

Verification of Distributed Protocols Using Decidable Logic

Sharon Shoham



Tel Aviv University

Programming Languages Mentoring Workshop 2019




The research leading to these results has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement No [759102-SVIS])

Why verify distributed protocols?

- Distributed systems are everywhere

- Safety-critical systems
- Cloud infrastructure
- Blockchains



- Distributed protocols are notoriously hard to get right

- Even small protocols can be tricky
- Bugs occur on rare scenarios
- Testing is costly and not sufficient

Verifying distributed protocols is hard

Verify distributed protocols for any number of nodes and resources




- Infinite state-space


- unbounded #processes
- unbounded #messages
- unbounded #objects

- Asymptotic complexity of verification

- Rice theorem




Safety of Infinite State Systems



System S is **safe** if all the **reachable** states satisfy the property $P = \neg \text{Bad}$

Inductive Invariants



System *S* is **safe** if all the **reachable** states satisfy the property $P = \neg \text{Bad}$


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

Init \subseteq *Inv* (**Initiation**)

if $\sigma \in \text{Inv}$ and $\sigma \rightarrow \sigma'$ then $\sigma' \in \text{Inv}$ (**Consecution**)

$\text{Inv} \cap \text{Bad} = \emptyset$ (**Safety**)

Inductive Invariants



System *S* is **safe** if all the **reachable** states satisfy the property $P = \neg \text{Bad}$

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


Init \subseteq *Inv* (**Initiation**)

if $\sigma \in \text{Inv}$ and $\sigma \rightarrow \sigma'$ then $\sigma' \in \text{Inv}$ (**Consecution**)

$\text{Inv} \cap \text{Bad} = \emptyset$ (**Safety**)


Example

- N pairs of players pass a ball:
 - $1\uparrow$ will pass to $1\downarrow$
 - $1\downarrow$ will pass to $1\uparrow$
 - $2\uparrow$ will pass to $2\downarrow$
 - $2\downarrow$ will pass to $2\uparrow$...




Example

- N pairs of players pass a ball:
 - $1\uparrow$ will pass to $1\downarrow$
 - $1\downarrow$ will pass to $1\uparrow$
 - $2\uparrow$ will pass to $2\downarrow$
 - $2\downarrow$ will pass to $2\uparrow$...



Example

- N pairs of players pass a ball:
 - $1\uparrow$ will pass to $1\downarrow$
 - $1\downarrow$ will pass to $1\uparrow$
 - $2\uparrow$ will pass to $2\downarrow$
 - $2\downarrow$ will pass to $2\uparrow$...
- The ball starts at player $1\uparrow$
- Can the ball get to $2\downarrow$?



Example


- N pairs of players pass a ball:

- $1\uparrow$ will pass to $1\downarrow$
- $1\downarrow$ will pass to $1\uparrow$
- $2\uparrow$ will pass to $2\downarrow$
- $2\downarrow$ will pass to $2\uparrow$...

- The ball starts at player $1\uparrow$

- Can the ball get to $2\downarrow$?

- Is “**the ball is not at $2\downarrow$** ” an inductive invariant?



Example


- N pairs of players pass a ball:

- $1\uparrow$ will pass to $1\downarrow$
- $1\downarrow$ will pass to $1\uparrow$
- $2\uparrow$ will pass to $2\downarrow$
- $2\downarrow$ will pass to $2\uparrow$...

- The ball starts at player $1\uparrow$

- Can the ball get to $2\downarrow$?


- Is “**the ball is not at $2\downarrow$** ” an inductive invariant? **No!**
 - **Counterexample to induction**



Example

- N pairs of players pass a ball:

- $1\uparrow$ will pass to $1\downarrow$
- $1\downarrow$ will pass to $1\uparrow$
- $2\uparrow$ will pass to $2\downarrow$
- $2\downarrow$ will pass to $2\uparrow$...




- The ball starts at player $1\uparrow$
- Can the ball get to $2\downarrow$?
- Is “**the ball is not at $2\downarrow$** ” an inductive invariant? **No!**
 - **Counterexample to induction**
- Inductive invariant: “**the ball is not at $2\uparrow$ nor $2\downarrow$** ”

Logic-based verification

Provers/solvers for different logics made huge progress

- Propositional logic (SAT) – industrial impact for hardware verification
- Satisfiability modulo theories (SMT) – major trend in software verification
- Automated first-order theorem provers
- Interactive theorem provers
- Z3, CVC4, iProver, Vampire, Coq, Isabelle/HOL

Logic-based verification



Represent *Init*, *Tr*, *Bad*, *Inv* by logical formulas: **Formula** \Leftrightarrow **Set of states**

$Inv(V)$ is an **inductive invariant** if the **verification conditions** (VCs) are valid:

Initiation $Init(V) \Rightarrow Inv(V)$ $unsat(Init(V) \wedge \neg Inv(V))$

Cons. $Inv(V) \wedge TR(V, V') \Rightarrow Inv(V')$ $unsat(Inv(V) \wedge TR(V, V') \wedge \neg Inv(V'))$

Safety $Inv(V) \Rightarrow \neg Bad(V)$ $unsat(Inv(V) \wedge Bad(V))$

Challenges for logic-based verification

Formal specification

Modeling the system and its invariants

Deduction

Checking validity of the VCs

Inference

Finding an inductive invariant

Deduction


$\text{Inv}(V)$ is an **inductive invariant** if the following **verification conditions** are valid:

Initiation $\text{Init}(V) \Rightarrow \text{Inv}(V)$ $\text{unsat}(\text{Init}(V) \wedge \neg \text{Inv}(V))$

Cons. $\text{Inv}(V) \wedge \text{TR}(V, V') \Rightarrow \text{Inv}(V')$ $\text{unsat}(\text{Inv}(V) \wedge \text{TR}(V, V') \wedge \neg \text{Inv}(V'))$

Safety $\text{Inv}(V) \Rightarrow \neg \text{Bad}(V)$ $\text{unsat}(\text{Inv}(V) \wedge \text{Bad}(V))$

Church's Theorem



Are the logical VC's valid ?

Counterexample



Unknown /
Diverge



Proof




Deduction

Interactive theorem provers (Coq, Isabelle/HOL, LEAN)

- Programmer proves the inductive invariant
- Huge programmer effort (~10-50 lines of proof per line of code)

e.g. Verdi



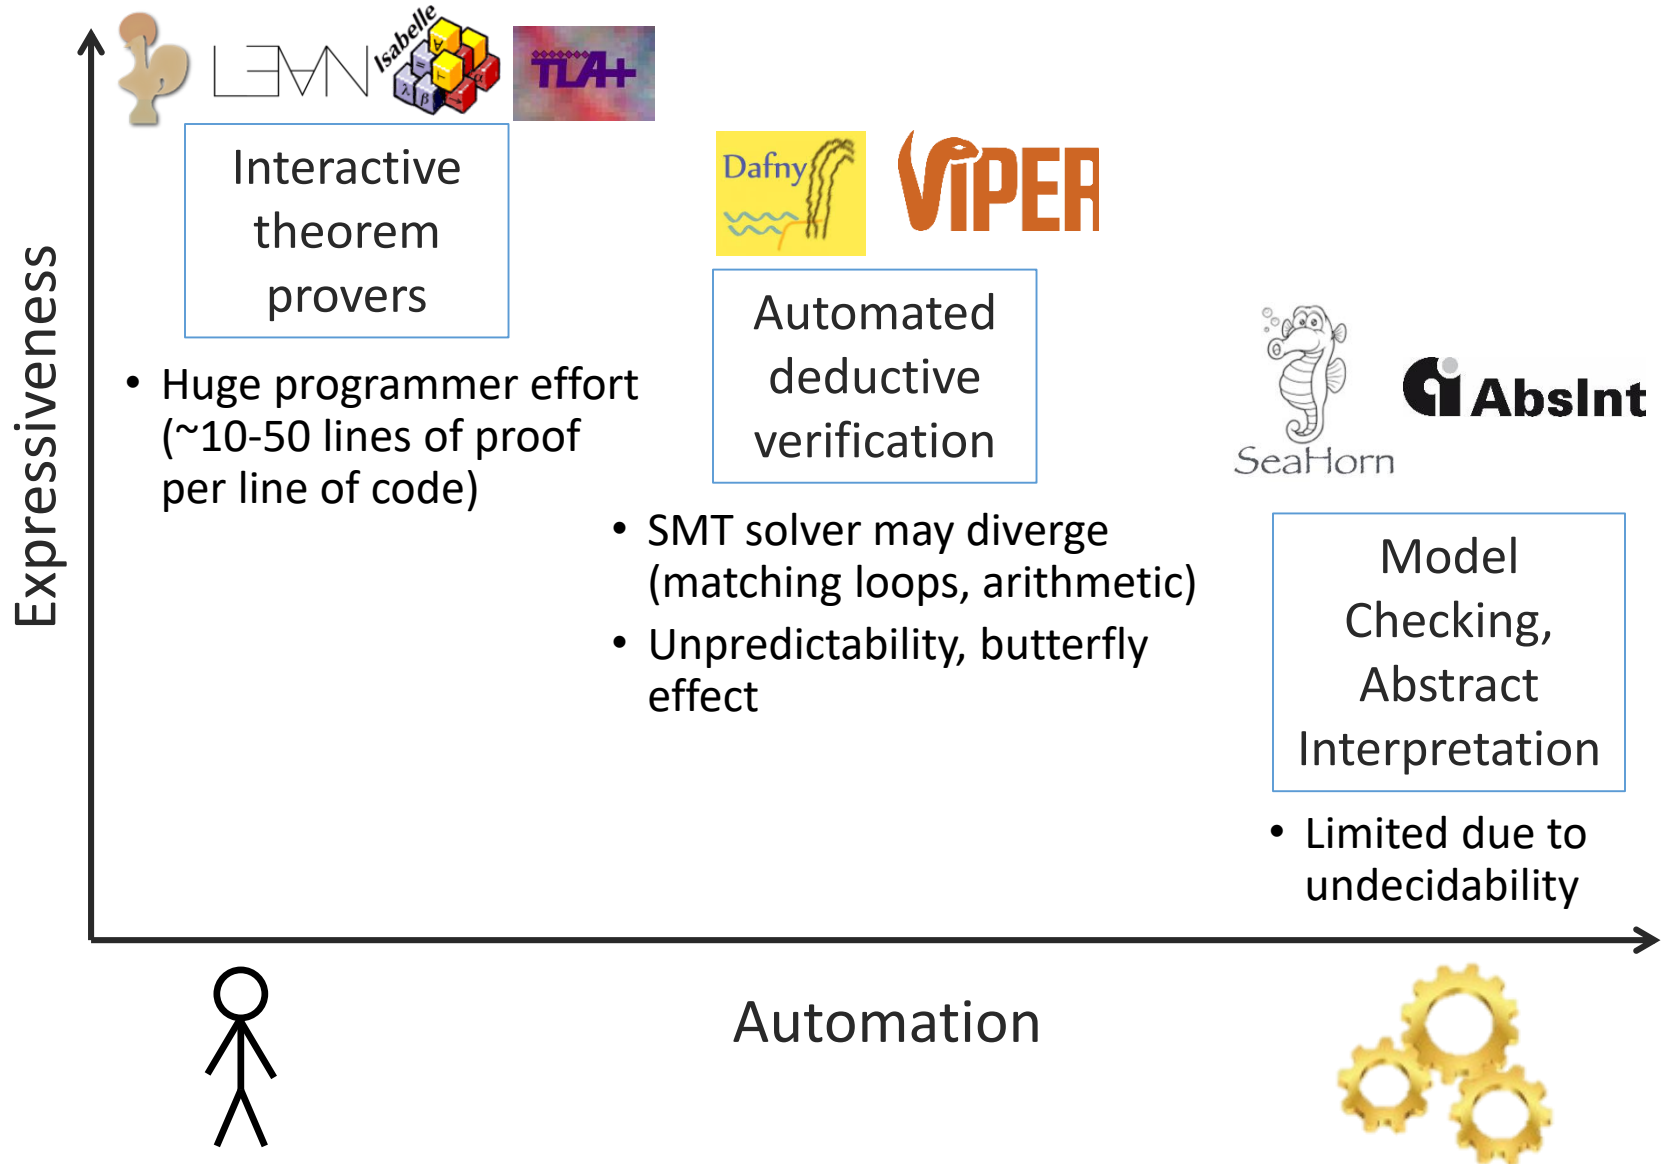
Automatic solvers/provers (e.g. Z3, CVC4, Vampire)

- VCs discharged automatically
- Tools may diverge (for SMT: matching loops, arithmetic)
- Unpredictability (butterfly effect)

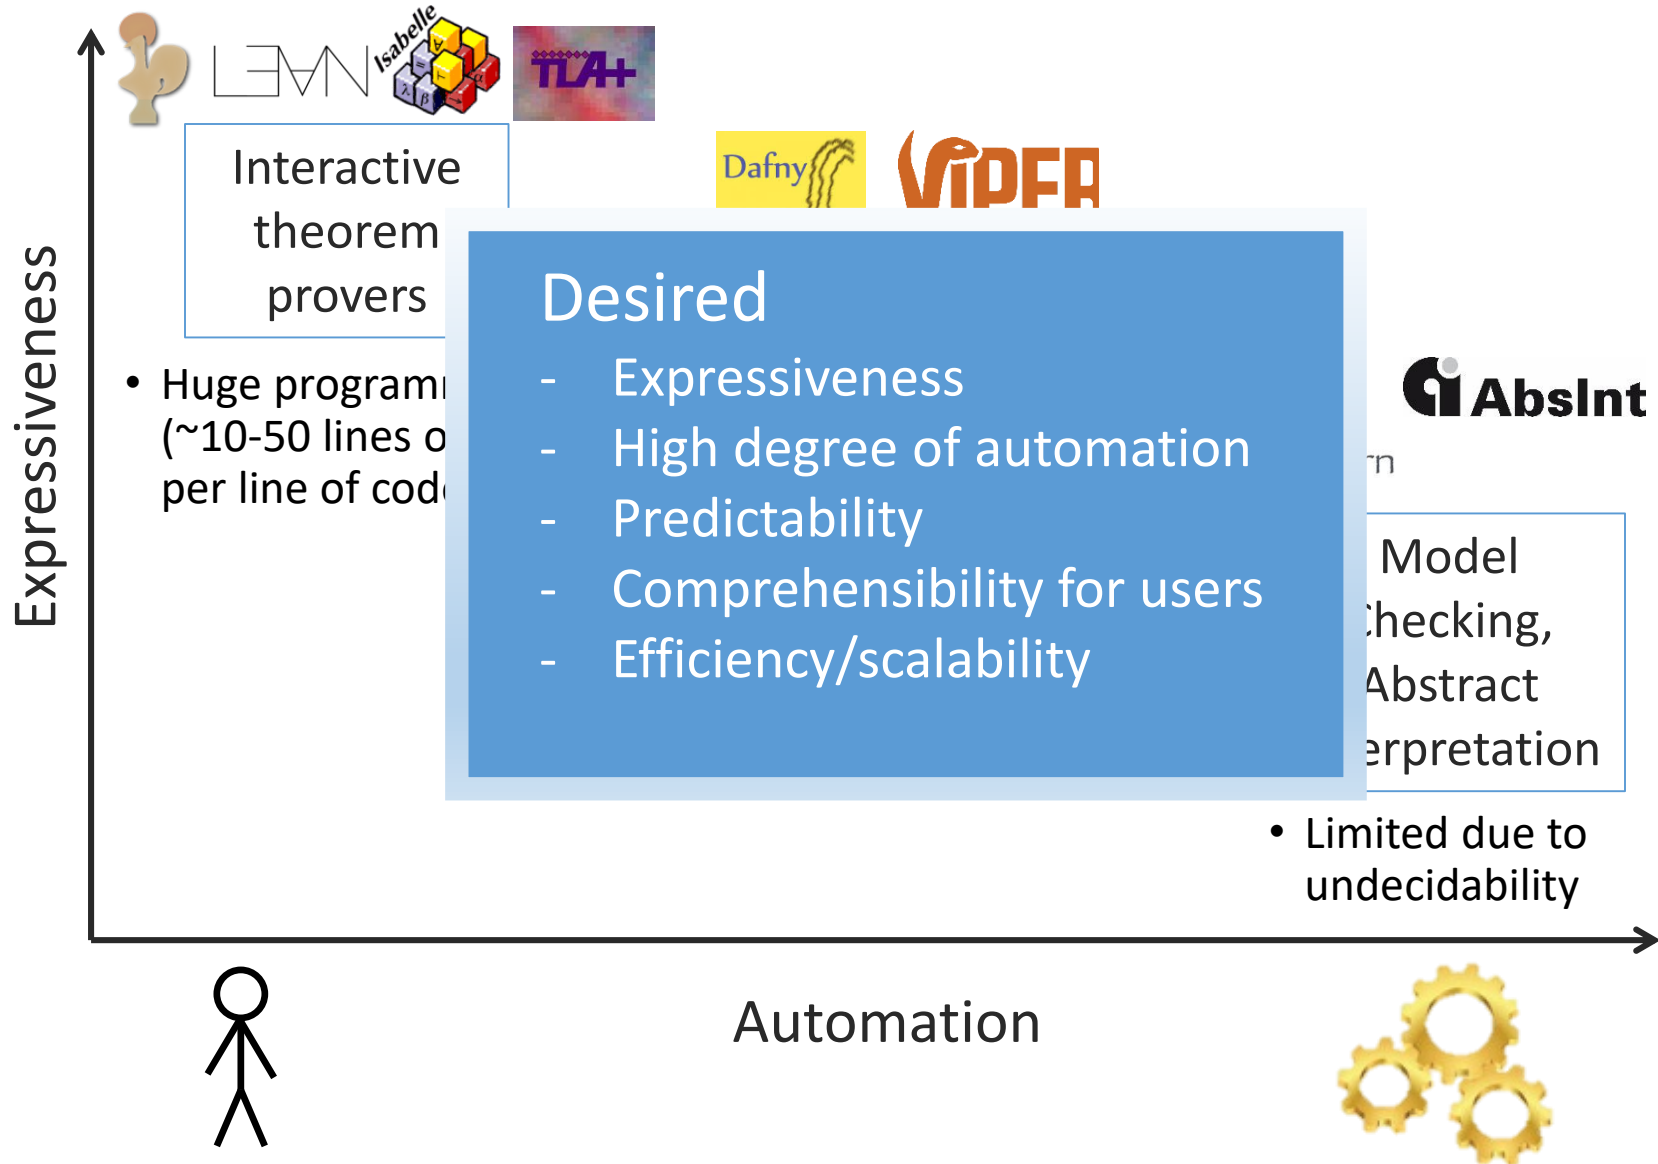
e.g. Ironfleet



Logic-based verification approaches



Logic-based verification approaches




This talk: Restrict VC's to decidable logic

$\text{Inv}(V)$ is an **inductive invariant** if the following **verification conditions** are valid:

Initiation $\text{Init}(V) \Rightarrow \text{Inv}(V)$ $\text{unsat}(\text{Init}(V) \wedge \neg \text{Inv}(V))$

Cons. $\text{Inv}(V) \wedge \text{TR}(V, V') \Rightarrow \text{Inv}(V')$ $\text{unsat}(\text{Inv}(V) \wedge \text{TR}(V, V') \wedge \neg \text{Inv}(V'))$

Safety $\text{Inv}(V) \Rightarrow \neg \text{Bad}(V)$ $\text{unsat}(\text{Inv}(V) \wedge \text{Bad}(V))$



Are the logical VC's valid ?

∈ Decidable logic

With good tool support

Counterexample



Proof



Challenges for verification with decidable logic

Formal specification

Modeling in a decidable logic

Deduction

Checking validity of the VC's



Invariant inference

Finding an inductive invariant

This talk

Logic: EPR – decidable fragment of first order logic

Formal specification

- Surprisingly expressive

Invariant inference

- Automatic (based on PDR)
 - Semi-algorithm: may diverge
- Interactive
 - Based on graphically displayed counterexamples to induction

Effectively Propositional Logic – EPR

Decidable fragment of first order logic

+ Quantification ($\exists^* \forall^*$) - Theories (e.g., arithmetic)

☺ Allows quantifiers to reason about unbounded sets

- $\forall x, y. \text{leader}(x) \wedge \text{leader}(y) \rightarrow x = y$

☺ Satisfiability is decidable \Rightarrow Deduction is decidable

☺ Small model property \Rightarrow Finite cex to induction

☺ Turing complete modeling language

☹ Limited language for safety and inductive invariants


➤ Suffices for many infinite-state systems

Successful verification with EPR

- Shape Analysis
[Itzhaky et al. CAV'13, POPL'14, CAV'14, CAV'15]
- Software-Defined Networks
[Ball et al. PLDI'14]
- Distributed protocols
[Padon et al. PLDI'16, OOPSLA'17, POPL'18, PLDI'18]
- Concurrent Modification Errors in Java programs
[Frumkin et al. VMCAI'17]

Example: Leader Election in a Ring

- Nodes are organized in a unidirectional ring
- Each node has a unique numeric id
- Protocol:
 - Each node sends its id to the next
 - A node that receives a message 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 the leader
- Theorem:
 - The protocol selects at most one leader




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


Modeling with EPR

- **State**: finite first-order structure over vocabulary V
 - \leq (ID, ID) – total order on node id's
 - **id**: Node \rightarrow ID – relate a node to its id
 - **btw** (Node, Node, Node) – the ring topology
 - **pending**(ID, Node) – pending messages
 - **leader**(Node) – leader(n) means n is the leader
- } Axiomatized in EPR

protocol state



structure



Modeling with EPR

- **State**: finite first-order structure over vocabulary V (+ axioms)
- **Initial** states and **safety** property: EPR formulas over V
 - $\text{Init}(V)$ – initial states, e.g., $\forall \text{id}, n. \neg \text{pending}(\text{id}, n)$
 - $\text{Bad}(V)$ – bad states, e.g., $\exists n_1, n_2. \text{leader}(n_1) \wedge \text{leader}(n_2) \wedge n_1 \neq n_2$
- **Transition relation**: expressed as EPR formula $\text{TR}(V, V')$, e.g.:
$$\exists n, s. "s = \text{next}(n)" \wedge \forall x, y. \text{pending}'(x, y) \leftrightarrow (\text{pending}(x, y) \vee (x = \text{id}[n] \wedge y = s))$$
$$\vee \exists n. \text{pending}(\text{id}[n], n) \wedge \forall x. \text{leader}'(x) \leftrightarrow (\text{leader}(x) \vee x = n)$$

...

Modeling with EPR

- **State**: finite first-order structure over vocabulary V (+ axioms)

Propose(n): send($\text{id}(n)$, next(n))

Recv(n ,msg): if msg = $\text{id}(n)$ then leader(n) := true

if msg > $\text{id}(n)$ then send(msg,next(n))

- **Transition relation**: expressed as EPR formula $\text{TR}(V, V')$, e.g.:

$\exists n, s. "s = \text{next}(n)" \wedge \forall x, y. \text{pending}'(x, y) \leftrightarrow (\text{pending}(x, y) \vee (x = \text{id}[n] \wedge y = s))$

$\vee \exists n. \text{pending}(\text{id}[n], n) \wedge \forall x. \text{leader}'(x) \leftrightarrow (\text{leader}(x) \vee x = n)$

...

Modeling with EPR


- **State**: finite first-order structure over vocabulary V (+ axioms)
- **Initial** states and **safety** property: EPR formulas over V
 - $\text{Init}(V)$ – initial states, e.g., $\forall \text{id}, n. \neg \text{pending}(\text{id}, n)$
 - $\text{Bad}(V)$ – bad states, e.g., $\exists n_1, n_2. \text{leader}(n_1) \wedge \text{leader}(n_2) \wedge n_1 \neq n_2$
- **Transition relation**: expressed as EPR formula $\text{TR}(V, V')$, e.g.:
$$\exists n, s. "s = \text{next}(n)" \wedge \forall x, y. \text{pending}'(x, y) \leftrightarrow (\text{pending}(x, y) \vee (x = \text{id}[n] \wedge y = s))$$
$$\vee \exists n. \text{pending}(\text{id}[n], n) \wedge \forall x. \text{leader}'(x) \leftrightarrow (\text{leader}(x) \vee x = n)$$

...

Modeling with EPR

- **State**: finite first-order structure over vocabulary V (+ axioms)
- **Initial** states and **safety** property: EPR formulas over V
 - $\text{Init}(V)$ – initial states, e.g., $\forall \text{id}, n. \neg \text{pending}(\text{id}, n)$
 - $\text{Bad}(V)$ – bad states, e.g., $\exists n_1, n_2. \text{leader}(n_1) \wedge \text{leader}(n_2) \wedge n_1 \neq n_2$

Specify and verify the protocol for **any** number of nodes in the ring



Using EPR for Verification

- System model $\text{Init}(V), \text{Bad}(V), \text{TR}(V, V') \in \text{EPR}$
- Inductive invariant $\text{Inv}(V) \in \forall^*$
- Verification conditions

Initiation	$\text{Init}(V) \Rightarrow \text{Inv}(V)$	$\text{unsat}(\text{Init}(V) \wedge \neg \text{Inv}(V))$
Cons.	$\text{Inv}(V) \wedge \text{TR}(V, V') \Rightarrow \text{Inv}(V')$	$\text{unsat}(\text{Inv}(V) \wedge \text{TR}(V, V') \wedge \neg \text{Inv}(V'))$
Safety	$\text{Inv}(V) \Rightarrow \neg \text{Bad}(V)$	$\text{unsat}(\text{Inv}(V) \wedge \text{Bad}(V))$

Verification conditions $\in \text{EPR}$


➔ Decidable to check

Inductive Invariant for Leader Election

Safety property:

$$I_0 = \neg \text{Bad} = \forall x, y: \text{Node}. \text{leader}(x) \wedge \text{leader}(y) \rightarrow x = y$$

Inductive?
No!



- \leq (ID, ID) – total order on node id's
- **btw** (Node, Node, Node) – the ring topology
- **id**: Node \rightarrow ID – relate a node to its id
- **pending**(ID, Node) – pending messages
- **leader**(Node) – leader(n) means n is the leader

Inductive Invariant for Leader Election

Safety property:

$$I_0 = \neg \text{Bad} = \forall x, y: \text{Node}. \text{leader}(x) \wedge \text{leader}(y) \rightarrow x = y$$

Inductive invariant: $\text{Inv} = I_0 \wedge I_1 \wedge I_2 \wedge I_3$

$$I_1 = \forall n_1, n_2: \text{Node}. \text{leader}(n_1) \rightarrow \text{id}[n_2] \leq \text{id}[n_1]$$

The leader has the highest id

$$I_2 = \forall n_1, n_2: \text{Node}. \text{pnd}(\text{id}[n_1], n_1) \rightarrow \text{id}[n_2] \leq \text{id}[n_1]$$

Only highest id can be self-pnd

$$I_3 = \forall n_1, n_2, n_3: \text{Node}. \text{btw}(n_1, n_2, n_3) \wedge \text{pnd}(\text{id}[n_2], n_1) \rightarrow \text{id}[n_3] \leq \text{id}[n_2]$$

Cannot bypass higher nodes

- \leq (ID, ID) – total order on node id's
- **btw** (Node, Node, Node) – the ring topology
- **id**: Node \rightarrow ID – relate a node to its id
- **pending**(ID, Node) – pending messages
- **leader**(Node) – leader(n) means n is the leader

The reason for using “**btw**” instead of “**next**”

Inductive Invariant for Leader Election

Safety property:

$$I_0 = \neg \text{Bad} = \forall x, y: \text{Node}. \text{leader}(x) \wedge \text{leader}(y) \rightarrow x = y$$

Inductive invariant: $\text{Inv} = I_0 \wedge I_1 \wedge I_2 \wedge I_3$

$$I_1 = \forall n_1, n_2: \text{Node}. \text{leader}(n_1) \rightarrow \text{id}[n_2] \leq \text{id}[n_1]$$

The leader has the highest id


$$I_2 = \forall n_1, n_2: \text{Node}. \text{pnd}(\text{id}[n_1], n_1) \rightarrow \text{id}[n_2] \leq \text{id}[n_1]$$

Only highest id can be self-pnd

$$I_3 = \forall n_1, n_2, n_3: \text{Node}. \text{btw}(n_1, n_2, n_3) \wedge \text{pnd}(\text{id}[n_1], n_1) \rightarrow \text{id}[n_3] \leq \text{id}[n_1]$$

Cannot bypass higher nodes

I can decide EPR!



$\text{Init}(V) \wedge \neg \text{Inv}(V)$
 $\text{Inv}(V) \wedge \text{TR}(V, V') \wedge \neg \text{Inv}(V')$
 $\text{Inv}(V) \wedge \text{Bad}(V)$

EPR Solver

Proof



Axioms: Leader Election Protocol

- \leq (ID, ID) – total order on node id's
- **btw** (a: Node, b: Node, c: Node) – the ring topology
- **id**: Node \rightarrow ID – relate a node to its unique id
- **pending**(ID, Node) – pending messages
- **leader**(Node) – leader(n) means n is the leader

	Intention	EPR Modeling
Node ID's	Integers	$\forall i: \text{ID}. i \leq i$ Reflexive $\forall i, j, k: \text{ID}. i \leq j \wedge j \leq k \rightarrow i \leq k$ Transitive $\forall i, j: \text{ID}. i \leq j \wedge j \leq i \rightarrow i = j$ Anti-Symmetric $\forall i, j: \text{ID}. i \leq j \vee j \leq i$ Total $\forall x, y: \text{Node}. \text{id}(x) = \text{id}(y) \rightarrow x = y$ Injective
Ring Topology	Next edges + Transitive closure	$\forall x, y, z: \text{Node}. \text{btw}(x, y, z) \rightarrow \text{btw}(y, z, x)$ Circular shifts $\forall x, y, z, w: \text{Node}. \text{btw}(w, x, y) \wedge \text{btw}(w, y, z) \rightarrow \text{btw}(w, x, z)$ Transitive $\forall x, y, w: \text{Node}. \text{btw}(w, x, y) \rightarrow \neg \text{btw}(w, y, x)$ Anti-Symmetric $\forall x, y, z, w: \text{Node}. \text{distinct}(x, y, z) \rightarrow \text{btw}(w, x, y) \vee \text{btw}(w, y, x)$
		$\text{"next}(a) = b"$ $\equiv \forall x: \text{Node}. x \neq a \wedge x \neq b \rightarrow \text{btw}(a, b, x)$

So far

Formal specification with EPR

- Surprisingly expressive
 - Integers: numeric id's expressed with \leq
 - Transitive closure: ring topology expressed with btw
 - Network semantics: pending messages
 - Sets and cardinalities (for consensus protocols) [OOPSLA'17]
 - Liveness properties [POPL'18, FMCAD'18]
 - Implementations [PLDI'18]

Not in
this talk

Next

Invariant inference: finding inductive invariants

(1) Automatically

- Adapt techniques from finite-state model checking (PDR)





(2) Interactively

- Based on graphically displayed counterexamples to induction

How can we find a **universally quantified** inductive invariant?

Inductive Invariant for Leader Election

I_0 $\neg \text{Bad}$	$\forall n_1, n_2 : \text{Node}. \text{leader}(n_1) \wedge \text{leader}(n_2) \rightarrow n_1 = n_2$ $\neg \exists n_1, n_2 : \text{Node}. \text{leader}(n_1) \wedge \text{leader}(n_2) \wedge n_1 \neq n_2$	At most one leader elected
I_1	$\forall n_1, n_2 : \text{Node}. \text{leader}(n_1) \rightarrow \text{id}[n_2] \leq \text{id}[n_1]$ $\neg \exists n_1, n_2 : \text{Node}. \text{leader}(n_1) \wedge \text{id}[n_2] > \text{id}[n_1]$	The leader has the highest id
I_2	$\forall n_1, n_2 : \text{Node}. \text{pnd}(\text{id}[n_1], n_1) \rightarrow \text{id}[n_2] \leq \text{id}[n_1]$ $\neg \exists n_1, n_2 : \text{Node}. \text{pnd}(\text{id}[n_1], n_1) \wedge \text{id}[n_2] > \text{id}[n_1]$	Only highest id can be self-pnd
I_3	$\forall n_1, n_2, n_3 : \text{Node}. \text{btw}(n_1, n_2, n_3) \wedge \text{pnd}(\text{id}[n_2], n_1) \rightarrow \text{id}[n_3] \leq \text{id}[n_2]$ $\neg \exists n_1, n_2, n_3 : \text{Node}. \text{btw}(n_1, n_2, n_3) \wedge \text{pnd}(\text{id}[n_2], n_1) \wedge \text{id}[n_3] > \text{id}[n_2]$	Cannot bypass higher nodes





Construct Inv by excluding “bad” states

1. How to find these states?
2. How to generalize into conjectures?

Generalization using Diagram

Use **diagrams** as generaliza


- state σ is a **finite** first-or

$$\begin{aligned} \text{Diag}(\sigma) = \exists \mathbf{x} \mathbf{y}. & \mathbf{x} \neq \mathbf{y} \wedge L(\mathbf{x}) \wedge \neg L(\mathbf{y}) \\ & \wedge \leq(\mathbf{x}, \mathbf{y}) \wedge \neg \leq(\mathbf{y}, \mathbf{x}) \\ & \wedge \leq(\mathbf{x}, \mathbf{x}) \wedge \leq(\mathbf{y}, \mathbf{y}) \end{aligned}$$

$\sigma' \models \text{Diag}(\sigma)$ iff σ is a substructure of σ'

σ is obtained from σ' by removing elements and projecting relations on remaining elements

$$\text{exclude}(\sigma) = \neg \text{Diag}(\sigma)$$




Generalization using Diagram :

Generalize even more if σ is a **partial** structure


$$\begin{aligned} \text{Diag}(\sigma) = \exists \textcolor{red}{x} \textcolor{green}{y}. & \textcolor{red}{x} \neq \textcolor{green}{y} \\ & \wedge \leq(\textcolor{red}{x}, \textcolor{green}{y}) \wedge \neg \leq(\textcolor{green}{y}, \textcolor{red}{x}) \\ & \wedge \leq(\textcolor{red}{x}, \textcolor{red}{x}) \wedge \leq(\textcolor{green}{y}, \textcolor{green}{y}) \end{aligned}$$

$$\text{exclude}(\sigma) = \neg \text{Diag}(\sigma)$$



\forall^* Invariant - excluded substructures


$$\text{Inv} \equiv \underbrace{\forall \bar{x}. (I_{1,1}(\bar{x}) \vee \dots \vee I_{1,m}(\bar{x})) \wedge \dots \wedge \forall \bar{x}. (I_{n,1}(\bar{x}) \vee \dots \vee I_{n,m}(\bar{x}))}_{\text{clause / conjecture}}$$




$$\text{Inv} \equiv \underbrace{\neg \exists \bar{x}. (\neg I_{1,1}(\bar{x}) \wedge \dots \wedge \neg I_{1,m}(\bar{x})) \wedge \dots \wedge \neg \exists \bar{x}. (\neg I_{n,1}(\bar{x}) \wedge \dots \wedge \neg I_{n,m}(\bar{x}))}_{\text{cube}}$$

Leader election example


At most
one
leader




The leader
has the
highest ID



Only the leader
can be self-
pending



Cannot bypass
higher nodes



substructure

How to find the (partial) states to generalize from?

(1) Automatic inference: UPDR

- Based on Bradley's IC3/PDR [VMCAI11,FMCAD11]
 - SAT-based verification of finite-state systems
- Abstracts concrete states using their logical diagram
- Backward traversal performed over diagrams
- Blocking of CTI excludes a *generalization* of its diagram → generates universally quantified lemmas

-
- [CAV'15, JACM'17] Property-Directed Inference of Universal Invariants or Proving Their Absence, A. Karbyshev, N. Bjorner, S. Itzhaky, N. Rinetzky and S. Shoham.
 - [VMCAI'17] Property Directed Reachability for Proving Absence of Concurrent Modification Errors, A. Frumkin, Y. Feldman, O. Lhoták, O. Padon, M. Sagiv and S. Shoham.

UPDR: Possible outcomes

- Universal inductive invariant found
 - System is safe


Used to infer **inductive invariants** / **procedure summaries** of:

- Heap-manipulating programs, e.g.
 - Singly/Doubly/Nested linked list
 - Iterators in Java - Concurrent modification error (CME)
- Distributed protocols
 - Spanning tree
 - Learning switch
 - ...

No need for
user-defined
predicates/
templates!

UPDR: Possible outcomes

- Universal inductive invariant found
 - System is safe
- Proof that no universal inductive invariant exists
 - Safety not determined*




* can use Bounded Model Checking to find real counterexamples

UPDR: Possible outcomes

- Universal inductive invariant found
 - System is safe
- Proof that no universal inductive invariant exists
 - Safety not determined*
- Divergence
 - In general, inferring universal ind. inv. is undecidable
 - For linked lists it is decidable, UPDR will also terminate
 - Proof uses well-quasi-order and Kruskal's tree theorem


-
- [POPL'16] Decidability of Inferring Inductive Invariants, O. Padon, N. Immerman, S. Shoham, A. Karbyshev, and M. Sagiv.

Automatic Inference (e.g., UPDR)



Ultimately limited by undecidability

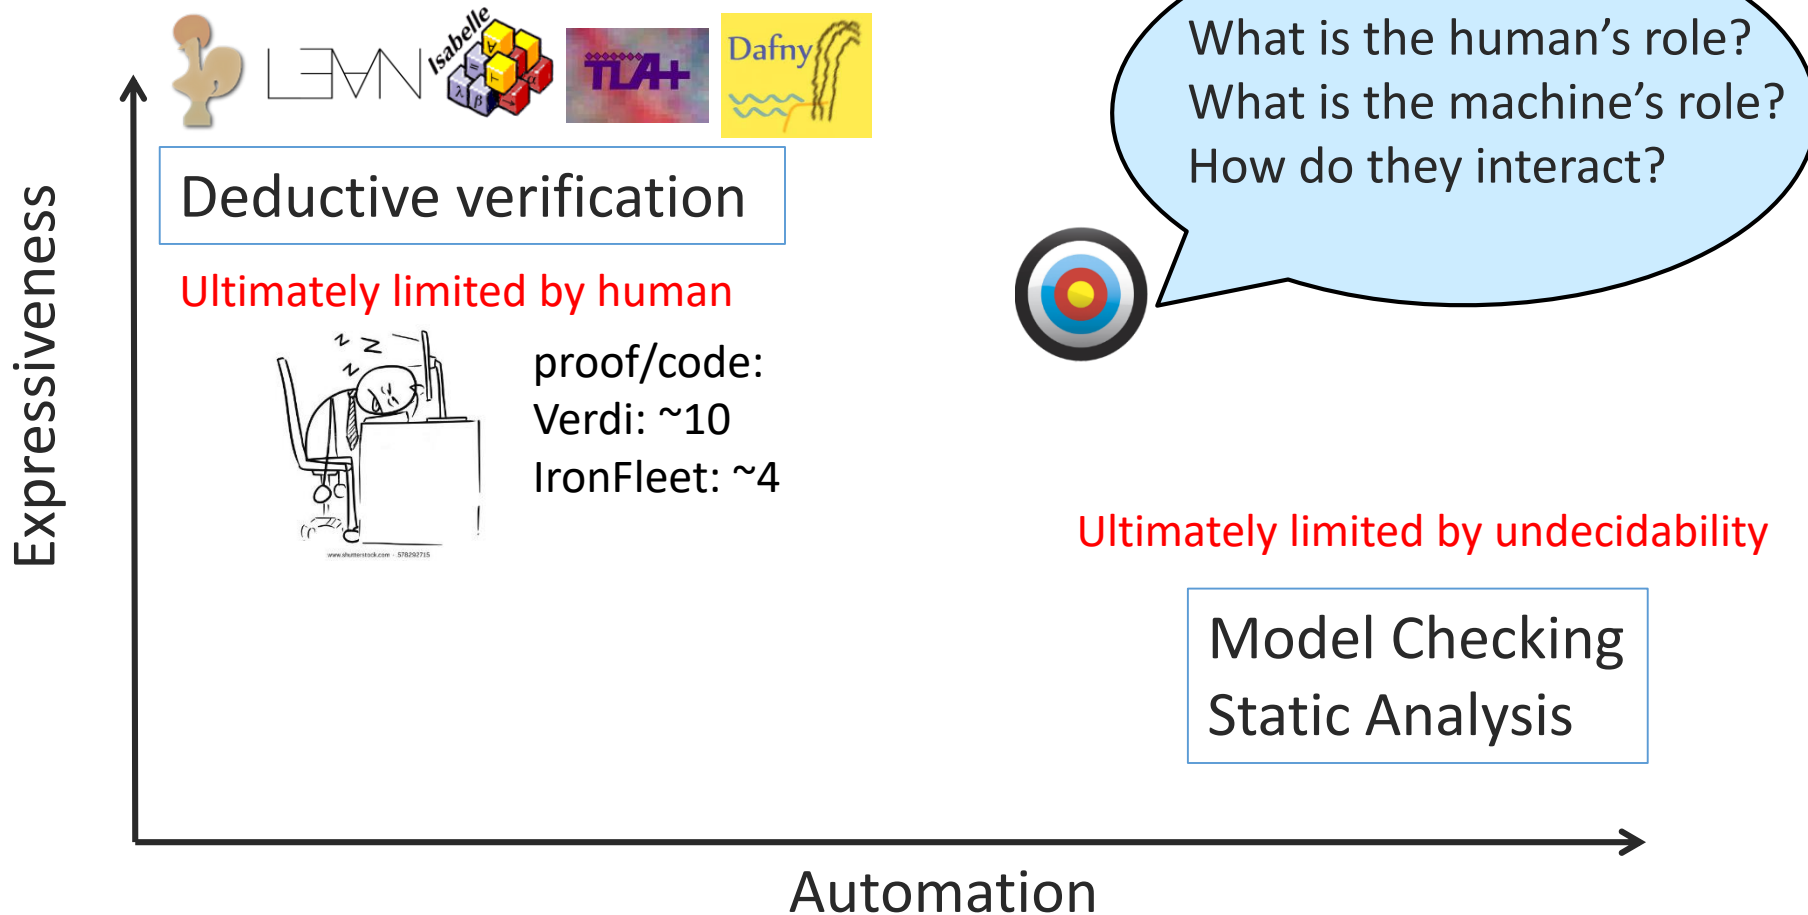
(2) Interactive Inference



- Let the user guide the tool
 - User has intuition about the essence of the proof
 - Computer is good at handling corner cases



Interactive Inference



erc

Supervised Verification of Infinite-State Systems

Ivy: Interactive Generalization

$$Inv = I_0 \wedge \cdots \wedge I_k$$



Displays “minimal” CTI to exclude

Generalizes to a partial state

- removes “irrelevant” facts (graphical interface - checkboxes)

Translates to universally quantified conjecture (via diagram)

Provides auxiliary automated checks:


1. BMC(K): uses SAT solver to check if conjecture is true up to K

- User determines the right K to use

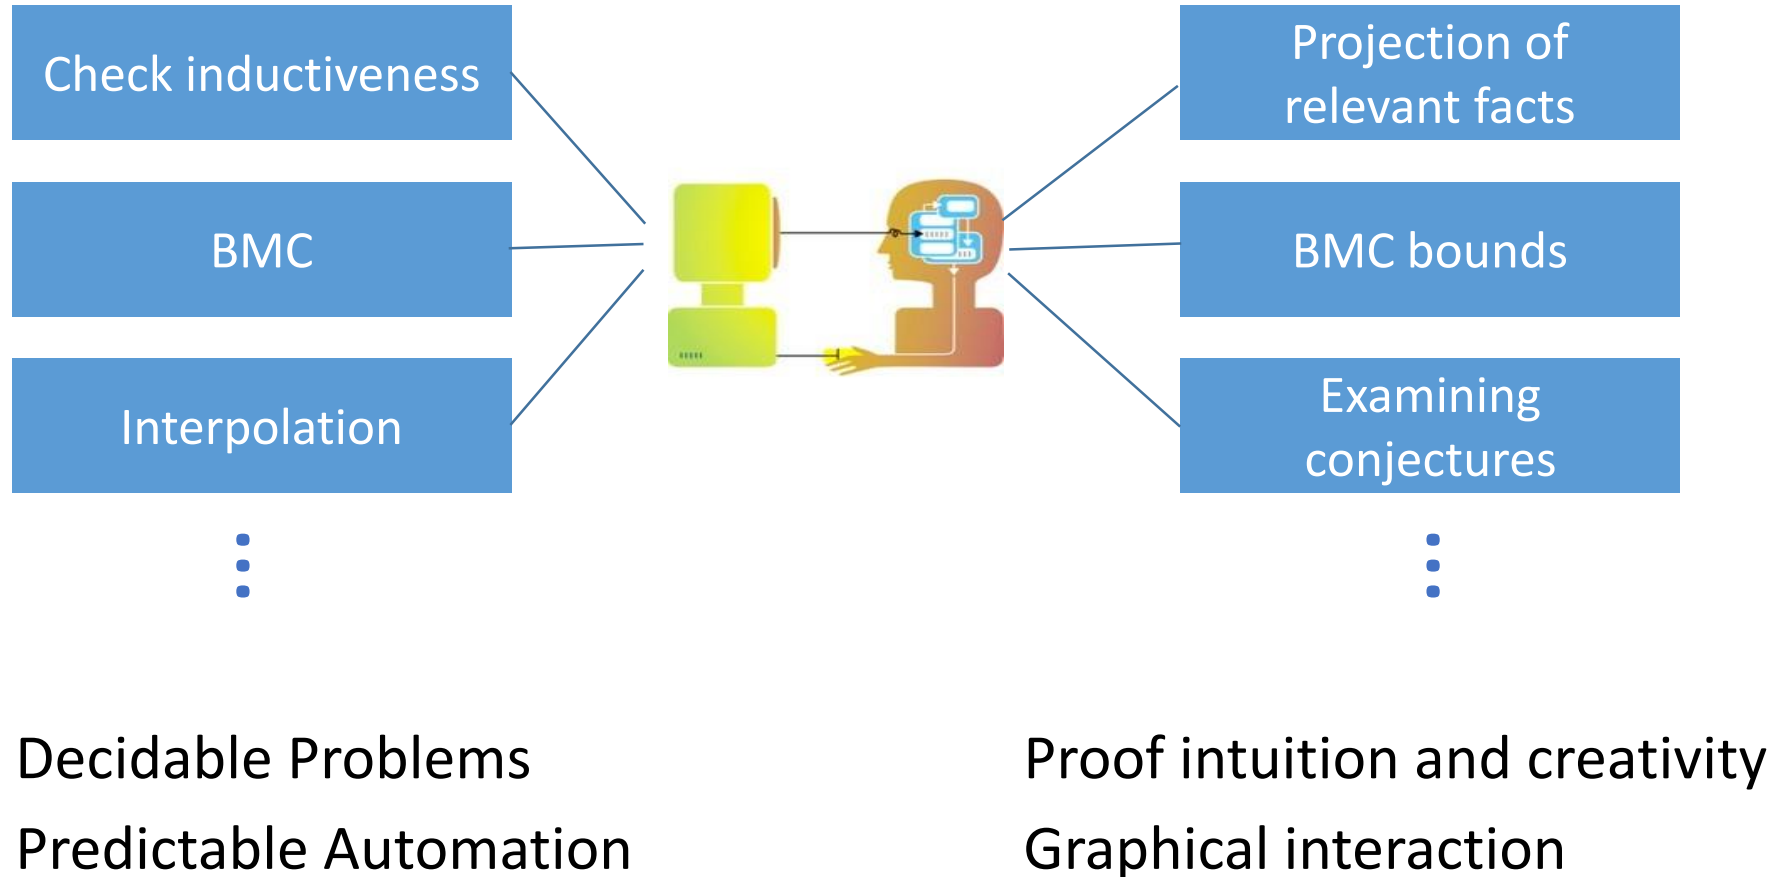
2. ITP(K): uses SAT solver to discover more facts to remove

Examines the proposed conjecture – it could be wrong

Adds I_{k+1}



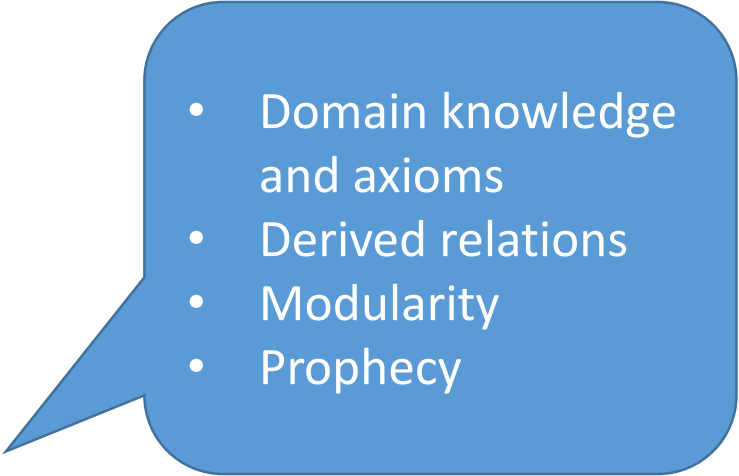
Interactive Verification in IVy



Summary 1

Verification with decidable logic

- EPR - decidable fragment of FOL
 - Deduction is decidable
 - Finite counterexamples
- Can be made surprisingly powerful
 - Transitive closure: linked lists, ring topology [PLDI'16]
 - Paxos, Multi-Paxos, [OOPSLA'17]
 - Liveness and Temporal Properties [POPL'18]
 - Developing verified implementations [PLDI'18]

- 
- Domain knowledge and axioms
 - Derived relations
 - Modularity
 - Prophecy

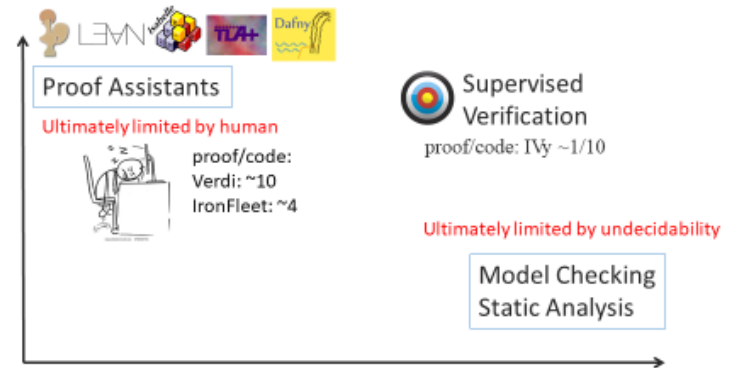
Summary 2

Invariant Inference

- Automatic inference: UPDR [CAV'15,JACM]
- Interactive inference: Ivy [PLDI'16]
- Use logical diagram to infer $\text{Inv} \in \forall^*$
- Can also prove absence of $\text{Inv} \in \forall^*$

Take away

- Decidable logic is useful
 - facilitates automation
- We need ways to guide verification tools
- How to divide the problem between human and machine?
- Different inference schemes
- Different Forms of interaction
- Other logics
- Theoretical understanding of limitations and tradeoffs



erc

Supervised Verification of Infinite-State Systems