

Flexible Consistency for Wide area Peer Replication

Sai Susarla and John Carter
{*sai, retrac*}@*cs.utah.edu*

UUCS-04-016

School of Computing
University of Utah
Salt Lake City, UT 84112 USA

November 18, 2004

Abstract

The lack of a flexible consistency management solution hinders P2P implementation of applications involving updates, such as read-write file sharing, directory services, online auctions and wide area collaboration. Managing mutable shared data in a P2P setting requires a consistency solution that can operate efficiently over variable-quality failure-prone networks, support pervasive replication for scaling, and give peers autonomy to tune consistency to their sharing needs and resource constraints. Existing solutions lack one or more of these features.

In this paper, we describe a new consistency model for P2P sharing of mutable data called *composable consistency*, and outline its implementation in a wide area middleware file service called *Swarm*¹. Composable consistency lets applications compose consistency semantics appropriate for their sharing needs by combining a small set of primitive options. *Swarm* implements these options efficiently to support scalable, pervasive, failure-resilient, wide-area replication behind a simple yet flexible interface. We present two applications to demonstrate the expressive power and effectiveness of composable consistency: a wide area file system that outperforms Coda in providing close-to-open consistency over WANs, and a replicated BerkeleyDB database that reaps order-of-magnitude performance gains by relaxing consistency for queries and updates.

¹*Swarm* stands for Scalable Wide Area Replication Middleware.