Optimizing ISDN to Give More Capacity at Less Cost

James Y. Bryce

Always On/Dynamic ISDN Network Architecture (AO/DI) is a proposed standard promulgated by the Vendors' ISDN Association (VIA). AO/DI creates the illusion of full-time connection while actually reducing connection times and costs to carriers and users. AO/DI drives down overall bandwidth usage and need for customer premises equipment (CPE) for Internet service providers. This seeming magic stems from effective use of D-channel packet switching to transmit routine data traffic such as security authentication, electronic mail, and requests for B-channel set-up and tear-down.

Perspective

ISDN (Integrated Services Digital Network) has been around in one form or another for more than a decade. As a creature of an historically intensely regulated collection of private sector monopolies and public sector postal, telegraph, and telephone government agencies, ISDN languished in its first attempt at market acceptance. Why? It was born too late. It was born too soon. It was born with technological lockjaw.

The First Time Around

Too late, because the time scale perspective of its creators was conditioned on capital recovery measured in decades and "services" conceived as a basis for filing tariffs with regulators. But, concurrently with its birth, technological developments began to reduce the value of capital to zero in a few years, and legal pressures began to foster competition and remove barriers to entry.[1] Profits derived from artificial construction of services that cost nothing to add, but that formed the groundwork for cash flow, came under pressure from competitors prepared to cut prices to the bone on the way to wresting market share.

Too soon, because the real need for its extension of digital connectivity all the way to subscribers was not in voice services, but in computer connectivity. When ISDN first appeared on the scene in the early 1980s, computers and computer networking were not ubiquitous. Today, they are. It was a hard and largely unsuccessful sale to get voice users to adopt ISDN; all too often, its complex tariffs and new equipment delivered less for more. Early attempts to sell its use for computer networking focused on providing local area networking using B-channel data rates of 64 kilobits per second (Kb/s). This attempt was ludicrous when Ethernet was delivering megabits to the desk at far less cost. ISDN promoters who thought computer users would flock to the technology simply did not see or understand the freight train thundering down the tunnel. They thought they were competing with 9.6 Kb/s and 19.2 Kb/s terminal-to-host communications; the LAN savvy laughed and installed another million Ethernet nodes.

With lockjaw, because its parents were determined to milk the maximum profit possible from the new baby through devising proprietary implementations designed to lock users into a vendor's products to the exclusion of competitors. It didn't take long for the user community with the most immediate big money—governments and large businesses—to pan ISDN as an expensive way to get nothing new. Instead, they ended up with dead-end hardware and software married to a single supplier charging outlandish rates that delivered upgraded functionality on a geological time scale, as

Mr. James Y. Bryce is a writer, speaker, and consultant on communications systems and future technologies. Jim practiced law for 10 years prior to studying computer science and installing the first commercial Ethernet using IBM PCs. He uses his professional education and experience to design and improve networks to optimally fulfill an organization's needs. As the author of *Special Edition, Using ISDN,* now in its second edition, Jim currently speaks, writes, and consults on the convergence of computer and telephone technologies.

1070-3683/97/$5 © 1997 Technology Futures, Inc.

well as spending most of its resources balkanizing tariffs and engineering around competitors. No wonder ISDN came to mean "It Still Does Nothing," among other monikers, some much less attractive and not printable in this journal.

In effect, the parents of ISDN did everything they could to ensure its failure the first time around.

Those marketing ISDN for voice totally missed another technology undergoing a metamorphosis... to a force so appealing, so useful, so powerful as to hold the potential to forever and dramatically change governments, economies, and lives in general.

> Being digital changes the character of the standards for machine-to-machine communications. People used to sit around tables in Geneva and other such places to hammer out (a telling metaphor from the industrial age) world standards for everything from spectrum allocation to telecommunications protocols. Sometimes this took so long, *as in the case of the telephone standard ISDN* , that it was obsolete by the time it was agreed upon [emphasis supplied].[2]

The Second Time Around

Fast forward to the late 1980s and early 1990s. Faced with the imminent demise of an aesthetically appealing technology, foster parents stepped in to save what had become a technological foundling. Their efforts were met with the awesome forces of strategically arrayed complex legislative/regulatory environments, tariffs, and habits going back a century. What to do? Proselytize potential users, equipment makers, government agencies, and carriers with the word that ISDN has risen, risen indeed!

In both Europe and North America, this effort began to bear fruit. In North America, the North American ISDN Users Forum[3] turned its sights on the crying need for standardization and, in conjunction with its rich uncle, the National Institute of Standards and Technology,[4] conducted countless boring standards meetings for similarly countless standards committees that hashed over countless minutia of technology and regulation to create National ISDN 1 (NI-1) and its successors. Participants drowned their ennui in after-hours revelry and socializing along with whatever tourist pleasures the host city provided; it's nice to jet set around on the company or agency.[5]

Thanks to standardization, both at the regional level (NI-x, etc.) and the international level through the International Telecommunications Union,[6] we started having a workable, interoperable, and reasonably economic ISDN around 1995. ISDN now delivers end-to-end standardized digital signals throughout much of the technologically savvy world.

ISDN's Unexpected Partner

But what on earth do we do with it? ISDN truly does provide a number of sophisticated voice options that are often very beneficial in particular business situations. Under certain tariff conditions, ISDN can actually be more economic for voice use than analog telephony; this was not true until standardization brought increased function and lower prices. As a result, the initial thrust of National ISDN-1 (NI-1) marketing by telcos in North America again emphasized voice.

Those marketing ISDN for voice totally missed another technology undergoing a metamorphosis from a technically obtuse academic and military research sandbox to a force so appealing, so useful, so powerful as to hold the potential to forever and dramatically change governments, economies, and lives in general. The Internet penetrated the world at a warp speed its most optimistic

Figure 1
Web Site Increases

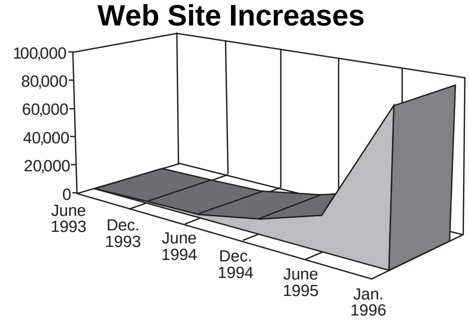

Source: Mathew Gray, MIT

proponents never foresaw, even in their most outrageous dreams (see Figure 1).

The Internet demands more, faster, for longer. Balaji Kumar analyzed this extensively in the last issue of *NTQ*.[7] Table 1 from that article provides a thumbnail sketch of the Internet's impact on the public switched network.[8]

Currently, the only generally-available telecommunications technology that can begin to ameliorate the growing Internet demand is ISDN.[9] Never mind your cable modems, xDSL, satellites, wireless, and other methods. They're in the (near?) future for most of us. They'll come, but ISDN is here now. So how can we optimize it?

ISDN as Today's Best Solution for the Internet

The worldwide public switched network is moving to digital technology. The overall design for that digital system is called Integrated Services Digital Network. But, current parlance overlooks this more general, and accurate, application of the term "ISDN." Most of the time, "ISDN" is understood to mean delivery of digital services over the local loop, the link between the telephone company's central office switch and the user's business or residence. Using the ISDN specifications, basic rate ISDN (BRI) services are designated "narrowband ISDN,"[10] while higher data rate ISDN that is imbedded within switches and switch-to-switch communications is designated "broadband ISDN."[11] This article concerns local loop communications; so, unless qualified, from here on, ISDN means narrowband ISDN or BRI.

ISDN Fundamentals

A quick review of the fundamentals of ISDN will help avoid potential ambiguities.

Figure 2 shows a typical way of thinking about the ISDN basic rate interface. This is a conceptual view of BRI; there is no wire or cable used for ISDN that actually looks like this. In reality, BRI is delivered to the end user on a single twisted pair that may already be installed from the telephone exchange. There's no need for the phone company to install new or special cable. Of course, should the user already have used up all available pairs with analog service, the phone company may have to install additional twisted pairs from the pole or underground cable to support a new ISDN connection. That's part of the deal in most tariffs: The phone company must install any additional outside wires whether they be POTS or ISDN when a user requests a tariffed service. Inside wiring design and costs are borne by the user, at least in the United States.

BRI has two *B-channels* that each carry 64,000 bits per second (64 Kb/s) *in each*

Currently, the only generally-available telecommunications technology that can begin to ameliorate the growing Internet demand is ISDN.

Table 1
Residential Voice & Internet Traffic Comparisons

For residential voice traffic:	
An average number of calls per hour per subscriber line	2 calls per hour
An average call holding time	3 minutes
Average usage on the line	2 × (3 × 60)/100 = 3.6 CCS*
For residential Internet traffic:	
An average number of calls per hour per subscriber line	2 calls per hour
An average call holding time	20 minutes
Average usage on the line	2 × (20 × 60)/100 = 24 CCS*

* Where the maximum possible usage per line is 36 CCS.

Source: B. Kumar

Figure 2
Basic Rate Interface

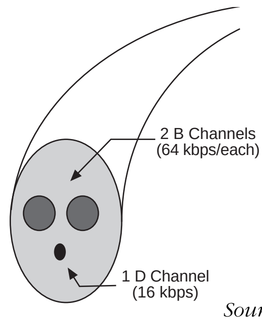

Source: J. Y. Bryce

> The D-channel is the element that defines ISDN.

direction (full duplex); so, while the user is sending 64 Kb/s to someone else, that receiving party can simultaneously be sending 64 Kb/s to the user. Since there are two B-channels, the user can be connected to two different places at the same time or can combine the two B-channels and connect to a single location at 128 Kb/s—*both ways.*

If the user has larger needs, the next step is the Primary Rate Interface (PRI) that is delivered on two twisted pairs. Figure 3 shows a PRI with 23 B-channels (64 Kb/s each) and one D-channel that, at 64 Kb/s, is four times larger than that of a BRI. This 23 B-channel design is true for the United States and Canada; in much of the rest of the world, a PRI has 30 B-channels and the 64 Kb/s D-channel. This is essentially the same difference as that between the T-1 line used in the United States and Canada and the E-1 line used elsewhere.

Figure 3
Primary Rate Interface

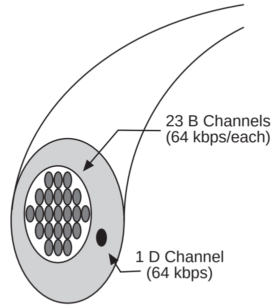

Source: J. Y. Bryce

Just one more detail—If the user needs multiple PRIs, the first will have 23 B-channels and one D-channel (European 30/1); several additional PRIs may be using the already existing D-channel in the first to control the additional B-channels in the other PRIs. So, subsequent PRIs will have 24 B-channels and no D-channel (European 31/0). Some system designers prefer each PRI to have its own D-channel to avoid the possibility of a single point of failure in one D-channel controlling the several PRIs.

D-Channel Power

The D-channel is the real power behind ISDN. In fact, the D-channel is the element that defines ISDN. T-1 service is digital, but T-1 does not provide a D-channel. While B-channels carry the users' usual traffic (they bear the traffic, hence they're called bearer channels), the D-channel controls everything. All the information to place the call, maintain it, and terminate it is carried on the D-channel. The D-channel never hangs up. It's always talking with the telephone company switch. The D-channel can be talking to a number of devices at the same time, all the time; the exact number is implementation dependent. Everything for calls can already be in place on the D-channel. All the user's equipment has to do is tell the D-channel when to finally tie available B-channels to the already-defined location.

D-channel signaling makes getting things going fast—really fast. It's the best-kept secret of ISDN. Everyone knows about the data rate on the B-channels, 64 Kb/s; most know about the BRI data rate on the D-channel, 16 Kb/s; but few know the D channel secret—it can set up or tear down a telephone connection in a fraction of a second (like 1/18)—or less! Figure 4 contrasts this connection time with that of an analog modem.

Rapid connection/disconnection means computers communicating with ISDN can work in a very different way from those using POTS. Most people think of ISDN as

Figure 4
Call Set-Up/Tear-Down Comparison

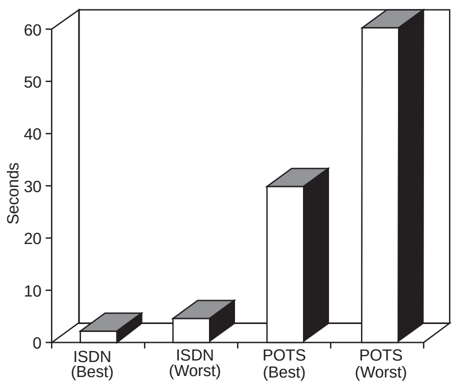

Source: J. Y. Bryce

providing a faster data rate, but they continue to think of dialing up to a location, say an Internet service provider (ISP), and staying connected for the duration of their use of the resource.

Think about this. If a user downloads a Web page and starts reading it, the download might take 10 seconds, while the user could be reading for a minute, two minutes, or more. Why stay connected? With POTS, the answer is easy; it could take 30 seconds or more to connect and maybe not much less to disconnect. Users don't want to go through all that every time they decide to check out a link; consequently, they stay connected, tying up their phone line, the phone line of the ISP, the ISP's modem, and the phone company's switch, while not sending or receiving anything most of the time.

With ISDN, the connection is made in far less time. The phone connection joining a B-channel from the user to a B-channel at the destination (e.g., an ISP) takes a fraction of a second. There are none of the synchronization problems experienced with POTS, so the user's system and the end system get connected quickly and simply. Then, the only remaining wait is authentication for security purposes; this is usually Challenge Handshake Authentication Protocol (CHAP) in the TCP/IP world.[12] If the security is to be optimized at all, the wait will be about five seconds, maybe less. Notice that most of the wait is *not due to ISDN*, but to the authentication (CHAP, PAP, or whatever) process.

Under the best of conditions, authentication would be performed only once when the user first connects; then, it would be maintained over the D-channel even when the B-channels are not connected. Future B-channel connections are effective within the fraction of a second needed to make the phone connection, because authentication is already taken care of.

Even with authentication's two- to five-second delay, it's quite reasonable for the user to break the connection while reading a Web page. If the user sees a desirable link, the user clicks, waits a few seconds, and things start again. Users pay the phone company and the ISP less connect time. The phone company and the ISP can service more subscribers with less lines and equipment. A win/win/win situation is at hand. The author has been using just this method of connecting to an ISP using ISDN for the last few years, even though the local exchange carrier involved does not have a metered charge for connection. Effective implementation of this method removes the frequent misperception that a *dedicated ISDN* connection is desirable.

In jurisdictions where there is no metered ISDN charge, many ISPs offer the so-called *dedicated ISDN,* encouraging users to sign up for ISDN connections in exchange for considerably higher ISP monthly charges and promises of higher-quality service. Dedicated ISDN is an oxymoron. From the perspective of the local exchange carrier, ISDN is a dial-up service; from the user and ISP perspective, however, it amounts to the same as a dedicated line when dialed up and left up. Such a line usually costs substantially less than a line tariffed as dedicated. This state of affairs is heading for disaster.

Always On/Dynamic ISDN

The stage is set with increasing Internet usage, telephone companies crying foul, and users waiting interminably for Web pages. The usual solution path in the public

In Condemnation of "Dedicated ISDN"

I am very concerned that widespread use of dedicated ISDN, with flat-rate-tariffs, will result in the blow that killed the goose that laid the golden egg. By that, I mean the telco will observe the substantial number of users that maintain continuous off-hook conditions on ISDN, a switched service. Then, the telco will move to change the rate structure. U S WEST filed for a flat rate of about $185 in some of its areas, citing the dedicated ISDN issue. One might expect a similar move by other telephone companies. Such companies might also be tempted to go to a time charge. These moves should be actively opposed.

But, in opposition, it would be best for users to go in with clean hands by showing they are using the resource in an efficient and responsible way. ISDN is especially well-designed for efficient use. The rapid set-up and tear-down of an ISDN connection fosters careful adjustment of connection time. I suggest all ISDN users explore this more efficient use and move away from the dedicated ISDN idea. In that way, we may prevent a move by the carriers to raise flat rates or go to metered rates. We will be in a good position to oppose any such moves by showing responsible efforts to limit usage to only that which is really useful for exchanging information.

—Excerpted from J. W. Bryce, *Special Edition, Using ISDN, 2nd Edition* (Indianapolis, IN: Que Corporation, 1996), pp. 505-506.

communications world would seek to set appropriate standards and mold regulations. But, traditional standards organizations move at the rate of frozen molasses, and regulatory bodies push rights and incentives around trying to balance everyone's interests while burning time and money. The computer industry doesn't function like that. It expects and demands rapid development and adoption of innovations, all the while spawning more competition, driving prices down, and functionality up. In the last issue of *NTQ,* Joan Van Tassel explored the dynamics of this atmosphere.[13]

The stage is set with increasing Internet usage, telephone companies crying foul, and users waiting interminably for Web pages.

Vendors' ISDN Association

Vendors of ISDN products have read the handwriting on the wall; here's the author's translation:

> The telecommunications establishment has no sense of urgency motivating it to optimize standards or charges for computer communications.

When this epiphany struck the makers of ISDN equipment, they realized their investments and profits were held hostage by an antiquated telecommunications world. They set out to change it in the characteristic American way recognized by Alexis de Tocqueville[14] over a century ago—they formed an association, the Vendors' ISDN Association. The VIA has identified three areas of technical improvement for ISDN:

- CPE diagnostics.
- Switch identification.
- Always On/Dynamic ISDN.

CPE diagnostics provides a means of testing the customer premises equipment, such as a router or other terminal adapter, in a standardized way to make sure it will operate properly with central office equipment and CPE of other vendors. Switch identification makes installation and configuration of ISDN CPE easier for users by automatically identifying the type of central office switch and configuring the CPE accordingly.

Thumbnail Sketch of AO/DI

While both of these initiatives are very important on their own, they do not address the problem of public switched network congestion brought on by intense Internet use. Always On/Dynamic ISDN (AO/DI) does. AO/DI takes a leap the author predicted a few years ago.[15] It uses the D-channel, which is *always on*, to *dynamically* control addition and removal of B-channels. It does some other tricks, too.

At first blush, many think this is already being done through the B-channel aggregation provided by techniques such as BONDING[16] and multilink PPP.[17] In fact, aggregation using such methods can only occur once a B-channel is already in place. AO/DI optimizes the situation by taking advantage of two D-channel aspects:

- The D-channel is always up.
- The D-channel can carry user data using X.25 packet switching.

Earlier sections discussed the idea that the D-channel is always up. But what about the D-channel carrying X.25 packets? There's been little mention of this capability because all the noise about ISDN has focused on the increases in data rate afforded by the B-channels compared with POTS. In the BRI world, the 16 Kb/s capacity of the D-channel looks very small compared with the 128 Kb/s capacity of two B-channels. In addition, the typical D-channel packet-switched offering affords only 9.6 Kb/s for packet traffic. Given those figures, the typical response of a computer jockey is, "Who cares?" Start caring.

That paltry 9.6 Kb/s is more than enough to set-up and maintain a session with an ISP at the other end. It's fine for continual exchange of ordinary, text-based electronic mail and news feeds. Here's the kicker: it can add B-channels whenever needed. That's worth caring about. According to Andy Kuzma, chairman of the AO/DI committee:

The **Vendors' ISDN Association (VIA)** is a nonprofit corporation and open group chartered with the express purpose of accelerating the deployment of ISDN through rapid convergence of end-user ISDN to public network interoperability specifications and industry-wide promotion of ISDN.

VIA's purpose is to expand and accelerate the deployment of ISDN products, services, and usage by providing an open forum for the exchange of ideas, user needs, and technical information regarding ISDN.

VIA Objectives

- Simplify ISDN CPE implementation for end users.
- Improve interoperability between CPE and the network.
- Simplify and/or automate ISDN CPE configuration, operation, and management.
- Propose specifications for development of enhanced ISDN capabilities.
- Promote open network interoperability standards.
- Promote uniform processes for ordering and implementation of ISDN by end users.
- Facilitate CPE/network interoperability testing and support.
- Stimulate demand for ISDN through communications, promotions, and market education programs.

VIA Initiatives

VIA will consider and evaluate proposals that will enhance ISDN by improving ease-of-use, reliability, interoperability, and functionality. The initial focus of VIA will be on simplified, automated ISDN configuration capabilities. VIA will actively solicit and evaluate proposals for additional ISDN interoperability enhancements and new capabilities. Visit them on-line at http://www.via-isdn.org.

The basic idea of AO/DI is that an ISDN D-channel X.25 call is placed from the subscriber to the packet data service provider. The multilink protocol and TCP/IP protocols are encapsulated within the X.25 logical circuit carried by the D-channel. The bearer channels are invoked as additional bandwidth is needed. The bearer channels use the multilink protocol *without* the Q.922 and X.25 encapsulation used on the D-channel.

Using the X.25 over the D-channel, while admittedly not the most efficient protocol stack, allows AO/DI to take advantage of the existing packet handlers at the central offices. The link associated to the D-channel X.25 packet connection is used as the primary link of the multilink protocol.

Because the D-channel is an always-available, connectionless, packet-oriented link between the CPE and the central office, it is possible to offer an always-available service based on it. Further, because the D-channel X.25 packets are handled at the central office by the X.25 packet handler, it is possible to route these packets without first crossing the time-division circuit-switched fabric of the switch, which reduces the impact to the telephony network.[18]

X.25 is a specification for exchanging information using packets. It works by using *packet switches*, which are devices that read an address carried in a stream of electrical bits interpreted by the equipment as a packet. With the address in hand, the switch forwards the packet to a destination based on what is read. The switch then examines the next packet which may be from the same source or from a different source and repeats the process. As a result, there is never a *continuous* connection between the source and destination using packet switching.

The term *connection* gets slippery, as there are also distinctions between *connection-oriented* and *connectionless* protocols. It would take several pages to technically define all this. The crux is the same: Packet switching is very different from the usual telephone call. In telephone calls, a path is set-up and dedicated to a particular call for the duration of that call. In packet switching, no such long-term path exists. The telephone call ties up resources for the entire

Using the X.25 over the D-channel, while admittedly not the most efficient protocol stack, allows AO/DI to take advantage of the existing packet handlers at the central offices.

time the call continues. Since most verbal calls probably have more time spent without sounds being transmitted, this is a waste of resources. In the world of data transmission, there is even less time devoted to moving actual information over an established circuit, hence more waste. Packet switching is more efficient simply because it uses communications resources only while something is really being exchanged; when a sender is silent, those same resources are available for others to use. Anything that can be done to move data transmission from the classic telephone model to the packet switched model will reduce waste.

AO/DI makes just that move. By placing traffic that can travel at lower data rates, such as electronic mail and news, within the packet switched system, it removes a substantial load from the usual model, the *circuit switched* telephone system. By providing a means for increasing and decreasing higher data rates through attachment of B-channels only as needed and controlling that attachment through the efficient packet switched D-channel, AO/DI overcomes its own data rate limitations.

By placing traffic that can travel at lower data rates, such as electronic mail and news, within the packet switched system, it removes a substantial load from the usual model, the *circuit switched* telephone system.

Problems Affecting Adoption of AO/DI

It looks like a rosy scene, but problems are waiting in the wings. These problems can be resolved if action is started now. For AO/DI to gain wide acceptance and realize its full potential, steps must be taken to:

- Modify or replace CPE.
- Provision central offices for X.25 packet switching.
- Change tariffs that do not favor X.25 packet switching.
- Change tariffs that discriminate against short connections.
- Provide for continuing authentication over the D-channel.

Modifying CPE

The first problem of CPE is probably not great. Initial prototypes implementing AO/DI are in the hands of testers. Delivery of the finished product to users is less than a year away. In many cases, existing CPE can be upgraded to AO/DI by firmware modifications, so the cost may be quite reasonable.

Provisioning Central Offices

Provisioning central offices for X.25 could be a problem in some places, but, on the whole, this capability is already in place, especially if the ISP is local to the switch. This is less true if the ISP must be reached through a long distance connection. When the telcos realize the substantial benefits to be derived from reductions in B-channel connection time, any additional investment will be justified.

Changing Tariffs that Do Not Favor D-Channel Packet Switching

Unfortunately X.25 packet switching tariffs will probably be the biggest stumbling block. These vary all over the place. In some areas, there is no or very minimal cost for use of X.25 packet switching; in others, the costs would rapidly become prohibitive with use of the sort AO/DI contemplates. For example, in Austin, Southwestern Bell offers 2B+D ISDN service for a cost to the end user, including all charges and taxes, of $70 per month; this is an unmetered service meaning no time charges. X.25 packet switched service over the D-channel has no added installation charge, but adds $2 per month to the user's bill. In addition, there is a *set-up charge* of $0.005 and a character transmission charge of $0.20 for each *kilosegment* of information.

This bears translation. The *set-up charge* is made each time a stream of packets is initiated to a destination. A *kilosegment* is one thousand segments; a *segment* is 64 octets. An *octet* is eight bits—that is, one byte or one character. If any reader understands this right away, that reader is ahead of otherwise very knowledgeable and experienced tariff experts at Southwestern Bell that helped the author plow through it.[19]

Here are the costs once everything is factored in. Every time AO/DI contacts the ISP, there is a set-up charge of $0.005. It is still unclear exactly what constitutes a contact, since this is a connectionless arrangement to start with and is not parallel to a telephone call. It appears, however, that these charges will stack up very quickly in the AO/DI model. The set-up charge can be eliminated by paying a *monthly virtual circuit charge* of $5. This applies only to one destination. So there is a $5 charge for the ISP, another $5 for the office, a third for an alternative ISP, etc. These charges are in addition to the $2 per month charge just for having the X.25 packet switching service.

Now the *kilosegment charge* kicks in. When all the elements are decoded, a kilosegment contains 64,000 bytes (characters). This article has about half that number—32,000. Quickly checking a random sample of electronic mail and routine business letters uncovered a rough average of about 1,000 characters each. The author receives and sends on the order of 50 such messages each day, and many are substantially longer than 1,000 characters.

It's pretty clear where all this is headed. A user who switched to AO/DI on this tariff scheme would spend $3.125 per month on only one megabyte of mail. Add to that the $5 charge for set-up and $2 to have packet service in the first place, and the user is paying $10.125 for the privilege of saving money for the phone company and the ISP. Then, what about all those news feeds and small Web hits that don't trigger a move to B-channel use? There is a negative incentive to use AO/DI under these conditions.

This disincentive can be added to mounting evidence of economic distortions produced by regulated monopoly prices based on little imagination about the prospects of technology. The tariff in question was really designed for credit card verification schemes using ISDN lines with no B-channels (0B+1D). The monthly charge for such service is $31, plus set-up, kilosegment, and taxes. In such cases, packets are sent to only one location, the banking system's verification center ($5 per month for the set-up); the number of bytes of information is very small: merchant number (15 characters?), card number (15 characters?), amount (5 characters?), and verification number (15 characters?). Throw in a few additional bytes for overhead, and the size averages 64 characters—1,000 transactions could be handled for about $0.20. This isn't a bad deal for the merchant. But the scheme is a horrible deal for the Internet user with AO/DI.

Clearly, any additional charge for X.25 packet switching when a user is already paying for B-channels mitigates against the use of AO/DI. This is especially true in a market such as that of Southwestern Bell in Texas, where there is no added cost for the abomination called dedicated ISDN by its proponents. What incentive is there for users to use AO/DI? None! Yet, AO/DI offers to markedly reduce switch congestion by routing low-volume traffic around the main switch, while limiting B-channel activity to shortened times of real need.

Similar anomalies exist in many, if not most, other tariffs regarding X.25 packet switching. It is mandatory that VIA and others work with regulatory bodies, local exchange carriers, and interexchange carriers to remove these barriers before introduction of AO/DI. In the Texas example, a rough first cut would eliminate any added costs for X.25 packet switching for all users who had BRI provisioned for 2B+D service. This would result in:

- Reduced network congestion to the phone companies.
- Increased Internet performance to users.
- Reduced fixed line and equipment costs to ISPs.

This one change in X.25 tariffs would go a long way toward getting users to use AO/DI when it becomes available. While it might affect performance, this may not be enough to encourage the effort to install and configure AO/DI. Something more is needed. The telephone companies could rig more tariffs to discourage non-use of AO/DI, but this is more of the same moss-encrusted thinking that's no longer practical in today's world.

A better alternative is within the control of the ISPs. In fact, the ISPs have the most to

This disincentive can be added to mounting evidence of economic distortions produced by regulated monopoly prices based on little imagination about the prospects of technology.

gain from introduction of AO/DI because the number of B-channels and the quantity of equipment to support them will drop dramatically with use of AO/DI. ISPs hold the cards to play so they, the phone companies, and the users all win. ISPs must design a rate structure that encourages use of AO/DI. The savings a user implementing AO/DI might realize could be half or more of the cost of "traditional" dial-in ISDN configuration. Rates for an AO/DI user could be so drastically low compared with dedicated ISDN that only the certifiably insane would continue with that abomination.

"Hold it! Hold it!" Some of the ISPs screech, "My dedicated ISDN users have Web and mail servers that have to be connected full time to the Internet." *Balderdash!* Such ISPs haven't gone to the trouble to configure dial-out ISDN from their end that accepts an Internet packet addressed to the user's equipment and triggers the ISP's router to place a call and pass the traffic to the user's servers.[20] It's time to stop looking at ISDN as just a somewhat higher data rate carrier for packets. Take advantage of the sophistication available on the D-channel and within the routers.

> Rates for an AO/DI user would be so drastically low compared with dedicated ISDN that only the certifiably insane would continue with that abomination.

Some ISPs with real courage will choose to push the numbers far enough to find that it costs less to support ISDN-AO/DI users than POTS users. Then, in a very wise competitive move, such ISPs will reduce their ISDN-AO/DI rates to less than those for POTS. Why not? It's true.

Changing High Tariffs for Short Connection Times

The preceding example posited a tariff environment without metered time charges for use of B-channels. While such a situation exists in Texas and a number of other venues, users on both coasts and in several other areas of the United States are plagued with the added complexity of determining the effect of time charges. The intent of this section is not to argue for or against such metering in general. Rather, the focus is on a single aspect all-too-frequently inherent in such tariffs and severely detrimental to the economics of AO/DI and even current time out efforts that seek to use B-channels efficiently.

Many customers are charged several times the ongoing per-minute metered charge for the first minute of connection. Typically, one finds the ongoing rate is $0.01 per minute, but the first minute of connection costs $0.02, or $0.03, or $0.05! This is ludicrous. The idea of charging more for the first minute was born when human operators had to set-up each call; clearly, the cost of set-up then was substantial. But, with current technology, the realistic cost of set-up is of no significance compared with the cost of encouraging longer connection times than absolutely necessary to pass the desired traffic.

Rules of thumb applied by users so unfortunate as to be subjected to these draconian measures from out of the distant past generally encourage holding connections for a period based on the ratio of the first minute charge to that of the succeeding minutes. For example, if the first minute charge is five cents and each additional minute is one cent, users would adjust their usage to hold the connection for six minutes. The first minute is the set-up (five-cent minute), and the remaining five minutes cost five cents, for a total of ten cents. The user is likely to want more information during that time and would prefer to avoid being hit for a nickel every two or three minutes. If the set-up fee were three cents, the connection holding time would be four minutes and the charge six cents. This is just one model; more sophisticated analysis of a user's actual pattern of use will yield a more accurate strategy, but most small users can't afford to pay for the analysis.

All of this is absurd. It results in overuse of the public switched network, whining from carriers that their resources are being swamped by the Internet, yells from ISPs that they need more lines and equipment, and carping from users that this Internet stuff sure costs a lot. It's time to consign the penalty for the first minute to the same bin where all the corded, tip, and ring switchboards have gone. We no longer need those switchboards, and they're the only justification for charging more for the first minute. Get rid of this imbalance in charges.

Considering Issues of Security and Authentication

What about authentication over the D-channel? Sadly, this is not in the VIA proposal. In fact, it's not fully in control of VIA. The Internet Engineering Task Force is in charge of Requests for Comment (RFC) that define issues surrounding authentication. The primary standard is RFC 1990,[21] "The PPP Multilink Protocol (MP)" dated August 1996. This RFC provides that requests for authentication may take place whenever a B-channel is added to a bundle of B-channels. However, this is *not* an absolute requirement.

> To *securely* [emphasis supplied] join an existing bundle, a PPP authentication protocol [RFC 1994, obsoleted RFC 1334 cited in the original] must be used to obtain authenticated information from the peer to prevent a hostile peer from joining an existing bundle by presenting a falsified discriminator option.
>
> This option is *not* [emphasis supplied] required for multilink operation. If a system does not receive the Multilink MRRU option, but does receive the Endpoint Discriminator Option, and there is no manual configuration providing outside information, the implementation MUST NOT assume that multilink operation is being requested on this basis alone.
>
> As there is also no requirement for authentication, there are four sets of scenarios:
>
> (1) No authentication, no discriminator: All new links MUST be joined to one bundle, unless there is manual configuration to the contrary. It is also permissible to have more than one manually configured bundle connecting two given systems.
>
> (2) Discriminator, no authentication: Discriminator match MUST join matching bundle, discriminator mismatch MUST establish new bundle.
>
> (3) No discriminator, authentication: Authenticated match MUST join matching bundle, authenticated mismatch MUST establish new bundle.
>
> (4) Discriminator, authentication: Discriminator match and authenticated match MUST join bundle, discriminator mismatch MUST establish new bundle, authenticated mismatch MUST establish new bundle.[22]

It's time to consign the penalty for the first minute to the same bin where all the corded, tip, and ring switchboards have gone.

These four "scenarios" with no or various forms of authentication provide the basis for options within the AO/DI technology. While AO/DI controls addition of the B-channel by using D-channel packets, current implementations and the judgment of its authors seem to run in favor of each added B-channel being authenticated with the CHAP procedures of RFC 1994. This puts performance at the mercy of the authentication processes running on the end machines each time a B-channel is added. While it takes a small fraction of a second for ISDN to process the request to add a B-channel, as already stated, it often takes several seconds for authentication to allow use of that B-channel.

It is suggested that the four scenarios outlined in RFC 1990 be provided as options for operation of AO/DI. Then, each user in conjunction with an ISP (or, in the case of remote office operation, in conjunction with MIS) determines the level of security needed for the personal or business environment at hand, recognizing the potential performance/security trade-offs and potential costs. Upon full consideration, additional options or security measures may appear and should also be incorporated. The goal here is to realize a system where authentication takes place once—when the D-channel becomes active in a packet switching session with the ISP or office.

Parting Shot

Fully implemented, Always On/Dynamic ISDN has the prospect of providing the most effective use of the public switched (and non-switched!) network to date. Its rapid adoption, concurrent with appropriate tariff modifications and security options, will go a long way toward realizing the full potential of ISDN.

The sophisticated control developed within ISDN itself will become the basis for the most effective use and integration of developing technologies offering higher bandwidths at lower cost—xDSL, cable, satellite, and wireless.[23] None of these technologies, so feared by the prophets of doom for ISDN, have the sophisticated power of control already embedded in the ISDN D-channel. None need develop that power.

Its rapid adoption, concurrent with appropriate tariff modifications and security options, will go a long way toward realizing the full potential of ISDN.

ISDN is here, ready, now. AO/DI provides the linchpin securing ISDN's place as the master control system choosing among all other technologies as needed to provide optimal communications at least cost second by second by second, benefiting service providers, carriers, and users. ntq

Ed. Note—Jim Bryce may be contacted through his Web site at [http://www.bryce.com], *which furnishes items of interest to those pursuing communications research.*

[1] Some trace the beginning of the end of total domination and control of telephony by "The Telephone Company" (AT&T) to Carterphone Device in Message Toll Telephone Services, 13 F.C.C. 2nd 420, *recon. denied*, 14 F.C.C. 2nd 571 (1968) that broke the dam erected against connection of any but telephone company supplied equipment to the public switched network. Others trace it back even earlier to Hush-A-Phone, 20 F.C.C. 391, *reversed*, 238 F2nd 266 (D.C. Cir. 1956) on remand, 22 F.C.C. 112 (1957) where the reversing court allowed the addition of a plastic ring to the transmitter of a telephone handset for the purpose of stopping sounds from traveling to the called party. Since no electrical connection was made in Hush-A-Phone, proponents of Carterphone for first-over-the-dam seem to have a good argument. In any case, the next chink was supplied by MCI, 18 F.C.C. 2nd 953 (1969), *recon. denied*, 21 F.C.C. 2nd 190 (1970) holding that MCI could carry telephone calls over its own microwave links and leading to competitive interexchange carriers. The final blow (and during the same time as the first wave of ISDN) was United States vs. AT&T, 552 F. Supp. 131 (D.D.C. 1982) that resulted in divestiture of AT&T and creation of the regional Bell operating companies (RBOCs). After more than a decade, it became apparent that all these changes called for a total overhaul of the Federal Telecommunications of Act 1934; this resulted in the Federal Telecommunications Act of 1996 incorporating sweeping changes designed to bring total competition throughout the telecommunications fabric of the United States. Similar changes are underway in many other countries and in international treaties.

[2] N. Negroponte, *Being Digital* (New York: Knopf, 1995).

[3] See [http://www.ocn.com/ocn/niuf/niuf_top.html].

[4] See [http://www.nist.gov].

[5] The Internet has opened a vast store of information that makes research far faster and easier than just a few years ago. For example, here are just a few Web sites that may be used as starting point for finding out much more about ISDN:

The author's site [http://www.bryce.com].
Bellcore: [http://www.bellcore.com/ISDN/ISDN.html].
Dan Kegels's ISDN page:
[http://alumni.caltech.edu/~dank/isdn].
Texas ISDN Users Group: [http://www.tiug.org].
California ISDN Users Group: [http://www.tiug.org].
Vendors' ISDN Association: [http://www.via-isdn.org].

[6] See [http://www.itu.ch]. Unfortunately, most of the standards maintained by the ITU are available only on a subscription basis at fairly substantial costs. The ITU is an arm of the United Nations that is financed by tax money. The statutes, regulations, and case law of many jurisdictions in the United States are now available at no charge over the Internet. Something is amiss in this picture.

[7] B. Kumar, "Impact of Internet Traffic on Public Telephone Networks," *New Telecom Quarterly*, Vol. 5, No. 1 (February 1997):41-50.

[8] Kumar, p. 43. CCS is an acronym for Centum Call Seconds. This is a measure developed by the telephone industry to determine system loading. One CCS is 100 seconds of connect time.

[9] Kumar, pp. 47-48.

[10] J. W. Bryce, *Special Edition, Using ISDN, 2nd Edition* (Indianapolis: Que Corporation, 1996), pp. 67-88. G. C. Kessler and P. Southwick, *ISDN—Concepts, Facilities, and Services, 3rd Edition* (New York: McGraw-Hill, 1997), pp. 37-334. W. Stallings, *ISDN and Broadband ISDN with Frame Relay and ATM, 3rd Edition* (Englewood Cliffs, NJ: Prentice-Hall, 1995), pp. 94-351.

[11] Bryce, pp. 521-538; Kessler, pp. 335-350; and Stallings, pp. 408-453.

[12] An Internet standard is generally defined by a Request for Comment (RFC). These are available on the Internet by beginning at the Internet Engineering Taskforce site [http://www.ietf.org] and selecting appropriate links. The site at [http://www.graphcomp.com/info/rfc] is one of a number of helpful indexes. CHAP is currently defined by RFC 1994. This made RFC 1334, the previous definition of CHAP, obsolete. The indexing system keeps track of such updates; it is important to make sure the most recent RFC is consulted.

[13] J. Van Tassel, "Rebooting the Regulatory Operating System: The Computer Industry Turns on the Power," *New Telecom Quarterly*, Vol. 5, No. 1 (February 1997). This article is available for free download at [http://www.ntq.com] where you merely register and may obtain either an html or pdf form of the material.
[14] A. De Tocqueville, *Democracy in America*, originally published in 1840 and considered an enduring classic. It is available in numerous editions currently in print. Lord James Bryce (collaterally related to the author of the instant article) wrote *The American Commonwealth* several decades later; it is also an enduring classic often cited alongside De Tocqueville. Unfortunately, Lord Bryce's work is not noted for its observations on the American obsession with associations.
[15] Bryce, pp. 499 and in the *1st Edition* published in 1995 at 439: "A Brief Excursion Into Virtual Cyberspace Intended to Blow Your Mind:

"ISDN provides control signals over the D-channel, and the D-channel is always up. If you don't believe this, look at Figure 10.35 in Chapter 10, "Putting an Adapter Inside Your Computer"; this is a graphic showing the D-channel up when you plug into ISDN regardless of whether or not you've made any B-channel connections.

A system could be set-up so users are always logged in, but B-channels are not connected until requested by the D-channel. In effect, this creates a full-time "virtual connection," but the resource that may have a metered charge, B-channels, is used only as needed. Security and authentication take place when the system is installed and the D-channel is up, not on each B-channel connection.

Now for the *coup de grâce*; the D-channel can maintain 255 virtual sessions simultaneously. So you could have your system virtually connected to 255 sites at the same time, popping among them and merrily delivering information from and to everywhere all at once. Yet, you'll only be paying for the actual B-channel when it really is passing traffic because the D-channel pops B-channels on and off to wherever needed at will within a fraction of a second. Think of it as a giant combination of call holding, call waiting, and conference call, with your equipment hopping all over and all the lines staying up, so it's real fast.

Can you get this ready to go today? Not as far as I know. Can it be built with existing technology. Yes. Will it be? Yes. When? Soon. What economic effect will it have? Big.
[16] A means of tying B-channels together created by the BONDING Consortium (Bandwidth ON Demand INteroperability Group) at [http://www.hep.net/ftp/networks/bonding/aaareadme-bonding]. Also, see Bryce, pp. 420-424.
[17] RFC 1990. See note 12 for instruction on obtaining an RFC over the Internet.
[18] A. Kuzma, "Always On/Dynamic ISDN," Vendors ISDN Association, October 1996 [ftp://via-isdn.org/pub/out/aoac_rfc.doc] for a description of AO/DI on the local loop. See also "Always On/Dynamic ISDN Network Architecture," December 1996 [ftp://via-isdn.org/pub/out/ao-dib~4.doc] for a description of AO/DI at the central office and ISP. See also "Always On/Dynamic ISDN," December 1996 [ftp://via-isdn.org/pub/out/aoac_r~4.doc] for the most recent and detailed information.
[19] The author wishes to express special thanks to Jorge Sedeño, attorney for Southwestern Bell specializing in issues of tariff. Mr. Sedeño went to great lengths at terrific speed to clarify these tariff issues. He has advised Southwestern Bell on the issues involved in the AO/DI technology in anticipation of a need for tariff modifications. Any errors in or misinterpretations of tariffs are the responsibility of the author.
[20] Bryce, pp. 498-514.
[21] See note 12 for instructions on obtaining an RFC over the Internet.
[22] K. Sklower, et al., "The PPP Multilink Protocol (MP)," RFP 1990 (August 16, 1996). Available on-line from Internet Engineering Taskforce at [http://www.ietf.org].
[23] For more information, see [http://www.bryce.com/adsletal.htm].