



THRESHOLD-BASED INFERENCE OF DEPENDENCIES IN DISTRIBUTED SYSTEMS

Anshul Rastogi

Mentors: Darby Huye, Max Liu, & Dr. Raja Sambasivan

Project partner: Tanmay Gupta

OVERVIEW

01 DISTRIBUTED SYSTEMS
& TRACING

02 *THE MYSTERY MACHINE*

03 *SCOOBY SYSTEMS*

04 FUTURE WORK

DISTRIBUTED SYSTEMS




Distributed systems are

Networks of devices/machines ("nodes"), such as computers or servers that communicate with one another to complete tasks

Distributed systems surround us.

Some examples:

- Google
- Facebook
- Cellular networks



WHAT IS DISTRIBUTED TRACING?

Distributed tracing is when a software (instrumentation) tracks the flow of service requests in a distributed system

This is useful for:

- Debugging
- Regulation
- Analyzing performance

Gives developers information about interactions in the system.

A *trace* is the data about the request path that results from tracing

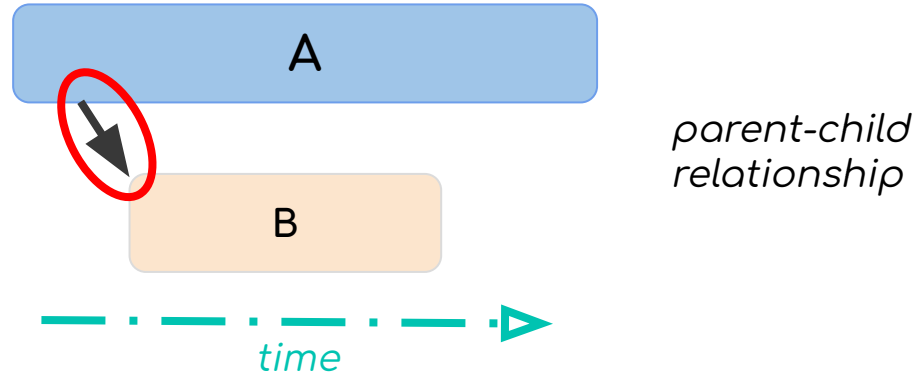


TRACING: SPANS

Span (definition):

“The ‘span’ is the primary building block of a distributed trace, representing an individual unit of work done in a distributed system.”

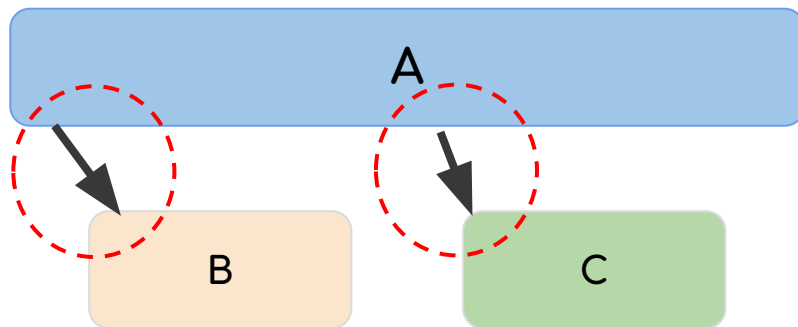
- OpenTracing.io






TRACING: CHALLENGES WITH SPANS

- Caller-callee relationships
 - B dependent on A
 - C dependent on A
- **Is C dependent on B?**
- Instrumentation could help
- Tedious, hard to do with heterogeneous systems
- How to infer relationship with minimal instrumentation?




TRACING: SPANS VS. SEGMENTS



Spans

- Many relationships between parts of A and B
 - E.g. asynchronous concurrency missed



Segments

- A₁ happens before A₂ and B; A₂ and B have no dependency on one another
- Requires more instrumentation

THE IDEAL TRACE: INTRODUCTION

Developers use tracing to determine causality within a system.


Therefore, causal relationships should be thoroughly represented throughout the entire system in a trace

To have enough information, more instrumentation can be used

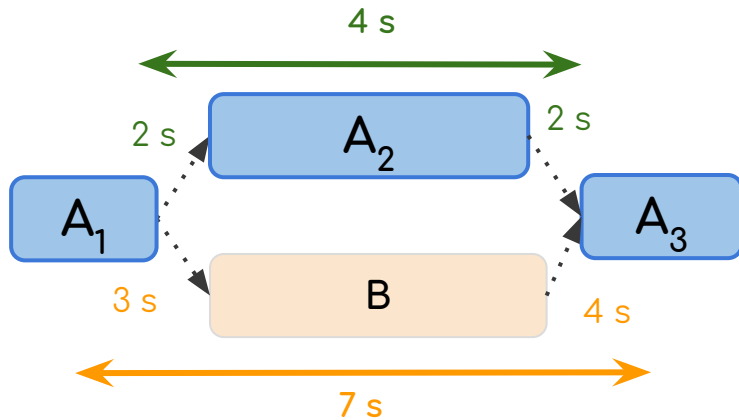


THE IDEAL TRACE: CAUSALITY

- **Happens-before** relationship
 - A_1 must occur for B to occur
- **Concurrent** relationship
 - Segments can happen in any order or simultaneously
 - No dependency between segments
 - e.g. B and A_2



THE IDEAL TRACE: APPLICATIONS



Critical Path: longest path in trace

Determines minimum request path latency, shows amount of “slack” in system

“Slack Analysis”

OVERVIEW

- 01 DISTRIBUTED SYSTEMS & TRACING
- 02 *THE MYSTERY MACHINE*
- 03 *SCOOBY SYSTEMS*
- 04 FUTURE WORK

THE MYSTERY MACHINE: GCM



The *Mystery Machine* produces a **Global Causal Model (GCM)**

- to approach the ideal trace, *The Mystery Machine* uses many traces rather than just one to infer dependencies while minimizing instrumentation
- uses a *segment-based model* – more clarity
- shows happens-before dependencies between every segment across the traces such that the dependencies hold for *every trace*
- presents inferred GCM of all system interaction to user

THE MYSTERY MACHINE: LIMITATIONS

- **Problem:** Assumes enough traces to capture every interaction in request paths
- **Problem:** Assumes that all traces are correct, leaving no room for error such as:
 - Clock skew
 - Anomalies in structure (caused by bugs)





OVERVIEW

- 01 DISTRIBUTED SYSTEMS & TRACING
- 02 *THE MYSTERY MACHINE*
- 03 *SCOOBY SYSTEMS*
- 04 FUTURE WORK

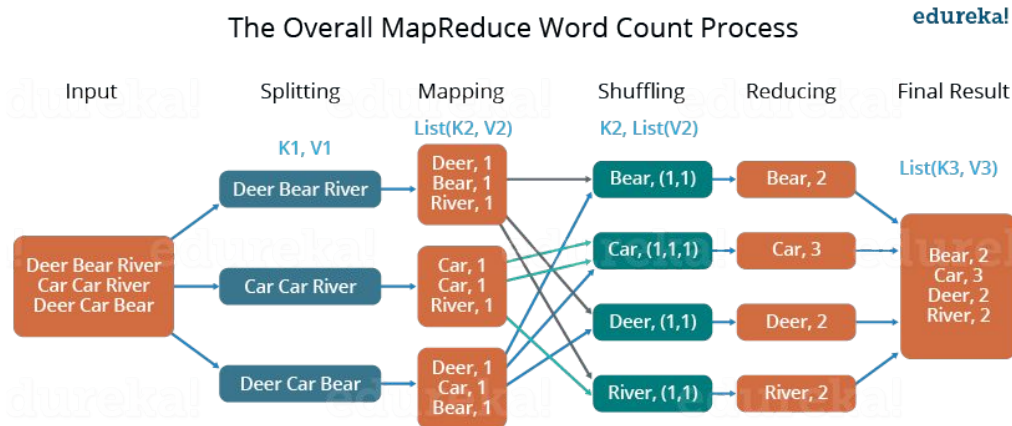
USER-DEFINED THRESHOLD

- Address rigidity of *The Mystery Machine*
 - Increase tolerance
 - Flukes should not affect GCM
- Introduce threshold – relationship is only removed from final model if violations to relationship exceeds threshold
- User has freedom to choose threshold amount



SCOOBY SYSTEMS: SCALABILITY

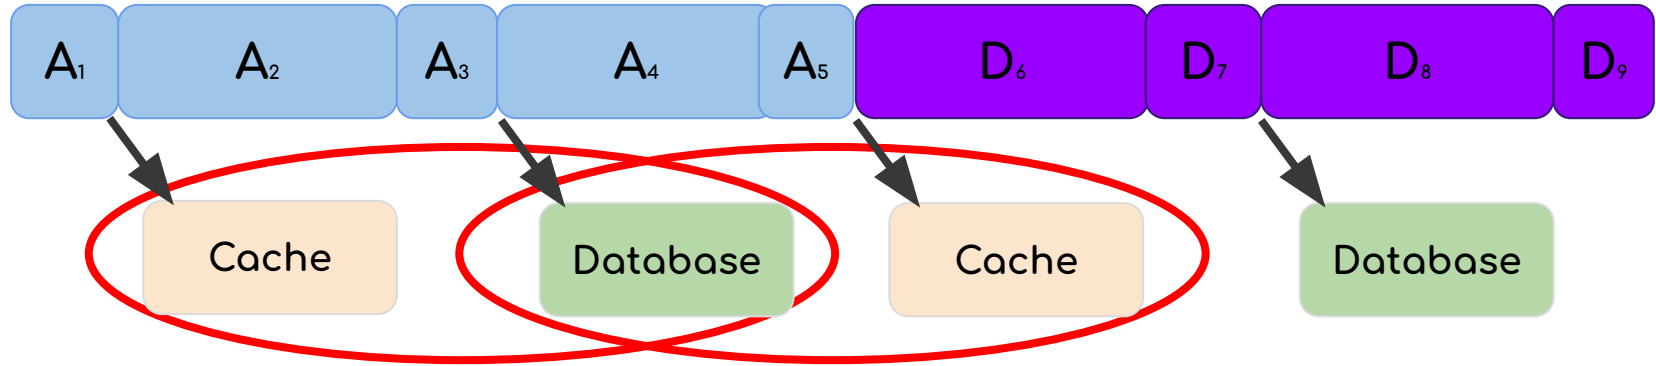
- Implemented in Hadoop MapReduce for scalability
- MapReduce allows the program to execute efficiently across a cluster of devices
- Have had successful runs of MapReduce implementation on small “pseudo-clusters” (simulated clusters) as we don’t currently have access to an actual distributed system



OVERVIEW

- 01 DISTRIBUTED SYSTEMS & TRACING
- 02 *THE MYSTERY MACHINE*
- 03 *SCOOBY SYSTEMS*
- 04 FUTURE WORK

ADDRESSING REPEATS



- **Problem:** *The Mystery Machine* assumes that there are no repeats of the same segment or event
- GCM concludes no dependency between B and C because *The Mystery Machine* sees both $B \rightarrow C$ and $C \rightarrow B$
- But what if $B \rightarrow C$ is the true structure?

NEXT STEPS: BACKTRACE DEPTH

Without further instrumentation, need a way to distinguish repeats

Add more information:

- Currently *Scooby Systems* increases specificity in instances of a segment by associating them with the process ID
- Backtrace is more accurate – e.g. A-cache and D-cache instead of repeats of just cache
- But what if A calls cache and database twice in the same trace?
 - Could increase backtrace depth – include A's parent information
 - However, at some point, would be unable to consolidate information about cache

MOVING FORWARD

- Implement backtrace information
- Proposing solutions / further analysis of *The Mystery Machine* limitations
- Evaluation
 - Traces from DeathStarBench

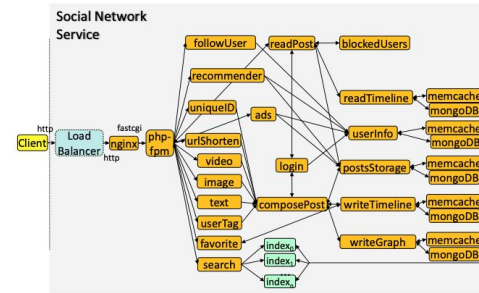


Figure 4. The architecture (microservices dependency graph) of Social Network.

ACKNOWLEDGEMENTS

Thanks to **Darby Huye, Max Liu, Raja Sambasivan**, & D.O.C.C. Lab for their guidance



Thank you to **Tanmay Gupta** for being my PRIMES partner for the first iteration of *Scooby Systems* in 2021.

CREDITS: This presentation template was created by **Slidesgo**, including icons by **Flaticon**, infographics & images by **Freepik**

Thanks to the **PRIMES Program** for providing this opportunity



**THANK YOU FOR YOUR
ATTENTION**



ANY QUESTIONS?