## Verification of Distributed Protocols: Decidable Modeling and Invariant Inference

Oded Padon VMware Research

Tutorial at FMCAD 2022 18 October 2022

# Verification of Distributed Protocols: Decidable Modeling and Invariant Inference

Part 1: Decidable Modeling

- The Ivy deductive verification system
- The many-sorted EPR fragment
- Main challenge: expressing interesting systems and properties in EPR
  - Expressing transitive closure
  - Expressing sets and cardinalities
  - Liveness and temporal verification

#### Part 2: Invariant Inference

#### • Problem setting

# Distributed protocols – excellent opportunity for verification

- Distributed systems are everywhere
  - Safety-critical systems
  - Internet scale services
  - Cloud infrastructure



- Distributed systems are notoriously hard to get right
  - Even small protocols can be tricky
  - Bugs occur in rare scenarios
  - Testing is costly and not sufficient





### Automatic verification of infinite-state systems



### Automatic verification of infinite-state systems



#### **Deductive verification**



#### **Deductive verification**



## Inductive invariants



System S is **safe** if all the reachable states satisfy the property  $\neg Bad$ 

## Inductive invariants



System S is safe if all the reachable states satisfy the property  $\neg Bad$ 

System S is safe iff there exists an **inductive invariant** Inv :

Init  $\subseteq$  Inv (Initiation)if  $\sigma \in$  Inv and  $\sigma \rightarrow \sigma'$  then  $\sigma' \in$  Inv (Consecution)Inv  $\cap$  Bad =  $\emptyset$  (Safety)

# Counterexample To Induction (CTI)

- States  $\sigma, \sigma'$  are a CTI of Inv if:
- **σ'** ∉ Inv
- σ 🗆 σ'
- A CTI may indicate:
  - A bug in the system
  - A bug in the safety property
  - A bug in the inductive invariant
    - Too weak
    - Too strong



x := 1; y := 2; while \* do { assert ¬even[x]; x := x + y; y := y + 2; TR





x := 1; y := 2; while \* do { assert ¬even[x]; TR | x := x + y; y := y + 2; }







$$\ln v = y^2 - 2y - 4x + 4 = 0$$

x := 1; y := 2; while \* do { assert ¬even[x]; TR | x := x + y; y := y + 2; }



even[x]

#### **Deductive verification**



# Verifying distributed systems is hard

#### Verdi

#### Verification of Raft in Coq 50,000 lines of manual proof



#### IronFleet

Verification of Multi-Paxos in Dafny 12,000 lines and 3.7 person-years Uses solver for undecidable SMT checks



[SOSP'15] Hawblitzel et al. IronFleet: proving practical distributed systems correct

[PLDI'15] Wilcox et al. Verdi: a framework for implementing and formally verifying distributed systems

# Effects of undecidability

- The verifier may fail on tiny programs
- No explanation when tactics fails
  - Counterproofs
- The butterfly effect





# The challenge of automated deduction

from: Ferraiulo, Baumann, Hawblitzel, Parno, "Komodo: Using Verification to Disentangle Secure-enclave Hardware from Software", SOSP '17

The most frustrating recurring problem was proof instability [...] Timeouts are challenging to debug, because the solver generally fails to provide useful feedback [...] even once fixed, the proof may easily timeout again due to minor perturbations. Worse, minor changes can trigger timeouts in seemingly unrelated proofs

Gap: deductive power of automated provers is not translating into verification productivity

"

# Rich logic, Poor logic

Rich Man Poor Man

#### **Rich SMT Theories**

- Linear arithmetic, Bitvectors, Arrays, Strings, Datatypes, ...
- Great tools: Yices, Z3, CVC4, Boolector, MathSAT, SMTInterpol, ...
- Essential in Dafny, Sage, Klee, F\*, ....
- Hides complexity from the user
- But solvers are heuristic, unpredictable, and not transparent



# The Ivy deductive verification system – design hypotheses

PST

- Predictability: we can reasonably predict solver performance
- Stability (or continuity): small, irrelevant changes do not greatly affect performance
- Transparency: failures have understandable explanations Hypotheses:



[PLDI'16] O.P., Kenneth L. McMillan, Aurojit Panda, Mooly Sagiv, Sharon Shoham.

Ivy: Safety Verification by Interactive Generalization

[CAV'20] O.P., Kenneth L. McMillan. Ivy: A Multi-modal Verification Tool for Distributed Algorithms

## Design space for formal verification

Proof Assistants proof/code

Ultimately limited by human



proof/code: Verdi: ~10 IronFleet: ~4



Decidable deduction Finite counterexamples proof/code: ~0.2

#### Ultimately limited by undecidability

Decidable Models Model Checking Static Analysis

Automation

# lvy's principles

- Specify systems and properties in decidable fragment of first-order logic
  - Allows quantifiers to reason about unbounded sets
  - Decidable to check inductiveness
  - Finite counterexamples to induction, display graphically or textually
- Interact with the user to find inductive invariants
- Challenge: use restricted logic to verify interesting systems
  - Transitive closure: network topology
  - Sets and cardinalities: Paxos, Byzantine Fault Tolerance, Reconfiguration
  - Liveness and temporal properties

#### Effectively Propositional Logic – EPR a.k.a. Bernays-Schönfinkel-Ramsey class

- Limited fragment of first-order logic without theories

  - Relations, but no functions r(a, b)• Restricted quantifier prefix:  $\exists^* \forall^* \varphi_{QF} \quad \exists x \forall y. r(x, y)$
- Finite model property
  - A formula is satisfiable iff it has a model of size # constant symbols + # existential variables
- Complexity:
  - NEXPTIME-complete
  - $\Sigma_2^P$  if relation arities are bounded by a constant
  - NP if quantifier prefix is also bounded by a constant

F. Ramsey. On a problem in formal logic. Proc. London Math. Soc. 1930



r(a)

r(x,y)





**EPR Satisfiability** 

Skole  

$$\exists x, y. \forall z. r(x, z) \leftrightarrow r(z, y)$$

$$=_{SAT} \forall z. r(c_1, z) \leftrightarrow r(z, c_2)$$

$$=_{SAT} (r(c_1, c_1) \leftrightarrow r(c_1, c_2)) \wedge (r(c_1, c_2) \leftrightarrow r(c_2, c_2))$$

$$=_{SAT} (p_{11} \leftrightarrow p_{12}) \wedge (p_{12} \leftrightarrow p_{22})$$

# Effectively Propositional Logic – EPR a.k.a. Bernays-Schönfinkel-Ramsey class

- Limited fragment of first-order logic without theories
  - Relations, but no functions
  - Restricted quantifier prefix:  $\exists^* \forall^* \varphi_{QF}$
- Finite model property
  - A formula is satisfiable iff it has a model of size:
     # constant symbols + # existential variables
- Complexity:
  - NEXPTIME-complete
  - $\Sigma_2^P$  if relation arities are bounded by a constant
  - NP if quantifier prefix is also bounded by a constant

F. Ramsey. On a problem in formal logic. Proc. London Math. Soc. 1930







# Decidable Fragments in Ivy

- Many-sorted EPR: allows acyclic functions and quantifier alternations
  - E.g.,  $f: A \to B$  but not  $g: B \to A$
  - Maintains finite model property
  - Finite complete instantiations, supported by Z3
- QFLIA Quantifier Free Linear Integer Arithmetic
- FAU Finite Almost Uninterpreted [CAV'07]
  - Allow limited arithmetic + acyclic quantifier alternations
  - Maintains finite complete instantiations

#### [CAV'07] Ge & de Moura: Complete Instantiation for Quantified Formulas in Satisfiabiliby Modulo Theories





#### Logic-based deductive verification in Ivy



# Challenge: encoding into EPR

- Many-sorted EPR: allows acyclic functions and quantifier alternations
  - E.g.,  $f: A \to B$  without  $g: B \to A$
  - Maintains small model property of EPR
  - Finite complete instantiations
- But what can you possibly express in such a restricted logic?
  - Transitive closure over deterministic paths
  - Set cardinalities
  - Avoiding quantifier alternations
  - Encoding liveness and LTL

# Example: Leader election in a ring

- Unidirectional ring of nodes, unique numeric ids
- Protocol:
  - Each node sends its id to the next
  - Upon receiving a message, a node passes it (to the next) if the id in the message is higher than the node's own id
  - A node that receives its own id becomes a leader
- Theorem: at any given time there is at most one leader
  - Inductive?





[CACM'79] E. Chang and R. Roberts. An improved algorithm for decentralized extrema-finding in circular configurations of

# Example: Leader election in a ring

- Unidirectional ring of nodes, unique numeric ids
- Protocol:
  - Each node sends its id to the next
  - Upon receiving a message, a node passes it (to the next) if the id in the message is higher than the node's own id
  - A node that receives its own id becomes a leader
- Theorem: at any given time there is at most one leader
  - Inductive? NO





[CACM'79] E. Chang and R. Roberts. An improved algorithm for decentralized extrema-finding in circular configurations of

# Example: Leader election in a ring

• Unidirectional ring of nodes, unique numeric ids

• Protocol:

- Each node sends its id to the next
- Upon receiving a message, a node passes it (to the next) if the ic *Proposition:* This algorithm detects one and only one
  A not highest number.
- Theorem and the consistent direction of messages, any message must meet all other processes before it comes back to its initiator. Only one message, that with the highest number, will not encounter a higher number on its way around. Thus, the only process getting its own message back is the one with the highest number.



[CACM'79] E. Chang and R. Roberts. An improved algorithm for decentralized extrema-finding in circular configurations of

## Leader election protocol – first-order logic

≤ (ID, ID) – total order on node id's
btw (Node, Node, Node) – the ring topology
id: Node □ ID – relate a node to its unique id
pending(ID, Node) – pending messages
leader(Node) – leader(n) means n is the leader

Axiomatized in first-order logic

first-order structure



protocol state



 $\langle n_5, n_1, n_3 \rangle \in I(btw)$ 

## Leader election protocol – first-order logic

•≤ (ID, ID) – total order on node id's

•btw (Node, Node, Node) – the ring topology

•id: Node 🗆 ID – relate a node to its unique id

•pending(ID, Node) – pending messages

•leader(Node) – leader(n) means n is the leader

protocol state

Axiomatized in first-order logic

first-order structure



## Leader election protocol – first-order logic

•≤ (ID, ID) – total order on node id's

btw (Node, Node, Node) – the ring topology
id: Node 

ID – relate a node to its unique id

•pending(ID, Node) – pending messages

•leader(Node) – leader(n) means n is the leader

action send(n: Node)
 "s := next(n)"
 pending(id(n),s) := true

action receive(n: Node, m: ID)
 requires pending(m, n)
 if \* then
 pending(m, n) := false
 if id(n) = m then
 leader(n) := true
 else if id(n) ≤ m then
 "s := next(n)"
 pending(m, s) := true

TR:

 $\exists n,s: Node. "s = next(n)" \land \forall x:ID,y:Node. pending'(x,y)↔(pending(x,y)∨(x=id(n)∧y=s))$ 

Bad:

assert I0 =  $\forall x,y$ : Node. leader(x)  $\land$  leader(y)  $\rightarrow$  x = y
### Leader election protocol – inductive invariant

**Safety property:** I<sub>o</sub>

 $I_{0} = \forall x, y: Node. leader(x) \land leader(y) \rightarrow x = y$ 

**Inductive invariant:**  $Inv = I_0 \land I_1 \land I_2 \land I_3$ 

 $I_1 = \forall n_1, n_2$ : Node. leader  $(n_2) \rightarrow id[n_1] \leq id[n_2]$  The leader has the highest ID

 $I_2 = \forall n_1, n_2$ : Node. pending(id[n\_2], n\_2)  $\rightarrow$  id[n\_1]  $\leq$  id[n\_2] Only the leader can be self-pending

 $I_3 = \forall n_1, n_2, n_3$ : Node.  $btw(n_1, n_2, n_3) \land pending(id[n_2], n_1) \rightarrow id[n_3] \leq id[n_2]$ 

Cannot bypass higher nodes

≤ (ID, ID) – total order on node id's
btw (Node, Node, Node) – the ring topology
id: Node □ ID – relate a node to its unique id
pending(ID, Node) – pending messages
leader(Node) – leader(n) means n is the leader

### Leader election protocol – inductive invariant

Safety property: I<sub>o</sub>



## Principle: first-order abstractions

Concept	Intention	First-order abstraction
Node ID's	Integers	
Ring Topology	Next edges + Transitive closure	

## Key idea: representing deterministic paths

[Shachar Itzhaky PhD, SIGPLAN Dissertation Award 2016]

#### Alternative 1: maintain n

- n<sup>\*</sup> defined by transitive closure of r
- not definable in first-order logic

#### A ternative 2: maintain $n^*$

- n defined by transitive reduction of  $n^*$
- Unique due to outdegree  $\leq 1$
- Definable in first order logic  $n(x,y) \equiv n^*(x,y) \land x \neq y \land$  $\forall z. n^*(x,z) \rightarrow z = y \lor z = x$





For every class C of finite graphs above:

- Axioms for path relation universally quantified
- Successor formula 1 universal quantifier
- Update formulas for node / edge addition and removal **Dyn-EPR**
- Soundness Theorem Every graph of class C satisfies the axioms of C Edges agree with successor formula
- Completeness Theorem Every finite structure satisfying the axioms of C is isomorphic (paths and edges) to a graph of class C

# Sound and complete\* axiomatization of

deterministic paths

LineRing $\leq (x, y)$ btw(x, y, z)

For every class C of finite graphs above

- Axioms for path relation universally 4
- Successor formula 1 universal quantifier

Universally quantified formulas

finite model property

+ Completeness Thm. for finite structures

Sound and complete automatic deductive verification

- Update formulas for node / edge addition and removal Dyn-EPR
- Soundness Theorem Every graph of class C satisfies the axioms of C Edges agree with successor formula

Ac

• Completeness Theorem Every finite structure satisfying the axioms of C is isomorphic (paths and edges) to a graph of class C

# Challenge: encoding into EPR

• But what can you possibly express in such a restricted logic?

• Transitive closure over deterministic paths

Set cardinalities

- Avoiding quantifier alternations
- Encoding liveness and LTL

### Paxos



- Family of distributed consensus protocols
  - let nodes maintain consistent state under crashes and packet loss
- Variants for different tradeoffs and extra features
- Active research and extensive industry use
  - Pervasive approach to fault-tolerant distributed computing

# Sets and cardinalities in EPR

- Consensus algorithms use set cardinalities
  - Wait for messages from more than N / 2 nodes
- Set Cardinalities + Arithmetic + Uninterpreted  $\gg$  EPR?
- Insight: set cardinalities are used to get a simple effect
   Can be modeled in first-order logic and EPR!
- Axiomatize quorums in first-order logic: **sort** Quorum
  - relation m(Node, Quorum)
  - set membership (2<sup>nd</sup>-order logic in first-order)

axiom  $\forall q_1, q_2$ : Quorum.  $\exists n$ : Node.  $\boldsymbol{m}(n, q_1) \land \boldsymbol{m}(n, q_2)$ 

action propose(r:Round) { require  $|\{n \mid \varphi(n)\}| > \frac{N}{2}$ ... } action propose(r:Round) { require  $\exists q. \forall n. m(n,q) \rightarrow \varphi(n)$ ... }

**V**EPR

## Principle: first-order abstractions

Concept	Intention	First-order abstraction
Messages	Network semantics: dropping duplication reordering	<pre>relation start_msg(Round) relation join_msg(Node, Round, Round, Value) relation propose_msg(Round, Value) relation vote_msg(Node, Round, Value)</pre>
Quorums	Majority sets	
Byzantine Quorums	2/3 Majority	

## Paxos in first-order logic

			the second se	41	# voting, and the correspondence
	1 sort node, quorum, round, value	22	action start_round(r : round) {	42	# v is arbitrary if the node
	2	23	assume r ≠ ⊥	43	local maxr, $v := max \{(r')\}$
	3 relation ≤ : round, round	24	<pre>start_round_msg(r) := true</pre>	44	
	4 axiom total_order(≤)	25		45	propose $msp(r, y) := true$
	5 constant 1 : round	26	action JOIN_ROUND(n : node, r : round) {	46	h dans making the same
	6	27	assume r ≠ ⊥	47	action vors(n : node, r : )
	7 relation member : node, quorum	28	assume start_round_msg(r)	48	assume r ≠ ⊥
	8 <b>axiom</b> $\forall q_1, q_2$ : quorum. $\exists n$ : node. member $(n, q_1) \land$ member $(n, q_2)$	29	assume $\neg \exists r', r'', v. r' > r \land join_ack_msg(n, r', r'', v)$	49	assume propose msp(r, v
	9	30	# find maximal round in which n voted, and the corresponding vote.	50	assume and r' n. r
1	0 relation start_round_msg : round	31	# maxr = $\perp$ and v is arbitrary when n never voted.	51	vote $msg(\mathbf{n}, \mathbf{r}, \mathbf{v}) := true$
1	1 relation join_ack_msg : node, round, round, value	32	local maxr, $v := \max \{(r', v') \mid vote_msg(n, r', v') \land r' < r\}$	52	1
1	2 relation propose_msg : round, value	33	join_ack_msg(n, r, maxr, v) := true	53	action LEARN(n : node, r :
1	3 relation vote_msg : node, round, value	34		54	assume r ± 1
1	4 relation decision : node, round, value	35	action PROPOSE(r : round, q : quorum) {	55	# 2b from quorum q
1	5	36	assume r ≠ ⊥	56	assume Vn member n
1	6 init ∀r. ¬start_round_msg(r)	37	assume $\forall v. \neg propose_msg(r, v)$	57	decision(n.r. v) -= true
1	7 init $\forall n, r_1, r_2, v. \neg joln\_ack\_msg(n, r_1, r_2, v)$	38	# 1b from quorum q	58	1
1	8 init $\forall r, v. \neg propose\_msg(r, v)$	39	assume $\forall n. member(n, q) \rightarrow \exists r', \upsilon. join ack msg(n, r, r', \upsilon)$	20	,
1	9 init $\forall n, r, v. \neg vote_msg(n, r, v)$	40	# find the maximal round in which a node in the quorum reported		
2	0 init $\forall n, r, v, \neg decision(n, r, v)$				

41 # voting, and the corresponding vote. 42 # v is arbitrary if the nodes reported not voting. 43 local maxr, v := max { $(r', v') | \exists n. member(n, q)$ 44  $\land fotn_ack_msg(n, r, r', v') \land r' \neq \bot$ } 45 propose\_msg(r, v) := true # propose value v 46 } 47 action vore(n : node, r : round, v : value) { 48 assume r  $\neq \bot$ 49 assume propose\_msg(r, v) 50 assume  $\neg \exists r', r'', v. r' > r \land fotn_ack_msg(n, r', r'', v)$ 50 yote\_msg(n, r, v) := true 52 } 53 action LEARN(n : node, r : round, v : value, q : quorum) { 54 assume r  $\neq \bot$ 55 # 2b from quorum q 56 assume  $\forall n. member(n, q) \rightarrow vote_msg(n, r, v)$ 57 decision(n, r, v) := true 58 }

 $\forall n_1, n_2: \text{node}, r_1, r_2: \text{round}, \upsilon_1, \upsilon_2: \text{value}. \ decision(n_1, r_1, \upsilon_1) \land \ decision(n_2, r_2, \upsilon_2) \rightarrow \upsilon_1 = \upsilon_2$ 

 $\forall r : round, v_1, v_2 : value. \ propose\_msg(r, v_1) \land propose\_msg(r, v_2) \rightarrow v_1 = v_2$ 

 $\forall n : \text{node}, r : \text{round}, v : \text{value}. \ vote\_msg(n, r, v) \rightarrow propose\_msg(r, v)$ 

$$\forall r: \text{round}, v: \text{value.}(\exists n: \text{node. } decision(n, r, v)) \rightarrow \exists q: \text{quorum.} \forall n: \text{node. } member(n, q) \rightarrow vote\_msg(n, r, v)$$

 $\forall n : \text{node}, r, r' : \text{round}, v, v' : \text{value}. \textit{join}\_ack\_msg(n, r, \bot, v) \land r' < r \rightarrow \neg \textit{vote}\_msg(n, r', v')$ 

 $\forall n : \text{node}, r, r' : \text{round}, v : \text{value}. \textit{ join}\_ack\_msg(n, r, r', v) \land r' \neq \bot \rightarrow r' < r \land \textit{ vote}\_msg(n, r', v)$ 

 $\forall n : \text{node}, r, r', r'' : \text{round}, v, v' : \text{value}.join\_ack\_msg(n, r, r', v) \land r' \neq \bot \land r' < r'' < r \rightarrow \neg vote\_msg(n, r'', v')$ 

 $\forall n : \text{node}, v : \text{value}. \neg vote\_msg(n, \bot, v)$ 

 $\forall r_1, r_2 : \text{round}, \upsilon_1, \upsilon_2 : \text{value}, q : \text{quorum}. \ propose\_msg(r_2, \upsilon_2) \land r_1 < r_2 \land \upsilon_1 \neq \upsilon_2 \rightarrow \sigma_1 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_1 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_2 \rightarrow \sigma_2 \neq \sigma_2 \rightarrow \sigma_2$ 

 $\exists n : \text{node}, r', r'' : \text{round}, v : \text{value}. \ member(n,q) \land \neg \textit{vote}\_msg(n,r_1,v_1) \land r' > r_1 \land \textit{join}\_ack\_msg(n,r',r'',v)$ 



# Challenge: encoding into EPR

• But what can you possibly express in such a restricted logic?

- Transitive closure over deterministic paths
- Set cardinalities

Avoiding quantifier alternations

• Encoding liveness and LTL

# Quantifier alternation cycles

• Axiom



# Modularity



[PLDI 2018] Marcelo Taube, Giuliano Losa, Kenneth L. McMillan, O. P., Mooly Sagiv, Sharon Shoham, James R. Wilcox, Doug Woos

Madularity for decidability of deductive verification with applications to distributed systems

## Inductive invariant of Paxos

# safety property

```
invariant decision(N1,R1,V1) & decision(N2,R2,V2) -> V1 = V2
```

# proposals are unique per round

```
invariant proposal(R,V1) & proposal(R,V2) -> V1 = V2
```

# only vote for proposed values

```
invariant vote(N,R,V) -> proposal(R,V)
```

```
# decisions come from quorums of votes:
```

invariant forall R, V. (exists N. decision(N,R,V)) -> exists Q. forall N. member(N, Q) -> vote(N,R,V)

#### # properties of one\_b\_max\_vote

```
invariant one_b_max_vote(N,R2,none,V1) & ~le(R2,R1) -> ~vote(N,R1,V2)
invariant one_b_max_vote(N,R,RM,V) & RM ~= none -> ~le(R,RM) & vote(N,RM,V)
invariant one_b_max_vote(N,R,RM,V) & RM ~= none & ~le(R,RO) & ~le(RO,RM) -> ~vote(N,RO,VO)
```

# property of choosable and proposal

invariant ~le(R2,R1) & proposal(R2,V2) & V1 ~= V2 -> exists N. member(N,Q) & left\_rnd(N,R1) & ~vote(N,R1,V1)
# property of one\_b, left\_rnd

invariant one\_b(N,R2) & ~le(R2,R1) -> left\_rnd(N,R1)

Protocol	Model [LOC]	Invariant [Conjectures]		
Paxos	85	11	2.2	0.1
Multi-Paxos	98	12	2.6	0.1
Vertical Paxos*	123	18	2.2	0.2
Fast Paxos*	117	17	6.2	1.6
Flexible Paxos	88	11	2.2	0
Stoppable Paxos*	132	16	5.4	0.9

#### \*first mechanized verification Transformation to EPR reusable across all variants!

Protocol	Model [LOC]	Invariant [Conjectures]		
Paxos	85	11	2.2	0.1
Multi-Paxos	98	12	2.6	0.1
Vertical Paxos*	123	18	2.2	0.2
Fast Paxos*	117	17	6.2	1.6
Flexible Paxos	88	11	2.2	0
Stoppable Paxos*	132	16	5.4	0.9

Proof / code ratio:

IronFleet: ~4 Verdi: ~10 Ivy: ~0.2

#### \*first mechanized verification Transformation to EPR reusable across all variants!

Protocol	Model [LOC]	Invariant [Conjectures]		
Paxos	85	11	2.2	0.1
Multi-Paxos	98	12	2.6	0.1
Vertical Paxos*	123	18	2.2	0.2
Fast Paxos*	117	17	6.2	1.6
Flexible Paxos	88	11	2.2	0
Stoppable Paxos*	132	16	5.4	0.9

 $\mu$  – mean  $\sigma$  – std. deviation

#### \*first mechanized verification Transformation to EPR reusable across all variants!

Protocol	Model [LOC]	Invariant [Conjectures]		
Paxos	85	11	2.2	0.1
Multi-Paxos	98	12	2.6	0.1
Vertical Paxos*	123	18	2.2	0.2
Fast Paxos*	117	17	6.2	1.6
Flexible Paxos	88	11	2.2	0
Stoppable Paxos*	132	16	5.4	0.9

Rounds			т.о.		
2	1.2	0.1	0		
4	1.8	0.4	0		
8	107	129	30%		
16	229	110	70%		
Multi Payos in EOL					

#### \*first mechanized verification Transformation to EPR reusable across all variants!

Protocol	Model [LOC]	Invariant [Conjectures]		
Paxos	85	11	2.2	0.1
Multi-Paxos	98	12	2.6	0.1
Vertical Paxos*	123	18	2.2	0.2
Fast Paxos*	117	17	6.2	1.6
Flexible Paxos	88	11	2.2	0
Stoppable Paxos*	132	16	5.4	0.9

	Rounds			Т.О.
	2	186	123	50%
	4	300	0	100%
	8	300	0	100%
	16	300	0	100%
S	toppa	ble F	<b>Paxos</b>	s in FOL

#### \*first mechanized verification Transformation to EPR reusable across all variants!

#### Appendix: The Proof of Correctness

We now prove that Stoppable Paxos satisfies its safety and liveness these For clarity and conciseness, we write simple temporal logic for with two temporal operators:  $\Box$  meaning *always*, and  $\diamond$  meaning *i* ally [13]. We use a linear-time logic, so  $\Diamond$  can be defined by  $\Diamond F$ for any formula F. For a state predicate P, the formula  $\Box P$  asser P is an invariant, meaning that it is true for every reachable state temporal formula  $\Diamond \Box P$  asserts that at some point in the execution, Ifrom that point onward We define a predicate P to be *stable* iff it satisfies the following cor-

if P is true in any reachable state s, then P is true in any state real from s by any action of the algorithm. We let stable P be the assert is state predicate P is stable. It is clear that a stable predicate is inva-it is true in the initial state. Because stability is an assertion only reachable states s, we can assume that all invariants of the algorith true in state s when proving stability. Our proofs are informal, but careful. The two complicated, mul

proofs are written with a hierarchical numbering scheme in which the number of the  $y^{ch}$  step of the current level-x proof [9]. Although appear intimidating, this kind of proof is easy to check and helps to

#### A.1 The Proof of Safety

We now prove that *Consistency* and *Stopping* are invariants of Sto Paxos. First, we define:

 $NotChoosable(i, b, v) \triangleq$  $(\exists Q: \forall a \in Q: (bal[a] > b) \land (vote_{i}[a][b])$  $\forall (\exists j < i, w \in StopCmd : Done2a(j, b, w))$ 

 $\forall$  (( $v \in StopCmd$ )  $\land$  ( $\exists j > i, w : Done2a(j, b$ We next prove a number of simple invariance and st.

algorithm

 $1, \forall i, b, v : \Box (Chosen(i, b, v) \rightarrow Done2a(i, b, v))$ 2.  $\forall i, b, v, w : \Box((Done2a(i, b, v) \land Done2a(i, b, v)))$ 3.  $\forall i, b, a, v : \Box((note, [a][b] = v) \Rightarrow Done2a(i, b, b)$ 

easy.

#### **Stoppable Paxos**

Dahlia Malkhi Leslie Lamport

Lidong Zhou

April 28, 2008

have been chosen as the  $j^{\text{th}}$  command for some j < i. Although the basic idea of the algorithm is not complicated, getting the details right was not

> (1):5. Sqfc.41(i, b, u)' PROOF: We assume c < b and  $w \neq v$  and prove NotCheosoble(i, c, u)'. By Lemma 1.7, it suffices to prove NotCheosoble(i, c, w). We split the proof into two (2)1 CASE: mol2a(i h O) - T PROOF: (1)4 (substituting j ← i) implies NotChoosable(i, c, w). (2)2. CASE: sval2a(i, b, Q)  $\neq \top$ PROOF: Since c < b, we can break the proof into two sub-cases. (3)1. CASE: mbal2a(i, b, Q) < c < b

NotChoosable(j, c, w).

 $c \leq mbal2a(k, b, Q)$ ; and (4)1 and case assumption (3)3 imply  $w \in StopCmd$ . Therefore, NoReconfigBefore(k, mbal2a(k, b, Q)) implies

тт

0

PROOF: Assumption (1)1.4 and Lemma 3 imply NotChoosable(i, c, w) Proto: Assumption (1) (1) which terms as important entropy of the entropy of the

(1)6. NoReconfigBefore(i, b) (1)6. Noticcomplication(i, o) PROOF: We assume j < i, w ∈ StopCnd, and c ≤ b and we prove NotChoosable(j, c, w). By Lemma 1.7, it suffices to prove NotChoosable(j, c, w). Since c ≤ b, we need consider only the following two cases.

(2)1. CASE: b = c PROOF: Assumption (1)1.3 implies Done2a(i, b, v)'. Since i > i and From the state of the second state of the second state of the state of the second state of the state of the

(232. CASE: c < b (3)2. CASE:  $sval2a(j, b, Q) \neq \top$ 

(2) C. Cain: real2a(j, k, Q) ≠ T. Fluori: Pluse assumption (2); we have the kilowing two sub-mases. (1): C. Cain: real2a(j, k, Q) < c < 1 MatChooseNet(j, m)). As here we scenarpiden, and Lemma 3 imply MatChooseNet(j, m). (1), there we can be approximately a sub-real scenario (1), and (1), and (1), and (1), and (1), and (1), (1); C. Casz: c ≤ real2a(j, k, Q) Fluori: Assumption (1), 3 implies DN(j, k, Q). The (5)? case scenario (1), and plies  $w \neq svol2a(i, b, Q)$ . By the (3)2 case assumption and the defi-19

nition of scal2a, we then have  $w \neq val2a(j, b, Q)$ . The (4)2 case assumption (which implies  $mval2a(j, b, Q) \neq -\infty$ ) and (1)3 then imply NotChoosable(j, c, w). (1)7. NoneChoosableAfter(i, b, v) PROOF: We assume v ∈ StepCmd, j > i, c < b, and w any command and we prove NotChoosable(j, c, w). By Lemma 1.7, it suffices to prove NotChoosable(j, c, w). We split the proof into two cases. (2)1. CASE: sval2a(i, b, Q) = T (2)1. CASE: real2a(t, b, Q) = ⊤ PRODC: Assemption: (1)1.3 implies: E4(i, b, Q, v), so the assumption v ∈ StopCrof implies: E4(i, b, Q, v). The case assumption, the assumption j > i, and E4b(i, b, Q, v) imply sea2la(j, b, Q) = ⊤. The assumption c < b and step (1)4 then imply NotChoosable(j, c, w). (2)2. CASE: scal2a(i, b, Q)  $\neq \top$  $\langle 3 \rangle$ 1. sval2a(i, b, Q) = val2a(i, b, Q) = v PROOF: Assumption (1)1.3 implies E3(i, b, Q, v), which implies sad2a(i, b, Q) = v. The case assumption and the definition of sval2a then implies val2a(i, b, Q) = v. implies surfact, b, Q = v,  $\beta (2)$ . *Dworlda*(i, b, b (2), v), PROOF: (3)1, assumption (1)1A, and the definition of val2a imply  $ort_{i} [a](mb2(a)(i, b, Q)] = v$  for some acceptor e in Q, which by Lemma 1.3 implies *Dworlda*(i, b (2)(a) = v for some acceptor e in Q, which by Lemma 1.8, M by the assumption v < k, its affinise to consider the following two cases. (i) introducing the set of the state o (3)4. CASE:  $mbal2a(i, b, Q) \le c < b$ (4)2. NotChoosable(j, c, w) PROOF: (4)1 and case assumption (3)4 imply mbal2a(j, b, Q) < c < b. By assumption (1)1.4, Lemma 3 implies NotChoosable(j, c, w). Theorem 1 Gonsistency A new remain k = 0 -denotements, it suffices to assume Chosen (i, b, v) and Chosen(i, c, w) and to prove v = w. Without loss of generality, we can assume  $k \leq c$ . We then have two cases. 1. CASE: b < c PROOF: We assume  $v \neq w$  and obtain a contradiction. Lemma 1.1 and Chosen(i, c, w) imply Done2a(i, c, w). By Lemma 4, this implies

 $b, w \in StopCmd$  : NotChoosable(j, c, w)fter(i b =) 🚔  $nd) \rightarrow \forall j > i, c < b, w : NotChoosable(j, c, w)$  $\triangleq$  Done2a(i, b, v)  $\Rightarrow$  SafeAt(i, b, v)  $\land$  NoReconfigBefore(i, b)  $\land$  NoneChoosableAfter(i, b, v) safety proof is the following proof that PropIng is invariant. i, b, v : PropInv(i, b, v)) $(v, v, v \to rroprid(1, b, 0))$  PropInv(i, b, v) is true in the initial state because Done2a(...)We therefore need only show that it is stable. We do this by i true in a state s and proving it is true in state <math>t. For any state  $^{1}$  be its value in state s.  $\forall j, c, w$  : PropInv(j, c, w) i is an instance number, b a ballot number, v a command, and Q a quorum.  $s \rightarrow t$  is a Phase2a(i, b, v, Q) step. E1(b, Q)SafeAt(i, b, v)' NoReconfigBefore(1, b) .1 ter(i, b, v) Inv(i, b, v))', it suffices to prove it for a partic-emma 1.7 (the stability of NotChoosable(...))  $efore(i, b) \land NoneChoosebleAfter(i, b, v)$ egere (i, b)  $\land$  inductive consistent equation (i, b, v) false ) true. We can therefore assume  $s \rightarrow t$  is a guorum Q. Formula E1(b, Q) holds because it ase2a(i, b, v, Q) action. HOVE" clause of (1)1 are proved as steps (1)5. se steps are used in their proofs.  $\Rightarrow Done2a(j, \ mbal2a(j, b, Q), \ val2a(j, b, Q))$ 

some more definitions, culminating in the key invariant

 $\stackrel{\text{\tiny b}}{=} \forall c < b, w \neq v : NotChoosable(i, c, w)$ 

e(i, b) 🚔

SafeAt(i, c, w). The assumptions b < c, an  $v \neq w$  then im-ply NotChoosable(i, b, v). By Lemma 2, this contradicts the assumption Chosen(i, b, v). 2. CASE: b = c

PROOF: Lemma 1.1 implies  $Done2a(i, b, v) \wedge Done2a(i, c, w)$ , which by Lemma 1.2 implies b = c.

Theorem 2 
Stopping

Theorem 2 C Stopping PROOF: By difficultion of Stopping, it suffices to assume Chesen(i, b, v), Channel(i, c, w) = 4 StopCode, and (>) and us obtain a contradiction. We aplit it is a sufficient of the suffi 2. CASE:  $c \ge b$ 

PROOF: Chosen(j, c, w) and Lemma 1.1 imply Done2a(j, c, w). Lemma 4 then implies NolleconfigBe(orc), c). The case assumption, the as-sumptions  $v \in StopCmd$  and j > i, and NoReconfigBe(orc), c) imply NotChoosable(i, b, v). The assumption Chosen(i, b, v) and Lemma 2 then provide the required contradiction

#### A.2 The Proof of Progress.

Theorem 3  $\forall b, Q$  : Progress(b, Q)PROOF: We assume P1(b, Q), P2(b, Q) and P3(b) and we must prove that there exists a v such that either  $\Diamond$  Chosen(i, b, v) or  $(v \in StopCmd) \land \Diamond$  Chosen(j, b, v). (1)1. ODE1(b. O)

(1) COR16.40 Phonor P11(a), Implies that the laftic & londer eventually constance as Phonor P11(b), mplies that the laftic & londer eventually receives the Planetic messages. Because heigh is set to a value c only by receiving a ballor, c message, assumption [24]) implies heigh [2] S. Heno, a mass eventually receive the Planetic message and escentre Planet1(c, b). By P2(b, Q), the Planet1 message) is essentially received by the loader.

(1)2.  $\forall i, w : \Box(Done2a(i, b, w) \Rightarrow \Diamond Chosen(i, b, w))$ PROOF: Done2a(i, b, w) means that a Phase2a(i, b, w) action has been executed Flower, Lowinz (1, b, w) means that it is also all (1, b, w) introduced to the down sector of section as covering a ( $^{2}\mathbb{Z}^{2}$ ), b, w)\_{1} measure to overy a coeptor. If it is in (2, then assumption  $P_{2}(b, q)$  implies that it eventually receives that measure, Assumption  $P_{3}(b)$  implies that it eventually receives  $P_{3}(b) = P_{3}(b, q)$  (in the section of the secti 21

vol2a(k, b, Q)) and (4)2 imply NoReconfiaBefore(k, m 18

PROOF: Assume  $mbal2a(j, b, Q) \neq -\infty$ . By definition of mbal2a, this implies val2a(j, b, Q) is a command (and not T). Since E1(b, Q) holds by assump-tion (1)1.4, the definitions of mbal2a and val2a imply that some acceptor a

tion (1)1.4, the definitions of mold2a and val2a imply that some acceptor a in Q has sent a (134' a, b, dm2da2h, b, Q); weight, b, Q); weights work, [dm2da2h, b, Q] with the message was sent, Lemma 1.3 beta implied. Dw2da2h, (v, dQ) with the message was sent. Lemma 1.3 beta implied. Dw2da2h, (v, dQ) with the message was sent. (1),  $V_{1,2} < A_{1,2} < A_{1,2}$ 

(1)4.  $\forall j, c < b, w : (sval2a(j, b, Q) = T) \Rightarrow NotChoosable(j, c, w)$ PRODE: We assume c < b and sval2a(j, b, Q) = T and prove NotChoosable(j, c, w). We split the proof into two cases.

PROOF: The case assumption implies mbal2a(j, b, Q) < c, so assumption (1)1.4 and Lemma 3 imply NotChooselle(j, c, w).

PROOF: Since c < b, we can split the proof into the following three cases. (3)1. CASE: mbal2a(j, b, Q) < c < bPROOF: By assumption (1)1.4, the case assumption and Lemma 3 imply

(4) I. vel2a(j, b, Q) ∈ StopOnd and w = vel2a(j, b, Q) (4) I. vel2a(j, b, Q) ∈ StopOnd and we can choose k > j such that mbal2a(k, b, Q) ≥ mbal2a(j, b, Q). PROOP: We deduce that val2a(j, b, Q) ∈ StopOnd and such a k exists

by the (2)2 case assumption, the assumption  $sral2n(i, b, O) = \top$ , and

(4)2. Done2a(k, mbal2a(k, b, Q), val2a(k, b, Q)) PROOF: The (3)3 case assumption and (4)1 imply mbal2a(k, b, Q)  $\neq -\infty$ . Step (1)2 then proves (4)2.

(4)3. NotChoosable(j, c, w) PROOF: Assumption (1)1.1 (with j ← k, c ← mbal2a(k, b, Q), and w ←

Step (4)1 asserts i < k; case assumption (3)3 and (4)1 impl

(3)2. CASE:  $c \le mbal2a(j, b, Q)$  and  $w \ne mal2a(j, b, Q)$ 

(3)3. CASE:  $c \leq mbel2a(j, b, Q)$  and w = ml2a(j, b, Q)

(2)1. CASE: mbal2a(j, b, Q) = -∞

(2)2. CASE:  $mbal2a(j, b, Q) \neq -\infty$ 

the definition of seal2a.

NotChoosable(j, c, w)

PROOF: By (1)3.

NotChonsoble(1 c w)

57

# Challenge: encoding into EPR

• But what can you possibly express in such a restricted logic?

- Transitive closure over deterministic paths
- Set cardinalities
- Avoiding quantifier alternations

• Encoding liveness and LTL

### Liveness properties

- Liveness property: "something good eventually happens"
- Often depend on fairness assumptions
- Typically proven by ranking functions, well-founded relations
  - Well beyond EPR, or not?

[POPL'18] O. P., Jochen Hoenicke, Giuliano Losa, Andreas Podelski, Mooly Sagiv, Sharon Shoham.
 Reducing Liveness to Safety in First-Order Logic
 [FMCAD'18] O. P., Jochen Hoenicke, Kenneth L. McMillan, Andreas Podelski, Mooly Sagiv, Sharon Shoham.
 Temporal Prophecy for Proving Temporal Properties of Infinite-State Systems.

## Lasso & Dynamic Abstraction

Finite State Parameterized



Liveness  $\Leftrightarrow$  No Lasso

# Lasso & Dynamic Abstraction

Finite State Parameterized



Infinite State Finite Abstraction



 $\mathsf{Liveness} \Leftrightarrow \mathsf{No} \ \mathsf{Lasso}$ 

Liveness ← No Lasso Problem: Spurious Lasso

# Lasso & Dynamic Abstraction

Finite State Parameterized



Infinite State Finite Abstraction



**Dynamic Abstraction** 

All expressible in

**EPR** 



Liveness ← No Lasso Fewer Spurious Lassos

Liveness  $\Leftrightarrow$  No Lasso

Liveness ← No Lasso Problem: Spurious Lasso

# Challenge: encoding into EPR

- But what can you possibly express in such a restricted logic?
  - Transitive closure over deterministic paths
  - Set cardinalities
  - Avoiding quantifier alternations
  - Encoding liveness and LTL

#### Logic-based deductive verification in Ivy



# Part 1: Conclusion

- Verification in EPR
  - Deduction is decidable
  - Finite counterexamples
  - Transparent failures
- Powerful encodings
  - Transitive closure for deterministic paths
  - Set cardinalities
  - Modular decomposition to battle ∀∃ cycles
  - Liveness and temporal properties





# Part 1: Open Questions

- What cannot be proven with EPR?
  - Proof theory
- Why are solvers stable on EPR?
  - Can it be generalized?
  - Part of why DPLL/CDCL works in practice
- Can it be used beyond a small group of fans?
  - Transitive closure ☺, LTL ☺, Cardinalities ☺, Breaking ∀∃ cycles ⊗
- Can you get the benefit of EPR without the costs?
  - Bounded quantifier instantiations
  - Special case of a more general "simple proofs" principle?

# Verification of Distributed Protocols: Decidable Modeling and Invariant Inference

Part 1: Decidable Modeling

- The Ivy deductive verification system
- The many-sorted EPR fragment
- Main challenge: expressing interesting systems and properties in EPR
  - Expressing transitive closure
  - Expressing sets and cardinalities
  - Liveness and temporal verification

#### Part 2: Invariant Inference

#### • Problem setting

# **Problem** setting

Several Ivy papers provide:

- A collection of examples with manually written invariants
- Invariant checking is decidable, typically ~1s, complex protocols ~10s
- Counterexamples are finite
- Invariants aren't too large or too complex

Protocol	Model [LOC]	Invariant [Conjectures]	EPR [sec] $\mu \sigma$	
Paxos	85	11	2.2	0.1
Multi-Paxos	98	12	2.6	0.1
Vertical Paxos	123	18	2.2	0.2
Fast Paxos	117	17	6.2	1.6
Flexible Paxos	88	11	2.2	0
Stoppable Paxos	132	16	5.4	0.9

#### *# safety property*

invariant decision(N1,R1,V1) & decision(N2,R2,V2) -> V1 = V2 *# proposals are unique per round* invariant proposal(R,V1) & proposal(R,V2) -> V1 = V2 # only vote for proposed values invariant vote(N,R,V) -> proposal(R,V) # decisions come from quorums of votes: invariant forall R, V. (exists N. decision(N,R,V)) -> exists Q. forall N. member(N, Q) -> vote(N,R,V) # properties of one\_b\_max\_vote invariant one b max\_vote(N,R2,none,V1) & ~le(R2,R1) -> ~vote(N,R1,V2) invariant one b max vote(N,R,RM,V) & RM ~= none -> ~le(R,RM) & vote(N,RM,V) invariant one\_b\_max\_vote(N,R,RM,V) & RM ~= none & ~le(R,RO) & ~le(RO,RM) -> ~vote(N,RO,VO) # property of choosable and proposal invariant ~le(R2,R1) & proposal(R2,V2) & V1 ~= V2 -> exists N. member(N,Q) & left rnd(N,R1) & ~vote(N,R1,V1) # property of one\_b, left\_rnd invariant one b(N,R2) & ~le(R2,R1) -> left rnd(N,R1)



UPDR [CAV'15, JACM'17] Aleksandr Karbyshev, Nikolaj Bjorner, Shachar Itzhaky, Noam Rinetzky and Sharon Shoham. Property-Directed Inference of Universal Invariants or Proving Their Absence.

[SOSP'19] Haojun Ma, Aman Goel, Jean-Baptiste Jeannin, Manos Kapritsos, Baris Kasikci, and Karem Sakallah. I4: incremental inference of inductive invariants for verification of distributed protocols.

FOL-IC3 [PLDI'20] Jason Koenig, O. P., Neil Immerman, and Alex Aiken. First-Order Quantified Separators.

[NSDI'21] Travis Hance, Marijn Heule, Ruben Martins, and Bryan Parno. Finding Invariants of Distributed Systems: It's a Small (Enough) World After All.

[OSDI'21] Jianan Yao, Runzhou Tao, Ronghui Gu, Jason Nieh, Suman Jana, and Gabriel Ryan. DistAI: Data-Driven Automated Invariant Learning for Distributed Protocols.

[NSF'21] Aman Goel and Karem Sakallah.

14

On Symmetry and Quantification: A New Approach to Verify Distributed Protocols.

FOL-IC3 [TACAS'22] Jason Koenig, O. P., Sharon Shoham, and Alex Aiken. Inferring Invariants with Quantifier Alternations: Taming the Search Space Explosion.

[OSDI'22] Jianan Yao, Runzhou Tao, Ronghui Gu, and Jason Nieh. DuoAI: Fast, Automated Inference of Inductive Invariants for Verifying Distributed Protocols.

# Interesting ideas

- Generalize from finite instances
  - |4
  - IC3PO
- Explicitly enumerate candidate invariants
  - SWISS
  - DistAl
  - DuoAl
- Generalize IC3/PDR by adapting lemma generation
  - UPDR
  - FOL-IC3



[SOSP'19] Haojun Ma, Aman Goel, Jean-Baptiste Jeannin, Manos Kapritsos, Baris Kasikci, and Karem Sakallah. I4: incremental inference of inductive invariants for verification of distributed protocols.

# Generalize from finite instances (IC3PO)



[NSF'21] Aman Goel and Karem Sakallah.

On Symmetry and Quantification: A New Approach to Verify Distributed Protocols.
# Explicitly enumerate invariants (SWISS)

- Typically  $Inv = p_1 \wedge \dots \wedge p_n$
- Each  $p_i$  isn't too complex, maybe we can explicitly enumerate all of them? Would that be useful?
- If we find some p that is inductive, we can learn it as an invariant
  - May help make additional p's inductive
- Ultimately, we want to prove the safety property

[NSDI'21] Travis Hance, Marijn Heule, Ruben Martins, and Bryan Parno. Finding Invariants of Distributed Systems: It's a Small (Enough) World After All.

# Explicitly enumerate invariants (SWISS)



[NSDI'21] Travis Hance, Marijn Heule, Ruben Martins, and Bryan Parno. Finding Invariants of Distributed Systems: It's a Small (Enough) World After All.

# Explicitly enumerate invariants (SWISS)

#### Exploring the space of candidate invariant predicates for Paxos

	Candidate invariant space	Number of candidate invariants	Symmetries	Counter-exa mple 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.

[NSDI'21] Travis Hance, Marijn Heule, Ruben Martins, and Bryan Parno.

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

# Explicitly enumerate invariants (DistAl, DuoAl)



[OSDI'21] Jianan Yao, Runzhou Tao, Ronghui Gu, Jason Nieh, Suman Jana, and Gabriel Ryan. DistAI: Data-Driven Automated Invariant Learning for Distributed Protocols.

[OSDI'22] Jianan Yao, Runzhou Tao, Ronghui Gu, and Jason Nieh.

DuoAI: Fast, Automated Inference of Inductive Invariants for Verifying Distributed Protocols.

# Adapt IC3/PDR (UPDR, FOL-IC3)

- Core task for lemma learning in IC3/PDR:
  - Given  $F = p_1 \land \dots \land p_n$  and state s, find p s.t.  $s \not\models p$  and  $post([F] \cap [p]) \subseteq [p]$
- UPDR: for *p* a forall-only formula, the diagram of *s* can be used
- FOL-IC3:
  - Closely related problem is separation: given  $s_1^+, ..., s_n^+$  and  $s_1^-, ..., s_m^-$  find p such that  $s_1^+, ..., s_n^+ \models p$  and  $s_1^-, ..., s_m^- \models \neg p$
  - For forall-exists formulas, separation is NP-complete
  - Use a SAT solver!
  - Many more techniques to scale to complex protocols

[CAV'15, JACM'17] Aleksandr Karbyshev, Nikolaj Bjorner, Shachar Itzhaky, Noam Rinetzky and Sharon Shoham. Property-Directed Inference of Universal Invariants or Proving Their Absence.

[PLDI'20] Jason Koenig, O. P., Neil Immerman, and Alex Aiken.

First-Order Quantified Separators.

[TACAS'22] Jason Koenig, O. P., Sharon Shoham, and Alex Aiken. Inferring Invariants with Quantifier Alternations: Taming the Search Space Explosion.

# Induction Duality and Primal-Dual Houdini

#### Induction Duality

• Formal symmetric connection between execution traces and incremental induction proofs

#### • Primal-Dual Houdini



- Houdini + Dual Houdini = Primal-dual invariant inference
- Interesting theoretical properties
- Promising empirical results for distributed protocols

[POPL'22] O. P., James R. Wilcox, Jason R. Koenig, Kenneth L. McMillan, and Alex Aiken. Induction Duality: Primal-Dual Search for Invariants.

# **Execution vs Incremental Induction**

}



a, b, c := 0, 0, 0
while \* {
 assert a ≥ 0
 a, b, c := a+b, b+c, c+1

Loop invariant:  $a \ge 0 \land b \ge 0 \land c \ge 0$ 



Predicates and Incremental Induction

# Formalizing symmetric incremental induction

- Intuition: if P is invariant, Q is also invariant
- Formally:  $post([P] \cap [Q]) \cap [P] \subseteq [Q] \iff P \land Q \vDash wp(P \rightarrow Q)$
- Visually:



Contrast with IC3/PDR relative inductiveness:
 post([P] ∩ [Q]) ⊆ [Q]



Predicates and Incremental Induction

#### **Executions vs Incremental Induction:** Symmetry (-1, -1, a = 0 $a \ge 0$ (0, -1, 0)b = 0(0,1,2) $b \ge 0$ (0,0,1) $c \ge 0$ (0,0,0)**States and Transitions**

Predicates and Incremental Induction Steps

### **Executions vs Incremental Induction: Symmetry**



#### **Executions vs Incremental Induction: Symmetry**



#### Induction Duality Structure $a \ge 0$ b = 0V: finite sets of states V': finite sets of predicates $\begin{array}{l} a \geq 0 \land \\ b \geq 0 \land \end{array}$ **Induction Duality** $c \ge 0$ $c \geq 0$ Induction over bounded proofs Induction over executions $E'_{\omega} = \{ (P, P \cup Q) \in V' \times V' \mid \cdots \}$ $E_{\omega} = \{ (S, S \cup T) \in V \times V \mid \cdots \}$ **Dual-Inductive** Inductive 5 Bounded incremental induction Transitions $E' = \{ (P, P \cup Q) \in E'_{\omega} \mid |Q| \le k \}$ $E = \{(S, S \cup \{t\}) \in V \times V \mid t \in post(S)\}$ Reachable k-provable (-1, -1, 1) $a \ge 0$ (0, -1, 0)b = 0(0,1,2)(0,0,1) $c \ge 0$ (0.0.0)



Reachable

k-provable

Induction over bound $E_{\omega} = \{(S, S \cup T) \in Dual - I_{\omega} \in U \}$	led proofs $V \times V \mid \cdots$ } nductive	Induction over e $E'_{\omega} = \{(P, P \cup e)\}$	executions $Q) \in V' \times V' \mid \cdots \}$ <i>Inductive</i> UI		
Transitions $E = \{(S, S \cup \{t\}) \in$	Transitions $E = \{(S, S \cup \{t\}) \in V \times V \mid t \in post(S)\}$ Reachable		Bounded incremental induction $E' = \{ (P, P \cup Q) \in E'_{\omega} \mid  Q  \le k \}$		
I			provable		
Concept	Induction d	ual	$\stackrel{\gamma}{\underset{\alpha}{\longleftarrow}} dual$		
	k-provable		invariant		
	reachable		k-abstractly-reachable		
	k-abstractly	-reachable	reachable		
	Invariant		k-provable		
	dual-inducti	ive			
	inductive				



Concept	Induction dual	$\xleftarrow{\gamma}{\alpha} dual$
reachable: <b>E</b> *	k-provable	invariant
k-provable: $E'^*$	reachable	k-abstractly-reachable
invariant: $\alpha(E^*)$	k-abstractly-reachable	reachable
k-abstractly-reachable: $\gamma(E'^*)$	invariant	k-provable
inductive: $E'_{\omega}$	dual-inductive	⊇reachable
dual-inductive: $E_{\omega}$	inductive	⊇k-provable

# Houdini and Dual Houdini



**Theorem:** all the states in  $S_H$  cannot be ruled out by k-bounded incremental induction (no matter how long)

# Primal-Dual Houdini



# **Galois Connection and Induction Duality**

Lattice of sets of states





Lattice of sets of predicates

## **Example: Toy Consensus Protocol**

- **v** (Node, Value) votes
- b(Node) bit used by a node to remember if it voted already
- **d**(Value) decisions
- *m*(*Node*, *Quorum*) membership

action vote(n: Node, x:value)
 requires ¬b(n)
 v(n,x) := true
 b(n) := true

**axiom**  $\forall q_1, q_2: Quorum. \exists n: Node. \mathbf{m}(n, q_1) \land \mathbf{m}(n, q_2)$ 

action decide(x:value) requires  $\exists q \forall n. \mathbf{m}(n,q) \rightarrow \mathbf{v}(n,x)$  $\mathbf{d}(x) := true$ 

Safety specification:

 $\forall x, y. \mathbf{d}(x) \land \mathbf{d}(y) \to x = y$ 

Inductive invariant:

$$\forall x. \mathbf{d}(x) \to \exists q \ \forall n. \mathbf{m}(n, q) \to \mathbf{v}(n, x)$$
  
 
$$\forall n, x, y. \ \mathbf{v}(n, x) \land \mathbf{v}(n, y) \to x = y$$
  
 
$$\forall n, x. \ \mathbf{v}(n, x) \to \mathbf{b}(n)$$

## Induction Edges

- Induction edge from P to Q if:  $post([P] \cap [Q]) \cap [P] \subseteq [Q] \quad \Leftrightarrow \quad P \land Q \vDash wp(P \rightarrow Q)$
- Intuition: if P is invariant, Q is also invariant
  - If a trace violates Q, it must also violate P (no later than the violation of Q)

$$\begin{array}{c} \forall n, x. v(n, x) \rightarrow b(n) \\ \hline \forall n, x, y. v(n, x) \wedge v(n, y) \rightarrow x = y \\ \hline \forall n, x, y. v(n, x) \wedge v(n, y) \rightarrow x = y \\ \hline \forall x. d(x) \rightarrow \exists q \ \forall n. m(n, q) \rightarrow v(n, x) \\ \hline \forall n, x. \neg v(n, x) \\ \hline \forall v. \neg d(v) \end{array}$$



#### Primal Houdini





#### **Dual Houdini**



$$\forall v_1, v_2. \mathbf{d}(v_1) \land \mathbf{d}(v_2) \rightarrow v_1 = v_2$$

$$\forall n. \neg \boldsymbol{b}(n) \qquad \qquad \forall n, v. \neg \boldsymbol{v}(n, v)$$

## Primal Houdini





#### **Dual Houdini**



$$\forall v_1, v_2. \mathbf{d}(v_1) \land \mathbf{d}(v_2) \rightarrow v_1 = v_2$$

$$(\top) \qquad \qquad \forall n, v. v(n, v) \rightarrow b(n)$$

#### Primal Houdini







 $\forall n, v. \mathbf{v}(n, v) \rightarrow \mathbf{b}(n)$ 

#### **Dual Houdini**

$$\forall v_1, v_2. \mathbf{d}(v_1) \land \mathbf{d}(v_2) \rightarrow v_1 = v_2$$





$$(\mathsf{T}) \qquad \forall n, v. v(n, v) \rightarrow \mathbf{b}(n)$$
$$\forall v. \mathbf{d}(v) \rightarrow \exists q \ \forall n. \mathbf{m}(n, q) \rightarrow \mathbf{v}(n, v)$$

#### Primal Houdini







 $\boldsymbol{m}(n_1, q_1)$  $\boldsymbol{m}(n_1,q_1)$  $v(n_1, v_1), v(n_1, v_2)$  $\boldsymbol{v}(n_1, v_2), \boldsymbol{v}(n_1, v_2)$  $d(v_1)$  $d(v_1), d(v_2)$  $b(n_1)$  $b(n_1)$ 

#### **Dual Houdini**

$$\forall v_1, v_2. \mathbf{d}(v_1) \land \mathbf{d}(v_2) \rightarrow v_1 = v_2$$







#### Primal Houdini





# Primal-Dual Houdini: Theorems

• Exploration: every iteration discovers new states and predicates

- Possible Progress:
  - If k-provable, every iteration can discover a new useful predicate
  - If unsafe, every iteration can discover a new useful state

• Termination by Stratification: if states and predicates are discovered in a stratified manner, then the algorithm terminates for k-provable or unsafe cases

# Evaluation for Distributed Protocols (k = 1)

Example	pdH	UPDR	FOL-IC3	SWISS	IC3PO	DistAl
ring		<b>√</b>				1
cons		1			1	1
paxos	1	1		1		
spaxos					1	
paxos-h	1					
spaxos-h						
locksrv	1	1	1	1	1	1
skv		1			1	1
skvr		1				
cache					1	

[POPL'22] O. P., James R. Wilcox, Jason R. Koenig, Kenneth L. McMillan, and Alex Aiken. Induction Duality: Primal-Dual Search for Invariants.

# Conclusion (Primal-dual Houdini)

- Goal: symmetric primal-dual invariant inference
- Key idea: bounded incremental induction
- Result 1: Induction duality
- Result 2: Primal-Dual Houdini
- New synthesis task: check dual-inductiveness and find induction edges
- Future directions
  - Application to more domains
  - Primal-dual version of more advanced algorithms



# Conclusion

- Using undecidable reasoning is powerful, but at a price
- Decidable logic (EPR) offers a different tradeoff
  - More effort in encoding systems and properties
  - More reliable automation
- Reliable invariant checking opens the path to invariant inference
  - Many recent ideas (and new adaptations of old ideas)
  - We are exhausting the benchmark set progress would come from new benchmarks!
- Induction duality
  - Promising new idea
  - May be applicable in other domains
- Reach out to discuss more or collaborate!



