



ENHANCING DISTRIBUTED TRACING TO ORDER EVENTS

Tanmay Gupta and Anshul Rastogi
Mentor: Dr. Raja Sambasivan,

DISTRIBUTED SYSTEMS TODAY


Distributed systems surround us.

Some examples:

- Google
- Facebook
- Cellular networks*

Distributed systems are

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



Google Search

I'm Feeling Lucky




WHAT IS DISTRIBUTED TRACING?

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

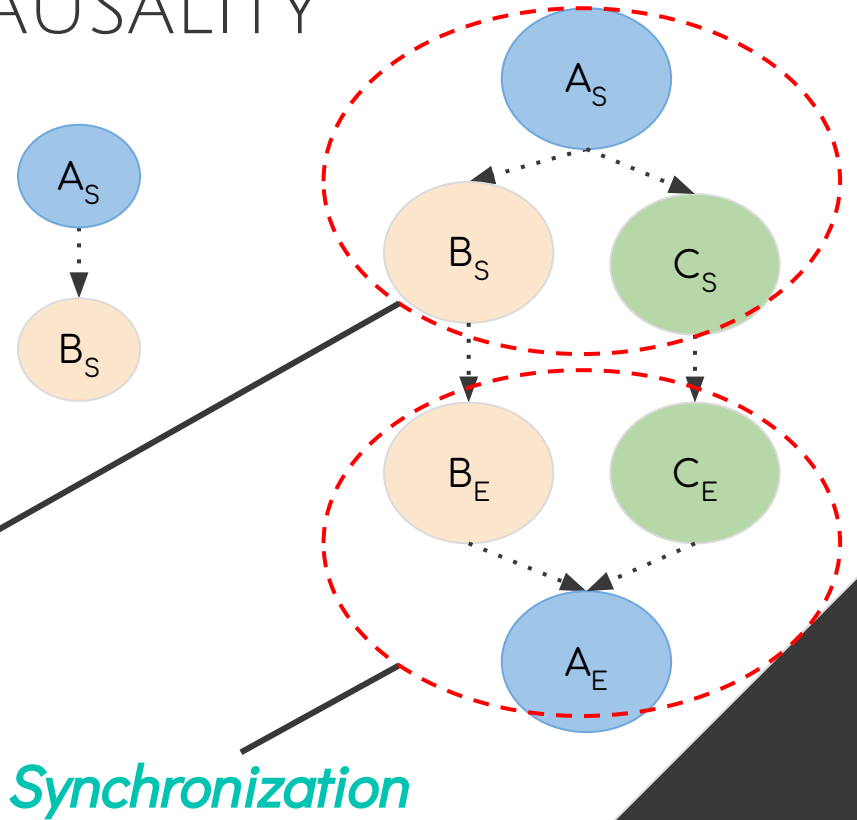
This is useful for:

- Debugging
- Regulation
- Analyzing performance




THE IDEAL TRACE: CAUSALITY

- **Happens-before** relationship
 - A_s must occur for B_s to occur
- **Concurrent** relationship
 - Events could happen in any order or simultaneously



THE IDEAL TRACE: APPLICATIONS



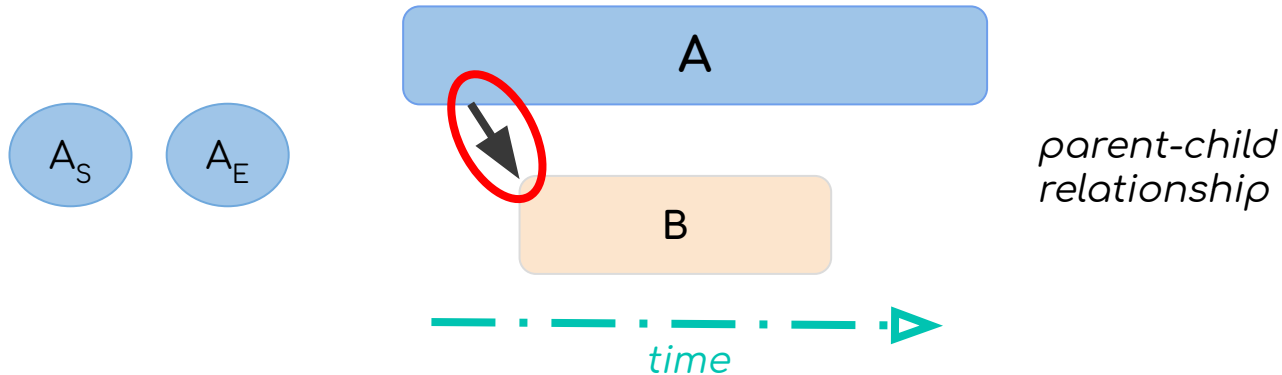
Critical Path: longest path in trace

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: THE CHALLENGES

- Caller-callee relationships
 - B dependent on A
 - C dependent on A
- Is C dependent on B?




TRACING: THE CHALLENGES

- Is C dependent on B?
 - Instrumentation could help
 - Fork in threads => concurrent
 - Hard to do with heterogeneous systems





CAN WE USE **BIG**
DATA TO GET
CLOSER TO THE
IDEAL TRACE?






OVERVIEW


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

SPANS, SEGMENTS, & EVENTS




Spans

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



Segments

- Traces missing logs



Events

THE MYSTERY MACHINE: GCM

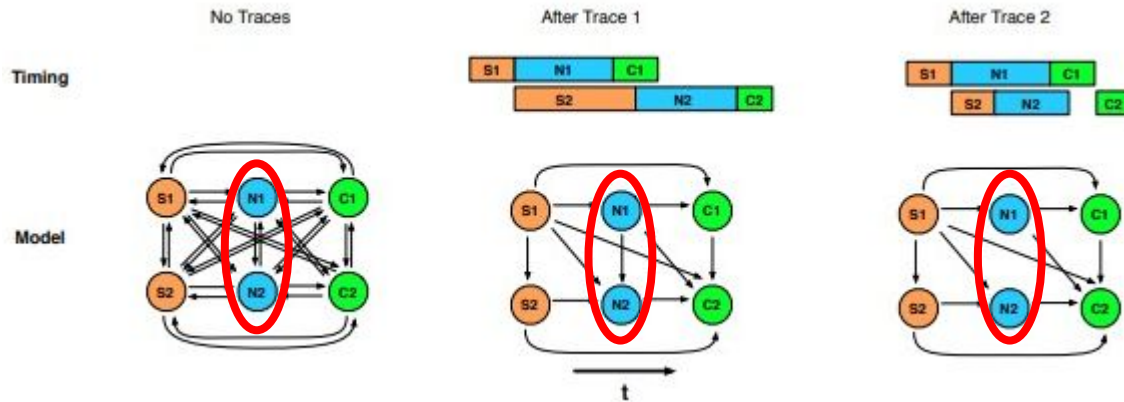


The *Mystery Machine* produces a **Global Causal Model (GCM)**

- Uses a *segment-based model*.
- shows happens-before dependencies between every segment across the traces such that the dependencies hold for *every trace*.

THE MYSTERY MACHINE: HOW IT WORKS

- First creates all hypothetical happens-before edges between all segments to create GCM
- Then takes happens-before relationships across the traces and removes violated edges from GCM
- Reaches final GCM once it has iterated through all traces

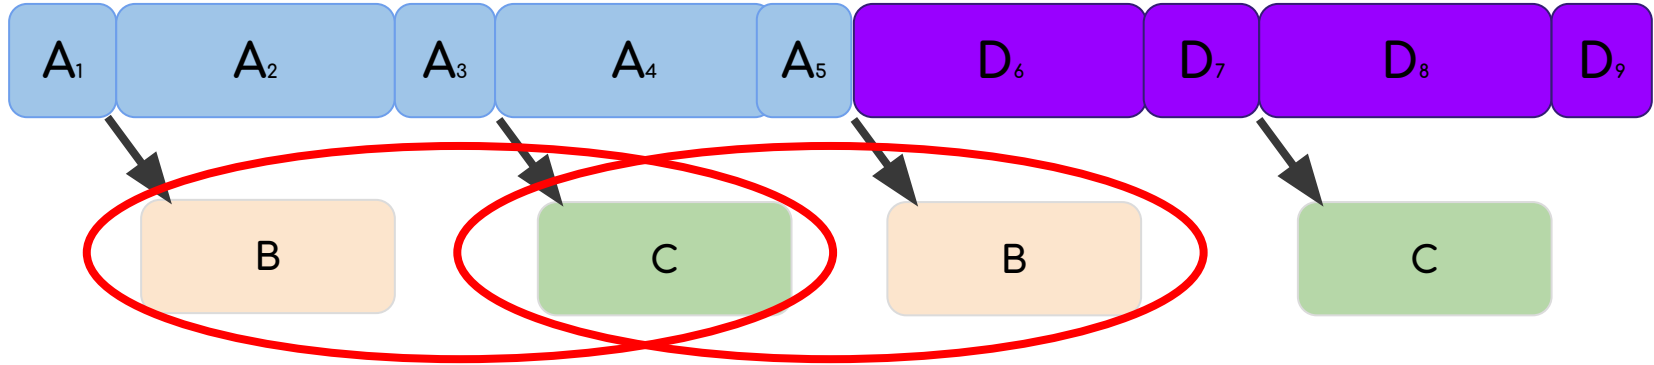


THE MYSTERY MACHINE: LIMITATIONS

- **Problem:** Assumes enough natural variation
- *The Mystery Machine* assumes Facebook-specific things
- **Problem:** Assumes that all traces are correct, leaving no room for error such as:
 - Clock skew
 - Anomalies in structure (caused by bugs)



THE MYSTERY MACHINE: LIMITATIONS



- **Problem:** Mystery Machine assumes that there are no repeats of the same segment or event
 - GCM starts with interconnected B and C
 - Sees **B** -> **C**; removes **C** -> **B** from model
 - Sees **C** -> **B**; removes **B** -> **C** from model
- But what if **B** -> **C** is the true structure?



OVERVIEW

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

USER-DEFINED THRESHOLD

- Address **rigidity** of *The Mystery Machine*
 - Don't eliminate edge after just one violation

$$\begin{aligned} s_{A \rightarrow B} &= \# \text{ of successes of } A \rightarrow B \\ v_{A \rightarrow B} &= \# \text{ of violations of } A \rightarrow B \\ u_{A \rightarrow B} &= A, B \text{ do not both occur} \end{aligned}$$

- Chose **success-based model** contrary to *The Mystery Machine*
 - Counts successes of each relationship
 - User has freedom to choose threshold in terms of $s_{A \rightarrow B}$ and $v_{A \rightarrow B}$


$$S_{A \rightarrow B} = S_{A \rightarrow B}$$

$$V_{A \rightarrow B} = S_{B \rightarrow A} + q_{A \rightarrow B} \quad \text{A and B at same time}$$

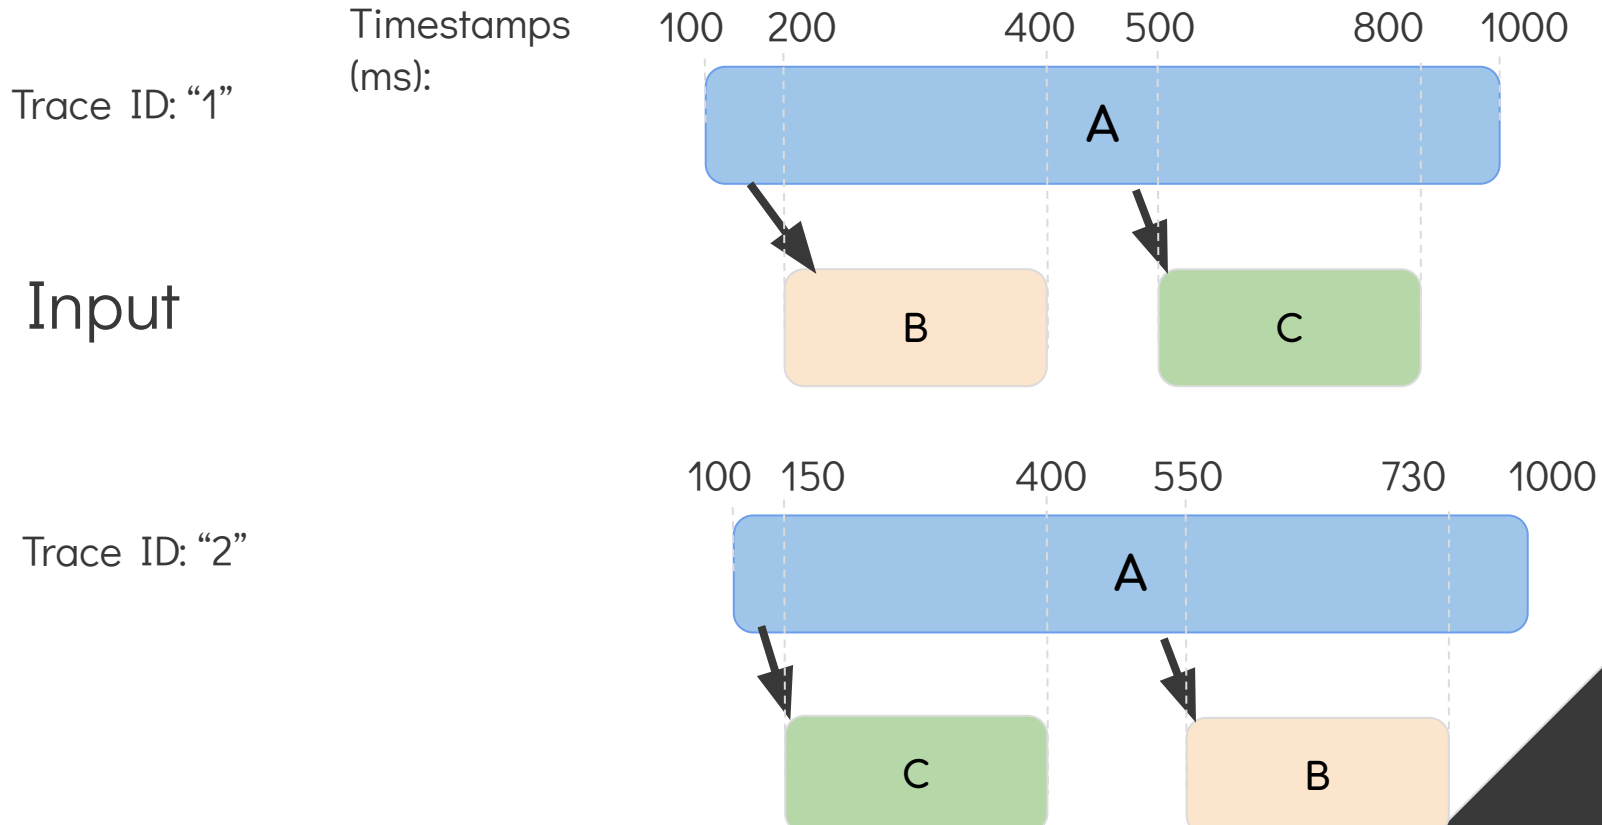
$$U_{A \rightarrow B} = T - S_{A \rightarrow B} - V_{A \rightarrow B} \quad \text{Total \# of occurrences}$$

SCOOBY SYSTEMS

- Preprocessing
 - Create Causal Model
- Main Algorithm
 - for each trace
 - for each pair of spans
 - update causal model
- Apply Threshold
 - For each relationship in causal model
 - if (passes threshold) add edge to final model

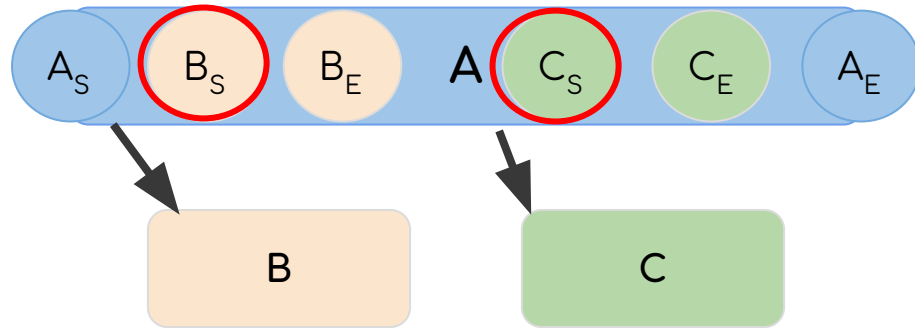


SCOOBY SYSTEMS IN ACTION



SCOOBY SYSTEMS IN ACTION

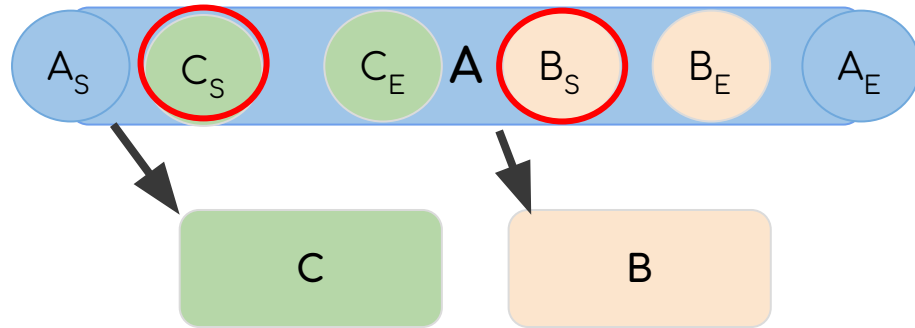
Trace 1



	A_S	A_E	B_S	B_E	C_S	C_E
A_S	0	1	1	1	1	1
A_E	0	0	0	0	0	0
B_S	0	1	0	1	1	1
B_E	0	1	0	0	1	1
C_S	0	1	0	0	0	1
C_E	0	1	0	0	0	0

SCOOBY SYSTEMS IN ACTION

Trace 2



	A_S	A_E	B_S	B_E	C_S	C_E
A_S	0	1	1	1	1	1
A_E	0	0	0	0	0	0
B_S	0	1	0	1	0	0
B_E	0	1	0	0	0	0
C_S	0	1	1	1	0	1
C_E	0	1	1	1	0	0

SCOOBY SYSTEMS IN ACTION


$$\begin{bmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 2 & 2 & 2 & 2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 2 & 1 & 1 \\ 0 & 2 & 0 & 0 & 1 & 1 \\ 0 & 2 & 1 & 1 & 0 & 2 \\ 0 & 2 & 1 & 1 & 0 & 0 \end{bmatrix}$$



THRESHOLD

$$\begin{bmatrix} 0 & 2 & 2 & 2 & 2 & 2 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 2 & 1 & 1 \\ 0 & 2 & 0 & 0 & 1 & 1 \\ 0 & 2 & 1 & 1 & 0 & 2 \\ 0 & 2 & 1 & 1 & 0 & 0 \end{bmatrix}$$


s/T > 90%


$$\begin{bmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$

GCM VISUALIZED



$$\begin{bmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}$$



OVERVIEW

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

NEXT STEPS

- Implement *Scooby Systems* in Hadoop
 - Scalable
- Threshold-based determination of edges
- Proposing solutions / further analysis of *Mystery Machine* limitations
- Evaluation
 - DeathStarBench

The Overall MapReduce Word Count Process

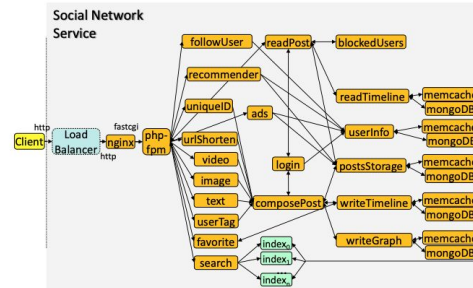
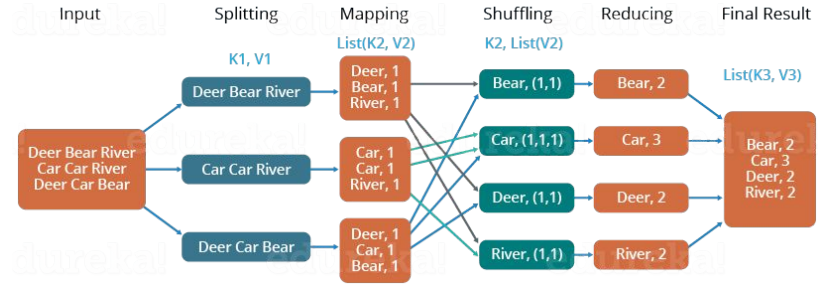


Figure 4. The architecture (microservices dependency graph) of Social Network.

ACKNOWLEDGEMENTS

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



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

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

Thanks to our **families**



Thank **YOU** for listening!

Any Questions?

