

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE
UNIVERSITY OF SOUTHERN DENMARK, ODENSE

COMPUTER SCIENCE COLLOQUIUM

An Optimal Implementation of Fetch-and-Increment

Faith Ellen

Department of Computer Science
University of Toronto, Canada

Tuesday, 20 May, 2014 at 14:15

Auditorium U81

Abstract:

A new wait-free implementation of a fetch-and-increment object shared by n processes from read-write registers and load-linked/store-conditional objects will be presented. The step complexity of each fetch-and-increment operation is $O(\log n)$, which is optimal. The implementation uses $O(m)$ objects, each of which stores $O(\log m)$ bits, where m is the number of fetch-and-increment operations that are performed. This can be extended to implementations of other objects, such as fetch-and-add, and gives a new algorithm with $O(\log n)$ step complexity for strong renaming of processes whose original names are from a universe of size n .

The implementation uses a new object, called a buffer. It supports an operation which, if successful, puts a value into its input area that can depend on the value that is currently there, an operation that copies the value in its input area to its output area, provided its output area is empty, and an operation that empties its output area. An implementation of a buffer from a small constant number of load-linked/store-conditional objects will be given in which all three operations have constant step complexity.

This is joint work with Philipp Woelfel and was presented at DISC 2013.

Host: Joan Boyar