

# Finding Invariants of Distributed Systems: It's a Small (Enough) World After All

Travis Hance

Marijn Heule

Ruben Martins

Bryan Parno

*Carnegie Mellon University*

# Distributed systems are hard to get right

The image shows three overlapping news article snippets, each with a black border, illustrating major distributed system outages:

- Top snippet:** "Buggy code update triggers Bing, Yahoo outage". It includes a date of "Jan 5, 2015 9:00 AM" and a "Like" button with a count of 9.
- Middle snippet:** "Google blames software bug for Friday night Gmail outage". It features a "Promoted Content" label and a partial image of a building with the number "10".
- Bottom snippet:** "Amazon Web Services suffers outage, takes down Vine, Instagram, others with it". It includes a "Topic: Amazon" label, a "Follow via: RSS" link, a "Summary" section, and author information: "By Zack Whittaker for Between the Lines | August 26, 2013 -- 13:22 GMT (06:22 PDT)". It also shows social sharing options for Facebook (0), Twitter (1), and LinkedIn.

# Solution: Verification?

- **Formal verification** can rule out bugs in:
  - Abstract protocol descriptions
  - Implementation
  - Liveness
- **Problem:** Verification is hard
  - Hand-crafted system invariants
  - Invariants must be checked for inductiveness

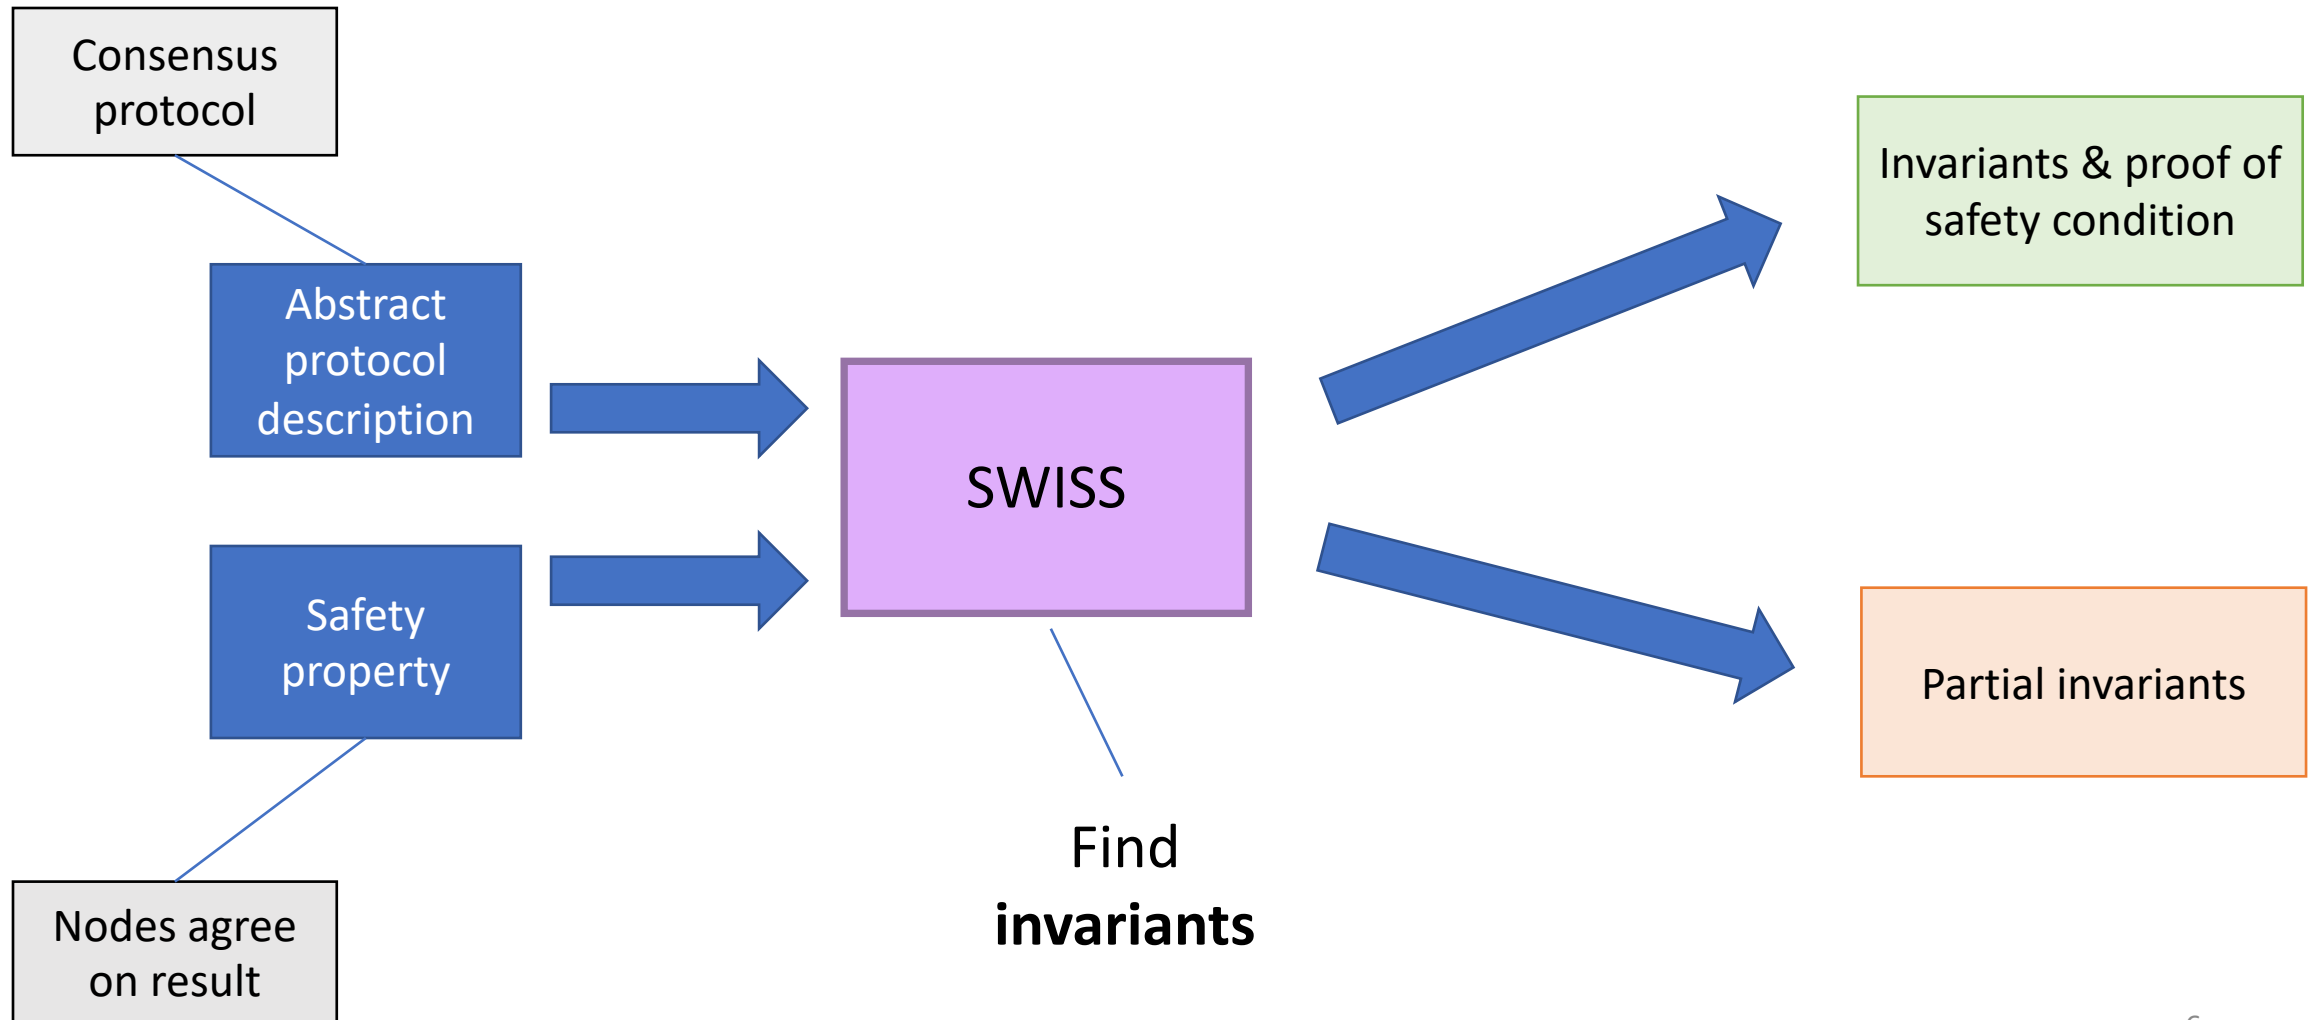
# Solution: Automated Verification

- Less painstaking manual proof work
- **Problem:** Automated Verification is also hard, often undecidable
- Prior work
  - Automate invariant checking (IVy)
  - Automate invariant *finding*
    - I4 (Ma et al.) (SOSP, 2019)
    - Separators algorithm (Padon et al.) based on IC3 (PLDI, 2020)
- Many useful protocols (e.g., Paxos) still out of reach of fully-automated solutions

# SWISS Contributions


- System to automatically prove safety for distributed protocols
- Scales to automate the verification of Paxos
- Handles universal & existential quantifiers
- Can accept additional user guidance—otherwise fully automated
- Produces partial invariants even when it doesn't complete

# SWISS Overview



# SWISS Overview: Invariants

- An **invariant** is a statement about the system which holds true at every point in the execution
- We need an invariant which is ...
  - **Useful** – it can be used to prove the safety condition
  - **Inductive** – it is itself strong enough to prove that it remains true



# An Example Invariant (Paxos)

$$\begin{aligned} & (\forall r, v_1, v_2. \text{proposal}(r, v_1) \wedge \text{proposal}(r, v_2) \implies v_1 = v_2) \\ \wedge & (\forall n, r, v. \text{vote}(n, r, v) \implies \text{proposal}(r, v)) \\ \wedge & (\forall r, v. (\exists n. \text{decision}(n, r, v)) \implies \exists q. \forall n. \text{member}(n, q) \implies \text{vote}(n, r, v)) \\ \wedge & (\forall n, v. \neg \text{vote}(n, \text{none}, v)) \\ \wedge & (\forall n, r, r_{\max}, v. \text{onebmaxvote}(n, r, r_{\max}, v) \implies \text{oneb}(n, r)) \\ \wedge & (\forall n, r, r_1, r_2. \text{oneb}(n, r_2) \wedge r_2 > r_1 \implies \text{lefttrnd}(n, r_1)) \\ \wedge & (\forall n, r_1, r_2, v_1, v_2. \text{onebmaxvote}(n, r_2, \text{none}, v_1) \wedge r_2 > r_1 \implies \neg \text{vote}(n, r_1, v_2)) \\ \wedge & (\forall n, r, r_{\max}, v. \text{onebmaxvote}(n, r, r_{\max}, v) \wedge r_{\max} \neq \text{none} \implies r > r_{\max} \wedge \text{vote}(n, r_{\max}, v)) \\ \wedge & (\forall n, r, r', r_{\max}, v. \text{onebmaxvote}(n, r, r_{\max}, v) \wedge r > r' \wedge r' > r_{\max} \implies \neg \text{vote}(n, r', v')) \\ \wedge & (\forall r_1, r_2, v_1, v_2, q. r_2 > r_1 \wedge \text{proposal}(r_2, v_2) \wedge v_1 \neq v_2 \implies \exists n. \text{member}(n, q) \wedge \text{lefttrnd}(n, r_1) \wedge \neg \text{vote}(n, r_1, v_1)) \end{aligned}$$



# SWISS Overview

SWISS

Abstract  
protocol  
description

Breadth

Finisher

Invariants & proof of  
safety condition

Cast a "wide net"  
Find *any* invariant

Find invariant to  
complete proof

Many small invariants

One big invariant

$I_1, I_2, I_3, \dots, I_n$

$I_{\text{last}}$

Safety  
property

Partial invariants

# SWISS Invariant Search

Exploring the space of candidate invariant predicates for Paxos

	Candidate invariant space	Number of candidate invariants	Symmetries	Counter-example filters	Removing redundant invariants	Invariant predicates
Finisher	6 terms	~ 99,000,000,000,000	~ 200,000,000,000	155	155	5
Breadth	3 terms	~ 820,000,000	~ 3,000,000	~ 900,000	2,250	801

100 ms on average  
Brute force is not feasible

**Counterexample-guided synthesis:**  
When one predicate fails to be inductive, use it to narrow your search space.

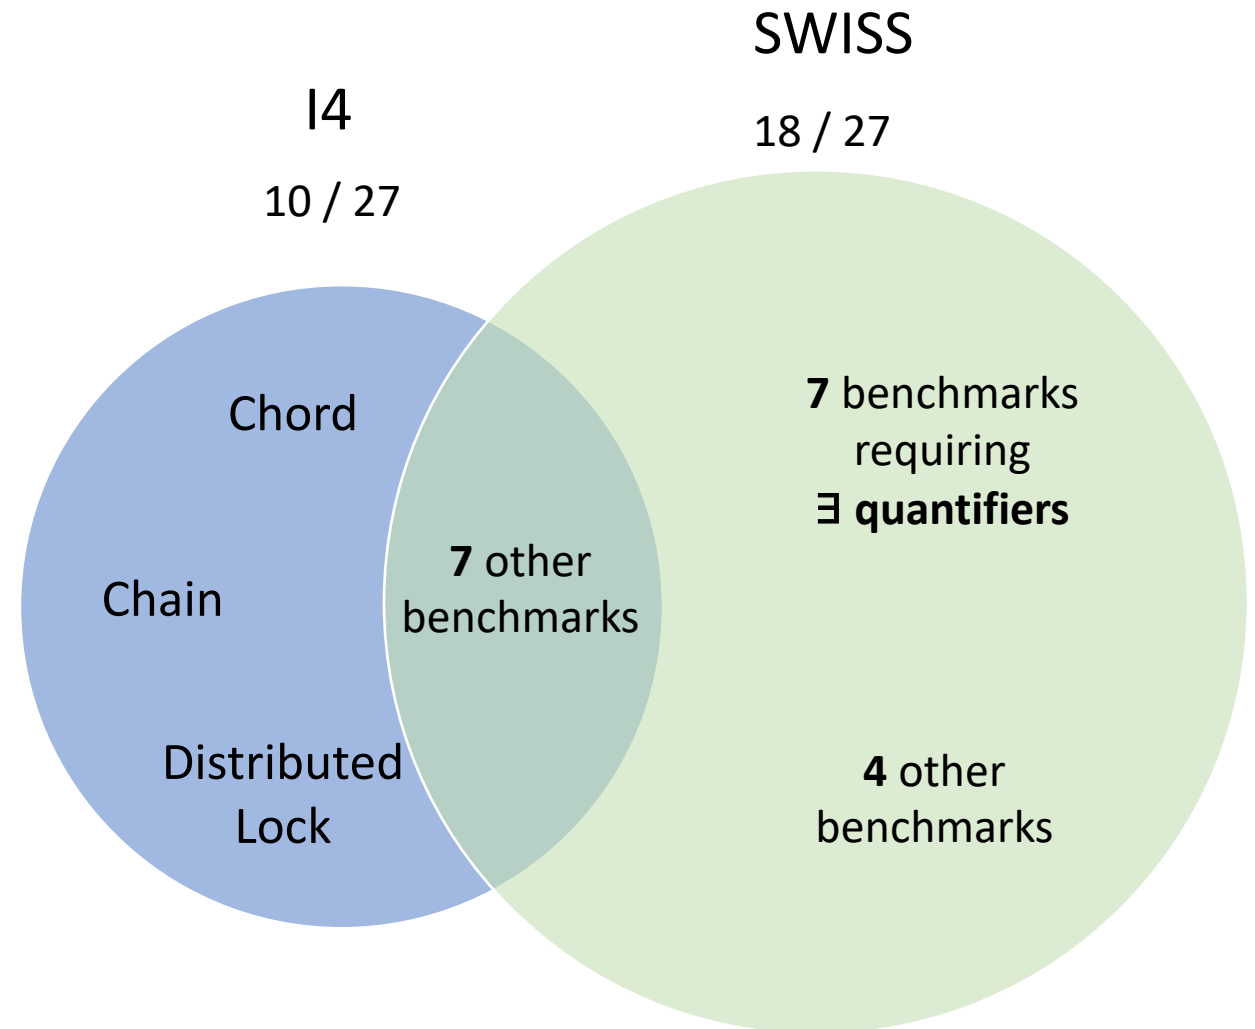
Finisher—which is directed by the desired safety property—is more effective at filtering a large space of candidate invariants.

# Evaluation

- Benchmark synthesis on 27 protocols, including 6 Paxos variants
  - SWISS solves 18 / 27 each within 6 hours
  - Includes **Paxos** and variant **Flexible Paxos**
  - Also solves **Multi Paxos** if given additional guidance on input search space

# Evaluation

I4 (2019) is usually the fastest, but doesn't handle existential quantifiers.



MA, H., GOEL, A., JEANNIN, J., KAPRITSOS, M., KASIKCI, B., AND SAKALLAH, K. A.

I4: incremental inference of inductive invariants for verification of distributed protocols.

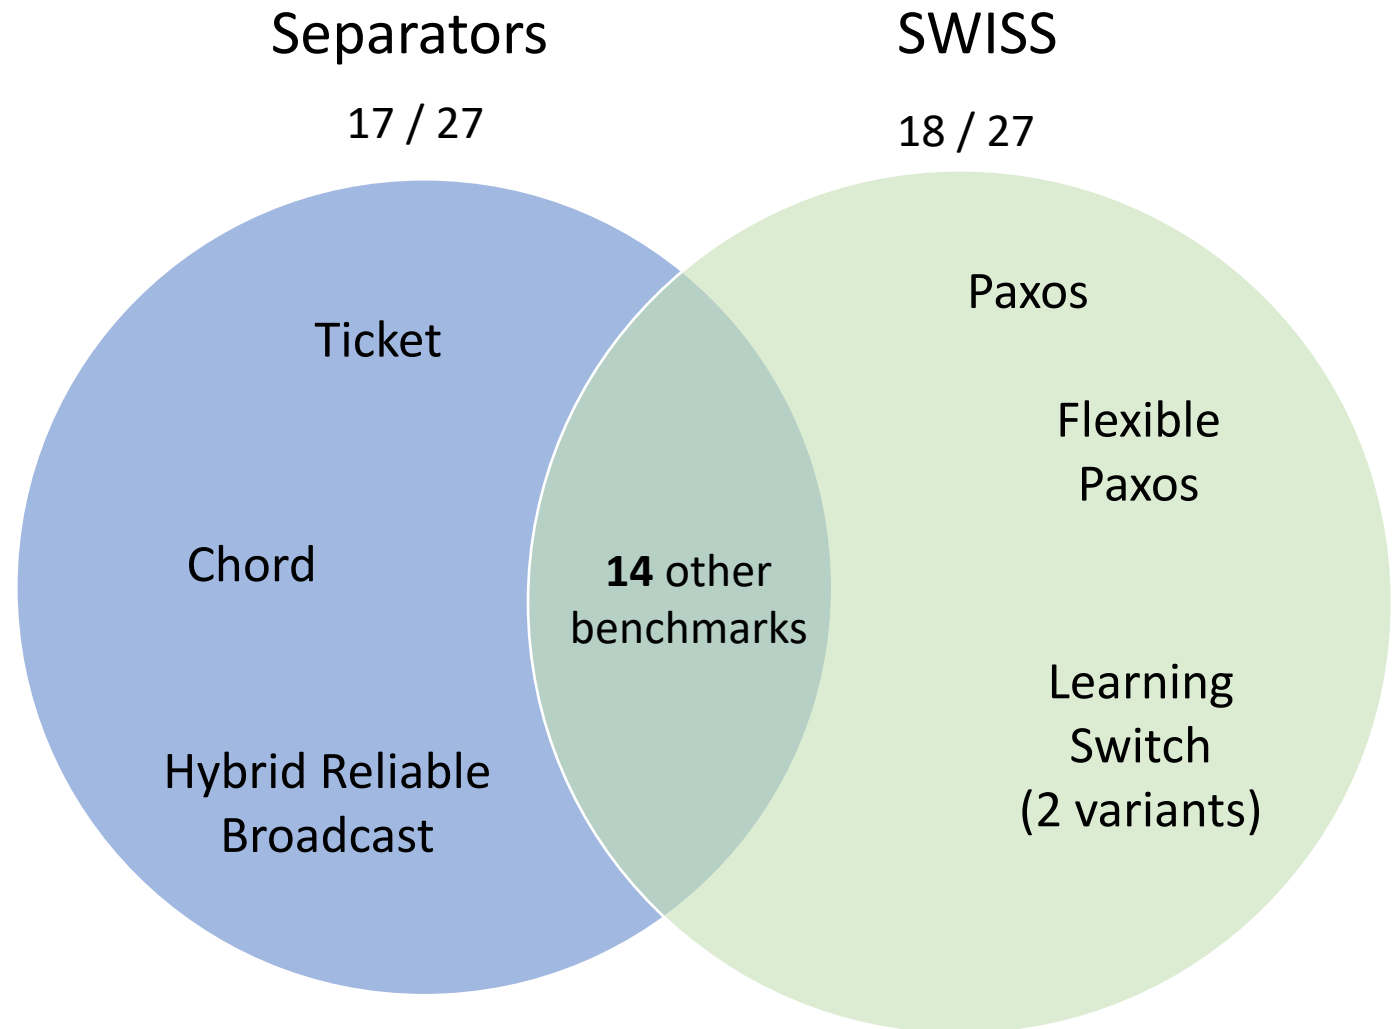
In *Proceedings of the ACM Symposium on Operating Systems Principles, (SOSP) (2019)*, T. Brecht and C. Williamson, Eds., ACM, pp. 370–384.

# Evaluation

Padon et al.'s Separators algorithm (2020) does not scale to Paxos, but is often faster on other benchmarks.

## Still out of reach

Fast Paxos  
Vertical Paxos  
Stoppable Paxos



# Further Evaluation in Paper

- Analysis of the sizes of invariants we expect harder protocols to require
- Benchmarks of individual optimizations
  - Optimizations that didn't help
- Parallelizability
- SMT bottlenecks
- Performance on restricted search spaces

# Conclusion

- SWISS scales invariant synthesis to protocols not tackled previously
- SWISS has differing strengths relative to prior approaches—suggests there are still ideas that can be combined and improved

*Finding Invariants of Distributed Systems: It's a Small (Enough) World After All*

Travis Hance

thance@andrew.cmu.edu