

Open Source Project

DataCentric Networking

Éireann Leverett

Cambridge University

March 9, 2011




**UNIVERSITY OF
CAMBRIDGE**

Hadoop:ZooKeeper

- Shared hierarchical namespace
- znodes
- Ephemeral nodes
- Watches
- File and directory simultaneously
- Replicated
- Atomic updates
- Ordered

- Good for reads/Bad for writes
- Keeping track of versions
- Keeping track of devices
- The ordered quality allows higher level abstraction

The ordering of transactions and version control on nodes allows us to think about synchronisation primitives. We are guaranteed sequential consistency, and we know that the client's view of data will be the same (within a certain given window of time). We can also imagine implementing ownership, group membership, and change notification using watches.



Remember, this names space acts like a file and a directory. So `/app1` can contain data, but also has children: `/app1/p_1`, `/app1/p_2`, `/app1/p_2`.

- Imagine your task was to create a company to keep every iPhone in the world running the latest updates. Not just the updates for the firmware, but all the updates of all the apps they had installed.
- To do this you would need fast read access to data about every phone. You would need to poll it to find out what versions it was currently running, and then store that data.
- You would need to access it quickly and find all devices of the same type to update. You also might optimise your bandwidth usage to download an update once, and then multicast it to all phones.
- So if you wanted to find out which patches would help the most phones simultaneously (in other words, what's the most popular app), you might also use zookeeper. Now this is an example application, but it approximates a real world problem I would like to examine.

QUESTIONS ANSWERED

SIMPLE	50¢
GUESSES	\$1.00
INTELLIGENT	\$2.00
HONEST	\$5.00

**DUMB LOOKS
ARE STILL FREE**

It's been nice learning with you!