SAND2011-3222C

Facilitating Efficient Parallelization of Information Storage and Retrieval on Large Data Sets

Damian Dechev
Scalable Computing R&D Department
Sandia National Laboratories
ddechev@sandia.gov

Abstract

The purpose of this work is to develop a lock-free hash table that allows a large number of threads to concurrently insert, modify, or retrieve information. Lock-free or non-blocking designs alleviate the problems traditionally associated with lock-based designs, such as bottlenecks and thread safety. Using standard atomic operations provided by the hardware, the design is portable and therefore, applicable to embedded systems and supercomputers such as the Cray XMT. Real-world applications range from search-indexing to computer vision. Having written and tested the core functionality of the hash table, we plan to perform a formal validation using model checkers.

Categories & Subject Descriptors: E.2 [Data]: Data Storage Representations---Hash Table Representations

General Terms: Algorithms, Design, Languages.

Bio

Steven and Pierre are undergraduate students at the University of Central Florida, majoring in Computer Science. They are researching under Dr. Dechev, studying lock-free, concurrent data structures. Steven's technical interests are in the efficient design of concurrent server applications. Pierre's interests include devising modern programming techniques.

ICS'11, May 31–June 4, 2011, Tuscon, Arizona, USA.
ACM 978-1-4503-0102-2/11/05.