CCIT Report #790 May 2011

CAFÉ: Scalable Task Pools with Adjustable Fairness and Contention

Dmitry Basin* Rui Fan† Idit Keidar* Ofer Kiselov* Dmitri Perelman*

Abstract

Task pools have many important applications in distributed and parallel computing. Pools are typically implemented using concurrent queues, which limits their scalability. We introduce *CAFÉ, Contention and Fairness Explorer*, a scalable and wait-free task pool which allows users to control the trade-off between fairness and contention. The main idea behind CAFÉ is to maintain a list of *TreeContainers*, a novel tree-based data structure providing efficient task inserts and retrievals. TreeContainers don't guarantee FIFO ordering on task retrievals. But by varying the size of the trees, CAFÉ can provide any type of pool, from ones using large trees with low contention but less fairness, to ones using small trees with higher contention but also greater fairness.

We demonstrate the scalability of TreeContainer by proving an $O(\log^2 N)$ bound on the step complexity of insert operations when there are N inserts, as compared to an average of $\Omega(N)$ steps in a queue based implementation. We further prove that get operations are wait-free. Evaluations of CAFÉ show that it outperforms the Java SDK implementation of the Michael-Scott queue by a factor of 30, and is over three times faster than other state-of-the-art non-FIFO task pools.

*Department of Electrical Engineering, Technion, Haifa, Israel {sdimbsn@tx,idish@ee,sopeng@t2,dima39@tx}.technion.ac.il
†School of Computer Engineering, Nanyang Technological University, fanrui@ntu.edu.sg