Myriad of different LAN technologies co-existing in a WAN. For example:

- Fast Ethernet (100 Mbps)
- Gigabit Ethernet (1000 Mbps); 10 and 100 GigE

  → Purdue CS backbone: 10 Gbps

  → AT&T (tier-1 provider)?
- WLAN (11, 54, 300 Mbps)
- 4G cellular (100 Mbps mobile, 1 Gbps stationary)

  → today: pre-4G

  → ITU-R
- WiMAX (tens of Mbps up to 1 Gbps)

  → wider area: several miles

- modem/DSL (cable and dial-up)

  $\rightarrow$ 12 Mbps (down)/2 Mbps (up), 50/10, higher
- FDDI (Fiber Distributed Data Interface), 100 Mbps
- ATM, etc.

Note: WAN is a collection of LANs

→ ultimately: everything happens at LANs

Each LAN, in general, speaks a different language

→ message format (syntactic)

→ behavioral (semantic)

Internetworking handles this problem by translating everything to IP (Internet Protocol)

→ technical definition of **I**nternet

→ collection of interconnected LANs speaking IP

But:

→ IP injects overhead

→ packet forwarding example: switches choose not to speak IP

Thus:

→ IP (layer 3) is not necessary

→ large systems of layer 2 (LAN) switches

→ size limit: heterogeneity and management headache

→ today: L2 + L3

⟶ common attitude: avoid IP if possible

⟶ most devices speak IP in case needed

⟶ management aspect: naming

IP provides naming flexibility:

→ configurable addresses

→ IP: v4 32-bit, v6 128-bit

→ note: 48-bit LAN addresses are hardwired and unique address per NIC

Common practice: assign similar addresses to network devices belonging to an organization

→ ARIN in the U.S.

→ blocks of contiguous addresses: makes routing easier

→ e.g.: Purdue 128.10.$*$.$\star$, 128.210.$*$.$\star$

→ router bottleneck: table look-up speed

Naming: IP is not enough

Communicating entities are *processes* residing on hosts A and B running some operating system (Linux/UNIX, Windows, etc.)

→ IP only specifies host

→ more accurately: one of the NICs attached to a host

→ host with multiple NICs: multiple IP addresses

→ multi-homed

Thus an name/address must also identify which process a message is destined for on a host

→ port number abstraction

→ 16-bit

→ destination address: tuple (host IP, port number)

Network Performance

In computer networks, speed is at a premium

→ if slow, typically not used in practice

→ e.g., cryptographic protocols are not turned on at routers

→ emphasis on lightweight network core

→ push heavyweight stuff toward the edge (i.e., host)

→ end-to-end paradigm

→ guided Internet design and evolution

Three yardsticks of performance:

- bandwidth: bps (bits-per-second)

  → throughput: includes software processing overhead

  → e.g., 802.11b WLAN: nominal bandwidth 11 Mbps, throughput around 6 Mbps

- latency: msec (millisecond)

  → signal propagation speed

  → approximately: speed of light

  → delay: includes software processing overhead and waiting time at routers (queueing)

  → delay at high speed routers: very small (μsec)

  → delay at WLAN AP: up to hundreds of millisecond

- jitter: delay variation

  → not good for real-time content (video, audio, voice)

Bandwidth vs. throughput:

bandwidth—maximum data transmission rate achievable at the hardware level; determined by signalling rate of physical link and NIC

throughput—maximum data transmission rate achievable at the software level; overhead of network protocols inside/outside OS is accounted for

reliable throughput—maximum reliable data transmission rate achievable at the software level; effect of recovery from transmission errors and packet loss accounted for

$\longrightarrow$ networks tend to be "leaky"

Meaning of "high-speed" networks:

- signal propagation speed is bounded by SOL (speed-of-light)

  → ∼300K km/s or ∼186K miles/s

  → optical fiber, copper: nearly same

- Ex.: latency: Purdue to West Coast

  → around 2000 miles: ∼10 msec (= 2000/186000)

  → lower bound

- Ex.: geostationary satellites: ∼22.2K miles

  → latency: ∼120 msec

  → end-to-end (one-way): ∼240 msec

  → round-trip (two-way): ∼480 msec

  → typically: ∼500 msec

- thus: a single bit cannot go faster

  → can only increase "bandwidth"

  → analogous to widening highway, i.e., more lanes

  → simulatenous transmission of multiple bits

  → hence "broadband" is a more accurate term

- interpretation: "high-speed" ⇔ "many lanes"

  → what does it buy?

  → completion time of large files faster

  → in this sense, "higher" speed

Some units:

Tbps, Gbps, Mbps, Kbps:

10^{12}, 10^9, 10^6, 10^3 bits per second; indicates data transmission rate; influenced by clock rate (MHz/GHz) of signaling hardware

$\longrightarrow$ communication rate: factors of 1000

$\longrightarrow$ data size: 1 KB means 1024 bytes

Common bit rates:

- 10, 100, 1000 Mbps (1 Gbps); 10 and 100 Gbps Ethernet
- 11 Mbps (and 5, 2, 1 Mbps) for 802.11b WLAN

  → 5, 2 and 1 Mbps: fallback rates
- 54 Mbps (and 48, 36, 24, 18, 12, 9, 6 Mbps) for 802.11g/a WLAN
- 300 Mbps for 802.11n WLAN
- 64 Kbps (toll quality digitized voice)

  → landline
- ∼10 Kbps (cell phone quality voice)
- 1.544 Mbps (T1), 44.736 Mbps (T3)
- popular backbone speeds: 1, 10, 100 Gbps

Purdue's backbone network: ITaP

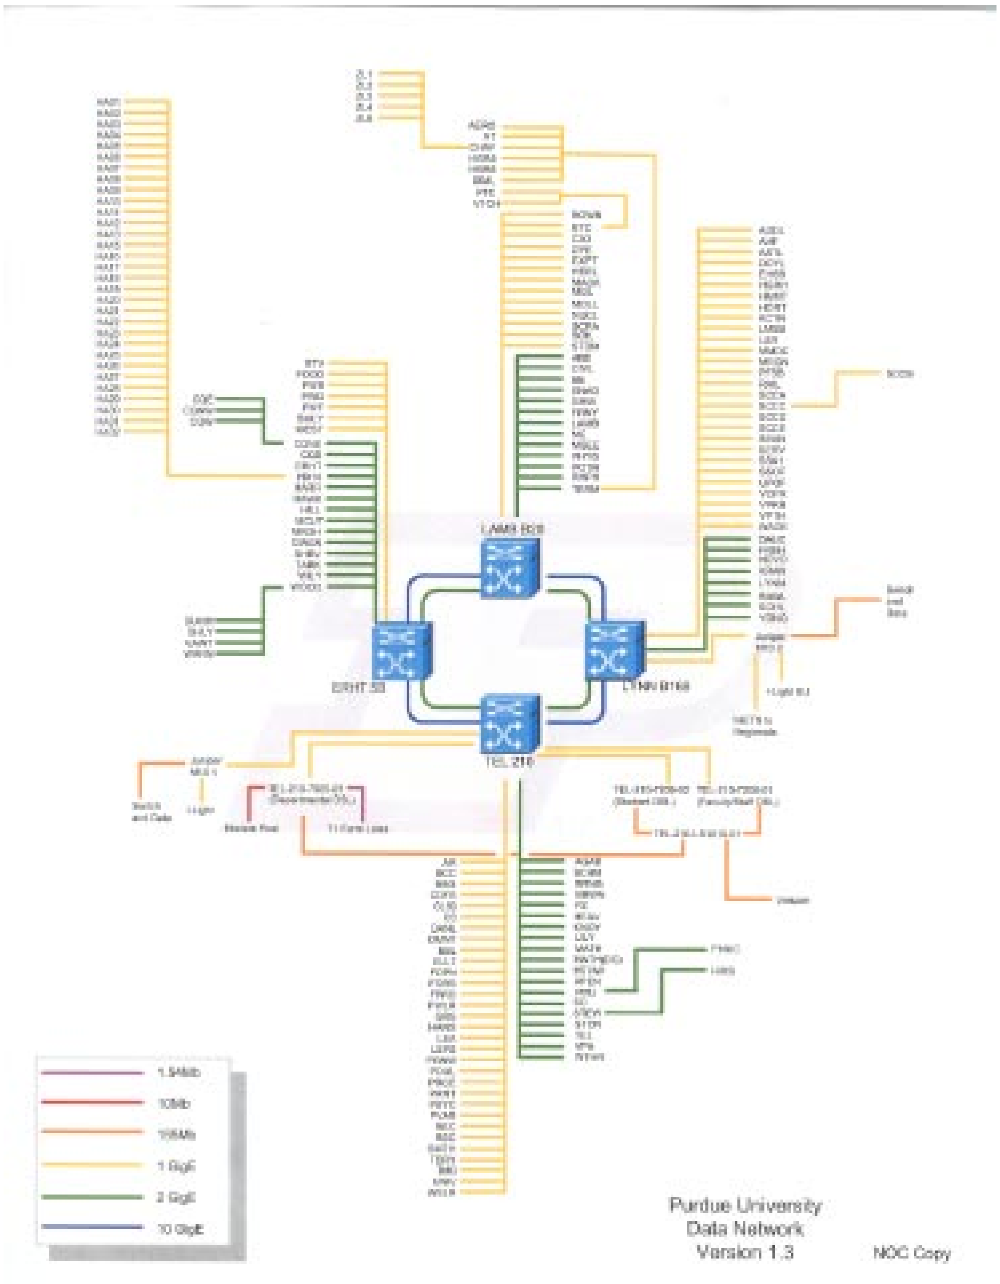

Level3 backbone network: www.level3.com

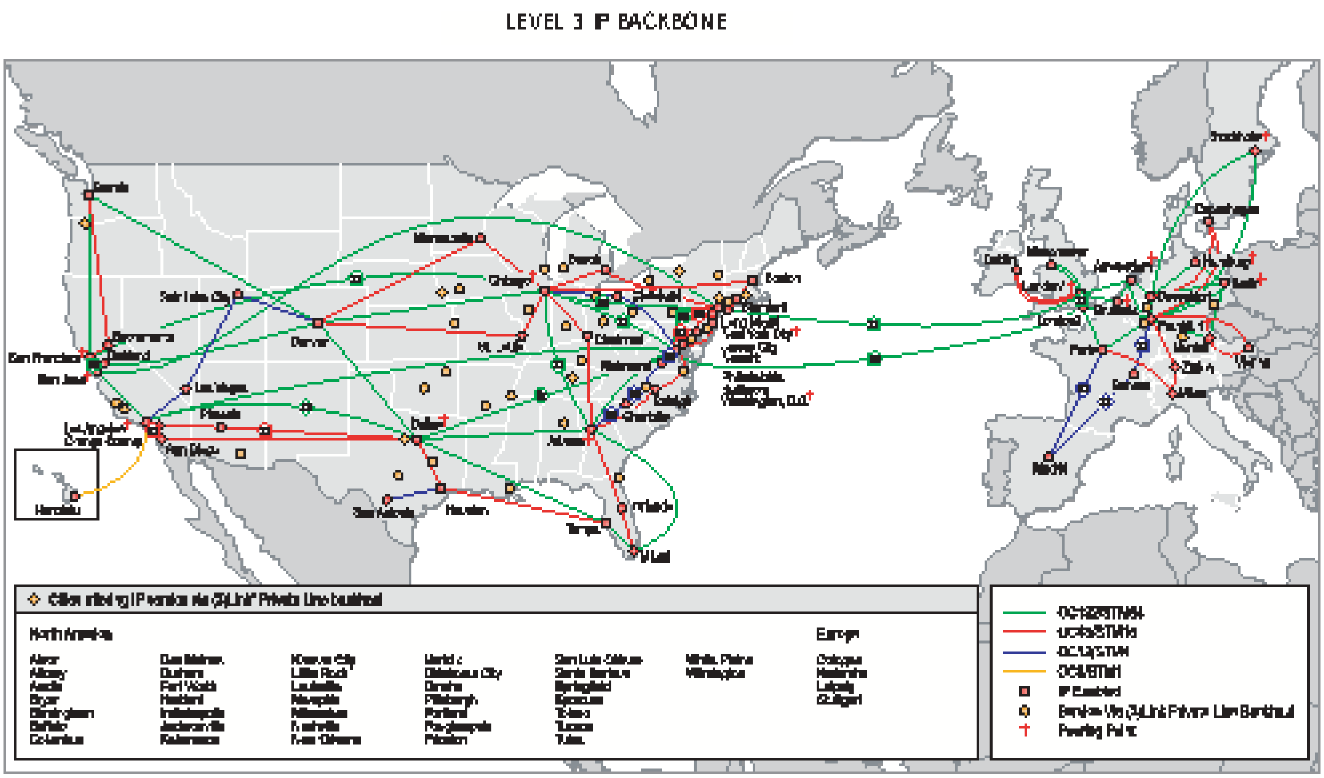

→ 10 Gbps backbone (green): same speed as Purdue

→ Purdue CS Dept.: 10 Gbps backbone

What is traveling on the wires?

Mixture of:

bulk data (data, image, video, audio files), voice, streaming video/audio, real-time interactive data (e.g., games), etc.

→ around 90% of Internet traffic is TCP file traffic

→ HTTP web and P2P

Multimedia (video/audio) streaming: on the rise

→ a minority but share is increasing

VoIP:

→ telcos: VoIP is dominant

→ non-telcos: increasing but relative volume is small

Internet traffic is "bursty": MPEG compressed real-time video

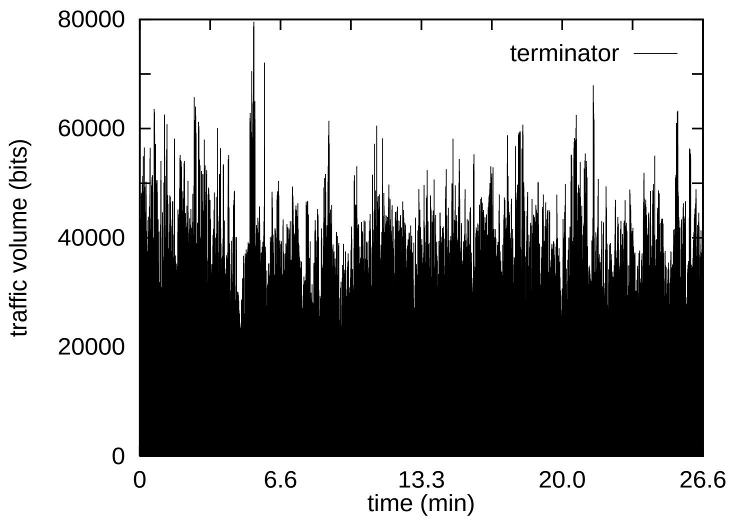

Reason:

- video compression

  → utilize inter-frame compression

- across scenes, significant scenary changes

  → e.g., action movies

- within scenes, few changes

Burstiness is not good for networks: why?

Main source of traffic burstiness:

→ skewed file size

- 90/10 rule
- 90% of files are small, 10% are very large

  → "many mice, few elephants"

  → the few elephants make up 90% of total traffic

  → same for disk space

Real-world is inherently skewed . . .

How to make sense of all this?

Study of networks has three aspects:

- architecture

  → system design, real-world manifestation

- algorithms

  → how do the components work

- implementation

  → how are the algorithms implemented

Key concern: performance, i.e., speed

→ security, fault-tolerance, etc. are important but second to performance

→ prevailing network system culture