CMSC 327 Distributed Systems

Project 5: Replication
Due November 29, 2010

Unlike the previous projects that were focused on implementation and evaluation, this project is more concerned with design. Specifically, we will examine the role of replication in a few commonly used software tools. Replication is one of the core principles used to build distributed systems. Replication is used for a variety of reasons including fault tolerance, supporting disconnected operation, and supporting concurrent access. In this project we will consider tools that approach replication in slightly different ways:

- **rsync** – http://samba.anu.edu.au/rsync/

  "rsync is a file transfer program for Unix systems. rsync uses the "rsync algorithm" which provides a very fast method for bringing remote files into sync."

- **unison** – http://www.cis.upenn.edu/~bcpierce/unison/

  "Unison is a file-synchronization tool for Unix and Windows. It allows two replicas of a collection of files and directories to be stored on different hosts (or different disks on the same host), modified separately, and then brought up to date by propagating the changes in each replica to the other."

- **mercurial** – http://mercurial.selenic.com/

  "Mercurial is a free, distributed source control management tool."

Deliverable

Explore the three replication tools by reading the documentation and experimenting with the software. Reflect on the tools similarities and differences. Be sure to comment on the following aspects:

- What problem is the tool addressing? What do people use it for?
- How is the tool used for (or use) replication? Exactly what is replicated?
- What type of concurrent operations does the tool support?
- Does the tool use a pessimistic or optimistic consistency strategy? Why?
- If the tool supports optimistic replication, how does it handle reconciliation?
- Does the tool use a push or pull based approach? Why?
- What are the tool's critical design choices?

Submit a PDF file of your report via moodle: `cmsc327_proj5_LASTNAME_FIRSTNAME.pdf`

Learning Objectives

- Explore the Uses of Replication
- Critically Reflect on the Design of Software Tools