

Distributed Computing

CMSC 433, March 27th

Is distributing computing different?

- What kinds of distributed computing environments exist?
- Ways in which distributed computing is different
 - Addressing objects
 - Latency
 - Partial failure
 - Concurrency

2

Distributed computing environments

- Usually refers to multiple CPU's
- Shared address space or message passing?
- On same chip, in same room, or across the internet?
 - Latency, failure modes

3

Existing environments

- Seti @ Home
- Server for search engine
- My laptop, PDA, cell phone, MP3 player and digital camera

4

Types of failure

- Machine sleeps
 - wakes up, recovers state
- Machine crash or failure
 - machine may reboot and rejoin
- Network partition
 - network may heal

5

Uniform view of distributed objects

- Some objects are remote, some are local
 - Doesn't really matter to user of object
 - Objects might transparently migrate
- Design doesn't have to take object distribution into account
- Failure and performance issues don't belong in the design
- The interface doesn't change if an object is remote

6

Uniform view

- not appropriate for
 - wide area networks,
 - consumer electronics,
 - portable devices
- appropriate for some local area networks
 - but robust distributed applications plan for failure
 - even if local

7

Memory access

- Can we make the fact that an object is remote transparent?
- Perhaps for objects
 - What about int's ?
 - What about char *'s?
- If you can't directly access fields and create pointers to them,
 - not transparent

8

Partial failure

- Computers fail
- OS's crash
- Networks fail
- PDA's get turned off or taken out of the room
- Often no warning

9

Queue example

- Want to add x to remote queue q
 - q.enqueue(x)
- Operation could fail
- Want to reliably enqueue x

10

Queue example

- ```
while (true) {
 try {
 q.enqueue(x);
 break;
 }
 catch (RemoteException e) {}
}
```

11

## Partial failure

- Object was enqueued, but failure occurred during return message
- Could enqueue x multiple times
- How to fix?
  - Need a request tag so that duplication enqueue requests can be detected

12

## Concurrency

- Distributed computations mandates concurrency

13

## Latency

- Remote calls are much more expensive than local calls

14

## Latency

- Making a call to an object on a remote machine is expensive

15