# From TVLA to IVY

Mooly Sagiv

REPS at 60

Edinburgh, September 11, 2016

1

#### The TVLA Team







































# The TVLA Principles

- Concrete semantics expressed as evolving relations/graphs in FO<sup>TC</sup>
  - Celebrate unboundedness
  - No arithmetic
- Abstract Interpretation with Canonical Partially Disjunctive Abstraction
- Effective heuristics for automatic postcondition calculation
  - Kleene evaluation
  - Focus
  - Coerce
  - Differencing



#### Example: Concrete Interpretation



#### **Example: Abstract Interpretation**



#### Lessons learned

- FO<sup>TC</sup> is powerful
- But FO<sup>TC</sup> reasoning is complicated
  - Decidability of implications
  - Scalability of disjunctive abstractions
  - TVLA heuristics are effective for experts and specialized domains
    - Effective for shape analysis of common data structures and graduate students
    - But hard to apply by non experts and complicated clients
    - CEGAR has limited power

# The IVY System

#### Kenneth McMillan





#### Oded Padon





http://microsoft.github.io/ivy/

# **IVY** Principles

- Concrete semantics expressed as evolving relations/graphs in Effectively Propositional Logic (EPR)
  - Explore locality of updates
  - Simulate FOTC via differencing
- Interactive interference of conjunctive invariants
- [Abstract interpretation is coming]

# Ivy: Safety Verification by Interactive Generalization



#### Oded Padon, Kenneth McMillan, Aurojit Panda, Sharon Shoham



#### PLDI 2016 http://microsoft.github.io/ivy/

## Ivy: Safety Verification by Interactive Generalization

- Verification of distributed systems
- Modeling infinite-state systems in a way which allows decidable automated reasoning (EPR)
- Interactive discovery of inductive invariants

# Safety of Transition Systems



System S is safe if no bad state is reachable System S is safe iff there exists an inductive invariant Inv s.t.:

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

### Challenges for Deductive Verification

- 1. Formal specification:
  - Modeling the system
  - Formalizing the safety property
- 2. Inductive Invariants
  - Hard to specify manually
  - Hard to infer automatically
- 3. Deduction Checking inductiveness
  - Undecidability of implication checking
    - Unbounded state, arithmetic, quantifier alternation

#### Existing Approaches for Verification of Infinite-State Systems

- Automated invariant inference
  - Abstract Interpretation
  - Ultimately limited due to undecidability
- Use SMT for deduction with manual program annotations (e.g. Dafny)
  - Requires programmer effort to provide inductive invariants
  - SMT solver may diverge (matching loops, arithmetic)
- Interactive theorem provers (e.g. Coq, Isabelle/HOL)
  - Programmer provides inductive invariant and proves it
  - Huge effort (10-100 lines of proof per line of code)

Usually opaque when failing

## Our Approach in Ivy

- Restrict the specification language for decidability
  - Deduction is decidable with SAT solvers
  - Challenge: verify complex systems using a restricted language
    - Solution: domain specific axioms
- Finding inductive invariants (still undecidable):
  - Combine automated techniques with human guidance
  - Graphical user interaction
  - Key: generalization from counterexamples to induction
  - Decidability allows reliable automated checks





### Relational Modeling Language (RML)

- Designed to make verification tasks decidable
  - Yet expressive enough to model systems
- Turing-Complete



- Universally quantified inductive invariants are decidable to check
- System state described by finite (unbounded) relations
- No numerics
- Simple (quantifier-free) updates
- Universally quantified axioms (domain specific)
  - Total orders, partial orders, lists, trees, rings, quorums, ...

#### Languages for verification

Language	Executable	Deduction	Expresiveness	
C, Java, Python	$\checkmark$	Undecidable	Turing-Complete	
Dafny	$\checkmark$	Undecidable	Turing-Complete	
SMV	X	Decidable Verification of Temporal properties	Finite-state	
lvy	in progress	Decidable (EPR)	Turing-Complete	
Alloy	X	Undecidable	Turing-Complete	
Coq, Isabelle/HOL	$\checkmark$	Manual	Turing-Complete	
TLA+	X	Manual	Turing-Complete	



I *can* decide

inductiveness!



$$Inv = \neg Bad$$





Inv = 
$$\neg$$
Bad  $\land \varphi(\sigma 1, \sigma 1')$ 





Inv = 
$$\neg$$
Bad  $\land \varphi(\sigma 1, \sigma 1')$ 





 $\mathsf{Inv} = \neg \mathsf{Bad} \land \varphi(\sigma 1, \sigma 1') \land \varphi(\sigma 2, \sigma 2')$ 

- Key challenge for invariant inference: generalization
- Ivy's approach: put the user in the loop *interactive generalization*





# Example: Leader Election in a Ring

- Nodes are organized in a 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

# Example: Leader Election in a Ring

- Nodes are organized in a ring
- Each node has a unique numeric id
- Protocol:

Proposition: This algorithm detects one and only one

- Each highest number.
- A no *Argument:* By the circular nature of the configuration if the id in the mes and the consistent direction of messages, any message must meet all other processes before it comes back to its
- A ncinitiator. Only one message, that with the highest num-
- Theore ber, will not encounter a higher number on its way
  - around. Thus, the only process getting its own message
  - The 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 processes



### Leader Election Protocol (RML)

- ≤ (ID, ID) total order on node id's
- btw (Node, Node, 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

```
action send(n: Node) = {
    "s := next(n)";
    pending(id(n),s) := true
}
```

```
action receive(n: Node, m: ID) = {
  requires pending(m, n);
  pending(m, n) := false;
  if id(n) = m then
    // found Leader
    leader(n) := true
  else if id(n) ≤ m then
    // pass message
    "s := next(n)";
    pending(m, s) := true
}
```

```
protocol = (send | receive)*
```

```
next(a)=b \leftrightarrow \forall x: Node. x=a \lor x=b \lor btw(a,b,x)
```

```
assert I0 = \forall x,y: Node. leader(x) \rightarrow id(y) \leq id(x)
```







rcv(2, id(3))

rcv(3, id(3))













### Inductive Invariant for Leader Election

- ≤ (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

#### Safety property: 10

$$IO = \forall x,y: Node. leader(x) \rightarrow id(y) \leq id(x)$$

Inductive invariant: Tnv – TO

I How can we find an inductive invariant without knowing it?

#### Ivy: Check Inductiveness (1)





1. Each node sends its id to the next

2. A node that receives a message passes it (to the next in the ring) if

the id in the message is higher than the node's own id

3. A node that receives its own id becomes the leader









#### Ivy: Check Inductiveness (2)





1. Each node sends its id to the next

2. A node that receives a message passes it (to the next in the ring) if

the id in the message is higher than the node's own id

3. A node that receives its own id becomes the leader



#### Ivy: Generalize from CTI (2)











#### Ivy: Check Inductiveness (3)



*Init* ⊆ *Inv* (*Initiation*) *if* σ ∈ *Inv* and σ → σ' *then* σ' ∈ *Inv* (Consecution) *Inv* ∩ Bad = Ø(Safety)

id

pnd

Tid

Id

 $\leq$ 

id

id

pnd

### Completeness and Interaction Complexity

- Any generalization from CTI adds one universally quantified clause
- A universally quantified invariant in CNF with N clauses, can be obtained by the user in N generalization steps
  - Assuming the user is optimal
- If the user is sub-optimal, backtracking (weakening) may be needed

#### Verified Protocols

Protocol	Model Types	Relations & Functions	Property (# Literals)	Invariant (# Literals)	CTI Gen. Steps
Leader in Ring	2	5	3	12	3
Learning Switch	2	5	11	18	3
DB Chain Replication	4	13	11	35	7
Chord	1	13	35	46	4
Lock Server 500 Coq lines [Verdi]	5	11	3	21	8 (1h)
Distributed Lock 1 week [IronFleet]	2	5	3	26	12 (1h)
Paxos	Mark in prograss				
Raft	work in progress				

#### Expressiveness vs. Automation



### Summary

Expressiveness

- RML modeling language that makes deduction decidable
  - Many systems can be verified (axioms for orders, trees, rings, ...)
- Interactive generalization for finding inductive invariants
- Application to the domain of distributed protocols
- User intuition and machine heuristics complement each other:
  - User has intuition that leads to *better generalizations*
  - Machine is better at finding bugs and corner cases
- Interactive process assists user to gain intuition about the protocol







Coq	Dafny	http://microsoft.github.io/ivy/
		Static Analysis Types
		Automation

### 4 Lessons Learned from Tom

- Spoonfeed the reader
- Look the other way
- Think deeply
- Dedication/Dedication/Dedication







#### Thanks





#### Xavier Rival

