Rethinking Consistency Management in Real-time Collaborative Editing Systems

Ph.D. Defense

Computer Science Department Conference Room (1432)
14th Floor, 34 Peachtree Street

Friday, May 18, 2007 at 10:30 AM

Jon A Preston
Department of Computer Science

Networked computer systems offer much to support collaborative editing of shared documents among users. Software Engineering is one of many fields that benefits from computer-assisted collaboration as a myriad of developers, project managers, testers, and designers work together to develop large, complex systems that consist of a multitude of process and product artifacts. Multi-discipline and geographically-distributed production and research teams collaborate and co-author documents for businesses and universities worldwide. Such collaborations occur asynchronously via access to shared document repositories often assisted by configuration management systems and occur synchronously via shared, real-time collaborative editing systems often assisted by awareness-enhancing technology.

Increasing concurrent access to shared documents by allowing multiple users to contribute to and/or track changes to these shared documents is the goal of real-time collaborative editing systems; yet concurrent access is either limited in existing systems that employ exclusive locking or concurrency control algorithms such as operational transformation (OT) may be employed to enable concurrent access. Unfortunately, such OT based schemes are costly with respect to communication and computation. Further, existing systems are often specialized in their functionality and require users to adopt new, unfamiliar software to enable collaboration.

We have developed a set of deadlock-free multi-granular dynamic locking algorithms and data structures that maximize concurrent access to shared documents while minimizing communication cost. These algorithms provide a high level of service for concurrent access to the shared document and may integrate merge-based or operational transformation (OT)-based consistency maintenance policies more locally among a subset of the users within a subsection of the document – thus reducing the communication costs in maintaining consistency.

Additionally, we have developed client-server and P2P implementations of our hierarchical document management algorithms. Simulations results indicate that our approach achieves significant communication and computation cost savings. We have also developed a hierarchical reduction algorithm that can minimize the space required of RTCES, and this algorithm may be pipelined through our document tree.

Further, we have developed an architecture that allows for a heterogeneous set of client editing software to connect with a heterogeneous set of server document repositories via Web services. This architecture supports our algorithms and does not require client or server technologies to be modified – thus it is able to accommodate existing, favored editing and repository tools.

Finally, we have developed a prototype benchmark system of our architecture that is responsive to users' actions and minimizes communication costs.

Dissertation Committee

Dr. Sushil K. Prasad, Chair

Dr. Xiaolin Hu, Member

Dr. Melody Moore Jackson, External Member

Dr. Rajshekhar Sunderraman, Member