

Random Testing for Distributed Systems with Theoretical Guarantees

Rupak Majumdar

Max Planck Institute for Software Systems, Kaiserslautern, Germany

rupak@mpi-sws.org

Abstract

Random testing has proven to be an effective way to catch bugs in concurrent and distributed systems. This is surprising, as the space of executions is enormous and conventional formal methods intuition would suggest that bad behaviors would only be found by extremely unlikely coincidences.

Empirically, many bugs in distributed systems can be explained by interactions among only a small number of features. Thus, one can attempt to explain the effectiveness of random testing under various “small depth” hypotheses. In particular, it may be possible to test all interactions of k features for a small constant k by executing a family of tests that is exponentially or even doubly-exponentially smaller than the family of all tests. Moreover, under certain conditions, a randomly chosen small set of tests is sufficient to cover all k -wise interactions with high probability.

I will describe two concrete scenarios. First, I will describe bugs in distributed systems caused by network partition faults. In many practical instances, these bugs occur due to two or three key nodes, such as leaders or replicas, not being able to communicate, or because the leading node finds itself in a block of the partition without quorum. In this case, I will show using the probabilistic method that a small set of randomly chosen tests will cover all “small partition” scenarios with high probability.

Second, I will consider bugs that arise due to unexpected schedules (interleavings) of concurrent events. Again, many bugs depend only on the relative ordering of a small number of events (the “bug depth” of the bug). In this case, I will show a testing algorithm that prioritizes low depth interleavings and a randomized testing algorithm that bounds the probability of sampling any behavior of bug depth k for a fixed k . The testing algorithm is based on combinatorial insights from the theory of partial orders, such as the notion of dimension and its generalization to d -hitting families as well as results on online chain partitioning.

Beyond the potential for designing or explaining random testing procedures, the technical arguments show the potential of combining “Theory A” and “Theory B” results to the important domain of software testing.

This is joint work primarily with Filip Nikić [1], and with Dmitry Chistikov, Simin Oraee, Burcu Kulahcioglu Özkan, Mitra Tabaei Befrouei, and Georg Weissenbacher. This work was partially funded by an ERC Synergy Award (ImPACT).

2012 ACM Subject Classification Theory of computation → Generating random combinatorial structures, Software and its engineering → Software testing and debugging

Keywords and phrases Random testing, Hitting families

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2018.1

Category Invited Paper

References

- 1 Filip Nikić. *Combinatorial Constructions for Effective Testing*. PhD thesis, University of Kaiserslautern, 2018.



© Rupak Majumdar;

licensed under Creative Commons License CC-BY

38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2018).

Editors: Sumit Ganguly and Paritosh Pandya; Article No. 1; pp. 1:1–1:1

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany