Internet

Kelly Rivers and Stephanie Rosenthal

15-110 Fall 2019

Announcements

- Homework 5 full due Monday

Big Picture

Unit 3 – scaling up computing for larger tasks

- Last week - how CPUs and computers can work together to run multiple programs
- This week – the largest distributed system invented – the Internet
  - How does it work?
  - Why/how does it enable good and bad actors?
  - How do you protect yourself online?

The Internet

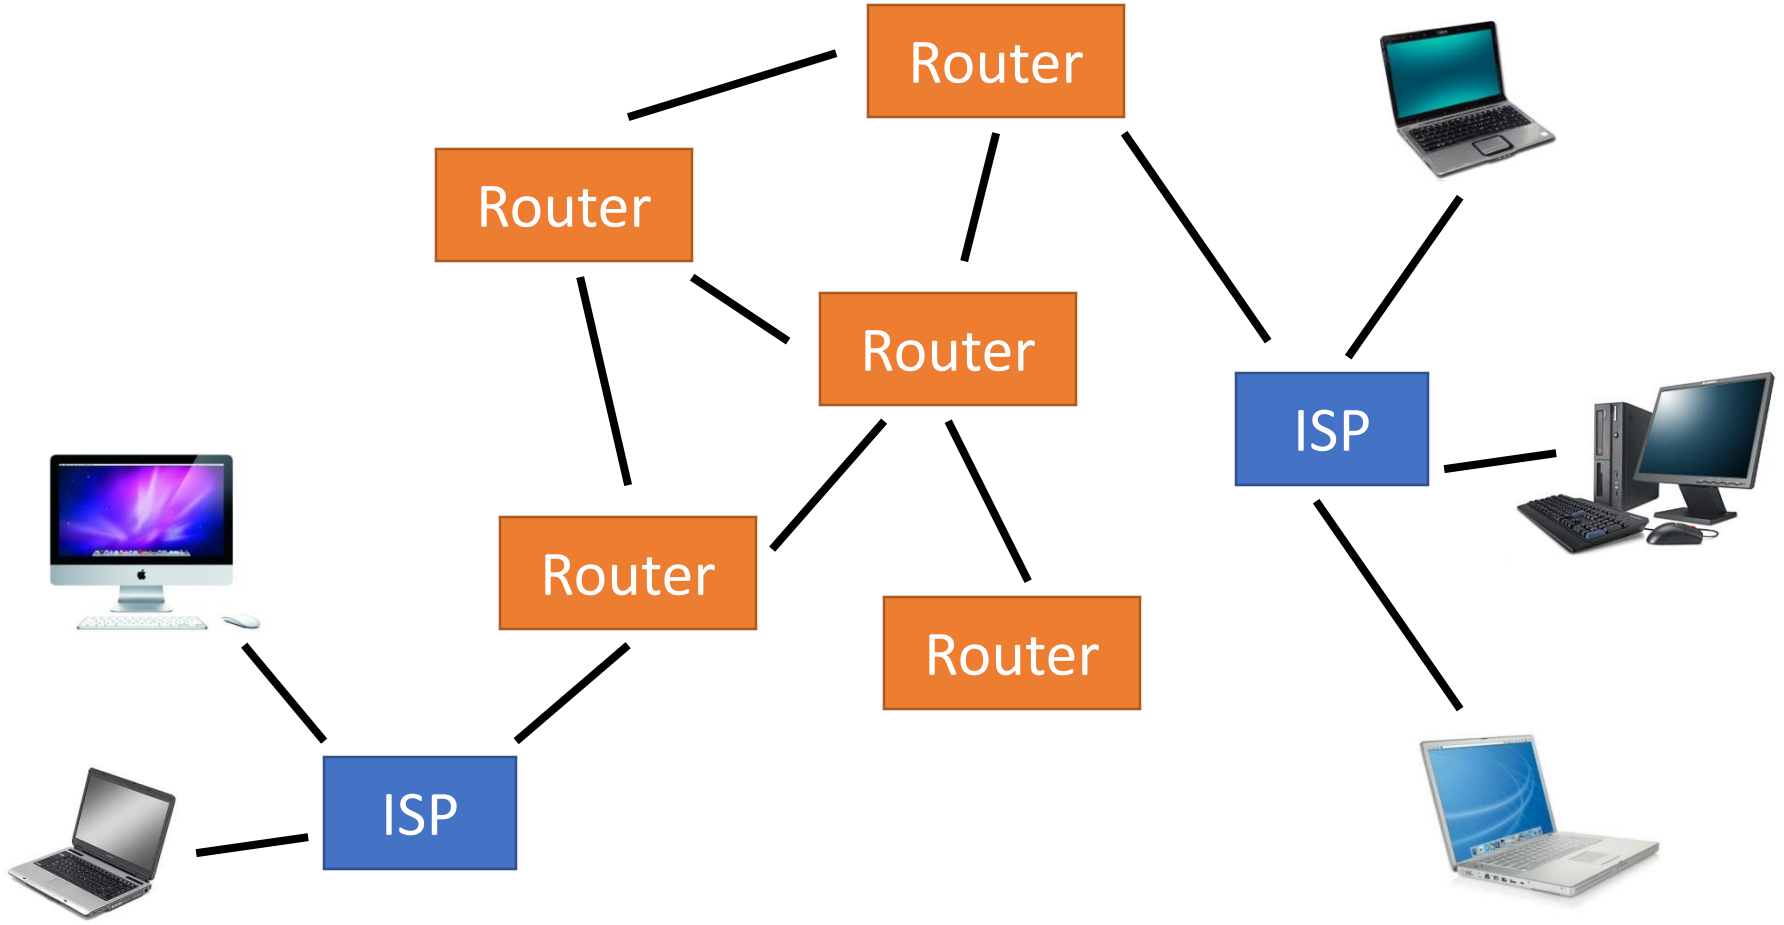

The Internet

Internet Service Providers (ISPs) provide the wired or wireless (WiFi) connections that allow computers to get onto the internet.

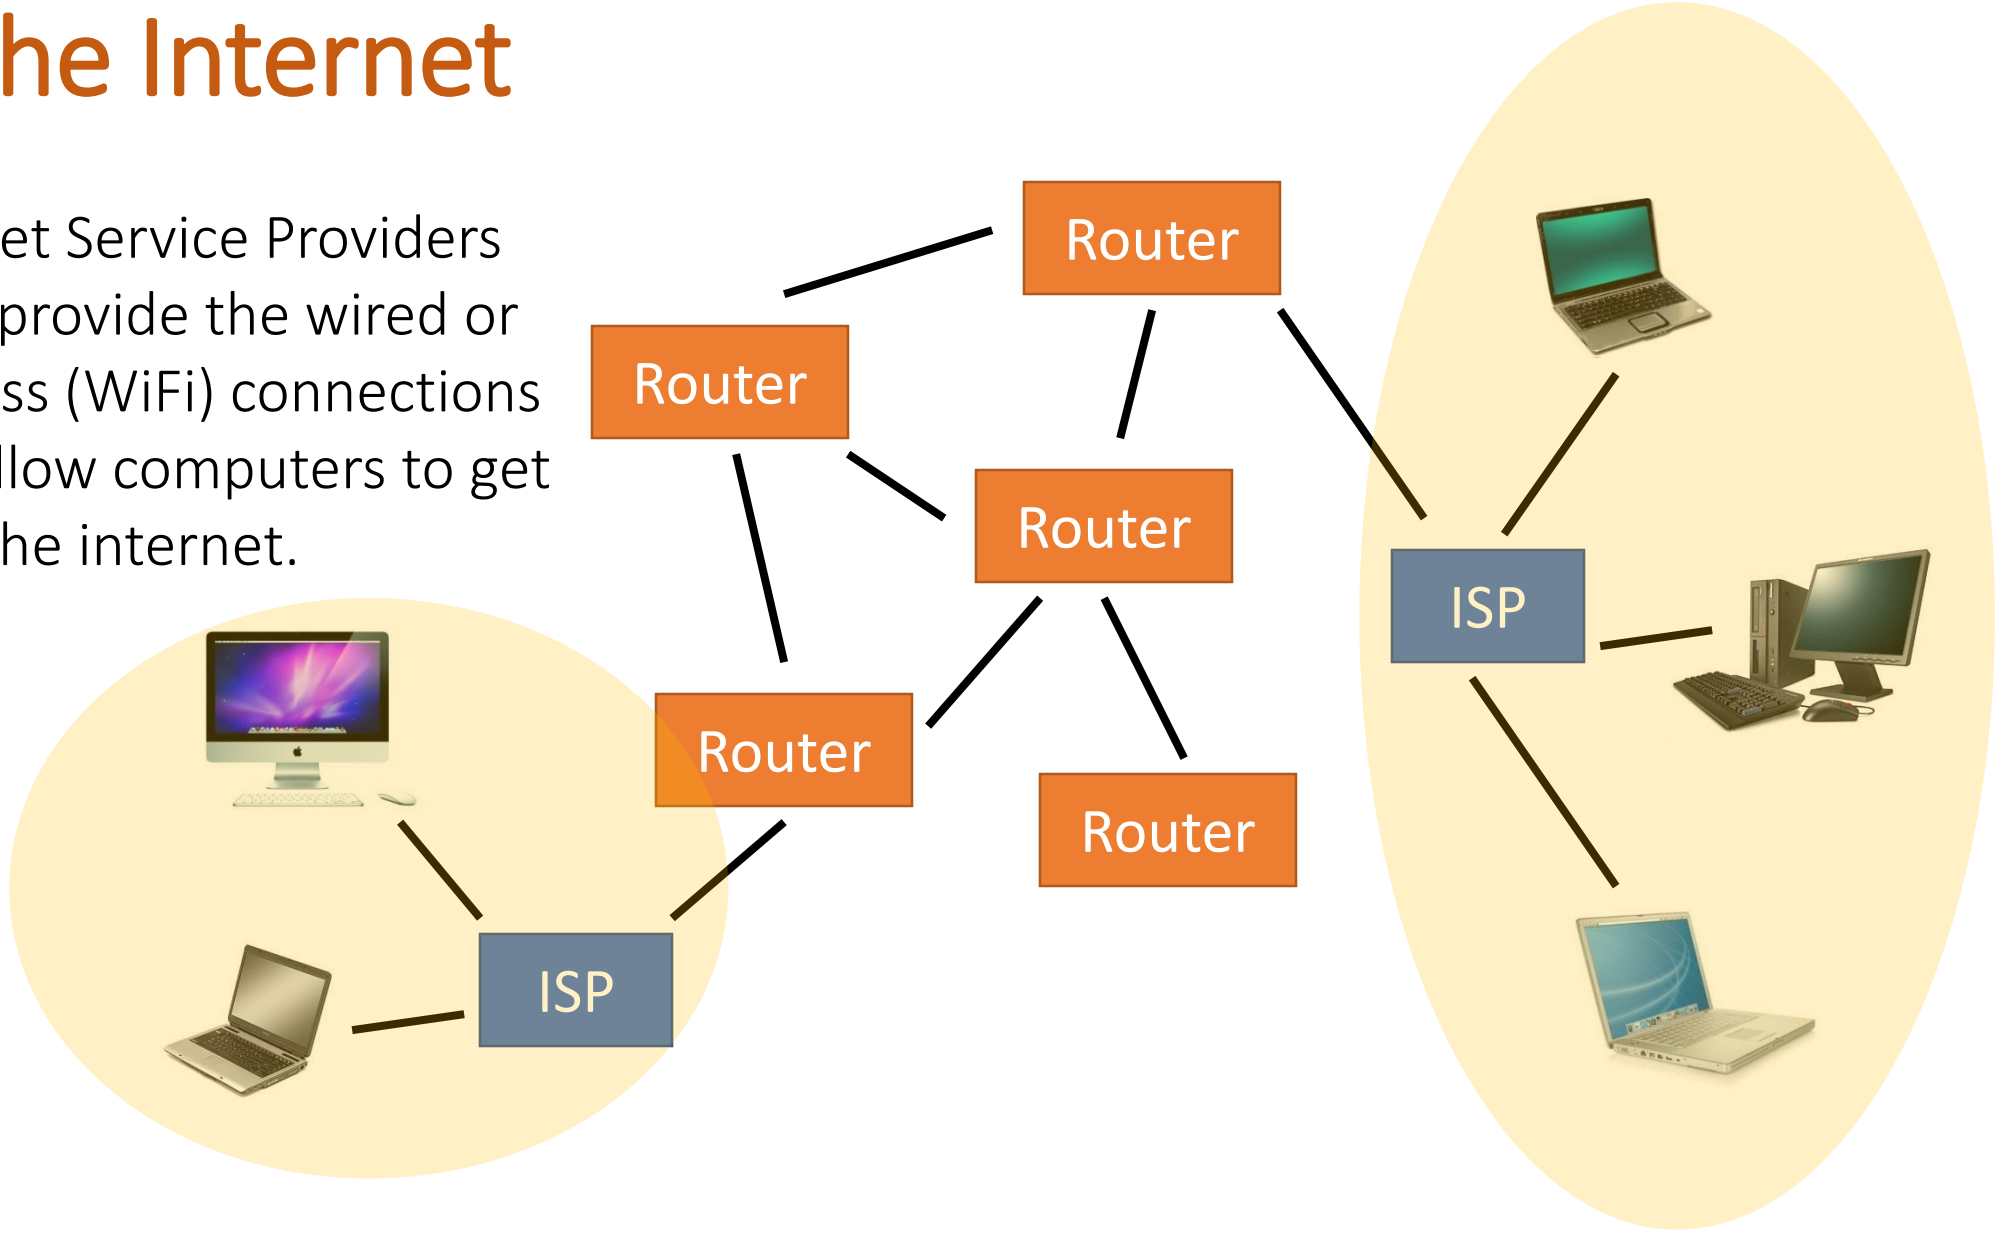

The Internet

Routers (computerized switches) connect the ISPs together. Lots of routers are available to help move data around from one place to another

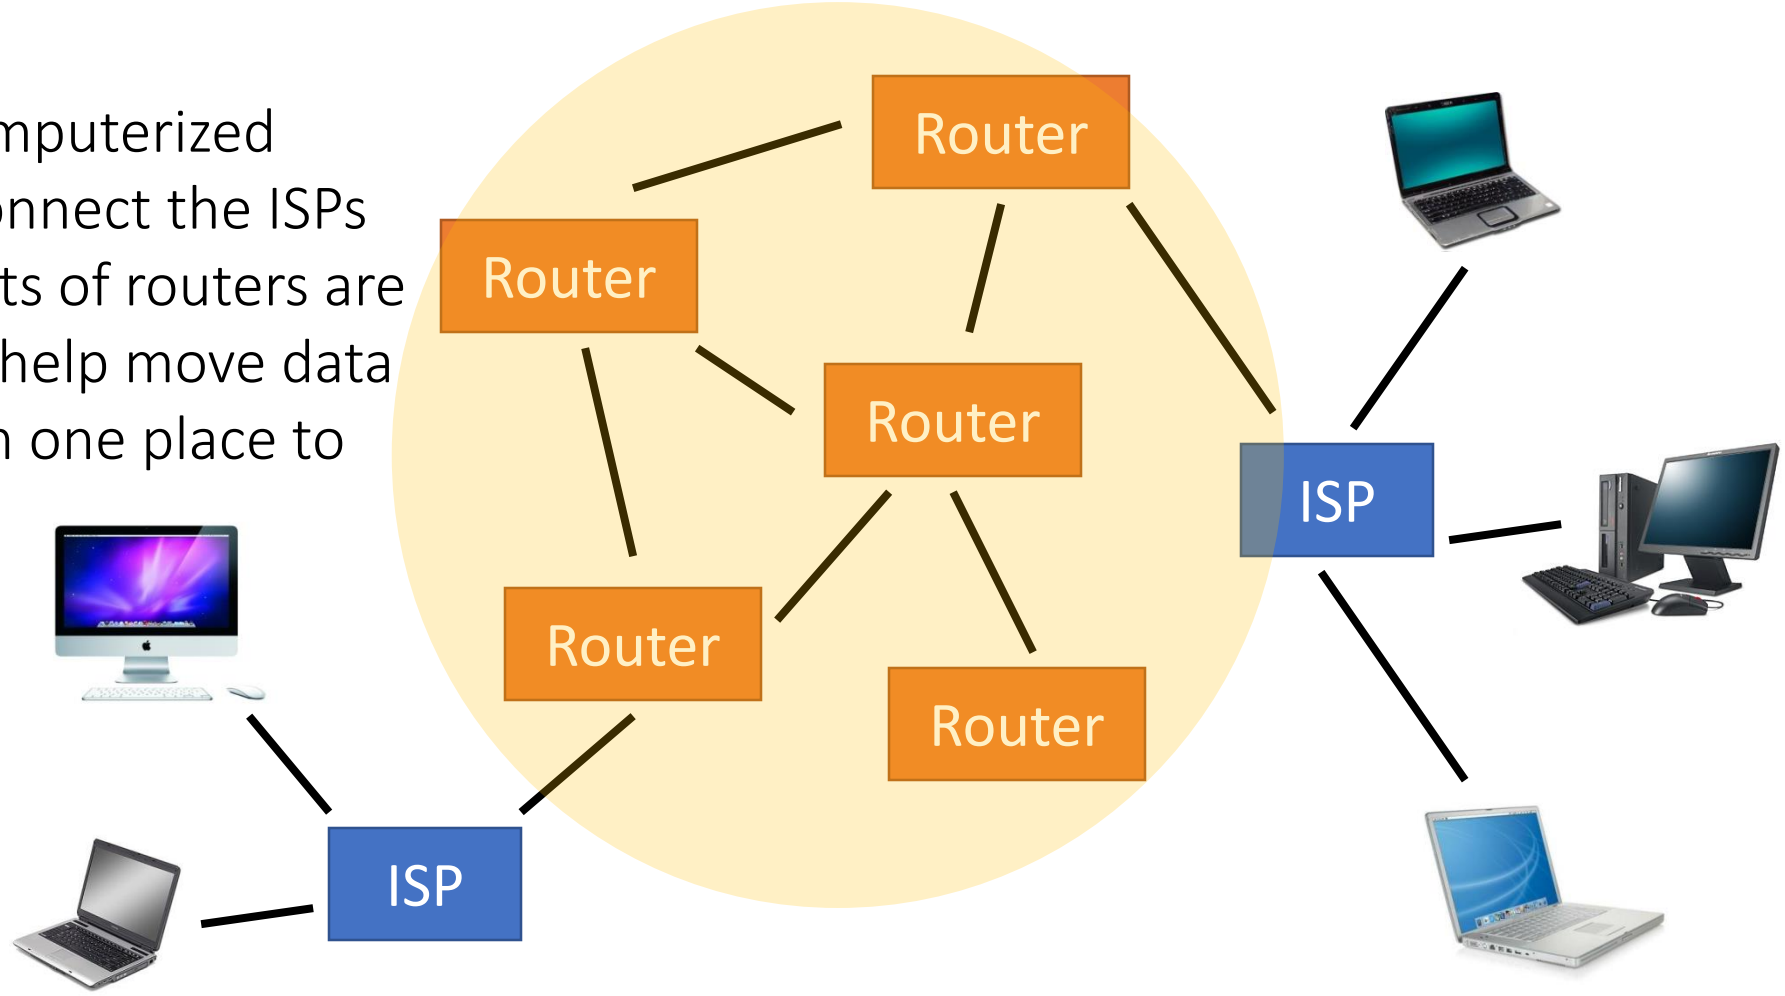

Historic Internet (ARPANET)

The original internet was called the ARPANET, funded by the US gov't. It included CMU (CARNEGIE) and several other universities around the country.

In the early 1980s, the gov't split it's secure network from the rest of the ARPANET, creating the start of the internet we have today

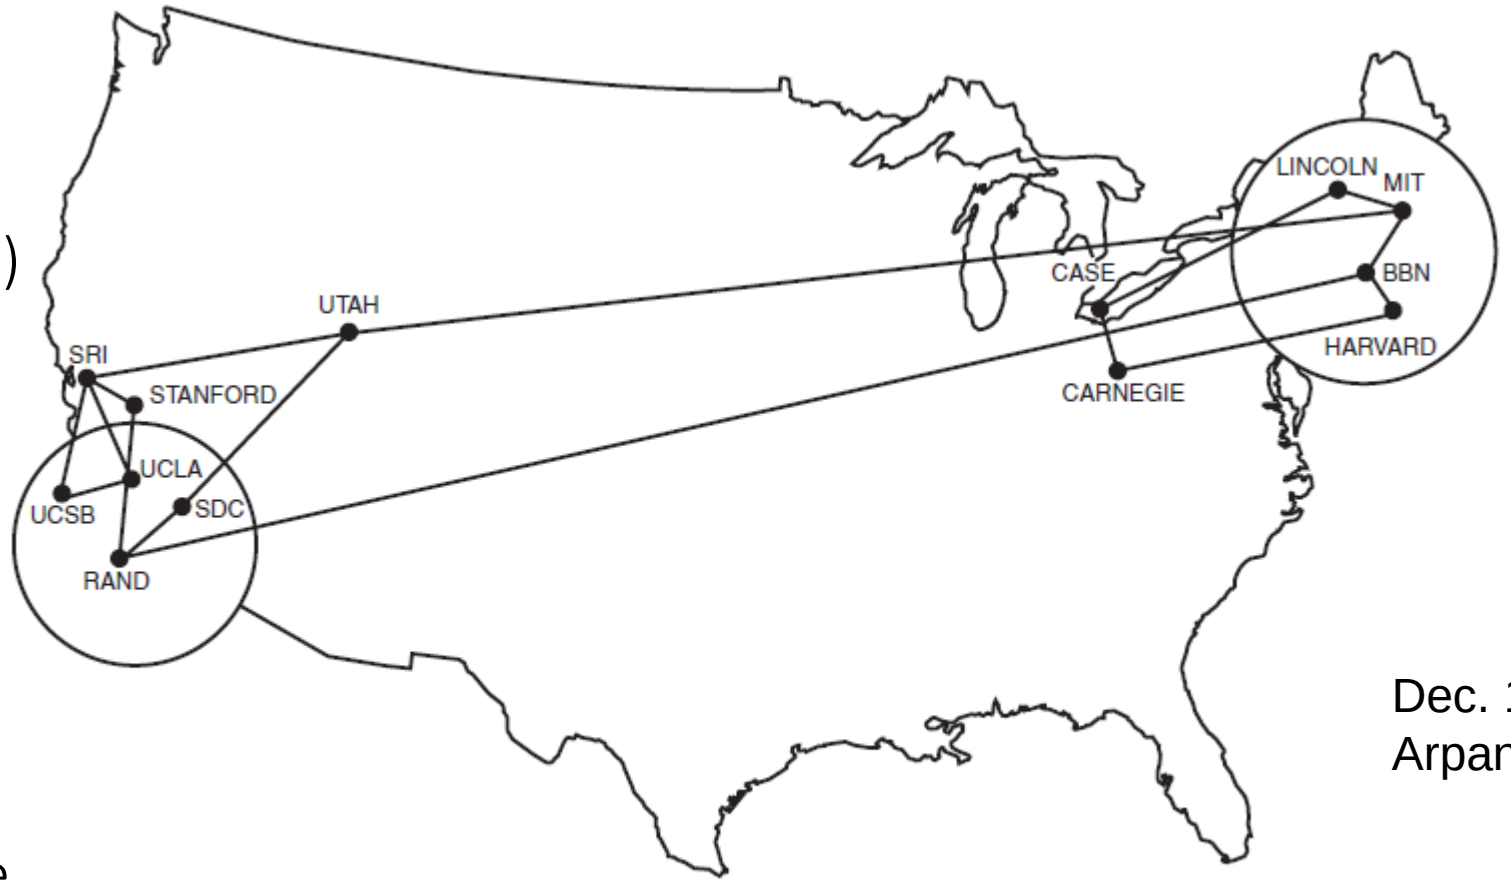

Dec. 1970
Arpanet

Source: Heart, F., McKenzie, A., McQuillian, J., and Walden, D., ARPANET Completion Report, Bolt, Beranek and Newman, Burlington, MA, January 4, 1978.

Today's Internet

This is a tree representing all of the computers on the Internet in the 2000's. Each computer is connected to an ISP. The brightest points represent nodes in the tree with many connections (the ISPs). Then ISPs are connected to each other through routers.

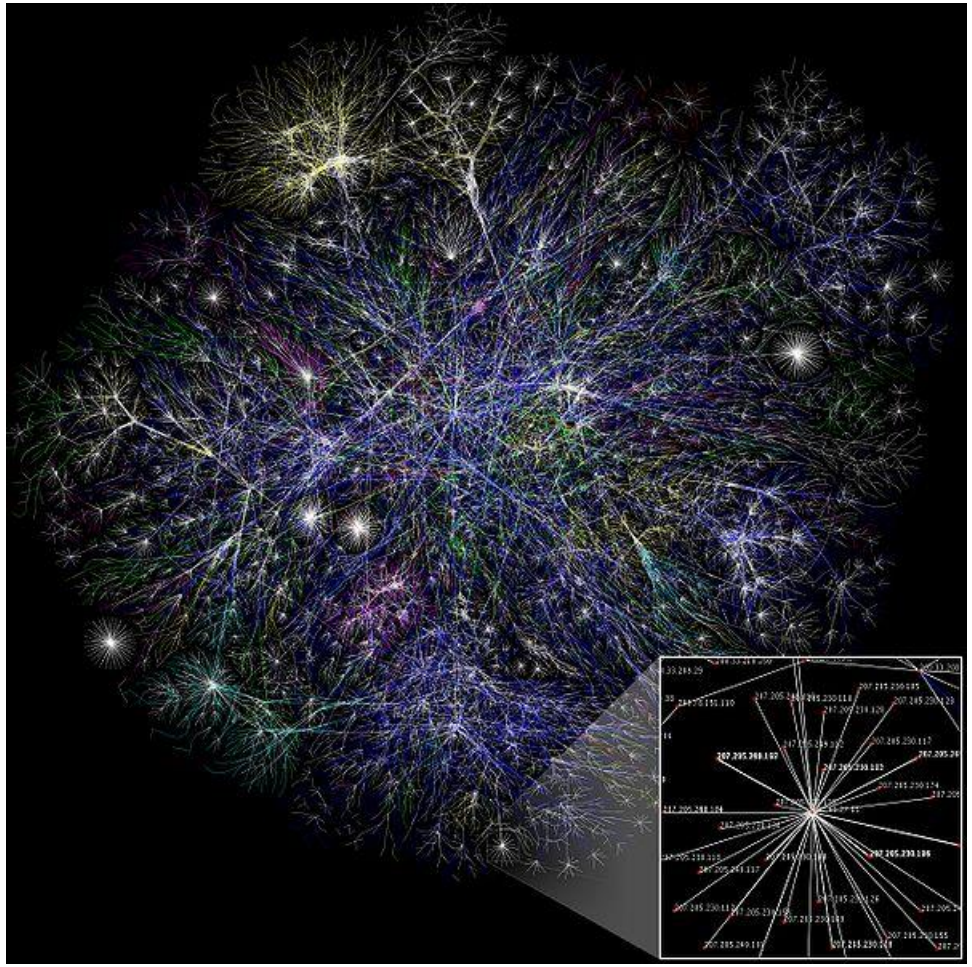

2000' s Internet Map (small section)

The Internet

The Internet has many standard ways of communicating across the network so that computers that have never talked before can do so.

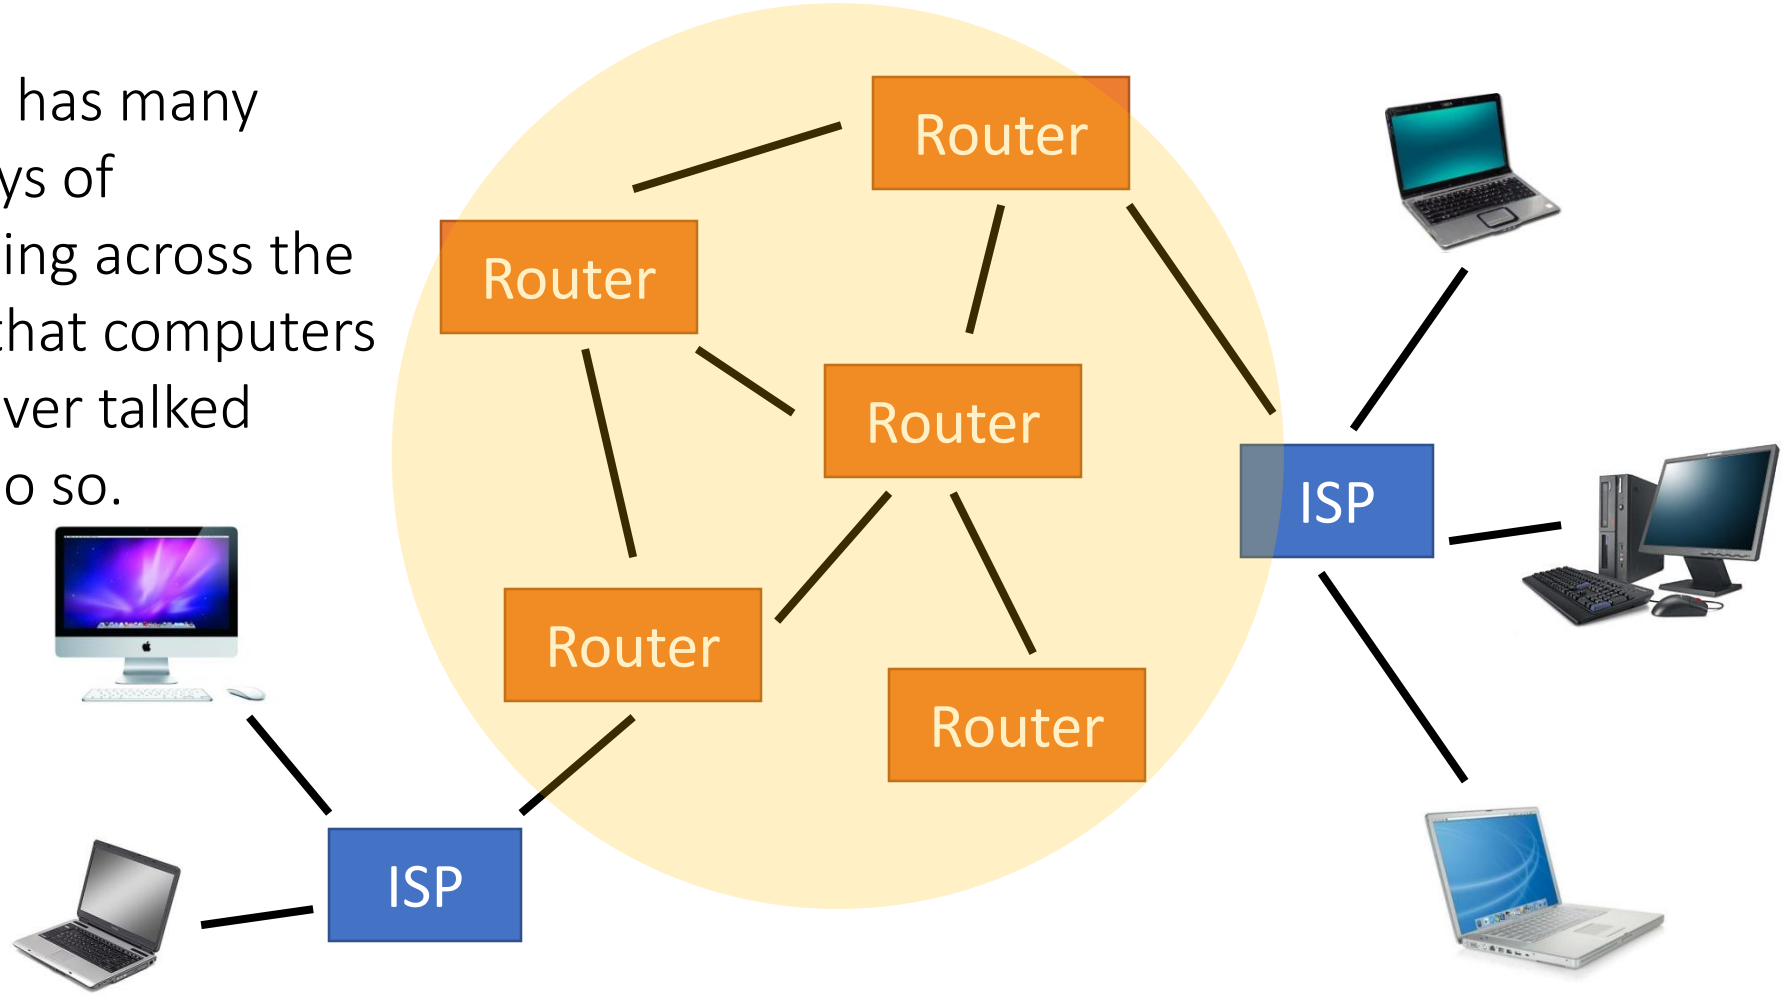

HTML is a language for displaying information visually on the screen, for example.

The Internet

It is hard for any person or organization control the data on the internet, because there isn't one bottleneck point for a person or organization to control.*

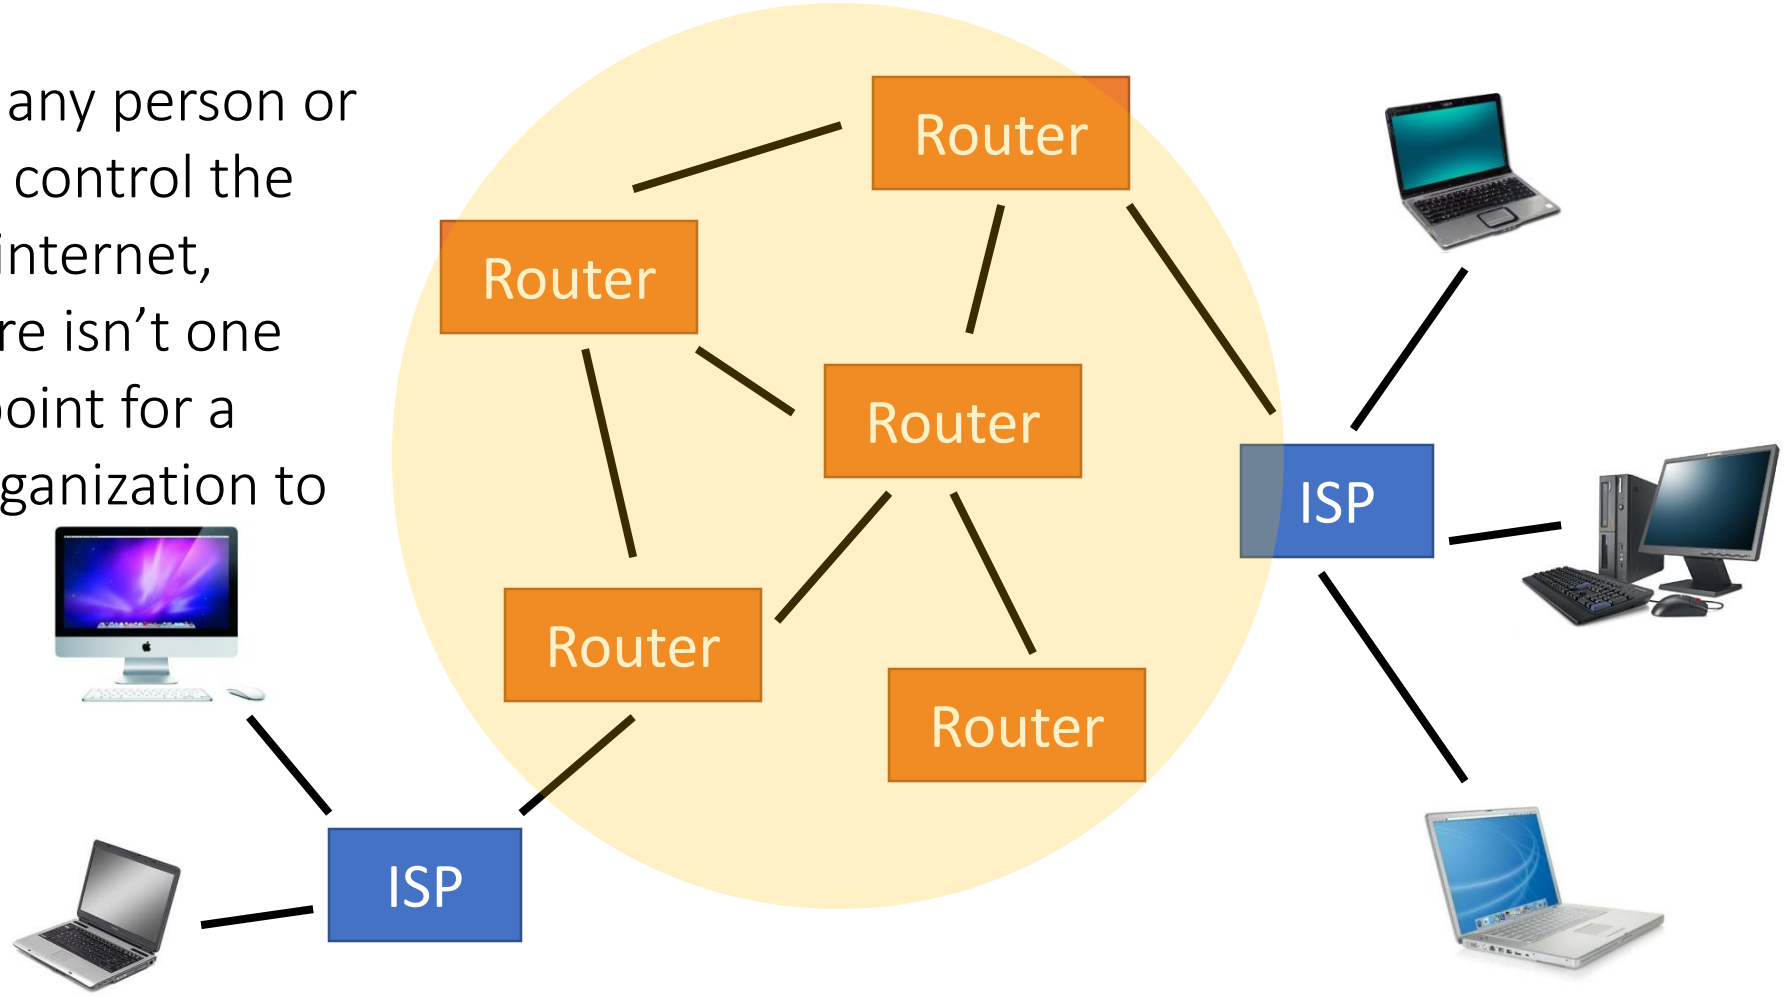

*This is not true for ISPs however. They can control the data coming and going to computers.

Net Neutrality

Because the ISPs can control what information you can get and how fast you get it, they have the power to charge money for the privilege of using specific websites.

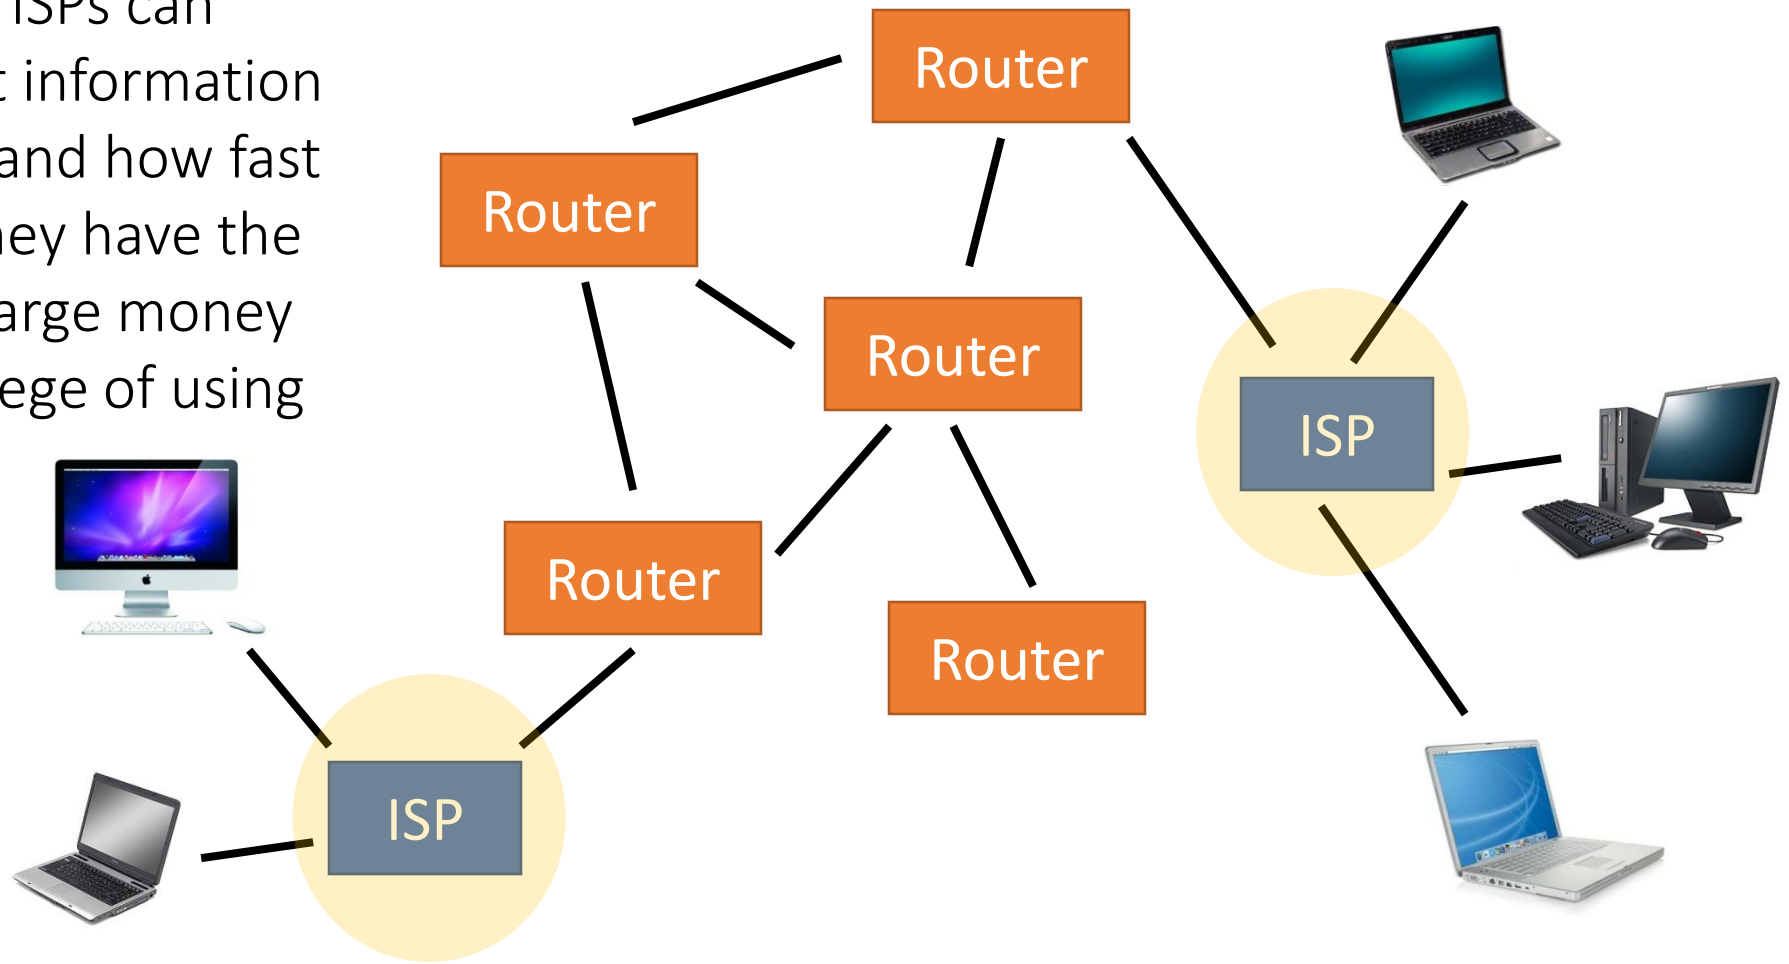

Net Neutrality

The central debate in net neutrality is whether or not the ISPs are allowed to do this or whether they must make every site available equally

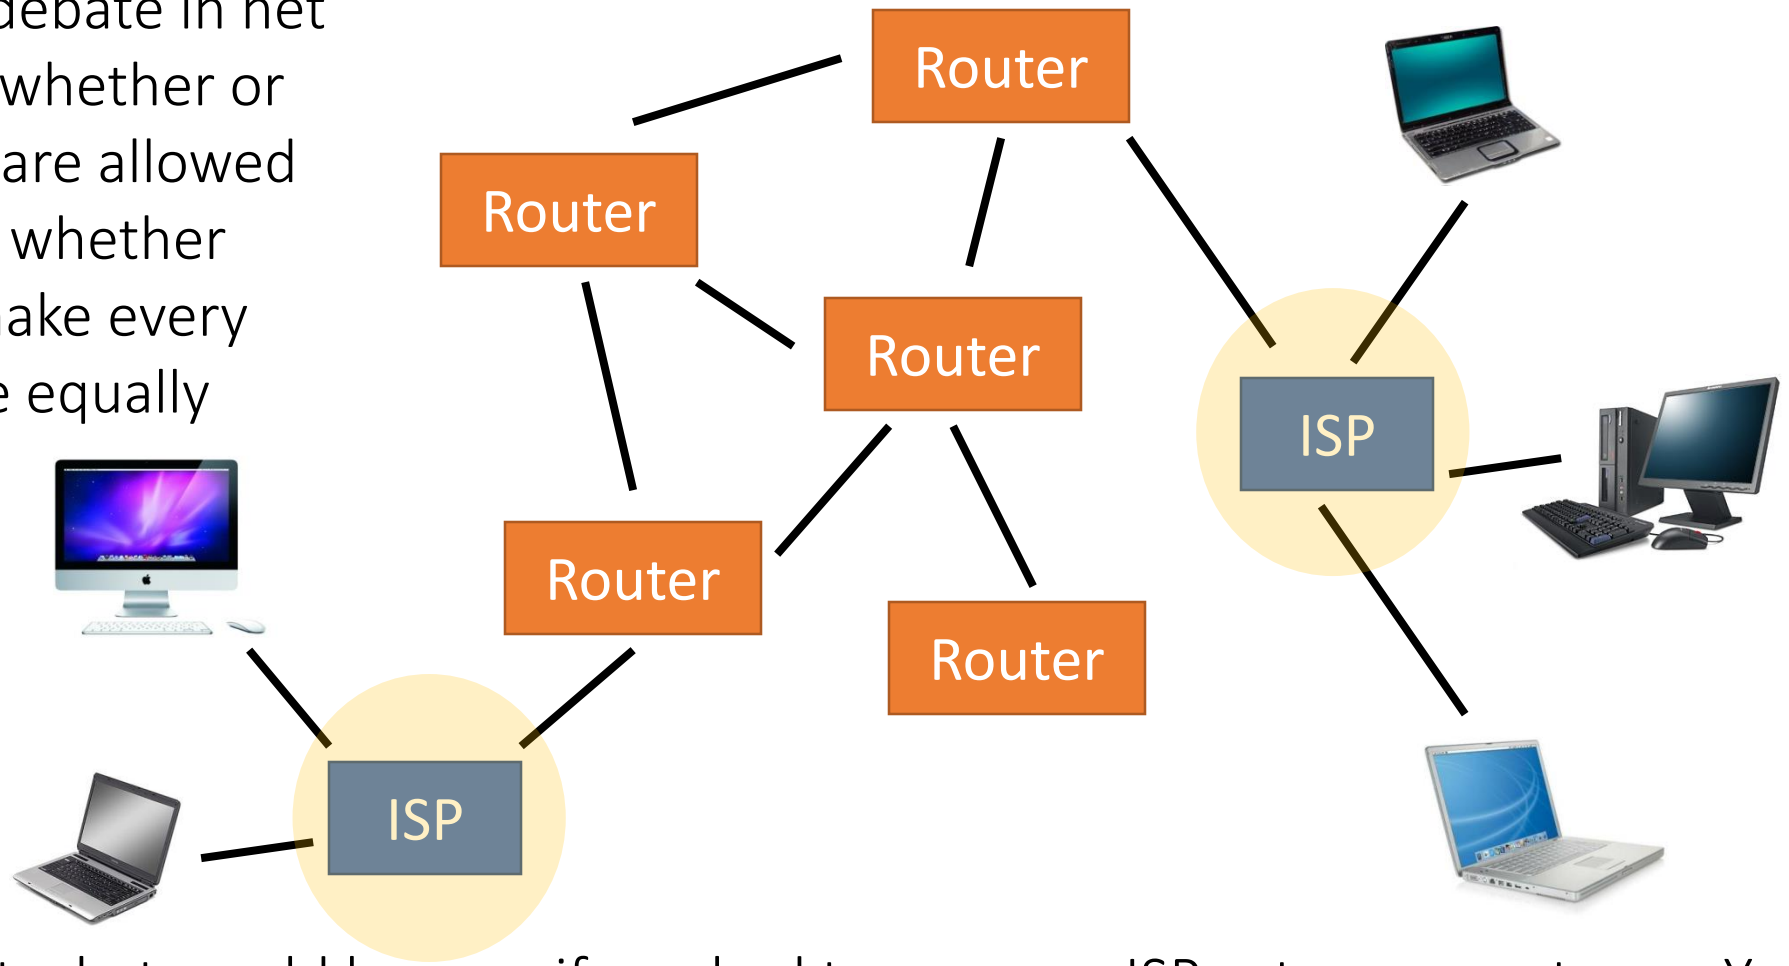

Think about what would happen if you had to pay your ISP extra money to use Youtube at all or if you had to pay so that it would run fast enough that you could watch a Netflix movie.

How do you download a website to view?

You type in a website address www.google.com on your web browser. The web address is just a nickname for the IP address, so your computer needs to find the IP address

IP Addresses

IP stands for "Internet Protocol"

Each computer on the Internet is assigned an IP Address consisting of four numbers between 0 and 255, inclusive.

Example: ____ . ____ . ____ . ____

128. 2. 13. 163

How many bits per address?
How many computers can be on the Internet at the same time?

Data sent on the Internet must always be sent *to some IP address*

Instead of making you remember IP addresses, we assign web addresses and a way to map between web address and IP address (more on that next)

Where do IP addresses come from?

- An IP address **isn't part of** a computer!
- Groups of addresses are allotted to various organizations by IANA (Internet Assigned Numbers Authority)

  These organizations assign addresses to computers.
- Sometimes the IP address is static, sometimes dynamic
  - static for important machines that you always want to find at the same place
  - dynamic for others (your phone, computer, etc) is assigned a different IP address each time it comes online

What does an IP address “say”

- Identifies a particular machine *at a particular time*
- Identifies (somewhat vague) geographic location based on the organization (ISP) that “owns” it
- What it doesn’t say
  - who is using the machine to do what
  - what kind of machine it is

How do you download a website to view?

You type in a website address www.google.com on your web browser. The web address is just a nickname for the IP address, so your computer needs to find the IP address

How do you download a website to view?

Your computer talks to a router within your ISP asking if it knows the IP address of the website. The router remembers recent IP addresses to save time searching for it.

www.google.com
172.217.7.206

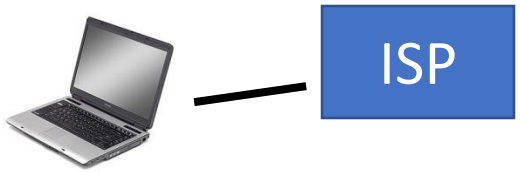

How do you download a website to view?

If it does, it will return it to the computer (slide 22).

If it does not, then the ISP will ask other routers (who would also send the response back to the computer).

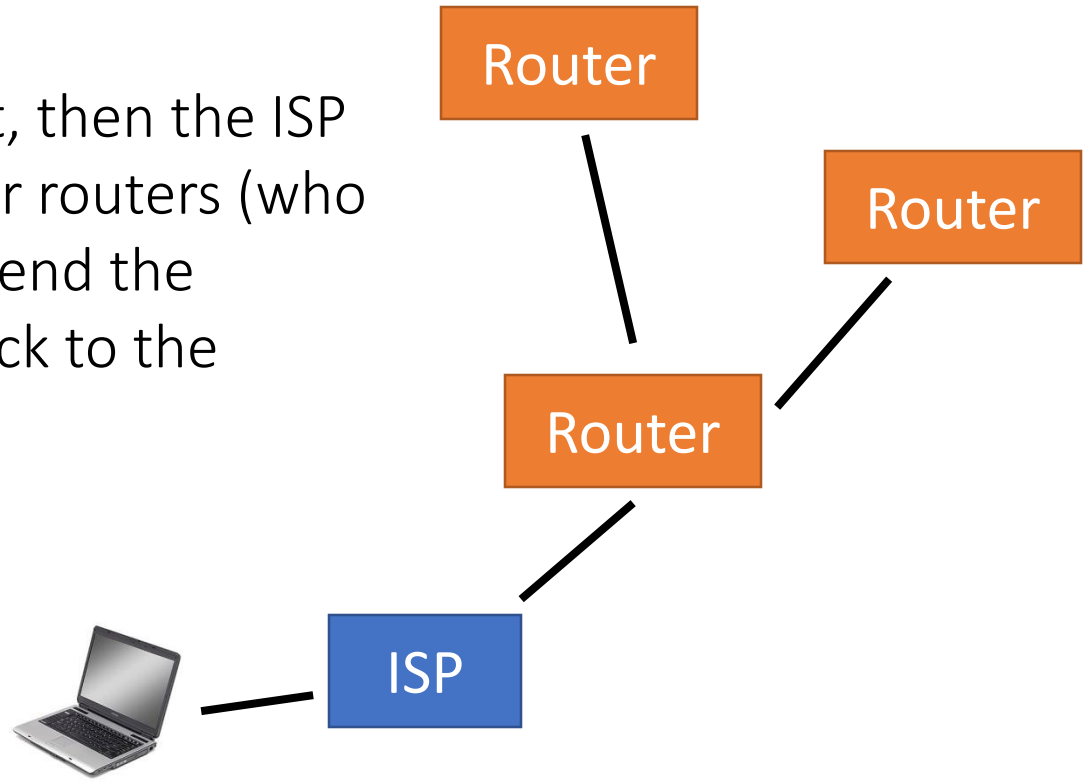

www.google.com
172.217.7.206

How do you download a website to view?

Eventually, it will ask a DNS server if no close router knows. A domain name server (DNS) is a separate computer that knows the addresses of all other computers. ISPs tell DNS servers what IP addresses map to what web addresses

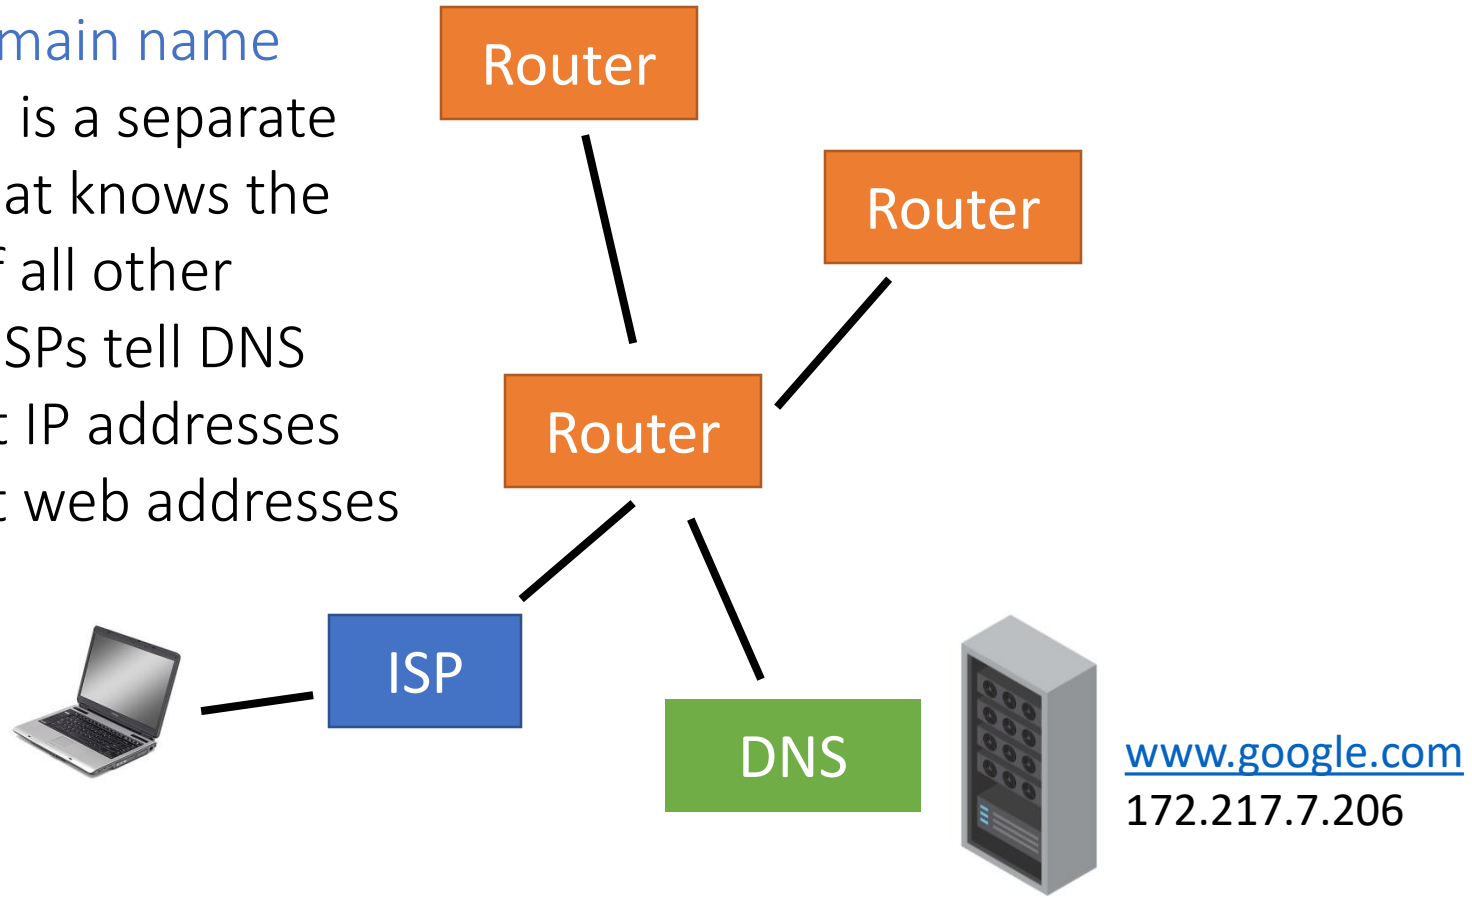

DNS Servers

A DNS server is like a phonebook. There are many DNS servers on the Internet so that if one breaks, people can still use the Internet.

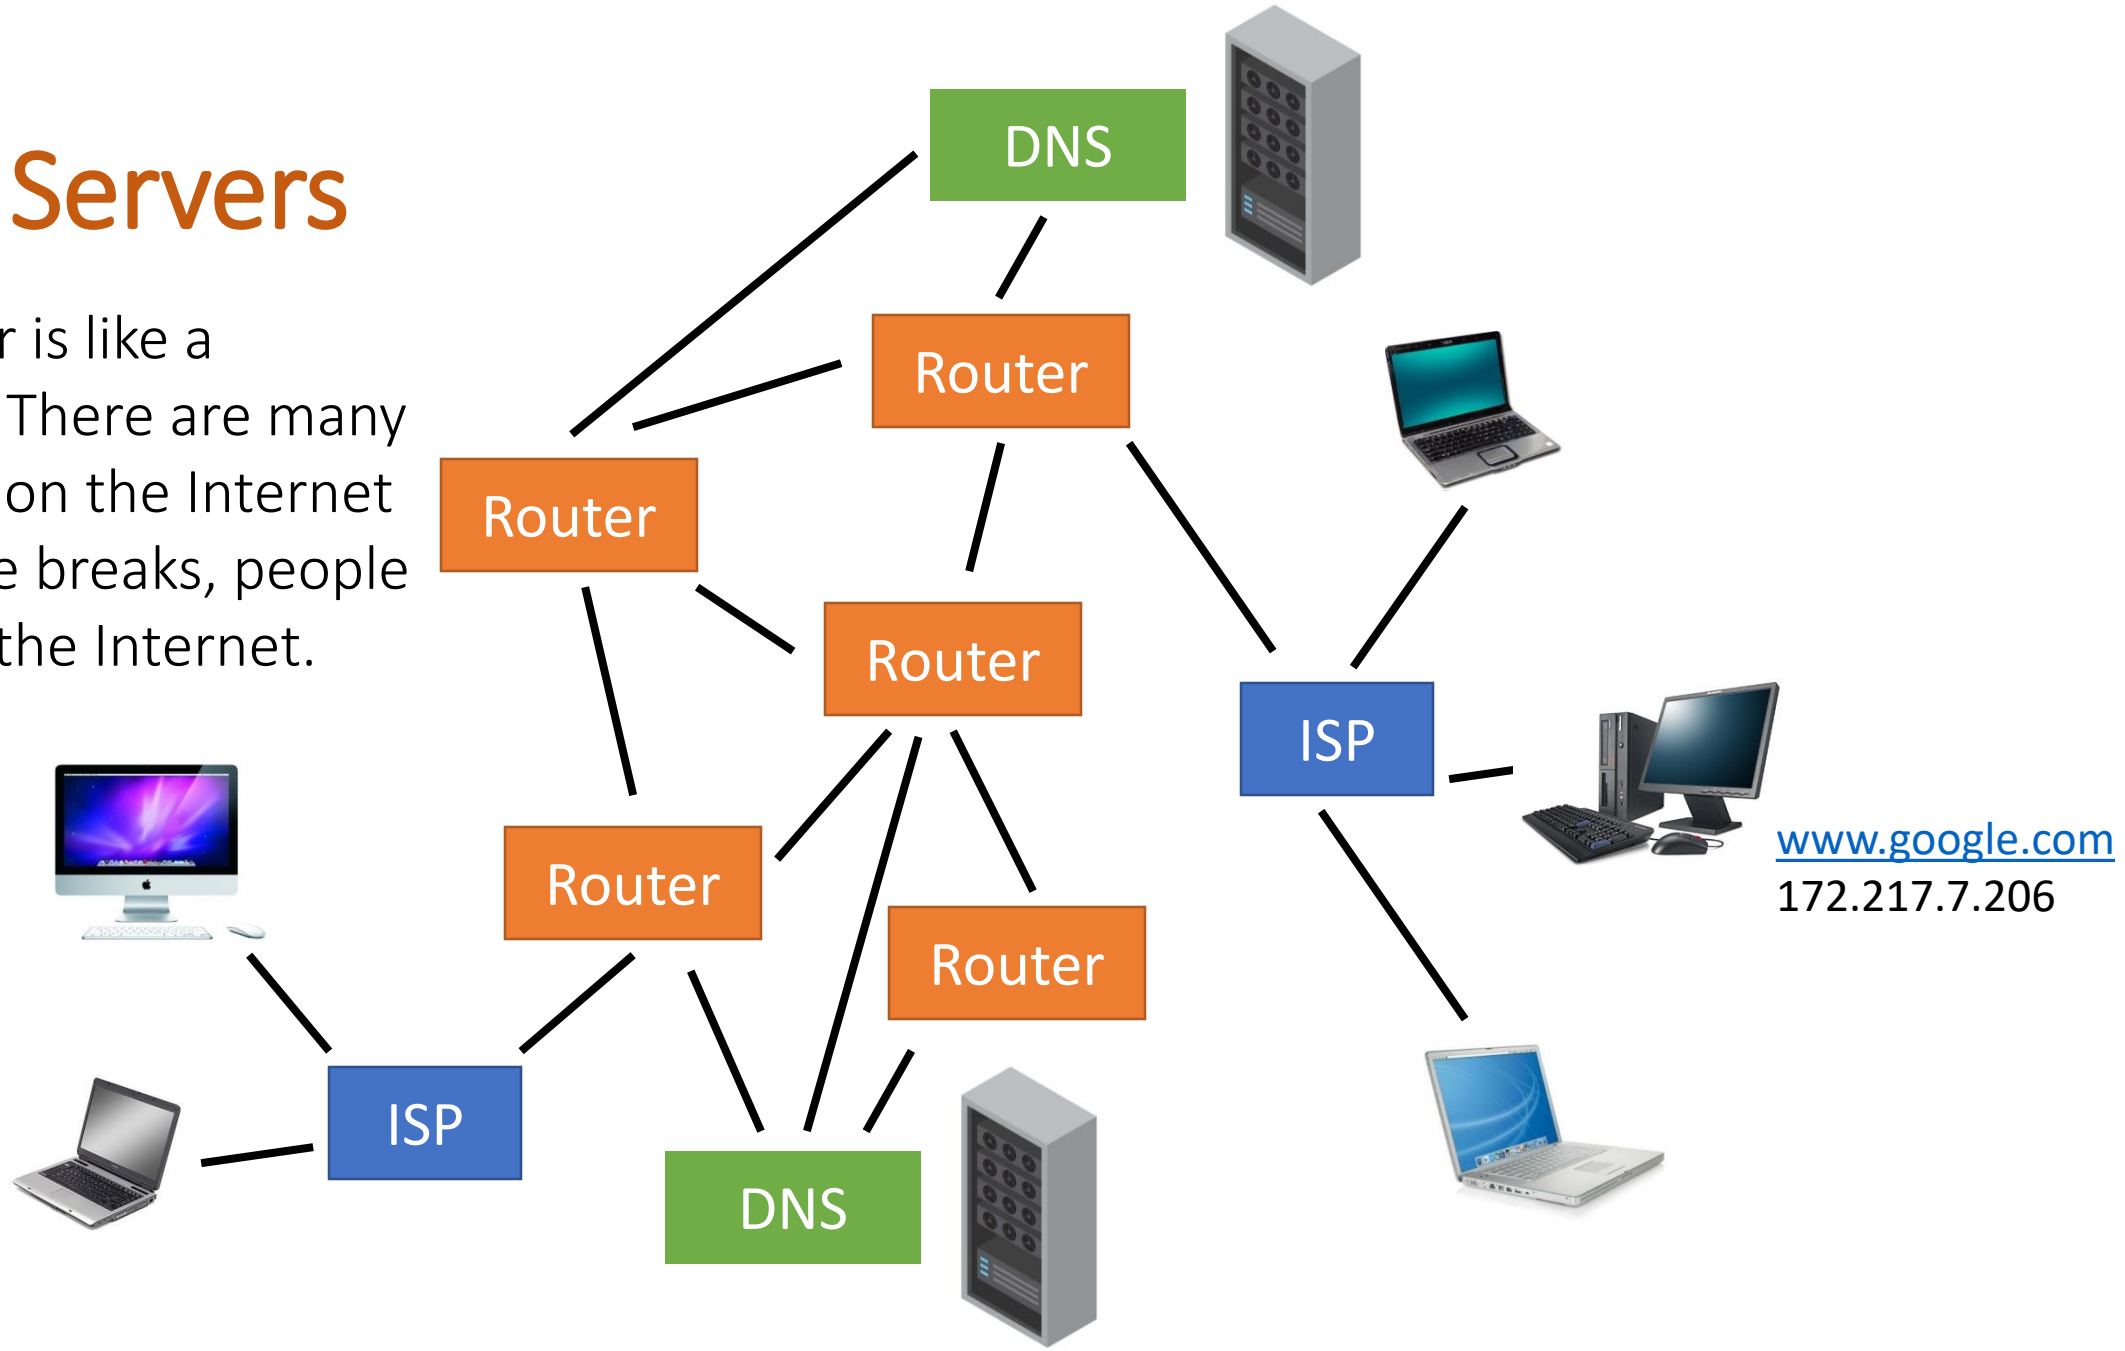

How do you download a website to view?

Once your computer knows the IP address, it can send a packet to that IP address with a request to send back a particular webpage.

www.google.com/index.html
OR
www.scholar.google.com

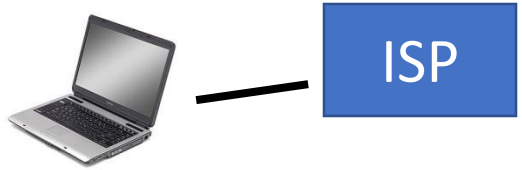

www.google.com
172.217.7.206

Packets

Packets are small messages that are addressed to particular IP addresses

Think of letters or postcards. They have a message and an address

> When you send a letter in the mail, you don't tell the post office what roads to take to deliver your message; you just tell the post office to get it there eventually.

Packets work the same way. The message tells the ISP to get your packets to the right place, and eventually they will.

Packets also have a parity bit which tells the receiver whether the message is corrupt or scrambled in some way

Many Paths to 172.217.7.206

Sending a packet to google.com

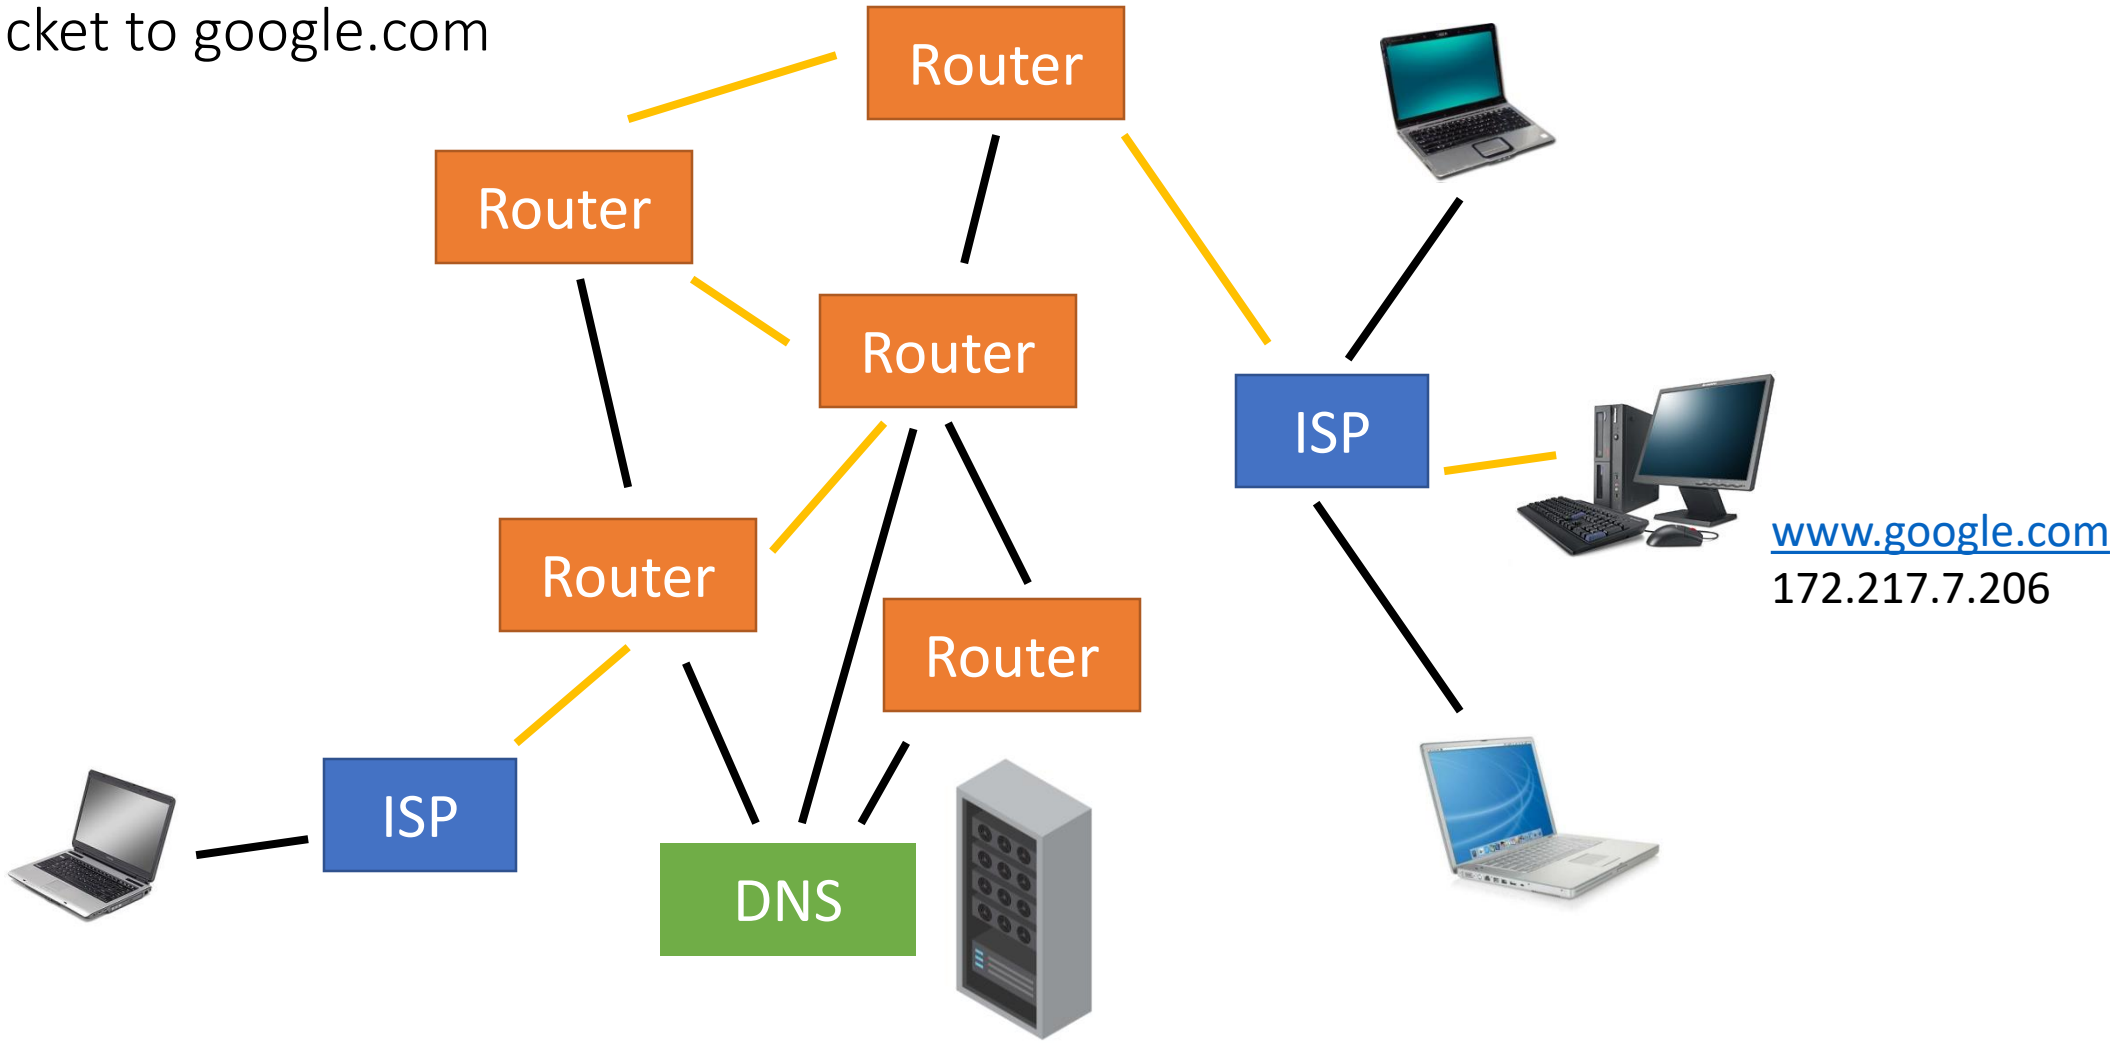

Many Paths to 172.217.7.206

Sending a packet to google.com

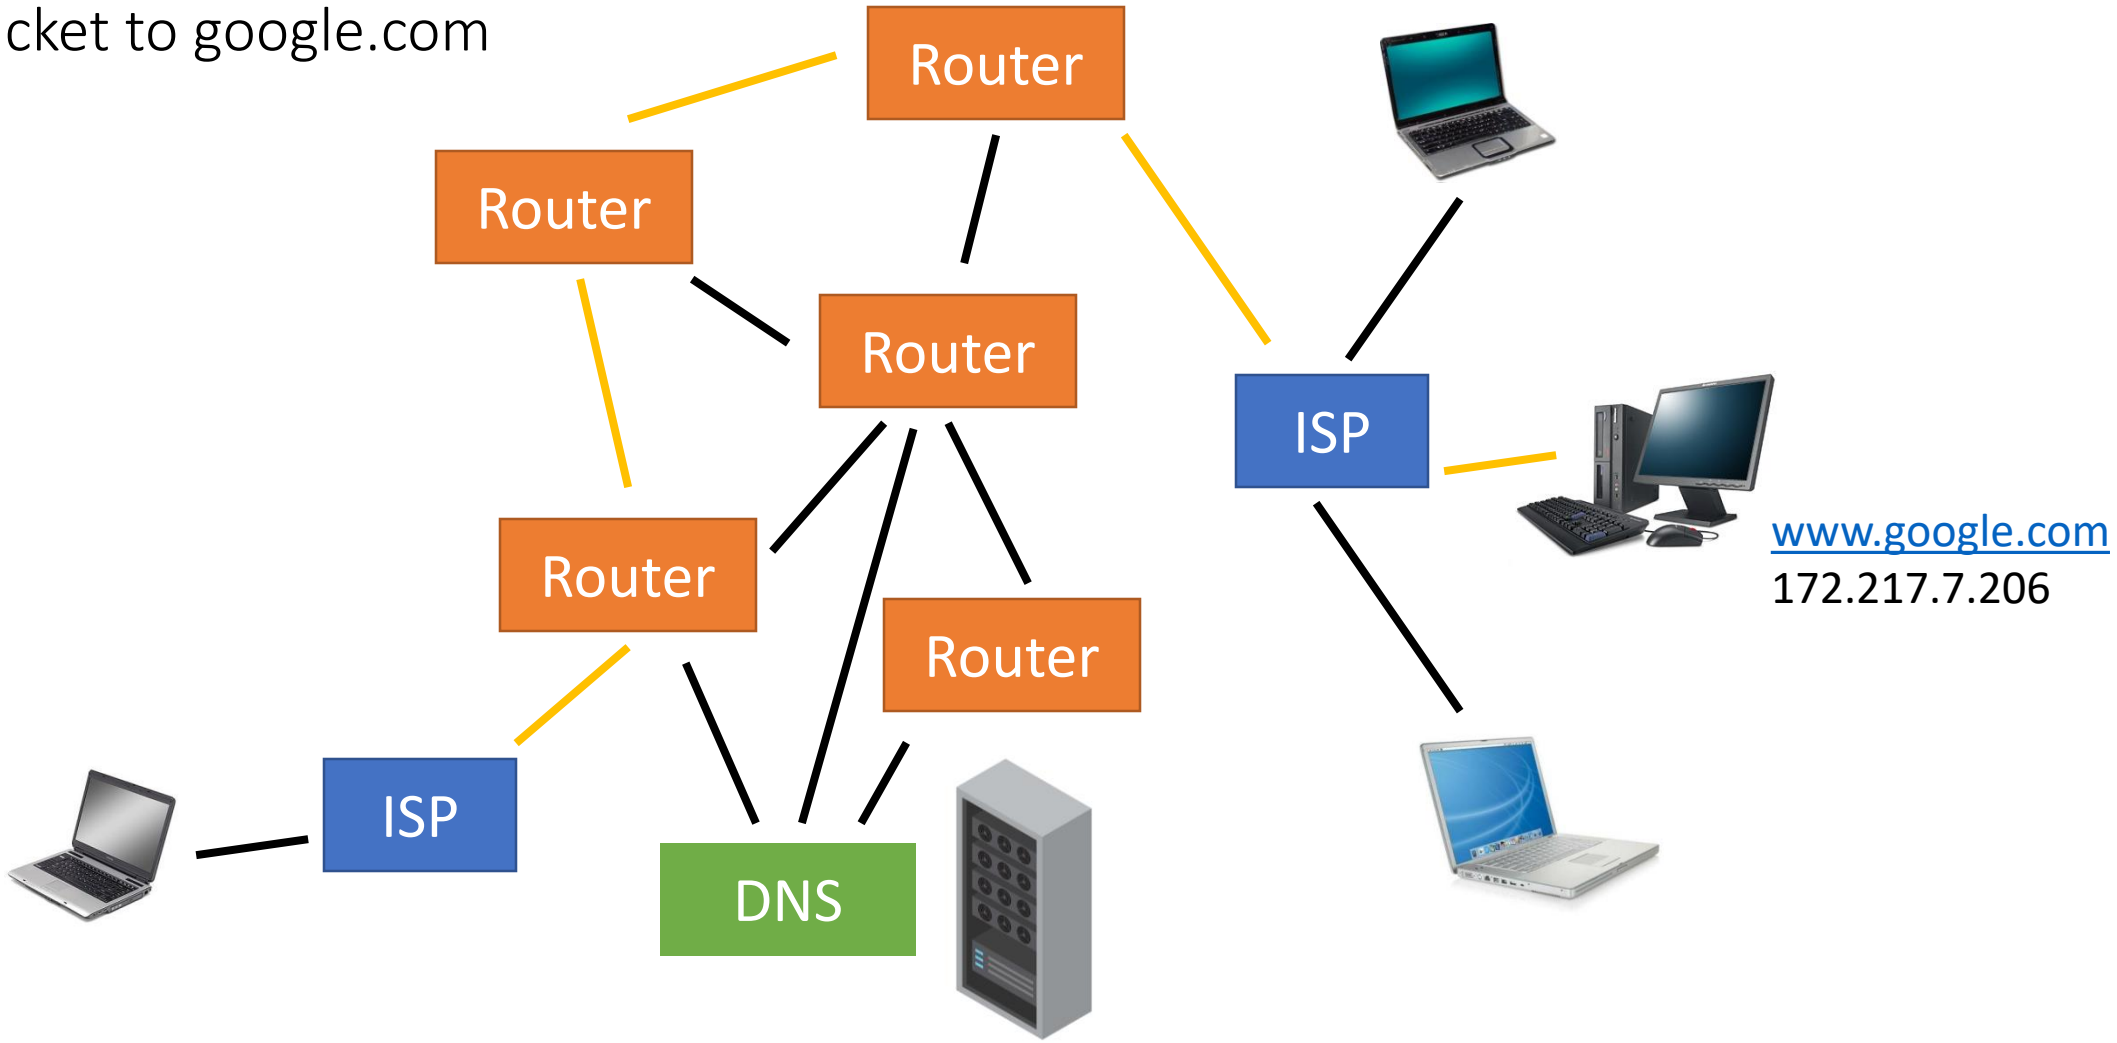

Many Paths to 172.217.7.206

Sending a packet to google.com

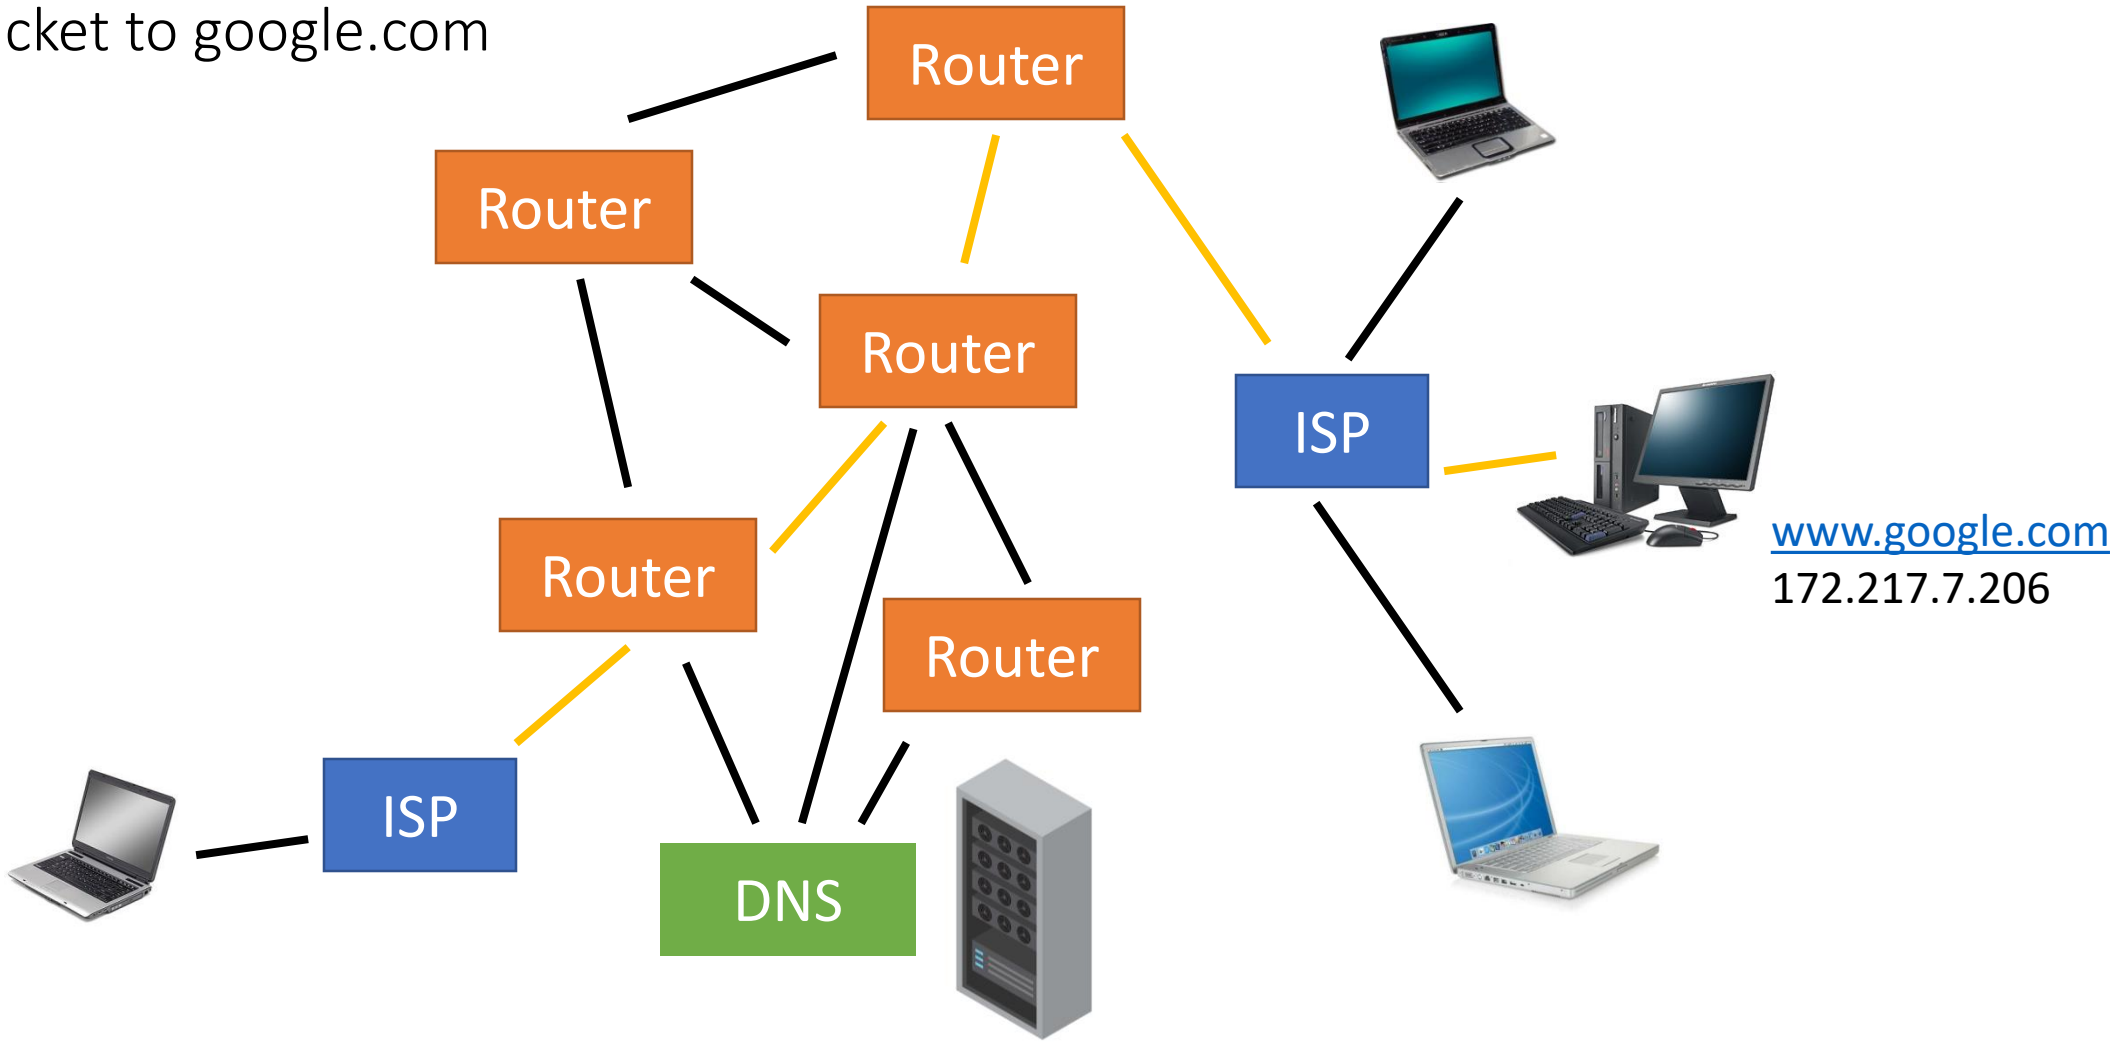

Packets

Packets are small messages that are addressed to particular IP addresses

Think of letters or postcards. They have a message and an address

> When you send a letter in the mail, you don't tell the post office what roads to take to deliver your message; you just tell the post office to get it there eventually.

Packets work the same way. The message tells the ISP to get your packets to the right place, and eventually they will.

If your message says to send data back (e.g., send a website), the other computer will send one or more packets back to you the same way.

Packet Switching

Each packet can take a different route from one location to another. This means they may arrive at different times or in different orders.

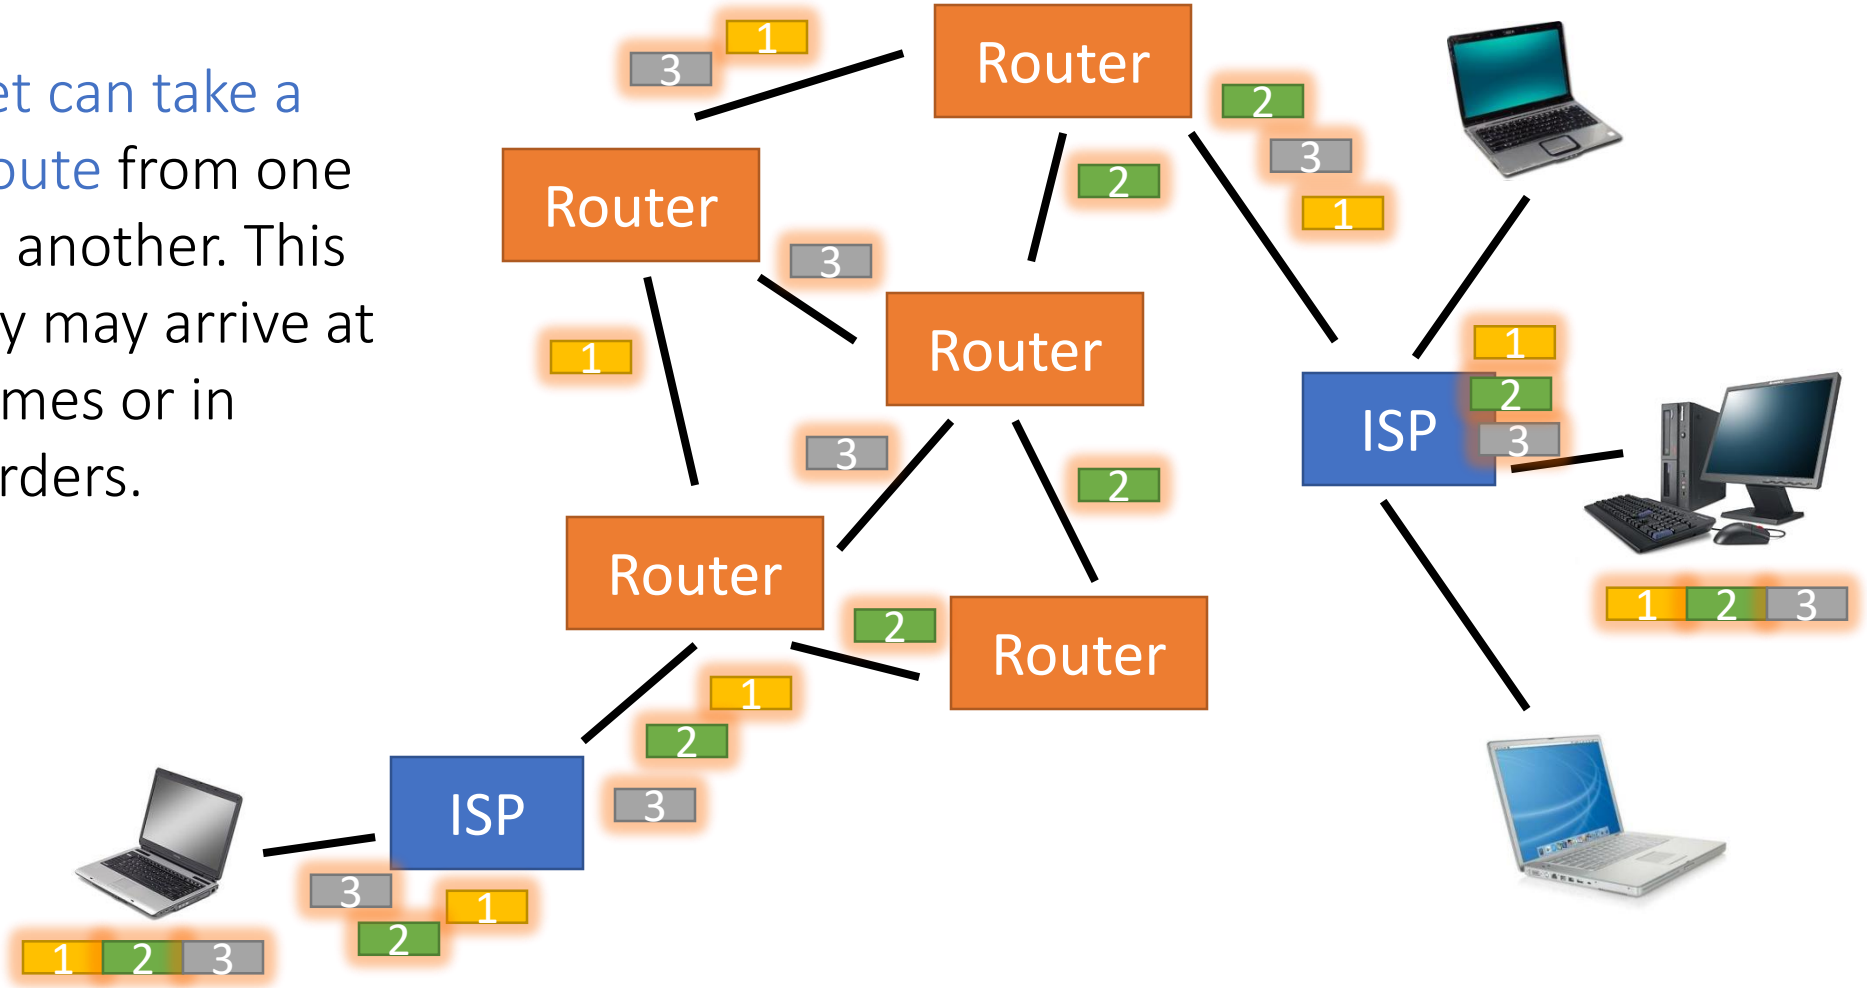

Packet Switching

- There are no guarantees that all packets come through the same wires
- There are no guarantees about the order that your computer will receive the packets
- There are no guarantees about all the packets arriving at your computer
- There are no guarantees that the packets were not corrupted (parity bit doesn’t check out)

Buffering

Buffering occurs when your browser can show some of the data without showing all of it.

If you receive the first packets with the first part of a movie, you can watch that part while it receives more.

If the browser is missing data, then it may wait until it receives the next correct packet.

This concept also works for websites. It may display the content before the images, for example.

Fault Tolerance of the Internet

If a computer that's on the internet goes down, what happens?

- Your computers go on and off the network all the time
- Nothing happens to the routers, ISPs, or anything else
- Your ISP assigns a new IP address (or the same one) each time any computer goes up and down
- Your ISP checks to see if your computer is off the network before it tries to send data to it

Fault Tolerance of the Internet

If a website (server) goes down, what happens?

- If it is the only server at a company, the same thing as if your computer goes down. The company doesn't get any web traffic because it is unreachable
- If it is not the only server to get to the same website, then the website's ISP reroutes the data to a new server or creates a new server on the network with the new IP address

- Activity: different people ping google.com
  - Google is so busy that multiple servers are available all the time.
  - If any one goes down, no big deal.

Fault Tolerance of the Internet

What happens if a DNS server goes down?

- Similar to a server, there isn't just one DNS server. They share data and talk to each other. If one goes down, your router finds another one.

Fault Tolerance of the Internet

How do countries turn off their internet?

- Typically data comes into each country through large wires and distributed to different ISPs (often country-run in these cases).
- If countries are the ISP, they can just prevent the data from reaching routers to direct the traffic
- If there are multiple ISPs, typically there is a router just inside the country that the government can shut off

Fault Tolerance of the Internet

How does data travel across an ocean?

- https://www.nytimes.com/interactive/2019/03/10/technology/internet-cables-oceans.html

What would happen if there was an earthquake or volcano that severed this connection?