What is traveling on the wires?

Mixed data:

bulk data, audio/voice, video/image, real-time interactive data, etc.

$\longrightarrow$ $> 85\%$ of Internet traffic is bulk TCP traffic

$\longrightarrow$ due to Web and http

$\longrightarrow$ barriers to streaming traffic implosion

$\longrightarrow$ technical and other

Tilting toward *multimedia* data; i.e., traffic with QoS requirements including real-time constraints.

Internet traffic is bursty:

$\longrightarrow$ multimedia: MPEG compressed video

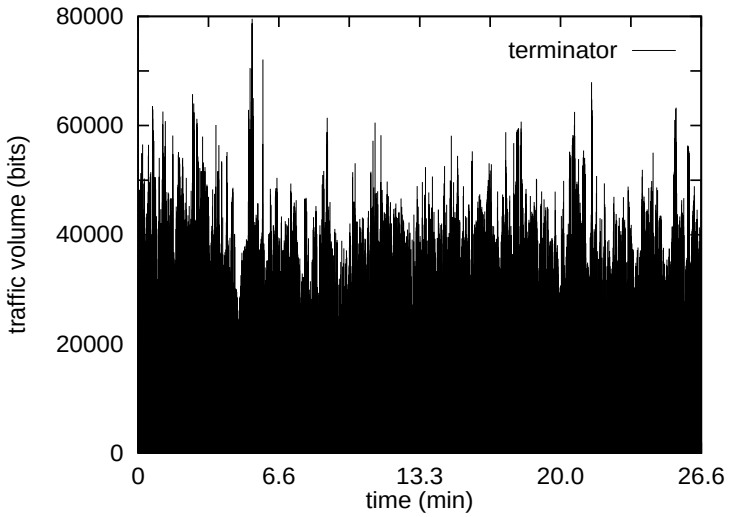

Why?

- pattern of scene changes in movies

  $\rightarrow$ within a scene few changes
- across scenes, significant "scenary" changes

  $\rightarrow$ "director's eye"
- video compression

  $\rightarrow$ utilize inter-frame compression

$\longrightarrow$ file sizes on file servers

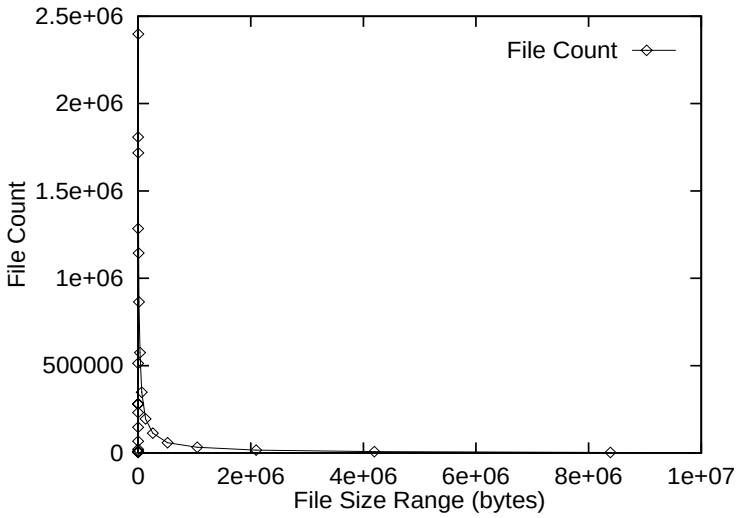

Why?

- bulk data: 80/20 rule-of-thumb
- majority of files are small, a few very large

  $\rightarrow$ disproportionate contribution to total traffic

  $\rightarrow$ "elephants and mice"

Usage pattern in the real-world: uneven or "unfair"

Given mixed payload:

Data networks capable of carrying diverse payload on the same network is a recent phenomenon.

Even today, much of voice traffic (telephony) is carried on an entirely separate communication network vis-à-vis data traffic, operating under different internetworking principles from the latter.

$\longrightarrow$ time-division multiplexing (TDM) for telephony

$\longrightarrow$ packet switching for data networks

How is time—viewed as a resource—shared?

Time-division multiplexing:

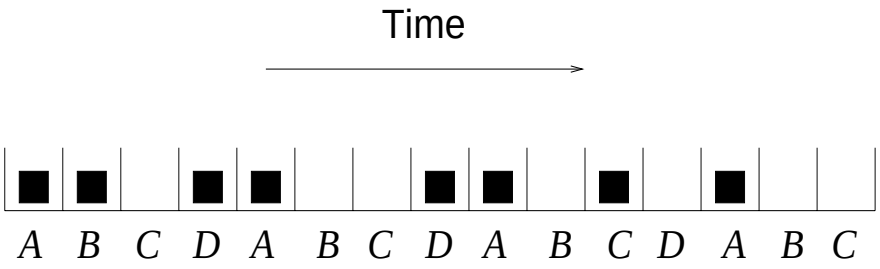

Packet switching:

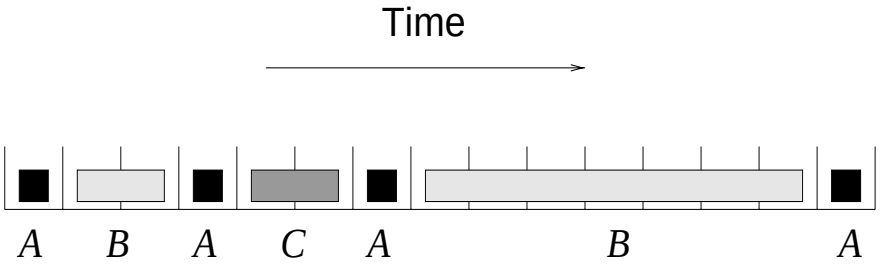

How is "real estate" shared?

Circuit switching: Virtual channel is established and followed during the lifetime of an end-to-end connection.

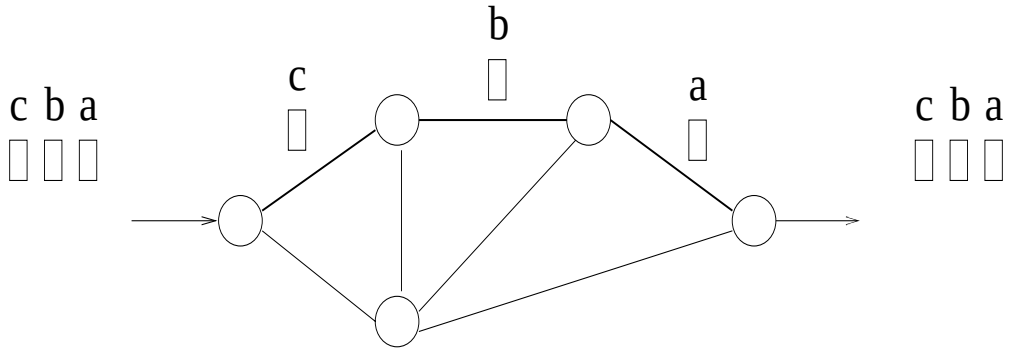

$\longrightarrow$ static route

$\longrightarrow$ in-order delivery

$\longrightarrow$ small routing table

Telephone networks (and ATM networks).

Packet switching: Every packet belonging to an end-to-end conversation is an independent entity; may take a different route from other packets in the same connection.

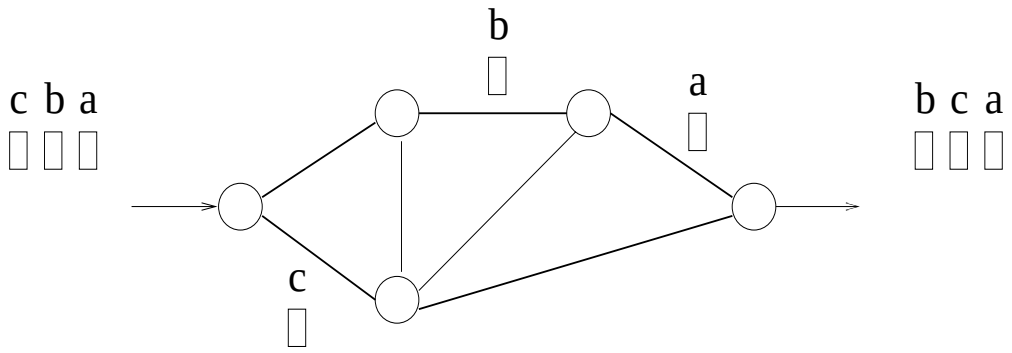

$\longrightarrow$ dynamic route

$\longrightarrow$ out-of-order delivery

$\longrightarrow$ larger route table

Trade-off between processing overhead and route goodness

Trend: convergence to packet-switched technology

- $\longrightarrow$ layer 2 switching in the backbone: VC
- $\longrightarrow$ move away from IP due to overhead
- $\longrightarrow$ IP critical at peering points

Yet another drawback of packet switching:

- $\longrightarrow$ "bully phenomenon"
- $\longrightarrow$ video: 24 frames-per-second (f/s)
- $\longrightarrow$ voice: 8000 samples-per-second (s/s)
- $\longrightarrow$ what to do?

Asynchronous transfer mode (*ATM*) :

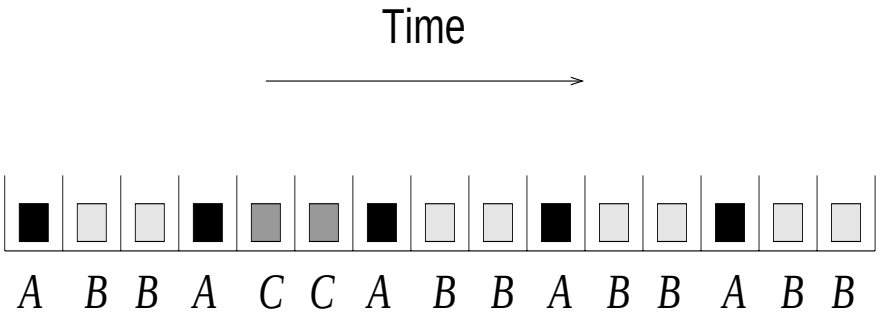

$\longrightarrow$ 53 byte packet or *cell*.

Synergy of all forms of data, audio, video, bulk, etc. One unified network with "integrated" services.

Addresses bully problem but ...

$\longrightarrow$ significant overhead (48 + 5)

$\longrightarrow$ why 48 bytes?

$\longrightarrow$ performs its own routing (VC based)

$\longrightarrow$ function duplication

$\longrightarrow$ very complex (overloaded with features)

$\longrightarrow$ feature $\neq$ "how to"

Much has migrated to new layer 2 switching technology

$\longrightarrow$ MPLS (multiprotocol label switching)

$\longrightarrow$ ATM community reincarnated as MPLS ...

$\longrightarrow$ after shrinkage

$\longrightarrow$ supporting role to IP

In the meantime, at routers receiving mixed payload ...

Try to avoid packet loss, but no loss comes at a cost:

- fast memory (buffer) is not cheap
- management overhead: ASIC vs. software vs. hybrid
- packets have to wait in line for their turn

  → queueing delay

  → who gets preference?

Depends on scheduling.

- FIFO (first-in-first-out)
- priority queue
- round robin + weighted fair queue

  → use TOS field of IPv4 to encode priority
- reservation

  → software-based "line leasing"

Is adding more and more buffer space a good solution?

$\longrightarrow$ no: related to "elephants and mice"

$\longrightarrow$ bandwidth is preferred (and, presently, cheaper)

When is it outright bad?

$\longrightarrow$ real-time multimedia payload