

Strategies for Exploring Large Scale Data

Joseph F. JaJa

*Institute for Advanced Computer Studies and
Department of Electrical and Computer Engineering
University of Maryland
joseph@umiacs.umd.edu*

Abstract

We consider the problem of querying large scale multidimensional time series data to discover events of interest, test and validate hypotheses, or to associate temporal patterns with specific events. This type of data currently dominates most other types of available data, and will very likely become even more prevalent in the future given the current trends in collecting time series of business, scientific, demographic, and simulation data. The ability to explore such collections interactively, even at a coarse level, will be critical in discovering the information and knowledge embedded in such collections. We develop indexing techniques and search algorithms to efficiently handle temporal range value querying of multidimensional time series data. Our indexing uses linear space data structures that enable the handling of queries in I/O time that is essentially the same as that of handling a single time slice, assuming the availability of a logarithmic number of processors as a function of the temporal window. A data structure with provably almost optimal asymptotic bounds is also presented for the case when the number of multidimensional objects is relatively small. These techniques improve significantly over standard techniques for either serial or parallel processing, and are evaluated by extensive experimental results that confirm their superior performance.