

Flexible I/O Monitoring for Distributed Systems

Tânia Esteves

tania.c.araujo@inesctec.pt

INESC TEC & University of Minho

Supervisors: João Paulo and Rui Oliveira (INESC TEC & University of Minho)



Universidade do Minho



Cofinanciado por:



Motivation

- Developing, configuring and managing distributed systems are difficult, costly and challenging tasks.
- Tracing and analysis frameworks provide insights into how the systems' state evolves over time.
- Key for performance, correctness and security analysis.

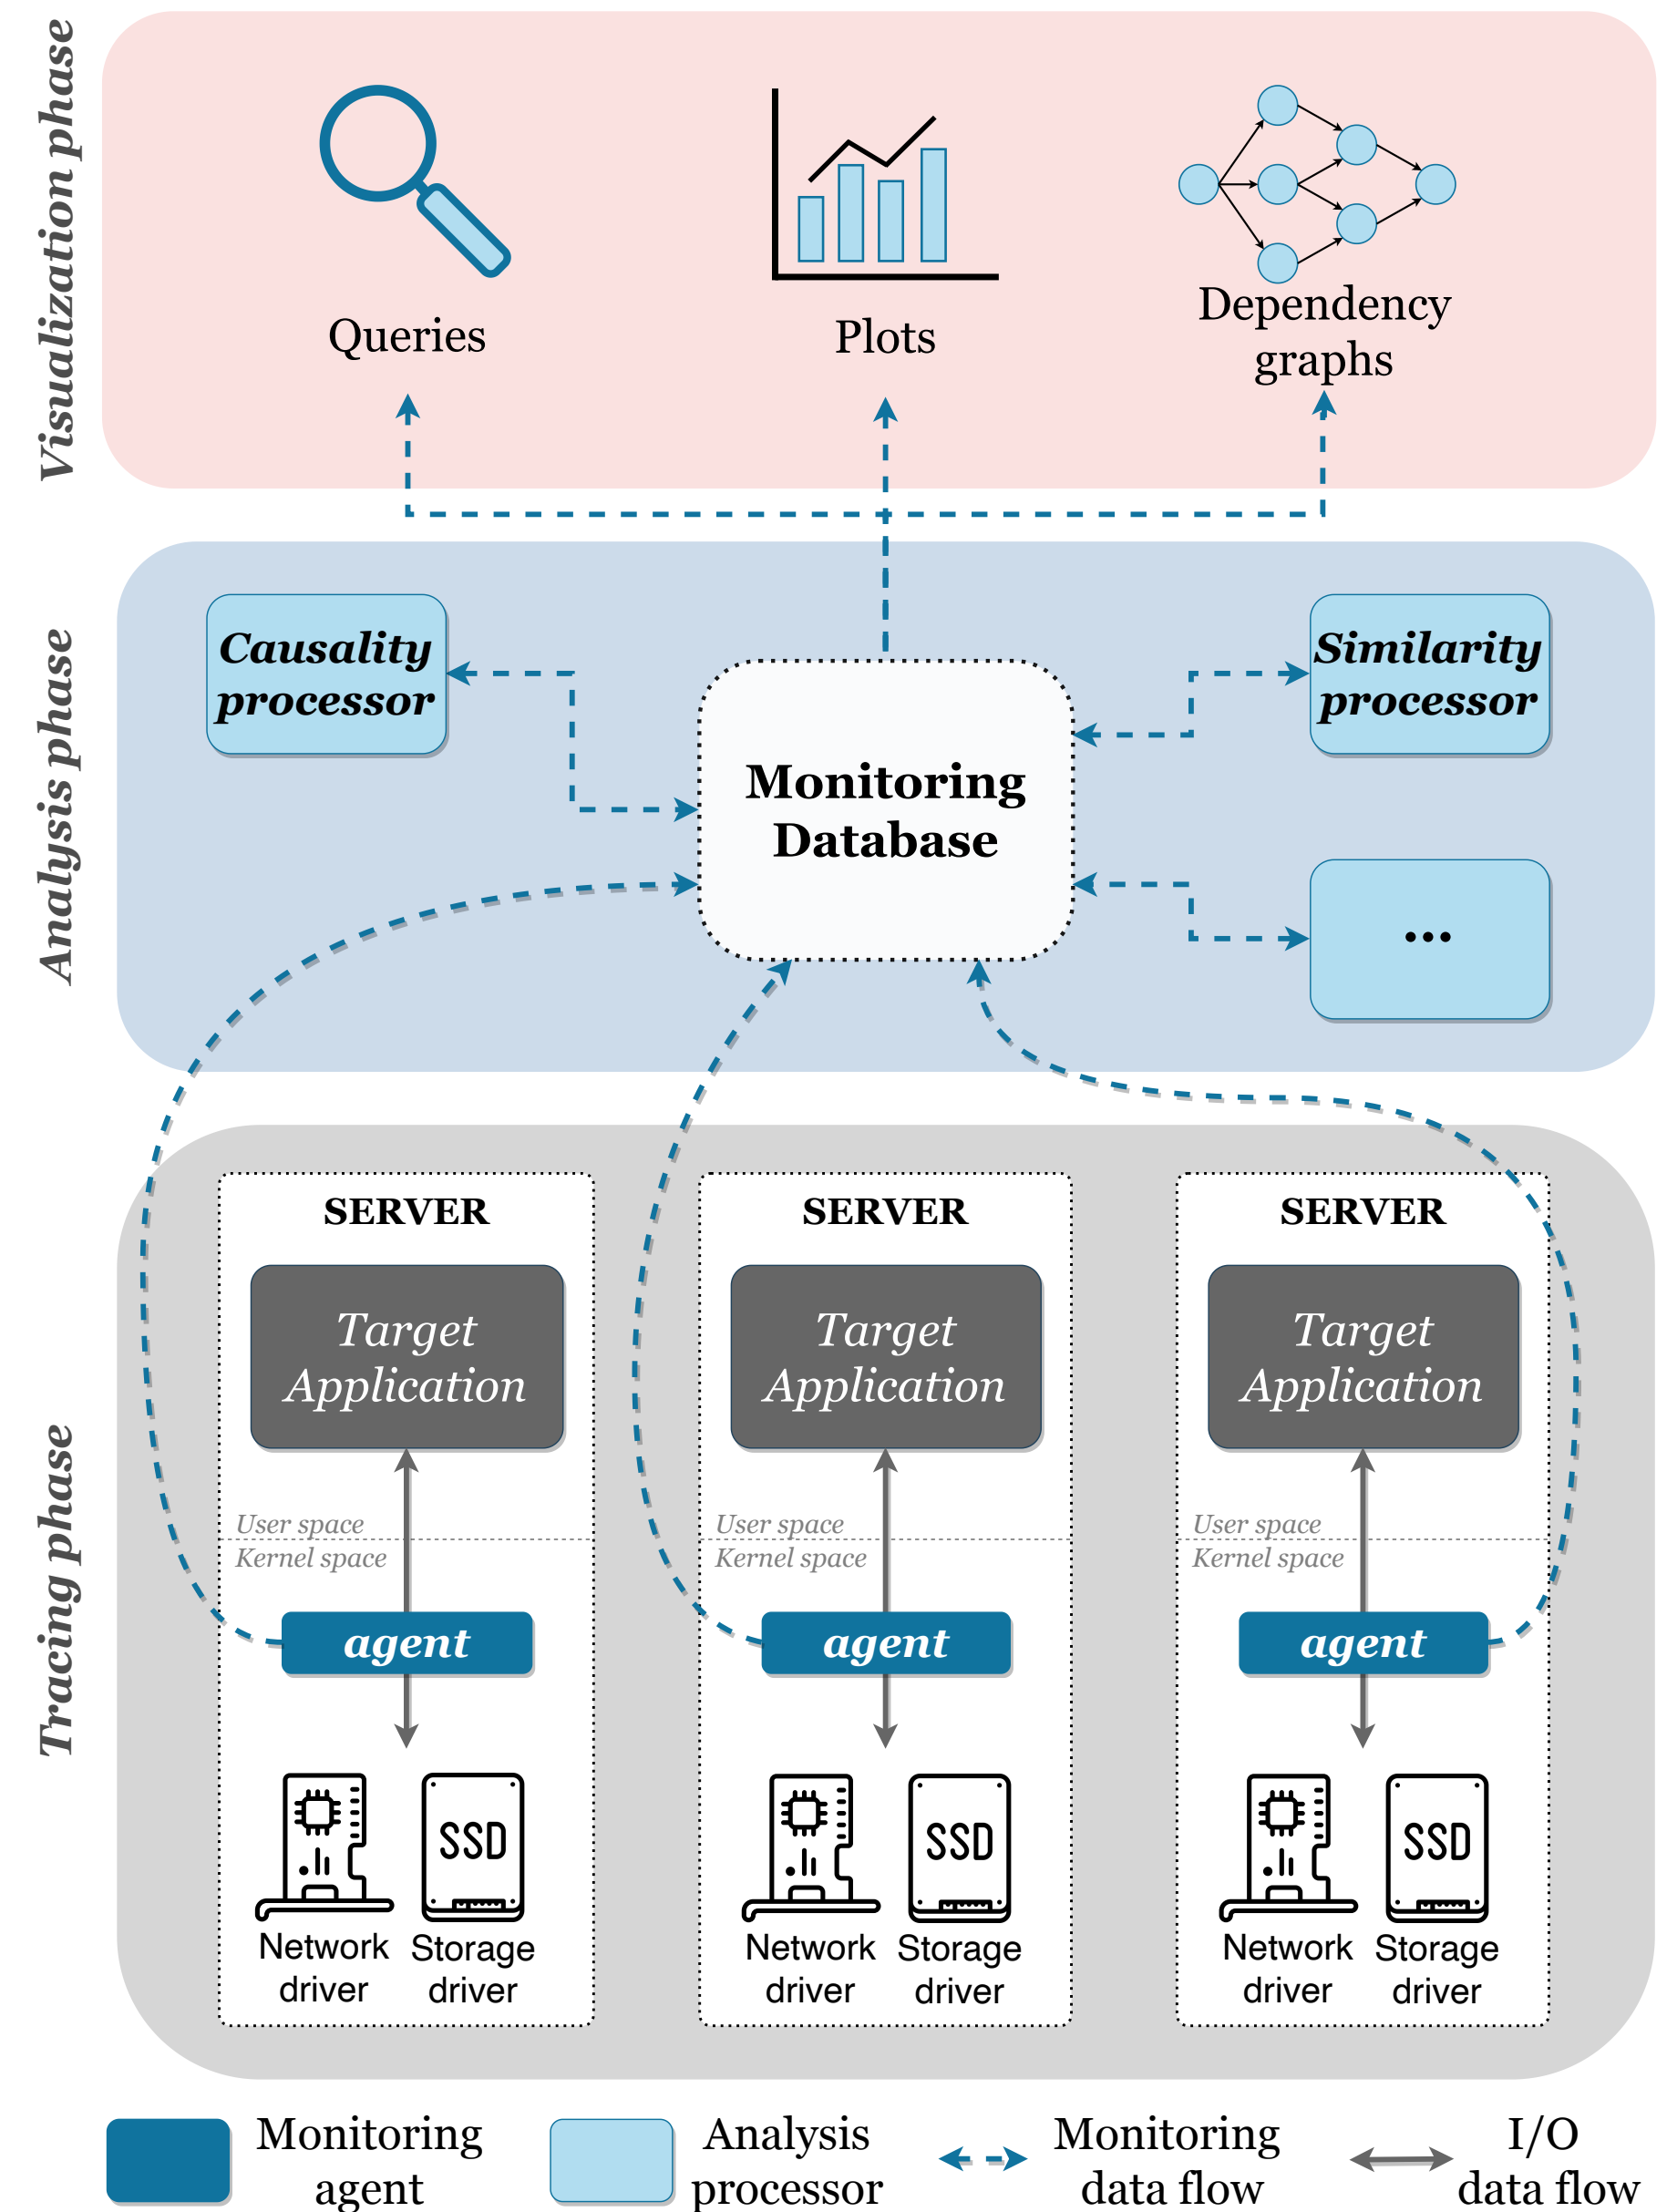
Problem

Current solutions suffer from several limitations:

- Intrusiveness (require instrumentation).
- Significant performance and storage overheads.
- Disregard of events' content.
- Lack of mechanisms to automatically analyze the tracing information.

Proposed solution

A novel content-aware monitoring solution to analyze the data flow of distributed systems in a non-intrusive way.



CAT: Content-aware tracing and analysis for distributed systems

- By supporting different tracing tools, CAT can balance the coverage of captured events with the impact on applications' performance.
- Depending on the applications' workload, full coverage of events can be attained with negligible performance and storage overheads.
- CAT can pinpoint correctness and dependability flaws that are not visible when using context-based state-of-the-art tools and are not detected by the integrity mechanisms of the applications.

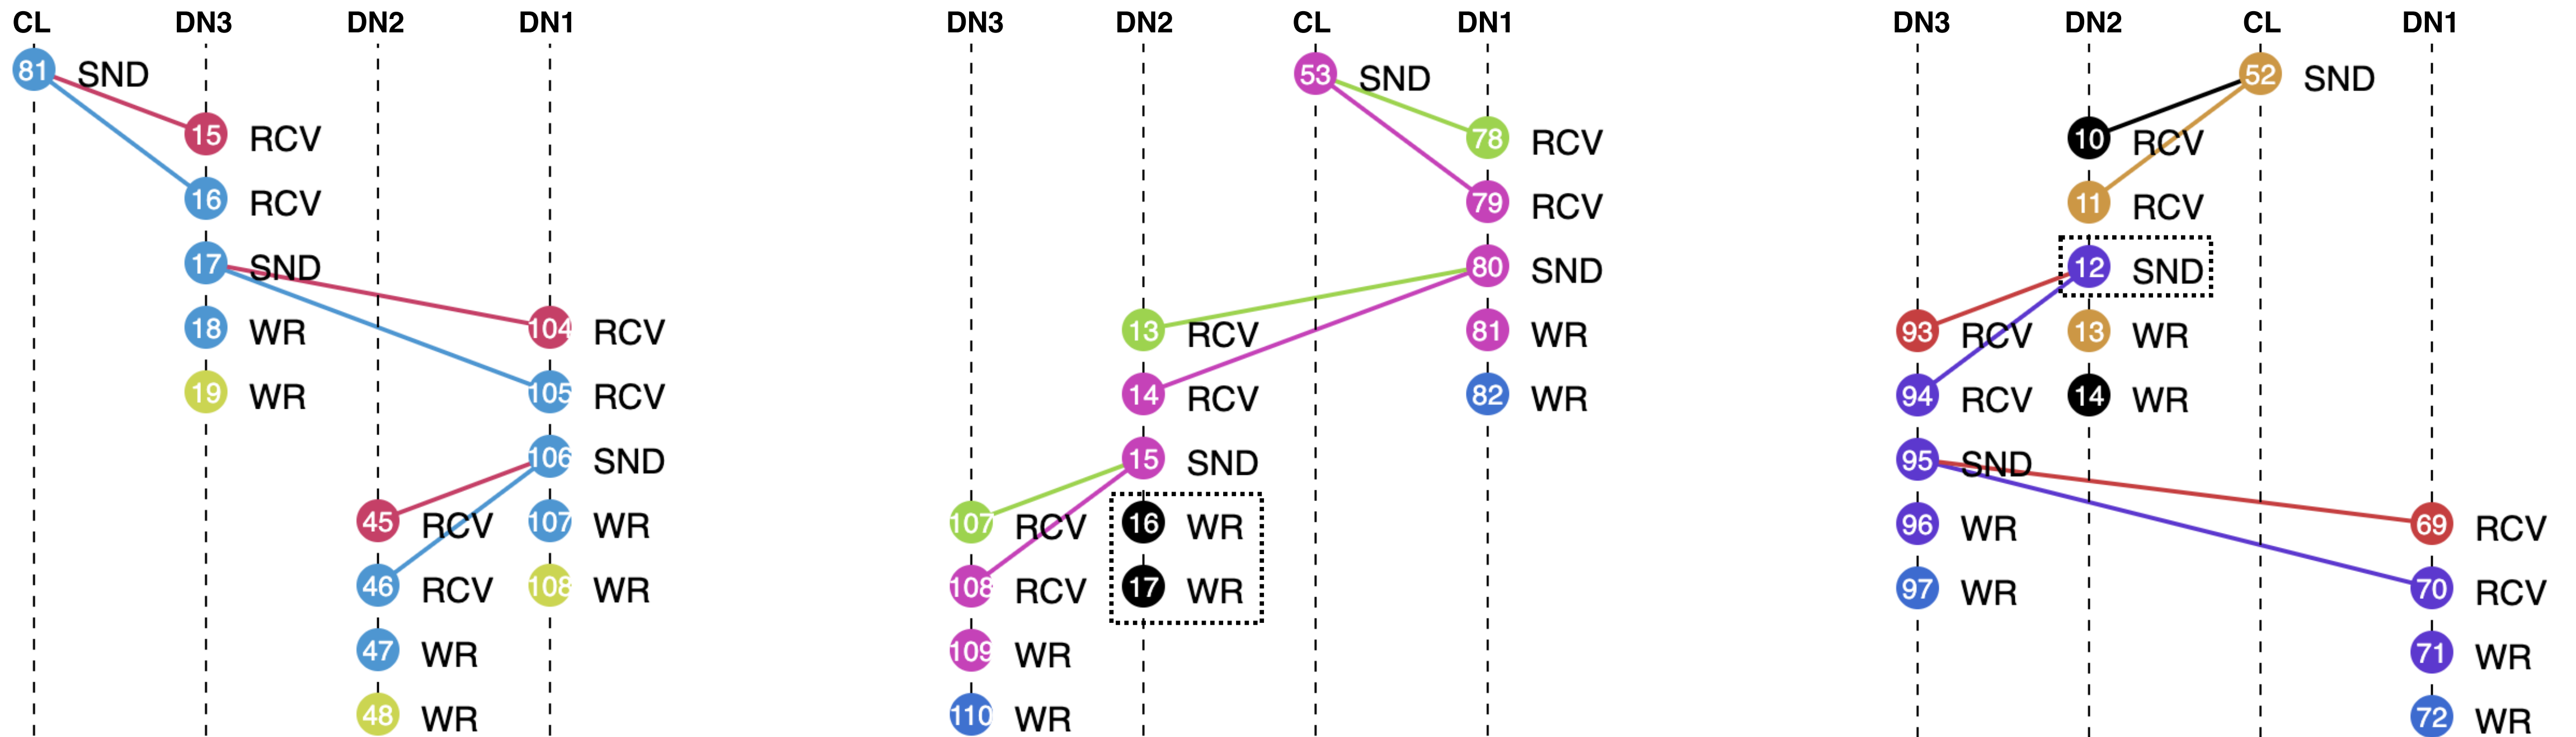


`dsrhaslab/cat`



<https://doi.org/10.1145/3464298.3493396>

HDFS file replication



(a) normal execution

(b) storage corruption

(c) network corruption

↑
Validation of correct behavior

↔
Identification of irregular behavior

Future work

- Improve the performance of the analysis processors.
- Provide a near-real-time solution.
- Allow the execution of queries and dynamic visual representations.