CSDS 500 Fall 2022 Colloquium

11:30 AM to 12:30 PM
Thursday, November 10, 2022
In-person: White 411 & Virtual
(Zoom Meeting ID: 935 2774 6217, Passcode: 946606)

"Exploiting Parallelism with In-memory Transactions: Lower bounds and Algorithms"

Abstract: Current general-purpose CPUs are multicores, offering multiple computing units within a single chip. The performance of programs on these architectures, however, does not necessarily increase proportionally with the number of cores. Designing concurrent programs to exploit these multicores emphasizes the need for achieving efficient synchronization among threads of computation. When there are several threads that conflict on the same data, the threads will need to coordinate their actions for ensuring correct program behavior. Traditional techniques for synchronization are based on locking that provides threads with exclusive access to shared data However, traditional techniques for synchronization are either too coarse-grained to exploit concurrency or require application-specific fine-grained synchronization.

The Transactional Memory (TM) abstraction is a synchronization mechanism that serves as a universal construction for multicore programming: it intends to combine an easy-to-use programming model with an efficient utilization of hardware concurrency. TM allows the programmer to speculatively execute sequences of shared-memory operations as atomic in-memory transactions with safe semantics: state witnessed by each transaction is consistent with some sequential execution. Thus, the programmer can design applications having largely sequential semantics in mind and let TM take care, at run-time, of dealing with problems associated with process asynchrony and adversarial failures.

In this talk, we focus on a model for hybrid TMs that exploits hardware extensions in prevalent CPU architectures. We present lower and upper bounds for implementing safe hybrid transactions. We then address the recent commercial availability of non-volatile memory which has prompted many researchers to also consider designing concurrent data structures for multicores that persist shared state (i.e., allowing the data structure to be recovered following a power failure). These so called persistent concurrent data structures further complicate the process of achieving safe and efficient implementations. We present some challenges and approaches for implementing safe persistent universal constructions and persistent transactional memory.

Srivatsan Ravi
University of Southern California

Bio: Srivatsan Ravi is a faculty member at the Department of Computer Science and a research lead at the Information Sciences Institute in University of Southern California. His research interests are centered around the theory and practice of concurrent and distributed computing as well as application of distributed techniques for networking and privacy-preserving computations. He received his bachelors degree from Anna University (India), his masters degree from Cornell University (New York) and doctoral degree from Technical University of Berlin (Germany) where he was a Marie Curie Ph.D Fellow.
He is a Principal Investigator and Key Personnel in several active DARPA and NSF programs.

This is to certify that _____attended this seminar. Certified by _____.
Certificates of attendance and other evidence of CPD activity should be retained by the attendee for auditing purposes.

CASE SCHOOL OF ENGINEERING
CASE WESTERN RESERVE UNIVERSITY