Computer Networks - Lecture 01: Introduction

Course Outline

- **Textbook:** Computer Networking: A Top-Down Approach, 8th ed., Kurose & Ross
- **Grading:**
    - Attendance & participation: 5-7%
    - Assignments & quizzes: 40%
    - Midterm: 15%
    - Final: 40%
- **Join code:** 142tcab
- **Course materials and discussions:** MS Teams
- **TA:** Eng. Mohamed Essam

Chapter 1: Computer Networks and the Internet

Outline

- What Is the Internet?
- The Network Edge
- The Network Core
- Delay, Loss, and Throughput in Packet-Switched Networks
- Protocol Layers and Their Service Models

What Is the Internet?

- The Internet is a **network of networks**.
- We will use the **public Internet** as our main example to learn about computer networks and their protocols.

A Nuts-and-Bolts Description of the Internet

- The Internet interconnects **billions of computing devices** throughout the world.
- These devices are called **hosts** or **end systems**.
- There were about 18 billion devices connected to the Internet in 2017, and the number will reach **28.5 billion by 2022.**
- End systems are connected by a network of **communication links** and **packet switches**.
- A **packet switch** takes a packet arriving on one of its incoming communication links and forwards that packet on one of its outgoing communication links.
- The **transmission rate** of a link is measured in bits/second (bps).
- The two most prominent types of packet switches in today's Internet are **routers** and **link-layer switches.**
- The sequence of communication links and packet switches traversed by a packet from the sending end system to the receiving end system is known as a **route** or **path** through the network.
- End systems access the Internet through **Internet Service Providers (ISPs).**
- Each ISP is in itself a **network** of packet switches and communication links.
- End systems, packet switches, and other pieces of the Internet run **protocols.**
- The **Transmission Control Protocol (TCP)** and the **Internet Protocol (IP)** are two of the most important protocols in the Internet.
- The Internet's principal protocols are collectively known as **TCP/IP.**

- **Internet standards** are developed by the **Internet Engineering Task Force (IETF).**
- The IETF standards documents are called **requests for comments (RFCs).**
- There are currently nearly 9000 RFCs.
- Other bodies also specify standards for network components, e.g. the IEEE 802 LAN Standards Committee.

A Services Description of the Internet

- The Internet is an **infrastructure** that provides services to **distributed applications.**
- Internet applications run on **end systems** - they **do not run** in the packet switches in the network core.
- End systems attached to the Internet provide a **socket interface** that specifies how a program asks the Internet infrastructure to deliver data to another end system.
- The Internet provides **multiple services** to its applications.

What Is a Protocol?

- It takes two (or more) communicating entities running the same **protocol** in order to accomplish a task.
- In a **human protocol**, there are specific messages we send, and specific actions we take in response to the received reply messages or other events (such as no reply within some given amount of time).
- Much of this course is about **computer network protocols.**
- A **protocol** defines the format and the order of messages exchanged between two or more communicating entities, as well as the actions taken on the transmission and/or receipt of a message or other event.

The Network Edge

- The Internet's **end systems** include desktop computers, servers, and mobile devices.
- End systems are also referred to as **hosts** because they host (that is, run) application programs.
- Hosts are sometimes further divided into two categories: **clients** and **servers**.
- Most of the servers reside in large **data centers.**
    - Google has 19 data centers on four continents, collectively containing several million servers.

Access Networks

- **Home access:** DSL, Cable, FTTH, and 5G Fixed Wireless
- **Access in the enterprise (and the home):** Ethernet and WiFi
- **Wide-area wireless access:** 3G and LTE 4G and 5G

Home Access: DSL

- When **digital subscriber line (DSL)** is used, a customer's telco is also its ISP.
- A DSL modem uses the existing telephone line to exchange data with a **digital subscriber line access multiplexer (DSLAM)** located in the telco's local central office (CO).
- The residential telephone line carries both data and traditional telephone signals simultaneously, which are encoded at different **frequencies:**
    - A high-speed downstream channel, in the 50 kHz to 1 MHz band
    - A medium-speed upstream channel, in the 4 kHz to 50 kHz band

    - An ordinary two-way telephone channel, in the 0 to 4 kHz band
- On the customer side, a **splitter** separates the data and telephone signals arriving to the home and forwards the data signal to the DSL modem.
- On the telco side, in the CO, the **DSLAM** separates the data and phone signals and sends the data into the Internet.
    - Hundreds or even thousands of households connect to a single DSLAM.
- **Downstream transmission rates:** 24 Mbps and 52 Mbps
- **Upstream rates:** 3.5 Mbps and 16 Mbps
- The newest standard provides for aggregate upstream plus downstream rates of 1 Gbps.
- DSL is designed for short distances between the home and the CO.
    - Located within 5 to 10 miles of the CO. (1 mile=1.6 km)

Other Home Access

- **Cable Internet access** makes use of the cable television company's existing cable television infrastructure.
    - It is often referred to as hybrid fiber coax (HFC) and is a shared broadcast medium.
    - Downstream bitrates of 40 Mbps and 1.2 Gbps, and upstream rates of 30 Mbps and 100 Mbps.
- **Fiber to the home (FTTH)** provides even higher speeds is that can potentially provide Internet access rates in the gigabits per second range.
- **5G fixed wireless** promises high-speed residential access, without installing costly and failure-prone cabling from the telco's CO to the home.

Access in the Enterprise/Home: Ethernet

- A **local area network (LAN)** is used to connect an end system to the edge router.
- **Ethernet** users use twisted-pair copper wire to connect to an Ethernet switch.
- **With Ethernet access:**
    - Users typically have 100 Mbps to tens of Gbps access to the Ethernet switch.
    - Servers may have 1 Gbps to 10 Gbps access.

Access in the Enterprise/Home: WiFi

- Wireless LAN access based on IEEE 802.11 technology (**WiFi**) is now just about everywhere.
- A wireless LAN user must typically be within a few tens of meters of the access point.
- 802.11 today provides a shared transmission rate of up to more than 100 Mbps.
- **Example:**
    - A home network with a roaming laptop, multiple home appliances, as well as a wired PC.
    - A base station (**WiFi access point**) that communicates with the wireless PC and other wireless devices in the home.
    - A **home router** that connects the wireless access point, and any other wired home devices, to the Internet.

The Network Core

- **Packet switching** is the dominant method used in the Internet.

Packet Switching

- In a network application, end systems exchange **messages** with each other.
- The source breaks long messages into smaller chunks of data known as **packets**.
- Each packet travels through communication links and **packet switches**.
    - Routers and link-layer switches
    - A router will typically have many incident links
- Most packet switches use **store-and-forward transmission** at the inputs to the links. That is, it must receive the entire packet before it can begin to transmit the first bit of the packet onto the outbound link.
- Each packet consisting of L bits; Transmission rate is R bits/sec.
- Sending one packet from source to destination over a path consisting of N links (N - 1 routers) each of rate R, the delay is:

dend-to-end = NL/R

- This equation ignores propagation delay.
- For each attached link, the packet switch has an **output buffer/queue**, which stores packets that the router is about to send into that link.
- In addition to the store-and-forward delays, packets suffer output buffer **queuing delays** that depend on the level of congestion in the network.
- The amount of buffer space is finite, therefore **packet loss** will occur - either the arriving packet or one of the already-queued packets will be dropped.
- **Forwarding tables and routing protocols:**
    - Every end system has an address called an **IP address** that has a hierarchical structure.
    - The **destination's IP address** is in the packet's header.
    - Each router has a **forwarding table**.
    - A router uses a packet's destination address to index a forwarding table and determine the appropriate outbound link.
    - The Internet has a number of special **routing protocols** that are used to automatically set the forwarding tables.

Circuit Switching

- Traditional **telephone networks** are examples of circuit-switched networks.
- In **circuit-switched networks**, the resources needed along a path (buffers, link transmission rate) are **reserved** for the duration of the communication session between the end systems.
- When two hosts want to communicate, the network establishes a dedicated **end-to-end connection** between the two hosts.
- The sender can transfer the data to the receiver at the **guaranteed** constant rate.
- The Internet makes its **best effort** to deliver packets in a timely manner, but it does not make any guarantees.
- **Multiplexing in circuit-switched networks:**
    - A circuit in a link is implemented with either **frequency-division multiplexing (FDM)** or **time-division multiplexing (TDM).**
    - **Example:**
        - FM radio stations use FDM to share the frequency spectrum between 88 MHz and 108 MHz, with each station being allocated a specific frequency band.

        - For a TDM link, time is divided into frames of fixed duration, and each frame is divided into a fixed number of time slots.
    - Circuit switching is wasteful because the dedicated circuits are idle during silent periods.
    - Establishing end-to-end circuits is complicated and requires **complex signaling software** to coordinate the operation of the switches along the end-to-end path.

Packet Switching vs. Circuit Switching

- Packet switching is **not suitable** for real-time services (telephone calls and video conference calls).
- Packet switching offers better **sharing** of transmission capacity and is **simpler**, more **efficient**, and **less costly** to implement.
- Circuit switching **pre-allocates** use of the transmission link regardless of demand, with allocated but unneeded link time going unused.
- Packet switching on the other hand allocates link use **on demand**.

A Network of Networks

- Over the years, the network of networks that forms the Internet has evolved into a very **complex structure.**
- Much of this evolution is driven by **economics and national policy**, rather than by performance considerations.

Network Structures

- **Naive approach:** Each access ISP directly connects with every other access ISP.
- **Network structure 1:** Interconnects all of the access ISPs with a single global transit ISP.
- **Network structure 2:** Hundreds of thousands of access ISPs and multiple global transit ISPs.
- **Network structure 3:**
    - In any given region, there is a regional ISP to which the access ISPs in the region connect.
    - Each regional ISP then connects to **tier-1 ISPs** that do not have a presence in every city in the world (a dozen tier-1 ISPs).
    - Each access ISP pays the regional ISP to which it connects, and each regional ISP pays the tier-1 ISP to which it connects.
    - There may be a larger regional ISP to which the smaller regional ISPs in that region connect.
- **Network structure 4:**
    - **Points of presence (PoPs):** A group of one or more routers in the provider's network where customer ISPs can connect into the provider ISP (not at the access level).
    - **Multi-homing:** Connect to two or more provider ISPs.
    - **Peering:** Pair of nearby ISPs at the same level of the hierarchy can directly connect their networks together.
    - **Internet exchange points (IXPs):** A third-party company can create an IXP that is a meeting point where multiple ISPs can peer together.
- **Network structure 5:** Builds on top of Network Structure 4 by adding **content-provider networks**.
    - **Example:** The Google data centers are all interconnected via Google's private TCP/IP network, which spans the entire globe but is nevertheless separate from the public Internet.

Delay, Loss, and Throughput in Packet-Switched Networks

- The physical laws introduce **delay and loss** as well as constrain **throughput**.
    - **Throughput** is the amount of data per second that can be transferred between end systems.
- The packet suffers from several **types of delays** at each node along the path:
    - **Processing delay:** (microseconds or less)
    - **Queuing delay:** (microseconds to milliseconds) (depend on the number of earlier-arriving packets)
    - **Transmission delay:** L/R (packet length L bits; R transmission rate in bps) (amount of time required to push all of the packet's bits into the link)
    - **Propagation delay:** (d/s distance between two routers divided by the propagation speed) (depends on the physical medium) (milliseconds)
        - 2 × 108 meters/sec to 3 × 108 meters/sec

Nodal Delay

- **dnodal = dproc + dqueue + dtrans + dprop**
- The contribution of these delay components can vary significantly.
    - **Example:**
        - LAN: dproc is negligible
        - Routers interconnected by a geostationary satellite link: dprop is hundreds of milliseconds (dominant).
- The processing delay, **dproc**, is often negligible.
    - However, it strongly influences a router's maximum throughput.

Queuing Delay and Packet Loss

- The queuing delay can vary from packet to packet (uses **statistical measures** such as average, variance, probability).
- **Queuing delay depends on:**
    - The rate at which traffic arrives (a packets/sec) (assume each is L bits).
    - The transmission rate of the link (R bps).
    - The nature of the arriving traffic (periodically or in bursts; or random).
- **Traffic intensity = La/R**
    - If La/R > 1, the queue will tend to increase without bound and the queuing delay will approach infinity!
- La/R < 1: The nature of the arriving traffic impacts the queuing delay.
    - If packets arrive periodically, then every packet will arrive at an empty queue and there will be **no queuing delay**.
    - If packets arrive in bursts but periodically, there can be a **significant average queuing delay**.
        - **Example:** N packets arrive simultaneously every (L/R)N seconds,

nth packet transmitted has a queuing delay of (n - 1)L/R seconds.

- A small percentage increase in the intensity will result in a much larger percentage-wise increase in delay.
- Performance at a node is often measured not only in terms of **delay**, but also in terms of the probability of **packet loss**.

End-to-End Delay

- Assume N - 1 routers, no queuing delay.
- **dend-to-end = N(dproc + dtrans + dprop)**
- **Traceroute** is a simple program. When the user specifies a destination hostname, the program in the source host sends multiple, special packets toward that destination (graphical interface PingPlotter).
    - The source sends 3 × N packets to the destination.
    - As these packets work their way toward the destination, they pass through a series of routers.
    - When a router receives one of these special packets, it sends back to the source a short message that contains the **name and address of the router**.
    - The source can **reconstruct the route** taken by packets flowing from source to destination, and the source can determine the **round-trip delays** to all the intervening routers.

Throughput

- Use the **speedtest** application to measure the end-to-end delay and download throughput between a host and servers.
- If a file consists of F bits and the transfer takes T seconds for Host B to receive all F bits, then the **average throughput** of the file transfer is F/T bits/sec.
- We may think of bits as **fluid** and communication links as **pipes**.
- In a simple two-link network, the throughput is **min{Rc, Rs}**, that is, it is the transmission rate of the **bottleneck link**.
- For a network with N links between the server and the client, with the transmission rates of the N links being R1, R2, ..., RN. The throughput for a file transfer from server to client is **min{R1, R2, ..., RN}**.
- When there is **no other intervening traffic**, the throughput can simply be approximated as the minimum transmission rate along the path between source and destination.
- Links in the **core** of the communication network have very high transmission rates.
- The constraining factor for throughput in today's Internet is typically the **access network**.

Protocol Layers

- Network designers organize protocols in **layers**.
- A protocol layer can be implemented in software, in hardware, or in a combination of the two.
    - **Application-layer protocols** are almost always implemented in software and so are **transport-layer protocols.**
    - The **physical layer and data link layers** are responsible for handling communication over a specific link, they are typically implemented in a network interface card.
    - The **network layer** is often a mixed implementation of hardware and software.
- Potential **drawbacks** of layering is that one layer may duplicate lower-layer functionality and the functionality at one layer may need information that is present only in another layer.

Protocol Layering

- The **application layer** is where network applications and their application-layer protocols reside.
    - With the application in one end system using the protocol to exchange packets of information (called **messages**) with the application in another end system.
    - **Example:** HTTP, SMTP, FTP, DNS
- The **transport layer** transports application-layer messages between application endpoints (a transport-layer packet is referred as a **segment**).
    - The **UDP protocol** provides a connectionless service to its applications.
    - **TCP protocol** provides a connection-oriented service to its applications: guaranteed delivery; flow control; congestion-control.
- The **network layer** (IP layer) is responsible for moving network-layer packets known as **datagrams** from one host to another.
    - IP protocol defines the fields in the datagram as well as how the end systems and routers act on these fields.
    - This layer contains **routing protocols** that determine the routes that datagrams take between sources and destinations.
- The **link layer** delivers the datagram to the next node along the route. At this next node, the link layer passes the datagram **up** to the network layer.
    - A datagram may be handled by **different** link-layer protocols at different links along its route.
    - The link-layer packets are referred as **frames.**
    - **Example:** Ethernet, WiFi
- The job of the **physical layer** is to move the **individual bits** within the frame from one node to the next.
    - Depends on the actual transmission medium of the link.
    - **Example:** Ethernet has many physical-layer protocols: one for twisted-pair copper wire, another for coaxial cable, another for fiber, and so on.

Encapsulation

- The transport layer takes the **message** and appends additional information. The **transport-layer segment** encapsulates the application-layer message.
- The network layer adds network-layer header information such as source and destination end system addresses, creating a **network-layer datagram**.
- The datagram is then passed to the link layer, which (of course!) will add its own link-layer header information and create a **link-layer frame**.

Summary

- What Is the Internet?
- The Network Edge
- The Network Core
- Delay, Loss, and Throughput in Packet-Switched Networks
- Protocol Layers and Their Service Models