University of Illinois at Urbana-Champaign
Dept. of Electrical and Computer Engineering

ECE 101: Exploring Digital Information Technologies for Non-Engineers

Distribution and Streaming

Everyone Connected at the Edge

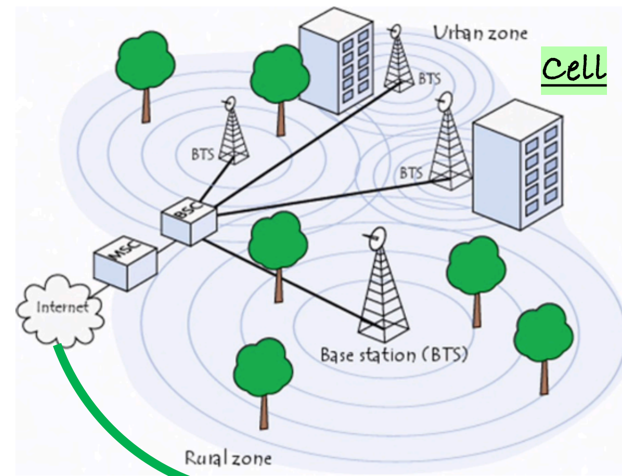

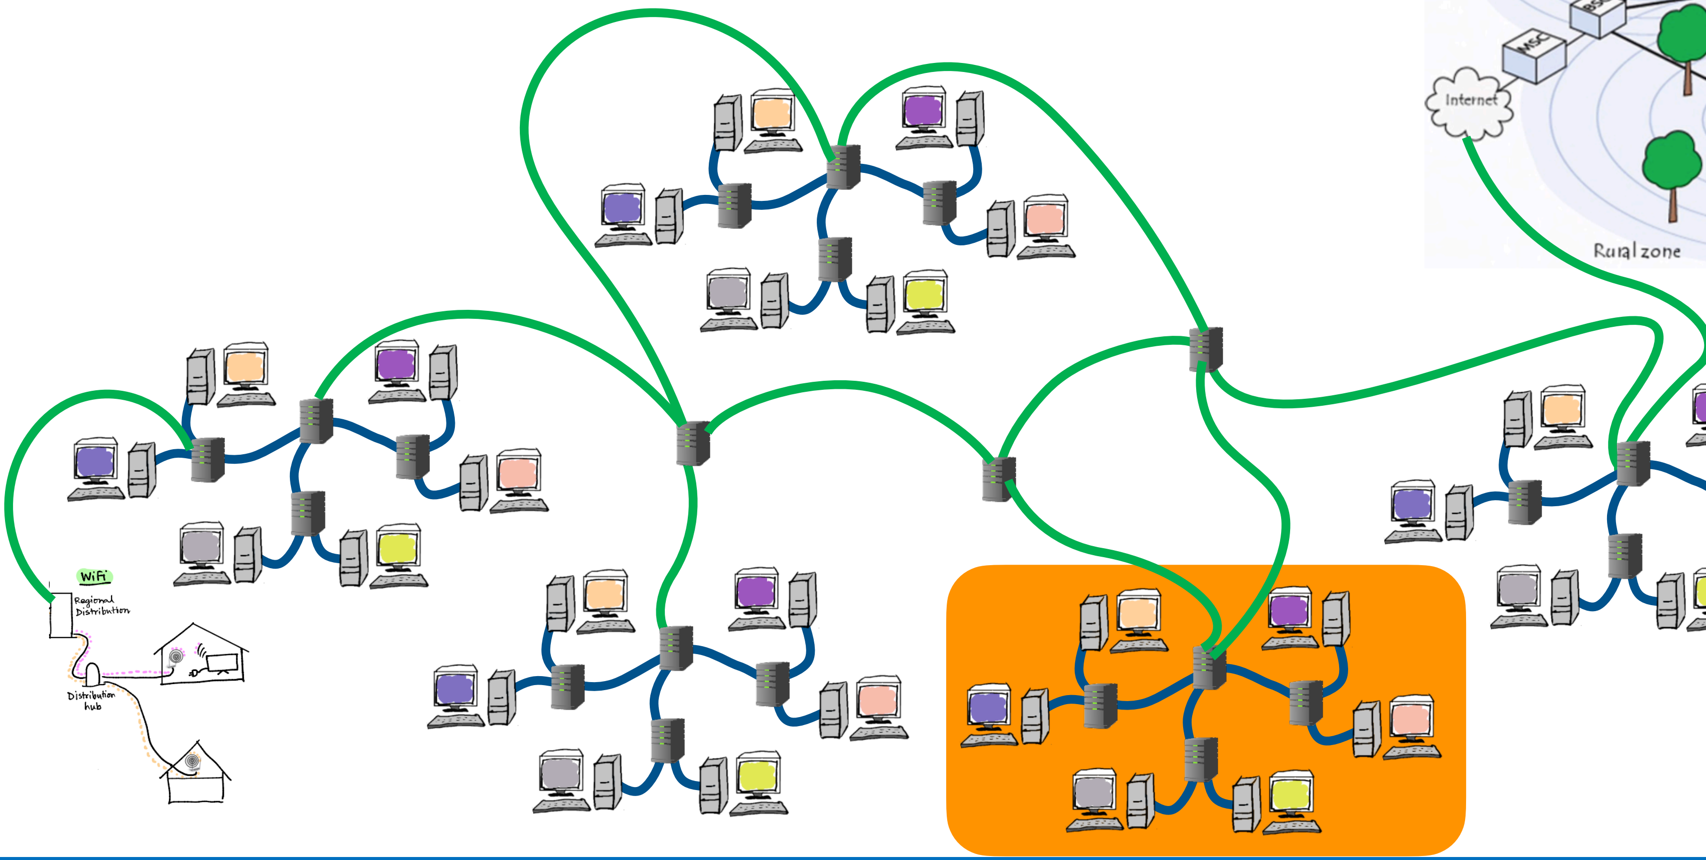

How does Netflix get to your screen

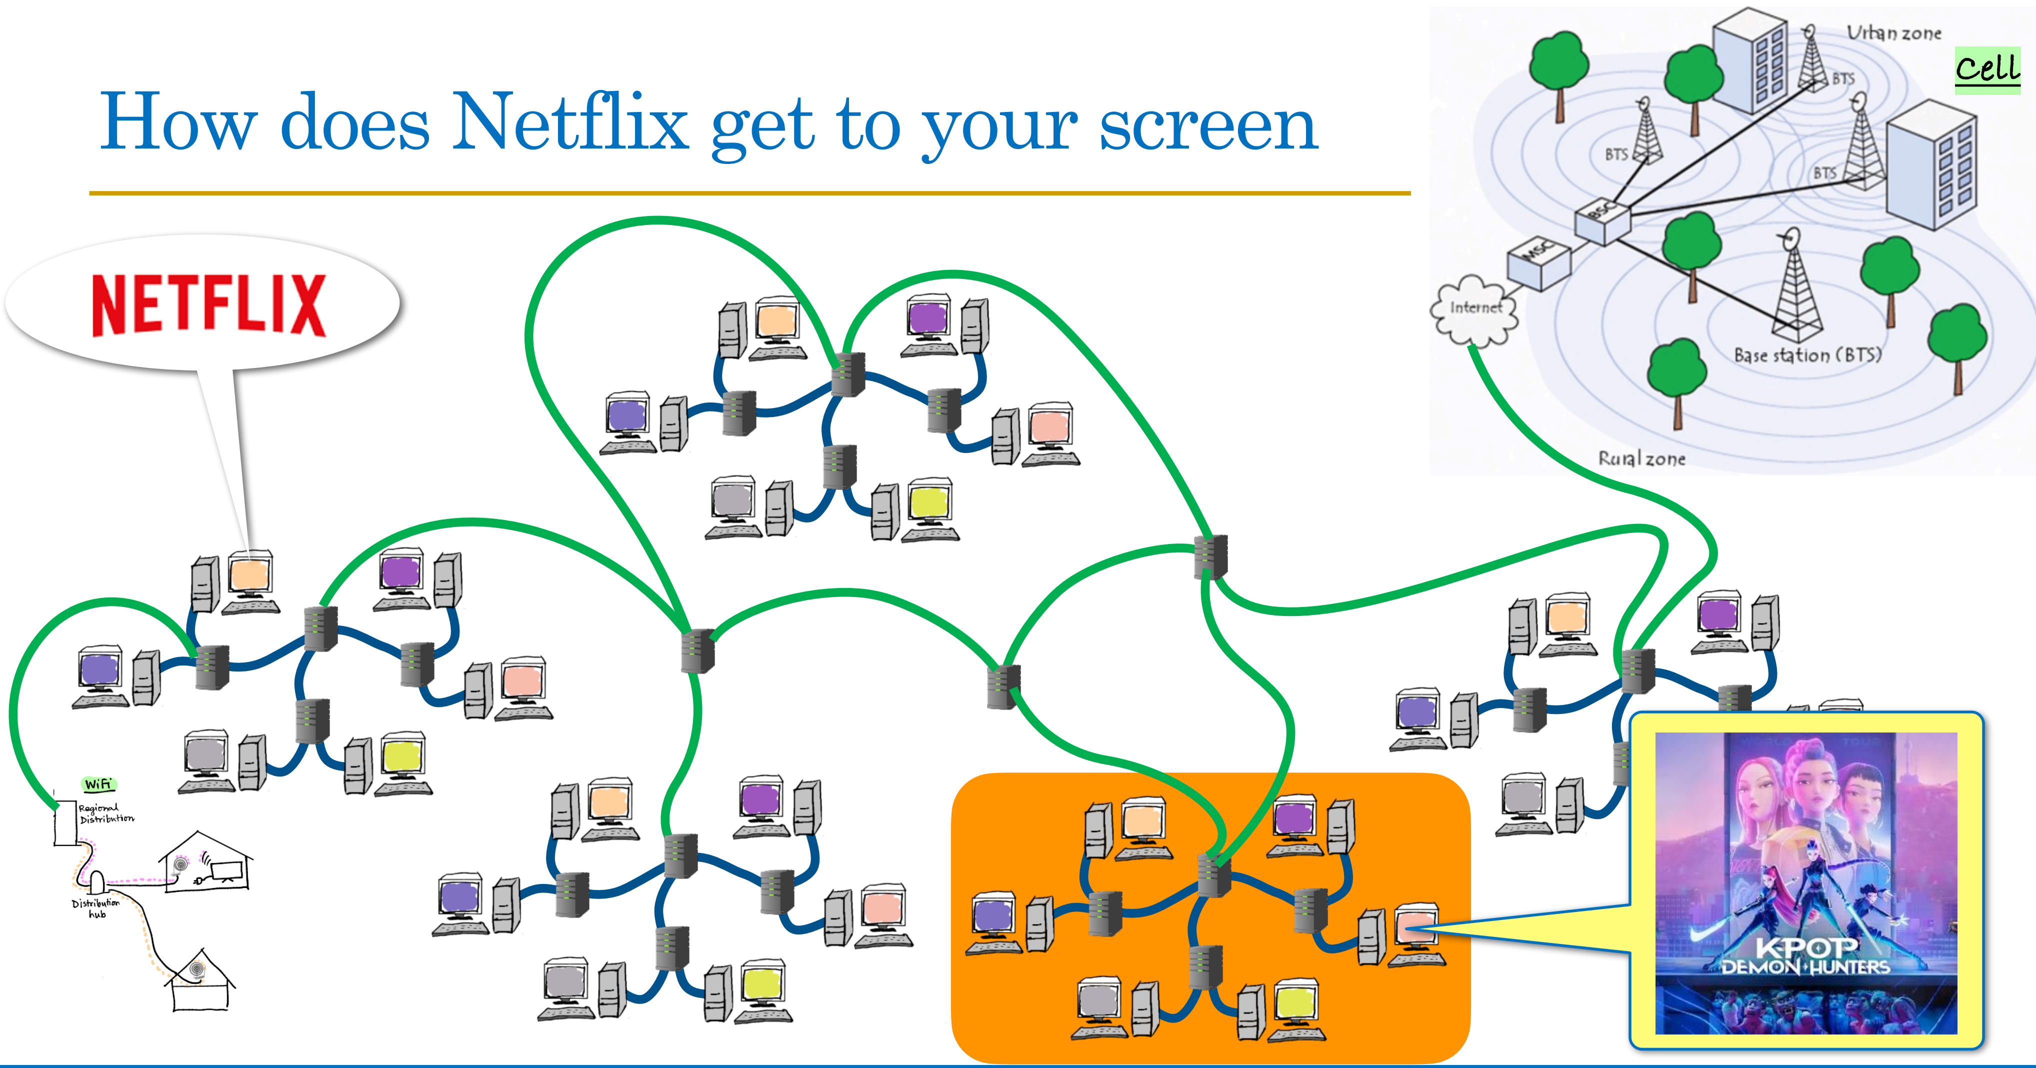

Remember Packets of Digital Data

Binary Foundation

Digital data consists of bits and bytes, representing all forms of online content using 1s and 0s.

Variety of Digital Content

Digital content includes text messages, emails, documents, photos, songs, movies, Zoom meetings, etc.

Data Size Examples

A single emoji in a text might be a few kilobytes, while HD movies can be several gigabytes in size.

Data Packets

Protocols like TCP-IP enable packaging bits and bytes into data packets and their seamless transmission to render digital content across platforms.

How Data Travels - the Internet as Plumbing

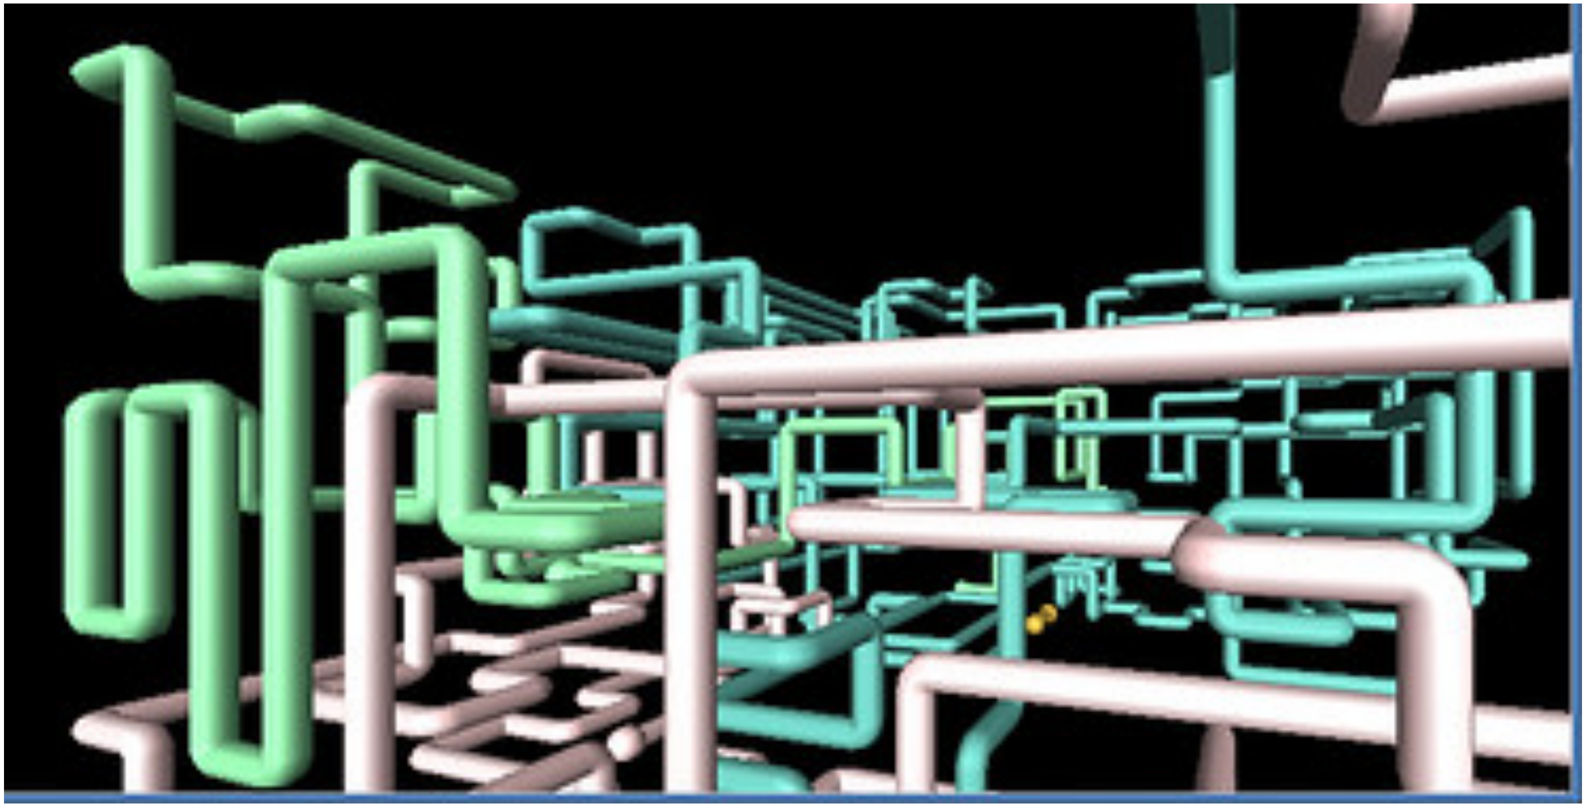

https://www.flickr.com/photos/jvmanna/2844048090

How Data Travels - One Individual Packet at a Time

Data Packets

Online content is divided into small packets for easier transmission across the internet.

Network Routing

Packets travel independently through routers and servers, similar to cars navigating highways.

Packet Reassembly

Packets are reassembled on the receiving device to reconstruct the original content like jigsaw puzzle.

Efficient Data Delivery

Packet-switching enables efficient, reliable transmission despite different routes taken by packets.

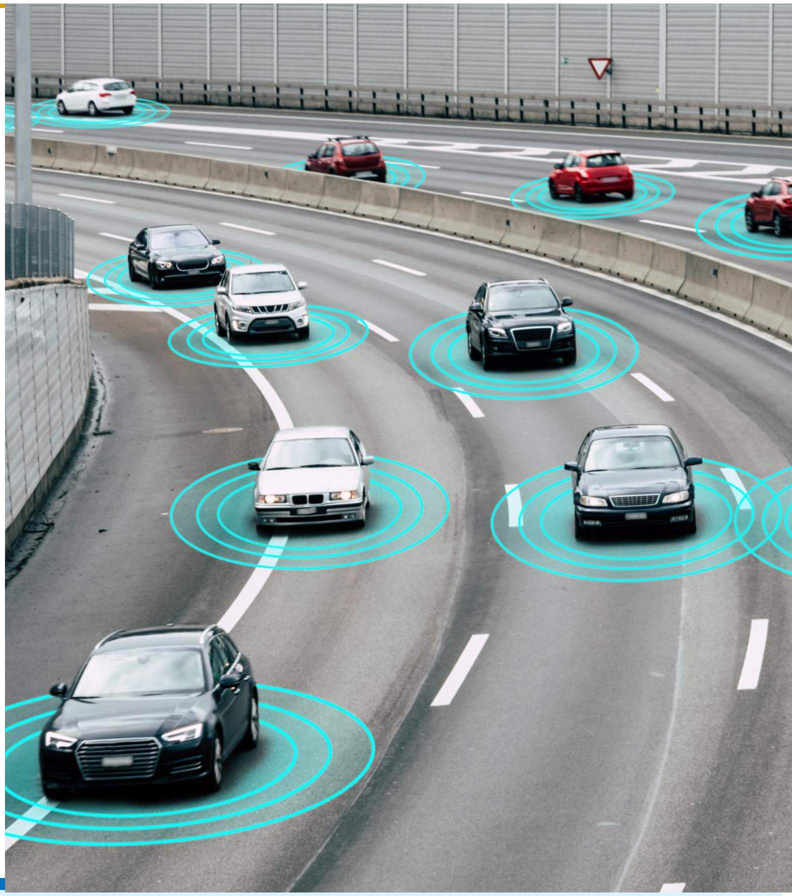

Network Links Can Be Viewed as Pipes for Data Packets

Imagine a pipe that carries water.

Now imagine that we have
- **10 gallons per minute**
- **flowing** through the pipe.

Next, add a **second pipe**
- **with** the **same flow** rate
- and put it above the first pipe.

No problem, right?

Water from the top pipe goes into the bottom pipe.

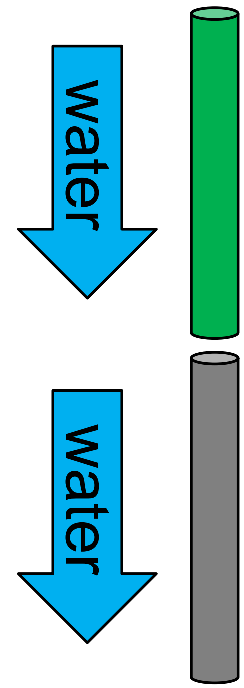

What Happens When a Pipe Can't Handle the Water?

Now let's **add a third pipe…**
with the same flow of water.

Now what happens?

Oops! **Water spills** everywhere!

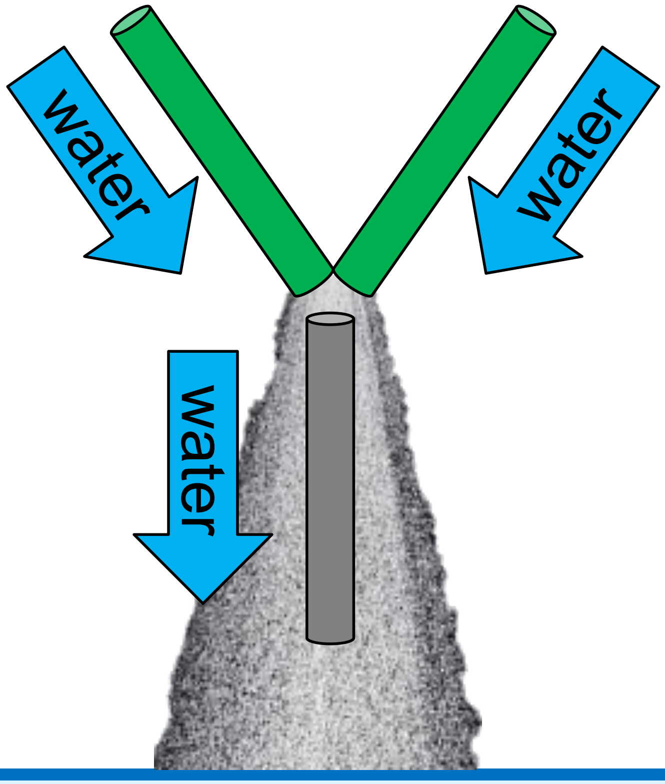

Similarly, Internet Routers May Discard Packets

The **same thing happens**

- **when** an Internet **router**
- **receives too many packets**
- for one outgoing link.

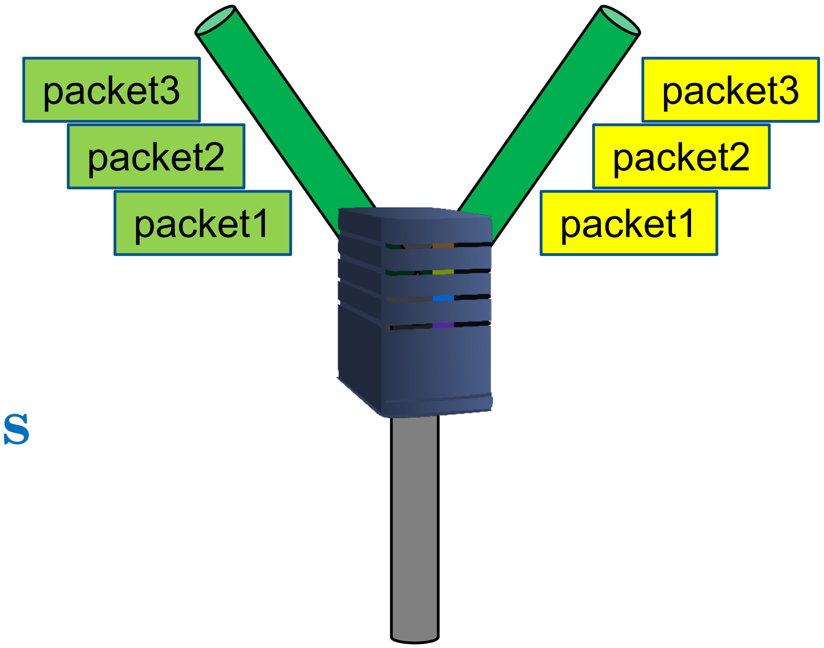

Remember that Internet Only TRIES to Deliver a Packet

That's one reason that
- Internet **packet delivery**
- **is unreliable**.

So that routers
- can discard packets
- without causing problems.

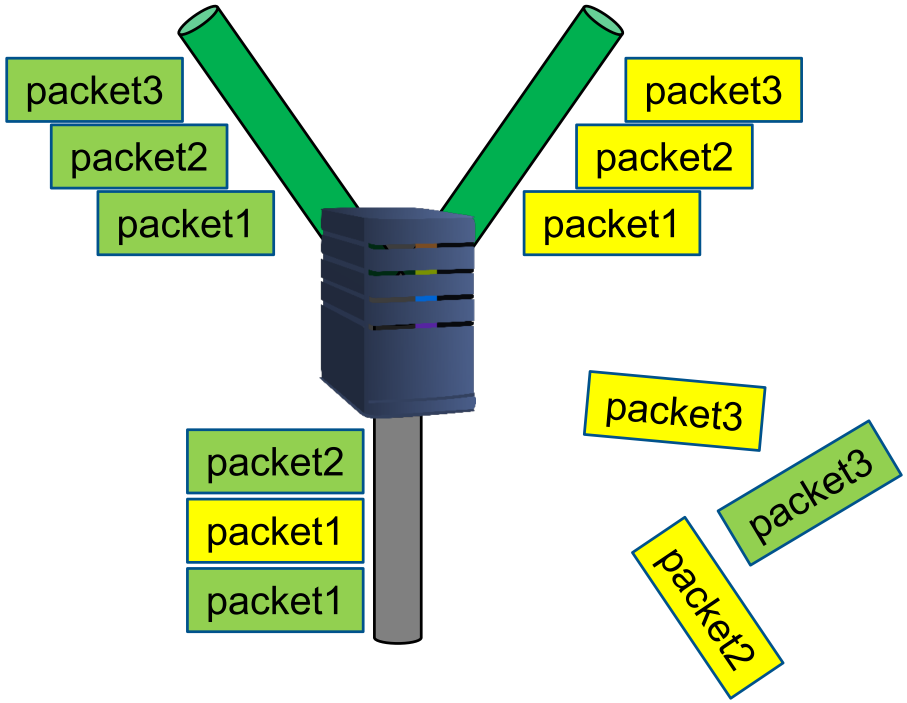

Internet Service Definition Does Not Suit All Needs

A brief aside…

Unreliability raises some major issues that we won't cover.

For example, the **Internet**

- **does NOT make a good substitute for 911**
- or any other campus/state/country's emergency telephone services.

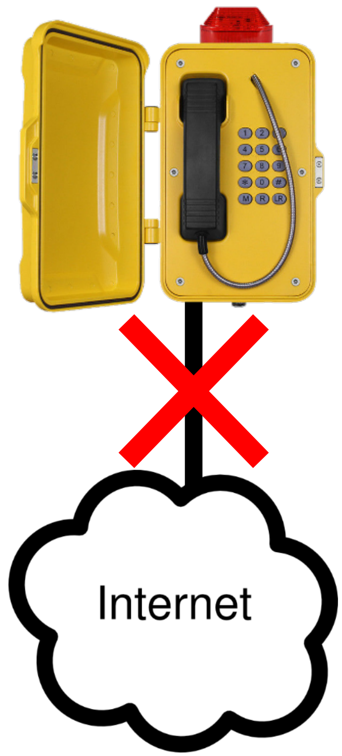

How Quickly Can a Client Communicate with a Server?

How much data can a client send through the Internet to a server each second?

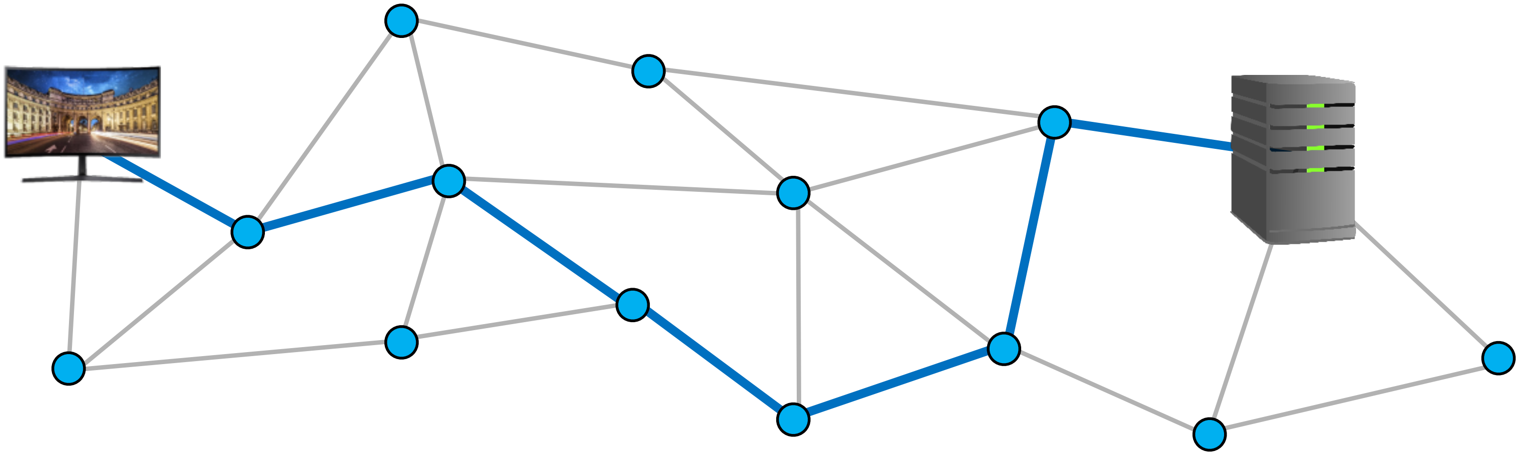

How Quickly Can a Client Communicate with a Server?

The answer changes constantly
as other computers (TCP connections)
use the Internet.

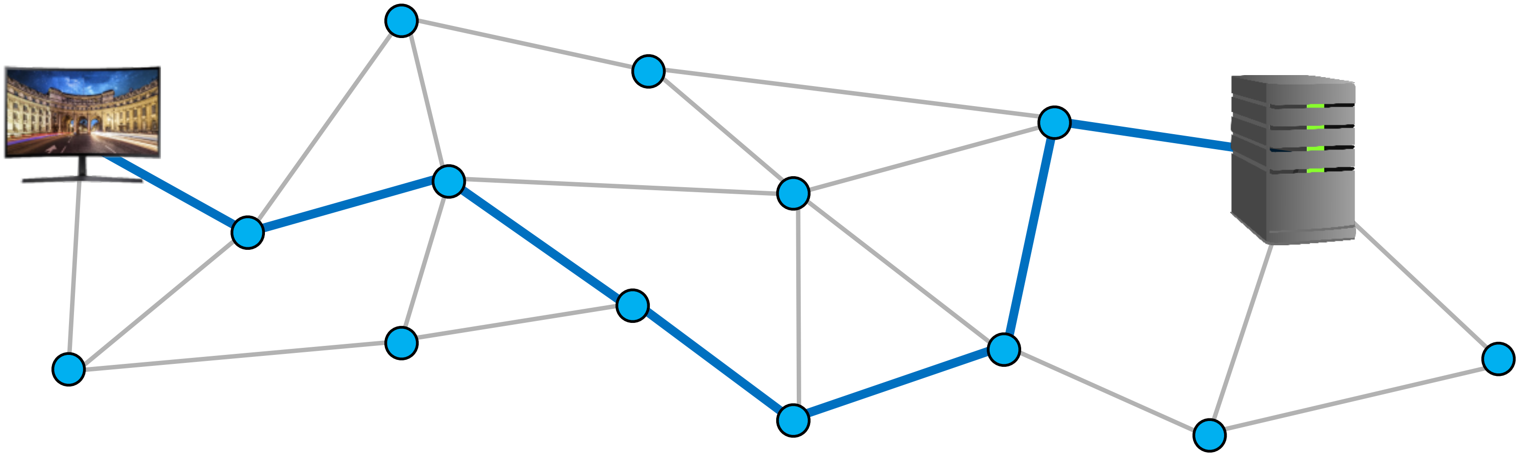

TCP Tries to Estimate How Much Data to Send

TCP constantly tries
- to **estimate** the **achievable data rate**
- **in** bits or **bytes per second**,
- **called bandwidth**.*

*Electrical engineers object to this usage because the term is also used to measure ranges of frequency.

Information moves at nearly the speed of light:

~**130 msec** to the other side of the Earth

- **0-5 Mbps:** Basic tasks like checking email or streaming music on one device.
- **5-40 Mbps:** Standard video streaming, social media, and other basic tasks.
- **40-100 Mbps:** HD video streaming, gaming, video conferencing, and downloading large files on a few devices.
  (Compressed 4K movie ~20 GB, will take 3 mat to download on 100Mbps)

TCP Tries to **Estimate** How Much Data to Send

More delay means a worse estimate!!!

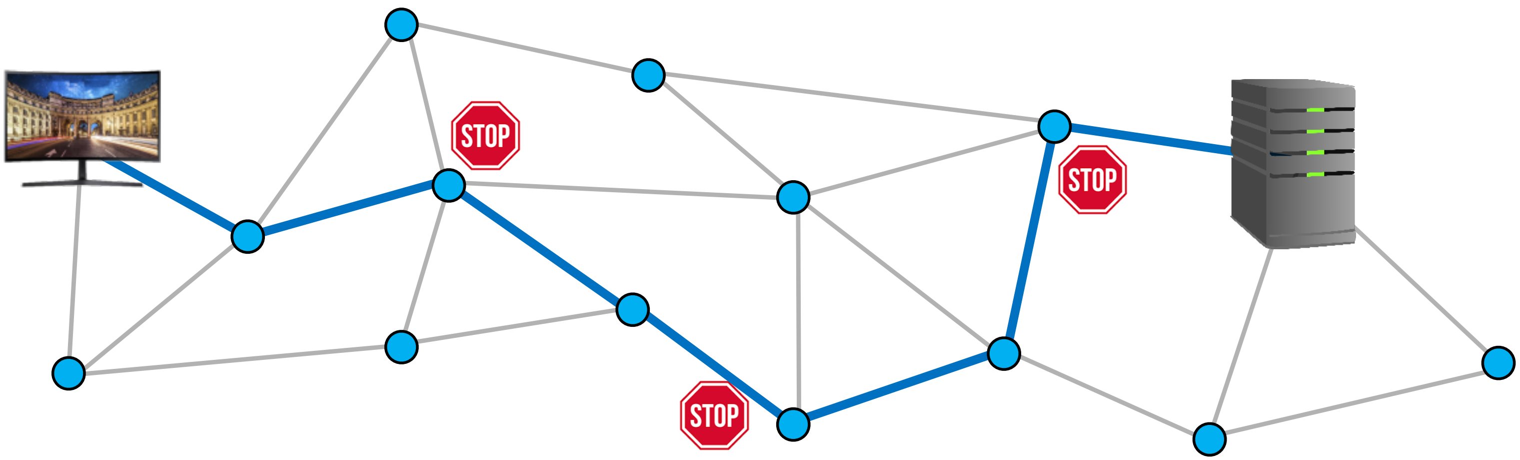

Let's Quantify a Single Transmission

Let's **quantify a simple problem**
with our friends Jan and Pat.

Jan wants to **send** Pat a **500 MB** video.

The route between them goes **over three links**.

Each link has limited **bandwidth**.

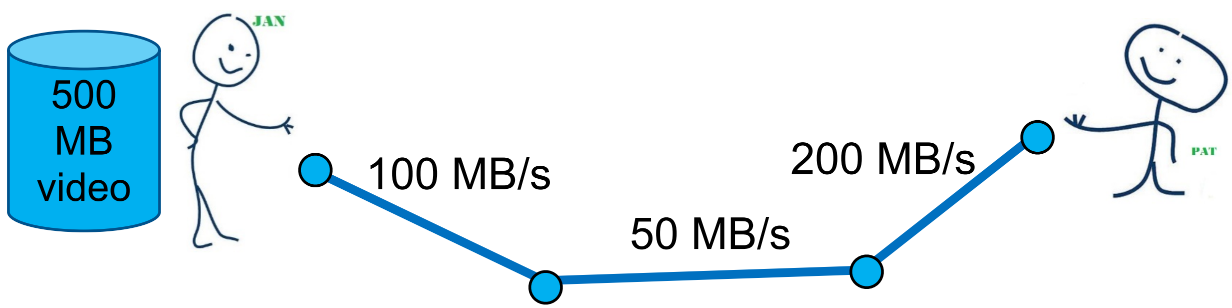

Let's Quantify a Single Transmission

What is the bandwidth from Jan to Pat?

50 MB/s – the **minimum of** the **link bandwidths** on the route through the Internet

So how long does sending the video take?

$$\frac{500\text{ MB}}{50\text{ MB/s}} = \textbf{10 seconds}$$

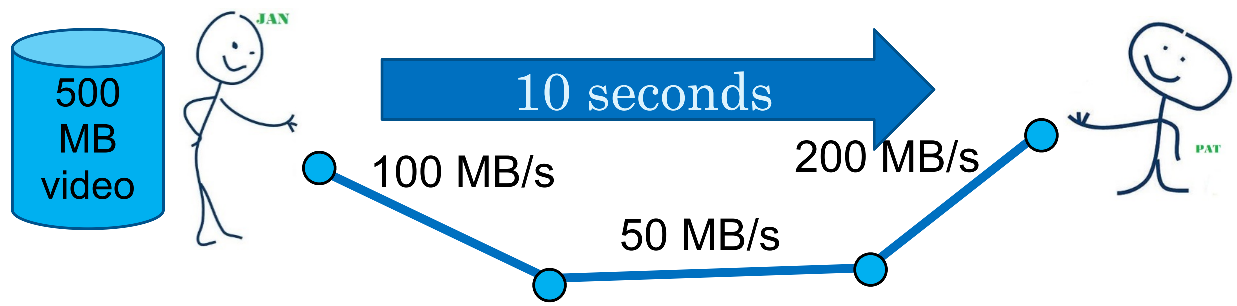

Connections Share Links in the Internet

What happens if Alice also sends Bob a video?

The **two transmissions** must **share** the middle link!

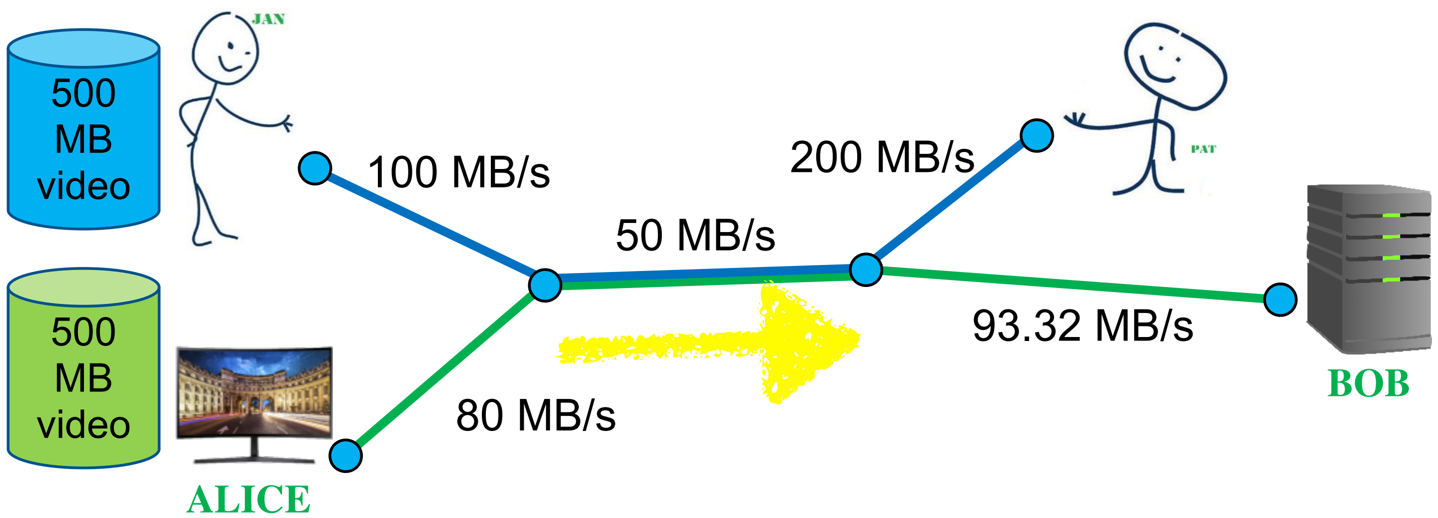

Each Connection May Use an Equal Share

If they share fairly, what is the link bandwidth for each connection?

$$\frac{50\ \text{MB/s}}{2\ \text{connections}} = \mathbf{25\ MB/s}$$

for each connection

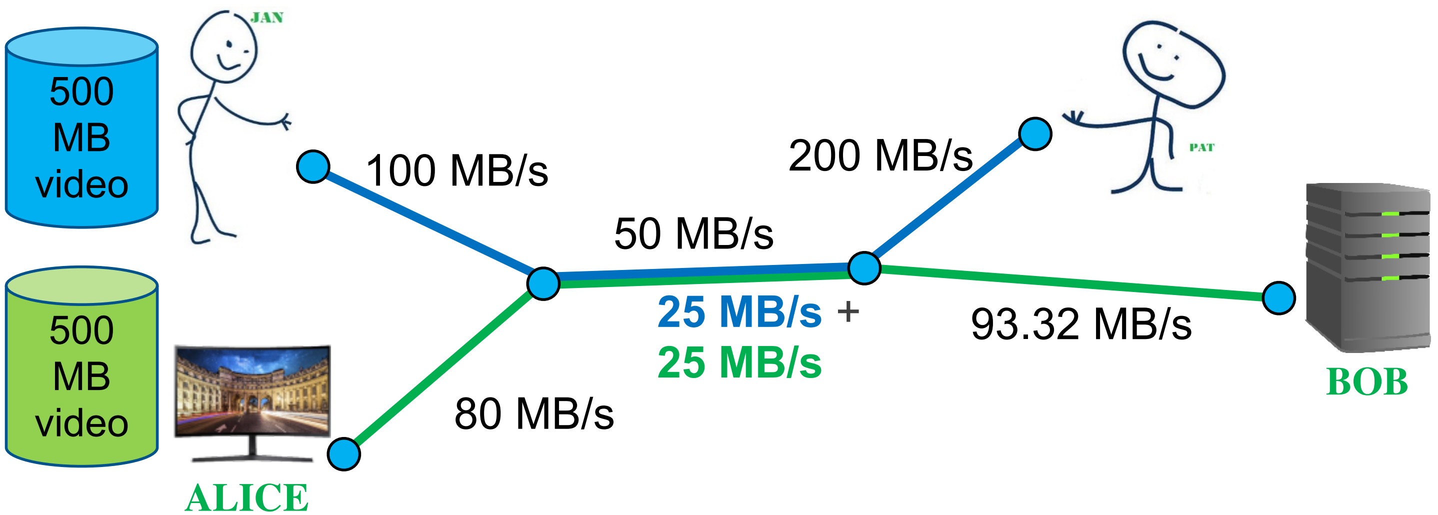

Shared Links Slow Down Transfers

What is Jan's connection bandwidth? **25 MB/s**

What about Alice's? **25 MB/s**

And how long to send either video (same size, same bandwidth)?

$$\frac{500\ \text{MB}}{25\ \text{MB/s}} = \mathbf{20\ seconds}$$

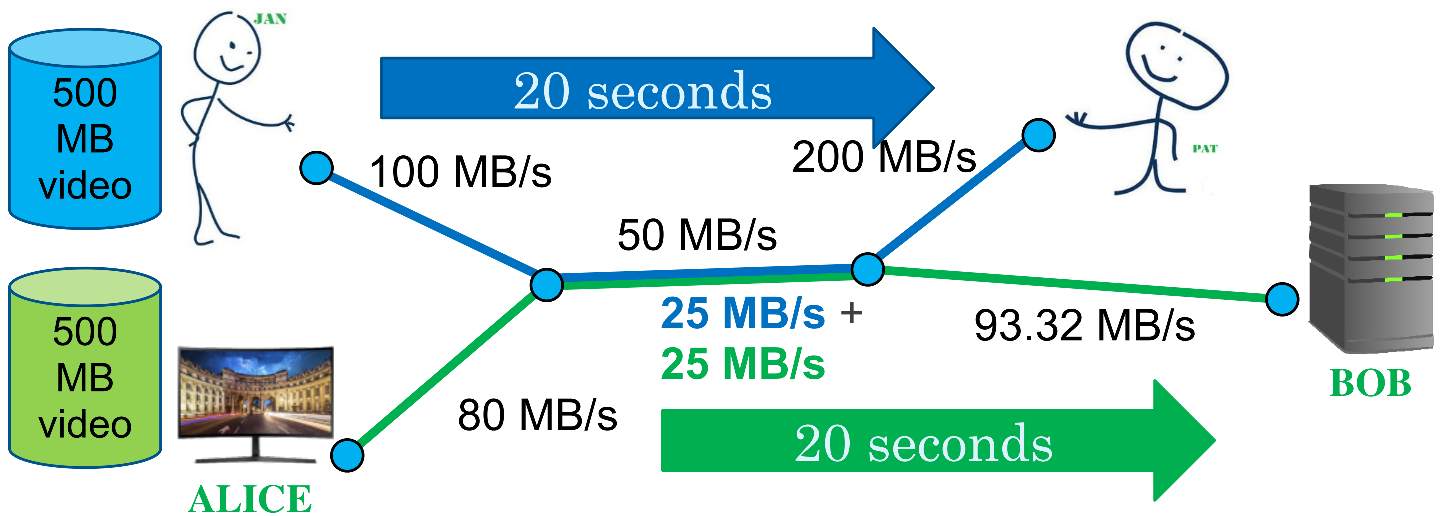

Another Scenario: a Low-Bandwidth Connection

Let's change the situation:
- Pat is now using his phone.
- The last Jan-to-Pat link carries only **10 MB/s**.

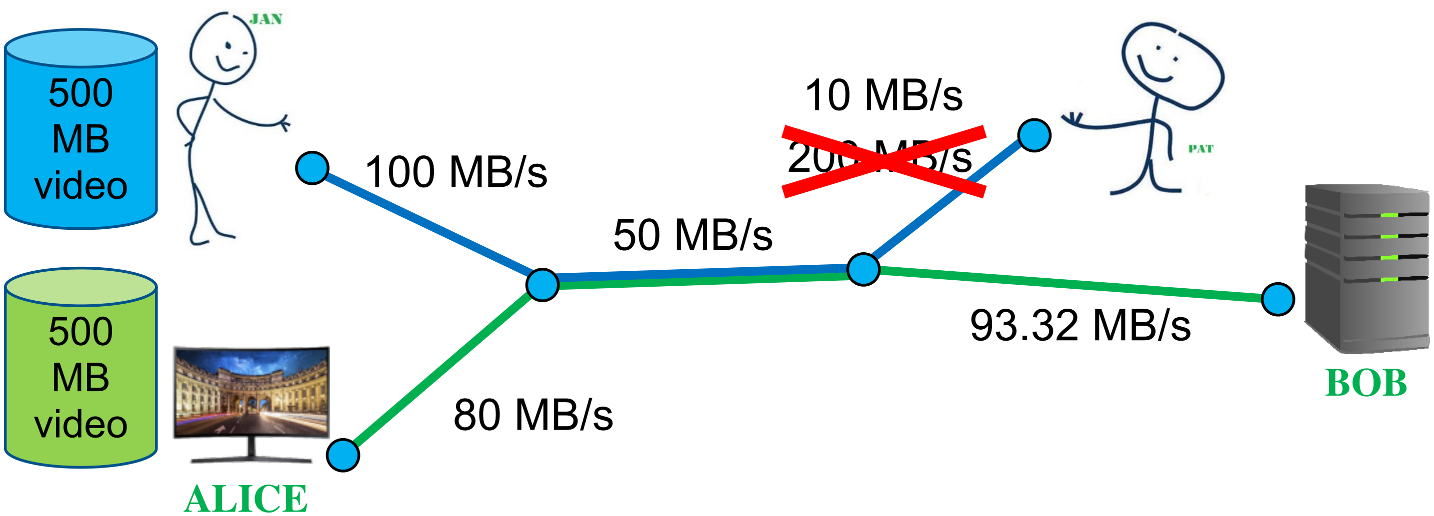

Another Scenario: a Low-Bandwidth Connection

What is Jan's connection bandwidth? **10 MB/s**

Alice can use the rest of the middle link's bandwidth!

So what is Alice's connection bandwidth? **40 MB/s**

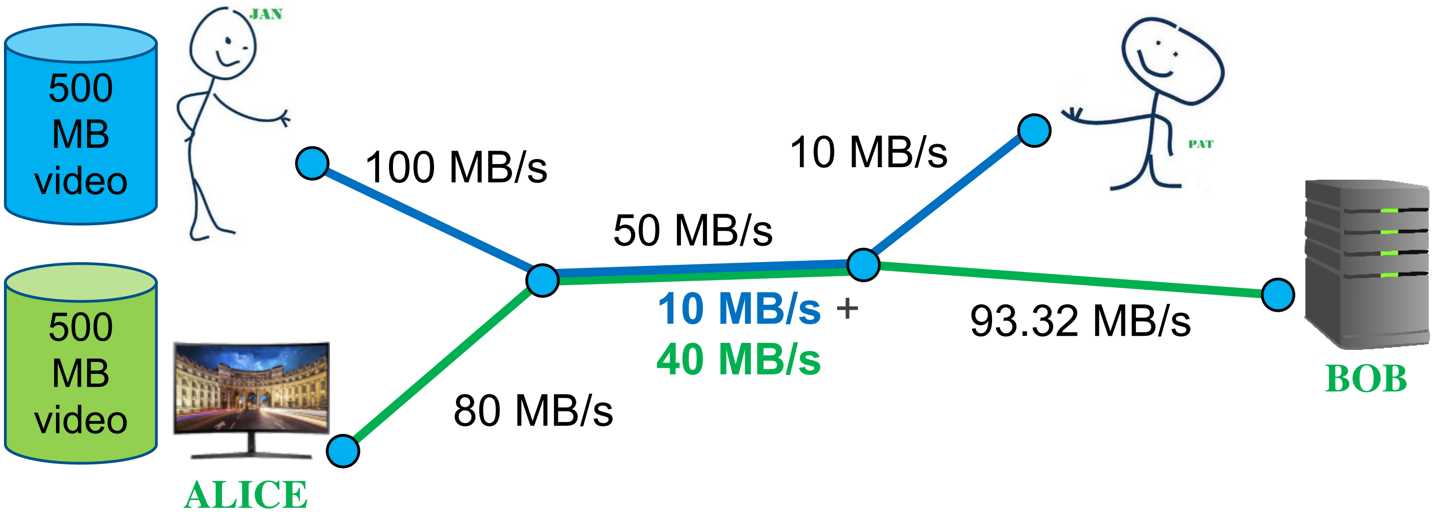

Another Scenario: a Low-Bandwidth Connection

How long does Jan's transfer require? $\frac{500\text{ MB}}{10\text{ MB/s}}$ **=50 seconds**

What about Alice's transfer? $\frac{500\text{ MB}}{40\text{ MB/s}}$ **=12.5 seconds**

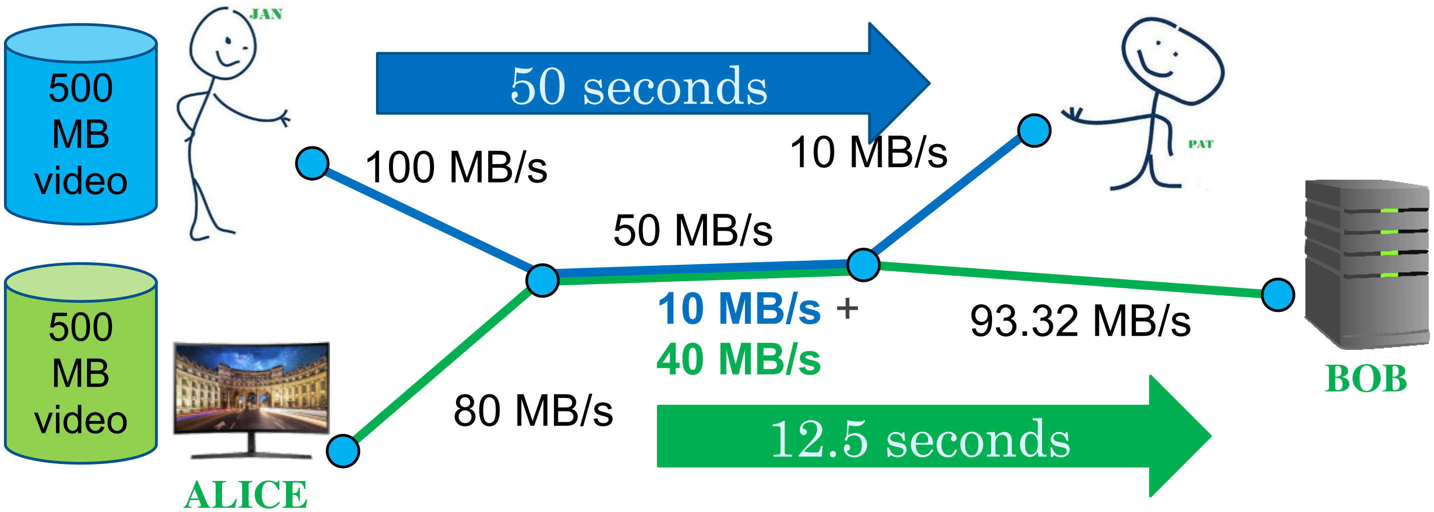

Streaming vs Downloading

Streaming Explained

Streaming delivers digital content in real-time, allowing users to access it instantly without full download.

Downloading Explained

Downloading saves the entire file locally before access, enabling offline use without internet connection.

User Experience Comparison

Streaming suits instant access and limited storage, while downloading is ideal for offline access and full control.

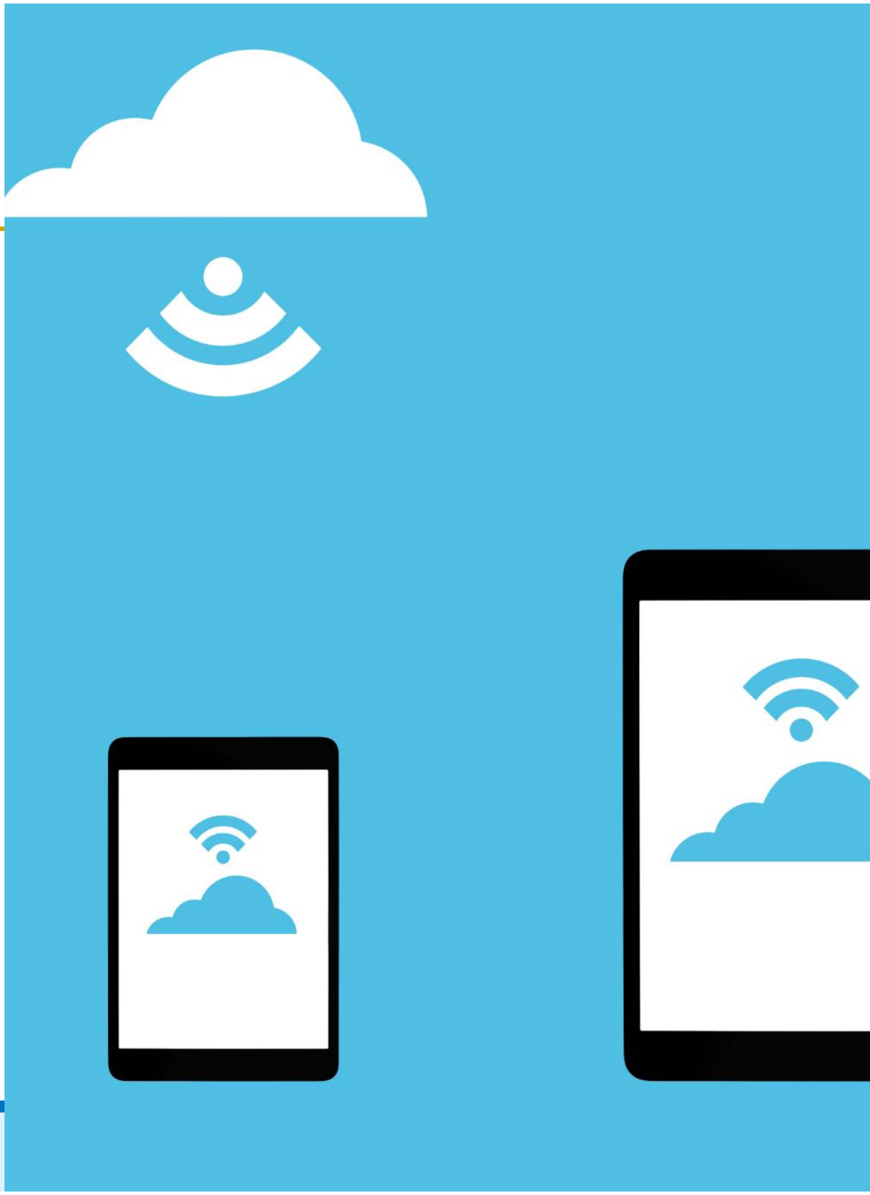

Class Participation Question: Answer on Canvas

Jan wants to **send** Pat a **2 GB** video.

How long does it take for Jan to send the video

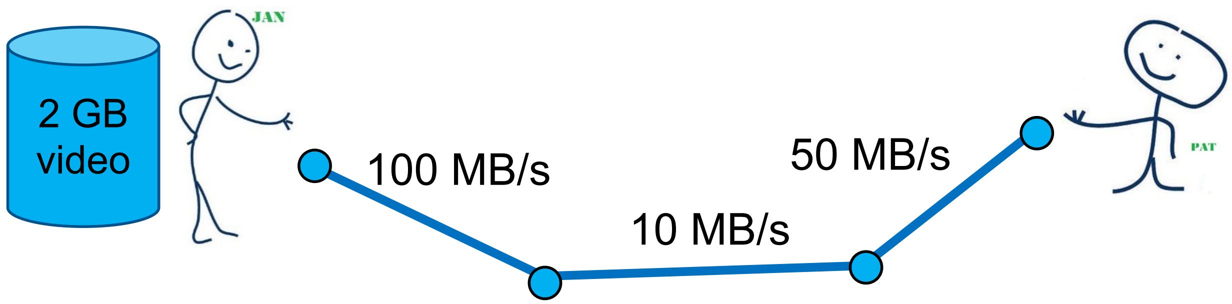

Both Video and Network Rates Vary Over Time

The (high) bandwidths in our example

- are more than enough
- to deliver a video in real-time.

However,

- the **rate of data needed** by a video **varies**,
- and the **bandwidth of** the **network varies**.

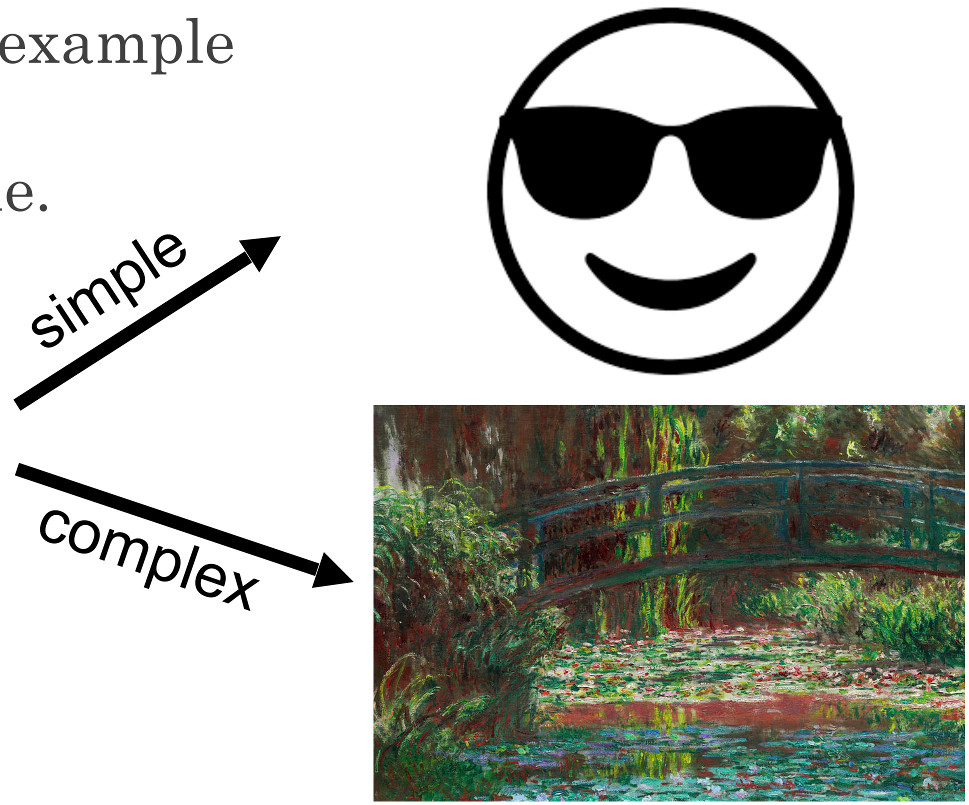

Streaming Technology and Performance

Things done to help streaming

- Compression
- Buffering
- CDNs

Compression (First week of class)

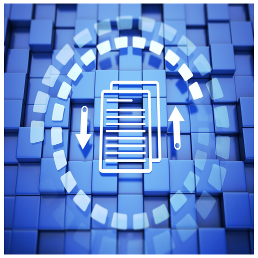

Definition of Compression

Compression reduces digital file sizes for faster transmission and easier storage without losing important data.

Example of Video Compression

A raw 10 GB video file compressed to 1 GB MP4 format, significantly reducing file size with minimal quality loss.

Compression Formats

Formats like MP4 for video and MP3 for audio keep quality while making files smaller, ideal for streaming and mobile use.

Buffering Helps Avoid Need for Video Stalls

To handle variability, video players **use** a technique called **buffering**.

- Before the video starts playing, your computer downloads the first second of video.
- While the first second plays, the computer downloads the second second.
- And so forth.

Buffering happens when data is not received fast enough to keep video playback smooth and continuous.

Buffering Can't Hide Inadequate Bandwidth

Often, a video player starts downloading before you press "Play."
That way, the video seems to start instantaneously.

But **you've probably noticed occasional delays**
- for ads, or even for videos,
- **when** network and/or video content **variability**
- **or insufficient** network **bandwidth**
- **made your computer run out of video** to show you!

Slow internet speed, network congestion, and server delays are primary causes.

CDN: Content Delivery Networks

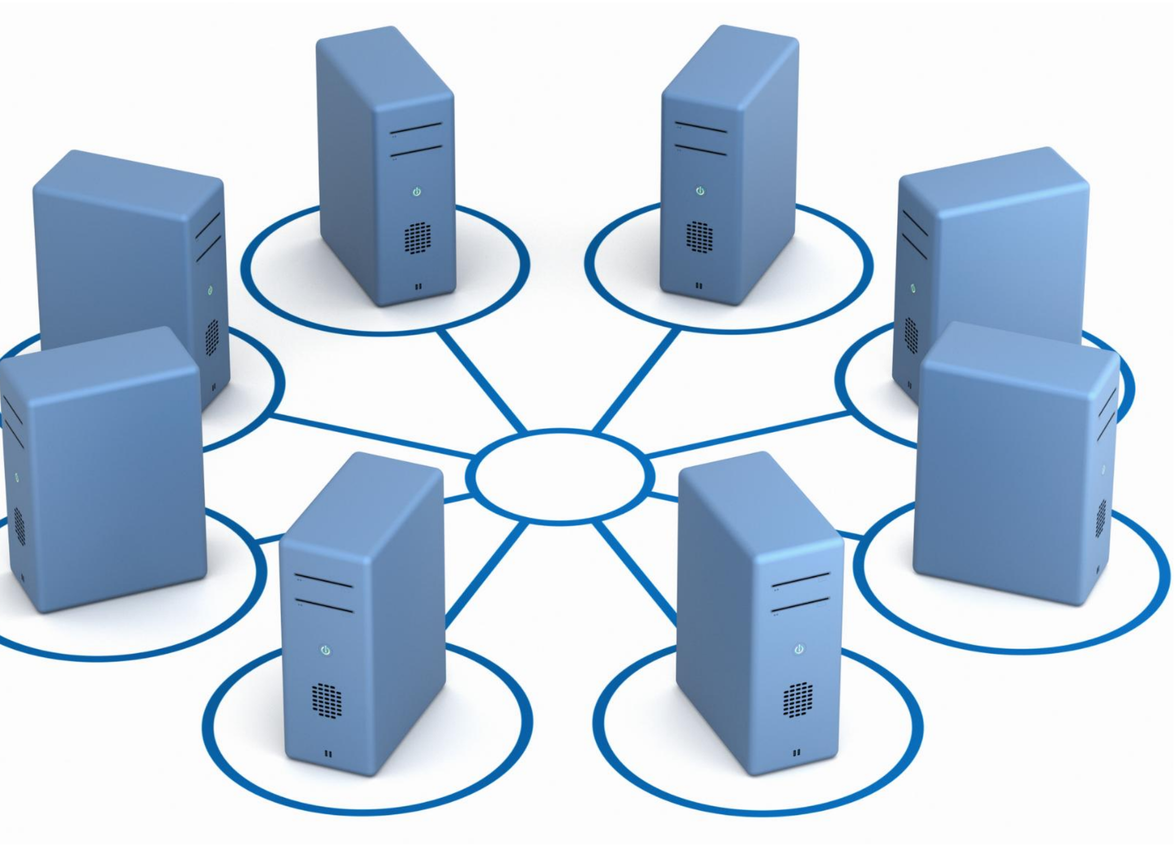

Distributed Server Network

CDNs use distributed servers across various locations to store copies of content close to users.

Reduced Latency and Faster Delivery

Reducing data travel distance speeds up content delivery and lowers latency for users.

Improved User Experience

CDNs ensure fast, reliable access to digital content even during peak usage times.

Locality is Also Important for Content Distribution

Locality plays an important role:

- TCP's achieved bandwidth drops with increased delay, so
- moving content closer to users improves their experience, and
- not sharing heavily-loaded backbone links improves end-to-end bandwidth for everyone.

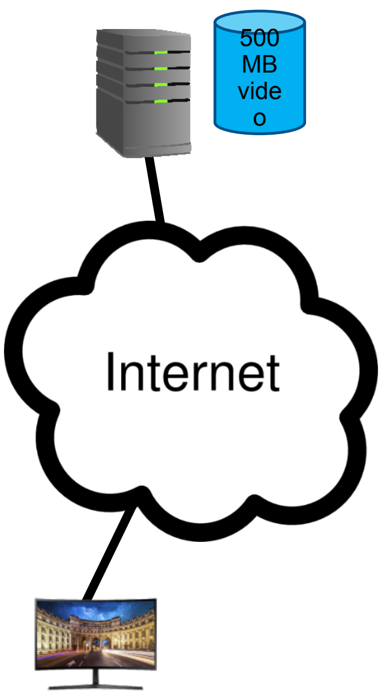

Distribution Reduces Latency and Increases Bandwidth

For example, Google has **datacenters** spread **around the world**.

Each datacenter **has access to all** of the company's videos (or other **content**).

Anything you receive **from Google comes from** a **nearby** datacenter.

Content Replicated for Availability and Locality

Actually, for any given video,
- **copies** are kept **at multiple datacenters**
- **in case of catastrophe**

Any **datacenter**
- serving a population **that**
- **frequently requests** a particular video
- **keeps a copy cached**.

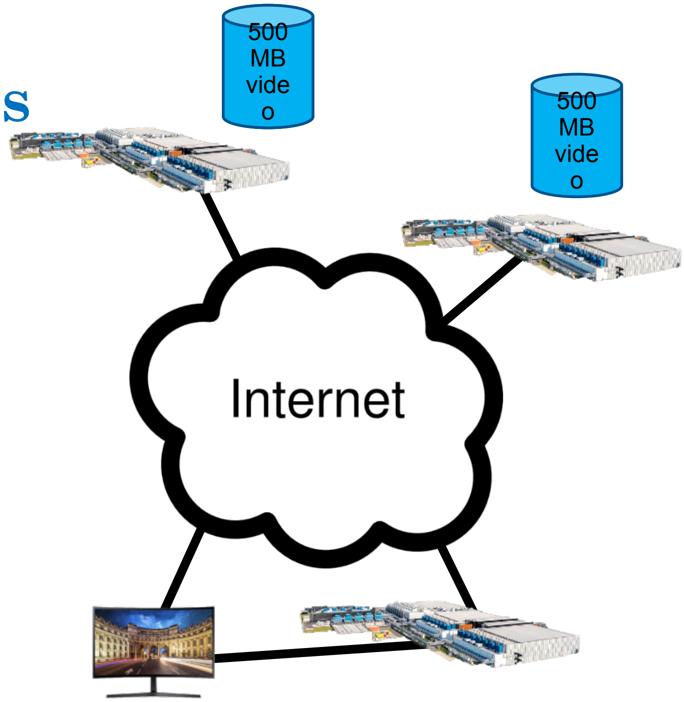

Datacenters are Connected with High-Bandwidth

How are these datacenters connected?

The answer varies:

- Google **built its own** optical network;
- other companies **rent fibers** (multiple Tbps) **or wavelengths** (10-100 Gbps), or
- P**ay another company** (such as Akamai) to handle content distribution.

Large Formats Take Longer to Send

Imagine that a company keeps all videos in 4K Ultra HD format (20 Mbps bandwidth).

What if a customer has only 10 Mbps?

Does the company transmit the 4K format,

- forcing the user to wait for buffering, or
- to stall every so often to catch up?

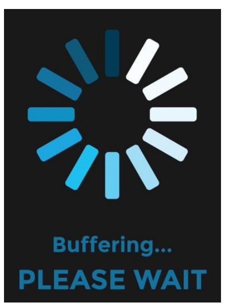

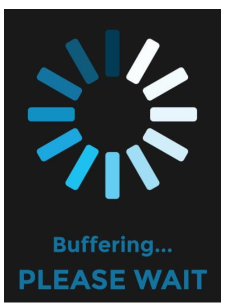

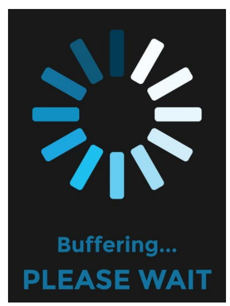

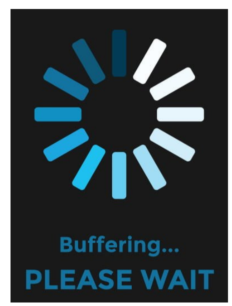

Changing Image Formats Takes Time and Energy

What if a customer is using an old phone with a small display?

Does the company transmit the 4K format,
- forcing the phone to use battery power
- to reduce the images to fit to the screen?

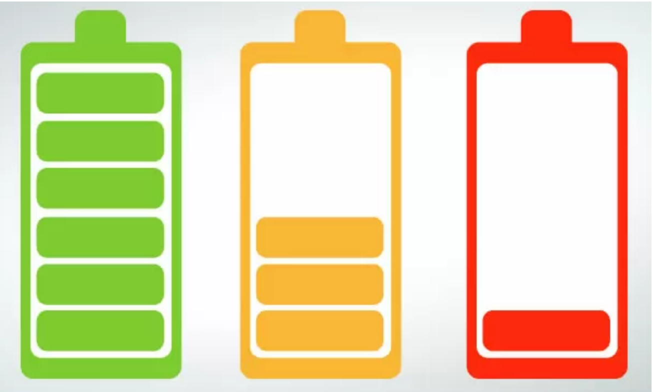

Companies Keep Multiple Formats to Serve Customers

Instead, many **companies**

- **keep video in multiple formats**
- and **send the** format that **best fit**s a customer's network bandwidth and display capabilities
- (some companies allow you to override the choice, if you enjoy pain).

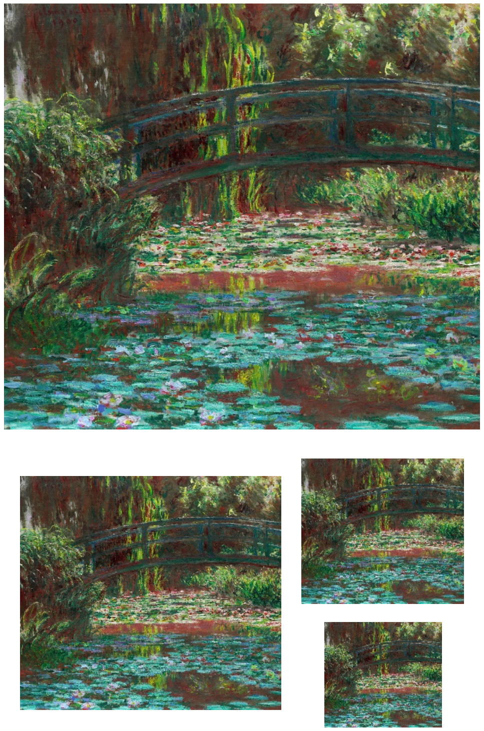

Transcoding Uses Significant Computational Resources

Converting video formats (**transcoding**)
- **requires a lot of computation**, thus
- costs a lot of money.

But **happy customers are worth money**!
- Unhappy customers have a bad habit
- of finding another company and
- no longer being a company's customers.

Transcoding also Used to Reduce Costs

Similar techniques are used for photos.

Compare, for example,
- a photo downloaded from Facebook
- with a photo from your camera's phone.

Facebook's app
- reduces the image size (on your phone/computer)
- before transmitting to the Facebook server,
- **reducing bandwidth/time as well as space**.

Real-World Examples

Adaptive Video Streaming

Netflix utilizes CDNs and adaptive streaming to deliver high-quality video efficiently to users worldwide.

Real-Time Video Conferencing

Zoom compresses video and audio in real-time to enable smooth and stable video conferencing experiences.

Dynamic Audio Streaming

Spotify streams audio in small chunks, adjusting quality dynamically based on user connection speed.

Low-Latency Live Streaming

Twitch employs real-time streaming protocols to minimize delay and support interactive live broadcasts.

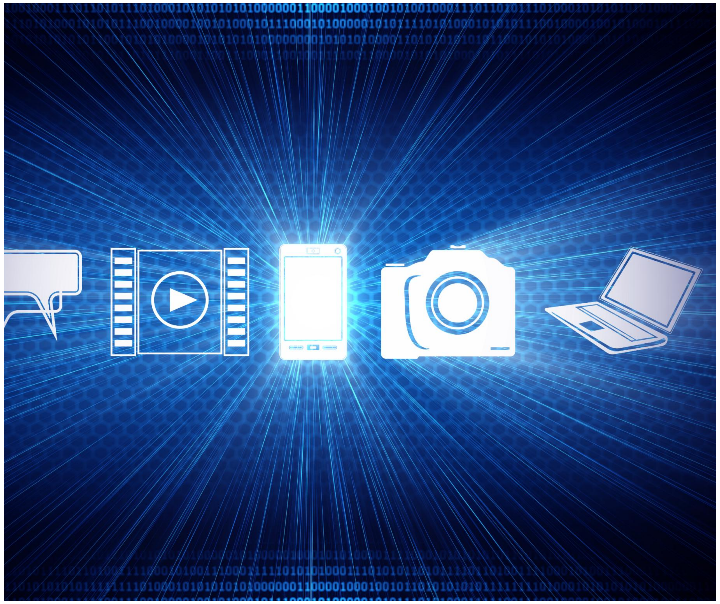

Terminology You Should Know from These Slides

- Bandwidth
- Streaming
- Downloading
- Compression
- Buffering
- CDNs
- datacenter
- locality
- transcoding

Concepts You Should Know from These Slides

- Routers can drop packets
- Connection bandwidth is the minimum over all links in a route
- Bandwidth is shared between connections
- More delay means less achieved bandwidth
- Transfer time = (# bytes) / (achieved bandwidth)
- Different technologies used for better streaming performance
- Buffering can help with variability in video data rate and low connection bandwidth
- Moving content closer to user is useful

Information about Exam

- Exam will be **50 mins long**. Booklet provided. Bring pen/ pencils.
- You are allowed 1 letter-size (8.5" x 11") **handwritten cheatsheet** (you may use both sides). NO PRINTOUTS
- The exam is closed book/notes, and calculators are not allowed.
- Fri, Oct 3. During your lab section but in ECEB 3081 (lecture classroom)
- This exam will test concepts covered in the Past & Present module.
- Students with accommodations please schedule at TAC