



WESTFÄLISCHE
WILHELMS-UNIVERSITÄT
MÜNSTER

Computer Networks, Winter Term 2009/2010

Requirements

- What do you (does an app) expect from a network?
- How can resources be shared?
 - (Recall OS)
- What is and how do you measure network performance?



- Connectivity
 - Last week: Internetworking, layered architecture, addressing, forwarding
- Cost-effective resource sharing
- Support for common services
- Performance
- (Security, separate lecture)
- (Scalability; entire class)

- Notice
 - Requirements depend on stakeholders' interests


- Connectivity
- **Cost-effective resource sharing**
- Support for common services
- Performance

- How to share the network among hosts?
- How to share the same link between several hosts?
- Solution: **Multiplexing**
 - System resource is shared among multiple users
 - Possibly with differing levels of Quality of Service (QoS)

- On-demand time-division (avoids idle time)
- Schedule link on a per-packet basis
- Packets from different sources interleaved on link
- Buffer packets that are contending for the link
- Buffer (queue) overflow is called congestion

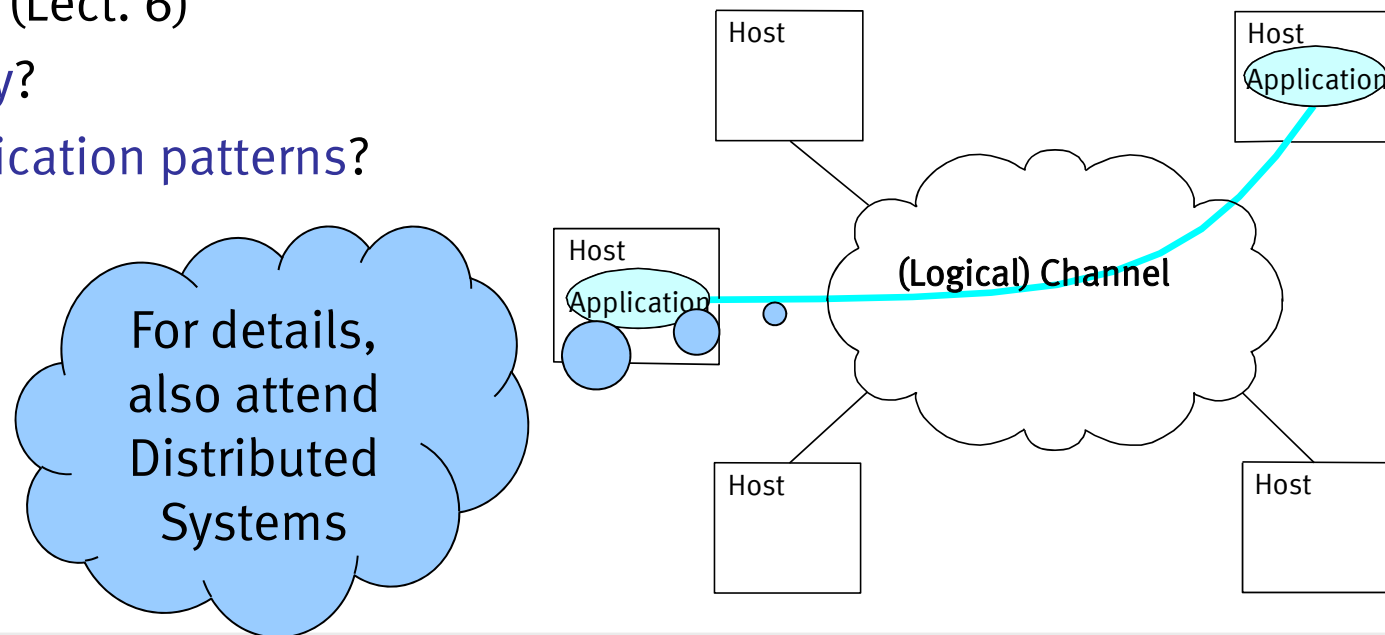
Compare with
Round Robin
scheduling in
OS

- Java applet
- http://media.pearsoncmg.com/aw/aw_kurose_network_2/applets/queuing/queuing.html



- Connectivity
- Cost-effective resource sharing
- **Support for common services**
- Performance

- Turn host-to-host connectivity into **process-to-process** communication
- Fill gap between what applications expect and what the underlying technology provides
 - Security? (Lect. 6)
 - Reliability?
 - Communication patterns?



- Bit-level errors (electrical interference)
 - Checksums (Lect. 10)
- Packet-level errors (congestion)
 - Timeouts and ACKs (e.g., TCP, see Lect. 5)
- Link and node failures
 - Redundancy + adaptive routing (Lect. 8, 9)

- Messages may be delayed
- Messages may be delivered out-of-order

Request/Reply

- WWW: Browser requests media objects from Web server
- Distributed file systems: client requests file access from server, e.g., FTP, NFS

Stream-Based

- Video: sequence of frames
- E.g., images with 320x240 pixels, 24b per pixel
 - $(320 \cdot 240 \cdot 24b) / 8 = 230,400B$
 - 30 fps = 6912 KBps \approx 55 Mbps
- Video applications
 - Video on-demand
 - Video conferencing

Dissemination

- Publish identifiable, authentic content
- Content propagates, multiple cached copies
- See Van Jacobsen: A New Way to look at Networking
 - <http://video.google.com/videoplay?docid=-6972678839686672840>

- Connectivity
- Cost-effective resource sharing
- Support for common services
- **Performance**

- B = byte; b = bit
- Prefixes
 - Frequent convention
 - MB = 2^{20} bytes (Intuition: Memory)
 - Mbps = 10^6 bits per second (Intuition: Clock)
 - Consequence: 1 MB takes more than 1 sec. on 8 Mbps link
 - Standards
 - Système international d'unités (SI)
 - Based on powers of 10
 - “Usual” prefixes: kilo/k (10^3), mega/M (10^6), milli/m (10^{-3}), ...
 - Binary units defined by the IEC (International Electrotechnical Commission)
 - Based on powers of 2
 - New prefixes: kibi/Ki (1024), mebi/Mi (1024^2), gibi/Gi (1024^3), ...
 - Consequence: 1 MiB distinguished from 1 MB

- Bandwidth

- Amount of data transmitted per time unit, e.g., “10 Mbps”
 - Analogy: Diameter of pipeline
- Link versus end-to-end

- Latency (delay)

- Time to send message from point A to point B, e.g., “24 ms”
- One-way versus round-trip time (RTT)
- Components
 - Latency = Propagation + Transmit + Queuing Delay
 - Propagation = Distance / c*Speed of Light
 - Transmit = Packet Size / Bandwidth
- Java applet (without queuing delay, 1 kB = 8000 b)
http://media.pearsoncmg.com/aw/aw_kurose_network_2/applets/transmission/delay.html
- Measurements: ping, traceroute

- Ping

- Measures end-to-end RTT

- Orders of magnitude

- ping www.uni-muenster.de (ca. 1 ms)
- ping www.google.de (ca. 10 ms)
- ping www.nwu.ac.za (ca. 210 ms)

- Ping World Demo (Java)

- <http://www-iepm.slac.stanford.edu/tools/pingworld/>

- PingER project: Distributed measurements

- <http://www-iepm.slac.stanford.edu/pinger/>

- Traceroute

- Measures RTT for each hop
- Background, e.g., at <http://navigators.com/traceroute.html>
- Core idea (details later)
 - Each IP packet carries time-to-live (TTL) field
 - Decrement by 1 by each router
 - If TTL=0 then router sends error message back to sender
 - Repeatedly send IP packets to target host
 - Start with TTL=1
 - Increment TTL by 1 when error message arrives
- `traceroute www.google.de`
- `traceroute-nanog -n -A -M -Q www.google.de`

- Traceroutes from different locations

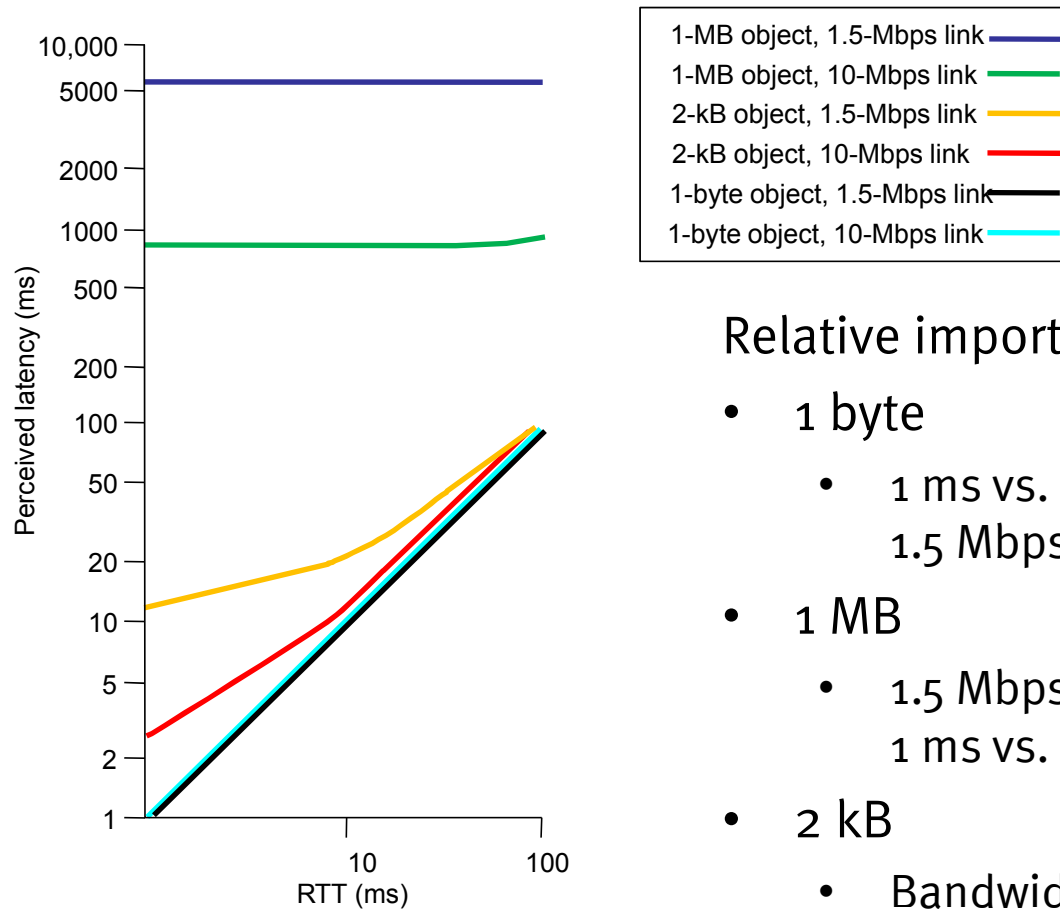
- <http://www.traceroute.org/>
- <http://www.slac.stanford.edu/comp/net/wan-mon/traceroute-srv.html#site>

- iperf

- Sets up TCP connection between two hosts
- Sends as much data as possible
- Sample results
 - <http://www-iepm.slac.stanford.edu/pinger/perfmap/iperf/sc2002/anim.gif>

- Alternative

- Download large file




Relative importance

- 1 byte
 - 1 ms vs. 100 ms dominates
 - 1.5 Mbps vs. 10 Mbps
- 1 MB
 - 1.5 Mbps vs. 10 Mbps dominates
 - 1 ms vs. 100 ms
- 2 kB
 - Bandwidth and latency relevant

- One-way example

- $50\text{ms} * 45\text{Mbps} = 2,250,000\text{b} \approx 275\text{KiB}$
(or 550KiB for RTT)




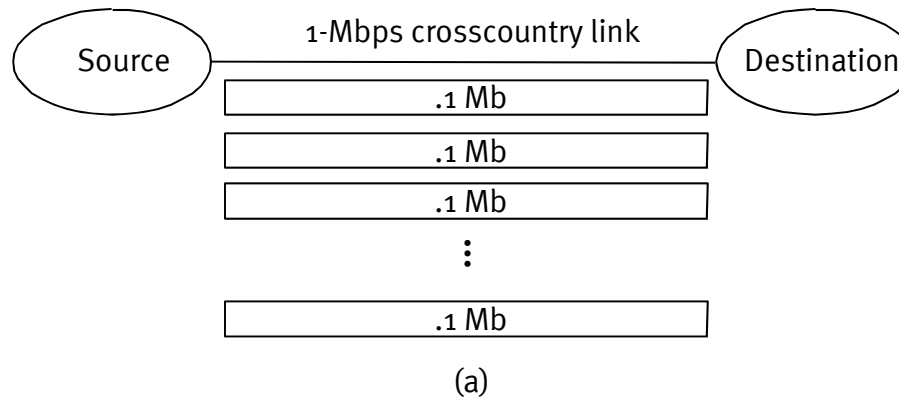
- Significance

- “The number of bits a sender must transmit before the first bit arrives at the receiver”
 - Amount of data to fill the pipe
- (RTT: “... before acknowledgement from receiver arrives”)

Bandwidth vs. Latency (cont'd)



Transfer of 10,000,000b with RTT of 100 ms



- A network must provide
 - general,
 - cost-effective,
 - robust,
 - high-performanceconnectivity among a large number of computers
- Resources (typically) shared on demand
 - Contention vs congestion
- Latency and bandwidth most important performance characteristics
 - Various tools for measurements

- Calculate and measure latencies
- Discuss relative importance of performance measures
- Explain significance of Delay x Bandwidth product