GCLab

# ANACIN-X

## Characterizing Non-Deterministic Communication

Dylan Chapp, Danny Rorabaugh, Heike Jagode, Sanjukta Bhomick, Michela Taufer

### Project Overview

Runtime non-determinism in High Performance Computing (HPC) applications presents steep challenges for computational reproducibility and correctness. These challenges are magnified in the context of complex scientific codes where the links between observable non-determinism and root causes are unclear. We apply a three-phase workflow to (1) build graph-structured models of non-deterministic communication in parallel applications; (2) identity windows of execution with maximum run-to-run variability; and (3) map runtime non-determinism to source code level root causes.

### Three-Phase Workflow for Characterizing Non-Determinism

**Phase 1:** Build event graph models from N execution traces of an MPI application, containing both communication and call-path data

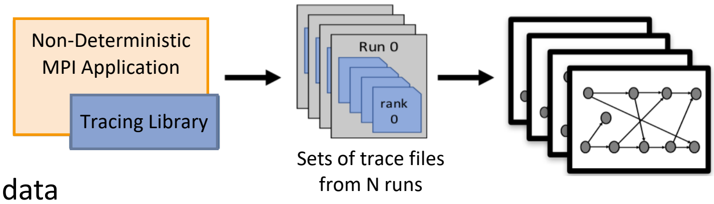


**Event Graphs:** Directed, labeled acyclic graphs that represent event ordering and associated call-paths

**Phase 2:** Compute pairwise graph kernel distances for a sliding window over the event graphs, quantifying the change in cross-run similarity over time

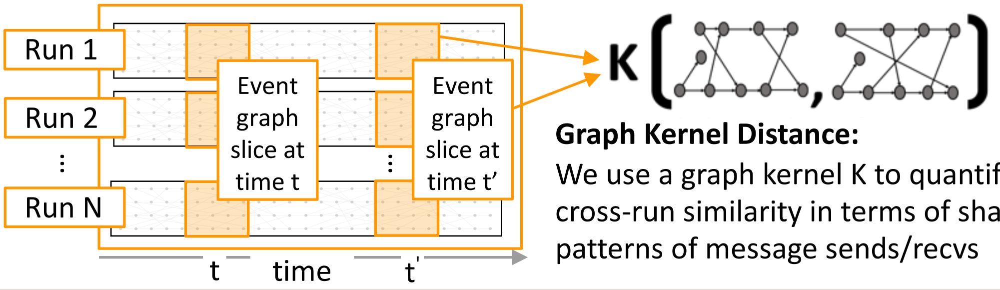


**Graph Kernel Distance:** We use a graph kernel K to quantify cross-run similarity in terms of shared patterns of message sends/recvs

**Phase 3:** Identify regions of high cross-run dissimilarity and use associated call-path data to map back to source-code level sources of non-determinism

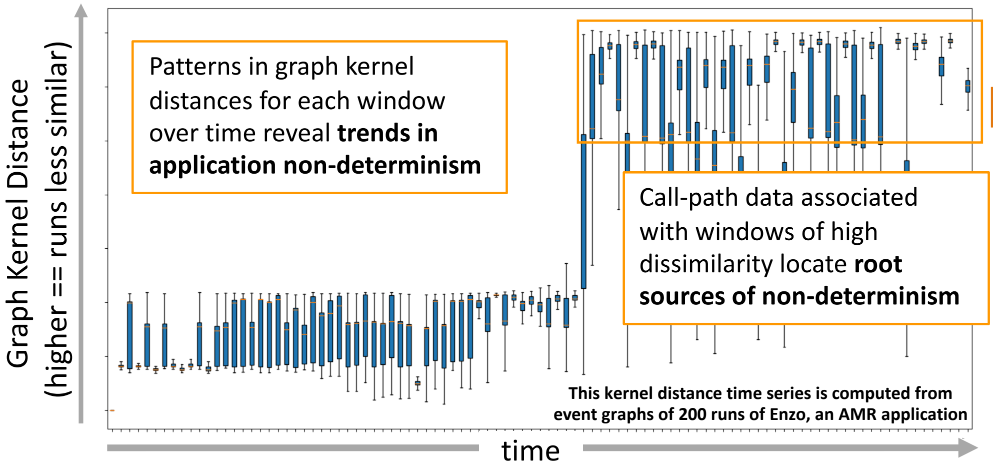


**This kernel distance time series is computed from event graphs of 200 runs of Enzo, an AMR application**

Relevant paths in call-graph identified

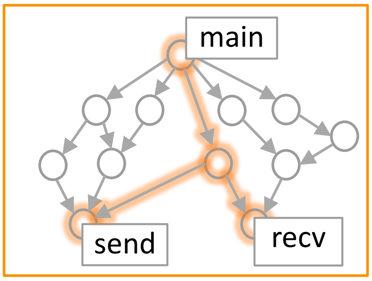


**Takeaways:**

- Patterns in non-deterministic communication found
- Root causes localized

THE UNIVERSITY OF TENNESSEE KNOXVILLE

NSF

GCL

Chapp, D., Rorabaugh, D., Sato, K., Ahn, D.H. and Taufer, M., 2019. A three-phase workflow for general and expressive representations of nondeterminism in HPC applications. *The International Journal of High Performance Computing Applications*, *33*(6), pp.1175-1184.

**Find out more at:**

https://globalcomputing.group