



Dissertation Defense

Vaibhav Gogte

Runtime Systems for Persistent Memories

Wednesday, December 4, 2019

11:30 am – 1:30 pm

3725 Beyster Bldg.



ABSTRACT: Persistent memory (PM) technologies, such as Intel and Micron's 3D XPoint, are here. Cloud vendors have already started public offerings with support for Intel's Optane DC persistent memory. Unlike traditional block-based storage devices, such as hard disks and SSDs, PMs can be accessed using a byte-addressable load-store interface, avoiding the expensive software layers required to access storage. This enables programmers to design data structures in PM that are accessed like memory and yet are recoverable upon crash or failure. However, several challenges remain in existing hardware, programming, and software systems that inhibit wide-scale PM adoption. This thesis focuses on building efficient mechanisms that span hardware and operating systems, and programming languages for integrating PMs in future systems.

First, this thesis proposes a persistency model for high-level languages to enable integration of PMs in to future programming systems. This thesis argues for persistency semantics that guarantee failure atomicity of synchronization-free regions (SFRs) --- program regions delimited by synchronization operations. The proposed approach provides clear semantics for the PM state that recovery code may observe and extends C++11's "sequential consistency for data-race-free" guarantee to post-failure recovery code.

Second, this thesis proposes StrandWeaver, a hardware persistency model that minimally constrains ordering on PM operations. StrandWeaver manages PM order within a strand, a logically independent sequence of PM operations within a thread. PM operations that lie on separate strands are unordered and may drain concurrently to PM. StrandWeaver implements these primitives in hardware to allow programmers to improve concurrency and relax ordering constraints on updates as they drain to PM. We demonstrate that StrandWeaver improves performance by up to 1.97x over existing ISA-level ordering mechanisms.

Chair: Prof. Thomas F. Wenisch