Page Replication — Variations on a Theme

Rudolf Fleischer* Steve Seiden*

Max-Planck-Institut für Informatik, Im Stadtwald, 66123 Saarbrücken.
E-mail: rudolf@mpi-sb.mpg.de, sseiden@acm.org

In the page replication problem, we are given a weighted graph and a start node s which initially contains a page. Other nodes which want to access this page can do so by sending a request to any other node holding the page. The cost of this access is then the distance between the two nodes. A node can also choose to copy the whole page at a cost of d times the distance, where d is usually a large constant (the page size). If the sequence of requests is not known in advance, we have an online problem so we use competitive analysis to measure the performance of algorithms, i.e., we compare the cost of our algorithm to the cost of the best possible offline algorithm. Various graph topologies (trees, rings, arbitrary graphs) have been studied in this model.

In this talk, we present several variants of the classical page replication problem and give some new upper and lower bounds. In the *continuous page replication problem* we allow requests at and replication to arbitrary points on any edge of the graph. In this model randomization does not help, and the deterministic algorithms are equivalent to randomized algorithms in the classical discrete model. In the *unequal cost model* we assume that the online algorithm has a replication factor different from d, the replication factor of the offline algorithm. We give optimal deterministic and randomized algorithms for the discrete and continuous variant of this model on trees. ¿From this we can derive much simpler proofs for known algorithms on rings.

*Both authors were partially supported by the EU ESPRIT LTR Project No. 20244 (ALCOM-IT), WP 3.2. The first author was also supported by a Habilitation Scholarship of the German Research Foundation (DFG).