



**Personal Broadband  
Services: DSL and ATM**



# Personal Broadband Services: DSL and ATM

*by Jim Lane*

**Virata**

*Enabling Personal Broadband Solutions*



#### ABOUT THE AUTHOR

Jim Lane is President of TRAC & Co., Inc., a consulting firm in Sagamore, Massachusetts, specializing in telecommunications and data communications. In addition to market research and analysis, product and business planning, Jim provides technical writing services, and is the author of such books as *ATM: Bandwidth for the Future for Telco Systems*, *ATM over SONET: Seamless LAN/WAN integration for integrated Telecom Technology*. Jim may be reached at [trac@capecod.net](mailto:trac@capecod.net).

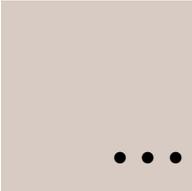


# Table Of Contents

	INTRODUCTION . . . . .	5
1	PERSONAL BROADBAND SERVICES . . . . . <i>A Realistic Scenario</i>	6
2	THE PROBLEMS WITH THE EXISTING NETWORK . . . . . <i>Analog Modems</i> <i>So what's wrong with ISDN?</i>	11
3	BUILDING THE NEW NETWORK . . . . . <i>Connectivity and the Telephone Network</i> <i>Getting Connected</i> <i>The Local Loop</i>	19
4	TAMING COPPER: DSL . . . . . <i>HDSL</i> <i>SDSL</i> <i>IDSL</i> <i>VDSL</i> <i>ADSL</i> <i>RADSL</i>	33
5	HOW THINGS WORK . . . . . <i>Baseband Systems</i> <i>Passband Systems</i> <i>CAP</i> <i>DMT</i>	39
6	LOW SPEED ADSL, G.LITE, AND OTHER LATE STARTERS . . . . . <i>Low Speed DSL and G.Lite</i> <i>Splitterless Implementation</i>	45



7	LAYERING IT ON . . . . .	54
8	ATM FOR THE LINK LAYER . . . . .	61
	<i>The ATM Layer</i>	
	<i>ATM Adaptation</i>	
	<i>The AAL Process</i>	
	<i>ATM'S Advantages</i>	
	<i>Summary</i>	
9	ATM OVER XDSL: PUTTING IT ALL TOGETHER . . . . .	72
	<i>Networking</i>	
10	REALITY CHECK . . . . .	81
11	VIRATA SOLUTIONS . . . . .	84
	<i>Virata's ISOS Solution</i>	
	<i>Application Specific Standard Products—ASSPs</i>	
12	THE NEXT STEPS . . . . .	88
	<i>Learning More</i>	
	<i>Making it Happen</i>	
	<i>Putting ATM to work over DSL</i>	
	GLOSSARY . . . . .	92
	INDEX TO FIGURES . . . . .	96



## Introduction

Welcome to this guide to the impact of a variety of technologies on the emergence of new personal broadband services. Virata is a leading supplier of ATM technology solutions to xDSL equipment vendors. In the following pages we set out on a journey to discover what xDSL technologies are, where they have come from, and the relationship between ATM and xDSL. This work starts shallow but runs deep, providing insights for the ATM and xDSL initiated as well as those new to either technology areas.

As is the case with most technologies there is a new alphabet soup of jargon to master. If you are patient enough to work your way through it you will be rewarded with the valuable ability to use the jargon and actually understand what it means!

Unlike many monographs of this kind, the presentation is not biased in favor of our company's particular product capabilities. We do not apologize for our expression of the enthusiasm we have for the power of ATM and xDSL. For those who are interested we have presented a summary of Virata's products and technology capabilities, carefully segregated in Chapter 11. Virata has assisted a number of leading LAN and xDSL companies in adding ATM technology to their products.

As with any technical exposition some of you will find errors or take exception to explanations of the technology. We welcome and encourage your comments. They will help us to improve. Enough said, let's get on with the task of understanding the power that comes from combining the multi-service delivery capability of ATM with the bandwidth creation of xDSL.

*Virata*

# 1

## Personal Broadband Services



### A Realistic Scenario

For the last fifteen years or so, telephone and cable TV providers have been trying to develop a rational business case for bringing broadband information services into the home. Most of these efforts have centered around video: interactive; high quality; switched digital; on-demand; near on-demand; almost on-demand; don't be so demanding, we'll get it to you, or whatever. Armed with the latest technology capable of dumping a torrent of information into the home, service providers felt that if they built it, the people would come. However, for the most part, they didn't—at least on a paying basis. While vendors were busy developing gigabit systems to allow instant access to thousands of movies, new sociological, demographic, and institutional trends were emerging. And out of these changes has developed the first compelling case for wholesale distribution of broadband services to businesses and homes everywhere.

In building this new business case, we just have to change our definitions a little. Rather than letting the availability of fiber optics and gigabit technology frame the services it is possible to provide, we have to look at what people are doing and how they are working as we approach the next century. When we do this we find that we don't need a gigabit fiber pipe into every home or neighborhood—a range between hundreds of kilobits and a few megabits will do. And we don't need to rebuild the entire local distribution infrastructure to provide the services people need and are interested in—the existing facilities will do for the most part. And as a result, using a revised definition of “broadband” services, we find a realistic business case can be built for supplying these services.

So what are these changes and how do we have to modify our definitions? Well, during the late eighties and early nineties, while service providers were focused on video entertainment, three separate forces started to redefine the environment of their customer base.



First, corporations began a massive program of downsizing, decentralization, and outsourcing. Responsibility and decision-making authority was pushed out from headquarters to divisions and branches scattered across the country. At the same time (as a result of, or to enable, decentralization) corporations were wiring themselves, both internally with LANs (Local Area Networks) and externally, linking their facilities with leased digital facilities like T1s or E1s and ISDN. PCs appeared on every desk, and everyone was gaining real-time access to information anywhere in the corporation.

At the same time, thousands of workers were leaving the conventional work force, in some cases abetted by the wave of downsizings. Work at home, at least on a part time basis has become an increasingly attractive option for many, especially two career families. Many companies are actively encouraging work at home, in some cases incentivized by local, state or national governments. There are approximately 10–12 million full or part time telecommuters today in the U.S. alone. As traffic congestion and air pollution increase, this trend can only grow.

Downsized workers in some cases have become independent contractors with either a home office or a small office nearby instead of an hour's commute away. The changing nature of work itself has made the option of working in a non-corporate setting available to professionals of every type. Hardware and software design for example are no longer tied to a company's lab, but may be performed on today's powerful desktop computers anywhere. Virtual companies are increasingly running with a lean core of full-time staff and outsourcing more and more tasks to independent contractors on an as-required basis. Hardware and software engineering, marketing, strategic planning all lend themselves to this new paradigm that allows these companies to tap the best resource for the particular task no matter where the resource is located. The home office worker population and telecommuting is expected to more than double by the beginning of the next century.

Finally, and perhaps most startlingly, the world discovered the personal computer en masse. For the past 10 years PCs have been pouring into homes by the millions. International Data Corporation estimates PCs had penetrated 36% of U.S. households in 1996, and expects penetration to pass 50% by 2001. What makes this interesting is that PCs are far from inexpensive. Conventional wisdom says an appliance (a VCR for example) has to break the \$500 point before it gains a mass market. Not so with PCs, where a complete system costs at least four times that amount. And of course, the key point is that after people discovered the PC, they fell in love with the Internet, and more recently the World Wide Web. The video buff of the eighties has turned into the Web surfer of the nineties. While the video

moguls were looking the other way, their audience found another toy. And it turns out this toy was quite accessible with kilobit modems, rather than multi-megabit feeds.

So we have three distinct and seemingly permanent trends: streamlined and decentralized corporations, increasing numbers of workers in unconventional locations, and millions of people experiencing the variety and intricacies of the Web on a recreational basis or as part of their professional activities.

The common thread connecting all these trends is naturally the need for communications. But where the telephone has stood us in good stead for a hundred years, today's environment requires data communications. And increasingly in this environment the workhorse of remote data communications, the analog modem, is not up to the task. Workers in remote and branch offices need to access corporate databases with the same speed as inside employees. They need to download complex documents, not just e-mail, and to be effective, they need to do it as quickly as if they were physically attached to the corporate LAN. Work-at-home staff and independent contractors have the same requirements. An engineer working on a new chip design will have to access large files on the engineering server. At analog modem rates, the download time can approach the commuting time avoided by working at home. Even the executive putting in a couple of hours after dinner needs crisp access to complex charts, graphs, tables, and presentations on corporate servers.

And everybody, either an independent business person working from home or remote office or even the recreational net surfer is becoming increasingly frustrated with using the Web. The Web is becoming dominated by complex graphics that can require a minute or more per page to download using an analog modem. New Web technologies including audio and video are on the horizon, promising to make using the Web even more interesting but more painful. The Web is on one hand becoming the principal way we gather information and on the other rapidly outgrowing the technology available to access this information.

So after several false starts, dating from ISDN, the real justification for digital communication services to the small office or home is emerging. The driver is not simultaneous voice and data, video dial tone or access to 500 channels of digital television, but is universal access to commercial and corporate data bases—connectivity to the public Internet and corporate intranets.

And this requirement is qualitatively and quantitatively different from what so many had envisioned as the motive for universal broadband services. Qualitatively, we are not dealing with high definition, full motion video, but for the most part static graphic images. We also are not dealing with time-sensitive information, but rather block transfers of data: a Web page, a CAD file, a sales presentation. We are



also dealing with bursts of information, not continuous streams. A web page is downloaded and then viewed for some time before the next page is accessed. A file is transferred, opened and used perhaps for hours before another file is required.

All of this has a very quantitative effect on bandwidth requirements. While the 28 or 56 kb/s analog modem is becoming inadequate for Internet/intranet connectivity, tens of megabits are not required. For everything but streaming video, a few hundred kilobits will make an enormous difference in file transfers and Web page response. As we approach T1 rates of 1.5 Mb/s or E1 rates of 2 Mb/s, we find the Web or access to corporate servers is as responsive in our remote office or home as it is in our corporate setting. And at these rates, streaming video is accessible for all but TV-quality full screen applications; for that four to six megabits would be required.

So after years of trials designed to bring tens or hundreds of megabits of data to every doorstep, the real application for broadband service has come out of hiding, and it's not so broad: ubiquitous access to public and private data networks at moderate bit rates. But unlike the "field of dreams" technology push of the video trials, this application is here, growing and pulling—hard! It is driven by the increasing desire of people to work in unconventional ways and places as effectively as they used to work in a structured corporate setting. Their need is complemented by the increasing numbers logging onto the Internet every night (15-20% of homes in the U.S. are on-line today) and the declining numbers clicking their cable TV remote controls. And it's demonstrated by the willingness of people to continuously replace modems they just purchased with incrementally faster models to extract the last bit of performance they can from their dial-up phone connections, order second or third phone lines for dedicated modem use (the number of homes with multiple lines in the U.S. has increased by a factor of eight since 1990), and in growing numbers, upgrade to ISDN service ... if they can get it.

In this context, "broadband" becomes something in the order of a T1 or E1, 1.5–2 Mb/s, available to every home, small or branch office, or school that wants one. We might call this new definition "personal broadband services" to complement our personal computers. While not the hundreds of megabits usually associated with broadband services, 1.5 Mb/s is 50 or more times as fast as most dial-up connections, and 25 times better than most dedicated lines used to tie branch offices into corporate intranets. On a personal level, this is broad enough for the next few years.

And just as Pentium or PowerPC chips were not necessary to launch the PC revolution, neither do we need to start our transition from analog modem technology with hundreds of megabits at our finger tips. We can start with our new definition

of personal broadband services in the T1 range, and just as PCs evolved upward in speed and capability, so will our access rates, not pushed by service provider's technology, but pulled as our need to access and share richer, more complex content with fellow employees over corporate intranets or with the world over the Internet grows with time.

But even if we redefine our requirements to much more modest data rates than proposed for previous broadband services, can we realistically expect success? The challenge is to provide nearly universal access to high speed digital services, and more important, provide it at a price point attractive to the consumer and profitable to the service provider. How, after so many unsuccessful attempts are we going to do this? Part of the answer lies in a technology called DSL, or Digital Subscriber Line, part lies in our redefinition of requirements down to a few megabits, and part lies in the ability of DSL to be furnished to individual subscribers on a basis that allows revenue to almost exactly track service provider investment. A brief look at some of the other technologies will illustrate why DSL is on the right track after so many false starts.

# 2

## The Problems with the Existing Network



### Analog Modems

One of the interesting things about the quest of bringing higher bandwidths to users in other than traditional corporate facilities is that while service providers have turned to one technically elegant scheme after another, the venerable analog modem just keeps getting better. At the dawn of the personal computer era, affordable modems chugged along at 300 bps and probably no one in their right mind believed that it would ever be possible (or necessary) to communicate over phone lines at speeds like 9600 bps. Yet by the end of the 80's, 9600 baud hardware was priced within the reach of most consumers, and it seems like rates have doubled almost every year since.\* Can't we just sit back and wait a few more years for modem technology to double its way up to T1 rates?

Unfortunately, we can't. Conventional analog modem technology has about hit the wall at today's current speeds of 28.8 kb/s or 33.6 kb/s. (But as we just said, in the late 70's no one believed we would be using 9600 bps modems, let alone three or four times that rate. Clever engineering has a way of obsoleting all prophecies.) All analog modems must force their data into the 4 kHz wide channel used to carry voice traffic over the telephone network. Rates like 28.8 kb/s obviously require more bandwidth than that, so modern high speed modems use sophisticated carrier modulation techniques and coding to overcome the bandwidth limitations of the telephone channel. These modems send symbols rather than the raw bit stream, each symbol representing some number of consecutive bits in the data stream. For example, V.32 modems take in the raw data four bits at a time, add in a fifth bit for error correction, and from the resultant five bit group generate one of 32 different symbols. Each symbol represents a different combination of amplitude and phase of the transmitted carrier. This process is called quadrature amplitude modulation (QAM), and as we shall see, is similar to schemes used for DSL. Since we generate a symbol every four data bits, the effect is to cut our bandwidth by a factor of four,

*\* Interestingly enough, since 1990, modems have been doubling in rate every 18 months, mimicking Moore's Law for microprocessor performance.*

and a 9600 bps data stream is reduced to a rate of 2400 symbols per second or 2400 baud, and fits into the 4 kHz voice channel. The fact that we add a fifth bit doesn't change this, it just means we have more symbols, and thus a harder time differentiating between them, but we still send them at a 2400 baud rate. Current V.34 modems use highly complex symbol constellations, codings, mappings, adaptive bandwidths and bit rates to wring the last drop of performance from a voice channel.

Figure 2.1  
Modulation  
Constellations

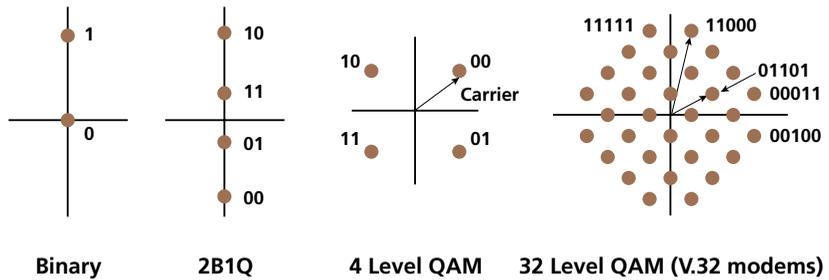


Figure 2.1 shows the constellations displayed by various modulation schemes. The left drawing shows simple binary modulation of a carrier. Here only the amplitude changes from a positive value to zero, and the transmission rate equals the data rate. The 2B1Q technique is a four level amplitude coding used by ISDN and HDSL. We gain a 2:1 efficiency in bandwidth but notice the levels are twice as close together as in simple binary coding. A four level QAM scheme is shown next. Here the amplitude remains constant and the phase of the carrier takes on one of four values, allowing us to transmit two bits of information per phase state. The transmission rate is halved because two bits produce one change in phase state. The final drawing is the 32-QAM process used for V.32 modems. In this case both amplitude and phase of the carrier change as in the transition shown from 11000 to 01101.

The trade-off is rather than sending only one of two possible pieces of information, a one or a zero of the original data stream, we are now sending one of 32 possible states. The receiving modem has to make the correct decision of which symbol was transmitted even though the difference between symbols is much smaller than for a simple one or zero signal. If the voice channel were perfect this would not be a problem, but all circuits in the network experience some form of noise. The twisted pair running from a home or business to the phone company



central office can pick up static from car ignitions, hair dryers, power lines, neon signs or other sources of electrical discharge. The age and quality of the cable also have a great deal to do with how much noise is present. Older, squirrel-chewed cables with significant moisture ingress will have more noise than a new cable. In addition, all the different pairs of wire in a cable, up to several hundred or a few thousand, couple together and leak their signals into adjacent pairs, a phenomenon called crosstalk.

All this interference creates a floor of noise in each channel with spikes popping up from that floor. If we are trying to differentiate between only two states, a one or a zero, the state can be highly corrupted before we make a wrong decision. If the choice is between one of two possible states 32 times closer together, only a little noise is required to corrupt one state and make it look like another. With the current generation of high speed modems, we are working right at the edge of the signal to noise ratio available in most phone lines.

That is, the signal to noise ratio required by theory to decode the symbol states in a 28.8 kb/s modem is what is available from good to very good phone lines. That we are bumping up against practical limits is evidenced by the fact that many 28.8 or 33.6 kb/s modems typically connect at lower rates. If the two modems at each end of the connection sense the signal to noise ratio is less than ideal, they back off to a lower rate that will insure a reliable connection. If everything is ideal; new phone cables, quiet neighborhood close to a central office, modern switching equipment; at both ends of the connection, 33.6 kb/s is possible. In reality, the best many modem connections can achieve is 26.4 kb/s.

But if high speed modem technology has hit the wall at 28.8 kb/s, how can the new 56K modems possibly work? They work by a little bit of sleight of hand that assumes most of the telephone network is digital, and by using a different modulation scheme. In fact, large portions of the network are digital, including switching equipment, interoffice transmission gear, and even portions of some of the local loops that connect subscribers to the central office.

When telephone conversations are digitized, an analog to digital encoder, usually in the central office, samples the conversation 8000 times a second and uses eight bits to encode the value of the sample. At the other end of the connection, the process is reversed, and a close approximation of the original voice is reproduced. However the encoding process introduces yet another type of noise, quantizing noise. This noise results because, when confronted with an analog sample which probably falls between two of the 256 quantization levels of the eight bit encoder, the encoder chooses the closest one. At the other end, the analog level produced is not that which was originally sampled, but the nearest level available to the finite

number of states of the decoder. The result is equivalent to noise in that the output level will be slightly different from the input as if displaced by a small spike of noise. For voice purposes, this slight difference is not noticeable, but for our over-worked high speed modem, it's just one more problem to deal with. If there were no other noise in the system, this quantization noise would set a floor limiting rates to 33.6 kb/s.

56K modems actually turn this quantization process to their advantage. The argument goes like this: quantizing noise is a function only of the encoding process. If we can get rid of the encoding step, we can eliminate the associated noise floor that ultimately limits us to 33.6 kb/s. If we can arrange things so that our digital data passes through only one decoder in the telephone network, the data will be very accurately converted into one of the 256 levels generated by the digital to analog converter. Our modem can be "tuned" to look for these specific levels which are well defined by standards. (Actually only 128 levels are sometimes available in the U.S. because the phone system periodically uses, or "robs" the least significant of the eight bits for channel supervision and signaling. Since it is not always "there" it cannot reliably be used for data communications—this is why we have leased 56 kb/s lines instead of 64 kb/s. In most of the rest of the world, this is not a problem because signaling is carried out of band.) Furthermore if that decoder is the one that drives the subscriber's modem, the only noise contributions introduced will be through the one piece of cable connecting the subscriber to the central office. The result should be a very high signal to noise ratio connection with a good probability of passing data at 56 kb/s.

The sleight of hand in this process is eliminating the encoding process and delivering our digital stream directly to the decoder. This requires that the connection, from the originating data source (an Internet service provider or a corporate access line), through the entire network to the ultimate digital to analog converter be entirely digital. There can be only one digital to analog conversion in the chain. If somewhere along the way, the digital signal is converted to analog, then subsequently re-digitized for the rest of the trip, then converted back to analog for delivery to the subscriber, that intermediate analog to digital conversion introduces quantizing noise and we are back to 33.6 kb/s.

In many cases this limitation of a single conversion is not too severe. For accessing local Internet service providers or a fairly near-by corporate intranet, the chances are high that the connection will be all digital (assuming the ISP or corporation has leased digital facilities to their serving central office). However, the longer the distance, the greater likelihood of a conversion back to analog over a portion of the distance, then back to digital. About 25% of the switching offices in the US are



still equipped with analog switches—the more offices the call passes through, the greater the probability of hitting an analog office. This may not impact the net surfer using a local ISP or the telecommuter, but legions of other independent contractors or workers in branch or remote offices in other states than the parent company may find themselves limited to connectivity in the order of 28.8 kb/s or less.

And even when everything goes right, we have only doubled our throughput. With the spiraling increase in content complexity, even 56 kb/s will not be enough. And unfortunately, barring a major technical breakthrough, 56 kb/s is the end of the road for analog technology. We are totally filling a digital telephone channel with 56 kb/s of data and converting it to analog only at the end of the process for delivery to the user. That channel just will not hold any more digits. If we are going to offer wider band connectivity we will have to look elsewhere than at the traditional telephone circuit.

## So What's Wrong with ISDN?

Actually, nothing, if you can get it. And one of the biggest problems with ISDN even today, 10 or so years after it was introduced, is its hit or miss availability to subscribers. Where it is widely available, Europe, for example, telecommuters and branch offices have been using it quite effectively for several years.

ISDN is a pure digital service, offering two 64 kb/s digital or bearer channels and one 16 kb/s signaling or data channel (2B+D) over one twisted pair for residential service. Notice right off the bat, we have all 64 kb/s available for data applications since the D channel handles the signaling out of band—no robbed bits here. However, ISDN was conceived as a way of providing two phone lines to the residence, one of which could be used for data (or as a second voice line) while the other would normally handle voice calls (an ISDN telephone would provide the analog to digital conversion in the residence). ISDN is not automatically a 128 kb/s service; it is two 64 kb/s circuits. As such, out of the box, it offers only 64 kb/s of bandwidth for our remote office or work at home applications. However, since end-to-end ISDN calls are always routed through digital switches, the circuit is always 64 kb/s of pure data, so a remote office can be equipped with ISDN and be assured of that bandwidth back to the parent facility, unlike the problems 56k modem technology can experience. Note that if you have ISDN service you can make a voice call to any telephone, anywhere, even if the phone you are calling is served by an analog switch; the network will take care of the digital to analog conversion at the appropriate point. But if the number you are calling is connected to an analog switch you cannot place a 64 kb/s digital data call. The voice D/A converters in the

network are not modems and would have no idea what to do with your data. To transmit data to a destination served by analog equipment you would have to connect a conventional analog modem to the analog port of your ISDN terminal adapter or ISDN telephone.

If we want to take advantage of the full 128 kb/s of the ISDN line, we have to buy special terminal adapters that “bond” the two B channels together, sending half our data down one channel and half down the other. A similar unit at the far end sorts everything out and insures that the data is reassembled in the proper order. The key word here is “similar.” Not every one has such bonding terminal adapters, especially local Internet service providers, so the service is hardly universal. Furthermore, many vendors of ISDN equipment have added their own proprietary twists to the basic B channel bonding standards, making interoperability somewhat problematic.

However, ISDN suffers from the same problems as some other technologies we will look at: the fork lift upgrade syndrome. ISDN is not a technology that can be applied on an individual subscriber basis—an entire central office must be equipped for ISDN service. The first requirement is that the office must have a digital switch. If the switch is analog, no ISDN. As we previously mentioned only about three-quarters of the switching offices in the US are digital and thus eligible for conversion to ISDN. Older analog offices are being converted to digital as they depreciate but at several million dollars a switch, conversion is constrained by investment resources of the phone companies. Even if an office is digital, expensive software and hardware add-ons are necessary to upgrade the switch to ISDN service. This is an all-or-nothing proposition for the phone company who must bet on sufficient subscriber acceptance to justify the investment. On the whole, acceptance of ISDN in the U.S. has been modest, making phone companies reluctant to forge ahead with wholesale conversion of the network.

ISDN also suffers from a chicken and egg problem. To recover the costs of ISDN conversion in the face of underwhelming customer acceptance, ISDN has been expensive. And of course, an expensive service is not apt to be adopted by the population at large. By 1997, it is estimated that a little more than one million subscriber lines are ISDN...out of a total of over 150 million access lines in the US.

And with fairly low acceptance, ISDN communication equipment like the terminal adapters that connect a personal computer to the network have been expensive. As a result, ISDN has required a significant financial commitment that most individuals have been reluctant to make. Corporations have been able to justify it in some cases for dedicated telecommuters and branch office applications, but private contractors and small offices who would expect to use it's high speed data capabilities only a small portion of the day have found the service too expensive.

ISDN is also getting squeezed by progress. In the days of 1200 bps modems, 64 kb/s of data looked pretty good. With the emergence of 56k modems that work over existing phone lines, ISDN, even at 128 kb/s of bonded capacity isn't quite so impressive. Even compared to more prevalent 28.8 kb/s connectivity, the expense of the service and terminal adapters compared to a \$100 modem gives pause to most potential users.

Another factor tending to marginalize ISDN is the Internet. ISDN is a switched service allowing a 64 kb/s data connection to be dialed like any phone call. In the early 80's when ISDN was developed, it was believed that people would be calling up others all over to transfer data—every time we needed to transfer data to or from a remote computer, we would have to call that particular computer. Then along came the Internet. Now, with one connection to the Internet, data can be transferred to any other computer similarly attached by using an e-mail address. We really don't need a networked switched service any more; the Internet provides it using routing. Given that ISDN is a measured service, at least for long distance calls, and the dial-up analog connection to the ISP is usually a local call covered by the subscriber's flat rate monthly calling package, and transport over the Internet is "free," even the pure economics of the situation are going against ISDN ever developing into a ubiquitous switched data service.

The final problem with ISDN at this late date is that it is just another burden on an already overloaded public switched telephone network. When ISDN was conceived, the Web didn't exist and planners thought ISDN users would call a computer, transfer some data (a file), and hang up: no different from a voice call. However, the Web and the Internet have fundamentally altered the dynamics of data communications. Using the Web is not about file transfers, it is about discovering and experiencing in real time an endless array of information, entertainment and data. This experience lasts not the few minutes of a file transfer, but tens of minutes (to quickly locate a few facts) to an hour or more. The average Internet household spends over six hours a week on-line; most homes don't spend six hours a month making voice calls.

The impact on the network of the long holding (connect) times of serious Internet users has received a lot of publicity lately. The Public Switched Telephone Network (PSTN) has been statistically tuned to support a relatively few short duration calls by each subscriber. An abnormal load of users and/or long holding time calls causes congestion in various parts of the network, resulting in busy signals as the network runs out of capacity—a phenomenon we all experience on Mother's

Day. Lately in parts of the country, especially since some large on-line service providers started offering flat rate access, every day has become Mother's day, as people log on and stay connected for hours instead of minutes.

From a network perspective, an ISDN call is no different from an analog modem call. Both tie up the same 64 kb/s of switch capacity and the same amount of bandwidth in local and interoffice trunking equipment. Shifting users to ISDN may provide a slight improvement in access speeds, but not enough to seriously shorten the on-line time of users who look for specific information and then log off. And it will have no impact on recreational users who will spend the same amount of time on line, and just access more information.

Telephone companies are coming to the realization that the only long term solution to network overload is to get as much Internet access traffic as possible off the public switched network altogether. Trying to reinforce the existing network to accommodate an exploding number of ever longer holding time calls is like building more freeways to relieve traffic congestion. It's terribly expensive, and somehow, you never quite catch up.

Can service providers be any more successful than highway engineers in relieving congestion? There is actually a good probability of success here because in stripping Internet access off the PSTN, providers also eliminate the very bottleneck holding back better access speeds: the 64 kb/s telephone channel. Starting with a fresh network so to speak, providers can design it to accommodate modern data communications access speeds from end to end, which in turn should cause heavy duty users to willingly give up their modems and terminal adapters for the new network.

In the next chapter we will look at some of the schemes proposed for bringing this network nirvana to the average consumer and apply the ultimate reality check: cost.

# 3

## Building the New Network



It is no secret that both cable TV and telephone companies have been trying to re-wire the countryside for at least a decade. We mentioned some of these attempts in Chapter 1. These efforts of these two service providers ultimately suffered from two major problems: each required a complete rebuild of its local distribution infrastructure before service could be offered to the first subscriber, and each was built on a somewhat dubious business case involving broad-based acceptance of enhanced video services as well as taking a significant share of the other's core business.

Without going into a lot of details, bringing true broadband services to residences and small businesses is a complex and expensive process. For over a hundred years the phone companies have been building a universal network based on twisted pair copper wire technology that today touches virtually every residence and business in the developed world. Telephone companies have been experimenting with Fiber To The Curb (FTTC) technology to bring broadband service via fiber optics to a node that serves up to 32 homes (or small businesses) using a combination of twisted pair and coax drops from the node to each subscriber. The current cost for an FTTC upgrade is about \$1000 per home or small business unit. In the U.S. alone, this equates to over 100 billion dollars for the whole country, or about one third the total current value of all local telephone company plant and equipment in service in the U.S.

In thirty years cable companies have achieved similar reach with their systems. They pass over 90% of the homes in the US and about 65% take service. Unlike the phone companies' twisted pair network, the cable system is inherently broadband, supporting several hundred megahertz of r.f. bandwidth. However, most cable plant still consists of one-way analog equipment unsuitable for two-way digital services. Cable companies are using Hybrid Fiber/Coax (HFC) technology to upgrade their networks. HFC uses fiber optics to bring broadband services to a distribution node serving about 500 homes, bypassing a portion of the in-place coax trunking equipment. Decreasing the amount of coax increases available bandwidth and improves the signal to noise ratio of the system. The remaining distribu-

tion portion of the coaxial cable plant is then upgraded to provide two-way capability and more downstream bandwidth to serve each residence. HFC upgrades cost \$250 per home passed or about \$25 billion for the entire U.S. cable system.

Driving the early efforts to develop and implement broadband local distribution systems was the assumption that subscribers would willingly pay for enhanced interactive video and video-based services: movies on demand with VCR-like control allowing pause or rewind, high quality digital video and audio, home shopping, home banking, or home travel agency services. All of this was in the face of the fact that TV watching of both basic and premium channels was essentially flat. And despite all the glitzy home-this and home-that features, apparently neither the phone or cable companies noticed that in the emerging two-earner family economic mode of the 80's, nobody was home anymore. Mix in the fact that Direct Broadcast Satellite services are approaching 10% market share, leaving a smaller pie for telcos and cable operators to battle over and the business case for broadband services isn't as compelling as it once was.

As a result, both cable and telephone companies have scaled back their broadband distribution programs. Several phone companies are moving beyond trials and embarking on programs to install FTTC local distribution systems in at least pockets of their franchises. However in the next few years deployment is expected to reach hundreds of thousands U.S. homes.

Cable operators have retrenched to using HFC to upgrade the downstream portions of their networks, leaving two-way activation for the future. By 1997, about half the homes passed have upgraded downstream bandwidth capabilities (delivering more revenue producing premium channels), but fewer than 25% are passed by two-way service. It is important to note that not everyone subscribes to cable, and cable is primarily a residential service; it is therefore not the best vehicle for bringing better data communications to small businesses.

So by the end of 1995, both the telephone companies and cable operators had drastically scaled back their timetables to re-wire the US with broadband technology. But about the same time a new service opportunity appeared on their collective radar screens: Internet Access. And as opposed to the build-it-and-they-will-come underpinnings for enhanced video services, Internet access was a market with a well established, demonstrated demand. People were logging onto the Internet by the millions and spending hours at it. Small business and residential subscribers alike were jumping in and demonstrating their enthusiasm in two ways. First, frequent and long holding time calls were being placed to either on-line service providers or local Internet Service Providers (ISPs). In addition, second line sales were booming. Surveys found the primary use for these lines was Internet access.

So consumers were already voting with their pocket books, a solid reality check. They were paying about \$20 a month for on-line services, and about the same for a second line to allow uninterrupted access. ISDN sales were showing an increase also, as rising numbers of telecommuters began moving beyond analog modems for speedier access to corporate networks.

In the face of this surging demand, both cable operators and telephone companies saw a renewed opportunity for their broadband distribution plans. However, as we have seen, both architectures, HFC and FTTC, require a fork lift upgrade to the infrastructure before service can be offered. Even with a proven demand for Internet access and telecommuting/home office services, can either the cable operators or telephone companies afford the massive investments necessary to provide them? And is the potential return worth the investment?

The cable industry is currently treating Internet access as “the next big thing,” giving it as much hype as their interactive video plans of the first half of the decade. Trials of “cable modems” are being conducted in several markets, with enthusiastic response from subscribers. The real question is whether or not the industry can turn these small groups of early adopters into a broad-based market. Just as with interactive video, unless a fairly secure market is visible, the investment in infrastructure is problematic.

One way around the problem of infrastructure investment is the telephony return cable modem. Cable modems were initially designed to use one of the upper channels in the cable TV spectrum for carrying downstream data traffic from the head end to the subscribers and the low frequencies (below channel 2) for traffic back. The downstream channel would have 10-27 Mb/s of capacity digitally modulated onto a subcarrier inserted into a 6 MHz TV channel. The return channel would provide about 10 MHz of upstream capacity. Since most Web access involves a lot of information returned for a few mouse clicks, the asymmetry is appropriate.

Given the low rate at which two-way conversions are being implemented, a new form of modem has been developed that uses an analog modem connected to a phone line for the return path. The telephony return cable modem allows operators with one-way systems to offer Internet access services without investing in the full two-way upgrade. As a vehicle for testing consumer demand for Internet access, the telephony return modem is interesting, but for all but the casual residential Web surfer, it, and cable modems in general, also have serious disadvantages.

First, all cable modem access systems are shared media networks. The physical cable system itself is a shared media network with hundreds of residences all tapping into the same piece of coaxial cable to receive TV programming. Internet access appears as just another TV channel, shared by all, just as an Ethernet is shared by

all its users—but in this case we have an “Ethernet” with potentially 500 users, all competing for bandwidth on one 10 or 27 Mb/s network. To be fair to the technology, the network can be segmented by using several channels for downstream bandwidth and assigning fewer users to each channel. However, this becomes a tradeoff with the number of TV channels the system is able to offer, and which will generate more revenue: another premium TV channel, or providing better throughput for essentially the same number of Internet users.

The other problem with telephony return systems is the telephony return. Here we are back again with analog modems when we are looking for more bandwidth. To be sure, the narrowband channel is upstream, but for serious work-at-home, telecommuting, or SOHO applications, it will prove inadequate. While the casual surfer can get away with a few mouse clicks upstream, business professionals have to periodically upload files as well as download, and uploading at 28.8 kb/s can be tedious. In addition, serious use of such a system will require a second phone line to support long sessions. So the user has to a) subscribe to cable service, b) order cable modem service at a premium rate, and c) pay for a second phone line on top of the cable bill. For the total price, the professional could probably have ISDN—and have a better chance of getting it if located in an industrial park, office building, or some other place not served by a cable system.

So, despite all the current hype surrounding cable modems, they may not be the answer for the serious business professional looking for better access to corporate intranets and the Internet.

If cable companies, with a broadband infrastructure (even if only one way) already passing 90% of homes in the US are going to have problems adequately serving professional users, can telephone companies do any better?

The answer very probably is yes. By focusing on a real demand, Internet and intranet access, and applying technology appropriate to the more modest data rates of this application (our personal broadband services definition), DSL, telephone companies are finding they are extremely well positioned to serve this market quickly and without the large up front investment previous broadband schemes required.

Digital subscriber line technology offers the following advantages:

- Re-use of the existing twisted pair infrastructure: no need to invest in new fiber optic distribution systems and distribution nodes
- Available anywhere: since DSL uses the existing twisted pair infrastructure, service can be supplied to virtually any subscriber with phone service, be they in a residential neighborhood, small office park, or downtown office tower.



- There is no need for costly central office switch upgrades as with ISDN; DSL traffic bypasses switches entirely
- No extra phone line is necessary; some DSL implementations can share the same phone line providing the subscriber's voice service
- DSL provides a dedicated high speed digital two-way access channel for each subscriber rather than using a shared media distribution architecture
- Pay as you go: investment tracks subscriber revenue

This last point is probably the most significant aspect of DSL technology. Investment matches revenue; it is not necessary to invest heavily in network upgrades before the first subscriber is turned up. And it is not necessary to worry about "take rate." For the price of one FTTC node that is dedicated to a cluster of only 32 homes, telephone companies can install the common equipment necessary to provide DSL service to any subscriber (remote or branch office, small office, home office or purely residential) anywhere in the serving area of that C.O. It does not matter how geographically scattered the demand is across the 10,000 or so lines typically served by a central office, since all lines terminate in the office, and any of them can be easily connected to the DSL equipment in the C.O. As demand builds, more modules are added to the common shelf, and more shelves are added as required. It is this pay-as-you-go aspect of DSL that finally will break the chicken-and-egg problem that has stalled broadband services in the past. DSL removes the risk of rolling out broadband services since the services can be targeted at individual users on a case by case basis rather than whole neighborhoods. And if for some reason, the service is not successful, hundreds of millions of dollars of investment in infrastructure upgrades are not stranded.

So, that wasn't so hard after all, was it? All we had to do is change our definition of broadband services from hundreds of megabits of video programming to a few megabits of Internet access to build a business case. But if the best many of us can expect of the telephone companies' twisted pair network is 28.8 kb/s (on a good day), how can it support the 1.5–2 Mb/s of even our more modest definition? Next we will take a closer look at the network and how people connect to corporate networks or ISPs today.

## Connectivity and the Telephone Network

As we have mentioned several times, telephone companies around the world have been building their networks for about one hundred years. For most of this time, the workhorse has been a simple pair of copper wires (19 to 26 gauge) loosely

twisted together. There are over 700 million such pairs or loops connecting subscribers to central offices world-wide. Originally, that's all there was anywhere in the network. A pair of wires connected a subscriber's phone to a central office and terminated on a jack on a cord board. When a subscriber wanted to talk to someone else, an operator used a patch cord to connect one subscriber's jack appearance to another's. More twisted pairs connected central offices together so subscribers could call others in the next town.

Because copper wire has a certain amount of electrical resistance, signals faded with distance, and "long distance" calling was pretty much limited to nearby towns. With the introduction of the vacuum tube amplifier, it became possible to boost the signal level of the call back up again, extending the range possible for a call. Placed every 50 miles on 19 gauge interoffice pairs, these amplifiers made true long distance calling possible. However, the build up of noise in successive chains of amplifiers eventually limits the long in long distance. In addition crosstalk between pairs in the same cable sheath and noise ingress also add noise and further limit performance. At the same time as amplifiers were introduced, the first automatic switches were beginning to replace operators and cord boards in the C.O., but calls still traveled over wire pair "trunks" interconnecting offices.

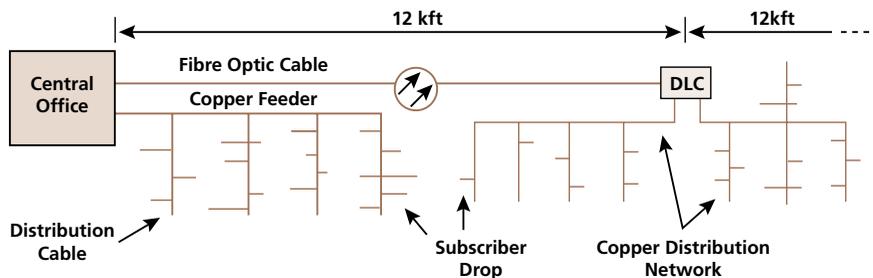
After the Second World War, the interoffice portion of the network underwent dramatic modernization as a surging volume of calls made continued reliance on copper wire technology impractical. Microwave radio and coaxial cable systems began to replace twisted pairs. Both provided better quality voice transmission over longer distances at a lower cost. In the 1960's digital transmission began to replace analog for many interoffice applications (but not long distance circuits). Digital transmission offered much quieter connections and could carry 24 or 30 conversations on the same twisted pair that supported one analog call. The downside to digital transmission (so called T-Carrier or E-Carrier circuits) was the signal needed regenerating (the digital equivalent of amplification) every mile or mile and a half. After about 50 regenerations, the bit error rate of the signal degenerates to the point where the background noise in the recovered analog signal is objectionable. However, the very low cost and stability of digital regenerators vs. analog amplifiers and the ability to get a 24 or 30 to 1 gain in traffic on the same twisted pair trunks ("pair gain," the term is much older than the company) made digital carrier extremely popular for regional interoffice use.

In the late 80's fiber optics made long distance digital circuits possible. The very low loss of fiber cable extended regenerator spacing out to a range similar to analog trunk amplifiers. Fiber cables also eliminated the crosstalk and noise ingress problems so even though more regenerators than analog amplifiers are necessary to reach

cross country, the signal quality at the far end is superior. Fiber optics also has the advantage of carrying thousands of phone calls on one glass strand rather than just 24 or 30 on wire as with digital carrier. Because of its superior performance and economic advantages, fiber has replaced wire (as well as radio and coax cable) in virtually all interoffice applications, local, regional and long distance.

Throughout this 100 years of evolution of the network, however, one factor has remained constant: your telephone is still connected to the network by a twisted pair of copper wires. The reason for this is simple economics. The cost of all of the improvements described could be amortized across tens of thousands of calls. Any cost of any change in the local loop, that portion of the network that connects you to the central office, has to be born by each subscriber. Copper wire is relatively inexpensive, it is in place, and it does the job. As we saw in the previous section, attempts to replace it with more modern technology cost more than the revenue from a basic phone bill could support.

Some change has come to the local loop, but usually not in a perceptible manner to the subscriber. In the U.S. in particular, digital technology and fiber optics have replaced portions of the loop, called the feeder (see Figure 3.1 below). This technology is called digital loop carrier, and is usually employed to serve new subdivisions beyond the range of pure copper wire or if the copper feeder cable is at capacity and there is demand for more service (unexpected second line growth, for example). DLC equipment is little more than a channel bank that can live outdoors. However even with fiber-based DLC systems, the connection from the subscriber to the DLC equipment is still twisted pair.



*Figure 3.1  
Copper and Digital  
Loop Carrier (DLC)  
Subscriber Feeds*

So that is the network today: a lot of very high tech digital fiber optic equipment connecting virtually every central office, and 100 year old technology reaching out and touching the subscriber (unless the subscriber is a good sized business, in which case the phone company may bring in one or two digital carrier circuits over fiber to serve its PBX and private line needs).

But the interesting thing is that in our quest for broadband service to the small office/home office environment, it isn't the 100 year old technology that limits us—it's all the fancy new stuff!

## Getting Connected

Suppose you want to dial into your ISP or corporate intranet. Your modem dials the number, the switch in the local CO routes your call over 2.4 Gb/s interoffice fiber optic transmission facilities, where it is passed to a local loop 155 Mb/s fiber ring and dropped off as a DS1 or E1 signal to the ISP's digital modems. Where in all this wideband digital network is the bottleneck? The answer is right in the first C.O. and in almost every C.O. After your call reaches your central office, it is probably digitized.

When digital telephone technology was introduced in the early 60's the only traffic to worry about was voice. Back in the 30's, when long distance equipment was being developed, Bell Laboratories determined that most speech energy lay in the frequency range below 3500 Hz, even though human hearing extended to almost 20,000 Hz. In fact the composite energy spectrum for a large range of normal talkers (male and female) peaks at 500 Hz and decreases almost linearly in power with increasing frequency. Highly emotional speech peaks at about 3200 Hz. Thus for faithful reproduction of speech, it is necessary to transmit only a narrow range of frequencies extending only to about 3500 Hz. In fact, the standard "4 kHz" voice channel universally used in telephone networks everywhere is designed to pass frequencies from about 300 Hz up to 3400 Hz.

This research impacted every portion of the network, including the microphones and speakers in handsets, but most importantly the design of interoffice transmission equipment. The original vacuum tube amplifiers could be band limited to this range, lowering the noise buildup in the circuit. Early interoffice trunking equipment and later microwave radios carried many simultaneous conversations on the same wire pair (or radio carrier) using frequency division multiplexing to "stack" 3400 Hz wide conversations one on top of the other. These channels are spaced in 4 kHz increments to provide filtering guard bands, hence the 4 kHz channel nomenclature. Obviously the narrower the channel, the more channels per wire or carrier, so there was a lot of motivation to band limit speech as much as practical. When digital carrier transport technology was developed, there was no reason to deviate from this 3400 maximum frequency which could handily be encoded using 8 bits with a sampling frequency of 8000 samples per second, or 64 kb/s per voice channel.

So the bottleneck is everything else in the entire network, built around the now universal 64 kb/s digitized voice channel, not the lowly twisted pair connecting us to the C.O. Therefore, if we are going to provide universal wideband services we have two problems: the 64 kb/s channel of the public switched network, and finding a method to send 1.5–2 Mb/s of data from a central office to every subscriber over the existing local loop network. Fortunately the first problem is well on the way to being solved. For the past 10 or 15 years, high speed fiber optic trunking equipment has been used to interconnect central offices, carrying multiple T1 or E1 circuits. Initially this equipment carried 28 DS1s (DS1 is the 1.544 Mb/s signal carried by T1 T-Carrier equipment) or 16 E1s; today this equipment transports over 1000 of these circuits.

In addition, due to the demand by large corporations for high speed private data networking capabilities, telephone companies have been installing wideband and broadband frame relay and ATM switching equipment in many central offices for the last five to seven years. This equipment operates with port speeds starting at DS1 or E1 and going up to 155 Mb/s or more. Thus there is already a robust, growing network designed for high speed data services in place that removes the bottleneck of the 64 kb/s public switched network.

## The Local Loop

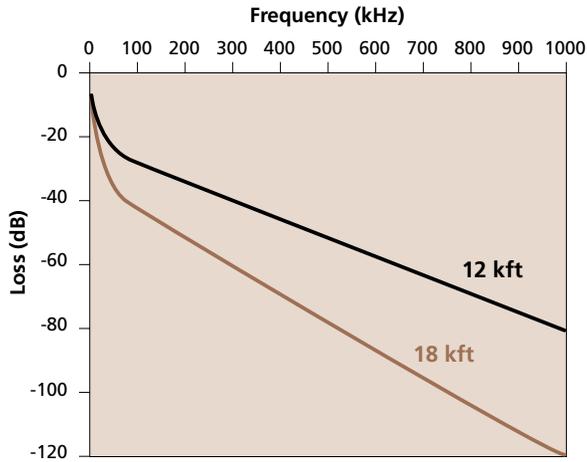
The second problem is delivering our high speed services over the existing loop infrastructure. Before describing how DSL technology accomplishes this feat, let's take a brief look at what DSL has to deal with. If we are going to reuse the existing loop network, we have to consider three factors: bandwidth, impairments caused by past installation practices, and the amount of copper available. All will set boundaries around any solution.

When contemplating using the existing twisted pair local loop for high speed service distribution, the first thought that comes to mind is bandwidth. We are used to thinking of the loop in terms of the 3 or 4 kHz voice channel, and it certainly isn't intuitively obvious that it could support several megabits of data traffic.

At first glance, things don't look too promising as the graph in Figure 3.2 shows.

We are targeting at least 1.5 Mb/s for our personal broadband services, usable over most of the local loops in the country; 6 Mb/s would be nice to have for real time video applications. At even 1 MHz, an 18,000 foot, 24 gauge loop has a whale of a lot of loss! 120 decibels represents a voltage drop of a factor of one million.

Figure 3.2  
Loss vs. frequency  
for 12 kft and  
18 kft of  
24 gauge wire

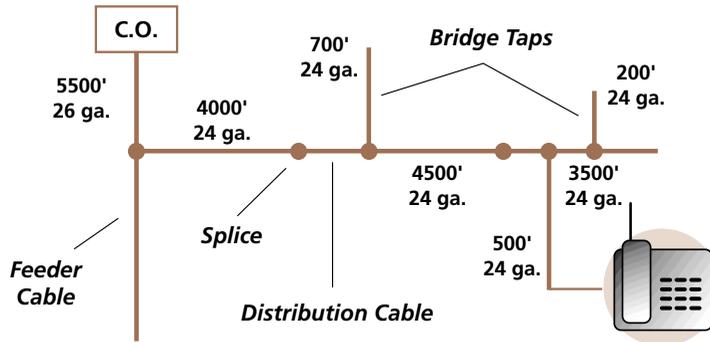


But loss isn't the only problem we have to deal with. We mentioned crosstalk above, and it is something our service has to live with. Crosstalk comes in two varieties: near end (NEXT) and far end (FEXT). NEXT is the result of a strong, near-by transmit source leaking into a receiver through the coupling between pairs; your next-door neighbor's DSL transmitter coupling into the pair feeding your home and interfering with your receiver. FEXT is the result of a source (or multiple sources) coupling into another pair and appearing at the far end along with the desired signal on that pair. For systems sensitive to NEXT, it is the more severe of the two, since a very strong signal, though weakly coupled is appearing along with the desired signal which has been attenuated by the entire length of cable (which as Figure 3.2 shows can be a lot).

It turns out that for some types of DSL equipment, NEXT actually limits the effective span length, rather than loss. FEXT is not expected to be a problem for most DSL systems except those operating at very high rates over short loop lengths.

In addition to the transmission characteristics, we have to deal with the systemic impairments in the loop caused by decades of installation practices centered around voice telephone service. As we mentioned earlier, the loop is divided into two parts: the feeder and the distribution segments. Feeder cables are large, high pair count cables that leave the C.O. and head down major corridors. Periodically, a certain number of pairs are dropped to a distribution frame and connected to distribution cables, which actually deliver service to the subscriber (See Figure 3.3). The distribution cables travel up and down every street, and as they pass a home or small business, the drop wire is attached to the cable.

Figure 3.3  
Typical Copper  
Feeder/Distribution  
Loop



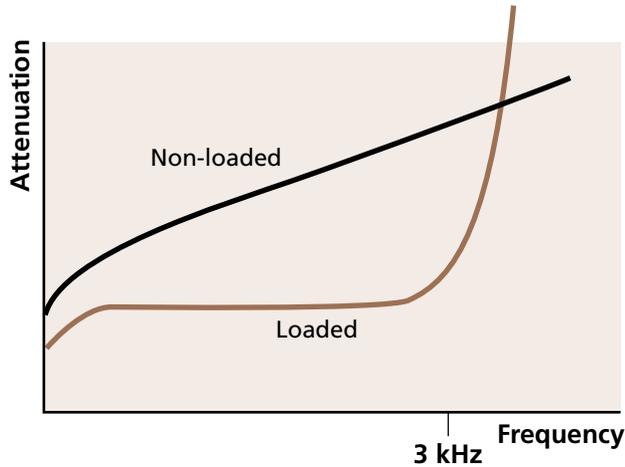
One thing is important to note: the pairs in a cable are never cut. When a subscriber orders service, the drop wire is bridged onto the passing distribution cable. Similarly, pairs from distribution cables may be bridged onto the feeder cable at the distribution frame. Thus provisioning service may consist of bridging an unassigned distribution pair onto an unassigned feeder pair, then bridging on the drop wire. The theory is that if a subscriber later terminates service, both the feeder and distribution pair can be used deeper into the loop for other subscribers since neither was cut. The problem is that often the original bridge taps were not removed when service was terminated. As subscribers add and delete lines, a feeder pair may acquire several bridge taps. This then is one impairment that any high speed distribution technology has to deal with. A twisted pair looks like a transmission line, and with bridge taps, looks like a transmission line with one or more stubs of random length connected. While not a particular problem at voice frequencies, at the rates necessary to support high speed data services, these stubs can cause frequency sensitive reflections or nulls in the line response. Since finding and removing bridge taps is an expensive and time consuming process, DSL technology must deal with them in place if it is to be successful.

Another impairment is gauge changes and splices. Often, existing loops use two different gauge wires for feeder and distribution. Since feeder cables have a higher pair count, they often use finer wire to keep overall size, weight, and cost down. Each gauge pair may have a different characteristic impedance and joining the two again creates more impedance anomalies in our transmission line. Splices create minor discontinuities in the line—yet another source of reflections.

Another major problem in using the local loop for high speed data services is loading. Loading is a process developed early in the days of the network to extend the useful range of a loop. Loading coils (small inductors) are periodically inserted

in series with the loop. Working with the natural mutual capacitance of the twisted pair of wires, the coils form a tuned circuit that reduces attenuation in the vf range at the expense of trashing the frequency response outside that range (See Figure 3.4). Loading coils can extend the reach of a pair up to 50%. Loading is required only on loop lengths greater than 18,000 feet for the 24 or 26 gauge wire typical of the loop.

Figure 3.4  
Effect of Loading  
on Twisted Pair  
Frequency  
Response

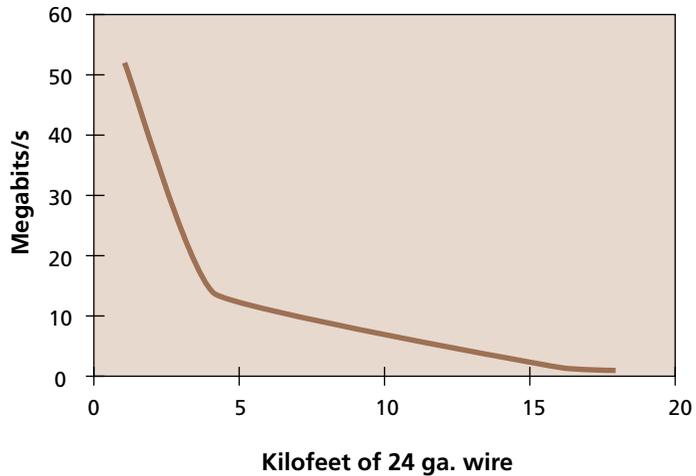


Unfortunately, this is just too much for DSL; the frequency response of the loop is bad enough as it is. Loading coils have to be removed. However, in the U.S., 85% of local loops are under 18 kft long, and the application of digital loop carrier equipment to reach subscribers far from a C.O. continues to decrease the amount of loaded loop plant in the system. It is estimated that only about 15% of the loops in place today use loaded cable. Unfortunately, just because a loop is shorter than 18 kft there is no guarantee that it is unloaded. Loaded cable has crept into the loop here and there for shorter lengths.

The result after the effect of transmission characteristics and impairments are considered is a capacity versus loop length curve for DSL systems, shown in Figure 3.5.

This chart tells us that we can have our 1.5 Mb/s service over unloaded 24 gauge loops up to 18 kft long or 2 Mb/s over a 5 km (16,000 foot) loop, and for 12 kft loops 6 Mb/s is possible. This is not necessarily a typical loop, since as we pointed out, feeder cables are often 26 gauge. This decreases our ability to reach outlying subscribers somewhat. On the other hand, telephone companies continue to work diligently to bring the maximum loop length down to 12,000 feet using digital loop carrier (DLC) equipment. About 10 or 15 years ago in the U.S.,

**Figure 3.5**  
*Data Capacity*  
*of a Local Loop*



phone companies adopted a planning strategy that said anyone within 12 kft of a C.O. would be served by copper. Beyond this distance a DLC unit would be placed 24 kft from the C.O. and serve subscribers 12 kft around it (refer back to Figure 3.1). The DLC was driven first by T-Carrier lines, and later by fiber optics. The objective was to reduce both high pair count feeder cables and the number of central offices. The serendipitous benefit for personal broadband services is that more and more subscribers will be eligible for 6 Mb/s service as time passes.

Finally we need to consider the number of pairs available. When a new cable route is installed, it is sized with enough capacity (number of pairs) to provide about 15 years of growth. Unfortunately, like a lot of guidelines in the telephone network, growth statistics were developed long before the advent of data communications. Typically, the loop has been planned with a 20% growth factor for second line services over and above the total number of homes or future home sites passed. The feeder cable is sized with a factor of 1.2 and distribution cables with a factor of 1.5, reflecting the greater uncertainty in take rate in any one segment of the route.

Lately, these sizing guidelines have been increasing as second lines for home businesses or teenagers have become more popular, but there are an awful lot of loops in place designed to older standards. The point is that the loop plant is not an inexhaustible resource, and we cannot count on having a dedicated pair or pairs for our DSL service. Having to install whole new cable routes to support DSL services would be extremely expensive and probably turn the economics of providing such services negative. This is the take rate problem again, just what we're trying to avoid.

So with all this in mind, we can begin to put some parameters around DSL service. DSL should be able to provide service to any subscriber up to 18 kft from a central office over an unloaded twisted pair with bridge taps. Only one pair should be required and this pair should be the pair that supplies voice service to the small business or home. This definition encompasses about 85% of the potential market today, and as more digital loop carrier is installed this number will grow. Equally important, it allows us to provide service over almost all loops installed for the past 20 or 30 years, not just relatively new construction.

# 4

## Taming Copper: DSL



Having put the digital subscriber line concept into both a business and physical context, it is time to take a close look at what DSL is and how turns your plain old phone line into a high speed Internet/intranet gateway. DSL is often written “xDSL”, indicating that it is a (growing) family of related standards and technologies, all designed to provide high speed datacom over long spans of twisted pair wire. The “x” can stand for H, S, I, V, A, or RA, , depending on the type of service offered by the particular flavor of DSL. First, we’ll provide an overview of these different technologies, then (for those interested) an in-depth description of how they work.

At its simplest, DSL technology is just next generation modem technology. It takes advantage of the fact that the loop is a reasonably wide band medium, and all band limiting that holds conventional modems to 33.6 kb/s or less occurs in the central office or the core network. The biggest difference between DSL and analog modems is not in technology or data rates, but in application. Analog modems are always physically located at the origination and destination of the data traffic—a subscriber and an ISP facility for instance. With DSL, one of the modems must (almost always) be located in the telephone company’s central office. It becomes the telco’s responsibility to recover the subscriber’s data and transfer it to a pure datacom network for delivery to the destination, perhaps via another DSL link.

We said “almost always” above, because some enterprising ISPs have gotten out front of the telephone companies and figured out a way to offer DSL services themselves. They lease what are called “burglar alarm” or dry pairs from the telephone company—one pair from the C.O. to the subscriber joined in the C.O. to a similar pair going to the ISP. These pairs were originally intended for use by alarm companies to provide a direct current loop to a customer to sense a continuity break caused by a door opening or glass shattering. They are called dry because they are pure copper wire end to end; no transformers, hybrids, filters, analog to digital converters or any other nasty thing in the middle to provide band limiting. Thus they are exactly what’s needed for DSL service, and still tariffed by telco’s even though most alarm companies use more sophisticated equipment these days.

In this special application the terminating DSL modem is located on the ISP's premises rather than in the C.O. While demonstrating a lot of initiative on the part of some ISPs, this application has a couple of drawbacks. First, unless the ISP can locate terminating equipment quite close to the C.O., a good portion of the overall DSL reach may be consumed by the span from the ISP location to the C.O., severely limiting the portion of the addressable customer base. Second, the ISP has to locate near every C.O. to reach the maximum number of its customers. Currently, one ISP remote modem pool can serve several central offices. Spreading equipment around tends to drive up operating expenses. Finally, at least one telephone company in the U.S. is taking a dim view of this competitive threat and is attempting to totally withdraw the tariffs on dry copper circuits.

At any rate, no matter who supplies it and how it gets to you, there are a lot of different kinds of DSL services. The following is a brief overview of each.

## HDSL

High speed Digital Subscriber Line was the first version of DSL introduced. HDSL provides a full duplex DS1 over two twisted pair up to 12,000 feet long. Developed by Bellcore in the late 80's, it was intended to be an economical method of satisfying the exploding corporate demand for DS1 services. Before HDSL, DS1s could be supplied only by either installing T-Carrier repeaters in the loop (and many loop pairs will not support the 1.5 Mb/s rates of DS1) or using fiber optics and perhaps installing new fiber cable. Both approaches were costly and time consuming. HDSL was designed to make provisioning of these T1 connections fast and inexpensive, by using existing loop pairs and requiring equipment only in the central office and on customer premises.

HDSL is of interest only because it was the pioneering high speed loop service. It is not a good candidate for providing universal personal broadband services for a couple of reasons. First, it requires two pairs to provide a full 1.5 Mb/s of service, although there are versions that offer half this rate over a single pair. However, neither the full or half rate versions can co-exist with voice telephone services on the same pair. HDSL is a well-defined technical standard, supported by several equipment vendors and is widely deployed in networks around the world.

## SDSL

SDSL stands for either single-line DSL or symmetric DSL, depending on who you talk to. It stands distinct from HDSL (which is also symmetric) in that it does operate over a single twisted pair, and in addition, allows transport of normal voice

service over the same line. SDSL is being offered in a variety of data rates ranging from 160 kb/s up to E1 at 2.048 Mb/s.

SDSL would be a good candidate for our personal broadband services except that it is limited to spans of about 10 kft or less at higher data rates; at these span lengths, other services will support significantly higher rates downstream. Because SDSL transmits and receives in the same band of frequencies in both directions, it is NEXT limited—this is why the 10 kft limitation. SDSL is kind of a narrowly targeted service where a user is close to the central office (or a remote DLC node), twisted pairs are scarce, and the up stream bandwidth is as important as downstream. If the conservation of pairs is not important, HDSL will work just as well. However, the relatively wide bi-directional bandwidth may appeal to many branch offices that have high upload requirements and well as download.

There is no approved standard for SDSL at this time, although it uses the same, proven technologies as other DSL systems. It is supported by two chip vendors and several equipment manufacturers.

## ISDL

ISDL is another vendor-driven entry into the DSL sweepstakes. It means ISDN DSL which is actually redundant since ISDN was the original digital subscriber line technology. ISDL provides 128 kb/s of pure datacom transport capacity. In most respects it is identical to ISDN, except rather than terminating on an ISDN switch, it terminates on a router and passes traffic typically to the Internet. ISDL thus provides dedicated access rather than switched service, and cannot support voice traffic as can ISDN (since it terminates on a router rather than a digital switch). It offers telephone companies the advantages of deloading long holding time traffic from the C.O. switches, as well as being totally compatible with the embedded telco tools for provisioning, administering and maintaining ISDN loops. Like ISDN, it requires only one pair, and can cover spans of up to 18 kft. However, it is not compatible with analog voice service, since the technology ISDN uses to put digits on the line assumes voice is carried as one of the digital B channels, and makes no provision for a 4 kHz analog channel.

For subscribers, on the other hand, it offers only 128 kb/s of symmetrical bandwidth, which as we said earlier, doesn't look so special in this day of 56 kb/s analog modems. ISDL meets two of our criteria for personal broadband services, 18 kft spans and single line service, but misses in its incompatibility with voice service and falls far short of the bandwidths most of us will need in the future.

## VDSL

Very high bit rate DSL is an emerging technology that promises to deliver data rates as high as 52 Mb/s (a SONET standard STS-1) downstream to the subscriber over short spans of copper wire, and lesser rates over longer spans. 52 Mb/s can be supported over a 1000 foot pair (recall Figure 3.5), while the rate drops to about 15 Mb/s for 3000 foot spans. Upstream rates are in the 1.5 to 2.3 Mb/s range.

VDSL was originally developed as part of the telephone companies' FTTC experiments. Once fiber delivered digital TV programming to the neighborhood node, an economical method was needed to reach the last 1000 feet or so to each of the 16-32 homes served. Operating at 52 Mb/s, VDSL could deliver several digitized programs to each residence as well as other services.

With the telephone companies slow-rolling FTTC projects, activity in the VDSL arena has also slowed. Several different VDSL formats have been proposed and trialed, but standardization is in the very early stages at this point.

VDSL offers far more data communications capability than most people need today and with its limited reach must be deployed in conjunction with FTTC or similar technology—not too many people live within 1000 feet of a central office. Given the price of FTTC it is not clear if our personal broadband services, even defined up into the tens of megabits, could support VDSL deployment as a stand-alone service. It will require video services to help pay the bills.

## ADSL

We have saved the best for last: Asymmetric Digital Subscriber Line. ADSL can deliver a full DS1 downstream to the subscriber over a single unloaded 24 gauge pair 18 kft long, and higher rates over shorter spans. 6 Mb/s is possible over 12,000 feet as we pointed out, and an 8 Mb/s E2 can be supported over about 2 km. Upstream rates presently are in the 64 to 640 kb/s range. Furthermore, ADSL can co-exist with standard voice service over the same pair. This certainly meets our requirements for personal broadband services: 1.5 to 6 Mb/s to the subscriber, reach to cover most loops, reasonable upstream rates of 640 kb/s to make occasional large file uploads convenient, and single pair operation without disrupting existing voice services.

ADSL is a relatively mature technology having been through several years of development and trials. It is currently in limited service deployment by some telephone companies and, for the moment, ISPs. There over a dozen different vendors offering ADSL equipment for both central office and customer premises applications.

Volume target pricing for user ADSL modems is in the same range as quality high speed analog modems, potentially making ADSL very affordable; service tariffs on the other hand are more problematic, starting in the \$200 per month range, putting it out of the reach of many small office/home office users and most casual net surfers. Wide deployment and more competition may change this picture in the future; otherwise ADSL could fall into the ISDN chicken and egg pricing trap.

Someone once remarked: “The nice thing about standards is that there are so many to chose from!” Unfortunately, this is true for ADSL. There are currently two competing, incompatible, standards for ADSL. In the U.S. one is the “official” standard, generated by the ANSI T1E1 committee, and the other is the marketplace de facto standard. The ANSI standard uses a complex, but robust, modulation technology called DMT (Discrete MultiTone—more in the next chapter); the other “standard” uses a simpler method called CAP (Carrierless Amplitude Phase—also more later).

CAP technology has been around several years and was the basis for many early trials, and despite its non-standard position, has moved into volume production by many vendors and deployment in several networks. CAP proponents argue it is more than adequate for the job, is lower cost and consumes less power than DMT products.

DMT advocates counter that their technology will provide better performance over longer, real-world impaired loop pairs, that cost and power consumption will decrease as next generation chip sets integrate more circuitry into fewer chips.

Customer reaction is mixed with some telephone companies choosing DMT products and others CAP. The ANSI T1E1 committee has been petitioned by some telcos and vendors to establish a separate CAP standard.

## **RADSL**

We conclude with Rate Adaptive DSL which is nothing more than an intelligent version of ADSL. Just as analog modems can “train down” to the rate necessary to establish a reliable connection, RADSL modems can automatically assess the condition of the twisted pair connecting them and optimize the line rate for a given line quality. This feature is important because the quality of spans varies widely depending on age, installation practices, proximity to external electrical interferers, and a variety of other factors. Line quality also varies with time of day, season and weather conditions. RADSL modems automatically compensate for these conditions, allowing full bandwidth under optimum conditions but lowering bandwidths (rather than allowing error rates to rise) if line quality degrades. RADSL allows the

service provider to provision service without having to measure a line and manually adjust or choose a modem to match. RADS� also allows a provider to provision via a management system a fixed line rate to match a particular service and tariff class, rather than having to inventory a number of modems of specific data rates.

So although there are many different versions of DSL technology, only one really meets the requirements for our personal broadband services: ADSL (along with RADS�). The others are or will be selectively deployed for specific applications, but it is clear that ADSL will be the winner overall because it can inexpensively supply a lot of bandwidth to almost any remote, small or home office, or residence that has phone service.

# 5

## How Things Work



DSL implementations can be divided into two broad categories, depending on the method they use to place the data on the twisted pairs: baseband or passband. ISDN, IDSL, and HDSL use a baseband approach, ADSL uses a passband system. Baseband systems have a frequency spectrum that extends down to zero frequency, while passband systems have a spectrum whose lower limit is at a non-zero frequency.

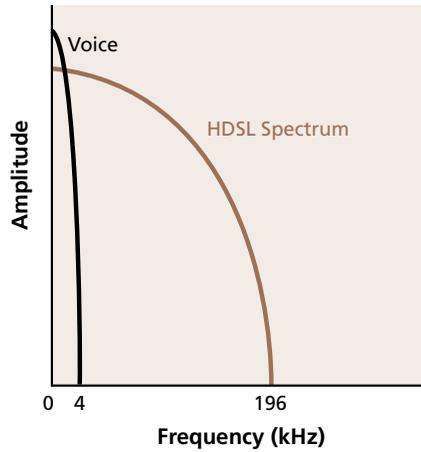
### Baseband Systems

ISDN, IDSL and HDSL all use a simple line coding technique to transport data, the 2B1Q format originally developed for ISDN. In this method the data traffic to be transmitted is accepted by the encoder two bits at a time (2B). Two bits permits four possible voltage levels (a quad of levels or Q) to be placed on the twisted pair, based on the value of the two bits as shown in back Figure 2.1.

Coding the data as above is used to halve the line rate, since the twisted pairs in the loop exhibit less loss at lower frequencies. This technology is called baseband because the spectrum of energy generated by the data stream after it is 2B1Q encoded lies between 0 Hz and some high value determined by the line rate. The spectrum of a 2B1Q encoded signal is shown in Figure 5.1. As you can see, voice service requires the 0 to 4 kHz portion of the spectrum, and this is why systems like HDSL cannot co-exist with voice service on the same pair.

In spite of their relatively simple coding scheme, implementing baseband systems can be complex. Realize that (in the case of ISDN) we are sending about 80 kb/s of encoded data simultaneously in both directions down the same pair of wires (HDSL is sending 392 kb/s bidirectionally down each of its two pair). We are in effect talking and listening continuously on the same pair. Baseband systems need circuits called hybrids to couple the transmitter and receiver to the twisted pair in such a manner that one does not interfere with the other. Since the transmitted signal is over 10,000 times stronger than the received the hybrid must be a fairly precise device to keep the transmitter from swamping out the much weaker received signal.

Figure 5.1  
HDSL Spectrum  
Use



The other complexity in implementing baseband systems is echoes. Remember the bridged taps and other discontinuities we discussed earlier? When a signal hits one of these discontinuities, a small portion of it is reflected back to the source. The problem is that in a two wire system, this reflected signal is indistinguishable from the incoming signal we are trying to receive and causes interference. In such systems, echo cancellers must be used. These devices detect the presence of echoes and subtract a time delayed, phase inverted, amplitude reduced version of the transmitted signal from the input to the receiver to exactly cancel out the incoming echo.

Both echo cancellers and hybrids have been a part of the network for years, and are not rocket science these days. But applications like HDSL are more sensitive to interference caused by hybrid imbalance and echoes than the ear listening to voice, and adding the appropriate equipment to HDSL to counter this interference adds expense and complexity. In addition to these problems, HDSL, being symmetric is also sensitive to NEXT.

## Passband Systems

Passband systems generate two or more channels well above the baseband that contain amplitude and phase modulated signals similar to those used by analog modems. Since all the data traffic is carried in these high frequency channels, the baseband portion of the spectrum is free to support voice service.

Passband systems are the preferred choice for our personal broadband service since one of our goals is to preserve the 0-4 kHz voice channel. But how does this technology manage to deliver 1.5 Mb/s down a highly lossy, impaired copper loop?

The answer is the same way an analog modem manages to put 33.6 kb/s over a 4 kHz voice channel—sending symbols instead of bits. In general both CAP and DMT ADSL modems use line coding to reduce the transmitted symbol rate to confine the spectrum in the lower frequency, lower loss regions of the response curve shown in Figure 3.2.

For shorter loops with less loss, we can extend the upper range of the downstream channel, allowing us to send the symbols faster, increasing the throughput. Or we could use more bits per symbol in the same bandwidth channel, which requires a better signal to noise ratio, which a shorter loop provides. Or we could do both, and this is what RADSL is all about: selecting the best combination of parameters for the particular loop in question.

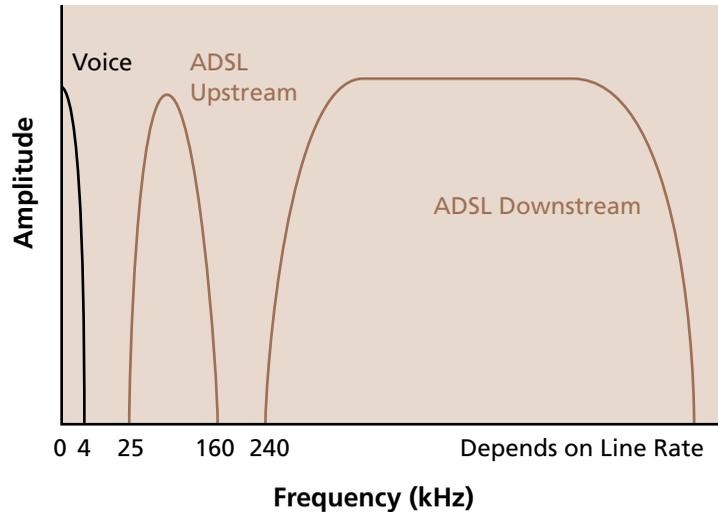
## CAP

For transmission in each direction CAP systems use two carriers of identical frequency (well above the 4 kHz required for voice), one shifted 90° relative to the other. Each of these carriers is multi-level amplitude encoded similar to the 2B1Q method above, but typically with more levels. The effect when these two carriers are summed is a modulation constellation similar to that shown in Figure 2.1 (right), with more or fewer stars depending on the number of levels. The result is a band of energy well above the  $\nu_f$  range, centered on the carrier frequency. The width of this channel is a function of the data rate to be transmitted and the number of coding levels used. This modulation is used for the up and downstream channels, using a lower frequency carrier for the upstream channel.

Depending on the data rate to be transmitted and the span length CAP systems may use anywhere from 4 to 512 amplitude and phase states to encode the data, and the upper frequency limit of the downstream channel will vary to a maximum of about 1.5 MHz (for high rates over short loops, this could extend to 8 MHz or so). The frequency use on a twisted pair with a CAP ADSL systems looks as shown in Figure 5.2.

Note that the voice channel is isolated from the data channels. With the wide gap between the top of the voice channel and the upstream channel, very simple filtering (by a device called the POTS [Plain Old Telephone Service] splitter) will separate normal phone service from the ADSL traffic. Also, each data channel is well isolated from the other. Thus hybrids and echo cancellers are not necessary for CAP implementations. The user modem is listening for the large downstream channel at the high frequency end of the spectrum; simple filtering will keep any energy from the upstream channel out of the receiver input—a precision hybrid is

Figure 5.2  
CAP ADSL Spectrum



not necessary. Similarly, an echo from the user generated upstream channel is not a problem since the receiver is not listening in that range—don't need echo cancellers either. And since the frequency bands for transmit and receive don't coincide, NEXT is not a problem—you are not listening in the same place your next door neighbor's modem is transmitting.

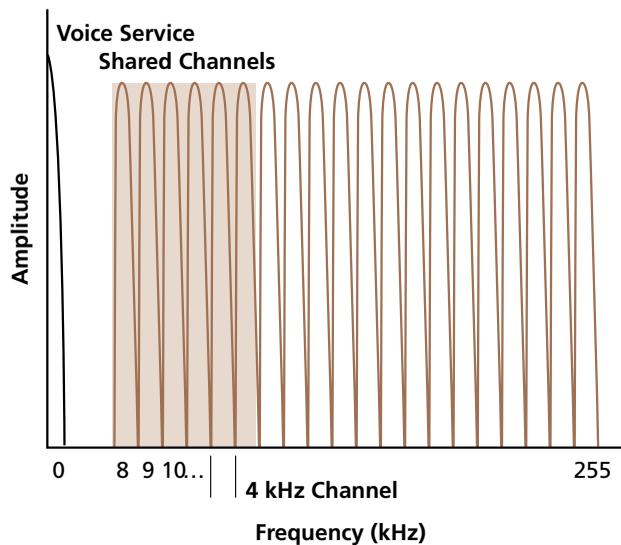
This is in fact one of the reasons ADSL is asymmetric. Separating the up and down stream channels keeps the modems simple and low cost. However, given the overall bandwidth available in a twisted pair, if you want to send a lot of data downstream, there is not a lot of spectrum available for upstream.

In the example before the start of this section, we emphasized loss. But we also have line imperfections like bridge taps, splices, and gauge changes, all of which put bumps and wiggles in both the frequency response and impedance of the line. Since the CAP spectrum is so broad, these line anomalies can distort the received signal, significantly impacting performance. CAP modems employ a process called adaptive equalization to minimize the impact of these defects. During start-up the modem pair measures the characteristics of the line and then uses equalizers that are the mirror image of the bumps and wiggles to compensate for the line's characteristics. The degree to which this is successful is a function of how many stages the equalizer has, defining how much compensation it can supply in how many places. This is in turn a function of complexity and cost. Thus although we avoid echo cancellers and hybrids, adaptive equalization takes its toll—no free lunch after all.

## DMT

Rather than using a single channel for each direction of transmission, DMT systems divide the spectrum above voice frequencies into as many as 256 very narrow channels, called bins. Each of these channels is 4 kHz wide, and again amplitude and phase modulation is used to place data into each channel. The overall effect is as if the data to be transmitted were divided into 247 separate streams and each stream fed to a modem and the modems were stacked in frequency. The spectrum of a DMT system is shown in Figure 5.3. Like CAP, the voice band is well spaced from the data channels, and a simple POTS splitter separates voice from ADSL.

*Figure 5.3  
DMT ADSL  
Spectrum*



Whereas CAP systems shift the bandwidth of the downstream channel, and hence the upper end of the spectrum depending on channel conditions, DMT operates with a fixed number of channels, 256. The DMT spectrum ranges from about 32 kHz to a little over 1 MHz. The rationale for this approach is that since many impairments are frequency sensitive, we can detect which of the channels are impaired, and shift the data to channels that are not. Each channel is individually monitored for signal to noise ratio and the data spread across best channels. Depending on the signal to noise ratio of that channel, it will be fed more or fewer bits out of the overall stream. DMT can modulate a subcarrier in each channel with an efficiency of up to 15 bits per symbol. If a channel degrades, bits are reassigned

elsewhere and if it degrades below a threshold, it is not used at all. Thus a DMT system constantly strives to optimize the signal to noise ratio of the system by continuously shifting data around among the best channels.

DMT systems are inherently rate adaptive, since if they run out of bins with acceptable signal to noise ratios, they automatically reduce the rate to match the available number of bins.

DMT systems use some of the lower numbered bins bidirectionally for up and downstream, so they need echo cancellers. This, together with the signal processing and computational power necessary to create and monitor 256 different channels, as well as shuffle bits around among them explains why early implementations of DMT equipment were significantly more expensive and power consuming than their CAP counterparts. However, because each channel is so narrow, little adaptive equalization is needed.

CAP or DMT—take your pick. Either way, we have a method of delivering our personal broadband services over up to 18 kft of twisted pair. However, the delivery choices all of a sudden got more complex in late 1997 as several variations of the basic themes of CAP and DMT appeared.

# 6

## Low Speed ADSL, G.Lite, And Other Late Starters

The information in the last chapter represents the state of the art up to the end of 1997. However, in 1997 a number of interested parties were studying the progress of a DSL technologies from R&D labs through service provider evaluations and field trials and decided it was time to put their collective oars in the water. The result was a spate of totally new product announcements and standards initiatives during the October 1997-January 1998 period designed to remove as many barriers as possible to the rapid adoption of ADSL in the consumer arena. This activity was obviously based on a significant amount of R&D activity and marketing investigation extending throughout 1997 by a lot of groups interesting in accelerating the rollout of DSL services.

1997 saw a number of field trials, and even initial service offerings by service providers across the U.S. But as these trials and service announcements mounted, a few things became apparent. First, service providers were not universally saluting the ANSI DMT standard; CAP-based equipment still had a heavy following. But more important, service providers were perceived as being somewhat conservative in the pace of their ADSL service rollouts. Standing on the sidelines was the entire computer hardware and software industry who were beginning to see ADSL as “The Next Big Thing.” Universal high speed Internet access could spur the next round of hardware upgrades to more powerful CPUs and enhanced graphics capabilities, as well as a wealth of new software applications. If MMX was the driver for the 1997 Christmas season, DSL-ready PCs had a lot of potential for ‘98.

Now the service providers had a certain degree of justification for their caution. They were worried about interference problems caused by mixing large amounts of Mb/s of ADSL with voice, ISDN, conventional T1 and HDSL in their cable plants. They were also reluctant to offer DSL service at high rates when their cable plants wouldn’t support it universally. Public Utilities Commissions generally like a uniform grade of service over an entire area. A service offering that says in effect “you’ll get whatever rate we can give you depending on how good our feeder cable

servicing you is and how far from the C.O. you are” is generally frowned upon, not to mention difficult to price. Service providers also saw the potential for low cost DSL services to cannibalize lucrative T1 tariffs. As a result, initial service offerings didn’t push the speed envelope, with most falling below 1.5 Mb/s.

A final problem service providers were wrestling with was installation. Early DSL systems required what is called a POTS (Plain Old Telephone Service) splitter to be installed at the customer premises service entrance. This device contained a simple high pass/low pass filtering system to keep the voice and data traffic from interfering with each other. Unfortunately installing this filter required a service person to visit the customer location and someone had to pay for this (phone companies estimate a service call costs in the vicinity of \$250–\$400). Initial tariffs for DSL included installation charges in the \$100 to \$200 range, creating a barrier to very broadbased consumer adoption. Installation in the central office was also a problem. DSL equipment needed new equipment shelves and racks as well as rewiring of the inside cable facilities to insert the C.O. DSL modem between the switch and the subscriber line. This too was labor intensive (expensive) as well as prone to error—even a small C.O. has several thousand pairs wired to the switch. Finding and rerouting the correct pair out of this mass (mess) of wires is tricky.

Even if the subscriber is willing to pay a premium for installation, the sheer numbers involved are daunting. The annual growth in phone lines in the U.S. is about three percent or five million new lines. With a rapid rollout and aggressive marketing push, DSL has the potential to attract several million subscribers during the first few years following introduction. Physically rewiring so many lines in the central offices and installing millions of splitters on the sides customer’s homes would be beyond the scope of existing service provider installation crews (on top of normal new line installations). DSL could easily become a victim of its own success, with demand exceeding the ability of phone companies to hire, train, and equip installation personnel. In addition to staffing up, meeting this challenge would also require significant capital outlays to furnish things like trucks and test equipment to the new field crews. Rapid acceptance of DSL by subscribers could easily result in long lead times for service, missed service dates and delays, all of which also create frowns of the faces of Public Utility Commissioners.

The bottom line is that service providers at the end of 1997 were enthusiastic about DSL service, but weren’t ready to jump in with both feet. Trials had met with highly positive subscriber response, but the path to the general availability stage appeared to cross a mine field or two. Hardware manufacturers were still battling over standards and no one was quite sure of the impact of filling feeder cables with high frequency DSL signals extending to 1 MHz or so. Installation

was shaping up to be a major headache, and DSL could have a negative impact profits from high ticket services like T1 just as capital was needed to train and equip installation crews. Many providers seemed to be leaning to rollout strategies that targeted limited offices (with known good outside plant), featuring modest access rates of 1.5 Mb/s or less (limiting impact on other services), at relatively high tariffs (limiting demand). This conservative strategy would allow providers to gain experience while controlling growth to levels that could be handled within existing budgets.

On the other side of the issue were the interested manufacturing parties. In addition to the computer industry's need for new capabilities to fuel growth, there were the dozen or so manufactures of DSL hardware who had invested heavily in development and needed to see a return in a reasonable time frame. Many of these companies were start-ups whose very life depended on rapid deployment of ADSL. None of these companies, large or small, were particularly pleased with the cautious deployment strategies of many of the large service providers.

The first steps toward moving the service providers off dead center came from Nortel and Lucent, whose central office switching equipment probably serves 85% of the lines in the U.S. In October, they both announced new versions of their line cards (the circuit card that connects the subscriber's twisted pair line to the switch or to digital loop carrier equipment) which had DSL modems built-in. Furthermore, these cards were backwards compatible with existing switch shelves and digital loop carrier equipment. This removed one of the big roadblocks: central office installation. No wires had to be touched and no new equipment shelves had to be installed. Backwards compatibility was especially significant for digital loop carrier equipment which is often installed in small cabinets with no room for additional shelves.

With its "Megabit Modem" plug-in line card replacement, Nortel went a couple of steps farther. Nortel's CAP-based solution offered 1Mb/s service downstream (and 120 kb/s upstream), more closely aligned with what service providers seemed to want to offer. Furthermore, because of the lower rate, the likelihood of reach problems for outlying subscribers was greatly diminished. Finally, Nortel's solution was "splitterless." No POTS splitter at the subscriber location and thus no service call was required ... another barrier down.

Lucent's October 1997 announcement incorporated conventional RADSL technology on the line card with 6 Mb/s down and up to 1.5 Mb/s upstream with a POTS splitter necessary. In January 1998, they announced "Wildwire," a splitterless implementation offering 1.5 Mb/s downstream.

About the same time Nortel and Lucent made their announcements, Rockwell unveiled CDSL or Consumer Digital Subscriber Line chips designed to implement

very low cost ADSL modems operating in the one megabit range. By November Rockwell and Nortel had agreed to rationalize their two specifications to guarantee interoperability between Nortel's switch cards and consumer modems using the Rockwell chip.

All in all, the last quarter of 1997 was a busy time. But January, 1998 was even better when the big guns came out to play. As we said above, the computer hardware and software industry was impatiently standing around waiting for the telecom folks to get it done. Finally, at the end of January, Compaq, Microsoft, and Intel (with the support of GTE and four of the Regional Bell operating companies) announced the formation of the Universal ADSL Working Group or UAWG. Loosely stated the goal is one standard, one chip, one motherboard for all 50 states (and the rest of the world.) The computer industry wants high speed Internet access, they want it on the PC motherboard, and they want it by Christmas. They don't want to build fourteen different versions of the mother board, depending on whether the consumer is served by RADSL or ISDL, CAP or DMT, Nortel or Lucent switches (who can blame them!). With a universal standard they hope to drive ADSL modem prices to analog levels from Day One. They are also hoping that by putting one standardized modem on the motherboards of a lot of PCs they can create enough consumer demand for DSL service access to overcome the perceived foot dragging of the service providers and jump-start the deployment of DSL technology.

Their goal is to rapidly develop a standard for submission to both ANSI and the ITU for a splitterless ADSL solution offering about 1 to 1.5 Mb/s of downstream bandwidth and more modest rates upstream. Their work is building on proposals by a small Boston-area company, Aware, that has been heavily involved in DMT modem development, and already has submitted early proposals to the ITU for a low speed, splitterless technology termed G.Lite.

The common threads running through all these new DSL proposals are low speeds and splitterless implementations. Let's take a look at how these new implementations differ from previous technologies discussed in the past chapter.

## Low Speed DSL and G.Lite

Low speed versions of DSL limit the downstream speed to 1 to 1.5 Mb/s and upstream rates to 100-200 kb/s or so. Either CAP or DMT modulation may be used, and in fact both are still very much in the picture with both Nortel and Lucent using CAP techniques for their low speed implementations while the UAWG and many stand-alone equipment vendors are using a version of DMT

called G.Lite\*. In either case, the major change from previous implementation is a big decrease in the high frequency spectra shown in Figures 5.2 and 5.3. In chapter 5 we noted that DMT uses up to 247 four kilohertz bins in the upstream direction, leading to a high end of about 1 MHz. The G.Lite version of the DMT standard proposes using bins up to only number 96, leading to a highest spectral component below 400 kHz. As we said in the previous chapter, CAP systems can trade bit rate, modulation complexity and bandwidth, so it is fairly easy for manufacturers to come up with implementations that carry 1.5 Mb/s or so in similar bandwidths. It is important to note that the impact of either scheme is on the high frequency end of the spectrum; the lower end of the DSL spectrum (the upstream information) remains unchanged, still lying in the 25–32 kHz region.

Lowering the line rate has several benefits to manufacturers and service providers. Most obviously, range is extended because we are operating in the lower loss region of the copper frequency vs. attenuation curve (Figure 3.2). Thus it becomes easier for service providers to offer uniform, predictable service over their range of loop lengths. Lower top end frequencies also theoretically mean less risk of interference with other traffic in the cable, especially conventional T1s or HDSL.

In addition, lower frequencies are less susceptible to line anomalies such as bridge taps and less sophisticated line equalization may be acceptable. In the case of DMT, fewer bins putatively means less processing. Some low speed DMT implementations could choose to separate upstream and downstream bins, eliminating the need for echo cancellation. All of these simplifications are intended to reduce complexity and hence cost...at least that is the hope of the UAWG.

## Splitterless Implementations

Recent press announcements tend to mix G.Lite, low speed, and splitterless technologies together with abandon. However, it is important to note that while G.Lite is a (proposed) splitterless standard, not all splitterless technologies are low speed. Indeed, no sooner than had low speed splitterless approaches to DSL been announced, manufacturers of full speed equipment announced splitterless versions of equipment delivering up to 7 Mb/s.

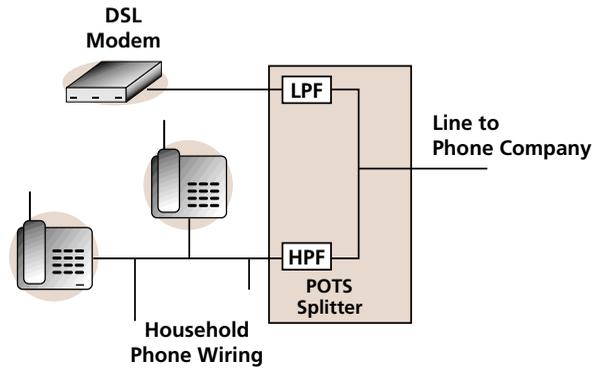
---

*\* G.Lite is a subset of the proposed DMT standard which is designated G.DMT (and in the U.S. as ANSI Standard T1.413). ITU standards all carry letter designations (such as X.25 or V.34) with the number following the letter assigned when the standard is ratified. G.Lite is designed to interoperate with any central office DMT modem adhering to the G.DMT/T1.413 standard as well as C.O.-based G.Lite Modems. Central office G.Lite modems are not required to interoperate with subscriber G.DMT/T1.413 modems.*

We mentioned above that low speed versions of DSL leave the bottom end of the spectrum in the 30 kHz region, just like full speed versions. The splitter is designed to separate the voice and DSL spectra, so what's different about splitter-less and conventional DSL or why did we ever need a splitter in the first place if the spacing between POTS and the lower end of the DSL passband is unchanged?

Conventional ADSL implementations used the POTS splitter for several purposes. The incoming phone line was attached to two filters, as shown in Figure 7.1, below.

*Figure 6.1  
POTS Splitter  
Application*



A low pass filter prevents DSL traffic starting at 30 kHz or so from entering the telephone instrument. Working with the high pass filter, the low pass filter provides sufficient attenuation of the local upstream DSL signals (the predominant source) to eliminate interference from the subscribers own modem. Now, telephones are not exactly hi-fi devices and only bats and young dogs can hear 30 kHz—so who cares? First, even though the ideal spectrum of DSL traffic is that shown in Figures 5.2 or 5.3, in the real world the modulation processes of either CAP or DMT produce lower amplitude side lobes that extend down into the *vf* region, creating a noise floor. Second, there are a lot of devices in a phone that are very non-linear. These devices, when presented with the 30 kHz+ high frequency signals of the DSL spectrum can actually act as demodulators and the difference products between pairs of high frequency signals can show up in the 0-4 kHz range. If all the signals in the DSL spectrum were allowed to interact in these non-linear elements, the result would be to create a significant amount of noise in the *vf* region. This effect will vary from phone to phone and is very unpredictable. The high pass filter takes care of the side lobe problem and the low pass insures the 30 kHz and above DSL signals are well attenuated. The filters also provide attenuation in the reverse

direction to reduce locally generated phone signals (dialing tones, off-hook pulses, fax modem tones, talking, etc.) from entering and saturating the very sensitive local DSL receiver.

The high pass filter is also necessary to shield the DSL receiver from normal incoming telephone signals put on the line at the central office. Ringing voltages are particularly nasty, being very high voltage and rich in harmonic content. Also, unlike the ear, DSL receivers can “hear” down to practically zero Hertz, so they would busily try to demodulate downstream voice, fax, or modem traffic if it was allowed to get in. The high pass filter keeps the DSL receiver and demodulator from getting distracted by normal *vf* traffic being sent from the central office to the subscriber.

The filters in the splitter also tend to isolate the receiver from the premises copper plant. Just as the local loop is full of anomalies, the local house inside wiring is a mixture of bridge taps which may or may not be terminated in working telephones, gauge changes, and changes in types of wire from twisted pair to flat cable. In addition the net impedance of the inside plant changes every time a phone goes off hook. Rather than presenting the input of the DSL modem with a wildly changing (and widely varying from home to home) impedance, the filters provide a more controlled impedance to the modem.

Splitterless implementations of DSL eliminate one of these filters, the low pass. Now the DSL modem is directly connected to the household wiring plant just as an analog modem was in the past. None of the above problems go away however, and the DSL modem must still be designed to cope with them. In particular, the high pass filter is retained and incorporated on the circuit card containing the DSL modem. This reduces DSL spectral side lobes to an acceptable level and continues to shield the modem from *vf* energy. However, the DSL frequencies from 30 kHz up can still play havoc with the nonlinear elements in the phone and the high pass filter obviously has no impact on them. Since the low pass filter is no longer present to help isolate the subscriber’s phones from the DSL modem, other steps must be taken to insure there are no interference problems. The easiest way is to just reduce the power of the modem’s upstream transmitter. Studies have shown that a 6 dB reduction in transmit power (reducing it to 25% of its former value) will allow splitterless implementations to meet the objectives set for *vf* band interference.

However, reducing the upstream transmit level imposes limitations on the upstream speed of splitterless DSL, since span length and bit rate are directly traded off against each other. In many installations the higher power downstream link will determine the maximum span length possible. In this case the lower power upstream signal must cope with this span length and the only option is to reduce

the upstream line rate. As a result splitterless DSL systems tend to have upstream rates in the 100-200 kbs range for spans in the 18 kft range.

Without the low pass filter to help it out, the high pass filter is now connected directly to the uncontrolled impedance profile of the home. The effect is that splitterless DSL modems will have to provide a little more and faster acting line equalization than before to react to conditions such as extensions going off hook. However, this is a “silicon” problem that is amenable to an on-chip solution and is not expected to add cost.

As you can see from the above discussion, the problems associated with splitterless implementations concern the low end of the spectrum which is independent of line speed. High speed splitterless ADSL is feasible and is being readied for market by several vendors. Splitterless technology, at any speed, is a significant advance in enabling widespread deployment of ADSL since it removes a large barrier: service provider premises installation.

## Summary

How big an impact this flurry of announcements will have on the ADSL market remains to be seen. Endorsements by the likes of Compaq, Intel and Microsoft are a very positive step. However, even if they manage to achieve their goal of a single standardized ADSL modem in every PC by Christmas 98, it is not clear if the consumer demand they hope to generate will energize the large service providers who control most of the copper. Absent serious, aggressive competition in the local loop, especially for residential access, the traditional measured, controlled deployments common to many new phone services are apt to prevail. It is also not clear whether the G.Lite standard will receive any more respect from key vendors than has the existing ANSI standard for DMT.

On the hardware side, splitterless ADSL is highly significant because it both eliminates many installation problems, and coupled with lower speeds, removes some of the risks of introducing new services into an already complex mixture carried by the feeder network. Equally significant is the incorporation of ADSL modems directly on switch and DLC line cards. This step further reduces the cost of installation and brings ADSL under the common management umbrella of the switch management system. These two new hardware approaches to DSL will probably do more to accelerate deployment than any consumer-based pressure the UAWG may be able to generate. They remove many real-world, cost-based roadblocks in the path to high speed Internet access. Current implementations were developed using CAP rather than DMT because mature, low cost, low power



consumption (very important in this application) CAP chips were available while DMT chip sets were still in early prototype stages. If the UAWG is successful with G.Lite, eventually DMT chips with competitive features will be available for switch line cards as well, allowing plug-and-play ADSL in both the home and central office.

Finally it is worth noting that the UAWG is focused on jump-starting the residential market. While low speed, splitterless implementations may be fine for residential and home office use, small offices and branch locations will need more bandwidth, especially upstream, than standards such as G.Lite offer. Conventional CAP and DMT solutions are more appropriate here and are available today. Since many small and branch offices already have local area networks in place, the stand-alone DSL modem model is also more appropriate than PC motherboard modems.

But in either case, lite or full speed DSL, and regardless of where the modem is, this is only the beginning. ADSL is just the first step in bringing personal broadband services to offices and residences. ADSL provides the physical link, but we still need protocols above it just as with any data communications system. Choosing the correct upper layers will be just as important as the correct physical transmission system if users are to realize the full benefit of personal broadband delivery to their doorstep.

# 7 Layering It On

Data communications between two entities takes place using a set of rules or protocols. These rules govern every step and aspect of the communications process. In this broad definition, communications implies more than just the transfer of data; it also encompasses the origination and destination user or process, insuring that the received data is interpreted and presented in the same format as originated.

Protocols tend to come in stacks (See Figure 7.1 below) made up of layers, with each layer dealing with a specific portion of the communications process. The lower layers deal with what we usually think of as communications, i.e. the reliable transfer of data from source to destination. The upper layers deal with interpretation and presentation via specific applications, and are not of particular concern in our discussion of personal broadband service delivery.

Protocols come in stacks with well defined layers to remove dependency of one layer on another. Thus Layer Three can be used with any Layer Two protocol that provides the services Layer Three expects. Therefore IP runs quite happily on Layer Two protocols like PPP or SLIP for point-to-point connections (like dial-up

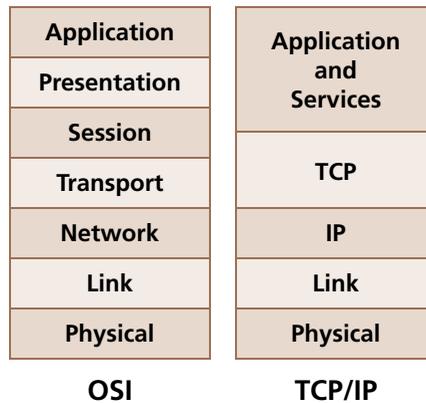


Figure 7.1  
Protocol Stacks

connections to an ISP), or on Layer Two LAN protocols like Ethernet, Token Ring, or FDDI. And Ethernet, a Layer Two Protocol, is equally at home on any of three Physical Layers: thick or thin coax or twisted pairs. The concept of stacks and layers lets us pull out one layer and slip in another without having to change any of the others above or below. So even though our personal broadband services might be based on ATM, and a corporate network could use frame relay, from Layer Three up they are totally interoperable and the only trick is to provide a conversion from ATM to frame relay to let a telecommuter access the corporation. More about this in Chapter 9.

The stack shown on the left is the classic OSI 7 Layer stack. Not all protocol suites use all seven layers explicitly. The TCP/IP suite tends to bundle the top three layers into one application and service layer sitting on top of TCP as shown; File Transfer Protocol (FTP) is an example. IBM's SNA uses a six layer stack combining the top two layers into one.

The lower layers are from bottom to top:

*Physical Layer*—this layer is concerned with the establishment and removal of a physical circuit between two adjacent communicating devices. ADSL is an example of a physical layer, and the specification of wire gauge, span lengths, acceptable total bridge tap length, transmitter levels, and modulation techniques, are all examples of things that are specified at this level. The physical layer may also make provisions for monitoring and reporting certain parameters, such as bit error rates or failures via a management system.

*Data Link Layer*—the link layer provides for reliable transport of data over the physical layer between two adjacent communicating devices. Protocols at this level are concerned with ordering and formatting data into identifiable frames, maintaining timing and synchronization between the ends of the link, detecting errors, and providing addressing to guide the frame to the appropriate link layer destination. Flow control is also provided at this layer so a source cannot overwhelm the device on the other end of the link with data.

*Network Layer*—the network layer is primarily concerned with routing data throughout an overall network, the Internet for example. Using a destination address built into the header of a network layer packet, devices in the network, typically routers or packet switches, relay the packet from source to destination.

*Transport Layer*—end to end communications is the responsibility of the transport layer. Higher layer information (which may be a continuous stream of data) is broken down into manageable chunks here for Layer Three. Layer Four also numbers these packets in sequence created. Layer Three is a “best

effort” service, and a packet delivered to Layer Three by Layer Four may be lost in the network or delayed, and arrive after later packets were sent. Layer Four is responsible for reordering packets and requesting retransmission of lost packets. Layer Four also provides end-to-end flow control so that a high speed server connected to the Internet by a DS3 doesn't choke the 28.8 kb/s modem on the receiving end.

For each of these layers, there are a large number of choices of individual protocols. For layers three and four, by far the most popular are Internet Protocol, IP, for layer three and Transmission Control Protocol, TCP, for layer four. Since both layers three and four are wedded to the overall network, they don't impact our system of delivery of personal broadband services. These layers will transparently pass through our delivery system.

We have already specified ADSL as our target physical layer, which leaves Layer Two, the link layer, to deal with. Since the user computer on one end of the ADSL connection and the service provider equipment on the other end form a well defined Layer Two link, Layer Two is something we have to worry about in specifying our personal broadband service. And as we said earlier, the protocol choice will have a great deal to do with the flexibility of service ultimately delivered.

We haven't discussed services in too much detail so far, so what kinds of services do we expect our personal broadband service to provide? Realistically, the answer is we aren't sure. Certainly we want to connect to the Internet and the Web, and to corporate intranets as well. But we are bringing at least 1.5 Mb/s of bandwidth into small offices, home offices and residences everywhere, with an upstream channel of over 600 kb/s. Video conferencing is a definite possibility. Specialty video programming is another, such as training or lectures, conference panel discussions and similar material. And as we have pointed out audio and video are becoming a factor even in browsing the Web. Dealing with audio and video in real time has to be included in our service mix along with file transfers and Web page downloads.

Another factor to consider is that these personal broadband services will find their way into multi-user environments. Small office certainly implies more than one person, and even a home office may have more than one Internet appliance. Our service should have the flexibility to handle the demands of several users simultaneously with one or more involved in a real-time application, like a conference.

Unfortunately, most of the classical layer two protocols used with TCP/IP or other higher layer protocols don't fit the bill. They were all developed when data communications was just that: transferring data, usually files or e-mail. None of them take into account the needs of real time information: timing sensitive transmission with low delay.

As data communications developed, the link layer was regarded as the first line of defense against a hostile physical world. In the 60's and 70's the physical network was made up of copper phone lines and microwave radio long distance links. These analog media were susceptible to noise from all kinds of external sources and occasional brief drop outs. Therefore as data communications protocols were developed, a lot of effort was put into error checking and recovery. The primary purpose of the link layer is to ensure reliable transfer of data across an assumed error prone channel. Little thought was put into what kind of data might be passed across the channel—back then bytes were bytes.

Generally, link layer protocols provide no operation on the data they transport. They accept a packet from Layer Three, encapsulate it in a frame and send it on its way. Short packets, long packets, it makes no difference; Layer Two takes them as they come. The frame typically consists of an address field, control field, the information packet from Layer Three, and a frame check sequence. The frame check sequence is a 16 or 32 bit error detecting code that allows the receiving end to determine whether the frame arrived error free.

In most Layer Two protocols, the receiving end of the link acknowledges receipt of frames, either on a frame-by-frame basis or in groups of frames called windows. If an error is detected by the receiver, it sends a negative acknowledgment, requesting retransmission of the frame. Under a windowing scheme acknowledgment is sent only every  $n$  frames, with  $n$  being the window size. Windows are more efficient than frame by frame acknowledgment. However, depending on the protocol, if an error is detected somewhere in the window, that frame, and all successive frames up to the end of the window must be retransmitted.

The fact that a Layer Two frame may be of variable length (depending on the size of packet handed down from Layer Three) and errors may cause retransmission of an arbitrary number of frames, most links have unpredictable delay characteristics. That is there is no guarantee on the elapsed time between the time a packet is passed to Layer Two on the transmit end and when it is passed back up to Layer Three on the receive end.

Furthermore, data communications protocols are not real-time processes; data can be accumulated as it arrives at a receiving PC and processed either after it all has arrived (opening a word processing file, for example) or after significant amounts have been received (as photos on a Web page are sometimes displayed). Either way, there is no assumed or necessary relationship between the rate at which the data is transferred and the rate at which is processed and displayed in the receiving system. There is no loss in utility or intelligibility of the data no matter how quickly it was transferred or if it came in bursts followed by pauses.

In this respect, data is like a book. When you begin to process (read) the book you have all the information in your hands and can process it at your leisure. Even if the book is serialized in a magazine, the data is just as useful if the magazine comes weekly as it is if it comes monthly—and it doesn't even matter if the magazine is delayed in the mail sometimes.

Contrast this with processes used to digitally encode voice. The encoder samples the voice wave form 8000 times a second and forms an eight bit word representing the voltage amplitude of the sample. These samples are sent in a continuous stream to the receiver. At the receiver, the word is decoded and that portion representing one eight-thousands of a seconds worth of voice is reproduced. The decoder immediately reproduces the next portion and the portion of that, building a smooth wave form.

Now suppose we introduce a couple of bytes of delay. The decoder reaches out for the next word to decode...and nothing's there! It has absolutely no idea what to do. It can't just put everything on hold—it's building a wave form in real time. At best it will look at those two bytes as 16 zeros and decode them as large negative swings of amplitude. The result will look something like that shown in Figure 7.2 below.

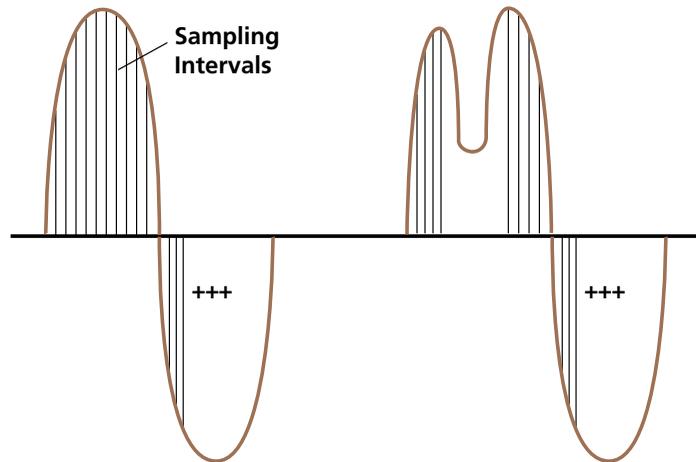


Figure 7.2  
Impact of  
Sample Delay on  
Reconstructing a  
Wave form

Realistically a brief glitch of 2 wrong 8 bit samples out of 8000 in the wave form as shown would show up as a click in the conversation which might not even be noticeable. Now suppose we throw in a delay equal to a three or four packets of more typical size of several hundred bytes each. Now we begin to approach a serious

drop-out of a quarter to half a second—certainly noticeable. (In reality, loss of two bytes or several kilobytes would induce all kinds of bad behavior in a typical T1 or E1 channel bank decoder; it would lose frame, probably give up entirely, disconnect all circuits, produce alarms and try to reframe.)

This is a major problem in trying to use traditional datacom protocols with real time processes. Data packets vary in length from a few tens of bytes to over a thousand bytes. As packets are routed around the network, they are buffered, interleaved with other packets and sent on with no regard (or knowledge) of what's in the packet. The source end of a real time process could insert packets into the network in a very methodical manner, but there is no guarantee that some of these packets won't get stuck in a queue somewhere behind one or more thousand byte long packets. The result is a large variation in arrival times; this delay variation is actually more of a problem than absolute delay.

Even if we gathered up many bytes of our digitized voice conversation and formed a large packet, so we could fill a buffer at the receive end to give the decoder a backlog to work off while it awaits the arrival of the next large packet, the network can't assure us it won't interject so much delay that the receiver will run out of data before the next packet arrives. Such a scheme would also introduce objectionably long absolute delay which would be a problem for two-way voice conversations. The result would be similar to holding a conversation over a satellite link. Even a half-second of absolute delay confuses speakers and makes them talk over each other. (Long absolute delays also upset datacom processes. TCP expects acknowledgments of transmissions, and long delays can cause timers to expire and unnecessary retransmissions.)

Digitally encoded video suffers the same problems. Video decoders run at a constant rate and expect the next byte of information to be there when they need it. If it's not, you get speckles, streaks or the entire picture freezes. In the real world, both audio and video systems can be built to tolerate a certain amount of random delay both by buffering a little data and filling in a brief gap by repeating the last sample received. For a one byte voice sample or one scan line of video, such repetition is not perceptible. But at some point these artifices become noticeable and objectionable, and again, the typical datacom network running traditional protocol stacks makes no guarantee about consistency or length of delay.

Such two-way processes place most stringent requirements on our personal broadband services. We require reasonably low absolute delay from source to destination with very little variation in delay from packet to packet.

In the past, none of this has been a problem, because real time processes like voice were handled on the public switched network which was designed expressly

to minimize overall delay and provide virtually no delay variation. Data was handled by a totally separate network designed to efficiently transfer large blocks of time insensitive information. The best way to accomplish this was to form large packets and flood them into the network on a first come, first served basis.

However, personal broadband services will be defined by a convergence of voice, video and data into a common channel feeding the small office or home. Therefore in choosing a Layer Two protocol, we would like to find one that gives us the flexibility to efficiently support classical datacom higher layer protocols like TCP/IP, but at the same time can handle the low delay and delay variation requirements of voice and video applications using entirely different higher layer protocols. In effect we need a way to tell the network what kind of traffic we are dealing with and how to handle it.

So is there a link layer protocol that meets our requirements for low delay and delay variation for voice and video and that can still cope with the long, variable length packets of classical datacom? The answer is, of course, ATM. This is exactly the task ATM was designed to accomplish. In the next chapter, we will give a brief overview of ATM and why it works so well with ADSL in providing our personal broadband services.

# 8

## ATM for the Link Layer

Asynchronous Transfer Mode is a relatively new packet technology that offers significant benefits over other data transfer methods such as frame relay, X.25 or IP routing. ATM was developed specifically to allow a single network to simultaneously handle the large frames bursted by datacom applications and to still provide the low delay requirements of real time processes. A link layer using ATM is divided into two sublayers: the ATM Layer and the ATM Adaptation Layer. The ATM layer is concerned with forming the basic unit of transport, the cell; routing it through the network; establishing and tearing down connections; and guaranteeing a defined level of performance, or Quality of Service, to the user. There are several different Adaptation Layers currently defined each of who's role is to map one of the many different types of traffic ATM can transport into the basic cell.

### The ATM Layer

ATM manages to bring the almost mutually exclusive requirements of real time processes and efficient data transport together because of two factors. First it is a link layer protocol, and second it uses a short, fixed length frame called a cell to hold data passed down from higher layers. A cell contains a five byte header (mostly address field) and a 48 byte information field. Figure 8.1 contrasts an ATM cell with a traditional Layer Two frame, such as HDLC.

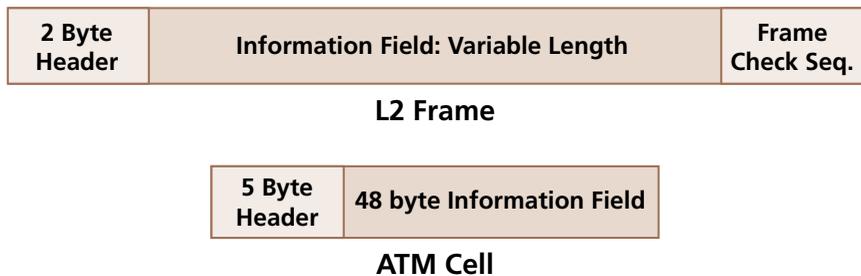


Figure 8.1  
HDLC Frame  
and ATM Cell

A couple of things are immediately obvious. First, the cell is much shorter and the data field is always 48 bytes as opposed to the variable length of the HDLC information field. It is this short, fixed length that allows ATM to transport both data and real time information. In an ATM-based network, real time traffic has to wait no more than one 53 byte cell for processing. As we mentioned above, in conventional data communications networks, a real time packet could easily find itself trapped behind several large data packets. With ATM this cannot happen, and furthermore, the delay variation can be made very small, since the cells are all the same size. Thus we can transmit our real time traffic very methodically, slipping it into a stream of datacom-carrying cells every third cell, tenth cell or at what ever rate is necessary.

We pay a price for the short cell size, though. Datacom packets are typically several hundred to over a thousand bytes long and they have to be chopped up into 48 byte fragments for transport through an ATM network. With its five byte header, ATM carries almost a 10% overhead. This contrasts with an HDLC frame which might only have 0.2% overhead. People concerned with data communications regard this 10% overhead as too high for efficient data networking. Voice and video types wanted an even shorter cell to reduce delay even further. The 53 byte cell was a compromise, and since ATM networks were designed to operate at very high rates, data transfers would be as fast or faster than in a pure datacom environment.

The other thing to notice from the above figure is that the cell has no error detection mechanism like the frame's frame check sequence (actually there is a error detection mechanism in the header to pick up address errors, but no mechanism to detect overall errors). ATM makes the valid assumption that going forward from the 90's, the error rate of the network is extremely low. Modern digital, fiber optic physical layer networks have orders of magnitude lower error floors than the older analog, copper wire networks. Therefore it is not necessary to check errors over every link in the network; an end-to-end check by higher layers is sufficient. Eliminating Layer Two error checking really speeds things up and lowers delay. We eliminate windowing and any retransmission of one or more frames altogether. The receiving entity doesn't have to buffer the frame until the checking process is complete; cells can be passed as soon as the address is read and verified.

ATM has another few tricks up its sleeve to provide a smooth ride through the network for real time processes. First, all switching is done at the link layer. This contrasts with most other datacom technologies which route at the network layer. An IP router for example receives a packet delivered over some physical medium. The router goes through the link layer error detection process, then strips off the framing information and examines the header of the Layer Three packet. The IP

header is 20 or more bytes long and contains a number of fields, including version number, header length, total length, an error detecting checksum, time to live, and an address field among others. Each of these fields must be processed to insure the packet is valid; if any of them is in error the packet is discarded. The router then looks up the address of the packet in a routing table to choose the proper physical output port to use to forward the packet to the next router. Before the packet is sent on, the header has to be rebuilt because some of the parameters in the header, such as time to live, change on a hop by hop basis. All this takes time and adds delay as packets sit in buffers while they are processed.

ATM switches traffic at the link layer based solely on the information contained in the header. Other than the header error checking, virtually no processing is done on the header or cell; only a single one bit parameter, the cell loss priority, may be changed if necessary by the switch. With so little processing of a shorter header, ATM zips its cells through a switch much faster than a router can process an IP packet.

ATM is also a very deterministic process. Cells always follow the same path through the network, and sequence of cells is guaranteed. In addition, ATM guarantees the quality of service on an end-to-end basis. IP on the other hand is a best effort, connectionless protocol. There is no pre-established path or connection; forwarding decisions are made on a packet-by-packet basis, based on factors such as network congestion or state of a link. If a link is down or momentarily congested, an alternate route will be used. Packets may be buffered in a router for varying amounts of time, or discarded if the router's buffers are full. There is no guarantee that a packet submitted to the network will emerge at the other end in a timely manner or emerge at all.

In addition, because IP chooses routes on a packet-by-packet basis, some packets may take a longer path than others, arriving out of sequence. It is up to the Transport Layer to reorder the packets before passing the entire message up to higher layers. The Transport layer therefore has to wait for all packets to arrive before it can process the message and then pass it up. This adds more delay. With ATM, cells can be streamed to higher layers as they arrive since preservation of cell sequence is guaranteed. For real time processes, this is critical.

ATM guarantees the route and quality of service by using signaling or pre-established connections to establish the path through the network before allowing traffic to flow. This makes ATM a connection-oriented protocol versus the connectionless nature of protocols like IP. When any kind of traffic is to be sent over an ATM network, the user delivers a number of parameters to the network either during the signaling portion of connection establishment or as part of a contract

between user and service provider for permanent connections (ATM “leased lines”). These parameters specify the bandwidth required, the error rate necessary, and the delay and delay variation that can be tolerated, along with the destination address. The network (or service provider for permanent connections) then checks to determine if a connection with these characteristics can be supported through the network from source to destination. If it can be, data is allowed to flow, and the characteristics are guaranteed for the duration of the connection. If the network cannot support the parameters the traffic is refused or the network may try to negotiate other values for the parameters with the user.

Every switch in the path keeps a record of the required parameters of the connection by cell address. If it is a low delay cell, it is given forwarding priority over other cells (say from a connection supporting a file transfer) that have little or no sensitivity to delay. With this mechanism, ATM can provide the low delay variation characteristics necessary for real time applications in a network that is carrying a mixture of traffic.

ATM offers (currently) five different classes of service (see the tables below), each defined by a combination of traffic and quality of service (QoS) parameters. Traffic parameters define “shape” of the data to be transferred in terms of burstiness, maximum and minimum acceptable rates, and cell delay variation tolerance. Traffic parameters are specified to the network at time of connection set-up. At the same time quality of service parameters are negotiated between the user and the network; these include peak to peak cell delay variation, maximum cell transfer delay, and cell loss ratio.

<b>Traffic Parameters</b>					
	CBR	rt-VBR	nrt-VBR	UBR	ABR
PCR	S	S	S	S	S
CDVT	S	S	S	S	S
SCR	NA	S	S	NA	NA
MBS	NA	S	S	NA	NA
MCR	NA	NA	NA	NA	S

*PCR: Peak Cell Rate, CDVT: Cell Delay Variation Tolerance,  
 SCR: Sustainable Cell Rate, MBS: Maximum Burst Size,  
 MCR: Minimum Cell Rate  
 S= specified, NA= not applicable.*



Quality of Service Parameters					
	CBR	rt-VBR	nrt-VBR	UBR	ABR
ptp CDV	S	S	NS	NS	NS
MCTD	S	S	NS	NS	NS
CLR	S	S	S	NS	opt.

*ptp CDV: peak to peak Cell Delay Variation, MCTD: Maximum Cell Transfer Delay, CLR: Cell Loss Ratio*

*NS: Not Specified, S: Specified*

The peak cell rate and the cell delay variation tolerance must be specified for all connections. These two parameters put an upper bound on the behavior of the source.

CBR, Continuous Bit Rate service is a service class that is used for ATM transfer of real time information that flows continuously, a 64 kb/s voice channel for example. The data-oriented parameters such as maximum burst size don't apply to such traffic and aren't specified. On the other hand, all the QoS parameters are, since cell delay variation and absolute delay are critical in such applications.

Real time Variable Bit Rate Service is intended for technologies like compressed, packetized audio and video which have real time concerns about delay, but tend to be bursty in nature rather than stream continuously. Variable bit rate in this context refers to the fact data is not sent regularly but in bursts.

The other three service classes are intended for data communications applications. Non real time VBR provides a robust connection for data because the user can specify the minimum acceptable error rate in terms of cell loss ratio.

At the other end of the spectrum is the Unspecified Bit Rate service where virtually nothing is specified. This is a best effort service similar to IP in that the network makes no guarantee of quality of service. Because it is so undemanding of network resources, UBR is expected to be a low cost way of handling bulk data transfer applications, e-mail, remote file access, and similar applications that are insensitive to delay or over all transfer time.

Available Bit Rate service is a new class that allows a user to take advantage of the dynamic nature of ATM. Like any traffic carrying system, the load on an ATM network varies from moment to moment. ABR allows the source to specify the maximum rate at which it is capable of transmitting and the minimum transfer rate acceptable. The network accepts the connection and then applies feedback on the source, allowing it to send faster when there is free bandwidth, but throttling it back when other connections with higher QoS requirements need capacity. This is

perhaps the closest any network has come to delivering dynamic bandwidth on demand. ABR is expected to be the work horse of data communications, providing interconnection of bursty sources like LANs and routers.

Another thing worth noting about the ATM connection process is that the parameters discussed above are negotiated for each direction of the connection. Thus symmetrical or highly asymmetrical connections can be established depending on the application. A file transfer would specify a large amount of bandwidth downstream and only a narrow return channel to handle acknowledgments. This conserves network resources, not to mention the user's bank account. It also matches the asymmetric nature of our personal broadband services.

## ATM Adaptation

A remaining issue with ATM is getting our traffic, real time or data, into the 48 byte payload of a cell. ATM places an adaptation layer directly above the ATM layer to carry out this task. The adaptation layer is the process whereby the many different kinds of traffic ATM can transport get segmented and ultimately reassembled into packets or data streams that can be handed off to upper layers in the protocol stack. This process involves more than just chopping long packets or continuous data streams into 48 byte chunks (called protocol data units or PDUs) and stuffing them into a cell. Real time processes usually need a sample of their timing source sent along with the data, or some other manner of synchronizing the source and sink clocks. At some level, error checking is desirable even in a low error environment like the fiber optic networks used to tie ATM equipment together. It might be desirable to multiplex several data streams onto one ATM connection; in that case a process is needed to assign unique IDs and sequence numbers to each such stream as it is segmented and placed into the cell. All of these processes are the responsibility of the adaptation layer.

During the early development of the ATM standards, five different AALs were sketched out, AALs 1 through 5. In the interim AAL2, intended for packet video, was not defined at all and AAL3 and 4 were combined into a common standard, AAL3/4. This adaptation layer is intended for transfer of data that may not have higher level error checking and/or data that consists of several different streams multiplexed together; it has found little application so far. The remaining two, 1 and 5 are the primary adaptation layers in use today. AAL1 provides for transport of real time information and contains a mechanism to carry along clock synchronization information as well as the data. AAL5 provides a straight forward method of segmenting very long data packets into 48 byte cell payloads.

## The AAL Process

The ATM adaptation layer is itself divided into two sublayers: the convergence sublayer (CS) and the segmentation and reassembly sublayer (SAR). The CS may be further divided into a service specific part (SSP) and a common part (CP) depending on the particular application. The SSP is meant to cater to any quirks unique to a specific service that are not covered in the CP, and the SSP will be defined in conjunction with that particular service. The SSP is often “null” meaning that data passes directly to the CP. The CP is defined as part of the AAL layer specification since it is a general process. We will cover only the CP in our discussion of the AAL.

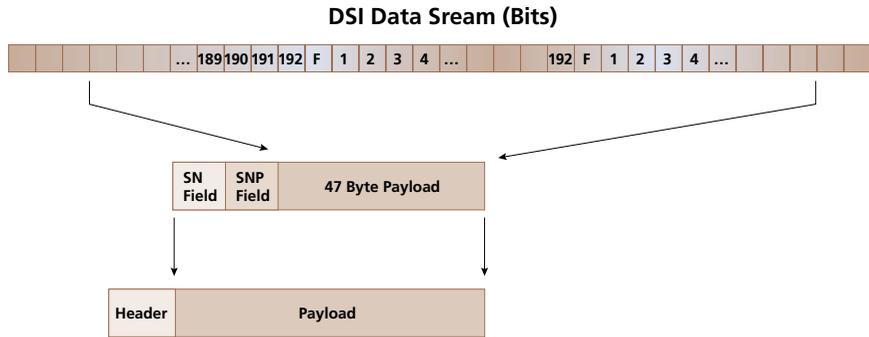
### *AAL1*

AAL1 is used for real time processes that produce continuous data streams: an E1 or a DS1 or a 64 kb/s digitally encoded voice channel for example, or digitally encoded video without compression or packetization. AAL1 provides for two modes of operation: unstructured and structured. The unstructured mode treats the incoming information as a bit stream with associated clock with no regard to framing or format and passes that stream and clock information intact through the ATM process to the other end of the connection. The structured process recognizes the underlying format of the bit stream, in effect providing E0 visibility into a E1 data stream for example. In dealing with E1s in the structured mode, the SSP would detect and remove framing of the E1, delivering just E0s to the CP. For unstructured operation the SSP is null and data passes directly to the CP.

Looking at the more common unstructured service, data and clock information are passed to the convergence sublayer which divides the data stream into 47 byte PDUs. At the same time the incoming clock is passed off to a different process, the Synchronous Residual Time Stamp generator. The SRTS process compares the clock of the incoming data to a network reference clock and encodes the difference in a four bit word. This difference is passed, one bit at a time with each 47 byte CS PDUs to the SAR sublayer.

The SAR sublayer generates a one byte header and adds it to the 47 byte CS PDU, to make the 48 byte SAR PDU, which is the payload of an ATM cell. The SAR header contains one of the SRTS bits and a three bit sequence number field (the SN portion of the header), a three bit CRC (cyclic redundancy code) error check on the first four bits of the header and a parity bit (the SNP portion). The 48 byte SAR PDU is then passed to the ATM layer, where the 5 byte cell header is added. This process is sketched out in Figure 8.2, below for a DS1.

Figure 8.2  
AAL1 Processing  
a DS1 Data  
Stream in the  
Unstructured  
Mode



Note that the bits of the DS1 are streamed into the SAR PDU payload without regard to the framing of the DS1. The SAR PDU payload holds 376 bits which is just shy of two 193 bit DS1 frames.

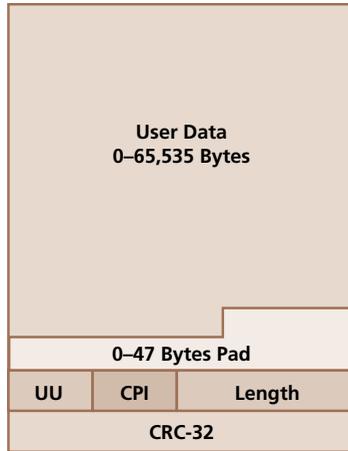
AAL1 places more emphasis on sequence numbering than might be thought necessary since ATM itself guarantees strict cell sequencing. Sequencing within AAL1 is not used for reordering, but to detect cell loss. In transporting continuous, real time information like DS1s or E1s there is normally no higher layer to detect lost data—bits are just streamed directly into the AAL process. Higher layer processes usually detect lost data and request retransmission. In continuous processes, there is no way to retransmit since the data isn't stored in memory somewhere, it's flowing in real time. With no way to recover, detecting a lost cell is more critical, since the external process must be alerted that a large block of data is missing, which will throw off framing and potentially cause mismatches in source and destination DS0s for example. The external process can apply what ever recovery mechanisms at its disposal (inserting 47 dummy bytes to maintain framing integrity for example). Loss of a cell also throws off the SRTS process and it must also be alerted to the fact that one of the SRTS bits has been lost.

### AAL5

The purpose of AAL5 is to provide the most streamlined transport possible for large blocks of data. AAL5 handles up to 65,535 bytes of data, and because of the overall assumptions ATM makes about error rates and sequencing, little overhead is added to the user data, making AAL5 very efficient. Like AAL1, AAL5 consists of convergence and segmentation and reassembly sublayers.

The convergence sublayer accepts user data, up to 64 kB at a time and forms one massive CS PDU as shown below in Figure 8.3

**Figure 8.3**  
**AAL5 Convergence**  
**Sublayer PDU**



There is no header applied to the data. At the end, 0-47 bytes of padding are added to bring the overall PDU length to an integral number of 48 byte cell payloads. A trailer of two 32 bit fields follows the padding. The trailer contains a one byte UU field, intended for user to user communications. Exact usage of this field is implementation specific. The CPI byte is undefined at this time and is included mainly to make the trailer 64 bits long. The next two bytes include the length of the user data field so padding can be deleted at the receiving end. The final 32 bits of the trailer make up a 32 bit CRC code which is calculated over the entire contents of the PDU that precede the CRC field. (We mentioned CRC coding in conjunction with AAL1 also. CRCs are useful because they can detect single, double, multiple and most burst errors in a data stream. As such they are much more likely to pick up errors than simple parity checks.)

For AAL5, the SAR sublayer is simplicity itself. The SAR sublayer accepts data from the CS, segments it into 48 byte cell payloads and passes it to the ATM layer for encapsulation in a cell. No SAR headers or trailers are generated or decoded at the receive end as with AAL1.

### **ATM's Advantages**

Beside carrying different service types, ATM offers several practical advantages for personal broadband services. In remote or small office environments, several users are apt to be using network based services simultaneously. Using ATM it is a simple matter to interleave or multiplex these different user data streams into a

common connection with no impact on any of the users. The small, uniform cell insures that one user's file transfer doesn't interfere with another's conference call. This is especially important in the small office environment where the price of bandwidth is important and such problems can't be solved by buying more or bigger connections to the network.

The same is true even for the single user in a home office. Several different tasks can be initiated and carried out simultaneously in the fore- and background, and again ATM insures equitable treatment for all. A real time video program can be viewed or recorded at the same time one or more file transfers are taking place in the background.

ATM, with its signaling protocols, also provides the ability to establish direct, high speed connections to other users or corporate locations as simply as making a phone call. In this respect it is similar to ISDN. Conventional datacom protocol suites such as TCP/IP on HDLC or PPP cannot provide this capability because they are inherently connectionless. While e-mail and Web-based applications and other IP-based connections will probably make up the bulk of our personal broadband service utilization, the ability to bypass the Internet altogether, and establish more secure, direct point-to-point connections with vendors, customers or the headquarters location is a critical, positive differentiator for ATM over ADSL.

The final advantage ATM offers over other options is that long term, it is where networking is headed. Increasingly the core of the Internet runs on ATM equipment, and ATM will gradually work it's way out to the network edges. Being directly compatible, rather than having to go through some interworking function, will make our personal broadband services more flexible, efficient and lower in cost.

## Summary

This has been a very brief overview of a rich and complex topic. We have not touched on areas such as switching, or management, and only briefly mentioned signaling. Even the areas discussed above only scratch the first layer of this powerful new technology. However, hopefully it is clear that with its ability to guarantee quality of service, provide low delay variation and strict data sequencing, and deliver speedy performance, ATM is the ideal choice for the voice and video applications that will become an increasingly important part of our personal broadband services.



At the same time, conventional datacom applications won't suffer in such an environment, and will probably receive better treatment than they are used to. The fact that ATM switches are used in rapidly increasing numbers in the Internet backbone is proof that ATM is more than adequate for data communications applications. And since the Internet is rapidly turning to ATM, using it as the link layer for personal broadband services will insure we receive the end-to-end quality of service required for the mixed bag of applications we will encounter in the near future.

# 9

## ATM Over xDSL: Putting ATM in the Loop



So how does this all work? Most people don't have a connector marked "ATM" on the back of their PCs; how do we connect to this new service? How does the service provider get service to you? How does the provider connect remote and small offices to the Internet and a corporate intranet using ATM/ADSL?

Connecting to personal broadband services will be the same as connection to Internet services today: either an external or plug-in internal modem will be used. If the modem is external, the interface to the PC will probably be Ethernet, implying an Ethernet adapter card in the PC unless it has Ethernet on the motherboard. During early trials of ADSL services, this has been the model. As a matter of fact, early trials have focused on Internet access substituting the above configuration for the typical 28.8 modem access. The Ethernet interface just serves as a convenient high speed port for the PC (serial modem ports are usually limited to little more than 100 kb/s).

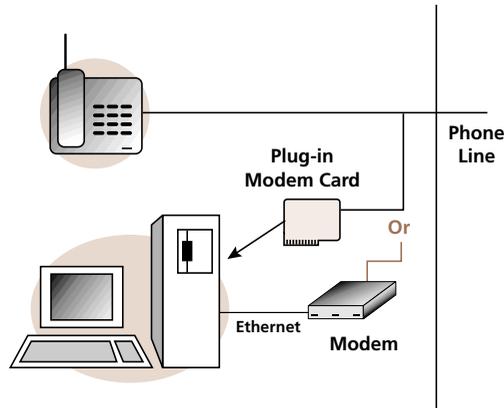
The problem with this configuration is that it leaves us in packet land. Short term, a packet-based protocol will be acceptable when emphasis will be on improved response for Internet access, general Web browsing, and connection to IP-based corporate intranets. Layer Two protocols like PPP can run over ADSL as easily as an analog modem. However, as we have pointed out, as real time multimedia content and applications grow, a packet-oriented Layer Two protocol has to give way to ATM. Making the transition to ATM over ADSL from today's packet access technologies for personal broadband services is easier than one might suspect. The key is the fact that large packets can be broken into multiple cells which can then co-exist with native cell-based real time traffic. Obviously, the converse is not true. And given the higher access speeds of ADSL and other cell-based networks, the inefficiencies of cell transport of packets, even when intermingled with real time cells, is unnoticeable.

In examining how we make this transition, there are two different applications to investigate: the single user home office/residential user, and small or branch

offices with multiple users. In either case, making the transition means bringing ATM directly into the PC.

For home office and single user applications, this involves outfitting the PC with an ATM/ADSL modem that plugs into the PC's PCI bus or USB, just as analog modems do today (or in the future buying a PC with a low speed ADSL modem on the motherboard). In such an installation, ATM is transparent to the user, contained entirely within the modem card. This is really the target market for splitterless implementations of ADSL such as G.Lite. As shown in Figure 9.1 below, the ADSL modem connects directly to the subscribers phone line just like the analog modem it replaces ... truly plug and play.

*Figure 9.1  
SOHO/Residential  
ADSL Installation*



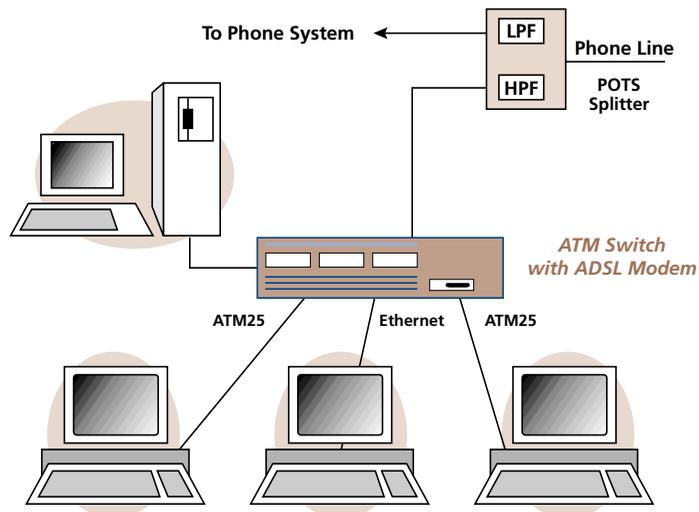
For a multi-user small office environment bringing ATM directly to the PC means replacing the existing LAN with an ATM-based LAN. Simply replacing the existing wide area network interface card in the office router with an ATM interface card is not enough. Trying to run real time traffic on the existing packet-based LAN in the office would defeat the purpose of converting to an ATM personal broadband service. Even though multimedia services would flow easily down the ADSL connection, they would get bogged down behind data packets on the LAN headed to a printer or coming from a local server. To realize the full benefits of ATM over ADSL, the small office or branch office local network must be upgraded.

Rather than an Ethernet or Token-Ring network interface card plugged into each PC, users will have an ATM NIC or better yet a combination NIC which does ATM and Ethernet or Token-Ring . An ATM switch will replace the router that provides the link between the office LAN and the outside world. Now before you think we've gone off the deep end here for a small office, we are not talking

about your typical ATM switch with OC-3 fiber interfaces to each work station and gigabit backplanes. The existing 25 Mb/s ATM to the desktop standard will work very nicely in this environment, using a couple of twisted pair to each PC.

A modest switch with rational backplane speeds will do just fine, and probably offer improved overall performance for the network. 25 Mb/s twisted pair ATM Network Interface Cards for PCs are priced quite competitively with Ethernet NICs today, and by 1998 small office ATM switches should be in the same price range as branch office routers.

Figure 9.2  
Remote Office/  
Branch Office  
ADSL Solution



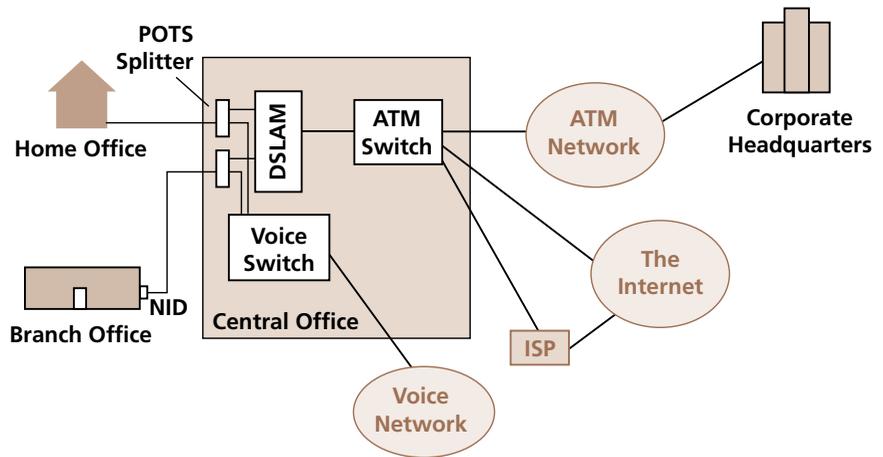
In the small office application the ADSL modem plugs into a port on the switch, just as routers have DS0/E0 or DS1/E1 plug-ins for a network interface. For small and branch office multi-user applications, a more robust set of protocols will probably be necessary than those that may ultimately be part of the more consumer-oriented Lite versions of ADSL. Higher downstream speeds than 1.5 Mb/s will often be necessary and more upstream bandwidth will be desirable. Here the more conventional versions of ADSL using a POTS splitter will be appropriate. This set-up is shown above in Figure 9.2.

This new small office topology has several advantages over what it replaces. The obvious is that our personal broadband service's cell format is preserved through to the desktop. In addition every user in this environment has a dedicated, full duplex connection at over twice the speed (in each direction) rather than hanging off a shared Ethernet or Token Ring. One user accessing a local server doesn't hang up

another using the ADSL access link—it's like the difference between party line and private line telephone service.

Now that we have your new office straightened out, let's take a look at what happens outside. The twisted pair from your network interface device runs to the C.O. Here another POTS splitter siphons off voice services and sends them to the local voice switch. Voice service works exactly as before in its own 0-4 kHz baseband slot with no knowledge that anything is going on above it. The ADSL output of the splitter passes to a port on an ATM concentration device, sometimes referred to as a DSLAM—Digital Subscriber Line Access Multiplexer (Some DSLAMs can also handle other types of traffic such as IP/PPP). This device concentrates the low rate upstream channels into one high speed ATM stream that matches the port speeds of network ATM switches, typically 34, 45, or 155 Mb/s. In the reverse direction, the DSLAM accepts a high speed cell stream from the switch (which may or may not be located in the same office), examines the address of each cell and sends it to the appropriate DSL modem plugged into the DSLAM. Figure 8.3 shows a diagram of the entire process.

Figure 9.3  
Data Networking  
with ADSL



Although Figure 9.3 shows them as separate pieces, the central office POTS splitter may be a part of the DSLAM, and in some cases, all this hardware may actually be contained on the voice switch line card itself with the concentration of the traffic from many ADSL modems performed by a new common card in the switch line card shelf.

Through its connection to the DSLAM, our personal broadband service is now connected to the world-wide true broadband data network. This lets us access the Internet or “dial” into corporate intranets. Remote access to intranets is quite similar to accessing them with analog modems, but faster. The cells originating in a PC are routed through ATM switches and delivered to a corporate access concentrator just as analog modem calls are routed through voice switches to reach the same point. Because this process bypasses the Internet entirely, the level of security is the same as offered by today’s methods.

Another interesting thing to note about the model shown in Figure 9.3 is that we are always connected. Accessing the Web is as easy as launching your browser; no more waits while modems dial, ISPs connect, no more busy signals. E-mail delivery is announced, not discovered when you happen to remember to look. From the user’s perspective, personal broadband services look just like being connected to a LAN in a physical corporate setting. And that was one of our goals: to allow branch and remote office workers to function as seamlessly as if they were located in the headquarters location.

## Networking

A concern of some may be the availability of ATM network services and interfaces either at the corporate or ISP interface. Despite initial excitement, ATM has not penetrated the network as quickly as many had predicted. What is the point of basing our personal broadband services on ATM if we cannot connect to our destinations or find a network to reach them?

First, we chose ATM as the Layer Two protocol to future-proof personal broadband services, not because ATM was the preferred networking paradigm today. However, every indication is that it will be in the future. As we have pointed out, the Internet increasingly runs on ATM; local service providers are building regional ATM networks and interexchange carriers all offer ATM service. ATM is not ubiquitous, but it is growing and there is nothing else on the horizon that looks like it will challenge it as the network of the next century.

So what do we do in the mean time? Short term, the initial applications for ADSL will be to provide much higher access speeds to the Internet and corporate intranets. The Internet backbone may be based on ATM, but most local ISPs aren’t—how do we connect to our favorite ISP? Corporate data networks run mostly on private leased lines like DS1s or E1s or public frame relay networks, not on ATM. The common denominator of all these networks is that they are packet based, yet our personal broadband service is cell based—how do we glue the two together?

The short answer is it is going to require a cooperative effort on the equipment in various parts of the network: on the user side, the service providers' portion, and the destination equipment. The important point is that ATM is only a Layer Two protocol, and so is significant only over one link, not end-to-end. We can shift from one Layer Two protocol to another at any point necessary to interface with a different network, ATM to frame relay for example. Conversions like these are made all the time in networks today, for example when a LAN packet is transferred to a DS1 to reach a server attached to a LAN at a different location; in this case the conversion might be from Ethernet to HDLC to FDDI, but the principal is the same. The only real issue is who does it, and where.

Another important thing to note is that we are changing nothing above Layer Two—this is the benefit of the well defined layers in a protocol stack. Even though we may be delivering our TCP/IP Internet Packets over ATM to an ISP, all he is really concerned with is the IP layer; if he can recover the IP packets he can do his job regardless of whether they came by ATM, dial-up point to point protocol (PPP), or Federal Express.

The same holds true for reaching corporate networks. Here we may be running IPX, and someone may have to interwork ATM and frame relay along the way but to the end to end process, accessing data on a corporate server from a branch office for example, the interworking is invisible.

Let's look at a couple of examples. Probably the most widespread application for ATM over ADSL will be Internet access. Very few ISPs have ATM networks—most are local operations that have modems on the customer side and lease a few high speed lines to a larger commercial ISP who connects them into the backbone Internet. Today this is a totally packet based operation using PPP as a Layer Two protocol on the user side almost exclusively. We will make the assumption that the user receives ADSL services from the telephone company, not the ISP. This application is diagrammed in Figure 9.4.

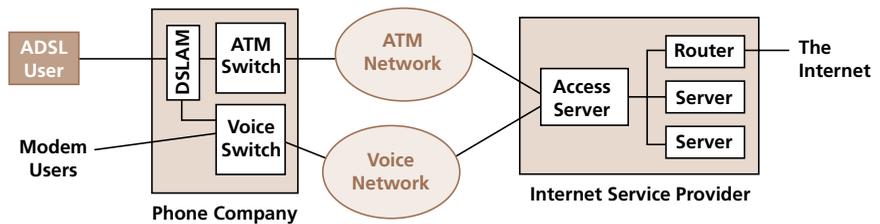


Figure 9.4  
Accessing an ISP

The user is connected to the DSLAM as usual. As part of the service agreement (especially in the early stages), the user designates an ISP, and the phone company establishes a permanent virtual circuit (PVC) between the user and the ISP. The user is given a specific address to place in the cell header that corresponds to the ISP. Each ATM switch between the user and the ISP is programmed with that address so that as user cells are generated with that address they are always switched to the ISP. Thus the user always has an open line to the ISP. The user's communications program in the PC delivers IP packets to the ATM/ADSL modem card which runs each packet through an AAL5 process then segments them into cells, each with the PVC address mentioned above.

The ISP is used to receiving user traffic from the voice switch over trunks connected to its access server. The access server holds the modems and recovers IP packets from the PPP frames used over the phone line and routes them to a local server or the ISP's backbone router and on to the Internet. The job of interworking probably falls with the ISP in this case (although if the telephone service provider and ISP are one and the same, the DSLAM could perform the function). The ISP must install an ATM interface card in the access server, or a whole new server which supports ATM interfaces. In either case, the interface card will recover packets from the ATM cells, reassembling them using AAL5. One of the more important functions of this card is to map between each cell address and the local Layer Two addressing used within the ISP LAN for each PVC. The incoming ATM stream from the phone company will typically have many simultaneous users embedded in it.

This same model represents a home office/telecommuter accessing a corporate network. We just change the label "Internet Service Provider" to "Corporate Communications Center" in Figure 9.4. Remote users used to dial in; now they use ATM. In this case the ATM network cloud may be a little larger. Whereas ISP access is local or regional, reaching a corporate network could involve several states and carriers. However, unlike dialing in with an analog modem or ISDN, the telecommuter is always on line to the corporate location; the PVC is always in place, the various ATM addresses always programmed in the switches, whether the user is actively communicating to the corporate offices or not. In effect the user is just as attached to the corporate LAN as any user resident in the company's physical location. This is handy for receiving e-mail as soon as it is sent.

In this model, the home office worker is not just connected to one corporate facility, but to the whole corporate network, even if it is based on high speed private lines and IP. In a corporate network, the router in Figure 9.4 would have several leased lines connecting to other locations. An IP packet originated by the home

office worker can reach that router and thus other locations using the LAN just as easily as a locally originated packet.

We have now given our home office worker two PVCs: one to the ISP and one to the corporate network. More can be added to reach important destinations as required, although once Switched Virtual Circuits become available, they will be preferred for reaching destinations that require only occasional access.

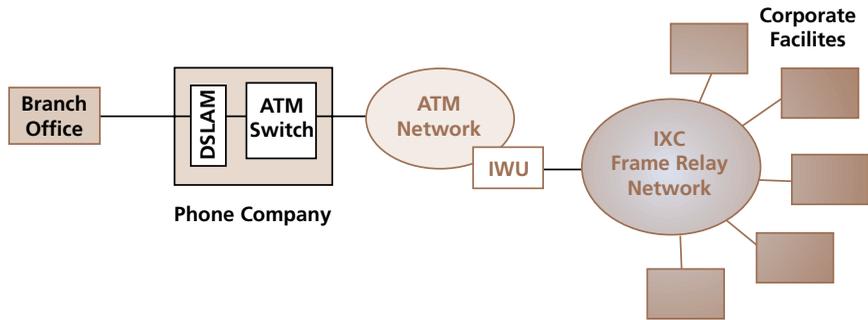
Actually, early on in the introduction of personal broadband services, the process will be a little more cumbersome than just substituting ATM for PPP for Layer Two. PPP provides some important adjunct services to the communications process that can't be ignored: authentication and security. Eliminating PPP makes it difficult for the ISP to authenticate users by password until a lot of software is rewritten to pass the security protocols over ATM instead of PPP. Since we want to make the introduction of personal broadband services as painless as possible for all concerned, a process called tunneling will probably be used until there is a significant enough body of users to justify the changes necessary to access servers. Cumbersome as it is, however, tunneling has the advantage that no software at either end of the connection has to change—that's pretty painless.

Tunneling involves taking a Layer Two frame, like a PPP frame, and transporting it inside a different Layer Two protocol like ATM. Normally the software in a PC processes information into Layer Two frames, usually using PPP, and passes them over the PCI bus to a modem card. The modem card accepts the data and modulates the modem. With tunneling, we would substitute an ATM/ADSL modem. The new modem would take the exact same PPP packets as a normal modem and run them through AAL5 to create an AAL5 PDU, then break it up into cells. At the other end, the process is reversed and the PPP frame retrieved. The net result is that what the access server sees coming out of the ATM interface card is exactly what it sees coming out of all the analog modem cards in the server. And the communications program in the user's PC does what it always did: wrap IP packets in PPP frames. The whole process is transparent and requires only hardware interface cards to servers on the part of the ISP or in a corporate network.

As a final example, let's look at a small branch office that needs to be connected directly to a corporate frame relay network as shown in Figure 9.5

In this situation, a PVC would be established from the branch through the ATM network. This would could involve both the local service provider and an interexchange carrier. At some point the ATM network touches the frame relay network, and here an Interworking Unit (IWU) is used to translate between the two network types. The IWU can be part of either the ATM switch or the frame relay switch. Interworking between the two services is not a trivial matter, and

Figure 9.5  
Tying into a  
Corporate Frame  
Relay Network



involves more than just address translations. However, the ATM Forum and the Frame Relay Forum have developed a service level interworking function that allows ATM end equipment to communicate transparently with frame relay end equipment without either having to map its messages into the other's protocol. The process is performed in the network by the carrier or carriers as traffic crosses the edge of the networks. Therefore the branch office using ATM/ADSL can be easily brought into the corporate frame relay network without sacrificing the forward looking features of our personal broadband services.

Thus starting our new service out on the right foot by using ATM does not make it incompatible with the bulk of the world running on different protocols. It will require some new hardware in specific parts of the network, but it is perfectly feasible to interwork the new service with existing network services. And as real time multimedia applications begin to emerge and ATM networks become prevalent, we are well positioned for both.

# 10

## Reality Check



All this sounds terrific ... but how real is it? Are we faced with another case of over-hype and under-delivery? Now, in early 1998, we certainly are in ADSL's hype phase but there are signs that follow-through will be fairly swift. Probably the most important fact that there is demonstrable demand for services such as we have described, if only for faster Internet access. There are hundreds of thousands of small businesses around the world that are ready, willing, and able to pay for better than analog modem access. Uniformly, anyone who has participated in an ADSL trial asks "when?" not "how much?" This is in marked contrast to many earlier video-based broadband service trials where enthusiasm evaporated when subscribers were asked to pay. Unlike a lot of other over-hyped technologies, personal broadband services are not a solution looking for a problem.

As far as "when" is concerned, the answer is now. Several ISPs, as large as UUNET and as small as HarvardNet (a local Boston concern, not affiliated with Harvard University) are delivering DSL service today in the U.S. using dry copper. Other enterprising types are using in-building copper in large office and apartment buildings to offer ADSL access speeds to small businesses or residents. The building's phone lines pass through a DSLAM located in the basement and a conventional high speed leased line returns the traffic to the ISP's operations center. These are very geographically limited offerings, targeted mainly at business users, but they are tariffed services, not trials. Meanwhile, most large telephone companies have conducted DSL trials, and some are moving beyond the trial phase. Several, including SBC, Ameritech, and US West have tariffed service offerings in selected cities with promises of expanded availability in 1998.

One of the most far-sighted deployments of ATM over ADSL is the Singapore ONE network of Singapore Telecom. This network supports broadband services to PCs and TV set-top converters over an ADSL local connection. ATM was selected as it allows multiple services to be delivered over a shared ADSL enabled access link.

The other factor that could foster early deployment of DSL services is increased competition in the local loop—assuming, in the case of the U.S., the Telecom Reform Act isn't in litigation for the next ten years. With competition for local service, embedded service suppliers will be looking for ways to wring more money out of the existing plant and ways to retain customers by offering value-added services. Competitors will be looking for ways to differentiate themselves from established telephone companies. Either way, given the penetration of the Internet into small businesses and residences, packages offering enhanced Internet access services are apt to be high on the agendas of all competitors.

Although DSL-based service is beginning to appear in 1998, the journey to universal personal broadband services will be slow and not without some pain. There are over 15,000 central offices in the U.S. alone; equipping them with DSLAMs is not going to happen overnight. Telephone companies will target initial deployment at central offices serving populations where they anticipate a reasonable take rate: cities and upscale suburban locations. Telecommuters and home office workers who chose to locate in more pastoral locations will probably experience a certain degree of frustration as more metropolitan areas receive higher priority.

Users fortunate enough to be located in early deployment areas may experience other frustrations. As they watch the little numbers in the lower left hand corner of their browser screen during a download, they may find values significantly less than the 150 or 200 Kbytes/sec their 1.5 Mb/s connection is supposed to deliver. Downloads will be much faster to be sure, but don't seem to be living up to expectations. The thing to remember is that the DSL-based connection is only one small part of the Internet. By bringing personal broadband service to millions, we are increasing their ability to sink information by a factor of 30-50. The Internet just isn't sized to support that kind of demand today.

A typical local ISP ties into its Internet service provider with one or two T1s or E1s and supports a few thousand analog modem customers with this link. Statistically only 10-15% of the customers are on line at any one time, they are receiving information at a low rate compared to the T1 links, and the varied applications (not everyone accesses just the Web) and bursty nature of Internet use supports this model. With DSL, the subscriber connection rate equals the rate of the ISP's link to the Internet. Mix in ISPs who may be using only a 10 Mb/s Ethernet to tie their equipment together in their operations centers, and probably aren't using top of the line high speed routers. Sprinkle in the fact that many Web and news group servers are also not necessarily the fastest in the world and are connected to the Internet by 56 kb/s or T1 lines themselves. Finally, strain this mixture through an Internet backbone that is growing ten times faster than the voice network but



still can't keep up with exploding demand. The result is a blend of the performance of the whole network, not just our DSL link. Subjectively, performance is going to look a whole lot better; objectively, many will probably see a net gain in throughput of five to ten in the early period of personal broadband services.

Now five or ten isn't bad, and it will only get better. If it doesn't, customers are going to find ISPs who are willing to lease higher speed links to their Internet provider and upgrade their internal networks. The major backbone Internet carriers already have massive network upgrades underway, shifting from 45 Mb/s DS3s to 155 Mb/s OC-3s or 622 Mb/s OC-12s. ATM switching is replacing routers in the core network. Vendors and content providers who want traffic at their Web sites will have to invest in faster servers and connections to the network.

What is important is not to get caught up in the hype and set your expectations too high. Reality is that the benefits of personal broadband services are going to evolve into a mature product over several years. Every link in the chain from user to content is part of the problem, not just our venerable analog modems. Everyone has their work cut out to reforge their links, but in an increasingly competitive environment with an exploding customer base, there is a lot of incentive.

But reality is also that for the first time, we can build a rational, profitable business case for delivering broadband services to branch and small offices as well as residences around the world. Given that reality, personal broadband services will come as quickly as service providers can deploy hardware.

# 11

## Virata Solutions



The technological challenges and significant time to market pressures facing telco and cable equipment suppliers have made it increasingly difficult for them to develop internally all of the required ATM components for their products in a timely and cost effective manner. As a result, many telco and cable equipment suppliers desire to acquire, from a third-party technology provider that has the requisite ATM, software, component and telecommunications system expertise, the ATM software and semiconductor devices that they need to enable their equipment to connect end users to the local standards based subscriber network.

### Introduction

Virata is the leading developer and supplier of Integrated Software on Silicon (ISOS™) solutions which enable the personal broadband solutions you have been reading about. The Company's ISOS solutions are employed by licensees to develop and market products aimed at the needs of the Digital Subscriber Line (xDSL), Fiber To The Curb (FTTC), Cable Access Television (CATV) and wireless markets. The two Application Specific Standard Product (ASSP) families currently shipping are the Proton chip set comprising four chips which operate with a separate Advanced RISC Machines (ARM™) microprocessor and the ATOM family, the members of which feature one or more embedded ARM microprocessors. In addition to these semiconductors, Virata licenses its extensive Universal Software offerings to its customers. Many of Virata's licensees find that this software and hardware combination provides up to 95% of the solution required for the initial development of a complete user-deliverable product. The Company believes that broadband local loop equipment suppliers using ISOS solutions are able to design, develop and market new products faster, with higher confidence and at lower design, development and product costs than would be possible any other way.



## Virata's ISOS Solution

Virata is the leading developer and supplier of Integrated Software On Silicon (ISOS) solutions enabling a new generation of cost-effective, high performance broadband equipment. These ISOS solutions combine high performance, programmable ASSPs with an integrated multi-protocol software stack that make it possible for broadband equipment vendors to quickly develop and market feature-rich, upgradeable, differentiated, low cost products suitable for application in the local loop. The Company believes that by designing-in ISOS solutions, broadband equipment vendors can simultaneously reduce development costs, accelerate time-to-market, reduce product costs, while differentiating their products with advanced features and functionality. Since a very large percentage of the required solution is delivered by Virata, there only remains a single code integration of the licensee's unique software with the Universal software stack in order to provide a product.

Virata's customers are employing ISOS solutions to design a wide range of equipment destined for application in xDSL, FTTC or CATV broadband environments, including: line cards for CO switches, channel banks and DSLAMs; DSLAM switching fabrics; CATV head-ends; LAN and WAN edge switches; single and multi-drop modems and multi-protocol NICs with enhanced features including auto-sensing for delivery of Ethernet packets or ATM cells; Optical Network Units (ONUs), Broadband Network Units (BNUs); cable modems and settop boxes.

A typical topology for broadband installations associated with the three technology markets is shown below.

	<b>Backbone Infrastructure</b>	<b>Distribution Hub</b>	<b>Customer Premises Equipment</b>
<i>xDSL</i>	Class 4/5 Switch	xDSL Line Card Distributed Loop Carrier DSLAM	ATU-R, Modem, Telephone, Hub, NIC, Broadband Gateway, NIM/NT
<i>FTTC</i>	ATM Switch	Optical Network Unit	BNU Multimedia PC
<i>CATV</i>	Satellite Dish	Cable Head End	Cable Modem Set Top Box

The products which comprise the ISOS solution range currently offered by the Company include: ASSPs from the Proton and ATOM families and the Universal software stack. The Proton four chip set, together with an ARM microprocessor, provides all the main components needed to construct an ATM switching fabric. Members of the ATOM family feature one or more embedded ARM processors and are ready for use in line cards, modems, NICs and settop boxes.

The Universal Software is closely coupled to work efficiently on each of the chips and includes ATMOS™, a real time ATM operating system, an comprehensive ATM driver, ATM Forum compliant signaling, ATM Forum compliant LAN Emulation software, Classic IP (RFC 1577), bridging to Ethernet (RFC 1483) and a light weight form of IP routing. Additional modules and various upgrades are offered as a standard part of the Company’s on-going support and maintenance programs. Documentation and debugging tools are included. In addition, as a result of the Company’s systems heritage, proven reference designs and development

platforms are available to accelerate development and testing of new products.

The overall package of hardware and software is depicted in Figure 11-1.

Since the needs of one customer may not be exactly the same as the needs of another, the ISOS solution has been packaged so that developers may pick and choose the set of application code and hardware platform that best fits their needs. Thus, ISOS has several modules, with more to come.

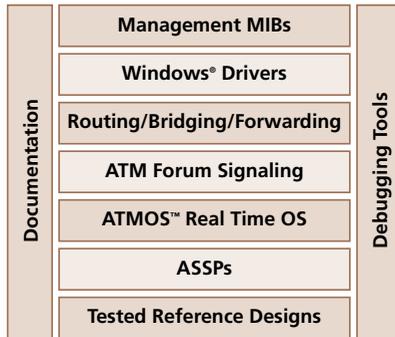


Figure 11.1  
Virata Solutions

The initial packages are depicted in chart 11-2.

Each of the modules shown above has the unique set of software inside that will provide the developer just the elements necessary to build the generic types of equipment listed in each module title. Of course, most licensees substantially enhance the capabilities of these modules as they develop their own unique value propositions. Thus, no two products built with the underlying ISOS technology set will ever look exactly alike.

On the other hand, each product constructed with an ISOS module will have the advantage of knowing it will interoperate with most of the world’s leading

Modem Package	NIM Package	Line Card Package	Switch Fabric Package
Signaling	Signaling	Switching & Signaling	Switching & Signaling
Bridging	Bridging	High Speed I/O	LANE Services
Routing	LAN Emulation	Traffic Shaping	IP Services
Classical IP	Winsock 2.0	Policing	Routing
USB	IP Services		Bridging
Ethernet	PPP		High Speed I/O
PPP/Tunneling	PCI		SNMP
	Windows® Drivers		

Hardware Reference Designs Linked to Licensed Packages Above

Figure 11.2  
Virata Licensing  
Modules

ATM and IP switched or routed environments, since these modules have been extensively tested in real networks throughout the world.

## Application Specific Standard Products—ASSPs

The company has two families of semiconductor devices. The Proton family, included a variety of semiconductor devices that provide SAR engines, Switching fabric, and I/O device support function on a standalone basis. The ATOM family achieves a high level of integration by including all non-physical layer functionality for a broadband local loop delivery subsystem into one integrated device.

The Proton ASSP family provides a complete ATM switching fabric, suitable for LAN, WAN edge, xDSLAM and CATV head-end switching applications.

### *Proton ASSP Family*

Product	Application	Date of Introduction
Gluon	ATM Forum CRC generator and ARM microprocessor support	April, 1995
Quark	ATM cell buffer and processor for switch or adapter network ports	April, 1995
Hadron	ATM cell address hasher for switching applications	January, 1996
TBX	ATM traffic shaping controller and buffer for switching applications	August, 1997

The ATOM family features one more embedded ARM RISC processors. Members of the family are summarized in the following table:

### *ATOM ASSP Family*

Product	Application	Date of Introduction
Hydrogen	Line cards, modems and NICs	August, 1997
Helium	Multi-Drop line cards, modems, WAN edge switches	September, 1998

At the heart of the Company's solutions lies a tradition of semiconductor design and integration. The current family of ASSPs are built using an extensive library of complex functional blocks invented for both this family as well as derived from two previous generations of ASSPs. The latest generation of ASSP solutions benefits from the Company's experience in the implementation of numerous end-to-end systems and its extensive knowledge of embedded microprocessors.

# 12

## The Next Steps



### Learning More

The following provides a starting point for learning more about xDSL and related subjects:

#### *Web Pages*

The ADSL Forum: [www.adsl.com](http://www.adsl.com)

A collection of technical white papers and tutorials on xDSL, member and vendor lists linked to home pages, FAQs, Forum membership and meeting information, and a terrific matrix of xDSL trials around the world.

TeleChoice's xDSL News Post: [www.telechoice.com/xdslnewz/](http://www.telechoice.com/xdslnewz/)

Latest news stories, press releases and analysis covering the entire DSL field, updated weekly or sooner; comprehensive list of vendors sorted by product type (chip, modem, C.O., system, etc.) technology (CAP vs DMT) linked to home pages; conference, trade show and seminar information.

The Universal ADSL Working Group: [www.uawg.org](http://www.uawg.org).

The new home page of the UAWG. Press releases, members list, and whatnot.

#### *Books*

Given the youth of ADSL, there is little in the way of full length texts yet available. A non-exhaustive search turned up only five, and four haven't been published as of February 1998!

Video Dialtone Technology; D. Minoli; McGraw-Hill, 1995

Good overall treatment of bringing video and other broadband services to the home; very light on DSL; more information on competing technologies such as FTTC, and HFC for those interested.

ADSL (Computer Communications); W. Goralski, McGraw-Hill, publication scheduled 2/98

ADSL/VDSL Principles, D. Rauschmayer, Macmillan, publication scheduled 12/97

Demystifying ATM/ADSL, M. Busby, Wordware, publication scheduled 6/98

Digital Subscriber Lines: Toward, Above, and Beyond ADSL; W. Chen; Macmillan, publication scheduled 1/98

DSL: ADSL, RADSL, SDSL, HDSL, VDSL; H. Hecht, J. Freeman, and M. Humphrey; McGraw-Hill, publication scheduled 1/98

### *Technical Literature*

The best source of material covering DSL is still papers in technical journals. Some of the following are very technical, some are of a more general nature. The list is hardly complete, but most of the articles are well referenced and can lead the reader further.

“How xDSL Supports Broadband Services to the Home,” M. Humphrey and J. Freeman, *IEEE Network Magazine*, Jan./Feb. 1997

#### *IEEE Communications Magazine*

“Systems Considerations for the Use of xDSL Technology for Data Access,” G. Hawley, Mar. 1997

“Frame Relay and ATM Interworking,” S. Dixit and S. Elby, June, 1996

“Asymmetric Digital Subscriber Line: Interim Technology for the Next Forty Years,” K. Maxwell, Oct. 1996

“ADSL: A New Twisted-Pair Access to the Information Highway,” P. Kyees, et al., April 1996

“Applicability of ADSL to Support Video Dial Tone in the Copper Loop,” W. Chen and D. Waring, May 1994

“Digital Subscriber Line Technology Facilitates a Graceful Transition from Copper to Fiber,” D. Waring, et al., Mar. 1991

“High Bit Rate Asymmetric Digital Communications Over Telephone Loops,” K. Kerpez and K. Sistanizadeh, *IEEE Transactions on Communications*, June, 1995

#### *IEEE Journal on Selected Areas in Communications*

“Digital Subscriber Line (HDSL and ADSL) Capacity in the Outside Loop Plant,” S. Ahamed, et al., Dec. 1995

“Broadcast Digital Subscriber Lines,” W. Chen, Dec. 1995

“A Discrete Multitone Transceiver for HDSL Applications,” J. Chow, et al., Aug. 1991

“The HDSL Environment” J. Werner, Aug. 1991

“Performance Evaluation of a Multichannel Transceiver System for ADSL and VHDSL Services,” P. Chow, et al., Aug. 1991

“High Bit Rate Digital Subscriber Lines: A Review of HDSL Progress,” J. Lechleider, Aug. 1991

## Making It Happen

### *The ADSL Forum*

If you are an equipment manufacturer or service provider a logical next step would be to join the ADSL Forum. This organization was formed late in 1994 and now numbers over 300 principal and auditing members. In addition to promoting and demystifying ADSL to the interested public, the ADSL Forum provides technical feedstock to relevant standards-making bodies, such as T1E1 in the U.S. and ETSI in Europe in the form of Technical Reports. Members work in committees to identify and recommend requirements in areas including interfaces, packet and ATM protocols, migration, and network management. More details about the Forum and membership information can be found on their Web page listed above.

### *ADSL Users*

Subscribers who would benefit from ADSL service should begin talking with their local telephone service providers and ISPs and express strong interest in obtaining service as soon as possible. Providers are still evaluating the market opportunity for ADSL and the more tangible demand they see, the sooner they will make a positive decision. Explain your applications and requirements to them, and show them how ADSL will make your company, and by extension most other businesses more productive. Press hard, and keep pressing; demonstrating sustained interest will be important.

### *Service Providers*

Just do it!



## Putting ATM to Work Over DSL

If you're ready to put ATM to work over ADSL, Virata is ready to help. Virata can deliver off-the-shelf chips from its growing family of highly integrated ATM controllers. We are also ready to discuss custom solutions to our partner's unique requirements.

For more information on our products, or to discuss your specific needs, please contact us.

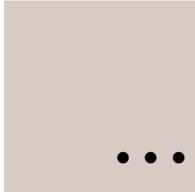
### **UNITED KINGDOM**

Virata Ltd.  
Mount Pleasant House  
2 Mount Pleasant  
Huntingdon Rd.  
Cambridge CB3 0BL  
Phone: +44 1223 566919  
Fax: +44 1223 566915

### **UNITED STATES**

Virata Corporation  
2933 Bunker Hill Lane, Suite 201  
Santa Clara, CA 95054  
Phone: +1 408 566 1000  
Fax: +1 408 980 8250

[www.virata.com](http://www.virata.com)



## Glossary

<b>2B1Q</b>	Modulation format used by ISDN
<b>ANSI</b>	American National Standards Institute
<b>ATMOS</b>	ATM operating system
<b>Baseband</b>	A range of frequencies starting at zero Hz and extending to an upper limit
<b>Baud</b>	Transmission rate of a multi-level coded system when symbols replace multiple bits. Baud rate is always less than bit rate in such systems
<b>Bps</b>	Bits per second—the raw data rate of a system
<b>CAP</b>	Carrierless Amplitude and Phase modulation (see Chapter 5)
<b>C.O.</b>	Central Office—where telephone subscriber lines meet the first switch
<b>crosstalk</b>	Undesired leakage of signals from one pair of wires into adjacent pairs
<b>distribution</b>	Portion of the telephone cabling plant that connects subscribers to feeder cables from the central office
<b>DLC</b>	Digital Loop Carrier—equipment used to concentrate many local loop pairs onto a few high speed digital pairs or one fiber optic pairs for transport back to the central office
<b>DMT</b>	Discrete MultiTone—a DSL modulation technique (see Chapter 5)
<b>Downstream</b>	References flow of information from the central office to the subscriber
<b>DS0</b>	Basic North American 64 kb/s digitized voice channel
<b>DS1</b>	First level in North American digital hierarchy; the 1.544 Mb/s signal consists of 24 DS0s multiplexed together
<b>DSL</b>	Digital Subscriber Line



<b>DSLAM</b>	Digital Subscriber Access Multiplexer—a central office device used to aggregate data traffic from many DSL subscribers into one high speed signal for hand-off to the datacom network.
<b>E0</b>	International basic 64 kb/s digitized voice channel
<b>E1</b>	First level in international digital hierarchy; the 2048 Mb/s E1 signal consists of 30 E0s carrying voice and two E0s used to carry signaling information for the voice channels
<b>ETSI</b>	European Telecommunications Standards Institute
<b>feeder</b>	That portion of the telephone cable plant that extends from the central office to distribution frames where distribution cables deliver traffic to subscribers
<b>FEXT</b>	Far end crosstalk—leakage of one or more foreign sources into the receiver of a system at the distant end of a transmission system
<b>FTTC</b>	Fiber To The Curb—a telephone company service delivery system that delivers voice and video programming to small clusters of residences using fiber optics as the feeder and either twisted pairs or coax cable as the distribution plant to each home
<b>HFC</b>	Hybrid Fiber Coax—a cable TV delivery system that uses fiber optics to feed a distribution node that delivers video traffic over coax cable to about 500 homes
<b>Hz</b>	Hertz—the basic unit of frequency measurement; one cycle per second
<b>ISP</b>	Internet Service Provider
<b>ISOS</b>	Integrated Software on Silicon
<b>ITU</b>	International Telecommunications Union
<b>Kb/s</b>	Kilobits per second
<b>KHz</b>	kilohertz—one thousand Hertz
<b>Loaded pair</b>	A twisted pair phone line with inductors, or loading coils, inserted periodically to flatten the frequency response in the 4 kHz voice band
<b>loop</b>	portion of the telephone network that connects the subscriber to the central office
<b>Mb/s</b>	Megabits per second
<b>NEXT</b>	Near End Crosstalk—leakage of undesired local signals into the local receiver; could be from the companion transmitter or other nearby sources
<b>passband</b>	A range of frequencies that has a non-zero lower limit and some upper limit

<b>PDU</b>	Protocol Data Unit—a segment of data generated by a specific layer in a protocol stack; usually consists of a block of data from a higher layer (the service data unit or SDU) encapsulated by the next lower layer with a header and trailer
<b>POTS</b>	Plain Old Telephone Service
<b>PPP</b>	Point to Point Protocol—a common Layer Two protocol used with Internet protocols and services
<b>PSTN</b>	Public Switched Telephone Network
<b>PVC</b>	Permanent Virtual Circuit—a dedicated path through a frame relay or ATM network intended for long term transport of information; the equivalent of a leased line
<b>QAM</b>	Quadrature Amplitude Modulation—a modulation technique used by modems and DSL equipment in which a carrier's amplitude and phase are simultaneously modulated by the digital traffic
<b>QoS</b>	Quality of Service
<b>SLIP</b>	Serial Line Internet Protocol—an older Layer Two protocol used for Internet traffic; much less sophisticated than PPP
<b>SOHO</b>	Small Office, Home Office
<b>T1</b>	Transmission system used for transport of DS1 signals; also called T-Carrier equipment; used interchangeably with DS1 by most people
<b>Upstream</b>	Referencing the flow of information from the subscriber to the central office
<b>vf</b>	Voice frequency. In telephony, typically the range is from zero to four kilohertz





## Index to Figures



<i>Figure 2.1</i>	<i>Modulation Constellations</i> . . . . .	<b>12</b>
<i>Figure 3.1</i>	<i>Copper and Digital Loop Carrier Subscriber Feeds</i> . . . . .	<b>25</b>
<i>Figure 3.2</i>	<i>Loss vs. frequency for 12kft and 18kft of 24 gauge wire</i> . . . . .	<b>28</b>
<i>Figure 3.3</i>	<i>Typical Copper Feeder/Distribution Loop</i> . . . . .	<b>29</b>
<i>Figure 3.4</i>	<i>Effect of Loading on Twisted Pair Frequency Response</i> . . . . .	<b>30</b>
<i>Figure 3.5</i>	<i>Data Capacity of a Local Loop</i> . . . . .	<b>31</b>
<i>Figure 5.1</i>	<i>HDSL Spectrum Use</i> . . . . .	<b>40</b>
<i>Figure 5.2</i>	<i>CAP ADSL Spectrum</i> . . . . .	<b>42</b>
<i>Figure 5.3</i>	<i>DMT ADSL Spectrum</i> . . . . .	<b>43</b>
<i>Figure 6.1</i>	<i>POTs Splitter Application</i> . . . . .	<b>50</b>
<i>Figure 7.1</i>	<i>Protocol Stacks</i> . . . . .	<b>54</b>
<i>Figure 7.2</i>	<i>Impact of Sample Delay on Reconstructing a Wave Form</i> . . . . .	<b>58</b>
<i>Figure 8.1</i>	<i>HDLC Frame and ATM Cell</i> . . . . .	<b>61</b>
<i>Figure 8.2</i>	<i>AAL1 Processing a DS1 Data Stream in the Unstructured Mode</i> . . . . .	<b>68</b>
<i>Figure 8.3</i>	<i>AAL5 Convergence Sublayer PDU</i> . . . . .	<b>69</b>
<i>Figure 9.1</i>	<i>SOHO/Residential ADSL Installation</i> . . . . .	<b>73</b>
<i>Figure 9.2</i>	<i>Remote Office/Branch Office ADSL Solution</i> . . . . .	<b>74</b>
<i>Figure 9.3</i>	<i>Data Networking with ADSL</i> . . . . .	<b>75</b>
<i>Figure 9.4</i>	<i>Accessing an ISP</i> . . . . .	<b>77</b>
<i>Figure 9.5</i>	<i>Tying into a Corporate Frame Relay Network</i> . . . . .	<b>80</b>
<i>Figure 11.1</i>	<i>Virata Solutions</i> . . . . .	<b>86</b>
<i>Figure 11.2</i>	<i>Virata Licensing Modules</i> . . . . .	<b>86</b>



Virata Corporation  
2933 Bunker Hill Lane, Suite 201  
Santa Clara, CA 95054 U.S.A.

*phone:* +1 408 566 1000  
*fax:* +1 408 980 8250  
*email:* [info@virata.com](mailto:info@virata.com)

Virata Limited  
2 Mount Pleasant House, Huntingdon Rd.  
Cambridge CB3 0BL, United Kingdom

*phone:* +44 1223 566919  
*fax:* +44 1223 566915  
*email:* [info@virata.com](mailto:info@virata.com)