

## ABSTRACT

Laguna, Ignacio Ph.D., Purdue University, December 2012. Probabilistic Error Detection and Diagnosis in Large-Scale Distributed Applications. Major Professor: Saurabh Bagchi.

As today's distributed applications increase in complexity, it becomes increasingly difficult to detect errors and performance anomalies in these applications. In addition, some faults only manifest when the application is deployed at large scale. Most of the existing debugging tools scale poorly and do not automate the process of finding the origin of failures. Although it is desirable to automatically predict impending failures, most of the existing error detection approaches do not predict failures.

This dissertation proposes scalable techniques for error detection, problem localization, and failure prediction for distributed applications. First, an error detection and diagnosis technique for scientific applications is presented. The technique summarizes historic control-flow and timing information of MPI tasks using semi-Markov models. When a failure occurs, it leverages the models to determine the parallel task(s) and code region(s) where a fault is first manifested. The isolation of a difficult-to-catch bug in a large scale molecular dynamics simulation code and fault injections demonstrate the effectiveness of the technique. Second, frameworks for problem-localization and failure-prediction for commercial distributed applications are proposed. The frameworks learn application's normal behavior by monitoring multiple performance metrics. They then infer normal correlations between the metrics to pinpoint the suspicious metric(s) and code region(s) where faults are manifested. Using time-series models, the frameworks can predict impending failures with up to 15–51 minutes in advance. The frameworks are demonstrated with bug cases in Apache Hadoop, HBase, Android OS, and a campus-wide Java EE application.