

Networking

Atoms vs. Bits -- Recap

Bits

Movies, TV, newspaper, music, college, socializing, flirting, books, magazines, people-watching, art -- it's easy to think of "bit" activities

Atoms

Food, shelter (non-optional)

SUVs, monster-homes, huge lawns, and other atom-intensive objects (optional)

How much of atom intensive stuff do you really need to be happy? On the other hand, it's easy to imagine unlimited consumption of bit activities.

Primitive Wealth Legacy

Have a pyramid built in your honor -- atoms -- Khufu (aka Cheops)

Modern Wealth Legacy

Create a scholarship fund to improve the education of a many poor people -- bits -- Bill Gates

Business Week

By happy coincidence, there's a business week article that echoes the atoms vs. bits discussion. The original article was entitled "What's the Matter?" and was originally published April 1999 -- thanks to Angus Dawson for the heads up). <http://www.business2.com/services/10principles.html>

Oil

It's been widely reported that the "new" bit economy is reducing the magnification of the recent increase in the price of oil.

Terms: Bytes, K, Meg ...

"1k" = ~1000 bytes = ~1000 characters

1 page text = 2k

small image = 10k

"1 Meg" = ~1000k (aka "1 MB")

Typical computer program 1-3 Megs

1 hour stereo CD = 600 Megs (not compressed)

"1 Gig" = 1000 Meg

Hard disk: \$200 / 6 Gig drive = \$0.05 / Meg (was 4 gig in 1999)

RAM memory: \$1.5 / Meg (was \$3 in 1999)

Network Background

"Bandwidth" -- bits per second (Kbps, Mbps) used by EE types.

Bytes per second (K per second, MB per second) used by software types.

Figure roughly 10 bits per byte when converting to allow for overhead and so the math is easy.

"Latency" -- the initial delay for anything to get to the destination. The time required for a "dialog" type protocol between two computers can be largely bound by latency rather than bandwidth, even though bandwidth is the more frequently quoted figure.

Serial vs. Parallel. Paradoxically, the fastest future directions appear to be serial based. (Cheaper cables, remove the difficulty of keeping the parallel signals in synch.)

Copper can go up to around 1000 Mbps. Wire quality, shielding, and length all matter. (There are 1000 Mbps ethernet standards for both copper and fiber, but copper versions can't span very far -- say 100 Meters)

Fiber optic can go to around (in the lab) 1000 Gbps. Or 10^{12} bps -- on the order of a million simultaneous channels of video. More expensive than copper, but getting cheaper. A little harder to install, has fancier connectors, can't go around sharp bends. Current fiber standards run around 5 Gbps.

Current Technologies

Typical phone line / modem type application = 53 Kbps \approx 5K / sec

Cable Modem -- 500kbps -- 50k / sec

Digital Subscriber Loop (DSL) 500kbps -- 50k / sec (The local phone company monopolies have not been a good thing for development in this area -- fortunately the Cable companies scared them.)

10T Ethernet connection. Uses a 10Mbps signaling. Ideal case 2 computers in a room. May approach 900k/sec. Around 500k/sec more likely, with reductions the more traffic (from other computers) the ethernet must carry at the same time. Latency can approach 1 thousandth of second -- 1 ms.

USB 12Mbps -- 1.2MB/sec (10x faster USB 2.0 is rumored)

Standard SCSI -- 5MB/ sec

Fast SCSI -- 10MB/sec (a current 9 GB hard drive can just about saturate this on sustained reads. Hard drives have latencies on the order of 10ms.)

Ultra Wide SCSI -- 40MB/sec

Ultra2 Wide SCSI -- 80MB/sec

Firewire aka 1394 400Mbps -- 50MB/sec (800 Mbps and 1.6Gbps are rumored)

Fast Ethernet 100T (5-10MB/sec) and 1000T (50-100MB/sec) are increasingly in use. Capacity depends very much on how much "sharing" is going on.

Data Rates

Uncompressed CD sound -- 160k/sec (16 bit samples, 44kHz, stereo)

Compressed sound (MP3 or other schemes)-- 15k / sec -- depends very much on quality

Uncompressed video is on the order of 20MB/sec -- video compresses very, very well

Compressed VHS video = ~1.5M bps = 150k bytes / second. (this is how those little satellite TV dishes work).

Central Switched Circuit

"Connection Oriented"

Old

Old Phone Company Way

Phone calls are not necessarily handled this way any more, but the switched circuit concept will come up later.

Call

The "call" operation sets up an end-to-end connection.

In the old days, this was a wire connection all the way between the phones.

Good

Latency and bandwidth available don't vary during the call -- important for some applications --e.g. live 2-way video stream.

If the call is long, the set up cost is relatively cheap

Bad

The switching logic is centralized into the "center"

If the call is short, the setup is relatively costly

The call is using those resources, even during periods of silence -- overall inefficient.

LANs

Local Area Network -- computers in physical proximity. 200 computers within 1/4 mile radius -- one floor of a building scale.

Problem: want to get bytes from one computer to another

Could have centralized phone system type setup - switching logic gets messy

Packets

A relatively small bundle of bytes (like a small file) — perhaps 1k in size.

Packet = "header" information for routing + "body" of bytes to be transmitted. The body can contain, text, voice, video....the transport system is independent of the purpose.

Communication is broken up into packets. A typical text email message takes up a few packets. Audio takes a stream of thousands of packets. Each packet starts at the source and is communicated hop by hop across the routers to the destination. The packets are re-assembled at the destination.

Packets Analysis

"bursty"

The overall flow of packets may be "bursty" and a little unreliable.

Efficient use of resources

Fine for text, web pages

Bad for real time telephone, video

Q: how do you do audio or video? Related Q: why is there a 1 second delay when you start up your CD player? One solution: buffer up enough data to smooth out the bursts. This is how RealAudio works.

vs. Connection

Compare to "connection oriented" such as a local phone call. Large startup cost to build a dedicated "connection".. Inefficient use of the silent periods in the conversation. Efficient if we are talking 100% of the time. Latency and bandwidth are guaranteed.

Routing

Route all packets independently (TCP/IP) vs. route all packets along per-connection assigned path (ATM). Disadvantage that each packet has its own routing computation. However, tends to use the overall resources efficiently since bursty traffic can interleave.

Point: Different from and typically more efficient than a telephone connection

Point: Efficient = **cheap**

QOS

Quality of Service (QOS) extensions to packet communication -- try to guarantee latency/bandwidth for some of the traffic. Complicates routing. TCP/IP does not do a good job currently with QOS.

Broadband Ethernet LAN technology --

One of the Cleverest things ever

(editorial) An admirably elegant way to cheaply connect a group of computers -- original design credited to Bob Metcalfe at Xerox PARC.

Original design was for a radio based network -- read the story in Accidental Empires, by Robert X Cringely.

All computers on the LAN have an address

There is one wire, and all the computers are connected to it -- they share the wire -- only one computer is supposed to transmit at once.

Much simpler than the "central hub" shape of a telephone type network.

Simpler = cheaper.

All computers listen to the wire at all times
 Ethernet uses a single cable as the shared medium
 Could use the same sort of strategy with any shared medium -- e.g. a radio frequency
 Sender divides their message into small "packets" of, say, around 1000 bytes. Every packet begins with the address of the recipient.

Ethernet Protocol

The sender listens, waiting for a period of silence.
 When there's a period of silence, sender broadcasts the packet
 Everybody listens all the time -- ignoring packets not for them
 Sometimes two transmission overlap, and so the packets "collide" and get garbled. This can happen because (speed of light and all that) both senders can start sending before each other's signals have propagated down the wire to each other. The network card can usually detect this collisions and so know to re-transmit.
 Have a "wait/re-transmit" protocol to re-send packets -- wait a random amount of time -- one of Metcalfe's breakthrough ideas. Also part of what makes ethernet a little unpredictable in terms of performance, especially latency.

Ethernet Conclusions

Shared

there's just the one wire and everybody uses it

Distributed and cooperative

no central control -- surprisingly decentralized

Insecure

not too hard to listen and pick up packets not intended for you

Random

Performance degrades as more computers use the shared medium
 A little unpredictable how long it will take to get a message through
 Compare to the telephone system "connection" oriented service -- connection services is more predictable but typically less efficient overall.

Works really well overall

The uncertainty of ethernet has proved to not be a big deal at all -- it works so well that it's irresistible. IBM's competing "Token Ring" technology is less uncertain -- the rumor goes that that's why "uncertainty" was getting so much play as an important feature.

WAN

Wide Area Network

Connect separate LANs

Routers

To understand the Internet, want to define the idea of a router -- a device that connects two LANs

The router must be able to communicate on both the LANs -- it will translate between the two

Listen to packets on both sides

Retransmit packets on one LAN from the other when necessary -- eg. when the sender is on one LAN, and the recipient is on the other

NOTE: the IP address of your local router is one of the things you frequently need to configure on your computer in order to be on the Internet.

TCP/IP Standard -- Internet Foundation

Like any WAN technology, TCP/IP is an overall naming and communication standard which can be used by routers and hosts. The difference is that TCP/IP is a public standard and that sits on top of heterogeneous LAN technology. Any vendor is free to implement it in their equipment, and that equipment will automatically interoperate with all the other TCP/IP equipment (in theory).

Internet Protocol (IP) -- fundamental addresses, packets, and routing
Transmission Control Protocol (TCP) -- controls the overall flow of a collection of packets as part of a larger message. End-to-end service.

Introduce standards for all the different LANs, computers, and routers so they may interconnect.

TCP/IP is a free, public standard developed cooperatively by groups of volunteer engineering committees.

Contrast to "profit motive" type standards

Point: if a device supports the public standard, then it should automatically "fit in" with all the other devices TCP/IP structure.

TCP/IP Structure

1. Namespace standard
2. Packet format standard
3. Routing standard