DSC204A: Scaleable Data Systems Winter 2025

Lecture 17: Networking Fundamentals

Instructor: Umesh Bellur *Scribe: Shreyash Kondakindi; Susmit Singh; Kavya Sridhar*

Midterm Examination Information

The midterm examination will focus primarily on conceptual understanding and theory, with no coding questions involved. Topics covered after the second week will have greater importance, notably CPU scheduling and virtual memory concepts. Questions will be multiple-choice or short-answer format designed to evaluate conceptual grasp rather than implementation specifics.

Overview

Networking fundamentally involves connections, which are communication channels established between two processes or endpoints. Each endpoint in a network is uniquely identified by a port number, associated IP address, and MAC address.

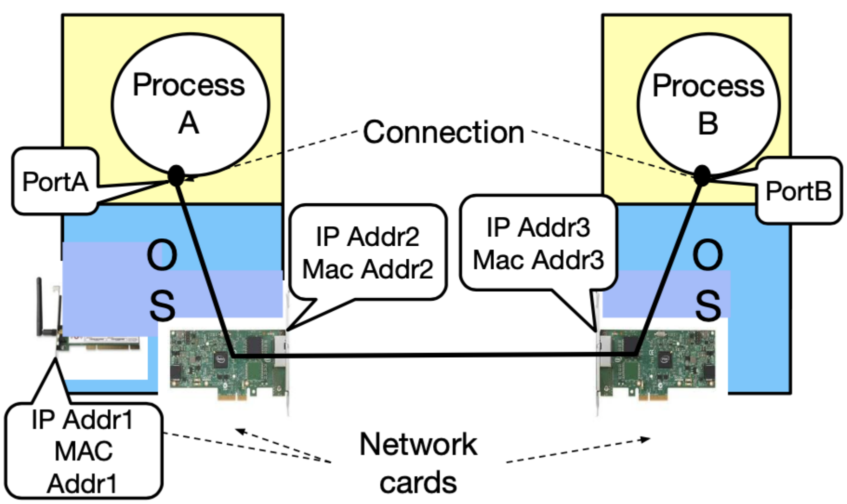

Why MAC addresses despite having IP addresses? MAC addresses are stable identifiers used especially in firewall setups. For example, to reliably block or permit network access irrespective of IP changes, MAC addresses are preferred due to their stability.

Network Core: Circuit Switching and Packet Switching

Networks operate primarily through two paradigms: Circuit Switching and Packet Switching.

Circuit Switching involves dedicated communication paths established via call setup and teardown phases. Bandwidth resources are specifically reserved for the entire communication duration, resulting in guaranteed performance, similar to traditional telephone systems. Resources are allocated either via frequency division (assigning specific frequency bands) or time division (allocating dedicated time slots).

Packet Switching, contrastingly, splits communication data into smaller packets. Resources in this system are shared dynamically, enabling efficient handling of bursty data, though it faces challenges like congestion leading to delays and packet loss. Packet switching leverages statistical multiplexing, meaning packets from different communications share the same network infrastructure, allowing optimal resource utilization.

An ongoing challenge with packet switching is ensuring circuit-like guarantees (such as consistent bandwidth) required by real-time audio and video applications. This issue remains largely unsolved, presenting a significant research opportunity.

Packet Switching: Store-and-Forward Mechanism

In packet-switched networks, data packets travel hop-by-hop between routers. Each packet must fully arrive at an intermediate router before being forwarded to the next (Store-and-Forward mechanism).

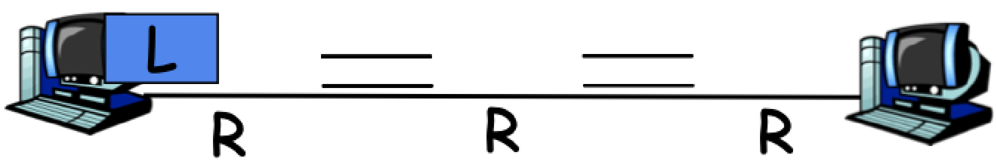

For instance, sending a packet of length L bits over a link with bandwidth R bits/sec takes L/R seconds per hop. For example, a 3-hop transmission introduces a delay of $3L/R$.

Packet Switching: Message Segmenting and Pipelining

Segmenting messages into smaller packets can significantly improve network efficiency.

Breaking a 7.5 Mbit message into 5000 smaller packets (each 1500 bits), using pipelining, reduces the transmission delay significantly from 15 seconds (single large packet) to around 5 seconds. The trade-off is packet overhead; very small packets incur more overhead.

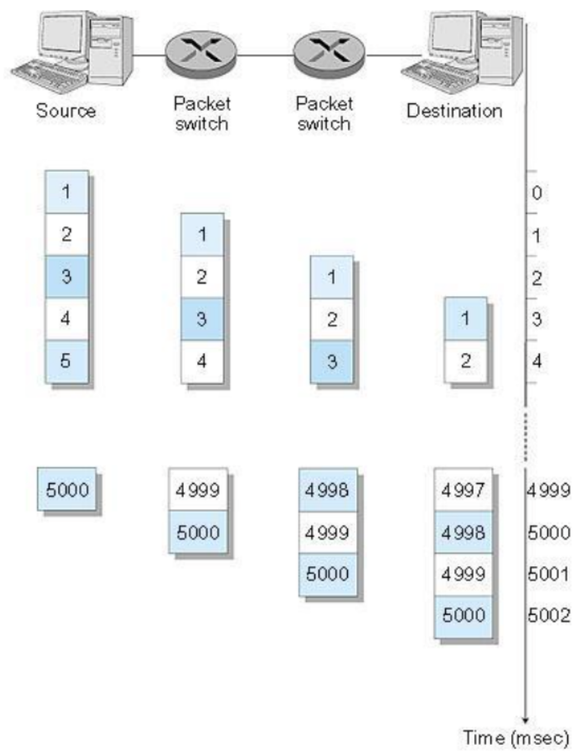

Characterizing Network Communication

Important parameters for characterizing network performance include:

- **Latency**: Time for the first bit to reach the destination.
- **Capacity (Bandwidth)**: Maximum rate of bits transmitted per second.
- **Jitter**: Variability in latency across packets.
- **Packet Loss and Reliability**: Likelihood and consequence of losing packets.

Sources of Delay in Packet-Switched Networks

Delays in packet-switched networks can be categorized into four types:

1. **Nodal Processing Delay**: Router checks packet headers, checks for errors, and determines the output link.
2. **Queuing Delay**: Time packets wait at routers due to congestion.
3. **Transmission Delay**: Time required to push packets onto the link, calculated by packet size divided by link bandwidth (L/R).
4. **Propagation Delay**: Time taken for a bit to travel across the link, calculated as distance divided by propagation speed in medium (usually light, so 2×10^8 m/s) = d/s.

The analogy of cars at toll booths effectively illustrates these concepts:

Imagine you have a caravan consisting of ten cars that needs to travel from one toll booth to another, separated by 100 miles. Each toll booth can only service one car at a time, taking a specific amount of time (transmission delay) before allowing the car onto the highway.

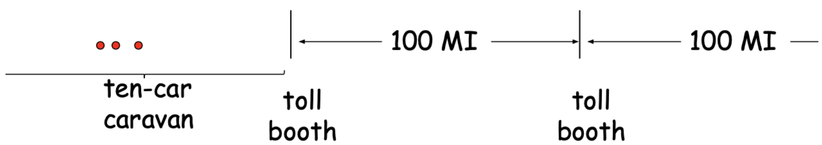

Scenario 1:

- Cars travel (propagate) at 100 miles per hour.
- Each car takes 12 seconds at the toll booth (transmission delay).
- The entire caravan (10 cars) thus takes $10 \times 12 = 120$ seconds (2 minutes) to pass through the first toll booth.
- After passing through, the last car has to travel 100 miles to the next booth. At 100 miles/hour, this takes exactly 1 hour.
- To calculate the total time until the caravan fully lines up at the second toll booth, add both delays:

```
Total Time = Transmission delay at first booth + Propagation delay to second booth
           = 2  minutes+ 60 minutes
           = 62 minutes
```

Thus, it takes a total of 62 minutes for the entire caravan to reach and line up completely at the second toll booth.

Scenario 2 (Higher Speed, Shorter Service Time):

Now let's consider a different scenario where the cars move faster, and the toll booths service quicker:

- Cars now travel at 1000 miles/hour.

- Each car is serviced at the toll booth in 1 minute.
- The entire caravan still takes 10 minutes (1 minute per car for 10 cars) to pass through the first toll booth.
- The first car will arrive at the second toll booth after just 6 minutes (100 miles at 1000 miles/hour).
- At 7 minutes, the first car has already reached the second toll booth, while 3 cars are still being serviced at the first toll booth.

In this scenario, the first car (analogous to the first bit of a packet) reaches the second toll booth before the entire caravan (entire packet) has finished being serviced at the first booth. This clearly demonstrates how in networks, a packet can start arriving at the next router before completely leaving the previous one.

Throughput and Bottlenecks

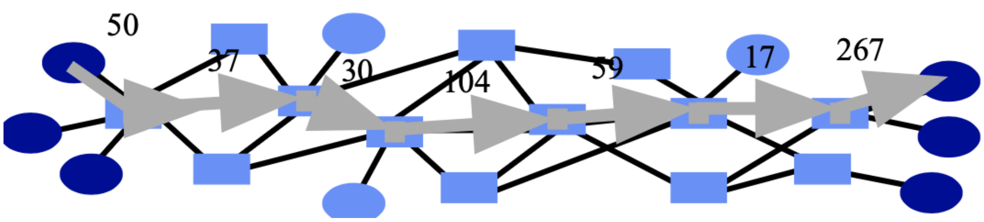

Throughput refers to the effective rate of successful message delivery over the network. It is constrained by the bottleneck, the slowest network segment. Factors influencing this bottleneck include low link bandwidth or congestion from multiple users sharing the link.

Real-world Latency and Route Tracing

Traceroute is a diagnostic tool used to measure delays across network hops. It sends multiple packets to intermediate routers along the path to a destination, measuring round-trip delays to each router, providing insights into network latency and identifying bottlenecks.

Use:

```
traceroute <any_domain_name>
```

on the terminal to find out more.

Congestion Control

When network buffers overflow due to excessive incoming data, packets are dropped. This mechanism signals the sender to reduce transmission rates, known as congestion control. This ensures network stability and reduces packet loss. Prioritization is also employed to manage resource allocation for critical streams (e.g., video conferencing).

Protocols for Data Ordering

Data ordering in packet-switched networks is typically managed by Transport Layer protocols like TCP (reliable, ordered delivery) and UDP (unreliable, unordered but faster).

Cross-country latency calculation example

Consider a network scenario involving data transmission across a country, approximately 5000 kilometers in distance. Let's break down this example step by step for clarity:

- **Distance and Speed**: The signal must travel 5000 kilometers (or 5 million meters). It travels at a speed roughly two-thirds the speed of light, which is about 2×10^8 meters per second.
- **Propagation Delay**: To calculate the time it takes just for the signal to travel this distance, we use the formula:

$$\text{Propagation Delay} = \frac{\text{Distance}}{\text{Speed}} = \frac{5 \times 10^6 \text{ m}}{2 \times 10^8 \text{ m/s}} = 0.025 \text{ seconds} = 25 \text{ milliseconds (ms)}$$

- **Round Trip Time (RTT)**: Acknowledgements (ACKs) must return to confirm packet receipt. Thus, the total RTT is twice the propagation delay:

$$\text{RTT} = 25 \text{ ms (forward)} + 25 \text{ ms (return)} = 50 \text{ ms}$$

- **Transmission Delay**: Suppose we have a packet of size 1250 bytes (equivalent to 10 kilobits) and a link speed (capacity) of 100 Mbps (megabits per second). The time to push the packet onto the link (transmission delay) is calculated as:

$$\text{Transmission Delay} = \frac{\text{Packet size}}{\text{Link speed}} = \frac{10 \text{ kbits}}{100 \text{ Mbps}} = 0.1 \text{ ms}$$

- **Effective Bandwidth**: Although the link speed is 100 Mbps, due to the high RTT, the effective data throughput is significantly reduced. To calculate this effective bandwidth, we divide the packet size by the total RTT:

$$\text{Effective Bandwidth} = \frac{10 \text{ kbits}}{50.1 \text{ ms}} \approx 200 \text{ kbps}$$

This example demonstrates that even with high link capacities, the effective bandwidth can be substantially lower due to latency (RTT), greatly influencing overall network performance.