L.R.I.

OPTIMIZATION OF THE NUMBER OF COPIES IN A DISTRIBUTED DATA BASE

E.G. COFFMAN JR., E. GELENBE, B. PLATEAU

Avril 1979 *Rapport de Recherche n° 34*

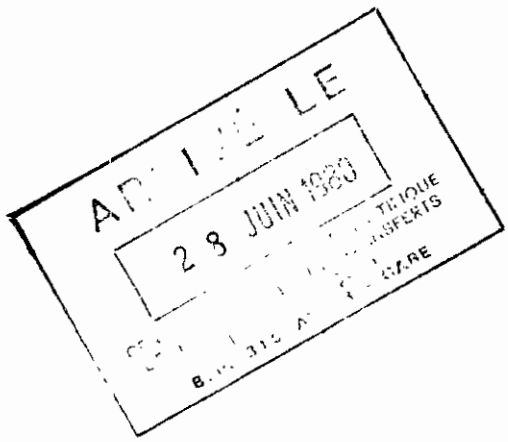

UNIVERSITE DE PARIS-SUD
Centre d'Orsay
LABORATOIRE DE RECHERCHE EN INFORMATIQUE
Bât. 490
91405 ORSAY (FRANCE)

OPTIMIZATION OF THE NUMBER OF COPIES IN A DISTRIBUTED DATA BASE

E.G. Coffman Jr.[+], E. Gelenbe [++], B. Plateau [++]

Abstract We consider the effect on system performance of the distribution of a data base in the form of multiple copies at distinct sites. The purpose of our analysis is to determine the gain in READ throughput that can be obtained in the presence of consistency preserving algorithms that have to be implemented when UPDATE operations are carried out on each copy. We show that READ throughput diminishes if the number of copies exceeds an optimal value. The theoretical model we develop is applied to a system in which consistency is preserved through the use of Ellis's ring algorithm.

+) Bell Laboratories ,Murray Hill,New Jersey (U.S.A).
++) Laboratoire de Recherche en Informatique, Université Paris-Sud Bât 490, 91405 Orsay, France.

INTRODUCTION

The widespread usage of large scale computer communication networks has made it possible to create loosely coupled, but integrated information processing and retrieval systems using multiple computers. This technological possibility has been further enhanced more recently by the wide-spread introduction of very small (personnal) computers, microprocessors etc. When used with local area networks or other local communication media, it becomes relatively easy (at least in principle) to construct loosely coupled multiple processor systems.

It therefore becomes of great interest to understand how these multiple processor systems (which we distinguish here from the traditional tightly coupled multiprocessors) can be used efficiently to handle the processing power necessary for large scale applications (such as numerical analysis, data bases, language processing, etc.).

There are several factors which render such multiple processor systems attractive. Among them we may mention the possibility of constructing variable performance systems using their natural extensibility obtained by adding more processors and their software to an existing system, enhanced reliability, system modularity, and the possibility of constructing clear hardware-imposed interfaces using the communication network.

However it is clear that these gains are not obtained without a substantial counterpart in the research effort necessary to understand the logical design of such systems. Furthermore, the performance aspects (throughput, response time and reliability enhancement) of multiple processor computers are not yet well understood. This paper is a first step in this direction.

It is difficult to fully understand the performance aspects of multiple processor architectures if one does not examine the problem in the context of a special class of applications. General studies have been carried out, but their use is limited to providing broad outlines of the major trade-offs involved. That is why in this paper we examine this problem in the context of an application area which is of importance : distributed data bases.

The creation of a distributed data base can be motivated by two types of reasons : a large decentralized enterprise can decide to unify its information systems on different machines by using computer networks as a common communication facility, or one can profit from the increased throughput and reliability of multiple processors in order to have a data base system of higher performance. In both cases, the increased potential throughput is reduced by the necessity to preserve consistency of the distributed information and this is one of the main sources of performance reduction.

In recent years many different proposals for consistency preserving algorithms have been published for multiple copy distributed data bases. In [1] several of these algorithms are briefly reviewed and some decentralized control schemes are proposed. The algoritnm proposed in [2] is typical of a class of asynchronous consistency control algorithms. A detailed survey and critic of most of the published algorithms can be found in [4].

Several reasons make multiple copy distributed data bases attractive : the reduction of communication costs for READ operations, which can proceed on a local copy, enhanced reliability due to the existence of a alternate site in case of failure, and improved READ throughput capacity obtained by sharing the READ load among several centers. Some of these advantages are offset by the necessity of implementing an elaborate procedure for preserving consistency.

Indeed, the consistency preserving mechanism that will be activated when updates take place, either at each update or periodically, will introduce an additional communication cost. Furthermore, the recovery procedures after failures will be more complicated because of the need to preserve consistency. Finally, the consistency control algorithm will lengthen the effective time necessary to carry out an update because of the coupling it introduces between different copies and will therefore reduce the processing capacity available for READ's.

In this paper we address this last issue. We consider first a system in which n centers contain fully or partially replicated copies of a data base. Under a general formulation, we obtain a general expression for effective available READ processing capacity in the presence of a stream of UPDATE's. This allows us to state an optimization problem : the determination of the optimum number of centers that maximize the READ processing capacity. An explicit expression for this quantity is obtained in the case of a fully duplicated system.