced by Tony Rutkowski of The Internet Society, which are available on the Internet by using the File Transfer Protocol (FTP) to connect to ftp.isoc.org as user “anonymous,” and downloading the file isoc/charts/internet3.ppt.

¹²From (Crocker, 1995), p. 1.

4. *Messaging access*: the user is restricted to mail-based access to Internet services.

For example, the major on-line services (such as Prodigy, CompuServe, and America OnLine) originally offered their subscribers only messaging access to the Internet, but are gradually moving toward the mediated access model. As will be discussed below, the analysis in this report considers only full and client access.

2.1.1. Internet services

A variety of services are available to the user whose computer has full or client access to the Internet, with new ones appearing frequently. Most Internet services can be classified into four categories:

- *Messaging* services enable people to communicate with other people via their computers. Examples include electronic mail (email), discussion groups (also called newsgroups or bulletin boards), and videoconferencing.
- *Remote computer access* allows people to use physically remote computers. For example, a software engineer might log into a computer at a customer site to observe a reported malfunction.
- *Information publication and retrieval* services are provided by a variety of programs that let people transfer computer data from one computer to one or many others. Given the generic nature of computer data, any type of digital information can be transferred: computer software, library catalogs, formatted research papers, graphics-laden product literature and price information, purchase orders, photographs, audio programs, video clips, etc. The Internet thus serves as an enormous, instantaneous library containing any information that someone somewhere has chosen to make available electronically to other Internet users. Anonymous file transfer (FTP) and the hypertext-linked pages of the World Wide Web are examples of such systems.
- *Navigation*: services help users find what they need in this vast library; they are the “phonebooks” of the Internet. For example, client programs for browsing the World Wide Web allow users to find and retrieve particular types of information (such as “photos of the Rocky Mountains”) from server databases anywhere on the Internet.¹³ Searchable email address directories provide another kind of navigation service.

While Internet messaging has always been popular, the “point-and-click” interfaces of Web browsers have extended the Internet’s appeal to a larger

¹³The server databases are called *Web pages*. Popular Web client programs, or *browsers*, at the time of this writing include Mosaic and Netscape.

number of computer users.¹⁴ They have also provided an obvious mechanism for making commercial information (such as restaurant listings or a florist's offerings) readily available to consumers, setting the stage for increased demand for connections from residential users.

2.1.2. Connecting to the Internet

Although the technical definition of the Internet is based solely on logical protocols that computers agree to use in common, some form of physical media, whether wire-based or wireless, must be used to connect each constituent user and network. This section defines a user connection to the Internet, starting with the media-independent aspects and moving on to a discussion of physical attributes.

Media-independent definition

At the virtual level, an Internet connection consists of protocol processing capability residing in two distinct computers. The first is the computer that connects to the network (the *host* computer); with full or client access to the Internet, this host is the Internet user's computer. The second is the switching computer (called an *IP router*) that connects the host to the rest of the Internet. A computer has full or client access to the Internet if:

- It has an Internet Protocol address that is unique among all other computers connected to the Internet;
- It can generate and parse data in the format, known as packets, expected by the Internet Protocol (in other words, it can send and receive IP packets);
- It is directly or indirectly connected to an IP router that knows how to forward IP packets to and from all other computers on the Internet. If it is always connected, then the computer has full access. If it is connected only sometimes, then the computer has client access.

Together, these capabilities give any one Internet-connected computer the ability to communicate with any other as a *peer*. For example, a host that is connected to the Internet according to the above definition can participate in any application that depends on the exchange of IP packets, including acting as a World Wide Web information server, running a World Wide Web client program such as Mosaic or Netscape, or using X-Windows to display the results of a computation taking place on another computer. **A computer that does not have both an IP address and the ability to exchange IP packets cannot**

¹⁴See (Markoff, 1993). At the beginning of 1994, 2.2 million computers were measured on the Internet, with this number growing approximately 8-10% per month. At the end of 1994, close to 4 million computers were connected and the growth rate had increased. See graph of "Internet Hosts Reachable" in (Rutkowski, 1994), p. 11, and "The Internet - Current Dynamics" in (Rutkowski, 1993), p. 3, as well as Rutkowski slides available via anonymous ftp as noted above.

participate directly in any of these distributed Internet applications. This is true even if physical wires—such as a dialup telephone connection to an Internet host—link the computer to the Internet.

This distinction illustrates a major limitation of much current residential access to the Internet. Most home users do not have IP addresses and thus do not connect their home computers directly to the Internet. Instead, they connect to a *proxy* Internet host, using their home computer as a remote terminal.¹⁵ The proxy Internet host can display Mosaic screens; the home computer cannot. The proxy host receives file transfers, email and netnews messages; the home user must download and read such transfers in a separate, relay step. In Crocker's terminology, the access is mediated.

The proxy computer may be an Internet host at the user's office, or a commercially-provided proxy.¹⁶ Typically, the home user connects to the proxy Internet host by dialing in to it over a residential telephone line. Although proxy connections are better than none, and an enterprising service provider can supply software to simplify the relay step or emulate the look and feel of the most popular Internet programs such as Mosaic, the effort required to provide this extra layer means that the services available to proxy customers will always lag behind and be more limited than those available to customers with full or client access. Consequently, this report considers only full and client levels of residential Internet access; cost models of mediated and messaging levels of access are beyond its scope.

Although most proxy Internet connections are provided over dialup telephone lines, it is important not to equate or confuse the connection model with the physical connection type. Protocols, such as the Serial Line Internet Protocol (SLIP) and the Point-to-Point Protocol (PPP), exist that allow Internet Protocol packets to be sent over dialup telephone lines.¹⁷ A computer running IP and these protocols can gain client access to the Internet using a dialup telephone line. Without these protocols, the same telephone line can provide only mediated or messaging levels of access.

¹⁵The term "proxy" was suggested by Professor David Tennenhouse of MIT's Laboratory for Computer Science.

¹⁶Many commercial providers of proxy Internet connections have appeared over the past few years. For example, one of the first and most widely used ones in the Boston area is "The World" (world.std.com), operated by Software Tool and Die, Inc. Delphi's Internet service also falls into this category at the time of this writing.

¹⁷SLIP is described in (Romkey, 1988); PPP in (Simpson, 1993). As an example of how a dialup circuit may still provide better than proxy access, consider that the Appletalk Remote Access program for Macintosh computers incorporates SLIP and allows remote Macintosh users to run Mosaic over a dialup link.

Physical connections

It was noted above that for a computer to be connected to the Internet, it must be directly or indirectly connected to an IP router that knows how to forward IP packets to and from all other computers on the Internet. Such a router is located at an *Internet Point of Presence (IPOP)*, a physical Internet access point operated by an organization commonly referred to as an Internet service (or access) provider.¹⁸ Generally hidden from the user, the physical connection between a user's computer and the IPOP may be quite complex and involve multiple types of equipment and transmission media.

Figure 2.1 illustrates the physical connection model for organizational Internet users.¹⁹ In the organizational model, the user's computer is assumed to be connected to an enterprise network, consisting of one or more interconnected Local Area Networks (LANs). With this architecture, the Internet provider only needs one physical connection to the enterprise network, instead of a separate connection to each individual computer.

¹⁸Most of the early Internet access providers were non-profits; some of these have transitioned to for-profit companies, and many new for-profit ventures have entered this industry. (Gillett, 1995) provides a brief discussion of the Internet access provider industry.

¹⁹The term "organizational" includes business users as well as users in non-profit organizations such as universities and government offices.

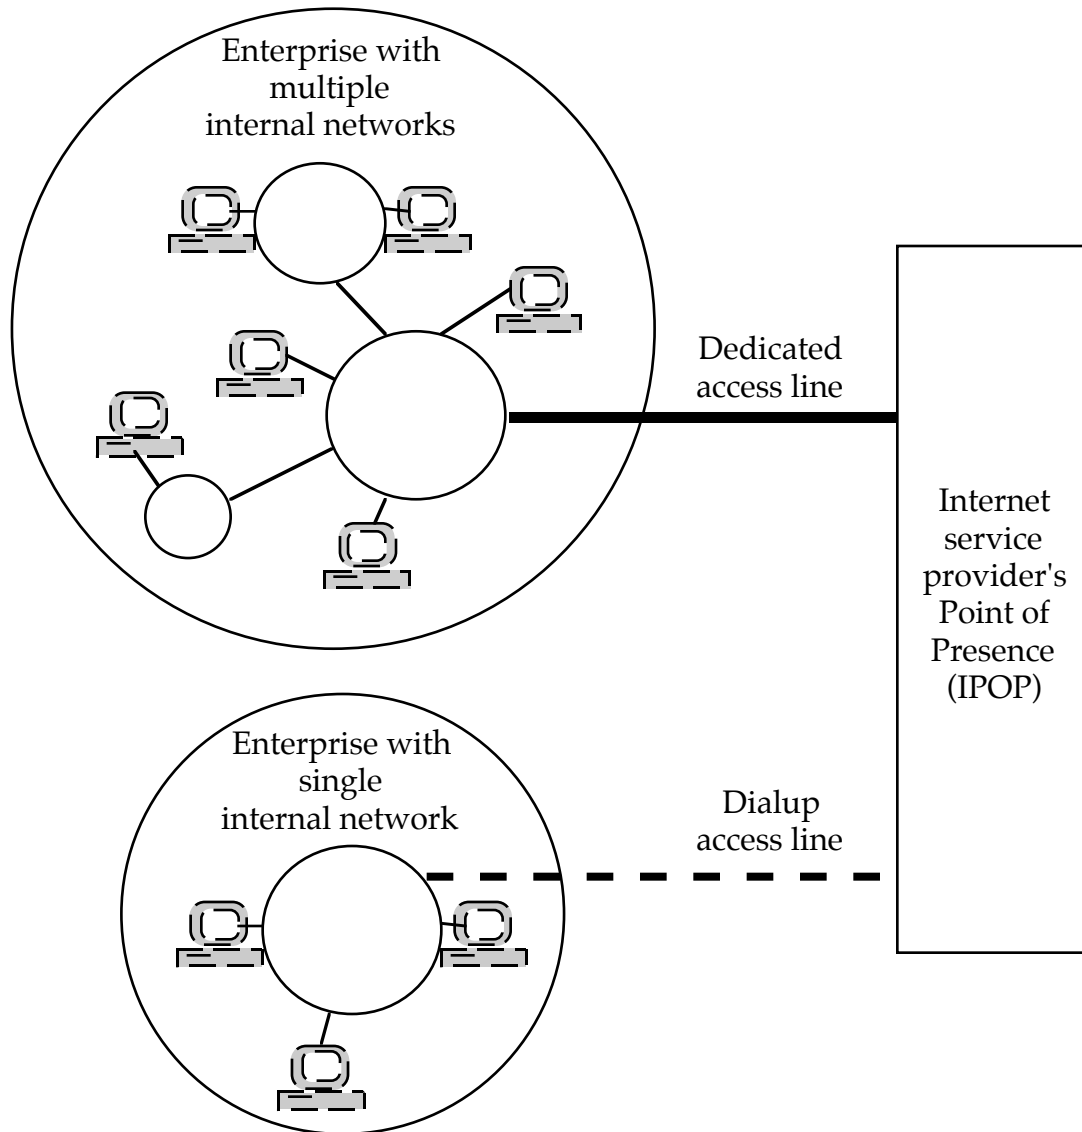


Figure 2.1: Model of connections to users within organizations

This physical connection, referred to as an *access line*, consists of a telecommunications circuit used for data transmission. It may be a dedicated circuit, always available for the exchange of Internet traffic.²⁰ Alternatively, it may be a dialup circuit, in which case it may only be available certain hours of

²⁰Dedicated telephone circuits may be leased from the Local Exchange Carrier (e.g. NYNEX) or from a Competitive Access Provider (e.g. MFS). Commonly available bandwidths include 56 or 64 Kbps and multiples thereof ("fractional T1" service), and 1.5 Mbps (T1 service), with higher bandwidth connections costing more. In some areas, 10 Mbps circuits can be purchased or leased from private microwave providers; fractional T3 services may also be available. The expense of dedicated circuits makes them cost-effective only for high-volume users, for whom the usage charges on a metered dial-up connection would cost more.

the day, or on an as-needed basis.²¹ For example, a call may automatically be dialed when packets are waiting to be sent to the Internet. In either case, both the bandwidth and the cost of the circuit are shared by all of the enterprise's users.

When the access line model is extended to individual users as shown in Figure 2.2, this sharing disappears. Each individual computer requires its own access line. In most cases, this line consists of a dialup telephone circuit.




Figure 2.2: Access line model of connections to individual residential users

A comparison of Figures 2.1 and 2.2 also illustrates a business difference between the two connection models. An organization may have internal technical support resources, providing a support interface for the Internet provider to deal with that is both geographically and organizationally centralized. In contrast, when serving individuals, there is no enterprise to provide a technical support buffer. Consequently, an Internet access provider that serves individual subscribers needs a much stronger—and possibly more

²¹The speed of an analog dialup circuit depends on the modems attached to each end of the circuit. The fastest modems available at the time of this writing can transfer data at 28.8 Kbps. If the modems at either end do not support the same data rates, the slower rate of the pair is used.

geographically dispersed—capability for offering technical support to customers.²²

2.2. Functions required by Internet technology

The discussion above has already touched upon two of the functions that must be performed by a service provider offering Internet access to residential customers: operating one or more IP routers as part of an Internet Point of Presence, and providing technical support to individual customers.

The routing function involves knowing how to forward Internet Protocol (IP) packets to and from all other computers on the Internet. This function, which is closely related to IP addressing, is discussed in more detail below. In addition to operating one or more IP routers, the provider must physically connect each individual customer to an IP router using the residential communications infrastructure. The details of this connection are described for cable in Chapter 3 and for ISDN in Chapter 4.

2.2.1. Technical support and network management

The technical support function includes answering individual customer's technical questions and assisting them with installation, configuration, and troubleshooting of their networking hardware and software. It also involves diagnosis of connectivity and loading problems. To support this function, providers typically operate network management equipment and monitoring tools. Diagnosing infrastructure connectivity or load problems requires the cooperation of the residential infrastructure provider; similarly, tracking down Internet problems requires cooperation between the Internet access provider and other constituent Internet network operators (i.e. other Internet providers).

2.2.2. Optional servers

The Internet access provider may also choose to provide servers to translate Internet names into addresses, and to support application services. These functions are now described in more detail.

Logical addressing

Internet computers use the *Domain Name Service* to translate between the names humans use to refer to their computers and numeric Internet Protocol (IP) addresses. Domain names are hierarchical, reflecting the organizational structure of the identified computer. For example, the domain name `clove.lcs.mit.edu` arises from a computer named "clove" that is part of the Laboratory for Computer Science (lcs), which is part of the Massachusetts Institute of Technology (mit), which is an educational (edu) institution.

²²The support needs of telephone users behind a PBX vs. an equivalent number of residential telephone subscribers provides an analogy for understanding the increased and more geographically dispersed support load.

Domain name service is implemented by a hierarchy of name servers. When a computer is added to the Internet, it is configured with the IP address of a domain name server. Then when the computer needs to communicate with a destination it knows by name but not by IP address, it asks that server to translate the name. The server performs the translation either based on its own table entries or by transparently requesting further information from other servers in the hierarchy.

Domain name translation is provided by a number of distributed servers that can be anywhere on the Internet; they do not have to be located at an Internet Point of Presence or organizationally affiliated with an Internet access provider. An Internet provider is therefore not required to supply its own domain name server. However, if a provider adds a large number of users to the Internet, increasing the number of registered domains and the load on existing servers, it may also need to add a domain name server if it is to be a good Internet citizen.

Applications

In a bulletin board system or commercial on-line service (such as Prodigy, CompuServe or America OnLine), supplying content is a large part of the service provider's job. The system's centralized authority structure ensures that users only have access to whatever content the service provider makes available on or through the system.

In contrast, the Internet's open organizational structure ensures that anyone who creates or uses an application that runs with the TCP/IP suite of protocols can supply content on the Internet. Therefore, supplying content is much less a part of the Internet service provider's job. The provider may choose to make particular content available just as any other Internet site may, but this function is not required of an Internet Point of Presence.

Two differences between the home and workplace environments, however, make it more likely that a residential Internet provider would assist at least some subscribers with access to content. First, home computers do not typically have as much disk storage as workplace servers. Network news feeds, for example, can require large amounts of storage: a feed of all newsgroups can generate on the order of a gigabyte of material per day.²³ Workplaces therefore configure servers to receive the news feed and assume that users will select what they want to read off the server. In the residential environment, the workplace server is available to telecommuters, but not all customers fall into that category. Some customers may choose to operate

²³Conversation with William Ang of Computing Resource Services, MIT Laboratory for Computer Science. Newsgroups offering executable code ("binaries") require the most space.

their own servers,²⁴ while others either “borrow” or pay for the use of servers operated anywhere on the network. The residential Internet provider is likely to try to serve the latter group of customers, since they represent an additional source of revenue.

Second, home computers are generally not powered on 24 hours a day. This can pose problems for the receipt of electronic mail, which can arrive at any time. Mail delivery will be more reliable if it is directed to a server that is always running; users then retrieve their mail from the server at their convenience, using a protocol such as the Post Office Protocol. Like a network news server, an email server can be operated by a workplace, an individual customer, or some form of service provider. As with network news service, the residential Internet provider is not required to provide email service, but it makes business sense to offer a complete package to the customers who require it, especially given the economy of scope of offering domain name service, network news, and email from the same server.²⁵

2.2.3. Addressing

Each computer on the Internet is assigned a unique Internet Protocol (IP) address. An Internet access provider typically manages a block of IP addresses from which it must assign individual addresses, or groups of addresses, to its customers. The details of how these assignments relate to customers’ physical network configurations can be complex;²⁶ they depend on the specific IP addressing scheme in use on the Internet. At the time of this writing, Internet addressing follows the scheme described in (Comer, 1991). However, a new addressing scheme is under development and is expected to be phased in shortly along with the transition to the next generation Internet Protocol.²⁷ Given the uncertainty surrounding the new scheme, address assignment is not discussed further here.

²⁴A server can be selective about which newsgroups it chooses to receive. For example, telecommuters may choose to subscribe to personal-interest newsgroups that their workplaces decline. Since a gigabyte of disk for a PC-class computer costs about under \$500 at the time of this writing, it is not out of the realm of possibility that a user would choose to store his or her own complete newsfeed, and certainly feasible for a user to store selected newsgroups.

Conversation with Mark Graham of Pandora Systems, San Francisco, CA, April 11, 1994.

²⁵For example, a residential Internet subscriber who is an employee of a large organization could choose to retrieve business-related mail from a workplace server, while relying on the residential Internet provider for personal email.

²⁶These details are described in (Comer, 1991). Address assignment can pose some challenges for the Internet provider when there are customers who operate multiple internal networks, or customers with part-time access requiring dynamic address assignment.

²⁷See (Metcalfe, 1995). At the time of this writing, many draft Requests for Comments (RFCs) describing various aspects of the IPng (IP Next Generation) protocol are available at <http://www.internic.net>.

2.2.4. Routing

This section describes the Internet routing functions that need to be performed by both the customer's networking equipment and the Internet access provider.

In its simplest form, internetworking of physically separate computer networks is accomplished as shown in Figure 2.3. A special computer called a router is physically connected to each of the separate networks. The router is supplied with intelligence that allows it to take traffic from Network 1 and deliver it to its intended destination on Network 2, and vice-versa. This intelligence takes the form of protocol processing software and tables of addresses and/or paths through the networks.²⁸




Figure 2.3: Conceptual view of internetworking

In practice, the networks to be interconnected are not always located in the same physical vicinity. In that case, the single router of Figure 2.3 is replaced by two routers: one connected to each network, and each connected to the other by some form of telecommunications line. Figure 2.4 illustrates the two-router model in the context of connecting simple networks of users to an Internet access provider's more complex network. (The Internet access provider's network is in turn connected, again by routers, to the rest of the Internet. These connections are assumed to take place within the Internet "cloud" shown in Figure 2.4.)

²⁸These tables are usually dynamically computed by the router using distributed algorithms such as Shortest Path First. See (Comer, 1991) for more details.




Figure 2.4: Internet access configuration

The customer and provider routers shown in Figure 2.4 have asymmetric duties. For example, consider what each type of router has to do to send traffic from the user to the rest of the Internet. The customer's router has only to follow a simple rule: send the traffic to the provider's router.²⁹ The provider's router has the harder job. It must know how to direct traffic from the customer's network to anywhere else on the Internet (including other customers' networks).³⁰

Because of the assumption of a simple customer network, the customer's router also has a simple job to deliver traffic to a customer computer. The provider's router, on the other hand, must be intelligent enough to choose which of the different customer routers should receive that traffic.

Hence when the customer is assumed to connect a simple network to the Internet, the customer's router can be quite simple. This observation is critical to making residential Internet access economical, since each residence

²⁹Were a subscriber to be connected to multiple competing Internet access providers, the problem gets harder. Some mechanism needs to be devised for the customer to specify to the router which provider should be used for any given packet. Since no such mechanisms are currently built into the standard Internet protocols, out-of-band mechanisms would need to be used. (Conversation with Dr. David Clark, January 28, 1994.)

³⁰As (Comer, 1991) describes, Internet routing is hierarchical. One way to be able to reach everyone else on the Internet is to participate in the top level of the hierarchical routing scheme. Alternatively, the access router can participate at a lower level, in which case it only knows how to direct traffic to another router that then knows how to reach everyone else on the Internet. The distinction affects the amount of table space needed in the access provider's router, and thus its cost. Also note that a single Internet access provider may operate more than one access router, and the different routers may operate at different levels of the routing hierarchy.

is a customer requiring a router. In fact, many residences have only a single computer, which is a degenerate case of a single, simple network.³¹

This section has described the Internet technologies involved in supporting the access line model. The next section discusses the limitations of that model, especially for residential Internet access.

2.3. Limitations of the access line model

Residential users are customers potentially connecting to the Internet from residential areas. These include:

- Small businesses, such as startup companies or professional offices (doctors, dentists, psychological counselors, attorneys, independent consultants, etc.);³²
- Telecommuters, whose costs are generally borne by their employers;
- Individuals in their non-workplace roles, such as citizen, consumer, parent, or hobbyist.

Table 2.1 characterizes the types of organizational and individual connections made possible by the access line model. As its two empty boxes illustrate, customers that need either the occasional ability to send traffic quickly, or a full-time, low-bandwidth connection, are not served by either type of connection model.

³¹Connecting individual computers (as opposed to networks of computers) also raises the possibility of using a customer bridge instead of a customer router. (Perlman, 1992) discusses in detail the distinctions and tradeoffs between bridging and routing. Briefly stated, a bridge would make multiple residences appear to be on the same physical network, while a router would make each residence appear to be on a separate physical network. Sharing a single physical network may be less than ideal given that there is no particular reason for trust among multiple unaffiliated households. This contrasts with the workplace situation, in which the people sharing a network are assumed to at least work for the same organization. The extent to which sharing a single network among households poses a security problem depends on the particular infrastructure, as discussed in later chapters.

³²In marketing jargon, this segment is referred to as SOHO, for "Small Office/Home Office."

Table 2.1: Characteristics of physical connections enabled by access lines

Peak Bandwidth	Availability	
	Part-time	Full-time
low to medium	organizational or individual model: dialup circuit	
medium to high		organizational model: dedicated circuit

Unfortunately, the missing entries correspond to many of the needs of residential users. For example, a telecommuter may need high-speed access to emulate the office LAN environment, but only one day each week. A home-based consultant may hold a regular video conference with a client, requiring high-bandwidth access only one hour each day. Alternatively, consider a startup business that wishes to make its product information available 24 hours a day to potential customers, by operating a World Wide Web information server that is always available. This customer would desire full-time Internet access, but would be unlikely to be able to afford the expense of a leased line. Additionally, such a dedicated resource would be largely wasted since such a customer is unlikely to generate large traffic volumes.

These needs are also unmet for non-residential users. For example, some organizations lease a dedicated Internet access line even though their aggregate traffic volume does not justify it, because they need either the full-time or the high-bandwidth aspect of the connection.³³

On average, residential users can be expected to generate less aggregate traffic volume than larger organizations with more users. Therefore, dedicating circuits to them is even more inefficient than dedicating circuits to low-volume organizations needing full-time access. Inefficiency aside, the expense of a dedicated circuit is simply beyond the budget of most residential users.

The fundamental problem with today's access models for residential data users is that they were designed for the voice telephony application. Dialed calls, with a relatively long connection establishment delay and all-or-nothing use of the access circuit, are a fine model for voice conversations between people. But computer data is much less predictable than human conversation. Computers typically generate data rapidly, in short bursts, then

³³For example, a company that expects its customer service representatives to respond quickly to electronic mail requests has a strong need for a full-time connection. Alternatively, an organization that needs to run only one high-bandwidth application can quickly exceed the capacity provided by dialup.

fall silent. They need to be able to receive bursts of data whenever they arrive, whether or not a person is sitting at the computer. Thus the all-or-nothing connection model is highly inappropriate for data traffic. This is especially true when only one computer, or a small number of computers, uses a communications line, and there is very little traffic aggregation to smooth out the individual bursty flows.

2.4. Incremental improvements to existing infrastructure

Given that the access line model is less than ideal for residential Internet access, what can augment or replace it? This section provides a brief overview of the enhanced technologies that the rest of the report focuses on for providing more satisfactory access. These technologies are discussed in detail in Chapters 3 and 4.

Although extending fiber optic links to all residences would alleviate today's bandwidth problems, two economic realities make this solution impractical in the short-term. The first problem is the high absolute cost of the required equipment, lasers in particular.³⁴ The second problem relates to the up-front nature of the investment: it is not feasible to upgrade one house at a time to fiber, as demand for new services materializes.³⁵ Instead, infrastructure providers must either make a large up-front capital investment in advance of proven customer demand, or wait for demand to materialize in the absence of high-bandwidth connection capability. This chicken-and-egg situation is a product of the non-incremental nature of the required investment, and has characterized much of the thinking about telecommunications infrastructure upgrades in the past decade.

There is an alternative, namely, to upgrade existing local access networks in a more incremental fashion. At least in the short term, incremental upgrades have the potential to satisfy the demand for Internet connections more quickly and economically.

What existing networks could be incrementally upgraded to provide the desired service? The ubiquitous residential telephone network is an obvious candidate, if it can be upgraded to provide higher-speed transmission of data. Another candidate is the coaxial cable network that passes more than ninety percent of U.S. homes and currently delivers Community Antenna Television, or CATV, service. Wireless networks are also a possibility but are beyond the scope of this report.

³⁴As Reed writes on page 1 of (Reed, 1993): "The high costs of optical systems have so far ruled out the widespread deployment of fiber-to-the-home and fiber-to-the-curb networks." See also (Calhoun, 1992), pp. 547-8.

³⁵Professor David Tennenhouse suggested thinking about the problem in this way. Flexibility of the cost structure of the needed investment is also discussed in (Calhoun, 1992), pp. 549-50; (Lu, Eiger and Lemberg, 1990), p. 1063; and (Faulkner, Payne, Stern and Ballance, 1989).

Figures 2.5 and 2.6 illustrate the architectures of the enhanced telephone and cable networks considered by this report. Figure 2.5 shows Internet access over local digital telephone circuits, as provided by the *Integrated Services Digital Network*, or *ISDN*, suite of technologies. ISDN consists of an upgrade to the local telephone plant in which bi-directional digital transmission is enabled over the existing pairs of copper wires that currently supply bi-directional analog telephone service. It relies on the installation of an ISDN-capable digital switch at the telephone company Central Office—an upgrade that has already happened in most of the U.S.³⁶

³⁶Although more advanced technologies exist for providing digital transmission over existing copper wire pairs, this thesis focuses on ISDN because of its near-term availability, its appropriateness in residential areas and its symmetric transmission bandwidth. High Bit Rate Digital Subscriber Line (HDSL) technology supports 1.5 Mbps over two copper wire pairs, but is limited to distances from the telephone company Central Office that are more typical of business than residential lines; telephone companies provision HDSL as a business service. See (Waring, Lechleider and Hsing, 1991), (Walkoe and Starr, 1991), (Lechleider, 1991), (Lindstrom, 1992), and (Werner, 1991). More experimental Asymmetrical Digital Subscriber Line (ADSL) technology has the potential to deliver much higher bandwidths (1.5–6 Mbps) to the home at the expense of bandwidth from the home; it is envisioned for delivery of on-demand video applications. See (Waring, 1991) as well as (Sutherland and Litteral, 1992), (Chen and Waring, 1994), (Northern Telecom Inc. and Amati Communications Corp., 1993), and (Angelos, Bosco, Compton, Gu et al., 1993). The research for this thesis found real deployments of Internet access over ISDN, but no announcements or deployments of Internet access over either HDSL or ADSL. Thus ISDN was selected as more appropriate for a case study.

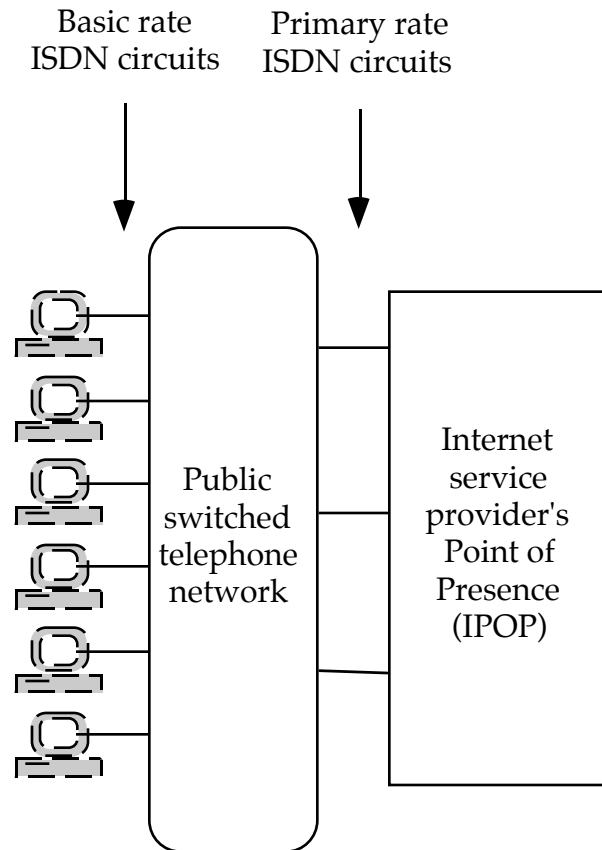


Figure 2.5: Architectural view of residential Internet access over ISDN

ISDN enables two types of digital circuits:

- Basic Rate Interface circuits provide one 9.6 Kbps signaling channel and two 64 Kbps channels that can be combined into a single 128 Kbps channel for user data. This is the type of ISDN circuit typically available to residential users.
- Primary Rate Interface circuits provide one 9.6 Kbps signaling channel and twenty-three 64 Kbps channels. These are typically used by businesses to combine traffic from multiple basic rate ISDN circuits.

Although ISDN provides faster data transmission than modems, it still follows the dialup access line model shown in Figure 2.2, enabling client access to the Internet but making full access difficult to achieve. Cable networks, however, follow a different model, as illustrated in Figure 2.6.

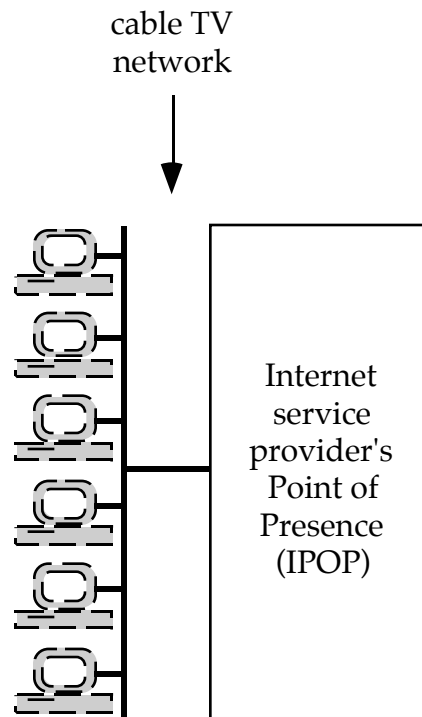


Figure 2.6: Architectural view of residential Internet access over cable

Instead of providing each individual computer with a point-to-point connection to the IPOP, cable's broadcast architecture provides a medium that can be shared by groups of individual computers. The star architecture of the access line model is replaced by a shared bus. The sharing is accomplished by protocols similar to those used in office computer networks (referred to as Local Area Networks, or LANs), but adapted to work over the different physical topology and longer distances that characterize the residential cable plant.³⁷

Just as the LAN-connected computers shown at the top of Figure 2.1 are always connected to their dedicated Internet access line, the cable LAN-connected computers of Figure 2.6 are always connected to their Internet service provider's Point of Presence. Thus, cable has the potential to provide full access to the Internet.

This architecture depends critically on the ability to transmit data in each direction—both from and to home computers—over the cable. This capability is often not provided in residential cable networks, since the typical application run over these networks is unidirectional broadcast of TV signals to the home. Providing the capability to send data in both directions over the cable is part of the infrastructure upgrade required to enable Internet access.

³⁷See (Karshmer and Thomas, 1992) for an overview of this technology and discussion of design issues. Chapter 3 describes the details of the cable LAN technology.

The speed of this access depends on the equipment used to modulate digital computer information onto cable's analog TV channels. At the time of this writing, such equipment provides bandwidths ranging from 500 Kbps to 10 Mbps.

The cable and ISDN approaches are described in greater detail in Chapters 3 and 4, respectively.

Chapter Three

Cable technology

This chapter reviews the present state and technical evolution of residential cable network infrastructure. It then discusses a topic not covered much in the literature, namely, how this infrastructure can be used to provide Internet access.³⁸ It concludes with a qualitative evaluation of the advantages and disadvantages of cable-based Internet access.

3.1. Cable networks

3.1.1. Topology of existing systems

Residential cable TV networks follow the *tree and branch* architecture shown in Figure 3.1. In each community, a *head end* is installed to receive satellite and traditional over-the-air broadcast television signals. These signals are then carried to subscriber's homes over *coaxial cable* that runs from the head end throughout the community.³⁹ Each 6 MHz TV channel is transmitted in analog form over (not necessarily the same) 6 MHz of enclosed spectrum on the cable. Multiple channels are sent over the same cable using *Frequency Division Multiplexing (FDM)*. Because different channels are sent at different frequency offsets (e.g. 54-60 MHz, 60-66 MHz, etc.), this form of transmission is referred to as *broadband*. (In contrast, *baseband* transmissions all take place in the "base" band starting at 0 Hz.)

³⁸(Berger, 1994), pp. 142-145, gives a brief overview of one form of cable-based Internet access.

³⁹The head end thus serves as an antenna for the community, giving rise to the name Community Antenna Television, or CATV. The tree and branch architecture is described in (Baldwin and McVoy, 1988), (Estrin, 1982), and (Reed, 1992).

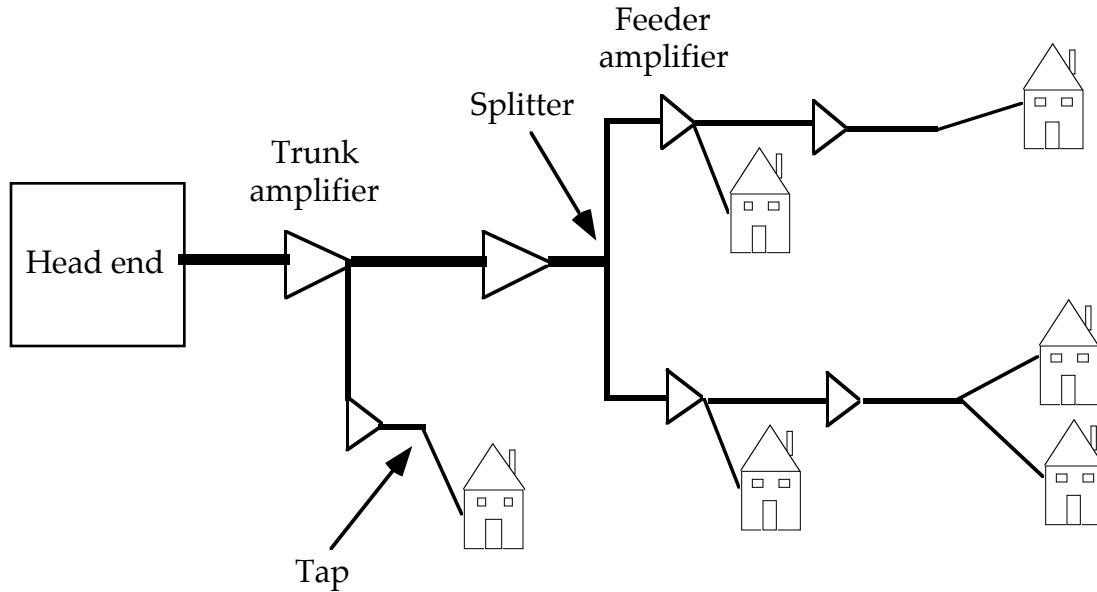


Figure 3.1: Coaxial cable tree-and-branch topology⁴⁰

To achieve geographical coverage of the community, the cables emanating from the head end are split (or “branched”) into multiple cables. When the cable is physically split, a portion of the signal power is split off to send down the branch. The signal content, however, is not split: the same set of TV channels reach every subscriber in the community. The network thus follows a logical *bus* architecture, as shown in Figure 3.2. With this architecture, all channels reach every subscriber all the time, whether or not the subscriber’s TV is on. Just as an ordinary television includes a tuner to select the over-the-air channel the viewer wishes to watch, the subscriber’s cable equipment includes a tuner to select among all the channels received over the cable.⁴¹




Figure 3.2: Logical bus architecture of cable TV network

⁴⁰Adapted from (Reed, 1992), Figure 2.4.

⁴¹This tuner is built into “cable-ready” television sets.

Because the signals attenuate as they travel several miles through the cable to subscribers' homes, *amplifiers* have to be deployed throughout the plant to restore the signal power. The more times the cable is split and the longer the cable, the more amplifiers are needed in the plant.

3.1.2. Technological evolution

The development of fiber-optic transmission technology has led cable network developers to shift from the purely coaxial tree-and-branch architecture to an approach referred to as *Hybrid Fiber and Coax*(HFC) networks. Transmission over fiber-optic cable has two main advantages over coaxial cable:

- A wider range of frequencies can be sent over the fiber, increasing the bandwidth available for transmission;
- Signals can be transmitted greater distances without amplification.

The main disadvantage of fiber is that the optical components required to send and receive data over it are expensive. Because lasers are still too expensive to deploy to each subscriber, network developers have adopted an intermediate *Fiber to the Neighborhood (FTTN)* approach.

Figure 3.3 shows the FTTN architecture. Various locations along the existing cable are selected as sites for *neighborhood nodes*. One or more fiber-optic cables are then run from the head end to each neighborhood node. At the head end, the signal is converted from electrical to optical form and transmitted via laser over the fiber. At the neighborhood node, the signal is received via laser, converted back from optical to electronic form, and transmitted to the subscriber over the neighborhood's coaxial tree and branch network.




Figure 3.3: Fiber to the Neighborhood (FTTN) architecture

In essence, FTTN replaces long coaxial cable runs with a long fiber and shorter cable runs. This replacement increases the bandwidth that the plant is capable of carrying.⁴² It also reduces both the total number of amplifiers needed and the number of amplifiers in *cascade* between the head end and each subscriber. The total number of amplifiers is an important economic component of the migration, because each amplifier must be upgraded or, more typically, replaced to pass the larger bandwidth that the fiber and shorter coaxial cable runs enable. The number of amplifiers in cascade is important to the quality of the service as seen by subscribers: since each amplifier is an active component that can fail, the fewer amplifiers in cascade, the lower the likelihood of failure. Fewer amplifiers and shorter trees also introduce less noise into the cable signal. These improvements translate into higher bandwidth, better quality service, and reduced maintenance and operating expense for the cable network provider.⁴³

The migration towards FTTN also creates a “cellular” architecture, in that a separate spectrum of cable channels may be delivered to each neighborhood. By effectively creating more channels, this architecture facilitates the addition of “personalcast”⁴⁴ services, in which different content may be delivered to (and/or possibly from) different subscribers simultaneously. Many of the envisioned services are based on digital content, such as compressed video to be viewed on demand, or catalog pages for home shopping. Some

⁴²A typical upgrade increases the upper spectrum limit from 550 to 750 MHz.

⁴³See (Stix, 1993), p. 102.

⁴⁴See (Gillett, Lampron and Tennenhouse, 1994), p.42, for Professor David Tennenhouse’s original use of this term.

personalcast services, such as telephony and Internet access, also require two-way transmission capability in the FTTN network.

FTTN has proved to be an appealing architecture for telephone companies as well as cable operators. Not only Continental Cablevision and Time Warner, but also Pacific Bell and Southern New England Telephone have announced plans to build FTTN networks.

3.1.3. Two-way transmission

Television, as its critics have long noted, is a one-way street: programs are broadcast to viewers, but no mechanism is built in for viewers to talk back to producers, or to send their own video signals into the network. This fact is reflected in typical residential cable TV network implementations, which transmit signals only in the direction from the head end to the subscriber (referred to as the *downstream* direction).

Many of the planned new services require transmission from the subscriber as well, with bandwidth requirements for such *upstream* transmission varying tremendously depending on the service. Pay-per-view movies could be ordered over the network with only a few low-bandwidth clicks of the remote control, while telephony would require a larger (but still relatively small) 4 KHz of bandwidth. Internet applications span a range of upstream requirements, from low-bandwidth email to high-bandwidth video sent from a home-based World Wide Web server.

Enabling upstream transmission requires three types of technical changes to the network:

- Spectrum must be allocated for the signals traveling in the upstream direction. Figure 3.4 shows a commonly-used spectrum map for the signals traveling over residential cable plants; the range from 5 to 42 MHz is typically dedicated to upstream transmission.⁴⁵ This range generally provides a maximum of 4 usable upstream channels.

⁴⁵As described in (Time Warner Cable Inc., 1993), Time Warner's "Full Service Network" trial in Orlando, Florida proposed to use the 850-1000 MHz frequency range for upstream transmission. The usual 5-42 MHz range was considered to be insufficiently wide for their purposes and, especially at the lower end of the range, known to be subject to interference from household appliances. The 5-42 MHz allocation is referred to as a "subsplit" channel assignment. Institutional cable networks ("I-nets") use mid-split or high-split (such as Orlando's) spectrum maps which allocate many more upstream channels, but these are not generally available in residential areas. See (Karshmer and Thomas, 1992), p. 36.




Figure 3.4: Sample cable spectrum map⁴⁶

- Amplifiers must include circuitry (diplex filters) to separate the upstream and downstream signals and amplify each direction separately, in the right frequency range. Rather than upgrading existing amplifiers in their tree and branch networks, most network providers have elected to incorporate upstream capability into the new amplifiers they need to purchase to pass the higher-bandwidth signals of hybrid fiber and coax networks.⁴⁷
- Downstream transmission from the head end is broadcast: the same signal is sent on all the wires. In contrast, upstream transmission is inherently personalcast: each subscriber is trying to place a different signal onto the network. As shown in Figure 3.5, when going up the tree these different signals must eventually share the same piece of transmission spectrum. Some form of *access method* is needed to arbitrate which signal is actually carried. Access methods come in different varieties, such as Time or Frequency Division Multiple Access (TDMA or FDMA), Carrier Sense Multiple Access (CSMA), etc. Which method is appropriate depends on the expected application.⁴⁸

⁴⁶After (Baldwin and McVoy, 1988), Figure 3-1. See also (Communications Futures Project, 1994), Figure 2.2.

⁴⁷Telephone conversation on March 11, 1994 with Mike Morris, sales representative for C-COR, a vendor of cable system products, including amplifiers.

⁴⁸See (Bertsekas and Gallager, 1992) for discussion and analysis of access methods.




Figure 3.5: Sharing of upstream bandwidth

Upstream transmission also presents a technical challenge because multiple noise sources come along with the multiple signal sources. Careful plant engineering is often required for the upstream signal to arrive in recognizable form.⁴⁹ The cellular aspect of the FTTN architecture helps it to address the noise ingress problem, by reducing the number of noise sources funneling into a single upstream channel.

3.1.4. Future evolution

Fiber to the neighborhood is one stage in a longer-range evolution of the cable plant. These longer-term changes are not necessary to provide Internet service today, but they might affect aspects of how Internet service is provided in the future.

Two trends are evident: connecting the head ends in a large metropolitan area together (typically via fiber), and adding digital switching to the system. Interconnected head ends enable remote operations; fewer resources need to be devoted to each head end, reducing the operating expense of a cable industry Multiple System Operator (MSO). They also create a regional market for advertising as well as enabling localized ad insertion. In addition, by placing expensive new equipment (such as a video server) at a regional hub, new services can be introduced to more communities at a lower capital cost.⁵⁰ Digital switches (typically using *Asynchronous Transfer Mode*, or ATM, technology) at each head end are envisioned to direct different traffic streams

⁴⁹See (Karshmer and Thomas, 1992) for further discussion of the noise ingress issues.

⁵⁰See (Dukes, 1992) for a discussion of the migration to regional interconnection and hubs.

to different neighborhood fibers, enabling Video on Demand (VOD) services.⁵¹ The implications of these trends for Internet service will be discussed below, after the mapping of Internet service to the cable architecture has been made clear.

3.2. Internet access over the cable plant

Recall that the key question in providing residential Internet access is what kind of network technology to use to connect the customer to the Internet Point of Presence. For residential Internet delivered over the cable plant, the answer is *broadband LAN* technology. This technology allows transmission of digital data over one or more of the 6 MHz channels of a CATV cable. Since video and audio signals can also be transmitted over other channels of the same cable, broadband LAN technology can co-exist with currently existing services. Broadband LANs were deployed in the 1980's by customers with large campuses, such as universities, manufacturing plants, and military bases, for whom integrated transmission made more economic sense than stringing multiple long wires.⁵²

The technology described here is the implementation used in the case that was studied for this report: Internet access offered over the Cambridge, Massachusetts cable plant by Continental Cablevision, a cable system operator, and Performance Systems International (PSI), an Internet service provider. The broadband LAN equipment used in this case is sold by Zenith Communication Products. This service was introduced in March, 1994.⁵³

The broadband LAN is built on top of two 6 MHz channels: one transmitting data from the head end to subscribers (downstream), and the other transmitting data from subscribers to the head end (upstream).⁵⁴ At the head end, a *frequency translator* connects the two channels into a *channel pair*, so that data can be sent from one customer to another.

The frequency translator at the head end can be either an analog or digital device. An analog device simply translates signals to a

⁵¹Cable network evolution is discussed in (Cook and Stern, 1994), including a diagram of how switching would be incorporated into the system.

⁵²See "Broadband LAN History" on p. 66 of the C-COR Commercial Systems Products Price List, August 1993; (Black, 1989); and (Stallings, 1987), pp. 80-95. Stallings describes a number of different broadband LAN implementations. In the 1990's, broadband LANs are gradually being supplanted by fiber-optic transmission.

⁵³I attended the press conference at the Museum of Science in Cambridge, Mass. on March 8, 1994. See also (Sandberg, 1994).

⁵⁴An alternative to this symmetric architecture has been proposed for use over cable plants that do not support upstream transmission: transmit data downstream over the cable plant, and upstream using standard modems over the telephone network. This "hybrid" architecture is supported by equipment available from Hybrid Networks Inc. Since the upstream bandwidth is significantly smaller than the downstream, this architecture is only suitable for asymmetric data flows. It is not considered further in this thesis.

new frequency and retransmits them. A digital device recovers the digital data from the analog signal and then retransmits the cleaned-up data on the new frequency.⁵⁵

Figure 3.6 illustrates the mapping of Internet service to the local cable plant that has been adopted by Continental and PSI. An IP router is installed at the cable head end and connected by a dedicated telecommunications line to the rest of the Internet. Since the cable transmission plant uses the Fiber-to-the-Neighborhood (FTTN) architecture, one or more channel pairs can be allocated to LAN service on each of the neighborhood transmission segments. (Figure 3.6 shows only one neighborhood segment.) The LANs on these transmission segments are connected to the IP router at the head end, thus connecting each subscriber to the Internet.

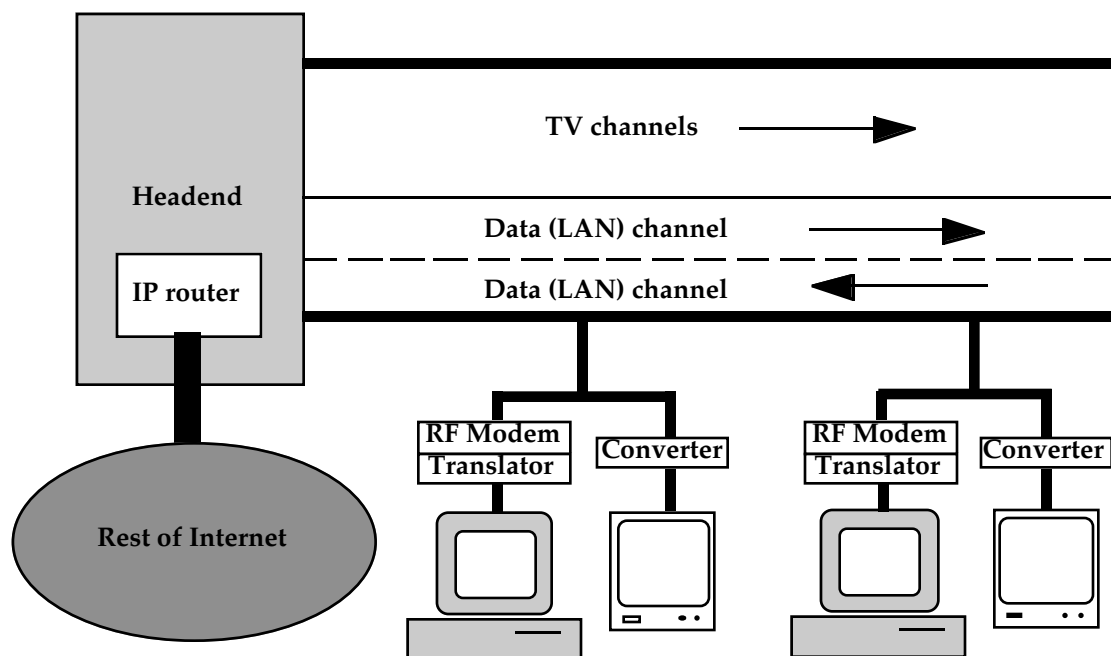


Figure 3.6: Internet access via cable

The remainder of this section explains in more detail how LAN service is provided over the channel pairs, covering the following issues:

- What protocol is used to arbitrate access to the upstream channel?
- How is digital data transmitted over the distances found in the local cable plant?
- What kind of equipment is needed to connect to the LAN?
- How is the IPOP router connected to the neighborhood LANs?
- What are the implications of cable plant evolution for residential Internet service?

⁵⁵(Stallings, 1987), p. 89. Analog frequency converters are used in the case studied.

3.2.1. Media access method

In LAN technology, the media access method refers to the protocol used to arbitrate access to the physical transmission channel among the different stations attached to it. The appropriateness of any given access method depends on the physical characteristics of the transmission media as well as the nature of the traffic that is to be sent.

The physical cable network is sufficiently different from the physical transmission media of ordinary LANs that standard media access protocols cannot be used as is. Ethernet, for example, uses *Carrier Sense Multiple Access with Collision Detection (CSMA/CD)*. With this method, each station “listens” to the transmission media. A station is able to transmit when it senses that the media is not in use (“carrier sense”). Because of the finite speed of light, signals do not propagate instantaneously and it is possible for two stations to begin transmitting simultaneously, in which case the transmissions will *collide* and be corrupted. Stations detect collisions by comparing the transmission they sent against what they receive over the broadcast bus. If a collision is detected, each sending station waits a random time interval and tries again.⁵⁶ Although the performance of this scheme is less predictable than a fixed scheme such as Time Division Multiple Access (TDMA), the dynamic bandwidth allocation works well for the unpredictable, bursty traffic patterns typically associated with data communication. As long as the media is not so heavily utilized that most transmissions result in collisions, CSMA/CD is quite efficient.⁵⁷

The effectiveness of CSMA/CD depends on its ability to avoid most collisions while detecting rapidly and efficiently those that do occur. The topology of the cable network, however, limits CSMA/CD’s effectiveness in both areas.

Collision avoidance

The window during which a collision can take place is the time it takes for one station on the bus to find out that another station has started transmitting, i.e. the time it takes for an electrical signal to travel from one station on the bus to another (called the *propagation delay*). In a short bus such as an Ethernet, this time window is very small, and most collisions are avoided. Since the cable network forms a much longer bus, this window is much wider.⁵⁸

The problem is further compounded by the “folded bus” topology of the cable plant, in which all transmissions are sent from the network stations on the upstream channel to the head end and back out the downstream channel. As

⁵⁶(Comer, 1991), p. 23.

⁵⁷See (Metcalfe and Boggs, 1976) and (Shoch and Hupp, 1980).

⁵⁸While an Ethernet’s length is limited to under 1 mile, Zenith’s product literature states that 5 to 30 miles is more typical of a cable network.

illustrated in Figure 3.7, the network stations send upstream and listen downstream; they do not listen on the upstream channel or send on the downstream channel.⁵⁹ Therefore, for example, the transmission from station 2 in Figure 3.7 must travel all the way to the head end and back before its neighboring station 1 notices that the bus is in use. Thus a physical separation of perhaps a few hundred yards translates into a signal propagation distance measured in miles.

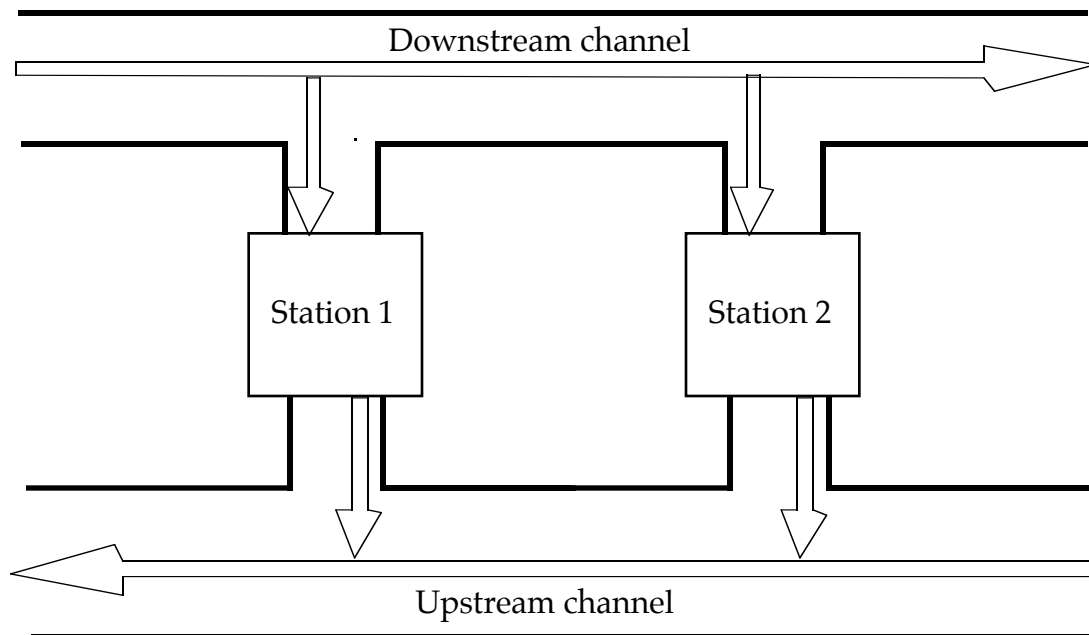


Figure 3.7: Network stations listen downstream and send upstream

The folded bus topology also creates unfairness in the CSMA/CD collision avoidance scheme. Stations closer to the head end experience the unfair advantage of a lower probability of collision, because their signals are physically on the upstream channel for less time, and therefore more quickly detected by the other stations' carrier sensing.

In summary, CSMA/CD-style collision avoidance is less effective on residential cable network topologies not only because the worst-case propagation delay is longer than with an ordinary Ethernet, but also because more stations appear "far away" from each other on the bus, increasing the likelihood of collisions.

⁵⁹Since the upstream and downstream channels operate in different frequency ranges, the station equipment would need twice as many tuners (and therefore be more expensive) if it had to listen upstream and send downstream as well as sending upstream and listening downstream.

Collision detection

The long signal propagation delays of the residential cable plant also make CSMA/CD collision detection less efficient. Ethernet collision detection relies on a station listening to its own transmission long enough to guarantee that it will not be corrupted by a collision. This time is twice the maximum propagation delay.⁶⁰ This guaranteed listening time is accomplished by setting a minimum packet length that takes at least twice the maximum propagation delay to clock out of the Ethernet interface. This limitation can be significant when propagation delays are long, as they are in the residential cable plant. If the minimum packet size is too long, many packets will have to be padded, wasting bandwidth.⁶¹ Alternatively, the interface clock rate may have to be reduced.

Because of these limitations, ordinary CSMA/CD is not used on residential cable LANs. Unfortunately, there is no standard replacement media access protocol; even the closest possibility—the modified CSMA/CD specified for broadband LANs (the IEEE 802.7 standard) is only designed for a maximum cable length of 3.8 Km (2.4 mi).⁶² Therefore, the choice of a media access method is up to the vendor. Zenith uses the 802.7 Ethernet for its LAN, but extends the length over which broadband CSMA/CD techniques can be used by having each station, when it is installed, measure its round trip propagation delay to the head end. The station then pads all packets as needed to guarantee a sufficiently long minimum packet size. In addition, interfaces are limited to a transmission clock rate of 4 Mbps, rather than the 10 Mbps of standard Ethernet.⁶³ Finally, Zenith's protocol forces stations closer to the head end to observe longer intervals of not sending after a transmission than stations further from the head end, to address the inherent unfairness introduced by the folded bus topology. The resulting protocol does

⁶⁰For the worst case, consider two stations, *A* and *B*, as far apart as possible on the bus. *A* starts to transmit, and just before its signal arrives at *B* (almost one propagation delay time later), *B* starts to transmit. The transmissions collide, and another propagation delay time is required for the corrupted transmission to travel back to *A*. Therefore, at most twice the maximum inter-station propagation delay time is required for a collision to be detected by the originator of the transmission.

⁶¹Analysis and simulation (beyond the scope of this thesis) are required to understand the real impact of this limitation on performance. If bandwidth efficiency only matters for high throughput applications such as file transfers in which most packets are large anyway, this limitation may not be very important.

⁶²These modified mechanisms are described in "8023HE Ethermodem Remodulator Theory of Operation," product literature from Chipcom for their broadband LAN equipment, July 1986. (The IEEE 802.3 standard specifies Ethernet CSMA/CD; see (Black, 1989), pp. 771-88.) This literature was provided by Ron Hoffman of MIT's Distributed Computing and Network Services.

⁶³Zenith also makes 500 Kbps LAN equipment. Each LAN runs in 1 MHz of a 6 MHz channel. With guard bands, 4 LANs can be frequency-division-multiplexed onto a single channel pair.

not appear to have been simulated or modeled theoretically; its efficiency and performance are not well understood.⁶⁴

Since very few vendors currently make station equipment for residential broadband LANs, the design and testing of a media access method that works well in this environment is an appropriate area for further research and standardization. Karshmer and Thomas provide an overview and brief bibliography of research in this area, suggesting methods that could be more efficient than Zenith's.⁶⁵ For example, the head end could do more than just translate frequencies; it could actually look at the packets and perform collision detection, thus avoiding part of the long signal propagation delay. Alternatively, given a Fiber to the Neighborhood architecture, LAN equipment could be installed at the neighborhood nodes, shortening the length of each cable LAN and allowing more standard CSMA/CD techniques to be used.

3.2.2. Transmission of digital data

To the user, a cable LAN looks like a digital channel that simply transmits the 1's and 0's of computer data. Because long wires attenuate, delay, and distort electrical signals with noise, the actual signal transmission is far more complex than it appears. First, sequences of bits are *coded*, or transformed into configurations of voltage levels that allow both efficient use of the channel and recovery of bit transitions even when sending and receiving clocks are not synchronized. Depending on the coding algorithm used, error detection and/or correction information may also be incorporated into the bit stream. Then the coded bits are *modulated* onto an analog waveform (a carrier signal), enabling frequency-shifted transmission.⁶⁶ When the signal reaches its destination, the receiver demodulates the analog waveform and decodes the bit stream, thus recovering the user's original sequence of 1's and 0's.

De/coding and de/modulation are the functions performed by high-speed modems such as the familiar 14.4 Kbps telephone modem; these topics have been extensively researched in the context of digital transmission over the local telephone plant.⁶⁷ The choice of coding and modulation schemes affects how successfully data can be transmitted over the cable plant in the presence of noise. Although many schemes are possible, a constraint is imposed by the requirement that senders and receivers must use matching schemes. In the

⁶⁴Zenith's access protocol is described in (Zenith Communication Products Inc., 1994). The brochure for Zenith's 500 Kbps "Homeworks" broadband cable modem products claims that the protocol can be used up to distances of 100 miles.

⁶⁵See (Karshmer and Thomas, 1992).

⁶⁶Coding and modulation are described in further detail in (Waring, Lechleider and Hsing, 1991) and (McNamara, 1982).

⁶⁷See (Waring, Lechleider and Hsing, 1991). The name "modem" derives from "modulator-demodulator".

cable LAN environment, this constraint translates into a requirement for the head end modem to share at least one coding and modulation scheme with each subscriber.

3.2.3. Subscriber equipment

Each customer's connection to the Internet access network is a station on the cable LAN, as is the head end's connection. Thus the total number of stations on each LAN is the number of subscribers plus one. Each station needs to perform the following functions:

- Implement the selected schemes for transmitting and receiving bits over the cable (coding and modulation);
- Implement the selected protocols for media access and address resolution;
- Support the upper layer protocols needed for Internet access: the TCP/IP suite. It is also helpful to the residential Internet service provider if the station is remotely manageable via the standard Simple Network Management Protocol (SNMP);⁶⁸
- If the station interfaces one or more customer LANs to the cable LAN, serve as a bridge or router.

There is a fundamental difference between the equipment needed to perform the first and last two of these functions. The first two functions are not standardized for cable LANs, but the equipment that provides them must be compatible among all stations using the channel. This combination can thus lock a provider into using a single vendor's equipment. In contrast, the last two functions are the same for cable LANs as for any other type of LAN, and can largely be implemented by off-the-shelf hardware or software available from a variety of vendors. The business implications of the non-standardized physical and media access layers will be discussed further in Chapter 8.

The first two functions are implemented by Zenith in two pieces of electronics: a broadband (or "RF", for radio frequency) modem and a network interface card. The modem determines the peak channel access rate; Zenith currently offers two models, capable of transmitting bits at 0.5 Mbps and 4 Mbps. A broadband modem is slightly more complex than an ordinary telephone baseband modem, because it has to shift the signal up to (in the case of a transmitting modem) or down from (in the case of a receiving modem) the transmission frequency band. Some broadband modems work only on particular channel pairs, but more flexible, "frequency-agile" models incorporate tuners that allow configuration of the modem for whatever

⁶⁸See (Comer, 1991), pp. 403-420 for a discussion and description of SNMP.

channel pair is chosen from the cable TV spectrum.⁶⁹ This flexibility allows the provider to move a subscriber over to a different channel pair via a simple configuration change, instead of requiring a site visit to exchange the customer's equipment. Such a move might be necessary to balance network loads, or to allow service to continue during network testing or maintenance. One could also imagine a more advanced modem that would configure itself, finding, for example, the least noisy or least busy channel to use for transmission.

The network interface card is quite similar to any standard Ethernet card except that it implements Zenith's variant of broadband CSMA/CD. Above the data link protocol layer, the same communications software that implements TCP/IP for a PC connected to a conventional LAN can be used by the residential Internet subscriber. This software needs to be configured with its own IP address (supplied by the residential Internet service provider) and the IP address of a DNS server, just as it must be for use over any LAN.

The residential Internet customer who has one or more LANs may need router functionality and may be able to fill this need with equipment available from a variety of vendors. Routers do not ordinarily come with cable LAN network interfaces; whether this is a problem for the customer depends on the customer's LAN configuration. This issue, and the question of what the network interface card plugs into, is discussed further in Chapter 5.

3.2.4. Connecting the IPOP router to the neighborhood LANs

The channel pairs carrying LAN service are frequency-division multiplexed together with TV channels over the same transmission media (fiber and coax). The IPOP router, however, only knows how to switch LAN data. The necessary filtering is achieved by placing the IPOP router at the head end and connecting it to each cable LAN through network stations located at the head end, as illustrated in Figure 3.8. Each station includes a broadband modem that serves the IPOP router in the same way it serves the customer. If multiple LANs are provided over a single neighborhood's cable plant (by devoting multiple channel pairs), the IPOP router needs a separate station for each LAN.

⁶⁹A number of vendors sell modems for broadband LANs, but very few of these are able to handle the distances found in the residential cable plant.

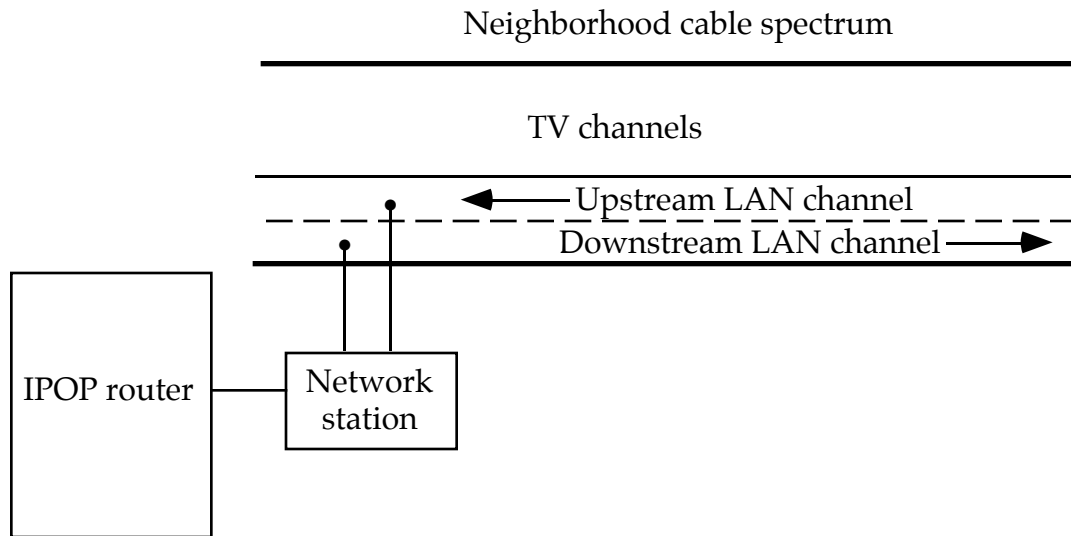


Figure 3.8: Head end router connection to cable LAN

3.2.5. Implications of cable plant evolution

In the longer term, cable operators are introducing Asynchronous Transfer Mode (ATM) switching into their head ends, and connecting multiple head ends together into fiber rings. ATM switches introduce an economy of scope. Instead of requiring a separate router for IP data, the small, fixed-length cells of ATM allow integrated switching of data, video, and voice. Internet service would be provided by running IP over ATM.

Fiber rings interconnecting head ends also introduce an economy of scope. Instead of requiring a top-level Internet router at every head end along with its own dedicated telecommunications line to the rest of the Internet, simpler routers at each head end are connected into a metropolitan area network, with its routing purely internal to the head end router operators. Only one of these routers needs to participate in full Internet routing and have a physical connection to the rest of the Internet, simplifying operations and reducing cost at the rest of the head ends.

3.3. Evaluation of cable LAN as Internet access network

In the cable Local Area Network (LAN) architecture, groups of subscribers connect to the Internet by sharing a single high bandwidth transmission medium. This architecture has a number of implications for the style of Internet access that can be provided over a cable LAN.

Bandwidth

The speed of a cable LAN is described by the bit rate of the modems used to send data over it. As this technology improves, cable LAN speeds may change, but at the time of this writing, cable modems range in speed from 500 Kbps to 10 Mbps, or roughly 17 to 340 times the bit rate of the familiar 28.8 Kbps telephone modem. This speed represents the **peak** rate at which a

subscriber can send and receive data, during the periods of time when the medium is allocated to that subscriber. It does **not** imply that every subscriber can transfer data at that rate simultaneously. The effective average bandwidth seen by each subscriber depends on how busy the LAN is. Therefore, a cable LAN will appear to provide a variable bandwidth connection to the Internet.

Still, the high peak data rates of cable LANs can make them a very effective Internet access network. For example, a 28.8 Kbps modem user who wants to view a simple graphical image (about 40 Kbytes) using the World Wide Web would have to wait more than 10 seconds for the image to download. In contrast, a cable LAN user could potentially view the image in under a second. People notice ten second delays; in contrast, a transfer taking less than a second may appear virtually instantaneous.

Full-time connections

Cable LAN bandwidth is allocated dynamically to a subscriber only when he has traffic to send. When he is not transferring traffic, he does not consume transmission resources. Consequently, he can always be connected to the Internet Point of Presence without requiring an expensive dedication of transmission resources.

At the Internet Point of Presence, the cable LAN is connected to an IP router. IP routing is *connectionless*: the router treats each IP packet as an independent entity. Therefore, no router resources need to be allocated to subscribers who are not transferring traffic.

In summary, cable LANs can provide full-time Internet connections because neither transmission nor switching resources have to be dedicated to a cable LAN subscriber who is not actively transferring traffic.

Security and network integrity

The shared medium of the cable LAN raises a security issue similar to one that arises in cellular telephony with its shared spectrum: eavesdropping on your neighbor becomes relatively easy. Encryption of data is required to solve this problem in both cases. Hence, support for encryption is likely to become an important requirement of cable modems.

Network integrity must also be managed carefully with a shared medium. For example, Ethernet broadcast packets may need to be disabled to ensure robust operations of residential cable LANs. If a subscriber's equipment were to break (or be maliciously altered) such that it was constantly sending Ethernet broadcast packets to all other network stations, it could effectively disable a cable LAN. While this same scenario can arise in a workplace LAN, finding and fixing the problem is much more difficult in the residential environment, in which the distance between network stations can be

measured in miles and the network manager may not be able to gain physical access to the offending equipment.

Multiple simultaneous services

Cable's broadband approach enables multiple network devices to use the network simultaneously: you can read your email while another member of your household watches the NBA Finals on TV. The ability to maintain a cable-based Internet connection while talking on the phone is perhaps even more useful, for example to discuss an email message with a colleague or customer, or look up information needed during a conversation. Although this capability comes "for free" now because telephone service is provided over a separate network, it would still be there if telephony were instead provided over broadband cable channels.

Scope of Internet Point of Presence

An Internet Point of Presence (IPOP) serves as a gateway to the rest of the Internet for any subscriber who can connect to it. If access to an IPOP is provided only through a local cable network, then the potential universe of subscribers who can connect to the Internet through that IPOP is limited to the homes passed by that cable system.

This limitation can pose a barrier to introduction of cable-based Internet access. For example, suppose a certain minimum number of subscribers is needed to recover the cost of the IPOP. The cable system may achieve this number by either passing a large number of homes, or by achieving high penetration levels for Internet service. The number of homes passed is a fact of the cable system's location, not a parameter that the cable operator can easily change. If it is too small, then the operator must rely on high penetration levels. These, however, are likely not to be achieved when the service is first introduced. If cable-based Internet access is not profitable at initial penetration levels, it may not be deployed even if it would be profitable at higher penetration levels that could be achieved later.




Figure 3.9: Greater subscriber population with interconnected head ends

This situation would change if cable networks ceased to have purely local scope. As Figure 3.9 shows, if multiple head ends were interconnected, an IPOP at one head end could serve a larger universe of cable subscribers. The same penetration percentage would represent a larger number of subscribers, making it more feasible to recover the investment needed to create the Internet Point of Presence.

Even regionally-interconnected head ends do not solve a different scope problem: how does a subscriber access the Internet while far from home? An out-of-region traveler has no way to connect to the IPOP that serves his home community. As the next chapter notes, the worldwide scope of the telephone network gives it an advantage in this regard.

Chapter Four

ISDN technology

While ISDN is extensively described in the literature, its use as an Internet access medium is less well-documented.⁷⁰ This chapter briefly reviews local telephone network technology, including ISDN and future evolutionary technologies. It then describes the technologies needed to provide Internet access over ISDN. It concludes with a qualitative evaluation of the advantages and disadvantages of ISDN-based Internet access.

4.1. Local telephone networks

Unlike cable TV networks, which were built to provide only local redistribution of television programming, telephone networks provide switched, global connectivity: any telephone subscriber can call any other telephone subscriber anywhere else in the world. A call placed from a home travels first to the closest telephone company *Central Office (CO) switch*. As shown in Figure 4.1, the CO switch routes the call to the destination subscriber, who may be served by the same CO switch, another CO switch in the same local area, or a CO switch reached through a long-distance network.

⁷⁰(Berger, 1994), pp. 137-140, gives a user-level overview of ISDN-based Internet access.

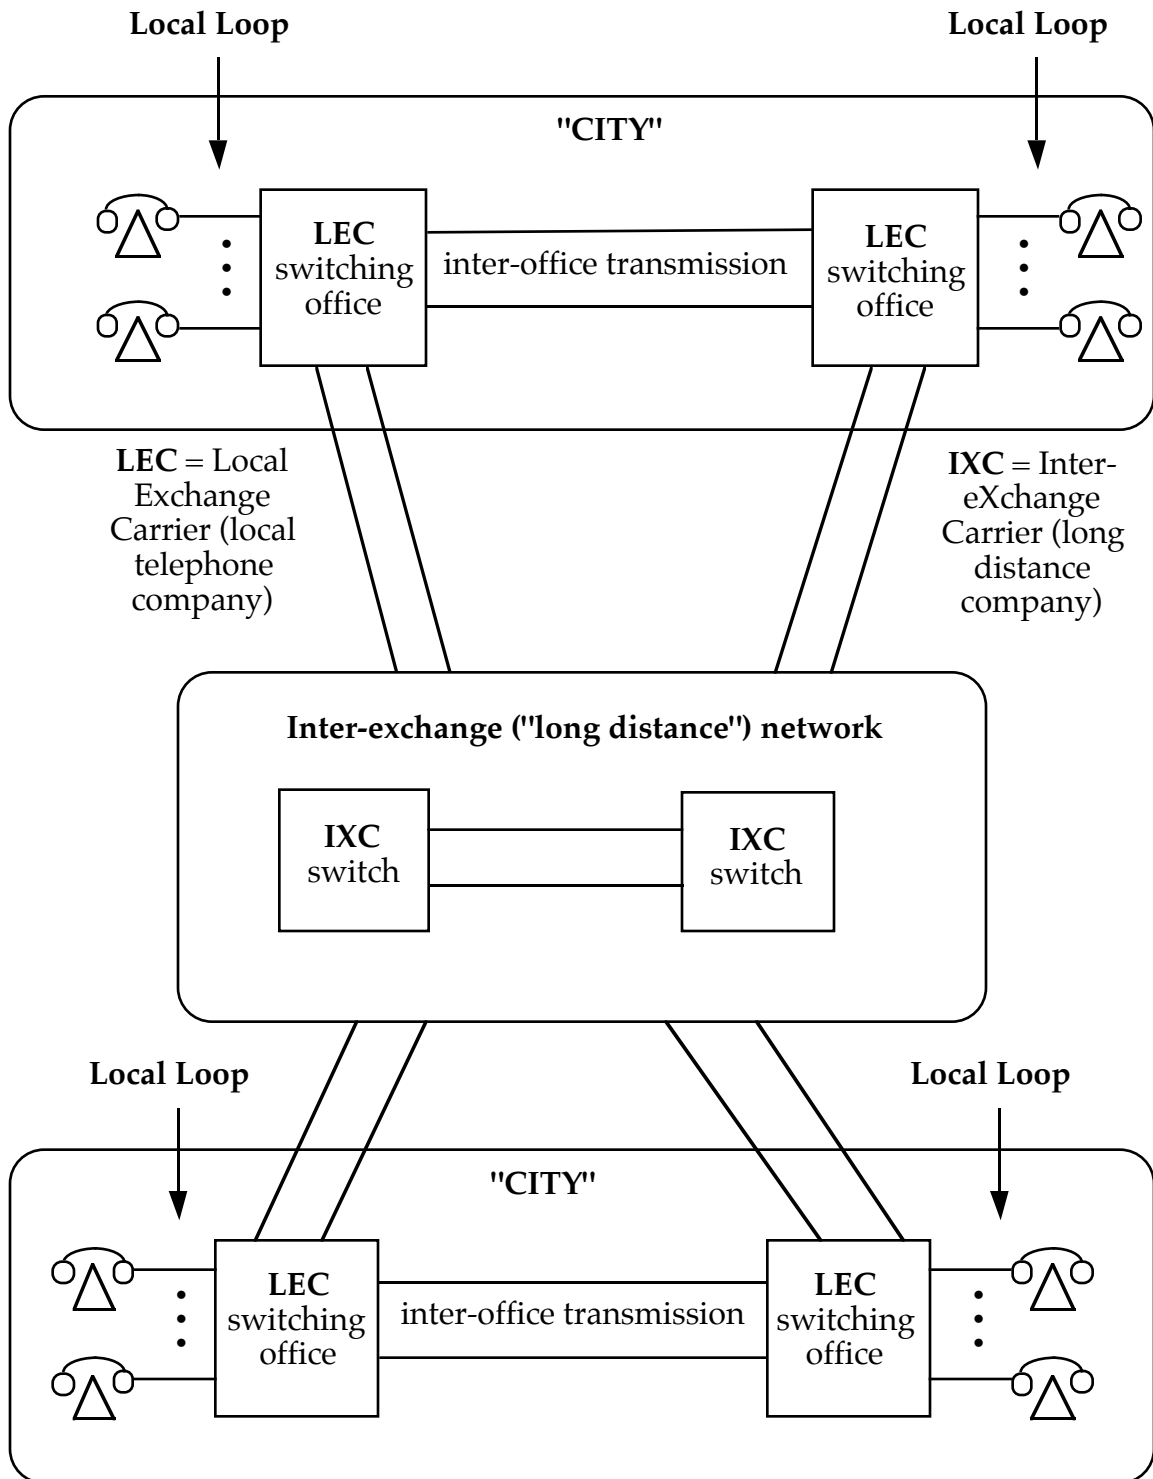


Figure 4.1: The telephone network

The portion of the telephone network that connects the subscriber to the closest CO switch is referred to as the *local loop*. The name derives from the pairs of copper wires, which form a "loop" between the CO and the customer

site, that have traditionally been used to connect subscribers to the CO.⁷¹ Since all calls enter and exit the network via the local loop, the nature of the local connection directly affects the type of service a user gets from the global telephone network.

4.1.1. Topology of existing networks

Most residential telephone service in the U.S. is provided over pairs of copper wire that run directly from a CO switch to each individual subscriber's home. Over each wire pair, referred to as a *subscriber line*, the baseband 4 KHz channel (i.e. 0-4 KHz) is used to carry the bi-directional, analog telephone signal. This traditional analog service is referred to as Plain Old Telephone Service, or POTS.

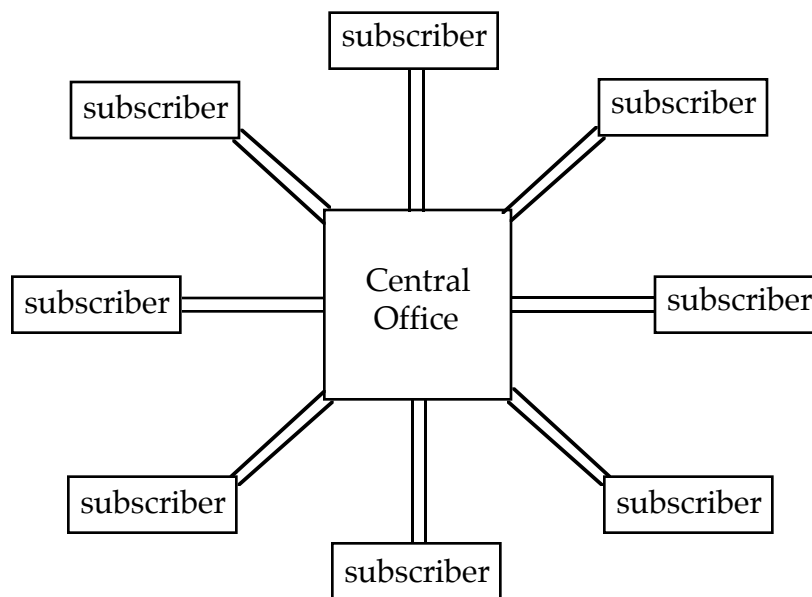


Figure 4.2: Logical star architecture of local telephone network

With a separate pair of wires to serve each subscriber, the local telephone network follows a logical *star* architecture, as shown in Figure 4.2. Since a Central Office typically serves thousands of subscribers,⁷² it would be unwieldy to string wires individually to each home. Instead, the wire pairs are aggregated into groups as shown in Figure 4.3. The largest groups of wires are *feeder cables*. At intervals along the feeder portion of the loop, *junction boxes* are placed. In a junction box, wire pairs from feeder cables are spliced to wire pairs in *distribution cables* that run into neighborhoods. At each

⁷¹It is also sometimes referred to as the “last mile” of the telephone network, although typically it covers several miles.

⁷²For example, Cambridge, MA has 40,000 households and 2 Central Offices.

subscriber location, a *drop* wire pair (or pairs, if the subscriber has more than one line) is spliced into the distribution cable.⁷³

⁷³The local transmission plant is described in (Gillett, Lampson and Tennenhouse, 1994), (Calhoun, 1992), and Chapter 4 of (Noll, 1991). The latter also contains photographs of a number of the items described in this section, including distribution cables, underground conduit, and a Main Distribution Frame.

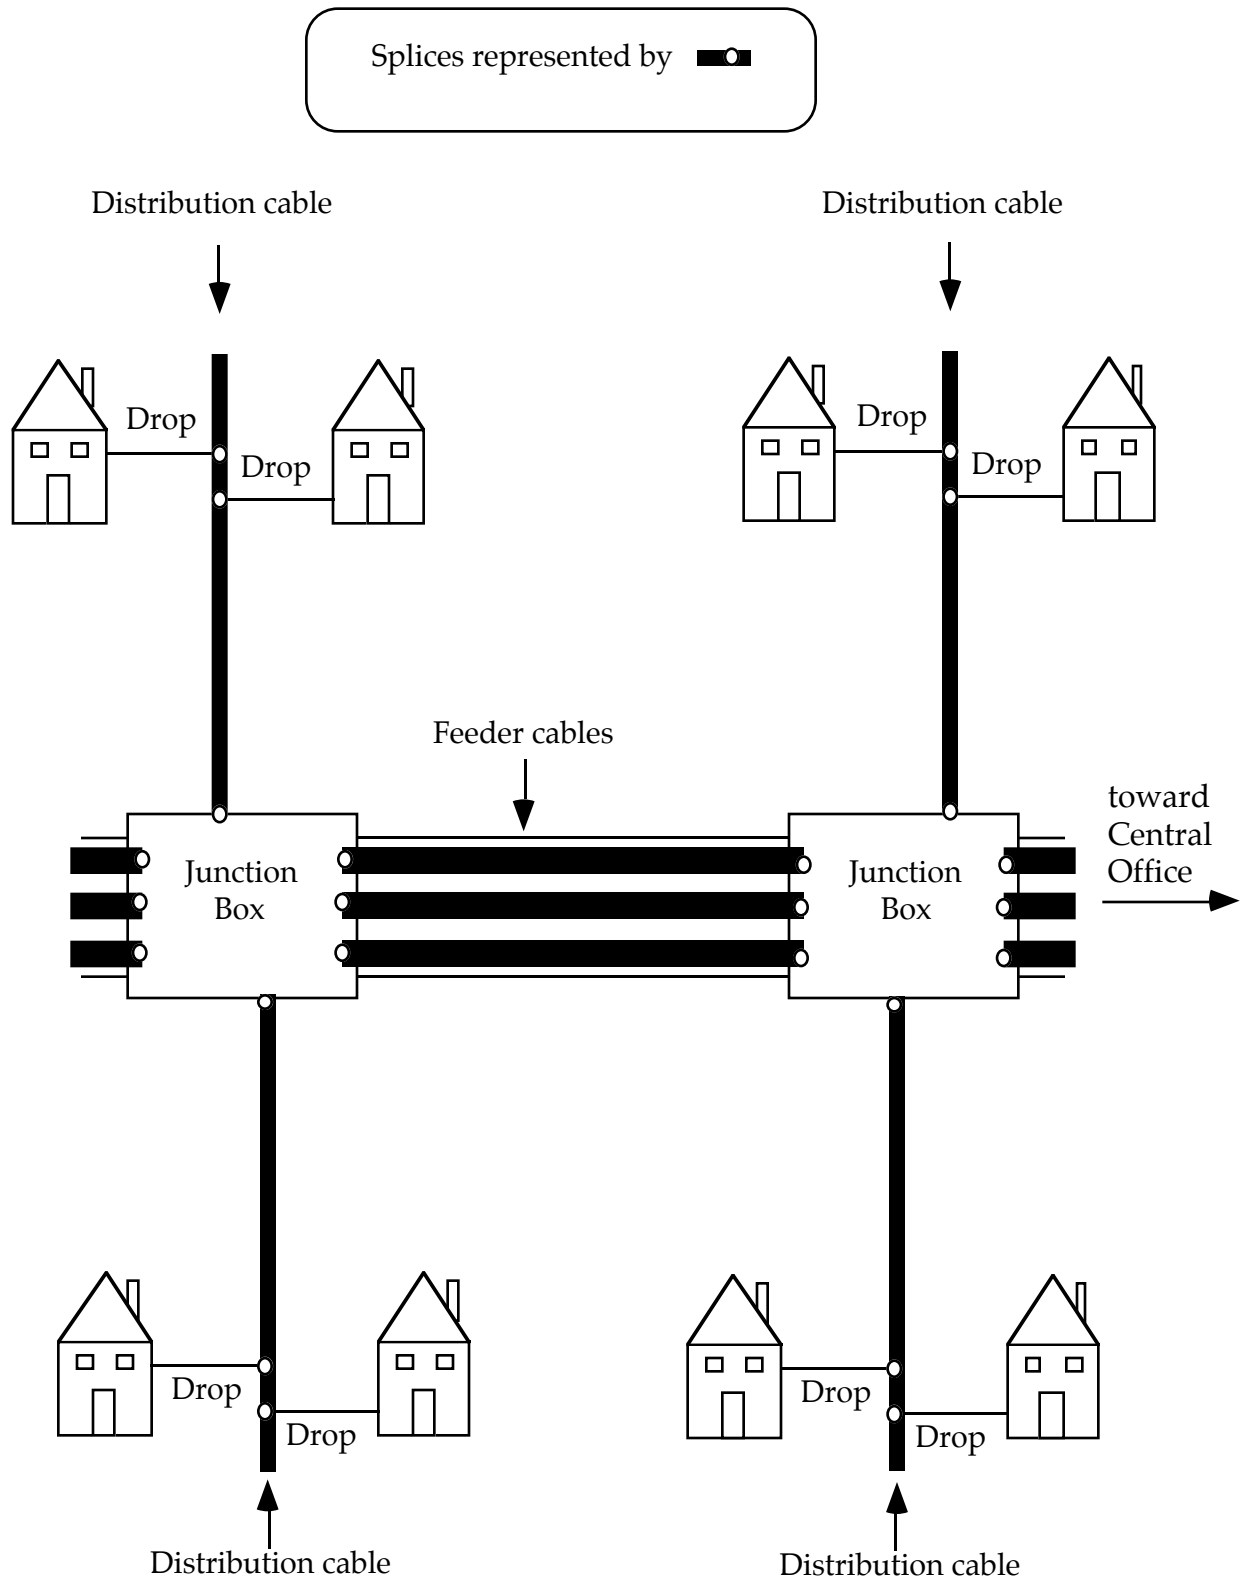


Figure 4.3: The local loop

The layout shown in Figure 4.3 is analogous to the tree-and-branch splits of the CATV network, except that instead of splitting off signal power at each junction, the telephone network splits off physical wire pairs. The aggregation of wire pairs in the local transmission plant constitutes space-division multiplexing, which creates a physical bus layout implementing a logical star architecture.

Since distribution cables are either buried or aerial (i.e. hanging from telephone poles), they are disruptive and expensive to change. Consequently, a distribution cable usually contains as many wire pairs as a neighborhood might ever need, in advance of actual demand. In contrast, feeder cables (each containing a large number of wire pairs) are generally run through underground conduits. Because it is relatively easy to string cable through conduit, new feeder cables are typically added as demand for new wire pairs warrants them.

For a new subscriber line to be activated, the subscriber's drop wire pair must be spliced to a free pair in the distribution cable, which must in turn be connected to a free pair in a feeder cable. Managing the configuration of all these connections thus becomes a major operational activity for local loop service providers.

The challenges inherent in managing thousands of individual pairs of wire are not limited to the local transmission plant. At the CO, each subscriber's line must somehow connect to the CO switch. Considerable space in the CO is taken up by the *Main Distribution Frame*, an enormous wiring rack that connects wire pairs to *line cards* that interface subscriber lines to the periphery of the CO switch. To understand the function of a line card, consider that the wire pair dedicated to each subscriber can be in one of three states:

- idle (when the subscriber is not using the phone, which is actually most of the time);
- establishing a connection (from the subscriber lifting the handset "off hook," through dialing and ringing);
- connected (a call has been successfully placed and is in progress).

With POTS, the transitions among these states are managed by electrical signals sent over the wire pair between the CO switch and the subscriber's equipment (which is usually a telephone, but could also be a fax machine or modem). One line card manages these signals for one subscriber line, supplying power to the subscriber equipment, watching for on/off hook

changes, etc.⁷⁴ Since most CO switches now use digital switching, the line card must also convert the analog signals received over the subscriber line to the digital format expected by the rest of the CO switch.

At the heart of the CO switch is a *switch fabric* that dynamically connects input ports to output ports leading to the appropriate destination for each call. Line cards are not directly connected to switch fabric input ports for two reasons:

- Building a switch fabric to physically accommodate tens of thousands of input ports would be quite a feat of engineering;
- Since not all subscribers place calls at the same time, not all line cards need to connect to the switch fabric at the same time.

For the first reason, the signals from subscriber lines are *multiplexed* together onto single, higher-capacity cables before being fed to the switch fabric (typically using Time Division Multiplexing, or TDM). For the second reason, subscriber lines are *concentrated*, that is, only a few connections to the switch fabric are provided for a group of subscriber lines. A logical diagram of the resulting switch architecture is shown in Figure 4.4.

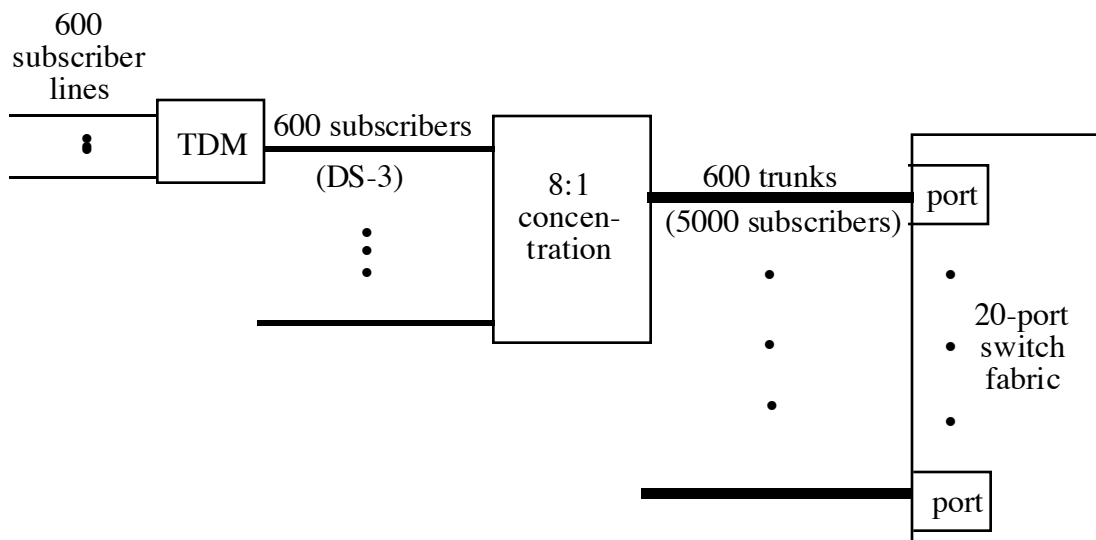


Figure 4.4: Central Office switch architecture⁷⁵

4.1.2. Technological evolution

Two overlapping trends are evident in local loop technology:

⁷⁴The functions of a line card are described in more detail on pp. 132-133 of (Noll, 1991) and in (Gillett, Lampson and Tennenhouse, 1994), p. 15.

⁷⁵From (Gillett, Lampson and Tennenhouse, 1994), p. 94.

- Deployment of fiber-optic cable in part or all of the local plant;
- Migration toward digital transmission in part or all of the local plant.

Fiber in the loop

Fiber is an attractive transmission medium because it can provide high-bandwidth, digital transmission more economically than large collections of copper pairs. Deployment of fiber throughout the loop has occurred mainly in business areas, both because these areas have contributed the most to the growth in demand for new telephone lines, and because they tend to be dense and use more advanced communication services, resulting in higher aggregate bandwidth needs.⁷⁶

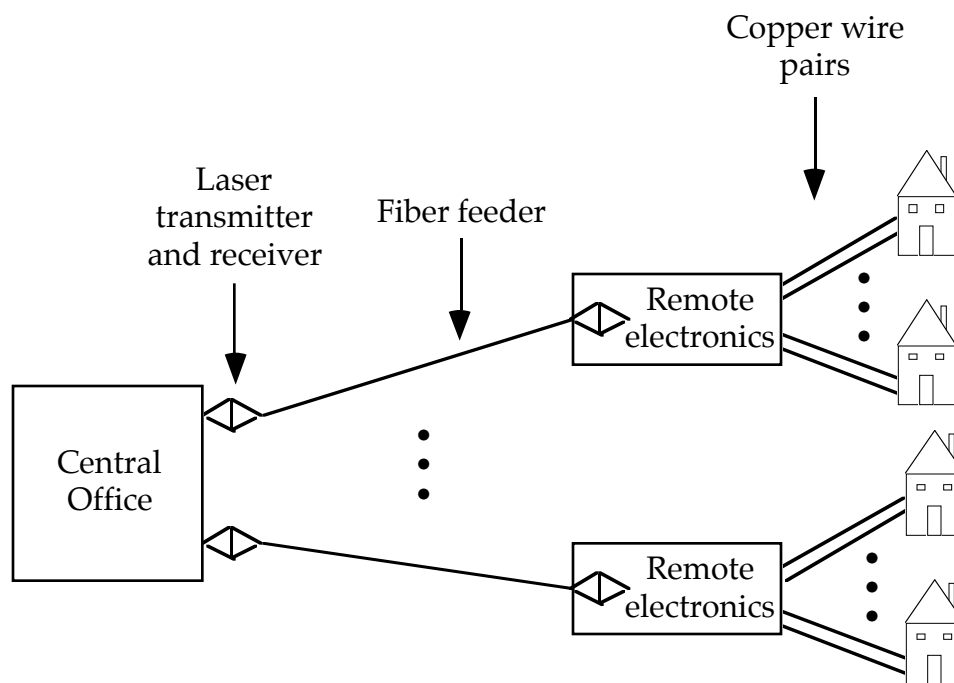


Figure 4.5: Fiber in the feeder portion of the loop

Given the lack of equivalent demand to justify the up-front costs of laying fiber in residential areas, deployment in these areas has been limited to date. Most residential deployments use fiber as a substitute for the feeder portion of the loop, as shown in Figure 4.5. The result is similar to the Fiber to the Neighborhood networks of cable TV, with two important differences:

- The “last mile” of plant is one copper wire pair for each subscriber, not a shared coaxial cable;

⁷⁶(Cook and Stern, 1994), pp. 78-79.

- The logical architecture is still a star: although the fiber carries all the traffic for a neighborhood, it is channelized into telephone call-sized chunks, each of which is dedicated to an individual subscriber line. A line card is still required for each subscriber line, and an idle line still consumes bandwidth on the fiber.⁷⁷

The net result is that while such a deployment saves operating expense for the telephone company,⁷⁸ the presence of fiber in the feeder is essentially invisible to subscribers, whose service options are largely unchanged as a result of the deployment.

The migration toward digital transmission over the copper wire plant is more significant for residential subscribers.⁷⁹ In the near term, *Integrated Services Digital Network (ISDN)* subscriber lines are being deployed. Further out, *Asymmetrical Digital Subscriber Line (ADSL)* technology has been proposed.

ISDN

As digital technology began to penetrate into long-distance transmission and switching systems in the 1960's and 70's, telephone companies conceived of a network that would offer end-to-end digital transmission. Digital bits would enter, traverse, and exit the network; in contrast to POTS service, the network would not have to convert calls from analog to digital internally. Such a service offering is "integrated" because, by handling signals in digital form, it could transparently handle different applications: for example, data bits, generated by computers, fax machines, etc. could be sent directly; voice bits would be generated by special telephones able to digitize voice signals at the boundary of the network.

Implementation of ISDN is hampered by the irregularity of the local loop plant. Referring back to Figure 4.3, it is apparent that loops are of different lengths, depending on the subscriber's distance from the Central Office. Different gauges of wire may be used within a single loop, and the number of splices may also vary. Chen and Waring describe one reason why the local loop plant ends up with unterminated wire pairs, referred to as *bridged taps*:

⁷⁷In contrast, Passive Optical Networks (PONs), which have been proposed for future deployments of fiber throughout the entire residential loop, use a shared bus architecture completely analogous to that of cable TV networks. See (Cook and Stern, 1994) or (Lu, Eiger and Lemberg, 1990) for a description of PONs.

⁷⁸For example, to add a new subscriber it is much easier to activate an additional channel on a fiber feeder than it is to run a new wire pair through conduit.

⁷⁹In theory, this migration is not restricted to all-copper loops; it can also take place over the copper portion of a loop that has been migrated to a fiber feeder. The practical limitation is only that the remote electronics (called a *Digital Loop Carrier*, or DLC) at the end of the fiber feeder must be able to support digital transmission over the copper wires attached to it; many existing DLCs do not.

Since loop plant construction usually occurs ahead of customer service requests, distribution cables are usually made available to all potential customer sites. Hence, it is a common practice to connect a twisted pair from a feeder cable with more than one distribution cable to maximize the probability of reaching a potential customer. The unused distribution cables result in bridged taps.⁸⁰

Within such an environment, achieving reliable transmission of bits is a challenge. For the same digital transmission equipment to work across the variety of loops found in the plant, either a plethora of configuration settings must be provided and manually adjusted, or the equipment must be able to adapt dynamically to the conditions found on the particular line. Since the former approach is economically infeasible, the advent of ISDN had to wait for the development of self-adaptive filtering techniques implementable with VLSI technology in the late 1980's.⁸¹ The result, standardized by the American National Standards Institute (ANSI) in 1988, is the *Basic Rate Access Digital Subscriber Line*.

Basic Rate Access provides a total of 144 Kbps of bidirectional digital transmission over the baseband channel from 0-50 KHz.⁸² Because it overlaps with POTS' baseband spectrum (0-4 KHz), ISDN cannot co-exist with POTS on the same wire pair. Once a subscriber loop has been upgraded to ISDN service, existing POTS equipment is no longer usable as is.

ISDN cannot be provided over loops with loading coils or loops longer than 18,000 feet (5.5 km), with bridged taps included in the loop length. Based on loop surveys conducted by Bellcore, Waring estimates that loading coils still exist on 25% of subscriber loops, and that the loops with loading coils are mainly the same as the loops longer than 18,000 feet.⁸³ By this estimate, Basic Rate Access can be provided over no more than 75% of North American subscriber loops.

The 144 Kbps provided by Basic Rate Access is subdivided into two *bearer*, or B channels, and one *delta*, or D channel. Each B channel transmits bits at 64 Kbps and is expected to carry voice or data. The D channel, which transmits at

⁸⁰(Chen and Waring, 1994; Chen and Waring, 1994), p. 104. Service disconnections and rearrangements also result in bridged taps; see (Waring, Lechleider and Hsing, 1991), p. 97.

⁸¹Consider that the adjustments must be made not only at installation time but also whenever moisture gets into the cables, the temperature changes significantly, etc. These effects, and the development of Digital Subscriber Line (DSL) technology, are described in (Waring, Lechleider and Hsing, 1991), pp. 98-100.

⁸²(Waring, Lechleider and Hsing, 1991), p. 103.

⁸³See (Waring, Lechleider and Hsing, 1991), p. 97. He states that loading coils were "used in the early days of telephony to boost and flatten the frequency response at the upper edge of the voice band." No digital transmission techniques can be used in the presence of loading coils.

16 Kbps, was envisaged for out-of-band signaling—for example, setting up a call by exchanging digital “packets” of instructions, instead of detecting on- and off-hook signal transitions. The result is referred to as $2B+D$, or an ISDN “U” interface.⁸⁴

Switching is also different in ISDN. With POTS, setting up a call consists of allocating resources to dedicate to that call for as long as it remains in progress. These resources include a connection through the CO switch fabric, bandwidth on inter-office trunks, etc. which together form a guaranteed-delay *circuit* through the telephone network. This technique of dedicating resources to connections is referred to as *circuit switching*.

Since data traffic is typically bursty (as discussed in Chapter 1), dedicating network resources to a single data source can be quite wasteful. Instead of circuit switching, computer networks use a technique called *packet switching* in which computer data are divided into small chunks, called packets. Each packet is given a *header* containing control information such as the packet’s destination, and is sent individually through the network when resources are available.⁸⁵

Since ISDN is designed for both voice and data, it incorporates both circuit and packet switching. A Central Office switch that offers ISDN service can be configured to switch either circuits or packets on either B channel independently.⁸⁶ The D channel can only be packet-switched. To provide this functionality as well as digital transmission to the subscriber, an ISDN line card must be used. A matching card, called an NT1, is also needed at the subscriber end of the line.⁸⁷ This card not only handles digital transmission but also formats and interprets the signaling packets sent over the “D” channel.

Most Central Office switches in North America are now digital and many of those are capable of providing ISDN service. However, not all of these switches follow the recent National ISDN standards,⁸⁸ hampering compatibility between ISDN equipment available from vendors and ISDN service available from the local telephone company.

⁸⁴(Waring, Lechleider and Hsing, 1991), p. 99.

⁸⁵Packet switching is described in (Comer, 1991), p. 18.

⁸⁶The two B channels are not switched as a unit. *Inverse multiplexing* subscriber equipment is available to aggregate the two channels, emulating a single 128 Kbps channel (Duncanson, 1994). For reasons that will be discussed later in this chapter, packet switching is rarely used for the B channels.

⁸⁷The acronym NT1 derives from the positioning of the card as the first point of ISDN Network Termination on the subscriber’s premises.

⁸⁸See (Bellcore, 1991) and (Bellcore, 1993).

Subsequent to the definition of Basic Rate Access (2B+D), Primary Rate ISDN was standardized. A *Primary Rate Interface (PRI)* provides 23 B channels and one D channel to the customer, typically a business. It is provided over the same types of loops suitable for digital T1 transmission: either fiber, or 4-wire copper lines. Since neither of these are generally available in residential areas, PRI is not considered further in this report.

4.1.3. Future evolution

In 1991, the FCC issued its "Video Dialtone" ruling allowing the local telephone companies to distribute video signals over their networks.⁸⁹ This opportunity has motivated plans for different forms of local network upgrades. Several local telephone companies have announced plans to completely rebuild their networks according to a Hybrid Fiber and Coax architecture.⁹⁰

For a more evolutionary approach, other telephone companies have turned to the Digital Subscriber Line technology that was originally developed for ISDN, allowing more re-use of the existing copper loop plant.⁹¹ This technology was developed into the High Bit Rate Digital Subscriber Line (HDSL), used to provide "repeaterless T1" lines to businesses as described in the previous chapter. For residential areas, this technology continues to evolve in the form of Asymmetric Digital Subscriber Line (ADSL).

Asymmetric Digital Subscriber Line

Asymmetric Digital Subscriber Line (ADSL) was originally conceived as a means to distribute video from the CO to the customer over a twisted pair of copper wires. In contrast to ISDN, ADSL adopts a *passband* approach. High and low pass filters divide the line into three frequency channels: bidirectional analog POTS, narrow digital upstream, and wide digital downstream, as shown in Figure 4.6. This technique, which is referred to as *Data Over Voice*, enables compatibility with existing POTS equipment.

⁸⁹See (Federal Communications Commission, 1991), (Federal Communications Commission, 1992), and (Chen & Waring, 1994b), p. 102.

⁹⁰See, for example, "Southern New England Telephone Buys AT&T Fiber/Coax Transport" in Cable Regulation Digest, April 25, 1994, Vol. 1 No. 17. (Cable Regulation Digest is a summary of regulatory news from Multichannel News.)

⁹¹Bell Atlantic and Ameritech have both announced trials of Video on Demand over existing copper loops. See (Markoff, 1994).

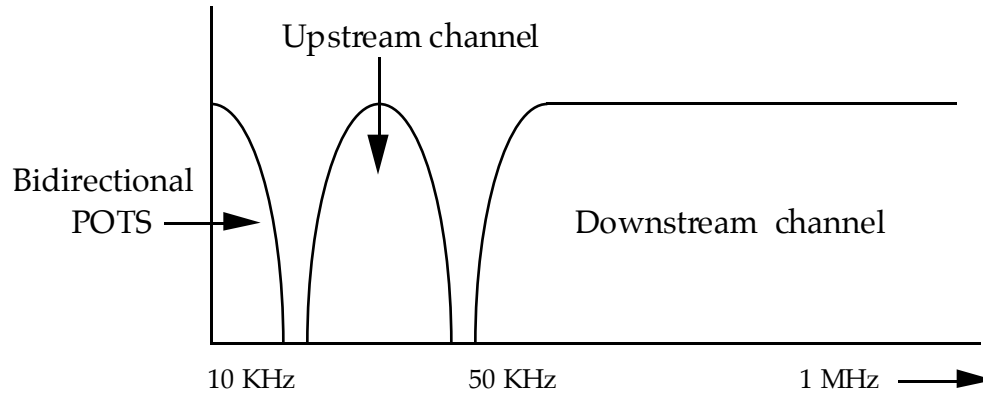


Figure 4.6: Spectrum allocation for ADSL transmission⁹²

The name ADSL arises from the asymmetry between the upstream and downstream channel sizes. Bidirectional low-bandwidth control, using the same format as the 16 Kbps ISDN “D” channel, is provided over the upstream channel and some portion of the downstream channel. The remainder of the downstream channel is expected to carry “video on demand” to the subscriber at 1.5 Mbps. The bidirectional POTS channel continues to provide analog telephone service to the subscriber in a transparent fashion.⁹³

Although ADSL transmission technology is fairly well established, other aspects of ADSL usage are much less well defined. For example, the interface to the subscriber was an open question before the ANSI standards committee responsible for ADSL (Committee T1E1.4) in the spring of 1994.⁹⁴ Some form of in-home demultiplexing will be necessary to split the three ADSL channels coming into the home, with corresponding multiplexing needed in the opposite direction. Similar functions will need to be performed by the ADSL line card that interfaces each subscriber to the CO switch. Finally, whether the digital channels will be geared toward television (“set-top box”) or computer-style connections is an open question.

Switching of ADSL is also unspecified. Just as cable companies would deploy Asynchronous Transfer Mode (ATM) switches to provide pay-per-view service, telephone companies are likely to deploy ATM switches to direct the “demanded” video to an ADSL video-on-demand subscriber. Presumably, the POTS and control channels will be concentrated—perhaps at first by the existing ISDN Central Office switch—and handled by the ATM switch as well.⁹⁵

⁹²From (Northern Telecom Inc. and Amati Communications Corp., 1993), p. 2.

⁹³(Waring, 1991), p. 1980.

⁹⁴Conversation with Ron Hunt of Amati Communications Corp., Palo Alto, CA, on February 2, 1994.

⁹⁵Conversation with Professor David L. Tennenhouse, March 16, 1994.

Despite several years of advances in VLSI technology and modulation, coding, and error correction schemes,⁹⁶ the heterogeneous, noisy loop environment still poses problems for high-bandwidth digital transmission over long copper pairs. For example, an adjacent pair carrying a digital T1 can make ADSL unusable over a long loop. Although not many T1's exist in residential areas, those that do further limit ADSL's applicability, because they mean that ADSL technology cannot be guaranteed to succeed in all installations. The suggested solutions to this problem are "ADSL-2" and "ADSL-3," which run ADSL over progressively shorter distances by employing a fiber to the neighborhood architecture. ADSL line cards must be included in the neighborhood node's electronics. With this architecture, higher digital transmission speeds (up to 6 Mbps downstream and 384 Kbps, or "H0", upstream) can be supported over ADSL.⁹⁷

4.2. Internet access over the local telephone network

On the cable network, one wide channel is shared among a number of customers to connect each of them to the IPOP router. The local telephone network gives the residential Internet provider a different topology to work with: narrow channels dedicated to each customer individually. With this topology, it does not make sense to use a shared-bus LAN approach to connect the IPOP router to customers. Instead, point-to-point connections are used. An architectural view of these connections is shown in Figure 4.7.




Figure 4.7: Architectural view of point-to-point Internet access

⁹⁶See, for example, (Northern Telecom Inc. and Amati Communications Corp., 1993), (Chow, Tu and Cioffi, 1991); (Cioffi, 1991); (Bingham, 1990); (Sistanizadeh, 1991); (AngelosBoscoComptonGu et al., 1993); (Kerpez, 1991); (Jones, 1991); and (Hsing and Lechleider, 1991).

⁹⁷See (Chen and Waring, 1994), especially Figure 1, p. 103. ADSL-2 and -3 make the telephone industry's approach look more and more like cable's Fiber to the Neighborhood approach, without the dynamic bandwidth allocation of cable's bus architecture.

In contrast to the shared-bus architecture of a cable LAN, the telephone network requires the residential Internet provider to maintain multiple connection ports in order to serve multiple customers simultaneously. Thus, the residential Internet provider faces problems of multiplexing and concentration of individual subscriber lines very similar to those faced in telephone Central Offices. Not surprisingly, similar types of solutions are adopted to deal with these issues.

The next sections explain in more detail how Internet service is provided over point-to-point telephone network connections, covering the following issues:

- How is digital data transmitted through the local telephone plant?
- How are IP packets exchanged over point-to-point connections?
- How are individual subscriber connections concentrated and multiplexed before being connected to the residential Internet provider's IP router?
- How effectively can the point-to-point architecture provide full-time Internet connections?

The final section identifies the open issues involved in providing Internet connections over ADSL lines.

4.2.1. Transmission of digital data

ISDN transmission is accomplished by the NT1 device described previously. This device incorporates modem functionality, encoding and modulating digital data onto an analog signal for transmission to the Central Office. The NT1 also provides protocol support, such as the call signaling (Q.931) used over the "D" channel.

NT1 devices are available from a variety of vendors. The NT1's encoding and modulation protocols need to match those used by the ISDN line card in the local Central Office switch. In theory at least, these protocols are standardized and compatibility should not be a big problem for customers.

4.2.2. Exchange of IP packets

The point-to-point approach is familiar to computer users accustomed to dialing in to workplace or commercial servers. As discussed in Chapter 2, however, most such connections today provide only terminal emulation, not full Internet connectivity. To provide peer Internet service, IP must be run all the way to the subscriber's computer. This computer must have an Internet IP address, run the TCP/IP protocol suite, and exchange IP packets over the point-to-point connection.

Since bandwidth is not shared over these point-to-point links, no media access protocol is needed. However, a protocol is still needed to map between IP and the physical reality of a point-to-point telephone line. The physical line delivers a stream of bits, while upper layer protocols expect to work with IP packets.⁹⁸ In between, protocol functionality is needed to “frame” the bytes and IP packets, i.e. identify the start and end of each byte and each packet. An intermediate protocol also needs to identify and resolve the address of the connecting station, providing a mapping between a telephone number and an IP address. This function is not necessarily trivial, particularly on a dialup line, on which the connecting station’s IP address may be dynamically or temporarily assigned.

Serial Line IP (SLIP) was an early, limited point-to-point protocol that was never officially standardized but was widely implemented anyway (Romkey, 1988). One of its limitations was that it assumed the use of IP; no other protocols could be multiplexed over the SLIP connection. To address this and other deficiencies, the Point-to-Point Protocol (PPP) has been developed and is being standardized by the Internet community (Simpson, 1993). PPP is supported by most vendors of Internet-compatible ISDN equipment at the time of this writing.

A version of PPP (called MP) that can treat the two “B” channels of ISDN as a single link is under development (Sklower, Lloyd, McGregor and Carr, 1994). Some equipment vendors also offer proprietary protocol solutions incorporating data compression functionality (Derfler, 1994). Ensuring compatibility between customer and Internet provider equipment, however, is a key consideration favoring the selection of standard protocols.

4.2.3. Concentration and multiplexing of connections

Figure 4.8 illustrates the concentration and multiplexing that are used by residential Internet providers. Concentration happens because the provider connects to the local telephone network with fewer lines than the total number of subscribers: not all subscribers are expected to connect to the Internet at the same time. The total number of **active** subscribers is thus limited by the number of connections the residential Internet provider maintains to the local telephone network.

⁹⁸See (Tanenbaum, 1988) for explanation of the functions performed by different communications protocol layers.

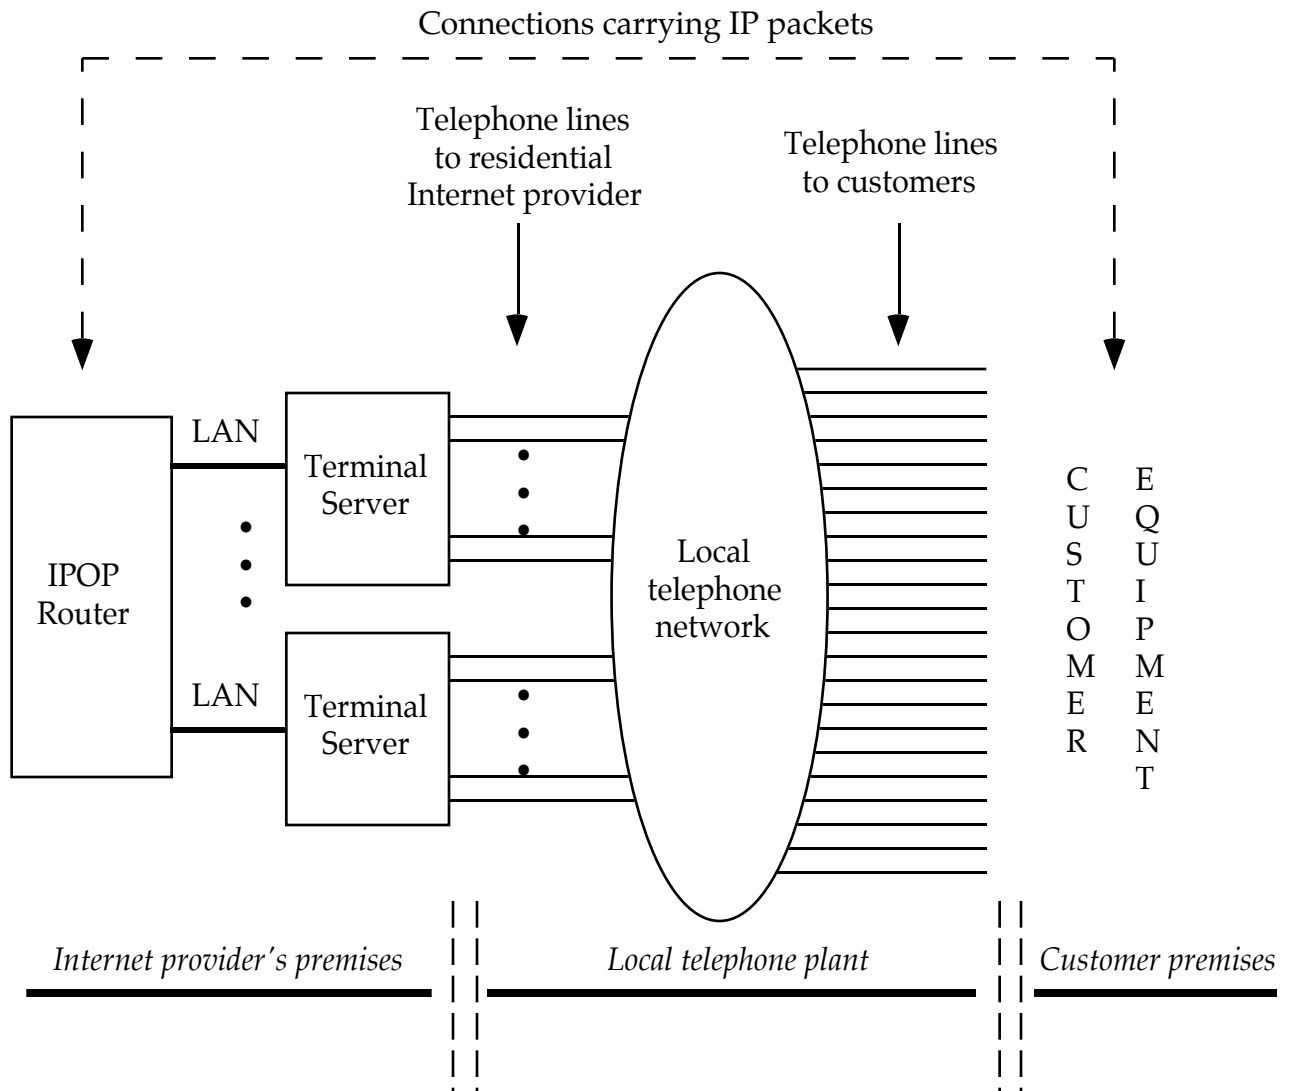


Figure 4.8: Detailed view of point-to-point Internet access

Even with this concentration, it does not make sense to plug each such connection directly into a router port; an impractically large number of physical ports would be needed, and they would be used inefficiently, since each line is not expected to be transferring data all the time that it is in use. In addition, router ports are generally built to talk to networks, not individual users. For these reasons, the individual lines are multiplexed (and sometimes concentrated as well) onto LANs by devices called *terminal servers*. These devices physically connect a large number of telephone lines to a smaller number of LAN (typically Ethernet) ports.⁹⁹ They also convert, or translate, between the LAN protocol used to talk to the IP router (e.g. Ethernet) and the point-to-point protocol (such as SLIP or PPP) used to

⁹⁹Terminal servers are usually built in a modular fashion, with the ability to add 8 or 16 phone ports at a time.

exchange IP packets with the subscriber over the phone line. Since the LAN protocol typically runs at much higher speed than the individual phone lines (e.g. 10 Mbps compared to 20 Kbps), terminal servers have the potential to concentrate a large number of phone lines even if all lines are transferring data simultaneously.

Although most terminal servers are built to handle either individual POTS lines (with modems) or ISDN lines, some can also handle aggregated lines, such as an ISDN Primary Rate Interface (digital T1, or 23B+D) carrying 23 individual circuits. Whether a provider chooses to connect to the local telephone network via individual or aggregate lines will depend on a number of factors, including:

- Availability of aggregate circuits or multiple individual circuits in the provider's location;¹⁰⁰
- Pricing policy of the local telephone company;¹⁰¹
- The number of subscribers the provider expects to serve.

When it can be used, the approach of using a few T1 concentrators to multiplex individual subscriber circuits into an IP router is simpler than the use of a larger number of smaller terminal servers, because it uses fewer wires and boxes. This approach has been adopted by Internex, Inc., the residential Internet provider studied to develop the Internet over ISDN model described in Chapter 4.¹⁰² Notice that if the number of T1 lines into the Internet provider grows large enough, an economic incentive is created for the Internet provider to co-locate its facilities with a telephone company Central Office, to minimize distance-sensitive T1 tariffs.

4.3. Evaluation of ISDN as Internet access network

The point-to-point telephone network gives the residential Internet provider an architecture to work with that is fundamentally different from the cable plant. Instead of multiplexing the use of LAN transmission bandwidth as it is

¹⁰⁰For example, a provider whose premises are located in a residential area may have difficulty procuring multiple individual circuits. As noted in (Waring, 1991), p. 1979, the telephone plant in residential areas may have as few as 1.3 lines per residence. Given an Internet provider's large expected demand for circuits, T1 circuits are likely to be a more economical option for both the telephone company and the Internet provider.

¹⁰¹For example, is the price of a Primary Rate Interface (PRI) circuit equal to, more or less than the price of 23 individual Basic Rate Interface (BRI) ISDN circuits? As (Network Express Inc., 1994) shows, this relation varies across different parts of the U.S. The distance from the provider's premises to the closest telephone Central Office may affect this tradeoff, since PRI pricing may depend on distance while BRI pricing generally doesn't.

¹⁰²Telephone conversation with Robert Berger of Internex, Inc., April 28, 1994. Notice that the T1 terminal server must perform the same protocol translation function as the POTS or ISDN terminal server, on 24 separate channels at a time.

needed, subscribers multiplex the use of dedicated connections to the Internet provider over much longer time intervals. As with ordinary phone calls, subscribers are allocated fixed amounts of bandwidth for the duration of the connection. Each subscriber that succeeds in becoming active (i.e. getting connected to the residential Internet provider instead of getting a busy signal) is guaranteed a particular level of bandwidth until hanging up the call.

Bandwidth

Although the predictability of this connection-oriented approach is appealing, its major disadvantage is the limited level of bandwidth that can be economically dedicated to each customer. At most, an ISDN line can deliver 144 Kbps to a subscriber, roughly four times the bandwidth available with POTS.¹⁰³ This rate is both the average and the peak data rate. A subscriber needing to burst data quickly, for example to transfer a large file or engage in a video conference, may prefer a shared-bandwidth architecture, such as a cable LAN, that allows a higher peak data rate for each individual subscriber.¹⁰⁴

Part-time connections

Another implication of the telephone network's connection-oriented approach is that full-time Internet access via ISDN would consume many more infrastructure resources than full-time Internet access via cable. Recall that a cable LAN's packet-based approach requires no shared resources to be dedicated to a subscriber who is not sending or receiving data, making full-time Internet connections economical even for low-volume customers. The same cannot be said for telephone network-based connections between subscribers and residential Internet providers.

First, recall from Figure 4.8 that each connected subscriber uses up both a terminal server input port and the telephone line connecting that port to the phone network.¹⁰⁵ Thus, a subscriber who needs a full-time connection requires a dedicated port on a terminal server. This is an expensive waste of resources when the subscriber is connected but not transferring data.

Resources must also be dedicated within the telephone network. Inside the "Local Telephone Network" cloud shown in Figure 4.8, individual subscriber calls traverse telephone company transmission facilities and one or more Central Office (CO) switches. When those calls are circuit-switched, shared switch resources (e.g. connections through the switch fabric) are dedicated to individual calls, whether or not data is flowing.

¹⁰³ISDN technology does not incorporate a migration path to a faster bit rate.

¹⁰⁴For example, today's broadband LAN modem models offer a minimum burst rate of 500 Kbps and a maximum of 10 Mbps. The average per-subscriber bandwidth is lower than the burst rate because the LAN bandwidth is shared among multiple subscribers.

¹⁰⁵If the terminal server handles a multiplexed format such as T1, each subscriber uses up one of the guaranteed slots on the multiplexed line and port.

Full-time Internet connections over circuit-switched lines are impractical. Too many of them would overload the switching and transmission facilities of the local telephone network, which is sized on the assumption that the “average” telephone call lasts only a few minutes.¹⁰⁶ They could also result in huge charges for subscribers subject to metered service, such as businesses, or subscribers located outside the residential Internet provider’s local calling area.

ISDN lines would seem to solve this problem by giving the subscriber the option of using packet instead of circuit switching on the 64 Kbps “B” channels, or using the packet-switched 16 Kbps “D” channel for data transfer in addition to call signaling. In practice, however, these options are not used, for the following reasons:

- The required packet switching protocol is X.25, an early protocol that was designed for use with error-prone lines.¹⁰⁷ The implementation complexity of X.25 makes it difficult to engineer for high performance. This issue is especially significant on the already low-bandwidth “D” channel. While public data networks have moved on to lighter-weight protocols such as Frame Relay, the telephone network is slower to adopt these types of advances.
- The protocol is implemented in an end-to-end fashion: the same protocol must be supported by both ends of the ISDN connection. Thus packet switching can only be used when compatible ISDN implementations are available at both the source and destination CO. Given that ISDN is still not ubiquitously deployed in North American central offices, and that variations still exist between different switch vendors’ X.25 implementations, such availability cannot be guaranteed.

In summary, most subscribers choose to use the “B” channels in circuit mode for data transfer.¹⁰⁸ Therefore, full-time Internet connections are not practical with ISDN.

One approach that is used to work around this restriction is *dialing on demand*. For example, a subscriber connecting a LAN to the Internet over a dialup telephone line can use a router that places a call over its serial (i.e. telephone) port only when data needs to be sent from the attached LAN. But dialing on demand is still not a perfect solution. First, dialing introduces a call setup delay, which may be tolerable when the purpose of the call is a

¹⁰⁶For example, the concentration illustrated in Figure 4.4 rests on voice telephony’s “3-3-3” assumption: the average call lasts three minutes, customers make three call attempts during the busy hour, and the channel is 3 KHz wide (White, 1993).

¹⁰⁷See (Berger, 1994), pp. 137-140, for more details about ISDN protocols.

¹⁰⁸More favorable circuit-mode tariffs are also be a consideration in some cases. See (Lindstrom, 1994).

remote login session that lasts half an hour, but unacceptable when the call is initiated on behalf of a user who has unknowingly accessed a remote World Wide Web server by clicking on a hypertext link.¹⁰⁹ Second, the customer can only be expected to dial the Internet provider when data needs to be sent; the customer has no way to tell when data needs to be received. Therefore, dialing on demand is not a satisfactory solution if the customer needs more than client access to the Internet (for example, the customer wishes to make her World Wide Web server available 24 hours a day).

Multiple simultaneous services

Recall that cable's broadband approach enables multiple devices, such as a computer and a TV, to use the network simultaneously. By providing two 64 Kbps channels, an ISDN line is supposed to enable the simultaneous use of a single telephone line for both the telephone and the computer.

In practice, to plug a telephone into a digital ISDN line, the phone must either be a special ISDN phone (one that digitizes the voice signal before it leaves the phone), or plugged into a Terminal Adapter (a device that can digitize the signal from an analog POTS phone). To avoid the purchase of new telephone equipment, many subscribers simply keep their existing POTS line for their phone service, while adding an ISDN line for their computer.¹¹⁰ This approach has the added benefit of allowing the subscriber who purchases appropriate inverse multiplexing equipment to use the additional "B" channel's bandwidth as needed for the computer.¹¹¹

In summary, two lines are required for most subscribers who wish to use the phone and the Internet at the same time. Because of the limited number of residential loops, this approach will not scale well to ubiquitous residential Internet connectivity.¹¹²

Scope of Internet Point of Presence

Once an Internet Point of Presence (IPOP) is connected to the telephone network, it can serve as a gateway to the Internet for anyone who can call it.

¹⁰⁹The extent of this problem depends on the length of the call setup delay. At the MIT Communications Forum on April 27, 1995, Robert Berger of Internex estimated that ISDN call setups take 1-10 seconds, depending on the details of the physical configuration over which the ISDN circuit is delivered. This compares favorably to the 4-10 seconds needed to establish a POTS connection plus the 30-60 seconds needed for a POTS modem pair to negotiate the use of a common speed and protocol (see also (Berger, 1994), p. 131 and 138).

¹¹⁰See notes on ISDN from lecture by Mitchell Kapor in (Gillett, Lampson and Tennenhouse, 1994), pp. 39-40.

¹¹¹See (Duncanson, 1994). The inverse multiplexing function is often built into ISDN computer networking equipment.

¹¹²As discussed above, ISDN's ubiquity is subject to other limitations as well. For example, subscribers whose loops are longer than 18,000 feet may be limited to Switched 56 service instead, in which data is transmitted at 56 Kbps and the additional 8 Kbps are used for in-band signaling. Loops with loading coils cannot run any form of digital service.

Given the worldwide scope of the telephone network, that is quite a large number of potential subscribers. In practice, the IPOP's subscriber population is determined by telephone network pricing policies. If enough telephone subscribers can call the Internet provider's access number without incurring steep ISDN usage charges, then it may be feasible to recover the IPOP investment even at very low Internet service penetration rates. This factor gives the ISDN approach a business advantage for getting started.

The telephone network's worldwide scope also addresses the traveler's problem identified in the previous chapter. The traveling subscriber can access the Internet by dialing his "home" IPOP, if he chooses to incur the long-distance toll charges (or the Internet provider offers an 800 number).

4.3.2. Asymmetric Digital Subscriber Line

Since Asymmetric Digital Subscriber Line (ADSL) technology has not been deployed beyond a few trials, many of the details of how it would be used and the equipment needed to support it have not yet been decided. This section discusses some of the issues that arise in considering how to provide Internet connections over ADSL; it does not describe a real implementation of residential Internet over ADSL, since none was found by the research undertaken for this thesis.¹¹³

Over a single pair of copper wires, ADSL provides a POTS channel, an ISDN D channel, and a unidirectional high-speed digital channel. ADSL's designers envisioned control packets (such as commands entered by a customer to order a video) flowing over the D channel, with movies flowing to the customer over the high-speed channel (the *video channel*). In currently available implementations, the D channel provides bidirectional 16 Kbps, while the video channel provides 1.5 Mbps toward the subscriber.

The POTS channel can be used to support the customer's existing telephone, but it is less obvious what to use for an Internet connection. We saw above that a 16 Kbps D channel is a low-performance solution for IP packets. But what if the D channel were only used to carry data from the customer, with the video channel used to carry data to the customer? This solution could be viable if most Internet traffic follows such an asymmetric pattern, but

¹¹³Between January and May of 1994, I contacted individuals in several organizations active in ADSL development, including Ameritech (ADSL standards committee chairperson), Amati Communications Corp. (development engineer), and Bell Atlantic (spokesperson). These discussions did not turn up any Internet over ADSL trials. Walter Chen, an ADSL researcher at Bellcore, and Robert Berger of Internex, Inc. discussed the possibility of Internet connections over ADSL (see (Berger, 1994), pp. 141-142). Bell Atlantic's ADSL press kit mentions numerous Video on Demand trials but no computer networking trials. Similarly, (Markoff, 1994) lists major interactive network trials currently in progress, showing television/video trials taking place over cable and telephone networks, but computer networking trials over cable networks only.

whether it does is the \$64,000 question. It is conceivable that client programs can be engineered to make do with limited upstream bandwidth (an X Windows client, for example, can send only mouse clicks back to its server, while receiving full screen bitmaps in return). It is less believable that access to residential servers, which could easily involve uploading of files (for example, viewing of digitized photographs), would follow such a pattern.

Recall that switching of the different services provided over ADSL is still an open question. If ATM cell switches are deployed to switch ADSL video channels, will the POTS and D channels also be cell-switched, or will they continue to be handled (in an evolutionary fashion) by existing ISDN-capable Central Office switches? The answer has implications for the type of customer equipment that would be needed to interface a computer to an ADSL line, as well as the equipment used to concentrate ADSL lines for switching, whether by IP router or ATM switch. For example, with the D and video channels both used to transfer IP packets, both would need to support a protocol (such as PPP or X.25) to map IP onto the physical channel. If the D channel used X.25 over ISDN while the video channel used ATM, more complex equipment would be needed to support the two different methods of IP data transfer.

In either case, ADSL does have the potential to support full-time Internet connections, since presumably neither the D nor the video channel would be circuit-switched. Also, since the POTS channel is separate, ADSL allows a subscriber to use the phone at the same time as the video channel. The X.25 packet switching that is already standard for the D channel would easily allow the computer and the television (more likely, the remote control) to share the use of this channel dynamically. To what extent the downstream channel can be shared by video and computer data, however, is an open question, with three possible answers depending on the nature of the protocols adopted for use on this channel:

1. The protocols build in the assumption that only video traffic will be sent; they do not incorporate the flexibility to describe different types of data. Under this scenario, the computer cannot use the downstream channel for non-video data traffic.
2. The protocols are flexible enough to support transmission of different types of data, but switching among the different types of data requires some form of call setup. This scenario might be adopted if the full downstream channel bandwidth is required to handle video successfully. Under this scenario, a video session would have to end (or at least be interrupted) before other types of data could be received; for example, the computer could not receive email while a video is being watched.

3. The protocols allow data and video traffic to share the channel dynamically, for example by placing both types of bits into Asynchronous Transfer Mode (ATM) cells. This method would provide the greatest flexibility to the subscriber by allowing multiple applications to appear to run simultaneously.

Since ADSL has not yet been widely deployed, its capability has evolved rapidly with advances in digital technology. Amati Communications Corp., for example, is developing ADSL equipment that they expect will be able to transmit 6 Mbps downstream and either 2B+D (144 Kbps) or H0 (384 Kbps) upstream, within the next year or two.¹¹⁴ While a 16 Kbps maximum upstream rate obviously limits ADSL's usefulness for computer networking, 384 Kbps could open up many more possibilities. Still, the asymmetric transmission capability will always limit ADSL's applicability. Computer applications show a remarkable ability to use as many resources as are available to them. With more resources always available to data flowing toward the subscriber, applications that would naturally cause significant data flow from the subscriber will always appear limited.

This chapter and the preceding one have discussed in general terms the technologies involved in providing Internet connections over the residential cable and ISDN infrastructures. The next two chapters move from generalities to the specific technology details and capital costs of the two case studies, first for cable and then for ISDN.

¹¹⁴(Northern Telecom Inc. and Amati Communications Corp., 1993), p. 1; conversation with Ron Hunt of Amati, February 2, 1994.

Chapter Five

Cost model of Internet over cable

How much does it cost to connect a residential computer to the Internet over the local cable plant? The answer to this question depends on a number of factors. Some parameter values are obvious, such as the maximum speed of the connection. Others are less clear: how many customers can share a neighborhood network?

Through a case study of the “PSICable” Internet service introduced for Cambridge, MA on March 8, 1994 by Continental Cablevision, Inc. and Performance Systems International, Inc. (PSI), I have developed a model of the capital investment required to provide residential Internet service over the cable plant. The model incorporates the effect of multiple factors on the cost to provide service. This chapter explains the cost items included in the model and describes the parameters that affect the model’s results. The following chapter develops a similar model for Internet connections provided over the local telephone network using ISDN technology.

Both models look only at the marginal cost of providing Internet connections; the costs of providing traditional cable TV and telephone service are not included. And neither model accounts for technology investments that are used for Internet service but are being made for other purposes. For example, before offering Internet service, Continental Cablevision upgraded its Cambridge plant from coaxial tree and branch to a Fiber to the Neighborhood (FTTN) architecture. However, this upgrade was primarily carried out to improve reliability and reduce maintenance costs.¹¹⁵ Only the portions of the upgrade that were specifically required to support Internet service are included in the model.

¹¹⁵See (Stix, 1993), p. 102 for a description of such upgrades. In that author’s view, “After the election of Clinton and Gore, some of these capital upgrade programs were politically corrected to become information superhighways.”

The cost figures used in the model reflect what it would cost an Internet service provider to purchase equipment that is currently available from both computer networking and cable transmission equipment vendors. Equipment costs are based on price lists published by vendors, with volume discounts factored in where appropriate. Since the market for computer networking equipment is extremely competitive and digital technology continues to advance rapidly, cheaper and more capable products are constantly appearing. Therefore, the costs reported here should be considered as conservative estimates. Incidental costs (under \$100), such as cables and connectors needed at the head end, are not included in the model.

5.1. Overview

Figure 5.1 shows a top-level view of the cable cost model. This model was generated using the program Demos (“DEcision MOdeling Software”) from Lumina Decision Systems, Inc. Demos incorporates an object-oriented graphical user interface to spreadsheet-style data manipulation, providing a straightforward way to express in equations the relationships among the different components of the model. The diagrams of the models shown in this report are extracted from Demos’ graphical user interface.

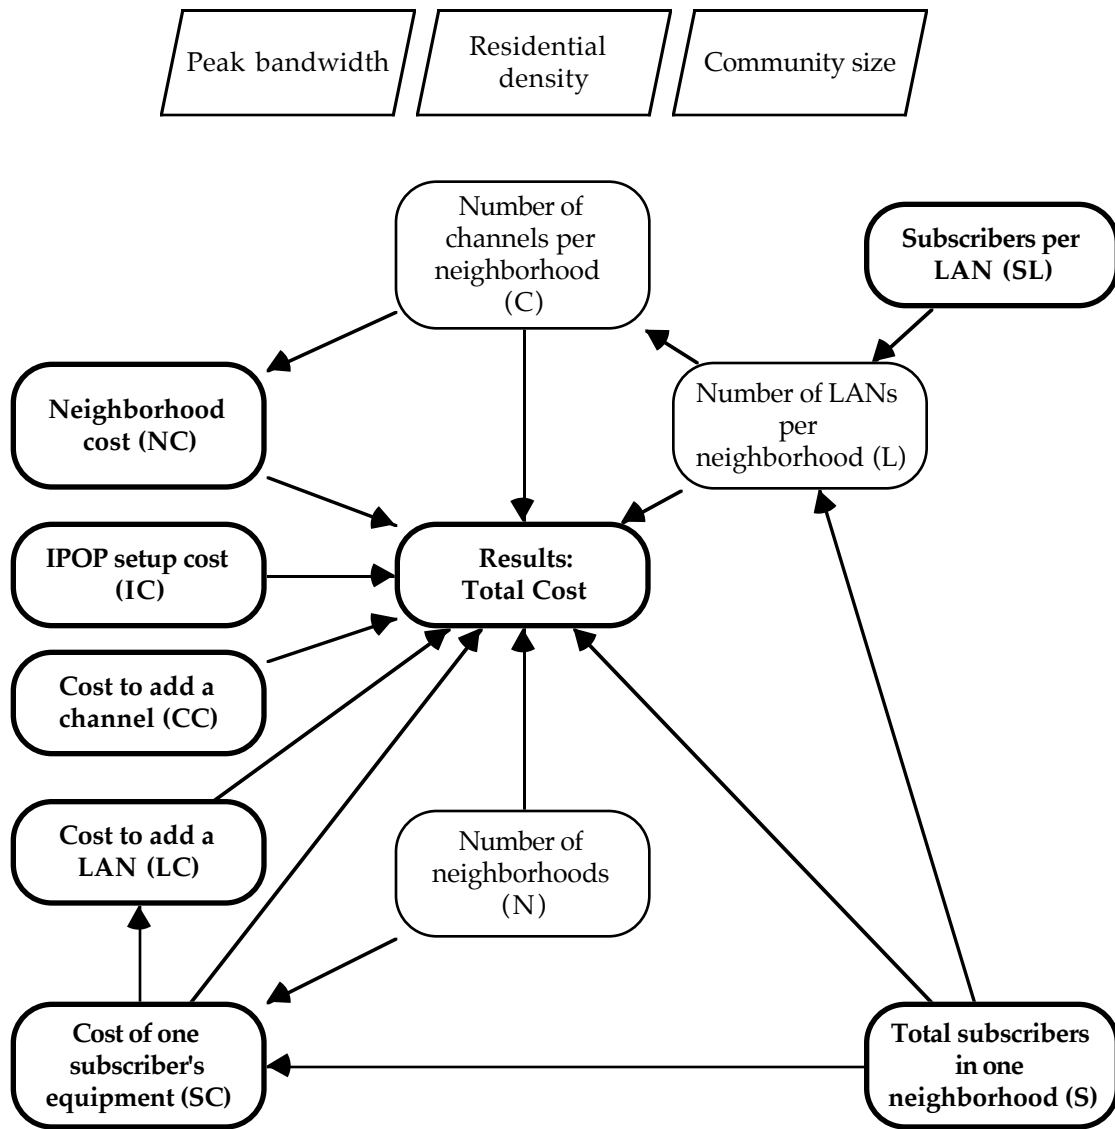


Figure 5.1: Capital cost model of Internet over cable

The five nodes on the left side of the figure comprise the technology reference model. The three rhomboids across the top of the figure represent parameters whose values can vary, influencing the resulting cost to provide service. On the right side of the figure are two sub-models: one incorporating a variety of possible traffic patterns (“Subscribers per LAN”) and the other modeling market demand (“Total subscribers in one neighborhood”). Results are shown by the node in the center of the figure.

This chapter details the model’s inputs and variables, discussing first the technology reference model, then the models of market demand and traffic levels. Results are described in Chapter 7.

5.2. Technology reference model

At a conceptual level, the technology needed to connect residential computers to the Internet can be thought of as having three components:

- An Internet Point of Presence, or IPOP, consisting of an Internet Protocol (IP) router physically connected to, and able to route traffic appropriately to and from, the rest of the Internet;
- Some form of community network that physically connects the Internet Point of Presence to individual subscribers' computers;
- Hardware and software allowing those subscriber computers to connect to the community network and use Internet services.

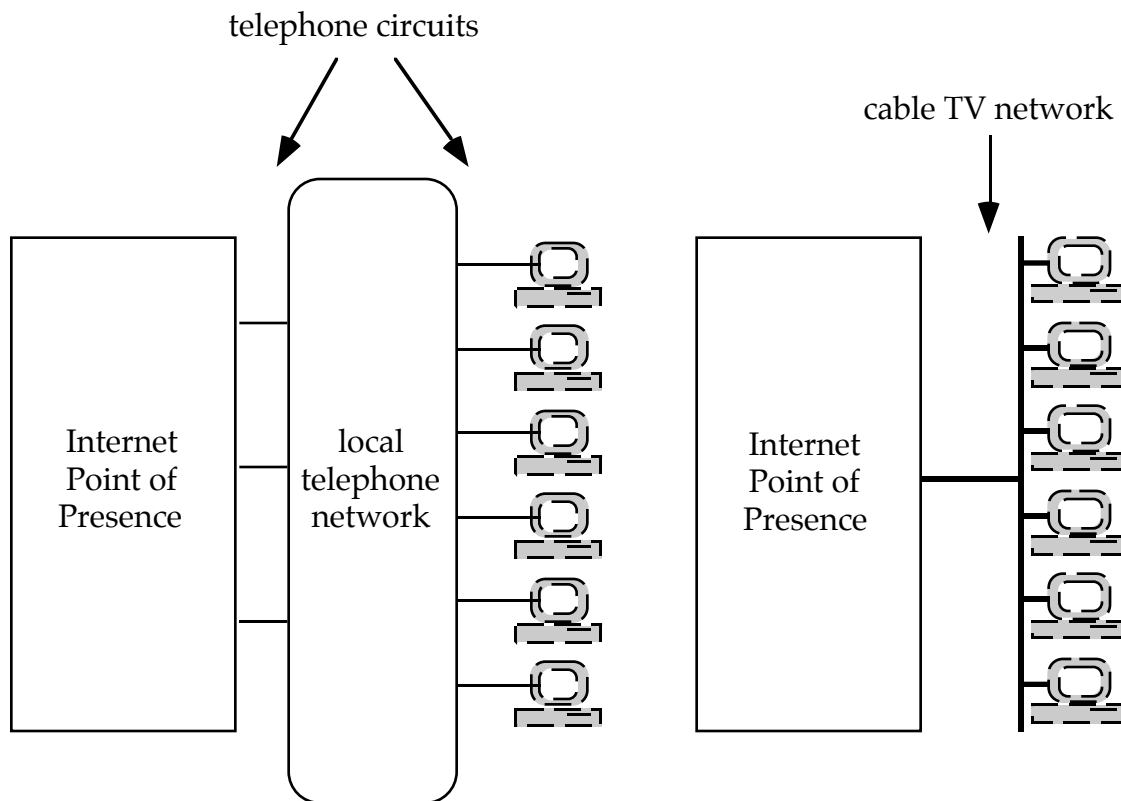


Figure 5.2: Components of residential Internet connections

Figure 5.2 illustrates how this conceptual breakdown maps to local telephone and cable networks. The Internet Point of Presence (IPOP) component is the same for the two approaches; what differs is the nature of the community network and the associated equipment required for subscribers and the IPOP to connect to it. The telephone network approach is described in more detail in Chapter 6, while this chapter details the cable approach.




Figure 5.3: Residential Internet connections over cable TV plant

Figure 5.3 shows how residential computers are connected to an Internet Point of Presence, using a pair of cable TV channels to serve as a Local Area Network (LAN). Setting up this system requires the following five steps:

1. Each neighborhood's transmission plant must be upgraded to pass traffic to the head end, enabling two-way transmission of data;
2. For each neighborhood, one or more channel pairs must be dedicated to carrying the traffic of one or more community Local Area Networks (LANs);
3. A community-wide Internet Point of Presence (IPOP) must be set up to serve all the community LANs;
4. Each subscriber must be connected to a community LAN;
5. Each community LAN must be connected to the IPOP.

The technology and cost of each of these steps is now discussed in detail.

5.2.1. Upgrade neighborhood transmission plant

This component is represented by the node labeled "Neighborhood cost (NC)" in Figure 5.1. A neighborhood in this context is the group of residences served by a single fiber and its attached coaxial tree and branch network.¹¹⁶

¹¹⁶The cable operator may choose to run multiple fibers to a single neighborhood node, and attach separate coaxial networks to each fiber. The model considers each fiber as serving a separate neighborhood.

Given that the initial motivation for migration of the cable plant to a Fiber to the Neighborhood architecture is reduced maintenance expense, the model incorporates the assumption that Internet connectivity offered today is the first service to require the FTTN cable plant to transmit signals in the “upstream” direction, that is, from the customer to the head end. Although other two-way services are also envisioned for the upgraded plant, most (such as video-on-demand) do not require as much upstream bandwidth as Internet service. Therefore, allocating the full cost of the two-way transmission upgrade to Internet service is a reasonable, if conservative, assumption.

With the FTTN architecture, each neighborhood's transmission plant can be upgraded separately. This upgrade has two components: one for the fiber portion of the plant, and one for the coaxial cable (“coax”). All costs listed here are based on the assumption that upstream transmission is accomplished over the channels typically allocated for this purpose by equipment vendors: 5-42 MHz.¹¹⁷

Upstream transmission over the fiber requires the installation of an upstream laser transmitter (at the neighborhood node) and matching receiver (at the head end). The model uses \$9100 for the cost of this equipment, based on prices published by C-COR, a cable plant equipment vendor.¹¹⁸

Upgrading the coax portion of the plant is more complex, since every amplifier must be configured with the appropriate modules to amplify upstream traffic. The incremental cost of two-way amplifiers is about \$225.¹¹⁹ The total amplifier upgrade cost then depends on how many amplifiers exist in the coaxial cable portion of the plant. The number of amplifiers needed is a function of both the length of the cable (since signals attenuate with distance) and the number of times the cable is physically split (since each split results in some signal loss). The model makes the simplifying assumption that the cable length is always one mile, while the number of cable splits rises as housing gets more dense.

¹¹⁷ The channel at 5 MHz is considered hard to use because of interference from common household equipment. Vendor equipment (such as that available from Zenith Communication Products and C-COR Corporation) typically provides for a maximum of 4-5 upstream channels, and only uses the channels between 12 and 42 MHz.

¹¹⁸ C-COR Corporation, August 1993 Commercial Systems Products price list, pp. 11-12, lists matching mid-price lasers for \$5400 (transmitter) and \$3700 (receiver). According to John Tucker, sales engineer for C-COR (1/27/95 telephone conversation), lower-priced models are not of sufficient quality to be used for data transmission.

¹¹⁹ According to the C-COR “Cable TV Products” price list of December 1993 and Mike Morris, C-COR sales representative, this additional cost is the same whether purchased as a modular upgrade to an existing amplifier or as an additional feature of a new amplifier. This cost includes \$25 for installation, which requires about one-half hour of a technician's time to re-balance the upstream channel.

The reasoning behind this assumption is based on the observation that Fiber to the Neighborhood is essentially a cellular approach, in the sense that each fiber re-uses the same channel frequencies.¹²⁰ Thus each neighborhood can be thought of as a cell of a particular size, measured both in terms of geography and number of homes served. The determination of actual cell boundaries

reflects a complex trade-off of three factors:

- Real-world obstacles such as the location of rivers, railroad tracks, etc.;
- The desire to limit the length of coaxial cable runs to a mile or less, so that more expensive, high-powered amplifiers are not needed;¹²¹
- The desire to keep the number of homes in each cell within a range that, at the low end, keeps cell cost overhead to a reasonable percentage, while at the high end, allows sufficient bandwidth for applications such as Pay Per View that depend on re-use of bandwidth among cells. The model assumes that this range is 500-1500 homes.¹²²

Based on these tradeoffs, the model assumes that cell sizes are set to maximize the number of homes per cell (up to the limit of 1500 homes) while simultaneously limiting geographic cell sizes to under one mile in radius.¹²³ Under this assumption, the main consideration driving the number of amplifiers is the extent of cable splitting, i.e. the residential density. The more subscribers served per cell, the more times the cable must be split and therefore the larger the number of amplifiers needed. The model assumes that four amplifiers are needed per cell in suburban neighborhoods (i.e. 500 homes per cell), and ten in urban neighborhoods (1500 homes per cell).¹²⁴

Table 5.1 summarizes the resulting total capital cost of a neighborhood transmission upgrade, including the up and downstream lasers for the fiber, at different ends of the range of residential densities.

¹²⁰For broadcast TV, we would expect the same information to be placed on each re-used channel. The cellular approach makes it possible for applications such as Pay Per View and LAN service to carry different information on each re-used channel in each cell.

¹²¹In (Stix, 1993), p. 102, this rule is summarized as follows: "In the upgraded networks, optical cables reach a point a mile or less from a group of homes. Then existing coaxial cable provides the rest of the link."

¹²²This assumption is based on estimates provided by Will Biedron of Continental Cablevision in March of 1994. Other operators of Fiber to the Neighborhood networks may use a different range; for example, Scott Wilcox of Southern New England Telephone reported a target cell size of 200 homes, consisting of a fiber plus 4 coaxial cable trees serving 50 homes each, with 1-3 amplifiers needed for each tree (telephone conversation, March 9, 1994). Chapter 7 reports the results of sensitivity analysis for the cell size (and thus number of amplifiers) parameter.

¹²³Given the high residential density of Cambridge, MA (the town involved in the PSICable case study), one would not expect any such cells to end up with fewer than 500 homes passed. This result might not hold in less densely populated areas.

¹²⁴This assumption is based on the need for an amplifier every quarter-to-third mile of cable length in suburban areas ((Baldwin and McVoy, 1988) and discussions with Will Biedron of Continental Cablevision in March, 1994).

Table 5.1: Cost of neighborhood transmission upgrade

Type of neighborhood	Homes passed	Total upgrade cost
suburban	500	\$10,000
urban	1500	\$11,350

5.2.2. Convert channel pair to support a LAN

This component is represented by the node labeled “Cost to add a channel (CC)” in Figure 5.1.

A cable LAN is constructed out of a pair of separate channels, each transmitting data in a different direction. For one LAN subscriber to send data to another, signals must be transmitted to the head end on the upstream channel, converted from the upstream to the downstream channel’s frequency range, and sent to the receiving subscriber on the downstream channel. The conversion is accomplished by a device called a frequency translator or converter, which logically creates a “channel pair” out of two independent 6 MHz channels. The function performed by this device is independent of the type or speed of LAN(s) for which the channel is used.¹²⁵

Current models of frequency translators cost about \$3,000.¹²⁶ This expense can be viewed as a per-channel portion of the transmission upgrade cost.

5.2.3. Set up an Internet Point of Presence

This component is represented by the node labeled “IPOP setup cost (IC)” in Figure 5.1. It represents the fixed cost of equipment required to connect a community’s neighborhood LANs to the Internet, so that subscribers can access Internet users and services beyond the local community. This equipment consists of one or more Internet routers, one or more external communication circuits physically connecting that router to the rest of the Internet, and network management equipment.

The number and bandwidth of neighborhood LANs to be connected to the Internet Point of Presence (IPOP) determines the number and capacity of routers and circuits needed, resulting in the total IPOP cost. Figure 5.4 illustrates this relationship.

¹²⁵Since the frequency translator effectively “binds” spectrum together, and the same spectrum is re-used among neighborhood cells, one frequency translator could actually be used for more than one cell, if the different cell’s signals were combined together with some form of custom multiplexing equipment. This would allow a single LAN to be shared among subscribers in more than one neighborhood, an approach that might be practical for bootstrapping usage on higher bandwidth LANs such as the 4 Mbps service. The model does not include this non-standard option.

¹²⁶Appendix B to Commercial Price List, Network Products for Cable TV, Zenith, March 1, 1994; LANCity Corporation product literature, October 1993.

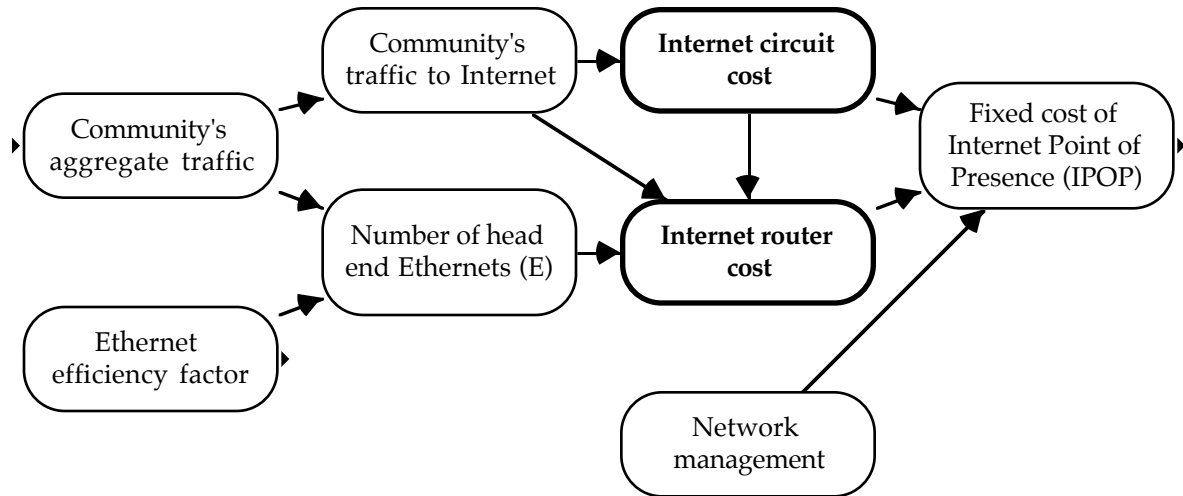


Figure 5.4: Model of Internet Point of Presence cost

The aggregate traffic generated by the community on all the neighborhood LANs consists of two components: traffic destined within the community and traffic destined to the rest of the Internet. The model incorporates the conservative assumption that 100% of the LAN traffic is destined to the rest of the Internet.¹²⁷ Under this assumption, the bandwidth of external circuits connecting the router to the rest of the Internet must be greater than or equal to the aggregate traffic generated by the community.

While the aggregate community traffic level grows smoothly in the model as LANs are added, the speed of external circuits follows a step function, reflecting the price points available for purchasing telecommunications lines of different capacities. Because the underlying telephone network uses T1 (1.5 Mbps) and T3 (45 Mbps) as basic units of capacity, “multiple T1” or “fractional T3” services require extra multiplexers, which, depending on the number required, may not be cost-effective relative to T3 services—especially competitive T3 services such as microwave links.¹²⁸ The model therefore assumes that aggregate bandwidth levels below 1.5 Mbps can be served by

¹²⁷This assumption is likely to be true when the service is first deployed and has few users, so it is valuable for determining fixed entry cost. The true value of this parameter could depend on factors such as the location of electronic mail servers for cable LAN customers, the location of telecommuters’ employers, etc. The chosen value represents a worst-case assumption.

¹²⁸This analysis is based on a conversation with Dan Long, NEARNet Network Analysis Manager, April 20, 1994. For example, multiplexing 4 T1 circuits together turns out not to make economic sense, because the total circuit and multiplexer cost is close to the cost of a T3 circuit. Although 10 Mbps microwave circuits can be purchased, their cost is also sufficiently close to the cost of a microwave T3 circuit that the model assumes a T3 would be purchased instead. The T3 circuit estimates are based on local microwave service providers, which generally offer much lower prices than the local telephone company.

purchasing a T1 circuit, while higher levels require purchasing one or more T3 circuits.¹²⁹ In addition to the circuit itself, a high-speed digital adapter is required—at each end of the Internet circuit—to send and receive digital data. This device is called a CSU/DSU, for Channel Service Unit/Data Service Unit. The resulting total circuit costs are summarized in Table 5.2.¹³⁰

Table 5.2: Capital cost of Internet circuit¹³¹

Equipment	Type of circuit	
	T1	T3
Purchase of circuit	\$24,000	\$55,000
CSU/DSU (modem)	\$2,000	\$20,000
TOTAL	\$26,000	\$75,000

The aggregate community bandwidth level also affects the type and number of routers required for the Internet Point of Presence. Routers switch packets among two types of ports: *network* ports, which LANs plug into, and *serial* ports, used for the point-to-point communication circuits that connect routers to other routers. The necessary router configuration (and hence cost) is determined by the number and speed of each type of port required. For example, when the aggregate community traffic requires a T3 external circuit, a high-end router is needed to support a High Speed Serial Interface (HSSI) port. A requirement for large numbers of network ports can similarly drive a need for one or more high-end routers.

The number of network ports needed is not necessarily the same as the number of cable LANs. Router network ports accept standard LANs such as Ethernet (10 Mbps) or FDDI (100 Mbps). Since the cable LANs in the model

¹²⁹The aggregate bandwidth level is based on adding together individual user's **average** bandwidth requirements. It is therefore possible for users on LANs that offer 4 Mbps of peak bandwidth to consume less than 1.5 Mbps of average aggregate bandwidth and require only a T1 external circuit. A provider offering such a configuration would have to set user expectations accordingly, since the peak 4 Mbps bandwidth would only be available within the community and not to the rest of the Internet. Such a configuration is likely to be sufficient only for very small numbers of subscribers.

¹³⁰Notice that the cost to link the head end to the rest of the Internet might be quite different if the connection were provided over fiber linking multiple Internet-capable head ends.

¹³¹Circuit and CSU/DSU prices are based on quotes from Peter Caparso, NYNEX sales representative (circuit installation prices, 4/22/94); Bill Ally of Allcom (CSU/DSU prices, 4/27/94); and Jim Naro and Dan Long of NEARNet (circuit and CSU/DSU costs). All costs assume local (as opposed to long-distance) circuits. T3 costs are based on the purchase price of a microwave T3 circuit. T1 circuit costs are confirmed by (Berger, 1994), p. 141. To compare a leased T1 accurately against a purchased T3, the T1 cost in the table includes not only the installation fee (about \$2000) but also the monthly fee (about \$600 based on typical mileages in the Cambridge area) capitalized over 3 years. Three years was chosen since it is the depreciation cycle used for Internet service provider equipment (electronic mail from Robert Berger of Internex Inc.).

operate at lower bandwidths (500 Kbps or 4 Mbps), they can be multiplexed together before being plugged into router network ports, thus reducing the number of network ports required.¹³² The model assumes that the IPOP router offers 10 Mbps Ethernet ports and that traffic can be multiplexed onto those ports at 60% efficiency.¹³³ Thus each 6 Mbps of aggregate LAN bandwidth requires one 10 Mbps Ethernet port.

In sum, the aggregate LAN bandwidth determines the speed and number of external circuits to the Internet as well as the number of Ethernet ports required on the IPOP router(s). Based on these three variables, the actual router configuration can be derived from the options offered by router vendors. The model incorporates configuration information from three commonly-deployed routers: the 3Com Netbuilder II low-end and mid-price models, and the Cisco 7000 model at the high end.¹³⁴ Sample configurations and prices are summarized in Table 5.3.

Table 5.3: Capital cost of one Internet Protocol router¹³⁵

Speed of Internet circuit(s)	Number of 10 Mbps Ethernet ports		
	1-3	4-7	8-24
T1 (1.5 Mbps)	\$15,000	\$30,000	\$60,000
T3 (45 Mbps)	\$60,000	\$60,000	\$60,000

¹³²For example, two 4 Mbps cable LANs could be connected to a single 10 Mbps Ethernet segment using intelligent bridges that would ensure that each LAN's traffic is not seen on the other LAN.

¹³³This factor is based on the analysis of Ethernet's multiplexing efficiency presented in (Metcalfe and Boggs, 1976).

¹³⁴3Com product literature for the Netbuilder II low- and mid-price routers shows that they can support either 3 or 7 Ethernet ports (respectively) and 1 T1 serial port, for \$15,000 or \$30,000 respectively. The Cisco Systems Products Catalogue (February 1994, Cisco 7000 Configuration Worksheet pp. 39-40) shows that the Cisco 7000 high-end router supports 5 slots, each of which can contain either an HSSI or 6 Ethernet ports. Other types of ports are also available, e.g. one FDDI port per slot. These ports are more expensive than Ethernet ports. According to Ron Hoffman of MIT's Distributed Computing and Network Services, Ethernet ports are inexpensive enough that \$60,000 is a reasonable approximation to the cost of the router with any number of Ethernet ports populated. (MacKie-Mason and Varian, 1993), p. 6, lists a cost of \$100,000 for such a router, but this figure includes amortized operating and service costs.

¹³⁵The algorithm for determining the router model is as follows: If the aggregate traffic to the Internet requires a T3 circuit, then use a Cisco 7000. Otherwise, use a 4-port 3Com Netbuilder II if the aggregate traffic can be handled by 1-3 Ethernet ports, else use an 8-port Netbuilder II. More than one IP router is necessary if the combined number of slots needed for Ethernet ports (6 per slot) and (possibly multiple) serial lines exceeds 5.

The final component of IPOP cost is network management equipment, estimated at \$10,000.¹³⁶ The model assumes that all LAN and IPOP equipment is compatible with the Internet standard Simple Network Management Protocol (SNMP). Given this compatibility, any small computer (such as a low-end workstation or PC) can run openly available SNMP monitoring software. Since most of this software is targeted to Unix platforms, the model's estimate includes the cost of a Unix-capable computer in addition to the software cost. This computer can also serve as a console for head end equipment.

Since all the neighborhood cable LANs physically connect to the community's head end, and the neighborhood LANs must physically connect to the IPOP router, it is most convenient to locate the IPOP router at the head end. Therefore, even an established Internet service provider must create a new IPOP to serve cable LAN customers.¹³⁷ However, servers to provide network news and electronic mail do not have to be co-located with the head end; they can be anywhere on the Internet and operated by any entity, not just the cable Internet service provider. In fact, some customers (especially telecommuters) will prefer to use servers they already have free access to (their corporate or university netnews server, for example). Others may prefer to leave their home computer on all the time, so that they can receive email directly, instead of gathering mail in bunches from an email server. In any case, provision of these types of services is optional for at least some of the customer base, and could be purchased independently of the Internet service provider by the rest of the customers. Therefore, server costs are not included as part of the IPOP cost.¹³⁸

Costs to operate equipment and publicize the service (i.e. operational and marketing costs) are beyond the scope of this report, and are not in the capital cost model.

Table 5.4 summarizes the minimum cost to set up an Internet Point of Presence depending on the level of aggregate traffic generated by the community. At high penetration levels and/or in large communities, this level may be high enough that multiple routers and connections are needed and the actual IPOP cost is higher than the totals shown in Table 5.4.

¹³⁶This estimate is based on discussions with Christopher Lindblad of MIT's Telemedia, Network and Systems Group (Laboratory for Computer Science) and Mark Graham of Pandora Systems.

¹³⁷The IPOP's router and Internet circuit may also serve traditional dialup or leased telephone line Internet subscribers, creating an economy of scope.

¹³⁸Providing these types of services is likely to be profitable, especially given the economy of scope of providing multiple services from the same server. Since these services are not essential to providing connectivity, however, they are accounted for separately.

Table 5.4: Minimum capital cost of Internet Point of Presence

Equipment	Community's aggregate traffic	
	Under 1.5 Mbps	Over 1.5 Mbps
Router	\$15,000	\$60,000
External connection	\$26,000	\$75,000
Network management	\$10,000	\$10,000
TOTAL	\$51,000	\$145,000

5.2.4. Connect Subscriber to the LAN

This component is represented by the node labeled "Cost of one subscriber's equipment (SC)" in Figure 5.1. This cost must be incurred for each customer who subscribes to the service.

For a subscriber to connect to the cable LAN, equipment is needed to perform the following functions:

- Provide a physical interface to the subscriber's computer.
- Provide a physical interface to the cable, including broadband modem functionality.
- Support the cable LAN's media access protocol. If the subscriber has more complex equipment than a single computer, support for routing protocols may also be needed.
- Support features needed for the residential Internet provider to manage the network effectively. Such features might include traffic filtering capabilities (e.g. to prevent the insertion of traffic that might somehow harm the network, such as an Ethernet broadcast packet), support for the Simple Network Management Protocol (SNMP), or other features designed to aid the provider in troubleshooting network or subscriber equipment problems.

This collection of functions could be implemented in a variety of ways, ranging from complete integration into the subscriber's equipment to a separate, standalone box. An example of complete integration would be a single plug-in card containing a cable modem chip set and networking software. The advantage of this approach would be its potential for a lower cost product. Getting to that point, however, would require expensive and time-consuming development, since as of 1994, cable modem protocols are not standardized and no such chip sets exist. This approach also runs into two practical problems: first, different subscribers have different types of computers (e.g. IBM PC vs. Apple Macintosh vs. Unix workstation), so one type of plug-in card will never be enough; and second, requiring a card to be

plugged into each subscriber's computer demands much more skill and staff support than requiring a wire to be plugged into a box. In contrast, the separate box approach gains economies of scale in production, inventory management, and customer support by allowing the service provider to use the same type of equipment for all customers. On the other hand, the separate box also requires a separate interface on the subscriber's computer, which may require additional expense.

The PSICable service offering studied for this report adopted the separate box approach. The box design was constrained by the following practical matters:

- The box needed to be produced quickly, since the service was slated for availability only 7 months after it was announced.¹³⁹ Therefore, existing components were preferable.
- The interface to the subscriber computer needed to be something standard and widely available; Ethernet was selected for this purpose.¹⁴⁰
- The cable LAN interface components needed to be affordable to residential subscribers. At the time the PSICable service was being developed, the major vendors of symmetric cable LAN equipment were Zenith Communication Products and LANCity Corporation. The LANCity equipment provides a 10 Mbps LAN, incorporating an RF modem and Ethernet bridge into one \$7100 product, intended for connecting institutions (such as bank branches) together.¹⁴¹ This high price put the LANCity equipment out of reach for residential Internet service. In contrast, Zenith offered two cable LAN products, the 0.5 Mbps "Homeworks" and the 4 Mbps "LAN4000," for \$500 and \$900, respectively.¹⁴²

The Zenith products, which physically consist of a broadband modem attached by a wire to a plug-in "Z-LAN" card designed only for IBM PC architecture computers, dictated the IBM PC architecture used in PSICable's subscriber box. Figure 5.5 shows this box, consisting of a PC with a Z-LAN

¹³⁹Development of the service was announced in August of 1993 (1993; Carnevale and Keller, 1993) while the service itself was announced in March of 1994 (Sandberg, 1994).

¹⁴⁰It is not clear how common LAN interfaces are among the installed base of home computers, but Ethernet is generally assumed to be the most common among those that do exist.

¹⁴¹LANCity Corporation product literature, October 1993. Conversation with Anthony DiSessa, LANCity sales representative, March 7, 1994. The LANCity product is re-sold by Digital Equipment Corporation as its "Channelworks" product.

¹⁴²Appendix B to Commercial Price List, Network Products for Cable TV, Zenith, March 1, 1994.

card to talk to the cable modem and an Ethernet card to talk to the subscriber's computer.¹⁴³ The use of this box has several implications:

- Because the box provides an Ethernet port, the subscriber who already has an Ethernet need only plug it in. The subscriber who has only a single computer must also have an Ethernet card to talk to the box, even though the single computer is not attached to a LAN within the residence.
- The subscriber equipment includes a minimal PC, providing a convenient platform to support the network management needs of the residential Internet provider.¹⁴⁴

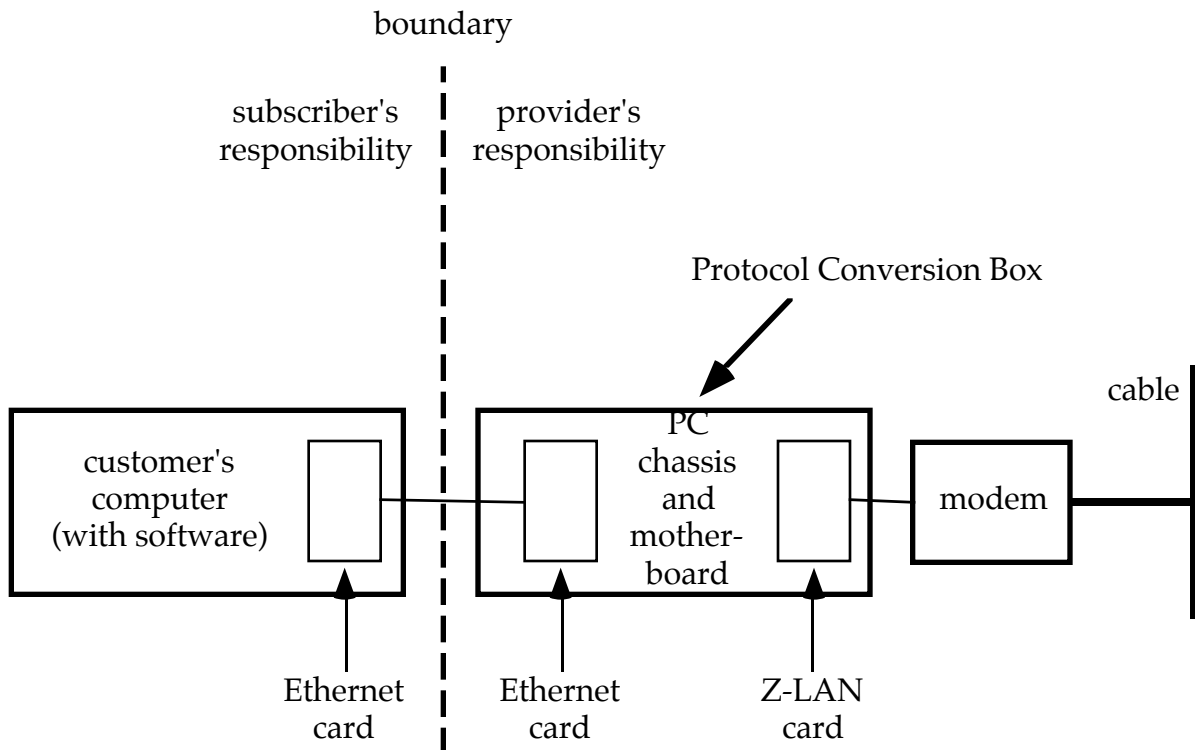


Figure 5.5: Subscriber equipment in the PSICable case study

¹⁴³My thanks to Will Biedron of Continental Cablevision for discussing the architecture and approximate cost of this equipment with me.

¹⁴⁴PSICable's alternative would have been to insert the Z-LAN card directly into the subscriber's computer. Although this approach eliminates the need for the subscriber to have a standard Ethernet card, it would also limit the provider to serving only customers with IBM PC architecture computers. The network management capabilities of the separate PC box were also important in the case study, since the Z-LAN cards available at the time the service was developed were limited in this area (for example, they did not initially support SNMP). Notice that including a PC chassis in the subscriber equipment also leaves the door open to serving customers who do not already own computers, although this does not appear to have part of the motivation for the approach.

The subscriber equipment costs used in the model reflect the boundary shown in Figure 5.5, representing the division of equipment responsibilities between the provider and the subscriber in the case studied (i.e. the PSICable service offering). This division is not the only possibility; one could also imagine a full-service provider supplying everything including the subscriber's Ethernet card and software, or alternatively, in a world with more standardized cable access protocols, a subscriber procuring her or his own cable LAN interface equipment from a variety of vendors.

In the case study, the subscriber's responsibility includes the home computer, a standard Ethernet card for that computer, and software including the TCP/IP protocols and network applications (e.g. Mosaic). This software can be purchased commercially but is also available for free (Rickard, 1994). The cost of an Ethernet card ranges from \$60 to about \$200, depending on the type and number of computers to be connected.¹⁴⁵ The model does not include the cost of any equipment that is the subscriber's responsibility in the case study.

The costs that are included in the model are those of the LAN interface equipment (the Zenith Z-LAN card plus broadband cable modem) plus those of the PC chassis, motherboard, and Ethernet card (part of the standalone box, and collectively referred to in the model as the "Protocol Converter"). As noted above, the Zenith equipment is available for \$500 (0.5 Mbps Homeworks) or \$900 (4 Mbps LAN4000). In addition, volume discounts are available. The model assumes a 25% discount on purchases of 1000 or more Homeworks interfaces, and 100 or more LAN4000 interfaces.¹⁴⁶ The model further assumes that with wholesale, volume purchases of PC equipment, the Protocol Converter can be built for \$400.¹⁴⁷

The Protocol Converter box shown in Figure 5.5 performs the function of an extremely low-end, fixed configuration router: it routes packets between the cable LAN on one side and standard Ethernet on the other. If the subscriber has either a single computer or a single Ethernet, then this protocol converter is all that is needed. Some subscriber configurations, however, require more sophisticated equipment. For example, if the subscriber's LAN uses technology other than Ethernet, or incorporates multiple LANs, then a more

¹⁴⁵This price range is based on Ethernet adapters listed in the DECdirect Hardware 1993 summer catalog and the April 1994 Global Computer Supplies mail-order catalog. Ethernet plugs come in 3 flavors: AUI (Thicknet), BNC/10Base2 (Thinnet), and RJ45/10BaseT (twisted pair wiring). If the customer card's Ethernet plug is different from the plug on the provider equipment, a transceiver, costing under \$100, is needed to convert between the two. Generally, more expensive Ethernet cards won't need the transceiver, and cheaper ones will, so the overall cost to the consumer is about the same in all cases.

¹⁴⁶Based on a telephone conversation with Tim Frahm of Zenith on April 13, 1994.

¹⁴⁷For example, the May 1994 issue of Byte Magazine contains an ad from Jameco in Belmont, CA, for a "486 bare bones system includ[ing 40 MHz] motherboard, computer case and power supply" for \$420 (quantity one); and an ad (p. 256) from Pacific Coast Micro, Inc., San Diego, CA, for a 40 MHz 486 motherboard for \$240, also quantity one.

sophisticated router may be needed to provide and route among more and/or different types of ports.

The model assumes that this need can be filled by adding a cable LAN interface to a low end router typically used to perform these types of functions and available for about \$1600 (Rickard, 1994).¹⁴⁸ It further assumes that all subscribers to the 4 Mbps service need this more sophisticated equipment, while all 500 Kbps subscribers can use the \$400 protocol converter. These assumptions reflect PSICable's actual service offering. Differentiating subscribers by the speed of their service instead of by their internal LAN configuration is not ideal, however. The model therefore includes an alternative scenario (referred to as "Individual 4 Mbps") in which the assumption that all 4 Mbps subscribers have complex LANs is relaxed, and 4 Mbps subscribers are thus able to use the \$400 protocol converter.

Table 5.5 summarizes the cost of subscriber equipment supplied by the Internet service provider in the case study. For comparison, the model also includes the price of the 10 Mbps equipment offered by LANCity but not part of the case study.

Table 5.5: Estimated cost of subscriber equipment

Equipment	Maximum LAN Bandwidth		
	0.5 Mbps	4 Mbps	10 Mbps
LAN Interface	\$500	\$900	
Protocol Converter	\$400	\$1,600	
TOTAL (quantity one)	\$900	\$2,500	\$7,100
TOTAL (vol. discount)	\$775	\$2,275	\$7,100

5.2.5. Connect Community LAN to Internet router

This component is represented by the node labeled "Cost to add a LAN (LC)" in Figure 5.1. Once a channel pair has been dedicated to LAN service, all that is required to make the LAN functional is to attach the IPOPOP router to the channel pair—in other words, to make the head end router a station on the LAN.

Connecting the head end to the LAN is just like connecting any other subscriber to the LAN, with one exception: simple protocol conversion is always sufficient, because the connection is not to a subscriber computer, but to the IPOPOP router. Thus no additional routing is needed, even for the 4 Mbps

¹⁴⁸Since cable LAN interfaces are non-standard and not generally offered by router vendors, PSI had Morning Star Technologies, Inc., one of their router suppliers, customize their Morning Star Express router in just this way. (Telephone conversation with Kate Murphy of Morning Star Technologies, Inc., April 18, 1994.)

service. Table 5.6 summarizes the cost to connect the head end’s Internet router to a cable LAN.¹⁴⁹

Table 5.6: Cost of equipment to connect IPOP router to cable LAN

Equipment	0.5 Mbps	4 Mbps	10 Mbps
LAN Interface	\$500	\$900	
Protocol Conversion	\$400	\$400	
TOTAL (quantity one)	\$900	\$1,300	\$7,100
TOTAL (vol. discount)	\$775	\$1,075	\$7,100

5.3. Subscriber demand model

With the FTTN architecture, the market demand for Internet service can be thought of as the product of two parameters: the number of subscribers in each neighborhood, and the number of neighborhoods in the community (the “Community Size” parameter at the top right of Figure 5.1). The number of neighborhoods is modeled for three values: one (minimal entry into a town, serving only one neighborhood), six (a small community), and twenty (a large town).¹⁵⁰ Deriving a value for the number of subscribers is more complex and is illustrated in Figure 5.6.

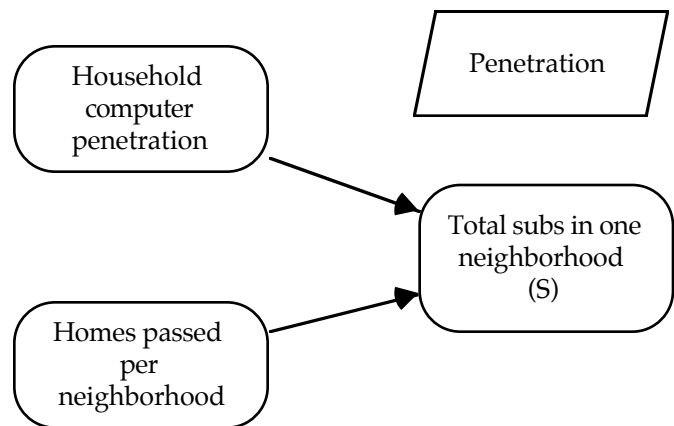


Figure 5.6: Model of subscriber demand

Market demand is represented by the variable “Total subscribers in one neighborhood (S).” This variable is the product of three parameters, working clockwise from the lower left of Figure 5.6:

¹⁴⁹As noted above, with special multiplexing equipment, each cable LAN could serve more than one neighborhood. This approach would reduce the number of cable LAN connections required.

¹⁵⁰The range of 6-20 is based on information supplied by Will Biedron of Continental Cablevision.

1. The maximum number of subscribers in a neighborhood—in other words, the number of homes passed by the neighborhood cable. This number is a function of the residential density: 500 is typical of suburban cable installations, while 1500 represents a dense urban scenario.
2. The fraction of households that have computers, estimated at about $1/3$.¹⁵¹ The maximum number of subscribers is multiplied by this factor to derive the *potential* subscriber base, according to the following assumptions:
 - a. Households that do not already have a computer are unlikely to become Internet subscribers; the availability of Internet service is not enough to persuade them to spend the \$1500 - \$3000 necessary to purchase a home computer.¹⁵²
 - b. Households that have a computer but are not already cable subscribers will be willing to pay the \$25 - \$50 installation fee to become cable subscribers.¹⁵³
3. The market penetration, defined as the percentage of the *potential* subscriber base that actually purchases the service. The results are modeled for a variety of penetration values between 0 and 100%, as described in Chapter 7.

The total number of subscribers in a neighborhood affects the number of LANs that need to be provided in that neighborhood. It also indirectly affects the cost of subscriber equipment, because of volume discounts.

¹⁵¹See (Pearl, 1994), in which the fraction is estimated at 0.33 based on data from Link Resources Corp. and the White House. In an April 20, 1994 electronic mail message, Jack Rickard, editor of Boardwatch Magazine, estimated the fraction at 0.27. Household survey results reported in (Ziegler, 1995) estimate the fraction at 0.31 in January of 1995, up from 0.27 in July of 1994.

¹⁵²Note that the subscriber equipment described in the previous section essentially puts a PC-compatible computer into the home, but in the case study, this computer is not configured or marketed for customer use; the customer is assumed to have a computer already.

¹⁵³These are the installation fees quoted by a Continental Cablevision sales representative for Cambridge, April 18, 1994. In the case study, Internet service cannot be purchased independently of CATV service; the minimum CATV subscription is \$10/month. According to Henry James, Continental Cablevision's Director of Communications, the Kagan media index reported that in February 1993, about 62% of U.S. homes passed subscribed to cable.

5.4. Traffic model

Figure 5.7 illustrates the traffic model used to estimate how many subscribers can reasonably share a single cable LAN. The value of this variable affects how many LANs need to be provided to each neighborhood and thus directly affects the cost to provide Internet service. The relation can also be viewed in reverse: if a provider enters the business by offering one LAN to a neighborhood, what is the maximum number of subscribers that can reasonably be served, and thus what is the cost per subscriber? Clearly, the more subscribers that can share a LAN, the lower each subscriber's cost.

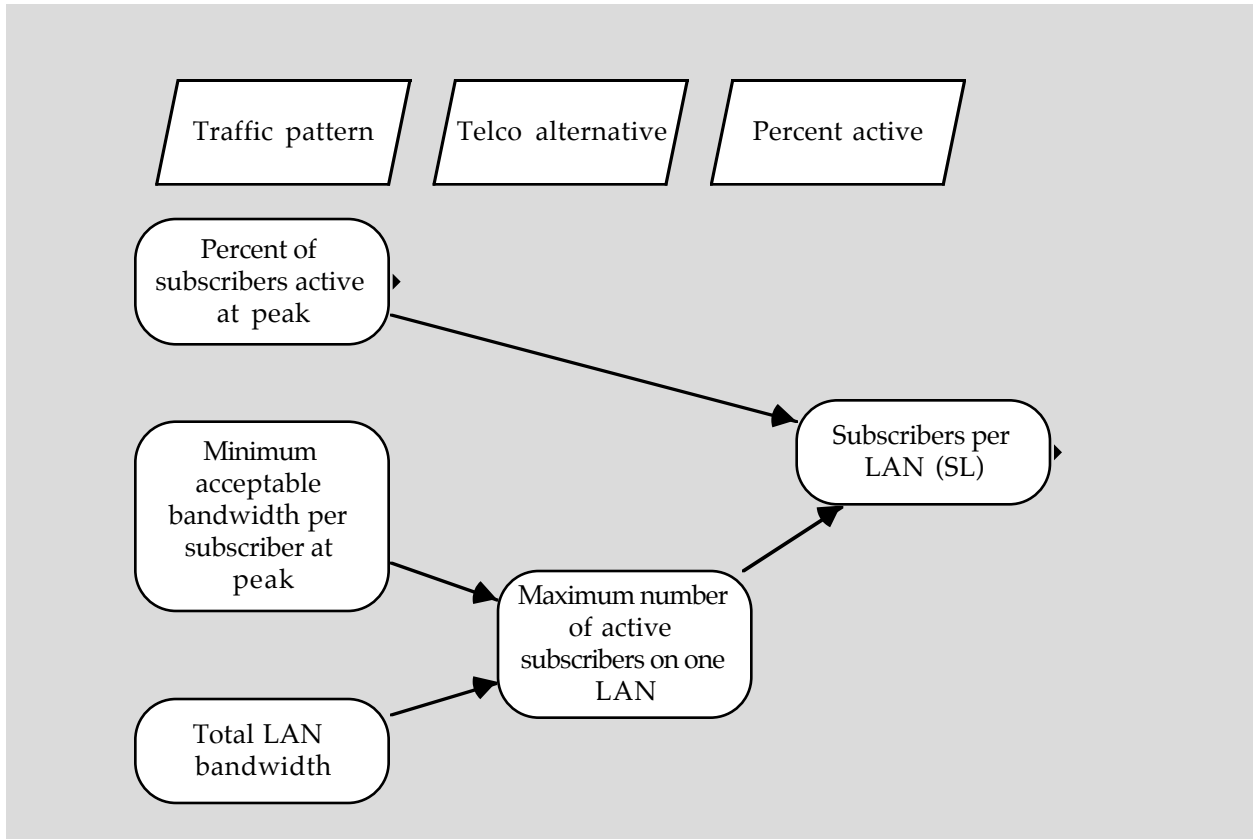


Figure 5.7: Traffic model

The number of subscribers that can share a LAN and still experience acceptable performance depends on the *traffic pattern*, comprising both the *mean level of traffic*, or *load*, generated by each subscriber as well as the distribution of that load in time. The traffic pattern to expect over a LAN connecting homes to the Internet is not well understood. A connection to the Internet enables not only a wide variety of applications, but also a mix that changes over time as new applications evolve. Each application offers different distributions and levels of load to the network. Since few homes are directly connected to the Internet today, it is difficult to predict what the usage mix will look like for this new service. Analogies to current systems—office

LANs and residential dial-in systems such as bulletin boards or workplace terminal servers—may not be appropriate, for the following reasons:

- Home usage may be more recreational than workplace usage, changing the application mix and associated pattern of load;
- Current dial-in systems are limited by the relatively low speed of modem connections. If users could get more bandwidth, would their usage patterns change? For example, would they use more graphical user interfaces (such as X Windows and Mosaic) and thus increase the average packet size?

The question of how many home subscribers can share a LAN requires further study, especially as residential Internet service is made available and used. Given the current absence of data, however, the cost model does not attempt to answer this question or to select an appropriate statistical model for the traffic pattern. Instead, it assumes that demand can be characterized by a single parameter: the minimum average bandwidth each subscriber will consider acceptable, based on the alternatives available to him or her. For example, if the network is so heavily shared that at times of peak usage, a file transfer is slower than it would be over a standard telephone modem, a user who does not need full-time connectivity might prefer a dial-up connection.

The model incorporates several alternative values for the average bandwidth, represented by the “Telco alternative” variable in Figure 5.7. The values are selected to approximate a variety of real options that could be available to customers over the local telephone network depending on their location and budget. These values are:

1. 30 Kbps. This value models the 28.8 Kbps V.34 dialup modems that have begun to appear in early 1995.¹⁵⁴ It also reflects the limit, to under 32 Kbps, on bit rates achievable through modems with the current telephone system architecture, as described in Chapter 4.
2. 60 Kbps, or “1B ISDN”. This value models both the 56 Kbps available through “Switched 56” service (a hybrid analog/digital service which is usually available to subscribers unable to get ISDN service), and the 64 Kbps delivered by a single ISDN “B” channel.¹⁵⁵

¹⁵⁴See (Mossberg, 1994).

¹⁵⁵An ISDN subscriber might limit usage to only one of the two “B” channels for several reasons. First, the Internet provider might not offer “2B” service (1994). Second, the subscriber may not wish to pay for the extra inverse multiplexing equipment needed to treat the two channels as a single pipe. Third, some ISDN tariffs are higher for “2B” service (Lindstrom, 1994).

3. 120 Kbps, or “2B ISDN”. This value models the full 128 Kbps of ISDN service, using the two B channels for circuit-switched data as described in Chapter 4. The 120 Kbps approximation also captures the 112 Kbps ISDN service (i.e. two 56 Kbps channels) which may be the only option available to some subscribers far from their local telephone Central Office.¹⁵⁶

Each such value represents the minimum average bandwidth each active subscriber ever expects to see—even during a period of peak utilization. Thus, it can also be thought of as the average active subscriber’s usage during a period of peak utilization (the box labeled “Minimum acceptable bandwidth per subscriber at peak” in Figure 5.7). By parceling out the total LAN bandwidth among active subscribers demanding this average level of bandwidth, the “Maximum number of active subscribers on one LAN” is derived. Since not all subscribers are expected to be actively using the LAN at any given moment, however, the total number of subscribers able to share each LAN is actually larger; it is derived by dividing the maximum number of active subscribers by the percent of all subscribers expected to be active during peak usage periods. The model tests values of 10, 20 and 30 percent of subscribers active at peak.¹⁵⁷

Average bandwidth during the peak period is not a perfect proxy for the actual traffic distribution, but it is a reasonable approximation given that the goal is to determine the peak level of network resources needed. Furthermore, this parameter tells us something about the types of traffic flows for which bandwidth really matters. Computer traffic can be broadly divided into two classes:

- *Interactive* traffic is generated by applications that transfer data at seemingly random intervals. Examples include a remote login, or a “chat” session with a user typing at a keyboard.

¹⁵⁶The model also includes alternatives of 15 Kbps (representing 14.4 Kbps dialup modems) and 1.5 Mbps T1 service. These alternatives are not discussed further in this thesis: in early 1995, 14.4 Kbps modems are being rapidly supplanted by 28.8 Kbps models, while T1 service is not a realistic alternative for most residential subscribers. Although a few ADSL deployments have been announced, these provide T1 speeds only in the direction to the home.

¹⁵⁷The range for this “percent active” parameter is derived from a figure of 5% measured for the Laboratory for Computer Science (LCS) dial-in server (data supplied by William Ang of LCS Computing Resource Services), and the 25% estimated by the planners of Time Warner’s Orlando Full Service Network (see (Time Warner Cable Inc., 1993), pp. 2-3). The LCS figure of 5% is an average, and includes periods in which all dial-in lines are in use. Therefore, the actual demand is probably higher than what can be measured. Time Warner’s estimate is based on expectations for video-on-demand, a service which is (presumably) only in use when a customer is actively watching a movie. In contrast, a subscriber’s computer may use Internet service (for example, to receive email) without the subscriber’s active involvement. Because of these factors, I use a range of 10-30% instead of 5-25%.

- *Streaming* traffic is generated by file transfers. These may be initiated directly by users—for example, a user downloads a binary file from a software archive. They may also be caused indirectly, for example by a user downloading email from a server, or clicking on a graphical image in a Web browser, causing that image to be downloaded.

Unlike interactive traffic, streaming traffic will use as much bandwidth as is available until the data transfer is complete. If N users are sharing a cable LAN and all of them are doing file transfers, and if the LAN sharing algorithm is working effectively and fairly, each user's transfer should receive approximately $(1/N)$ th of the total bandwidth. As N rises, the file transfers slow down, and each transfer takes more time. Assuming the pattern of interactive traffic and the intervals between file transfers remain unchanged, streaming traffic will begin to dominate the overall traffic pattern. Thus, using the average bandwidth allocation to determine how many users can share the network is equivalent to considering a worst-case scenario in which all users are simultaneously transferring files.

Table 5.7: Subscribers per cable LAN

Minimum average bandwidth per subscriber at peak	Per 500 Kbps LAN		Per 4 Mbps LAN	
	Maximum number of active subscribers	Total subscribers (assuming 20% active)	Maximum number of active subscribers	Total subscribers (assuming 20% active)
30 Kbps	16	80	133	665
60 Kbps	8	40	66	330
120 Kbps	4	20	33	165

Table 5.7 summarizes the traffic model at the mid-point of the percent active parameter range (i.e. 20% active). The number of subscribers able to share each cable LAN rises with the total LAN bandwidth, but falls as the percent active parameter value and the per-user average bandwidth requirement rises. Sharing a LAN among more users clearly has economic advantages. For example:

- The provider can spread fixed capital costs across a larger number of subscribers, lowering the average per-subscriber cost to introduce the service and thus increasing the service's appeal to potential subscribers (assuming this lower cost is reflected in the price of the service);
- The risk of entry is lower since the provider has a larger base of potential customers from which to recoup fixed costs.

Chapter 7 explores these points in more detail, examining the effects of varying the average bandwidth and percent active parameters on the feasibility of offering either the 500 Kbps or the 4 Mbps cable LAN services.

5.5. Summary

This chapter has described the details of an engineering economic model of residential Internet connections provided over a cable television network. It presented the costs of each element in a technology reference model, consisting of an Internet Point of Presence (IPOP), Local Area Networks (LANs) built out of pairs of up- and downstream cable channels, and subscriber interface equipment. It then described the models of subscriber demand and network load that were employed to complement the technology model.

The following chapter describes a similar model for Internet connections offered over ISDN telephone circuits. The results of both models are presented in Chapter 7.

Chapter Six

Cost model of Internet over ISDN

This chapter describes in detail an engineering economic model of residential Internet connections provided over ISDN telephone circuits. This model is based on a case study of Internex, Inc. a commercial provider of Internet access over ISDN in the San Francisco Bay (“Silicon Valley”) area.

Like the Internet over cable model developed in the previous chapter, the Internet over ISDN model considers the marginal capital cost of providing Internet service over ISDN lines. It assumes that the local telephone company has already invested in ISDN-ready Central Office switching equipment, and therefore does not allocate this cost to Internet service. It does, however, assume that the subscriber will purchase an ISDN line and equipment in order to connect to the Internet, so all such per-subscriber capital costs are included.

The model described in this chapter is based on Internex’s offering but does not match it in every detail. In particular, in order to provide a valid comparison with the PSICable service, telephone line charges applicable in Cambridge, MA are used instead of California rates.

6.1. Overview

Figure 6.1 shows a top-level view of the Internet over ISDN model. The box on the left models subscriber demand. The four boxes down the middle of the figure represent the technology reference model, including the Internex customer’s choice of a 1B or 2B ISDN service offering (“Telco service”). The following sections describe the demand and technology sub-models in more detail. Results, represented by the box on the right, are discussed in Chapter 7.

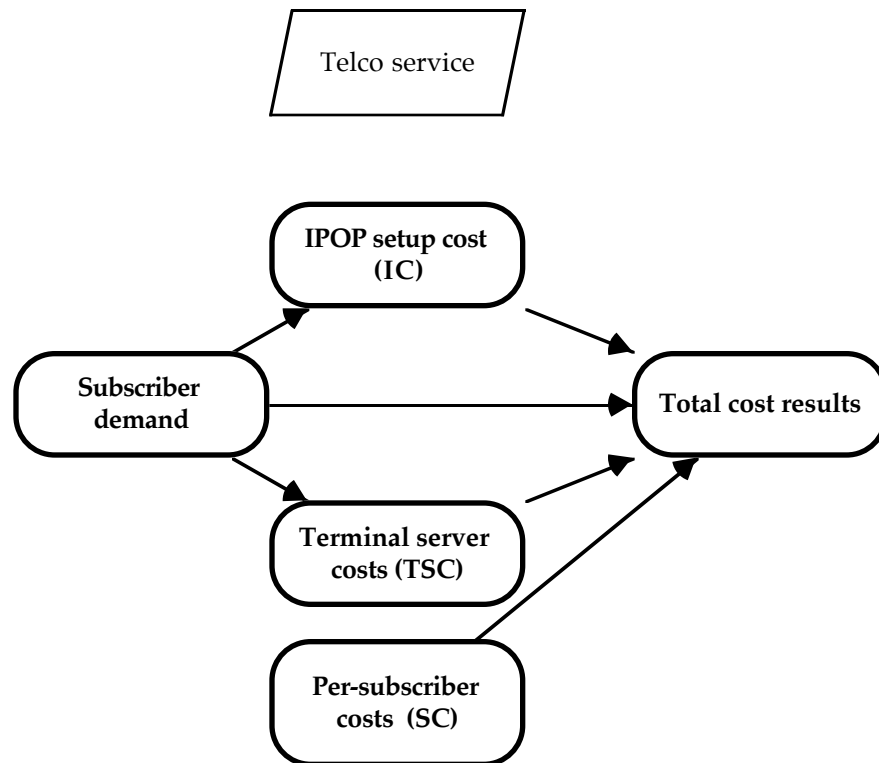


Figure 6.1: Capital cost model of Internet over ISDN

6.2. Technology reference model

Chapter 5 outlined the components of the technology needed to connect residential computers to the Internet:

- An Internet Point of Presence, or IPOP, consisting of an Internet Protocol (IP) router physically connected to, and able to route traffic appropriately to and from, the rest of the Internet;
- Some form of community network that physically connects the Internet Point of Presence to individual subscribers' computers;
- Hardware and software allowing those subscriber computers to connect to the community network and use Internet services.

Figure 6.2 shows these components in more detail, based on the technology described in Chapter 4. The Internet Point of Presence component, consisting of an IP router and circuit to the rest of the Internet, is identical in the ISDN and cable models; see the detailed description of this component in Chapter 5. The community network and subscriber equipment components, however, are quite different for the ISDN approach. Instead of a cable Local Area Network shared by a group of subscribers, ISDN connections consist of point-to-point circuits terminating in multiplexers located at the Internet Point of Presence. These multiplexers, shared by groups of subscribers, are represented

by the “T1 Terminal Server” boxes shown in Figure 6.2. The point-to-point connections physically consist of ISDN circuits from each residential subscriber to his or her nearest telephone company Central Office, and T1 (PRI) circuits from the Internet service provider to its closest telephone company Central Office.¹⁵⁸ The ISDN and T1 circuits are interconnected into point-to-point connections by the public switched telephone network.

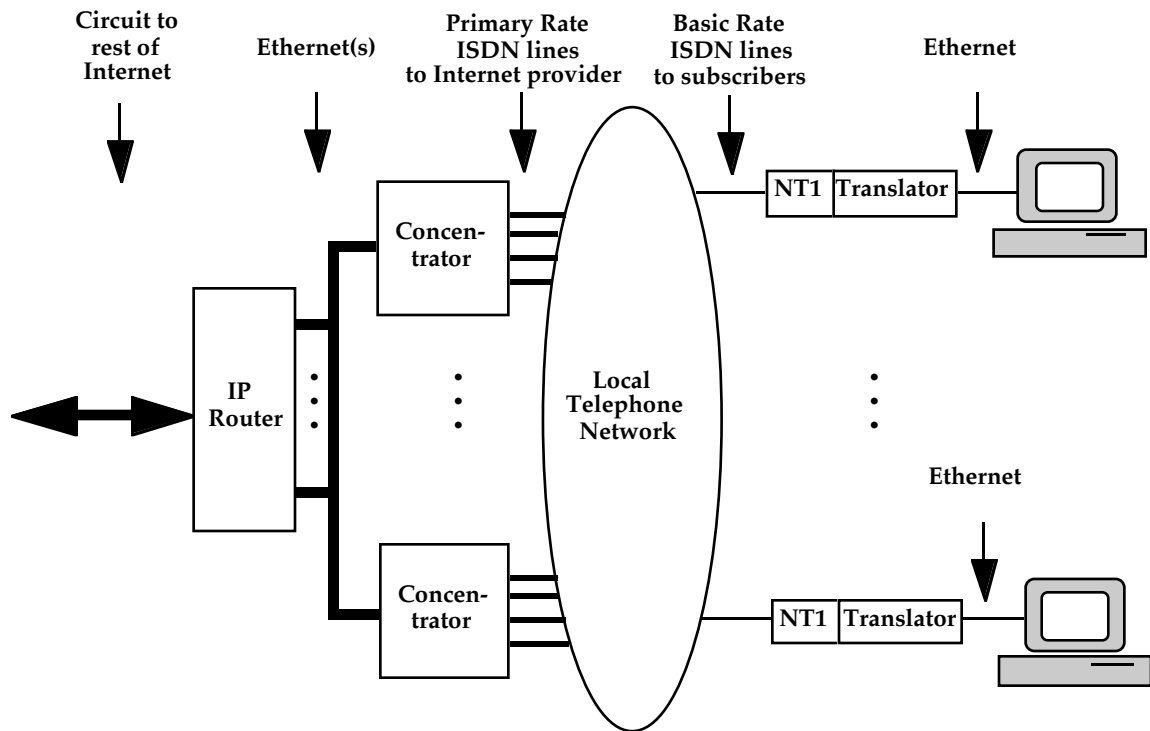


Figure 6.2: Internet over ISDN configuration

6.2.1. Terminal servers

Figure 6.3 shows the model of terminal server equipment cost. The total cost of terminal server(s) consists of the per-port cost multiplied by the number of input ports needed.

¹⁵⁸The use of T1 circuits to connect the Internet Point of Presence to the public telephone network reflects the implementation adopted in the case study. It is not, however, the only possibility. As (Network Express Inc., 1994) describes, in some regions, tariffs favor the use of multiple ISDN (i.e. Basic Rate Interface, or BRI) lines and associated multiplexing equipment, instead of PRIs.




Figure 6.3: Model of terminal server equipment cost

The estimated cost of this equipment is based on the terminal servers used by Internex, Inc.¹⁵⁹ These devices multiplex up to four T1 (i.e. ISDN Primary Rate Interface, or PRI) lines, or input ports, onto a single Ethernet output port, including the CSU/DSU functionality needed for each T1 line.

The number of terminal servers needed is thus a function of the number of T1 lines (input ports) needed to support the user population. Since each PRI supports 23 “B” channels, it can support up to 23 users of a single ISDN B channel, or 11 users of dual B channel service. Thus a terminal server with 4 input ports can support up to 92 “1B ISDN” or 44 “2B ISDN” users.

The cost of a terminal server input port has two components: equipment cost and line cost. The equipment cost is the fraction of the terminal server cost allocated to each port. Although in practice, equipment pricing is not linear in the number of ports (models with larger numbers of ports have a lower price per port), it is close enough to linear to be able to approximate the equipment cost per port by a constant \$4000.¹⁶⁰

The line cost is the 3-year capitalized cost to lease a T1 line from the Internet service provider’s premises to the closest telephone company Central Office. The model estimates this cost at \$19,400, based on approximate fees of \$1,400 for installation and \$500 per month.¹⁶¹ These fees, and the resulting capitalized cost, are slightly lower than those used to estimate the cost of leasing a T1 circuit to connect the IPOP to the rest of the Internet. The difference arises because T1 circuit fees increase with distance, and the

¹⁵⁹Telephone conversation with Robert Berger, president of Internex, Inc., April 29, 1994.; Ascend Communications Inc. Pipeline access router product literature.

¹⁶⁰ This figure is based on Robert Berger’s estimate of \$15-25,000 for a 4-port “access router” from Ascend Communications, Inc., and October 14, 1994 price quotes from David Sugarman of DSC International, a reseller of Ascend’s Pipeline access router products, of \$4-5000 for a model supporting one PRI line, and \$9-10,000 for a model supporting four PRI lines.

¹⁶¹These fees are based on quotes from Peter Caparso, NYNEX sales representative, April 22, 1994. (Network Express Inc., 1994) also lists representative fees for all seven Regional Bell Operating Companies, showing how widely they can vary (e.g. by a factor of 2).

connection to the rest of the Internet is assumed to be longer than the connection to the closest Central Office. In addition, the Internet circuit is assumed to go between telephone company Central Offices, which may result in an additional charge.

6.2.2. Subscriber equipment

Like the cable model, the Internet over ISDN model assumes that the subscriber's responsibility includes a home computer with an Ethernet interface and Internet software; therefore, these costs are not included. As shown in Figure 6.4, the model does include the cost of equipment needed to convert transmissions between Ethernet and ISDN protocols, as well as telephone Central Office equipment needed to activate a new ISDN subscriber. This boundary was selected to be comparable to the cable model; it does not reflect the Internex case study, in which all subscriber equipment purchases are paid for by the subscriber.

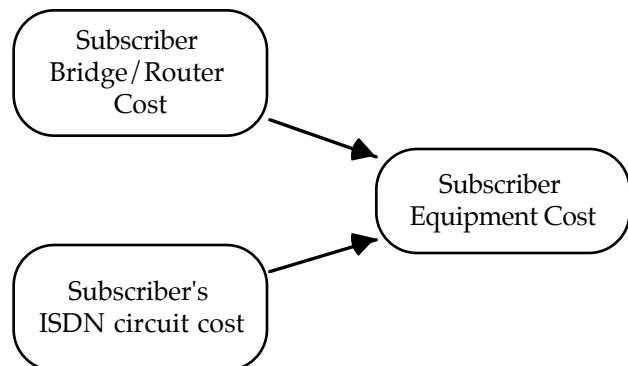


Figure 6.4: Model of subscriber equipment cost

Subscriber interface equipment

A wide array of equipment is available to connect computers to ISDN circuits. The choices include both standalone boxes and plug-in cards, offering a wide variety of prices and feature sets.¹⁶² Prices range from roughly \$500 to \$3000, with variation accounted for by feature differences such as:

- The type of computer interface supported. For example, serial port interfaces are cheapest but may limit the connection speed to less than the ISDN circuit's speed, depending on the subscriber's computer model. Ethernet interfaces are faster (and therefore preferred); 10BaseT Ethernet interfaces cost less than 10Base2, for example.

¹⁶²For a description of vendor offerings at the end of 1994, see (Derfler, 1994) and (Robertson, 1994). See also Newsbytes story, November 16, 1994, "Motorola ISDN Terminal Adapter 'Modem' Under \$500."

- The number of ISDN channels supported; for example, the \$500 plug-in cards only support the use of 1 ISDN “B” channel for data. More sophisticated devices use the second “B” channel only when it is needed, including associated autodial and hangup features. In addition, devices that include the NT-1 card needed to terminate the ISDN circuit generally reflect its cost (\$175-200) in their price.
- The complexity of the LAN configuration. For example, the most expensive devices support ten IP addresses, while lower-end equipment supports only one.
- Support for multiple higher-level protocols; in particular, support for the non-proprietary Point-to-Point Protocol (PPP) and/or Multichannel Point-to-Point Protocol (MP), allowing equipment from multiple vendors to interoperate.¹⁶³

For the same reasons discussed in Chapter 5 for the cable model, the ISDN model assumes that a standalone box connecting to the Ethernet port on the subscriber’s computer is the preferred solution. Such a box essentially serves as a two-port bridge or router, converting between the computer’s Ethernet protocol and the telephone line’s ISDN protocol. With appropriate PPP or MP software supported by this box, TCP/IP can run transparently over the ISDN line, allowing the line to serve as an interoperable “LAN extension cord”—with the LAN supplied by the Internet service provider, and connected to the rest of the Internet through the IPOP router.

In late 1994, the lowest-priced device with this minimum feature set listed for \$1190 (this price includes an NT-1 and Terminal Adapter).¹⁶⁴ Given the rapidly falling prices and discounts off list price commonly available for such computer equipment, the model uses a cost of \$1100 for this component of subscriber interface equipment. Since these devices are sold directly to consumers, any economies that may be derived from manufacturing large volumes of these devices should already be reflected in their market prices; accordingly, the model does not incorporate any additional volume discounts on the \$1100 price.

Subscriber’s ISDN circuit

Since a subscriber is assumed to purchase ISDN service in order to connect to the Internet, the model includes the capital cost to equip a subscriber’s ISDN line at the telephone company Central Office (which is assumed to already

¹⁶³Electronic mail from Robert Berger of Internex, Inc. See (1994) and (Simpson, 1993).

¹⁶⁴This price is based on the Combinet 160 model. The Combinet 150 lists for \$990 but does not include an NT-1.

have an ISDN-capable switch).¹⁶⁵ The model estimates this cost at \$400, based on the following components:

- An ISDN line card, matching the modulation and coding functions performed by the subscriber's NT-1, and costing approximately \$140;
- A slot on the switch for the ISDN line card to slide into, costing approximately \$180, plus labor to configure the software (approximately \$30);¹⁶⁶
- Any labor needed to prepare and install the physical circuit. The model assumes that on average, this cost is covered by the ISDN installation fee, which in Massachusetts is the same for ISDN as POTS, about \$60.

The model does not include monthly or usage fees charged by the telephone company for ISDN service, which would have to be paid by the subscriber. This exclusion ensures that the ISDN model remains comparable with the cable model, which covers the capital cost to provide Internet service, independent of how the providers involved choose to bill for it, and independent of operational costs that might be covered by monthly fees.

6.3. Subscriber demand model

In the cable model described in Chapter 5, subscriber demand is discussed in terms of the number of subscribers that can share each neighborhood cable Local Area Network (LAN). Demand is modeled as a percentage of the maximum number of potential subscribers in each neighborhood.

The telephone model does not include a concept equivalent to the neighborhood in the cable model. Instead, the individual subscriber provides the smallest unit of analysis. A residential Internet provider connected via leased T1 to its nearest Central Office switch can serve customers from anywhere else on the telephone network, not just those in its neighborhood. While a penetration level of 10% represents 10% of each neighborhood in the cable model, the ISDN model admits any geographic distribution of the same number of subscribers.

So that the two models would produce results that could be compared meaningfully, the amount of equipment needed to provide residential Internet service over ISDN at each penetration level is determined by the actual **number** of subscribers represented by the subscriber demand

¹⁶⁵The cable model does not have an item comparable to the ISDN line card; no capital equipment is required at the cable head end to add a cable subscriber.

¹⁶⁶The line card, slot, and configuration cost estimates were supplied by Walter Mansell of MIT Telecommunications Services, November 21, 1994. The slot cost is derived by dividing the \$90,000 cost of the complete shelf (called an "Integrated Services Line Unit") by the 512 slots it can support.

assumptions in the cable model. The maximum possible number of subscribers across all neighborhoods is assumed to be the same in each model. This is represented by the “Number of potential homes passed (cable equivalent)” shown in Figure 6.5.



Figure 6.5: Subscriber demand model

For example, in a large community (20 neighborhoods) with an urban population density (1500 potential subscribers per neighborhood), the number of potential homes served is 30,000. The “Total number of subscribers” is then derived from this number by taking only the fraction of households with personal computers (assumed to be one-third, as explained in Chapter 5) and multiplying the result by the penetration parameter. The number of active subscribers is computed based on assumptions of 10, 20, and 30% of subscribers actively connecting to the Internet at any given time, as justified in Chapter 5.

In summary, a penetration level of 10% may represent a different geographic distribution in the two models, but it always represents the same number of total and active subscribers. To provide a feel for the potential size of the customer base, Table 6.1 shows the total number of subscribers (100% penetration) for different types of communities.

Table 6.1: Number of subscribers

Residential density	Community size		
	One neighborhood	Six neighborhoods	Twenty neighborhoods
suburban	167	999	3330
urban	500	2997	9990

6.4. Summary

This chapter has described the details of an engineering economic model of residential Internet connections provided over ISDN circuits. It has presented the costs of the elements of the technology reference model that are different from the Internet over cable model. These elements include ISDN subscriber interface equipment as well as multiplexers used to concentrate

multiple ISDN circuits onto router Ethernet ports. Once concentrated, the ISDN traffic can be handled by the same Internet Point of Presence configuration as in the cable model.

The following chapter describes the results of both the Internet over cable and Internet over ISDN models.

Chapter Seven

Results

This chapter presents the analysis of the cable and ISDN Internet access cost models described in Chapters 5 and 6. First, it defines the model's cost result variables and briefly reviews the model's input parameters. Then a series of graphs are presented and discussed, showing the average capital cost, number of cable LANs needed per neighborhood, and components of average capital cost. The first set of average cost graphs is based on an original set of input parameter values, while subsequent graphs present the results of sensitivity analyses that explore the effect of changes in each of the different input parameter values. The chapter concludes with a discussion of absolute entry costs under two different entry strategies.

7.1. Result variables

Figure 7.1 shows the result variables for both models. As shown on the left side of the figure, cost results are separated into three components—fixed, sharing and subscriber cost—that are added together to get total cost. This result represents the complete cost to deploy an Internet over cable or ISDN system to a given number of subscribers, including the cost of each subscriber's equipment as delimited by the model.

Dividing total costs by the number of subscribers at any given penetration level produces the per-subscriber cost values shown on the right side of Figure 7.1.

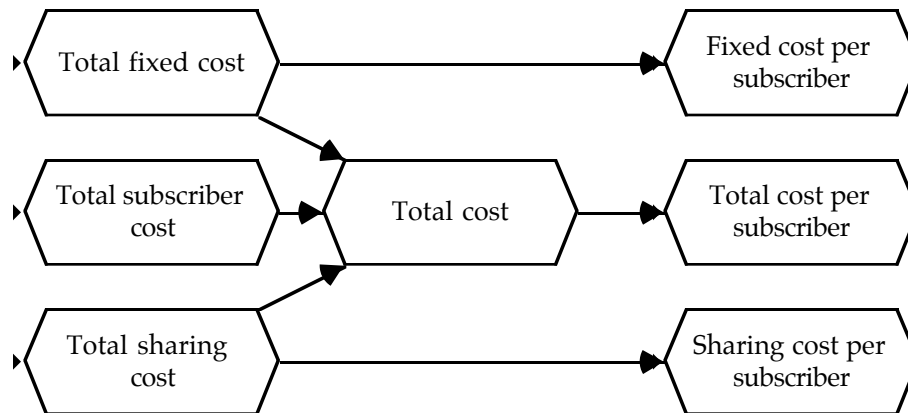


Figure 7.1: Model of cable and ISDN results

Total fixed cost is the fixed cost of preparing the system for the first subscriber. This cost is spread across all subscribers using the system. In the cable model, it consists of the cost to put in an Internet Point of Presence (IPOP), plus an upstream transmission upgrade for each neighborhood. In the ISDN model, it consists of only the IPOP cost.

Total subscriber cost is the variable cost of supplying each customer with his or her own equipment, or short-run marginal cost. In both models, it is calculated by multiplying the cost of each individual subscriber's equipment by the number of subscribers. For cable, each individual's subscriber equipment consists of a protocol conversion device (bridge or router) translating between the subscriber computer's Ethernet and the cable Local Area Network (LAN) protocols, plus a cable modem for physical transmission over the cable. For ISDN, it consists of a protocol conversion device translating between Ethernet and ISDN protocols, including a built-in NT-1 to handle the physical transmission over the telephone line. Additionally, ISDN subscriber equipment includes the capital cost to equip the subscriber's ISDN circuit at the telephone company Central Office.

Total sharing cost is a longer-run marginal cost. It represents the cost of equipment shared by subsets (i.e. more than one but fewer than all) of the total number of subscribers. For cable, it is computed as the cost of equipping channels to run LANs and connecting those LANs to the head end router, for all neighborhoods.¹⁶⁷ For ISDN, it is the cost of the terminal servers and T1 lines that connect the Internet provider to the local telephone network.

¹⁶⁷All cable cost results are based on the assumption that at least one LAN is required for each neighborhood. Cable's sharing cost might be reduced if multiple neighborhood's spectrum were "glued together" into a single LAN, as noted in Chapter 5. As we shall see below, cable's sharing cost is small enough that such a change is unlikely to have a significant effect on average cost.

7.2. Input parameters

Table 7.1 summarizes the input parameter values used in the initial analysis below, and the alternative values explored in subsequent sensitivity analyses.

Table 7.1: Summary of input parameters

Parameter	Original Value(s)	Sensitivity Analysis
Peak bandwidth	500 Kbps (cable) 4 Mbps (cable) 128 Kbps (ISDN)	Not applicable
Residential density	1500 homes per cell	500 homes per cell
Community size	6 cells	1 cell, 20 cells
Percent active	20%	10%, 30%
Average bandwidth	120 Kbps	30 Kbps, 60 Kbps
Cable LAN efficiency	100%	60%, 80%

The capital cost of Internet access over cable or ISDN depends on the following factors:

- The peak bandwidth of the service offered: 500 Kbps or 4 Mbps for cable, vs. 128 Kbps for ISDN. Quantitative results are shown below for all 3 cases.
- The residential density of the community, ranging from 500 to 1500 homes per neighborhood cell as defined in Chapter 5. The analysis begins with the higher density value, since that is more typical of the town in the case study (Cambridge, Massachusetts). Sensitivity analysis then explores the effect of reducing this parameter's value to 500 homes per cell.
- The total size of the community, represented by the number of neighborhood cells supported. The model tests values of one, six, and twenty cells for this parameter. Initial quantitative results are presented for the intermediate value (six cells), with sensitivity analysis performed on the higher and lower values.
- The percentage of Internet service subscribers expected to be active simultaneously during peak usage periods. The model tests values of 10, 20 and 30% for this parameter. Initial quantitative results use the intermediate value (20%); results for lower and higher values are presented in the sensitivity analysis.
- The average bandwidth expected by each active subscriber. Since the objective of the thesis is to compare cable and ISDN networks, and ISDN bandwidth is fixed, results are presented for the conservative comparison

case in which the average bandwidth is the same as the maximum ISDN bandwidth. The initial value used is 120 Kbps, representing an average of the 112 or 128 Kbps that may be available to ISDN subscribers depending on their exact location. Sensitivity analysis tests smaller average bandwidths of 60 and 30 Kbps.

- How efficiently cable Local Area Networks (LANs) allow subscribers to share their bandwidth. Since little is known about the sharing protocol used in the case study, the model begins with the assumption of perfect efficiency. Sensitivity analysis explores what happens if this assumption is substantially violated.
- The number of subscribers, represented by the *penetration level*: This parameter forms the horizontal axis of all the result graphs presented below. Its value can range from 0 to 100%, by increments of 5%. As explained in Chapter 5, it represents the percentage of computer-owning households in a service area that choose to subscribe to Internet service. For example, given the model's assumption that one-third of homes have computers, a service area consisting of one 1500-home neighborhood would have 500 subscribers at 100% penetration. Alternatively, a service area consisting of six 500-home neighborhoods would have 1000 subscribers at 100% penetration.

The analysis begins with an examination of the cost per subscriber resulting from the original parameter settings.

7.3. Average capital cost

Figure 7.2 compares the cost per subscriber to provide the 500 Kbps and 4 Mbps cable services and the 128 Kbps ISDN service, as a function of the penetration level.¹⁶⁸ The model parameters are fixed at the original values shown in Table 7.1. Given the assumption of 1500 homes per cell, six cells (i.e. 9000 homes), and one-third of homes with PCs (i.e. 100% penetration = 3000 subscribers), each 5% increment in penetration represents an additional 150 subscribers.

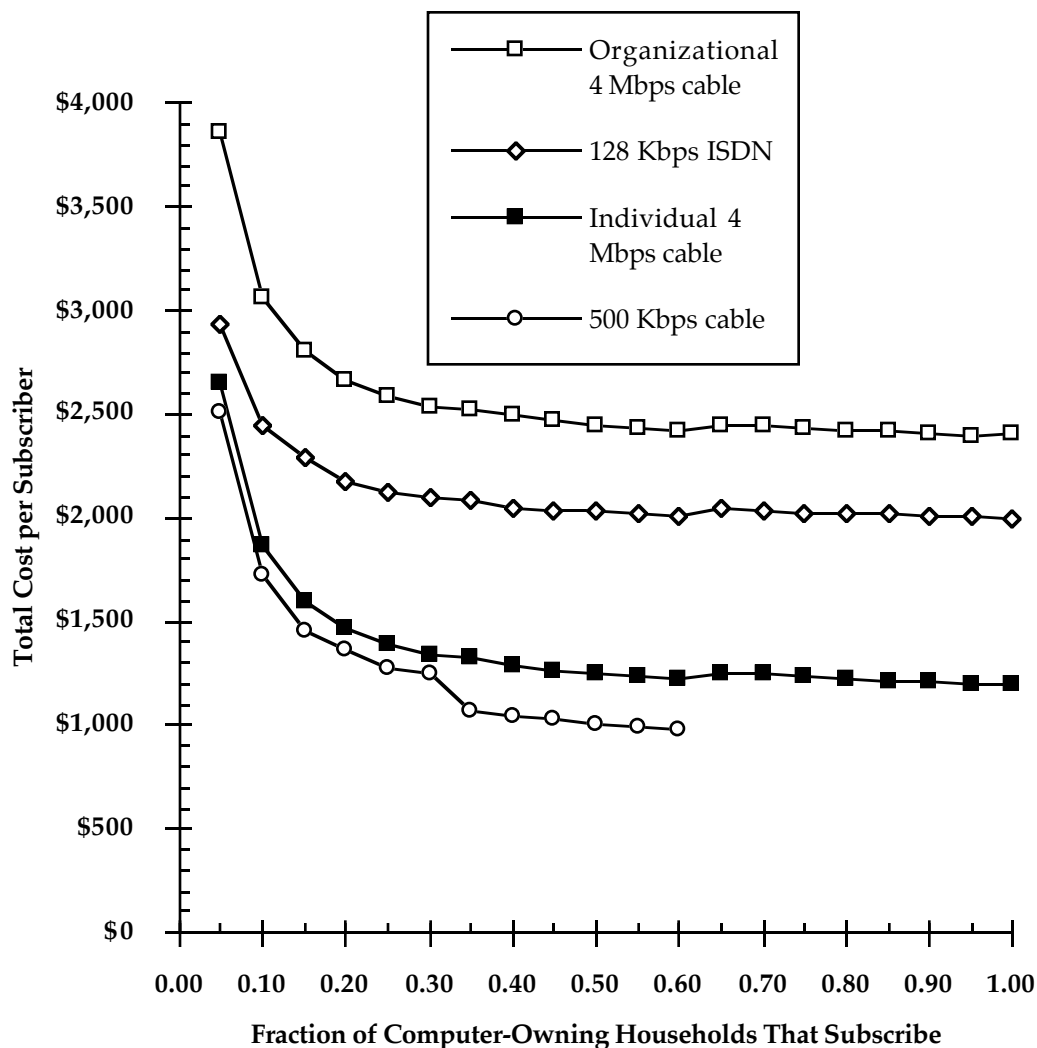


Figure 7.2: Average cost with original parameter settings

The greater economy of the 500 Kbps cable service relative to ISDN is immediately evident in Figure 7.2. Beyond initial penetration levels, the 500

¹⁶⁸ Each penetration level represents the same number of subscribers across all 3 services. Note that the vertical axis represents average cost for any given penetration level; the curves do not represent trajectories, i.e. the marginal cost to move from one penetration level to another.

Kbps cable service can provide the same average bandwidth and four times the peak bandwidth of the ISDN service for less than half the cost per subscriber. This result holds true up to the maximum penetration level achievable with the 500 Kbps cable service—60% under the initial parameter assumptions. (This penetration limit is discussed further in the next section.)

Under these same assumptions, the 4 Mbps cable service is not subject to penetration limits. The average cost of the organizational 4 Mbps service of the case study is more expensive than the ISDN service by about \$400, or 20% more per subscriber to provide the same average bandwidth. However, the alternative, individual 4 Mbps service considered in this report (discussed below as part of the sensitivity analysis) is \$800, or 40%, less expensive than ISDN. For thirty-two times the peak bandwidth, however, even 20% higher cost is a terrific bargain. This result is shown in Figure 7.3, which plots the per-subscriber cost of each service normalized by the number of peak Kbps of bandwidth that the service provides. (In Figure 7.3 and all following graphs, “4 Mbps cable” refers to the Organizational 4 Mbps cable service.)




Figure 7.3: Average cost per Kbps of peak bandwidth

Figures 7.2 and 7.3 dramatically demonstrate the central hypothesis of this thesis: **the shared cable architecture can provide the same average bandwidth and higher peak bandwidths much more economically than the dedicated ISDN architecture.** With the number of subscribers sharing each cable LAN chosen to provide each active subscriber with an average bandwidth of no less than ISDN's dedicated bandwidth, 500 Kbps Internet access over cable can provide four times the peak bandwidth of ISDN access for less than half the cost per subscriber. In addition, 4 Mbps cable access can provide thirty-two times the peak bandwidth of ISDN for 20% more cost per subscriber when targeted at organizations, and 40% less cost when targeted at individuals. The economy of the shared bandwidth approach is most evident in Figure 7.3, which compares the per-subscriber cost per bit of peak bandwidth. These costs level out as penetration increases, to about \$0.60 for the 4 Mbps and \$2 for the 500 Kbps cable services, versus close to \$16 for ISDN. (The comparable figure for the Individual 4 Mbps service is \$0.30.)

7.4. Penetration limitations

This section discusses in more detail the penetration limit mentioned in the previous section. Throughout this chapter, cost results are presented only for the penetration ranges that can actually be achieved.

With the point-to-point ISDN approach described in Chapter 6, increasing the number of subscribers served means providing ISDN telephone lines to more customers. As (Waring, 1991, p. 1979) observes, the telephone plant in residential areas may have as few as 1.3 lines per residence. Therefore, if enough households subscribe to ISDN in addition to their existing analog telephone service, strong growth in demand for ISDN could exhaust the existing installed plant. There is no technical reason, however, why new lines could not be installed; in fact, this would be a likely outcome of strong growth in demand.¹⁶⁹ Therefore, no limits to ISDN penetration appear in the graphs below.

In contrast, the spectrum map of existing cable systems creates a technological limit. As discussed in Chapter 3, the most common residential cable spectrum map provides only four upstream channels; remapping to gain more upstream bandwidth would require replacement of all amplifiers. Since each cable LAN requires symmetric up- and downstream bandwidth, the limited upstream bandwidth constrains the number of LANs that can be provided. As a consequence, in some scenarios 100% penetration may not be achievable.

¹⁶⁹If the telephone company were to have to string new wires, a large capital cost would have to be added to the ISDN service. This cost is not reflected in the ISDN model, which assumes existing circuits can be utilized. With only one-third of households (i.e. computer owners) considered as potential subscribers, this assumption is reasonable. It might come into question, however, if the fraction of households with computers were to grow substantially.

Recall from Chapter 5 that the cable model assumes that each neighborhood cell has its own separate spectrum. Figure 7.4 illustrates the number of cable LANs needed in each cell assuming the original configuration of input parameter values (as summarized in Table 7.1). It shows that for the 500 Kbps service, once penetration in each cell reaches 60% of households with computers, further subscribers cannot be supported without degradation in service quality.¹⁷⁰ In other words, if penetration were extended beyond 60%, the number of subscribers sharing each LAN would become too large to provide each active subscriber with an average bandwidth of 120 Kbps. Because of this result, cost results in following sections are not plotted beyond 60% penetration, for the 500 Kbps service under the original parameter assumptions.




Figure 7.4: Number of cable LANs required per neighborhood

Figure 7.4 also shows that the 4 Mbps service is not subject to penetration limits under the original parameter assumptions.¹⁷¹ This result reflects the greater spectral efficiency of the 4 Mbps service. Each 6 MHz channel can provide either one 4 Mbps LAN, or four 500 Kbps LANs (for a total of 2

¹⁷⁰Recall from Chapter 5 that each cable channel-pair can support either four 500 Kbps LANs or one 4 Mbps LAN. Since there are four upstream channels, there is a maximum of four cable LAN channel pairs. Therefore, penetration limits kick in when the number of 500 Kbps LANs needed exceeds sixteen, or the number of 4 Mbps LANs needed exceeds four.

¹⁷¹It is also evident from Figure 7.2 that the first 4 Mbps LAN enables a higher potential penetration level (30%) than the first 500 Kbps LAN (approximately 2%). Thus the provider attempting a minimal entry strategy of “one of everything” can reach a larger subscriber base with the 4 Mbps service, providing a larger base over which to recover fixed costs.

Mbps).¹⁷² As a result, the 4 Mbps equipment can serve more aggregate bandwidth, and therefore higher penetration, with a single channel.¹⁷³

Sensitivity analyses below explore the effect of changes in the input parameter assumptions on cable penetration limits, for both the 500 Kbps and 4 Mbps services.

7.5. Components of capital cost

This section explores the drivers behind the shapes of the different cost curves shown in Figure 7.2. It examines in greater detail the components that make up total cost, continuing with the original parameter assumptions shown in Table 7.1.

Figures 7.5, 7.6, and 7.7 show the contribution of each cost component to the average cost of the 500 Kbps and 4 Mbps cable and 128 Kbps ISDN services, respectively. Two key observations apply to all three graphs:

- Subscriber cost, as defined at the beginning of this chapter, is the most significant component of average cost.
- Average cost quickly approaches its minimum value as penetration increases. For all three services, average cost at 15% penetration (450 subscribers) is less than 150% of average cost at the full penetration level that can be supported.

¹⁷²Each 500 Kbps LAN uses 1 MHz, with the remaining 2 MHz used for guard bands.

¹⁷³Technological advances in digital transmission may allow more bits per second to be encoded onto a single 6 MHz channel. Such an advance would ameliorate the penetration limitations described here, assuming the average bandwidth demanded stays constant. Alternatively, if household computer penetration were to rise, the number of potential subscriber households would increase. In that scenario, a technological advance would be necessary to avoid worsening the penetration limits shown here.

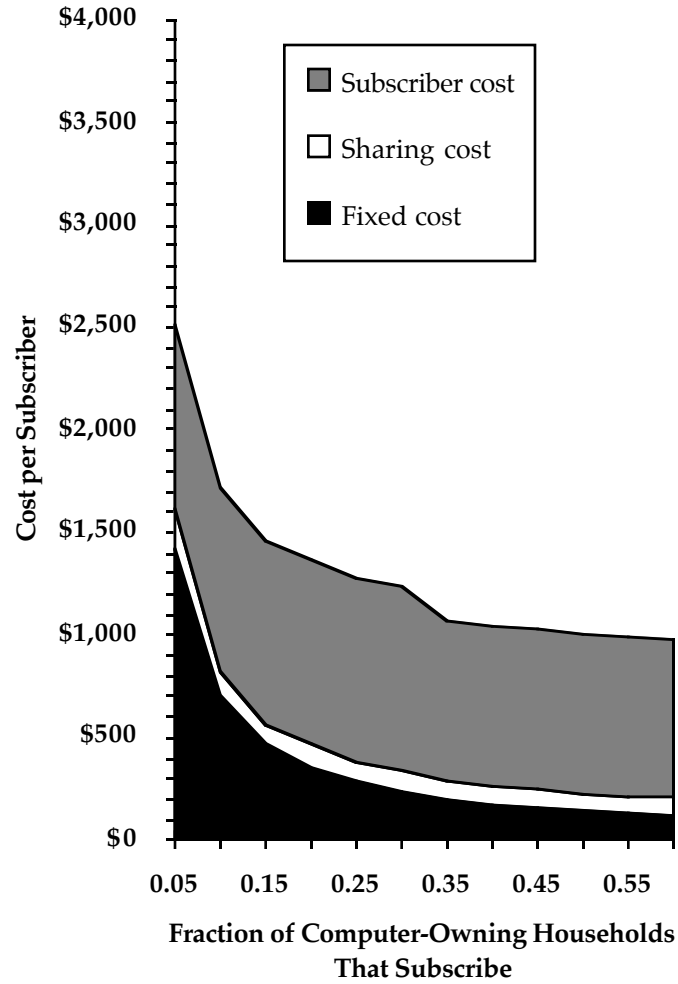


Figure 7.5: Components of average cost, 500 Kbps cable

The logarithmically-decaying shape evident in these graphs is caused by the geometrically slowing rate of increase in the denominator of average cost. From the first (not shown) to the 150th subscriber (5% penetration), the denominator increases by a factor of 150; from 5 to 10% penetration (150 to 300 subscribers), it doubles; from 10 to 15% penetration, it goes up by a factor of 0.5; etc.




Figure 7.6: Components of average cost, Organizational 4 Mbps cable




Figure 7.7: Components of average cost, 128 Kbps ISDN

7.5.1. Subscriber cost

Figures 7.5, 7.6 and 7.7 show that subscriber cost quickly becomes the dominant cost for all three services. Subscriber cost is already more than 50% of total cost per subscriber at 10% penetration (300 subscribers) for the 500 Kbps cable service, and at 5% penetration for the other two services. Furthermore, subscriber equipment's fraction of the average cost continues to rise as penetration increases.

The effect of the volume discount for the 500 Kbps cable service is visible in Figure 7.5 as a reduction in slope of subscriber cost between 30 and 35% penetration (900–1050 subscribers). Since the volume discount for the 4 Mbps cable service has already kicked in at 5% penetration (150 subscribers), its effect is present but not visible in Figure 7.6. Even with this discount, subscriber

cost is still impressively significant for the 4 Mbps cable service. This result reflects the case study's assumption that each 4 Mbps subscriber needs the equivalent of a branch-office router. Sensitivity analysis later in this chapter explores the effect of relaxing this assumption.

Given how important these results show subscriber cost to be, one would expect changes (presumably, reductions) in subscriber equipment costs over time to have a significant effect on the total cost to provide Internet access over local residential infrastructure. This result has several business implications, which are discussed in more detail in Chapter 8:

- Improvements in subscriber equipment are a key area on which to concentrate R&D energy;
- The telephone industry's current approach of letting subscribers buy their own equipment reduces capital risk significantly for the Internet and/or infrastructure provider. This point is reinforced by Figure 7.8, which illustrates how low average cost falls without subscriber equipment.




Figure 7.8: Average cost without subscriber equipment

7.5.2. Sharing cost

For all three services, the sharing cost remains a small fraction of the average cost for all penetration levels. Table 7.2 shows, however, that it is a much more significant fraction for the ISDN service than for the two cable services.

Table 7.2: Sharing cost's contribution to average cost

Service	Typical percentage
500 Kbps cable	8
4 Mbps cable	1
128 Kbps ISDN	21

ISDN's comparatively high sharing cost reflects the high cost of adding a terminal server port and the necessity of doing so frequently. The increment is a T1 line plus another port on a T1 terminal server, which are modeled as costing \$19,400 and \$4000, respectively.¹⁷⁴ Each T1 line is assumed to be running the Primary Rate Interface service of ISDN, which provides 23 "B" channels and can therefore handle no more than 11 active "2B" ISDN subscribers, regardless of the actual bandwidth used over each "B" channel. Thus sharing cost averages to about \$2130 (i.e. \$23,400/11) per active subscriber.

In contrast, cable's sharing cost consists of the cost to add a channel and a LAN, which are modeled respectively as \$3000 and \$900 (500 Kbps cable) or \$1300 (4 Mbps cable). Recall from Chapter 5 that each channel can provide one 4 Mbps LAN or four 500 Kbps LANs. The number of active subscribers each LAN can serve depends on the assumed bandwidth requirements. With the original assumption of 120 Kbps per subscriber, each 500 Kbps LAN can serve 4 active subscribers; thus each channel can serve 16 active subscribers. With a 4 Mbps LAN, each channel can serve 33 active subscribers. The 4 Mbps LAN shares more efficiently: sharing cost averages to about \$130 (i.e. \$4300/33) per active 4 Mbps subscriber, and \$415 (i.e. \$6600/16) for each active 500 Kbps subscriber.¹⁷⁵

The significant cost of the leased T1 lines needed to connect the Internet service provider to the local telephone network highlights another business and policy implication: ISDN Internet service would cost less to provide if these lines were not needed. Figure 7.9 illustrates the effect of removing this cost element from the model. The basic result of the model remains unchanged; the shared cable services are still more economical than the dedicated ISDN service. However, Figure 7.9 shows that an ISDN Internet provider who can eliminate T1 line charges has a cost advantage of approximately \$350 per subscriber.

¹⁷⁴Recall from Chapter 6 that T1 lease costs are amortized over 3 years to derive capital costs.

¹⁷⁵This difference, combined with the greater expense of 4 Mbps subscriber equipment, explains why the percentage of sharing cost is lower for the 4 Mbps service than for the 500 Kbps service.




Figure 7.9: Average cost with co-located ISDN Internet Point of Presence

One way to eliminate (or reduce) these T1 line charges is to co-locate the Internet service provider with the local telephone company Central Office (just as the cable Internet service provider expects to co-locate with the cable head end). In that scenario, an external T1 circuit is replaced with an intra-office wire. Alternatively, the Internet service provider might be the local telephone company itself. These possibilities are discussed further in Chapter 8.

7.5.3. Fixed cost

Figures 7.5, 7.6, and 7.7 demonstrate that fixed cost quickly becomes a small proportion of per-subscriber cost as penetration increases. Recall that fixed cost comprises the cost of an IP router, Internet circuit, and network

management for all three services, as well as the cost of each neighborhood's upstream transmission upgrade for the cable services.

The fixed cost thus depends on the size of the IP router and Internet circuit needed. This is a function of how much traffic is aggregated at the Internet Point of Presence (located at the cable head end or ISDN Internet provider's premises). Given the assumption of 120 Kbps average bandwidth per subscriber, the aggregate traffic level for any given penetration level is the same for each service, independent of the peak bandwidth.¹⁷⁶

With six neighborhoods, enough traffic has already been aggregated that a high-speed circuit and correspondingly high-end router are required even to serve 5% penetration. This configuration remains sufficient until 65% penetration, at which point fixed cost steps up by \$75,000 because another high-speed circuit (but not another IP router) must be added to handle the aggregate traffic volume.¹⁷⁷

Since the fixed cost stays constant at low penetrations while the subscriber base increases quickly, fixed cost per subscriber declines quickly (for example, it is halved by the doubling of the subscriber base between 5 and 10% penetration). By the time the step up is experienced at 65% penetration, the subscriber base is large enough to make its effect on average cost insignificant. In summary, fixed cost becomes a small fraction of the total once penetration levels reach approximately 20%, and it stays a small fraction for all penetration levels above that.

7.5.4. Discussion

The amortization of a low fixed cost over a large subscriber base and the more or less constant proportion of sharing cost together cause the average cost curves shown in Figures 7.6 and 7.7 to flatten out around 50% penetration.¹⁷⁸ Thus there appears to be a minimum level of overhead (i.e. fixed plus sharing cost) for each service, such that average cost is limited in how close it can come to subscriber cost.

Table 7.3 explores this minimum level of overhead by comparing the minimum average cost—i.e. the average cost at full penetration—against the

¹⁷⁶Differences in the fixed cost per subscriber between the cable and ISDN services are thus caused solely by the allocation of the cost of the upstream cable upgrade to Internet service. If this assumption were relaxed (for example because the cable operator simultaneously chooses to offer interactive TV services that also require upstream capability), the fixed costs would be identical for the two approaches.

¹⁷⁷This step up is not seen in Figure 7.5, since the 500 Kbps cable service cannot serve more than 60% penetration under the original set of parameter assumptions.

¹⁷⁸Presumably the same effect would be clear in Figure 7.5, if the 500 Kbps service could be extended beyond 60% penetration.

cost of subscriber equipment (also at full penetration, reflecting volume discounts).¹⁷⁹

Table 7.3: Minimum overhead levels

Service	Minimum average cost	Subscriber equipment cost	Average cost is above subscriber cost by:
500 Kbps cable	\$985	\$775	27%
4 Mbps cable	\$2,404	\$2,275	6%
128 Kbps ISDN	\$2,003	\$1,500	34%

The lower overhead of the 4 Mbps cable service (\$129) compared to the 500 Kbps cable service (\$210) reflects the greater channel efficiency of the 4 Mbps service as well as the slight reduction in overhead (i.e. further amortization of fixed cost) caused by extension of the service to 100%, rather than 60%, penetration. Of course, the 4 Mbps service's smaller overhead percentage also reflects its higher subscriber cost.

The high overhead (\$503 per subscriber) of the ISDN service reflects the large sharing cost of this approach, as discussed above. This overhead seriously limits the extent to which average cost can approach subscriber cost for this service.

7.6. Sensitivity analysis

Up to this point, all results have been examined for the original parameter values listed in Table 7.1. This section explores how these results are affected by variations in the input parameters, noting not only the direction but also the magnitude of the effects.

In each of the following subsections, the results of a sensitivity analysis are presented in which one parameter value is varied while the rest are held constant at their original values. In order, the parameters changed are: subscriber density, the number of neighborhoods served, the percent of subscribers assumed to be active simultaneously, the average bandwidth required by each subscriber, and the efficiency of the cable LAN bandwidth sharing algorithm. In addition, the results of a "what if" analysis are presented for reductions in subscriber equipment costs.

7.6.1. Residential density

The analysis above has assumed that subscribers live in a dense urban area like Cambridge, Massachusetts, in which it is possible to create neighborhood

¹⁷⁹Note that full penetration is 60% for the 500 Kbps cable service and 100% for the other two services.

cells such that up to 1500 subscribers can be served in each cell. As discussed in Chapter 5, this assumption may not hold in a more suburban or rural community. The area required to reach 1500 subscribers may be too geographically large to be practically served by a single cell; 500 homes per cell may be a more reasonable target.¹⁸⁰

It is not immediately obvious whether average cost will rise or fall when the residential density parameter's value is reduced from 1500 to 500 homes per cell. For example, with fewer subscribers per unit area, the cable is split less often, and the plant needs fewer amplifiers. As a result, the upstream transmission upgrade to these amplifiers costs less, and fixed cost declines, contributing to a decline in the average cost. On the other hand, sharing the fixed cost among a smaller number of subscribers increases the average cost.

Figure 7.10 plots the average cost of the three services with the residential density parameter value changed to 500 homes per cell. Comparing Figure 7.10 against Figure 7.2 shows that the overall effect of reducing the residential density is to increase the cost per subscriber, by about 5% at 100% penetration.¹⁸¹

¹⁸⁰Since most communities have a mix of different densities in different neighborhoods, the all-500 and all-1500 cell size assumptions should be considered boundary values.

¹⁸¹Average cost of the 500 Kbps service displays a greater differential, mainly because volume discounts apply to subscriber equipment at the full penetration level (60%) in the case of 1500 homes per cell, but do not apply at full penetration (100%) with 500 homes per cell.




Figure 7.10: Average cost with 500 homes per cell

Other differences between Figures 7.10 and 7.2 result from the observation that a lower subscriber density makes any given penetration level represent fewer subscribers. In Figure 7.10, each 5% increment in penetration represents 50 subscribers, instead of the 150 of Figure 7.2.

Since the level of traffic aggregated at the IPOP is lower for any given penetration level, the fixed cost function steps up differently. With 1500 homes per cell, a second Internet circuit is required at 65% penetration. In contrast, with 500 homes per cell, one Internet circuit is sufficient for all penetration levels. In fact, at very low levels (under 10% i.e. under 100 subscribers), a T1 circuit is sufficient to handle the load. This results in reduced fixed cost at those penetration levels and a step up to a T3 circuit and correspondingly more expensive router at 10% penetration.¹⁸²

Fewer subscribers for any given penetration level also means that fewer cable LANs or terminal server ports are needed. This effect not only reduces sharing cost, but also raises the penetration level that can be served by a limited number of cable LANs. Recall from Figure 7.4 that in the urban case, the 500 Kbps cable service is unable to serve beyond 60% penetration because the limit of 16 cable LANs per neighborhood is reached at that point. With one-third as many subscribers, nine cable LANs can serve 100% penetration, so the 500 Kbps cable graph extends all the way out in Figure 7.10.

Since the number of subscribers grows more slowly with increasing penetration, volume discounts take effect at higher penetration levels (10% for the 4 Mbps cable service, and 100% for the 500 Kbps). This effect raises the subscriber cost at lower penetration levels.

In summary, reducing the residential density from 1500 to 500 homes per cell increases the average cost by a small amount (5%). It does not change the basic results of Figures 7.2 and 7.3: the cable services still provide higher peak bandwidth per dollar than the ISDN service.

7.6.2. Community size

The analysis thus far has been restricted to an intermediate community size of six cells. Returning to the assumption of 1500 homes per cell, this section explores variation in the number of cells from one to twenty (1,500 - 30,000 homes).¹⁸³

¹⁸²This step up is particularly evident in Figure 7.10 for the ISDN service. For ISDN, the Internet Point of Presence (IPOP) is the only fixed cost; therefore, an increase in IPOP cost has a greater percentage effect on average fixed cost. The effect on the cable services appears in Figure 7.10 as a slightly slower rate of cost decrease between 5 and 10% penetration than between 10 and 15%.

¹⁸³Notice that the six- and twenty-cell analyses are based on the idealized assumption that the same penetration level applies in all cells. A different distribution of subscribers would not necessarily have much effect on the results, however. As long as the real community-wide average matches the model's, fixed and subscriber costs will be accurate, because the number of subscribers and the total bandwidth into the IPOP remain unchanged, regardless of the distribution of subscribers among cells. As we have seen, these costs comprise the bulk of overall cable cost. For the ISDN service, sharing cost is more important, but it too would be unaffected

Figure 7.11 presents average cost for the one-cell case. Note that with only one cell of 1500 homes, each 5% increment in penetration represents 25 subscribers.

by this assumption. The ISDN architecture does not depend on cells; sharing cost applies to all subscribers in the community, not per-cell.




Figure 7.11: Average cost with one neighborhood

The average cost to serve a single cell's worth of subscribers is higher than to serve six cells' worth. At 60% (i.e. full) penetration, the average cost of the 500 Kbps cable service is more than 50% (\$500) higher than with six cells. At 100% penetration, the average costs of the other two services are about 10% (\$200)

higher than their six-cell counterparts. The main factor contributing to this increase is the spreading of fixed costs over one-sixth as many subscribers.

As with the change to suburban subscriber density, the decrease in the number of cells reduces the aggregate IPOP traffic and correspondingly alters the fixed cost function. This alteration is the source of the jumps in average cost at low penetration levels; with one cell, a single T1 circuit and low-cost router are sufficient until 15% penetration (75 subscribers). From 15% to 100%, a single T3 circuit and high-end router are sufficient. Thus at 100% penetration, total fixed cost is lower than in the six-neighborhood case, which requires a second T3 circuit. It is not six times lower, however, so average cost increases.

A lack of volume discounts also contributes to this increase. For the 500 Kbps cable service, the volume discount at 1000 units never takes effect since the number of subscribers is six times lower for any given penetration level. The 4 Mbps discount at 100 units takes effect at 25% penetration (in contrast to 5% for the six-cell case), which contributes to higher average costs at penetrations below 25%.

The opposite effect is seen in Figure 7.12, which shows that the average cost to serve twenty cells is very similar to the average cost to serve six cells at full penetration.




Figure 7.12: Average cost with twenty neighborhoods

With 500 subscribers represented by each 5% increment in penetration, volume discounts already apply to the 4 Mbps service at 5% penetration. For the 500 Kbps service, they take effect between 10 and 15% penetration, resulting in a steeper drop in average cost at that point.

With such a large number of subscribers, the fixed cost function experiences more steps up as further Internet circuits and routers are added.¹⁸⁴ Because these increases are amortized over many subscribers, however, the effect on average cost is scarcely noticeable.

¹⁸⁴T3 circuits must be added at 20, 40, 75 and 95% penetration. Additional routers are required at 40 and 75% penetration to handle the extra ports.

These combined effects result in an asymptotic average cost that is very similar to (\$10 less than) the six-cell case, although it is reached at lower penetration levels (30% in Figure 7.12, vs. 50% in Figure 7.2). Thus, although a comparison of Figures 7.11 and 7.2 reveals clear economies of scale as the number of neighborhoods is increased from one to six, the increase from six to twenty cells does not buy as much scale economy. In other words, there are diminishing returns to scale.

This result is not surprising when we consider that average subscriber and sharing costs are only slightly affected by the number of subscribers (i.e. volume discounts) and therefore show very little economies of scale. While average fixed cost is greatly reduced as the first two thousand or so subscribers are added, beyond that point IPOP capacity must be increased to handle additional load, and average fixed cost stays fairly constant.¹⁸⁵ In other words, there are limits to how far fixed cost can be spread.

In conclusion, this analysis shows that the minimum efficient scale for providing residential Internet service is two thousand subscribers, under the original set of traffic assumptions (i.e. 20% of subscribers active and 120 Kbps average bandwidth). This result implies that reasonable entry strategies must serve multiple neighborhoods.

7.6.3. Percent active

This section presents the results of varying the assumption that during the peak busy period, 20% of subscribers are simultaneously active. "Active" is defined as sending traffic (cable model) or connected to the Internet Point of Presence (ISDN model). The rest of the models' parameters are set at their original values. Values of 10 and 30% active are tested.

Changing the number of subscribers assumed active simultaneously changes aggregate traffic requirements, but it does not change the total number of subscribers. Thus subscriber cost is unaffected by variations in this parameter, and volume discounts take effect at the same penetration levels as in the original case (between 0 and 5% for the 4 Mbps cable service, and between 30 and 35% for the 500 Kbps cable service).

Sharing cost, however, is affected. For the cable services, changing the aggregate level of traffic changes the number of subscribers able to share each LAN, and thus changes the number of LANs required to serve any given penetration level. This effect is illustrated in Figures 7.13 and 7.14 for the 500 Kbps and 4 Mbps cable services, respectively. The faster rise of the higher

¹⁸⁵Two thousand is derived from the jump up in fixed cost at 65% penetration for the 6-cell case, and 20% penetration for the 20-cell case. This number could change if the architecture were shifted to one of multiple head ends tied together, e.g. with a fiber ring, and sharing a single Internet Point of Presence among them. Such a shift represents a likely trend but is beyond the scope of this model.

percent active curves translates directly into greater sharing cost per subscriber, since sharing cost comprises per-LAN and per-channel equipment.




Figure 7.13: Number of LANs needed, 500 Kbps cable service




Figure 7.14: Number of LANs needed, 4 Mbps cable service

Figures 7.13 and 7.14 show that, in contrast to Figure 7.4, with 10% of subscribers assumed active, the limit on the number of cable LANs is no longer reached for the 500 Kbps cable service; while with 30% assumed active, not only is the limit of 16 500 Kbps cable LANs reached at a lower penetration

level (40%), but the 4 Mbps service is also limited in this case, to 85% penetration.

Sharing cost is also affected for the ISDN service, because changing the percent active assumption changes the number of subscribers expected to connect to the IPOP simultaneously. With a larger percentage of subscribers assumed active, each increment in penetration represents a larger number of simultaneously-connected subscribers. Therefore, more terminal server ports are needed for any given penetration level, and sharing cost per subscriber is correspondingly higher.¹⁸⁶

Finally, fixed cost is affected by the percent active assumption, because of its dependence on the aggregate level of traffic going out to the Internet. Recall that with 20% of subscribers assumed active, an additional circuit to the Internet is required at 65% penetration, causing a step up in fixed cost as seen in Figures 7.6 and 7.7. Varying the percent active assumption shifts the penetration level at which this step up is experienced. When only 10% of subscribers are assumed active, 100% penetration can be served with a single circuit. In contrast, with 30% of subscribers assumed active, the second circuit is needed at 45% penetration.

The combined effect of changes to fixed and sharing cost are shown in Figures 7.15 and 7.16, which plot the average total cost for 10 and 30% active, respectively.

¹⁸⁶More active subscribers might also increase the load on the telephone switch and inter-office trunking. Such an effect is beyond the scope of this thesis, which assumes that existing infrastructure is adequate to handle the longer hold times typical of data subscribers.




Figure 7.15: Average cost with 10% of subscribers simultaneously active




Figure 7.16: Average cost with 30% of subscribers simultaneously active

Figures 7.15 and 7.16 illustrate that the most significant effect of changing the percent active assumption is on the maximum penetration level that can be served with the cable LAN technology. For any given penetration level that can be supported, the average cost in Figures 7.15 and 7.16 is not significantly different from its counterpart in Figure 7.2.¹⁸⁷ This result is not surprising, since traffic loading has no effect at all on subscriber equipment cost, and as we have seen, the latter is a significant portion of total cost.

In summary, the fact that the expected traffic levels for residential Internet are largely unknown is slightly less important than it seemed in Chapter 1, since

¹⁸⁷The cable results are within 5% of the results in Figure 7.3, while the ISDN results are within 15%. The greater effect of this change on the ISDN service reflects the greater importance of ISDN sharing cost.

these levels scarcely affect average cost. They are, however, directly correlated to the maximum penetration level that can be served with the cable LAN technology. Given the assumption that each subscriber needs an average bandwidth equivalent to full ISDN service, the 4 Mbps cable service is much less limited than the 500 Kbps cable service in this regard. The following section explores the effect of varying the average bandwidth assumption.

7.6.4. Average bandwidth

This section explores the effect on average cost and penetration limits if the average bandwidth required by active users is less than the full bandwidth available over an ISDN connection (120 Kbps). Average bandwidth requirements greater than 120 Kbps are not investigated; while either cable service could support such a need, ISDN could not. Such a service would be unique to cable, and cannot be compared equivalently to ISDN.

The original 120 Kbps value is based on the assumption that at a minimum, the active cable user wants to receive an average bandwidth equal to the maximum dedicated bandwidth of the 2B ISDN circuit she could get as an alternative. If the user's alternative were instead a POTS modem or an ISDN circuit limited to a single data channel (1B ISDN), average bandwidth values of 30 and 60 Kbps (respectively) would be more appropriate. These values are tested in the sensitivity analysis.

These reduced values can also be thought of as representing a cable user whose alternative is 2B ISDN, but who is willing to accept a minimum average bandwidth less than the full dedicated bandwidth of ISDN. For example, users probably do not spend 100% of their time online transferring files. A value of 60 Kbps might be more appropriate for a user who transfers files half the time and spends the rest of the time on very low-bandwidth activities.

Figures 7.17 and 7.18 plot the average cost for assumptions of 30 and 60 Kbps of average bandwidth, respectively. The results are similar to those of the previous section: the main differences from Figure 7.2 are reductions in the number of cable LANs needed and shifts in the fixed cost step function. Like the percent active parameter, assumptions about average bandwidth have no effect on subscriber cost.




Figure 7.17: Average cost with 30 Kbps average subscriber bandwidth




Figure 7.18: Average cost with 60 Kbps average subscriber bandwidth

With each subscriber demanding one-fourth or one-half as much bandwidth, the number of subscribers that can share each cable LAN is increased by a factor of four or two, respectively. Thus, not only is sharing cost per subscriber reduced, but the number of cable LANs needed is smaller and the limits seen in Figure 7.4 are no longer reached. Full penetration can be supported for both the 500 Kbps and 4 Mbps cable services with either the 30 or 60 Kbps average bandwidth assumption.

Second, with each active subscriber demanding less bandwidth on average, the aggregate bandwidth into the Internet Point of Presence is correspondingly reduced. With average bandwidth reduced to 30 Kbps (Figure 7.17), a T1 circuit to the rest of the Internet is sufficient up to 10% penetration, thus explaining the reduced cost of this case at low penetrations.

At 10% penetration, a T3 circuit is required, but it remains sufficient even at 100% penetration, so no further steps up in fixed cost are seen. With average bandwidth of 60 Kbps (Figure 7.18), a T3 circuit is required immediately, but like the 30 Kbps case, this single circuit is sufficient throughout the penetration range and no steps up in fixed cost are seen.

ISDN sharing costs are unaffected by variations in the average bandwidth parameter. These costs depend only on the number of subscribers needing to connect to terminal server ports; they are independent of how much of the connection's bandwidth the subscriber is actually using.

Because of the dominance of subscriber equipment cost, changes to the average bandwidth assumption have very little effect on average total cost beyond initial penetration levels, especially for the ISDN service. For example, the shift in the fixed cost function reduces the average cost of the ISDN service by less than 1% at full penetration.

In summary, traffic loadings, as represented by the combination of percent active and average bandwidth assumptions, have little effect on total cost, especially for the ISDN service. For the cable services, their main effect is on the number of cable LANs needed to support the expected load. When the expected load is high, more cable LANs are needed than can be supported by the small number (4) of upstream channels. This effect is much more pronounced for the 500 Kbps service than for the 4 Mbps service.¹⁸⁸ Given the greater riskiness of the 500 Kbps service, a good entry strategy is to begin by deploying the 4 Mbps service, use this deployment to observe the traffic loadings that result, and use the observed data to determine whether loadings are sufficiently low that deployment of the 500 Kbps service makes sense.

7.6.5. Ethernet efficiency factor

Another uncertainty identified earlier in this report is the efficiency of the media access protocol that allows multiple subscribers to share a cable LAN. Since this protocol does not appear to have been analyzed or simulated, its actual efficiency is unknown. This section explores the effects of reducing sharing efficiency from its originally-assumed value of 100% down to 80% and 60%.¹⁸⁹ The resulting average costs are shown in Figures 7.19 and 7.20, respectively.

¹⁸⁸This result is not surprising, given the 4 Mbps service's greater spectral efficiency. The 4 Mbps modem converts each 6 MHz channel into one 4 Mbps LAN. In contrast, with the 500 Kbps modems, each 6 MHz channel supports a total of only 2 Mbps, because some of the bandwidth is used to separate the four LANs. Thus the 4 Mbps service can support twice as many subscribers on any given number of channels.

¹⁸⁹(Metcalfe and Boggs, 1976), p. 402 shows that the standard Ethernet sharing algorithm (CSMA/CD) is highly efficient (at least 98%) for reasonably-sized packets (at least 512 bytes). Therefore, the packet padding used to make standard CSMA/CD work over the long




Figure 7.19: Average cost with 80% cable LAN efficiency

cable distances (see Chapter 3) can impose overhead and still leave efficiency reasonably high.




Figure 7.20: Average cost with 60% cable LAN efficiency

Reducing the efficiency of the LAN effectively reduces the total bandwidth each LAN can provide, while the bandwidth demanded by each subscriber remains unchanged. As efficiency is reduced, the number of subscribers able to share each LAN is reduced and the number of LANs needed to serve any given level of penetration is increased. Thus sharing cost per subscriber increases. As before, the effect of this increase on total cost is very small because of the dominance of subscriber cost.

Reducing cable efficiency has exactly the opposite effect of reducing the average bandwidth demanded by each subscriber. Instead of the full penetration achieved in the plots of the previous section, Figures 7.19 and 7.20 show that cable LANs hit limits at ever lower penetration levels as their

efficiency is reduced. As expected, the limits are more severe for 500 Kbps LANs.

This result reinforces the strategy outlined in the previous section: start with the 4 Mbps cable service and use it to determine actual loadings and LAN efficiency. The next section discusses a critical issue that arises with this strategy: the high cost of the 4 Mbps subscriber equipment.

7.6.6. Subscriber equipment cost

The analysis in previous sections has shown the importance of subscriber equipment costs. This section examines the effect of alternative assumptions about the choice of subscriber equipment. Figure 5.21 demonstrates the potential impact of less expensive ISDN and 4 Mbps cable subscriber equipment.

ISDN cost reduction

A large component of the ISDN subscriber equipment cost is the cost of the ISDN-compatible bridge or router needed to mesh the subscriber into the TCP/IP network run by the Internet service provider. Such equipment is available in a variety of configurations from an ever-widening array of vendors.¹⁹⁰ Since this market is relatively new and highly competitive, appropriate feature sets are still in flux and prices are quite variable. The original model uses an intermediate value (\$1100, selected from a range of \$400-\$2500) for the ISDN bridge or router's cost, representing several different commonly-used vendors' equipment with an intermediate feature set. However, lower-cost alternatives do exist. The plot of 128 Kbps ISDN service in Figure 7.21 uses the lowest price currently available, \$600, for the subscriber's ISDN bridge or router.¹⁹¹

4 Mbps cable cost reduction

In contrast to the growing variety of ISDN equipment, few choices exist for cable-compatible data networking equipment, and competitive pressure on price appears less intense. While Zenith has recently announced standalone box versions of their previously card-based Homeworks products, they have declined to announce prices.¹⁹² Therefore, Figure 7.21 uses the same cable

¹⁹⁰See (Derfler, 1994) and (Robertson, 1994).

¹⁹¹This price is based on \$500 products available from Combinet and Intel, each of which requires a separate NT1; information based on product literature from Combinet and Intel, and personal communication from Professor David Farber, University of Pennsylvania. Since \$500 is a list price, and an NT-1 costs between \$100 and \$200, I am assuming a discount of between 0 and \$100 to arrive at a total price of \$600. Even cheaper products are available; see, for example, "Motorola Introduces Bitsurfer(TM) Digital Modem for ISDN", January 6, 1995, PRNewswire), which announces a \$500 list price (discounts implied) device including an NT1. However, such devices can only use one B channel (56 or 64 Kbps). Since such equipment cannot carry an average bandwidth of 120 Kbps, it cannot provide a service comparable to the other 3 shown in Figure 7.21.

¹⁹²See "Zenith Develops New Cable Modem," November 30, 1994, PRNewswire.

modem prices as Figure 7.2. It does, however, alter one important assumption about the router or bridge required to interface each subscriber's computing equipment to the Internet service provider's network.

Recall that the PSICable service assumes that a 4 Mbps subscriber is an organization (e.g. a bank branch), not an individual, and therefore bundles with the 4 Mbps service the equivalent of a branch-office router to interface the customer's LANs to the cable network. This bundling is a business decision; technically, the choice of modem speed and degree of router sophistication are independent.¹⁹³ Figure 7.21 tests an assumption different from PSICable's marketing decision: it assumes that the 4 Mbps customer is an individual and can use the same specially-engineered router as the 500 Kbps customer interfacing a single computer instead of a LAN. Therefore, the subscriber router cost is reduced from \$1600 to \$400, reducing the total subscriber equipment cost substantially.

¹⁹³Internex's pricing makes this independence clear; pricing is based on the matrix formed by the choice of connecting either an individual computer or a LAN, via either 1B or 2B ISDN service. See <http://www.internex.net> on the Internet for more details.)




Figure 7.21: Average cost with reduced subscriber equipment costs

The dramatic effect of this different strategy is evident in Figures 7.21 and 7.22. Even assuming the cheapest possible ISDN equipment, the 4 Mbps service is still more economical. It is 20% less expensive in absolute terms than the 128 Kbps ISDN service, and nearly forty times less expensive per bit of peak bandwidth.




Figure 7.22: Average cost per Kbps of peak bandwidth, with reduced subscriber equipment costs

In conclusion, this result reinforces the strategy outlined earlier of deploying the 4 Mbps service instead of the 500 Kbps service. Not only is the 4 Mbps service less prone to penetration limitations if traffic loadings turn out to lie towards the high end of the range, but it can be deployed to individuals for not much more capital cost than the 500 Mbps service—and provide much more bang for the buck.

7.7. Entry cost

While the discussion in this chapter has so far focused on the cost per subscriber, it is also interesting to consider the absolute levels of capital required to serve the high-bandwidth residential Internet market. This section presents total entry costs for two different entry strategies: “barebones” and “minimum efficient scale.”

With the barebones entry strategy, the provider installs the absolute minimum configuration of each piece of shared equipment. If subscriber demand grows beyond what that configuration can serve, additional

equipment has to be added. The alternative is to enter at the minimum efficient scale identified above (2,000 subscribers). While total entry cost is higher with this strategy, per-subscriber cost is more reasonable and therefore total costs may be more recoverable.

The numbers in this section should be interpreted with some caution. They are estimates of the cost of particular technical reference models; other detailed design decisions may produce different results. They represent capital costs only; operational costs, such as technical support for naive residential users or ISDN circuit usage fees, may be significant but are not included. Despite these caveats, the order of magnitude of these results should match reality reasonably well.

7.7.1. Cost of barebones entry strategy

Table 7.4 presents the costs of the barebones entry strategy. In this case, a single cell of 1500 homes is served by a single IP router and Internet circuit, with either a single cable channel pair and LAN or a single T1 terminal server and T1 line, and traffic parameters set to their original values. The rationale behind this strategy is to minimize risky up-front capital investment, rather than to provide a configuration that can accommodate further growth that may not happen. Alternatively, this strategy allows the provider to defer investment in equipment that is rapidly evolving, enabling the later purchase of potentially better and cheaper equipment.

Table 7.4: Barebones entry strategy: capital costs in thousand \$

		500 Kbps cable	Organizational 4 Mbps cable	Individual 4 Mbps cable	ISDN
Fixed cost	IP router	15	60	60	15
	Internet circuit	26	75	75	26
	Network management	10	10	10	10
	Upstream upgrade	11	11	11	–
<i>subtotal</i>		62	156	156	51
Sharing cost	cable: channel+LAN	4	4	4	23
	ISDN: Terminal server+T1				
<i>subtotal</i>		4	4	4	23
Subscriber cost	Subscribers supported	20	165	165	55
	Cost of each subscriber	0.9	2.275	1.075	1.5
<i>subtotal</i>		18	375	177	83
Total cost		84	535	337	157

The barebones 4 Mbps cable configuration is able to handle many more subscribers (165) than the 500 Kbps cable or ISDN services (20 and 55, respectively). This result implies that volume discounts can apply even to a minimal entry configuration for the 4 Mbps service, while they do not apply to the 500 Kbps entry configuration. It also means that a T3 circuit to the rest of the Internet and correspondingly more expensive router are needed right off the bat, raising the fixed cost of the 4 Mbps service.¹⁹⁴

7.7.2. Cost of minimum efficient scale entry strategy

Table 7.5 presents the results of the minimum efficient scale entry strategy, in which the configuration is determined by the smallest number of subscribers

¹⁹⁴Earlier in this chapter we saw that the T1 to T3 break happens at approximately 75 subscribers given the original traffic assumptions. Although it would theoretically be possible for a provider to serve 150 4 Mbps customers with a T1 circuit, customers would experience degraded service. They would not be able to communicate with the Internet any faster than 1.5 Mbps, despite their faster equipment. This might be a business decision the provider chooses to make, but it would need to set customer expectations and price the service accordingly. Given the low penetration levels at which the T3 configuration becomes necessary for all 3 services, starting with a T3 makes sense if the provider expects the number of subscribers to grow much at all.

(about 2,000) that produces the lowest possible cost per subscriber. Since the rationale behind this strategy is to reach a large number of potential subscribers, a 20-neighborhood configuration is assumed, raising the upstream upgrade cost.¹⁹⁵ With two thousand subscribers, volume discounts apply to both the 500 Kbps and 4 Mbps cable services.

Table 7.5: Minimum efficient scale entry strategy: capital costs in thousand \$

		500 Kbps cable	Organizational 4 Mbps cable	Individual 4 Mbps cable	2B ISDN
Fixed cost	IP router	60	60	60	60
	Internet circuit	75	75	75	75
	Network management	10	10	10	10
	Upstream upgrade	220	220	220	–
<i>subtotal</i>		365	365	365	145
Sharing cost	cable: channel+LANs	160	80	80	850
	ISDN: Terminal server+T1s				
<i>subtotal</i>		160	80	80	850
Subscriber cost	Subscribers supported	2,000	2,000	2,000	2,000
	Cost of each subscriber	0.775	2.275	1.075	1.5
<i>subtotal</i>		1,550	4,550	2,150	3,000
Total cost		2,075	4,995	2,595	3,995

Notice in both Tables 7.4 and 7.5 that the absolute capital cost of entry is actually quite low with either strategy, especially when subscriber costs are excluded. For example, even with the minimum efficient scale strategy, the 4 Mbps cable service requires only \$445,000 in fixed and sharing costs, of which half is needed for upgrading the cable plant to pass traffic upstream.¹⁹⁶ To put this cost in perspective, it is comparable to the fully loaded salaries of 2-3 professionals planning the service for a year. The ISDN service is more expensive, but can still be started for under \$1 million in fixed and sharing costs.

¹⁹⁵With 20 neighborhoods, the minimum efficient scale is already reached at 20% penetration.

¹⁹⁶The entry cost would be \$220,000 less (i.e. \$225,000) if that upgrade had already been performed on behalf of some other application.

While the subscriber equipment cost subtotal in Tables 7.4 and 7.5 is certainly part of total cost, it does not necessarily have to be included when calculating capital requirements for an Internet and/or infrastructure provider to enter the residential Internet market, for two reasons. First, Internex has shown that having subscribers pay for their own equipment is a viable business model. Second, even if the provider does purchase the subscriber equipment, it can choose to buy “just-in-time”, so that expenditures do not have to occur far in advance of market demand.¹⁹⁷

The resulting low entry costs are partly a consequence of assuming away the cost of the underlying infrastructure, i.e. the existing telephone network and cable TV plant. While Reed has addressed that topic in detail (Reed, 1992; Reed, 1993), this report has treated infrastructure cost as sunk. Given that assumption, it has shown that the economic barriers to entry are quite low for incrementally adding Internet service to existing infrastructure.

7.8. Summary

This chapter has presented the results of an engineering cost model of residential Internet access over cable and ISDN. It explored in detail the following key findings:

- For the same average bandwidth, cable-based access is more economical than ISDN and can provide higher peak bandwidth.
- The potential penetration of 500 Kbps cable-based access is constrained in many cases by limited upstream channel bandwidth. Constraints on 4 Mbps access appear only under worst-case parameter assumptions and are less severe (i.e. higher penetration values can be reached) than for the 500 Kbps service under those assumptions.
- Subscriber equipment is the most significant component of average cost.
- Sensitivity analysis on the model’s parameters did not change these basic results. In particular, changing traffic assumptions had very little effect on average cost, although it did affect penetration limitations.

The following chapter concludes the report with a discussion of the technical, business and policy barriers to residential Internet access over cable and ISDN. It also outlines further research in this area.

¹⁹⁷A just-in-time approach would force the provider to forego volume discounts, but would clearly have the financial advantage of not tying up capital, especially in the face of demand that may never materialize.

Chapter Eight

Conclusions

This chapter summarizes the results of the report, comparing cable and ISDN Internet services. It discusses technology, business and policy barriers to the diffusion of cable and ISDN Internet access, and recommends changes to corporate strategy and public policy to overcome these barriers. After suggesting further research, the chapter concludes with a summary of the report's key points.

8.1. Summary comparison of services

Table 8.1 summarizes the comparisons between the cable and ISDN services discussed in this report. A brief discussion of each point follows the table.

Table 8.1: Comparison of cable vs. ISDN Internet services

Basis for comparison	Cable	ISDN
Economics		
Average cost per average Kbps	Individual 500 Kbps: \$8.20 Individual 4 Mbps: \$10.03 Organizational 4 Mbps: \$20.00	Individual 128 Kbps: \$16.70
Average cost per peak Kbps	Individual 500 Kbps: \$2.00 Individual 4 Mbps: \$0.30 Organizational 4 Mbps: \$0.60	Individual 128 Kbps: \$15.65
Bandwidth	Peak of 500 Kbps or 4 Mbps	Peak of 128 Kbps
Connectivity	Full-time	Part-time
Penetration scalability	Penetration limits, especially for 500 Kbps service	Limited free wire pairs and switch capacity for additional ISDN service
Potential subscriber population	Local cable system	Mainly local calling area, but access from worldwide telephone network
Commercial availability	Cambridge, MA	Several regions in U.S., Europe

8.1.1. Economics

The central result of the modeling exercise reported in this document is that the equipment needed to provide Internet access over cable is more economical than that for ISDN. As Figure 7.3 showed, for the same level of average bandwidth, cable's shared-bandwidth approach can provide much higher peak bandwidth for a much lower price per bit. For example, under the original set of assumptions discussed in Chapter 7—in which the degree of sharing is set to provide the same level of average bandwidth (120 Kbps) for both cable and ISDN access—500 Kbps Internet access over cable can provide four times the peak bandwidth of ISDN access for less than half the cost per subscriber. In addition, 4 Mbps cable access can provide thirty-two times the peak bandwidth of ISDN for only 20% more cost per subscriber. When 4 Mbps subscribers are assumed to be individuals instead of organizations, this result changes to 40% **less** cost than ISDN. The economy of the shared bandwidth approach is most evident when comparing the per-subscriber cost per bit of peak bandwidth: \$0.30 for Individual 4 Mbps, \$0.60 for

Organizational 4 Mbps, and \$2 for the 500 Kbps cable services—versus close to \$16 for ISDN.

Both the ISDN and cable approaches incur low fixed costs. Other than the allocation of cable's upstream upgrade to Internet service, these fixed costs are identical.¹⁹⁸

Other overhead costs (called “sharing” costs in Chapter 7) are higher for ISDN than for cable. A major driver of this difference is that, unlike the cable model, the ISDN model does not assume co-location of the Internet Point of Presence with infrastructure facilities, and the Internet provider's connections to the telephone network are expensive. Even with co-location, however, the ISDN approach still requires equipment to multiplex and concentrate individual subscriber lines, while the shared cable approach does not, creating a cost advantage.

With low infrastructure (i.e. fixed plus sharing) costs, the cost of subscriber equipment is a very important component of average cost. This result is encouraging for several reasons. First, it means that cost reduction efforts focused on subscriber equipment will have a significant impact on overall costs. Second, subscriber equipment does not have to be an upfront investment; it can be purchased as demand—and the revenue to go with it—materializes. (The deferred purchase approach also has the advantage of automatically incorporating cost reductions created by technological progress in subscriber equipment.) Subscriber equipment may even be paid for directly by the customer, as in the ISDN case. In either case, the level of upfront infrastructure investment needed is low. Therefore, the risk of making that upfront investment in the presence of uncertain demand is correspondingly low. As Cook and Stern note:

In the residential area...low initial penetrations are likely for many services and this will call for the networks that minimize upfront infrastructure costs and permit 'Just-in-Time' (JIT) provisioning at the customer end.¹⁹⁹

Table 8.2 compares the relative magnitude of upfront infrastructure and per-subscriber costs if all computer-owning households in the U.S. were to be provided with Internet access over cable or ISDN today. Cable's infrastructure cost is about \$8 billion, half that of ISDN. In both cases, the subscriber equipment cost is much more significant. In other words, residential Internet access delivered over either existing infrastructure meets the criteria

¹⁹⁸Professor David Tennenhouse points out that the telephone network's supposed advantage in already being switched does not change this result, since voice switching is not useful for IP. This result might be different if ATM switching, which could be used for both voice and data, were already in place.

¹⁹⁹(Cook and Stern, 1994), p. 85.

identified by Cook and Stern—given that the investment that was required to create that existing infrastructure (e.g. ISDN Central Office switching) is treated as a sunk cost.

Table 8.2: Capital cost to serve all computer-owning households in U.S. (in billion dollars)²⁰⁰

Type of cost	Cable			ISDN
	Individual 500 Kbps	Individual 4 Mbps	Organizational 4 Mbps	Individual 128 Kbps
Infrastructure cost	9	8	8	17
Subscriber cost	26	36	76	50
Total cost	35	44	84	67

Cable's economic advantage is not the only difference between the two approaches. The following sections explore functional differences between Internet over cable vs. ISDN.

8.1.2. Bandwidth

The peak bandwidth of the cable services, 500 Kbps or 4 Mbps, is obviously much higher than ISDN's 128 Kbps. This allows cable Internet customers to run applications, such as videoconferencing or viewing and presenting multimedia World Wide Web pages, with higher quality and without wasteful and tiresome delays waiting for graphics to appear. Higher peak bandwidth may also enable some future application that cannot run at all within ISDN's limit of 128 Kbps.

Higher peak bandwidth, however, does not guarantee average bandwidth. Since cable LANs are a shared media, the number of subscribers sharing each LAN and their usage patterns will determine how often the peak bandwidth is available to any individual subscriber, and what fraction of the peak is available on average. The model's results are based on setting the number of subscribers per LAN such that during periods of peak usage, the bandwidth

²⁰⁰This table extrapolates average costs based on the original parameter settings of Chapter 7 (including 1500 households per cable "cell") and assuming all areas to be served are of at least the minimum efficient scale. Since these assumptions may not hold in rural areas, the estimates in Table 8.2 should be considered as lower bounds. The relative magnitudes should remain similar, however. As Chapter 5 describes, one-third of U.S. households (i.e. about 33 million households) are assumed to have computers. Neither cable nor ISDN presently reaches all U.S. households. This effect is ignored in computing Table 8.2, since the percentage of unreached households with computers is presumed to be small (under 5-10%). Unfortunately, it is probably the same households unreached by both networks: homes in remote rural areas that would require uneconomically long wire lengths. For these homes, wireless Internet access may prove to be a better alternative as it develops in the future.

available on average to each active subscriber amounts to the same bandwidth as ISDN. Reality, however, can diverge from this model in two ways.

First, the service provider may overload the LANs. The analysis in Chapter 7 has shown little economic incentive to do this (sharing costs are very low even with the model's assumption), but it has also shown that the number of available cable LANs may in some cases be too small to serve higher penetration levels, potentially leading to overload. If cable modems are able to provide higher bandwidth before these higher levels of penetration are reached, this problem becomes less likely.

The second divergence is that the actual performance observed at any given moment may be quite different from the average, because of the statistical nature of traffic loadings. With small numbers of subscribers, one LAN may be overloaded while the next one is empty.²⁰¹ The risk of the cable approach is that the perceived bandwidth—i.e. the subscriber's perception of application performance in the presence of a variable data rate infrastructure—will be either too low or too variable. It might be too low if too many users share each LAN. Alternatively, if only a few users share each LAN, perceived bandwidth will be highly variable.²⁰² Clearly, it is important to set customer expectations reasonably, so that a customer's inability to capture the peak bandwidth 100% of the time does not come as an unpleasant surprise.

8.1.3. Connectivity

As discussed in Chapter 3, because the cable approach is connectionless, it can deliver full-time Internet access. Calls do not have to be established for traffic to flow, and only when traffic is flowing is bandwidth actually used. Thus traffic can be sent or received at any time, and the Internet is always accessible.

Full-time access is less realizable for the ISDN approach. Although ISDN can set up circuits much more quickly than POTS, it still requires connection establishment before any data can flow. While in theory, full-time access could be achieved by leaving ISDN calls connected 24 hours a day, this approach runs into several practical problems. First, the Internet provider would have to dedicate a terminal server port and transmission facilities to each individual user. Second, ISDN circuit mode tariffs generally incorporate a per-minute usage fee which makes such an approach prohibitively expensive.²⁰³ Such tariffs may be set deliberately to discourage long call hold times, since these are not incorporated into the teletraffic models used by

²⁰¹Sophisticated cable LAN equipment that could dynamically redistribute subscribers to LANs based on load would help with this problem.

²⁰²This risk is greater for the 500 Kbps than the 4 Mbps service, since fewer users can share each lower-bandwidth LAN (see Table 5.7).

²⁰³InterneX's Web page (<http://www.interneX.net>) suggests the use of Frame Relay instead of ISDN for those needing full-time access.

telephony planners to decide how much switching and trunking capacity is needed. Whether or not permanent circuits are tariffed unattractively, they waste resources when no traffic is flowing, leading to greater expense in the system as a whole (e.g. by requiring more telephone switching and trunking resources).

Practically speaking, then, ISDN is a dial-up approach. Thus ISDN access is not only part time, but also less practical for incoming traffic. The subscriber's equipment may be intelligent enough to dial when traffic needs to be sent, but it cannot know when traffic is waiting to be received. Internet provider's equipment would have to incorporate a dial-out capability to provide that type of service. In either case, dialing incurs delays that may be visible to users.

8.1.4. Penetration and service scalability

Chapter 7 showed that the constraint of four upstream cable channels can limit penetration levels in some cases, and is more significant for the 500 Kbps cable service than the 4 Mbps.

If improvements in digital technology increase LAN bandwidth so that the 6 MHz channel spectrum is used more efficiently for digital data, maximum penetration levels can increase. Of course, so will peak bandwidths. 10 Mbps cable modems currently exist, and higher speeds are planned for the future. Cable's broadband architecture allows a service provider to migrate to faster LANs on one channel at a time, preserving backward compatibility for satisfied customers while providing enhanced service for early adopters.

ISDN also faces penetration limits but they are more hidden. While ISDN was originally conceived as a replacement for POTS, in reality most subscribers purchase it as an additional service. If enough residences do this, new pairs will have to be installed in the local telephone distribution plant, raising the cost of ISDN service so high that alternatives such as fiber might be considered instead.²⁰⁴

ISDN does not provide as graceful a migration path as cable does to higher bandwidths. The envisioned follow-on to ISDN, ADSL, requires large investments such as new switching equipment, that cannot easily be made for a small number of early adopters.

²⁰⁴The distribution plant may be buried cable (not in conduit) and therefore very expensive to modify. In some cases, even subscribers replacing their POTS service with ISDN may force the use of new pairs. For example, some wire pairs may be in too poor condition to support ISDN effectively.

8.1.5. Availability

Internet over ISDN is currently more widely available than Internet over cable. However, we are a long way from the day when you can put an ISDN or cable interface in your laptop computer and use it anywhere in the U.S.—let alone the world.

At the time of this writing, Internet over cable is not available as a commercial service anywhere outside PSICable's offering in Cambridge, MA.²⁰⁵ Even in Cambridge, only the organizational 4 Mbps service is being sold; the 500 Kbps service for individuals still cannot be purchased, despite its announcement on March 8, 1994.²⁰⁶

While ISDN is in theory available throughout the U.S., in practice it may be harder or easier to get depending on the Regional Bell Operating Company (RBOC). Pricing is also an important aspect of availability. In RBOC regions with tariffs that encourage data usage (i.e. low or no per-minute usage charges), several Internet service providers have begun offering ISDN connections.²⁰⁷ In other regions, tariffs effectively prevent such developments.

8.1.6. Network scope

An Internet service provider who has established an Internet Point of Presence (IPOP) attached to the telephone network can serve any customer who can call the IPOP's phone number. In contrast, an IPOP attached to an isolated cable system can provide access to the Internet from only that cable system's subscribers. A cable Internet subscriber who plans to access the Internet while traveling will additionally have to subscribe to a telephone-network based Internet service provider.

The smaller potential subscriber population of Internet over cable serves as a barrier to entry. Since the same penetration level represents fewer subscribers, it is more difficult to reach minimum efficient scale with the

²⁰⁵Two other trials were identified during the course of this research. Jones InterCable worked with Advanced Networks and Services, Inc. (ANS) to deliver an Internet connection to a school library in Alexandria, VA. Interestingly, this service was delivered over a completely coax network, not a hybrid fiber and coax architecture. Residential service was being trialed to a small number of Jones employees. Through the National Science Foundation's Global Schoolhouse Project, Media General Cable worked with SprintLink, also to bring an Internet connection to a school library. Several cable companies have announced plans to offer access to the Internet and commercial on-line services, including TCI (the largest U.S. cable company), suggesting that Continental may not stay the lone cable company offering Internet access. See (Robichaux, 1994) and (Robichaux, 1995).

²⁰⁶Conversation with PSI telephone sales representative, December 12, 1994; comments of David Fellows of Continental Cablevision at MIT Communications Forum, April 27, 1995.

²⁰⁷A scan in December, 1994 of provider listings on the Internet turned up Internet over ISDN offerings in BellSouth, Pacific Bell, and Ameritech territories, as well as in Europe. These three RBOCs are reputed to have the most encouraging ISDN tariffs.

cable approach. The cable industry trend toward interconnection of cable systems should help with this problem.

8.2. Technology, business, and policy barriers

Since the summer of 1994, Internex, Inc. has grown steadily selling Internet connections over ISDN.²⁰⁸ Other Internet over ISDN providers have appeared in other regions, including PSI.²⁰⁹ During the same period, the PSICable service has attracted few if any subscribers, and other commercial offerings of Internet over cable have not materialized. This section explores the technical, business, and policy barriers that have made the diffusion of ISDN Internet access lag that of standard dialup access, while completely blocking the deployment of Internet access over cable to date.

8.2.1. Equipment immaturity

The diffusion of Internet over residential cable today is severely hindered by the inadequacy of currently available cable modems. While a reasonably large number of vendors sell equipment for broadband office LANs, the research for this report identified only three vendors who currently have a product for sale that can provide symmetrical LAN service over the distances found in residential cable plants.²¹⁰ While each vendor's product has its own strengths and weaknesses, no one product combines all the features that are needed to make large-scale deployment feasible: a price aimed at residential customers, adequate security and traffic filtering (i.e. not only bridging but also routing), frequency agility, and reliable operation with automatic monitoring and maintenance capabilities. In addition, each vendor uses its own proprietary "Ethernet-like" access method. Since any one vendor's equipment is not guaranteed to interoperate with any other's, an Internet over cable provider today would be locked in to a single vendor, with all the potential negative consequences of such an arrangement. Thus there are both technological and business reasons why a provider would not rush to offer this service today.²¹¹

²⁰⁸In November of 1994, Internex had 175 customers and was adding 30 to 40 per month (email from Robert Berger of Internex, Inc., November 17, 1994). Many of Internex's customers are small businesses.

²⁰⁹See (Lindstrom, 1994). (Wagner, 1995), p. 1 describes "a 38 percent increase in the number of U.S. Internet service providers offering ISDN connections in the past four months—66 of them [in January 1995], compared with 48 in September [1994], according to listings maintained by InterNIC Information Services, a branch of General Atomics in San Diego. About a dozen more said they planned to bring out ISDN services in early 1995..."

²¹⁰Zenith Communication Products, Glenview IL, sells the cable modems used in the case study. LANCity Corp., Andover, MA, sells a 10 Mbps cable modem plus bridge. Cactus Computer, Austin TX, also sells cable modems. As new products continue to be introduced in this industry, subscriber equipment costs can be expected to fall. Note also that several prominent vendors, most notably General Instruments and Intel, have announced asymmetric cable modems that do not provide LAN-style networking. Instead, they support fast downstream transmission (e.g. 10 Mbps) but very low upstream bandwidths (e.g. 9.6 Kbps).

²¹¹An effort to standardize the cable access method has recently gotten underway in the IEEE 802.14 standards committee. If it succeeds in producing a useful standard, interoperability

The technological challenges involved in designing a reliable cable modem should not be underestimated. Cable plants have generally been built and engineered to carry TV signals in one direction. Carrying data in two directions is harder for two reasons:

- If a noisy channel causes errors, accurate data transmission requires correction of those errors through technological mechanisms such as coding. With TV signals, human viewers do the error correction.
- Upstream transmission exercises parts of the plant that may never have been tested or maintained.

Given that engineering the plant to be more reliable is expensive, the burden has typically fallen on modems to cope with a wide variety of transmission problems that are only now being characterized. Some of these problems are quite different from those found in the office environments that ordinary LAN equipment vendors are accustomed to; for example, outdoor temperature changes affect cable amplifier gain and thus transmission characteristics.

Although similar environmental factors apply to ISDN transmission, their effects on data-over-ISDN equipment are somewhat less significant. Since telephony is a bi-directional service, upstream transmission capability is better maintained and understood, placing less burden on the customer's ISDN transmission equipment. In addition, the effects of these environmental factors on data transmission are better understood, based on 20 years of industry experience with POTS modems.

These environmental effects do still matter, however. ISDN only became feasible when adaptive techniques were developed to cope with dynamically variable impairments. ISDN transmission equipment has thus started the race with a head start relative to cable modem technology. However, the same advances in adaptive techniques that are being developed for ADSL can be applied to cable modems as well.²¹² Therefore, we can expect cable modem technology to narrow ISDN's lead.

8.2.2. Market immaturity

In addition, the market for data-over-ISDN equipment is more mature than data-over-cable, as evidenced by a larger number of vendors, a wider variety

problems could be mitigated in the future. See "IEEE Computer Society Works On Two-Way Cable-TV Protocol Standard, 802.14," IEEE Project 802.14 press release, November 11, 1994. Progress may also follow from a Request for Proposal (RFP) for "High-Speed Cable Data Service (HSCDS)" issued by CableLabs, April 26, 1995.

²¹²For example, the Discrete Multi Tone approach described in (Northern Telecom Inc. and Amati Communications Corp., 1993) for ADSL can also be used for cable modems.

of price points and feature sets (see (Derfler, 1994)), and at least some interoperability among different vendors' equipment. Two factors contribute to this difference: business usage of ISDN for data networking, and the ability of individual buyers to purchase data-over-ISDN equipment.

Not all data-over-ISDN equipment is aimed only at the home market; businesses can (and do) use ISDN to connect, for example, LANs in remote sales offices to the corporate LAN. As with the diffusion of personal computers, businesses with a real economic need have of course proven to be more reliable customers than home users purchasing equipment for discretionary purposes, providing the necessary volume to jump-start product development and sales.

This ability to bootstrap a residential product off of a business base (or alternatively, a public networking product off of a private networking base) arises from the greater scope of the telephone network. A business can use ISDN to connect remote offices to its enterprise network. In contrast, a business cannot use the cable system in this fashion today (except where multiple offices within the same community need to be interconnected). Thus there is little private network product base from which to bootstrap public cable network products.

A second factor contributing to the lesser maturity of the data-over-cable equipment market is that the buyer is a cable operator, not an individual or business user. This structure creates a chicken-and-egg problem: before committing to buying a large batch of cable modems, the operator wants evidence of a reliable, affordable product; but before investing the necessary R&D to produce that, the equipment vendor wants to see a reasonable probability of orders. If the unit of purchase were smaller (as it is in the ISDN market, where each user purchases his own equipment), each purchase would entail less risk and therefore be more likely to happen. Thus the ISDN structure allows multiple entrepreneurial vendors to bet on different types of customers, while the cable structure enforces more uniformity as it constructs a barrier to innovation.

The ISDN market structure also has room for improvement, however. As long as different regions of the U.S. and world adopt radically different ISDN pricing schemes, demand remains fragmented. More consistent pricing would broaden equipment and services markets, allowing producers to take advantage of economies of scale that could further lower prices and boost demand.

Summary

In summary, for Internet over cable to flourish, more R&D and standardization efforts need to be focused on cable modems. In addition, if individual customers are allowed to purchase their own cable modems and

connect them to the network, market forces would help to improve both the price-performance and the inter-vendor compatibility of this equipment.

8.2.3. Marketing strategy

The models described in this report are based on a barebones Internet access product that simply connects the customer to the Internet. In reality, a variety of other features—such as electronic mail and network news storage, training services, etc.—could be bundled with the product or offered as higher-priced options.²¹³ Before product and pricing strategies can succeed, however, providers need a clearer idea of customer needs than they currently display.

Both PSICable and Internex's ISDN service are higher-priced than POTS dialup services. Therefore, to succeed, they must provide the additional level of service that is needed by the customer who is willing to pay more. Internex's pricing reflects such an understanding: prices and features (such as number of Internet Protocol (IP) addresses) differ depending on whether the customer is connecting an individual computer, a small LAN (2-5 computers), or a larger LAN.

In contrast, several mismatches are apparent in the PSICable service. For example, PSICable's 'Organizational' service offers 4 Mbps bandwidth over the cable LAN, but only 1.5 Mbps of bandwidth into the Internet for any one customer. Thus, since there are no high-bandwidth services available directly on the cable LAN, the 4 Mbps service essentially provides the same level of service as a leased T1 line. Since the 4 Mbps service is priced only slightly less than T1 Internet connections available from other Internet service providers, business customers do not have much incentive to adopt the new service.

Part of the high price of the 4 Mbps service reflects its targeting of business customers, assuming that they require full branch-office routers. Sensitivity analysis in Chapter 7 showed that if this assumption were relaxed, average costs would fall so much that prices could not only compete with but even undercut those of 128 Kbps ISDN service. When positioned as a service for individuals, competing against 128 Kbps ISDN, the 4 Mbps service provides great value for customers and has much more potential for success.

Another area of potential mismatch is the number of IP addresses provided to PSICable subscribers. The special box PSICable developed to substitute for a full router provides no more than 3 IP addresses. While this is more than enough to connect an individual computer, it seems likely that the early adopters of full-featured, more expensive Internet access will include technophile households already containing more than one computer and

²¹³Discussions of pricing strategies—as well as their implications for potential return on investment, taking into account the rapid depreciation of the equipment involved—are important but beyond the scope of this thesis.

possibly even a LAN. A little bit of market research would go a long way in helping to match the service to the likely customer needs.

Finally, moving this service beyond technophiles will involve other marketing issues such as ease of use, price, and brand awareness. An understanding of the different market segments is essential to deciding where to focus first. For example, telecommuters may be the largest group of early adopters. With costs mainly borne by their employers, they would be much less price sensitive than community networkers (e.g. parents using the service to send email to teachers). Alternatively, ease of use may be more important. For example, the Small Office/Home Office (SOHO) user may not want to have to understand multitudinous service options enough to choose intelligently among them; simplified, bundled feature sets targeted at typical SOHO needs will be more appealing.²¹⁴ Brand awareness, on the other hand, is likely to be important to all customer segments.

8.2.4. Corporate strategy

Delivery of residential Internet access has an obvious dependence on residential communications infrastructure. Thus monopoly infrastructure providers—cable and telephone companies—are in a position of power with respect to the offering of residential Internet service. Depending on how they use this power, three different business arrangements may emerge. The infrastructure provider may choose to:

1. Acquire the skills needed to offer the service, thus becoming an Internet service provider themselves;
2. Partner with one or more Internet service providers to offer the service jointly;
3. Not offer the service at all.

Current high-speed residential Internet offerings include examples of all three arrangements. The Regional Bell Operating Company Ameritech has pursued the first strategy, having won a contract from the National Science Foundation to operate one of the Network Access Points critical to the new Internet architecture and thus acquired Internet expertise.²¹⁵ PSICable and Internex are examples of the second strategy; PSICable is a contractual arrangement between Continental Cablevision and Performance Systems International (an Internet service provider), while Internex is a Pacific Bell ISDN reseller.²¹⁶ The third strategy, however, is by far the most common; it has been adopted, consciously or not, by all other cable operators and many telephone companies.

²¹⁴Research for this thesis uncovered confusion even among PSICable service representatives about IP addressing. Telephone company ISDN representatives revealed similar confusions over data pricing.

²¹⁵See (Messmer, 1994).

²¹⁶The Jones InterCable and Media General Cable trials also fall into this category.

Dependence on infrastructure assets

Two critical differences explain why the widespread choice of infrastructure providers not to offer Internet service (to date) has completely blocked the development of Internet over cable services but not stood much in the way of telephone network-based Internet services. The first is a policy difference. Telephone companies are common carriers; they must “furnish...communication service upon reasonable request” in a non-discriminatory fashion.²¹⁷ If the customer is an entrepreneur who wants to use telephone lines to offer residential Internet access, so be it; the entrepreneur does not have to convince the telephone company of the wisdom of his business plan, and the telephone company is assured of rents whether or not the business ultimately succeeds. Thus dialup Internet service has been able to flourish without anyone in the telephone company even having to know what it is, let alone deciding to offer it themselves.²¹⁸

The cable situation is different. Cable operators are not common carriers. Federal cable law contains provisions for the lease of a small number of channels “for commercial use by persons unaffiliated with the operator,”²¹⁹ but weak interpretation of the law by the Federal Communications Commission (FCC) and the lack of common carrier status have rendered these provisions largely ineffective.²²⁰ Offering Internet over cable requires the dedication of a pair of channels to data service. If all channels are currently utilized, offering Internet thus requires taking something else away—a definite entry hurdle when existing channels generate known revenue.²²¹ If a monopoly cable operator does not believe that selling Internet access is the most profitable use of a channel, there is little opportunity for an entrepreneur to prove him wrong. The flavor of the hold-up is captured well by this quote from a recent newswire story:

“It takes a lot of cable bandwidth...to run a great deal of data quickly down the pipeline,” said...a product manager at...one of

²¹⁷(Johnson, 1994), pp. 59-60 discusses this and other language from 47 USC. §201(a), §202(a), and §203 (The Communications Act of 1934).

²¹⁸Of course, this analysis depends on the assumption that the telephone company is able to provide telephone circuits to the entrepreneur. While this is true for POTS lines, we saw above that high ISDN pricing in several RBOC territories effectively makes ISDN service unusable for would-be entrepreneurs in those regions.

²¹⁹47 USC §532 (b) (1). In (b) (5) (B), “commercial use” is then defined as “the provision of video programming, whether or not for profit.” Whether Internet access would qualify as “commercial use” under this definition is obviously subject to interpretation. Would a Star Trek preview available at Paramount’s World Wide Web site be considered video programming?

²²⁰See (Johnson, 1994), pp. 60-63, for further explanation of this point. Essentially, the law as interpreted by the FCC allows the cable operator complete discretion in setting the price for these channels.

²²¹This problem should lessen over time as cable operators add more channel capacity to their systems. However, note that the new channels are generally downstream-only; upstream channels remain scarce.

the nation's largest cable system operators. "You could use that more profitably for a video channel."²²²

In contrast, the Internet provider getting started with the telephone company starts with a much more modest and familiar request: supply me with a few telephone circuits, just like any other small business. By the time the entrepreneur needs special help from the telephone company, her business is a going concern and in a better negotiating position.²²³

It should not be surprising that a cable industry manager would think as indicated in the quote above. Video is his traditional business; Internet is an unfamiliar innovation.²²⁴ Compared to the \$20 billion cable industry (and the \$80 billion telephone industry), Internet access, at most a \$0.5 billion industry, appears puny.²²⁵ However, as the Internet gets ever more press and continues to grow at 10% per month—tripling in size each year—it becomes harder to ignore.

Summary

With growing awareness of the value of the Internet access business, more infrastructure providers are likely to consider offering Internet connections. Given the economic value demonstrated in Chapter 7 of co-locating the Internet Point of Presence with a telephone switch, telephone companies have a strong incentive to offer Internet service themselves, as Ameritech is doing. Between that incentive, the lack of common carriage in cable, and the technical problems with the cable plant and modems, Internet over ISDN appears poised to grow more quickly than Internet over cable, despite the

²²²From November 30, 1994 Bloomberg newswire, reporting on the Western Cable Show in Anaheim, CA.

²²³For example, the growing needs of Software Tool and Die, Inc., a provider of dialup Internet service in Brookline, MA that currently has 10,000 subscribers, induced NYNEX to install fiber optic cable along its service route. This expensive installation, however, was not required when the company commenced operations. Conversation with Chetty Ramanathan of Software Tool and Die, Inc., December 9, 1994.

²²⁴This problem is exacerbated by the fact that the complete suite of technologies needed to offer Internet over cable is drawn from a number of different technical communities. Cable TV engineers have different technical specialties than computer networking folks, and the two communities don't overlap much. Thus when system-level problems arise with broadband data networking, they cannot be resolved by any individual engineer, and there is a lot of potential for finger-pointing. A cable operator or equipment vendor that has been involved in private broadband LANs in addition to residential TV infrastructure would have a big advantage.

²²⁵(Flynn, 1995) estimates the Internet access business at \$135 million based on data from Goldman Sachs. The Maloff Company, a Michigan consultancy focusing on business and the Internet, estimates that the Internet access market grew from \$119 to \$521 million from March 1994 to January 1995, not including access provided by the major commercial on-line services. This data is derived from summaries of Maloff's "Internet Service Provider Marketplace Analysis" (1993-4 and 1994-5) which are available on the Internet at <http://www.trinet.com/maloff>.

many potential advantages to the cable-based approach identified in this report.

8.3. Policy recommendations

The analysis above has highlighted immature cable modem technology, lack of access to cable channels for data services, and ISDN pricing as barriers to diffusion of higher-speed residential Internet access. This section recommends policies to help reduce these barriers.

8.3.1. Allow open subscriber access to cable networks

Since the early 1980's, telephony customers have been able to purchase their own subscriber equipment and attach it to the telephone network. This "open subscriber access" policy has resulted in a wide variety of benefits, including broader customer choice and lower equipment prices.²²⁶ One particular benefit has been the rapid pace of innovation and price reduction for POTS modems. This development has in turn been a critical driver of the success of ordinary dialup access to the Internet as well as commercial on-line service businesses.

Cable subscriber equipment markets still follow the earlier telephony model of closed subscriber access, in which cable operators buy large batches of equipment from a small number of vendors, then lease this equipment to subscribers over the long term. As discussed above, large purchase blocks create frictions in the marketplace. It is not a coincidence that today there is a much broader array of computer networking equipment sold for ISDN than for cable.

The cable industry's arguments against open subscriber access include concerns about unauthorized use of services and protection of the network in the face of malicious, malfunctioning, or noisy user equipment.²²⁷ Although the technological details are different, these arguments are similar in spirit to those that were made by the telephone industry when it fought against open subscriber access. In the case of telephony, the predicted problems did not come to pass. Clearly, mechanisms other than provider ownership of subscriber equipment—such as standardization and FCC equipment certification—can successfully address network safety concerns.²²⁸ Given this history, it is difficult to believe that innovative solutions could not be found

²²⁶See (National Telecommunications and Information Administration, 1991).

²²⁷(Large, 1995), pp. 62-3, provides a summary of such arguments. In several of the cases he describes, problems arise because a technology is assumed to be "locked in." One feasible approach to solving these problems is to perform the necessary functions in software instead of hardware.

²²⁸Professor David Tennenhouse points out that the electricity network provides open subscriber access even though network safety is a much bigger issue for it than for either the cable or telephone networks. Also, the equipment ultimately ends up in the custody of the subscriber regardless of who owns it.

to the technical problems the cable industry identifies with open subscriber access.

Another cable industry argument revolves around the operator's need to control subscriber equipment so that it can be upgraded frequently enough. This argument makes particularly little sense when applied to data services, whose subscribers would be owners of computer equipment. Such consumers are accustomed to a wide choice of equipment and a much more rapid pace of upgrades than they perceive as typical of cable set-top equipment. In addition, some fraction of data-over-cable subscribers would be using the service for business purposes. Instead of an individual consumer, a small business or a telecommuter's employer would be paying for the equipment and choosing to upgrade it. In this context, there is no reason to expect the customer to lag behind the cable operator in frequency of upgrades.

In summary, open subscriber access would benefit the cable industry in the same ways it turned out to benefit telephony. As *Business Week* observed:

In the end, cable's proprietary stance is not even in its own interest: Pursuing an open system would unleash a flood of innovative products and boost the traffic on the Information Highway. As one of its primary toll-takers, cable can only benefit. (Landler and Coy, 1994)

Data-over-cable services could be one of those innovative products. Open subscriber access would accelerate the technological development of cable modems. It would also lower the business risk for the provider of an Internet over cable service, by distributing a large portion of the necessary capital cost to the service's subscribers.

8.3.2. Improve access to cable channels for non-video services

Cable operators are not common carriers; through discriminatory pricing, they exert tight control over what content is transmitted through their channels. The analysis above has shown how this situation hinders the development of cable networks as Internet access media. Both Pool and Johnson have argued that application of common carrier status to the already-developed cable industry is an impractical goal.²²⁹ Federal cable law does, however, provide two alternative forms of independent access to channels: "commercial use" and "public, educational, or governmental use," or so-called PEG access.²³⁰ Both of these forms are limited to small numbers of channels and may not be available in all areas, since they are subject to state and local pre-emption.

²²⁹See (Pool, 1983), pp. 155-88, and (Johnson, 1994), pp. 60-63. Opposition comes both from the courts and the cable industry.

²³⁰PEG and commercial access are described in 47 USC §531 and §532, respectively.

Where these alternative forms of access are available, their use for residential Internet service should be allowed. In some cases, achieving this policy can be a simple matter of interpretation.²³¹

Commercial use

While PEG access would require an Internet service provider to work through cable franchising authorities, “commercial use” access would seem to give the provider the opportunity to lease channels directly from the cable operator. In practice, however, commercial leased cable channels are still much less open to an Internet service provider than telephone lines. First, commercial leased cable channels can be scarce; the number varies from 0 to 15 percent of a system’s total cable channels.²³² Second, cable operators maintain considerable latitude for price discrimination on these channels.²³³ In other words, the use of leased channels still depends on the cable operator’s judgment of the desirability and revenue potential of the service. This poses a barrier to entry, particularly for small new Internet service providers.

The FCC should fix this problem by strengthening its interpretation of the commercial leased access provisions of the 1992 Cable Television Consumer Protection and Competition Act. For example, the FCC should require public disclosure of price schedules and non-discriminatory pricing of commercial leased channel access. This re-interpretation is especially important for data services, since they are unfamiliar to most cable operators and therefore unlikely to be implemented if left entirely to the cable operator’s judgment.

PEG access

Under current law, the cable franchising authority (typically a municipality) has the right to negotiate with the cable operator for the use of channels for public, educational, or governmental (PEG) purposes. The law, however, does not define these terms precisely. Certainly use of the Internet would seem to meet these purposes: it provides public internetworking as well as access to government information repositories and educational institutions (universities as well as K-12 schools), including those in the local community. Thus a franchising authority is within its rights to request use of PEG channels for Internet access.

In practice, PEG channel access takes one of two forms. The municipality may use its negotiating rights to get the cable operator to provide it with a separate

²³¹The current rules governing these forms of access are relatively new, since The Cable Television Consumer Protection and Competition Act of 1992 changed the relevant law.

²³²See 47 USC §532 (b).

²³³See (Johnson, 1994), p. 63. The cable operator sets the prices, subject to the “maximum reasonable rates” the FCC allows them to charge for these channels (47 USC §532 (c)(4)(A)). The commission’s interpretation lets the maximum rate for commercial use channels be the same as the maximum rate charged a programming affiliate.

network (a so-called Institutional Network, or "I-net") linking local public institutions such as schools, libraries, and government offices. Many towns already use data networking over I-nets; the connection of these networks to the Internet is a logical next step.²³⁴

Alternatively, the town may use PEG access for community television broadcasting to residential subscribers.²³⁵ Both franchising authorities and cable operators should be aware that a bi-directional residential cable computer network would provide an alternative form of community broadcasting, possibly at lower cost and with a more diverse array of local content. Once such a network is in place, connecting it to the Internet becomes an additional, possibly commercially-provided option, one that would provide access to an even more diverse array of information.

Summary

In summary, the goal of this recommendation is **not** to prevent cable operators from vertically integrating into Internet service provision. The goal is to make the leased access mechanism a realistic option for both sides. As the National Cable Television Association has commented:

Congress intended to preserve the traditional editorial and packaging functions of cable operators, while crafting a leased access mechanism that could *co-exist* with those functions.²³⁶

As the opening section of the law describing this mechanism states:

The purpose of this section is to assure that the widest possible diversity of information sources are made available to the public from cable systems in a manner consistent with growth and development of cable systems. (47 USC §532 (a))

Strengthening the leased access mechanism is critical to furthering this goal for information sources outside the traditional video domain.

8.3.3. Price ISDN to encourage its adoption

For all of the deregulation that has taken place in the telecommunications industry, the Federal Communications Commission (FCC) and state Public Utility Commissions (PUCs) maintain some measure of control over the pricing of local telephone services. Where possible, these regulatory agencies

²³⁴Presentation by Mark Wheeler of Apple Computer, Inc. to MIT Cambridge Roundtable, March 15, 1995.

²³⁵Some towns operate their own local TV studio. This model has proved more difficult to implement, given the expense of studio equipment and the labor involved in program generation.

²³⁶National Cable Television Association, Comments, MM Docket No. 92-266, at 89 (January 27, 1993), as quoted in (Johnson, 1994), p. 63.

should insist that ISDN service is priced so as not to discriminate against its use among residential customers. This means reducing or eliminating usage charges that are much higher than those for residential POTS. It also means that the FCC's recent ruling increasing the Subscriber Line Charge (SLC) for ISDN service charts a policy course 180 degrees in the wrong direction—assuming the FCC supports the goal of encouraging higher-bandwidth residential access to data services such as the Internet.²³⁷

8.4. Suggestions for further research

This report has highlighted several areas in which further research is warranted. Among them are traffic patterns, protocols, architectures, and standards for data networking over residential cable plant, as well as the cost of achieving reliable data transmission over that plant.

8.4.1. Model extensions

The models described in this report can be extended to investigate additional research issues. For example, operational costs could be incorporated into each model. As discussed earlier in this report, providing technology-intensive Internet service to individual customers will demand significant customer support resources. Does either the cable or ISDN approach offer an advantage in this regard, whether because of greater reliability, or because of easier installation, upgrade, and everyday use?

A parallel demand and pricing model could be developed, based on an investigation of specific user needs and willingness to pay for the different types of Internet access described in this report.

The number of U.S. households with computers can be expected to grow. The cable model can easily be extended to investigate the impact of such growth on Internet-over-cable penetration limits. The dynamism of the computer technology marketplace could be incorporated into the model through revised equipment cost and capability assumptions, and/or numerical or functional estimates of future price declines.

²³⁷From "Bell Atlantic Files Waiver to Prevent ISDN Costs From Skyrocketing", press release, February 17, 1995: "On Jan. 11 [1995], the FCC refused to allow NYNEX to restructure ISDN subscriber line charges (SLCs) to be more competitive with other carriers. The SLC is federally mandated and is set annually...The Commission ruled that separate SLCs should be billed for each ISDN channel; ISDN lines have up to 24 channels...Several other carriers are billing SLCs per ISDN line, not channel...All common carriers must now comply with the FCC's rule interpretation in the NYNEX case unless they obtain a waiver." In other words, the FCC is mandating an increase in the monthly charge for all carrier's ISDN lines, based on the rationale that ISDN service can support more than one simultaneous connection (e.g. one B channel can be used for data transmission while the other is used for a voice or fax call). Of course, if the two connections are "bonded" together into a single 128 Kbps data circuit—a usage model that is not only technically possible but quite likely—then this rationale makes little sense.

At a conceptual level, each model consists of two components: an Internet Point of Presence plus a data network connecting subscribers to it. This report compares two different types of data networks—cable and ISDN—each connected to an Internet Point of Presence component. By changing either model component, numerous other services can be explored. For example:

- Replacing the Internet Point of Presence with an intra-community interconnection scheme would provide a model of an intra-community network. Such a network might allow neighbors to exchange email, or parents to exchange email with teachers, but not allow users to exchange email beyond their local community or find resources on the Internet. The model would allow exploration of the cost vs. functionality tradeoffs of such an approach, in which Internet access would be a separate undertaking.
- The economics of local access to commercial on-line services (such as America OnLine, Prodigy and CompuServe) could be explored by either replacing or augmenting the Internet Point of Presence component with a similar function for one or more commercial on-line services.
- The economics of bundled services could be explored by relaxing the assumption that the Internet provider does not provide content or other services beyond raw access.
- Different data network technologies could be explored, such as: 10 Mbps (or higher) cable; wireless services; dedicated T1 circuits such as those provided by HDSL or ADSL technologies; or switched technologies such as Frame Relay and Asynchronous Transfer Mode.

8.4.2. Protocols and traffic patterns

In Chapter 7 we found that changing assumptions about the traffic pattern had little effect on the system cost, but significant effects on the maximum number of subscribers that could be supported. In order to understand the scalability of Internet-over-cable systems, it is important to gain more knowledge of what actual traffic patterns look like, both in terms of average load and its distribution. In addition, it is clear that technical progress is needed to achieve higher peak bit rates over the 6 MHz channel. Such progress would also have the desirable consequence of raising the penetration limits experienced for any particular level of traffic.

Measurement of actual traffic patterns would also provide input to an important issue that has arisen in cable modem development: should bandwidth be symmetric (i.e. the same in both directions) or asymmetric (i.e. more bandwidth to the user than from him)? What happens to the aggregate bandwidth needs in each direction as different users, running different applications, share the same network? This report has investigated only a

symmetric approach, but asymmetric approaches are also possible and are being pursued by the cable industry.²³⁸

Chapter 7 demonstrated that assumptions about the efficiency of the media access protocol also affect penetration limits. Gaining a better understanding of cable media access protocol performance—whether through analysis or simulation—is essential to an understanding of how far the data over cable solution can take us. Such analysis could also result in improvements to existing cable LAN media access algorithms.

Other improvements might come from further research into architectural changes, such as placement of routers at the neighborhood node sites so that the cable LAN is physically shorter and standard broadband Ethernet equipment can be used.²³⁹

8.4.3. Cable system reliability for data

Earlier in this chapter we saw that lack of cable plant reliability for data transmission currently places a large burden on cable modems. This problem suggests several fruitful areas of inquiry. Building a reliable cable modem requires either a more reliable plant or compensation for the problems of a less reliable plant. To make this tradeoff intelligently, we need to know:

- What is involved in making the plant more reliable and how much would it cost? How much improvement is provided by the upgrade of coax-only systems to hybrid fiber and coax? What are the operational expenses?
- What are the impairments of a less reliable plant? Can subscriber equipment compensate for them, and if so, how? How much does that functionality add to equipment cost? How does the net result change cable's positioning vis-a-vis ISDN?
- What kinds of monitoring and test equipment are needed? To what extent can these functions be integrated into the cable plant itself, e.g. self-monitoring and diagnosing amplifiers? How much do these functions cost?

It is also worth exploring what would be required to develop understanding of these issues at an industry level, given the fragmented ownership and varying technologies in the nation's 11,000 cable systems. For example, cable plant surveys, in the spirit of the "loop surveys" undertaken by the pre-

²³⁸See, for example, (McCoy, 1995).

²³⁹See (Karshmer and Thomas, 1992) for further discussion of this and other ideas. Notice that even with a shorter wire, the folded bus topology can still cause fairness problems with standard Ethernet contention algorithms.

divestiture Bell System,²⁴⁰ could be extremely helpful to the industry. These could provide aggregate data, but protect the confidentiality of individual systems as necessary.

8.5. Summary and conclusion

This section summarizes the key findings of this report and concludes with thoughts on the future of Internet access over cable and ISDN.

- **The shared cable architecture can provide the same average bandwidth and higher peak bandwidths much more economically than the dedicated ISDN architecture.**

See Figure 7.2. 500 Kbps Internet access over cable can provide the same average bandwidth and four times the peak bandwidth of ISDN access for less than half the cost per subscriber. In addition, 4 Mbps cable access can provide the same average bandwidth and thirty-two times the peak bandwidth of ISDN for 40% less cost when targeted at individuals, and 20% more cost per subscriber when targeted at organizations. The economy of the shared bandwidth approach is most evident in Figure 7.3, which compares the per-subscriber cost per bit of peak bandwidth. These costs level out as penetration increases, to about \$0.60 for the 4 Mbps and \$2 for the 500 Kbps cable services, versus close to \$16 for ISDN. (The comparable figure for the 4 Mbps service targeted at individuals is \$0.30.)

- **Cable can provide full-time Internet connections, while ISDN is limited to part-time access.**

Applications built for the Internet's peer-to-peer architecture work more naturally over a full-time Internet connection. For example, a subscriber with only part-time access would have difficulty hosting his own home page on the World Wide Web, or navigating the Web without encountering dialing delays. ISDN's client-only Internet access provides a second-class service compared to cable's full access.

- **Assumptions about traffic loading have very little effect on cost.**

They do, however, affect the maximum penetration levels that can be served with the cable approach.

- **Cable-based Internet access is less risky to deploy at 4 Mbps than at 500 Kbps.**

The 4 Mbps service can deliver higher maximum penetration levels with the same limited upstream channel bandwidth. Because it can share

²⁴⁰See, for example, (Batorsky and Burke, 1984); also see (Chiu, Wu, Hwang and Wei, 1991).

bandwidth among a larger aggregate number of users, service quality should appear less variable. Technological developments in cable modems are expected to raise the bit rates achievable over cable channels higher than today's maximum of 10 Mbps, providing even better service.

- **Most of the cost of either approach is contributed by subscriber equipment.**

See Figure 7.8. This favorable result means that much of the cost of providing residential Internet access can be deferred until demand materializes. This minimizes the chicken-and-egg problem that is so characteristic of local access infrastructure deployments. It is also encouraging because subscriber equipment costs are based on digital technology and can be expected to decline.

- **The business model of "open subscriber access," in which subscribers purchase their own cable interface equipment, would reduce a cable operator's financial risk in offering Internet access.**

This model is especially apropos for the telecommuter and small business markets, in which the cost of equipment is generally borne by employers, not consumers. Open subscriber access would also foster technological development of data-over-cable equipment, including cable modems.

- **The larger scope of the telephone network gives ISDN an advantage over cable at serving small penetration levels economically.**

Potential subscribers can be drawn from a larger geographic area than that served by an isolated municipal cable system. The potential inability of a cable Internet provider to recoup its investment in Internet access facilities at low penetration levels serves as a barrier to entry. The size of this barrier can be reduced if multiple cable systems within a metropolitan area are interconnected, as shown in Figure 3.9.

- **Residential Internet access would be encouraged by laws providing greater access to cable channels by non-video information providers unaffiliated with cable operators.**

Within the confines of existing law under which cable companies are not common carriers, "commercial use" and "public, educational and governmental" access are the next best option for Internet service providers who are unable to interest a cable operator in carrying their service.

Thoughts on the future

Despite the many potential advantages to the cable-based approach identified in this report, Internet over ISDN appears at the time of this writing to be

poised to grow more quickly than Internet over cable. This situation could easily change, however. Contingencies that would affect it include:

- Legislative changes either imposing common carrier obligations on cable operators, or diminishing or enhancing PEG and commercial access to cable channels.²⁴¹
- Strong exercise of municipal authority in cable franchise renewals, specifically directed toward provision of Internet access by the cable operator.
- Acceleration of the pace at which isolated cable systems are interconnected. This outcome could potentially result from the wave of recent cable industry consolidations.
- Highly vocal demand from cable subscribers desiring faster Internet access.
- Successful implementation of the plans announced in the past six months by several of the leading cable companies to offer access to commercial on-line services—including the planned Microsoft Network—and the Internet.²⁴²

In summary, Continental Cablevision's offering of commercial Internet access over the cable plant in Cambridge, Massachusetts is unlikely to remain unique much longer.

²⁴¹Federal legislation proposed May 3, 1995 (HR1555) by Reps. Bliley, Dingell and Fields would continue PEG access to "common carrier video platform" operators.

²⁴²See (Robichaux, 1994) and (Robichaux, 1995).

Bibliography

References

- (1993). "Continental Plans to Connect Its Cable TV System With Internet." The Wall Street Journal August 24, 1993: A1.
- (1994). "PPP smooths path to ISDN." Computerworld 28(31)August 1, 1994: 50.
- (1994). "Some Baby Bells to Offer Internet Access Services: Netcom to Launch ISDN Pilot Program." The Internet Letter September 1, 1994: 1,6-7.
- Angelos, A., P. Bosco, K. Compton, Y. Gu and R. P. Johnson (1993). Critique of ADSL/HDSL for the Local Loop. Putting Telecommunications on the Technology Curve: Architecture and Economics Ed. S. E. Gillett, B. Lampson and D. Tennenhouse. Cambridge, MA, MIT/LCS/RSS-23. 167-173.
- Bailey, J. P. and L. W. McKnight (1995). Internet Economics: What Happens When Constituencies Collide. INET'95, Honolulu, HI, Internet Society.
- Baldwin, T. F. and D. S. McVoy (1988). Cable Communication: Second Edition. Englewood Cliffs, NJ, Prentice Hall.
- Batorsky, D. V. and M. E. Burke (1984). "1980 Bell System Noise Survey of the Loop Plant." AT&T Bell Laboratories Technical Journal 63(5): 775-818.
- Bellcore (1991). National ISDN. Bell Communications Research. Special Report SR-NWT-002006.
- Bellcore (1993). National ISDN-2. Bell Communications Research. Special Report SR-NWT-002120.

- Berger, R. J. (1994). High-Speed Internet Connections. in The Internet Unleashed: Everything you need to master the Internet Ed. S. Bang. Indianapolis, IN, SAMS Publishing (Prentice Hall). 129-154.
- Bertsekas, D. and R. Gallager (1992). Data Networks: Second Edition. Englewood Cliffs, NJ, Prentice Hall.
- Bingham, J. A. C. (1990). "Multicarrier Modulation for Data Transmission: An Idea Whose Time Has Come." IEEE Communications Magazine 28(5): 5-14.
- Black, U. (1989). Data Networks: Concepts, Theory, and Practice. Englewood Cliffs, New Jersey, Prentice Hall.
- Calhoun, G. (1992). Wireless Access and the Local Telephone Network. Norwood, MA, Artech House.
- Carnevale, M. L. and J. J. Keller (1993). "Cable Company is Set to Plug Into Internet." The Wall Street Journal August 24, 1993: B1.
- Chen, W. Y. and D. L. Waring (1994). ADSL Noise Environment and Potential System Performance. Supercomm/ICC, New Orleans, LA.
- Chen, W. Y. and D. L. Waring (1994). "Applicability of ADSL to Support Video Dial Tone in the Copper Loop." IEEE Communications Magazine 32(5): 102-109.
- Chiu, W.-J., M.-J. Wu, W.-K. Hwang and S.-W. Wei (1991). "Loop Survey in the Taiwan Area and Feasibility Study for HDSL." IEEE Journal on Selected Areas in Communications 9(6): 801-9.
- Chow, P. S., J. C. Tu and J. M. Cioffi (1991). A Multichannel Transceiver System for Asymmetric Digital Subscriber Line Service. IEEE Globecom, Phoenix, AZ, IEEE.
- Cioffi, J. (1991). A Multicarrier Primer. ANSI T1E1.4: HDSL and/or ADSL. Submission 91-157 to standards committee.
- Claffy, K. C., G. C. Polyzos and H.-W. Braun (1992). Traffic Characteristics of the T1 NSFNET Backbone. University of California at San Diego. Technical Report UCSD Technical Report CS92-252/SDSC Technical Report GA-A21019.
- Comer, D. E. (1991). Internetworking with TCP/IP, Volume I: Principles, Protocols, and Architecture. Englewood Cliffs, NJ, Prentice Hall.

- Communications Futures Project (1994). Costing New Residential Communications Networks. Australian Bureau of Transport and Communications Economics. Work in Progress Paper No. 5.
- Communications Futures Project (1994). Diffusion of Communications, Entertainment and Information Services. Australian Bureau of Transport and Communications Economics. Attachment 1 to Work in Progress Paper No. 4.
- Communications Futures Project (1994). Networked Communications Services to the Home: Future Demand Scenarios. Australian Bureau of Transport and Communications Economics. Work in Progress Paper No. 4.
- Communications Futures Project (1994). Towards the networked home: the future evolution of residential communications networks in Australia. Australian Bureau of Transport and Communications Economics. Work in Progress Paper No. 6.
- Computer Science and Telecommunications Board of the National Research Council (1994). Realizing the Information Future: The Internet and Beyond. Washington, DC, National Academy Press.
- Cook, A. and J. Stern (1994). "Optical Fiber Access - Perspectives Toward the 21st Century." IEEE Communications Magazine 32(2): 78-86.
- Crocker, D. (1995). To Be "On" the Internet. Internet Engineering Task Force (IETF) Network Working Group. Informational Request for Comments RFC 1775.
- Derfler, F. J., Jr. (1994). "Betting on the Dream." PC Magazine October 25, 1994: 267-287.
- Dukes, S. D. (1992). "Photonics for Cable Television System Design: Migrating to Regional Hubs and Passive Networks." Communications Engineering and Design May 1992: 34-48.
- Duncanson, J. (1994). "Inverse Multiplexing." IEEE Communications Magazine 32(4): 34-41.
- Estrin, D. L. (1982). Data Communications Via Cable Televisions Networks: Technical and Policy Considerations. Massachusetts Institute of Technology thesis.

- Faulkner, D. W., D. B. Payne, J. R. Stern and J. W. Ballance (1989). "Optical Networks for Local Loop Applications." Journal of Lightwave Technology 7(11): 1741-51.
- Federal Communications Commission (1991). Telephone Company-Cable Television Cross-Ownership Rules, Sections 63.54-63.58. Federal Communications Commission. Further Notice of Proposed Rulemaking First Report and Order and Second Further Notice of Inquiry CC Docket No. 87-266.
- Federal Communications Commission (1992). Telephone Company-Cable Television Cross-Ownership Rules, Sections 63.54-63.58. Federal Communications Commission. Second Report and Order, Recommendation to Congress, and Second Further Notice of Proposed Rulemaking CC Docket No. 87-266.
- Flynn, L. (1995). "Internet Server Takes a Big Step." The New York Times February 5, 1995: 10F.
- Fowler, H. J. and W. E. Leland (1991). "Local Area Network Traffic Characteristics, with Implications for Broadband Network Congestion Management." IEEE Journal on Selected Areas in Communications 9(7): 1139 - 1149.
- Gillett, S. E. (1995). Public Policies to Encourage Cable and ISDN Residential Internet Access. INET'95, Honolulu, HI, Internet Society.
- Gillett, S. E., B. Lampson and D. Tennenhouse (1994). Putting Telecommunications on the Technology Curve: Architecture and Economics. MIT Laboratory for Computer Science. Research Seminar Series MIT/LCS/RSS-23.
- Hsing, T. R. and J. W. Lechleider (1991). Digital Signal Processing for High-Speed Digital Transport Technology in the Copper Loop Plant. IEEE Globecom, Phoenix, AZ, IEEE.
- Johnson, L. L. (1994). Toward Competition in Cable Television. Cambridge, MA, The MIT Press.
- Johnson, L. L. and D. P. Reed (1990). Residential Broadband Services by Telephone Companies? Technology, Economics, and Public Policy. Rand. R-3906-MF/RL.
- Jones, D. C. (1991). A New Parallel Adaptive Digital Filter Architecture for High Speed Digital Subscriber Line Application. IEEE Globecom, Phoenix, AZ, IEEE.

- Karshmer, A. I. and J. N. Thomas (1992). "Computer Networking on Cable TV Plants." IEEE Network Magazine 6(6): 32-40.
- Kerpez, K. J. (1991). Forward Error Correction for Asymmetric Digital Subscriber Lines (ADSL). IEEE Globecom, Phoenix, AZ, IEEE.
- Landler, M. and P. Coy (1994). "Why Cable Companies Can't Keep Tying Up The Box." Business Week March 28, 1994: 32.
- Large, D. (1995). "Creating a Network for Interactivity." IEEE Spectrum 32(4)April 1995: 58-63.
- Lechleider, J. W. (1991). "High Bit Rate Digital Subscriber Lines: A Review of HDSL Progress." IEEE Journal on Selected Areas in Communications 9(6): 769-784.
- Lindstrom, A. (1994). "Bells Frustrate Users With ISDN Data Fees." Communications Week April 25, 1994: 1,92.
- Lindstrom, A. (1994). "PSI Offering Internet-Access Services." Communications Week May 23, 1994: 4, 45.
- Lindstrom, A. H. (1992). "HDSL: A simple solution that offers complex choices." Supplement to Telephony: Transmission Special (October 5, 1992): 34-38.
- Lu, K. W., M. I. Eiger and H. L. Lemberg (1990). "System and Cost Analyses of Broad-Band Fiber Loop Architectures." IEEE Journal on Selected Areas in Communications 8(6): 1058-1067.
- MacKie-Mason, J. K. and H. R. Varian (1993). Some Economics of the Internet. Tenth Michigan Public Utility Conference, Western Michigan University, University of Michigan.
- Markoff, J. (1993). "A Free and Simple Computer Link." The New York Times December 8, 1993: D1.
- Markoff, J. (1994). "Potent PC Surprises Those Betting on Interactive TV." The New York Times May 8, 1994: F1, F8.
- McCoy, C. (1995). "Intel, AT&T Unit, Start-Up Set Pact To Link PCs to Networks via TV Cable." The Wall Street Journal May 5, 1995.
- McNamara, J. E. (1982). Technical Aspects of Data Communication: Second Edition. Bedford, MA, Digital Press.

- Messmer, E. (1994). "Ameritech to launch IP services in 5-state region." Network World 11(12)March 21, 1994: 62.
- Metcalfe, B. (1995). "Get ready for the IPng, the next Internet Protocol: Think of it as Internet 95." Infoworld May 8, 1995: 53.
- Metcalfe, R. M. and D. R. Boggs (1976). "Ethernet: Distributed Packet Switching for Local Computer Networks." Communications of the ACM 19(7): 395-403.
- Mitchell, B. M. (1990). Incremental Cost of Telephone Access and Local Use. Rand. R3909-ICTF.
- Mossberg, W. S. (1994). "Personal Technology: When Shopping for a Modem, Faster Isn't Always Better." The Wall Street Journal September 29, 1994: B1.
- National Telecommunications and Information Administration (1991). The NTIA Infrastructure Report: Telecommunications in the Age of Information. National Telecommunications and Information Administration, U.S. Department of Commerce. NTIA Special Publication 91-26.
- Network Express Inc. (1994). NE ISDN InterHub: THE Central Site Solution for Remote ISDN LAN Access. Network Express, Inc. Product Literature.
- Noll, A. M. (1991). Introduction to Telephones & Telephone Systems: Second Edition. Boston, MA, Artech House.
- Northern Telecom Inc. and Amati Communications Corp. (1993). ADSL and Discrete Multitone (DMT). Amati Communications Corp.
- Pearl, D. (1994). "Debate Over Universal Access Rights Will Shape Rules Governing the Future of Communications." The Wall Street Journal January 14, 1994: A12.
- Perlman, R. (1992). Interconnections: Bridges and Routers. Reading, MA, Addison-Wesley.
- Pool, I. d. S. (1983). Technologies of Freedom. Cambridge, MA, Harvard University Press.
- Reed, D. P. (1992). Residential Fiber Optic Networks: An Engineering and Economic Analysis. Norwood, MA, Artech House.

- Reed, D. P. (1993). The Prospects for Competition in the Subscriber Loop: The Fiber-to-the-Neighborhood Approach. Twenty-First Annual Telecommunications Policy Research Conference, Solomons, Maryland, Telecommunications Policy Research Conference.
- Rickard, J. (1994). "From Windows to the World Via Internet - Not for the Faint of Heart." Boardwatch Magazine: Guide to Online Information Services and Electronic Bulletin Boards March 1994: 57-63.
- Robertson, B. (1994). "Choose an ISDN Flavor: NIC or Digital Modem." Network Computing September 1, 1994: 174-176.
- Robichaux, M. (1994). "Comcast Plans to Send On-Line Services To PCs Over Its Vast Cable-TV Network." The Wall Street Journal December 7, 1994: B4.
- Robichaux, M. (1995). "TCI, W.R. Hearst Set Venture to Provide On-Line Computer Services via TV Cable." The Wall Street Journal May 4, 1995: B7.
- Romkey, J. (1988). A Nonstandard for Transmission of IP Datagrams Over Serial Lines: SLIP. Internet Engineering Task Force (IETF) Network Working Group. Request for Comments RFC 1055.
- Rutkowski, A. M. (1994). "A Year in the Life of the Internet." Internet Society News 2(4)Winter 1994: 6-11.
- Rutkowski, T. (1993). "The Internet Is Its Own Revolution." Internet Society News 2(3)Autumn 1993: 2-4.
- Sandberg, J. (1994). "Cable Service Tying PCs to Internet to be Unveiled." The Wall Street Journal March 8, 1994: B10.
- Shoch, J. F. and J. A. Hupp (1980). "Measured Performance of an Ethernet Local Network." Communications of the ACM 23(12): 711-721.
- Simpson, W. (1993). The Point-to-Point Protocol (PPP). Internet Engineering Task Force (IETF) Network Working Group. Request for Comments RFC 1661.
- Sirbu, M., D. Reed and F. Ferrante (1989). "An Engineering, Economic and Policy Analysis of Residential Broadband Networks." Journal of Lightwave Technology (November 1989): 1876-1884.

- Sistanizadeh, K. (1991). Spectral Compatibility of Asymmetrical Digital Subscriber Lines (ADSL) with Basic Rate DSLs, HDSLs, and T1 Lines. IEEE Globecom, Phoenix, AZ, IEEE.
- Sklower, K., B. Lloyd, G. McGregor and D. Carr (1994). The PPP Multilink Protocol (MP). Internet Engineering Task Force (IETF) Network Working Group. Request for Comments RFC 1717.
- Stallings, W. (1987). Local Networks (Second Edition). New York, Macmillan.
- Stix, G. (1993). "Domesticating Cyberspace." Scientific American August 1993: 100-110.
- Sutherland, J. and L. Litteral (1992). "Residential Video Services." IEEE Communications Magazine 30(7): 36-41.
- Tanenbaum, A. S. (1988). Computer Networks. Englewood Cliffs, NJ, Prentice Hall.
- Tennenhouse, D., B. Lampson, S. E. Gillett and J. S. Klein (1995). Virtual Infrastructure: Putting Information Infrastructure on the Technology Curve. MIT Laboratory for Computer Science. Technical Memorandum MIT/LCS/TM-523.
- Time Warner Cable Inc. (1993). Development of a Full Service Network. Time Warner Cable, Denver.
- Wagner, M. (1995). "Internet Providers Flock to ISDN." Open Systems Today January 23, 1995: 1, 50.
- Walkoe, W. and T. J. J. Starr (1991). "High Bit Rate Digital Subscriber Line: A Copper Bridge to the Network of the Future." IEEE Journal on Selected Areas in Communications 9(6): 765-768.
- Waring, D. L. (1991). The Asymmetrical Digital Subscriber Line (ADSL): A New Transport Technology for Delivering Wideband Capabilities to the Residence. IEEE GLOBECOM '91, Phoenix, AZ, IEEE.
- Waring, D. L., J. W. Lechleider and T. R. Hsing (1991). "Digital Subscriber Line Technology Facilitates a Graceful Transition from Copper to Fiber." IEEE Communications Magazine 29(3): 96-108.
- Werner, J. J. (1991). "The HDSL Environment." IEEE Journal on Selected Areas in Communications 9(6): 785-800.

White, P. E. (1993). "The Changing Role of Switching Systems in the Telecommunications Network." IEEE Communications Magazine 31(1): 12-13.

Zenith Communication Products Inc. (1994). Z-LAN Protocol.

Ziegler, B. (1995). "Share of Homes with PCs Rises to 31% in Poll." The Wall Street Journal February 6, 1995: B5.

Background reading

(1991). "A Short Course in Cable." In The Broadcasting Yearbook 1991. Washington, D.C.: D-3.

(1991). The World's Telephones: A Statistical Compilation as of January 1990. International Marketing, Morristown, NJ.

(1992). "The Top 20 Market Cable Universe Comparison." Cablevision (April 5, 1992) 42-46.

(1993). "Cable Company Plans a Data Superhighway." The New York Times April 12, 1993: D1.

(1994). "Ascend Pipeline 50 breaks \$1500 barrier for remote office routers." Infoworld 16(14) April 4, 1994: 43.

(1994). "Boardwatch 100 Readers' Choice Contest - Preliminary Results." Boardwatch Magazine: Guide to Online Information Services and Electronic Bulletin Boards March 1994: 39-41.

(1994). "Switching Hubs Deserve Consideration." Communications Week May 9, 1994: 52.

(1995). "ISDN: Back to the future." The Economist January 7, 1995: 54-55.

Akgun, M. B. (1980). "Comparison of Technology and Capital Costs for New Home Services." IEEE Transactions on Cable Television CATV-5(3): 124-138.

Andrews, F. T. (1991). "The Evolution of Digital Loop Carrier." IEEE Communications Magazine 29(3): 31-35.

- Axner, D. H. (1992). "Why 10BaseT Has Been So Successful." Business Communications Review 22(3): 43-45.
- Bailey, J., P. Harlee, S. Matthews and A. Tear (1992). Incorporating New Ideas Into the Residential Loop. Carnegie Mellon University.
- Baker, B., M. Horan, R. Solomon, R. Neuman and E. H. Prodromou (1992). Competition in the Local Loop. M.I.T. Communications Forum, Cambridge, MA,
- Baran, P. (1993). The Universal Communications System of the Future: Telephone or TV Cable? NCTA Cable '93, San Francisco, CA,
- Bell, C. G. (1991). High-Tech Ventures: The Guide for Entrepreneurial Success. Reading, Massachusetts, Addison-Wesley.
- Bellcore (1991). Generic Requirements for High-Bit-Rate Digital Subscriber Lines. Bell Communications Research. Technical Advisory TA-NWT-001210.
- Carnevale, M. L. (1993). "Begun as a Government Experiment, What is Now Called the Internet Has Taken on a Life of Its Own." The Wall Street Journal November 15, 1993: R7.
- Carnevale, M. L. (1994). "U.S. Court Decision Could Slow Pace of Competition in Local Phone Service." The Wall Street Journal June 13, 1994: B3.
- Cerf, V. G. (1991). "Networks." Scientific American 265(3): 72-81.
- Cho, Y.-S. and E. F. Carr (1982). "Application of Digital Pair-Gain System SLC(TM)-96, in Developing Countries." IEEE Transactions on Communications COM-30(9): 2016-2021.
- Claffy, E., Braun, H. and Polyznous, G. (1993). "Measurement Considerations for Assessing Unidirectional Latencies." Internetworking Research and Experience.
- Coathup, L. and M. McNabb (1991). Fiber-to-the-Curb Narrowband and Broadband Assessment. Telecommunication Access Networks: Technology and Service Trends (Proceedings of the IX International Symposium on Subscriber Loops and Services: ISSLS '91), Elsevier Science Publishers B.V. (North-Holland), Amsterdam, Netherlands.
- Comer, D. E. and D. L. Stevens (1993). Internetworking with TCP/IP, Volume III: Client-Server Programming and Applications (BSD Socket Version). Englewood Cliffs, New Jersey, Prentice Hall.

- Copeland, J. A. (1989). "Data Compression Techniques for PC Communications." IEEE Journal on Selected Areas in Communications 7(2): 246-8.
- Cruz, G. C., R. S. Hisiger and R. S. Wolff (1989). "Strategic Telecommunications Network Planning in the Context of Emerging Technologies, Architectures and Services." IEEE Journal on Selected Areas in Communications 7(8): 1198-1206.
- Dertouzos, M. L. (1991). "Communications, Computers and Networks." Scientific American 265(3): 62-69.
- DeSimone, A. (1991). "Generating Burstiness in Networks: A Simulation Study of Correlation Effects in Networks of Queues." Computer Communication Review 21(1): 24-31.
- Dunn, D. A. and M. G. Johnson (1992). "Cable-Telco Competition." Telematics and Informatics 9(2): 73-86.
- Dutcher, W. (1994). Connecting a LAN to the Internet. The Internet Unleashed: Everything you need to master the Internet Ed. S. Bang. Indianapolis, IN, SAMS Publishing (Prentice Hall). 112-128.
- Ebert, I. and P. Richards (1993). "Technology's Role in Switching Systems Evolution." IEEE Communications Magazine 31(1): 26-32.
- Egan, B. (1991). Information Superhighways: The Economics of Advanced Public Communication Networks. Norwood, MA, Artech House.
- Eigen, D. J. (1990). "Narrowband and Broadband ISDN CPE Directions." IEEE Communications Magazine 28(4): 39-46.
- Eiger, M. and C. Harrington (1991). Present Value of Expenditures for Phased Fiber in the Loop Deployment. IEEE GLOBECOM '91, Phoenix, AZ, IEEE.
- Electronic Frontier Foundation (1992). The Open Platform: A Proposal by the Electronic Frontier Foundation for a National Telecommunications Infrastructure. The Electronic Frontier Foundation, Washington D.C.
- Estrada, S. (1993). Connecting to the Internet. Sebastopol, CA, O'Reilly & Associates, Inc.
- Ethem, M. (1992). "Migrating from Ethernet to FDDI: A Case Study." Business Communications Review 22(3): 38-42.

- Faulkner, D. W., D. B. Payne and J. R. Stern (1988). Passive Optical Telephony Networks and Broadband Evolution. IEEE GLOBECOM '88, Hollywood, FL, IEEE.
- Feder, B. (1992). "Optical Fiber (Almost) at Home." New York Times March 4, 1992.
- Ferguson, C. H. and C. R. Morris (1994). Computer Wars: The Fall of IBM and the Future of Global Technology. New York, Times Books.
- Fischer, W., O. Fundneider, E.-H. Goeidner and D. A. Lutz (1991). "A Scalable ATM Switching System Architecture." IEEE Journal on Selected Areas in Communications 9(8): 1299-1307.
- Fox, J. R. and S. T. Jewell (1990). "A Broad-Band Distributed-Star Network Using Subcarrier Fiber Transmission." IEEE Journal on Selected Areas in Communications 8(7): 1223-8.
- Fraser, A. G. (1983). "Towards a Universal Data Transport System." IEEE Journal on Selected Areas in Communications SAC-1(5): 803-16.
- Fraser, A. G. and P. S. Henry (1992). "Transmission Facilities for Computer Communication." Computer Communication Review 22(5): 53-61.
- Fraser, A. G. and W. T. Marshall (1989). "Data Transport in a Byte Stream Network." IEEE Journal on Selected Areas in Communications 7(7): 1020-33.
- Fuller, L. K. (1994). Community Television in the United States: A Sourcebook on Public, Educational, and Governmental Access. Westport, CT, Greenwood Press.
- Gabel, D. (1991). "An Application of Stand-alone Costs to the Telecommunications Industry." Telecommunications Policy 15(1): 75-84.
- Gall, D. L. (1991). "MPEG: A Video Compression Standard for Multimedia Applications." Communications of the ACM 34(4): 47-58.
- Gasbarrone, G., J. Swenson, J. Fanti and C. Gianfranco (1990). Metropolitan Area Networks: A Suitable Step Towards Broadband-ISDN. International Conference on Computer Communication, New Delhi, India, Springer-Verlag.

- Grover, W. D., W. A. Krzymien and A. Shen (1991). "Simulation and Experimental Studies on the Concept of a Rate-Adaptive Digital Subscriber Loop." IEEE Journal on Selected Areas in Communications 9(6): 941-950.
- Habib, I. W. and T. N. Saadawi (1992). "Multimedia Traffic Characteristics in Broadband Networks." IEEE Communications Magazine 30(7): 48-54.
- Hafez, H. M., A. K. Elhakeem and S. A. Mahmoud (1985). "Spread Spectrum Access in Two-Hop CATV Data Networks." IEEE Journal on Selected Areas in Communications SAC-3(2): 312-321.
- Hamilton, C. M. and L. M. McGlynn (1988). Preliminary Reliability Considerations for Fiber Optic Telecommunications Systems to the Home. IEEE GLOBECOM '88, Hollywood, FL, IEEE.
- Hampson, P. H. and J. M. Anderson (1990). Telecommunications Network Planning and Modernization for a Rural Environment. Second International Conference on Rural Telecommunications.
- Hart, J. A., R. R. Reed and F. Bar (1992). "The building of the Internet: Implications for the future of broadband networks." Telecommunications Policy November 1992: 666-689.
- Hatamian, M. and E. G. Bowen (1985). "Homenet: A Broadband Voice/data/video Network on CATV Systems." AT&T Technical Journal 64(2): 347-367.
- Hawley, G. T. (1991). "Historical Perspectives on the U.S. Telephone Loop." IEEE Communications Magazine 29(3)24-28.
- Heldman, R. K. (1988). ISDN in the information marketplace. Blue Ridge Summit, PA, Tab Professional and Reference Books.
- Hill, A. M., D. B. Payne, K. J. Blythe, D. S. Forrester, A. Silvertown, J. W. Arkwright, D. W. Faulkner and J. W. Ballance (1990). An Experimental Broadband and Telephony Passive Optical Network. IEEE GLOBECOM '90, San Diego, CA, IEEE.
- Hsing, T. R., C.-T. Chen and J. A. Bellisio (1993). "Video Communications and Services in the Copper Loop." IEEE Communications Magazine 31(1): 62-68.
- Hura (1992). "Computer Networks (LAN, MAN and WAN): An Overview." In Computer Engineering Handbook, Chapter 20.

- Irvin, D. R. (1993). "Making Broadband-ISDN Successful." IEEE Network Magazine 7(1): 40-45.
- Ishikawa, H. (1992). "Evolving from Narrowband." IEEE Communications Magazine 30(8): 32-36.
- Itoh, A., W. Takahashi, H. Nagano, M. Kurisaka and S. Iwasaki (1991). "Practical Implementation and Packaging Technologies for Large ATM Switching Systems." IEEE Journal on Selected Areas in Communications 9(8): 1280-88.
- Jayant, N. (1992). "Signal Compression: Technology Targets and Research Directions." IEEE Journal on Selected Areas in Communications 10(5): 796-818.
- Johnson, C. (1990). "Managing an Ethernet Installation: Case Studies from the Front Lines." Computer Communication Review 20(5): 8-11.
- Johnson, L. L. (1992). Telephone Company Entry into Cable Television: Competition, Regulation, and Public Policy. Rand.
- Kapor, M. and J. Berman (1992). Building the open road: the NREN as a test-bed for the National Public Network. Cambridge, MA, Electronic Frontier Foundation.
- Karpinski, R. (1992). "FITL market comes into focus." Telephony (Transmission Supplement) October 5, 1992: 42-48.
- Karpinski, R. (1992). "The Top 10 Reasons CATV Won't Overrun LEC Franchises Overnight." Telephony May 25, 1992: 7,16.
- Karpinski, R. (1992). "US West to consider combining fiber and coax in the loop." Telephony September 1992:
- Karshmer, A., J. Phelan and J. Thomas (1988). TVNet: An Image and Data Delivery System using Cable TV Facilities. Elsevier Science Publishers. 135-151.
- Kerpez, K. J. (1992). The Range of Baseband ADSLs as a Function of Bit Rate. IEEE Globecom, Orlando, FL, IEEE Service Center.
- Klein, W. (1992). "Operations and Management of a Campus Packet Network." IEEE Network Magazine 5(2): 16-22.
- Kleinrock, L. (1992). Technology Issues in the Design of NREN. Building Information Infrastructure Ed. B. Kahin. 116-140.

- Kovarik, K. D. and P. Maveddat (1994). "Multi-Rate ISDN." IEEE Communications Magazine 32(4): 48-54.
- Lampson, B. (1983). Hints for Computer System Design. 9th ACM Symposium on Operating Systems Principles, Bretton Woods, NH, Operating Systems Review vol. 17 no. 5, 33-48.
- Lattner, P. D., R. L. Fike and G. A. Nelson (1991). "Business and Residential Services for the Evolving Subscriber Loop." IEEE Communications Magazine 29(3): 109-14.
- Leslie, I. M., D. R. McAuley and D. L. Tennenhouse (1993). ATM Everywhere? University of Cambridge Computer Laboratory and MIT Laboratory for Computer Science.
- Lidinsky, W.P. (1991). "Network Requirements for Data Traffic." In Perspectives on Packetized Voice and Data Communications Eds. W.P. Lidinsky and D. Vlack. New York, IEEE Press.
- Limb, J. O. and L. E. Flamm (1983). "A Distributed Local Area Network Packet Protocol for Combined Voice and Data Transmission." IEEE Journal on Selected Areas in Communications SAC-1(5): 926-934.
- Lindstrom, A. H. (1992). "AT&T Paradyne Expands Copper Pipeline." Telephony 222(22): 17(2).
- Lottor, M. (1992). Internet Growth (1981-1991). SRI International, Network Information Systems Center. Request for Comments 1296.
- Lutz, C. (1992). "DCSs lead the way to SONET." Telephony (Transmission Supplement) October 5, 1992: 48-50.
- MacKie-Mason, J. K. (1994). Pricing Congestible Network Resources. University of Michigan.
- MacKie-Mason, J. K. and H. Varian (1994). "Economic FAQs About the Internet." Journal of Economic Perspectives 8(3): 75-96.
- MacKie-Mason, J. K. and H. R. Varian (1993). Pricing the Internet. Public Access to the Internet, JFK School of Government, Harvard University, University of Michigan.
- Malley, D. P. and O. K. Tonguz (1992). "Fiber in the Loop: Where and When is it Feasible?" IEEE Journal on Selected Areas in Communications 10(9): 1523-1544.

- Maxemchuck, N. F. and A. Netravali (1985). "Voice and Data on a CATV Network." IEEE Journal on Selected Areas in Communications SAC-3(2): 300-311.
- Mazzei, U., C. Mazzetti and G. Roso (1988). Economics of Digital Carriers and Fiber Optic Systems in Subscriber Loop. ISSLS '88 (International Symposium on Subscriber Loops and Services), Boston, MA, IEEE.
- McGarty, T. (1983). "Cable Based Metropolitan Area Networks." IEEE Journal on Selected Areas in Communications SAC-1(5): 816-831.
- McQuillan, J. (1991). "Understanding Network Innovation." Business Communications Review 21(4)April 1991: 14.
- McQuillan, J. (1992). "The Killer Application for ATM." Business Communications Review 22(9): 12.
- Miller, M. D. (1994). "A scenario for the deployment of interactive multimedia cable television systems in the United States in the 1990's." Proceedings of the IEEE 82(4):
- Moondra, S. L. (1989). "Impact of Emerging Switching-Transmission Cost Tradeoffs on Future Telecommunication Networks." IEEE Journal on Selected Areas in Communications 7(8): 1207-1218.
- Mudd, D. R. and M. Z. Starkey (1992). "Economic implications of collocation for the exchange access market." Telecommunications Policy 16(6): 511-526.
- Musser, J. M., T. T. Liu, L. Li and G. J. Boggs (1983). "A Local Area Network as a Telephone Local Subscriber Loop." IEEE Journal on Selected Areas in Communications SAC-1(6): 1046-1054.
- Niedzwiecki, A., R. T. Matsunaga and P. V. Coverdale (1989). "All-Digital Network Transmission Planning and Terminal Requirements." IEEE Journal on Selected Areas in Communications 7(8): 1166-71.
- Nutt, G. J. and D. L. Bayer (1982). "Performance of CSMA/CD Networks Under Combined Voice and Data Loads." IEEE Transactions on Communications COM-30(1, pt. 1): 6-11.
- Okada, K. and H. Shinohara (1992). "Fiber Optic Subscriber Systems." IEEE LTS 3(4):6-11.

- Pappalardo, D. (1995). "ISDN Still Costly, Hard to Get." Communications Week March 13, 1995: 28.
- Personick, S. D. (1993). "The Evolving Role of Telecommunications Switching." IEEE Communications Magazine 31(1): 20-24.
- Rendleman, J. and E. Roberts (1995). "BellSouth Teams Up To Promote ISDN Use." Communications Week April 24, 1995: 1, 118.
- Rockwell, M. (1995). "Prodigy to Offer ISDN-Access Service Package." Communications Week April 24, 1995: 118-119.
- Rowbotham, T. R. (1991). "Local Loop Developments in the U.K." IEEE Communications Magazine 29(3): 50-59.
- Rutkowski, A. M. (1994). "A Year in the Life of the Internet." Internet Society News 2(4)Winter 1994: 6-11.
- Sazegari, S. A. (1992). "The Shape of Competition in the Local Loop." Business Communications Review 22(3): 47-51.
- Schmidt, A. and R. Campbell (1992). Internet Protocol Traffic Analysis with Applications for ATM Switch Design. University of Illinois Urbana-Champaign. Department of Computer Science. Technical Report UIUCDCS-R-92-1735.
- Schroeder, M. D., A. D. Birrell, M. Burrows, H. Murray, R. M. Needham, T. L. Rodeheffer, E. H. Satterthwaite and C. P. Thacker (1991). "Autonet: A High Speed, Self Configuring Local Area Network Using Point-to-Point Links." IEEE Journal on Selected Areas in Communications 9(8): 1318-1335.
- Shumate, P. W. and R. Snelling (1991). "Evolution of Fiber in the Residential Loop Plant." IEEE Communications Magazine 29(3): 68-74.
- Smith, R. (1992). "Cautious—Not Complacent." Telephone Engineer & Management (January 1992).
- Sosnosky, J. and T.-H. Wu (1991). "SONET Ring Applications for Survivable Fiber Loop Networks." IEEE Communications Magazine 29(6): 51-58.
- Srinagesh, P. (1994). Internet Cost Structures and Interconnection Agreements. Telecommunications Policy Research Conference, Solomons, MD, October 1-3, 1994.

- Starr, T. J. J. (1991). "Bridge copper's future to fiber." Telephone Engineer and Management **October 1, 1991**: 49-51.
- Stewart, A. (1992). "CAP-italize on copper." Telephone Engineer & Management **96(19)**: 53-57.
- Sumida, M. and H. Tsuji (1986). "Optical Transmission Modules for Intraoffice Transmission Networks." IEEE Journal on Selected Areas in Communications **SAC-4(9)**: 1468-1473.
- Swinehart, D. C., L. C. Stewart and S. M. Ornstein (1984). Adding Voice to an Office Computer Network. Xerox PARC. CSL-83-8.
- Taylor, T. M. (1991). "Power and Energy in the Local Loop." IEEE Communications Magazine **29(3)**: 76-83.
- Tenzer, G. (1991). "The Introduction of Optical Fiber in the Subscriber Loop in the Telecommunication Networks of DBP TELEKOM." IEEE Communications Magazine **29(3)**: 36-38.
- Terry, J. (1992). "Alternative Technologies and Delivery Systems for Broadband ISDN Access." IEEE Communications Magazine **30(8)**: 58.
- Teske, P. and J. Gebosky (1991). "Local Telecommunications Competitors: Strategy and Policy." Telecommunications Policy **15(5)**: 429-436.
- Testi, N. and J. Naegle (1992). "Tuning Performance on TCP/IP FDDI LANs." Business Communications Review **22(3)**: 16-26.
- Toba, H. and N. Shibata (1992). "Advanced Technology for Fiber Optic Subscriber Systems." IEEE LTS **3(4)**: 12-18.
- Trondoli, A. and G. A. Vago (1988). A Techno-Economical Assessment of a Fiber Optic Local Loop Architecture in Relation to Network Implementation Technology. ISSL '88: International Symposium on Subscriber Loops and Services, Boston, MA, IEEE.
- Tumolillo, A. (1991). Alternative Loop Technologies-Cost Structures and Competitive Analysis. Telecommunication Access Networks: Technology and Service Trends (Proceedings of the IX International Symposium on Subscriber Loops and Services: ISSLS '91), Elsevier Science Publishers B.V. (North-Holland), Amsterdam, Netherlands.
- Tumolillo, A. (1992). "What Do Telcos Know About the Subscriber Loop?" Telephony January 20, 1992: 26-29.

- Verma, D. C. (1991). Guaranteed Performance Communication in High Speed Networks. University of California, Berkeley; Computer Science Division thesis.
- White, P. E. (1991). "The Role of the Broadband Integrated Services Digital Network." IEEE Communications Magazine 29(3): 116-119.
- Whorwood, B. (1986). Traffic Theory. Local Telecommunications 2 : Into the Digital Era Ed. J. M. Griffiths. London, Peter Peregrinus Ltd. on behalf of the Institution of Electrical Engineers (IEE).
- Wilder, C. (1995). "Testing Internet Access: MCI teams with utility for 2-Mbps service." Information Week March 6, 1995: 64.
- Wilson, C. (1992). "ADSL: A High-Bandwidth Hope for Copper." Telephony (Transmission Supplement) October 5, 1992: 40-41.
- Wilson, C. (1992). "Copper or Fiber? Place Your Bets Now." Telephony 223(20): S12(4).
- Wu, T.-H. and M. E. Burrowes (1990). "Feasibility Study of a High-Speed SONET Self-Healing Ring Architecture in Future Interoffice Networks." IEEE Communications Magazine 28(11): 33-51.
- Wu, T.-H., D. J. Kolar and R. H. Cardwell (1989). High-Speed Self-Healing Ring Architectures for Future Interoffice Networks. IEEE GLOBECOM '89, Dallas, TX, IEEE.
- Wu, T.-H. and R. C. Lau (1992). "A Class of Self-Healing Ring Architectures for SONET Network Applications." IEEE Transactions on Communications 40(11): 1746.
- Yamada, H. and S. Sumita (1991). "A Traffic Measurement Method and its Applications for Cell Loss Probability Estimation in ATM Networks." IEEE Journal on Selected Areas in Communications 9(3): 315-324.
- Zachary, G. P. and S. K. Yoder (1993). "Order from Chaos: Computer Industry Divides into Camps of Winners and Losers (Balance of Power is Tilted by Steady Price Declines and Linking of Machines)." Wall Street Journal January 27, 1993: A1 and A4.
- Zervos, N. A. (1991). High-Speed Carrierless Passband Transmission Over the Local Cable Network. IEEE GLOBECOM '91, San Diego, CA, IEEE.