The Mystery Machine: End-to-end performance analysis of large-scale Internet services

Michael Chow
David Meisner, Jason Flinn, Daniel Peek,
Thomas F. Wenisch

UNIVERSITY OF MICHIGAN

facebook

Internet services are complex

Tasks

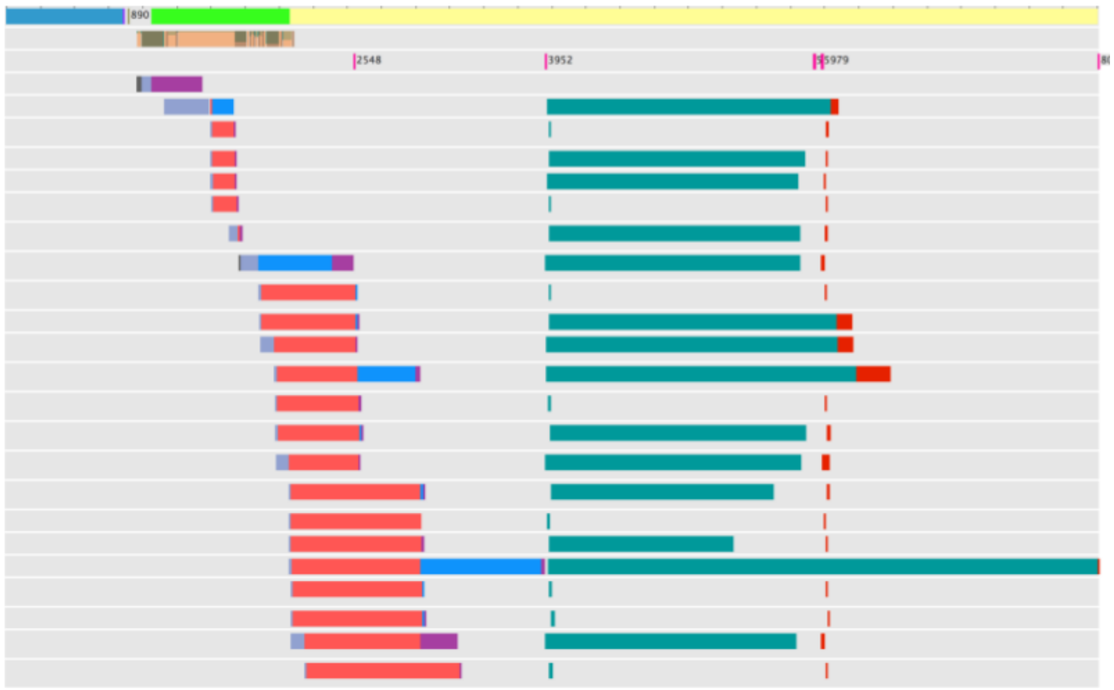

Time

Scale and heterogeneity make Internet services complex

Internet services are complex

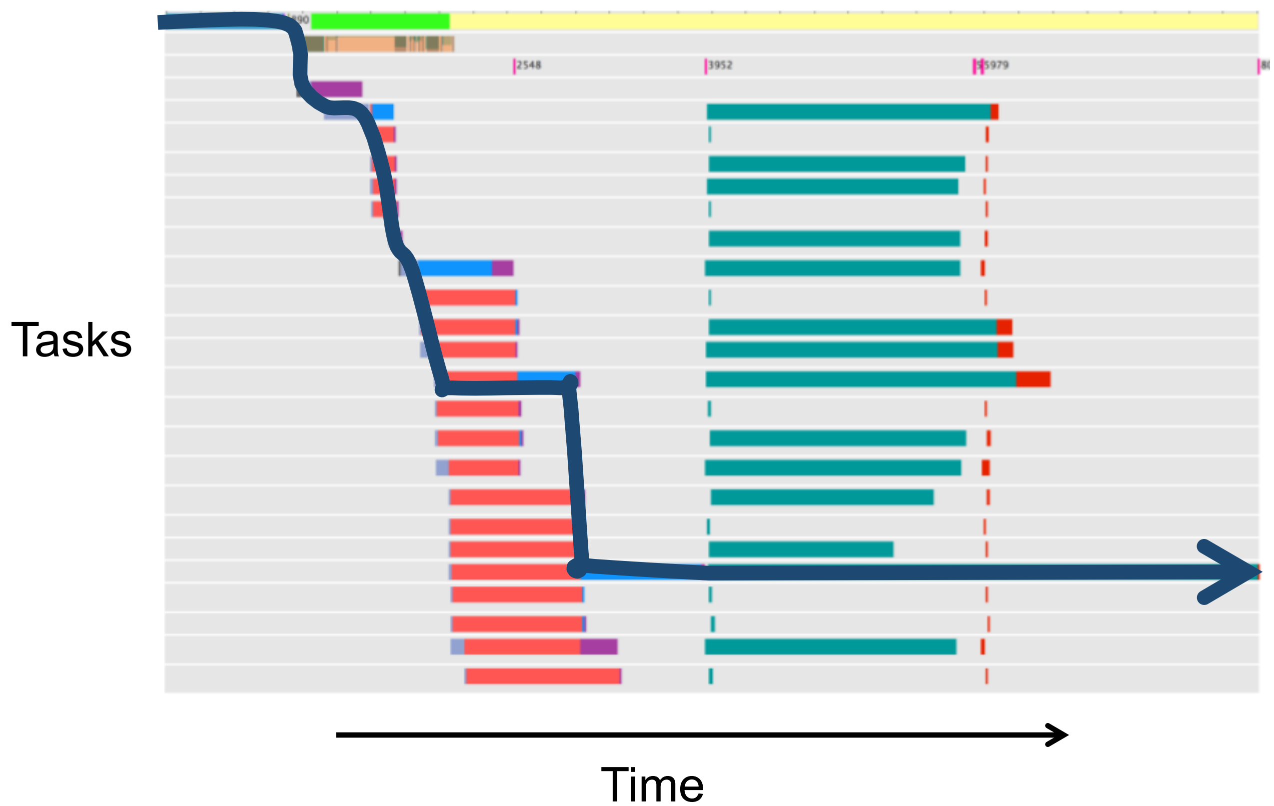

Scale and heterogeneity make Internet services complex

Analysis Pipeline

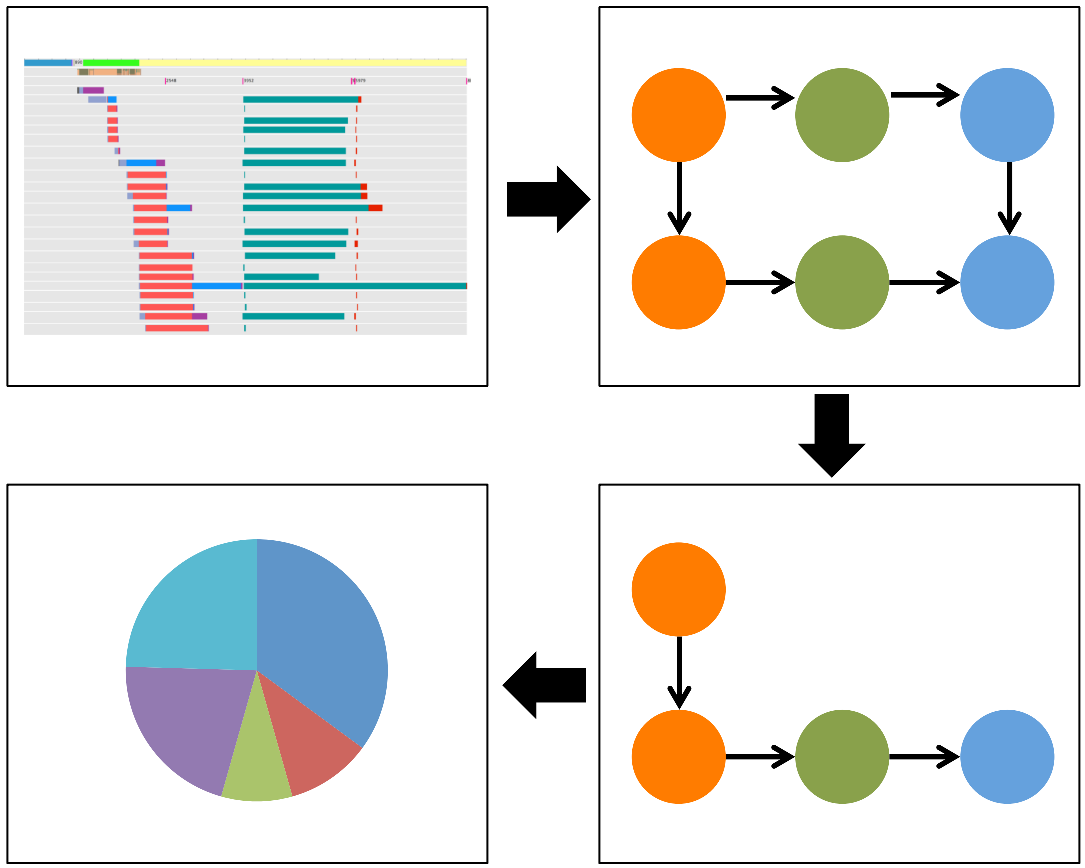

Step 1: Identify segments

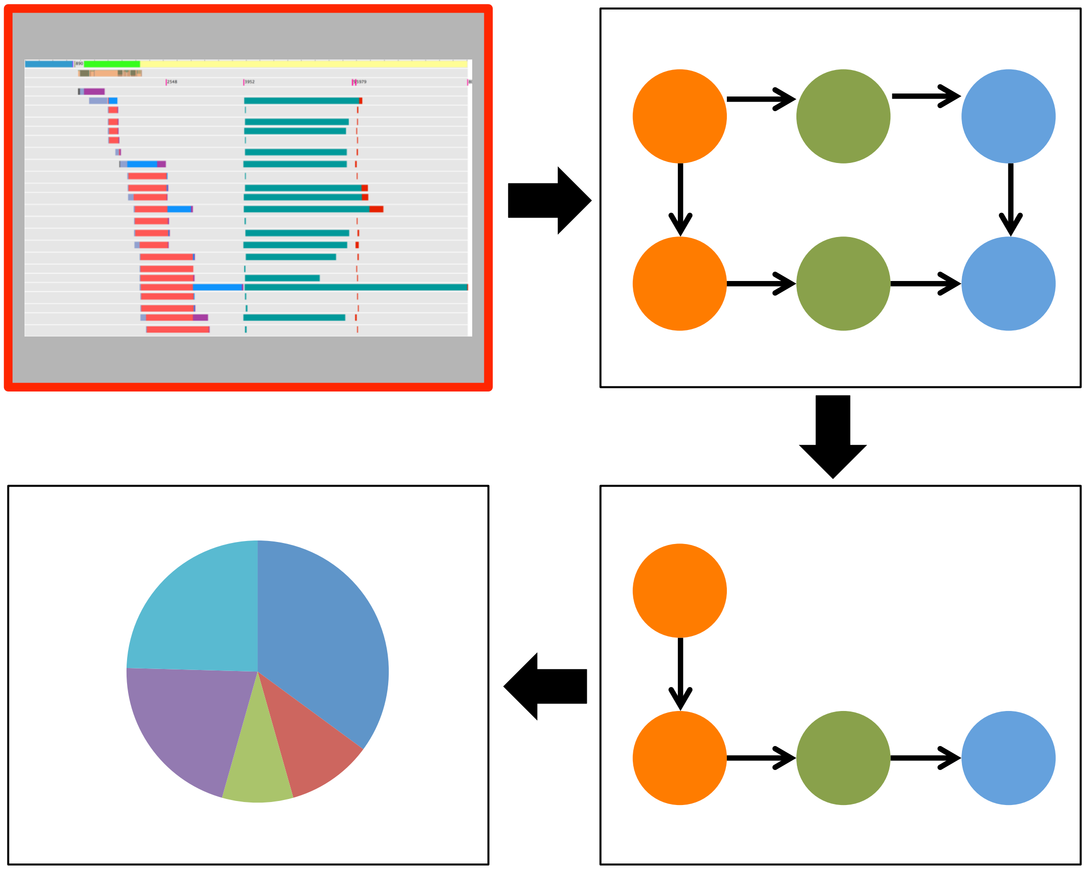

Step 2: Infer causal model

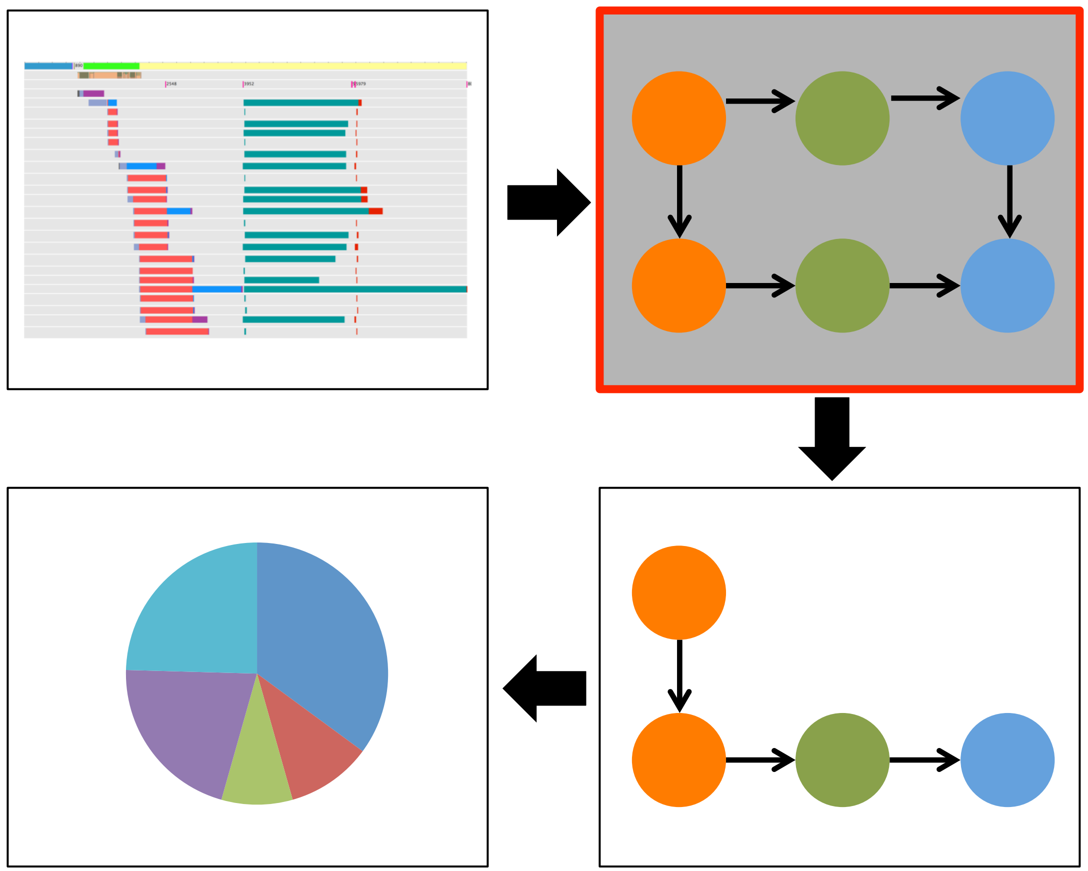

Step 3: Analyze individual requests

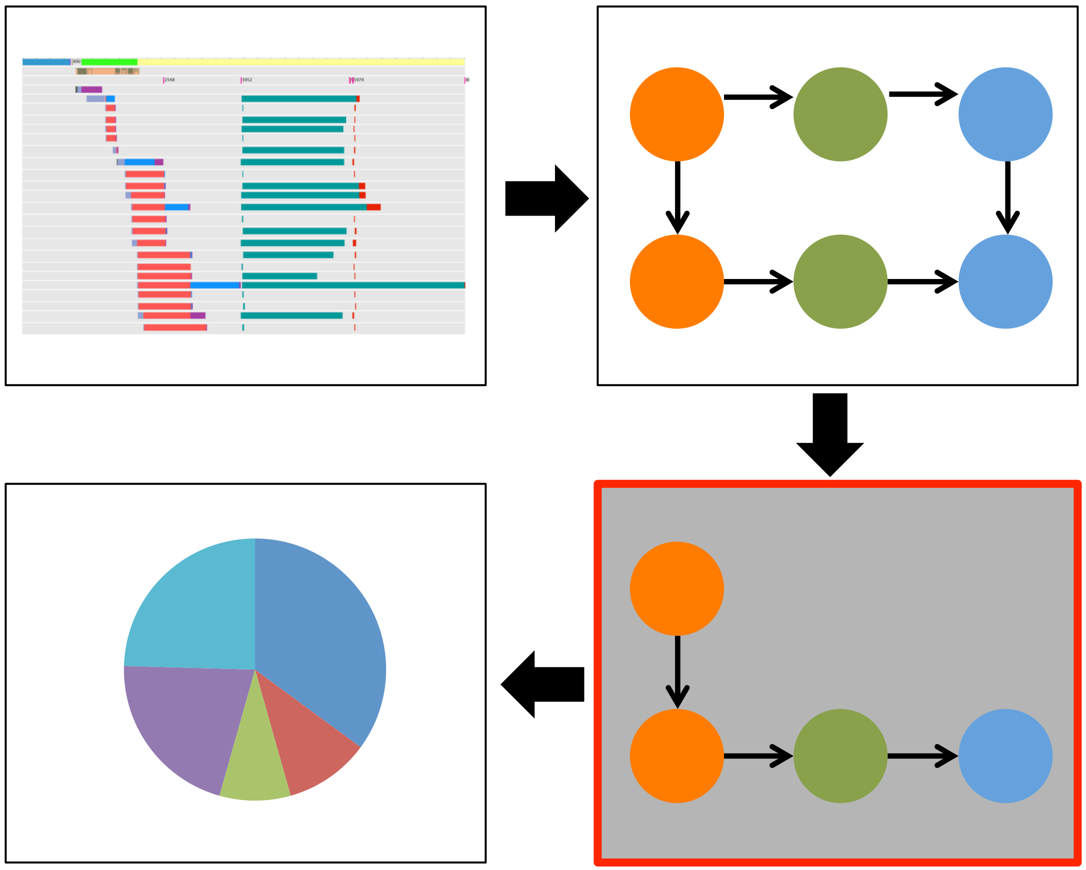

Step 4: Aggregate results

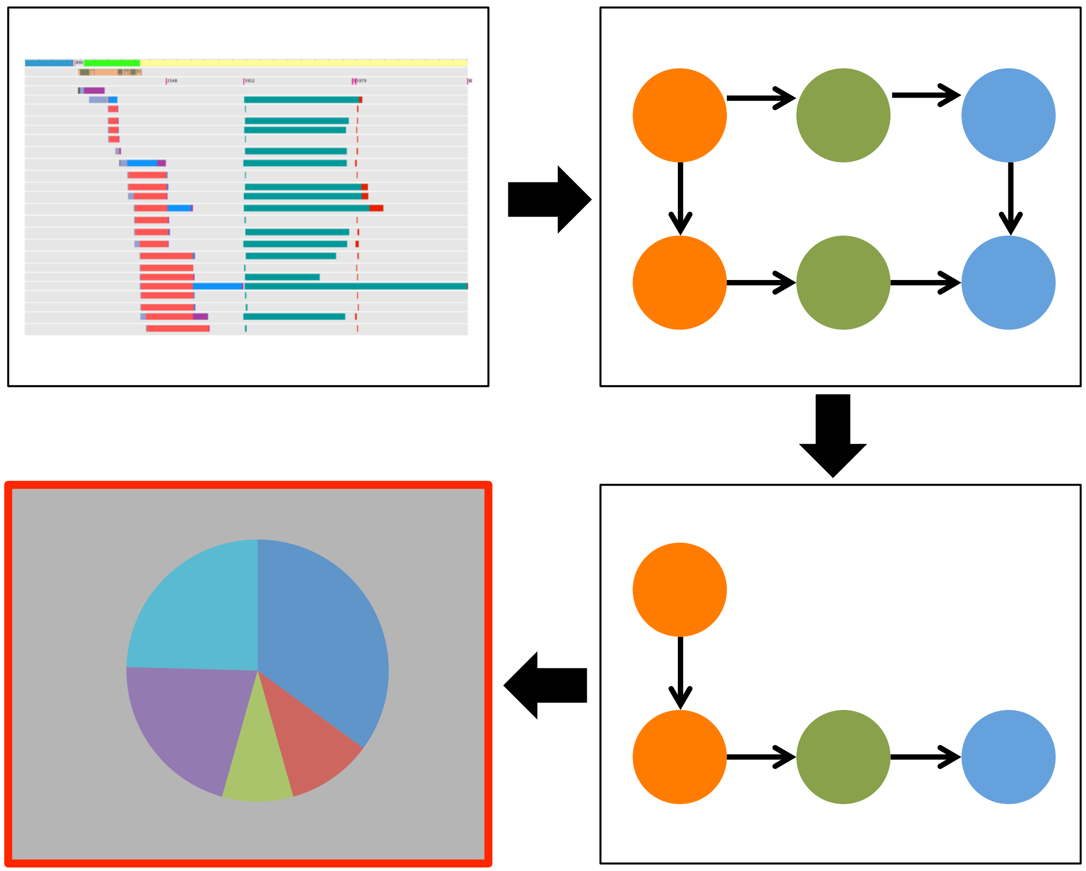

Challenges

- Previous methods derive a causal model
  - Instrument scheduler and communication
  - Build model through human knowledge

Need method that works at scale with heterogeneous components

Opportunities

- Component-level logging is ubiquitous

  Tremendous detail about a request's execution

- Handle a large number of requests

  Coverage of a large range of behaviors

The Mystery Machine

1) Infer causal model from large corpus of traces
   - Identify segments
   - Hypothesize all possible causal relationships
   - Reject hypotheses with observed counterexamples
2) Analysis
   - Critical path, slack, anomaly detection, what-if

Step 1: Identify segments

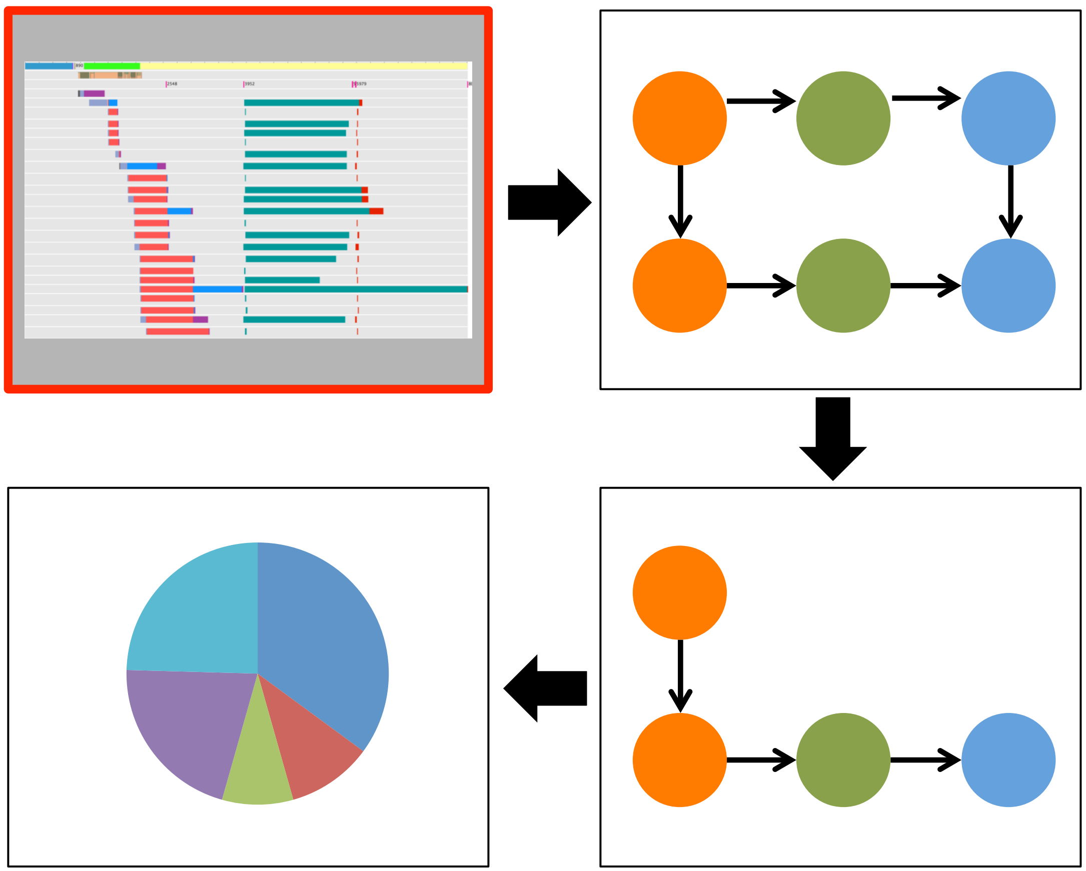

Define a minimal schema

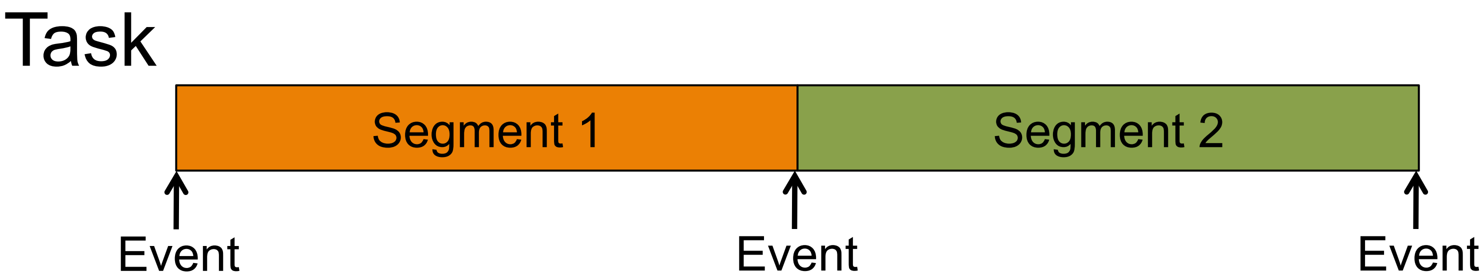

Define a minimal schema

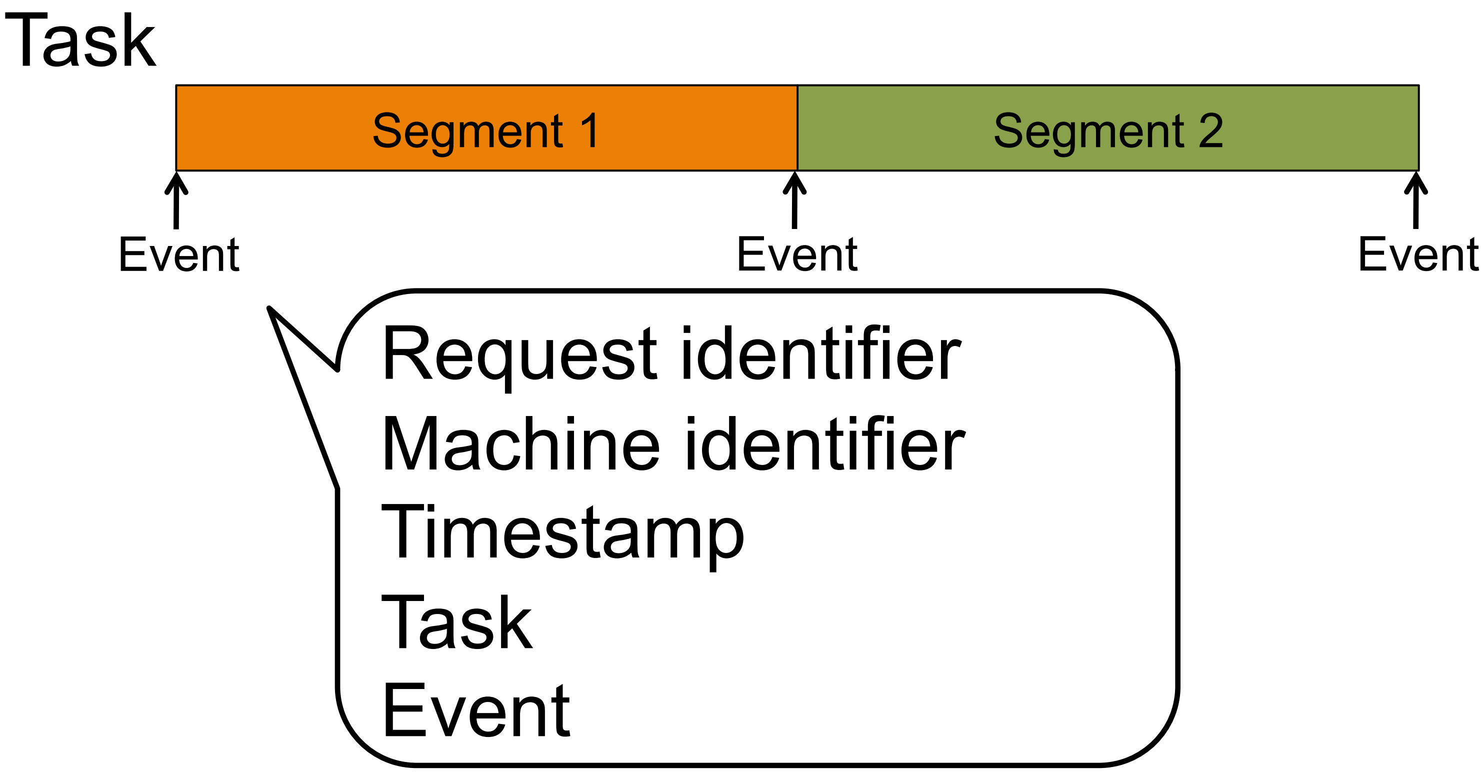

Aggregate existing logs using minimal schema

Step 2: Infer causal model

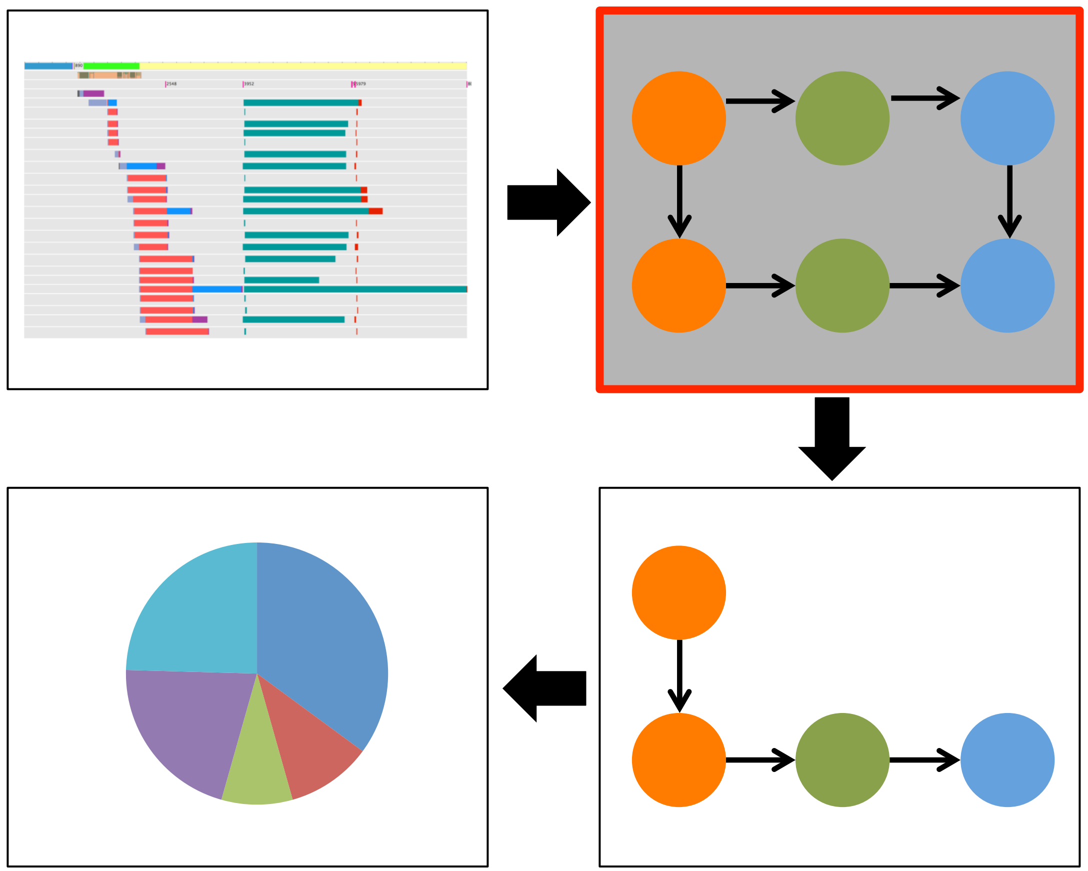

Types of causal relationships

Relationship	Counterexample
Happens-Before 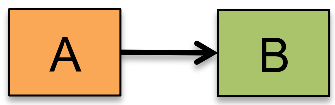	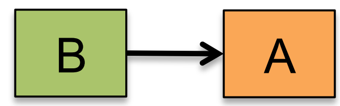
Mutual Exclusion 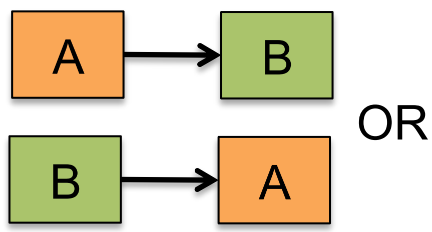	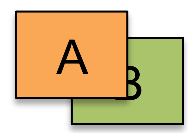
Pipeline 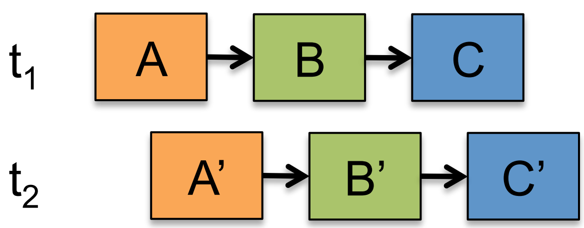	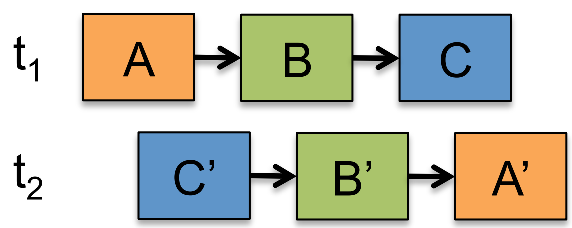

Producing causal model

Causal Model

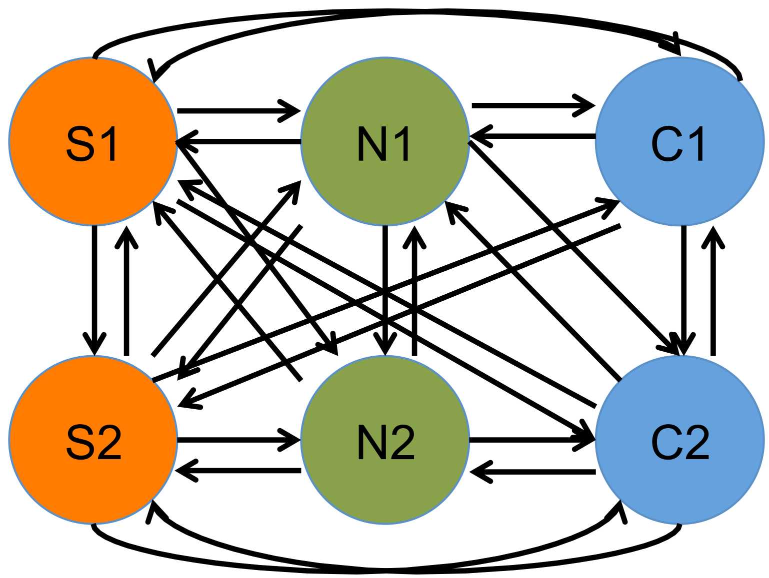

Producing causal model

Causal Model

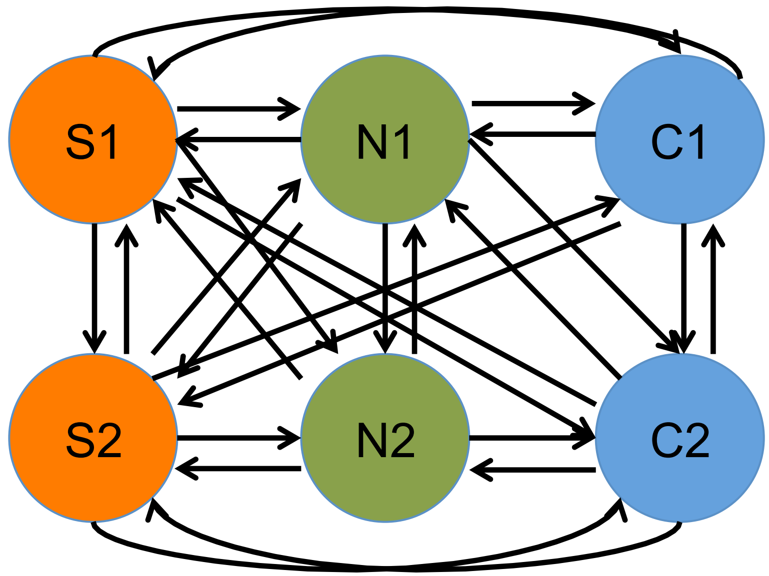

Trace 1

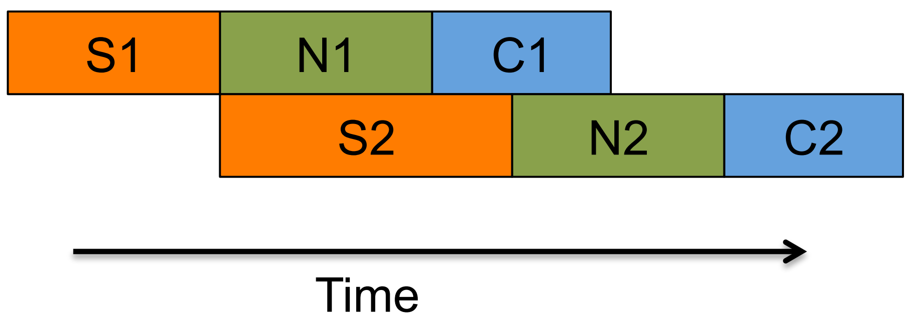

Producing causal model

Causal Model

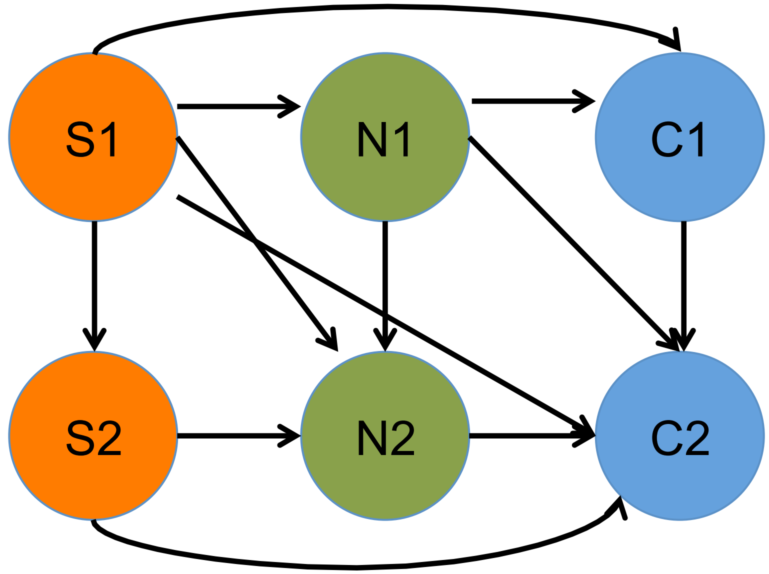

Trace 1

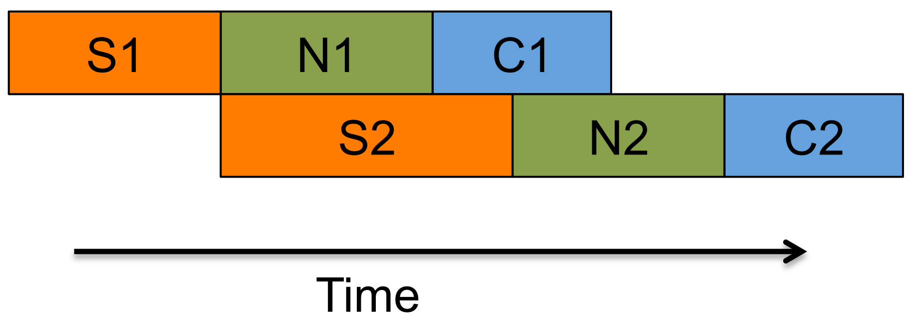

Producing causal model

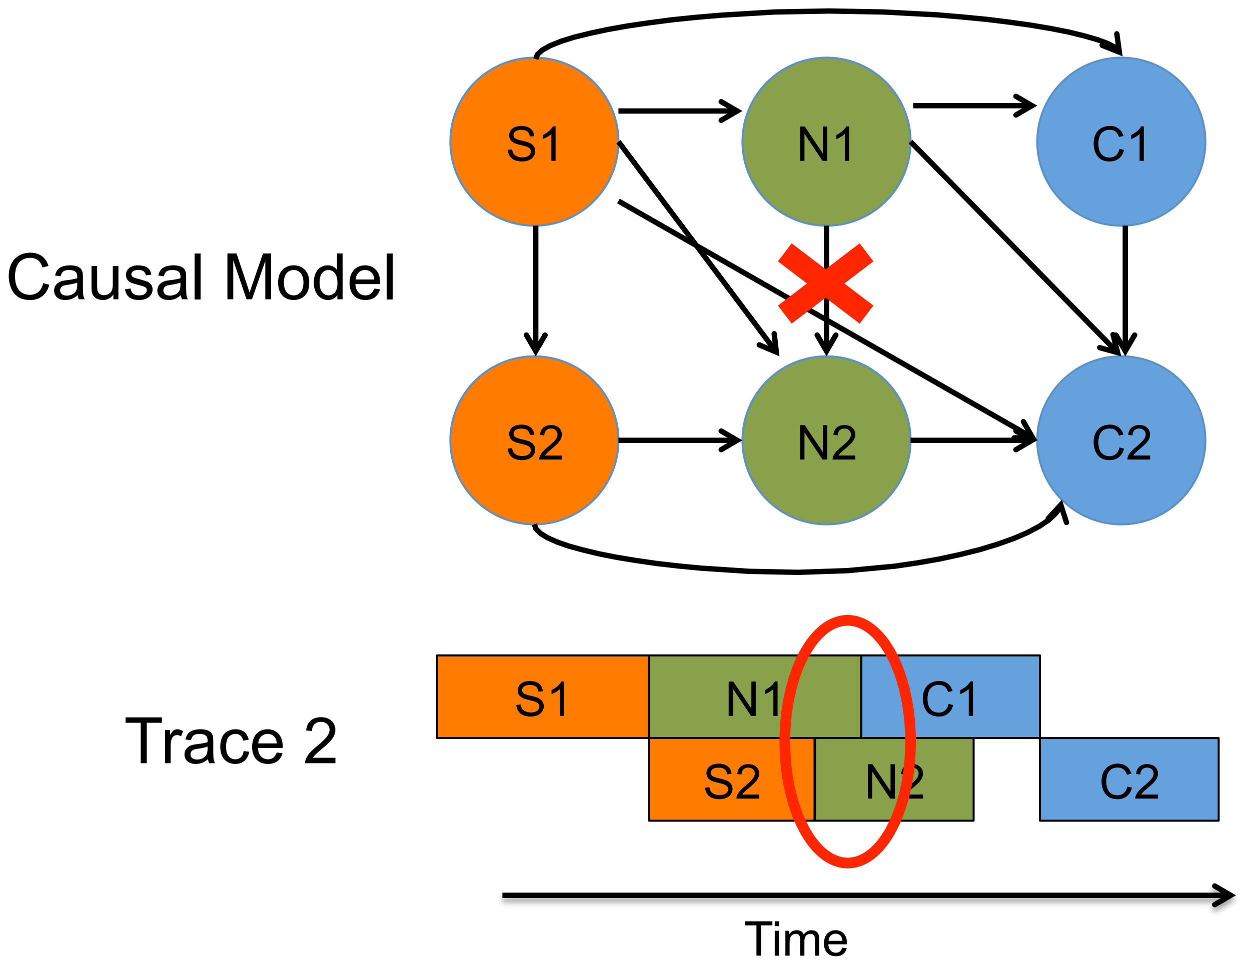

Producing causal model

Causal Model

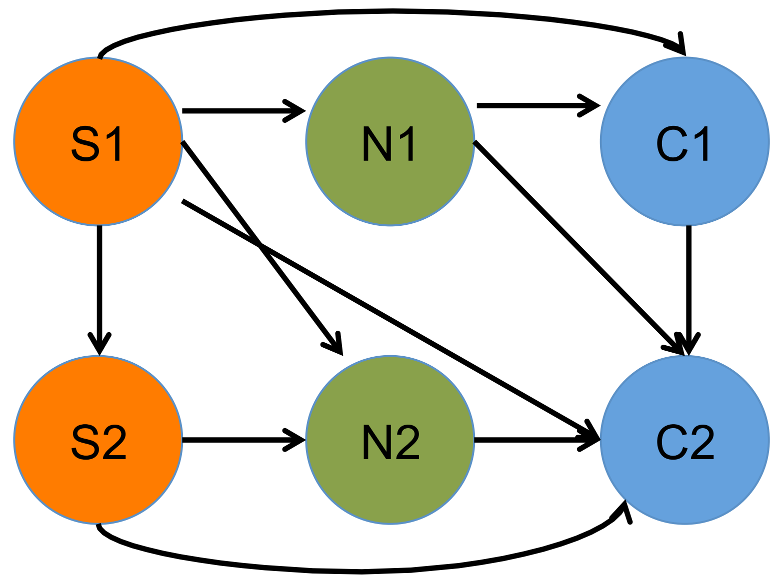

Trace 2

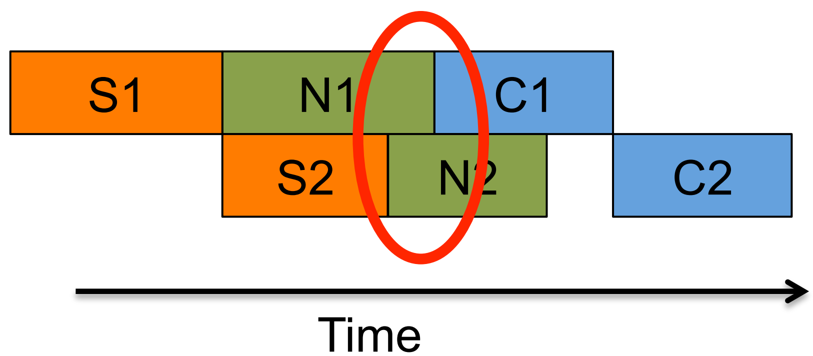

Producing causal model

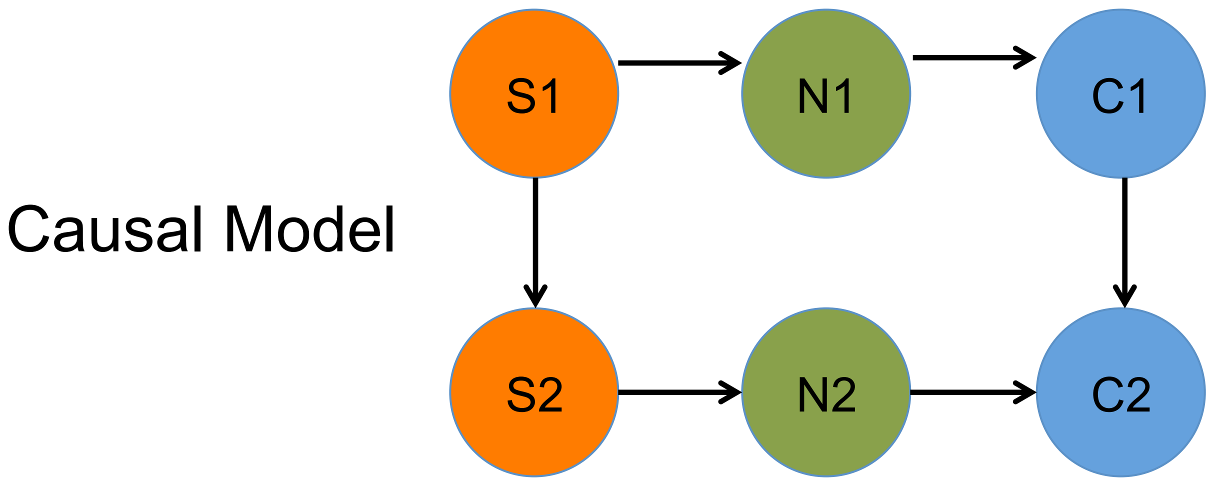

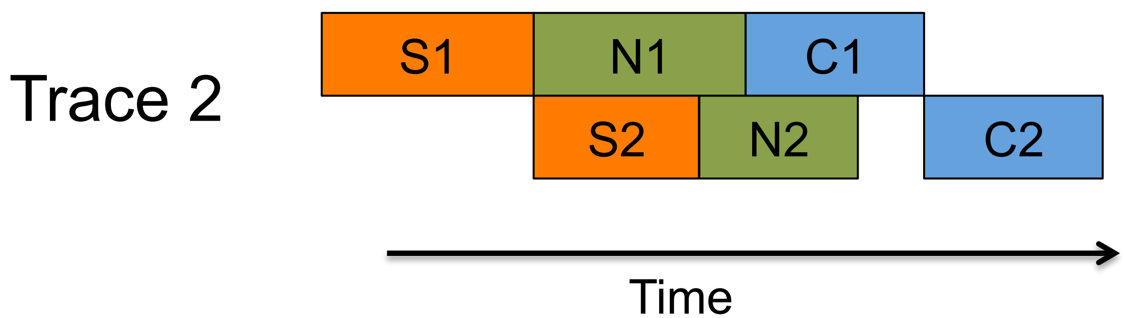

Step 3: Analyze individual requests

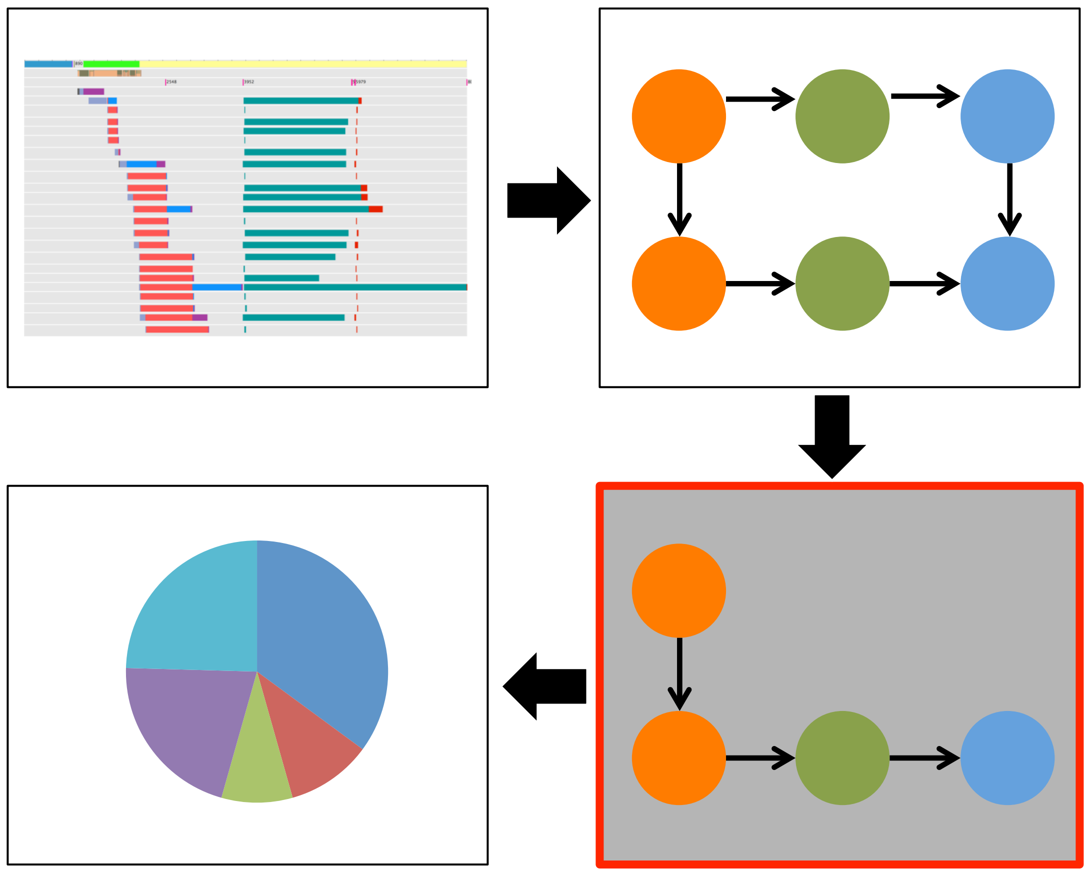

Critical path using causal model

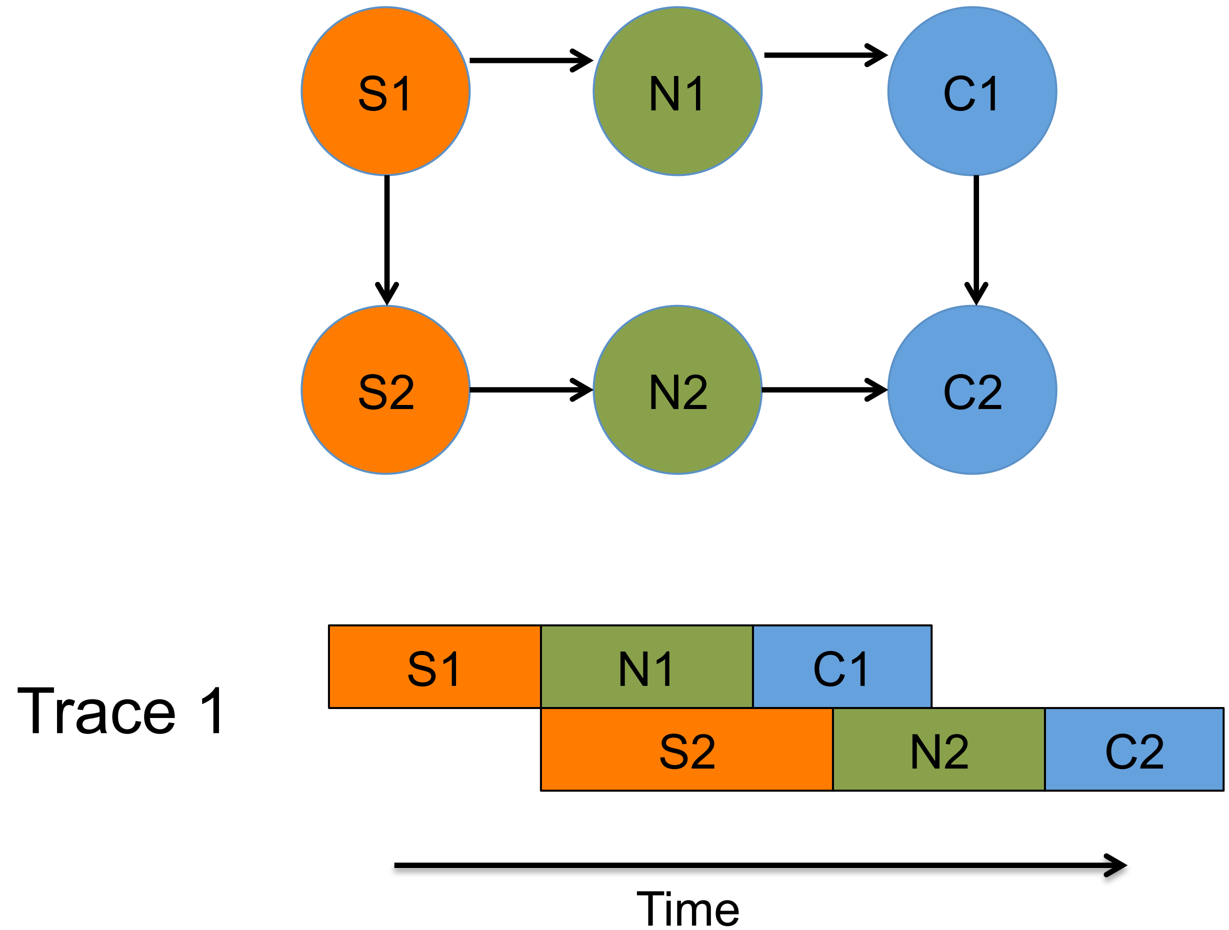

Critical path using causal model

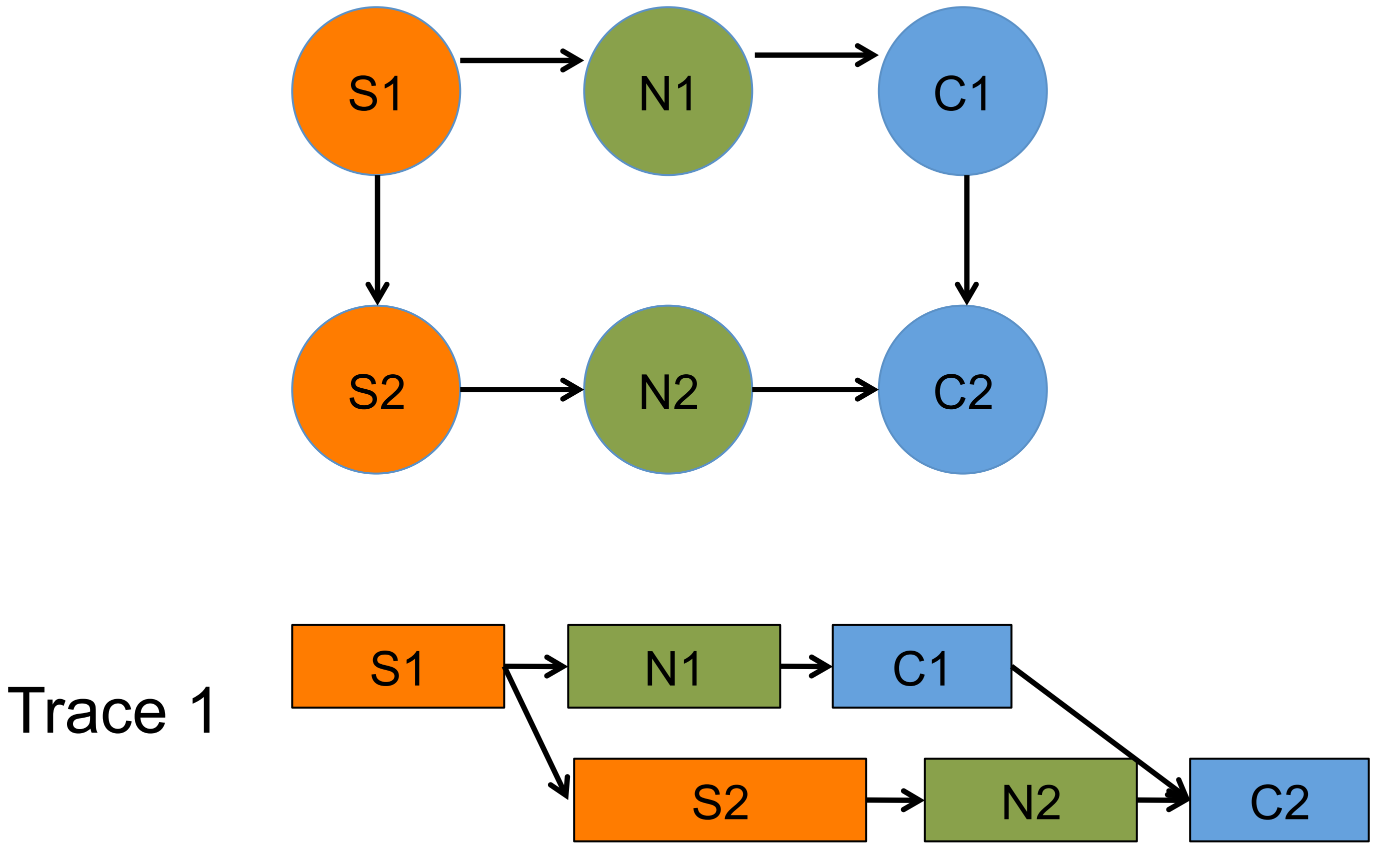

Critical path using causal model

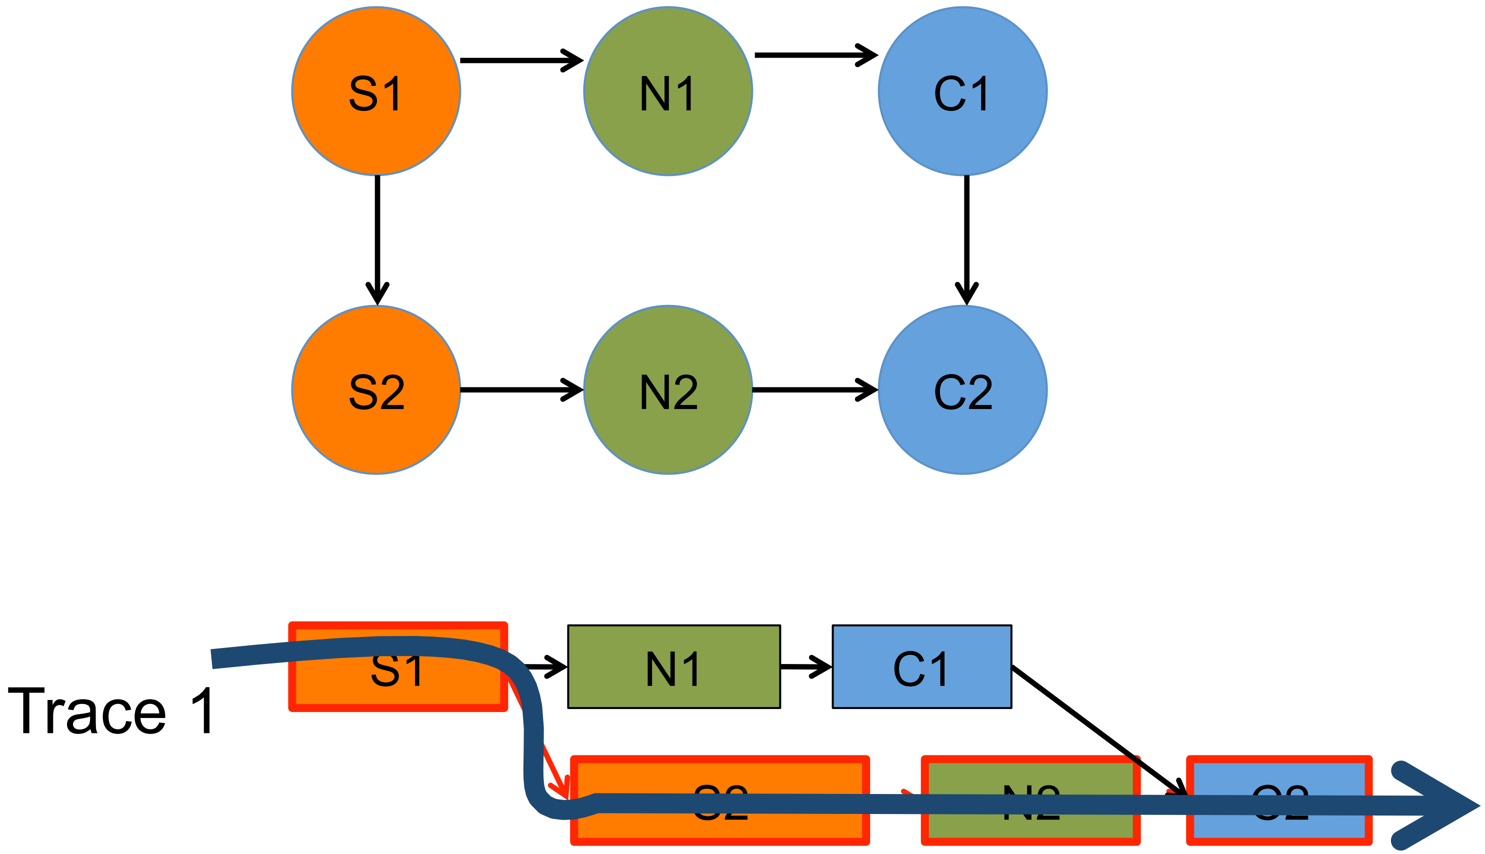

Step 4: Aggregate results

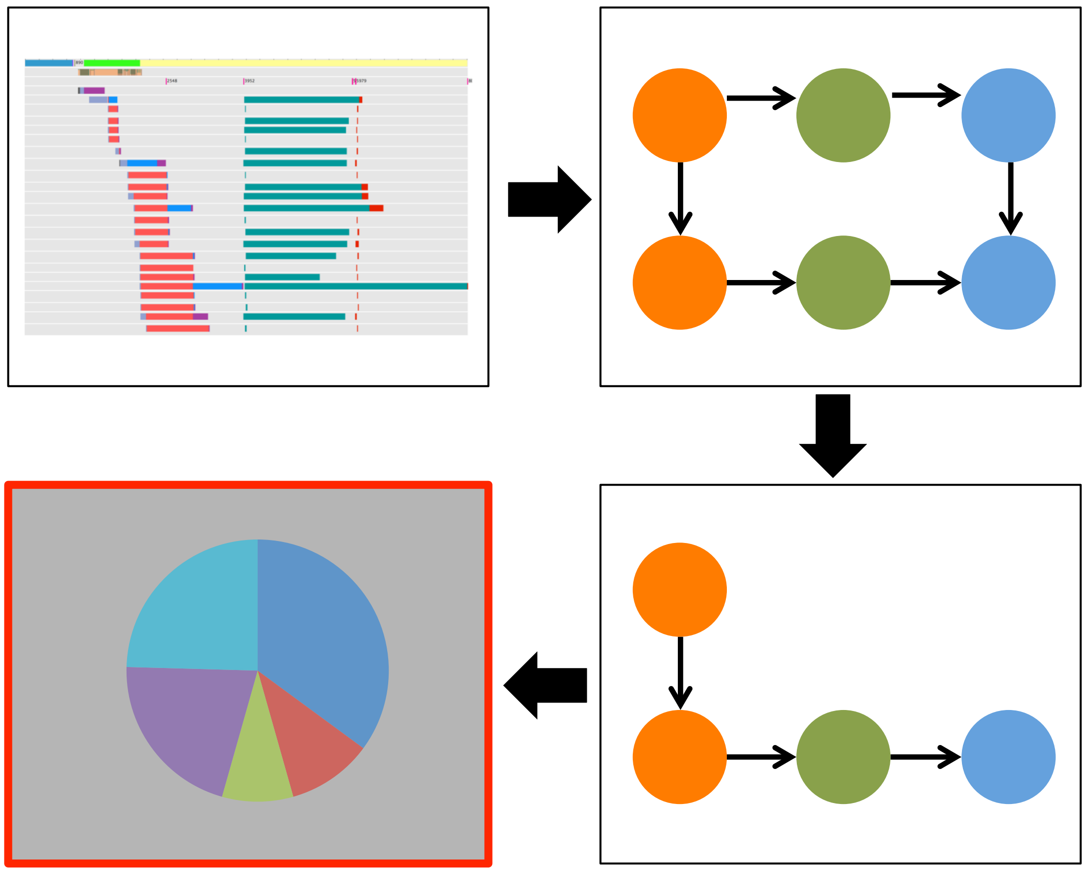

Inaccuracies of Naïve Aggregation

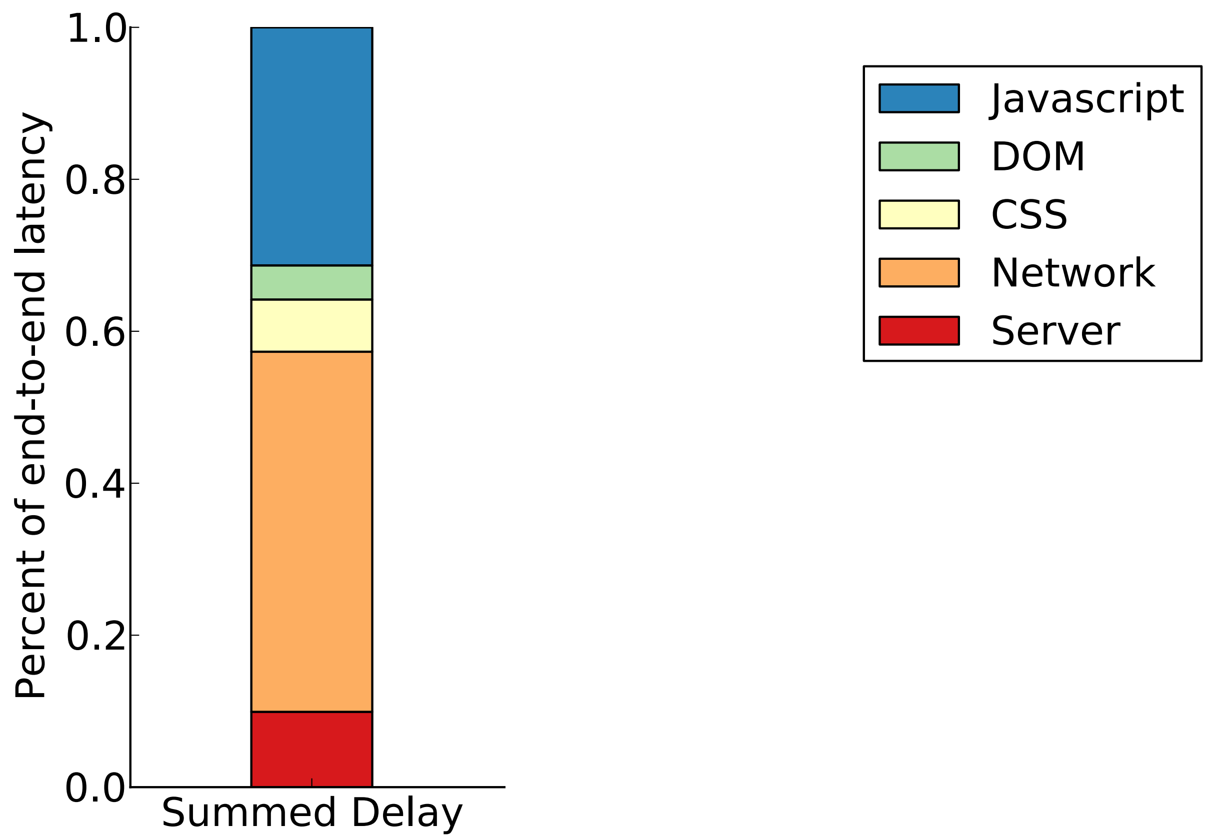

Inaccuracies of Naïve Aggregation

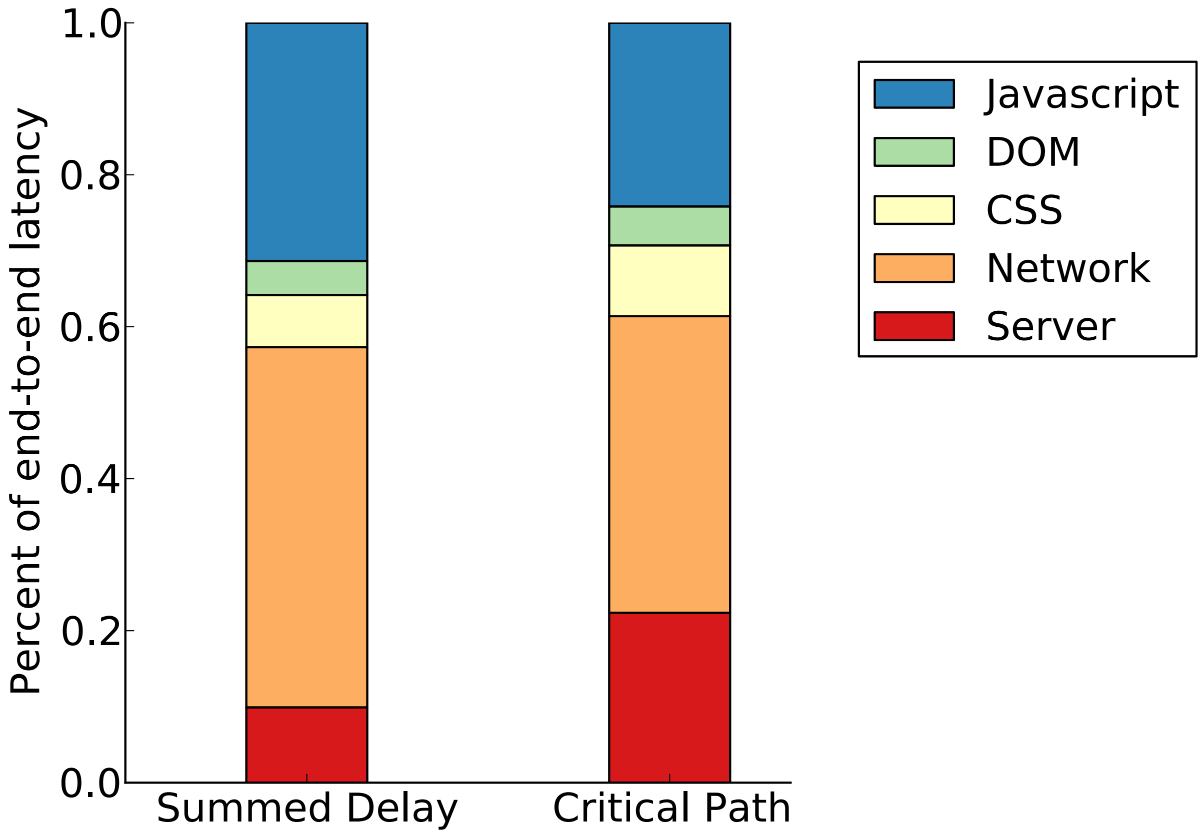

Inaccuracies of Naïve Aggregation

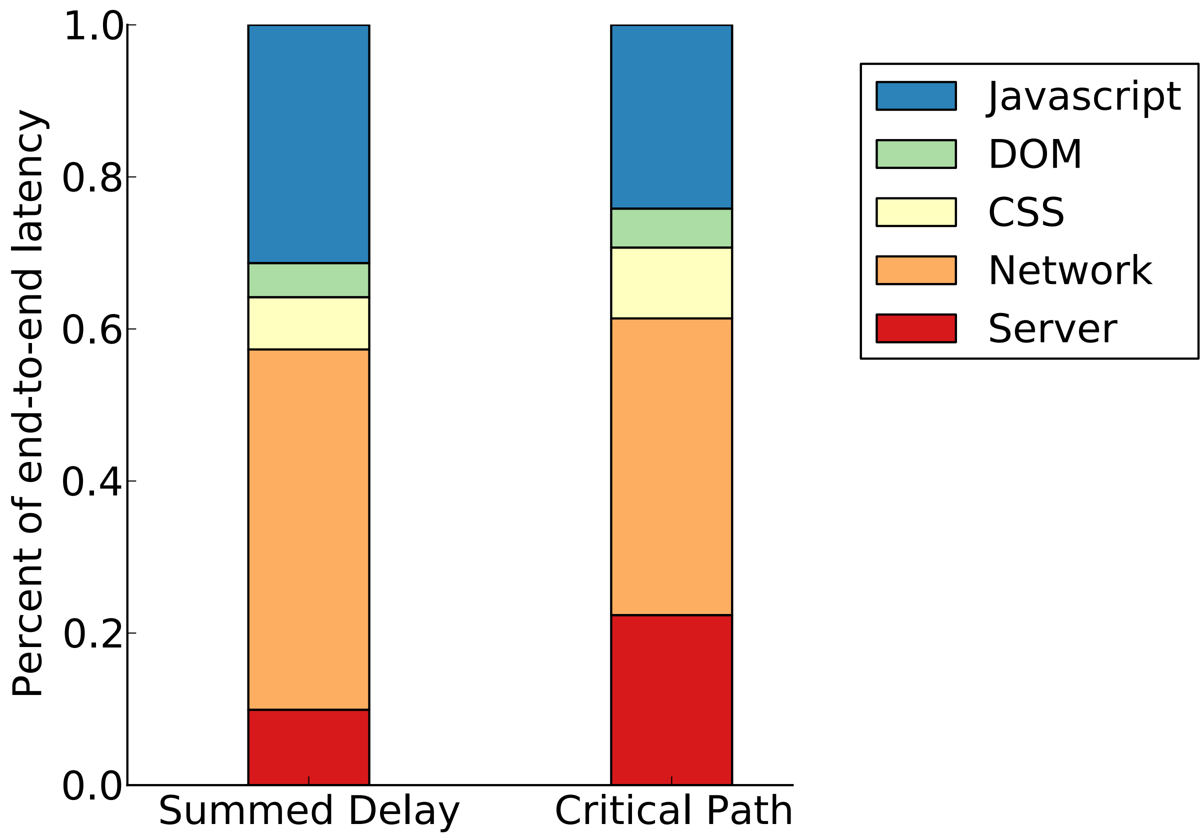

Need a causal model to correctly understand latency

High variance in critical path

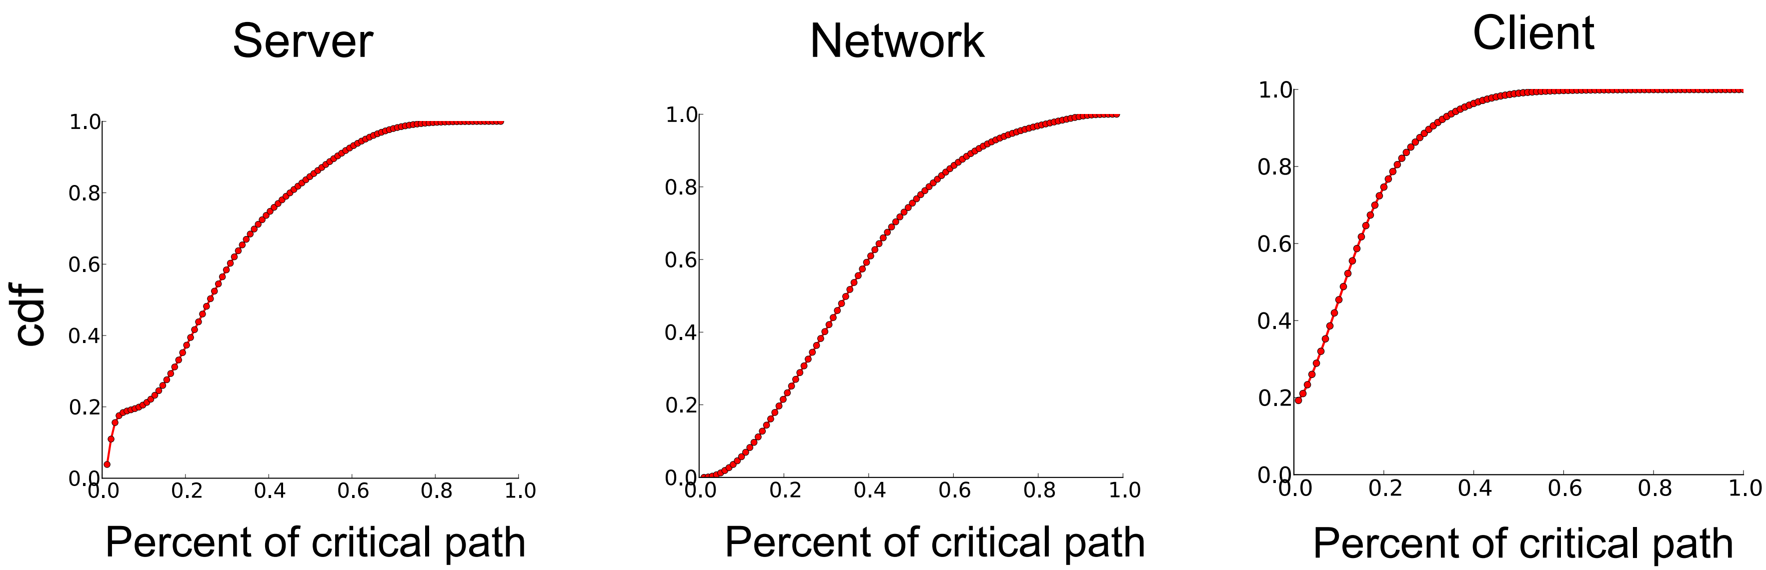

- Breakdown in critical path shifts drastically
  - Server, network, or client can dominate latency

High variance in critical path

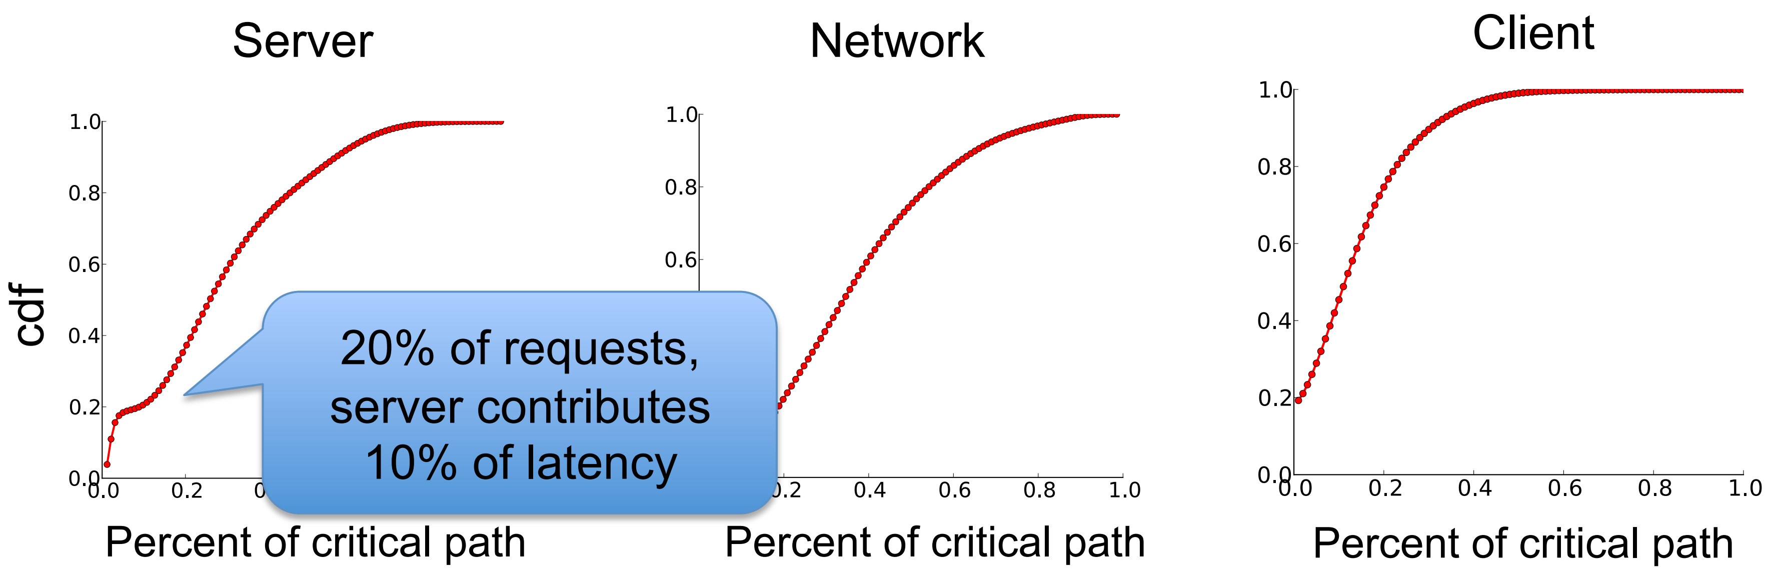

- Breakdown in critical path shifts drastically
  - Server, network, or client can dominate latency

High variance in critical path

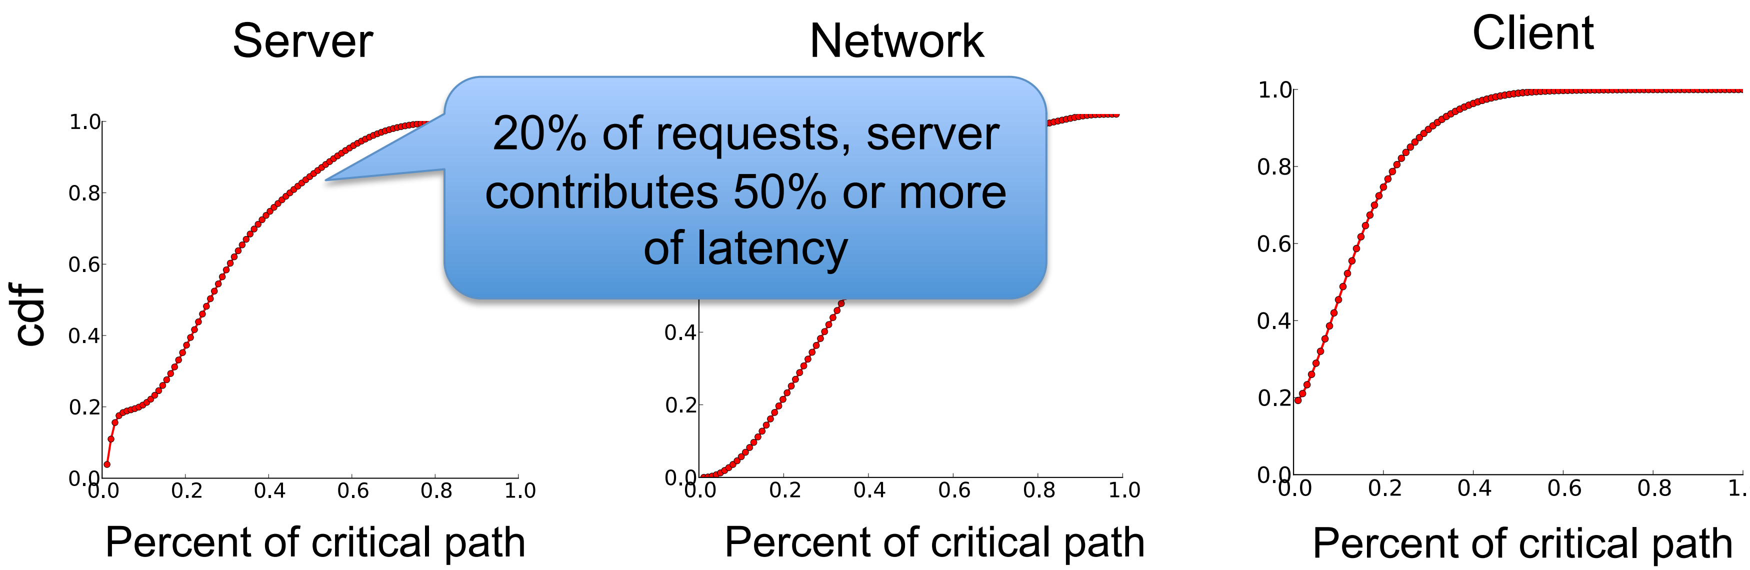

- Breakdown in critical path shifts drastically
  - Server, network, or client can dominate latency

Diverse clients and networks

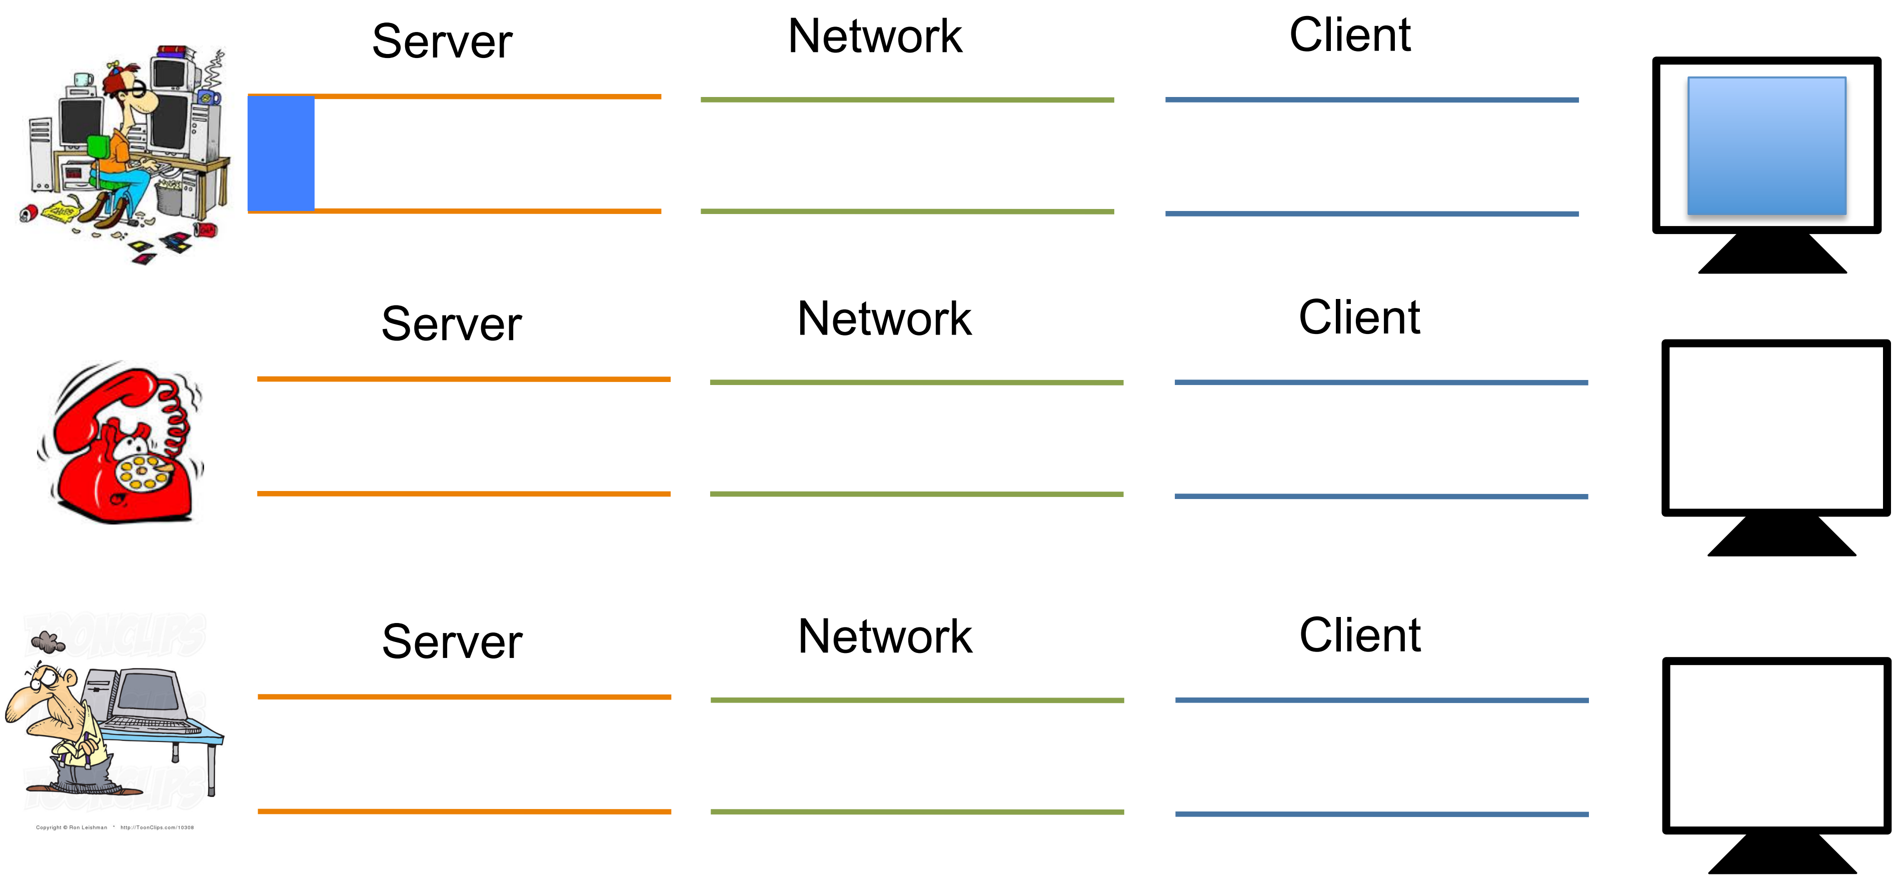

Diverse clients and networks

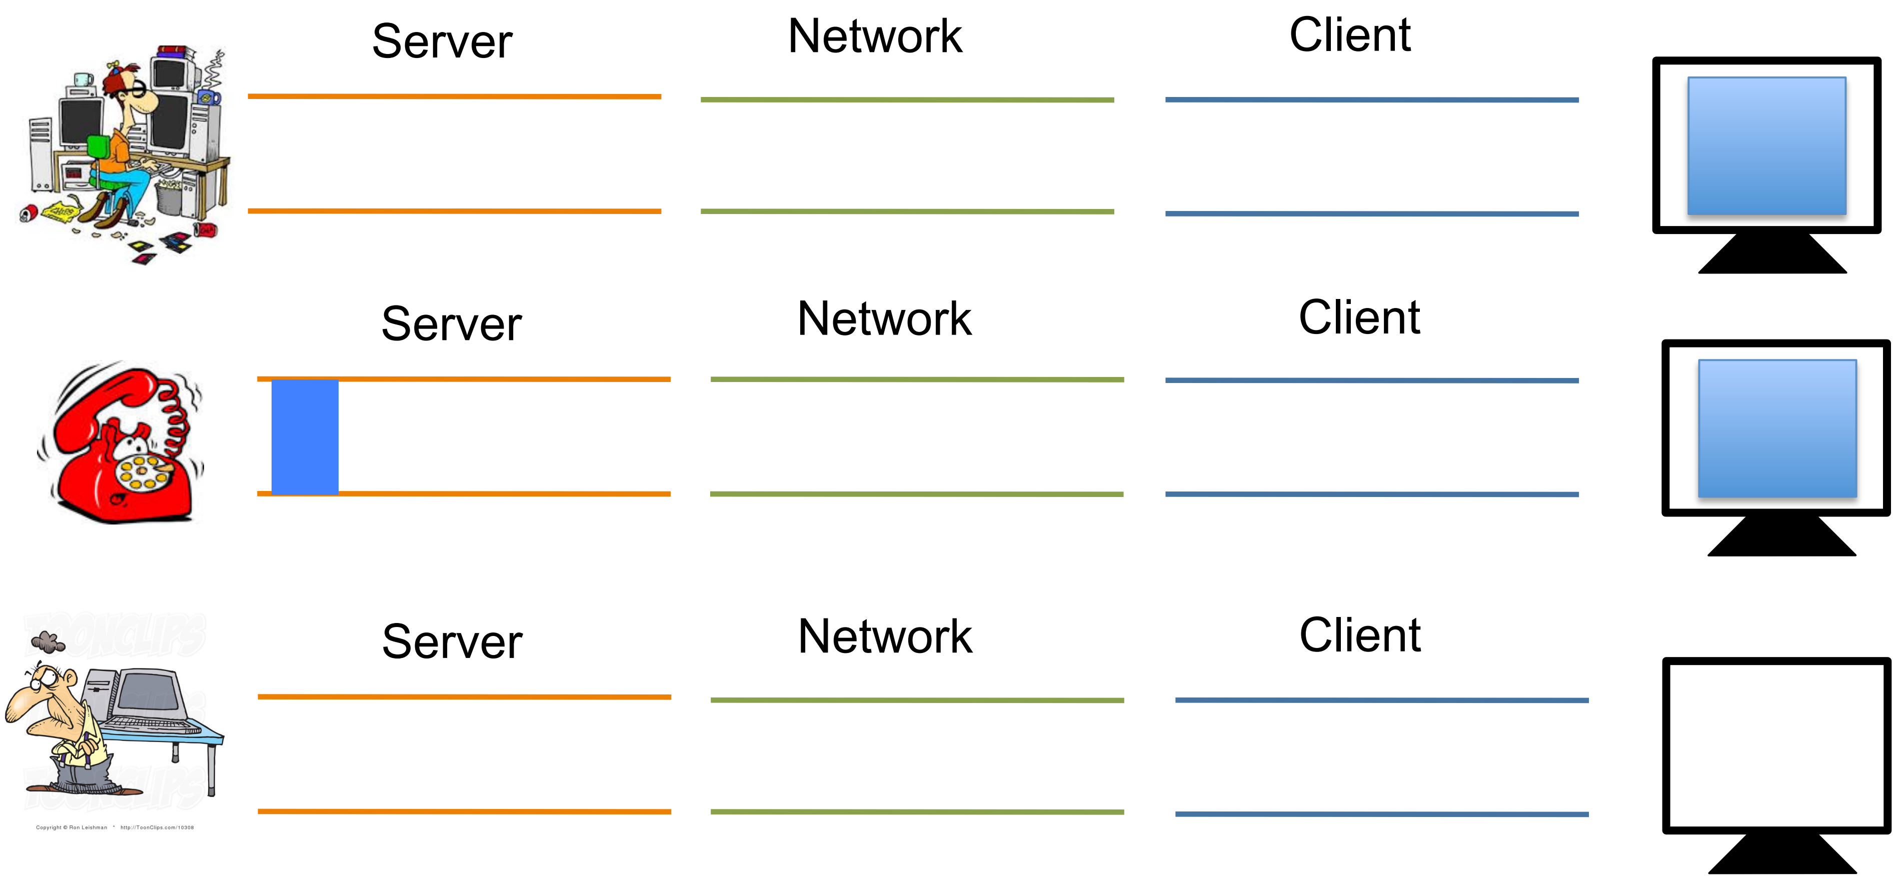

Diverse clients and networks

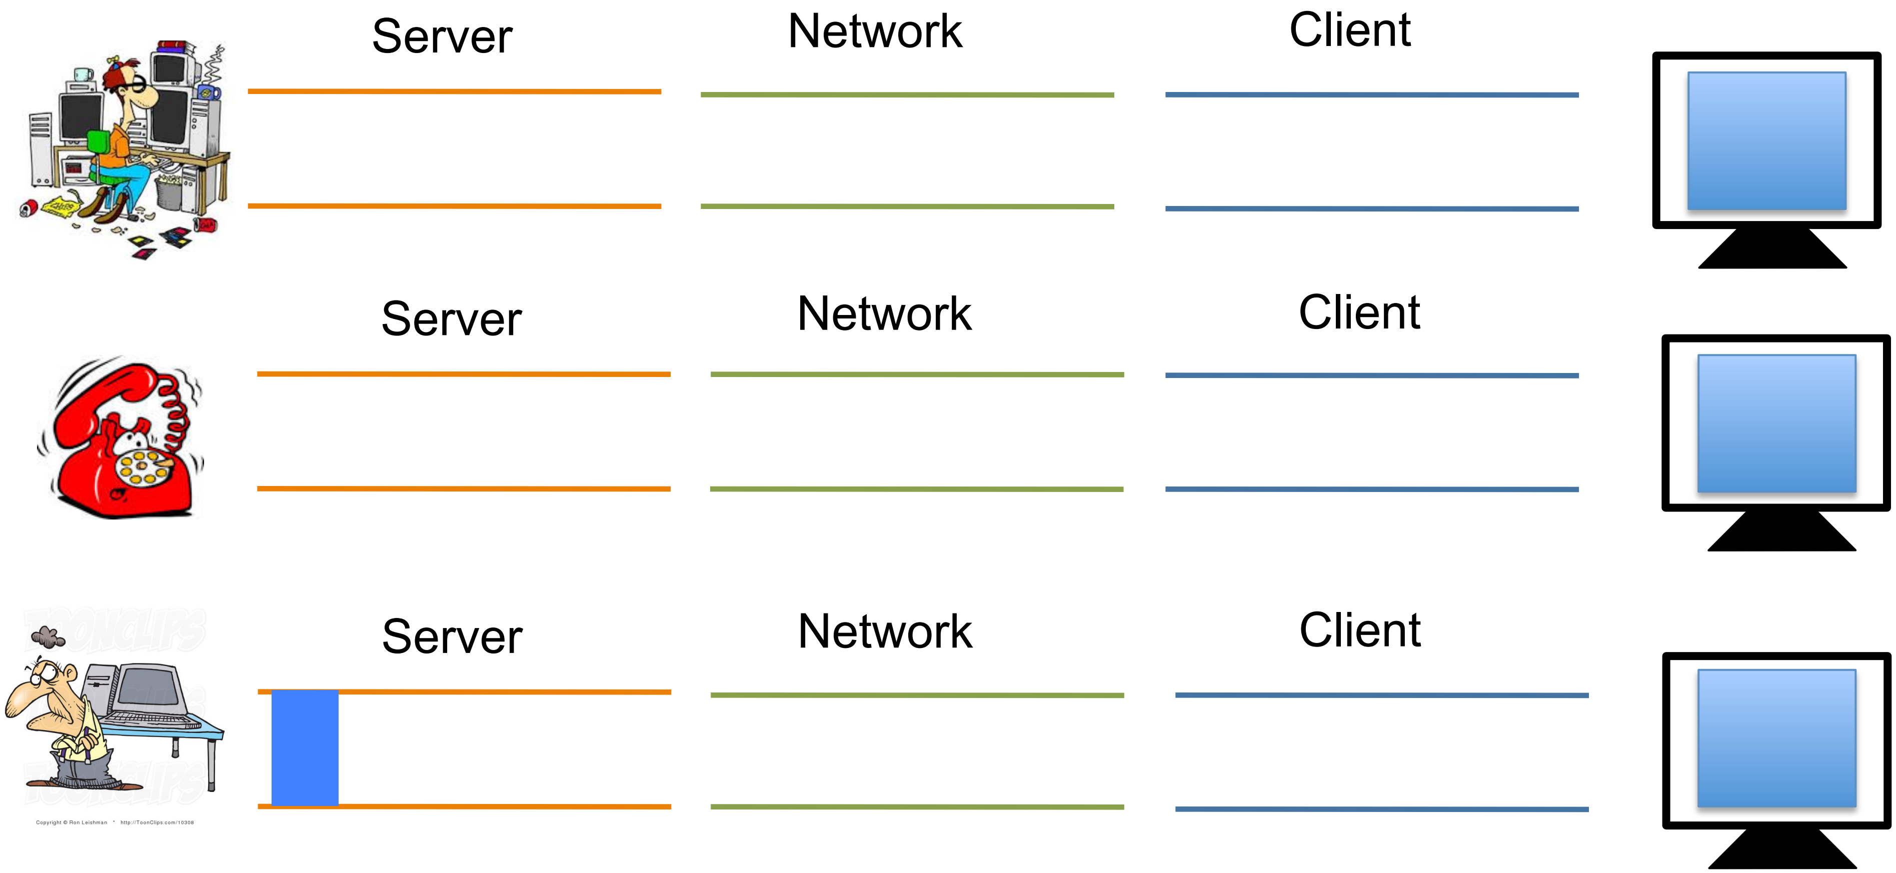

Differentiated service

Slack in server generation time
Produce data slower
End-to-end latency stays same

No slack in server generation time
Produce data faster
Decrease end-to-end latency

Deliver data when needed and reduce average response time

Additional analysis techniques

- Slack analysis

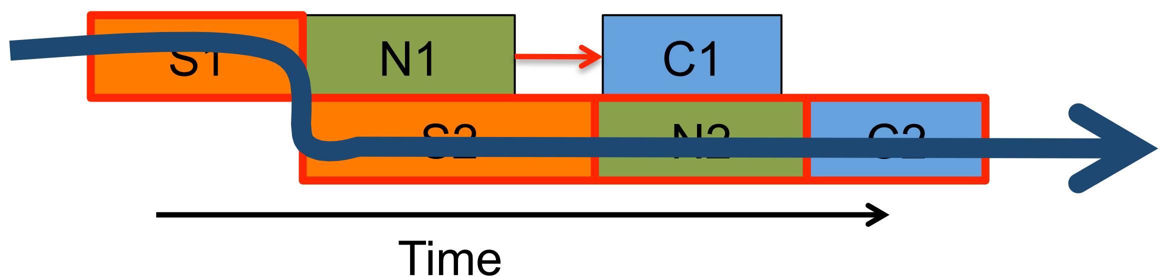

- What-if analysis
  - Use natural variation in large data set

What-if questions

- Does server generation time affect end-to-end latency?
- Can we predict which connections exhibit server slack?

Server slack analysis

Slack < 25ms

Slack > 2.5s

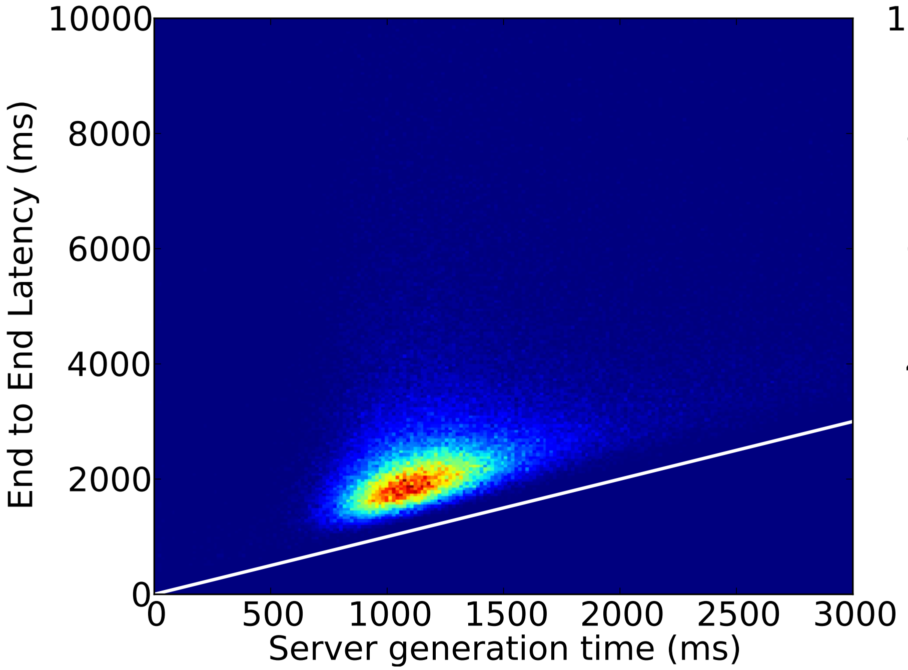

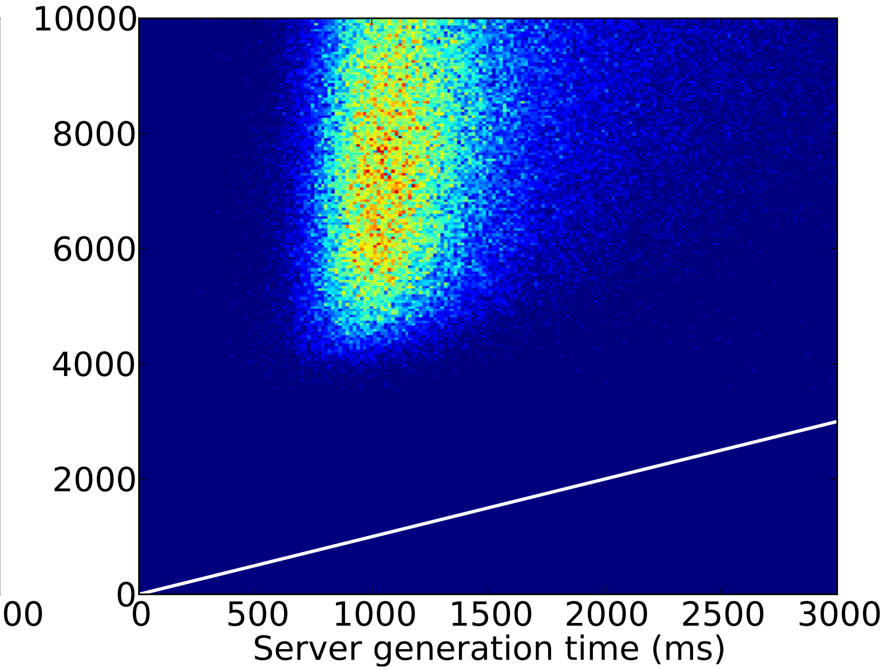

Server slack analysis

Slack < 25ms

Slack > 2.5s

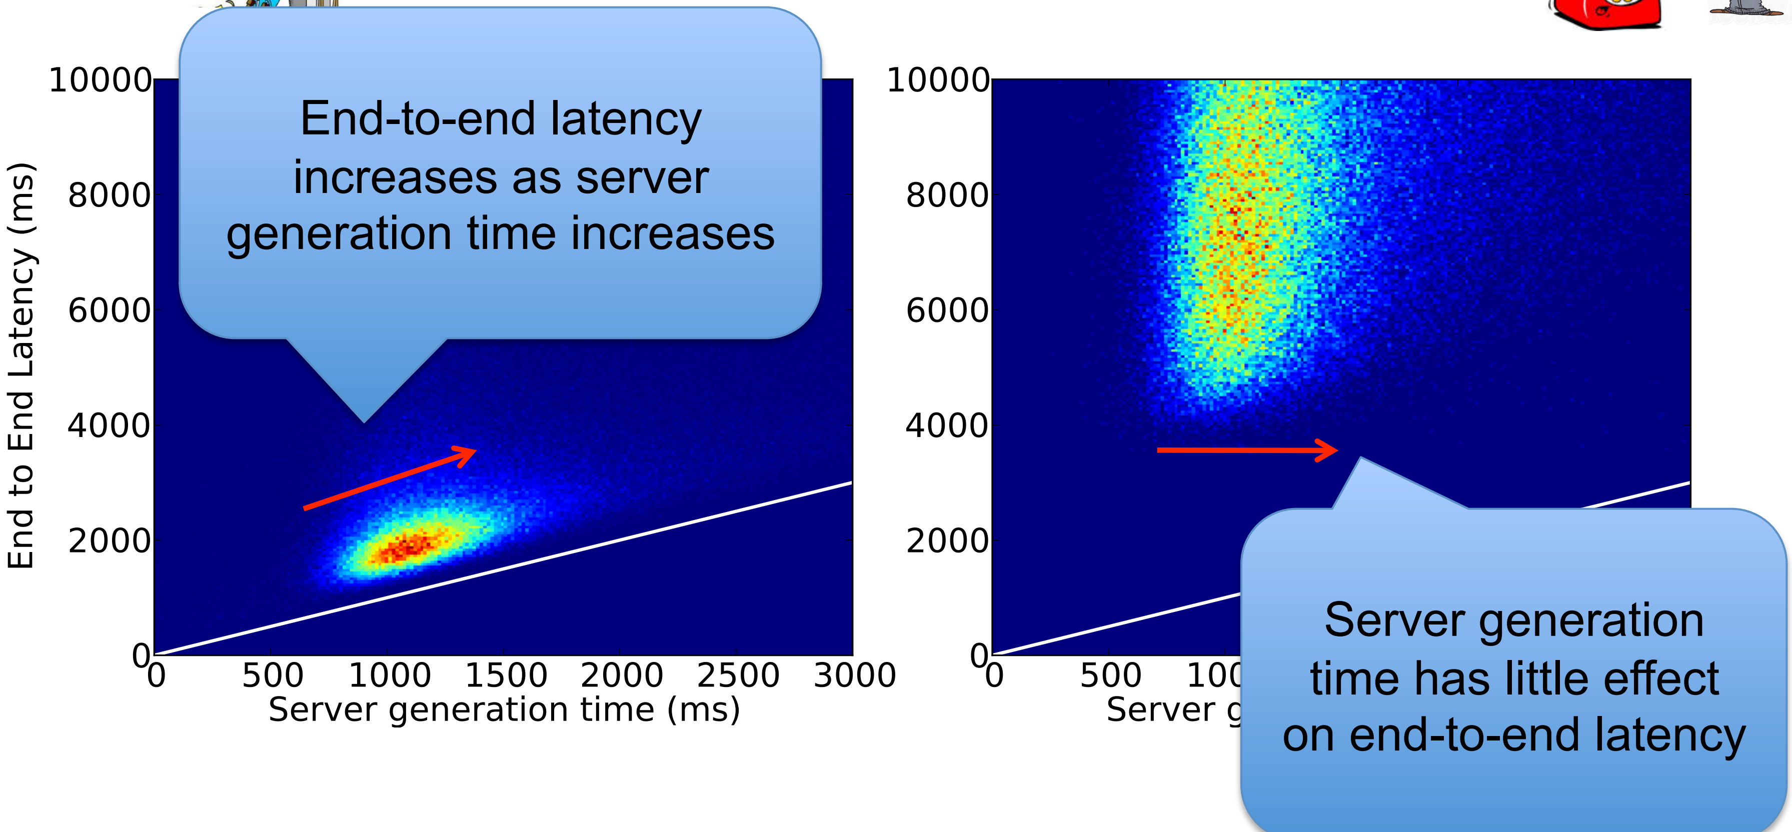

Predicting server slack

- Predict slack at the receipt of a request
- Past slack is representative of future slack

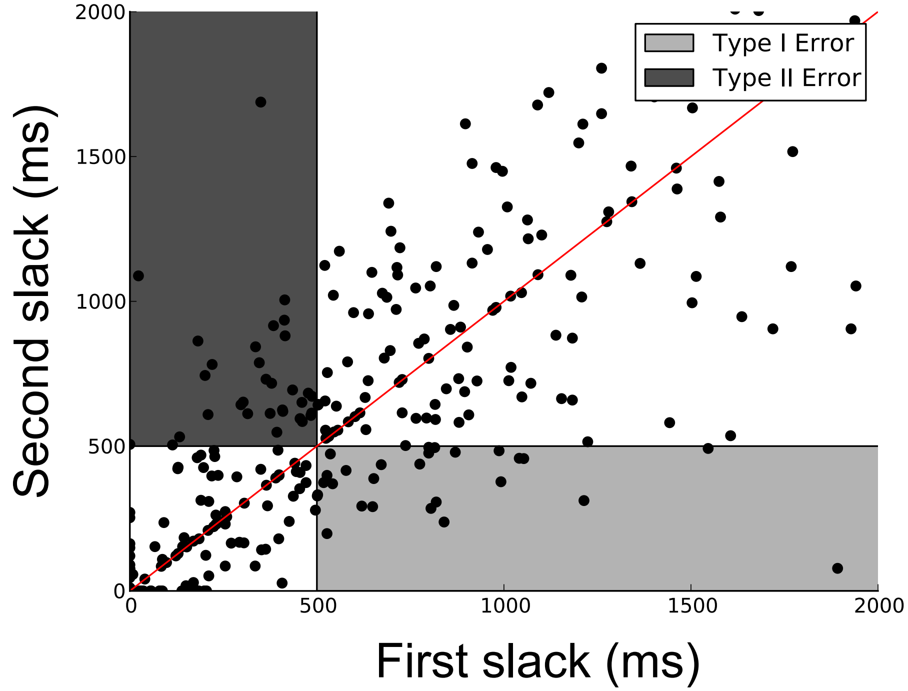

Predicting server slack

- Predict slack at the receipt of a request
- Past slack is representative of future slack

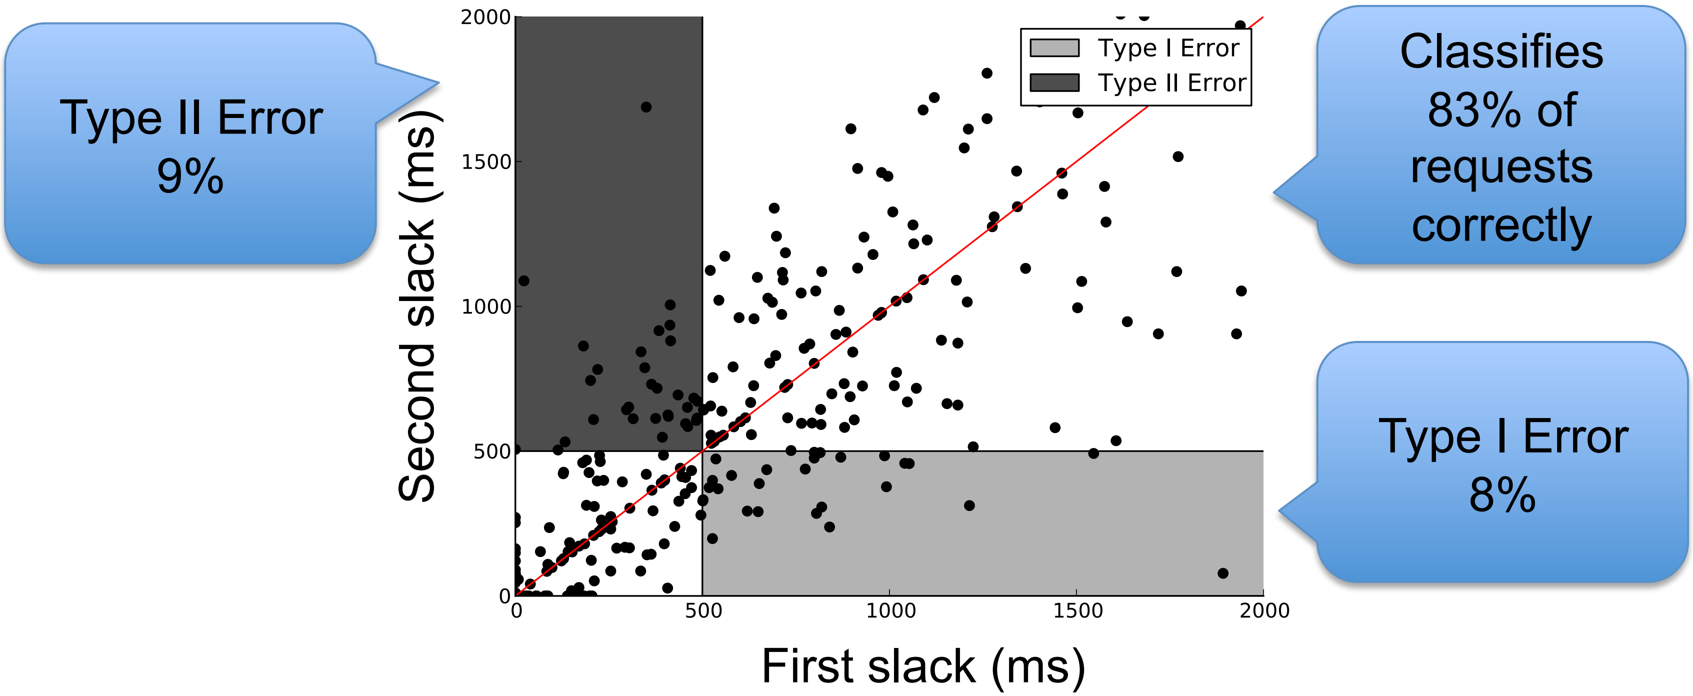

Conclusion

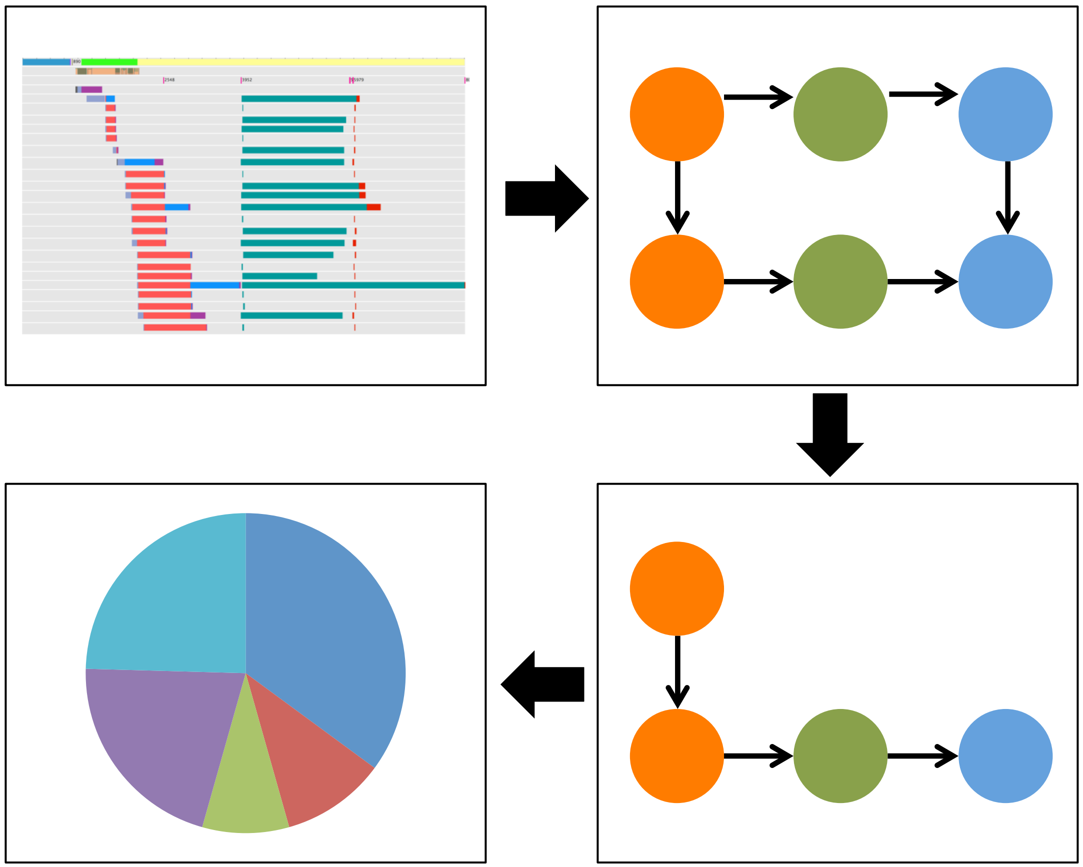

Questions