TabulaROSA: Tabular Operating System Architecture

Jeremy Kepner[1–4], Ron Brightwell[5], Alan Edelman[2,3], Vijay Gadepally[1,2,4], Hayden Jananthan[1,4,6], Michael Jones[1,4], Sam Madden[2], Peter Michaleas[1,4], Hamed Okhravi[4], Kevin Pedretti[5], Albert Reuther[1,4], Thomas Sterling[7], Mike Stonebraker[2]
[1]MIT Lincoln Laboratory Supercomputing Center, [2]MIT Computer Science & AI Laboratory, [3]MIT Mathematics Department, [4]MIT Lincoln Laboratory Cyber Security Division, [5]Sandia National Laboratories Center for Computational Research, [6]Vanderbilt University Mathematics Department, [7]Indiana University Center for Research in Extreme Scale Technologies

Overview

- The rise in computing hardware choices is driving a reevaluation of operating systems
- An operating system can be viewed as software that brokers and tracks the resources of the compute engines and is akin to a database management system
- To explore the idea of using a database in an operating system role, this work defines key operating system functions in terms of rigorous mathematical semantics that are directly translatable into database operations

Compute/Data Intensive Operations

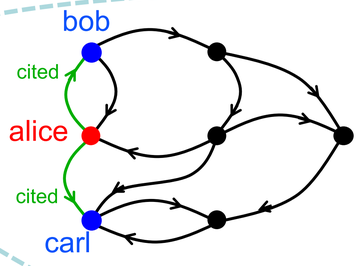

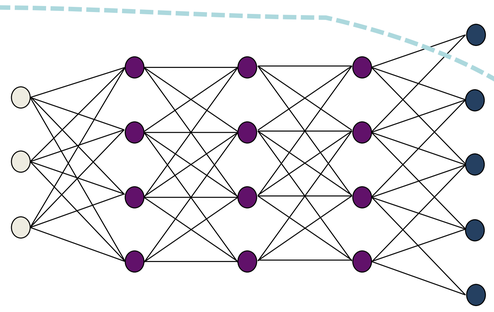

$$C = A \oplus B \qquad C = A \otimes C$$

$$C = A \oplus.\otimes B$$

Tabular Instruction Set

TabulaR Operating System Architecture

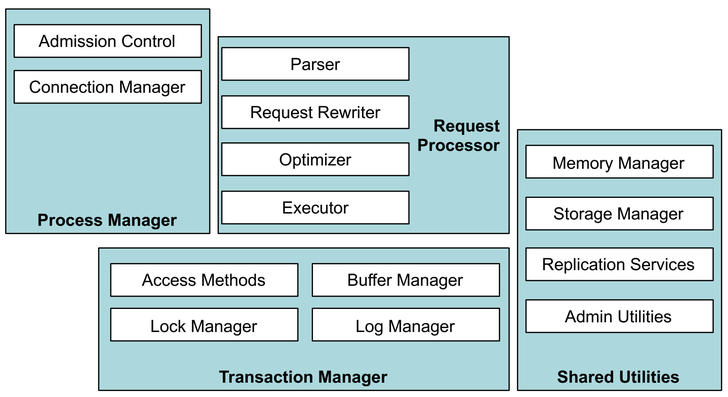

Dynamic Binding of Instructions, Data & Authorizations to Hardware

Hardware Isolation

Massively Parallel Heterogeneous Hardware

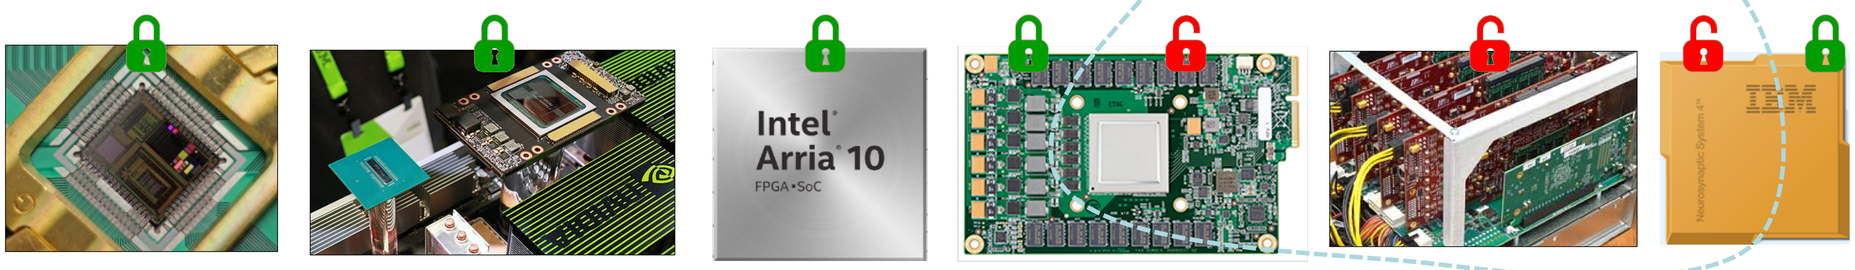

Simulation

- Operating system equations provide a mathematical specification for a Tabular Operating System Architecture (TabulaROSA) that can be implemented on any platform
- Simulations of forking in TabularROSA are performed using an associative array implementation using D4M[1] and compared to Linux on a 32,000+ core supercomputer[2]
- The TabulaROSA simulations show 20x higher performance as compared to Linux while managing 2000x more processes in fully searchable tables

Performance Results

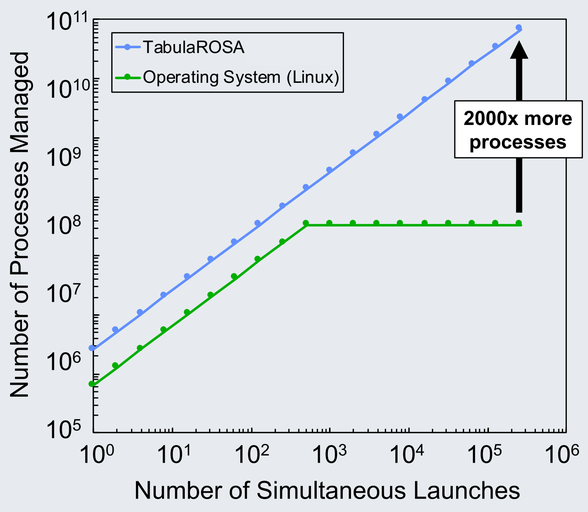

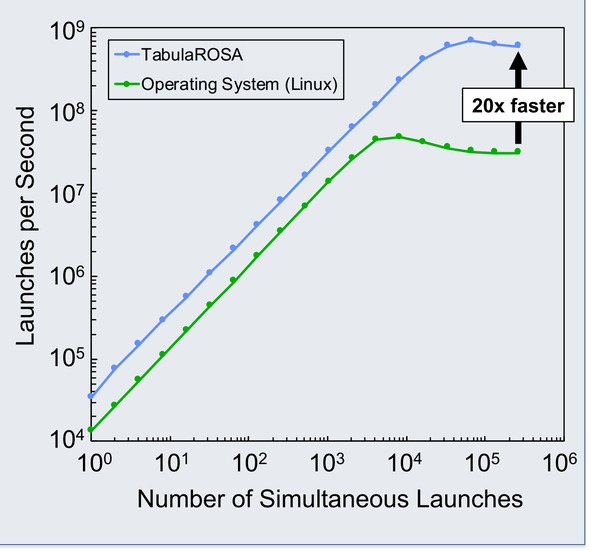

References:
[1] Dynamic distributed dimensional data model (D4M) database and computation system, Kepner et al, ICASSP 2012.
[2] *TabulaROSA: Tabular Operating System Architecture for Massively Parallel Heterogeneous Compute Engines*, Kepner et al, IEEE HPEC 2018

This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8702- 15-D-0001 and National Science Foundation grants DMS-1312831 and CCF- 1533644. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering or the National Science Foundation.