Routing Architecture

1 Abstract

Digital developed the intermediate system-to-intermediate system (IS-IS) intradomain routing information exchange protocol for the DECnet Phase V network layer architecture. This protocol, which has been adopted by the International Organization for Standardization, is based on a link state routing algorithm. The benefits derived from the IS-IS protocol include a self-stablizing method for reliable link state packet distribution, a hierarchical network structure to support larger networks, protocols for efficiently utilizing local area networks, and simultaneous support for multiple network layer protocols.

The network layer architecture has three basic components. The first concerns the transmission of data packets from one end system (a host) to a remote end system, regardless of whether or not these packets are sent by way of routers. The main features of this component are packet formats and addressing. Standards for these features are defined in the connectionless network layer protocol (CLNP), adopted by the International Organization for Standardization (ISO), and in the internet protocol (IP), the equivalent standard in the transmission control protocol/internet protocol (TCP/IP) suite.[1,2]

The second component relates to handshaking between neighbors (i.e., directly connected systems) and mapping network layer addresses to data link layer addresses. The ISO protocol that performs this function is the end system-to-intermediate system (ES-IS) protocol.[3] The address resolution and internet control message protocols provide most of the same functionality in the TCP/IP protocol suite.[4,5]

The third component of the network layer architecture pertains to routing. The routing protocol developed for Digital's DECnet Phase V network architecture and adopted by the ISO is the intermediate system-to-intermediate system (IS-IS) intradomain routing information exchange protocol.[6]

The architecture for DECnet Phase V allows support of many network layer protocols, i.e., CLNP, IP, Novell NetWare, and AppleTalk.[7] Each network layer suite has its own protocols for the first two components of the network layer architecture. DECnet Phase V support for a particular network layer suite implies support for such protocols. Consequently, end systems that implement an existing network layer protocol need not be modified to operate with DECnet Phase V routers (i.e., intermediate systems). This paper briefly discusses data packet formats, types of routing control packets, and neighbor handshaking protocols and then focuses on the third

component of the network layer architecture, concentrating on the IS-IS routing protocol.

Support for any network protocol suite can be added easily to the IS-IS routing protocol. DECnet Phase V routing products currently support the DECnet Phase IV, CLNP/DECnet Phase V, and the IP protocols. Support for the Novell NetWare, XNS, and AppleTalk protocols is under investigation.

2 Data Packet Formats

A network layer data packet carries data, usually generated by higher-layer protocols, between host systems. The purpose of the network routing layer is to correctly deliver data packets to their destinations. To accomplish this task, additional pieces of information are required; these are carried in the header of the data packet. The most important function of the header is addressing. Each data packet must uniquely identify the source and destination addresses for the packet. Other important functions include: checksumming, to ensure that transmission errors are detected; fragmentation and reassembly, to allow the transmission of large packets over links that can support only smaller packets; error reporting, to notify someone should an error occur; security, to identify special security requirements of packets; quality of service maintenance, to ensure that the correct level of service is provided; and congestion notification, to notify the source and destination should congestion occur along the path of a data packet.

The DECnet Phase IV architecture uses a proprietary packet format for data exchange. The DECnet Phase V architecture continues to support this format to allow compatibility with existing Phase IV systems. However, DECnet Phase V uses the ISO CLNP standard for communication between DECnet and open systems interconnection (OSI) systems. Use of this standard protocol also permits DECnet Phase V systems to communicate with other vendors' end systems that implement the ISO standard. In addition, communication using IP is possible with systems that implement the TCP/IP suite.

DECnet Phase IV employs a 16-bit network layer addressing scheme. When using the CLNP, the addresses, known as network service access point (NSAP) addresses, vary in length up to 20 octets. Defining a common mapping procedure allows a DECnet Phase IV address to be expressed as an equivalent ISO NSAP address. Similarly, an ISO NSAP address thus derived, and therefore Phase IV compatible, may be converted back to the original Phase IV address. Converting the source and destination addresses and the packet formats enables any DECnet Phase IV packet to be translated into a CLNP packet and back again. Therefore, two DECnet Phase IV systems can communicate over a portion of a network that supports only the CLNP. Similarly, two DECnet/OSI (or even pure OSI) systems can communicate over a portion of the network that supports DECnet Phase IV, provided that the addresses chosen are Phase IV compatible.

3 Overview of Routing Control Packets

The IS-IS protocol uses three basic types of packets:

1. Hello Packet. The protocol uses Hello packets to keep track of neighbors. Routers determine the identity of neighbors and periodically check the status of the link to that neighbor by exchanging Hello packets.

2. Link State Packet. Link State Packets (LSPs) list, for each neighbor of the node issuing the LSP, the ID of that neighbor and the cost of the link to it. This list includes both router neighbors and end-system neighbors. The cost of the link is assigned by the network manager to reflect the desirability of using that link. A number of factors determine the cost, including throughput capacity and the monetary cost associated with using the link.

3. Sequence Number Packet. Sequence Number Packets (SNPs) are used to ensure that neighboring routers have the same notion of what is the most recent LSP from every other router. There are two types of SNPs: the Complete Sequence Number Packet (CSNP) and the Partial Sequence Number Packet (PSNP).

   The CSNP lists all LSPs present in the issuing router's LSP database, together with their sequence numbers, and is used to synchronize LSP databases. The CSNP is transmitted upon link start-up on point-to-point links and periodically on a local area network (LAN). This use of the CSNP to ensure LSP database consistency of all routers on the LAN is described in more detail in the section Efficient Use of LANs.

   The PSNP lists only a few LSPs and is used to explicitly acknowledge or request one or more LSPs.

4 Neighbor Handshaking Protocols

The architecture for DECnet Phase V uses the ES-IS protocol to enable routers and end systems on a LAN to learn about each other's presence. Every end system periodically multicasts an End System Hello protocol data unit (PDU) to the multicast address "All Intermediate Systems." This PDU contains the end system's NSAP address and permits the receiving routers to create an entry that maps the NSAP address to the corresponding data link address from which the PDU was received. The routers use this information to deliver data PDUs to the end systems and also to communicate the existence of the end systems to other routers by means of the routing protocols.

In a similar manner, all routers periodically multicast an Intermediate

System Hello to the multicast address "All End Systems." This data permits the end systems to determine the data link addresses of all routers on the LAN. In the absence of other information, an end system will transmit any data PDUs destined for another system to one of the routers it has discovered. However, the router to which the data PDU is sent may not be the best path. Indeed, direct transmission of the data PDU to the

destination system may be possible, if the source and destination systems are on the same LAN. In such cases, the router concerned sends a Redirect PDU back to the source end system. The Redirect contains the data link address to use for this NSAP address, which the end system can then use for subsequent transmissions.

The ES-IS protocol replaces the proprietary DECnet Phase IV initialization protocol for use between the DECnet and OSI systems. However, operation of the DECnet Phase IV protocol is still necessary to enable handshaking between DECnet Phase IV and DECnet Phase V systems. To avoid confusion, the Phase IV initialization messages transmitted by Phase V systems have a version number that is acceptable to only Phase IV systems. Such messages are ignored by other Phase V systems.

5 Routing Protocols, with Emphasis on the IS-IS Protocol

Routing protocols are used to calculate the path, i.e., the route, that a data packet will take through a network. Typically, a routing protocol dynamically adjusts to network problems, such as failed links or routers, to ensure that the network continues to operate in a robust manner. Use of dynamic routing protocols also eases installation and configuration, because routes are calculated by means of the algorithm, not the user.

The two main types of dynamic routing protocols are distance vector and link state. Many routing protocols are based on distance vector routing, for example, DECnet Phase III, DECnet Phase IV, and the routing information protocol (RIP).[8] In a distance vector protocol, each router is responsible for keeping track of and informing its neighbors about its distance (i.e., total cost) to each destination. The router computes its distance to each destination based on its neighbors' distances to each destination. The only information a router has to know a priori is its own ID and the cost of its links to each neighbor.

Consider the distance vector routing example shown in Figure 1. Suppose a router R with five ports is configured with costs c(1), c(2), c(3), c(4), and c(5) for each of the ports, respectively. Further suppose that the neighbor on port 1 informs R that it is d(1) from some destination D, the neighbor on port 2 informs R that it is d(2) from D, and so forth. R can then figure out its own distance to destination D. If the destination is R itself, then R's distance to D is 0. Otherwise, R's distance to D is the minimum value of c(i) + d(i), for i = 1 through 5. If R receives a packet addressed to destination D, R should forward the packet through the port with minimum total cost to D.

Because of their slower convergence rate, distance vector protocols generally provide lower performance than link state protocols. Distance vector protocols adapt to changes in topology less quickly than link state

protocols, and until the protocol adapts to such a change, routing can be disrupted. The main reason for this convergence problem stems from incorrect information. When changes such as link failures occur in the network, the information that each node transmits to its neighbors is only

that node's current impression of the distance to each destination, which may be incorrect information. Consequently, the distance vector algorithm may take several iterations to converge to the correct routes.

The first deployed link state routing protocol was developed by Bolt Beranek and Newman (BBN) for the Advanced Research Projects Agency Network (ARPANET).[9,10]. In link state routing, each router determines its local status and then constructs an LSP, defined earlier in the section Overview of Routing Control Packets. This LSP is transmitted (or "flooded") to all the other routers, which are responsible for storing the most recently generated LSP from each router.[11] (If the large size of the network makes it impractical for the LSP database to contain information for every other router, the network can be made hierarchical, as described in the Hierarchy section.) All routers (or all routers in an area, when hierarchical routing is used) then compute routes based on a complete topology. Figure 2 illustrates an example of link state routing, with a router R determining the state of its neighbors and then broadcasting this information by means of Hello Neighbor messages.

Link state algorithms respond rapidly and consistently to changes in networks, as compared with distance vector algorithms. Once the LSPs have been distributed, each router can calculate routes without further reference to the other routers. The results are more stable routing and lower consumption of link bandwidth and router CPU. Therefore, the design of the IS-IS routing algorithm was based on the original BBN link state routing algorithm, which used an algorithm known as the shortest path first (SPF) to calculate the routes.[12]

The IS-IS protocol corrected many deficiencies and added extra functionality.

1. The IS-IS protocol provides a more stable method for reliably distributing LSPs. The ARPANET method was an early algorithm that used excessive overhead and was unstable in rare circumstances. The IS-IS protocol design uses a self-stabilizing protocol for LSP distribution that requires much less bandwidth.

2. The IS-IS protocol can be used in a hierarchical manner to support larger networks.

3. The ARPANET method assumed all connections were point-to-point links. Many nodes can be connected with a LAN. Modeling a LAN as a fully connected set of nodes attached with point-to-point links would be extremely inefficient. The IS-IS routing protocol incorporates protocols for efficiently utilizing LANs.

4. Given that a router has limited memory, the network can grow beyond a

size that the router can support. If the router failed simply because its LSP database overflowed the available space, network management could not be used to reconfigure the router. If the router continued to operate and based the routing on an incomplete database, loops might

form and adversely affect routes that traverse that router. The IS-IS protocol has mechanisms that enable overloaded routers to remain reachable for network management.

5. Certain control packets can get very large. The IS-IS protocol has mechanisms for ensuring that fragments of a control packet can be dealt with independently rather than required to be fully reassembled first.

6. The IS-IS routing protocol can support many network layer protocols simultaneously. This support is known as Integrated IS-IS.[13]

6 Hierarchy

As a network grows, several factors may overload the routing protocol: the LSP database may become too large to fit into memory; computing routes may require too much CPU; the task of keeping the LSP databases up-to-date may consume too much bandwidth; or the network may be unstable because link changes are frequent. To deal with these factors, the IS-IS protocol allows the network to be partitioned into areas. Within an area, the level 1 routers keep track of all the nodes and links. Level 2 routers keep track of the location of the areas but are not concerned with the detail inside the areas. A level 2 router can also act as a level 1 router in one area.

To use the IS-IS protocol in a hierarchical way, it is convenient for the network layer addresses to be topologically hierarchical. Figure 3 illustrates the structure of an IS-IS address. All nodes in a particular area have the same value for the area address field of their address. A level 1 router looks at the area address portion of the destination address in a packet. If this field matches the router's area, the router assigns the packet a path based on the ID portion of the address. Otherwise, the router routes the packet toward a level 2 router, which directs the packet to the correct area.

The IS-IS protocol treats the last octet of the address as a selector, which is used only for demultiplexing multiple network users within the destination system. The selector field can therefore be ignored with respect to IS-IS routing.

In general, the area address itself is hierarchically subdivided. This structure is useful for address administration and for routing between routing domains, for example, different corporations, which may be interconnected by means of a public network. However, from the point of view of IS-IS operation, the entire area address is a single identifier for the area.

In a network of global dimensions, possibly comprising millions of addresses, the ability to use hierarchical addressing is essential to help

provide some of the topological information. This addressing scheme is analogous to the use of country codes in international telephone numbers, which allows calls to be routed to other countries without complete knowledge of the internal structure of all the telephone systems in the world.

7 Efficient Use of LANs

All routers connected to a LAN are neighbors. If the routing protocol was simply to consider all pairs of nodes on the LAN as neighbors, then each router on the LAN would issue an LSP listing every node on the LAN. In addition, the LSP distribution would be inefficient if each router had to transmit every LSP to all other routers on the LAN and then receive acknowledgments from all these same routers.

The IS-IS protocol dramatically reduces the required size of the LSP database by considering the LAN as a pseudonode. Each router then claims to have one link to the pseudonode, rather than a link to every other router on the LAN. Only the pseudonode claims to have links to all the end systems on the LAN.

This approach requires that an LSP be transmitted for the pseudonode itself, and thus some router on the LAN has to take on the responsibility for transmitting the packet for the pseudonode. The router with the numerically highest priority (or, in the event of a tie, the highest data link address) is elected the designated router (DR). The DR gives a name to the LAN by appending an octet to its own ID.

For example, assume a LAN has 5 routers and 100 end systems, as shown in Figure 4. Let R5 be the elected DR. R5 might name the LAN R5.17. In that case, R1, R2, R3, R4, and R5 each issue an LSP listing the neighbor R5.17. R5 will issue a second LSP, from source R5.17, listing R1, R2, R3, R4, R5, and all the end systems (E1 through E100) as neighbors.

The IS-IS protocol also contains special features to allow efficient distribution of LSPs on the LAN. IS-IS does not require explicit acknowledgments to LSPs on the LAN. Instead, a router that has an LSP to forward to the LAN simply multicasts the LSP to the other routers. A router that receives an LSP on the LAN will not multicast the same LSP on the LAN. Theoretically, if no packets get lost, only a single router would issue an LSP on the LAN.

However, packets do get lost, so the detection of lost LSPs is important. IS-IS detects lost LSPs by having the DR periodically broadcast a summary of the LSP database in a CSNP. Based on the CSNP, a receiving router can determine whether it has missed an LSP (in which case it will explicitly request the LSP from the DR), or it has a more recent LSP than the DR has (in which case the receiving router will multicast the LSP on the LAN to the other routers).

8 Database Overload

An implementation of a router typically has a finite amount of storage for the LSP database. Therefore, the router could receive an LSP and not be able to store it. The space may be inadequate for two reasons. First, the network may experience a static overload, i.e., the network may have become so large that the router cannot store the LSP database. Second, an ordering of events can temporarily make the LSP database larger than necessary, causing a temporary overload. For example, the DR on a large LAN may fail. The DR's previous pseudonode LSP is still in the other routers' databases. The new DR on the LAN will give the LAN a new ID and attempt to purge the previous pseudonode LSP. However, until the purge is complete, other routers will have to temporarily store twice as much information about that LAN.

Without considering this storage problem, a router implementation might employ any of the following strategies: the router might fail and recover only with operator intervention; the router might fail and reboot; or the router might ignore the temporary overload and perform routing in the best way possible.

Each of these possible strategies is undesirable. If a router fails and needs human intervention to recover, routing will be disrupted longer than necessary if the problem is only temporary. Crashing and automatically rebooting is desirable if the overload is very short-lived (so the overload condition is corrected before the router has rebooted). Otherwise, this strategy can cause long-term instability, since after rebooting, the router starts to exchange routing information with neighbors, only to eventually overload and fail again. Routing based on an incomplete LSP database can be dangerous and can cause widespread misrouting and/or routing loops.

IS-IS solves the storage problem by requiring a router that cannot store its LSP database to set an overload flag in its own LSP. Other routers then treat that router as an end system and route to that router but not through that router. Thus, the overloaded router is available through network management. If the router has not needed to refuse an LSP from a neighbor for a period of a minute (or as configured by network management), the router will clear the flag in its LSP. Thus, if the problem is temporary, the network will recover without human intervention. An important feature of this solution is that changing the flag does not change the size of the LSP database and hence does not lead to oscillation of the overloaded condition.

9 Limiting the Size of Routing Control Packets

Some IS-IS packets (specifically, LSPs and CSNPs) may become too large to be transmitted as single packets. Consequently, the packets may split into

several packets for transmission.

An LSP can become very large if a router has many neighbors. However, this situation is rarely an issue, except for the pseudonode LSP for a LAN. The IS-IS protocol avoids such large LSPs, which would need to be fragmented for transmission across each link and then reassembled at each router. The protocol has the LSP source break the LSP into individual fragments, each with its own unique ID and sequence number. The ID of the LSP is no longer simply the ID of the router issuing the LSP but has an additional octet appended to the router's (or pseudonode's) ID indicating the fragment number. Each fragment is independently flooded to the other routers. Only in the route computation is any connection made between the fragments of a router's LSP.

A CSNP can become large as well, since it includes the range of source addresses of LSPs to which it refers. If the range indicates x through y, then all LSPs with source IDs between x and y will be included and only those LSPs. Absence of an LSP that lies within the range implies that the issuing router has no knowledge of that LSP. Therefore, the IS-IS protocol can take action based on a CSNP fragment without waiting for all fragments. If a CSNP fragment is lost, then a lost LSP in that fragment's source address range might not be detected until the next time a CSNP fragment listing the ID of the lost LSP is transmitted.

10 Support of Multiple Protocols with IS-IS

Extending the IS-IS protocol to support multiple protocol suites is relatively straightforward. The OSI version of the IS-IS protocol supports routing for OSI CLNP, which also implies support for DECnet Phase V (since Phase V user data packets are identical to CLNP packets at the network layer). DECnet Phase V routing extends IS-IS to allow support for DECnet Phase V and for Phase IV-Phase V interoperability. Also, Digital worked on the Internet Engineering Task Force (IETF) to define the extension to IS-IS for support of IP.[13]

To understand how the OSI IS-IS protocol can be extended to support multiple protocol suites, consider what the IS-IS protocol provides. For example, consider a level 1 router within an area. The IS-IS routing protocol allows this router to know the identity and up/down status of the other routers and links in the area and which routers in the area are level 2 routers. IS-IS calculates routes to all other routers in the area. IS-IS also provides a number of important background functions, such as allowing information to be reliably broadcast between the routers in the area and allowing up/down status to be periodically checked. In addition, IS-IS allows each router to know which OSI addresses are reachable by means of each other router. (At level 1, the router would list the NSAPs of all its end-system neighbors; at level 2, the router would list all the areas and address prefixes it can reach.) IS-IS therefore already has most of the information needed to calculate routes for additional routing protocols.

To add routing support for another protocol suite such as IP, the IS-IS protocol is updated to announce the addresses that are reachable by means of that protocol suite. For example, to add IP support to IS-IS, a new field is defined in the LSPs to announce IP addresses, expressed in ordered pairs of the form (IP address, subnet mask). This allows IP addresses and OSI (i.e., DECnet Phase V) addresses to be assigned independently, while still allowing most of the overhead functions required by a routing protocol, such as checking link status and propagating the information, to be performed only once for all supported protocol suites.

If all routers support a particular protocol, the data packets for that protocol can be transmitted in native mode, i.e., no additional header is required. If some routers do not support a particular protocol, then the packet must be encapsulated in a network layer header for a network layer protocol that all the IS-IS routers do support. In DECnet Phase V, all the routers support both IP and CLNP, so these two protocols are transmitted in native mode. However, if support for another protocol is added, for instance AppleTalk support, then the routers that have AppleTalk neighbors need to be able to parse AppleTalk packets. However, other routers will not need to be modified. To facilitate knowing when to encapsulate, IS-IS routers announce which protocols they support in their IS-IS packets. Also, routers that support the AppleTalk protocol and have AppleTalk neighbors list in their LSPs that they can reach certain AppleTalk destinations.

The IS-IS packets are encoded such that a router can ignore information pertaining to protocol suites that the router does not support but can correctly interpret the rest of the IS-IS packet. Assume that R1 and R2 are the only two routers in an area that support the AppleTalk protocol. R1 and R2 therefore announce in their LSPs which AppleTalk destinations they can reach. R1 and R2 use a format for including AppleTalk information in IS-IS LSPs that other routers in the same area can forward but will otherwise ignore. Assume R2 receives an AppleTalk packet for forwarding with destination D3, reachable through R1. Then R2 encapsulates the packet as data inside a CLNP (or IP) packet with destination R1. When R1 receives the packet, it removes the CLNP header and forwards the packet to D3. If R1 and R2 are adjacent, or if all the routers along the path from R2 to R1 support the AppleTalk protocol, then encapsulation of AppleTalk packets inside CLNP packets would not be necessary. Thus, encapsulation occurs automatically only when needed for transmission through routers that do not support the protocol of the data packet to be forwarded.

Using one integrated routing protocol to route packets from multiple protocol suites has significant advantages over using a separate routing protocol for each suite. Probably the most important advantage is that only one routing protocol needs to be managed and configured. A single coordinated routing protocol can respond to network problems, such as link failures, in an efficient manner, improves bandwidth utilization,

and minimizes the CPU and memory requirements in routers. Also, integrated routing allows automatic encapsulation and eliminates the need for manual configuration of where and when to encapsulate.

11 Summary

IS-IS is a powerful and robust routing protocol. Many aspects are innovative and have been copied by other routing protocols. When looked at as a whole, the algorithms may appear complex, but when examined individually, the designated router election, the LSP propagation, and the LSP database overload procedure, for example, are all quite simple. IS-IS provides efficient routing for a variety of protocol suites, currently including DECnet Phase IV, CLNP/DECnet Phase V, and IP.

12 References

1. Information Procesing Systems, Data Communications: Protocol for Providing the Connectionless-Mode Network Service , ISO 8473 (Geneva: International Organization for Standardization, 1988).

2. J. Postel, "Internet Protocol," Internet Engineering Task Force RFC 791 (September 1981).

3. Information Processing Systems, Telecommunications and Information Exchange between Systems: End System to Intermediate System Routeing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (ISO 8473), ISO 9542 (Geneva: International Organization for Standardization, 1988).

4. D. Plummer, "Ethernet Address Resolution Protocol," Internet Engineering Task Force RFC 826 (November 1982).

5. J. Postel, "Internet Control Message Protocol," Internet Engineering Task Force RFC 792 (September 1981).

6. Information Technology, Telecommunications and Information Exchange between Systems: Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service (ISO 8473), ISO /IEC 10589 (Geneva: International Organization for Standardization /International Electrotechnical Commission, 1992).

7. G. Sidhu, R. Andrews, and A. Oppenheimer, Inside AppleTalk, Second Edition (Reading, MA: Addison-Wesley, 1990).

8. C. Hedrick, "Routing Information Protocol," Internet Engineering Task Force RFC 1058 (June 1988).

9. J. McQuillan, I. Richer, and E. Rosen, "ARPANET Routing Algorithm Improvements, First Semiannual Technical Report," BBN Report 3803 (April 1978).

10. E. Rosen et al., "ARPANET Routing Algorithm Improvements, Volume 1,") BBN Report 4473) (August 1980).

11. R. Perlman, Interconnections: Bridges and Routers (Reading, MA: Addison-Wesley, 1992).

12.E. Dijkstra, "A Note on Two Problems in Connection with Graphs," Numerical Mathematics, vol. 1 (1959): 269-271.

13.R. Callon, "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments," Internet Engineering Task Force RFC 1195 (December 1990).

13 Biographies

Ross W. Callon As a member of Digital's Network Architecture Group from 1988 to 1993, Ross Callon worked on routing algorithm and addressing issues. He was a primary author of the Integrated IS-IS protocol and of the guidelines for using NSAP addresses in the Internet. Previously, he was employed by Bolt Beranek and Newman as a senior scientist and helped develop the ISO CLNP protocol. Ross received a B.Sc. (1969) in mathematics from MIT and an M.Sc. (1977) in operations research from Stanford University. He is currently employed as a consulting engineer at Wellfleet Communications.

Radia J. Perlman As a member of the Network Architecture Group, Radia Perlman has been designing protocols for bridges and routers since joining Digital 13 years ago. She designed the spanning tree algorithm used by all standardized forms of bridges, as well as many of the protocols in IS-IS. Radia authored the book Interconnections: Bridges and Routers and has more than 20 patents filed or pending in the areas of bridging, routing, and network security. She holds S.B. and S.M. degrees in mathematics and a Ph.D. in computer science, all from the Massachusetts Institute of Technology.

I. Michael C. Shand Consulting engineer Michael Shand of Networks Engineering is responsible for the DNA Phase V network routing layer architecture. Prior to this, he worked on the Phase V X.25 access and HDLC architectures. He represents Digital on the ISO network layer committee and was a major contributor to the standardization of the IS-IS routing protocol (ISO/IEC 10589). Mike came to Digital in 1985 from Kingston Polytechnic (U.K.). He has an M.A. (1971) in natural sciences from the University of Cambridge and a Ph.D. (1975) in surface chemistry from Kingston Polytechnic.

14 Trademarks