Accepted Manuscript

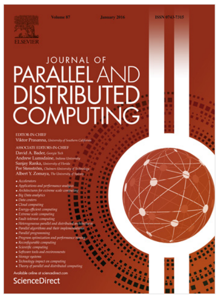

A contention adapting approach to concurrent ordered sets

Konstantinos Sagonas, Kjell Winblad

PII: S0743-7315(17)30305-2
DOI: https://doi.org/10.1016/j.jpdc.2017.11.007
Reference: YJPDC 3779

To appear in: *J. Parallel Distrib. Comput.*

Received date : 10 August 2016
Accepted date : 9 November 2017

Please cite this article as: K. Sagonas, K. Winblad, A contention adapting approach to concurrent ordered sets, *J. Parallel Distrib. Comput.* (2017), https://doi.org/10.1016/j.jpdc.2017.11.007

A Contention Adapting Approach to Concurrent Ordered Sets☆

Konstantinos Sagonas[a,b], Kjell Winblad[a,*]

[a]*Department of Information Technology, Uppsala University, Sweden*
[b]*School of Electrical and Computer Engineering, National Technical University of Athens, Greece*

Abstract

With multicores being ubiquitous, concurrent data structures are increasingly important. This article proposes a novel approach to concurrent data structure design where the data structure dynamically adapts its synchronization granularity based on the detected contention and the amount of data that operations are accessing. This approach not only has the potential to reduce overheads associated with synchronization in uncontended scenarios, but can also be beneficial when the amount of data that operations are accessing atomically is unknown.

Using this adaptive approach we create a contention adapting search tree (CA tree) that can be used to implement concurrent ordered sets and maps with support for range queries and bulk operations. We provide detailed proof sketches for the linearizability as well as deadlock and livelock freedom of CA tree operations. We experimentally compare CA trees to state-of-the-art concurrent data structures and show that CA trees beat the best of the data structures that we compare against by over 50% in scenarios that contain basic set operations and range queries, outperform them by more than 1200% in scenarios that also contain range updates, and offer performance and scalability that is better than many of them on workloads that only contain basic set operations.

Keywords: concurrent data structures, ordered sets, linearizability, range queries

1. Introduction

With multicores being widespread, the need for efficient concurrent data structures has increased. This need has lead to an intensification of research in this area. For example, a large number of concurrent data structures for ordered sets have recently been proposed. To enable parallel operations in the data structure, some of them use fine-grained locking [1, 2, 3, 4] while others use lock-free techniques [5, 6, 7, 8, 9, 10, 11, 12].

☆Research supported in part by the European Union grant IST-2011-287510 "RELEASE: A High-Level Paradigm for Reliable Large-scale Server Software" and the Linnaeus centre of excellence UPMARC (Uppsala Programming for Multicore Architectures Research Center).

*Corresponding author

Email addresses: kostis@it.uu.se (Konstantinos Sagonas), kjell.winblad@it.uu.se (Kjell Winblad)