Digital Copper for Delivering Converged Services

Donald B. Skipwith
Vice President, Business Development
Symmetricom

Kishan Shenoi
Chief Scientist
Symmetricom

Introduction

Convergence—or the ability to manage multiple applications over a single broadband facility—is the focus of today's telecom industry. The popularity of the Internet, and potential use of Internet protocol (IP) not just for data but also real-time services such as voice and video, promised new converged service offerings. However, the "killer application" that has emerged so far is electronic mail, a service offering that can be provided, for the most part, over dial-up links at kbps speeds. For the multitude of proposed applications, including multimedia entertainment, on-line gaming, file swapping, video and other services on demand, videoconferencing, distance learning, and telemedicine, it is essential that high-speed transmission and switching, at Mbps speeds, be made available in an economical manner.

Until recently, carriers and investors alike swarmed to fiber optics as the ubiquitous solution. Equipment vendors such as Alcatel, Cisco Systems, Cienna, Juniper Networks, Lucent Technologies, Nortel Networks, and Tellabs, to name just a handful, developed optical solutions to make the backbone networks truly ready for high-speed services and applications requiring Mbps bit rates. The Achilles heel, however, has been, and remains, the access network. Optical fiber, having found its home in network backbones, was recently emerging in the metro and last-mile networks. The dot-com implosion and associated telecom meltdown has led to many company closings and dramatic reductions in capital expenditure (CAPEX) by the surviving service providers. Carriers are now searching for other techniques to provide broadband connections over the last-mile network segment. As the ephemeral vision of access networks transitioning from copper to fiber has fades from view, the local-exchange carriers (LECs) must, out of necessity, turn their focus to leveraging existing assets to provide new high-value services. Copper, the telephone company's largest asset, and associated central office (CO)–based digital subscriber line access multiplexers (DSLAMs), offers a simple and economical solution for high-bandwidth converged services in the network access segment. Approaches requiring "forklift" upgrades are likely to be shunned, whereas incremental deployments are favored.

In the following, we pose the problem of access, a promising solution, and some representative applications.

The Problem

Access, or the link from the telephone company's ("telco's") CO or remote digital loop carrier (DLC) to the end user, continues to be the choke point in end-to-end broadband networking. A surfeit of fiber capacity exists in the long-distance, network backbone. It is how to get on to these core fiber networks using copper cable access that must be addressed.

The majority of residential customers use dial-up facilities to access the Internet. This poses a different kind of problem to the telco. Circuit-switched networks were designed for handling telephone (voice) calls, typically minutes in duration. Internet-directed calls tend to be of much longer durations, tying up valuable switching resources. The deployment of asymmetric digital subscriber line (ADSL) technology has gone a long way toward providing "always-on" connectivity for residential users and is the telco technology of choice for serving the residential market, where cable companies provide stiff competition.

Both ADSL and cable-modem approaches are suitable for residential customers, characterized by high downstream bit rates (network to end- user) and low upstream bit rates (end user to network). Business customers, on the other hand, have a different set of needs. These include symmetric bit rates, high quality of service (QoS), high reliability, and enforceable service-level agreements (SLAs). Moreover, the access mechanism must support the variety of voice and data needs of the business. Fiber-to-the-premises has long been the panacea, providing unlimited access speeds, but given the current economic environment, fiber "everywhere" may well be the promise of the next decade.

RHK, a leading market analysis firm based in California, believes that copper is the key to providing ubiquitous and economical delivery of high-speed services—2 to 15 Mbps—to businesses, particularly in the small-to-medium enterprise (SME) community. In a recent multi-client study, RHK determined that out of the 765,000 office buildings in the United States, 622,000—more than 81%—were not be served by fiber at the end of 2002. RHK estimates that by 2005, more than 63% of the buildings will remain unserved by fiber. Furthermore, while fiber deployment is more prevalent in "greenfield" situations, copper loops are available in 100% of buildings, both old and new.

The "problem" is thus one of devising methods that leverage the existing telco infrastructure and provide access bit rates in the 2 Mbps to 15 Mbps range supporting a variety of traffic types.

Today versus Tomorrow

Today, E1 (2.048 Mbps) and T1 (1.544 Mbps) services are the workhorses of corporate networking in the European Community (E1) and North America (T1). E1/T1—provided by copper cables—is as fast as it gets unless a direct fiber facility is available. However, fiber connections only make sense in very large business buildings, so serving the SMEs can be a problem. The older, smaller buildings that typically house the SMEs may be "passed by fiber" (fiber is available out on the curb), but neither the economic incentives nor competitive pressures exist to drive the LECs to deliver fiber directly to these SMEs. Furthermore, E1/T1 services provided by the LECs are very lucrative. The answer given to the SMEs that only have copper access and need additional access bandwidth is "buy another E1/T1 circuit." Consequently, more bandwidth is available, but the burst rate is still limited to 2/1.5 Mbps. Individual T1s (E1s) typically limited to a particular service—such as local telephony, long-distance telephone, and/or Internet access—are nice, but multiple pipes provide increased bandwidth, at a high cost, that is not shared across applications. Nevertheless, there is increasing competition in the access space with the incumbent LEC (ILEC) being challenged by other licensed carriers and competitive LECs (CLECs) to offer new, high-margin, converged services to the SMEs that depend on copper access.

Whereas fiber may indeed be the ultimate solution for high-speed access, the near-term approach that achieves the desired goals in a cost-effective manner is a relatively new DSL technology—G.shdsl and associated multi-link techniques to provide copper-based "fat pipes."

G.shdsl

A new international standard for single-pair, high-speed DSL (ITU–T G.991.2, aka G.shdsl) was approved in 2001. Unlike currently deployed ADSL, G.shdsl delivers a symmetric (same speed, both directions—required by the business applications) 2.3 Mbps bit rate up to 10,000 feet, or approximately three kilometers, from the telephone serving office on a single loop. It is rate-adaptive, and the standard allows for variable rates optimized for distance. By design, G.shdsl is spectrally friendly and quite tolerant to line impairments, allowing it to coexist in the same access cable binder groups with residential ADSL service. Unlike ADSL, the G.shdsl also spells out the requirements for repeaters (regenerators). These devices are deployed in the outside plant and a single repeater can, effectively, double the reach at a given bit rate. Whereas the standard allows more, it is generally accepted that deployment of more than three repeaters in a loop is unlikely. With three repeaters the reach for 2.3 Mbps service can, in principle, be extended to nearly 12 kilometers from the serving CO using a single subscriber loop.

All digital transmission schemes exhibit a "rate-reach" profile, and thus fixed-rate transmission schemes, such as high-bit-rate DSL (HDSL), require the deployment of repeaters/regenerators every so often (approximately every 9,000 feet for HDSL). Unfortunately, repeatered DSL has a deleterious impact on ADSL services carried in the same binder group. Consequently, E1/T1 and HDSL deployment requires some administrative oversight in the form of binder management. With G.shdsl, multiple loops (without repeaters) can be "bonded" together to get the requisite bit rate. Such multi-link techniques can be used to provide "fat pipes" on the one hand or longer reach at a given bit rate on the other. This obviates the need for binder management.

From a deployment perspective, all modern DSLAMs are capable of supporting a mix of ADSL and G.shdsl by installing the appropriate port blade. This goes a long way toward having a unified management system for operations, administration, maintenance, and provisioning (OAM&P).

In the past, telephone service providers have focused on providing ADSL services to the residential and small office/home office (SOHO) market. The emergence of G.shdsl that can coexist with the ADSL infrastructure now allows a symmetrical business DSL service burstable to 2.3 Mbps on a single loop. Additionally, the ability to combine several G.shdsl physical links into a single, high-speed logical link provides burstable bandwidth upward of 15 Mbps. Applications such as high-speed Internet access, transparent local-area network (LAN) services, and converged voice and data services can now be serviced using high-speed, multi-link, G.shdsl access.

Multi-Link G.shdsl Access

Three schemes for bonding multiple G.shdsl links are subsequently described. The advantages derived from creating a "fat pipe" from multiple "thin pipes" are self-evident. However, these schemes introduce a modicum of administrative overhead that is pointed out in each case. It will be noted that two of the methods are asynchronous transfer mode (ATM)–centric and that a brief explanation of why is provided first.

ATM, as a technology, has been much maligned as being "inefficient." However, it has survived because it is arguably the best compromise for handling a variety of traffic types, including voice and video, that require low-latency, real-time transmission. Furthermore, it is well understood, widely deployed, and able to support a range of QoS parameters. While it does not have the panache of IP, it is the "wheel dog" of today's networks. The majority of deployed DSLAMs expect that the xDSL transmissions conform to ATM at the Layer-2 level. While other network segments might migrate to IP or morph to multiprotocol label switching (MPLS), the access segment is more than likely to remain ATM–centric for a long time to come. Thus our initial premise of "no forklift upgrades" has been adhered to.

The first two methods described address the access segment only. That is, they are methods for providing fat pipes between a customer premises and the serving DSLAM and are referred to as techniques for *access inverse multiplexing*. This is depicted generically in *Figure 1*.

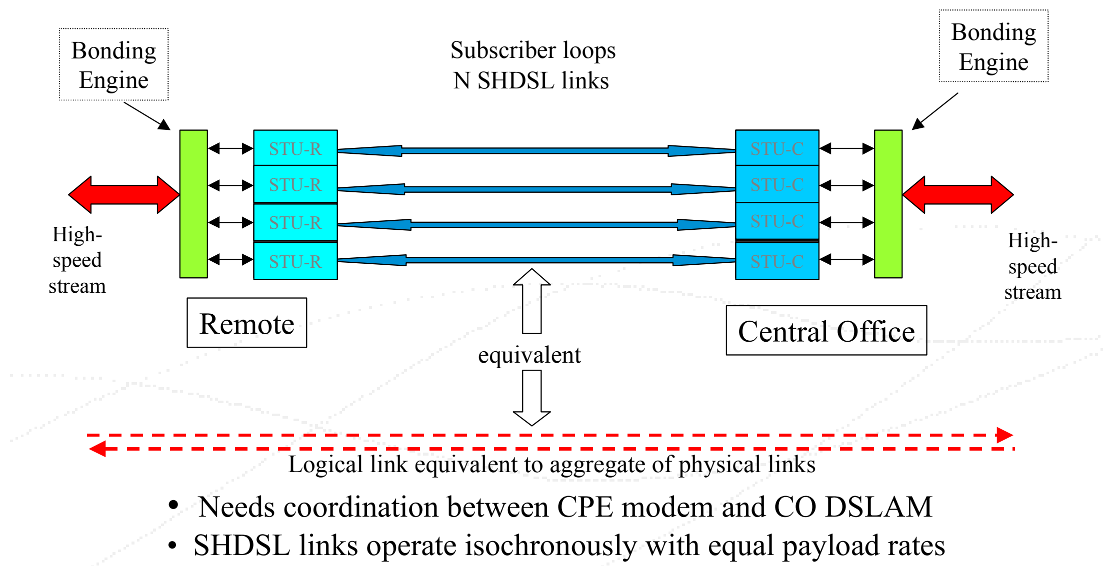

Figure 1: *Access Inverse Multiplexing*

The 4.6 Mbps Mode

The International Telecommunication Union–Telecommunication Standardization Sector (ITU–T) G.shdsl standard allows for the bonding of two G.shdsl links, effectively doubling the bit rate relative to a single loop—hence the terminology "4.6 Mbps mode." Bonding is achieved by operating the two loops synchronously and splitting the bearer traffic between the two in an octet interleaved fashion. The ITU–T standard is thus a Layer-1 bonding method and is upper-layer agnostic. There is a small constraint on the differential delay between the two loops that is easily met, considering that the loops have identical end points (the modem or STU–R, and the DSLAM or STU–C). The 4.6 Mbps mode provides twice the bit rate at a given distance or about 40 percent improvement in reach at a given bit rate when compared with single loop G.shdsl. A pictorial representation of the 4.6 Mbps mode is shown in *Figure 2*.

There is a minor administrative constraint as well, since the two loops must terminate in "adjacent" ports of the same DSLAM blade. A variant of this philosophy has been promulgated as "HDSL-4" for carrying T1 rates over two loops and taking advantage of the spectral friendliness and rate-reach afforded by the trellis coded–pulse amplitude modulation (TC–PAM) of G.shdsl. One drawback of the 4.6 Mbps mode is that if one loop goes out of service, the link is lost.

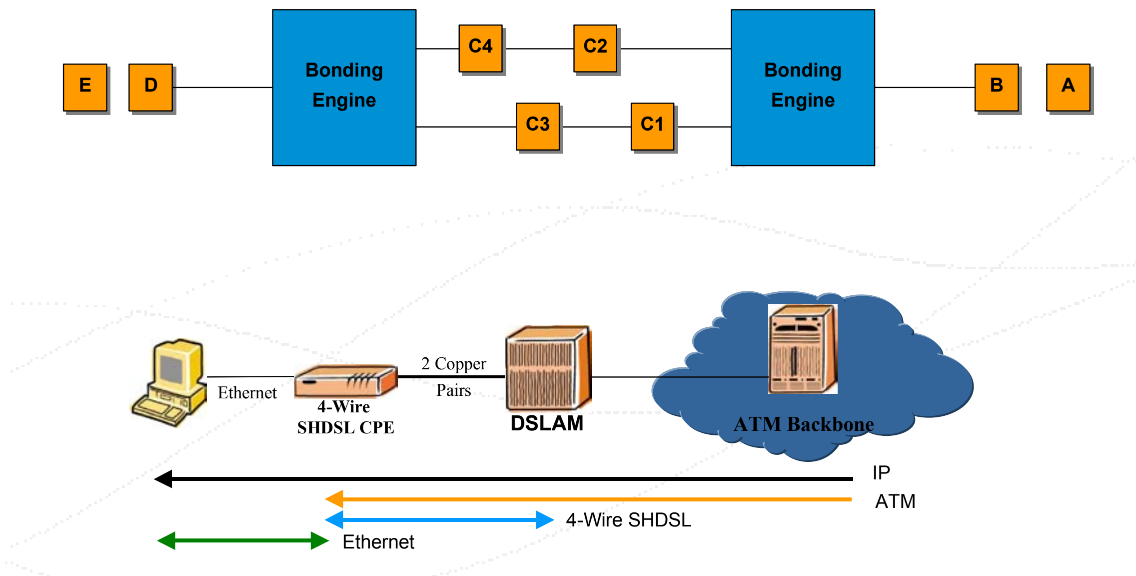

Figure 2: Four-Wire SHDSL

The IMA Mode

Inverse multiplexing over ATM (IMA) is a standard promulgated by the ATM Forum (AF-PHY-0086.001, March 1999) for splitting an ATM cell-stream over multiple ATM cell-streams of equal rate and is illustrated in *Figure 3*. However, IMA is a point-to-point methodology, implying that the STU–R (the customer-located modem) and STU–C (the DSLAM) both be equipped with IMA functionality. The aggregation method does introduce a small amount of overhead, so the aggregate bit rate is a tad less than the sum of the bit rates of the component pipes. One advantage of IMA is that if one loop goes out of service, the link remains up, albeit at a reduced bit rate.

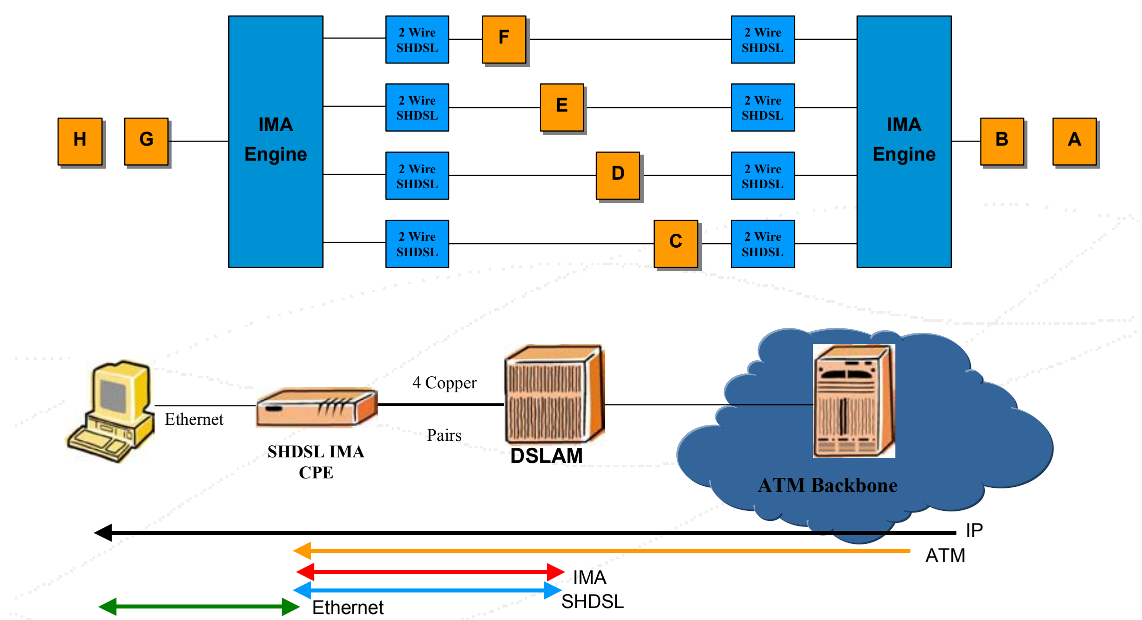

Figure 3: SHDSL IMA

For example, an STU–R with four loops can support an aggregate bit rate of about 9.2 Mbps, since each loop can support 2.3 Mbps. A four-loop IMA can provide four times the bit rate at a given distance, or 100% improvement in reach for a given bit rate, when compared with single loop G.shdsl. There is a small constraint on the differential delay between the four loops that is easily met, considering that the loops have identical end points (the STU–R and the STU–C). There is a minor administrative constraint as well, since the four loops must terminate in "adjacent" ports of the same DSLAM blade.

The Bonded–PVC Mode

Not all DSLAMs support IMA. However, all DSLAMs do support multiple permanent virtual circuits (PVCs) in a single G.shdsl link. The bonded–PVC method is an end-to-end bonding technique whereby a single "fat-pipe" PVC is split over a multiplicity of "thin-pipe" PVCs that are not necessarily of equal rate but transported across the network individually to the far end, where the recombination is performed. The method requires a slightly higher overhead than does IMA. The buffering requirement for accommodating different speeds adds delay, making the bonded–PVC mode appropriate for data "fat-pipe" PVCs. Bearer channels requiring low latency, such as voice channels, are typically "thin pipes" to begin with and are assigned regular PVCs that can be contained within a single G.shdsl link. Since this is an end-to-end bonding method, we refer to it as *network inverse multiplexing*. A pictorial representation is provided in *Figure 4*.

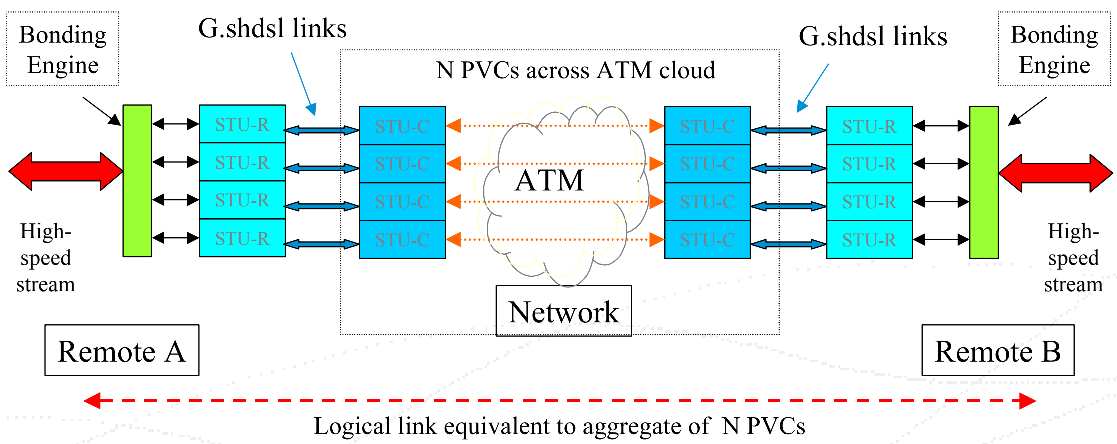

- Needs coordination between CPE Modems
- Individual links can be of different rates
- Network appears as N individual PVCs
- End-to-end differential delay between PVCs must be within limits

Figure 4: *Network Inverse Multiplexing*

For example, an STU–R with eight loops can support "fat-pipe" PVCs of a bit rate just shy of 18.4 Mbps, since each loop can support a PVC of (just below) 2.3 Mbps. There is a small constraint on the differential delay across the eight PVCs on an end-to-end basis that is easily met, considering that the loops have identical end points (the STU–R and the STU–C) and the individual PVCs can be assigned identical routes through the ATM cloud. Note that the eight loops need not terminate in "adjacent" ports of the same DSLAM blade (as in the other two bonding methods), since the DSLAM is unaware that the constituent PVCs are indeed linked to each other.

Applications

By using G.shdsl, in particular multi-link G.shdsl, a wide variety of applications can be serviced. A few are provided here by way of example.

Direct Internet Connections

SMEs have applications that more directly approach those of the large enterprise, vastly different from the SOHO segment. SOHOs are happy with an asymmetrical connection well served by ADSL technology. SMEs need symmetrical solutions for high-speed Internet browsing, Web hosting, secure commerce, file backup, and LAN extensions. ADSL technology—limited to 384 kbps symmetrical—simply cannot provide the bandwidth necessary to serve the SME segment. New direct Internet connectivity services offering symmetrical service at 2 and 4 Mbps are being introduced using G.shdsl. Some carriers estimate that the overall cost of providing and maintaining 2 Mbps G.shdsl access is somewhere between one-half to one-quarter the cost of existing E1/T1 access.

G.shdsl services may be provided from the same DSLAM as ADSL services. 2 Mbps service—comparable to existing E1/T1 leased-line services—may be provided using a single port on the G.shdsl card and a single wire pair. If the DSLAM supports the G.shdsl four-wire feature, symmetrical 4.6 Mbps service using two pairs can be provided to facilitate higher-speed applications such as Web hosting, file backup, and off-site application service providers (ASPs). This is depicted in *Figure 2*. Again, provided that the DSLAM supports the feature, four loops can be deployed to provide a 9.2 Mbps service. From a different viewpoint, the two-loop and four-loop configurations can be used to provide 2 Mbps service at much greater distances than is possible with a single loop sans repeaters. In *Figure 5*, the rate-reach achievable with one, two, and four loops is provided. *Figure 6* shows the rate-reach by considering some common bit rates and also includes the reach improvement using one repeater (per loop).

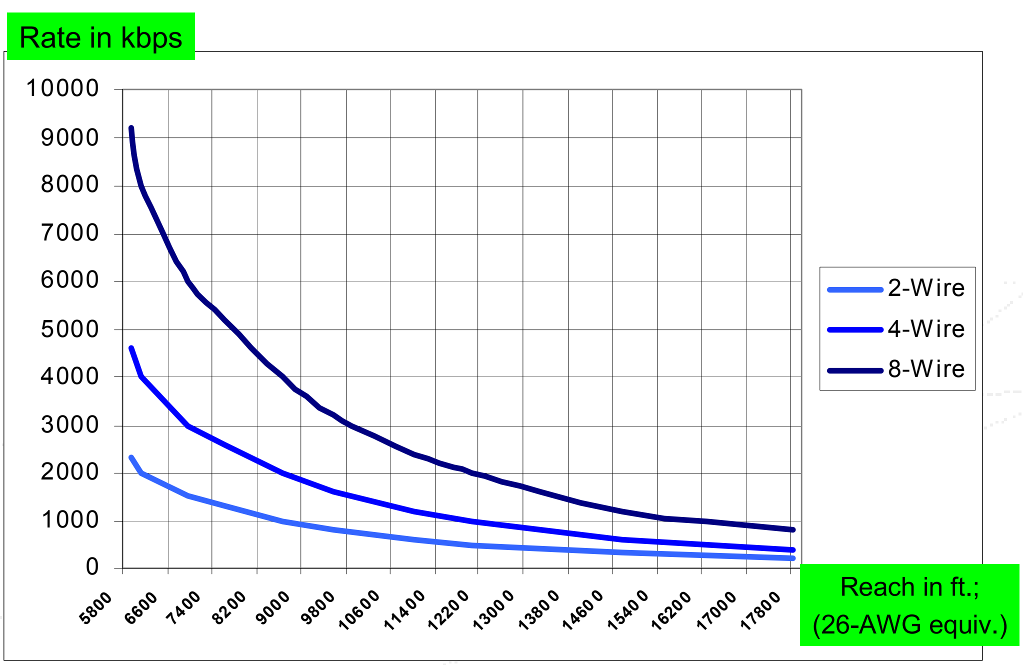

Figure 5: *Rate versus Reach Comparision*

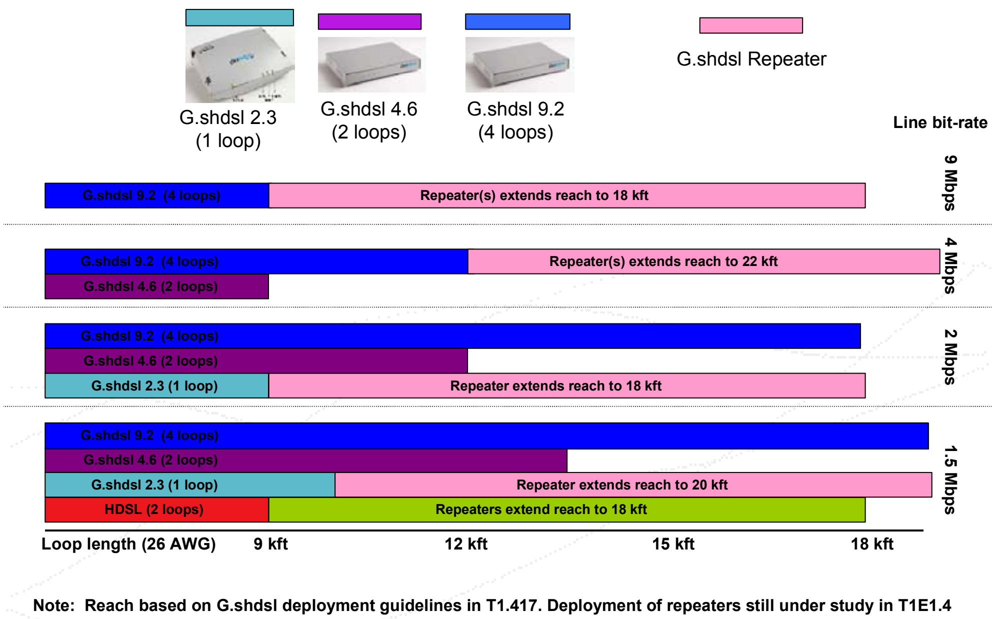

Figure 6: *G.shdsl Rate-Reach*

Transparent LAN Services

LANs are now migrating from their 10 Mbps origins to 100 Mbps. Wide-area network (WAN) interconnection of LAN segments are now being stressed due to the limitations of copper fed E1/T1 services operating at 2/1.5 Mbps. Ideally, LANs could be served—via copper pairs—in incremental steps of 2 Mbps up to 10 Mbps.

G.shdsl transport using IMA is an ideal solution for burstable bandwidth—over copper—from 2 Mbps to 10 Mbps. In fact, 9.2 Mbps access services may be provisioned over four copper loops. The logical "fat pipe" allows shared applications across the 9.2 Mbps available bandwidth. This provides an appropriate solution for interconnecting LAN segments in geographically dispersed locations to create what has been called a "World Wide LAN". *Figure 7* depicts the access arrangement between one customer location and the associated serving CO. Obviously, the same access scheme is appropriate for providing direct Internet access.

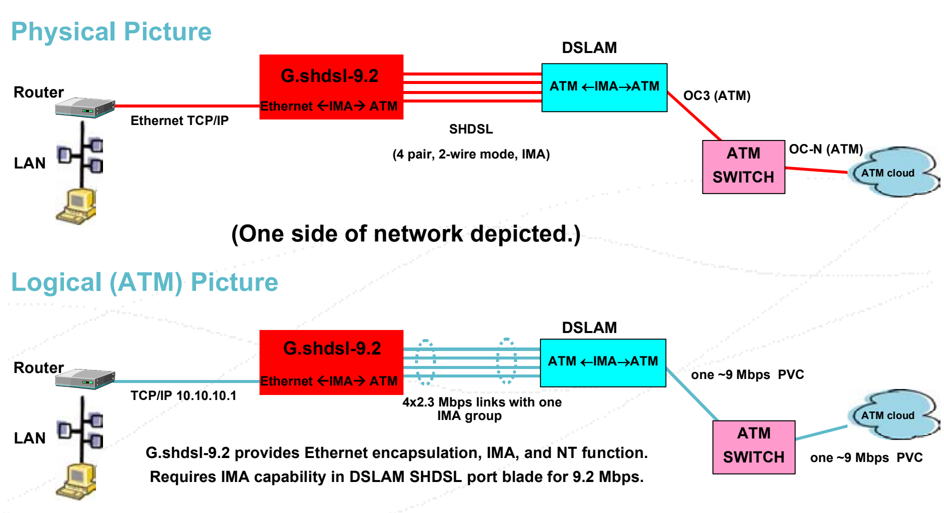

Figure 7: *Transparent LAN Solution*

Converged Services

Access to the Internet is mandatory for an SME. However, if an SME is geographically dispersed, it is not mandatory for each location to have Internet access directly. The security implications are significant. From the viewpoint of operations and maintenance, information technology (IT) professionals prefer to have the locations linked via a private network with one location, usually the SME "headquarters," providing Internet access for the entire company. Transparent LAN service using a logical fat pipe as discussed in the previous section is a suitable way to interconnect LAN segments at the various customer locations. What is one to do if the desire is to interconnect existing telephone switches (e.g., private branch exchange [PBX]) as well as data routers? One solution rests in end-to-end bonding (*network inverse multiplexing*) for linking multiple PVCs to provide one logical, fat-pipe PVC for data, with additional PVCs for the voice trunks. In fact, such an integrated access device (IAD) with eight G.shdsl links could provide 18.4 Mbps access bit rates, as well as provide legacy (drop-side) interfaces—E1/T1 and E3/T3—in addition to 10/100 BaseT (Ethernet).

Many corporate networks are composed of larger regional headquarters networked with remote offices. Today's technology allows networking of all "on-net" communications via carrier-owned ATM fiber core networks. Unlike the public Internet, these private ATM clouds allow the QoS necessary to meet or beat SLAs currently in place for traditional circuit-switched, time division multiplexed (TDM) networks. Voice from legacy applications is typically transported via T1/E1 connections using channelized TDM. The IAD allows attachment to T1/E1 legacy voice applications converting the voice to ATM for transport either in ATM adaptation layer–2 (AAL–2) end-to-end PVC provisioning mode or AAL–1 circuit-emulation service. As illustrated in *Figure 8*, remote offices may now be connected providing converged services via G.shdsl access. Bandwidth delivered to the remote locations may be provisioned from a single E1/T1 rate—2 Mbps—all the way up to 15 Mbps. The result is a unified corporate office experience with high-speed LANs. Furthermore, four-digit voice calling can now be replicated at remote office locations. Since the private network is secure, only one single point of access to the Internet, with the associated firewall, is required. Needless to say, each individual location could have backup facilities for public switched telephone network (PSTN) and Internet access.

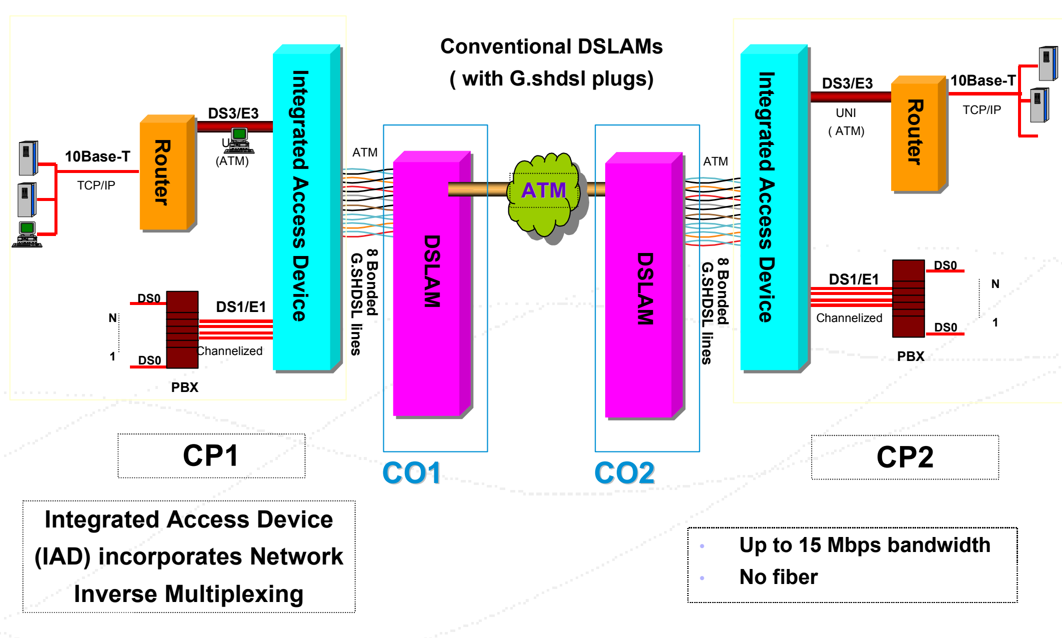

Figure 8: *Customer to Customer*

Summary

Today's telephone service providers know that by combining diverse services on to a single broadband access line, much money may be saved in infrastructure costs, while new, high-margin, high-revenue business services can be introduced. Carriers can now span the broadband business gap connecting the SMEs with burstable bandwidth up to 15 Mbps over multi-loop G.shdsl. The result is that carriers can add high-value, high-margin business services with a simple upgrade to existing copper infrastructure and CO DSLAM equipment.

The term "digital copper" is quite descriptive of subscriber loops carrying digital information and serves to make a firm distinction between the access means discussed here as well as other access methods such as fiber or wireless. The authors would like to thank Dr. Avi Vaidya of Charles Industries who coined the term.