

Talk Title: Automated Reasoning for Replicated Systems

Date: 26th July 2019

Time: 11:30 AM

Location: RM 101

Abstract

For applications with geo-distributed client base such as Twitter, Amazon, Wikipedia, etc., in order to provide an always-on service with uniform low latency, application data needs to be replicated in multiple locations across the world. Such replicated systems satisfy a number of desirable properties including scalability and fault tolerance, but writing correct applications for such systems is inherently hard due to unbounded concurrency and the lack of sequential consistency. To make such systems easier to program, a number of so-called weakly consistent replicated data stores have emerged in the last few years, providing a tradeoff between consistency and performance. It is still a monumental task for programmers to find the weakest consistency level under which their applications can run correctly. In this talk,

I will present my recent work which addresses this problem by proposing automated techniques to reason about the correctness of programs in replicated systems. Notably, the proposed techniques are parametric in the consistency model, and hence given an application and a specification of the consistency model, they can automatically find correctness bugs or verify that the application will run correctly. In particular, I will focus on a special class of programs called Convergent Replicated Data Types (CRDT), and present an automated approach to verify an important property called convergence for such programs on replicated systems.

Speaker Bio

Dr. Kartik Nagar is a postdoctoral research associate at Purdue University working with Prof. Suresh Jagannathan. His research interests are in Program Analysis and Automated Verification. He completed his PhD from Indian Institute of Science under the guidance of Prof. Y N Srikant, and his PhD thesis was on static timing analysis of programs for Real-time systems.
