Taking Formal into areas unaccustomed to it

Ganesh Gopal
Utah University

Abstract

Every area of concurrent programming requires formal correctness frameworks within which its designers can safely conduct performance optimizations. Our interest is in contributing formal methods for high performance computing. The main set of challenges that HPC faces in this regard is the use of multiple concurrency models and APIs to support a mixture of legacy and new software, millions of threads running on heterogeneous cores, and an impoverished base of formal methods to start from. We show that by attempting to formalize even age-old APIs such as MPI, one learns about the formal model underlying message passing concurrency, and this can not only help build correctness tools for MPI programs, but understand similar issues sure to arise in future HPC programming proposals. We show that understanding MPI's `matches-before' order is crucial to building formal dynamic verification tools, and resource-dependent correctness outcomes. We show how unmodified MPI applications can be dynamically analyzed on 1000-CPU platforms. We also briefly cover another line of work that involves modeling and analyzing SIMD concurrency in GPU-based programs through SMT-based methods.

Biography

Ganesh Gopalakrishnan earned his PhD in 1986 from Stony Brook.
Over the last five years, his group has tailored Formal Methods to benefit High Performance Computing practitioners. Specific outcomes include: ISP a dynamic verification tool for MPI programs; GEM: Graphical Explorer for MPI programs which is an Eclipse/Parallel Tools Platform tool integration of ISP; DAMPI: a Distributed Analyzer for MPI programs; PUG: a Prover for User GPU kernels; and XUM: an eXtensible Utah Multicore, which is a message passing multicore written in VHDL/Verilog for FPGAs.

At Utah, the Center for Parallel Computing at Utah (CPU, http://www.parallel.utah.edu) helps his group and external collaborators to address technical challenges in building large-scale high-performance and reliable high performance computing systems.