robust cyclic redundancy check. We then moved on to the topic of multiple access protocols. We identified and studied three broad approaches for coordinating access to a broadcast channel: channel partitioning approaches (TDM, FDM), random access approaches (the ALOHA protocols and CSMA protocols), and taking-turns approaches (polling and token passing). We saw that a consequence of having multiple nodes share a single broadcast channel was the need to provide node addresses at the link layer. We learned that physical addresses were quite different from network-layer addresses and that, in the case of the Internet, a special protocol (ARP—the Address Resolution Protocol) is used to translate between these two forms of addressing. We then examined how nodes sharing a broadcast channel form a LAN and how multiple LANs can be connected together to form larger LANs—all *without* the intervention of network-layer routing to interconnect these local nodes.

We also covered a number of specific link-layer protocols in detail—Ethernet and PPP. We ended our study of the link layer by focusing on how MPLS networks provide link-layer services when they interconnect IP routers.We wrapped up this chapter (and indeed the first five chapters) by identifying the many protocols that are needed to fetch a simple web page. Having covered the link layer, *our journey down the protocol stack is now over*! Certainly, the physical layer lies below the data link layer, but the details of the physical layer are probably best left for another course (for example, in communication theory, rather than computer networking). We have, however, touched upon several aspects of the physical layer in this chapter and in Chapter 1 (our discussion of physical media in Section 1.2). We'll consider the physical layer again when we study wireless link characteristics in the next chapter.

Although our journey down the protocol stack is over, our study of computer networking is not yet at an end. In the following four chapters we cover wireless networking, multimedia networking, network security, and network management. These four topics do not fit conveniently into any one layer; indeed, each topic crosscuts many layers. Understanding these topics (billed as advanced topics in some networking texts) thus requires a firm foundation in all layers of the protocol stack—a foundation that our study of the data link layer has now completed!

Homework Problems and Questions

Chapter 5 Review Questions

SECTIONS 5.1–5.2

R1. Consider the transportation analogy in Section 5.1.1. If the passenger is analagous to a datagram, what is analogous to the link layer frame?

R2. If all the links in the Internet were to provide reliable delivery service, would the TCP reliable delivery service be redundant? Why or why not?

R3. What are some of the possible services that a link-layer protocol can offer to the network layer? Which of these link-layer services have corresponding services in IP? In TCP?

SECTION 5.3

R4. Suppose two nodes start to transmit at the same time a packet of length L over a broadcast channel of rate R. Denote the propagation delay between the two nodes as d_{prop}. Will there be a collision if $d_{prop} < L/R$? Why or why not?

R5. In Section 5.3, we listed four desirable characteristics of a broadcast channel. Which of these characteristics does slotted ALOHA have? Which of these characteristics does token passing have?

R6. Describe polling and token-passing protocols using the analogy of cocktail party interactions.

R7. Why would the token-ring protocol be inefficient if a LAN had a very large perimeter?

SECTION 5.4

R8. How big is the MAC address space? The IPv4 address space? The IPv6 address space?

R9. Suppose nodes A, B, and C each attach to the same broadcast LAN (through their adapters). If A sends thousands of IP datagrams to B with each encapsulating frame addressed to the MAC address of B, will C's adapter process these frames? If so, will C's adapter pass the IP datagrams in these frames to the network layer C? How would your answers change if A sends frames with the MAC broadcast address?

R10. Why is an ARP query sent within a broadcast frame? Why is an ARP response sent within a frame with a specific destination MAC address?

R11. For the network in Figure 5.19, the router has two ARP modules, each with its own ARP table. Is it possible that the same MAC address appears in both tables?

SECTION 5.5

R12. Compare the frame structures for 10BASE-T, 100BASE-T, and Gigabit Ethernet. How do they differ?

R13. Suppose a 10 Mbps adapter sends into a channel an infinite stream of 1s using Manchester encoding. The signal emerging from the adapter has how many transitions per second?

R14. In CSMA/CD, after the fifth collision, what is the probability that a node chooses $K = 4$? The result $K = 4$ corresponds to a delay of how many seconds on a 10 Mbps Ethernet?

SECTION 5.6

R15. Consider Figure 5.26. How many subnetworks are there, in the addressing sense of Section 4.4?

SECTION 5.7

R16. What is the r
supporting th

R17. Suppose that
a trunking pr
Justify your

Problems

P1. Suppose the
1101 and an
containing th
Your answer

P2. Show (give
parity check
of a double

P3. Suppose th
bytes consi
"Link Laye

P4. Consider th
numbers 0

a. the bina

b. the ASC

c. the ASC

Compute t

P5. Consider t
10101010

P6. Consider t

a. 10010

b. 10100

c. 01010

P7. In this pro
erator G (

a. Why c

b. Can th

P8. In Sectio
slotted A

a. Recal
is $Np($

SECTION 5.7

R16. What is the maximum number of VLANs that can be configured on a switch supporting the 802.1Q protocol? Why?

R17. Suppose that N switches supporting K VLAN groups are to be connected via a trunking protocol. How many ports are needed to connect the switches? Justify your answer.

Problems

P1. Suppose the information content of a packet is the bit pattern 1110 1011 1001 1101 and an even parity scheme is being used. What would the value of the field containing the parity bits be for the case of a two-dimensional parity scheme? Your answer should be such that a minimum-length checksum field is used.

P2. Show (give an example other than the one in Figure 5.6) that two-dimensional parity checks can correct and detect a single bit error. Show (give an example) of a double-bit error that can be detected but not corrected.

P3. Suppose the information portion of a packet (D in Figure 5.4) contains 10 bytes consisting of the 8-bit unsigned binary ASCII representation of string "Link Layer." Compute the Internet checksum for this data.

P4. Consider the previous problem, but instead of containing the binary of the numbers 0 through 9 suppose these 10 bytes contain

a. the binary representation of the numbers 1 through 10.

b. the ASCII representation of the letters A through J (uppercase).

c. the ASCII representation of the letters a through j (lowercase).

Compute the Internet checksum for this data.

P5. Consider the 7-bit generator, G=10011, and suppose that D has the value 1010101010. What is the value of R?

P6. Consider the previous problem, but suppose that D has the value

a. 1001000101.

b. 1010001111.

c. 0101010101.

P7. In this problem, we explore some of the properties of the CRC. For the generator G (=1001) given in Section 5.2.3, answer the following questions.

a. Why can it detect any single bit error in data D?

b. Can the above G detect any odd number of bit errors? Why?

P8. In Section 5.3, we provided an outline of the derivation of the efficiency of slotted ALOHA. In this problem we'll complete the derivation.

a. Recall that when there are N active nodes, the efficiency of slotted ALOHA is $Np(1-p)^{N-1}$. Find the value of p that maximizes this expression.

b. Using the value of p found in (a), find the efficiency of slotted ALOHA by letting N approach infinity. *Hint*: $(1 - 1/N)^N$ approaches $1/e$ as N approaches infinity.

P9. Show that the maximum efficiency of pure ALOHA is $1/(2e)$. *Note*: This problem is easy if you have completed the problem above!

P10. Consider two nodes, A and B, that use the slotted ALOHA protocol to contend for a channel. Suppose node A has more data to transmit than node B, and node A's retransmission probability p_A is greater than node B's retransmission probability, p_B.

a. Provide a formula for node A's average throughput. What is the total efficiency of the protocol with these two nodes?

b. If $p_A = 2p_B$, is node A's average throughput twice as large as that of node B? Why or why not? If not, how can you choose p_A and p_B to make that happen?

c. In general, suppose there are N nodes, among which node A has retransmission probability $2p$ and all other nodes have retransmission probability p. Provide expressions to compute the average throughputs of node A and of any other node.

P11. Suppose four active nodes—nodes A, B, C and D—are competing for access to a channel using slotted ALOHA. Assume each node has an infinite number of packets to send. Each node attempts to transmit in each slot with probability p. The first slot is numbered slot 1, the second slot is numbered slot 2, and so on.

a. What is the probability that node A succeeds for the first time in slot 5?

b. What is the probability that some node (either A, B, C or D) succeeds in slot 4?

c. What is the probability that the first success occurs in slot 3?

d. What is the efficiency of this four-node system?

P12. Graph the efficiency of slotted ALOHA and pure ALOHA as a function of p for the following values of N:

a. $N=15$.

b. $N=20$.

c. $N=30$.

P13. Consider a broadcast channel with N nodes and a transmission rate of R bps. Suppose the broadcast channel uses polling (with an additional polling node) for multiple access. Suppose the amount of time from when a node completes transmission until the subsequent node is permitted to transmit (that is, the polling delay) is d_{poll}. Suppose that within a polling round, a given node is allowed to transmit at most Q bits. What is the maximum throughput of the broadcast channel?

P14. Consider three LANs interconnected by two routers, as shown in Figure 5.38.

a. Assign IP addresses to all of the interfaces. For Subnet 1 use addresses of the form 192.168.1.xxx; for Subnet 2 uses addresses of the form 192.168.2.xxx; and for Subnet 3 use addresses of the form 192.168.3.xxx.

Figure 5.38 •

b. Assign

c. Consid
the AR
single-

d. Repeat
(and th

P15. Consider
with a swi

a. Consic
router
tainin
addres

b. Suppo
ARP c
query
the IP
source

c. Suppo
ARP c
MAC
entrie

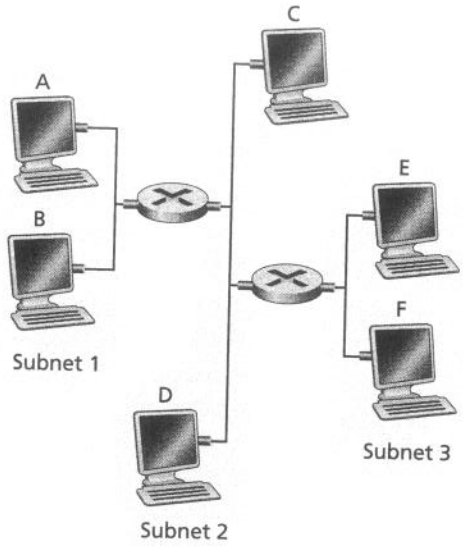

Figure 5.38 ♦ Three subnets, interconnected by routers

b. Assign MAC addresses to all of the adapters.

c. Consider sending an IP datagram from Host E to Host B. Suppose all of the ARP tables are up to date. Enumerate all the steps, as done for the single-router example in Section 5.4.2.

d. Repeat (c), now assuming that the ARP table in the sending host is empty (and the other tables are up to date).

P15. Consider Figure 5.38. Now we replace the router between subnets 1 and 2 with a switch S1, and label the router between subnets 2 and 3 as R1.

a. Consider sending an IP datagram from Host E to Host F. Will Host E ask router R1 to help forward the datagram? Why? In the Ethernet frame containing the IP datagram, what are the source and destination IP and MAC addresses?

b. Suppose E would like to send an IP datagram to B, and assume that E's ARP cache does not contain B's MAC address. Will B perform an ARP query to find B's MAC address? Why? In the Ethernet frame (containing the IP datagram destined to B) that is delivered to router R1, what are the source and destination IP and MAC addresses?

c. Suppose Host A would like to send an IP datagram to Host B, and neither A's ARP cache contains B's MAC address nor does B's ARP cache contain A's MAC address. Further suppose that the switch S1's forwarding table contains entries for Host B and router R1 only. Thus, A will broadcast an ARP request

message. What actions will switch S1 perform once it receives the ARP request message? Will router R1 also receive this ARP request message? If so, will R1 forward the message to Subnet 3? Once Host B receives this ARP request message, it will send back to Host A an ARP response message. But will it send an ARP query message to ask for A's MAC address? Why? What will switch S1 do once it receives an ARP response message from Host B?

P16. Consider the previous problem, but suppose now that the router between subnets 2 and 3 is replaced by a switch. Answer questions (a)–(c) in the previous problem in this new context.

P17. Recall that with the CSMA/CD protocol, the adapter waits $K \cdot 512$ bit times after a collision, where K is drawn randomly. For $K = 100$, how long does the adapter wait until returning to Step 2 for a 10 Mbps Ethernet? For a 100 Mbps Ethernet?

P18. Suppose nodes A and B are on the same 10 Mbps Ethernet bus, and the propagation delay between the two nodes is 325 bit times. Suppose node A begins transmitting a frame and, before it finishes, node B begins transmitting a frame. Can A finish transmitting before it detects that B has transmitted? Why or why not? If the answer is yes, then A incorrectly believes that its frame was successfully transmitted without a collision. *Hint*: Suppose at time $t = 0$ bit times, A begins transmitting a frame. In the worst case, A transmits a minimum-sized frame of 512 + 64 bit times. So A would finish transmitting the frame at $t = 512 + 64$ bit times. Thus, the answer is no, if B's signal reaches A before bit time $t = 512 + 64$ bits. In the worst case, when does B's signal reach A?

P19. Explain why a minimum frame size is required for Ethernet. For example, 10Base Ethernet imposes a minimum frame size constraint of 64 bytes. (If you have done the previous problem, you might have realized the reason). Now suppose that the distance between two ends of an Ethernet LAN is d. Can you derive a formula to find the minimum frame size needed for an Ethernet packet? Based on your reasoning, what is the minimum required packet size for an Ethernet that spans 2 kilometers?

P20. Suppose that you would like to increase the link speed of your Ethernet cable, how would this upgrade affect the minimum required packet sizes? If you upgrade your cable to a higher speed and realize that you cannot change packet size, what else can you do in maintain correct operation?

P21. Suppose nodes A and B are on the same 10 Mbps Ethernet bus, and the propagation delay between the two nodes is 245 bit times. Suppose A and B send frames at the same time, the frames collide, and then A and B choose different values of K in the CSMA/CD algorithm. Assuming no other nodes are active, can the retransmissions from A and B collide? For our purposes, it suffices to work out the following example. Suppose A and B begin transmission at $t = 0$ bit times. They both detect collisions at $t = 245$ bit times. They finish transmitting a jam signal at $t = 245 + 48 = 293$ bit times. Suppose $K_A = 0$ and

$K_B = 1$. At wh
A begin transr
returning to S
Does B refrai

P22. Consider a 1C
connected to
maximum dis
1000 bytes ar
ensure that a
node transmi
maximum di
the signal pr

P23. Suppose fou
Ethernet cab
300m, 400m
tocol is used
$2 \cdot 10^8$ m/se
a. What is t
required
b. If all frai

P24. In this probl
protocol. In
slots. Unlike
less than a fi
Suppose all
of the chanr
infinite nur
can detect a
- If, for a
contenc
probabi
possess
entire f
- If some
transm
mitting
conten

Note that
which las
random n
where x i

$K_B = 1$. At what time does B schedule its retransmission? At what time does A begin transmission? (*Note*: The nodes must wait for an idle channel after returning to Step 2—see protocol.) At what time does A's signal reach B? Does B refrain from transmitting at its scheduled time?

P22. Consider a 100 Mbps 100BASE-T Ethernet with all nodes directly connected to a hub. To have an efficiency of 0.50, what should be the maximum distance between a node and the hub? Assume a frame length of 1000 bytes and that there are no repeaters. Does this maximum distance also ensure that a transmitting node A will be able to detect whether any other node transmitted while A was transmitting? Why or why not? How does your maximum distance compare with the actual 100 Mbps standard? Assume that the signal propagation speed in 100BASE-T Ethernet is 1.8×10^8 m/sec.

P23. Suppose four nodes, A, B, C, and D, are all connected to a hub via 10Mbps Ethernet cables. The distances between the hub and these four nodes are 300m, 400m, 500m, and 700m, respectively. Recall that the CSMA/CD protocol is used for this Ethernet. Assume that the signal propagation speed is $2 \cdot 10^8$ m/sec.

a. What is the minimum required frame length? What is the maximum required frame length?

b. If all frames are 1500 bits long, find the efficiency of this Ethernet.

P24. In this problem you will derive the efficiency of a CSMA/CD-like multiple access protocol. In this protocol, time is slotted and all adapters are synchronized to the slots. Unlike slotted ALOHA, however, the length of a slot (in seconds) is much less than a frame time (the time to transmit a frame). Let S be the length of a slot. Suppose all frames are of constant length $L = kRS$, where R is the transmission rate of the channel and k is a large integer. Suppose there are N nodes, each with an infinite number of frames to send. We also assume that $d_{prop} < S$, so that all nodes can detect a collision before the end of a slot time. The protocol is as follows:

- If, for a given slot, no node has possession of the channel, all nodes contend for the channel; in particular, each node transmits in the slot with probability p. If exactly one node transmits in the slot, that node takes possession of the channel for the subsequent $k - 1$ slots and transmits its entire frame.
- If some node has possession of the channel, all other nodes refrain from transmitting until the node that possesses the channel has finished transmitting its frame. Once this node has transmitted its frame, all nodes contend for the channel.

Note that the channel alternates between two states: the productive state, which lasts exactly k slots, and the nonproductive state, which lasts for a random number of slots. Clearly, the channel efficiency is the ratio of $k/(k + x)$, where x is the expected number of consecutive unproductive slots.

a. For fixed N and p, determine the efficiency of this protocol.

b. For fixed N, determine the p that maximizes the efficiency.

c. Using the p (which is a function of N) found in (b), determine the efficiency as N approaches infinity.

d. Show that this efficiency approaches 1 as the frame length becomes large.

P25. Suppose two nodes, A and B, are attached to opposite ends of an 800 m cable, and that they each have one frame of 1,500 bits (including all headers and preambles) to send to each other. Both nodes attempt to transmit at time $t = 0$. Suppose there are four repeaters between A and B, each inserting a 20-bit delay. Assume the transmission rate is 100 Mbps, and CSMA/CD with backoff intervals of multiples of 512 bits is used. After the first collision, A draws $K = 0$ and B draws $K = 1$ in the exponential backoff protocol. Ignore the jam signal and the 96-bit time delay.

a. What is the one-way propagation delay (including repeater delays) between A and B in seconds? Assume that the signal propagation speed is $2 \cdot 10^8$ m/sec.

b. At what time (in seconds) is A's packet completely delivered at B?

c. Now suppose that only A has a packet to send and that the repeaters are replaced with switches. Suppose that each switch has a 20-bit processing delay in addition to a store-and-forward delay. At what time, in seconds, is A's packet delivered at B?

P26. In the Ethernet standard, a sender pauses 96 bit times between sending consecutive frames. This pausing time is referred to as inter-frame gap, and it is used to allow a receiving device to complete the processing of a received frame and to prepare for the reception of the next frame. Since the Ethernet standard was specified, there has been a tremendous improvement in technology including the speed of processors, memory, and the Ethernet rates. If the standard were to be rewritten, how would these improvements impact the inter-frame gap?

P27. Consider Figure 5.38 in problem P14. Provide MAC addresses and IP addresses for the interfaces at Host A, both routers, and Host F. Suppose Host A sends a datagram to Host F. Give the source and destination MAC addresses in the frame encapsulating this IP datagram as the frame is transmitted *(i)* from A to the left router, *(ii)* from the left router to the right router, *(iii)* from the right router to F. Also give the source and destination IP addresses in the IP datagram encapsulated within the frame at each of these points in time.

P28. Suppose now that the leftmost router in Figure 5.38 is replaced by a switch. Hosts A, B, C, and D and the right router are all star-connected into this switch. Give the source and destination MAC addresses in the frame encapsulating this IP datagram as the frame is transmitted *(i)* from A to the switch, *(ii)* from the switch to the right router, *(iii)* from the right router to F. Also give the source and destination IP addresses in the IP datagram encapsulated within the frame at each of these points in time.

P29. Consider Fi
mum total a
2 servers in
to any other

P30. Suppose the
All links ar

P31. Suppose th
are 100 Mb

P32. Let's consi
5.24. Supp
(iii) A send
initially en
these event
mitted fran

P33. In this prol
cations. O
link bandv
packet cor

a. Consid
source
entirel
time re
determ

b. Packet
unplea
(rough
50 (co

c. Calcu
$R = 6$

d. Comr

P34. Consider
router is
hosts an
the link
reread th

P35. Conside
R5 and
enginee
R6-R4-
R5-R4-
table in

P29. Consider Figure 5.26. Suppose that all links are 100 Mbps. What is the maximum total aggregate throughput that can be achieved among the 9 hosts and 2 servers in this network? You can assume that any host or serrver can send to any other host or server. Why?

P30. Suppose the three departmental switches in Figure 5.26 are replaced by hubs. All links are 100 Mbps. Now answer the questions posed in problem P29.

P31. Suppose that *all* the switches in Figure 5.26 are replaced by hubs. All links are 100 Mbps. Now answer the questions posed in problem P29.

P32. Let's consider the operation of a learning switch in the context of Figure 5.24. Suppose that *(i)* B sends a frame to E, *(ii)* E replies with a frame to B, *(iii)* A sends a frame to B, *(iv)* B replies with a frame to A. The switch table is initially empty. Show the state of the switch table before and after each of these events. For each of these events, identify the link(s) on which the transmitted frame will be forwarded, and briefly justify your answers.

P33. In this problem, we explore the use of small packets for Voice-over IP applications. One of the drawbacks of a small packet size is that a large fraction of link bandwidth is consumed by overhead bytes. To this end, suppose that the packet consists of P bytes and 5 bytes of header.

a. Consider sending a digitally encoded voice source directly. Suppose the source is encoded at a constant rate of 128 kbps. Assume each packet is entirely filled before the source sends the packet into the network. The time required to fill a packet is the **packetization delay**. In terms of L, determine the packetization delay in milliseconds.

b. Packetization delays greater than 20 msec can cause a noticeable and unpleasant echo. Determine the packetization delay for $L = 1{,}500$ bytes (roughly corresponding to a maximum-sized Ethernet packet) and for $L = 50$ (corresponding to an ATM packet).

c. Calculate the store-and-forward delay at a single switch for a link rate of $R = 622$ Mbps for $L = 1{,}500$ bytes, and for $L = 50$ bytes.

d. Comment on the advantages of using a small packet size.

P34. Consider the single switch VLAN in Figure 5.30, and assume an external router is connected to switch port 1. Assign IP addresses to the EE and CS hosts and router interface. Trace the steps taken at both the network layer and the link layer to transfer an IP datagram from an EE host to a CS host (*Hint:* reread the discussion of Figure 5.19 in the text).

P35. Consider the MPLS network shown in Figure 5.36, and suppose that routers R5 and R6 are now MPLS enabled. Suppose that we want to perform traffic engineering so that packets from R6 destined for A are switched to A via R6-R4-R3-R1, and packets from R5 destined for A are switched via R5-R4-R2-R1. Show the MPLS tables in R5 and R6, as well as the modified table in R4, that would make this possible.

P36. Consider again the same scenario as in the previous problem, but suppose that packets from R6 destined for D are switched via R6-R4-R3, while packets from R5 destined to D are switched via R4-R2-R1-R3. Show the MPLS tables in all routers that would make this possible.

P37. In this problem, you will put together much of what you have learned about Internet protocols. Suppose you walk into a room, connect to Ethernet, and want to download a web page. What are all the protocol steps that take place, starting from powering on your PC to getting the web page? Assume there is nothing in our DNS or browser caches when you power on your PC. (*Hint:* the steps include the use of Ethernet, DHCP, ARP, DNS, TCP, and HTTP protocols.) Explicitly indicate in your steps how you obtain the IP and MAC addresses of a gateway router.

Discussion Questions

You are encouraged to surf the Web in seeking answers to the following questions.

D1. Roughly, what is the current price range of a 10/100 Mbps adapter? Of a Gigabit Ethernet adapter? How do these prices compare with a 56 kbps dial-up modem or with an ADSL modem?

D2. Switches are often priced by number of interfaces (also called *ports* in LAN jargon). Roughly, what is the current per-interface price range for a switch consisting of only 100 Mbps interfaces?

D3. Many of the functions of an adapter can be performed in software that runs on the node's CPU. What are the advantages and disadvantages of moving this functionality from the adapter to the node?

D4. Search the Web to find the protocol numbers used in an Ethernet frame for an IP datagram and for an ARP packet.

D5. Read references [Xiao 2000, Huang 2002, and RFC 3346] on traffic engineering using MPLS. List a set of goals for traffic engineering. Which of these goals can only be met with MPLS, and which of these goals are met by using existing (non-MPLS) protocols? In the latter case, what advantages does MPLS offer?

Wireshark Lab

At the companion Web site for this textbook, http://www.awl.com/kurose-ross, you'll find a Wireshark lab that examines the operation of the IEEE 802.3 protocol and the Wireshark frame format.

A second Wireshark lab examines packet traces taken in a home network scenario similar to that in Figure 5.37.

AN INTERV

Simon S.

Simon S. Lam is f
at the University c
with the ARPA Ne
worked on satelli
group that invent
secure sockets la
won the 2004 A
are in design an
es. He received
MS and PhD fro
of Engineering i

Why did you deci

When I arrived at
control theory. Th
impressed by him.
sible thesis topic. I
(later called Packe
we did was to intr
col. Shortly there
ility problem and

You were active i
student days at U
the Internet woul

The atmosphere
industry and aca
to provide access
tists could use th
the Packet Radio
was building thr
Vint Cerf and B
Back then,
as logical rather
power of person
things in perspe
cards for batch p
typically housed
dent, I had only