A LOCKING PROTOCOL FOR RESOURCE COORDINATION IN DISTRIBUTED DATABASES*‡

Daniel A. Menasce†
Gerald J. Popek
Richard R. Muntz

Computer Science Department
University of California
Los Angeles, California 90024

EXTENDED ABSTRACT

A locking protocol to coordinate access to a distributed database and to maintain system consistency throughout normal and abnormal conditions is presented in this paper. The proposed protocol is robust in the face of crashes of any participating site, as well as communication failures. Recovery from any number of failures during normal operation or any of the recovery stages is supported. Integration of virtually any locking discipline including predicate lock methods is permitted by this protocol. The locking algorithm operates, and operates correctly, when the network is partitioned, either intentionally or by failure of communication lines. Each partition is able to continue with work local to it, and operation merges gracefully when the partitions are reconnected.

It is shown in this paper that for many topologies of interest, the delay introduced by the protocol is not a direct function of the size of the network. The communication cost is shown to grow in a relatively slow, linear fashion with the number of sites in the network. An informal proof of the correctness of the algorithm is also presented in this paper.

The algorithm has as its core a centralized locking protocol with distributed recovery procedures. A centralized controller with local appendages at each site coordinates all resource control, with requests initiated by application programs at any site. Recovery is broken down into three disjoint mechanisms; for single node recovery, merge of partitions and reconstruction of the centralized controller and tables.

The paper concludes with a proposal for an extension aimed at optimizing operation of the algorithm to adapt to highly skewed distributions of activity. The extension applies nicely to interconnected computer networks.

KEYWORDS AND PHRASES: concurrency, crash recovery, distributed databases, locking protocol, consistency.

---

‡ In accordance with ACM publication policy, only an abstract of this paper appears in the Conference Proceedings, since the full paper will appear in the ACM Transactions on Database Systems. Pre-prints of the paper were distributed to Conference participants.

* This research was supported by the Advanced Research Projects Agency of the Department of Defense under Contract MDA 903-77-C-0211.

† Partially supported by the Conselho Nacional de Desenvolvimento Cientifico e Tecnologico, CNPq, Brazil.