gives

$$\frac{x-2}{x-1}\cdot\frac{x-3}{x-2}\cdots\frac{x-(n+1)}{x-n}=\frac{x-(n+1)}{x-1}$$

where $x=2/p$.

(b) From the result of previous question,

$$\alpha=\frac{x-(n+1)}{x-1}$$

Therefore,

$$x=\frac{(n+1)-\alpha}{1-\alpha}=2/p$$

Accordingly,

$$p=\frac{2(1-\alpha)}{(n+1)-\alpha}$$

46. At every second, the bucket volume must not be negative. For a given bucket depth D, and token rate r, we can calculate the bucket volume $v(t)$ at time t seconds, and enforce $v(t)$ to be nonnegative.

$$v(0)=D-5+r=D-(5-r)\geq 0$$
$$v(1)=D-5-5+2r=D-2(5-r)\geq 0$$
$$v(2)=D-5-5-1+3r=D-(11-3r)\geq 0$$
$$v(3)=D-5-5-1+4r=D-(11-4r)\geq 0$$
$$v(4)=D-5-5-1-6+5r=D-(17-5r)\geq 0$$
$$v(5)=D-5-5-1-6-1+6r=D-6(3-r)\geq 0$$

We define the functions $f_1(r), f_2(r), \ldots, f_6(r)$ as follows.

$$f_1(r)=5-r$$
$$f_2(r)=2(5-r)=2f_1(r)\geq f_1(r) \quad (\text{for } 1\leq r\leq 5)$$
$$f_3(r)=11-3r\leq f_2(r) \quad (\text{for } r\geq 1)$$
$$f_4(r)=11-4r<f_3(r) \quad (\text{for } r\geq 1)$$
$$f_5(r)=17-5r$$
$$f_6(r)=6(3-r)\leq f_5(r) \quad (\text{for } r\geq 1)$$

First of all, for $r \geq 5$, $f_i(r) \leq 0$ for all i. This means if the token rate is faster than 5 packets per second any positive bucket depth will suffice (i.e., $D \geq 0$). For $1 \leq r \leq 5$, we only need to consider $f_2(r)$ and $f_5(r)$, since other functions are less than these functions. One can easily find $f_2(r) - f_5(r) = 3r - 7$. Therefore, the bucket depth D is enforced by the following formula:

$$D \geq \begin{cases} f_5(r) = 17 - 5r & (r = 1, 2) \\ f_2(r) = 2(5 - r) & (r = 3, 4, 5) \\ 0 & (r \geq 5) \end{cases}$$

Chapter 7

2. `4 M A R Y 4377 7 J A N U A R Y 7 2002 2 90000 150000 1`

8.

INT	4	15
INT	4	29496729
INT	4	58993458

10.
```
15 be 00000000 00000000 00000000 00001111
15 le 00001111 00000000 00000000 00000000

29496729 be 00000001 11000010 00010101 10011001
29496729 le 10011001 00010101 11000010 00000001

58993458 be 00000011 10000100 00101011 00110010
58993458 le 00110010 00101011 10000100 00000011
```

GLOSSARY

3DES: Triple DES, a version of DES that uses three keys, effectively increasing the key size and robustness of the encryption.

3G: Third-generation mobile wireless, a class of cellular wireless technologies based on CDMA.

4B/5B: A type of bit-encoding scheme used in FDDI, in which every 4 bits of data are transmitted as a 5-bit sequence.

802.3: IEEE Ethernet standard.

802.5: IEEE token ring standard.

802.11: IEEE wireless network standard.

802.17 IEEE resilient packet ring standard.

822: Refers to RFC 822, which defines the format of Internet email messages. See *SMTP*.

AAL: ATM Adaptation Layer. A protocol layer, configured over ATM. Two AALs are defined for data communications, AAL3/4 and AAL5. Each protocol layer provides a mechanism to segment large packets into cells at the sender and to reassemble the cells back together at the receiver.

ABR: (1) Available bit rate. A rate-based congestion-control scheme being developed for use on ATM networks. ABR is intended to allow a source to increase or decrease its allotted rate, based on feedback from switches within the network. Contrast with *CBR*, *UBR*, and *VBR*. (2) Area border router. Router at the edge of an *area* in a link-state protocol.

ACK: An abbreviation for *acknowledgment*. An acknowledgment is sent by a receiver of data to indicate to the sender that the data transmission was successful.

additive increase/multiplicative decrease: Congestion window strategy used by TCP. TCP opens the congestion window at a linear rate, but halves it when losses are experi-

enced due to congestion. It has been shown that additive increase/multiplicative decrease is a necessary condition for a congestion-control mechanism to be stable.

AES: Advanced Encryption Standard. A cryptographic cipher that has been proposed to supersede DES.

AF: Assured forwarding. One of the per-hop behaviors proposed for Differentiated Services.

ALF: Application Level Framing. A protocol design principle that says that application programs better understand their communication needs than do general-purpose transport protocols.

AMPS: Advanced mobile phone system. Analog-based cell phone system. Currently being replaced by digital system, known as PCS.

ANSI: American National Standards Institute. Private U.S. standardization body that commonly participates in the ISO standardization process. Responsible for SONET.

API: Application programming interface. Interface that application programs use to access the network subsystem (usually the transport protocol). Usually OS-specific. The socket API from Berkeley Unix is a widely used example.

area: In the context of link-state routing, a collection of adjacent routers that share full routing information with each other. A routing domain is divided into areas to improve scalability.

ARP: Address Resolution Protocol. Protocol of the Internet architecture, used to translate high-level protocol addresses into physical hardware addresses. Commonly used on the Internet to map IP addresses into Ethernet addresses.

ARPA: Advanced Research Projects Agency. One of the research and development organizations within the Department of Defense. Responsible for funding the ARPANET as well as the research that led to the development of the TCP/IP Internet. Also known as DARPA, the *D* standing for Defense.

ARPANET: An experimental wide-area packet-switched network funded by ARPA and begun in the late 1960s, which became the backbone of the developing Internet.

ARQ: Automatic repeat request. General strategy for reliably sending packets over an unreliable link. If the sender does not receive an ACK for a packet after a certain time period, it assumes that the packet did not arrive (or was delivered with bit errors) and retransmits it. Stop-and-wait and sliding window are two example ARQ protocols. Contrast with *FEC*.

ASN.1: Abstract Syntax Notation One. In conjunction with BER, a presentation-formatting standard devised by the ISO as part of the OSI architecture.

ATM: Asynchronous transfer mode. A connection-oriented network technology that uses small, fixed-size packets (called *cells*) to carry data.

ATMARP: Address Resolution Protocol as enhanced for ATM networks.

ATM Forum: A key ATM standards-setting body.

authentication: Security protocol by which two suspicious parties prove to each other that they are who they claim to be.

autonomous system (AS): A group of networks and routers, subject to a common authority and using the same intradomain routing protocol.

bandwidth: A measure of the capacity of a link or connection, usually given in units of bits per second.

Bellman-Ford: A name for the distance-vector routing algorithm, from the names of the inventors.

BER: Basic encoding rules. Rules for encoding data types defined by ASN.1.

best-effort delivery: The service model of the current Internet architecture. Delivery of a message is attempted but is not guaranteed.

BGP: Border Gateway Protocol. An interdomain routing protocol by which autonomous systems exchange reachability information. The most recent version is BGP-4.

BISYNC: Binary Synchronous Communication. A byte-oriented link-level protocol developed in the late 1960s by IBM.

bit stuffing: A technique used to distinguish control sequences and data on the bit level. Used by the HDLC protocol.

block: An OS term used to describe a situation in which a process suspends execution while awaiting some event, such as a change in the state of a *semaphore*.

Bluetooth: A short-range wireless standard used to connect computers, mobile phones, and peripheral devices, among other things.

bridge: A device that forwards link-level frames from one physical network to another, sometimes called a LAN switch. Contrast with *repeater* and *router*.

broadcast: A method of delivering a packet to every host on a particular network or internet. May be implemented in hardware (e.g., Ethernet) or software (e.g., IP broadcast).

CA: Certification authority (also known as certificate authority). An entity that signs security certificates, thereby promising that the public key contained in the certificate belongs to the entity named in the certificate.

CBC: Cipher block chaining. A cryptographic mode in which each plaintext block is XORed with the previous block of ciphertext before encryption.

CBR: Constant bit rate. A class of service in ATM that guarantees transmission of data at a constant bit rate, thus emulating a dedicated transmission link. Contrast with *ABR*, *UBR*, and *VBR*.

CCITT: The now defunct *Comité Consultif International de Telegraphique et Telephonique*, a unit of the International Telecommunications Union (ITU) of the United Nations. Now replaced by ITU-T.

CDMA: Code Division Multiple Access, a form of multiplexing used in wireless networks.

CDN: Content distribution network. A collection of surrogate web servers, distributed across the Internet, that respond to web HTTP requests in place of the server. The goal of widely distributing the surrogate servers is to have a surrogate close to the client, making it possible to respond to requests more quickly.

cell: A 53-byte ATM packet, capable of carrying up to 48 bytes of data.

certificate: A document digitally signed by one entity that contains the name and public key of another entity. Used to distribute public keys. Also see *CA*.

channel: A generic communication term used in this book to denote a logical process-to-process connection.

checksum: Typically a ones complement sum over some or all of the bytes of a packet, computed and appended to the packet by the sender. The receiver recomputes the checksum and compares it to the one carried in the message. Checksums are used to detect errors in a packet and may also be used to verify that the packet has been delivered to the correct host. The term *checksum* is also sometimes (imprecisely) used to refer generically to error-detecting codes.

chipping code: Random sequence of bits that is XORed with the data stream to implement the direct sequence technique of spread spectrum.

CIDR: Classless interdomain routing. A method of aggregating routes that treats a block of contiguous Class C IP addresses as a single network.

circuit switching: A general strategy for switching data through a network. It involves establishing a dedicated path (circuit) between the source and destination. Contrast with *packet switching*.

client: The requester of a service in a distributed system.

CLNP: Connectionless Network Protocol. The ISO counterpart to the Internet's IP.

clock recovery: The process of deriving a valid clock from a serially transmitted digital signal.

concurrent logical channels: Multiplexing several stop-and-wait logical channels onto a single point-to-point link. No delivery order is enforced. This mechanism was used by the IMP-IMP protocol of the ARPANET.

congestion: A state resulting from too many packets contending for limited resources (e.g., link bandwidth and buffer space on routers or switches), which may force the router (switch) to discard packets.

congestion control: Any network resource management strategy that has, as its goal, the alleviation or avoidance of congestion. A congestion-control mechanism may be implemented on the routers (switches) inside the network, by the hosts at the edges of the network, or by a combination of both.

connection: In general, a channel that must be established prior to use (e.g., by the transmission of some setup information). For example, TCP provides a connection abstraction that offers reliable, ordered delivery of a byte stream. Connection-oriented networks, such as ATM, are often said to provide a *virtual circuit* abstraction.

connectionless protocol: A protocol in which data may be sent without any advance setup. IP is an example of such a protocol.

context switch: An operation in which an operating system suspends the execution of one process and begins the execution of another. A context switch involves saving the state of the former process (e.g., the contents of all registers) and loading the state of the latter process.

controlled load: One of the service classes available in the Internet's Integrated Services architecture.

CRC: Cyclic redundancy check. An error-detecting code computed over the bytes composing a packet and then appended to the packet by the network hardware (e.g., Ethernet adaptor). CRC provides stronger error detection than a simple checksum.

crossbar switch: A simple switch design in which every input is directly connected to every output and the output port is responsible for resolving contention.

CSMA/CD: Carrier Sense Multiple Access with Collision Detect. CSMA/CD is a functionality of network hardware. "Carrier sense multiple access" means that multiple stations can listen to the link and detect when it is in use or idle; "collision detect" indicates that if two or more stations are transmitting on the link simultaneously, they will detect the collision of their signals. Ethernet is the best-known technology that uses CSMA/CD.

cut-through: A form of switching or forwarding in which a packet starts to be transferred to an output before it has been completely received by the switching node, thus reducing latency through the node.

datagram: The basic transmission unit in the Internet architecture. A datagram contains all of the information needed to deliver it to its destination, analogous to a letter in the U.S. postal system. Datagram networks are connectionless.

DCE: Distributed Computing Environment. An RPC-based suite of protocols and standards that support distributed computing. Defined by OSF.

DDCMP: Digital Data Communication Message Protocol. A byte-oriented link-level protocol used in Digital Equipment Corporation's DECNET.

DDoS: Distributed denial of service. A DoS attack in which the attack originates at a set of nodes. Each attacking node may put only a marginal load on the target machine, but the aggregate load from all the attacking nodes swamps the target machine.

DECbit: A congestion-control scheme in which routers notify the endpoints of imminent congestion by setting a bit in the header of routed packets. The endpoints decrease their sending rates when a certain percentage of received packets have the bit set.

decryption: The act of reversing an *encryption* process to recover the data from an encrypted message.

delay bandwidth product: The product of a network's RTT and bandwidth. Gives a measure of how much data can be in transit on the network.

demultiplexing: Using information contained in a packet header to direct it upward through a protocol stack. For example, IP uses the ProtNum field in the IP header to

decide which higher protocol (i.e., TCP, UDP) a packet belongs to, and TCP uses the port number to demultiplex a TCP packet to the correct application process. Contrast with *multiplexing.*

demultiplexing key: A field in a packet header that enables demultiplexing to take place (e.g., the ProtNum field of IP).

dense mode multicast: PIM mode used when most routers or hosts need to receive multicast packets.

DES: Data Encryption Standard. An algorithm for data encryption based on a 64-bit secret key.

DHCP: Dynamic Host Configuration Protocol. A protocol used by a host as it boots or when it is connected to a network, to learn various network information, such as its IP address.

DHT: Distributed hash table. A technique by which a message is routed toward a machine that supports a particular object, based on the object's name. The object is hashed to a unique identifier, with each intermediate node along the route forwarding the message to a node that is able to interpret a larger prefix of this ID. DHTs are often used in peer-to-peer networks.

Differentiated Services: A new architecture for providing better than best-effort service on the Internet. It has been proposed as an alternative to Integrated Services.

direct sequence: A spread spectrum technique that involves XORing the data stream with a random bit sequence known as a chipping code.

distance vector: A lowest-cost-path algorithm used in routing. Each node advertises reachability information and associated costs to its immediate neighbors, and uses the updates it receives to construct its forwarding table. The routing protocol RIP uses a distance-vector algorithm. Contrast with *link state.*

DMA: Direct memory access. An approach to connecting hosts to I/O devices, in which the device directly reads data from and writes data to the host's memory. Also see *PIO.*

DNA/DECNET: Digital Network Architecture. An OSI-based architecture that supports a connectionless network model and a connection-oriented transport protocol.

DNS: Domain name system. The distributed naming system of the Internet, used to resolve host names (e.g., cicada.cs.princeton.edu) into IP addresses (e.g., 192.12.69.35). The DNS is implemented by a hierarchy of name servers.

domain: Can refer either to a context in the hierarchical DNS namespace (e.g., the "edu" domain) or to a region of the Internet that is trea ed as a single entity for the purpose of hierarchical routing. The latter is equivalent to *autonomous system*.

DoS: Denial of service. A situation in which an attacking node floods a target node with so much work (so many packets) that it effectively keeps legitimate users from accessing the node, hence, they are denied service.

DS3: A 44.7-Mbps transmission link service offered by the phone company. Also called T3.

DSL: Digital subscriber line. A family of standards for transmitting data over twisted pair telephone lines at multimegabit-per-second speeds.

duplicate ACK: A retransmission of a TCP acknowledgment. The duplicate ACK does not acknowledge any new data. The receipt of multiple duplicate ACKs triggers the TCP *fast retransmit* mechanism.

DVMRP: Distance Vector Multicast Routing Protocol. Multicast routing protocol originally used in the MBone.

DWDM: Dense wavelength division multiplexing. Multiplexing multiple light waves (colors) onto a single physical fiber. The technique is "dense" in the sense that a large number of optical wavelengths can be supported.

ECN: Explicit congestion notification. A technique by which routers inform end hosts about congestion by setting a flag in packets they are forwarding. Used in conjunction with active queue management algorithms like RED.

EF: Expedited forwarding. One of the per-hop behaviors proposed for Differentiated Services.

EGP: Exterior Gateway Protocol. An early interdomain routing protocol of the Internet, which was used by exterior gateways (routers) of autonomous systems to exchange routing information with other ASs. Replaced by BGP.

encapsulation: The operation, performed by a lower-level protocol, of attaching a protocol-specific header and/or trailer to a message passed down by a higher-level protocol. As a message travels down the protocol stack, it gathers a sequence of headers, of which the outermost corresponds to the protocol at the bottom of the stack.

encryption: The act of applying a transforming function to data, with the intention that only the receiver of the data will be able to read it (after applying the inverse function,

decryption). Encryption generally depends on either a secret shared by the sender and receiver or on a public/private key pair.

Ethernet: A popular local area network technology that uses CSMA/CD and has a bandwidth of 10 Mbps. An Ethernet itself is just a passive wire; all aspects of Ethernet transmission are completely implemented by the host adaptors.

exponential backoff: A retransmission strategy that doubles the timeout value each time a packet is retransmitted.

exposed node problem: Situation that occurs on a wireless network where two nodes receive signals from a common source, but each is able to reach other nodes that do not receive this signal.

extended LAN: A collection of LANs connected by bridges.

fabric: The part of a switch that actually does the switching, that is, moves packets from input to output. Contrast with *port*.

fair queuing (FQ): A round-robin-based queuing algorithm that prevents a badly behaved process from capturing an arbitrarily large portion of the network capacity.

fast retransmit: A strategy used by TCP that attempts to avoid timeouts in the presence of lost packets. TCP retransmits a segment after receiving three consecutive duplicate ACKs, acknowledging the data up to (but not including) that segment.

FDDI: Fiber Distributed Data Interface. A token ring networking technology designed to run over optical fiber.

FEC: **1** Forward error correction. A general strategy for recovering from bit errors introduced into data packets without having to retransmit the packet. Redundant information is included with each packet that can be used by the receiver to determine which bits in a packet are incorrect. Contrast with *ARQ*.

2 Forwarding equivalence class. A set of packets that are to receive the same forwarding treatment at a router. MPLS labels are normally associated with FECs.

Fibre Channel: A bidirectional link protocol commonly used to connect computers, peripherals, and storage devices. Originally had a bandwidth of 100 MBps but since enhanced to GBps speeds.

firewall: A router that has been configured to filter (not forward) packets from certain sources. Used to enforce a security policy.

flow control: A mechanism by which the receiver of data throttles the transmission rate of the sender, so that data will not arrive too quickly to be processed. Contrast with *congestion control.*

flowspec: Specification of a flow's bandwidth and delay requirements presented to the network to establish a reservation. Used with RSVP.

forwarding: The operation performed by a router on every packet: receiving it on an input, deciding what output to send it to, and sending it there.

forwarding table: The table maintained in a router that lets it make decisions on how to forward packets. The process of building up the forwarding table is called *routing*, and thus the forwarding table is sometimes called a *routing table*. In some implementations, the routing and forwarding tables are separate data structures.

fragmentation/reassembly: A method for transmission of messages larger than the network's MTU. Messages are fragmented into small pieces by the sender and reassembled by the receiver.

frame: Another name for a packet, typically used in reference to packets sent over a single link rather than a whole network. An important problem is how the receiver detects the beginning and ending of a frame, a problem known as framing.

Frame Relay: A connection-oriented public packet-switched service offered by the phone company.

frequency hopping: A spread spectrum technique that involves transmitting data over a random sequence of frequencies.

FTP: File Transfer Protocol. The standard protocol of the Internet architecture for transferring files between hosts. Built on top of TCP.

GMPLS: Generalized MPLS. Allows IP to run natively over optically-switched networks.

GPRS: General Packet Radio Service. A packet transmission service provided by cellular wireless networks.

GSM: Global System for Mobile communication. Digital cellular phone system being deployed throughout the world (less so in the United States and Canada). Similar to PCS, which is being deployed throughout the United States and Canada.

gopher: An Internet information service.

H.323: Session control protocol often used for Internet telephony.

handle: In programming, an identifier or pointer that is used to access an object.

hardware address: The link-level address used to identify the host adaptor on the local network.

HDLC: High-Level Data Link Control protocol. An ISO-standard link-level protocol. It uses bit stuffing to solve the framing problem.

hidden node problem: Situation that occurs on a wireless network where two nodes are sending to a common destination, but are unaware that the other exists.

hierarchical routing: A multilevel routing scheme that uses the hierarchical structure of the address space as the basis for making forwarding decisions. For example, packets might first be routed to a destination network and then to a specific host on that network.

HiPPI: High Performance Parallel Interface. An ANSI-standard network technology capable of Gbps transmission rates, typically used to connect supercomputers to peripheral devices. Used in same way as *Fibre Channel.*

host: A computer attached to one or more networks that supports users and runs application programs.

HTML: HyperText Markup Language. A language used to construct World Wide Web pages.

HTTP: HyperText Transport Protocol. An application-level protocol based on a request/reply paradigm and used in the World Wide Web. HTTP uses TCP connections to transfer data.

IAB: Internet Architecture Board. The main body that oversees the development of the Internet architecture.

IBGP: Interior BGP. The protocol used to exchange interdomain routing information among routers in the same domain.

ICMP: Internet Control Message Protocol. This protocol is an integral part of IP. It allows a router or destination host to communicate with the source, typically to report an error in IP datagram processing.

IEEE: Institute for Electrical and Electronics Engineers. A professional society for engineers that also defines network standards, including the 802 series of LAN standards.

IETF: Internet Engineering Task Force. The body responsible for the specification of standards and protocols related to the Internet.

IMAP: Internet Message Access Protocol. An application layer protocol that allows a user to retrieve her email from a mail server.

IMP-IMP: A byte-oriented link-level protocol used in the original ARPANET.

Integrated Services: Usually taken to mean a packet-switched network that can effectively support both conventional computer data and real-time audio and video. Also, a name given to a proposed Internet service model that was designed to supplement the current best-effort service model.

integrity: In the context of network security, a service that ensures that a received message is the same one that was sent.

interdomain routing: The process of exchanging routing among different routing domains. BGP is an example of an interdomain protocol.

internet: A collection of (possibly heterogeneous) packet-switching networks interconnected by routers. Also called an internetwork.

Internet: The global internet based on the Internet (TCP/IP) architecture, connecting millions of hosts worldwide.

interoperability: The ability of heterogeneous hardware and multivendor software to communicate by correctly exchanging messages.

interrupt: An event (typically generated by a hardware device) that tells the operating system to stop its current activity and take some action. For example, an interrupt is used to notify the OS that a packet has arrived from the network.

intradomain routing: The exchange of routing information within a single domain or autonomous system. RIP and OSPF are example intradomain protocols.

IP: Internet Protocol (also known as IPv4). A protocol that provides a connectionless, best-effort delivery service of datagrams across the Internet.

IPng: Internet Protocol—Next Generation (also known as IPv6). Proposed version of IP that provides a larger, more hierarchical address space and other new features.

IPSEC: IP Security. An architecture for authentication, privacy, and message integrity, among other security services to the Internet architecture.

IRTF: Internet Research Task Force. A sibling body to the IETF, responsible for charting direction in research and development for the Internet.

IS-IS: A link-state routing protocol, similar to OSPF.

ISDN: Integrated Services Digital Network. A digital communication service offered by telephone carriers and standardized by ITU-T. ISDN combines voice connection and digital data services in a single physical medium.

ISO: International Standards Organization. The international body that drafted the seven-layer OSI architecture and a suite of protocols that has not enjoyed commercial success.

ITU-T: A subcommittee of the International Telecommunications Union, a global body that drafts technical standards for all areas of international analog and digital communication. ITU-T deals with standards for telecommunications, notably ATM.

jitter: Variation in network latency. Large jitter has a negative impact on the quality of video and audio applications.

JPEG: Joint Photographic Experts Group. Typically used to refer to a widely used algorithm for compressing still images that was developed by the JPEG.

Kerberos: A TCP/IP-based authentication system developed at MIT, in which two hosts use a trusted third party to authenticate each other.

key distribution: Mechanism by which users learn each others' public keys through the exchange of digitally signed certificates.

LAN: Local area network. A network based on any physical network technology that is designed to span distances of up to a few thousand meters (e.g., Ethernet or FDDI). Contrast with *SAN, MAN,* and *WAN.*

LANE: Local area network emulation. Adding functionality to ATM to make it behave like a shared-media (i.e., Ethernet-like) LAN.

LAN switch: Another term for a *bridge*, usually applied to a bridge with many ports. Also called an Ethernet switch if the link technology it supports is Ethernet.

latency: A measure of how long it takes a single bit to propagate from one end of a link or channel to the other. Latency is measured strictly in terms of time.

LDAP: Lightweight Directory Access Protocol. A subset of the X.500 directory service that has recently become a popular directory service for information about users.

LER: Label edge router. A router at the edge of an MPLS cloud. Performs a complete IP lookup on arriving IP packets, and then applies labels to them as a result of the lookup.

link: A physical connection between two nodes of a network. It may be implemented over copper or fiber-optic cable or it may be a wireless link (e.g., a satellite).

link-level protocol: A protocol that is responsible for delivering frames over a directly connected network (e.g., an Ethernet, token ring, or point-to-point link). Also called link-layer protocol.

link state: A lowest-cost-path algorithm used in routing. Information on directly connected neighbors and current link costs are flooded to all routers; each router uses this information to build a view of the network on which to base forwarding decisions. The OSPF routing protocol uses a link-state algorithm. Contrast with *distance vector*.

LSR: Label-switching router. A router that runs IP control protocols, but uses the label switching forwarding algorithm of MPLS.

MAC: Media access control. Algorithms used to control access to shared-media networks like Ethernet and FDDI.

MACA: Multiple access with collision avoidance. Distributed algorithm used to mediate access to a shared media.

MACAW: Multiple access with collision avoidance for wireless. Enhancement of the general MACA algorithm to better support wireless networks. Used by 802.11.

MAN: Metropolitan area network. A network based on any of several new network technologies that operate at high speeds (up to several Gbps) and across distances wide enough to span a metropolitan area. Contrast with *SAN*, *LAN*, and *WAN*.

Manchester: A bit-encoding scheme that transmits the exclusive-OR of the clock and the NRZ-encoded data. Used on the Ethernet.

MBone: Multicast backbone. A logical network imposed over the top of the Internet, in which multicast-enhanced routers use tunneling to forward multicast datagrams across the Internet.

MD5: Message Digest version 5. An efficient cryptographic checksum algorithm commonly used to verify that the contents of a message are unaltered.

MIB: Management information base. Defines the set of network-related variables that may be read or written on a network node. The MIB is used in conjunction with SNMP.

MIME: Multipurpose Internet Mail Extensions. Specifications for converting binary data (such as image files) to ASCII text, which allows it to be sent via email.

Mosaic: A once-popular and free graphical World Wide Web browser developed at the National Center for Supercomputing Applications at the University of Illinois.

MP3: MPEG Layer 3. Audio compression standard used with MPEG.

MPEG: Moving Picture Experts Group. Typically used to refer to an algorithm for compressing video streams developed by the MPEG.

MPLS: Multiprotocol Label Switching. A collection of techniques used to effectively implement IP routers on top of level 2 (e.g., ATM) switches.

MSAU: Multistation access unit. A device used in token ring networks to connect several stations to the ring and remove them in the event of failure.

MSDP: Multicast Source Discovery Protocol. A protocol used to facilitate interdomain multicast.

MTU: Maximum transmission unit. The size of the largest packet that can be sent over a physical network.

multicast: A special form of broadcast in which packets are delivered to a specified subgroup of network hosts.

multiplexing: Combining distinct channels into a single, lower-level channel. For example, separate TCP and UDP channels are multiplexed into a single host-to-host IP channel. The inverse operation, *demultiplexing*, takes place on the receiving host.

name resolution: The action of resolving host names (which are easy for humans to read) into their corresponding addresses (which machines can read). See *DNS*.

NAT: Network address translation. A technique for extending the IP address space that involves translating between globally understood IP addresses and local-only addresses at the edge of a network or site.

NDR: Network Data Representation. The data-encoding standard used in the Distributed Computing Environment (DCE), as defined by the Open Software Foundation. NDR uses a receiver-makes-right strategy and inserts an architecture tag at the front of each message.

network-level protocol: A protocol that runs over switched networks, directly above the link level.

NFS: Network File System. A popular distributed file system developed by Sun Microsystems. NFS is based on SunRPC, an RPC protocol developed by Sun.

NIST: National Institute for Standards and Technology. The official U.S. standardization body.

node: A generic term used for individual computers that make up a network. Nodes include general-purpose computers, switches, and routers.

NRZ: Nonreturn to zero. A bit-encoding scheme that encodes a 1 as the high signal and a 0 as the low signal.

NRZI: Nonreturn to zero inverted. A bit-encoding scheme that makes a transition from the current signal to encode a 1 and stays at the current signal to encode a 0.

NSF: National Science Foundation. An agency of the U.S. government that funds scientific research in the United States, including research on networks and on the Internet infrastructure.

nv: Network video. A videoconferencing application.

OC: Optical carrier. The prefix for various rates of SONET optical transmission. For example, OC-1 refers to the SONET standard for 51.84-Mbps transmission over fiber. An OC-*n* signal differs from an STS-*n* signal only in that the OC-*n* signal is scrambled for optical transmission.

ONC: Open Network Computing. A version of SunRPC that is being standardized for the Internet.

optical switch: A switching device that forwards optical lightwaves from input port to output port without converting to electrical format.

OSF: Open Software Foundation. A consortium of computer vendors that have defined standards for distributed computing, including the NDR presentation format.

OSI: Open Systems Interconnection. The seven-layer network reference model developed by the ISO. Guides the design of ISO and ITU-T protocol standards.

OSPF: Open Shortest Path First. A routing protocol developed by the IETF for the Internet architecture. OSPF is based on a *link-state* algorithm, in which every node constructs a topography of the Internet and uses it to make forwarding decisions. Today known as Open Group.

overlay: A virtual (logical) network running on top of an existing physical network. Overlay nodes communicate with each other through tunnels rather than over physical links. Overlays are often used to deploy new network services since they do not require the cooperation of the existing network infrastructure.

packet: A data unit sent over a packet-switched network. Also see *frame* and *segment*.

packet switching: A general strategy for switching data through a network. Packet switching uses store-and-forward switching of discrete data units called packets, and implies *statistical multiplexing*.

participants: A generic term used to denote the processes, protocols, or hosts that are sending messages to each other.

PAWS: Protection against wrapped sequence numbers. Engineering transport protocol with a large enough sequence number space to protect against the numbers wrapping around on a network where packets can be delayed for a long period of time.

PCS: Personal Communication Services. New digital cellular phone system being deployed throughout the United States and Canada. Similar to GSM, which is being deployed throughout the rest of the world.

PDU: Protocol data unit. Another name for a packet or frame.

peer: A counterpart on another machine that a protocol module interoperates with to implement some communication service.

peer-to-peer networks: A general class of applications that integrate application logic (e.g., file storage) with routing. Popular examples include Napster and Gnutella. Research prototypes often use distributed hash tables.

PEM: Privacy Enhanced Mail. Extensions to Internet email that support privacy and integrity protection. See also *PGP*.

PGP: Pretty Good Privacy. A collection of public domain software that provides privacy and authentication capabilities using RSA and that uses a mesh of trust for public key distribution.

PHB: Per-hop behavior. Behavior of individual routers in the Differentiated Services architecture. AF and EF are two proposed PHBs.

physical-level protocol: The lowest layer of the OSI protocol stack. Its main function is to encode bits onto the signals that are propagated across the physical transmission media.

piconet: Wireless network spanning short distances (e.g., 10m). Used to connect office computers (laptops, printers, PDAs, workstations, etc.) without cables.

PIM: Protocol Independent Multicast. A multicast routing protocol that can be built on top of different unicast routing protocols.

Ping: A Unix utility used to test the RTT to various hosts over the Internet. Ping sends an ICMP ECHO_REQUEST message, and the remote host sends an ECHO_RESPONSE message back.

PIO: Programmed input/output. An approach to connecting hosts to I/O devices, in which the CPU reads data from and writes data to the I/O device. Also see *DMA*.

poison reverse: Used in conjunction with *split horizon*. A heuristic technique to avoid routing loops in distance-vector routing protocols.

port: A generic term usually used to mean the point at which a network user attaches to the network. On a switch, a port denotes the input or output on which packets are received and sent.

POTS: Plain old telephone service. Used to specify the existing phone service, in contrast to ISDN, ATM, or other technologies that the telephone companies offer now or may offer in the future.

PPP: Point-to-Point Protocol. Data link protocol typically used to connect computers over a dial-up line.

process: An abstraction provided by an operating system to enable different operations to take place concurrently. For example, each user application usually runs inside its own process, while various operating system functions take place in other processes.

promiscuous mode: A mode of operation for a network adaptor in which it receives all frames transmitted on the network, not just those addressed to it.

protocol: A specification of an interface between modules running on different machines, as well as the communication service that those modules implement. The term is also used to refer to an implementation of the module that meets this specification. To distinguish between these two uses, the interface is often called a *protocol specification*.

proxy: An agent sitting between a client and server that intercepts messages and provides some service. For example, a proxy can "stand in" for a server by responding to client requests, perhaps using data it has cached, without contacting the server.

pseudoheader: A subset of fields from the IP header that are passed up to transport protocols TCP and UDP for use in their checksum calculation. The pseudoheader contains source and destination IP addresses and IP datagram length, thus enabling detection of corruption of these fields or delivery of a packet to an incorrect address.

public key encryption: Any of several encryption algorithms (e.g., RSA) in which each participant has a private key (shared with no one else) and a public key (available to everyone). A secure message is sent to a user by encrypting the data with that user's public key; possession of the private key is required to decrypt the message, and so only the receiver can read it.

QoS: Quality of service. Packet delivery guarantees provided by a network architecture. Usually related to performance guarantees, such as bandwidth and delay. The Internet offers a best-effort delivery service, meaning that every effort is made to deliver a packet but delivery is not guaranteed.

RED: Random early detection. A queuing discipline for routers in which, when congestion is anticipated, packets are randomly dropped to alert the senders to slow down.

rendezvous point: A router used by PIM to allow receivers to learn about senders.

repeater: A device that propagates electrical signals from one Ethernet cable to another. There can be a maximum of two repeaters between any two hosts in an Ethernet. Repeaters forward signals, whereas *bridges* forward *frames*, and *routers* and *switches* forward *packets*.

REST: Representational State Transfer. An approach to building web services that uses HTTP as the generic application protocol.

reverse-path broadcast (RPB): A technique used to eliminate duplicate broadcast packets.

RFC: Request for Comments. Internet reports that contain, among other things, specifications for protocols like TCP and IP.

RIO: RED with In and Out. A packet drop policy based on RED, but involving two drop curves: one for packets that have been marked as being "in" profile and one for packets that have been marked "out" of profile. Designed to be used to implement differentiated services.

RIP: Routing Information Protocol. An intradomain routing protocol supplied with Berkeley Unix. Each router running RIP dynamically builds its forwarding table based on a *distance-vector* algorithm.

router: A network node connected to two or more networks that forwards packets from one network to another. Contrast with *bridge*, *repeater*, and *switch*.

routing: The process by which nodes exchange topological information to build correct forwarding tables. See *forwarding*, *link state*, and *distance vector*.

routing table: See *forwarding table*.

RPC: Remote Procedure Call. Synchronous request/reply transport protocol used in many client/server interactions.

RPR: Resilient Packet Ring. A type of ring network that is mostly used in metropolitan area networks. See *802.17*.

RSA: A public-key encryption algorithm named after its inventors: Rivest, Shamir, and Adleman.

RSVP: Resource Reservation Protocol. A protocol for reserving resources in the network. RSVP uses the concept of *soft state* in routers and puts responsibility for making reservations on receivers instead of on senders.

RTCP: Real-time Transport Control Protocol. Control protocol associated with RTP.

RTP: Real-time Transport Protocol. An end-to-end protocol used by multimedia applications that have real-time constraints.

RTT: Round-trip time. The time it takes for a bit of information to propagate from one end of a link or channel to the other and back again; in other words, double the latency of the channel.

SAN: Storage area network. A network that spans the components of a computer system (e.g., display, camera, disk). Includes interfaces like HiPPI and Fibre Channel. Contrast with *LAN*, *MAN*, and *WAN*.

schema: A specification of how to structure and interpret a set of data. Schema are defined for XML documents.

scrambling: The process of XORing a signal with a pseudorandom bitstream before transmission to cause enough signal transitions to allow clock recovery. Scrambling is used in SONET.

SDP: Session Description Protocol. An application layer protocol used to learn about the available audio/video channels. It reports the name and purpose of the session, start and end times for the session, the media types (e.g., audio, video) that comprise the session, and detailed information needed to receive the session (e.g., the multicast address, transport protocol, and port numbers to be used).

segment: A TCP packet. A segment contains a portion of the byte stream that is being sent by means of TCP.

semaphore: A variable used to support synchronization between processes. Typically a process *blocks* on a semaphore while it waits for some other process to signal the semaphore.

server: The provider of a service in a client/server distributed system.

SHA: Secure Hash Algorithm. A family of cryptographic hash algorithms.

signalling: At the physical level, denotes the transmission of a signal over some physical medium. In ATM, signalling refers to the process of establishing a virtual circuit.

silly window syndrome: A condition occurring in TCP that may arise if each time the receiver opens its receive window a small amount, the sender sends a small segment to fill the window. The result is many small segments and an inefficient use of bandwidth.

SIP: Session Initiation Protocol. An application layer protocol used in multimedia applications. It determines the correct device with which to communicate to reach a particular user, determines if the user is willing or able to take part in a particular communication, determines the choice of media and coding scheme to use, and establishes session parameters (e.g., port numbers).

sliding window: An algorithm that allows the sender to transmit multiple packets (up to the size of the window) before receiving an acknowledgment. As acknowledgments are returned for those packets in the window that were sent first, the window "slides" and more packets may be sent. The sliding window algorithm combines reliable delivery with a high throughput. See *ARQ*.

slow start: A congestion-avoidance algorithm for TCP that attempts to pace outgoing segments. For each ACK that is returned, two additional packets are sent, resulting in an exponential increase in the number of outstanding segments.

SMDS: Switched Multimegabit Data Service. A service supporting LAN-to-WAN connectivity, offered by some telephone companies.

SMTP: Simple Mail Transfer Protocol. The electronic mail protocol of the Internet. See *822*.

SNA: System Network Architecture. The proprietary network architecture of IBM.

SNMP: Simple Network Management Protocol. An Internet protocol that allows the monitoring of hosts, networks, and routers.

SOAP: A component of the web services framework for specifying and implementing application protocols.

socket: The abstraction provided by Unix that provides the application programming interface (API) to TCP/IP.

soft state: Connection-related information contained in a router that is cached for a limited period of time rather than being explicitly established (and requiring explicit teardown) through a connection setup.

SONET: Synchronous Optical Network. A clock-based framing standard for digital transmission over optical fiber. It defines how telephone companies transmit data over optical networks.

source routing: Routing decisions performed at the source before the packet is sent. The route consists of the list of nodes that the packet should traverse on the way to the destination.

source-specific multicast: A mode of multicast in which a group may have only a single sender.

sparse mode multicast: A mode used in PIM when relatively few hosts or routers need to receive multicast data for a certain group.

split horizon: A method of breaking routing loops in a distance-vector routing algorithm. When a node sends a routing update to its neighbors, it does not send those routes it learned from each neighbor back to that neighbor. Split horizon is used with *poison reverse*.

spread spectrum: Encoding technique that involves spreading a signal over a wider frequency than necessary, so as to minimize the impact of interference.

SSL: Secure Socket Layer. A protocol layer that runs over TCP to provide authentication and encryption of connections. Also known as Transport Layer Security (TLS).

statistical multiplexing: Demand-based multiplexing of multiple data sources over a shared link or channel.

stop-and-wait: A reliable transmission algorithm in which the sender transmits a packet and waits for an acknowledgment before sending the next packet. Compare with *sliding window* and *concurrent logical channels*. See also *ARQ*.

STS: Synchronous Transport Signal. The prefix for various rates of SONET transmission. For example, STS-1 refers to the SONET standard for 51.84-Mbps transmission.

subnetting: The use of a single IP network address to denote multiple physical networks. Routers within the subnetwork use a subnet mask to discover the physical network to which a packet should be forwarded. Subnetting effectively introduces a third level to the two-level hierarchical IP address.

SunRPC: Remote procedure call protocol developed by Sun Microsystems. SunRPC is used to support NFS. See also *ONC*.

switch: A network node that forwards packets from inputs to outputs based on header information in each packet. Differs from a *router* mainly in that it typically does not interconnect networks of different types.

switching fabric: The component of a switch that directs packets from their inputs to the correct outputs.

T1: A standard telephone carrier service equal to 24 ISDN circuits, or 1.544 Mbps. Also called DS1.

T3: A standard telephone carrier service equal to 24 T1 circuits, or 44.736 Mbps. Also called DS3.

TCP: Transmission Control Protocol. Connection-oriented transport protocol of the Internet architecture. TCP provides a reliable, byte-stream delivery service.

TDMA: Time Division Multiple Access. A form of multiplexing used in cellular wireless networks. Also the name of a particular wireless standard.

Telnet: Remote terminal protocol of the Internet architecture. Telnet allows you to interact with a remote system as if your terminal is directly connected to that machine.

throughput: The observed rate at which data is sent through a channel. The term is often used interchangeably with *bandwidth*.

TLS: Transport Layer Security. Security services that can be layered on top of a transport protocol like TCP. It is often used by HTTP to perform secure transactions on the World Wide Web. Derived from *SSL*.

token bucket: A way to characterize or police the bandwidth used by a flow. Conceptually, processes accumulate tokens over time, and they must spend a token to transmit a byte of data and then must stop sending when they have no tokens left. Thus, overall bandwidth is limited, with the accommodation of some burstiness.

token ring: A physical network technology in which hosts are connected in a ring. A token (bit pattern) circulates around the ring. A given node must possess the token before it is allowed to transmit. 802.5 and FDDI are examples of token ring networks.

TP4: OSI Transport Protocol Class 4. The most powerful OSI transport protocol. TP4 is the ISO equivalent of TCP.

transport protocol: An end-to-end protocol that enables processes on different hosts to communicate. TCP is the canonical example.

TTL: Time to live. Usually a measure of the number of hops (routers) an IP datagram can visit before it is discarded.

tunneling: Encapsulating a packet using a protocol that operates at the same layer as the packet. For example, multicast IP packets are encapsulated inside unicast IP packets to tunnel across the Internet to implement the MBone. Tunneling will also be used during the transition from IPv4 to IPv6.

two-dimensional parity: A parity scheme in which bytes are conceptually stacked as a matrix, and parity is calculated for both rows and columns.

Tymnet: An early network in which a *virtual circuit* abstraction was maintained across a set of routers.

UBR: Unspecified bit rate. The "no frills" service class in ATM, offering best-effort cell delivery. Contrast with *ABR, CBR,* and *VBR.*

UDP: User Datagram Protocol. Transport protocol of the Internet architecture that provides a connectionless datagram service to application-level processes.

UMTS: Universal Mobile Telecommunications System. Cellular wireless standard based on wideband CDMA that offers relatively high data rates.

unicast: Sending a packet to a single destination host. Contrast with *broadcast* and *multicast.*

URI: Uniform Resource Identifier. A generalization of the URL. Used for example, in conjunction with SIP to set up audio/visual sessions.

URL: Uniform Resource Locator. A text string used to identify the location of Internet resources. A typical URL looks like http://www.cisco.com. In this URL, http is the protocol to use to access the resource located on host www.cisco.com.

vat: Audioconferencing tool used on the Internet that runs over RTP.

VBR: Variable bit rate. One of the classes of service in ATM, intended for applications with bandwidth requirements that vary with time, such as compressed video. Contrast with *ABR, CBR,* and *UBR.*

VCI: Virtual circuit identifier. An identifier in the header of a packet that is used for virtual circuit switching. In the case of ATM, the VPI and VCI together identify the end-to-end connection.

vic: Unix-based videoconferencing tool that uses RTP.

virtual circuit: The abstraction provided by connection-oriented networks such as ATM. Messages must usually be exchanged between participants to establish a virtual circuit (and perhaps to allocate resources to the circuit) before data can be sent. Contrast with *datagram*.

virtual clock: A service model that allows the source to reserve resources on routers using a rate-based description of its needs. Virtual clock goes beyond the best-effort delivery service of the current Internet.

VPI: Virtual path identifier. An 8-bit or 12-bit field in the ATM header. VPI can be used to hide multiple virtual connections across a network inside a single virtual "path," thus decreasing the amount of connection state that the switches must maintain. See also *VCI*.

VPN: Virtual private network. A logical network overlaid on top of some existing network. For example, a company with sites around the world may build a virtual network on top of the Internet rather than lease lines between each site.

WAN: Wide area network. Any physical network technology that is capable of spanning long distances (e.g., cross-country). Compare with *SAN*, *LAN*, and *MAN*.

weighted fair queuing (WFQ): A variation of *fair queuing* in which each flow can be given a different proportion of the network capacity.

well-known port: A port number that is, by convention, dedicated for use by a particular server. For instance, the Domain Name Server receives messages at well-known UDP and TCP port 53 on every host.

WSDL: Web Services Description Language. A component of the web services framework for specifying and implementing application protocols.

WWW: World Wide Web. A hypermedia information service on the Internet.

X.25: The ITU packet-switching protocol standard.

X.400: The ITU electronic mail standard. The counterpart to SMTP in the Internet architecture.

X.500: The ITU directory services standard, which defines an attribute-based naming service.

X.509: An ITU standard for digital certificates.

XDR: External Data Representation. Sun Microsystems' standard for machine-independent data structures. Contrast with *ASN.1* and *NDR*.

XML: Extensible Markup Language. Defines a syntax for describing data that may be passed between Internet applications.

XSD: XML Schema Definition. A schema language for defining the format and interpretation of XML objects.

zone: A partition of the domain name hierarchy, corresponding to an administrative authority that is responsible for that portion of the hierarchy. Each zone must have at least two name servers to field DNS requests for the zone.

BIBLIOGRAPHY

[Bat68] K. E. Batcher. "Sorting Networks and Their Applications." *Proc. 1968 Spring AFIPS Joint Computer Conference*, Vol. 32, pp. 307–314, 1968.

[BBC+98] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. "An Architecture for Differentiated Services." *Request for Comments* 2475, December 1998.

[BCS94] R. Braden, D. Clark, and S. Shenker. "Integrated Services in the Internet Architecture: An Overview." *Request for Comments* 1633, September 1994.

[BDMS94] C. M. Bowman, P. B. Danzig, U. Manber, and M. F. Schwartz. "Scalable Internet Resource Discovery: Research Problems and Approaches." *Communications of the ACM*, 37(8):98–107, August 1994.

[Bel00] S. M. Bellovin. "ICMP Traceback Messages, March 2000." Work in Progress, Internet Draft draft-bellovin-itrace-00.txt.

[BF93] N. Borenstein and N. Freed. "MIME (Multipurpose Internet Mail Extensions) Part One: Mechanisms for Specifying and Describing the Format of Internet Message Bodies." *Request for Comments* 1521, September 1993.

[BG92] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, Englewood Cliffs, NJ, second edition, 1992.

[BG93] M. Bjorkman and P. Gunningberg. "Locking Effects in Multiprocessor Implementations of Protocols." In *Proceedings of the SIGCOMM '93 Symposium*, pp. 74–83, September 1993.

[Bha03] S. Bhattacharyya. "An Overview of Source-Specific Multicast." *Request for Comments* 3569, July 2003.

[Bla87] R. E. Blahut. *Principles and Practice of Information Theory*. Addison-Wesley, Reading, MA, 1987.

[BLFF96] T. Berners-Lee, R. Fielding, and H. Frystyk. "Hypertext Transfer Protocol—HTTP/1.0." *Request for Comments* 1945, May 1996.

[BLNS82] A. Birrell, R. Levin, R. Needham, and M. Schroeder. "Grapevine: An Exercise in Distributed Computing." *Communications of the ACM*, 25:250–273, April 1982.

[BM95] S. Bradner and A. Mankin, editors. *IPng: Internet Protocol Next Generation.* Addison-Wesley, Reading, MA, 1995.

[Boo95] P. Boorsook. "How Anarchy Works." *Wired*, 3(10):110–118, October 1995.

[BP95] L. S. Brakmo and L. L. Peterson. "TCP Vegas: End-to-End Congestion Avoidance on a Global Internet." *IEEE Journal of Selected Areas in Communication (JSAC)*, 13(8):1465–1480, October 1995.

[BPY90] M. Bowman, L. L. Peterson, and A. Yeatts. "Univers: An Attribute-based Name Server." *Software—Practice and Experience*, 20(4):403–424, April 1990.

[Bri95] T. Brisco. "DNS Support for Load Balancing." *Request for Comments* 1794, Rutgers University, New Brunswick, NJ, April 1995.

[BS88] L. Bic and A. C. Shaw. *The Logical Design of Operating Systems.* Prentice Hall, Englewood Cliffs, NJ, 1988.

[BS01] D. Barrett and R. Silverman. *SSH: The Secure Shell.* O'Reilly, 2001.

[Buf94] J. F. Koegel Buford. *Multimedia Systems.* ACM Press/Addison-Wesley, Reading, MA, 1994.

[BZ96] J. C. R. Bennett and H. Zhang. "Hierarchical Packet Fair Queueing Algorithms." *Proceedings of the SIGCOMM '96 Symposium*, pp. 143–156, August 1996.

[CCITT92a] Comité Consultif International de Telegraphique et Telephonique. "Open Systems Interconnection: Specification of Abstract Syntax Notation One (ASN.1)." CCIT Recommendation X.208, 1992.

[CCITT92b] Comité Consultif International de Telegraphique et Telephonique. "Open Systems Interconnection: Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1)." CCIT Recommendation X.209, 1992.

[CF98] D. D. Clark and W. Fang. "Explicit Allocation of Best-Effort Packet Delivery Service." *IEEE/ACM Transactions on Networking*, 6(4):362–373, August 1998.

[CFFD93] D. Cohen, G. Finn, R. Felderman, and A. DeSchon. "ATOMIC: A Low-Cost, Very-High-Speed, Local Communications Architecture." *Proceedings of the 1993 Conference on Parallel Processing*, August 1993.

[Cha93] A. L. Chapin. "The Billion Node Internet." In *Internet System Handbook*, D. C. Lynch and M. T. Rose, editors, Chapter 17, pp. 707–716. Addison-Wesley, Reading, MA, 1993.

[CJRS89] D. D. Clark, V. Jacobson, J. Romkey, and H. Salwen. "An Analysis of TCP Processing Overhead." *IEEE Communications*, 27(6):23–29, June 1989.

[Cla82] D. D. Clark. "Modularity and Efficiency in Protocol Implementation." *Request for Comments* 817, July 1982.

[Cla85] D. D. Clark. "The Structuring of Systems Using Upcalls." *Proceedings of the Tenth ACM Symposium on Operating Systems Principles*, pp. 171–180, December 1985.

[Cla88] D. D. Clark. "The Design Philosophy of the DARPA Internet Protocols." *Proceedings of the SIGCOMM '88 Symposium*, pp. 106–114, August 1988.

[Cla97] D. D. Clark. "Internet Cost Allocation and Pricing." *Internet Economics*, pp. 215–253. MIT Press, Cambridge, MA, 1997.

[CLNZ89] S. K. Chen, E. D. Lazowska, D. Notkin, and J. Zahorjan. "Performance Implications of Design Alternatives for Remote Procedure Call Stubs." *Proceedings of the Ninth International Conference on Distributed Computing Systems*, pp. 36–41, June 1989.

[Com00] D. E. Comer. *Internetworking with TCP/IP. Volume I: Principles, Protocols, and Architecture.* Prentice Hall, Upper Saddle River, NJ, fourth edition, 2000.

[CP89] D. E. Comer and L. L. Peterson. "Understanding Naming in Distributed Systems." *Distributed Computing*, 3(2):51–60, May 1989.

[CPVR97] J. Cohen, N. Phadnis, V. Valloppillil, and K. W. Ross. "Cache Array Routing Protocol v1.1." http://ds1.internic.net/internet-drafts/draft-vinod-carp-v1-01.txt, September 1997.

[Cro82] D. Crocker. "Standard for the Format of ARPA Internet Text Message." *Request for Comments* 822, August 1982.

[CS96] D. E. Comer and D. L. Stevens. *Internetworking with TCP/IP. Volume III: Client-Server Programming and Applications, BSD Socket Version.* Prentice Hall, Englewood Cliffs, NJ, second edition, 1996.

[CS97] D. E. Comer and D. L. Stevens. *Internetworking with TCP/IP. Volume III: Client-Server Programming and Applications, Windows Sockets Version.* Prentice Hall, Englewood Cliffs, NJ, 1997.

[CS00] D. E. Comer and D. L. Stevens. *Internetworking with TCP/IP. Volume III: Client-Server Programming and Applications, Linux/Posix Sockets Version.* Prentice Hall, Upper Saddle River, NJ, 2000.

[CV95] G. P. Chandranmenon and G. Varghese. "Trading Packet Headers for Packet Processing." *Proceedings of the SIGCOMM '95 Symposium*, pp. 162–173, October 1995.

[CZ85] D. R. Cheriton and W. Zwaenepoel. "Distributed Process Groups in the V Kernel." *ACM Transactions on Computer Systems*, 3(2):77–107, May 1985.

[Dan98] P. Danzig. "NetCache Architecture and Deployment." *3rd International WWW Caching Workshop*, June 1998.

[DBSP97] M. Degermark, A. Brodnik, S. Carlsson, and S. Pink. "Small Forwarding Tables for Fast Routing Lookups." *Proceedings of the SIGCOMM '97 Symposium*, pp. 3–14, October 1997.

[DCB+02] B. Davie, A. Charny, J. C. R. Bennett, K. Benson, J. Y. Le Boudec, W. Courtney, S. Davari, V. Firoiu, and D. Stiliadis. "An Expedited Forwarding PHB (Per-Hop Behavior)." *Request for Comments* 3246, March 2002.

[DEF+96] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C. Liu, and L. Wei. "The PIM Architecture for Wide-Area Multicast Routing." *ACM/IEEE Transactions on Networking*, April 1996.

[DH98] S. Deering and R. Hinden. "Internet Protocol, Version 6 (IPv6) Specification." *Request for Comments* 2460, December 1998.

[DP93] P. Druschel and L. L. Peterson. "Fbufs: A High-Bandwidth Cross-Domain Transfer Facility." *Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles*, pp. 189–202, December 1993.

[DPD94] P. Druschel, L. L. Peterson, and B. S. Davie. "Experience with a High-Speed Network Adaptor: A Software Perspective." *Proceedings of the SIGCOMM '94 Symposium*, pp. 2–13, August 1994.

[DR00] B. S. Davie and Y. Rekhter. *MPLS: Technology and Applications*. Morgan Kaufmann Publishers, San Francisco, CA, 2000.

[DY75] R. L. Drysdale and F. H. Young. "Improved Divide/Sort/Merge Sorting Networks." *SIAM Journal on Computing*, 4(3):264–270, September 1975.

[Eas05] D. Eastlake III. "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)." *Request for Comments* 4305, December 2005.

[EFH+98] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification." *Request for Comments* 2362, April 1998.

[Eis06] M. Eisler. "XDR: External Data Representation Standard." *Request for Comments* 4506, May 2006.

[EWL+94] A. Edwards, G. Watson, J. Lumley, D. Banks, C. Calamvokis, and C. Dalton. "User-Space Protocols Deliver High Performance to Applications on a Low-Cost Gb/s LAN." *Proceedings of the SIGCOMM '94 Symposium*, pp. 14–23, August 1994.

[FB96] N. Freed and N. Borenstein. "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies." *Request for Comments* 2045, November 1996.

[FGM+99] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, and T. Berners-Lee. "Hypertext Transfer Protocol—HTTP/1.1." *Request for Comments* 2616, June 1999.

[FHPW00] S. Floyd, M. Handley, J. Padhye, and J. Widmer. "Equation-based Congestion Control for Unicast Applications." *SIGCOMM*, pp. 43–56, Stockholm, Sweden, 2000.

[Fie00] R. T. Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. Ph.D. thesis, University of California, Irvine, Irvine, CA, 2000. http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm.

[Fin88] R. A. Finkel. *An Operating Systems Vade Mecum*. Prentice Hall, Englewood Cliffs, NJ, 1988.

[FL06] V. Fuller and T. Li. "Classless Inter-Domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan." *Request for Comments* 4632, August 2006.

[FM03] B. Fenner and D. Meyer. "Multicast Source Discovery Protocol (MSDP)." *Request for Comments* 3618, October 2003.

[Gar00] L. Garber. "Technology News: Denial-of-Service Attacks Rip the Internet." *Computer*, 33(4):12–17, April 2000.

[GG94] I. Gopal and R. Guerin. "Network Transparency: The plaNET Approach." *IEEE/ACM Transactions on Networking*, 2(3):226–239, June 1994.

[Gin99] D. Ginsburg. *ATM: Solutions for Enterprise Internetworking*. Addison-Wesley, Reading, MA, second edition, 1999.

[GVC96] P. Goyal, H. Vin, and H. Chen. "Start-Time Fair Queueing: A Scheduling Algorithm for Integrated Services Packet Switching Networks." *Proceedings of the SIGCOMM '96 Symposium*, pp. 157–168, August 1996.

[Har00] A. Harrison. "Cyber Assaults Hit Buy.com, eBay, CNN, and Amazon." *Computerworld*, February 2000.

[HC99] M. Handley and J. Crowcroft. "Internet Multicast Today." *The Internet Protocol Journal*, 2(4), December 1999.

[Hed88] C. Hedrick. "Routing Information Protocol." *Request for Comments* 1058, June 1988.

[HFPW03] M. Handley, S. Floyd, J. Padhye, and J. Widmer. "TCP Friendly Rate Control (TFRC): Protocol Specification." *Request for Comments* 3448, January 2003.

[HMPT89] N. C. Hutchinson, S. Mishra, L. L. Peterson, and V. T. Thomas. "Tools for Implementing Network Protocols." *Software—Practice and Experience*, 19(9):895–916, September 1989.

[HP91] N. Hutchinson and L. Peterson. "The *x*-kernel: An Architecture for Implementing Network Protocols." *IEEE Transactions on Software Engineering*, 17(1):64–76, January 1991.

[HP95] G. J. Holzmann and B. Pehrson. *The Early History of Data Networks*. IEEE Computer Society Press, Los Alamitos, CA, 1995.

[HP06] J. L. Hennessy and D. A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann, San Francisco, CA, fourth edition, 2006.

[Huf52] D. A. Huffman. "A Method for the Construction of Minimal-Redundancy Codes." *Proceedings of the IRE*, 40(9):1098–1101, September 1952.

[Jac88] V. Jacobson. "Congestion Avoidance and Control." *Proceedings of the SIGCOMM '88 Symposium*, pp. 314–329, August 1988.

[Jaf81] J. M. Jaffe. "Flow Control Power Is Nondecentralizable." *IEEE Transactions on Communications*, COM-29(9):1301–1306, September 1981.

[Jai89] R. Jain. "A Delay-based Approach for Congestion Avoidance in Interconnected Heterogeneous Computer Networks." *ACM Computer Communication Review*, 19(5):56–71, October 1989.

[Jai91] R. Jain. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. John Wiley & Sons, New York, NY, 1991.

[Jai94] R. Jain. *FDDI Handbook: High-Speed Networking Using Fiber and Other Media*. Addison-Wesley, Reading, MA, 1994.

[JBB92] V. Jacobson, R. Braden, and D. Borman. "TCP Extensions for High Performance." *Request for Comments* 1323, May 1992.

[Kau05] C. Kaufman. "Internet Key Exchange (IKEv2) Protocol." *Request for Comments* 4306, December 2005.

[KC88] H. Kanakia and D. R. Cheriton. "The VMP Network Adaptor Board (NAB): High-Performance Network Communication for Multiprocessors." *Proceedings of the SIGCOMM '88 Symposium*, pp. 175–187, August 1988.

[Ken05a] S. Kent. "IP Authentication Header." *Request for Comments* 4302, December 2005.

[Ken05b] S. Kent. "IP Encapsulating Security Payload (ESP)." *Request for Comments* 4303, December 2005.

[KHF06] E. Kohler, M. Handley, and S. Floyd. "Datagram Congestion Control Protocol (DCCP)." *Request for Comments* 4340, March 2006.

[KHR02] D. Katabi, M. Handley, and C. Rohrs. "Congestion Control for High Bandwidth-Delay Product Networks." *ACM SIGCOMM '02*, pp. 89–102, August 2002.

[Kle75] L. Kleinrock. *Queuing Systems. Volume 1: Theory*. John Wiley & Sons, New York, NY, 1975.

[Kle79] L. Kleinrock. "Power and Deterministic Rules of Thumb for Probabilistic Problems in Computer Communications." *Proceedings of the International Conference on Communications*, June 1979.

[Kle01] J. Klensin. "Simple Mail Transfer Protocol." *Request for Comments* 2821, April 2001.

[KM87] C. Kent and J. Mogul. "Fragmentation Considered Harmful." *Proceedings of the SIGCOMM '87 Symposium*, pp. 390–401, August 1987.

[KP91] P. Karn and C. Partridge. "Improving Round-Trip Time Estimates in Reliable Transport Protocols." *ACM Transactions on Computer Systems*, 9(4):364–373, November 1991.

[KPS02] C. Kaufman, R. Perlman, and M. Speciner. *Network Security: Private Communication in a Public World*. Prentice Hall, Englewood Cliffs, NJ, 2002.

[LAAJ00] C. Labovitz, A. Ahuja, A. Abose, and F. Jahanian. "Delayed Internet Routing Convergence." *Proceedings of the SIGCOMM 2000 Symposium*, Stockholm, Sweden, August 2000.

[Lin93] H.-A. (Paul) Lin. "Estimation of the Optimal Performance of ASN.1/BER Transfer Syntax." *Computer Communications Review*, 23(3):45–58, July 1993.

[LM97] D. Lin and R. Morris. "Dynamics of Random Early Detection." *Proceedings of the SIGCOMM '97 Symposium*, pp. 127–136, Cannes, France, October 1997.

[LMKQ89] S. J. Leffler, M. K. McKusick, M. J. Karels, and J. S. Quarterman. *The Design and Implementation of the 4.3BSD UNIX Operating System*. Addison-Wesley, Reading, MA, 1989.

[LPW+02a] S. Low, F. Paganini, J. Wang, S. Adlakha, and J. Doyle. "Dynamics of TCP/AQM and a Scalable Control." In *IEEE INFOCOM*, June 2002.

[LPW02b] S. Low, L. Peterson, and L. Wang. "Understanding TCP Vegas: A Duality Model." *Journal of the ACM*, 49(2):207–235, March 2002.

[LS98] T. V. Lakshman and D. Stiliadis. "High-Speed Policy-Based Packet Forwarding Using Efficient Multidimensional Range Matching." *Proceedings of the SIGCOMM '98 Symposium*, pp. 203–214, September 1998.

[LTWW94] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. "On the Self-Similar Nature of Ethernet Traffic." *IEEE/ACM Transactions on Networking*, 2:1–15, February 1994.

[Mas86] J. Mashey. "RISC, MIPS, and the Motion of Complexity." *UniForum 1986 Conference Proceedings*, pp. 116–124, 1986.

[MD90] J. Mogul and S. Deering. "Path MTU Discovery." *Request for Comments* 1191, November 1990.

[MD93] P. E. McKenney and K. F. Dove. "Efficient Demultiplexing of Incoming TCP Packets." *Proceedings of the SIGCOMM '92 Symposium*, pp. 269–280, August 1993.

[MD98] C. Madson and N. Doraswamy. "The ESP DES-CBC Cipher Algorithm with Explicit IV." *Request for Comments* 2405, November 1998.

[MG98a] C. Madson and R. Glenn. "The Use of HMAC-MD5-96 within ESP and AH." *Request for Comments* 2403, November 1998.

[MG98b] C. Madson and R. Glenn. "The Use of HMAC-SHA-1-96 within ESP and AH." *Request for Comments* 2404, November 1998.

[Min93] D. Minoli. *Enterprise Networking: Fractional T1 to SONET, Frame Relay to BISDN*. Artech House, Norwood, MA, 1993.

[MJ95] S. McCanne and V. Jacobson. "vic: A Flexible Framework for Packet Video." *ACM Multimedia '95*, pp. 511–522, 1995.

[Mor68] D. Morrison. "PATRICIA — A Practical Algorithm to Retrieve Information Coded in Alphanumeric." *Journal of the ACM*, 15(4):514–534, October 1968.

[Moy98] J. Moy. "OSPF Version 2." *Request for Comments* 2328, April 1998.

[MP85] J. Mogul and J. Postel. "Internet Standard Subnetting Procedure." *Request for Comments* 950, August 1985.

[MP94] J. Mogul and J. Postel. "RIP Version 2—Carrying Additional Information." *Request for Comments* 1723, November 1994.

[MPBO96] D. Mosberger, L. L. Peterson, P. G. Bridges, and S. O'Malley. "Analysis of Techniques to Improve Protocol Latency." *Proceedings of the SIGCOMM '96 Symposium*, pp. 73–84, August 1996.

[MPFL96] J. L. Mitchell, W. B. Pennebaker, C. E. Fogg, and D. J. LeGall. *MPEG Video: Compression Standard.* Chapman Hall, 1996.

[MPS99] K. McCloghrie, D. Perkins, and J. Schoenwaelder. "Structure of Management Information Version 2 (SMIv2)." *Request for Comments* 2578, April 1999.

[Mul90] S. Mullender. "Amoeba: A Distributed Operating System for the 1990s." *IEEE Computer*, 23(5):44–53, May 1990.

[MvOV96] A. J. Menezes, P. C. van Oorschot, and S. A. Vanstone. *Handbook of Applied Cryptography.* CRC Press, Boca Raton, FL, 1996.

[MVS01] D. Moore, G. Voelker, and S. Savage. "Inferring Internet Denial of Service Activity." *Proceedings of 2001 USENIX Security Symposium*, August 2001.

[Nel92] M. Nelson. *The Data Compression Book.* M&T Books, San Mateo, CA, 1992.

[Nol97] P. Noll. "MPEG Digital Audio Coding." *IEEE Signal Processing Magazine*, pp. 59–81, September 1997.

[NRC94] National Research Council, Computer Science and Telecommunications Board. *Realizing the Information Future: The Internet and Beyond.* National Academy Press, Washington, DC, 1994.

[NRC01] National Research Council. *Looking Over the Fence at Networks.* National Academy Press, Washington DC, 2001.

[NYKT94] E. M. Nahum, D. J. Yates, J. F. Kurose, and D. Towsley. "Performance Issues in Parallelized Network Protocols." *Proceedings of the First USENIX Symposium on Operating System Design and Implementation (OSDI)*, pp. 125–137, November 1994.

[OCD+88] J. K. Ousterhout, A. R. Cherenson, F. Douglis, M. N. Nelson, and B. B. Welch. "The Sprite Network Operating System." *IEEE Computer*, 21(2):23–36, February 1988.

[OP91] S. O'Malley and L. Peterson. "TCP Extensions Considered Harmful." *Request for Comments* 1263, October 1991.

[OPM94] S. W. O'Malley, T. A. Proebsting, and A. B. Montz. "Universal Stub Compiler." *Proceedings of the SIGCOMM '94 Symposium*, pp. 295–306, August 1994.

[OSF94] Open Software Foundation. *OSF DCE Application Environment Specification.* Prentice Hall, Englewood Cliffs, NJ, 1994.

[PACR02] L. Peterson, T. Anderson, D. Culler, and T. Roscoe. "A Blueprint for Introducing Disruptive Technology into the Internet." *Proceedings of HotNets-I*, October 2002.

[Pad85] M. A. Padlipsky. *The Elements of Networking Style and Other Essays and Animadversions on the Art of Intercomputer Networking.* Prentice Hall, Englewood Cliffs, NJ, 1985.

[Par94] C. Partridge. *Gigabit Networking.* Addison-Wesley, Reading, MA, 1994.

[Par98] C. Partridge, et al. "A 50 Gb/s IP Router." *IEEE/ACM Transactions on Networking*, 6(3):237–247, June 1998.

[PB61] W. W. Peterson and D. T. Brown. "Cyclic Codes for Error Detection." *Proceedings of the IRE*, Vol. 49, pp. 228–235, January 1961.

[Per00] R. Perlman. *Interconnections: Bridges, Routers, Switches, and Internetworking Protocols.* Addison-Wesley, Reading, MA, second edition, 2000.

[Pet88] L. L. Peterson. "The Profile Naming Service." *ACM Transactions on Computer Systems*, 6(4):341–364, November 1988.

[PF94] V. Paxson and S. Floyd. "Wide-Area Traffic: The Failure of Poisson Modeling." *Proceedings of the SIGCOMM '94 Symposium*, pp. 257–268, London, UK, August 1994.

[PFTK98] J. Padhye, V. Firoiu, D. Towsley, and J. Kursoe. "Modeling TCP Throughput: A Simple Model and Its Empirical Validation." *ACM SIGCOMM '98 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, pp. 303–314, Vancouver, British Columbia, 1998.

[PG94] A. Parekh and R. Gallagher. "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Multiple Node Case." *IEEE/ACM Transactions on Networking*, 2(2):137–150, April 1994.

[Pie84] J. Pierce. "Telephony—A Personal View." *IEEE Communications*, 22(5):116–120, May 1984.

[PL01] K. Park and H. Lee. "On the Effectiveness of Route-Based Packet Filtering for Distributed DoS Attack Prevention in Power-Law Internets." *Proceedings of ACM SIGCOMM'01*, August 2001.

[PM97] D. Perkins and E. McGinnis. *Understanding SNMP MIBS.* Prentice Hall, Upper Saddle River, NJ, 1997.

[Pos81] J. Postel. "Internet Protocol." *Request for Comments* 791, September 1981.

[Pos82] J. Postel. "Simple Mail Transfer Protocol." *Request for Comments* 821, August 1982.

[Pre02] R. Preshun. "Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)." *Request for Comments* 3416, December 2002.

[QPP02] X. Qie, R. Pang, and L. Peterson. "Defensive Programming: Using an Annotation Toolkit to Build DoS-Resistant Software." *Proceedings of OSDI '02*, December 2002.

[Ram93] K. K. Ramakrishnan. "Performance Considerations in Designing Network Interfaces." *IEEE Journal of Selected Areas in Communication (JSAC)*, 11(2):203–219, February 1993.

[RDR+97] Y. Rekhter, B. Davie, E. Rosen, G. Swallow, D. Farinacci, and D. Katz. "Tag Switching Architecture Overview." *Proceeedings of the IEEE*, 82(12):1973–1983, December 1997.

[Res01] P. Resnick. "Internet Message Format." *Request for Comments* 2822, April 2001.

[RF89] T. R. N. Rao and E. Fujiwara. *Error-Control Coding for Computer Systems*. Prentice Hall, Englewood Cliffs, NJ, 1989.

[RFB01] K. Ramakrishnan, S. Floyd, and D. Black. "The Addition of Explicit Congestion Notification (ECN) to IP." *Request for Comments* 3168, September 2001.

[RHE99] R. Rejaie, M. Handley, and D. Estrin. "RAP: An End-to-End Rate-Based Congestion Control Mechanism for Realtime Streams in the Internet." *INFOCOM* (3), pp. 1337–1345, 1999.

[Rit84] D. Ritchie. "A Stream Input-Output System." *AT&T Bell Laboratories Technical Journal*, 63(8):311–324, October 1984.

[RLH06] Y. Rekhter, T. Li, and S. Hares. "A Border Gateway Protocol 4 (BGP-4)." *Request for Comments* 4271, January 2006.

[Rob93] T. G. Robertazzi, editor. *Performance Evaluation of High-Speed Switching Fabrics and Networks: ATM, Broadband ISDN, and MAN Technology*. IEEE Press, Piscataway, NJ, 1993.

[Ros86] F. E. Ross. "FDDI—A Tutorial." *IEEE Communications*, 24(5):10–17, May 1986.

[ROY00] I. Rhee, V. Ozdemir, and Y. Yi. "TEAR: TCP Emulation at Receivers—Flow Control for Multimedia Streaming." NCSU Technical Report, 2000.

[RR06] E. Rosen and Y. Rekhter. "BGP/MPLS IP Virtual Private Networks (VPNS)." *Request for Comments* 4364, February 2006.

[RS01] R. Ramaswami and K. Sivarajan. *Optical Networks: A Practical Perspective*. Morgan Kaufmann Publishers, San Francisco, second edition, 2001.

[RS02] M. Rabinovich and O. Spatscheck. *Web Caching and Replication*. Addison Wesley, Reading, MA, 2002.

[Sal78] J. Saltzer. "Naming and Binding of Objects." *Lecture Notes on Computer Science*, 60:99–208, 1978.

[SB89] M. D. Schroeder and M. Burrows. "Performance of Firefly RPC." *Proceedings of the Twelfth ACM Symposium on Operating Systems Principles*, pp. 83–90, December 1989.

[SC03] H. Schulzrinne and S. Casner. "RTP Profile for Audio and Video Conferences with Minimal Control." *Request for Comments* 3551, July 2003.

[SCFJ03] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. "RTP: A Transport Protocol for Real-Time Applications." *Request for Comments* 3550, July 2003.

[Sch95] B. Schneier. *Applied Cryptography: Protocols, Algorithms, and Source Code in C.* John Wiley & Sons, New York, NY, 1995.

[SCH+99] S. Savage, A. Collins, F. Hoffman, J. Snell, and T. Anderson. "The End-to-End Effects of Internet Path Selection." *Proceedings of the ACM SIGCOMM Conference*, Cambridge, MA, September 1999.

[SCJ+02] H. Schulzrinne, G. Camarillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler. "SIP: Session Initiation Protocol." *Request for Comments* 3261, June 2002.

[Sha48] C. Shannon. "A Mathematical Theory of Communication." *Bell Systems Technical Journal*, 27:379–423, 623–656, 1948.

[Sho78] J. Shoch. "Inter-Network Naming, Addressing, and Routing." *Seventeenth IEEE Computer Society International Conference (COMPCON)*, pp. 72–79, September 1978.

[SHP91] J. Spragins, J. Hammond, and K. Pawlikowski. *Telecommunications: Protocols and Design.* Addison-Wesley, Reading, MA, 1991.

[SKPG01] T. Spalink, S. Karlin, L. Peterson, and Y. Gottlieb. "Building a Robust Software-based Router Using Network Processors." *Proceedings of the 18th ACM Symposium on Operating Systems Principles (SOSP)*, pp. 216–229, Chateau Lake Louise, Banff, Alberta, October 2001.

[SP99] O. Spatscheck and L. L. Peterson. "Defending Against Denial of Service Attacks in Scout." *Proceedings of OSDI '99*, February 1999.

[SPS+01] A. C. Snoeren, C. Partridge, L. A. Sanchez, C. E. Jones, F. Tchakountio, S. T. Kent, and W. T. Strayer. "Hash-based IP Traceback." *Proceedings of ACM SIGCOMM'01*, August 2001.

[SRC84] J. Saltzer, D. Reed, and D. Clark. "End-to-End Arguments in System Design." *ACM Transactions on Computer Systems*, 2(4):277–288, November 1984.

[SS98] D. Sisalem and H. Schulzrinne. "The Loss-Delay Based Adjustment Algorithm: A TCP-Friendly Adaptation Scheme." *Proceedings of NOSSDAV*, Cambridge, UK., 1998.

[SSZ98] I. Stoica, S. Shenker, and H. Zhang. "Core-Stateless Fair Queuing: A Scalable Architecture to Approximate Fair Bandwidth Allocations in High-Speed Networks." *ACM SIGCOMM '98*, August 1998.

[Sta00] W. Stallings. *Local and Metropolitan Area Networks*. Prentice Hall, Upper Saddle River, NJ, sixth edition, 2000.

[Sta03] W. Stallings. *Cryptography and Network Security*. Prentice Hall, Upper Saddle River, NJ, third edition, 2003.

[Sta07] W. Stallings. *Data and Computer Communications*. Prentice Hall, Upper Saddle River, NJ, eighth edition, 2007.

[Ste94a] P. A. Steenkiste. "A Systematic Approach to Host Interface Design for High Speed Networks." *IEEE Computer*, 27(3):47–57, March 1994.

[Ste94b] W. R. Stevens. *TCP/IP Illustrated. Volume 1: The Protocols*. Addison-Wesley, Reading, MA, 1994.

[SVSW98] V. Srinivasan, G. Vargese, S. Suri, and M. Waldvogel. "Fast Scalable Level Four Switching." *Proceedings of the SIGCOMM '98 Symposium*, pp. 191–202, September 1998.

[SW95] W. R. Stevens and G. R. Wright. *TCP/IP Illustrated. Volume 2: The Implementation*. Addison-Wesley, Reading, MA, 1995.

[SWKA00] S. Savage, D. Wetherall, A. Karlin, and T. Anderson. "Practical Network Support for IP Traceback." *Proceedings of the 2000 ACM SIGCOMM Conference*, August 2000.

[SZ97] I. Stoica and H. Zhang. "A Hierarchical Fair Service Curve Algorithm for Link-Sharing and Priority Services." *Proceedings of the SIGCOMM '97 Symposium*, pp. 29–262, October 1997.

[Tan01] A. S. Tanenbaum. *Modern Operating Systems*. Prentice Hall, Upper Saddle River, NJ, second edition, 2001.

[Tan03] A. S. Tanenbaum. *Computer Networks*. Prentice Hall, Upper Saddle River, NJ, fourth edition, 2003.

[Ter86] D. Terry. "Structure-free Name Management for Evolving Distributed Environments." *Sixth International Conference on Distributed Computing Systems*, pp. 502–508, May 1986.

[TL93] C. A. Thekkath and H. M. Levy. "Limits to Low-Latency Communication on High-Speed Networks." *ACM Transactions on Computer Systems*, 11(2):179–203, May 1993.

[TS93] C. B. S. Traw and J. M. Smith. "Hardware/Software Organization of a High-Performance ATM Host Interface." *IEEE Journal of Selected Areas in Communications (JSAC)*, 11(2):240–253, February 1993.

[VL87] G. Varghese and T. Lauck. "Hashed and Hierarchical Timing Wheels: Data Structures for the Efficient Implementation of a Timer Facility." *Proceedings of the Eleventh ACM Symposium on Operating Systems Principles*, pp. 25–38, November 1987.

[Wat81] R. Watson. "Identifiers (Naming) in Distributed Systems." In *Distributed System—Architecture and Implementation*, B. Lampson, M. Paul, and H. Siegert, editors, pp. 191–210. Springer-Verlag, New York, NY, 1981.

[WC91] Z. Wang and J. Crowcroft. "A New Congestion Control Scheme: Slow Start and Search (Tri-S)." *ACM Computer Communication Review*, 21(1):32–43, January 1991.

[WC92] Z. Wang and J. Crowcroft. "Eliminating Periodic Packet Losses in 4.3-Tahoe BSD TCP Congestion Control Algorithm." *ACM Computer Communication Review*, 22(2):9–16, April 1992.

[Wel84] T. Welch. "A Technique for High-Performance Data Compression." *IEEE Computer*, 17(6):8–19, June 1984.

[Wil00] B. Williamson. *Developing IP Multicast Networks. Volume I.* Cisco Press, Indianapolis, IN, 2000.

[WM87] R. W. Watson and S. A. Mamrak. "Gaining Efficiency in Transport Services by Appropriate Design and Implementation Choices." *ACM Transactions on Computer Systems*, 5(2):97–120, May 1987.

[WMB99] I. H. Witten, A. Moffat, and T. C. Bell. *Managing Gigabytes: Compressing and Indexing Documents and Images.* Morgan Kaufmann Publishers, San Francisco, CA, 1999.

[WPP02] L. Wang, V. Pai, and L. Peterson. "The Effectiveness of Request Redirection on CDN Robustness." *Proceedings of OSDI '02*, December 2002.

[WVTP97] M. Waldvogel, G. Varghese, J. Turner, and B. Plattner. "Scalable High-Speed Routing Lookups." *Proceedings of the SIGCOMM '97 Symposium*, pp. 25–36, October 1997.

[YHA87] Y.-S. Yeh, M. B. Hluchyj, and A. S. Acampora. "The Knockout Switch: A Simple, Modular Architecture for High-Performance Packet Switching." *IEEE Journal of Selected Areas in Communication (JSAC)*, 5(8):1274–1283, October 1987.

[ZDE+93] L. Zhang, S. Deering, D. Estrin, S. Schenker, and D. Zappala. "RSVP: A New Resource Reservation Protocol." *IEEE Network*, 7(9):8–18, September 1993.

[Zim80] H. Zimmerman. "OSI Reference Model—The ISO Model of Architecture for Open Systems Interconnection." *IEEE Transactions on Communications*, COM-28(4):425–432, April 1980.

[ZL77] J. Ziv and A. Lempel. "A Universal Algorithm for Sequential Data Compression." *IEEE Transactions on Information Theory*, 23(3):337–343, May 1977.

[ZL78] J. Ziv and A. Lempel. "Compression of Individual Sequences via Variable-Rate Coding." *IEEE Transactions on Information Theory*, 24(5):530–536, September 1978.

INDEX

B

C

E

F

G

H

I

J

K

L

M

N

O

P

S

T

U

V

W