

Content Distribution Network

Kai Shen
Dept. of Computer Science, University of Rochester

On Assignment #2

Important: run each test a large number of times to get rid of timer cost and random errors.

- Question #1:
 - the cost of function call is around 6-7ns.
 - common error: timing each function call.
- Question #2:
 - the cost of a context switch is around 0.8us.
- Question #3:
 - process switching cost is around 3-4us based my suggested pipe-based tests
- Question #4:
 - thread switching cost is around 7us based my suggested "condition variable"-based tests.
 - sort of strange since thread switching should be faster; probably because the locking overhead is significant

3/18/2003

URCS 573 - Spring 2003

2

On Assignment #2 (cont.)

- Question #5:
 - UDP roundtrip latency is hundreds of microseconds
 - there should be a big jump between 1,000-byte message and 10,000-byte message
- Question #6:
 - TCP roundtrip latency is hundreds of microseconds
 - connection setup/teardown should add considerable latency
- Question #7:
 - CPU overhead for UDP packet sending is around a few to 10s of microseconds
 - there should be a big jump between 1,000-byte message and 10,000-byte message

3/18/2003

URCS 573 - Spring 2003

3

On Assignment #3

- Measure the throughput and mean response time of a concurrent server.
- The server listens at port 5678 and it accepts simple requests.
- Run it in a Linux machine in the graduate software lab and test it using your client from a different machine.
- Your job is to write your client and do the measurement.
- I will tell you what the server does after the assignment is due.
- Due next Thursday but it should be a simple one.

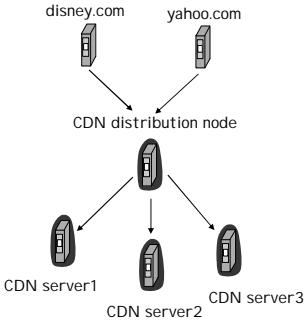
3/18/2003

URCS 573 - Spring 2003

4

Content Distribution Network

- CDNs sit between content providers and content consumer.
- A CDN contains hundreds of servers throughout Internet
- CDN replicates its customers' content in CDN servers. When provider updates content, CDN updates servers



3/18/2003 URCS 573 - Spring 2003 5

How CDN works?

Content provider

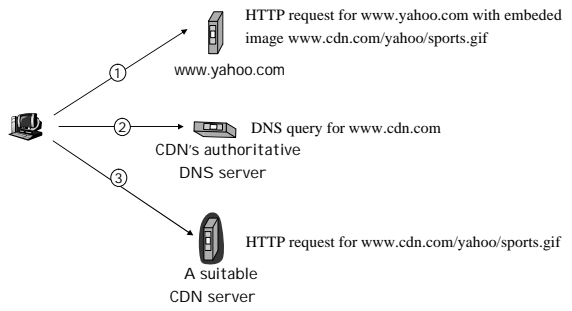
- www.yahoo.com
- Replaces www.yahoo.com/sports.gif with www.cdn.com/yahoo/sports.gif

CDN company

- replicates content at CDN servers
- uses its authoritative DNS server to redirect requests to a suitable CDN server

3/18/2003 CSC 257/457 - Fall 2002 6

How CDN works? (An Illustration)



3/18/2003 CSC 257/457 - Fall 2002 7

Determining a Suitable CDN Server

Heuristics for "suitability":

- Network proximity
- Runtime workload
- Locality

3/18/2003 URCS 573 - Spring 2003 8

Network Proximity

- CDN creates a “map”, indicating distances from local networks and CDN nodes
- when a DNS query arrives at the authoritative DNS server:
 - server determines local networks from which query originates
 - uses “map” to determine a nearby CDN server
- how to get this “map”
 - offline: “ping” each network several times and take the average
 - online: continuously “ping” each network at certain intervals and take the exponentially-weighted moving average

3/18/2003

URCS 573 - Spring 2003

9

Runtime Workload

- monitoring workload at request distributors
 - continuously poll each CDN server for workload at certain intervals and maintain the EWMA
 - randomly poll a few CDN servers for just-in-time load information
- request distribution policy is based on a combination of load and network proximity

3/18/2003

URCS 573 - Spring 2003

10

Locality

- Reduce the working set of each CDN server such that all pages can be directly served from memory.
- How?
 - Control the size of replication server set for each Web object.
 - e.g., for a CDN with 100 servers, each object is only replicated at 10 servers. The exact choice of this 10 servers depends on a hash value of the object URL.
- Does this conflict with network proximity?

3/18/2003

URCS 573 - Spring 2003

11

Summary

Somehow conflicting heuristics for “suitability”:

- Network proximity
- Runtime workload
- Locality
- many possible strategies based on a combination of the above.

How to measure their effectiveness?

- based on simulation
- simulating both the network and the servers
 - we need network topology and link properties
 - we need to simulate server memory/disk subsystem
- we need workload – traces?

3/18/2003

URCS 573 - Spring 2003

12