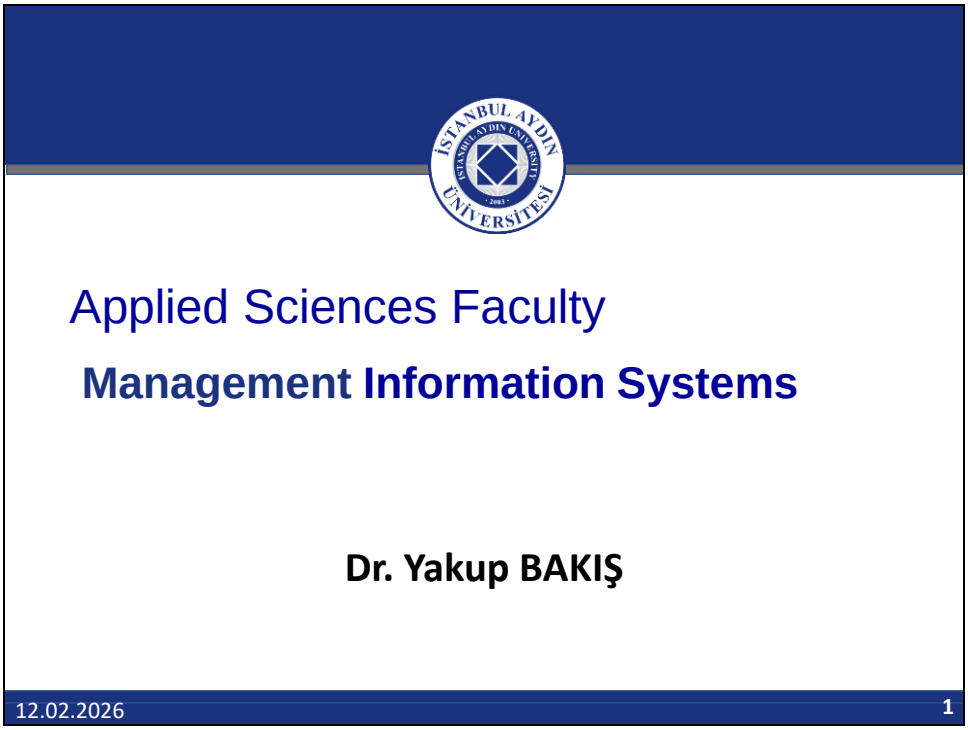

Last week, we talked about how the Internet works in general.

Today, we move one layer higher.
We will begin to understand the **Application Layer**,
which is the layer closest to the user.

When you open a website, send an email, or use Instagram, you are using the Application Layer.

This is the part of networking that you actually see and use every day.

Chapter 1

Computer Networks and the Internet

How many of you use the Internet every day?
Probably all of you.
But today, we will start learning how to build the web, not just use it.
That is the difference between a user and an engineer.
Users consume technology.
Engineers understand and create technology.

introduction

Chapter goal:

- Get "feel," "big picture," introduction to terminology
  - more depth, detail *later* in course
- Approach:
  - use Internet as example

Overview/roadmap:

- What *is* the Internet?
- What *is* a protocol?
- Network edge: hosts, access network, physical media
- Network core: packet/circuit switching, internet structure
- Performance: loss, delay, throughput
- Security
- Protocol layers, service models
- History

The goal of this chapter is to give you a big picture.
We are not going very deep today.
We just want to understand the terminology and the general structure.
Later in the course, we will go into much more detail.
Today is about understanding the map.
Later, we will explore each street.

1.1 What Is the Internet?

- There are a couple of ways to answer this question.
- First, we can describe the nuts and bolts of the Internet, that is, the basic hardware and software components that make up the Internet.
- Second, we can describe the Internet in terms of a networking infrastructure that provides services to distributed applications.
- Let's begin with the nuts-and-bolts description, using Figure 1.1 to illustrate our discussion.

There are two main ways to define the Internet.
First, we can describe its hardware and software components.
This is what we call the nuts-and-bolts view.
Second, we can describe the Internet as a service platform for applications.
Today, we will start with the nuts-and-bolts view.

The Internet: a “nuts and bolts” view

Billions of connected computing *devices*:

- *hosts = end systems*
- running *network apps* at Internet’s “edge”

Packet switches: forward packets (chunks of data)

- *routers, switches*

Communication links

- fiber, copper, radio, satellite
- transmission rate: *bandwidth*

Networks

- collection of devices, routers, links: managed by an organization

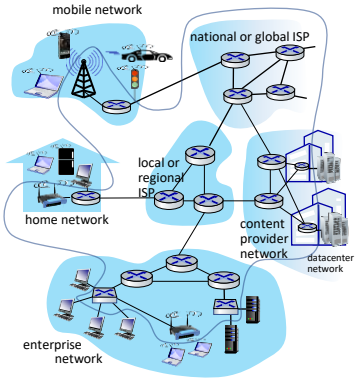

The Internet connects billions of devices.
These devices are called hosts or end systems.
They run applications like web browsers, email, and streaming apps.
These hosts are connected by communication links.
Links can be fiber, copper, radio, or satellite.
The devices that forward data are called packet switches.
The most important packet switch is the router.

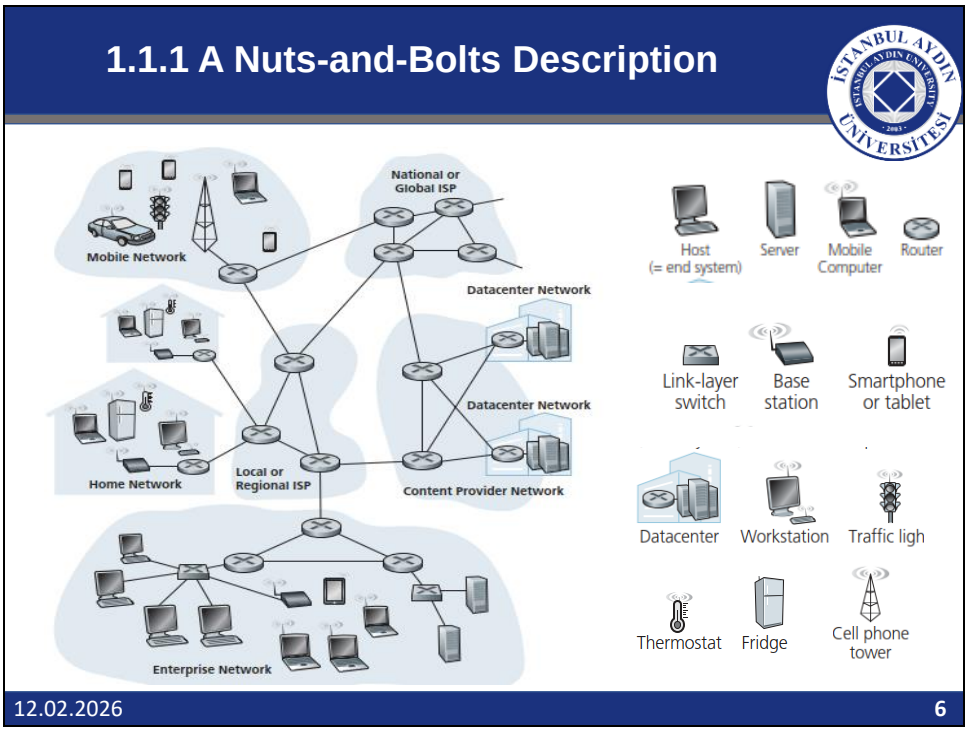

The Internet is a computer network that interconnects billions of computing devices throughout the world.

The Internet is no longer just computers.

Increasingly, users connect to the Internet with smartphones and tablets—today, close to half of the world's population are active mobile Internet users with the percentage expected to increase to 75% by 2025.

Furthermore, nontraditional Internet "things" such as TVs, gaming consoles, thermostats, home security systems, home appliances, watches eye glasses, cars, traffic control systems, and more are being connected to the Internet.

In Internet jargon, all of these devices are called hosts or end systems.

The Internet is becoming a network of smart devices.

The Internet: a "nuts and bolts" view

- *Internet:* "network of networks"
  - Interconnected ISPs

ISP = Internet Service Provider

- *protocols* are *everywhere*
  - control sending, receiving of messages
  - e.g., HTTP (Web), streaming video, Skype, TCP, IP, WiFi, 4G, Ethernet
  - *Internet standards*
    - RFC: Request for Comments
    - IETF: Internet Engineering Task Force

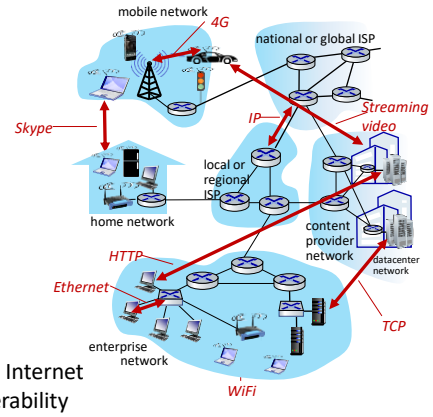

- Internet = decentralized system
- No central authority controls the whole Internet
- Standardization enables global interoperability

Introduction: 1-7

The Internet is not a single network.

An ISP is an Internet Service Provider.
For example, Turk Telekom or Vodafone or Superonline.
But no single ISP covers the whole world.
So ISPs must connect to each other ISPs.

They exchange traffic and share data.
This connection between ISPs is called **interconnected ISPs.**
That is why the Internet is called a network of networks.

Many independent networks are connected together to form one global system.

ISPs connect homes, companies, universities, and data centers.
When you open a website, your ISP communicates with other ISPs to deliver the data to you.

Now, how do all these networks understand each other?
The answer is: **Protocols.**

Protocols are everywhere in the Internet.
They control how messages are sent and received.
They define the format and the rules of communication.

For example:

HTTP is used for the Web.
TCP and IP are used for reliable data delivery.
WiFi and 4G are used for wireless communication.
Streaming protocols are used for video services.

Without protocols, devices could not communicate.
It would be chaos.

But who defines these protocols?
This is where Internet standards come into play.
The Internet works not because it is controlled by one company,
but because everyone follows the same rules.

Internet standards are written in documents called **RFCs — Request for Comments.**
These standards are developed by an organization called the **IETF — Internet Engineering Task Force.**

Thanks to these standards, devices made by different companies can communicate with each other.

1.1.1 A Nuts-and-Bolts Description

- In Internet jargon, all of these devices are called **hosts** or **end systems.**
- End systems are connected together by a network of **communication links** and **packet switches.**
- Different links can transmit data at different rates, with **the transmission rate** of a link measured in **bits/second.**
- The resulting packages of information, known as **packets**, which are **reassembled** into the original data.
- Packet switches come in many shapes and flavors, but the two most prominent types in today's Internet are **routers** and **link-layer switches.**
- a packet from the sending end system to the receiving end system is known as a **route** or **path** through the network.
- **Packets are sent independently**
- **The Internet is a packet-switched network.**
- **The user does not see the complexity behind data transmission.**
- **This packet-switching architecture increases efficiency and allows multiple users to share the same network resources.**

12.02.2026 8

When data is sent through the Internet, it is not transmitted as one large file. Instead, it is divided into small pieces called **packets**.

Each packet contains not only a portion of the data, but also control information such as the source address and destination address.
This additional information allows the network to deliver the packet to the correct destination.

These packets travel independently through the network.
They pass through devices called **routers** and **link-layer switches**, which are collectively known as packet switches.

The sequence of communication links and packet switches that a packet follows from the sender to the receiver is called a **route** or **path**.

It is important to understand that different packets belonging to the same file may take different routes through the network.
The Internet does not guarantee that all packets will follow the same path.

At the destination end system, all packets are collected and **reassembled** into the original data.
If everything arrives correctly, the user sees the complete file without noticing this complex process.

End systems are connected together by a network of communication links and packet switches.

These communication links are built from different types of physical media, such as coaxial cable, copper wire, optical fiber, and wireless radio spectrum.

Each communication link has a transmission rate, which is measured in **bits per second**.
Some links are very fast, such as fiber-optic connections, while others may be slower, such as wireless connections.

Packet switches come in many forms, but the two most important types in today's Internet are **routers** and **link-layer switches**.

Routers typically operate in the network core and make forwarding decisions based on IP addresses.
Link-layer switches usually operate within local networks and forward data based on MAC addresses.

Together, communication links and packet switches form the infrastructure that allows data to move across the global Internet.

1.1.2 A Services Description

The Internet: a "service" view

- *Infrastructure* that provides services to applications:
  - Web, streaming video, multimedia teleconferencing, email, games, e-commerce, social media, inter-connected appliances, …
- provides *programming interface* to distributed applications:
  - "hooks" allowing sending/receiving apps to "connect" to, use Internet transport service
  - provides service options, analogous to postal service

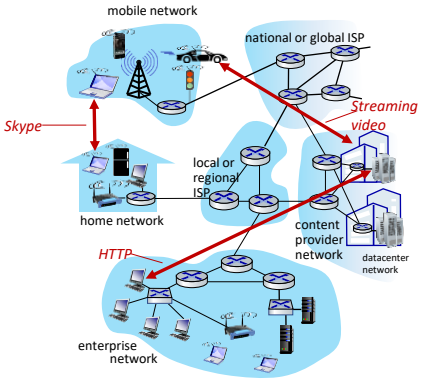

This perspective helps us understand that the Internet is not just hardware
it is a programmable platform for distributed applications.

Another way to define the Internet is as a service platform.
Instead of focusing on cables, routers, and hardware, we can think of the Internet as an infrastructure that provides services to applications.

These applications include the Web, email, video streaming, online games, video conferencing, social media, and many more.

The important point is this:
The Internet itself does not run these applications.
The applications run on end systems.
The Internet simply provides the communication service between them.

So how do applications communicate over the Internet?
They use something called a **socket interface**.

A socket is like a communication door between a program and the Internet.
It allows a program to send and receive data using Internet transport services.

In this sense, the Internet is similar to a postal service.
When you send a letter, you follow certain rules.
You put the letter in an envelope, write the address, and drop it in a mailbox.

Similarly, applications must follow certain rules to send data across the Internet.
The Internet provides different types of services, just like the postal service offers standard delivery, express delivery, and tracking options.

In the same way, the Internet offers different transport services, such as reliable delivery or faster but less reliable delivery.

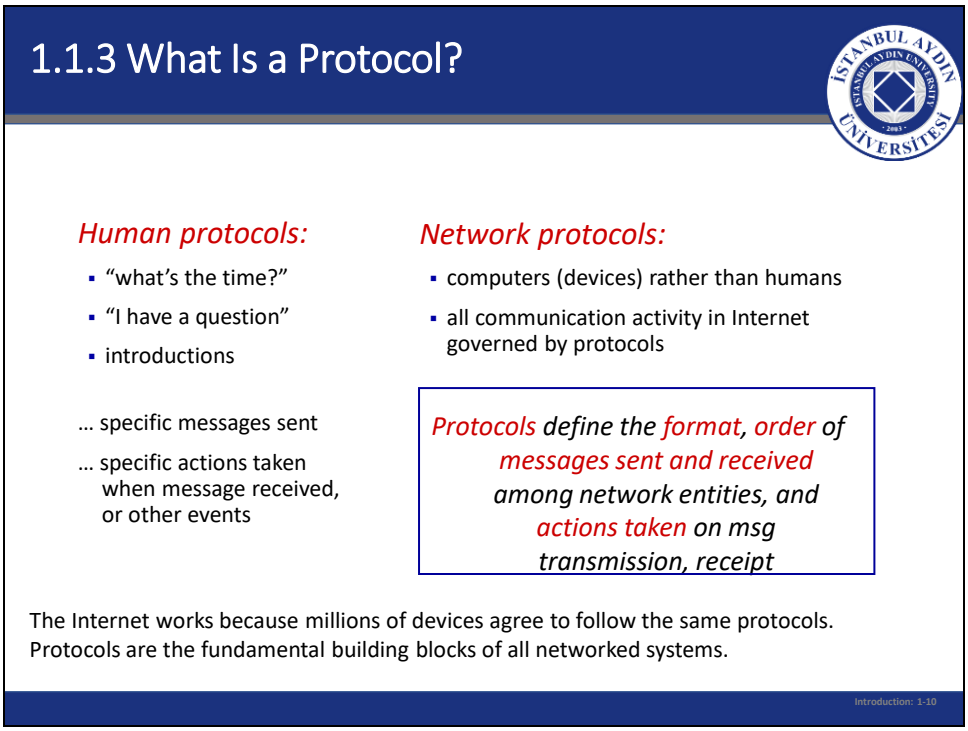

In the previous slide, we defined the Internet as a service platform.
We said that applications use a socket interface to communicate.

But here is the key question:
How does the Internet know what to do with the data sent through a socket?

The answer is: **Protocols.**

A protocol defines the rules of communication between network entities.
It defines the format of messages.
It defines the order of messages.
And it defines what actions are taken when a message is sent or received.

In other words, protocols tell devices how to talk to each other.

Without protocols, communication would be chaos.
Devices would send data in different formats, in random order, without understanding each other.

Just like humans follow social rules when speaking, computers follow protocols when communicating.

For example:
When you say "What's the time?",
there is an expected reply.

Similarly, in networking, when a client sends a request,

the server must send a response in a specific format. The Internet works because millions of devices agree to follow the same protocols.

This is how applications, sockets, and the Internet infrastructure work together.
Applications request services.
Sockets provide the interface.
Protocols define the communication rules.

And this is what makes the global Internet possible.

What's a protocol?

A human protocol and a computer network protocol:

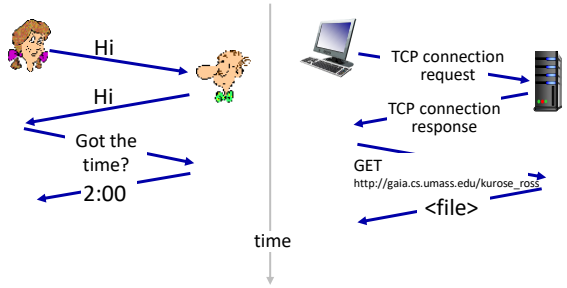

Q: other human protocols?

Think about a human conversation.
First, we say "Hi."
Then the other person replies.
Then we ask a question.
This is a protocol.
Computers do the same thing.
For example, a browser sends a TCP connection request.
The server replies.
Then the browser sends a GET request.

Explain important points

- **distributed** entities, exchanging messages (governed by protocols)
- Time going down
- go over definition of protocol (showing format, order of messages sent and received, and actions taken)

We'll see these kinds of diagrams a lot

A closer look at Internet structure
1.2 The Network Edge

Network edge:

- hosts: clients and servers
- servers often in data centers

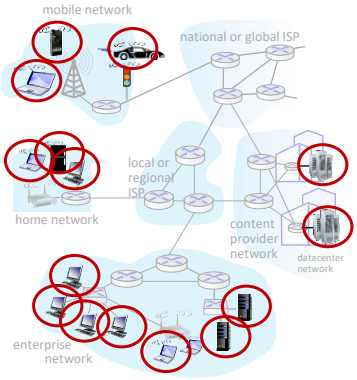

Now we move to a closer look at the Internet structure.
We start with something called the *network edge*.

The network edge contains the *hosts*.
In Internet terminology, hosts are also called end systems.

Hosts can be:

•Clients

•Servers

Clients request data.
Servers provide data.

For example, when you open a website, your computer is the client.
The web server in a data center is the server.

Today, many servers are located in large data centers around the world.

A closer look at Internet structure
1.2 The Network Edge

Network edge:

- hosts: clients and servers
- servers often in data centers

Access networks, physical media:

- wired, wireless communication links

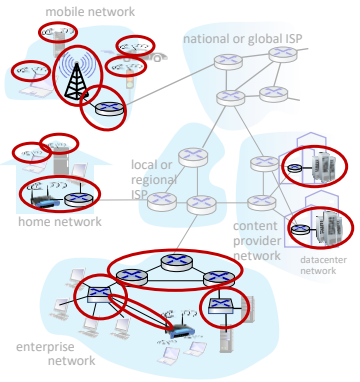

Now, how do these hosts connect to the Internet?

They use something called an *access network*.

Access networks connect end systems to the first router of the ISP.

These connections use physical media.
Physical media can be:

•Wired links

•Wireless links

For example:

•Fiber optic cables

•Copper cables

•WiFi

•Cellular networks

A closer look at Internet structure
1.2 The Network Edge

Network edge:

- hosts: clients and servers
- servers often in data centers

Access networks, physical media:

- wired, wireless communication links

Network core:

- interconnected routers
- network of networks

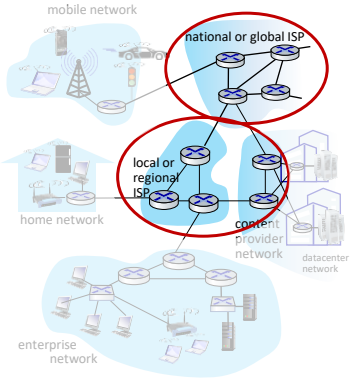

After the network edge, we have the *network core*.

The network core is made of interconnected routers.

This is the part of the Internet that moves packets across long distances.

Remember:
The Internet is called a *network of networks*
because many networks are connected together through routers.

1.2.1-2 Access networks and physical media

Q: How to connect end systems to edge router?

- residential access nets
- institutional access networks (school, company)
- mobile access networks (WiFi, 4G/5G)

What to look for:

- transmission rate (bits per second) of access network?
- shared or dedicated access among users?

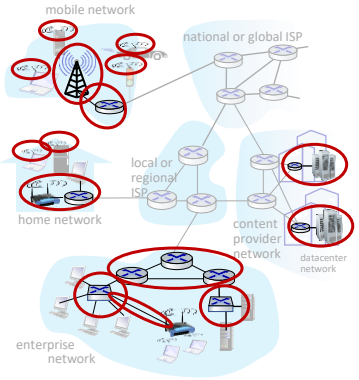

Now the important question:
How do end systems connect to the edge router?

There are three main types:

1.Residential access networks

2.Institutional access networks

3.Mobile access networks

When we evaluate an access network, we look at:

•Transmission rate

•Whether access is shared or dedicated

Transmission rate is measured in bits per second.
And some networks are shared among many users,
while others are dedicated per user.

1.2.1-2 Access networks and physical media

Access networks: cable-based access

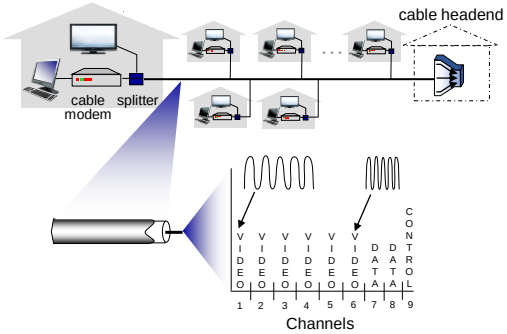

frequency division multiplexing (FDM): different channels transmitted in different frequency bands

Cable Internet uses frequency division multiplexing, or FDM.

This means different signals use different frequency bands.

For example:

•Some frequencies carry TV signals

•Some frequencies carry Internet data

All of this travels over the same cable.

1.2.1-2 Access networks and physical media

Access networks: cable-based access

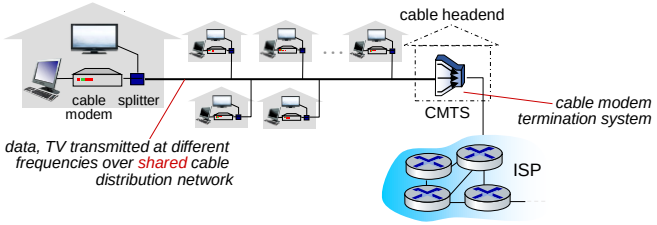

- HFC: hybrid fiber coax
  - asymmetric: up to 40 Mbps – 1.2 Gbs downstream transmission rate, 30-100 Mbps upstream transmission rate
- network of cable, fiber attaches homes to ISP router
  - homes *share access network* to cable headend

HFC stands for Hybrid Fiber Coax.

It uses fiber in the backbone and coaxial cable to homes.

It is asymmetric.

That means:

•Download speed is higher

•Upload speed is lower

Homes share the same cable network to the cable headend.
So bandwidth is shared.

1.2.1-2 Access networks and physical media

Access networks: digital subscriber line (DSL)

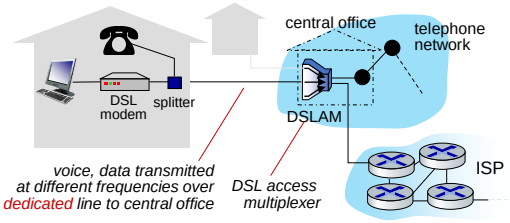

- use *existing* telephone line to central office DSLAM
  - data over DSL phone line goes to Internet
  - voice over DSL phone line goes to telephone net
- 24-52 Mbps dedicated downstream transmission rate
- 3.5-16 Mbps dedicated upstream transmission rate

DSL stands for Digital Subscriber Line.

It uses existing telephone lines.

Voice and data use different frequencies.

The DSLAM at the central office separates voice and Internet data.

DSL usually provides dedicated access to each home.

Cable is shared.
DSL is typically dedicated.

1.2.1-2 Access networks and physical media

Access networks: home networks

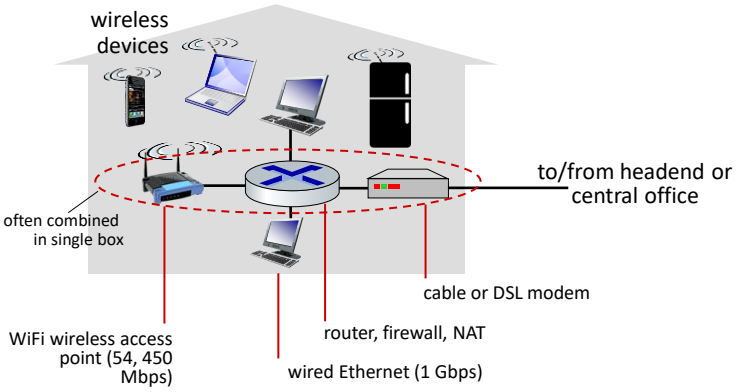

Inside your home, you usually have:

•A modem

•A router

•A firewall

•NAT

•Ethernet

•WiFi access point

Often, all of these are combined into one device.

How many of you have a WiFi router at home?

1.2.1-2 Access networks and physical media

Wireless access networks

Shared *wireless* access network connects end system to router

- via base station aka “access point”

Wireless local area networks
(WLANs)

- typically within or around building (~100 ft)
- 802.11b/g/n (WiFi): 11, 54, 450 Mbps transmission rate

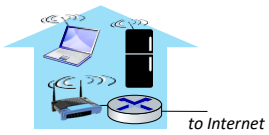

Wide-area cellular access networks

- provided by mobile, cellular network operator (10’s km)
- 10’s Mbps
- 4G cellular networks (5G coming)

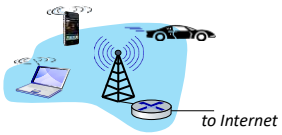

Wireless networks connect devices to routers through access points.

WiFi works within buildings.
Cellular networks cover wide areas.

WiFi is usually faster locally.
Cellular networks provide mobility.

1.2.1-2 Access networks and physical media

Access networks: enterprise networks

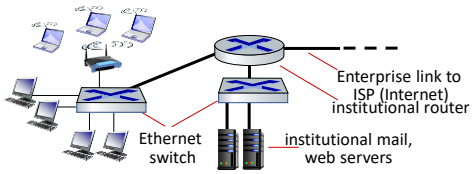

- companies, universities, etc.
- mix of wired, wireless link technologies, connecting a mix of switches and routers (we'll cover differences shortly)
  - Ethernet: wired access at 100Mbps, 1Gbps, 10Gbps
  - WiFi: wireless access points at 11, 54, 450 Mbps

Companies and universities use enterprise networks.

These networks combine:

•Ethernet switches

•Routers

•WiFi access points

They connect many users and servers internally
and then connect to an ISP.

1.2.1-2 Access networks and physical media

Host: sends *packets* of data

host sending function:

- takes application message
- breaks into smaller chunks, known as *packets*, of length *L* bits
- transmits packet into access network at *transmission rate R*
  - link transmission rate, aka link *capacity, aka link bandwidth*

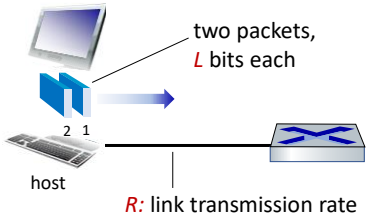

$$\text{packet transmission delay} = \text{time needed to transmit } L\text{-bit packet into link} = \frac{L \text{ (bits)}}{R \text{ (bits/sec)}}$$

When a host sends data, it:

1.Takes the application message

2.Breaks it into packets

3.Sends each packet into the access network

If a packet has L bits
and the transmission rate is R bits per second

The transmission delay is:
L divided by R.

Transmission delay equals packet length divided by link rate.

1.2.1-2 Access networks and physical media

Links: physical media

- bit: propagates between transmitter/receiver pairs
- physical link: what lies between transmitter & receiver
- guided media:
  - signals propagate in solid media: copper, fiber, coax
- unguided media:
  - signals propagate freely, e.g., radio

Twisted pair (TP)

- two insulated copper wires
  - Category 5: 100 Mbps, 1 Gbps Ethernet
  - Category 6: 10Gbps Ethernet

Physical links can be:

•Guided media

•Unguided media

Guided media includes:

•Copper

•Fiber

•Coaxial cable

Unguided media includes:

•Radio waves

1.2.1-2 Access networks and physical media

Links: physical media

Coaxial cable:

- two concentric copper conductors
- bidirectional
- broadband:
  - multiple frequency channels on cable
  - 100's Mbps per channel

Fiber optic cable:

- glass fiber carrying light pulses, each pulse a bit
- high-speed operation:
  - high-speed point-to-point transmission (10's-100's Gbps)
- low error rate:
  - repeaters spaced far apart
  - immune to electromagnetic noise

Coaxial cable has two copper conductors.
It supports broadband transmission.

Fiber optic cable uses light pulses.

It provides:

•Very high speed

•Very low error rate

•Long-distance transmission

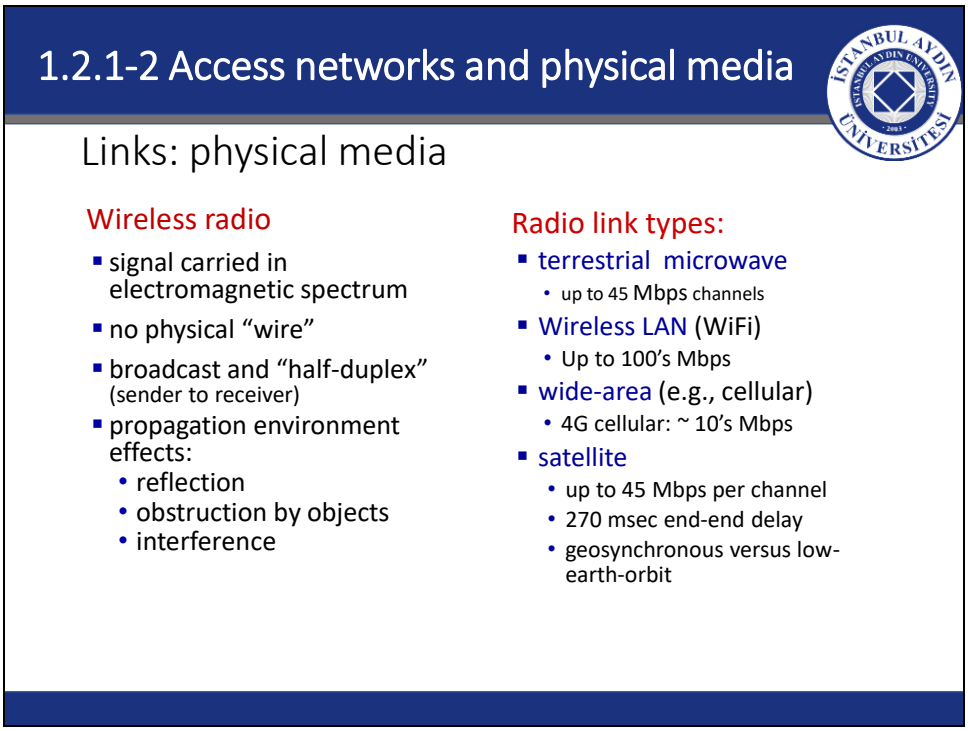

Wireless radio signals travel through electromagnetic spectrum.

There is no physical wire.

But wireless is affected by:

•Reflection

•Obstacles

•Interference

Satellite links have high delay
because the signal travels very long distances.

1.3 The Network Core

- mesh of interconnected routers
- packet-switching: hosts break application-layer messages into *packets*
  - forward packets from one router to the next, across links on path from source to destination
  - each packet transmitted at full link capacity

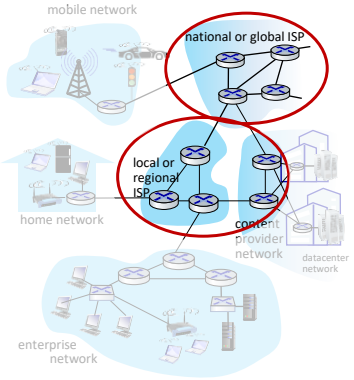

Now we move from the network edge to the **network core**.
The network core is the **mesh of interconnected routers and links** that connects all end systems.

The key idea here is **packet switching**.
Hosts break application-layer messages into smaller pieces called **packets**.
Then routers forward these packets from one router to the next, across links, until they reach the destination.

In packet switching, each packet is transmitted at the **full link capacity** on each link it uses.

So the core's job is simple to say, but complex to do:
move packets efficiently and reliably across the Internet.

1.3.1 Packet Switching

Packet-switching: store-and-forward

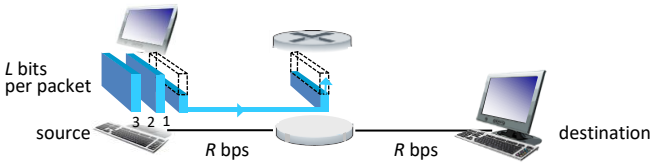

- Transmission delay: takes *L*/*R* seconds to transmit (push out) *L*-bit packet into link at *R* bps
- *Store and forward:* *entire* packet must arrive at router before it can be transmitted on next link
- *End-end delay:* 2*L*/*R* (above), assuming zero propagation delay (more on delay shortly)

One-hop numerical example:

- *L* = 10 Kbits
- *R* = 100 Mbps
- one-hop transmission delay = 0.1 msec

Let's understand a very important concept: **store-and-forward**.

First, transmission delay:
If a packet has length **L bits** and the link rate is **R bits per second**,
the time to push the packet into the link is **L divided by R** seconds.

Store-and-forward means the router must receive the **entire packet** before it can send it on the next link.

In a simple one-router path, ignoring propagation delay,
the end-to-end transmission delay becomes **2L/R**:
L/R to reach the router, and L/R to reach the destination.

The key message is:
routers do not forward "half packets"—they typically wait for the full packet first.

1.3.1 Packet Switching

Packet-switching: queueing delay, loss

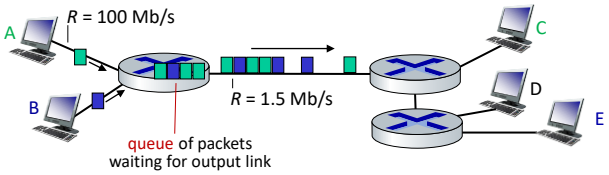

Packet queuing and loss: if arrival rate (in bps) to link exceeds transmission rate (bps) of link for a period of time:

- packets will queue, waiting to be transmitted on output link
- packets can be dropped (lost) if memory (buffer) in router fills up

Now we add a real-world effect: **queueing**.

Routers have output buffers, also called queues.
If packets arrive faster than the link can send them, they must wait in the queue.

In this figure, hosts send data at high speed, but the outgoing link is slower.
So packets build up in the queue.

If the buffer becomes full, packets are **dropped**.
That is packet loss.

A good analogy is traffic:
if cars arrive faster than the road can handle, a traffic jam forms;
if the road is completely blocked, some cars must take another action.

So packet switching gives flexibility, but it can also create delay and loss when there is congestion.

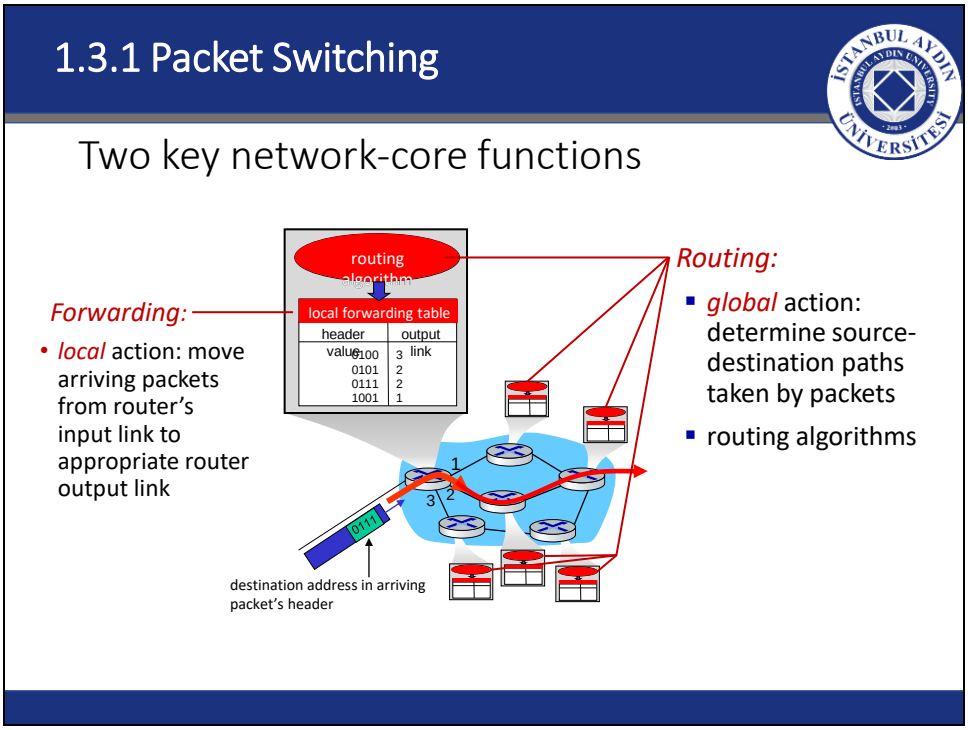

The network core has two key functions: **forwarding** and **routing**.

Forwarding is a local action inside one router.
The router looks at the destination address in the packet header and sends the packet to the correct output link using a forwarding table.

Routing is a global action.
It decides the end-to-end paths that packets take from source to destination. Routing algorithms compute these paths and help routers build their forwarding tables.

So in simple words:
Routing chooses the road map;
forwarding moves the packet at each intersection.

1.3.1 Packet Switching

Alternative to packet switching: circuit switching

end-end resources allocated to, reserved for "call" between source and destination

- in diagram, each link has four circuits.
  - call gets 2nd circuit in top link and 1st circuit in right link.
- dedicated resources: no sharing
  - circuit-like (guaranteed) performance
- circuit segment idle if not used by call (no sharing)
- commonly used in traditional telephone networks

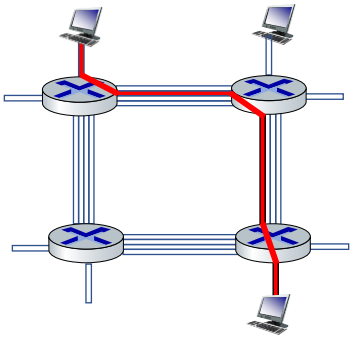

Packet switching is the main idea of the Internet,
but there is another approach called **circuit switching**.

In circuit switching, end-to-end resources are **reserved** for a call or session. This means a fixed part of the link capacity is dedicated to you for the whole communication.

The advantage is predictable performance.
The disadvantage is inefficiency: if you are silent or not sending data, that reserved capacity is wasted.

Traditional telephone networks are classic examples of circuit switching.

1.3.2 Circuit Switching

Circuit switching: FDM and TDM

Frequency Division Multiplexing (FDM)

- optical, electromagnetic frequencies divided into (narrow) frequency bands
- each call allocated its own band, can transmit at max rate of that narrow band

Time Division Multiplexing (TDM)

- time divided into slots
- each call allocated periodic slot(s), can transmit at maximum rate of (wider) frequency band, but only during its time slot(s)

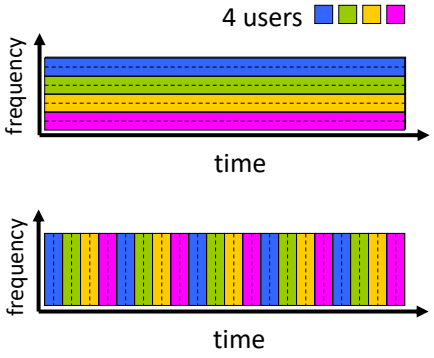

Now we will briefly go deeper into circuit switching,
mainly to compare it with packet switching.
In circuit switching, the link resources can be divided in two common ways: **FDM** and **TDM**.

With **FDM**, the frequency spectrum is divided into bands.
Each call gets its own band continuously.

With **TDM**, time is divided into repeating slots.
Each call gets a time slot in each frame.

So FDM shares by frequency, and TDM shares by time.

1.3.2 Circuit Switching

Packet switching versus circuit switching

packet switching allows more users to use network!

Example:

- 1 Gb/s link
- each user:
  - 100 Mb/s when "active"
  - active 10% of time

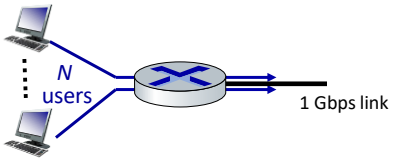

- *circuit-switching:* 10 users
- *packet switching:* with 35 users, probability > 10 active at same time is less than .0004 *

Q: how did we get value 0.0004?

Q: what happens if > 35 users ?

* Check out the online interactive exercises for more examples: http://gaia.cs.umass.edu/kurose_ross/interactive

Here we see an important argument for packet switching.

In real life, many users are not active all the time.
People send data in bursts: sometimes active, sometimes idle.

Packet switching can statistically support more users because it allocates capacity on demand.

The main idea to remember is:
packet switching is efficient for bursty traffic,
but it can also create congestion.

1.3.2 Circuit Switching

Packet switching versus circuit switching

Is packet switching a "slam dunk winner"?

- great for "bursty" data – sometimes has data to send, but at other times not
  - resource sharing
  - simpler, no call setup
- excessive congestion possible: packet delay and loss due to buffer overflow
  - protocols needed for reliable data transfer, congestion control
- *Q:* How to provide circuit-like behavior?
  - bandwidth guarantees traditionally used for audio/video applications

Q: human analogies of reserved resources (circuit switching) versus on-demand allocation (packet switching)?

So, is packet switching always better? Not always.

Packet switching is great because:

•it shares resources efficiently

•it is simpler, no call setup

But if the network becomes congested, we can get:

•packet delay

•packet loss

That is why networking needs protocols for:

•reliable data transfer

•congestion control

Later in the course, we will learn how TCP helps provide reliability and control congestion.