

Outline

- Review for mid term
- Midterm will cover everything covered so far
- **Open book, open notes, individual effort**
- **All forms of electronic communications strictly prohibited**
- Answer briefly and succinctly. Rampant use of buzzwords will receive no credit.



Sample Question

A local TV station wants to provide accurate weather and traffic reports. They install custom built electronic devices in listener's cars. These devices measure the present location (latitude and longitude using GPS), temperature, direction of movement and velocity every so often. These devices have limited storage space to store past measurements. These devices possess a radio device that can communicate with other radio units at a bandwidth of 11 Mbps. The nominal range of these radio devices are restricted to 150 feet. These devices can display traffic conditions to the user



Sample Question (cont)

You are hired as the CTO for this station. How would you utilize these devices to design a system that can display the local traffic conditions?

Answer: Each device stores the measurements for the past n time units. When a device comes in radio contact with another device, it performs a pair-wise anti-entropy to reconcile the observed traffic conditions. After a successful anti-entropy, both devices will know the traffic conditions along the routes taken by the two devices. “Eventually”, the devices will know about the local traffic conditions



Sample question

2. Storage space is of premium in these devices.
Discuss a scheme that will efficiently utilize the available storage space to display the prevailing local traffic conditions accurately

Answer: We utilize the applications notion of “local” to truncate the measurements logs.



Ubiquitous Computing Vision

- *The Computer for the Twenty-First Century*, Mark Weiser, Scientific American, Sep 1991
 - Virtual reality vs embodied reality
- *The Coming Age Of Calm Technology*, Mark Weiser and John Seely Brown, Oct 1996
 - Calm technology that moves from periphery into the center of attention and back to the periphery



Ubiquitous Computing Vision

- *People, Places, Things: Web Presence for the Real World. Cooltown Project at HP*
 - Bridge the electronic and physical world using the Web
 - Places, people, things
 - Discover URLs
 - Place specific identifying attributes



Ubiquitous Computing Vision

- *Next Century Challenges: Data-centric networking for invisible computing. The Portolano Project at the University of Washington* Mike Esler, Jeffrey Hightower, Tom Anderson and Gaetano Borriello In Mobicom '99
 - User interfaces: multiple interface, invisible interface
 - Distributed Services: Agent based approaches, service deployment
 - Resource discovery?
- Data should marshal, authenticate, adapt and pay for services as it proceeds
- Intermittent connectivity, power consumption



Ubiquitous Computing Vision

- *Pervasive Computing: Vision and Challenges*, M. Satyanarayanan, IEEE Personal Communications, August 2001



Systems design philosophy

- End-to-End Arguments in System Design – J. H. Saltzer, D. P. Reed and D. D. Clark – MIT (1980)
 - KISS principle (Keep It Simple, Stupid)
 - Modular vs end-to-end



Distributed Systems Architectures

- Centralized
 - Web
- Hierarchical
 - DNS
- Peer-to-Peer
 - Napster, gnutella



Distributed Systems Architectures

- Oceanstore: An architecture for Global-Scale Persistent Storage – University of California, Berkeley. ASPLOS 2000
 - Nomadic data access
 - Promiscuous caching
 - Updates – application level conflict resolution (similar to Bayou)
 - Untrusted infrastructure
 - Clients can be trusted, servers are not
 - Self certifying keys – secure hash



Distributed Systems Architecture

- Feasibility of a Serverless Distributed File System deployed on an Existing set of Desktop PCs – Microsoft research. ACM SIGMETRICS 2000
 - Not fully trusted
 - Disks are not that free
 - Compress data in storage
 - Files in directory are replicated together



Naming and Location Management

- **The Anatomy of a Context-Aware Application** – Andy Harter, Andy Hopper, Pete Steggles, Andy Ward and Paul Webster. AT&T Labs, Cambridge, UK
 - Users application should be available where-ever the user goes, in a suitably adapted form
 - Bats for location
 - Context aware application is one which adapts its behavior to a changing environment
 - E.g. Follow-Me applications
 - Context aware applications need to know the location of users and equipment, and the capabilities of the equipment and networking infrastructure
 - Modeling the environment
 - Containment relationships



Naming and Location Management

- *Active Names: Flexible Location and Transport of Wide-Area Resources*. Amin Vahdat, Michael Dahlin, Thomas Anderson, and Amit Aggarwal. In Proceedings of the Second USENIX Symposium on Internet Technologies and Systems, October 1999
 - Naming intent
 - Server selection, client customization, server customization
 - Resolvers to deal with active names



Replication Services

- *The Dangers of Replication and a Solution, Jim Gray, Pat Helland, Patrick O'Neil, and Dennis Shasha. In Proceedings of the ACM SIGMOD international conference on Management of data, 1996*
 - Availability and scalability: Provide high availability and scalability through replication
 - Mobility: Allow mobile nodes to read and update the database while disconnected from the network
 - Serializability: Provide single-copy serializable transaction execution
 - Convergence: Provide convergence to avoid system delusion
 - Group, master, two-tier vs lazy, eager



Synchronization and consistency

- Replication in the Harp File System, Barbara Liskov, Sanjay Ghemawat, Robert Gruber, Paul Johnson, Liuba Shrira, Michael Williams, MIT
 - Provides highly available, reliable storage for files
 - Guarantees atomic file operations in spite of concurrency and failure
 - Primary copy replication (Eager master)
 - Master server authoritative
 - Replicas – backup servers
 - Updates are sent to “enough” replicas to guarantee fail-safe behavior
 - Log structured updates – CP, AP, LB



Synchronization and consistency

- *The Case for Non-transparent Replication: Examples from Bayou* Douglas B. Terry, Karin Petersen, Mike J. Spreitzer, and Marvin M. Theimer. IEEE Data Engineering, December 1998
 - Transparent replication system:
 - Allow systems that were developed assuming a central file system or database to run unchanged on top of a strongly-consistent replicated storage system (e.g. Harp)
 - Non-transparent replication system:
 - Relaxed consistency model – access-update-anywhere
 - Applications involved in conflict detection and resolution. Hence applications need to be modified (e.g. Bayou, Coda file system etc)



Synchronization and consistency

- *Epidemic Algorithms for replicated database maintenance* [Alan Demers](#), Dan Greene, [Carl Hauser](#), Wes Irish, John Larson, Scott Shenker, Howard Sturgis, Dan Swinehart, and [Doug Terry](#). In Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing
 - Randomized algorithms for maintaining consistency for updates to replicas
 - Direct mail
 - Anti-entropy (push, pull, push-pull)
 - Rumor-mongering
 - Deletion and death certificates



Cont.

- Blind – $1/k$ probability of losing interest regardless if recipient is susceptible
- Feedback – $1/k$ probability only if recipient is infective
- Counter – lose interest after k unnecessary contacts
- Coin – k cycles regardless if susceptible
- Push and Pull
- Minimization – counters on both ends
- Connection limit – limits the number of connections
- Hunting – if a connection is rejected, choosing site can hunt for alternate sites



Synchronization and consistency

- Managing Update Conflicts in Bayou, a Weakly Connected Replicated Storage System [Douglas B. Terry](#), [Marvin M. Theimer](#), Karin Petersen, [Alan J. Demers](#), Mike J. Spreitzer and [Carl H. Hauser](#). In ACM Symposium on Operating Systems Principles (SOSP '95)
 - Primary commit – pair-wise anti-entropy
 - Session guarantees



Synchronization and consistency

- *Time, clocks and the ordering of events in a Distributed System* Leslie Lamport
 - Happens before
 - Partial ordering
 - Total ordering
 - Physical clocks



Replication and consistency

- *Exploiting Weak Connectivity for Mobile File Access*
 - Lily B. Mummert, Maria R. Ebling, M. Satyanarayan. In ACM Symposium on Operating Systems Principles (SOSP '95)
 - Successor to AFS
 - Half way between harp and bayou
 - Replication mostly transparent to end user
 - Hoard profiles to specify objects for the road

