Preface

The field of computer networks has seen explosive growth in technologies and products over the past decade. More and more users and networking systems have joined an interconnected world of communications, unconstrained by time and distance. Such explosive growth has led to many different solutions to the inevitable complexity of large networks, and those solutions are not always compatible. In 1992, IBM introduced the Networking Blueprint, followed in 1994 by the more comprehensive and encompassing Open Blueprint*, as a way to absorb the differences and create for users and systems a network that appears seamless, homogenous, and directly accessible, while the networks remain open, heterogeneous, and distributed.

The Open Blueprint is an architecture and framework, rather than a product or system. It defines the ways in which component products and services can be built that are compatible and provide users with a seamless view of data and networks. The architectural scheme must be filled in with real products and services, cooperating to provide the view specified by the architecture. This issue contains ten papers on a variety of technologies and systems that are compatible with the Open Blueprint. We are indebted to G. D. Schultz of the IBM Networking Hardware Division in Research Triangle Park, North Carolina, for his considerable efforts in creating, developing, and preparing this issue.

One aspect of heterogeneous networks that can be hidden from users and networks of users is the plethora of network protocols that have been developed and are now ingrained in applications and systems. These protocols—semantic and syntactic rules for the behavior of communicating functional units—have spread across environments and media with widely differing characteristics and technical requirements. Hess, Lorrain, and McGee use the Open Blueprint as a framework for discussing multiple protocols and how they can be joined in a seamless way through standardization, convergence, and coexistence. They also provide a description of asynchronous transfer mode (ATM), an advanced high-speed multiprotocol technology.

Christensen et al. describe another aspect of networks—the domain of local area networks (LANs). Unbeknownst to many users, LANs have progressed through three generations and are embarking on a fourth, leaving the simple image of a connected loop of machinery far behind, while preserving the appearance of a loop for users and connected machines. The authors discuss the evolution of LANs in order to set the stage for describing the fourth generation: LAN switching and the emergence of virtual LANs, which greatly increase apparent resources and allow LANs to cover large distances, without changing the user's hardware.

Among the latest developments in harnessing the power of large and distributed networks is the asynchronous transfer mode (ATM). ATM networks provide common support for data, voice, video, and other network traffic, while supporting both wide area and local area networks. The paper by Sultan and Basso is a tutorial and a description of ATM technology and relevant IBM products.

One example of a newer and more flexible network medium and protocol is described by Bauchot and Lanne in their paper on the IBM Wireless Radio Frequency (RF) LAN and its medium access control (MAC) protocol. High performance, short latency, and frequency hopping are provided by the adaptive nature of the MAC protocol, without sacrificing security, data compression, interference management, and attachment to existing LAN environments.

The complexity of some protocols, such as IBM's Systems Network Architecture (SNA) and Network Basic Input/Output System (NetBIOS*), has led to specialization and a subsequent lack of standardization among network routers that handle these popular protocols. Recently, proposals have been made for standardizing around data link switching (DLSw), which is already being implemented by various vendors. Gayek describes DLSw as a wide area network standard, compares it with earlier technologies, and speculates on future extensions and directions.

Bird et al. explore the history, recent enhancements, and likely future of the Advanced Peer-to-Peer Networking* (APPN*) architecture. High-performance routing is cited and described as an example of how APPN is evolving to meet the needs of ever faster networks and expanding customer requirements, under the watchful gaze of an industry consortium of APPN vendors. The paper also discusses such advanced areas as central directories, border nodes, dependent logical unit (LU) requesters, high-performance routing, and switched technology.

Underlying one of today's best known and largest networks—the Internet—is Transmission Control Protocol/Internet Protocol (TCP/IP). Britton, Tavs, and Bournas present the implications of the rapid growth in use of the Internet for the future of TCP/IP, while relating TCP/IP to the direction and technology of Open Blueprint. As the authors state, TCP/IP already supports 50 700 networks and 4.8 million host systems in 89 countries. So, its future and that of the Internet are of considerable interest to its current users, industry, academia, and IBM. Capacity, speed, and function must grow to anticipate future waves of networks, hosts, and users.

One of the inhibiting factors in the growth and flexibility of software systems in complex environments (like current networks) is the dependencies that are designed into software and then persist and are hard to remove. The better model for flexible and portable software is application-independent interfaces and protocols that allow software to operate without regard for the underlying structures. Pozefsky et al. explore this concept in the context of networks, focusing on the Multiprotocol Transport Networking (MPTN) architecture. One significant provision of MPTN allows current applications to execute unmodified on any existing protocol. MPTN is implemented today in the IBM AnyNet* family of products.

For the network application programmer, access to a portable and effective common programming interface is vital to portability and reuse of software across network architectures and protocols. Arnette et al. present the Common Programming Interface for Communications (CPI-C), describe the conversational services available to programs, and provide examples of CPI-C use for program-to-program communication.

Aldred et al. focus on an area that still presents a significant challenge for network architects: multimedia for real-time collaboration across networks. The authors discuss the challenge and requirements, along with existing standards, and describe the IBM Lakes architecture, which is designed to address these needs and provide for significant future technical enhancement.

*The next issue of the **Journal*** will be a special issue related to this one, covering aspects of Networking BroadBand Services (NBBS).

Gene F. Hoffnagle
Editor