Lecture 25 – Content Delivery Networks (CDN) and the Edge

Lecturer: Venkat Arun

Datacenters continued briefly: Datacenter Tax

- Datacenter networks are extremely fast
  - In the upper layers,100-400 Gbit/s links are common
  - Individual servers can have links with capacity up to 100 Gbit/s
- If you try to open a TCP socket (say, using code in assignment 1), you will have a hard time sending at these speeds even over localhost!
  - Live demo using iperf – could only reach ~140 Gbit/s. It uses 100% of a CPU just doing TCP/IP! And this reduced since last year!
- This is called datacenter tax. It is estimated that 30% of CPU cycles are spent on just the network stack (TCP/IP, serialization/deserialization (e.g. JSON, protobufs), HTTP, and TLS)
- SmartNICs are a way to offload some of this workload to dedicated hardware, since you do not need the full power of a modern CPU

Fun experiment running iperf to understand datacenter tax

- On my Mac I only got 15 Gbit/s when iperf started a TCP connection over localhost (that is, from my machine to itself). Last year, I got 140 Gbit/s. I do not know what changed
- On Linux, I got more. I tried starting multiple TCP streams and the *total* throughput looked like this. What do you think is going on?

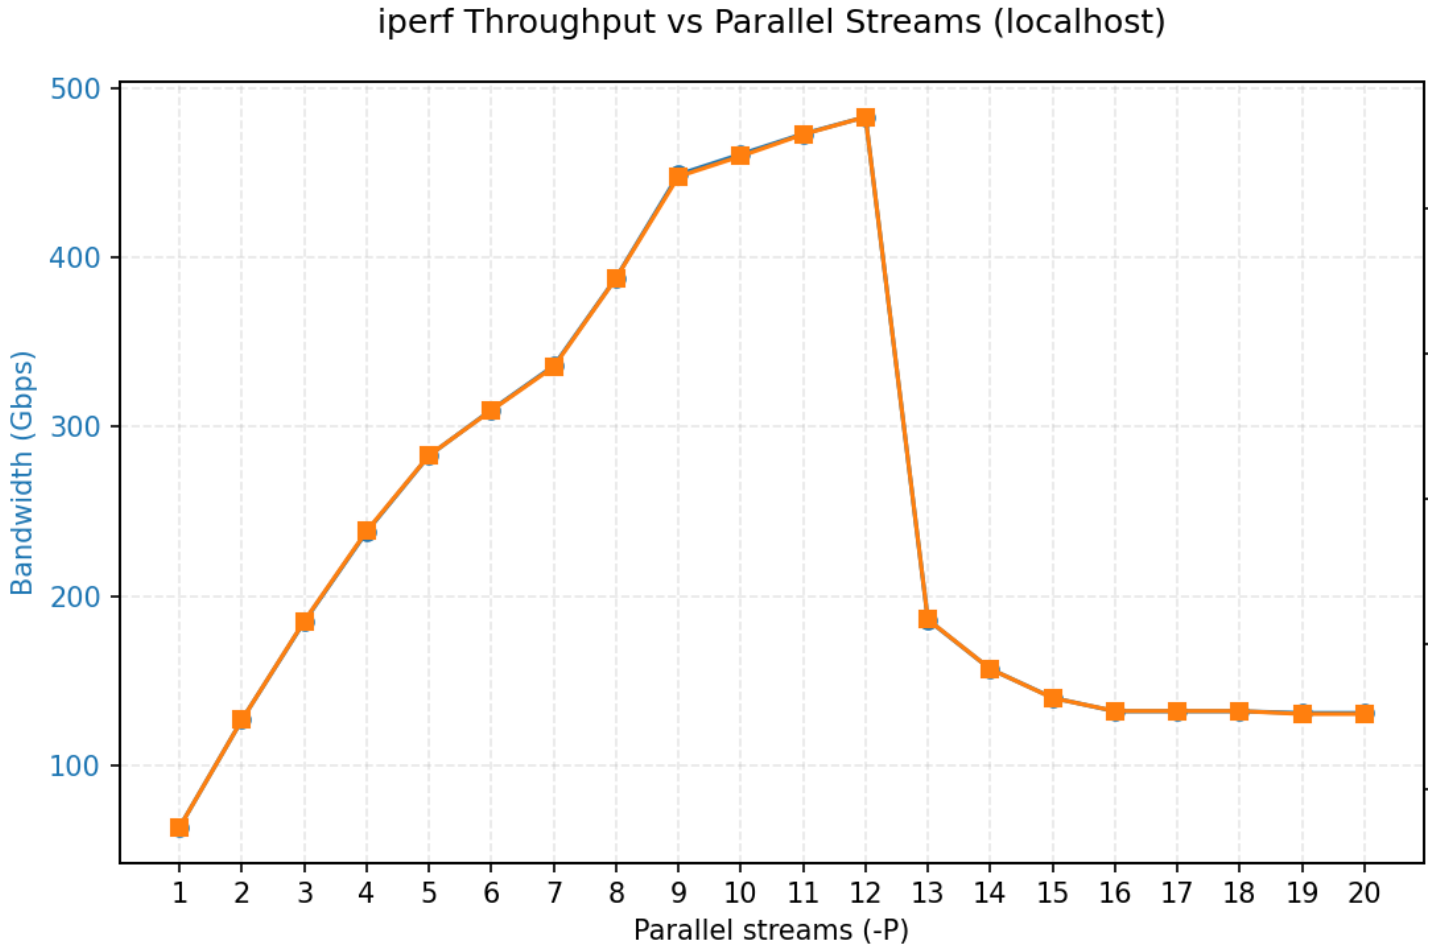

Datacenters are in a lot of places, but that still isn’t enough

Amazon’s locations seen in our class on datacenters

Akamai’s locations. Source: Akamai

What are CDNs/Edge computing?

- TL;DR - Have more, smaller datacenters
- Advantages:
  - Lower latency to client
  - Cheaper bandwidth to client, since it costs less to transmit data over shorter distances
    - Q: We don’t see this effect in our internet bill. Why?
- Disadvantages:
  - Less data storage capacity per location. Desired data might not be available
  - Less compute and memory capacity per location. Might have to redirect to another nearby location
  - Harder to manage.
    - E.g. some claim that, in theory, all the text content of Twitter can be managed from a single large server with a super-fast in-memory database. This is much easier to manage

Discussion: When should you use the edge vs the datacenter?

Example use case: delivering static content

- Websites have a lot of static files
  - Examples: Javascript code, static images like logos, CSS
  - The commonly used ones can be pre-distributed to the edge locations
  - HTML is loaded from a server in the main datacenter. HTML contains URLs for the other resources to be loaded. These can point to the CDN locations

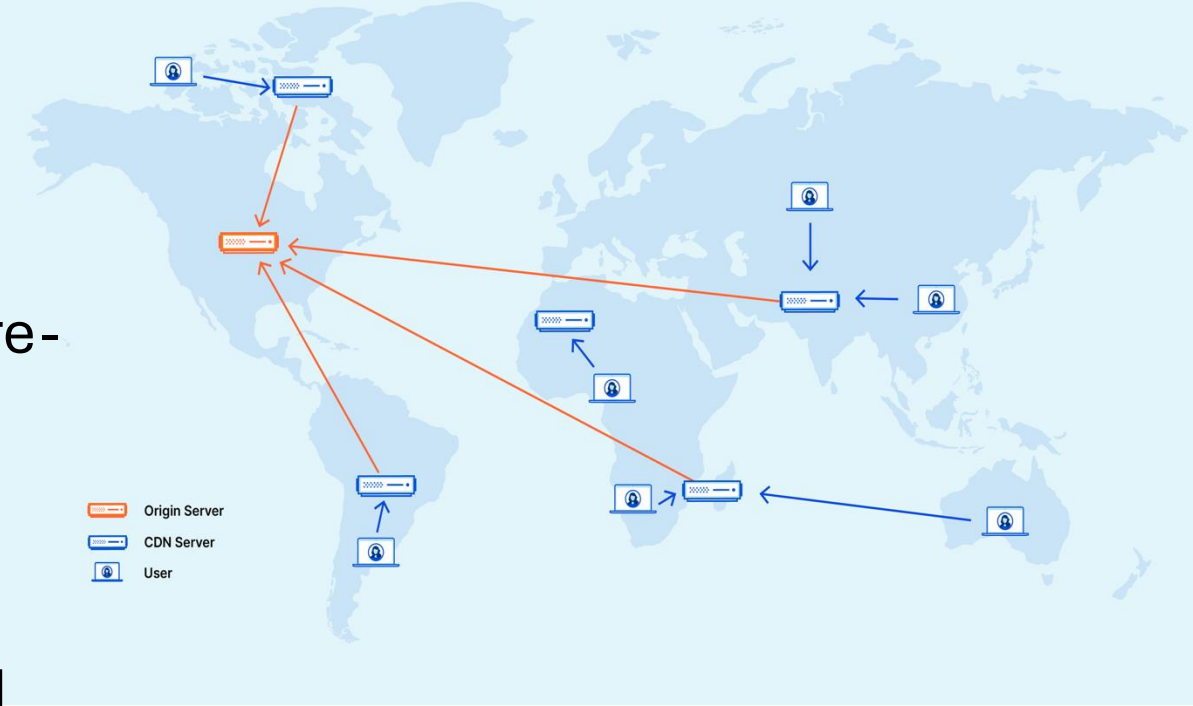

Source: https://www.cloudflare.com/learning/cdn/what-is-a-cdn/

Live Demo: Course website

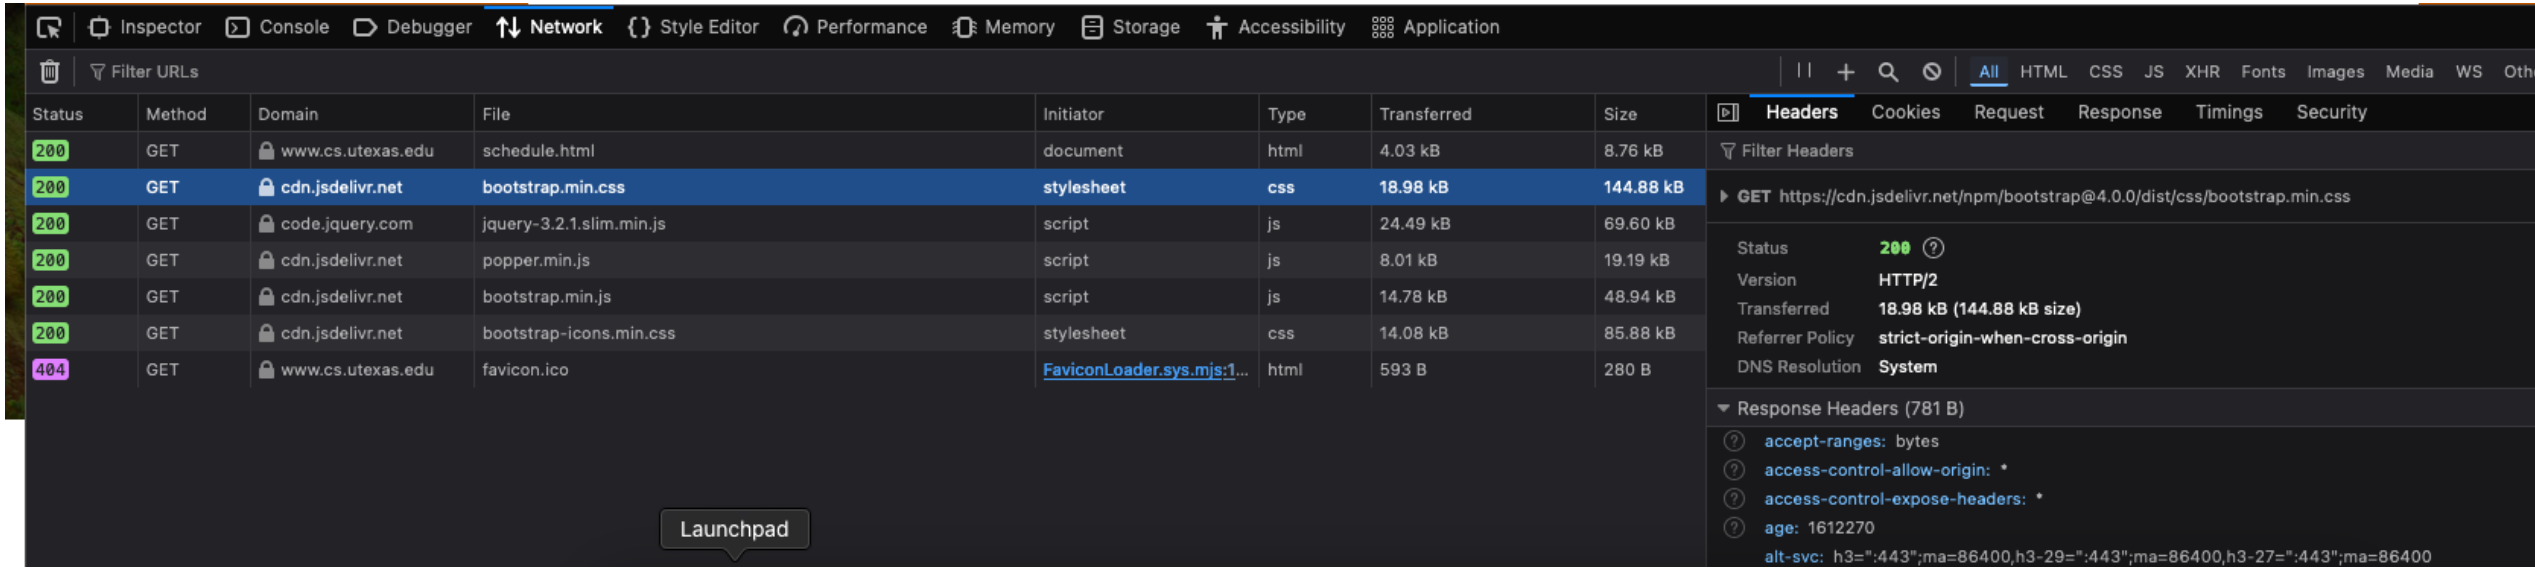

Image: “Network” tab of “Inspect element” in firefox for some the course website

- I used some standard libraries used by a lot of websites. Rather than serving those files to the user directly, I used a link to CDNs other people build and maintain.
- An additional advantage is that the web browser is likely to have these files already cached, in which case an entire request can be avoided

How does the browser know which CDN replica to access?

- We want the browser to go to the nearest replica
- However, the URL looks like: https://cdn.jsdelivr.net/npm/bootstrap@4.0.0/dist/css/bootstrap.min.css
- This contains no information about where to go
- Instead, people configure the DNS to give different IP addresses for the same domain depending on where the user is connected
- This is called **DNS load balancing**
  - It is also used to route clients to the closest origin server (i.e. large datacenter)

Live demo: Do a DNS lookup from various locations

- I will use the command “nslookup cdn.jsdelivr.net 8.8.8.8” after connecting to various locations using a VPN
- The VPN makes all packets from laptop come from a different (configurable) location. This means that a DNS query made over a VPN will cause the Domain Name System (DNS) to return the result that it would have returned to a laptop that was actually in that other location
- Note: 8.8.8.8 is the IP address of a globally accessible DNS server hosted by google.com. If I had not specified it, my laptop would have use a cached value. Another option was to use the command: “nslookup -querytype=NS cdn.jsdelivr.net”

Live demo: useful website - builtwith.com

- For example, this page lists the websites that use the Akamai CDN in the US

https://trends.builtwith.com/websitelist/Akamai/United%20States

- It also guesses what services other websites are using, revealing the true complexity of today's internet services

Live demo: Walmart.com

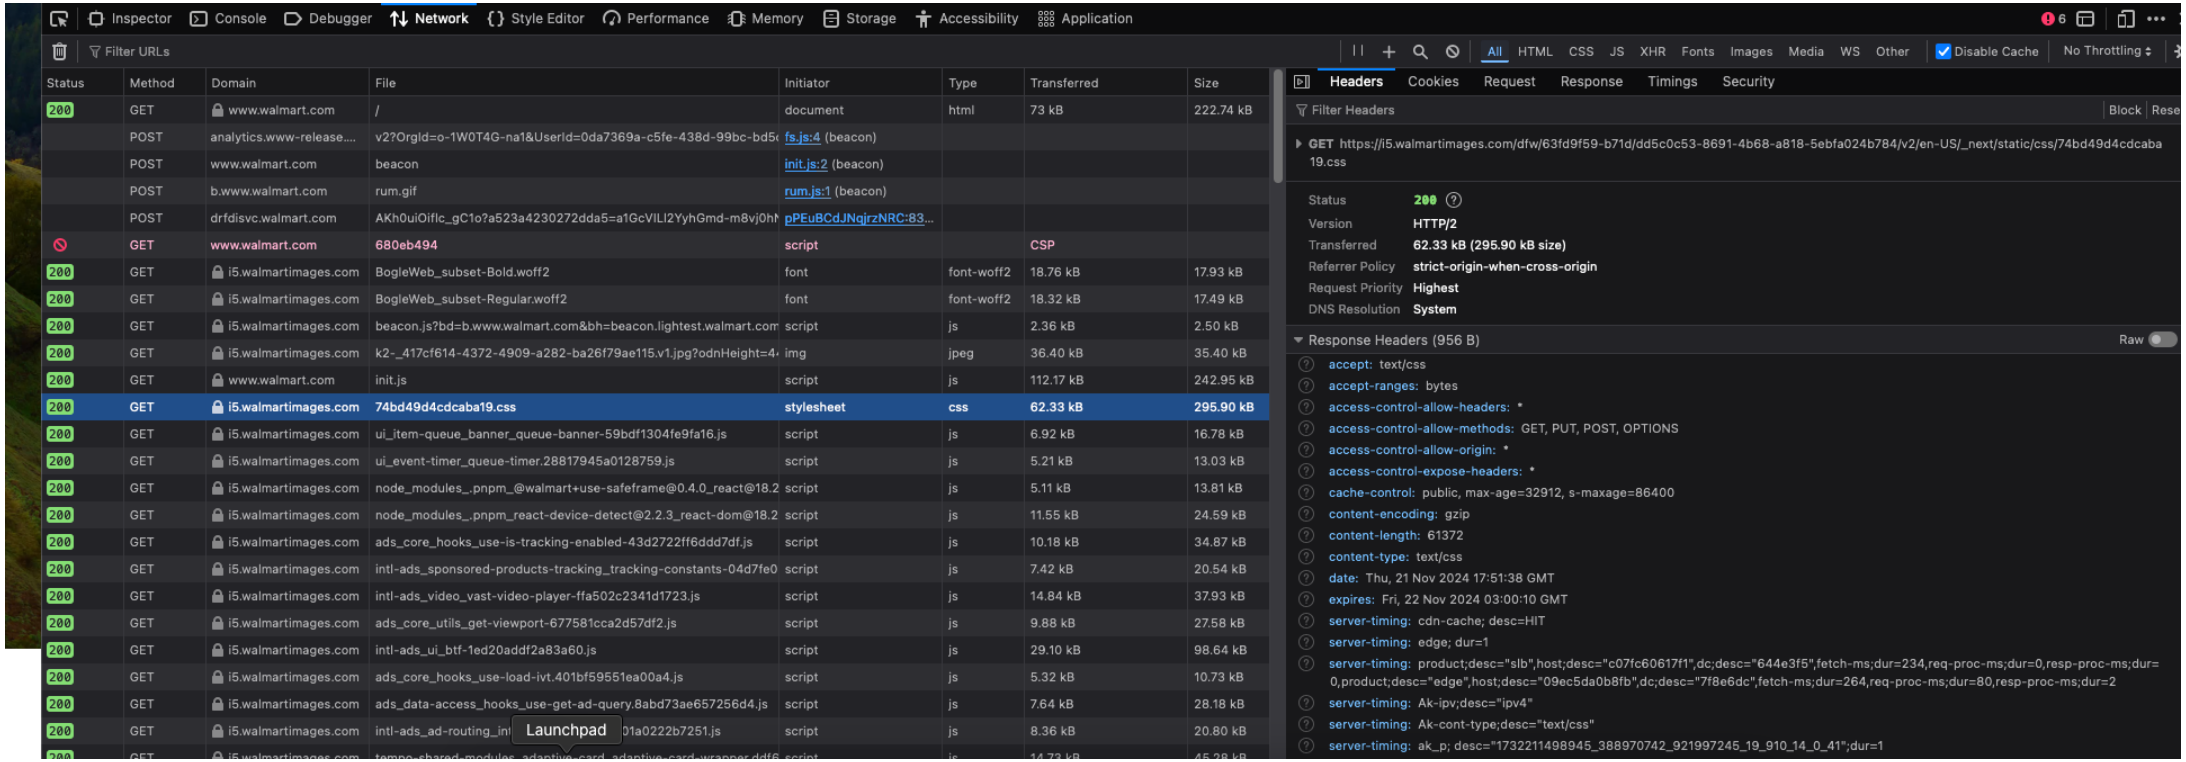

- The complex URL tells the browser which domain to download an image from
- Based on nslookup, it seems to use akamai and Fastly as its CDN depending on the location

Instapoll question

- Suppose, I think that DNS load balancing is too complex.
  - That is, I do not want to make the DNS return different IPs depending on the client's location
- Instead, I want the origin server to look at the client's IP address and send a URL that will work without DNS load balancing
- I have two options in mind. Which one works?
  - <edge-location-id>.mydomain.com/<image-id>
  - mydomain.com/<edge-location-id>/<image-id>
- Why?

Instapoll answer

- <edge-location-id>.mydomain.com/<image-id>
- Why? In this case, the DNS receives <edge-location-id>.mydomain.com, which means even if DNS returns the same IP address independent of the client's location, the origin server cans end different <edge-location-id>s to different clients based on their locaiton

- In the other case, the DNS only receives "mydomain.com", to which it will be forced to return the same IP address to all the requests.

Other uses for the edge

- Edge storage:
  - In addition to what we discussed, a common “static file” served from CDNs is stored video (E.g. youtube, Netflix). They are large and frequently accessed(thus, easy to cache)
- Edge compute is an emerging trend (we won’t discuss this in detail in this course) :
  - Cloud gaming puts a lot of the compute on a remote server. Latency requirements force the compute to be close to the user
  - RANs in 5G networks