

Tracking Behavioral Changes in Distributed Systems using Distalyzer

Karthik Nagaraj, Charles Killian, Jennifer Neville
Purdue University



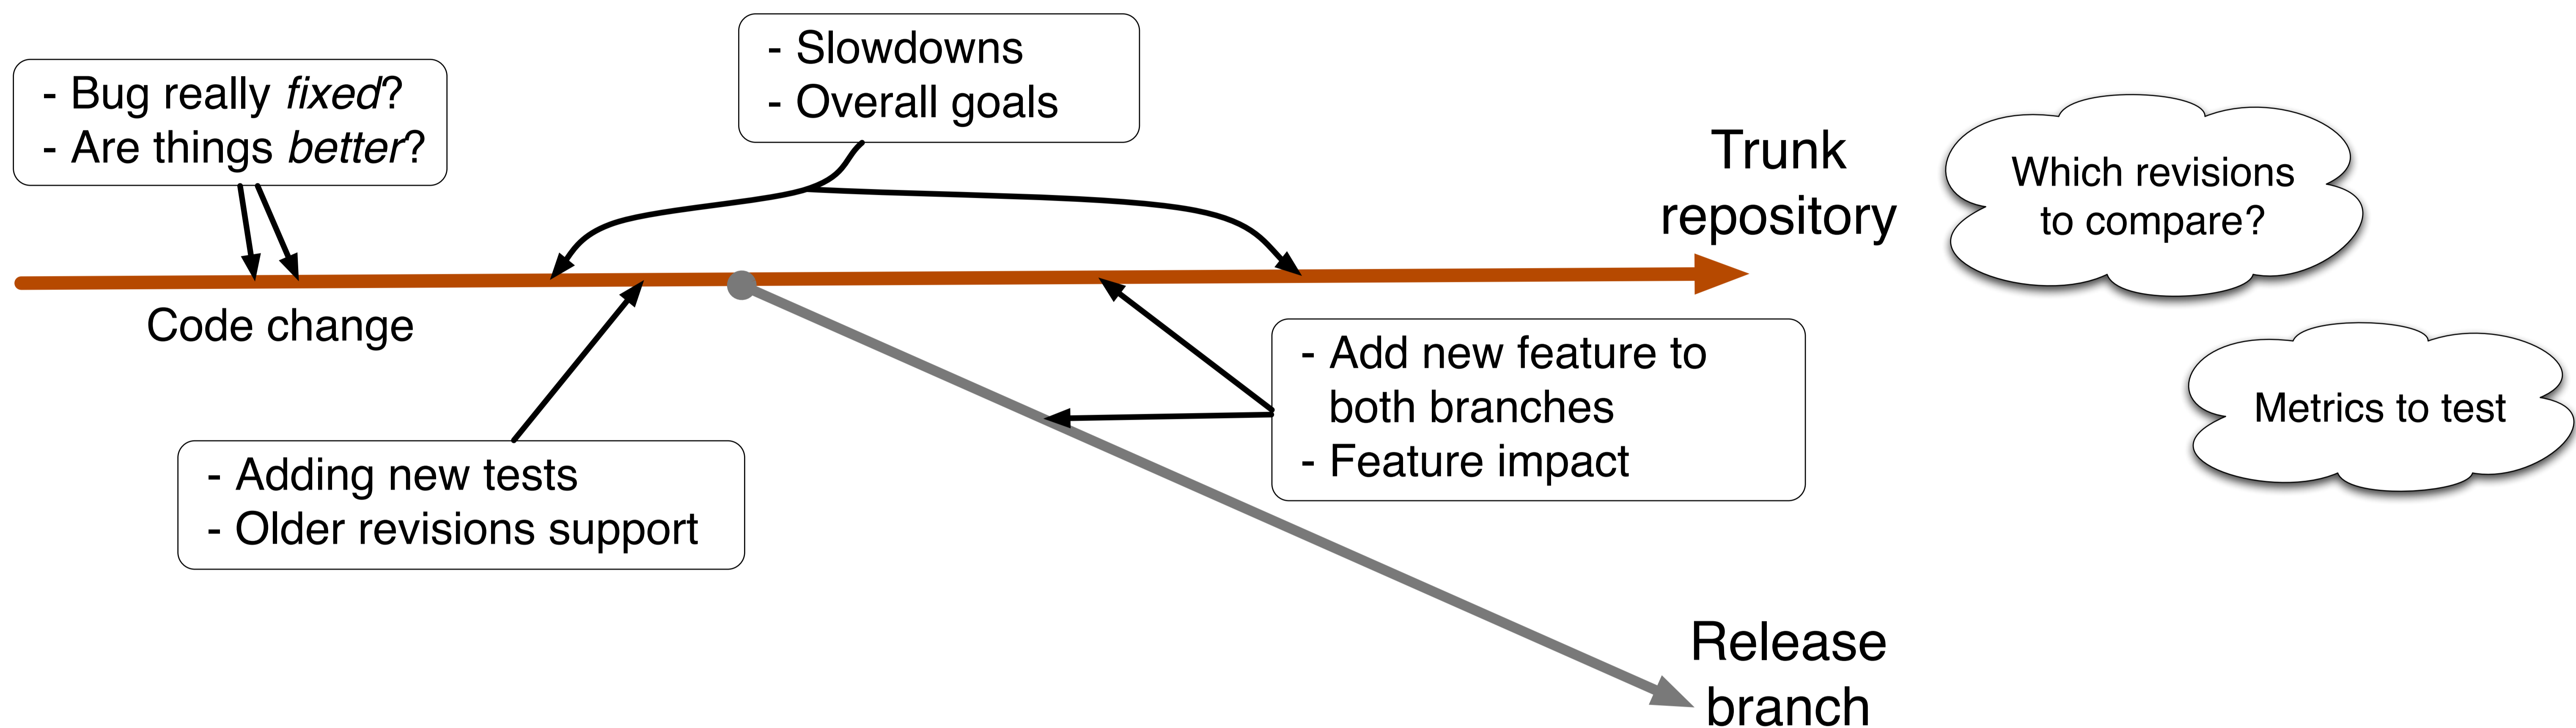
{knagara, ckillian, neville}@cs.purdue.edu

Problem

- Tracking concurrent updates to large systems projects, to analyze the effectiveness of bug fixes, and long-term impacts on the code-base.
- Project management after unit-tests is currently manual and tedious.

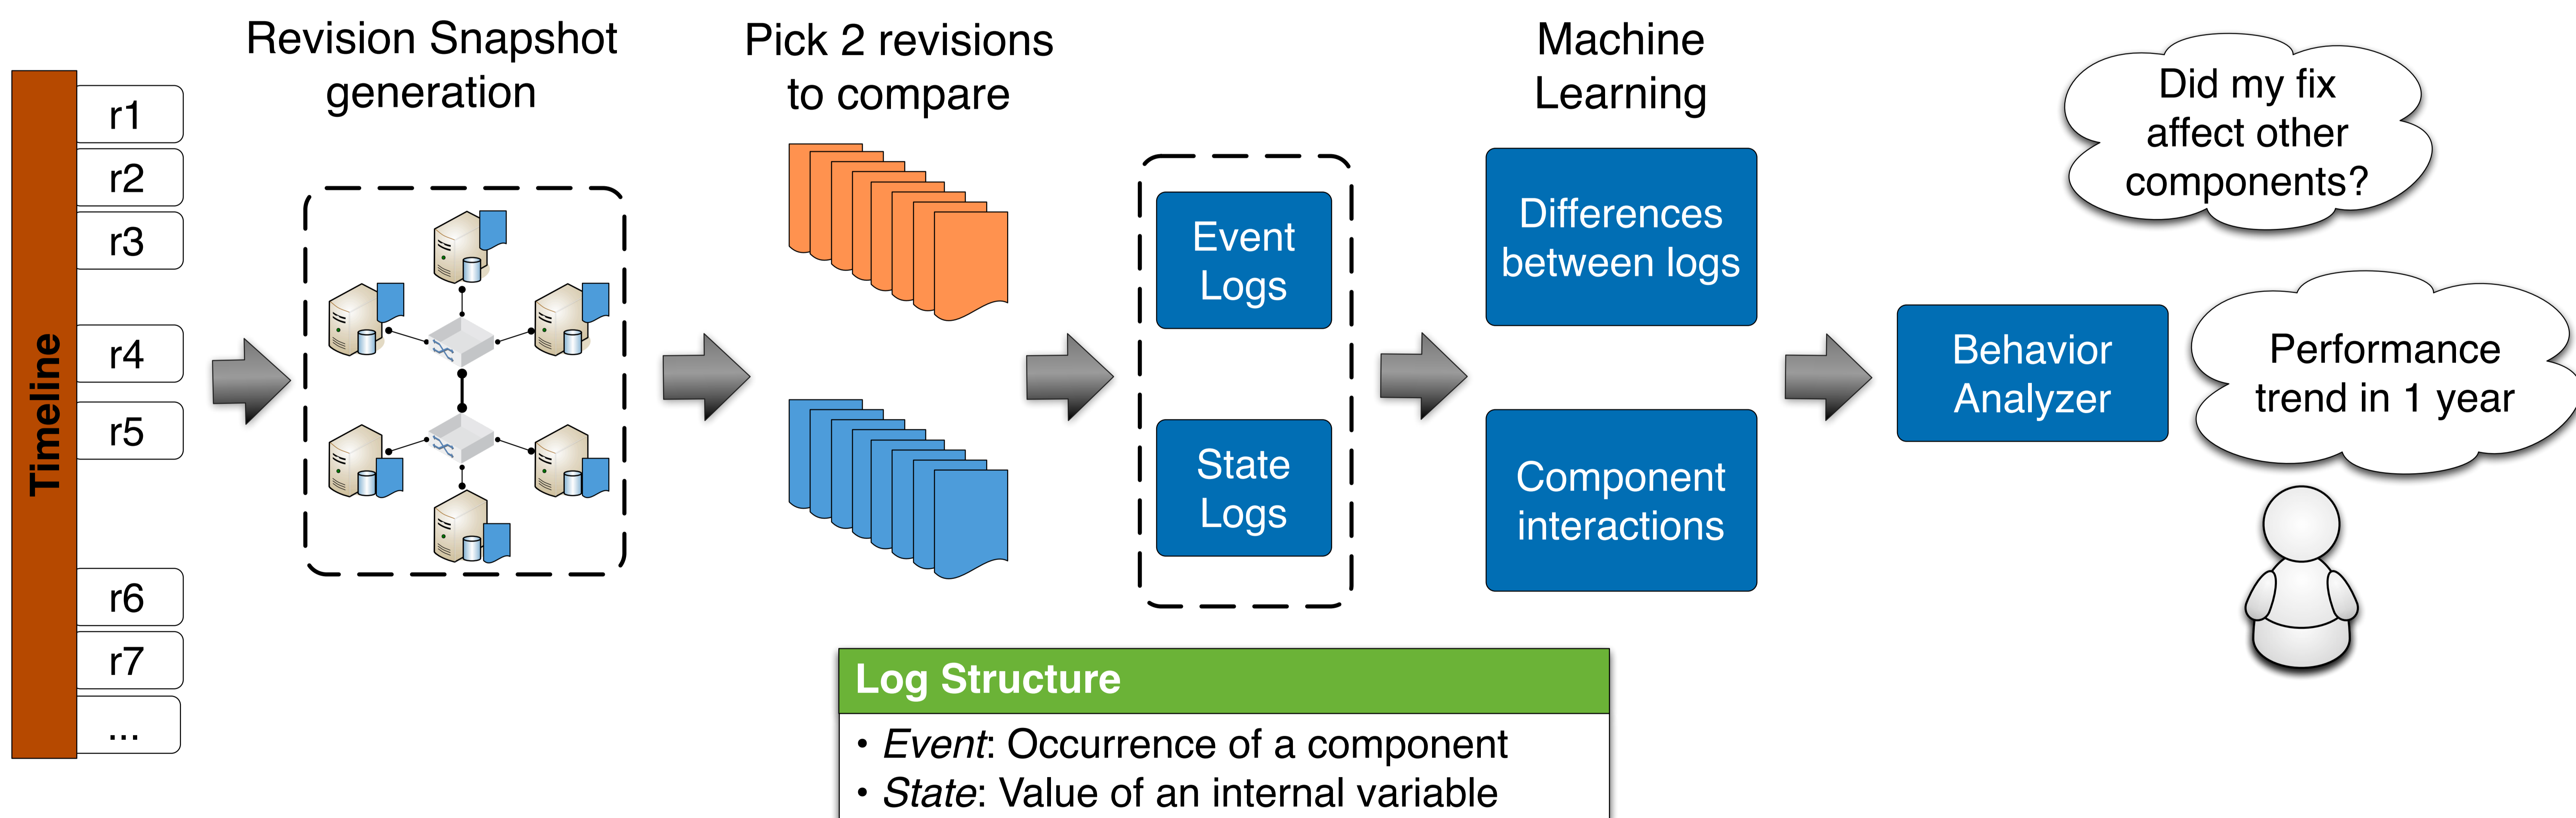
Challenge

- Code changes have unknown side-effects, due to complexity of distributed systems.
- Code & Tests change over time across multiple revision branches.



Approach

- Automated code monitoring infrastructure that leverages existing tests
- Minimal structure from abundant logs to systematically compare revisions.



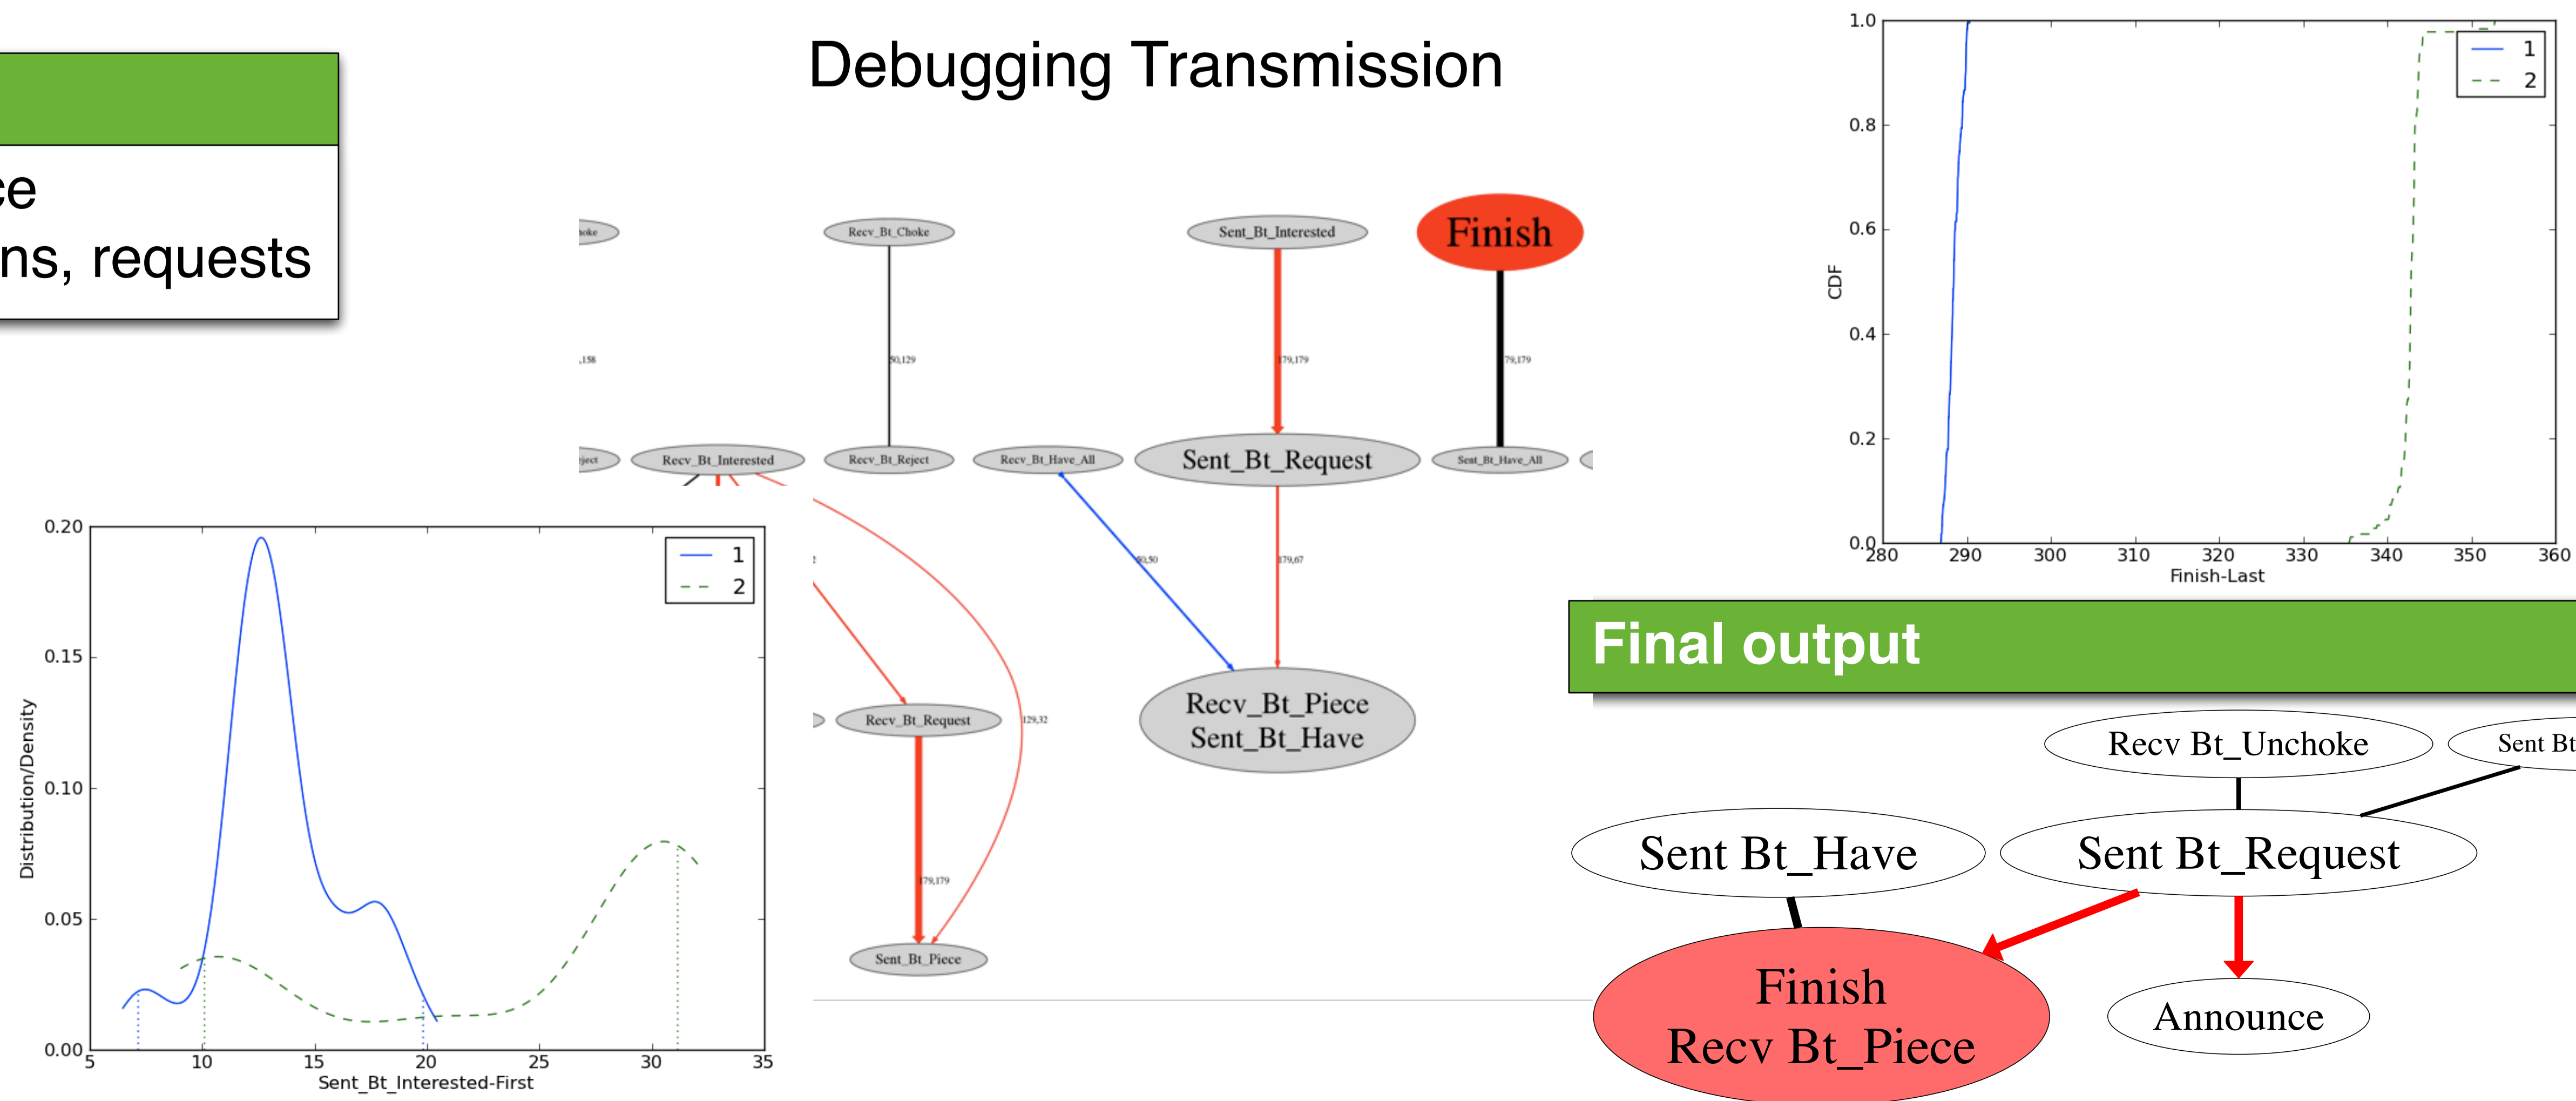
Phase One Results

Diagnosing overall performance problems

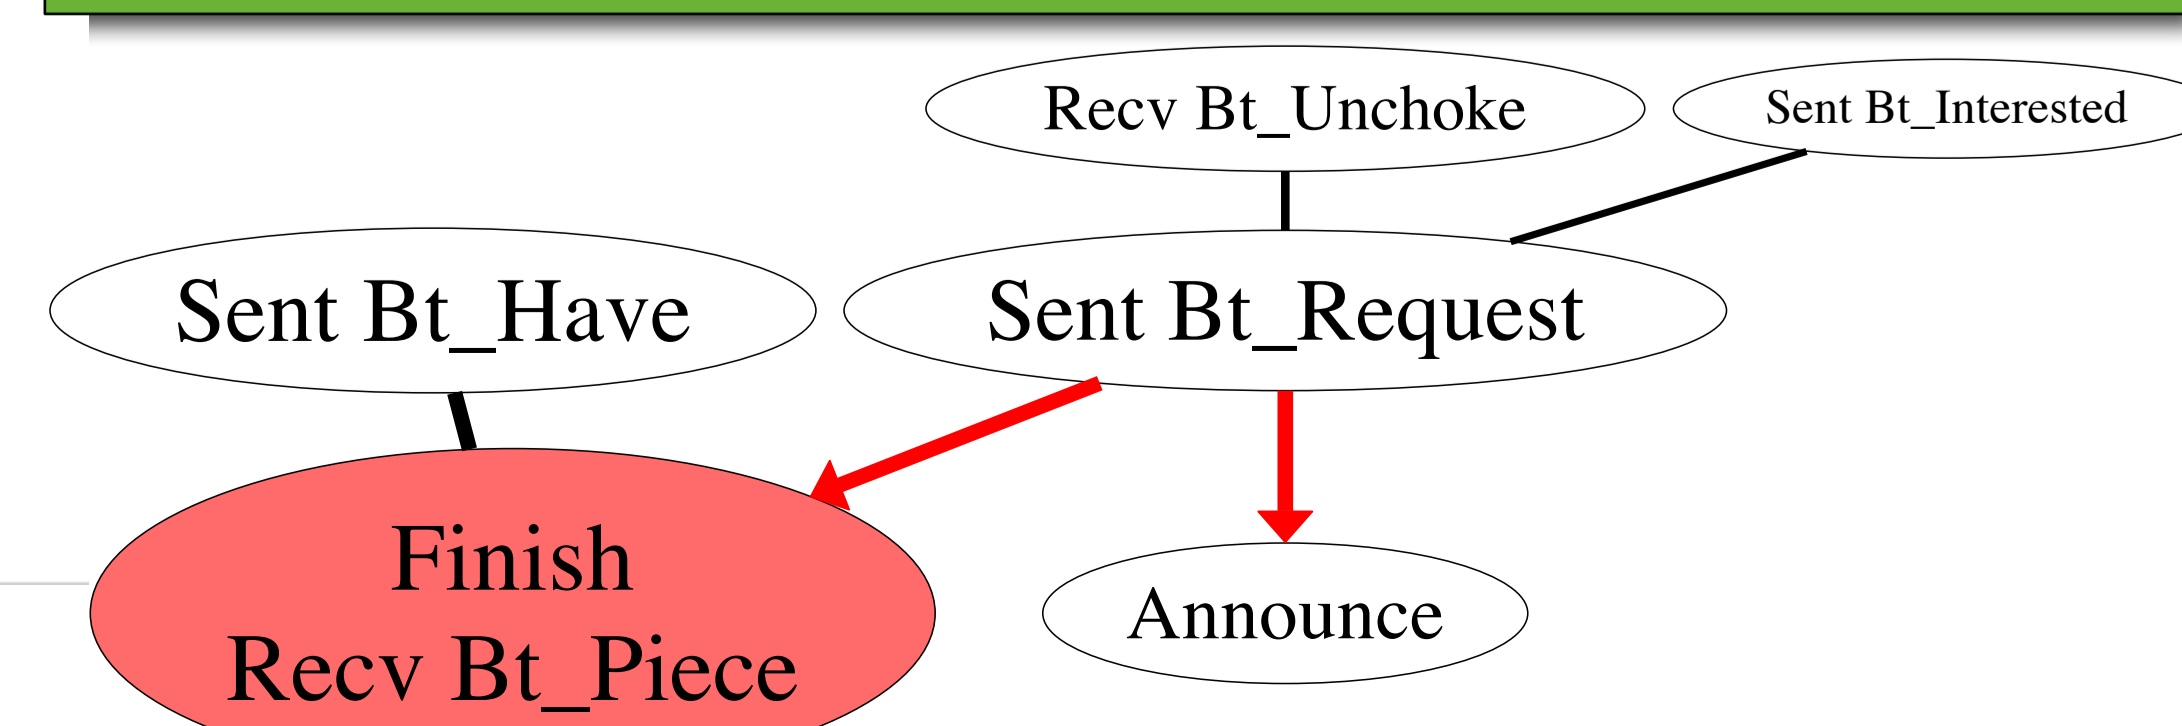
Overview

- Identify root cause of performance
- Compare different implementations, requests

Debugging Transmission



Final output



BitTorrent (Transmission vs. Azureus):
• 2 perf. bugs in Transmission

BigTable (HBase):
• 3 perf. issues in heavy tail

This research was supported in part by the National Science Foundation (NSF) grant #1054567-CNS.