PART ONE

INTRODUCTION TO OPEN SYSTEMS

1 INTRODUCTION

Books that discuss computer communications invariably begin by drawing analogies between computer networking and earlier, landmark inventions that have had a profound impact upon, perhaps even "revolutionized," society. Andrew S. Tanenbaum compares the impact of computer networks to the mechanical systems accompanying the Industrial Revolution, while Douglas E. Comer likens digital communications networks to the great railroads of the nineteenth century. But neither the Industrial Revolution nor the railroad has made as great an impact on human civilization as "the marriage of the engineering of telecommunications to that of the computer industry" (Martin 1976, 2). Why? No previous technology has advanced quite so rapidly and with such unbounded horizons as the computer, and no previous technology has achieved anything close to the ubiquity of the modern telecommunications system.

James Martin accurately predicted that through this union, the telecommunications system would aid distributed processing, and the computer would facilitate telephony switching. The actual chronology of events in fact exceeded Martin's expectations, for shortly after his speculation in the mid-1970s, information processing was delivered to the desktop. A decade marked by increased computer speed, memory, and storage, accompanied by a proliferation of useful and distributed applications, has fundamentally changed the way in which much of society works and interacts: we now send mail, do our banking, and exchange documents electronically, from our business places and our homes. This change in human behavior has affected the telecommunications system more profoundly than Martin forecast when he suggested that computers would merely facilitate switching. It adds a level of sophistication to the equipment attached to the telephone network that could never have been achieved by a telephone handset with a 12-digit keypad. In addition to placing a voice call to conduct business, we increasingly seek to exchange images—files of immense size—and to animate them in the process, and we expect do so in *milliseconds*. In many respects, information has become as important a commodity to switch as voice. The traditional voice and data networks will undergo profound changes in the next decade, as both seek to integrate the services of the other.

Therein lies a problem with the marriage. As in the Houses of Montague and Capulet, the parents of voice and data don't get along. Rarely have computer and communications providers shared a common set of beliefs and purposes. In the House of Telephony, data switching has historically been viewed as a second-tier service, incapable of ever achieving the "cash-cow" status of voice, and therefore much less important to the "bottom line." In the House of Data, telephony providers have been criticized as being intolerably slow to respond to the increasing demand for bandwidth, willing only to focus on "dataphony,"[1] and the data services offered by "common carriers" have historically been much less powerful and flexible than on-premises, local area networking alternatives.

Never have these differences of culture and philosophy been more obvious than during the development of Open Systems Interconnection, during which the debates between the Houses of Voice and Data were often more religious and political than technical.[2] This is perhaps because, by the mid-1970s, the networking of computers had begun to look like a lucrative new market opportunity rather than an amusing academic toy. The notion of open systems networking became *interesting* to both the voice and the data worlds at nearly the same time, for profit's sake and no other; and both the House of Data and the House of Voice wanted to secure as big a slice of the new pie as possible.

But what exactly is "open systems networking"? There are, of course, many ways to answer this question. The answer certainly does not lie strictly within the reference model for Open Systems Interconnection (ISO 7498, 1984), because architectures and protocols other than OSI

1. *Dataphony* is a term coined by Christine Hemrick, presently with Cisco Systems, to distinguish low-bandwidth, terminal-to-mainframe networking applications from high-bandwidth, distributed-processing applications—i.e., *real* data networking.

2. OSI wasn't the first pretext for these debates. The initial experience of the conflict for one of the authors can be traced back to the first time a Bell Telephone employee marked the area surrounding a data access arrangement in a computer laboratory at a Burroughs development facility with red tape, plopped down a Bell modem and telephone, and said "Don't touch!" Shortly thereafter, it was necessary to move the entire wall without disturbing the tape.

are widely acknowledged as bases for open systems networking. The publication of the OSI reference model is noteworthy primarily because it represents an internationally recognized effort at codifying what constitutes “openness.” What is recorded in the OSI reference model as the definition of an open system is in fact far less significant than the events that motivated—in the minds of some, provoked—an international interest in open systems networking.

Even at this late stage in the evolution of open systems networking, *any* attempt at defining *open systems* is highly subjective. For the purposes of this book, however, *open systems networking* implies or suggests the following: multivendor, interoperable hardware and software systems, based on internationally recognized and publicly available documentation (“standards”), which can be acquired “off the shelf” (as a standard rather than special-order product).

Why is everyone so excited about open systems? Some are excited because the concept represents “safe networking”: protection from proprietary networking solutions that lock users into dependence on the products and services of a single vendor (and thereby place users at the mercy of that vendor, in both an economic and a product- or feature-availability sense). Especially among government agencies that have spent millions of dollars (or the equivalent) on custom networking equipment, it is widely perceived that the enhanced interoperability brought about by openness and standards leads to a (desirable) highly competitive market, which will greatly reduce the cost of networking. Others have an altogether different concern: single-vendor solutions are not inherently evil, but information technology and distributed processing today span so many markets that no single vendor provides hardware and software solutions for every conceivable information technology application, and by necessity, companies with diverse needs *must* purchase information technology products from many vendors. Finally, some believe that open systems networking is the only way to achieve the service ubiquity of telephony for data.

Open systems networking and its associated standardization processes are an enormous undertaking that encompasses far more than establishing guidelines for data communications and information technology. Open systems standards have widely varying political and economic ramifications for users, equipment manufacturers, and network providers. For the network consumer, two very desirable effects of open systems standardization are to enhance interoperability and to foster a highly competitive market. For the vendor of a product line that interconnects via a proprietary networking technology, however, open systems standardization represents yet another opportunity for competitors to pry customers away from the hard-won market share that it has nurtured on that proprietary networking solution; the competitors, of course, view this as a major benefit. Finally, for the communications carriers in certain countries, standards are quite literally *enforceable laws* that govern the way in which public network resources may be used; standards offer them the means to extend their control over voice and postal services to data.

Today's open systems have different origins as well. OSI was, from the beginning, intended to be *the* open systems networking solution. TCP/IP[3] was not originally designed for such a lofty purpose; on the contrary, it began as a private networking experiment conducted within the U.S. computer science research community and supported by the Department of Defense Advanced Research Projects Agency (ARPA), with a potential for military applications. The ARPANET may have been the first operational packet-switching network, but few of those who designed and installed the original four-node network[4] in the fall of 1969 anticipated that in just over two decades, from such a humble beginning, a global Internet of over 1.5 million computers and an estimated 5 million users would evolve. And yet practically everything we know about packet switching, and a good deal of what we know about distributed processing, has been affected by the research and experimentation associated with the Internet. (A complete description of the history of TCP/IP is inappropriate here; for our purposes, it is sufficient to identify the landmark achievements in the history of TCP/IP [see Table 1.1].)

TCP/IP evolved into an open systems networking alternative largely due to the inability of the OSI standards developers to deliver the promised goods in a timely fashion, for the standards kept coming, and more were promised, but interoperable OSI implementations were hard to find. By 1984, so much hype had preceded the delivery of actual OSI-

3. The term *TCP/IP* is commonly used to refer either specifically to the transmission control protocol (TCP) and internet protocol (IP) or generally to the entire suite of protocols that have been developed by the Internet community to operate in conjunction with TCP and IP in the capital-*I* Internet (the global interconnection of networks running the TCP/IP protocols) and in individual enterprise-specific "internets."

4. The four original sites were the University of California at Los Angeles (UCLA), the Stanford Research Institute (SRI), the University of California at Santa Barbara (UCSB), and the University of Utah. According to Stephen Crocker, one of the graduate students who connected the first host—a SIGMA VII—to the first ARPANET interface message processor (IMP) at UCLA, "An RFP was released in the summer of '68, and Bolt Beranek and Newman (BBN) won. The contract called for delivery of a four node network in fall '69 with 50 (not 56) kilobit trunks. IMP 1 was delivered to UCLA prior to its scheduled delivery date of 9/1/69. SRI, UCSB, and Utah followed at monthly intervals" (Stephen D. Crocker, personal correspondence, December 1991).

TABLE 1.1 Landmarks in the History of TCP/IP

1969	RFC 1*	Network-control program (NCP) development First Internet RFC ("Host Software") IMP 1 delivered to UCLA
1971	RFC 114	First FTP, Telnet
1972	RFC 318	Telnet
1973	RFC 475	FTP† and (first) network mail system TCP development begins at DARPA
1974	RFC 675	First TCP implementations: SRI, BBN, UCLA
1979		Ethernet is born
1981	RFCs 786, 788 RFCs 791, 793	First mail-transfer protocol, first SMTP TCP and IP become Internet standards
1982	RFC 821	SMTP Internet standard (9/80: RFC 772)
1983	RFCs 882, 883 MIL-STD-1777 RFC 854	Domain Names U.S. DOD military standard for TCP Telnet Internet standard (6/80: RFC 764)
1985	RFC 959	FTP Internet standard (7/72: RFC 354)
1990	RFC 1157	SNMP Internet standard (8/88: RFC 1067)

* The network control program was completed and working prior to the introduction of the request for comments (RFC) document series. According to Jon Postel, who has been the editor of the RFC series since its inception, it was agreed at that time that the RFC series would include only "working documents" and that the ARPANET Network Information Center (NIC) would begin a standards series with NCP; somewhere along the line, NCP disappeared from the documentation. It is rumored that copies still exist in Stephen Crocker's documentation archives.

† According to Jon Postel, the origins of what we know as Internet mail are found in the "mail" command in this FTP specification.

compliant equipment, and so few interoperable products existed, that industry observers began to speak of it as a paper tiger. Proprietary networking solutions continued to dominate the market, and although manufacturers' marketing representatives talked a great deal about "conforming to the OSI reference model," their development groups managed to deliver only token products. The need for genuine open systems remained unfilled and was growing. Gradually, TCP/IP ate OSI's mid-eighties lunch, abetted in no small part by Dan Lynch's highly successful TCP/IP Implementers' Workshops, which evolved into the even more successful Interoperability Conferences and Exhibitions, sponsored by Interop, Inc. The first Implementers' Workshop was held in Monterey, California, in August 1986; the first Interop conference was held in Monterey, in March 1987, and by 1992, the attendance at the now semiannual Interop exceeded 50,000. Interop allowed TCP/IP vendors to demonstrate real products operating in a multivendor environment on real networks, while the OSI community endlessly debated the arcane merits of formal description techniques, conformance statements, and protocol implementation conformance statements.

Interop was just one of many enabling vehicles for the success of TCP/IP; the foremost was and remains the Internet infrastructure—the actual *Internet,* consisting of real networks and real systems—which facilitates experimentation and research on a global scale.[5] A grass-roots level of cooperation permeates the Internet, linking academics, network providers, and even the fiercest of competitors in the manufacturing sector: egos and company biases are frequently set aside to bring useful new technology into the Internet. This is the essence of what makes TCP/IP successful today.

Still, OSI keeps coming. The promise is quietly, but inexorably, becoming a reality. All of the critical-path protocol standards have been completed, and OSI X.400 message handling and X.500 directory applications are today operated over both pure-OSI stacks and hybrid stacks in enterprise networks and across the Internet. Real OSI products are now demonstrated alongside TCP/IP products at Interop as the industry attempts to shape the multiprotocol morass of today into the multiprotocol Internet of tomorrow.

Why has OSI failed to meet expectations? Unfortunately, OSI had to be all things to all people. It had to accommodate the needs of the teletex and videotex services; integrated services digital network (ISDN); and the government agencies and postal, telephone, and telegraph (PTT)

5. According to Larry Landweber, who keeps track of Internet connectivity throughout the world, 109 countries (as of summer 1992) have some sort of connectivity to the Internet through IP, BITNET, UUCP, OSI, or FIDONET links; of these, 46 countries have direct IP connectivity (Landweber 1992).

agencies of 20 or 30 countries. Practically every innovation that came along in the early stages of OSI development had to be included, and this generous policy of inclusion played to the detriment of OSI, complicating it in some cases beyond reason, creating uncertainties among product planners, and most important, sapping valuable expertise that should have been devoted to implementing and testing it.

Despite these handicaps, OSI has managed to bring a variety of valuable new services to data networking: "white pages" directory services, a powerful network-programming language, multimedia messaging, and routing and addressing mechanisms that permit internets (and potentially, the Internet) to be scaled up to very large size. These may soon be appreciated as landmark achievements (see Table 1.2).

Open systems networking today is about OSI *and* TCP/IP, and perhaps other protocol stacks as well. Migration, evolution, and transition from TCP/IP to OSI—marching orders from the 1980s—are regarded now as irrelevant strategies; the operative words for the 1990s (and beyond) are *coexistence* and *integration*. The Internet is experimenting with OSI directory and message handling applications because they *add value*. Gateways are now provided between OSI and TCP/IP mail applications because they *serve the community*. OSI transport services provided by TCP and IP support OSI applications where OSI transport protocols have yet to be deployed, and transport service bridges are used where necessary because *it is a practical thing to do*. Backbone and regional networks switch OSI and TCP/IP datagrams, host implementations are becoming "dual-stack," and SNMP is run over OSI because *it all works*. In the Internet, conformance takes a back seat to interoperability, and OSI will be far more useful as part of the Internet than it has ever been on its own.

In *Open Systems Networking: TCP/IP and OSI*, the authors hope to provide an understanding of how the components of TCP/IP and OSI work, how they are similar and how they differ, how they came to be what they are today, and how they might play together in the future. There is much cause for optimism and enthusiasm, and the authors hope to impart some of this to the readers.

Organization of This Book

The remainder of Part One, "Introduction to Open Systems," describes the OSI and TCP/IP standards processes (and their key participants), and establishes the convention of examining OSI and TCP/IP in a fea-

TABLE 1.2 **Landmarks in the History of OSI**

1977		published articles on OSI RM
1978	 CCITT X.200	First OSI standards meeting (TC 97/SC 16) First CCITT OSI RM (Grey Book)
1980	CCITT T series	Teletex (telematic services)
1983		First NIST OSI workshop
1984	ISO 7498 X.400 series	OSI reference model Message handling (MHS)
1987	ISO 8326/8327	Session service and protocol
1988	ISO 8571 X.500 series ISO 8822/8823 ISO 8072/8073 ISO 8473/9542 FIPS 146	File transfer (FTAM) The directory (ISO 9594) Presentation service and protocol Transport service and protocol Internetwork protocol and routing First release of U.S. GOSIP (8/88)
1990	ISO 9595/9596	Management service and protocol

Note: The dates shown in the left-hand column are the years in which the ISO Central Secretariat formally published the corresponding OSI standards. In many cases, the final version of a standard was widely available many years earlier (as an approved draft international standard; see Chapter 2).

ture-by-feature, side-by-side manner.

Part Two, "Open Network Architectures," examines the architectural models for TCP/IP and OSI. OSI promulgates a formal (and formidable) reference model; a single, readily identifiable "reference model" description of the TCP/IP architecture does not exist, so the authors compare the collection of TCP/IP architectural "folklore" to the more formal specification of the OSI reference model. Key concepts such as layering, services, and protocols are introduced, and the descriptive techniques used in OSI standards—the service model, state machines, and time-sequence diagrams—are covered here as well. OSI's data-definition and network-programming language, Abstract Syntax Notation One, is described here, as are names and addresses and the roles they play in open systems networking.

An architectural characteristic of open systems networks (as opposed to proprietary networks) is the assumption of a set of generic, or generally available, applications that become building blocks (tools) for creating more complex distributed systems. In OSI, these are called *application service elements*; in TCP/IP, simply "applications." OSI and TCP/IP differ somewhat in the way in which applications are constructed. The differences between OSI and TCP/IP application "service" architectures are described in Part Three, "Upper Layers." Three "daily-use" application services that are common to TCP/IP and OSI—electronic mail, directories/information services, and network management—are presented here. An overview of the basic requirements of applications—synchronization, token control, connection management, activity management, remote operations, and reliable transfer—are introduced at a conceptual level here as well, so that readers have a general understanding of these capabilities to which they can refer when the specific mechanisms in the layers that provide these capabilities (presentation and session) are discussed later in the book.

Part Four, "Middle Layers," examines how end-to-end data transport, internetworking, and routing are performed in OSI and TCP/IP. The similarities and differences that exist between OSI and TCP/IP transport services, for example, are presented at a "bit level" of detail. The roots and history of the "connections versus datagrams" debate (which persists even today within the OSI community) are exposed here as well. Rather than include an exhaustive recapitulation of readily available information about existing point-to-point link and LAN technologies at the data-link layer, *Open Systems Networking: TCP/IP and OSI* focuses on *emerging digital technologies* that have been touted as broadband[6] platforms for advanced distributed applications: frame relay, FDDI, SMDS, and broadband ISDN.

Part Five, "The Future of Open Systems Networking," attempts a Hegelian *synthesis* of TCP/IP (*thesis*) and OSI (*antithesis*) by reviewing the guidelines for, and politics of, building a multiprotocol Internet. This part describes the status of the Internet activities that are directed at expanding the internetworking platform of the Internet to sustain its remarkable growth and examines issues related to evolving the Internet from its current TCP/IP core to a system that supports internetworking based on OSI, XNS/IPX, and AppleTalk® as well.

---

6. With apologies to electrical engineers—in particular, those who are familiar with the notion of broadband as it is applied in the world of local area networks—the term *broadband* is used here in the telephony sense of the word; i.e., transmission rates in excess of 1 megabit per second.