$AR\tilde{A}$: *a cooperatif system for locating mirrors over the Web*

2 janvier 2001

Over the last few years, large scale information sharing has risen to dominance throughout Internet with the advent of cooperative applications. The exponentially increasing number of users scattered all over the world highlights performance and information availability problems, due to access latency, high server load and network congestion. Information replication in the vicinity of end users is a commonly used technique for improving availability and boosting performance.

For replication to be effective, users should first be able to transparently locate copies that better suit their needs, i.e nearby for better access time or the most recent version for consistency reasons. Firstly, it is fundamental to have a system capable of dynamically discovering and advertising freshly created copies. Secondly, a naming system supporting replication is needed to ensure transparency. In this thesis, we are interested in persistent replication of mirrors, and more particularly in discovering, advertising and transparently locating replicas over the Web.

Our solution, $AR\tilde{A}$, uses cooperation by exchanging mirrors location messages between users. Cooperation is limited to groups of users having common interests and who are geographically close. Transparent and local redirection of users requests, using URLs prefixes and handling aliases for naming replicas guarantees access transparency. Finally, the system scales geographically due to the epidemic, asynchronous and lazy nature of the cooperation protocol, coupled with the principle of locality.

We propose an evaluation methodology we used to show the need for mirroring and to validate to proposed system. Performance measurements show encourouging results in favour of using $AR\tilde{A}$ for transparently locating mirrors in large scale information systems.