15

Optical Internets: Evolving to a 3G Architecture

The subject of IP/MPLS/optical interworking resurfaces in this chapter in relation to the migration of the Internet to a 3G transport network. A 3G transport Internet? This idea means the Internet takes on the characteristics of traditional transport networks such as those that are supported by the traditional carriers, such as AT&T, MCI Worldcom, BT, Mercury, France Telecom, etc.

At this point in the evolution to 3G, the Internet is not taking on the complete role of these transport networks. But it is likely to take on many of the characteristics of legacy transport networks, and it will surely evolve to interwork with the vital SS7 and AIN/IN resources that are owned and operated by the traditional carriers.

The first part of the chapter picks up on the IP/MPLS/optical interworking subject and provides some more examples of how this powerful protocol assembly will evolve. The next part deals with the issues of Internet and legacy transport network interworkings, and brings the IP-based call processing architectures into the discussion.

MIGRATION TO IP OPTICAL NETWORKING

Today's backbone networks are made up of SONET/SDH technology, with the access networks operating at OC-3 and OC-12, and the long-haul systems operating at OC-48 and OC-192. It is anticipated that

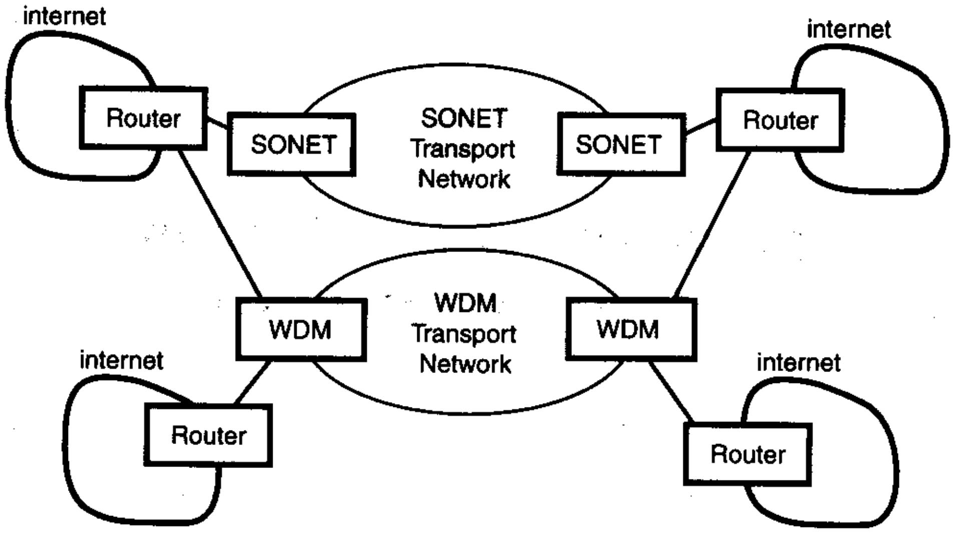

Figure 15–1 Migration to IP optical networking.

WDM will be deployed with IP operating directly over it in the near future. Figure 15–1 shows a probable migration scenario. An edge router rests between the internets and the optical networks. Based on the destination of the traffic and available technology, the traffic is either sent through the SONET TDM transport network or the WDM transport network. This approach is attractive because a migration can take place without disruption to the existing transport network.

IP AND THE OPTICAL BACKBONE

Figure 15–2 shows an example of an optical backbone network that supports IP. This technology is new and has not yet seen extensive deployment. Nonetheless, a number of companies are developing optical routers, and the U.S. Department of Defense Advanced Research Projects Agency (DARPA) is funding an effort among several companies and universities. It is called all-optical label swapping (AOLS).

The routers attached to the optical backbone routers are connected through copper or fiber. For this example, let's assume the router has its interface to the network edge node configured with WDM. The optical router is tasked with managing each WDM channel in relation to the

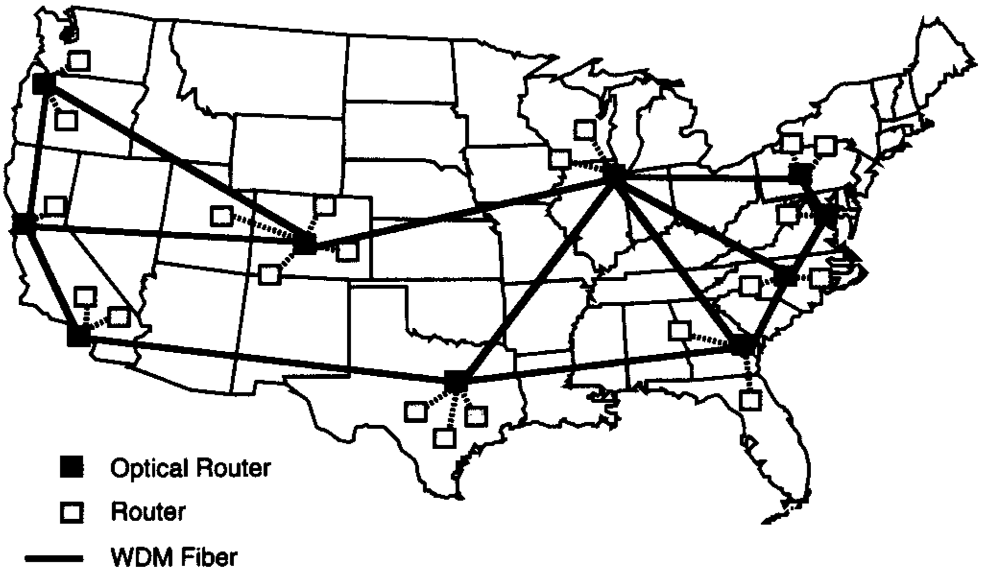

Figure 15–2 Example of an optical/IP backbone.

router interfaces as well as the WDM interfaces associated with the backbone links.

Channels must be added, dropped, bridged, and passed through, perhaps on all the interfaces. In addition, these operations should reflect the state of the network traffic.

Example of IP and λ Forwarding

Figure 15–3 shows a scheme for supporting IP traffic through the optical network. Be aware that all possible add, drops, inserts, and pass-throughs are not shown, nor are the channels coming from the Chicago node. Notwithstanding this general view, the depiction accurately reflects a IP router-to-λ configuration. These routers are configured to support WDM.

The following channel relationships exist:

- *Node A:* This router adds a WDM channel to the WADM interface to Chicago. The channel is also dropped at node B.
- *Node B:* This router adds a WDM channel to the WADM interface to Chicago. In turn, this channel is dropped at node A. Since the channels from nodes A and B are made available to these two

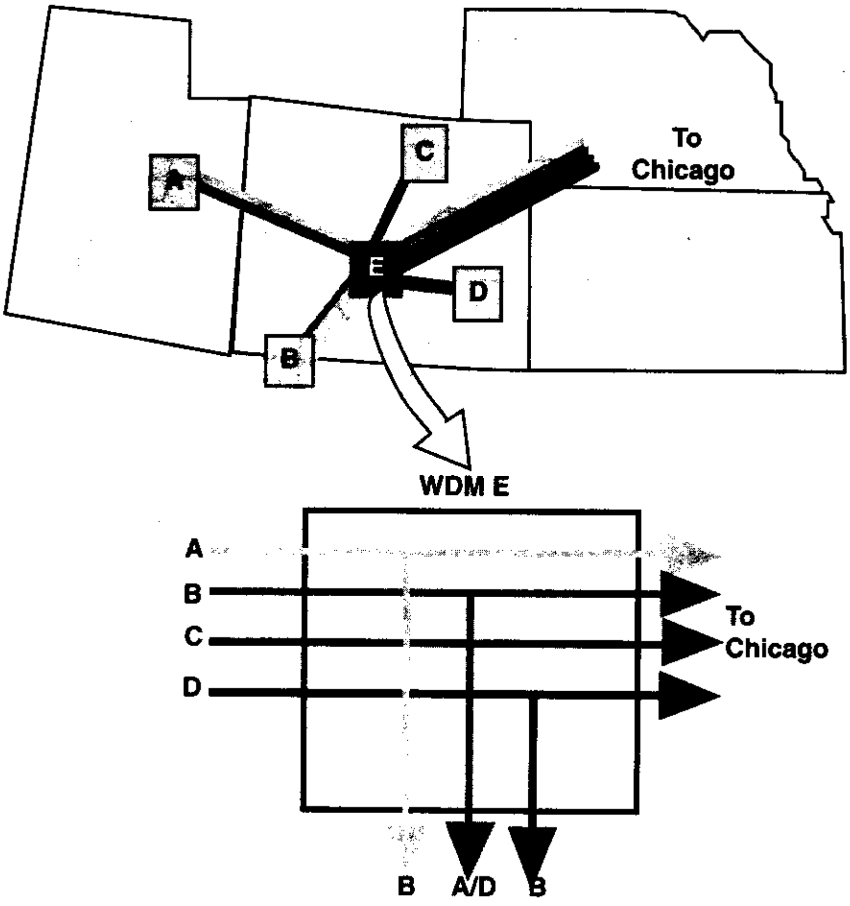

Figure 15–3 Adding, dropping, and cross-connecting IP traffic between routers.

nodes, we can assume they have a lot of traffic to exchange with each other. Node B also has a drop relationship with Node D.

- *Node C:* Node C has no WDM channel relationships with the other router nodes, only to Chicago. Perhaps this router connects to sites that process classified information and is not made available to the other routers.
- *Node D:* Node D has a WDM channel dropped to node B, as well to the interface to Chicago.

Therefore, the following WDM channel relationships exist between these four IP/λ nodes:

- A ⇒ B
- B ⇒ A
- B ⇒ D
- D ⇒ B

IP Subnets

IP and associated protocols, such as OSPF, TCP, etc., use IP addresses, and forwarding and route discovery operations are based on using these addresses. An IP optical network must be able to correlate destination addresses in the IP packet to a route in the optical network.

One possibility of how this service might be realized is offered in Figure 15–4, which shows that traffic has been forwarded from the optical router in Denver to the Chicago node. This node is tasked with forwarding the IP traffic to nodes G, H, I, and J. At these nodes are IP-based internets, identified with the IP address prefixes shown in the figure (some IP addresses are reserved and/or set aside for special use; this simple example assumes all these addresses are available).

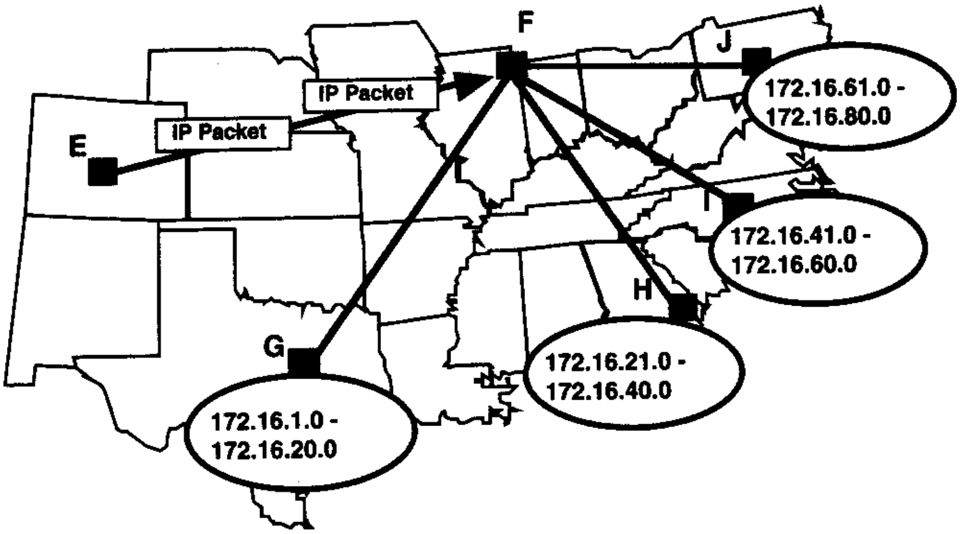

Note: For simplicity, some reserved addresses are used in this example.

Figure 15–4 IP subnets.

This address configuration presents some attractive implementation opportunities. Since the address prefixes are contiguous, it is possible to correlate a wavelength with a group of address prefixes.

However, this view is not one shared by AOLS, and I do not think mapping prefixes to a wavelength is the best long-range solution. A better approach is to use labels and map the labels to a wavelength (as defined in Chapter 12). We will examine this operation shortly.

Support of Non-optical Nodes

The previous example is altered slightly. Figure 15–5 shows that the optical node E in Colorado also supports native-mode IP on, say, DS3, or OC-3 interfaces. In order to route the IP traffic to the correct destination, node E is configured with a mapping table. This table associates IP address prefixes with a wavelength.

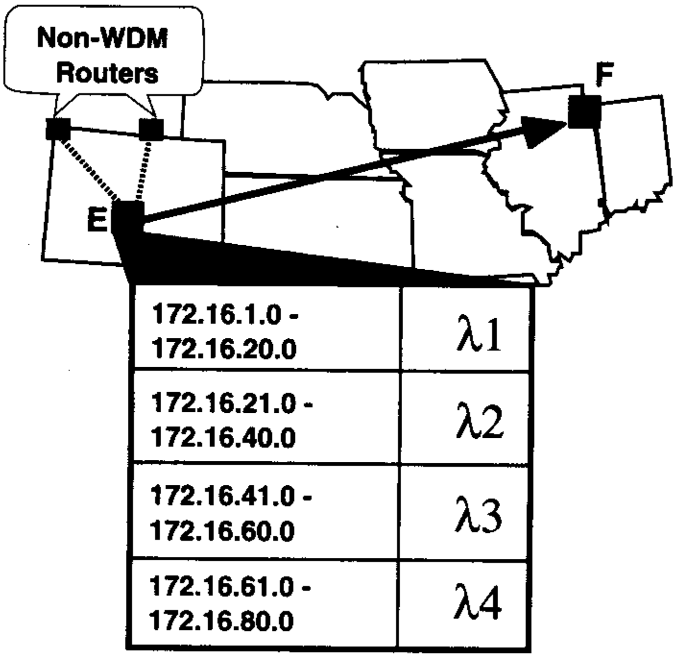

172.16.1.0 - 172.16.20.0	λ1
172.16.21.0 - 172.16.40.0	λ2
172.16.41.0 - 172.16.60.0	λ3
172.16.61.0 - 172.16.80.0	λ4

Figure 15–5 Node E accepts native IP packets and maps to WDM channels.

When an IP packet is received at node E, it examines the IP address in the packet header and maps the IP packet onto the associated WDM channel.

This process is performed at an edge device, that is, a node that sits at the boundary of the WDM network. In some of our previous examples, node E also acts as a node in the backbone WDM network.

The decision to use a node for both an edge and a backbone node should be weighed carefully. The IP address analysis and the associated IP-to-WDM channel mapping may consume more overhead than an optical backbone node can afford. In that case, the task of IP-WDM mapping can be pushed out to another router, assuming that the router is so designed. This scenario was shown earlier.

Some of the WDM channels on the fiber between nodes E and F are shown in Figure 15–6; they represent an expanded view of previous illustrations. We are using four wavelengths for the purpose of explaining IP-WDM concepts. The actual number of wavelengths, of course, is implementation-specific.

As shown in Figure 15–7, the incoming IP packets arrive at the link interfaces at node E. Through the use of its IP/λ mapping table, the node places the packets onto an outgoing queue that is associated with a wavelength on a specific outgoing interface. The next task is to transmit these packets out of the optical interface. This task entails an electro-optical conversion process wherein the electrical bits in the buffer are translated to the associated optical bits for transmission onto the fiber.

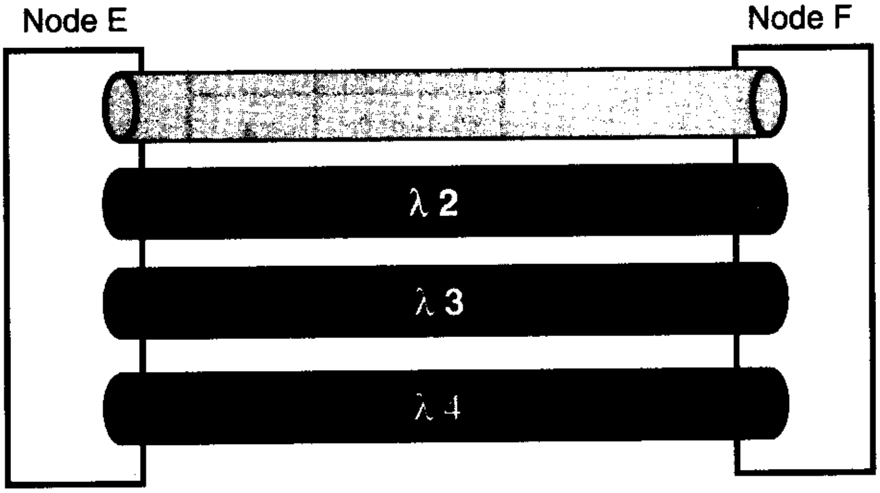

Figure 15–6 Some of the WDM channels.

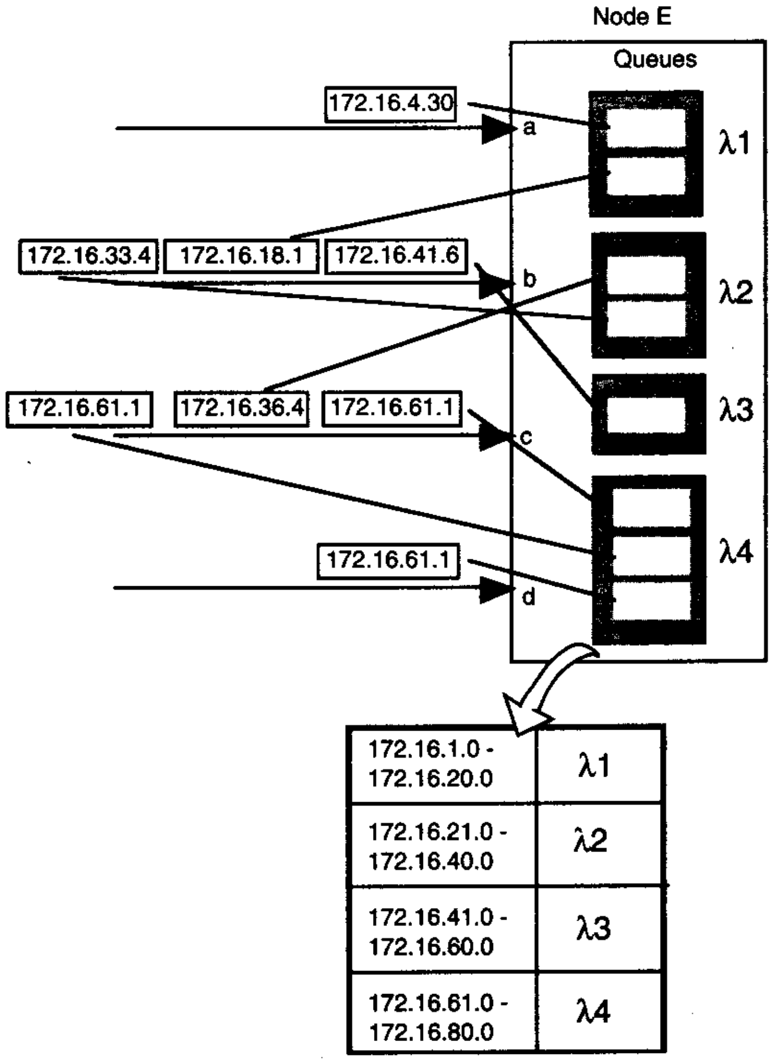

172.16.1.0 - 172.16.20.0	λ1
172.16.21.0 - 172.16.40.0	λ2
172.16.41.0 - 172.16.60.0	λ3
172.16.61.0 - 172.16.80.0	λ4

Note: Some IP addresses are reserved; this example assumes that all addresses shown are available for use.

Figure 15-7 Mapping the IP addresses to the WDM channels.

PLACING MPLS INTO THE PICTURE

Thus far, the examples show the mapping of IP prefix addresses onto wavelengths. This operation is certainly feasible, and some networks that do not need the features of MPLS can opt for this implementation. For other networks, another approach is to: (a) first map IP addresses to MPLS labels, and (b) then map the labels to wavelengths. The reason for bringing MPLS into the picture is to exploit the traffic engineering and scaling capabilities of MPLS, as well as other attributes explained in Chapter 9.

As shown in Figure 15–8, a set of IP addresses (or address prefixes) can be aggregated into a single MPLS label. The MPLS specification provides a number of rules on aggregation and on how the labels are set up between label switching routers (LSRs). For this discussion, perhaps you can see the opportunity to use WDM channels to support label switching. We will examine this idea next.

If label switching networks continue to grow (and they surely will), there will be opportunities to interwork and combine label switching routers (LSRs) with optical routers. As Figure 15–9 shows, native-mode IP networks will continue to exist for many years; in fact, there is little incentive to push label or WDM technology into local area networks or conventional point-to-point local loops. To do so would entail changing

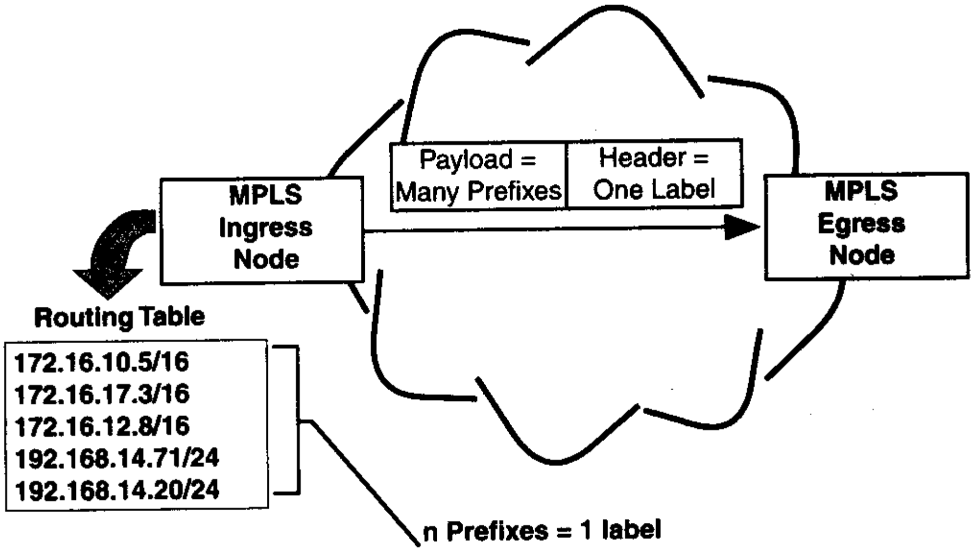

Figure 15–8 Aggregating addresses into a label.

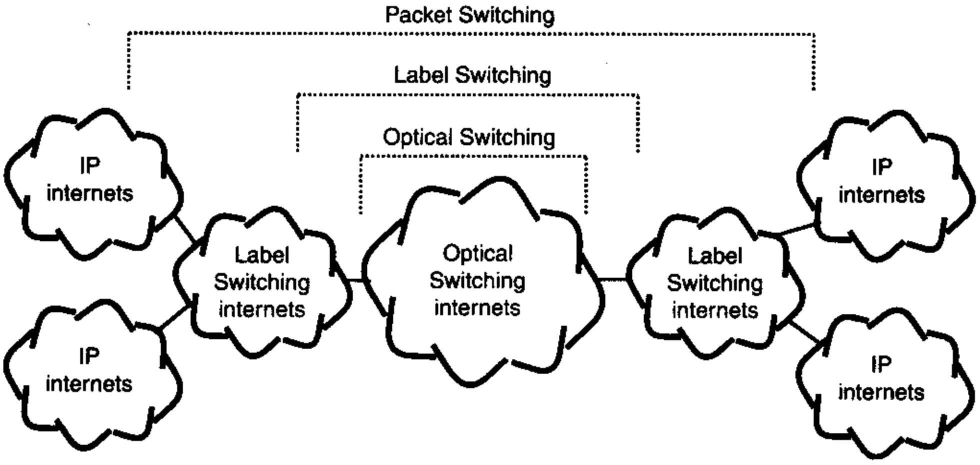

Figure 15–9 Interworking WDM and label switching internets.

the software and hardware architecture of user machines, such as PCs, palm units, etc.

So, this part of an internet stays the same, and the interfaces from a router back to the user computers are conventional Ethernet, PPP, DSL, V.90, cable modem, and so on. The router's interfaces out to the network will be MPLS, WDM, and probably a melding of the two.

This nested approach allows the network administrator to set up different switching domains in an organized fashion. In Figure 15–9, the packet switching domain uses IP in the data plane and OSPF, IS-IS, and BGP in the control plane. The label switching domain uses MPLS forwarding in the data plane and CR-LDP, OSPF (extensions), and RSVP-TE in the control plane. The optical switching domain uses lambda OSP switching in the data plane and GMPLS, LMP, and perhaps others (as they emerge) in the control plane.

Figure 15–10 shows how the different internet nodes interwork with each other. At node E in Denver, native IP packets are received from attached routers (not shown in this figure). The following events take place to relay the traffic to node F in Chicago for distribution to nodes G, H, I, and J.

- *Event 1:* Previously, OSPF, IS-IS, or BGP has discovered the addresses associated with hosts attached to nodes G, H, I and J. This information is stored in a routing table at node E.

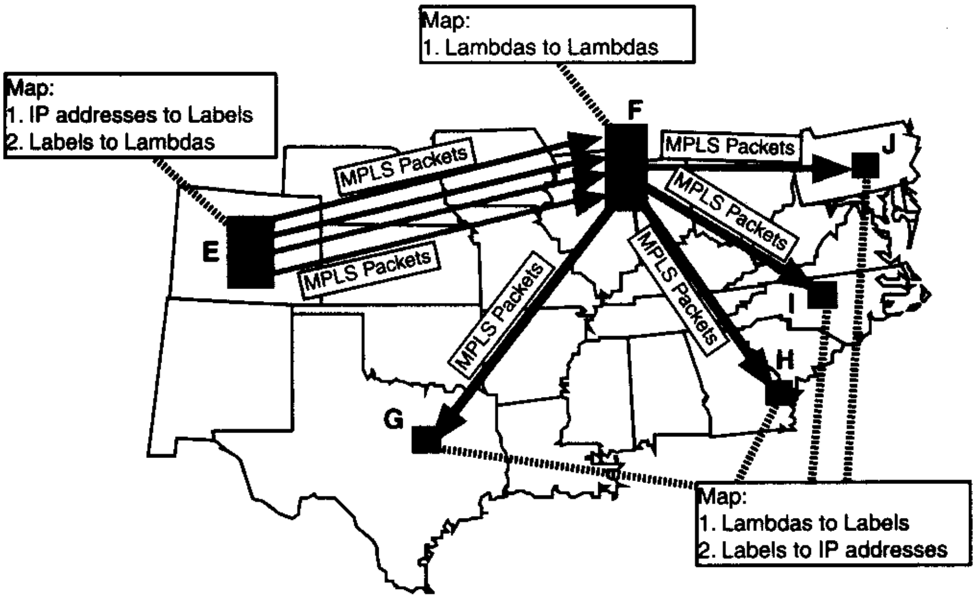

Figure 15–10 End-to-end operations.

- *Event 2:* Upon node E receiving an IP packet, it accesses the routing table to determine the next node that is to receive this packet. This next node is node F.
- *Event 3:* By previous configurations (using a label binding protocol, such as LDP or RSVP-TE), node E knows (a) the outgoing interface for this packet, and (b) the associated label. Therefore, node E appends the MPLS label header to the IP packet.
- *Event 4:* Furthermore, by previous configurations (using an λ "binding" protocol, such as GMPLS or LMP), node E also knows the specific wavelength that is to be used for this packet.
- *Event 5:* Node E maps the label to the appropriate wavelength on the appropriate interface, then sends the packet (to node F).
- *Event 6:* For this example, Node F is operating as a O/O/O PXC for data plane operations. Due to the prior execution of the optical control plane, node F knows that the wavelength associated with a specific fiber interface is to be optically cross-connected to one of its downstream neighbor nodes (that is, nodes G, H, I, or J).
- *Event 7:* Using the configured MEMS fabric (configured as a result of GMPLS or LMP), node F relays the packet to the appropriate

output interface to, say, node G. The λ used is the same on each OSP (say, λ1).

- *Event 8:* Node G in Dallas is a egress node to both the MPLS and optical routing domain. It receives the packet on a specific wavelength on its input interface with node F. Its optical and MPLS cross-connect tables (again, pre-configured earlier with the IP, MPLS, and optical control planes) reveal that the packet has reached the end of the label switching path, as well as the optical switching path.
- *Event 9:* Therefore, node G resorts to O/E/O operations, terminating the LSP and OSP, and passes the native-IP packet to the relevant subnet.

Figure 15–11 shows the operations for mapping the MPLS labels in the packet header to a WDM channel. Obviously, these concepts are similar to the IP address prefix mappings discussed earlier.

Which leads to the questions, "Why bother with the labels? Why not just map the IP addresses directly into the WDM channels?" The answer is that under some situations, it might be sufficient for the edge router to map an IP address prefix directly into the WDM channel. However, this situation may not be the same for all networks. For example, it is highly unlikely that WADM nodes will be prevalent in local and campus networks, at least not to any great extent (for a number of years). The same holds true in wide area backbones, but to a lesser degree.

The point is that label switching networks are prevalent and growing. When WDM nodes begin to be integrated into IP-based networks, we will then see a convergence of label switching and WDM.

Therefore, IP to wavelength mapping may have a niche in the industry. The prevalent view, however, is that the IP to label and then to wavelength mapping will emerge as the dominant technology.

Some additional thoughts about Figure 15–11 might prove helpful. The four wavelengths now have labels associated with them. If the optical network has the capability to negotiate end-to-end wavelength reservations, then the core optical nodes (node F in this example) do not have to process the MPLS label during data transfer. For example, λ1 is mapped from node E, through F, to G. Consequently, wavelength conversion is not required at node F.

The table at the bottom of Figure 15–11 shows that each wavelength is associated with one of the end nodes of an LSP, nodes G, H, I, or J. This information can be used by node E when it receives a GMPLS request from node F for wavelength allocations pertaining to certain sites. It can determine if the required wavelength is available for use.

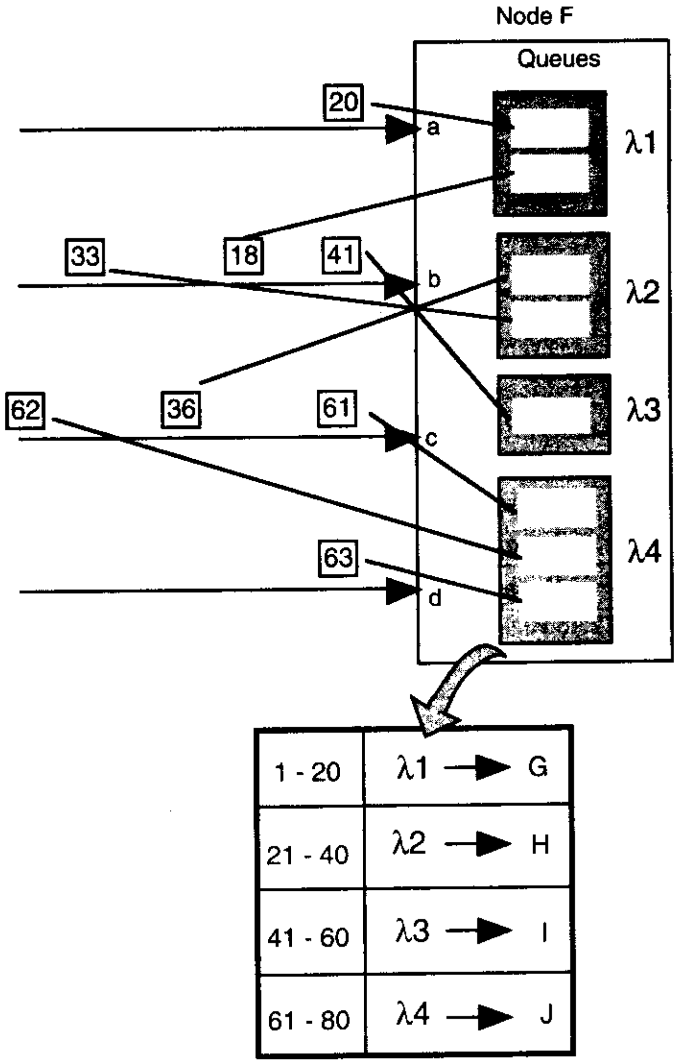

1 - 20	λ1 → G
21 - 40	λ2 → H
41 - 60	λ3 → I
61 - 80	λ4 → J

Note: Some MPLS labels are reserved; this example assumes that all labels are available for use.

Figure 15–11 Mapping MPLS labels to WDM channels.

PUTTING IT TOGETHER

To conclude the examples of optical internet operations, Figure 15–12 shows that the switching table entries provide sufficient information to move an IP datagram from node E to the next node in the path, which is node F.

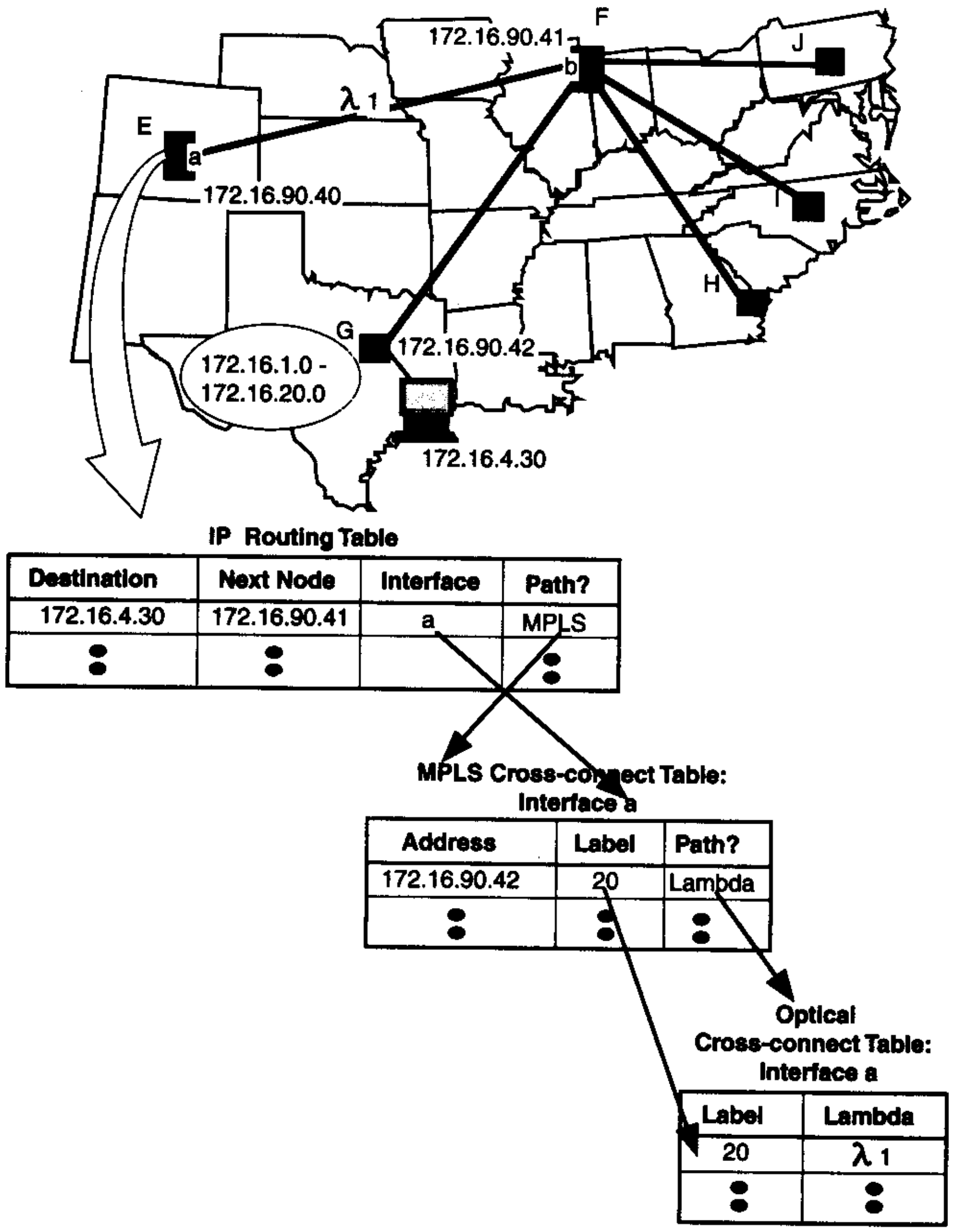

IP Routing Table

Destination	Next Node	Interface	Path?
172.16.4.30	172.16.90.41	a	MPLS
:	:		:

MPLS Cross-connect Table: Interface a

Address	Label	Path?
172.16.90.42	20	Lambda
:	:	:

Optical Cross-connect Table: Interface a

Label	Lambda
20	λ 1
:	:

Figure 15–12 The data plane tables at Node E.

The three tables are parts of the IP, MPLS, and λ data planes. Recall that these tables are created by the IP, MPLS, and λ control planes. In an operational system, the tables would contain more entries, but those shown are sufficient for our examination. The tables can be implemented in any fashion the node designer wants; for efficiency, they should be in the form of ASICs, or especially designed chip sets. Also, nodes E, F, and G have been assigned IP addresses to aid in explaining the operations in these examples.

The three tables (more formally known as forwarding information bases) perform the following services.

IP Routing Table

- *Destination:* Through OSPF, IS-IS, or BGP, node E knows about address 172.16.4.30.
- *Next Node:* Also courtesy of the IP routing protocols, node E knows that the next node to receive this IP datagram is node F, assigned address 172.16.90.41.
- *Interface:* Node E knows the output link (interface) to node F is a.
- *Path?:* The entry of "MPLS" alerts node E that the interface is configured for MPLS operations. Had it been coded as, say, "IP," then node E would have resorted to conventional IP forwarding. However, for this operation, an MPLS LSP is used. So, the Path? value states the path to node F is an MPLS path, and it acts as an index into the MPLS table for interface a. Also note that had the Path? value been coded as "Lambda," it would have alerted node E to resort to IP to λ mapping, which is not shown in this example. In any of these cases, the decisions are made by the network administrator.[1]

MPLS Cross-connect Table: Interface a

- *Address:* The address of 172.16.90.42 is the address for node G. This is the address for the next logical, adjacent MPLS node. As explained in earlier chapters, MPLS does not require an LSP to be

---

[1] Today's routers are not yet set up to allow these types of configuration commands to be entered into the router's operating system. Before long, routers will have these capabilities.

set up at each node in the path. If you have forgotten about this aspect of MPLS, review the material in Chapters 9, 10, and 12, for it is an important component in supporting the interworking of label and lambda switching.

- *Label:* Label 20 is associated with the LSP between nodes E and G. You may be wondering how this table identifies the LSP, since it contains only the end of the LSP tunnel's address (172.16.90.42). This table is created and stored by and at node E, so it surely knows its own address (172.16.90.40).
- *Path?:* This entry of "Lambda" alerts node E that the interface is configured for wavelength operations. Thus, this entry acts as an index into the optical cross-connect table for interface a.

Optical Cross-connect Table: Interface a

- *Label:* The label value of 20 is another index that points to the wavelength that is assigned to MPLS packets with labels of 20 in their headers.
- *Lambda:* The traffic can now be sent to node F on λ 1 from interface a.

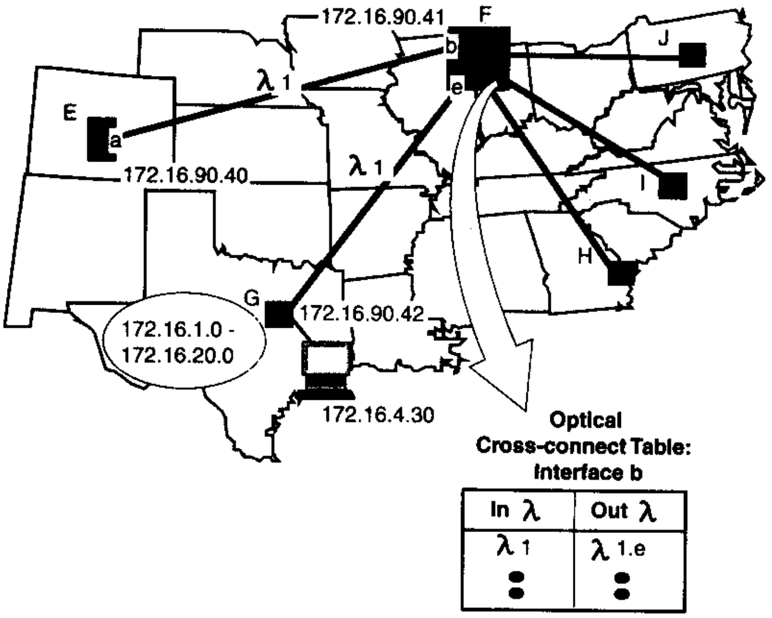

In λ	Out λ
λ 1 ⋮	λ 1.e ⋮

Figure 15–13 The data plane at Node F.

Figure 15–13 shows the optical cross connect table at node F for its interface b. The other tables are not needed for this operation, since prior control plane procedures mapped $\lambda 1$ from node E's a interface to $\lambda 1$ on node F's e interface. In this specific example, node F does not participate in the LSP between nodes E and G

The optical cross-connect table at node F, shown in Figure 15–13, is not necessarily a table. It could be the state of the MEMS mirrors in the switching fabric. To show this idea, the example is re-rendered in Figure 15–14.

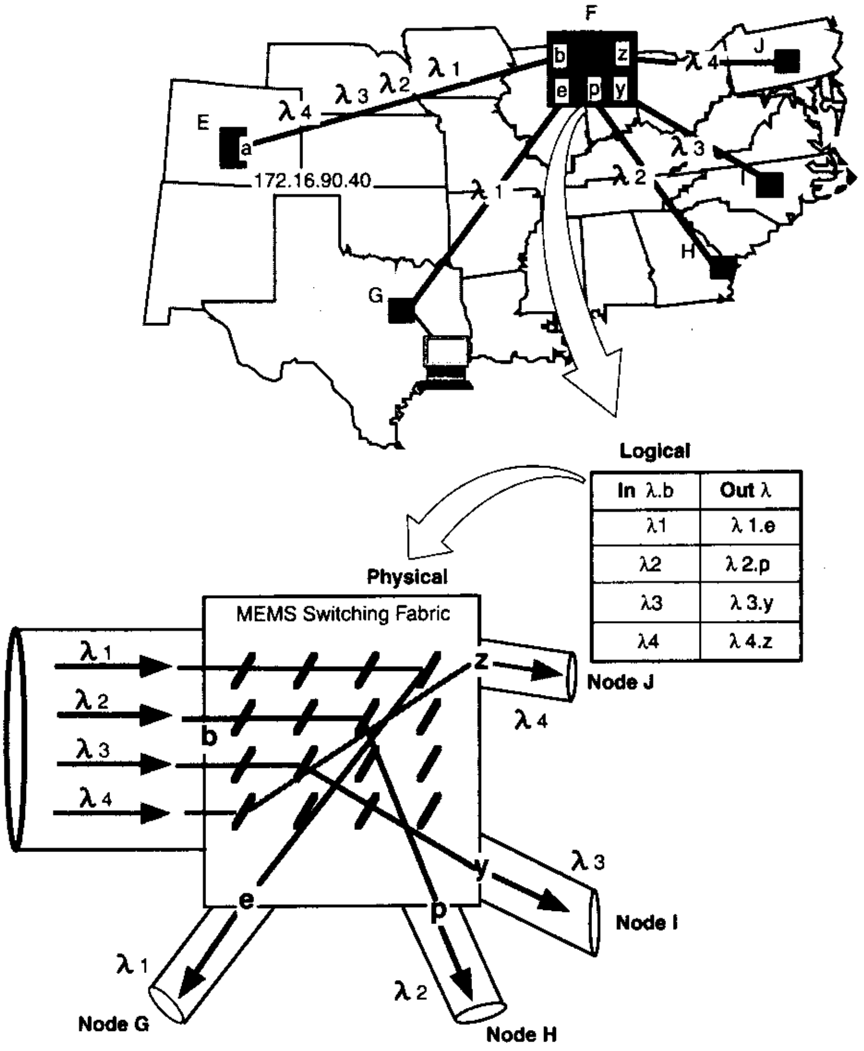

In λ.b	Out λ
λ1	λ 1.e
λ2	λ 2.p
λ3	λ 3.y
λ4	λ 4.z

Figure 15–14 The logical view and a MEMS physical implementation.

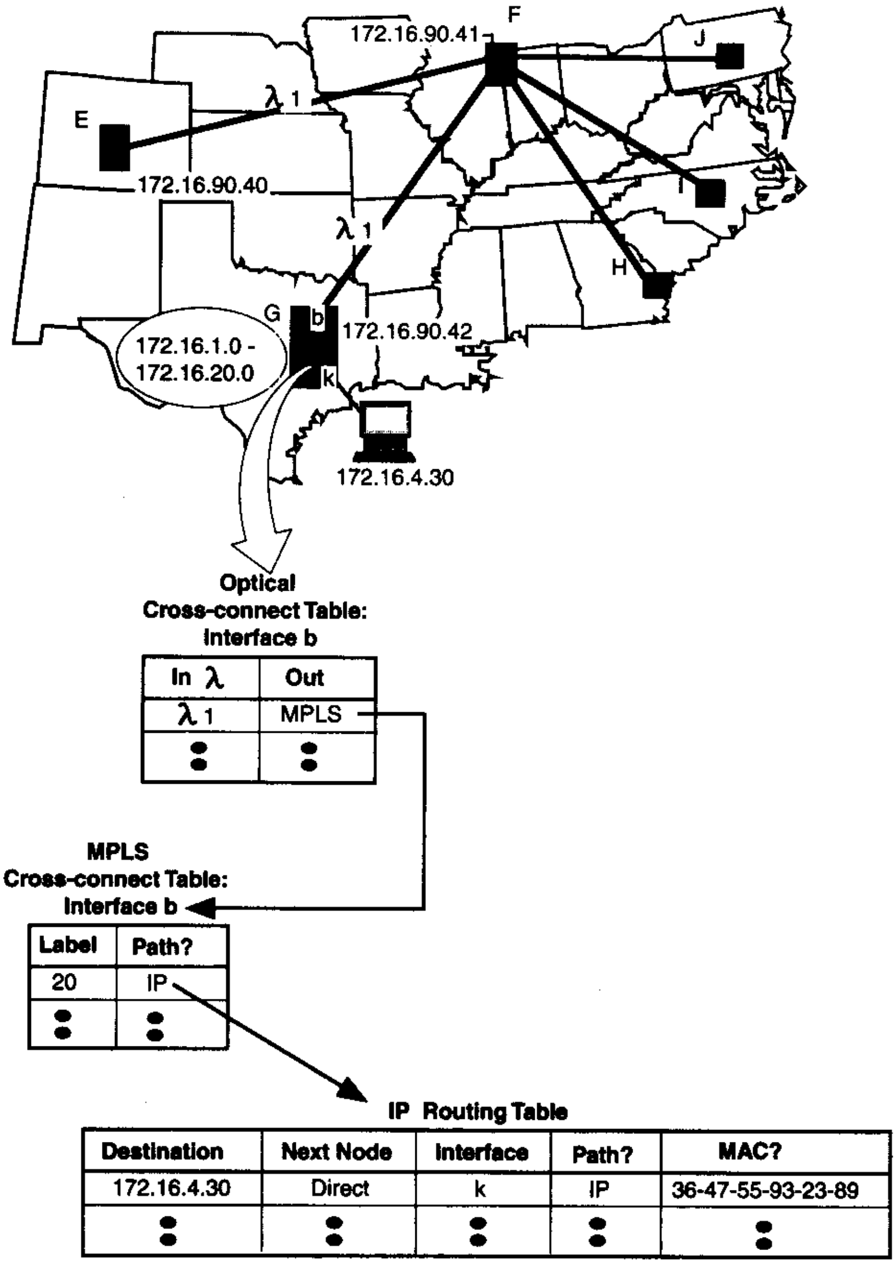

Optical Cross-connect Table: Interface b

In λ	Out
λ 1	MPLS
:	:

MPLS Cross-connect Table: Interface b

Label	Path?
20	IP
:	:

IP Routing Table

Destination	Next Node	Interface	Path?	MAC?
172.16.4.30	Direct	k	IP	36-47-55-93-23-89
:	:	:	:	:

Figure 15–15 The data planes at Node G.

Here, the optical cross-connect table shows the relationships of the incoming ports and wavelengths to the outgoing ports and wavelengths. In addition, the MEMS mirrors show how they reflect the switched OSPs from the input port b to four output ports: e, p, y, and z.

Figure 15–15 shows the data planes at node G. As we look at each entry, we see that the operations at node G are almost the reverse of those at node E. And this make good sense, since node E is the source of the MPLS LSP and node G is the end of the LSP. Also, it is possible that the tables at all nodes are similar in their structure. I have shown those entries that are relevant to this example.

Optical Cross-connect Table: Interface b

- *In λ:* Wavelength 1's OUT entry identifies the end of the lightpath.
- *Out:* The "MPLS" entry points to the MPLS table for interface b.

MPLS Cross-connect Table: Interface b

- *Label:* Label 20 coming into interface b identifies the next label to be used on the next LSP segment, or, in this example, IP procedures.
- *Path?:* The value of "IP" identifies the end of the MPLS LSP. Therefore, the node "pops" label 20 off the packet to reveal the IP header. The IP routing table is consulted.

IP Routing Table

- *Destination:* Destination address 172.16.4.30 is examined.
- *Next Node:* The next node entry reveals that there is no next node. This address is attached to a subnet on node G.
- *Interface:* The interface to which node 172.16.4.30 is connected is found on node G's interface k.
- *Path?:* Interface k is a native IP address, so it is not necessary to build an LSP to node 172.16.4.30. (If this entry were coded "MPLS," node G would create a new MPLS tunnel and index into

the MPLS tunnel for its appropriate outgoing interface, but this is not needed here.)

- *MAC?:* A LAN MAC address of 36-47-55-93-23-89 is associated with IP address 172.16.4.30. This entry means that the link on interface k is an Ethernet link. Therefore, node G can encapsulate the IP datagram into the Ethernet frame and deliver the traffic to node 172.16.4.30.

Issues in MPLS/Optical Plane Interworking

Some of the issues of interworking MPLS and the optical plane are discussed in Chapters 10 and 12. For this discussion, it is important to pick up on some of these thoughts in these chapters, and express some problems with label networks, specifically, MPLS interworking with PXCs [AWDU01] and [RAJA02].

There are no analogs of label merging in the optical domain. This implies that a PXC cannot merge several wavelengths into one wavelength. I do not view this situation as an insurmountable problem, but it will require label merging to be performed at the edge LSR. Then the label (or labels) can be mapped to a wavelength.

Most designers do not favor converting the wavelength back to electrical signals within the network, so the network operator must have a complete knowledge of the QOS requirements and (of course) the destination of the user traffic. Thus, constrained routing should be employed to set up the LSP and associated optical trails.

Another distinction is that a PXC cannot perform the equivalent of label push-and-pop operations in the optical domain. One again, this operation must be performed at an edge device and must be kept transparent to the PXC nodes within the network, as shown in these examples in relation to node F.

My final statement on this important issue is that it seems quite reasonable to expect that the core optical network will not process IP datagrams in its data plane, but it is equally reasonable to expect the core network to be able to support the interworkings of the IP, MPLS, and λ control planes. In this manner, the network is designed to give the best performance for payload processing and to have the capabilities to configure, maintain, protect, restore, and tear down the optical facilities, as they support IP and MPLS.

PROTOCOL STACK ALTERNATIVES

Figure 15–16 shows several protocol stacks that either exist now or will be deployed in the near future. My view is that all these implementations have a place in a communications network and it is a matter of deciding which protocol stack is appropriate for the specific network requirement.

There has been an increased awareness of the benefits of the SONET technology over the past few years. At the same time, it also recognized that the cost and overhead of SONET may not be warranted in some situations.

One situation that comes to mind is a simple point-to-point link between two buildings on a campus. The rich functionality and expense of SONET may be overkill for this situation. The network manager may not need all the diagnostics and alarms that go hand-in-hand with SONET.

But the issue goes further than the use or non-use of SONET for certain topologies. As we have examined in this book, the issue also involves the use or non-use of ATM. As noted earlier, some critics of ATM state that ATM has too much overhead and is too expensive for certain applications and topologies.

It is obvious that the deployment of SONET and ATM is not appropriate in many cases. But it should also be understood that an

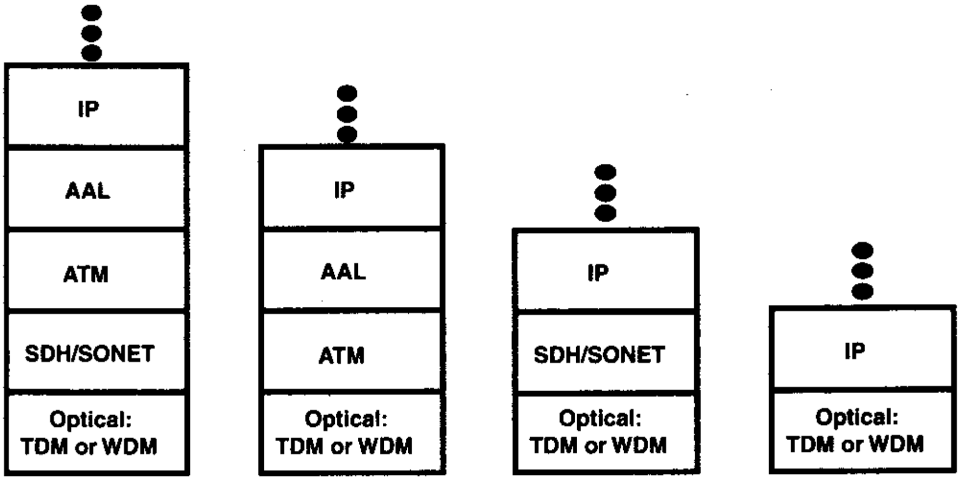

Notes: (1) IP may be running on top of PPP, not shown here.
(2) For PONs and metropolitan networks, Ethernet is replacing ATM.

Figure 15–16 Protocol stack alternatives.

implementation without SONET will not have the superior provisioning, backup, and OAM capabilities that are built into SONET. Likewise, an ATM-less implementation will not have ATM's traffic management capabilities. If they are not needed, then don't use them, but recognize the implications of such a scaled-down system and what it means to the network administrator and the customers who use the system.

Revisiting the Digital Wrapper

The protocol stack arrangements shown in Figure 15–16 are accurate depictions of the alternatives. However, as shown in Figure 15–17, there should be another layer in these protocols stacks for the following situations: (a) The SONET layer is removed, and (b) WDM is employed. If SONET is not used, there is no method to organize the traffic on the

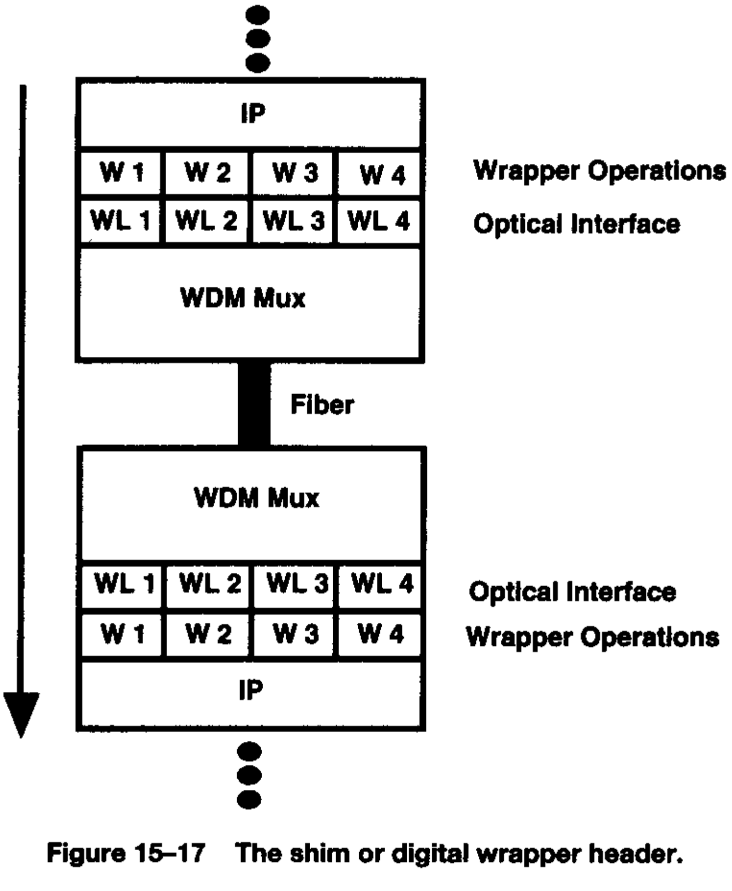

Figure 15–17 The shim or digital wrapper header.

fiber. Remember that SONET has headers and slotting arrangements that identify the virtual tributaries (VTs). If WDM is employed, it makes sense to add a header to the traffic as it enters the optical transmission system in order to carry information about the type of traffic and where the traffic is going.

Figure 15–17 shows these ideas. The header is called a shim header, or a digital wrapper. I am using the "lean stack" in this example: IP traffic runs directly over WDM. IP traffic has a shim header or digital wrapper placed around it at each input wavelength transmitter. In addition to containing information about the traffic, the information may also be used to switch the optical signals from one fiber to another in an optical switch.

INTERNETWORKING SS7 AND LEGACY TRANSPORT NETWORKS AND THE INTERNET

One of the common criticisms of the Internet that is voiced by some people is the absence of the many features that are common to telephony systems, such as high-quality voice calls, call forwarding, call screening, caller ID, and so forth. These features are quite important to many telephone users and are a vital part of the services that produce revenue for telephony service providers.

The initial Internet telephony products are called voice over IP (VoIP), and they are not intended to be a complete "telephone network." To be able to provide full-feature telephony services, a "transport Internet" must be able to avail itself to the SS7 technology, the lynchpin for telephony service features, and the foundation for the advanced intelligent network (AIN) services.

Reinvention or Use of the Telco Platform

If the telco platform is not used, the services that are part of SS7 must be "reinvented" by the Internet tasks forces—considered by many in the telephony industry to be a ridiculous alternative. Maybe so, maybe not: Several of the IP call processing protocols (such as SIP and Megaco) have begun to define service features, such as call waiting.

At any rate, the large SS7 vendors, such as Lucent, Nortel Networks, etc., are developing IP/SS7 gateways, and many products will be available in 2002.

The Integration of the Telephone and Internet Service Providers

In the United States, where deregulation legislation is taking effect, there is a lot of activity in the service providers' acquisitions and mergers. In addition, as voice-over data becomes more pervasive, the traditional data-only ISPs are increasing their interfaces into the traditional telephony architecture, such as the LEC equipment and SS7. Figure 15–18 shows the emerging architecture.

The customer still uses an LEC to connect to the ISP. The data traffic is exchanged through the ISP, the ISP's Internet and to the other user

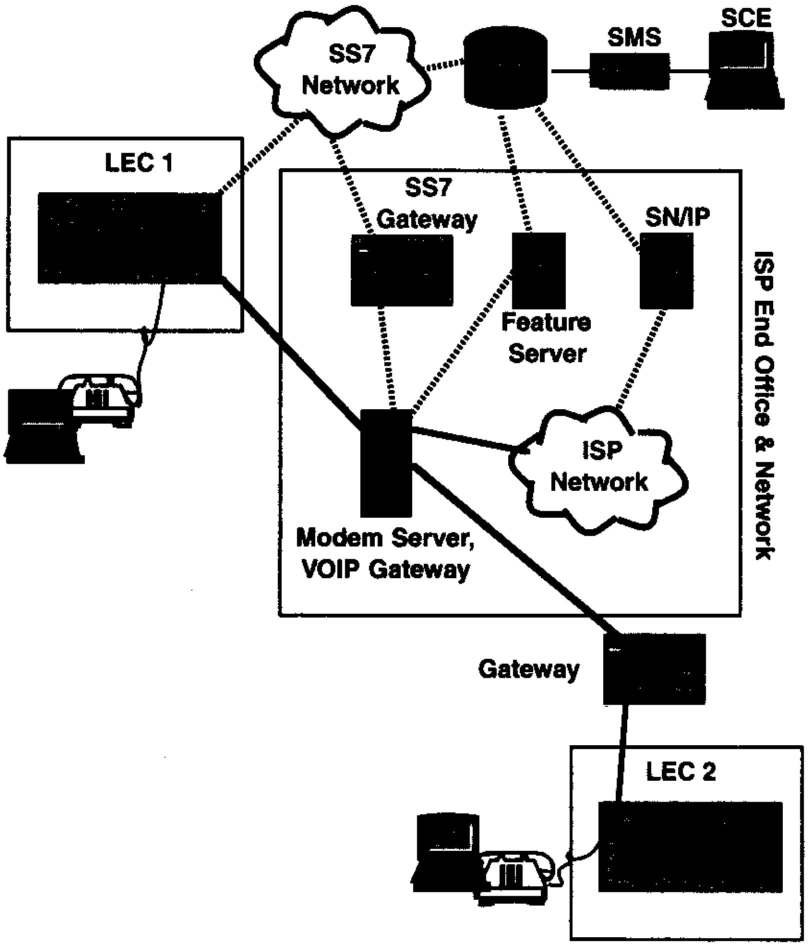

Figure 15–18 Integration of the telephone and ISPs.

(the "called" party). The ISP's network access server supports a modem pool and acts as a VoIP gateway.

In this example, the ISP is acting as an IXC by transporting the IP traffic over the ISP's network. The "ISP Network" in this figure also denotes the ISP's connection with the Internet.

Figure 15–18 also shows the interworking of the ISP with SS7 and the intelligent network (IN) components (shown at the top of the figure). The modem server connects to the IN components through an SS7 network or through a feature server, which, in turn, connects to the service control point (SCP). The SCP to feature server connection is authenticated with RADIUS. The job of the SCP is to allow the ISP to exploit the IN capabilities (billing, call screening, etc.), plus IP-specific features such as routing and billing by volume.

The IN service node and intelligent peripheral are also shown in this figure as part of the ISP architecture. This specific configuration will eventually be commonplace, but it is unusual at this stage of the evolution toward integrating voice and data networks. Moreover, the SN/IP operations may be part of the feature server node. The end offices are using the traditional circuit switches. Eventually, they will be replaced with packet switches.

THE INTERNET TRANSPORT NETWORK PROTOCOL STACK

Figure 15–19 shows a possible protocol stack for a transport network that is IP-based. By IP-based, I mean that the network uses IP and its companion protocols. Here is a brief description of the entites in this protocol stack:

- *Optical Layer:* This layer has been the focus of this book. With the use of 3G technologies, such as OTN architecture, ASON for dynamic services, etc. Obviously, this layer assumes a huge role in the 3G optical Internet.
- *Digital Wrapper:* The digital wrapper may be merged with the MPLS label.
- *MPLS:* This label switching technology will be used to implement constrained routes and other traffic engineering operations. For the foreseeable future, ATM and Frame Relay will provide services similar to the label switching services of MPLS. In addition, Diff-

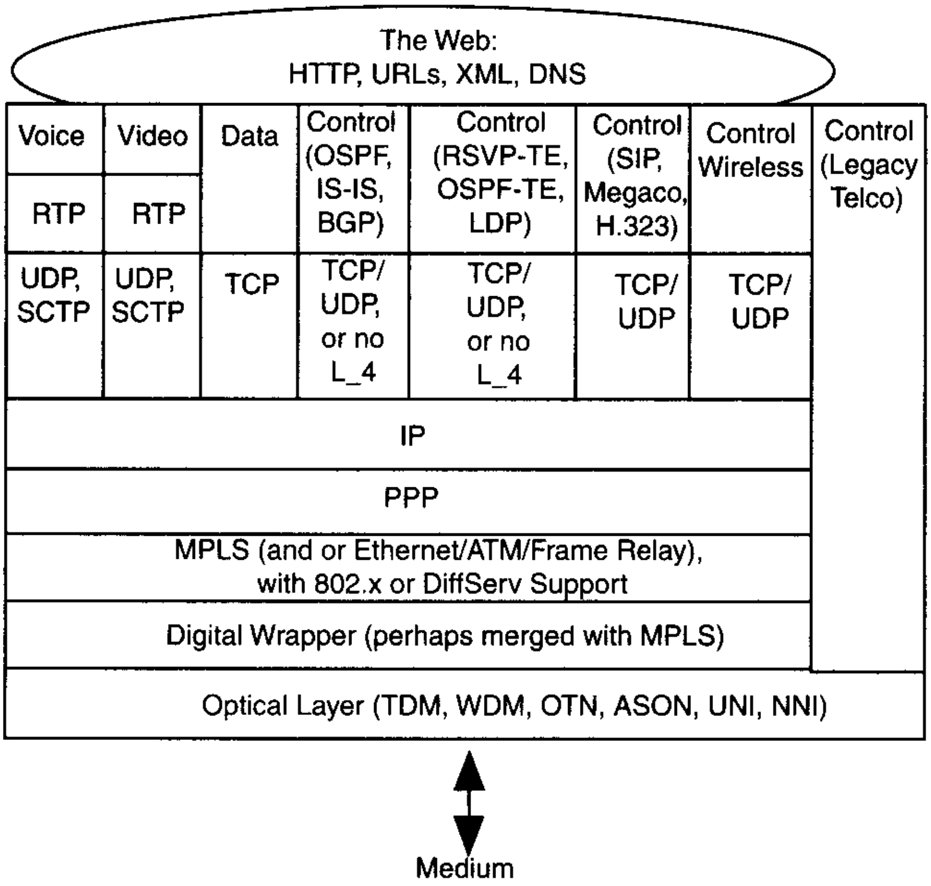

Figure 15–19 The Internet transport protocol stack.

Serv will eventually take over the traffic monitoring and policing functions of ATM and Frame Relay.

- *Note:* The role of Ethernet in PONs and metropolitan optical networks will be significant. Its presence in a wide area backbone transport is not a major issue at this time.
- *PPP:* This protocol with its many extensions (such as L2TP) will continue to play a key role in dial-up protocols, including interworking with some of the security aspects of the control planes.

- *IP:* This protocol forms the basis for addressing and packet forwarding. As Figure 15–19 suggests, all upper layer protocols (with the exception of the current legacy telco control plane) run on top of IP.
- *UDP/TCP/SCTP or No L_4:* Depending on the specific needs of an upper layer protocol, UDP, SCTP, or TCP may or may not be invoked. Figure 17–18 shows the most likely scenarios. (Note: OSPF does not operate with a L_4 protocol, but directly on top of IP.
- *Voice:* Packetized voice will operate over the Real Time Protocol (RTP).
- *Video:* Packetized video will operate over the RTP.
- *Data:* Most data applications will continue to use TCP.
- *Control (OSPF, IS-IS, BGP):* Conventional route advertising and route discovery will continue to use the IP-based routing protocols.
- *Control (RSVP-TE, OSPF-TE, LDP):* MPLS will use this control plane to advertise labels and addresses and to create bindings for an LSP, and a constrained route LSP.
- *Control (SIP, Megaco, H.323):* This emerging control plane will support VoIP and many service features associated with the telephone network.
- *Control (Wireless):* The network side of mobile, wireless networks will migrate to the use of the layer stack shown in Figure 15–19. But this migration will take quite some time. In the meantime, many of the control plane procedures will continue to be based on the legacy teleco control plane and of course the mobile-specific protocols, such as roaming, and registration.
- *Control (Legacy Telco):* This control plane will continue to dominate transport networks at the upper layers, pehaps indefinitely. There are those in the industry who would like to see the IP telephony control plane (SIP, Megaco, H.232, etc.) replace SS7 protocols such as ISUP. But don't hold your breath for this to occur any time soon.
- *The Web:* No, don't hold your breath, but all the IP-based upper layer protocols in Figure 15–19 are designed to interwork gracefully with the Web architecture, including the domain name system (DNS). This gives the IP-based protocols a decided advantage over the legacy telco technologies, which are not designed for Web interactions.

CONCLUSIONS

Well, that's it. Originally, my conclusions for this chapter and book were placed here, but one of my reviewers suggested that they be placed in the preface to this book. That's where they are, if you want to read them again.

I hope you have enjoyed reading this book and that you find it a useful addition to your library. I enjoyed writing it and I thank you for reading it.

Appendix A

The T1 Family

T1 LINE CONFIGURATIONS

Today, the majority of T1 offerings digitize the voice signal through a variety of analog-to-digital techniques. Whatever the encoding technique, once the analog images are translated to digital bit streams, the T1 system is able to time division multiplex (TDM) voice and data together in 24 user slots within each T1 frame.

Figure A–1 shows a T1 configuration. There is no typical configuration for these systems. They can range from a simple point-to-point topology shown here, wherein two T1 multiplexers operate on one link, or they can employ with digital cross-connect systems (DCS) that add, drop, and/or switch payload as necessary across multiple links.

Voice, data, and video images can use one digital "pipe." Data transmissions are terminated through a statistical time division multiplexer (STDM), which then uses the TDM to groom the traffic across the transmission line through a T1 channel service unit (CSU) or other equipment, such as a data service unit (DSU) or a combined DSU and CSU. The purpose of the CSU is to convert signals at the user device to signals acceptable to the digital line (and vice versa at the receiver). The CSU performs clocking and signal regeneration on the channels. It also performs functions such as line conditioning (equalization), which keeps the signal's performance consistent across the channel bandwidth, signal reshaping, which reconstitutes the binary pulse stream, and loop-back

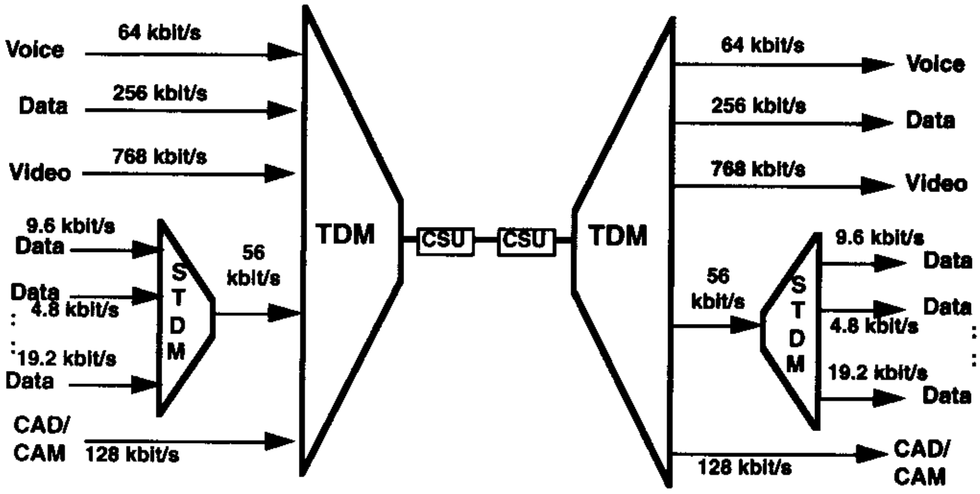

Where:

CAD/CAM	Computer assisted design/computer assisted manufacturing
CSU	Channel service unit
STDM	Statistical time division multiplexer
TDM	Time division multiplexer

Figure A–1 Possible topology for a digital carrier system.

testing, which entails the transmission of test signals between the DSU and the network carrier's equipment.

The bandwidth of a line can be divided into various T1 subrates. For example, a video system could utilize a 768 kbit/s band, the STDM, in turn, could multiplex various data rates up to a 56 kbit/s rate and perhaps a CAD/CAM operation could utilize 128 kbit/s of the bandwidth.

THE DIGITAL NETWORK

The U. S. T1-based digital network has been under development for over thirty years. During this time, a hierarchy of transmission levels (low speeds to high speeds) has been implemented through time division multiplexers, channel banks and digital cross-connects. These levels are designated by digital signal (DS) numbers ranging from DS0 to DS4.

DS1 Frame Format

Each basic channel operates at 64 kbit/s signal. This signal is called a digital signal level 0, or DS0. The 0 means that the signal is not multiplexed (digital signal, 0 level of multiplexing). The multiplexed 24 DS0 signals are collectively called DS1 (digital signal, first level of multiplexing). Let us see how DS1 is formed.

A few simple calculations are needed at this point in the discussion in order to understand the DS1 signal. After each of the 24 channels in a channel bank (terminal) has been sampled, quantized, and encoded, the resultant pulse train (bit stream) is called a *frame*. A frame has a time duration of 125 microseconds (μsec) (1 second/8000 samples = .000125). The bit duration is 648 nanoseconds (nsec): 125 μs/193 = 648 nsec. Further, each PAM sample is encoded into an eight-bit 5.184 nsec word: 648 nsec × 8 = 5.148 μsec.

The frame contains 24 eight-bit binary words, as depicted in Figure A–2. At the end of channel 24, an additional bit (the F bit) is appended to the frame. This bit becomes the 193rd bit of a frame and is used for framing (synchronization) and a variety of operations and maintenance services.

These calculations provide an insight into the DS1 bit rate. We just learned that the pulse code modulation (PCM) terminal produces 24 8-bit words, plus the F bit. The sampling rate of each channel in the system is 8000 times per second. Thus, 8000 x 193 bits per frame = 1,544,000 bits per second, or 1.544 Mbit/s, which is the DS1 line bit rate (see Table A–1). This DS1 signal is transmitted onto the T1 TDM cable facilities.

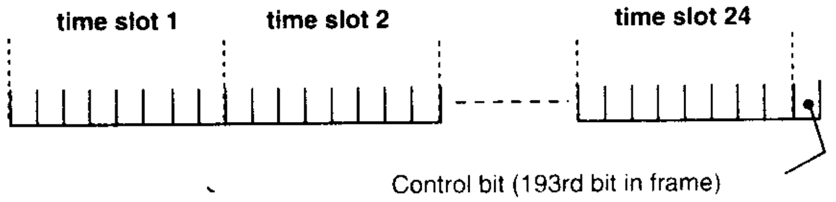

(a) 1.544 Mbit/s frame.

Figure A–2. The T1 frame and coding scheme.

Table A–1 Digital Signal at the First Level (DS1)

24	Channels or words
× 8	Bits per word
192	*Word bits* / frame
+ 1	*F* bit
193	Bits per complete frame
× 8000	Sampling rate/second
1,544,000	bit/s or 1.544 Mbit/s

NORTH AMERICAN ASYNCHRONOUS DIGITAL HIERARCHY

When the digital network was in its early stages of development, common clocks (such as a primary reference source) were not available. In order to synchronize the switches, terminals, and multiplexers to a common rate, bit stuffing was used to bring lower-rate signals into a common higher rate.

Five levels of multiplexing exist within the North American asynchronous digital hierarchy (see Figure A–3). Starting with the 1.544 Mbit/s DS1, each level of the hierarchy increases a facility's channel capacity. The current systems are capable of handling over 24,000 channels. There are several types of asynchronous digital multiplexers/demultiplexers that support this hierarchy, and they are described in this section.

A DS1 signal may be combined with another DS1 signal to produce a 3.152 Mbit/s signal containing 48 voice frequency (VF) channels. The multiplexer used for this operation is called an M1C mux, which means first level in, combined level out. This level is called a digital signal at the first level combined, or DS1C.

The combination of four DS1s to produce a 6.312 Mbit/s bit stream is called a digital signal at the second level, or DS2. It supports 96 voice channels. The multiplexers used for this operation are called M12, which means first level signals in, second level out.

Figure A–3 also shows the relationship of the digital signal cross-connect (DSX, also called a digital cross-connect) to the hierarchy. These components are equipment frames containing jack panels that serve as channel bank and multiplexer cross-connect interfaces in the telco office. The frames are named DSX-0, DSX-1, DSX-1C, DSX-3, and DSX-4 for each of the six DS rates (DSX-0 and DSX-4 are not shown in Figure A–3). Each frame connects equipment that operates at the respective DSn rate.

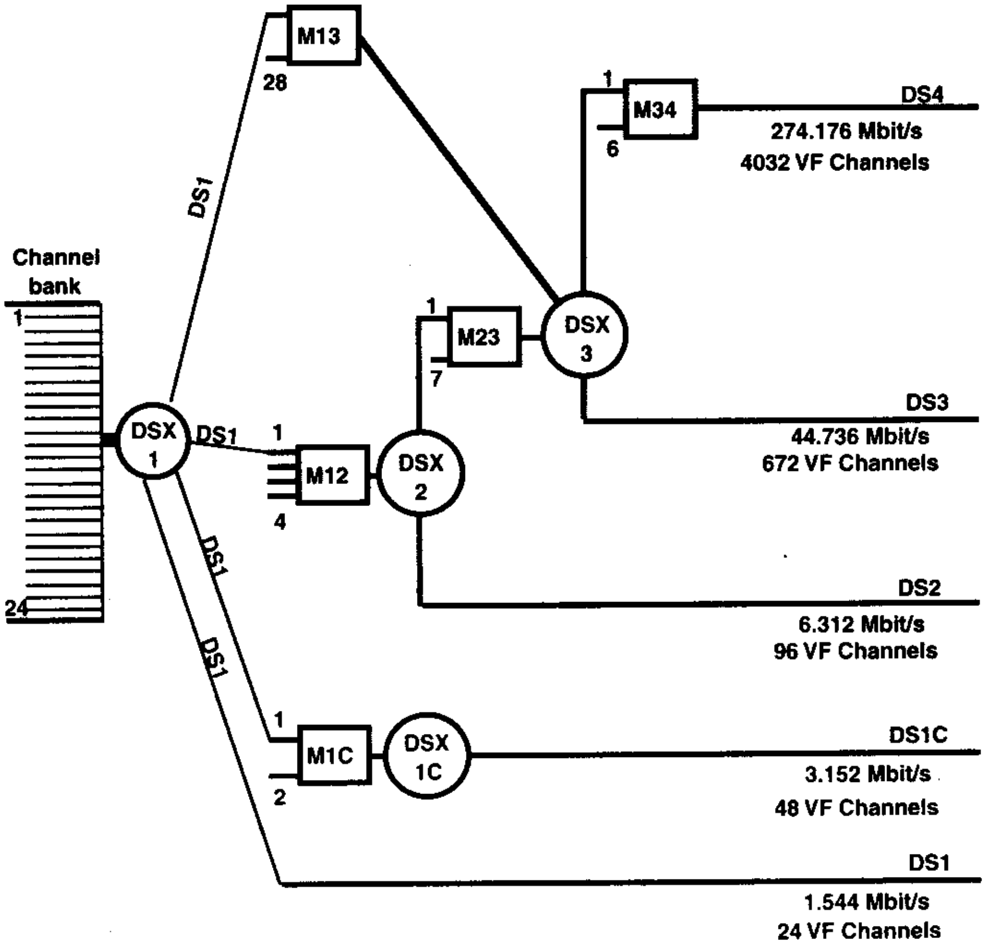

Figure A-3 The digital hierarchy.

The DSX-0 is employed for connecting and terminating digital data system (DDS) equipment.

The DSX-1, DSX-1C, and DSX-2 have several common features, such as monitoring and patching jacks, accommodations for conventional telephone plugs, order-wire terminations, and tracer lamps (to identify the two ends of a cross-connection). As shown in Figure A–3, the DSX-1 connects to a channel bank (and other equipment, explained later). The DSXs are 110-ohm balanced points. The DSX-3 and DSX-4 are similar to the DSX-1, DSX-1C, and DSX-2 except they are 75-ohm coaxial interfaces.

It is important to distinguish between a DSX and a DCS, which is also called a digital cross-connect. The DSC is considerably more intelligent. It is software controlled and uses time-slot interchange (TSI) to transfer slots between input and output lines. DCSs eliminate channel banks at interfaces where channels are transferred between carrier

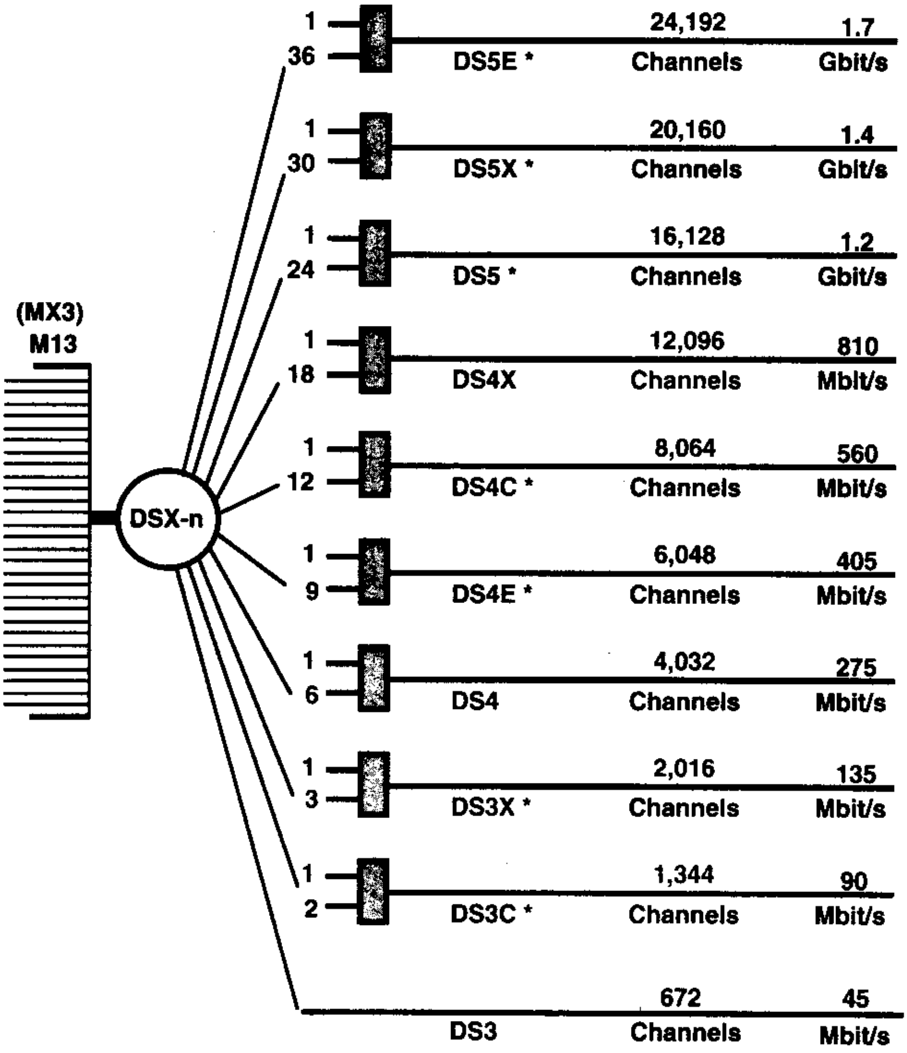

Note: * = not standardized

Figure A–4 North American Asynchronous Digital Hierarchy.

systems. DCS operations are explained in more detail later in this chapter and in Chapter 5.

One of the most common types of multiplexers is the MX3. It accepts up to 28 DS1 signals, 14 DS1C signals, or 7 DS2 signals as inputs, and creates a DS3 signal as its output. These multiplexers are called MX3 (with X designating the level in), meaning M13 for first level in and third level out, or M23 for second level in and third level out.

The final (formally defined) level within the North American asynchronous digital network is the DS4, which produces a 274.176 Mbit/s signal with 4,032 voice channels. An M34 mux is used for producing this digital level.

There are also several other common (but not standard) levels (shown in Figure A–4). These are: (a) the 90 Mbit/s DS3C, (b) the 135 Mbit/s DS3X, (c) the 405 Mbit/s DS4E, (d) the 560 Mbit/s DS4C, and (e) the 1.2, 1.4, and 1.7 Gbit/s systems. These systems accept DS3 as their inputs. The DS3C system accepts two DS3 inputs, the DS3X system accepts three DS3 inputs, and the DS4E system accepts nine DS3 inputs. The "E" indicates "extended." The signal is not high enough to be called a DS4C, which is a 560 Mbit/s system transmits 12 DS3 signals. The reader should check the vendors' offerings for equipment operating above DS4, due to the lack of standards at these levels.

SUBSCRIBER-TYPE SYSTEMS

Subscriber-type PCM systems are available that use the same quantizing and encoding processes as the D2, D3, and D4 systems explained previously (see Figure A–5). These systems are also software programmable for voice and data circuits just as the D4 and D5 channel banks are. They are capable of SF/ESF, AMI/B8ZS, and ADPCM (LBRV) operations. The main difference is that one terminal is located in the central office while the other is in the field near or on the customer's location. They may be referred also to as a *pair gain system*, a *digital loop carrier*, or a *subscriber loop carrier*. Some of them can also extend a leased DS1 and/or DS3 to the customer's premise for his own use.

Subscriber-type systems support a wide variety of applications by various operating companies. One of the more popular uses is providing service to developing areas for new subdivisions where an existing cable plant is insufficient. A system can provide the service immediately and permanently, or it can be moved to another location (if growth in the area eventually justifies a central office). Regardless of whether the service is

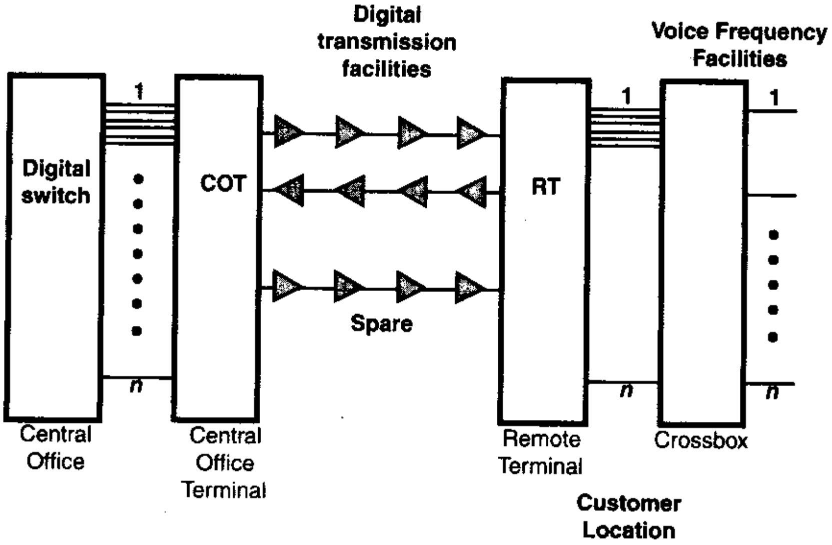

Figure A–5 Basic subscriber system arrangement.

permanent or temporary, a subscriber system is easy to engineer and install on short notice. An example is a new industrial park experiencing sudden and unexpected growth, resulting in demands for service exceeding the available loop plant. The system can be installed and operating within a few weeks. Also, many companies use these systems to provide for temporary service to large functions such as business conventions or sporting events.

There are other reasons to justify the placement of a subscriber loop carrier in the loop plant. First, the copper pairs serving the subscribers will be much shorter, thus overcoming distance limitations in providing the newer services. Second, shortening the customer loop decreases the exposure to power-line interference with its resultant degradation and noise impact on these circuits. Flexibility is further enhanced by the 1200-ohm loop capability of the remote terminal. Third, electronics allow the future ability to provide new services quickly. The distance from the central office to the remote terminal is limited only by the copper DS1 span line performance. Today, most of these systems employ fiber optics, so there is very little distance limitation. Subscriber loop carriers provide applications for videoconferencing and local area networks (LANs) links.

Subscriber carrier systems provide several functions which channel banks do not because they are in the local loop. These are as follows: ringing, coin collection, party lines, remote terminals, and subscriber line testing as well as batteries for backup power and fan units when not installed in a controlled environment.

References

[AGRA92] Agrawal, G. P., *Fiber-Optic Communication Systems*, John Wiley, New York, 1992.

[ASHW01] Ashwood-Smith, Peter, et al., "Generalized MPLS: Signaling Functional Description," draft-ietf-mpls-generalized-signaling-05.txt, July, 2001.

[AWDU99] Awduche, W., et al., "Requirements for Traffic Engineering Over MPLS," RFC 2702, September, 1999.

[AWDU01] Awduche, Daniel O., et al., "Multiprotocol Lambda Switching: Combining MPLS Traffic Engineering Control with Optical Cross-connects," draft-awduche-mpls-te-optical-03.txt, April, 2001.

[BELL99] Bell Labs, "Bell Labs Trends and Developments," Volume 3, Number 1, Spring, 1999.

[BELL01] Bellato, Alberto, et al., "Optical Transport Networks GMPLS Control Framework," draft-bellato-ccamp-g709-framework-00.txt, June, 2001.

[BERN01] Bernstein, L. et al, Optical InterDomain Routing Considerations, draft-bernstein-optical-bgp-01.txt, July, 2001.

[BISH01] Bishop, D. J., Giles, C. R., Das, S., "The Rise of Optical Switching," *Scientific American,* January, 2001.

[BLAC91] Black, Uyless, *OSI: A Model for Computer Communications Standards,* Prentice Hall, 1991.

[BLAC01] Black, Uyless, and Waters, Sharleen, *SONET and T1: Architectures for Digital Transport Networks,* Second Edition, Prentice Hall, 2001.

[BLAC02] Black, Uyless, *MPLS and Label Switching Networks,* Second Edition, Prentice Hall, 2002.

[BLUM01] Blumenthal, Daniel J., "Routing Packets with Light," *Scientific American,* January, 2001.

[BROR01] Brorson, S., et al., "Link Management Protocol (LMP) for DWDM Optical Line Systems," IETF draft-fredette-lmp-wdm-02.txt, July, 2001.

[CHIU01] Chiu, Angela, et al., "Impairments and Other Constraints on Optical Layer Routing," draft-ietf-ipo-impairments-00.txt, May, 2001.

[CHRA99] Chraplyvy, Andrew R., "High Capacity Lightwave Transmission Experiments," *Bell Labs Technical Journal,* January-March, 1999.

[DUTT98] Dutton, Harry, Jr., *Understanding Optical Communications,* Prentice Hall, 1998.

[FERN00] Fernandez, M., and Kruglic, E., "MEMS Technology Users in a New Age in Optical Switching," *Lightwave,* August, 2000.

[FREE00] Freeland, Daren, et al., "Considerations on the Development of an Optical Control Plane," draft-freeland-octrl-cons-01.txt, November, 2000.

[FULL01] Fuller, Meghan, "Startup's Switchless Switch Targets Network Edge," *Lightwave,* October, 2001.

[G.70901] "Interface for the Optical Transport Network," ITU-T Recommendation G.709, February, 2001.

[G.87200] "Architecture of Optical Transport Networks," ITU-T Recommendation G.872, 2000.

[GILE99] Giles, C. R., and Spector, Magaly, "The Wavelength Add/Drop Multiplexor for Lightwave Communication Networks," *Bell Labs Technical Journal,* January-March, 1999.

[JACK99] Jackman, Neil A., et al., "Optical Cross-connects for Optical Networking," *Bell Labs Technical Journal*, January-March, 1999.

[JOHN99] Johnson, Steven R., and Nichols, Virginia L., "Advanced Optical Networking—Lucent's MONET Network Elements," *Bell Labs Technical Journal,* January-March, 1999.

[KAMI97] Kaminow, I. P., and Koch, T. L. (Eds.), *Optical Fiber Telecommunications IIIA,* Academic Press, 1997.

[KOLE00] Kolesar, Paul, and Buck, Preston, "The 62.5 vs. 50 Micron Debate," *Lightwave,* April, 2000.

[LANG01] Lang, Jonathan, P., et al., "Link Management Protocol (LMP)," IETF draft-lang-mpls-lmp-02.text, September, 2001.

[LEE01] Lee, Young, et al., Protection Scheme for Optical Channel Concatention, draft-ylee-protection-occ-00.txt, June 2001.

[LIGH01a] "Optical Units Reference," Available from: www.lightreading.com.

[MANN01] Mannie, Eric, "GMPLS Extensions for SONET and SDH Control," draft-ietf-ccamp-gmple-sdh-os.txt, June, 2001.

[NORT98] Nortel Networks, "S/SMDS TransportNode: Best-in-Class Portfolio of High-Capacity Solutions," Document No. 56088.16/11-98, Issue 2, November 10, 1998.

[NORT99a] Nortel Networks, "Positioning OC-12 TBM Networks for the 21st Century," Document No. 56107.16/02-99, Issue 1, February 25, 1999.

[NORT99b] Nortel Networks, "Highways of Light," *Nortel Networks Telesis,* Volume 101, 1999.

[ODSI01] www.odsi-coalition.com provides the specifications and source code for ODSI's work on optical bandwidth on demand.

[OIF01] www.oiforum.com provides information on the OIF's work on an optical network user-network interface (UNI).

[PAPA00] Papadimitriou, D., et al., "Optical Network-to-Network Interface Framework and Requirements," draft-papadimitriou-onni-frame-01.txt, November, 2000.

[PODZ00] Podzimek, Cheri, and Harmon, Dan, "Expanding the Network Diameter," *Lightwave,* May, 2000.

[RAJA02] Rajagopalan, Bala, et al., "IP over Optical Internets: A Framework," draft-ietf-ipo-framework-00.txt, no create date provided, expiration date is 13/2002.

[RAMA98] Ramaswami, Rajiv, and Sivarajan, Kumar N., *Optical Networks: A Practical Perspective*, Morgan Kaufman Publishers, San Francisco, 1998.

[REFI99] Refi, James J., "Optical Fibers for Optical Networking," *Bell Labs Technical Journal,* January-March, 1999.

[RYAN01] Ryan, Jim, "ITU G.655 Adopts Higher Dispersion for DWDM," *Lightwave,* September, 2001.

[STRI01] Strix, Gary, "The Triumph of the Light," *Scientific American,* January, 2001.

[T1X1.501] ANSI T1X1.05, "Need for Concatenating Optical Channels to Create Transparent High Bandwidth Channels," 2001.

[XUE01] Xue, Yong, "Carrier Optical Services Framework and Associated UNI Requirements," draft-many-carrier-framework-uni-00.txt, May, 2001.

[XU01] Xu, Yangguang, et al., "Intra-Domain GMPLS Control Plane Architecture for Automatically Switched Transport Network," draft-xu-ccamp-gmpls-arch-intra-domain-00.txt, September, 2001.

Acronyms

A/D
A-NZDF
AAL: ATM adaptation layer
AAL5: ATM adaptation layer type 5
ADM: add-drop multiplexer
AIN: advanced intelligent network
AIS: alarm indication signal
ANSI: American National Standards Institute
AOLS: all-optical label swapping
APS: automatic protection switching
ASE: amplifier spontaneous emission
ASIC: application-specific integrated circuit
ASON: automatic switched optical network
ATM: asynchronous transfer mode
AU: administrative unit
AUG: administrative unit group
BER: bit error rate
BGP: border gateway protocol
BIP: bit interleaved parity
BITS: building integrated timing supply
BLSR: bi-directional line-switched ring
BOD: bandwidth of demand
BPSR: bi-directional path switched ring
BPV: bipolar violation
BSNT: bit stream without octet timing
BSOT: bit stream with octet timing
C: container
CAD/CAM: computer assisted design/computer assisted manufacturing
CATV: cable TV
CBR: constant bit rate
CCid: control channel id
CH: cell header
CMIP: common management information protocol
CNE: client network element
CNLS: connectionless network
CO: central office (also: connection-oriented)
CP: common part
CPCS: common part convergence sublayer
CPE: customer premises equipment
CPI: common part indication
CPU: central processing unit
CR-LDP: constraint-based LDP
CRC: cyclic redundancy check
CS: cross-connect
CSU: channel service unit

DARPA: Defense Advanced Research Projects Agency
DCC: data communications channel
DCS: digital cross-connect system (also: digital cross-connect)
DDS: digital data system
DLC: digital loop carrier
DNS: domain name system
DQDB: dual queue dual bus
DRAM: dynamic random-access memory
DS: digital signal
DS0: digital signal level 0
DSF: dispersion-shifted fiber
DSU: data service unit
DSX: digital signal cross-connect
DWDM: dense wave division multiplexing
DWHdr: digital wrapper and header
E/O: electrical-to-optical converter
EDF: erbium-doped fiber
EDFA: erbium-doped fiber amplifier
EMS: element management system
ESI: external synchronization interface
ETSI: European Telecommunications Standard Institute
FCS: frame check sequence
FDDI: fiber distributed data interface
FDM: frequency division multiplexing
FEC: forward error correction (also: forwarding equivalence class)
FERF: far end receive failure
FF: fractional frequency
FS: feature server
FSC: fiber-switch capable
FTTC: fiber to the curb
FTTH: fiber to the home
G-PID: generalized PID
GCC: general communication channel
GFP: generic framing procedure
GLR: generalized label request
GLSP: generalized label switched path
GMPLS: generalized MPLS
HDLC: high-level data link control
HEC: header error control
HO-POH: higher order path overhead
I: information
IaDi: intra-domain
ID: identifier
IEC: International Electrotechnical Commission
IETF: Internet Engineering Task Force
IGP: Internal gateway protocol
IN: intelligent network
IP: Internet protocol
IPX: Internet Protocol X (vender-specific IP)
IS-IS: Intermediate system-to-intermediate system
ISDN: Integrated services digital network
ISO: International Standards Organization
ISP: Internet service provider
ISUP: ISDN user part
ITU: International Telecommunication Union
ITU-T: International Telecommunication Union—Telecommunication Standardization Sector
L2TP: layer 2 tunneling protocol
LAN: local area network
LASER: light amplification by stimulated emission of radiation
LDP: label distribution protocol
LEC: local exchange carrier
LI: length indicator
LLC: logical link control
LMP: link management protocol
LO-POH: lower order path overhead
LOF: loss of frame
LOH: line overhead
LOL: loss of light
LOP: loss of pointer
LOS: loss of signal
LSC: lambda-switch capable
LSP: label switched path (also: label switching path)
LSR: label switching router
LTE: line terminating equipment
MEMS: microelectromechanical systems
MONET: multi-wavelength optical networking consortium
MOR: multi-wavelength optical repeater

MPλS: multiprotocol lambda switching
MPLS: multiprotocol label switching
MPLS TE: MPLS traffic engineering
MSOH: multiplex section overhead
MTU: maximum transmission unit size
MUX: multiplexer
NAS: network access server
NCC: number of contiguous components
NDP: neighbor discovery protocol
NHLFE: next hop label forwarding entry
NLPID: network level protocol id
NM: nanometers
NNI: network node interface (also: network-to-network interface)
non-BOD: non-bandwidth of demand
NVC: number of virtual components
NZDF: non-zero dispersion fiber
OADM: optical ADM
OAM: operations, administration, and maintenance
OC: optical carrier (also: optical channel)
OCC: optical channel carrier
OCG: optical carrier group
Och: optical channel
ODSI: open domain service interconnect coalition
ODSI: optical domain service interconnect
ODU: optical channel data unit
ODUk: optical channel data unit
OFA: optical fiber amplifier
OH: overhead
OIF: optical interworking forum
OIF IrDi: inter-domain
OL: optical layer
OLN: optical line terminal
OLS: optical link system (also: optical line system)
OMS: optical multiplex section
OMU: optical mux unit
ON: optical node
ONC: optical network controller
ONE: optical network element
ONU: optical network unit
OOS: overhead signal
OPS: optical physical section
OPUk: optical channel payload unit
OSC: optical supervisory channel
OSI: open systems interconnection
OSNR: optical signal-to-noise ratio
OSP: optical switching path (also: optical switched path)
OSPF: open shortest path first
OTM: optical transport module
OTN: optical transport network
OTS: optical transmission section
OTU: optical transport unit
OTUk: optical channel transport unit
OUI: organizationally unique id
OXC: optical cross-connect (also: optical/electrical cross-connect)
PAD: padding (also: padding field)
PBX: private branch exchange
PCM: pulse code modulation
PDH: pleischronous digital hierarchy
PDU: protocol data unit
PID: protocol id
PLI: PDU length indicator
PMD: polarization-mode dispersion
PNNI: private network-to-network interface
POH: path overhead bit
PON: passive optical network
POP: point of presence
PPP: point-to-point protocol
PRS: primary reference source
PT: payload type
PTE: path terminating equipment
PTT: Postal, Telephone, and Telegraph Ministries
PXC: photonic cross-connect (also: optical/optical cross-connect)
QOS: quality of service
RAI: remote alarm indication
RCC: requested contiguous concatenation
RDT: remote digital terminal
RESV: reservation
RFC: request for comments

RS: regenerator section
RSOH: regenerator section overhead
RSVP: resource reservation protocol
RSVP-TE: RSVP traffic engineering
RTP: real time protocol
SA: service adapter
SAR: segmentation and reassembly sublayer
SCE: service creation environment
SCP: service control point
SCTP: stream control transmission protocol
SD: signal degrade
SDH: synchronous digital hierarchy
SDU: service data unit
SIP: session initiation protocol
SLA: service level agreement
SMDS: switched multimegabit data service
SMS: serviced management system
SN: service node
SNA: systems network architecture
SNAP: subnetwork access protocol
SNMP: simple network management protocol
SOH: section overhead
SONET/SDH: synchronous optical network/synchronous digital hierarchy
SOP: state of polarization
SPE: synchronous payload envelope
SRLG: shared risk link group identifier
SS: signaling system
SS7: signaling system number 7
SSM: synchronization status message
ST: signal type
STDM: statistical time division multiplexer
STE: section terminating equipment
STM: synchronous transport module
STP: signaling transfer point
STS: synchronous transport signal
SVC: switched virtual circuit or call
T: transparency
TDM: time division multiplexing (also: time-division multiplexer; time-division multiplex capable)
TE: traffic engineer
TIA: Telecommunications Industry Association
TL: transaction language
TLV: type-length-value
TRDP: topology and resource distribution protocol
TSI: time-slot interchange
TU: tributary unit
TUG: TU group
UED: user edge device
UI: unit interval (also: unnumbered information)
UNI: user-network interface
UNI-SR: UNI sub-rate
USF: unshifted fiber (also: dispersion-unshifted fiber)
UU: user-to-user information
VC: virtual controller (also: virtual container)
VF: voice frequency
VoIP: voice over IP
VPN: virtual private network
VT: virtual tributary
VTG: virtual tributary group
WADM: wavelength ADM
WAN: wide area network
WDM: wave division multiplexing (also: wavelength division multiplexing)
WECO: Western Electric Co.
WFIB: wavelength forwarding information base

Index

Page numbers ending in "f" refer to figures. Page numbers ending in "t" refer to tables.