

Leveraging Existing Instrumentation to Automatically Infer Invariant-Constrained Models




University of Washington

Ivan Beschastnikh
Yuriy Brun
Sigurd Schneider
Michael Sloan
Michael D. Ernst




Saarland University

Synoptic: Mining



from



University of Washington

Ivan Beschastnikh
Yuriy Brun
Sigurd Schneider
Michael Sloan
Michael D. Ernst



Saarland University

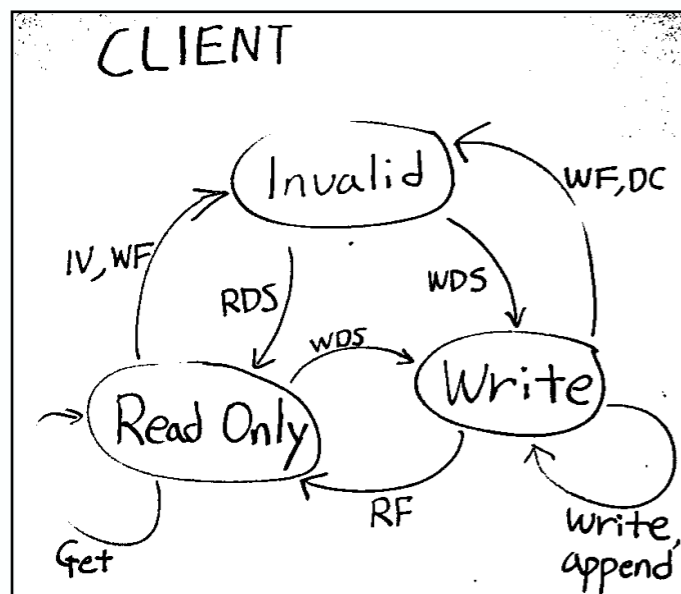
Motivating question

I am a developer.

**Why does my system
behave in a certain manner?**

One answer is to...

Study the relevant artifacts



```
while (transIter.hasNext()) {
    // Create new tc map for node m.
    if (!tc.containsKey(m)) {
        tc.put(m, new LinkedHashSet<EventNode>());
    }

    EventNode child = transIter.next().getTarget();

    if (!tcParents.containsKey(child)) {
        tcParents.put(child, new HashSet<EventNode>());
    }

    // Link m to c
    tc.get(m).add(child);
    tcParents.get(child).add(m);

    // Link m to all nodes that c is linked to in tc
    if (tc.containsKey(child)) {
        // m can reach nodes the child can reach trans
        tc.get(m).addAll(tc.get(child));
        // nodes that child can reach have m as a tc
        for (EventNode n : tc.get(child)) {
            if (!tcParents.containsKey(n)) {
                tcParents.put(n, new HashSet<EventNode>());
            }
            tcParents.get(n).add(m);
        }
    }
}
```

```
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
src : 2, dst : 0, timestamp : 20, type : tx_commit
src : 2, dst : 1, timestamp : 21, type : tx_commit
src : 0, dst : 2, timestamp : 22, type : ack
src : 1, dst : 2, timestamp : 23, type : ack
src : 2, dst : 0, timestamp : 0, type : prepare
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
src : 2, dst : 0, timestamp : 20, type : tx_commit
src : 2, dst : 1, timestamp : 21, type : tx_commit
src : 0, dst : 2, timestamp : 22, type : ack
src : 1, dst : 2, timestamp : 23, type : ack
src : 2, dst : 0, timestamp : 0, type : prepare
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
src : 2, dst : 0, timestamp : 20, type : tx_commit
src : 2, dst : 1, timestamp : 21, type : tx_commit
src : 0, dst : 2, timestamp : 22, type : ack
src : 1, dst : 2, timestamp : 23, type : ack
src : 2, dst : 0, timestamp : 0, type : prepare
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
```

Specification

Code

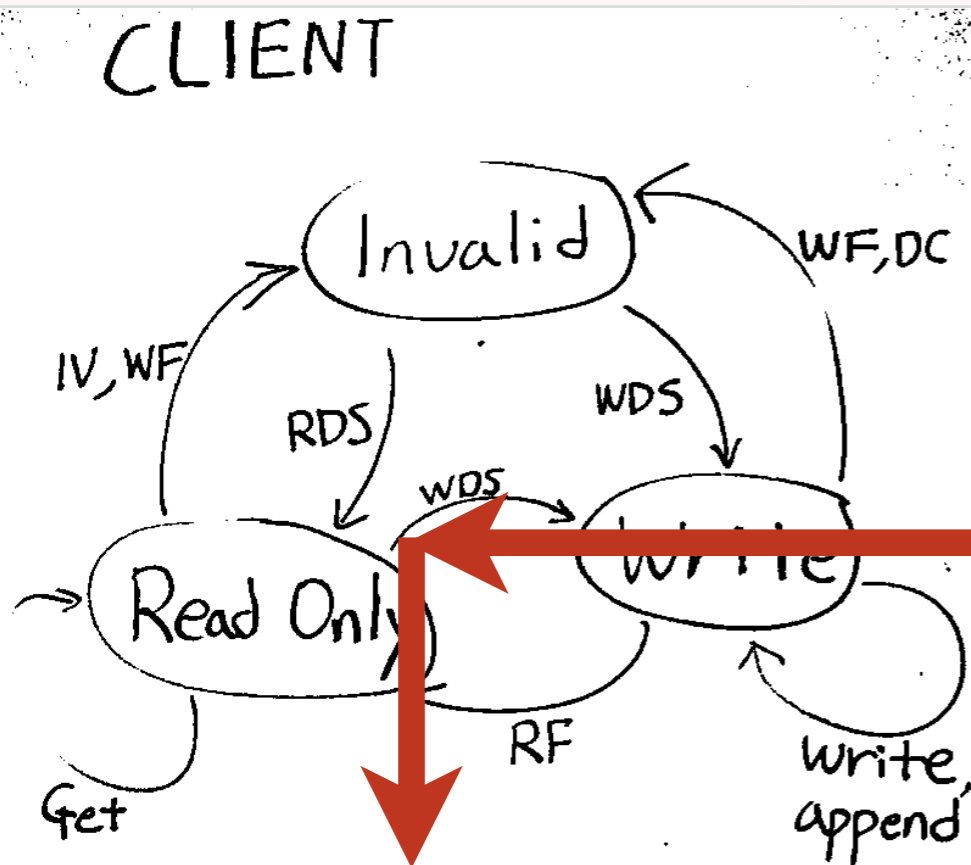
Log

Models

Logs

Manually defined

Hard to Create



Easy to Use

Concise and exact

Trivial instrumentation

Easy to Create

```
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
src : 2, dst : 0, timestamp : 20, type : tx_commit
src : 2, dst : 1, timestamp : 21, type : tx_commit
src : 0, dst : 2, timestamp : 22, type : ack
src : 1, dst : 2, timestamp : 23, type : ack
src : 2, dst : 0, timestamp : 0, type : prepare
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
src : 2, dst : 0, timestamp : 20, type : tx_commit
src : 2, dst : 1, timestamp : 21, type : tx_commit
src : 0, dst : 2, timestamp : 22, type : ack
src : 1, dst : 2, timestamp : 23, type : ack
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
src : 2, dst : 0, timestamp : 20, type : tx_commit
src : 2, dst : 1, timestamp : 21, type : tx_commit
src : 0, dst : 2, timestamp : 22, type : ack
src : 1, dst : 2, timestamp : 23, type : ack
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
```

Hard to Use

A low-level view

Use cases for Synoptic models

- Log summarization
- Mental model validation
- Test case generation
- Comparison to formal specifications
- Evaluate test suite against actual usage
- Verifying a code fix

Talk outline

- Motivation
- Synoptic's design
- Evaluation results
- Future work
- Conclusion

Synoptic inputs

1. Log file
2. Regular expressions to parse the log file

Two phase commit protocol log

1. Log file


Two phase commit protocol log

I. Manager **proposes** a transaction (TX), each replica replies with an **abort** or a **commit**.

II. Manager sends

- i) **TX commit** if all replicas commit
- ii) **TX abort** otherwise

Manager maintains a totally ordered log of events for all transactions in the system




```
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
src : 2, dst : 0, timestamp : 20, type : tx_commit
src : 2, dst : 1, timestamp : 21, type : tx_commit
src : 0, dst : 2, timestamp : 22, type : ack
src : 1, dst : 2, timestamp : 23, type : ack
src : 2, dst : 0, timestamp : 0, type : prepare
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
```

log.txt

Regular expressions

1. Log file (e.g., two phase commit log)
2. Regular expressions to parse the log file
 - Split log into **executions**
 - From each log line extract:
 - Total event relation (e.g., **time**)
 - Event **type**



```
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
src : 0, dst : 2, timestamp : 18, type : commit
src : 1, dst : 2, timestamp : 19, type : commit
src : 2, dst : 0, timestamp : 20, type : tx_commit
src : 2, dst : 1, timestamp : 21, type : tx_commit
src : 0, dst : 2, timestamp : 22, type : ack
src : 1, dst : 2, timestamp : 23, type : ack
src : 2, dst : 0, timestamp : 0, type : prepare
src : 2, dst : 1, timestamp : 1, type : prepare
src : 0, dst : 2, timestamp : 2, type : commit
src : 1, dst : 2, timestamp : 3, type : commit
src : 2, dst : 0, timestamp : 4, type : tx_commit
src : 2, dst : 1, timestamp : 5, type : tx_commit
src : 0, dst : 2, timestamp : 6, type : ack
src : 1, dst : 2, timestamp : 7, type : ack
src : 2, dst : 0, timestamp : 8, type : prepare
src : 2, dst : 1, timestamp : 9, type : prepare
src : 0, dst : 2, timestamp : 10, type : commit
src : 1, dst : 2, timestamp : 11, type : commit
src : 2, dst : 0, timestamp : 12, type : tx_commit
src : 2, dst : 1, timestamp : 13, type : tx_commit
src : 0, dst : 2, timestamp : 14, type : ack
src : 1, dst : 2, timestamp : 15, type : ack
src : 2, dst : 0, timestamp : 16, type : prepare
src : 2, dst : 1, timestamp : 17, type : prepare
```

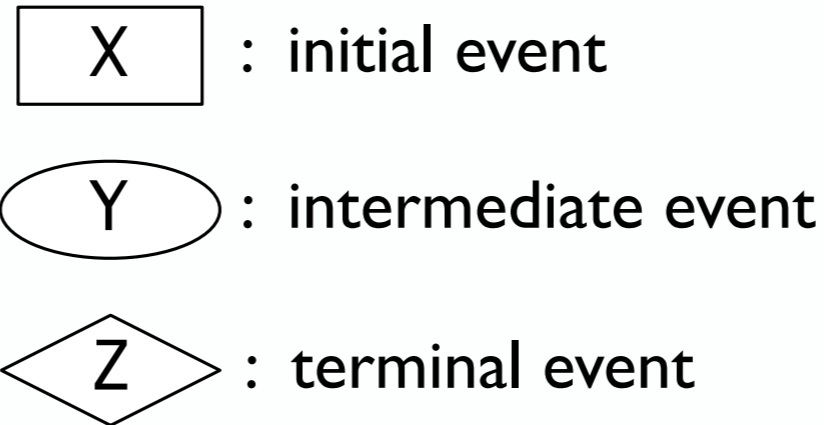
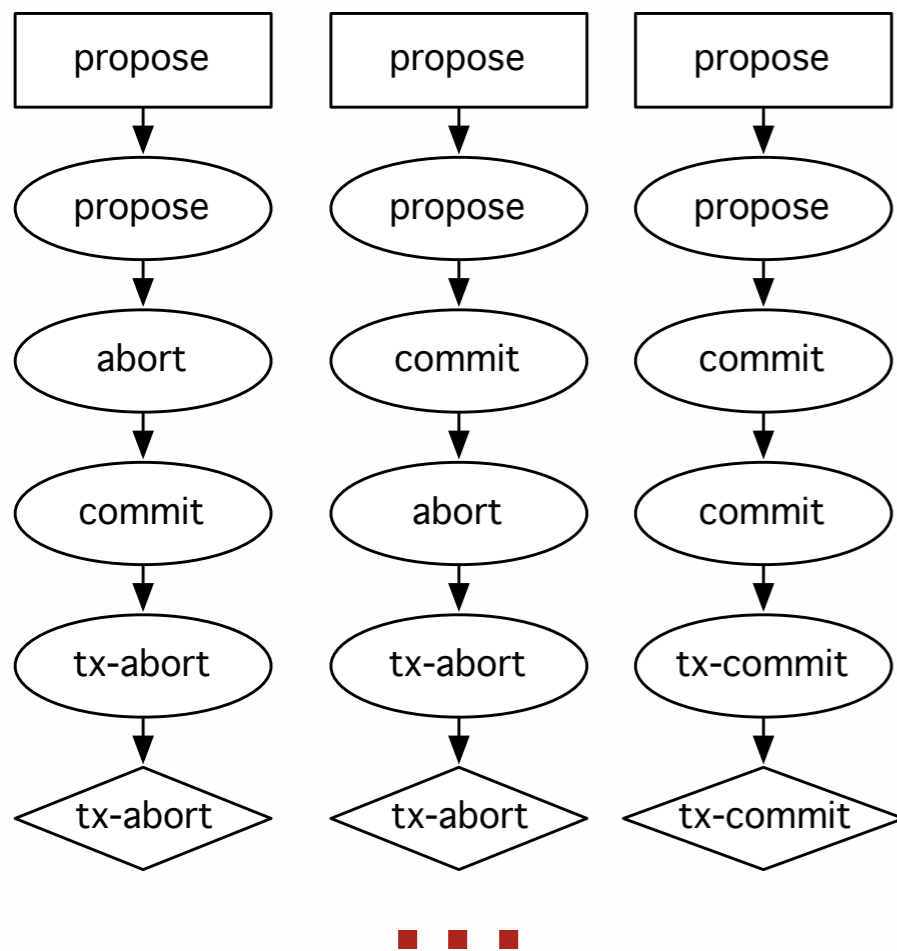
log.txt

Synoptic overview

1/5. Parse log into a trace graph

2/5. Construct the initial model

3/5. Mine temporal invariants


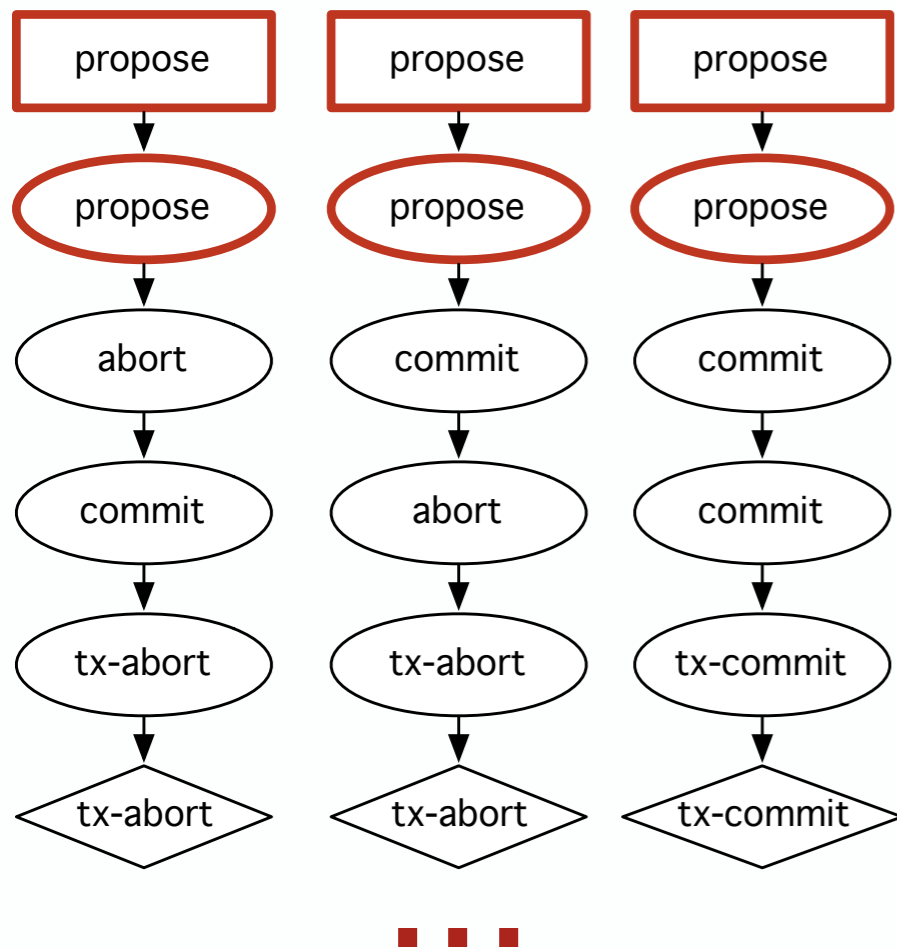


Synoptic overview

1/5. Parse log into a trace graph

2/5. Construct the initial model

3/5. Mine temporal invariants



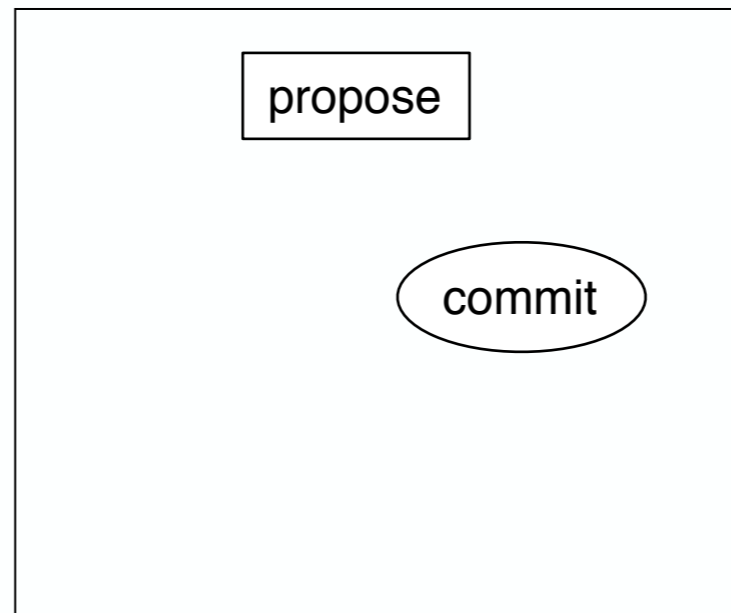
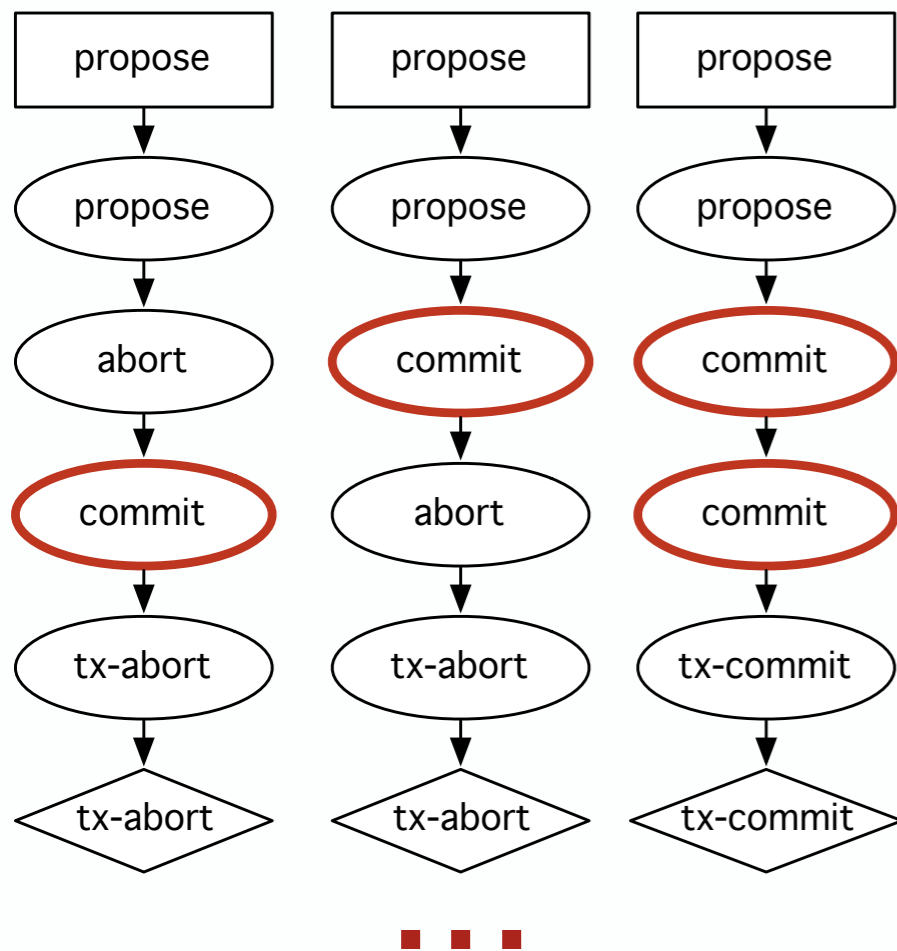
A compact model with one node per event type

Synoptic overview

1/5. Parse log into a trace graph

2/5. Construct the initial model

3/5. Mine temporal invariants





A compact model with one node per event type

Synoptic overview

1/5. Parse log into a trace graph

2/5. Construct the initial model

3/5. Mine temporal invariants



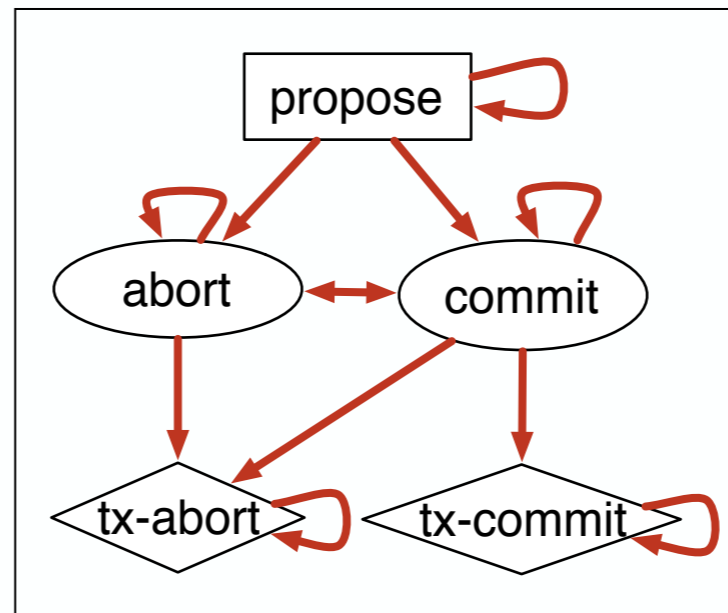

A compact model with one node per event type

Synoptic overview

1/5. Parse log into a trace graph

2/5. Construct the initial model

3/5. Mine temporal invariants




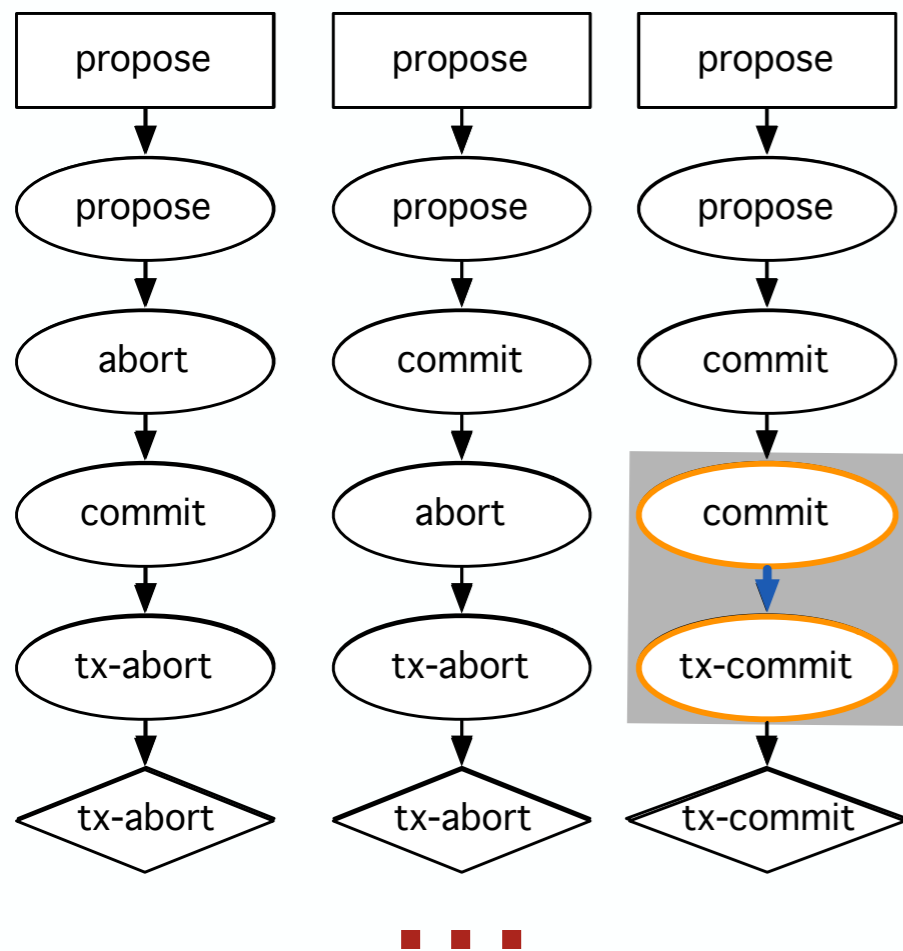
A compact model with one node per event type

Synoptic overview

1/5. Parse log into a trace graph

2/5. Construct the initial model

3/5. Mine temporal invariants



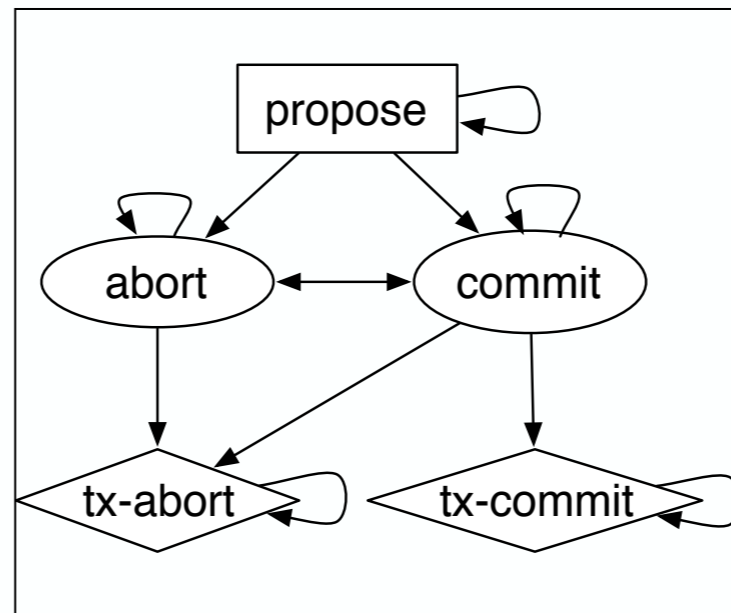
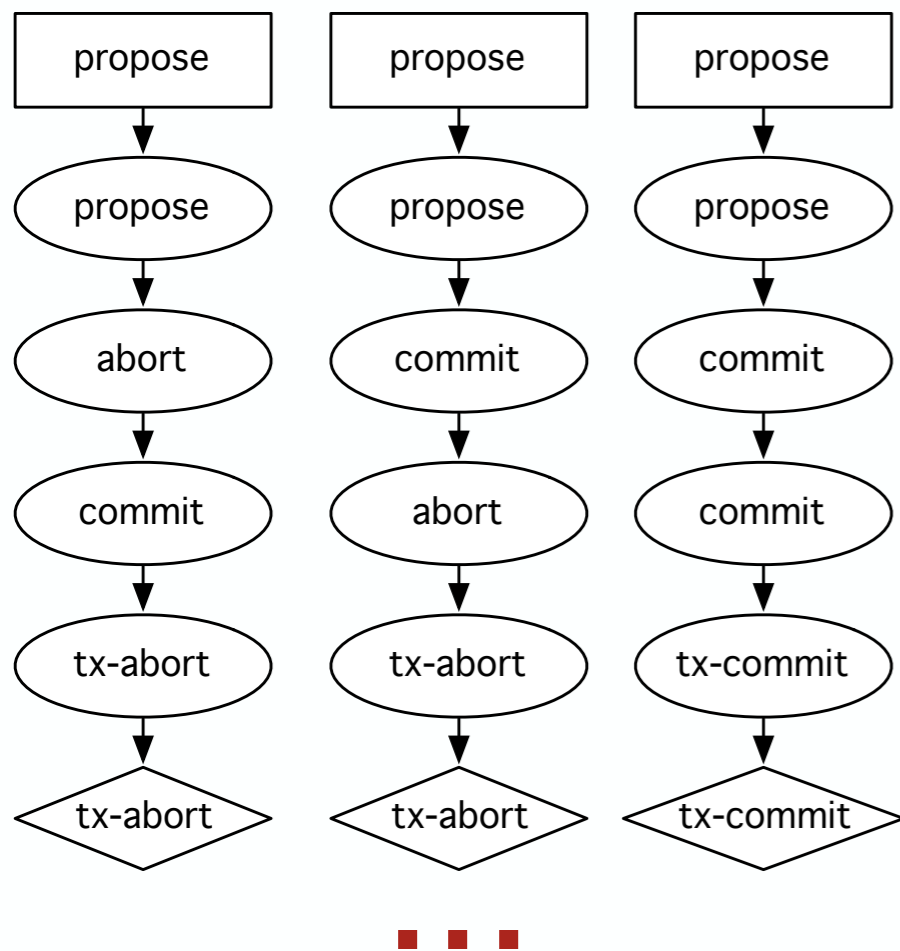
A compact model with one node per event type

Synoptic overview

1/5. Parse log into a trace graph

2/5. Construct the initial model

3/5. Mine temporal invariants



A compact model with one node per event type

abort → tx-abort

abort ↗ tx-commit

abort ← tx-abort

...

Synoptic overview

4/5. Refine the initial model until all invariants are satisfied

5/5. Coarsen model without unsatisfying any invariants

True for log, false for initial model

Choose an invariant invalid in the model
abort \rightarrow tx-abort

Synoptic overview

4/5. Refine the initial model until all invariants satisfied


5/5. Coarsen model without unsatisfying any invariants

Choose an invariant invalid in the model

abort \rightarrow tx-abort

Unsatisfied invariants exist

Refine a model node



Synoptic overview

4/5. Refine the initial model until all invariants satisfied

5/5. Coarsen model without unsatisfying any invariants


Choose an invariant invalid in the model

abort \rightarrow tx-abort

Unsatisfied invariants exist

All invariants satisfied

Refine a model node



Synoptic overview

4/5. Refine the initial model until all invariants satisfied

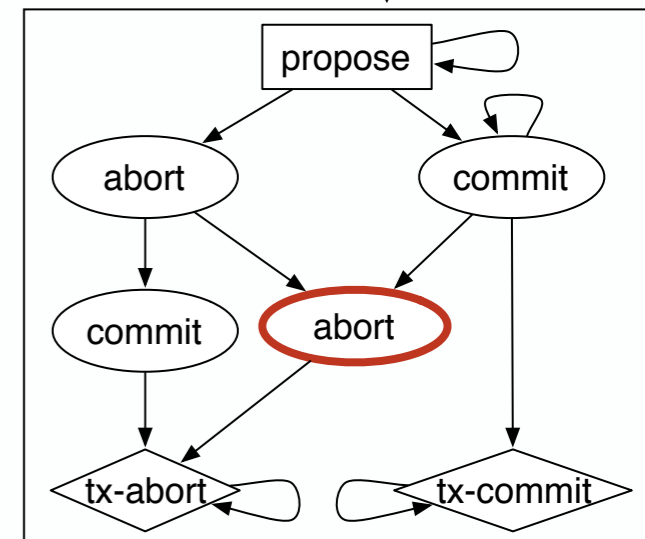
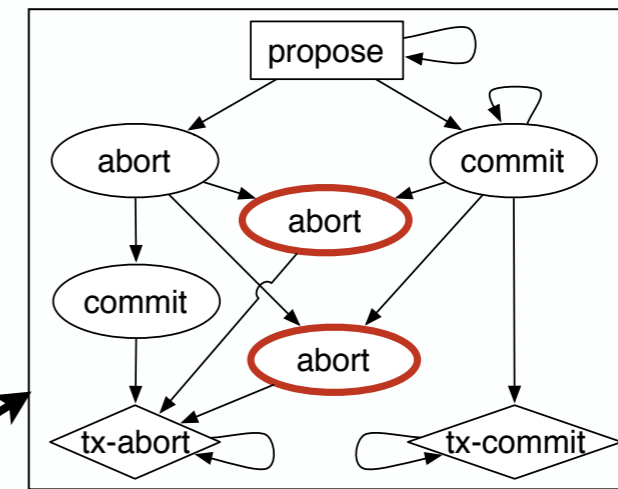

5/5. Coarsen model without unsatisfying any invariants

Choose an invariant invalid in the model
abort \rightarrow tx-abort

Unsatisfied invariants exist

All invariants satisfied

Refine a model node



Merge nodes that exhibit the same behaviors

Assumptions

Logs

- Multiple executions
- Events in an execution are totally ordered

System

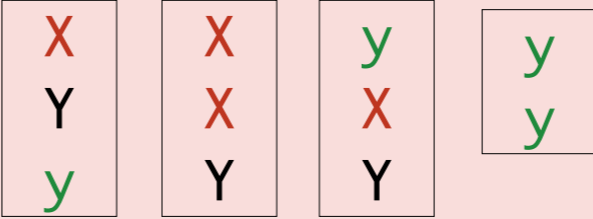
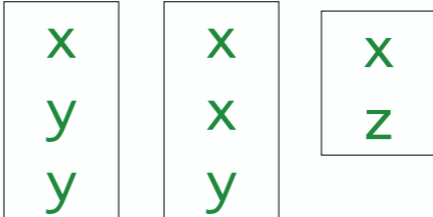
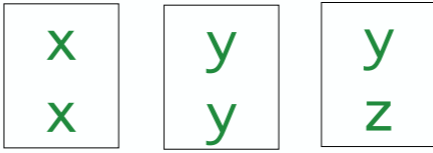
- Process generating the log can be modeled as an FSM
- Temporal system properties are evident in the orderings among events

Talk outline

- Motivation
- Synoptic's design
 - Overview
 - Assumptions
 - Temporal invariants
 - Model space exploration
- Evaluation results
- Future work
- Conclusion

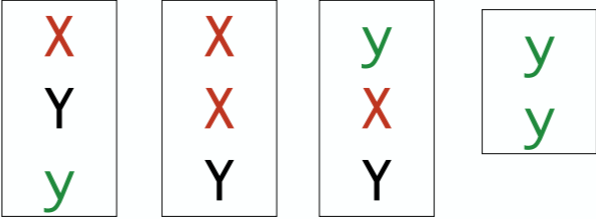
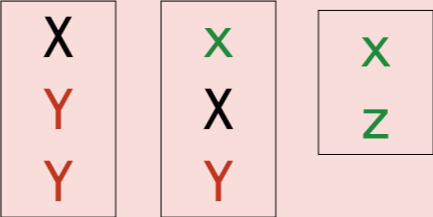
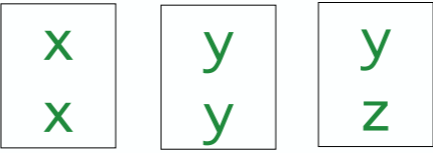
Temporal log invariants

- Synoptic mines three kinds of **temporal log invariants** that are true for all the logged executions

Invariant	Example	Type
$x \longrightarrow y$ always followed by		liveness
$x \longleftarrow y$ always precedes		safety
$x \not\longrightarrow y$ never followed by		safety

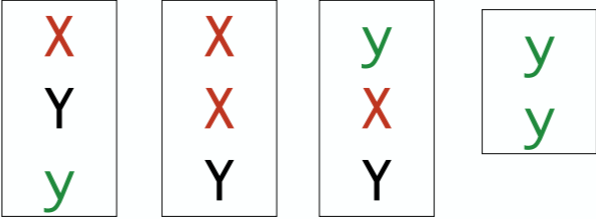
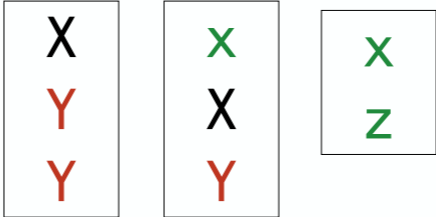
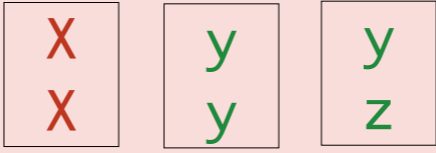
Temporal log invariants

- Synoptic mines three kinds of **temporal log invariants** that are true for all the logged executions

Invariant	Example	Type
$x \longrightarrow y$ always followed by		liveness
$x \longleftarrow y$ always precedes		safety
$x \not\longrightarrow y$ never followed by		safety

Temporal log invariants

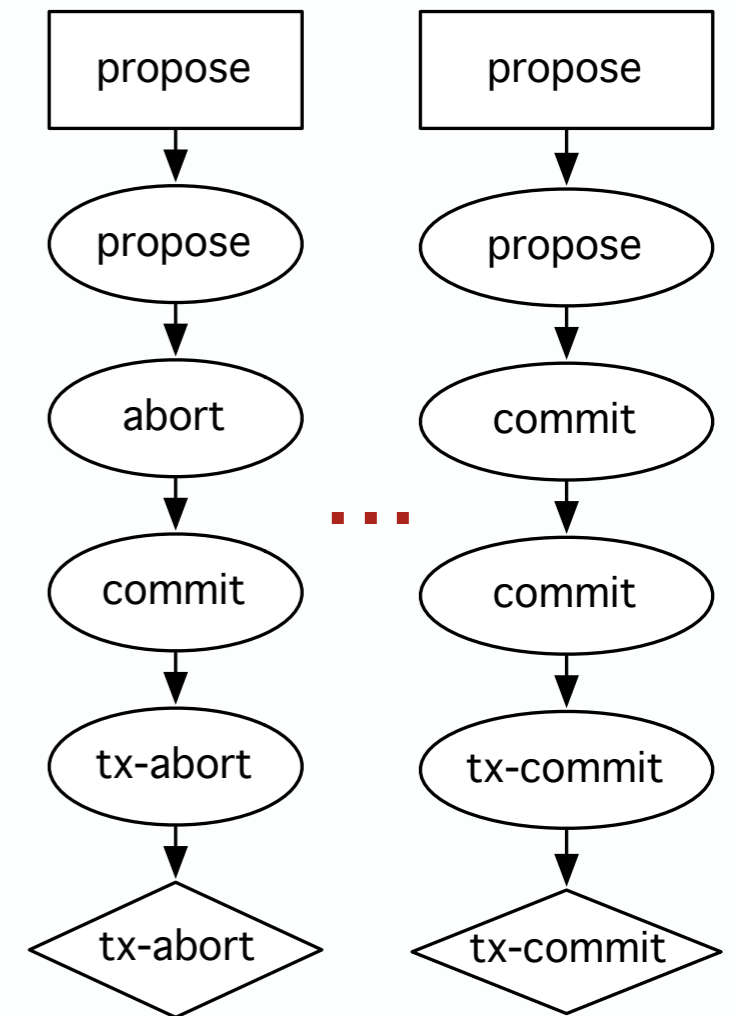
- Synoptic mines three kinds of **temporal log invariants** that are true for all the logged executions

Invariant	Example	Type
$x \longrightarrow y$ always followed by		liveness
$x \longleftarrow y$ always precedes		safety
$x \not\longrightarrow y$ never followed by		safety

Two phase commit invariants

- Mined invariants:

- abort \longrightarrow tx-abort
- abort $\not\longrightarrow$ tx-commit
- abort \longleftarrow tx-abort
- commit \longleftarrow tx-commit
- $\left\{ \begin{array}{l} \text{abort, commit,} \\ \text{tx-abort, tx-commit} \end{array} \right\} \not\longrightarrow$ propose
- propose \longleftarrow $\left\{ \begin{array}{l} \text{abort, commit,} \\ \text{tx-abort, tx-commit} \end{array} \right\}$




Satisfying invariants with refinement

- The initial graph violates 3 mined invariants:

- abort \longrightarrow tx-abort
- abort $\not\longrightarrow$ tx-commit
- abort \longleftarrow tx-abort
- commit \longleftarrow tx-commit
- $\left\{ \begin{array}{l} \text{abort, commit,} \\ \text{tx-abort, tx-commit} \end{array} \right\} \not\longrightarrow$ propose
- propose \longleftarrow $\left\{ \begin{array}{l} \text{abort, commit,} \\ \text{tx-abort, tx-commit} \end{array} \right\}$


One node per event type




- For each violation find a corresponding path
- Use refinement to eliminate such path counter-examples

Counter-example based refinement


Example invariant: abort \rightarrow tx-abort



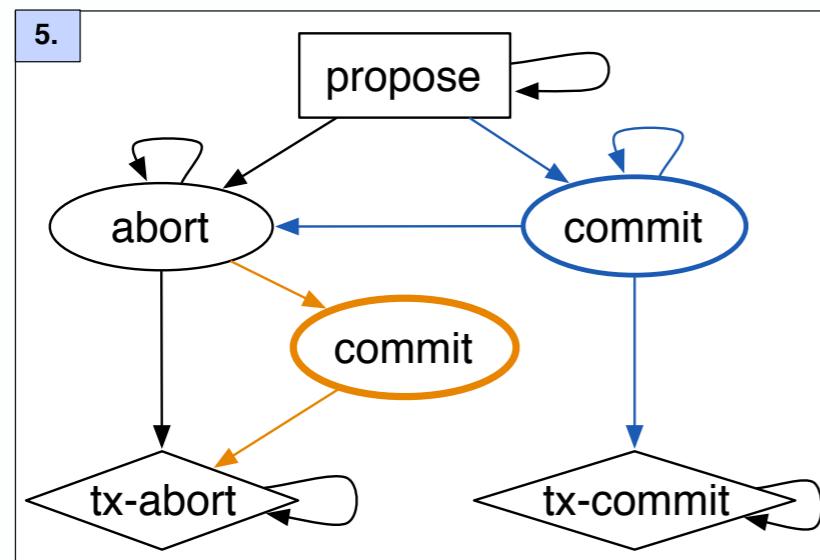
Initial model



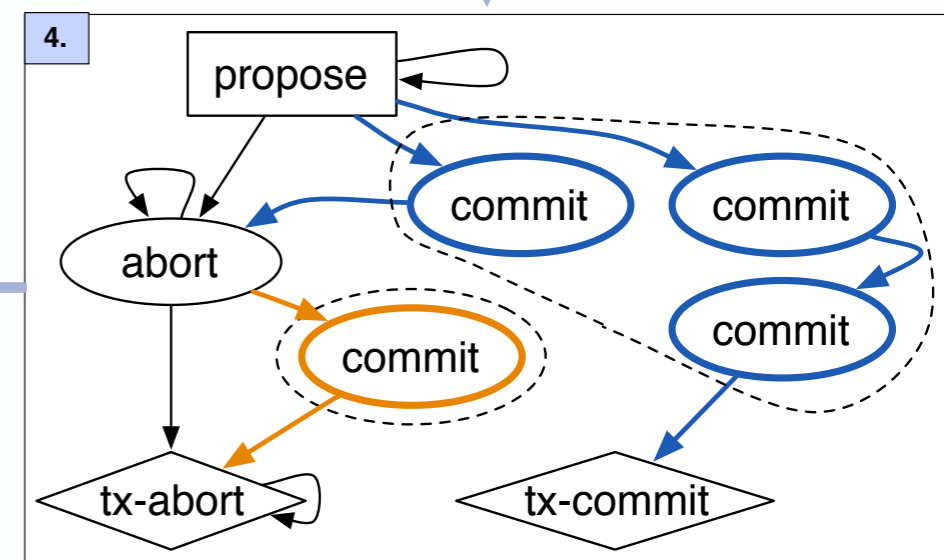
Invariant counter-example



Identify partition to refine



Refined model

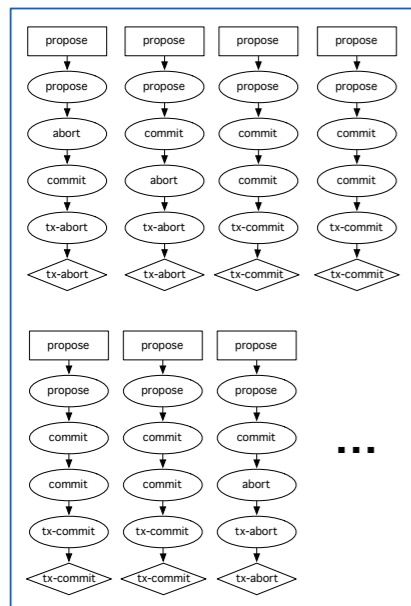


Eliminate counter-example

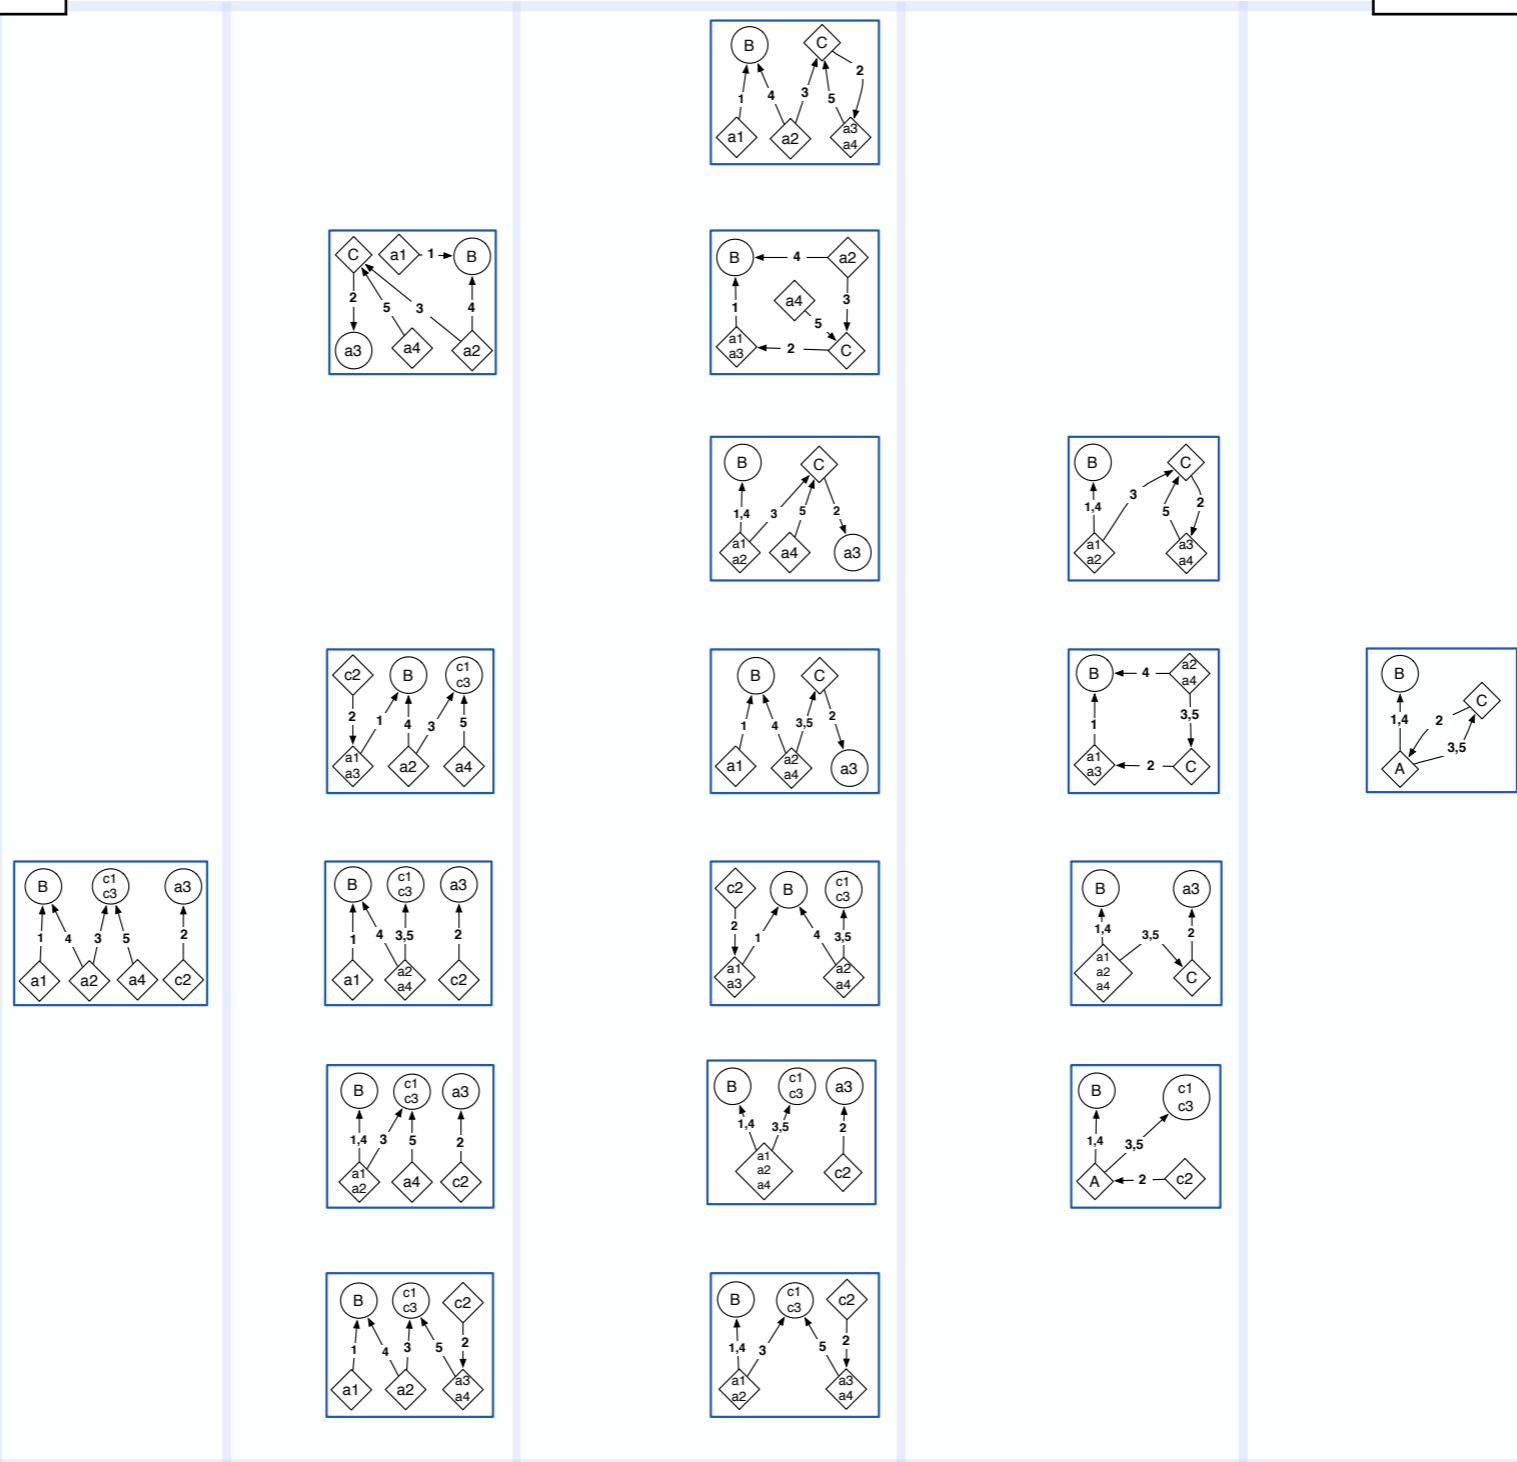
Model space

Larger models:
fewer behaviors

Smaller models:
more behaviors



Trace graph




Initial model

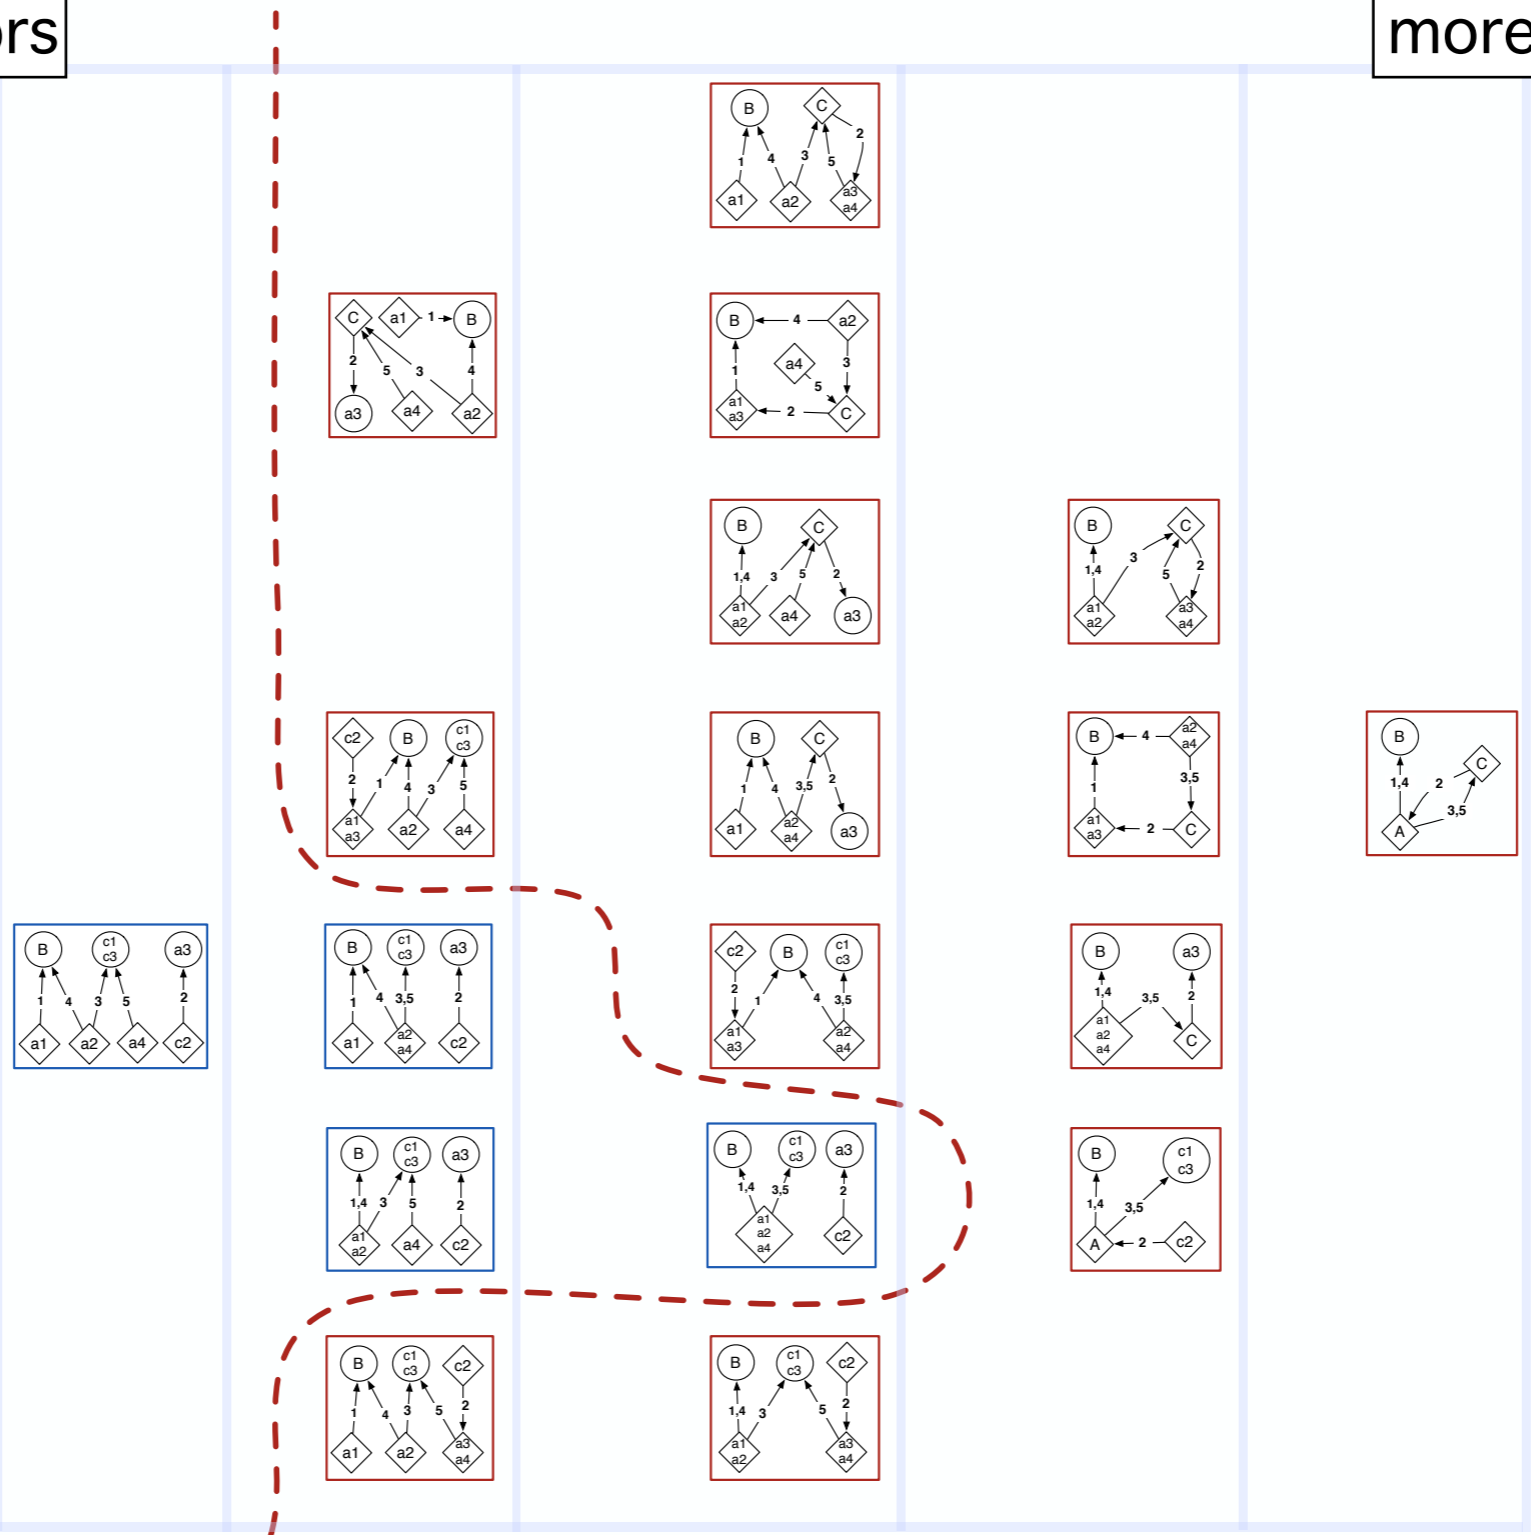
Model space

Larger models:
fewer behaviors

Smaller models:
more behaviors



Trace graph



Initial model

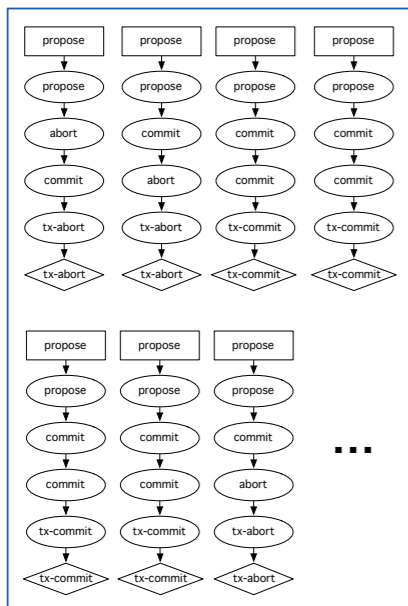
Model space

Larger models:
fewer behaviors

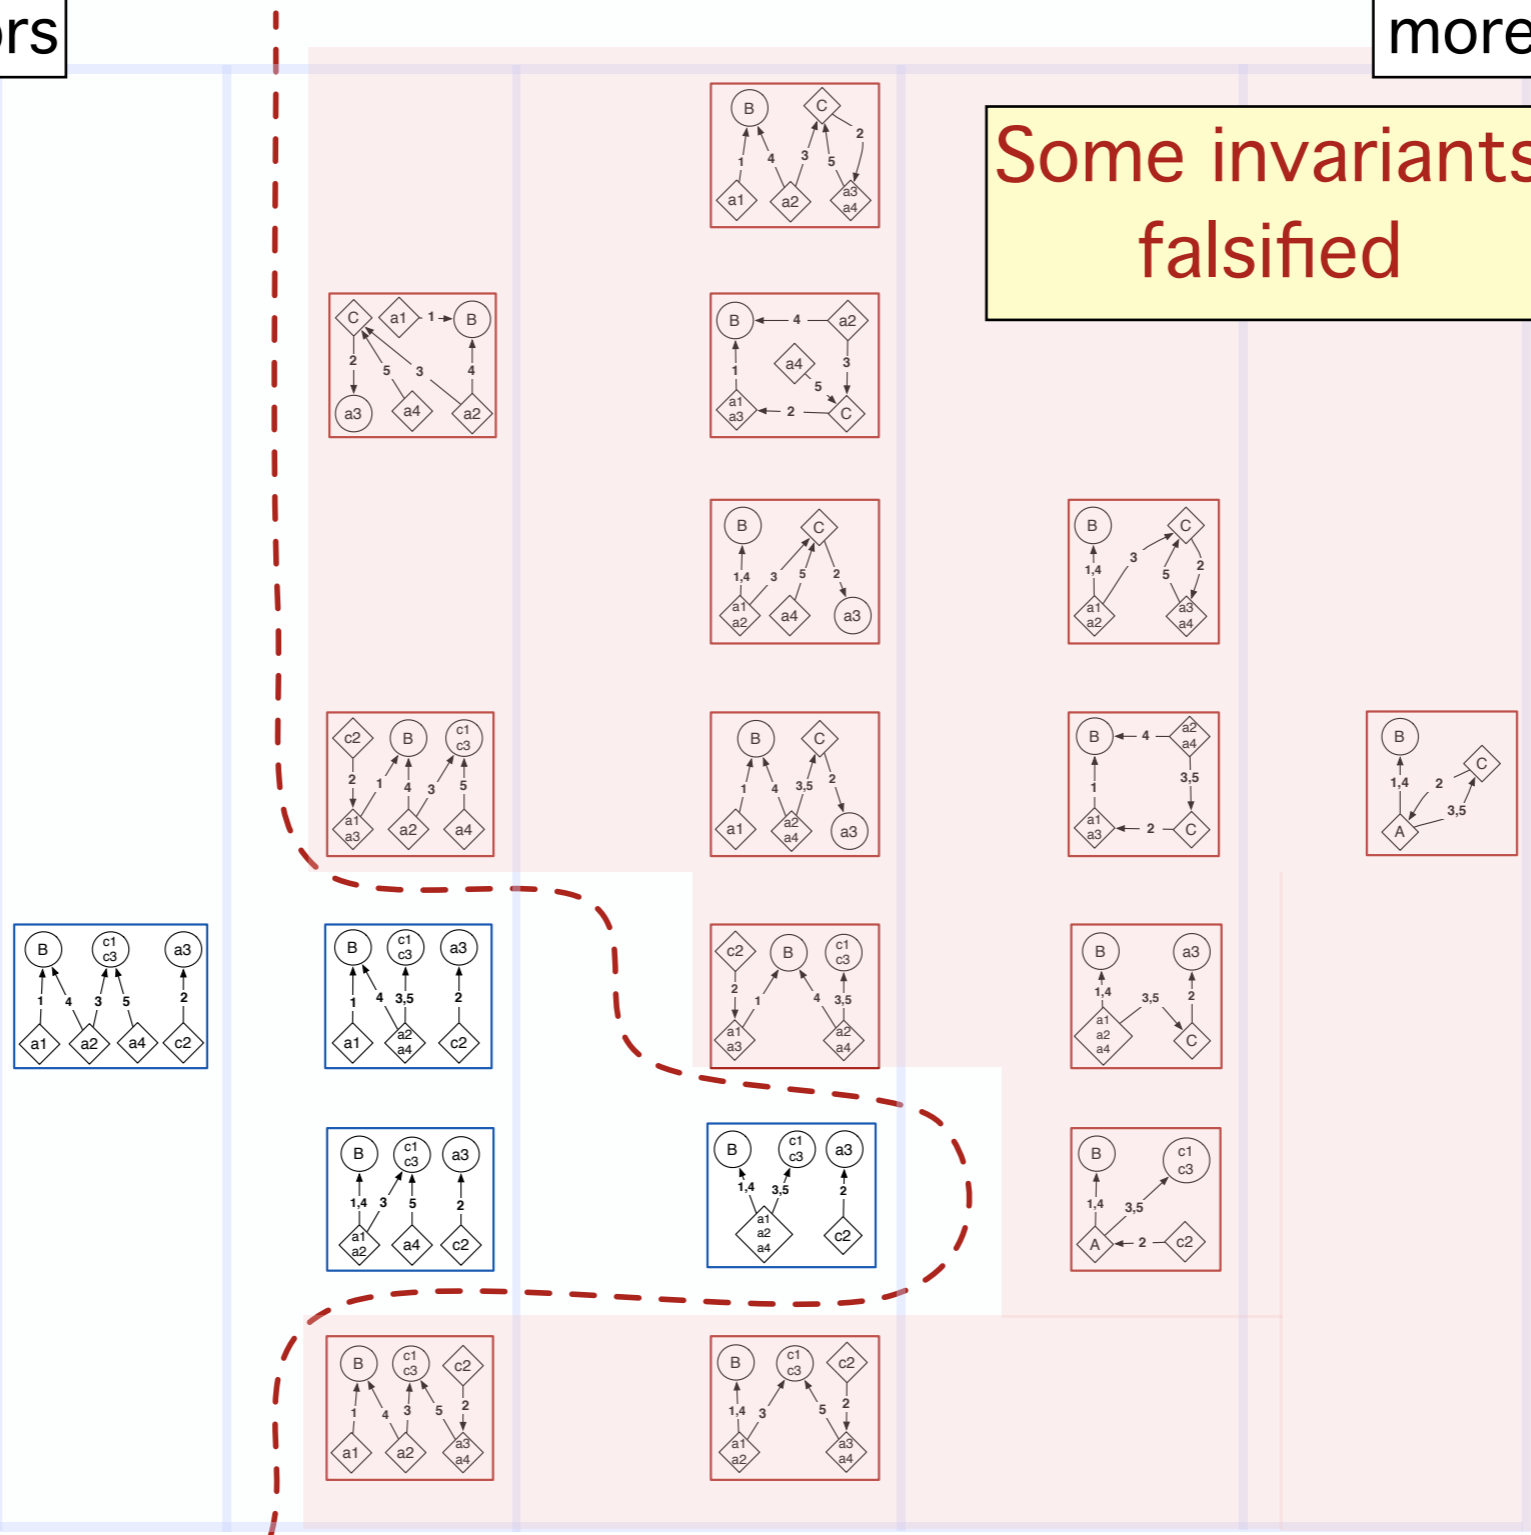
Smaller models:
more behaviors

Some invariants
falsified

Initial model



Trace graph



Model space

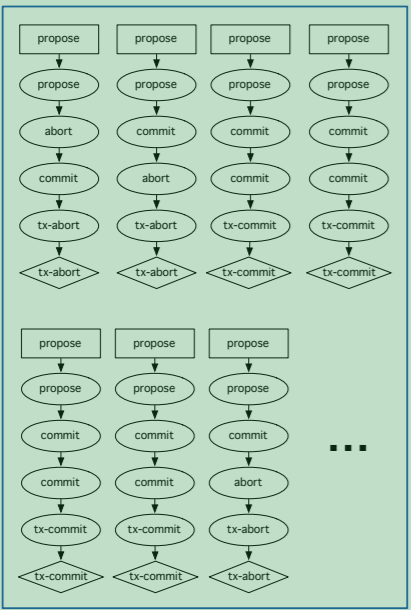
Larger models:
fewer behaviors

Smaller models:
more behaviors

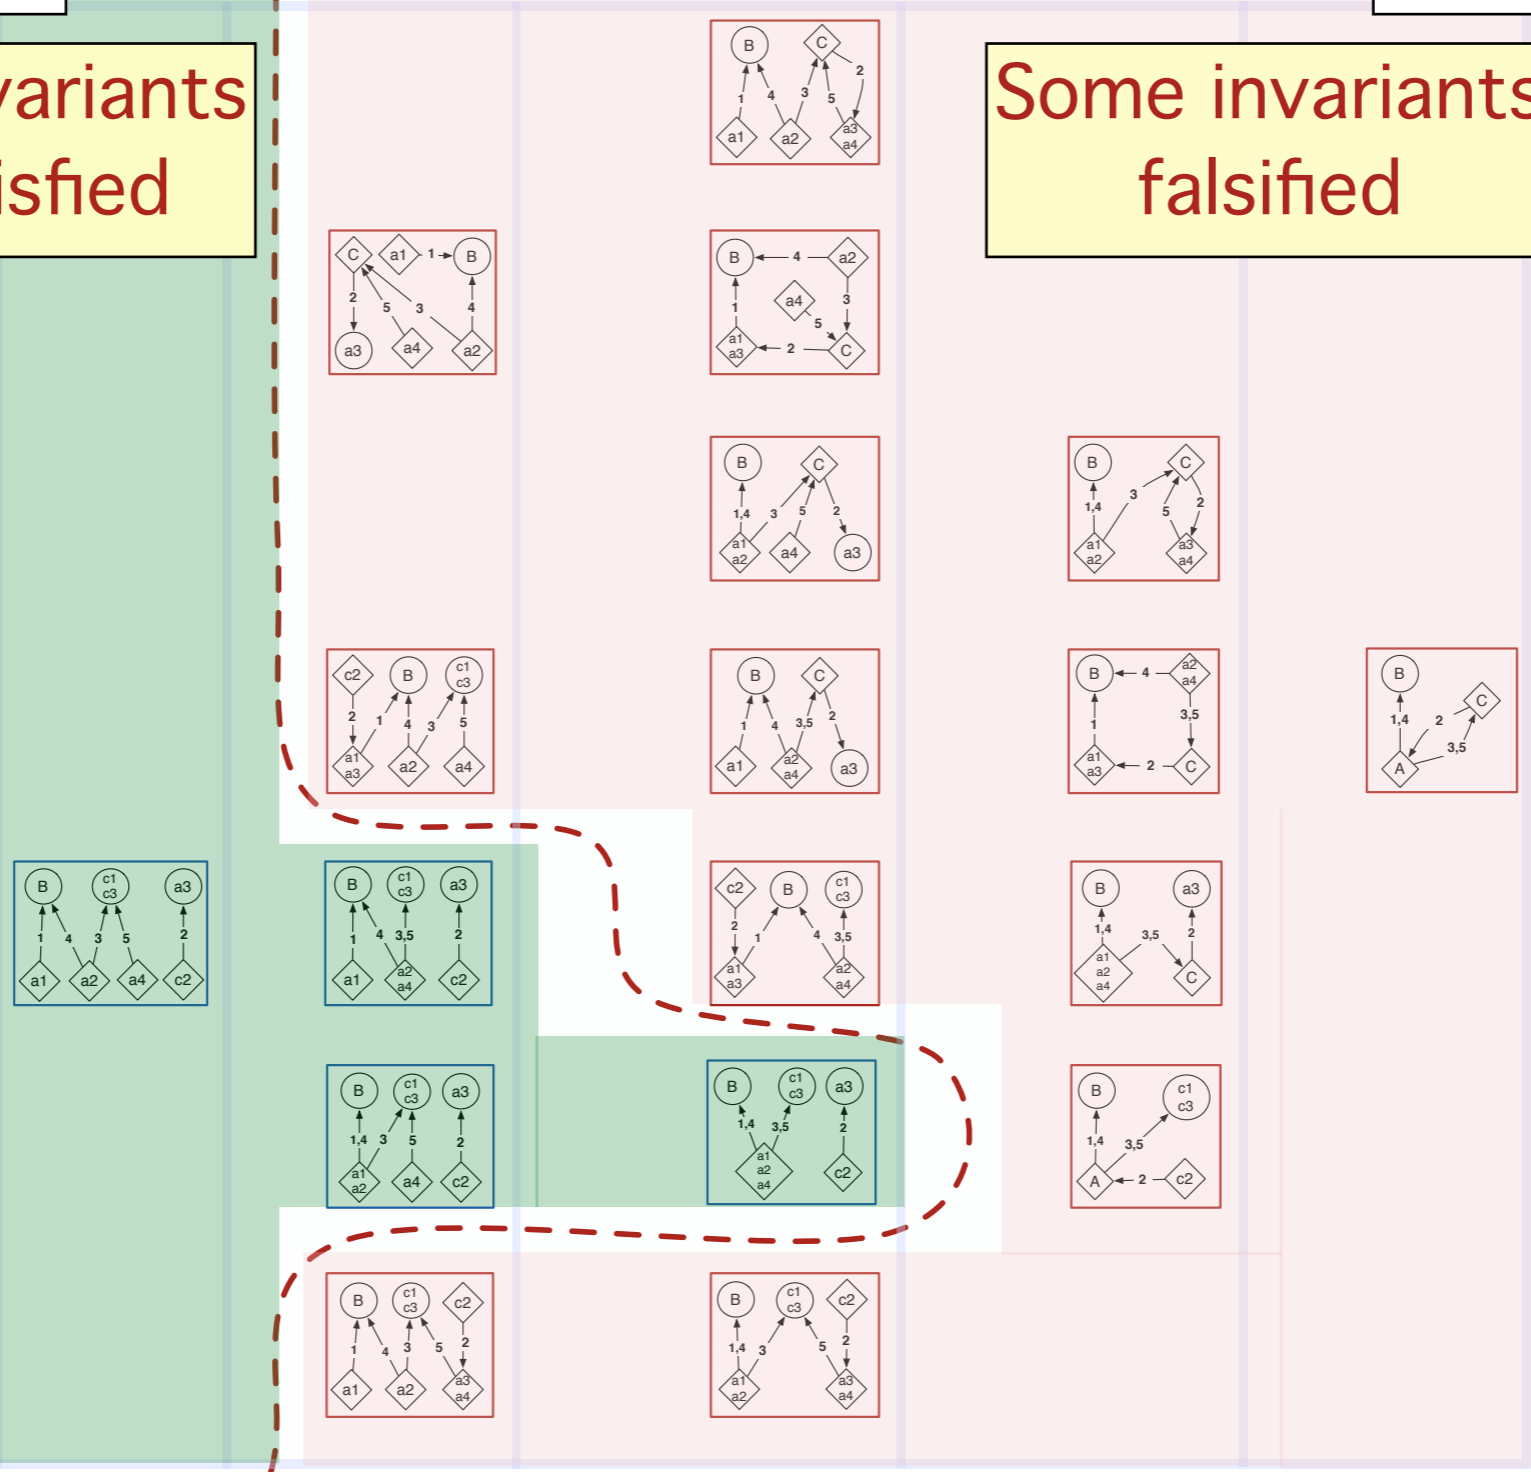
All invariants
satisfied

Some invariants
falsified

Initial model



Trace graph



Model space

Larger models:
fewer behaviors

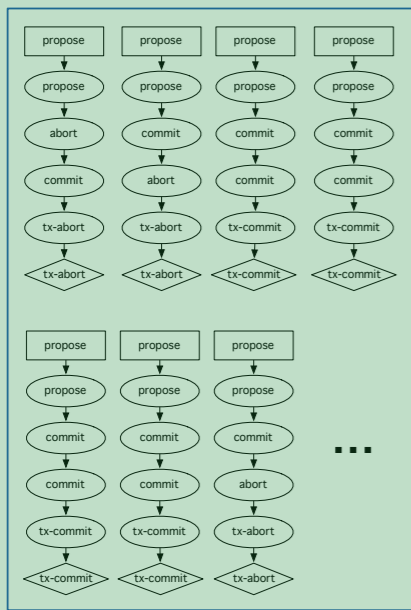
Smaller models:
more behaviors

All invariants
satisfied

















Some invariants
falsified

Initial model

Synoptic's goal:
Find the
smallest model
satisfying all
invariants



Trace graph



Model space

Larger models:
fewer behaviors

Smaller models:
more behaviors

All invariants
satisfied

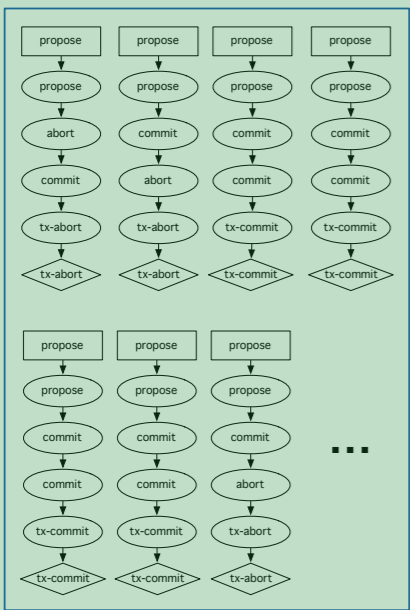
Some invariants
falsified

Coarsening

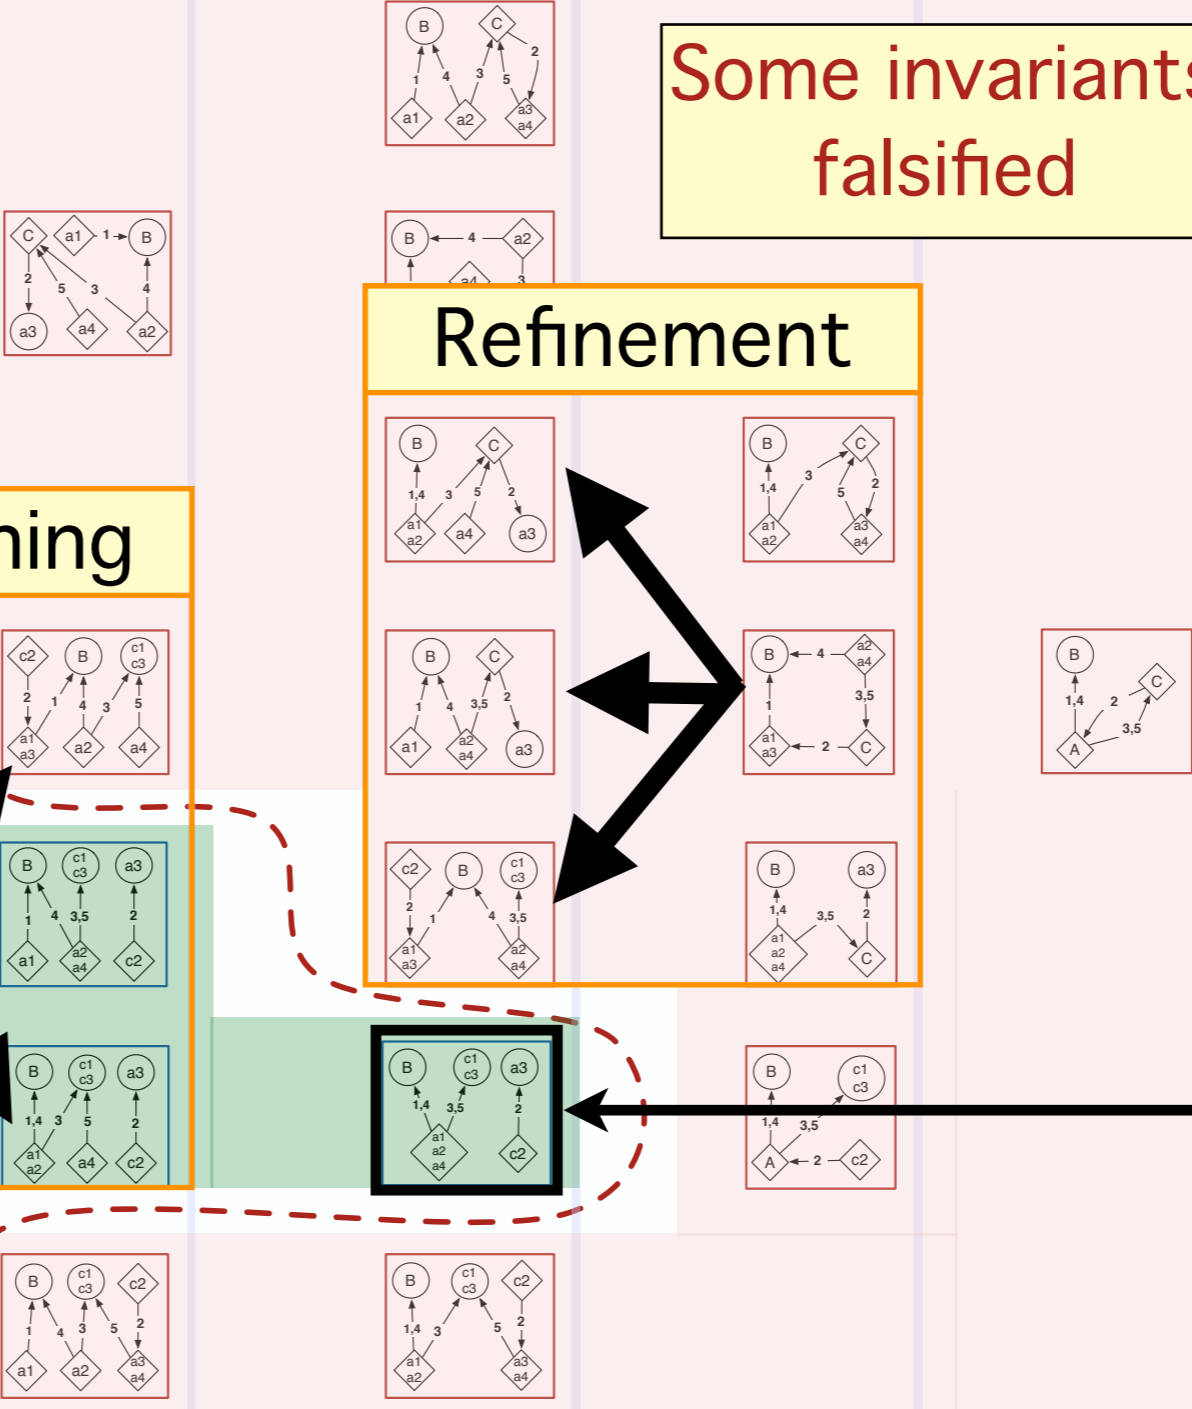
Refinement

Initial model

Synoptic's goal:
Find the
smallest model
satisfying all
invariants



Trace graph



















All invariants satisfied

Some invariants satisfied

Synoptic finds a local optimum, not a global one

Select a **coarsening** that maintains all the invariants

Select a **refinement** that satisfies an unsatisfied invariant



Initial model

Coarsen




Refine

Refine

Refine

All invariants satisfied

Some invariants satisfied







• Coarsening from initial traces is inefficient

- k-Tail [Biermann et al. 1972](#) and GK-Tail [Lorenzoli et al. ICSE 2008](#)





• Refinement is well known [Paige et al. 1987](#)

- Using mined invariant constraints is a new idea

Synoptic finds a local optimum, not a global one




Initial model





Refine

Select a coarsening that maintains all the invariants




Coarsen




Select a refinement that satisfies an unsatisfied invariant

Refine



Refine



Talk outline

- Motivation
- Synoptic's design
- Evaluation results
 - Formal evaluation
 - Reverse traceroute
 - Distributed systems course
- Future work
- Conclusion

Formal evaluation

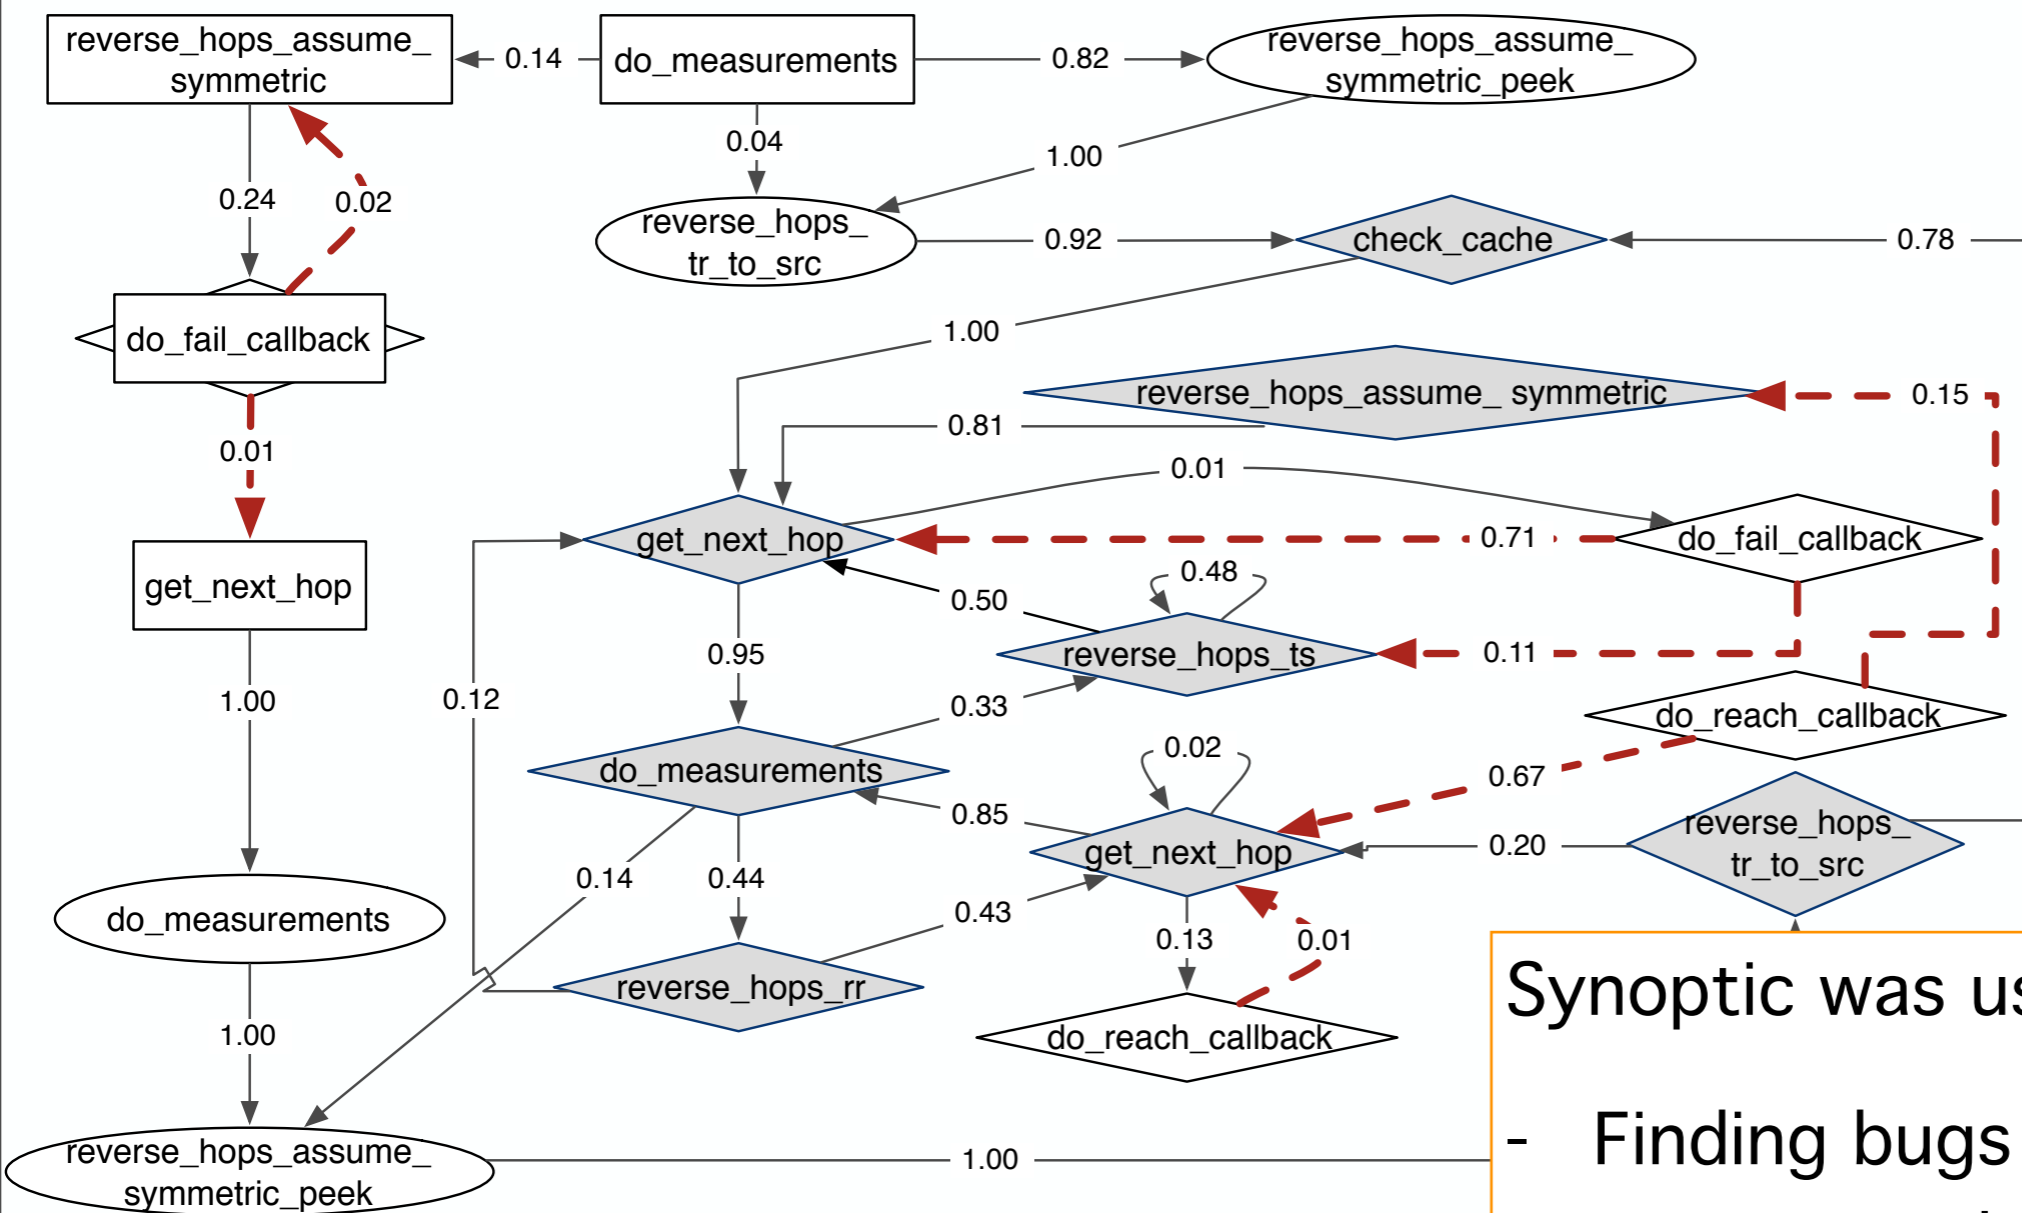
- Termination
 - Always finds a model
- Correctness
 - Satisfies all the mined invariants, and no others
- Conciseness
 - Local minimum
- Efficiency
 - Never violate a mined invariant
 - Make progress towards satisfying some mined invariant

See the paper
for more

Reverse traceroute case-study

- Reverse traceroute KATZ-BASSETT et al. NSDI 2010
 - Finds the likely reverse traceroute from an arbitrary destination to a source host
 - In live deployment for over a year
 - Deployed internally by a large Internet company
- Ran Synoptic on reverse traceroute server logs
 - Developer added 16 lines of logging code
 - Synoptic processed the 900,000 events in 12 minutes

Reverse traceroute case-study



2 bugs found:

- Shaded nodes should not be terminal
- Red edges should not exist

Synoptic was useful for:

- Finding bugs in code
- Increasing developer confidence
- Building system understanding

Synoptic limitations

- Quality and usefulness of the model depends on granularity of logging statements and user input
- Only considers three kinds of temporal invariants
- For some applications, the small model size may be irrelevant
- Online log analysis is not supported
- The input log must be totally ordered

Future work

- Handle logs generated by distributed systems
- Generalize invariants
 - Relax **always** and **never** constraints Lou et al. KDD 2010
 - Explore composition of temporal invariants Gabel et al. FSE 2008
 - Incorporating structural/data invariants Ernst et al. TSE 2001
- Tool improvements
 - Explicating model paths in the GUI
 - Eclipse plug-in for ease of use
- Automatically infer log format

Vaarandi et al. IFIP ICICS 2004	Zhu et al. OSR 2010
------------------------------------	------------------------

Contributions

Synoptic: a tool that extracts models from logs

- Specification mining algorithm
 - Uses **refinement** for **efficiency**
 - Mines and satisfies **temporal log invariants** for **accuracy**
- Proved important properties of refinement
- Showed that Synoptic helps developers to study large and complex logs, and can help find bugs

Try it!

<http://synoptic.googlecode.com>