













Anti-entropy works as follows: periodically, a server P randomly chooses a server Q and then the two exchange databases. Three modes of anti-entropy are possible. In *push* mode, P pushes its database to Q, which adds to its database everything in P that it doesn't have. In the *pull* mode, the reverse is done. *Push-pull* mode works in both directions.

This discussion comes from "Epidemic Algorithms for Replicated Database Maintenance," by A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry, Proceedings of the ACM SIGACT-SIGOPS 6<sup>th</sup> Annual Symposium on Principles of Distributed Computing, August, 1987.





We look at the fundamentals of epidemic theory in order to analyze the basic performance of the anti-entropy algorithms. Assume that any server can *infect* any other server with equal probability. Also assume that k members are already infected and infections occur in rounds: at each round, each server randomly picks another server.





If we compare the performance of the pull and push approaches; we observe that (for a very large number of servers) the pull approach converges (i.e., infects the entire population) faster than push does. Push-pull performs similar to pull.





A form of non-simple epidemic is called *rumor mongering*.



The math is taken from epidemic theory. The final equation tells us how many are still susceptible when there are no more infectious nodes. Thus these will never get the updates.





Recall that we have no time stamps, so it cannot be determined if an update is new or old.





There are many scenarios in which conventional version vector-based techniques fail to detect the conflicts correctly. In some applications, concurrent writes to the same object may not conflict as expected. Furthermore, concurrent writes to different objects may conflict. Bayou handles such cases by having the application define its own notion of conflict.



Papers on Bayou include:

"Session Guarantees for Weakly Consistent Replicated Data," Proceedings of the Third International Conference on Parallel and Distributed Information Systems, 1994.

"Managing Update Conflicts in Bayou, a Weakly Connected Replicated Storage System," Proceedings of the fifteenth ACM symposium on Operating systems principles, 1995 (http://delivery.acm.org/10.1145/230000/224070/p172-terry.pdf?key1=224070&key2=3718232721&coll=GUIDE&dl=GUIDE&CFI).

"Flexible Update Propagation for Weakly Consistent Replication," Proceedings of the sixteenth ACM symposium on Operating systems principles, 1997 (http://delivery.acm.org/ 1 0 . 1 1 4 5 / 2 7 0 0 0 0 / 2 6 6 7 1 1 / p 2 8 8 - p e t e r s e n . p d f ? key1=266711&key2=2038232721&coll=GUIDE&dl=GUIDE&CFID=86082163&CFTOKEN=970 29400).



Conflict detection is accomplished by using *dependency checks* — each write operation also supplies a dependency check that specifies a simple condition defined over the state of the database. The dependency check also defines what the application expects to see as a result when the condition gets evaluated. If indeed the expected result is observed then Bayou decides that there is no conflict, else it decides that there is a conflict.

If there is a conflict, then a *merge procedure* (i.e., a *conflict resolution procedure*), also supplied by the application, is executed to resolve the conflict. The merge procedure produces a new alternate update that is acceptable by the application and that would not create a conflict with the current database state (i.e., an update for which the dependency check will produce the expected result).



The slide shows the basic format of a Bayou write.



The slide shows an example Bayou write, including the dependency check and the merge procedure.





Note that "process" here means a client who is connecting to multiple servers.









We sketch an approach to an implementation of the constraints. Our first few ideas aren't entirely practical, but will lead to an approach that is.



If a server can't be found that satisfies the restriction of the second bullet, then the application must be notified that read-your-writes cannot be maintained.















