CSC2231: Akamai

http://www.cs.toronto.edu/~stefan/courses/csc2231/05au

Stefan Saroiu

Department of Computer Science

University of Toronto

Administrivia

- **Project proposals due today!!!**
- **No lecture on Monday:**
  - Eat well!
- **Shuffled the order of some lectures:**
  - Next Thursday:
    - *Web Caching and Zipf-like Distributions: Evidence and Implications*. Lee Breslau, Pei Cue, Li Fan, Graham Phillips, Scott Shenker, Infocom 1999.
- **No lecture on October 17th because of Cascon**

Web Performance Techniques

- **Four Common Categories**
  - Web Transport Optimization
  - Server Scalability
  - Document Relocation
  - Browser Optimization

Web Performance Techniques

- **Four Common Categories**
  - Web Transport Optimization
    - HTTP1.1 (Persistent Connections), TCP tweaks
  - Server Scalability
    - Clusters, OS enhancements
  - Document Relocation
    - Web caching, prefetching
  - Browser Optimization

Benefits of Proxy Caching

- **Proxy caching is the most commonly used method to improve Web performance**
  - Duplicate requests to the same document served from the cache
  - Hits reduce latency, network utilization, and server load
  - Introduces problems:
    - Misses increase latency (extra hops)
    - cache consistency

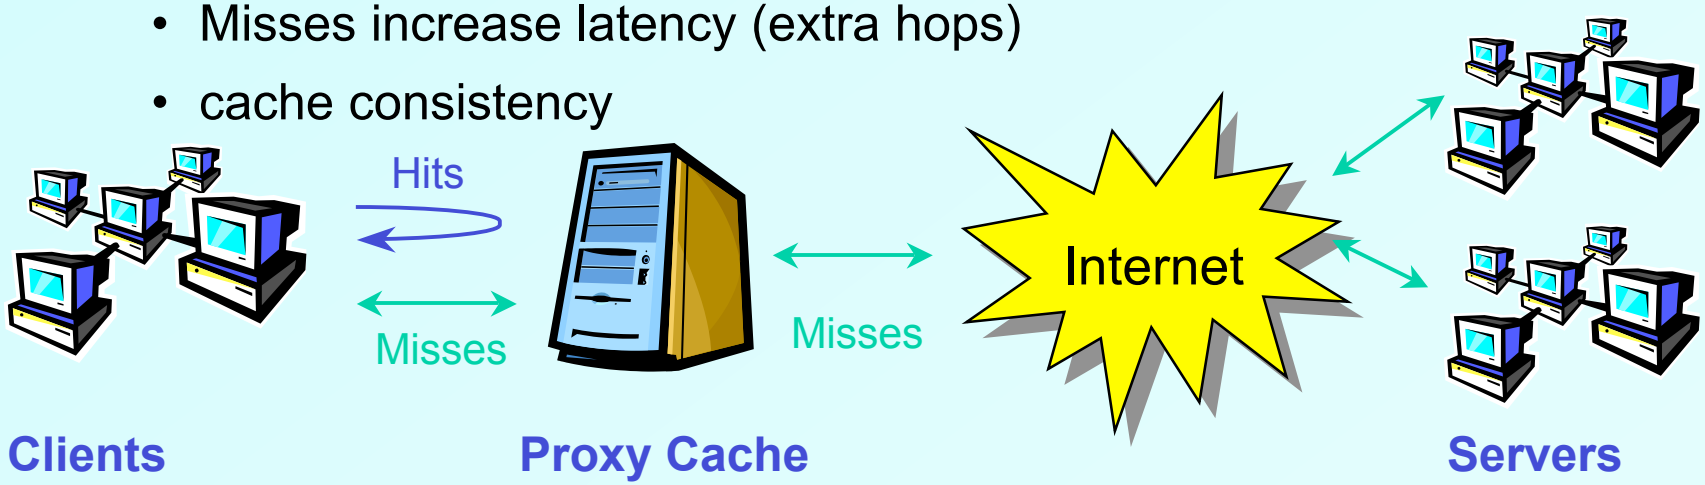

Cache Consistency

- **Fresh-enough is good-enough**
- **One writer,** many **readers**
  - Most content changes slowly wrt # reads

- **Cache consistency governed by standards**
- **"Expiration" based cache consistency**
  - Expires timestamp on each object
  - Cache revalidates content beyond that time
- **Why not callbacks?**

Data Update Propagation

- **IBM's Olympic Games website**
  - Back-end: database, content generation
  - Front-end: web caches (accelerators)
- **Cache consistency:**
  - Manager has callback list with <cache,object>
  - Developer annotates objects with data dependencies
- **Graph of data->object dependencies**
  - Invalidate caches and proactive regeneration of objects

Multiple Caches

- **Does it make sense to have caches upstream to local cache?**

Multiple Caches

- **Does it make sense to have caches upstream to local cache?**
- **Then, why do Internet servers make heavy use of reverse Web caches?**
  - Amazon, Google

Content Distribution

- Lots of excitement?
- Akamai, Digital Island/Sandpiper, Speedera
- What is a Content Distribution Network (CDN)?
  - Outsourced caching and replication services

Cache Deployments

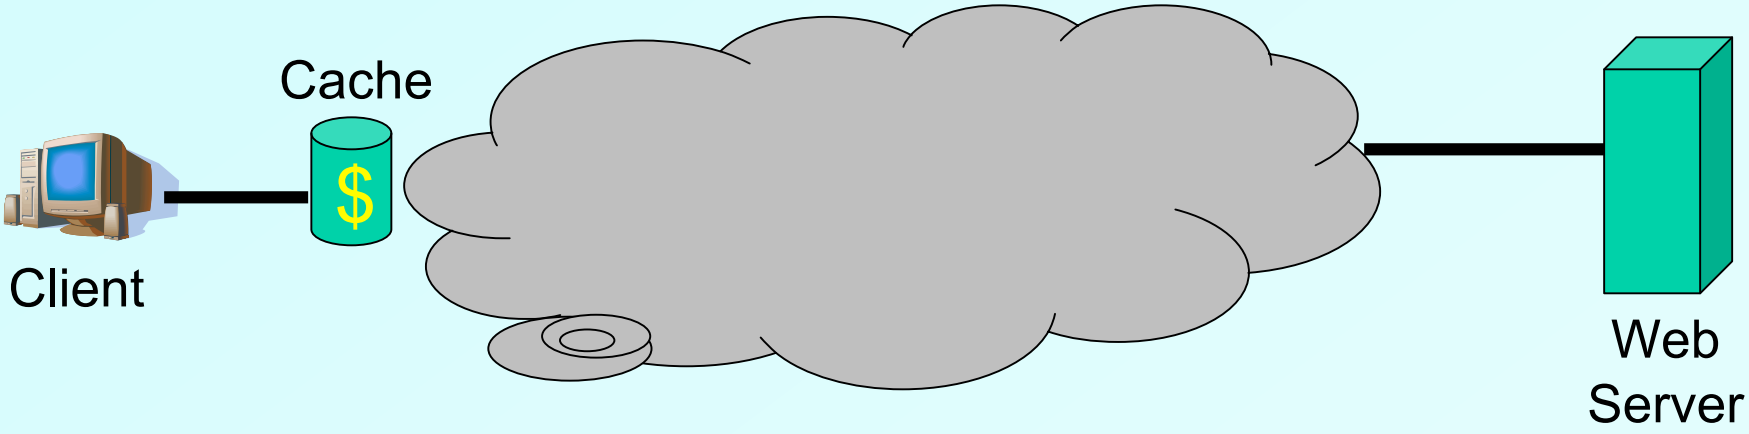

Where else?

Cache Deployments

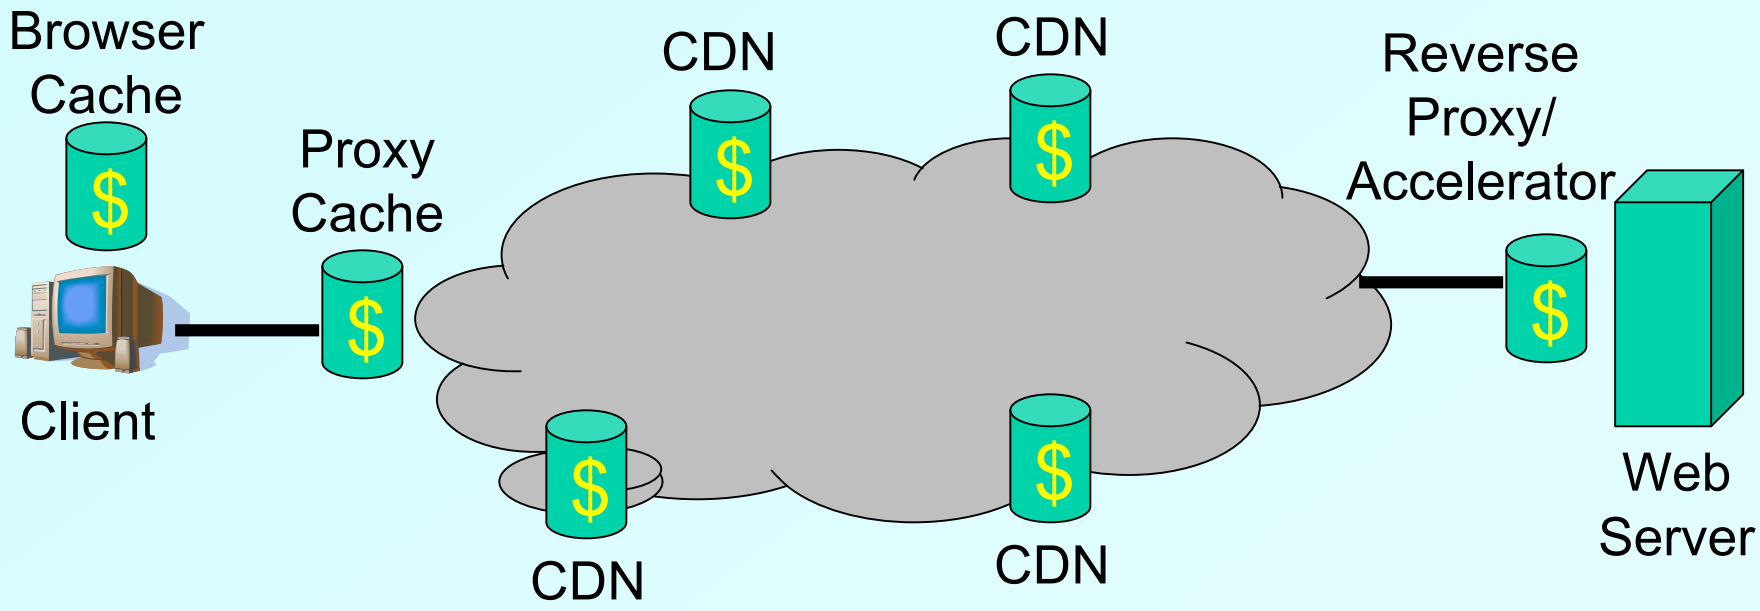

Content Providers' Advantages

- **CDN provider maintains networks and servers**
  - Capacity management
- **Sharing resources across a large number of sites**
  - Economy of scale
  - Control of content placement and routing
- **Protects content provider from unpredictable load bursts**
- **Communication between content provider and CDN network is not governed by standards**
  - Don't even need to use HTTP
  - Can cache "uncacheable" documents
  - Can deploy alternative cache consistency
  - Can place requirements on content providers

CDNs' Challenges

CDNs' Challenges

- **How to replicate content?**
- **Where to replicate content?**
- **How to find replicated content?**
- **How to choose among know replicas?**
- **How to direct clients towards replica?**

- **Akamai**

How Akamai Works

- **Clients fetch html document from primary server**
  - E.g. fetch index.html from cnn.com
- **URLs for replicated content are replaced in html**
  - E.g. <img src="http://cnn.com/af/x.gif"> replaced with <img src="http://a73.g.akamaitech.net/7/23/cnn.com/af/x.gif">
- **Client is forced to resolve aXYZ.g.akamaitech.net hostname**

How Akamai Works

- **Root server gives NS record for akamai.net**
- **Akamai.net name server returns NS record for g.akamaitech.net**
  - Name server chosen to be in region of client's name server
  - TTL is large
- **G.akamaitech.net nameserver choses server in region**
  - Should try to chose server that has file in cache
  - TTL is small

How Akamai Works

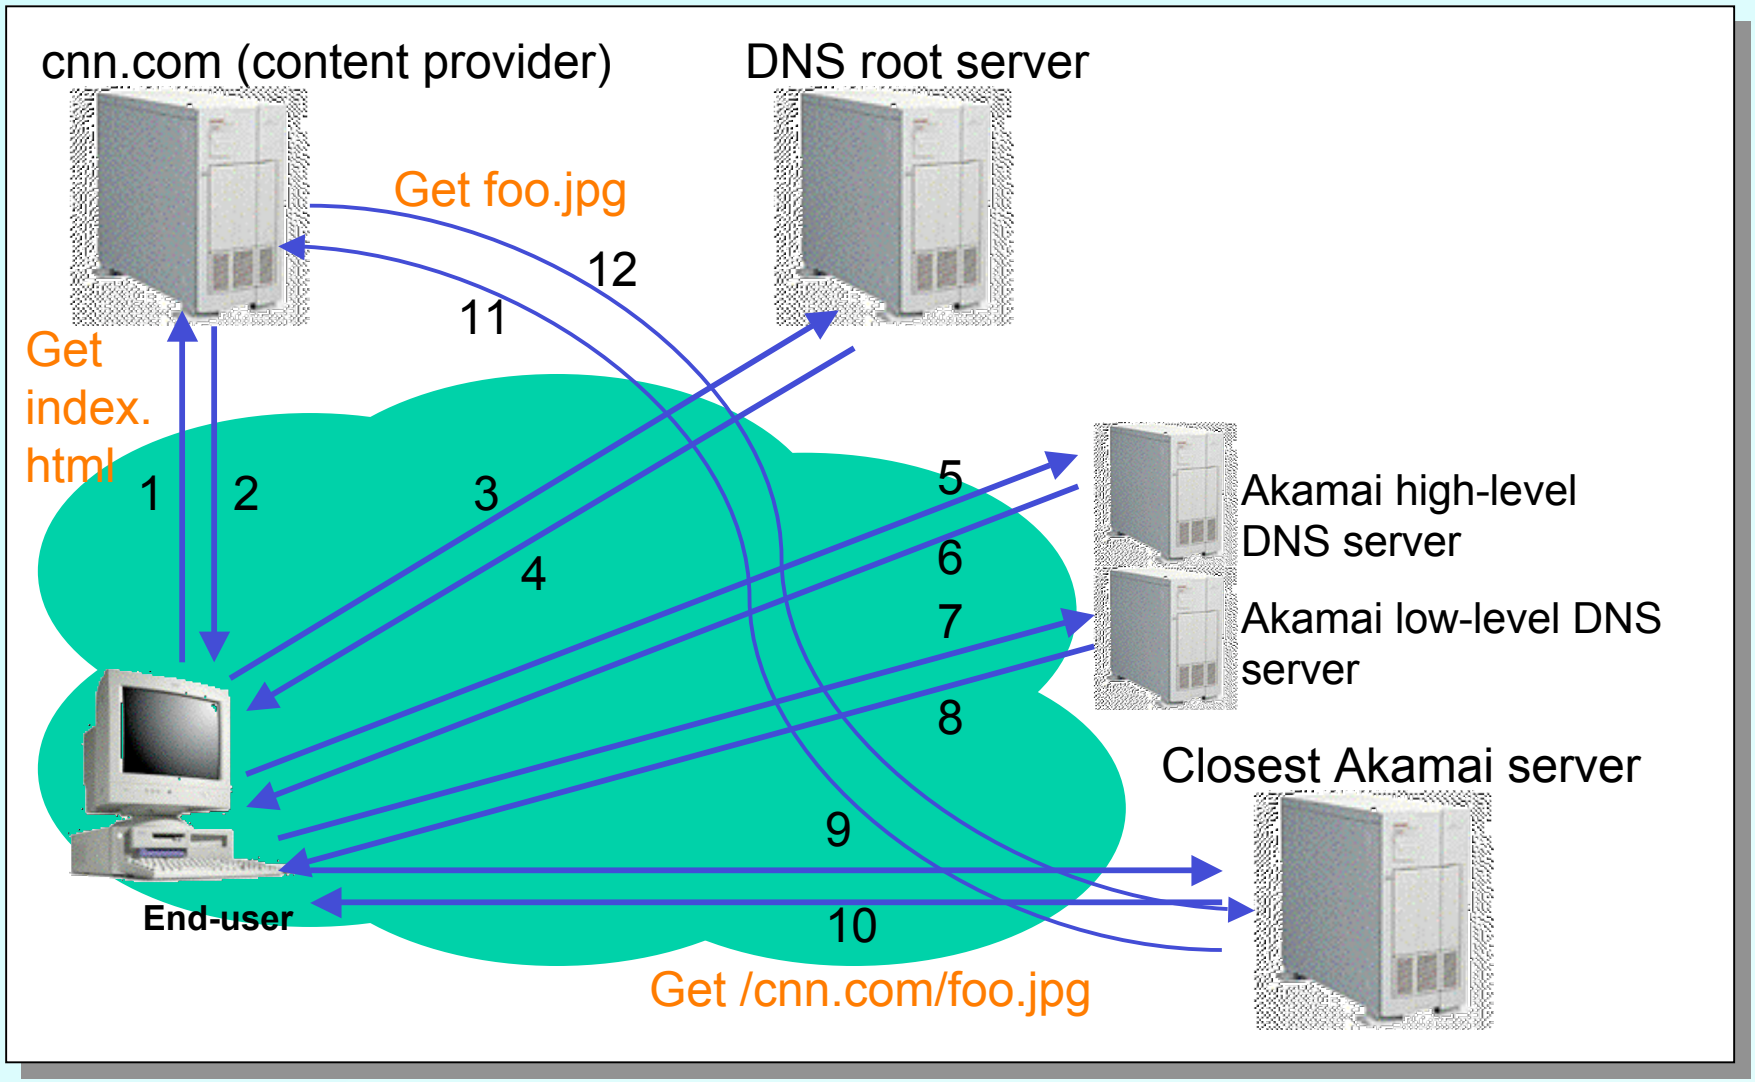

Akamai – Subsequent Requests

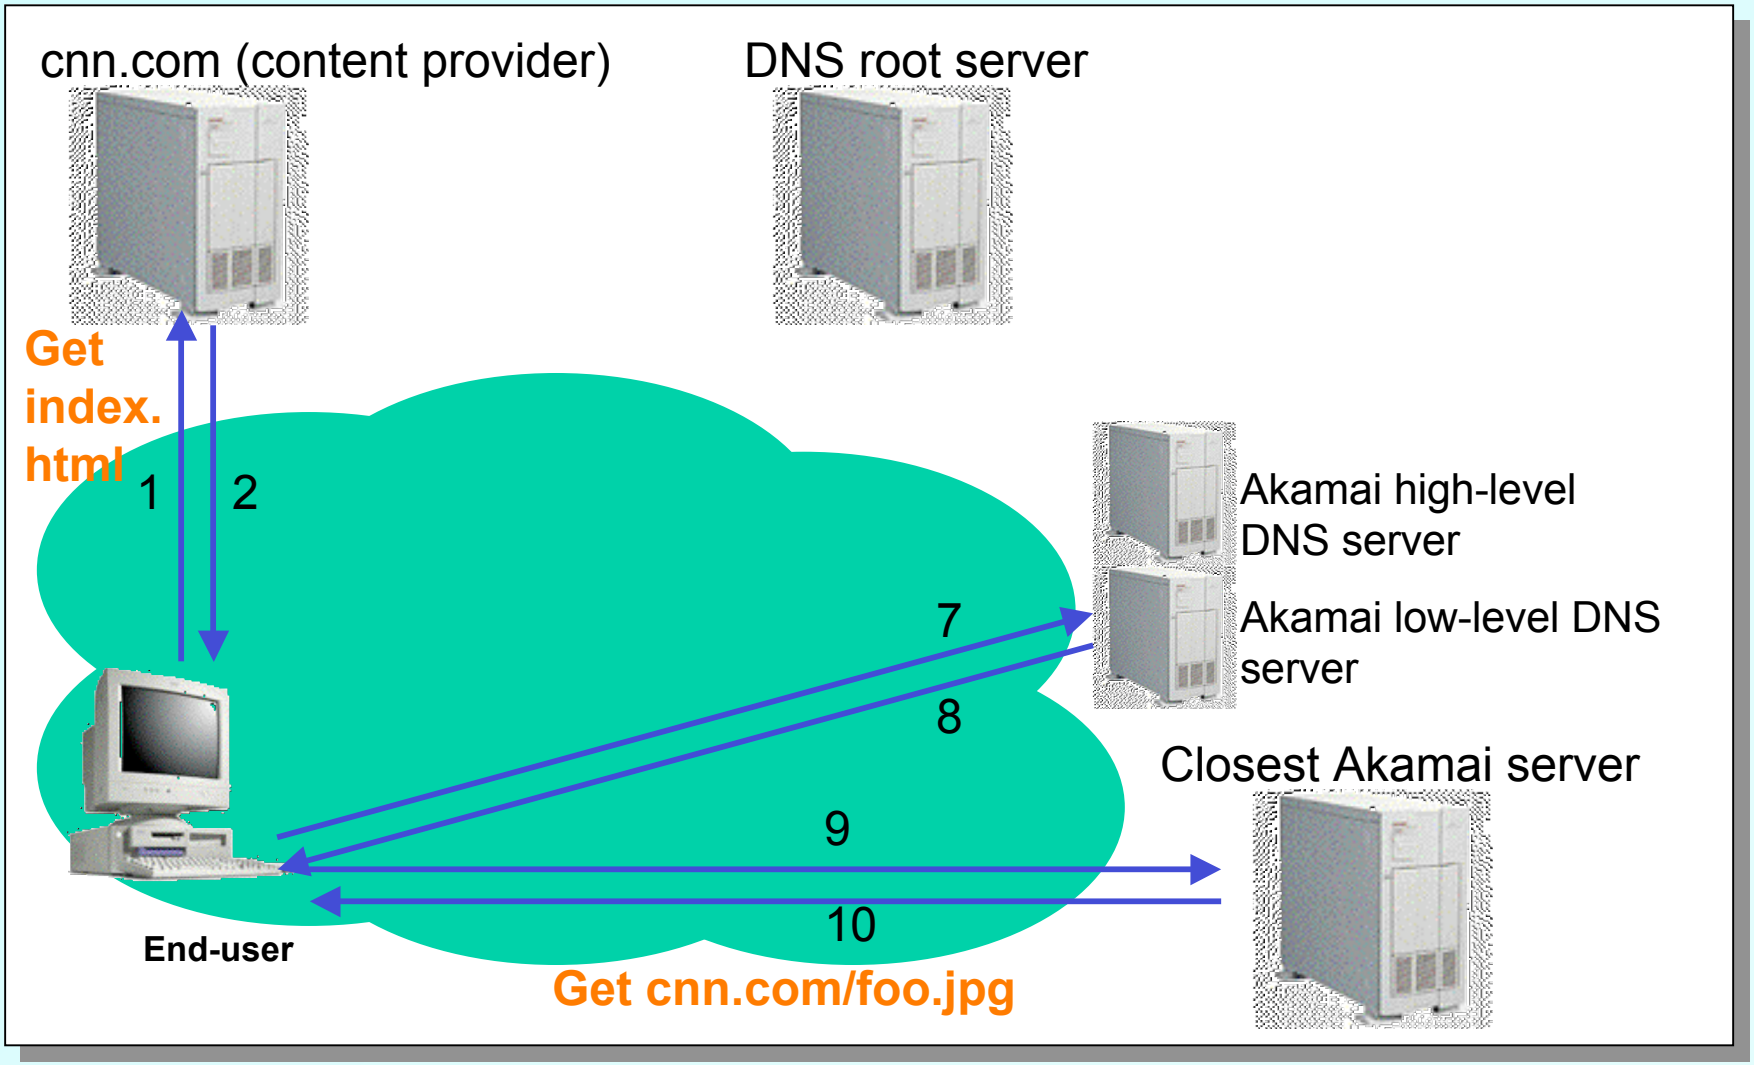

TCP Mambo-Jambo

- **Bandwidth-delay product:**
  - Number of outstanding (i.e., in-flight or unacknowledged) packets cannot exceed the bandwidth-delay product
- **A 1Mbps link**
  - RTT: 1 sec, cannot have more than 128 KB outstanding
  - RTT: 100ms, cannot have more than 1280 KB outstanding

CDN's Reduced Latency Benefits

- **DNS round-trip**
- **TCP handshake**
- **~ 8 RTTs to fill 1Mbps pipe**
- **Total: 128KB over 11 RTTs**
- **Coast-to-coast RTT: 60ms**
- **Toronto to Akamai RTT: 2-3ms**
- **Total RTT for filling-out pipe:**
  - Without Akamai: 600ms
  - With Akamai: 30ms

Lets look at a study

- **Zhang, Krishnamurthy and Wills**
  - AT&T Labs
- **Traces taken in Sept. 2000 and Jan. 2001**
- **Compared CDNs with each other**
- **Compared CDNs against non-CDN**

Methodology

- **Selected a bunch of CDNs**
  - Akamai, Speedera, Digital Island
- **Selected a number of non-CDN sites for which good performance could be expected**
  - U.S. and international origin
  - U.S.: Amazon, Bloomberg, CNN, ESPN, MTV, NASA, Playboy, Sony, Yahoo
- **Selected a set of images of comparable size for each CDN and non-CDN site**
  - Compare apples to apples
- **Downloaded images from 24 NIMI machines**

Response Time Results (II) Including DNS Lookup Time

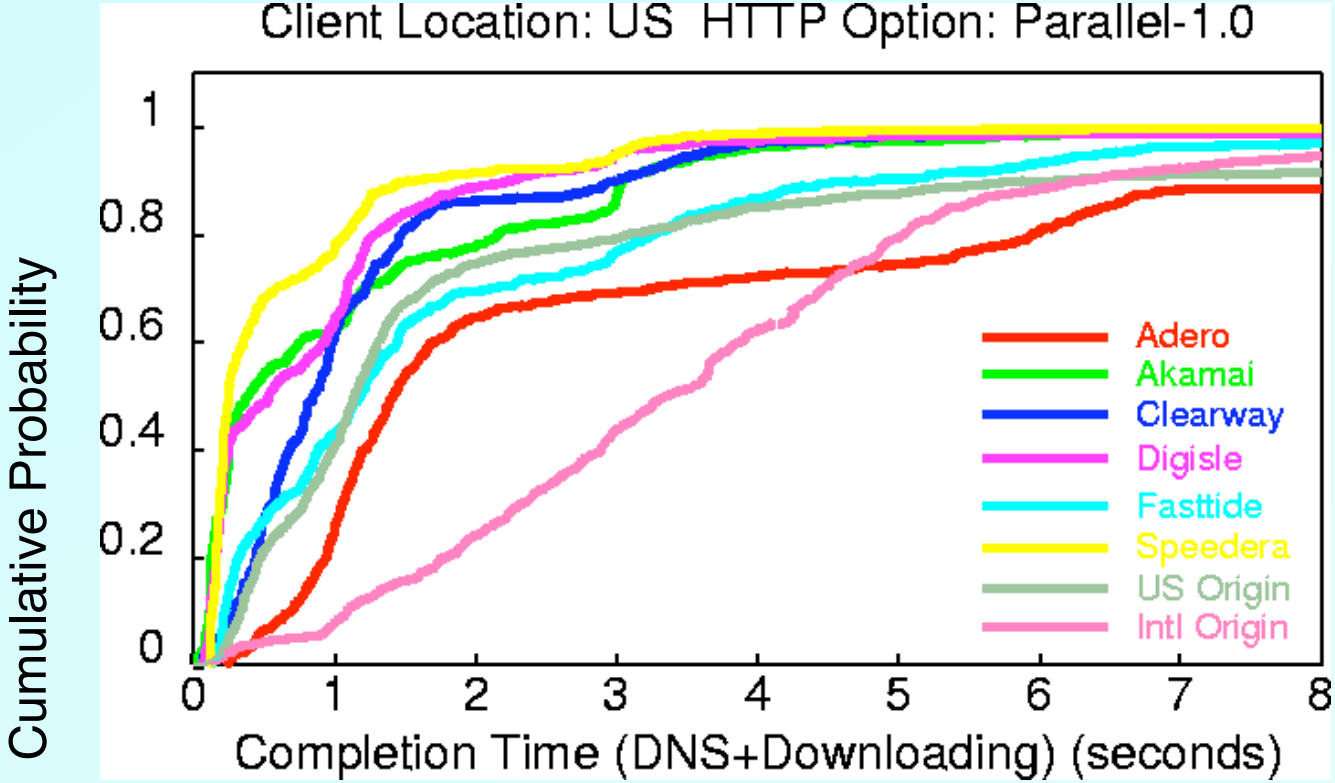

From Ken Birman's slides

Response Time Results (II) Including DNS Lookup Time

About one second

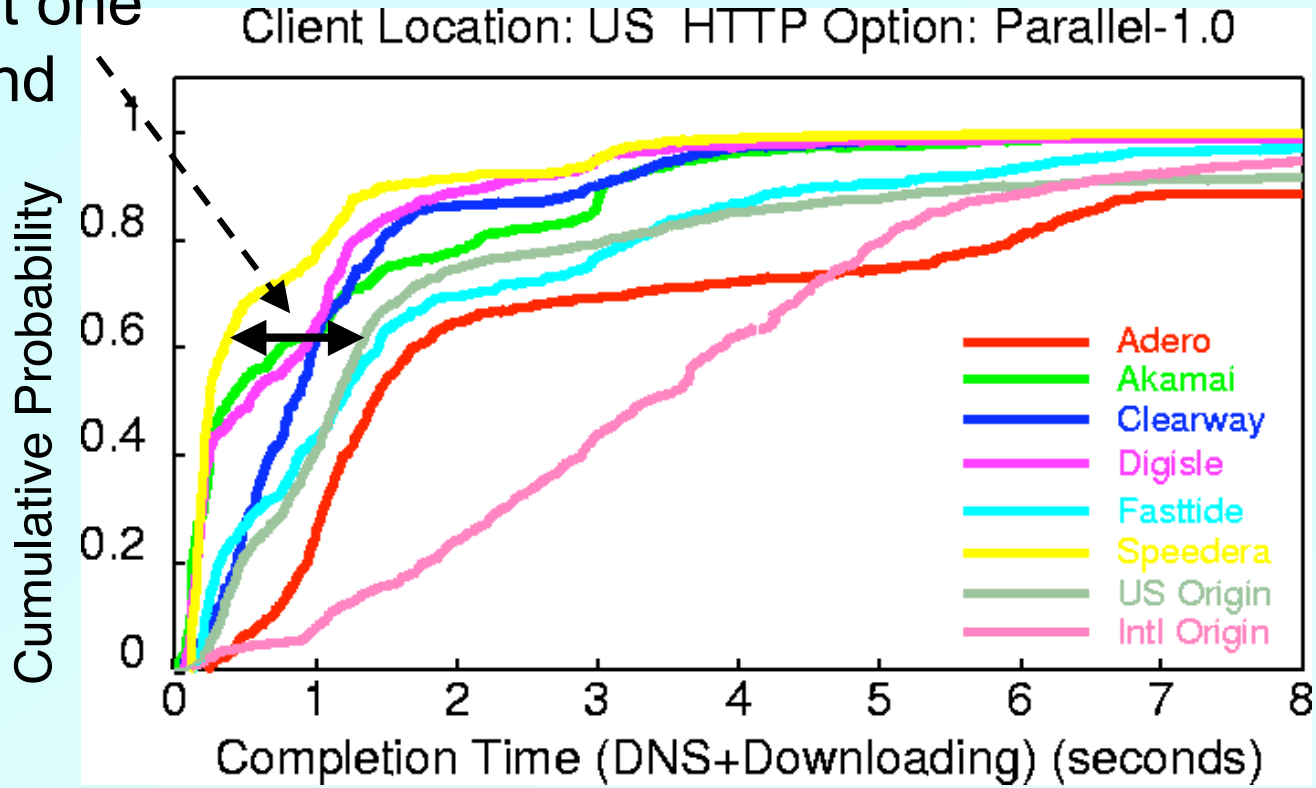

CDNs generally provide much shorter download time.

CDNs out-performed non-CDNs

- **Why is this?**
- **Lets consider ability to pick good content servers…**
- **They compared time to download with a fixed IP address versus the IP address dynamically selected by the CDN for each download**
  - Recall: short DNS TTLs

Effectiveness of DNS load balancing

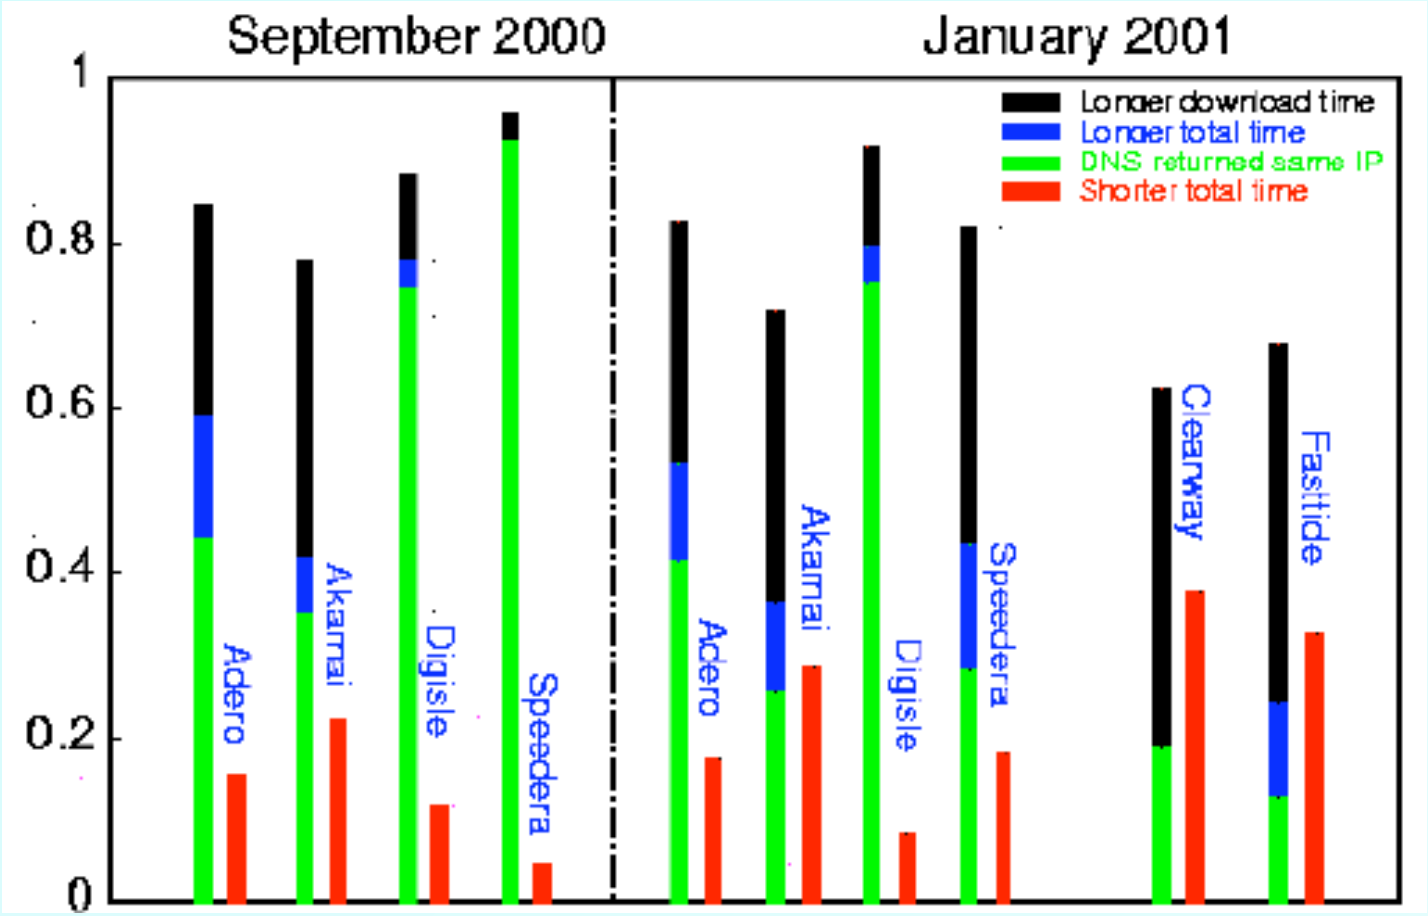

From Ken Birman's slides

Effectiveness of DNS load balancing

Black: longer download time

Blue: shorter download time, but total time longer because of DNS lookup

Green: same IP address chosen

Red: shorter total time

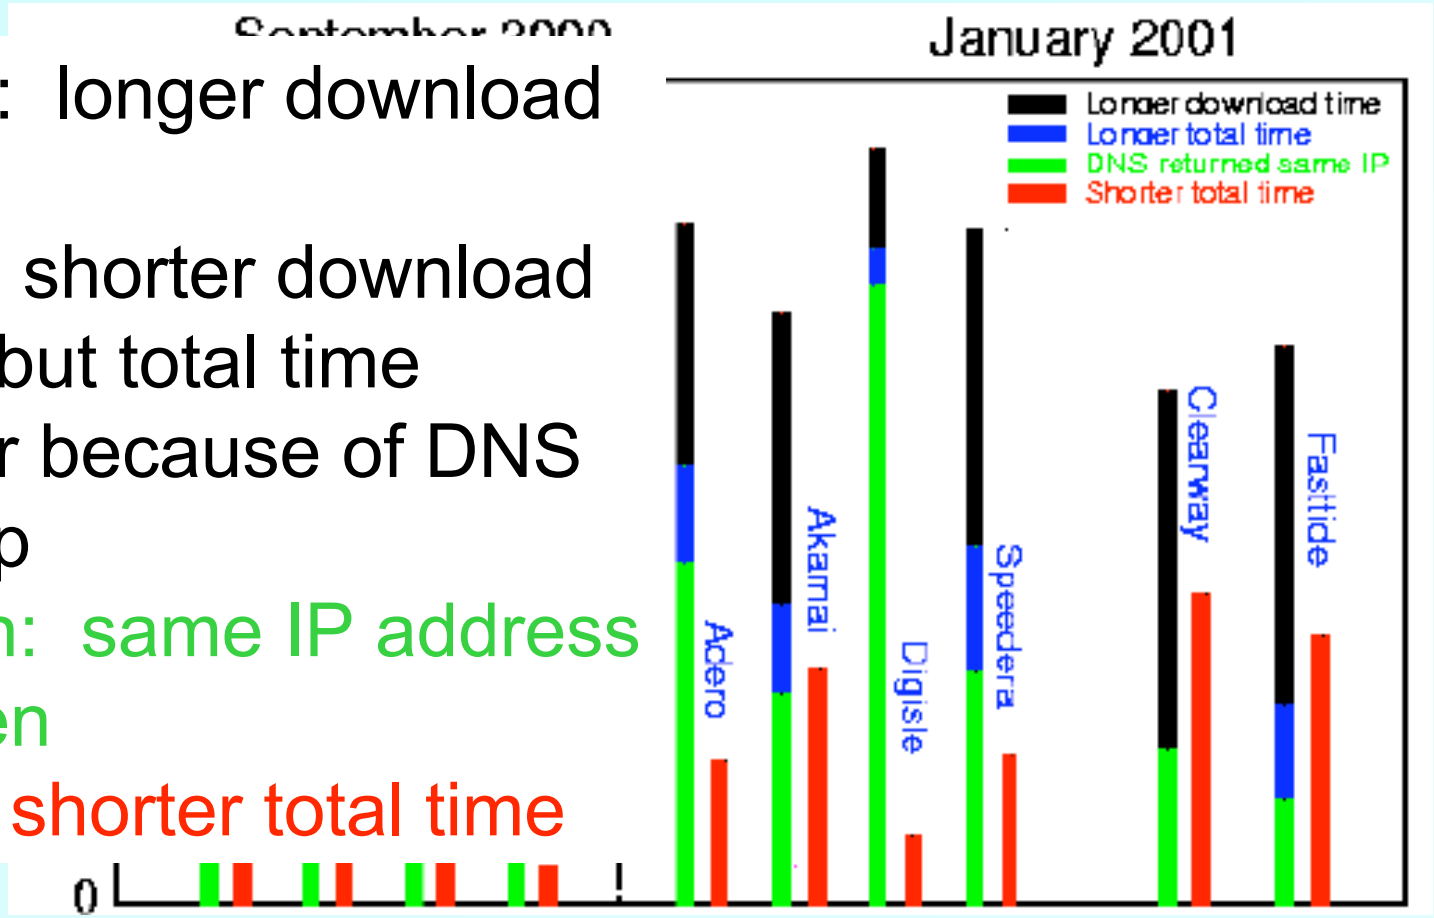

DNS load balancing not very effective

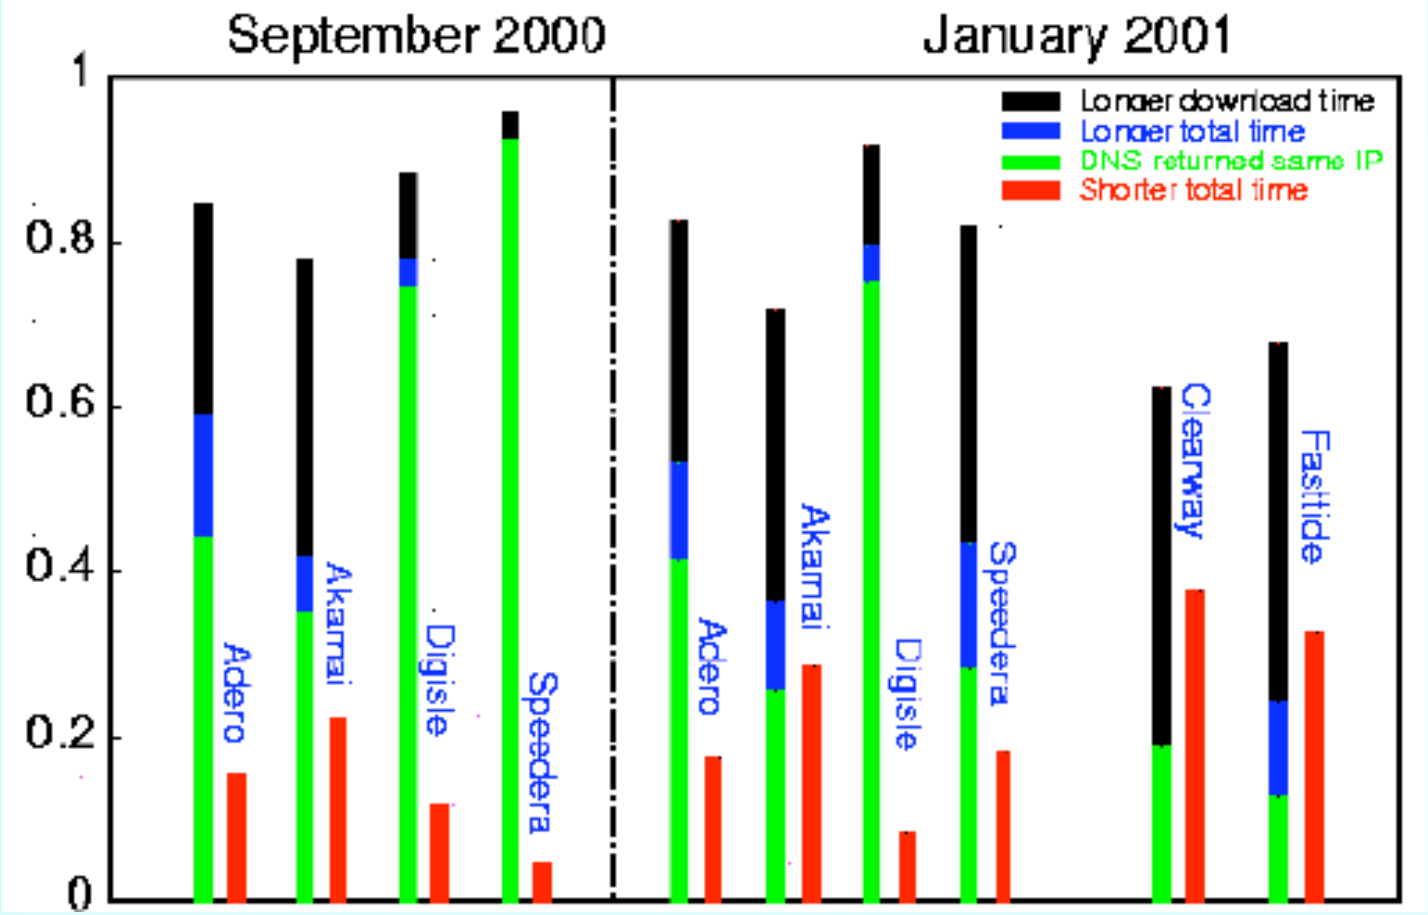

From Ken Birman's slides

Discussion

- **Networks are becoming bigger and faster**
  - Last-mile is the REAL problem
- **Do CDNs make any sense anymore?**

Discussion

- **CDNs for large content**
  - Video?

Discussion

- **CDNs as insurance against /. effect?**
  - Should we still use DNS-redirection for this?