

Correct-by-constuction, Attack-tolerant Critical systems

Automating Protocol Synthesis

Robert Constable, Robbert van Renesse,
Vincent Rahli, Nicolas Schiper, Rich Eaton
– Cornell University
Mark Bickford
– ATC-NY

Correct-by-construction

- We make formal proofs that high-level system requirements are achievable.
- We synthesize system code from the proofs.
- **Milestone:** We synthesized a fault-tolerant consensus algorithm and deployed it as a component of ShadowDB, a replicated database.

Attack tolerant

- **Innate Immunity**
 - We prove that the system tolerates certain kinds and numbers of failures (crash, send-omission, etc.) under some assumptions on environment
- **Population diversity**
 - to thwart attacks not covered by innate immunity, we
 - make variant proofs which synthesize variant algorithms
 - run the synthesized code in variant runtime evaluators (in various languages)
 - (planned) pro-actively reconfigure to use variants in a unpredictable way

Critical components

Empirical observation: There are crucial components in the "stack" of many real-world systems that only a few "gurus" understand and maintain.

Why? In a running system these components have many dynamically changing, loosely coupled parts that achieve their global requirements for subtle reasons.

The Problem:

Such components are difficult to get right in the first place, and cannot be quickly changed if and when a flaw or exploitable feature of their design is discovered.

Our Solution: (Semi-) **Automate** the reasoning the "guru" uses to understand how the complex component works and why it is correct.

Synthesize the code for the component from this reasoning.

Formalizing “guru” reasoning

- Process algebra? No
- Temporal logic? Not much
- Refinement maps? Sometimes
- Reason directly about interacting modules/actors/worker threads with input, output, and state.
 - Specify the state change and output for each module.
 - Define and prove local invariants.
 - Prove global invariants.
- I/O Automata?
 - Almost, but we need a better way to reason about properties of dynamically created processes

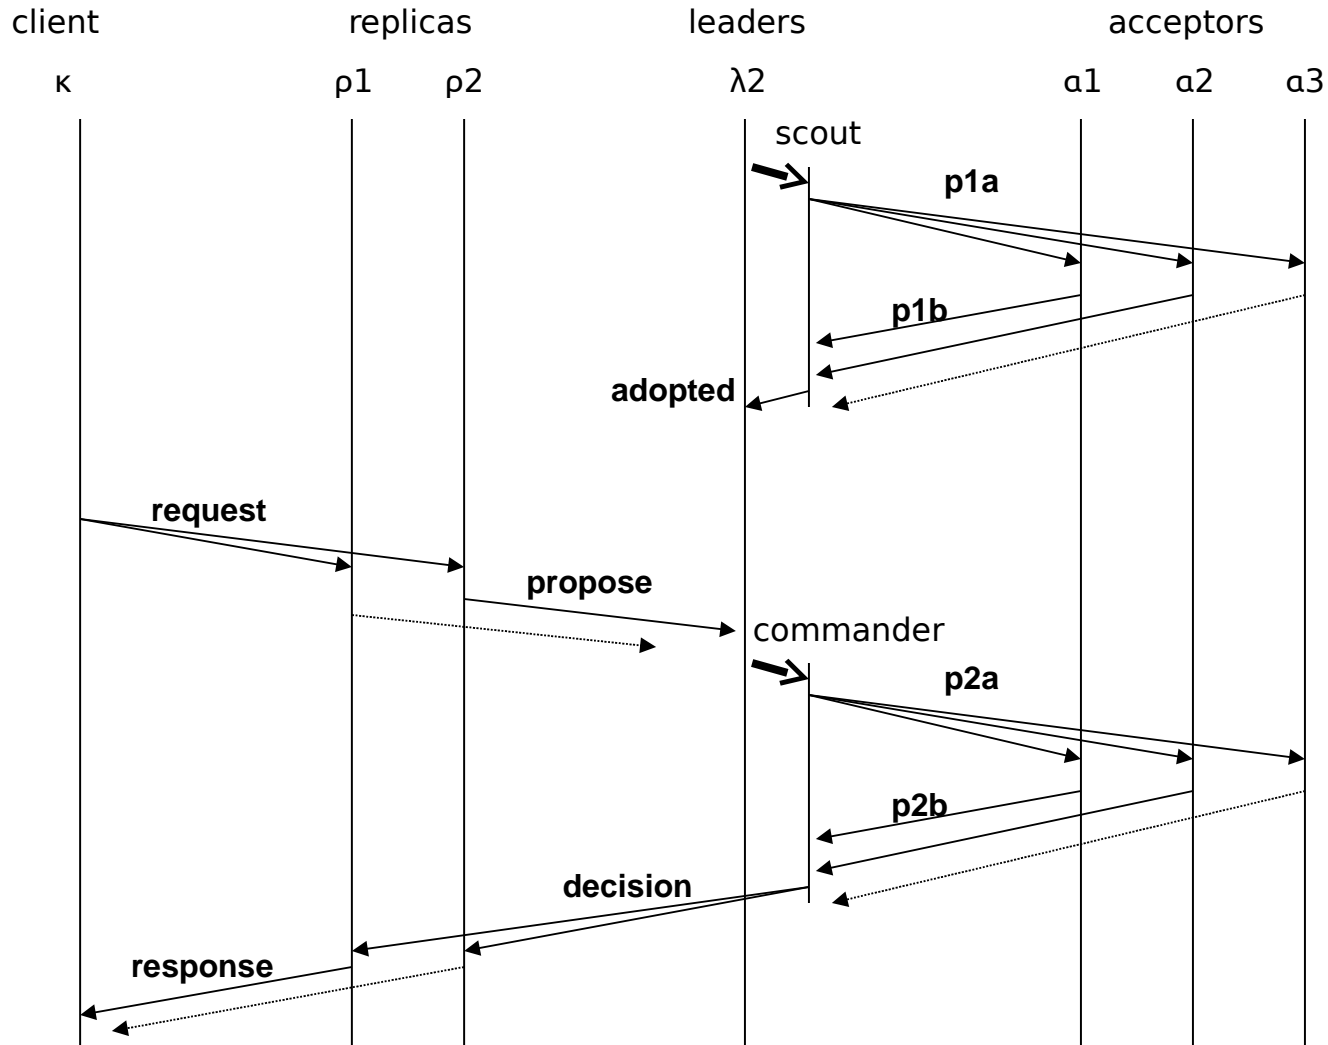
A Thread from Lamport's Paxos consensus algorithm

```
process Scout( $\lambda$ , acceptors, b)
var waitfor := acceptors, pvalues :=  $\emptyset$ ;
   $\forall \alpha \in \text{acceptors} : \text{send}(\alpha, p1a, \text{self}(), b)$ ;
for ever
  switch receive()
  case p1b,  $\alpha$ , b', r :
    if b' = b then
      pvalues := pvalues  $\cup$  r;
      waitfor := waitfor -  $\{\alpha\}$ ;
      if |waitfor| < |acceptors|/2 then
        send( $\lambda$ , adopted, b, pvalues);
        exit();
      end if;
    else
      send( $\lambda$ , preempted, b);
      exit();
    end switch;
end switch;
```

(From an explanation
of “multi-decree”
Paxos, in Robbert van
Renesse's

“Paxos made
moderately complex”)

Interacting processes in Paxos



Our Logical Method

- **Logic of events** = simple formal theory of mathematical structure corresponding to message sequence diagrams

Our Logical Method

- Logic of events = simple formal theory of mathematical structure corresponding to message sequence diagrams
- Processes described abstractly as “event classes” in EventML using “event class combinators”, $X \parallel Y$, $F \circ (X, Y)$, $\text{Prior}(X)$, $\text{Once}(X)$,
 - $X \gg = Y$ *delegation combinator* expresses dynamic process creation (classes form a monad)

EventML automatically synthesizes code (a set of process terms that execute in a message passing evaluator)

EventML automatically generates a logical form.

Nuprl then generates and proves a simplified “inductive logical form” (ILF)

EventML is both a programming and specification language.

Our Logical Method

- Logic of events = simple formal theory of mathematical structure corresponding to message sequence diagrams
- Processes described abstractly as “event classes” in EventML using “event class combinators”, $X \parallel Y$, $F \circ (X, Y)$, $\text{Prior}(X)$, $\text{Once}(X)$,
 - $X \gg Y$ *delegation combinator* expresses dynamic process creation (classes form a monad)
- **EventML automatically synthesizes code** (interpreted by a message passing evaluator)

Our Logical Method

- Logic of events = simple formal theory of mathematical structure corresponding to message sequence diagrams
- Processes described abstractly as “event classes” in EventML using “event class combinators”, $X \parallel Y$, $F \circ (X, Y)$, $\text{Prior}(X)$, $\text{Once}(X)$,
 - $X \gg Y$ *delegation combinator* expresses dynamic process creation (classes form a monad)
- EventML automatically synthesizes code (interpreted by a message passing evaluator)
- **EventML automatically generates a logical meaning.**
 - Nuprl then generates and proves a simplified “inductive logical form” (ILF)

Our Logical Method

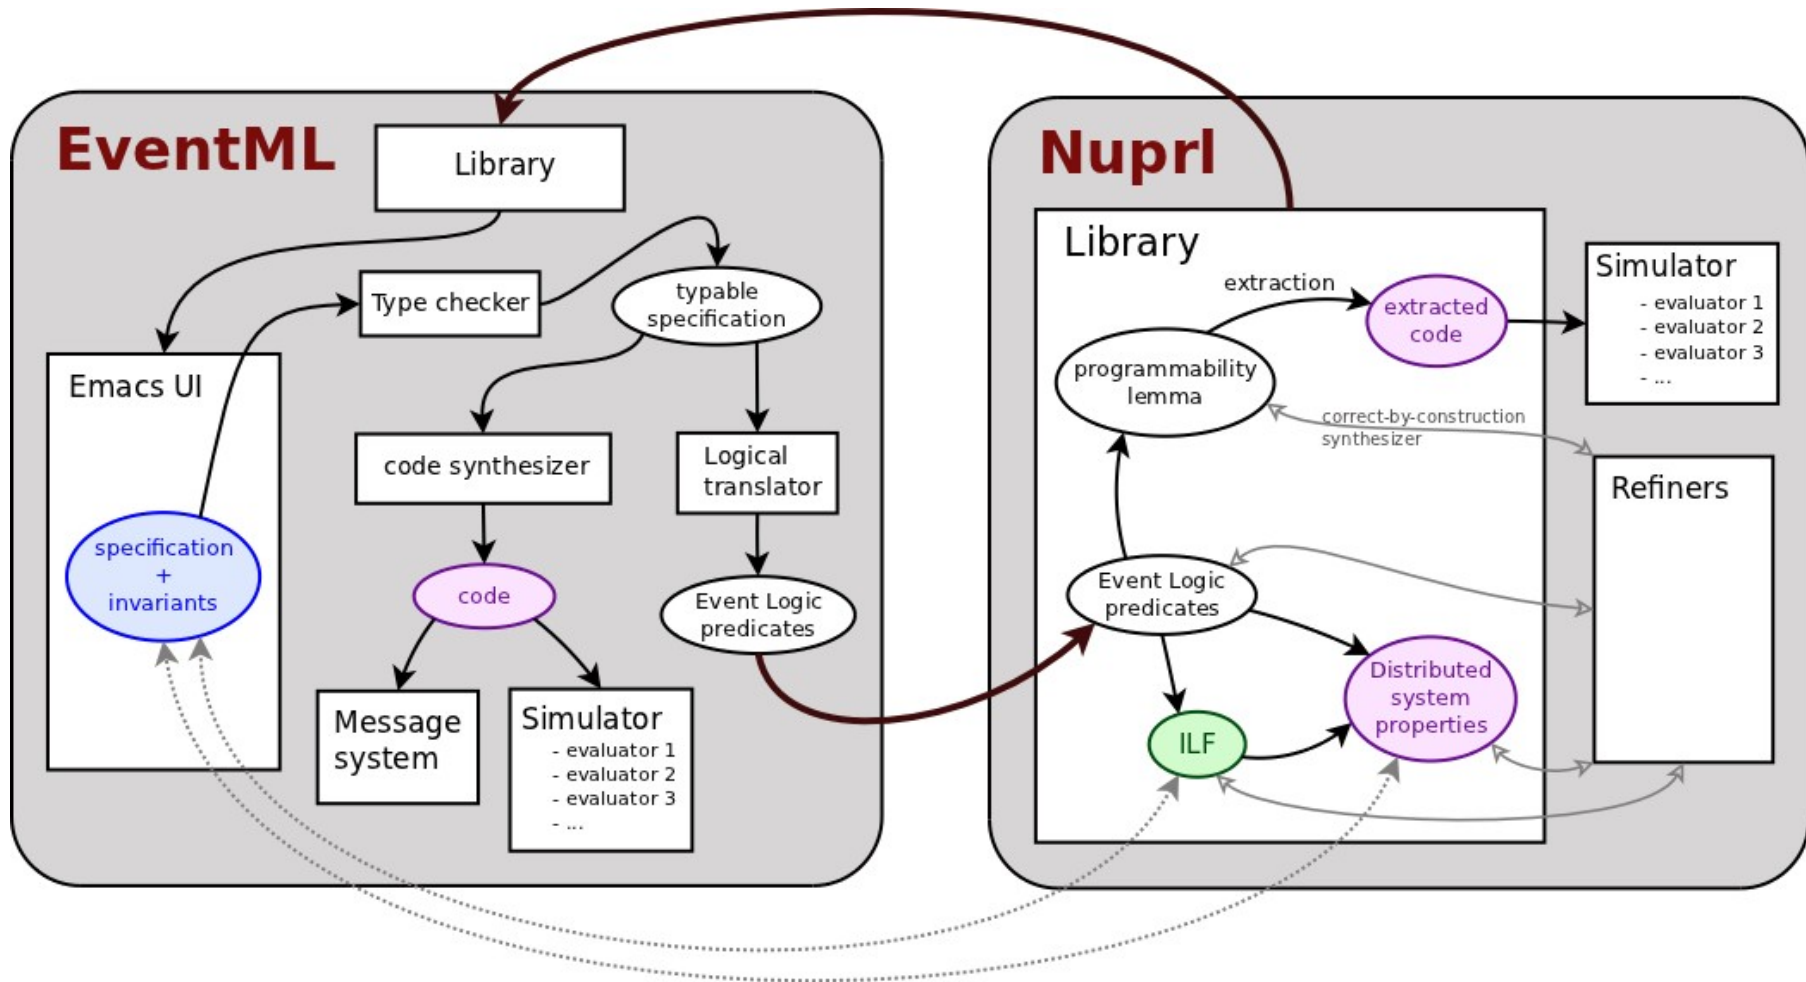
- Logic of events = simple formal theory of mathematical structure corresponding to message sequence diagrams
- Processes described abstractly as “event classes” in EventML using “event class combinators”, $X \parallel Y$, $F \circ (X, Y)$, $\text{Prior}(X)$, $\text{Once}(X)$,
 - $X \gg Y$ *delegation combinator* expresses dynamic process creation (classes form a monad)
- EventML automatically synthesizes code (a set of process terms that execute in a message passing evaluator)
- EventML automatically generates a logical form.
 - Nuprl then generates and proves a simplified “inductive logical form” (ILF)
- EventML is both a programming and specification language.

Progress since last PI meeting

- **Many enhancements to EventML**
 - Abstract data types
 - Invariant assertions, ordering properties
 - `classrec R p = X p || Y p >>= R`
- **Generation and simplification of ILF**
 - Using domain specific reasoners
 - Rewriting, quantifier elimination, etc. all proved by Nuprl tactics.
- **Synthesized code deployed**
 - Several versions of evaluators working
 - Consensus code being used in replicated database (ShadowDB).

Synthesized consensus protocols

- **3f+1 “simple” consensus algorithm**
 - Written in EventML with assertions
 - Most local invariants automatically proved
 - Using automatically generated ILF we proved the global consistency & validity properties in about **two days** (previous effort took two months)
 - Synthesized code is running in reconfiguration service of ShadowDB
- **Paxos nearly finished**



EventML (built by Vincent Rahli) cooperates with Nuprl at every stage of program development.

A Thread from Lamport's Paxos consensus algorithm

```
process Scout( $\lambda$ , acceptors, b)
var waitfor := acceptors, pvalues :=  $\emptyset$ ;
   $\forall \alpha \in \text{acceptors} : \text{send}(\alpha, p1a, \text{self}(), b)$ ;
for ever
  switch receive()
  case p1b,  $\alpha$ , b', r :
    if b' = b then
      pvalues := pvalues  $\cup$  r;
      waitfor := waitfor - { $\alpha$ };
      if |waitfor| < |acceptors|/2 then
        send( $\lambda$ , adopted, b, pvalues);
        exit();
      end if;
    else
      send( $\lambda$ , preempted, b);
      exit();
    end switch;
end switch;
```

(From an explanation
of “multi-decree”
Paxos, in Robbert van
Renesse's

“Paxos made
moderately complex”)

Part of EventML for Paxos

```
class ScoutNotify b = Output(\ldr.p1a'broadcast accpts (ldr, b));;

let on_p1b bnum loc (acloc,(b',pvals)) (waitfor,pvalues) =
  if eq_bnums bnum b'
  then let waitfor' = bag_remove (op =) waitfor acloc in
        let pvalues' = append_news same_pvalue pvalues pvals in
          (waitfor',pvalues')
  else (waitfor,pvalues) ;;

class ScoutState b = State1 (\loc.init_scout) (on_p1b b) p1b'base;;

let scout_output b ldr (a,(b',r)) (waitfor,pvalues) =
  if eq_bnums b b'
  then if bag_size waitfor < threshold
        then { adopted'send ldr (b,pvalues) }
        else {}
  else { preempted'send ldr b' };;

class ScoutOutput b = Once((scout_output b) o (p1b'base, ScoutState b));;

class Scout b = ScoutNotify b || ScoutOutput b ;;
```

Part of the Inductive Logical Form (ILF) for Paxos

$$\begin{aligned} & (\forall [\text{bnum}:\text{BNum}]. \forall [\text{accpts}:\text{bag}(\text{Id})]. \forall [\text{Op}, \text{Cid}:\{\text{T}:\text{Type} \mid \text{valueall-type}(\text{T})\}] . \\ & \forall [\text{eq_Cid}:\text{EqDecider}(\text{Cid})]. \forall [\text{es}:\text{E0}']. \forall [\text{e}:\text{E}]. \forall [\text{i}:\text{Id}]. \forall [\text{m}:\text{Message}]. \\ & \{ \langle \text{i}, \text{m} \rangle \in \text{paxos_scout_output}(\text{Cid}; \text{Op}; \text{accpts}) \text{ bnum@Loc } \circ (\text{Loc}, \\ & \quad \text{paxos_p1b}'\text{base}(\text{Cid}; \text{Op}), \text{paxos_ScoutState}(\text{Cid}; \text{Op}; \text{accpts}; \text{eq_Cid}) \text{ bnum})(\text{e}) \\ \iff & \downarrow (\text{header}(\text{e}) = \text{``paxos p1b''}) \\ & \wedge (\text{type}(\text{info}(\text{e})) = (\text{Id} \times \text{BNum} \times ((\text{BNum} \times \mathbb{Z} \times \text{Id} \times \text{Cid} \times \text{Op}) \text{List}))) \\ & \wedge (\text{i} = \text{loc}(\text{e})) \\ & \wedge (((\text{bnum} = (\text{fst}(\text{snd}(\text{body}(\text{info}(\text{e}))))))) \\ & \quad \wedge (\text{bag-size}(\text{fst}(\text{State of Scout bnum at e})) < \text{paxos_threshold}(\text{accpts})) \\ & \quad \wedge (\text{m} = \text{make-Msg}(\text{``paxos adopted''}; \\ & \quad \quad \text{BNum} \times ((\text{BNum} \times \mathbb{Z} \times \text{Id} \times \text{Cid} \times \text{Op}) \text{List}); \\ & \quad \quad \langle \text{bnum}, \text{snd}(\text{State of Scout (for bnum) at e}) \rangle))) \\ & \vee (((\neg(\text{bnum} = (\text{fst}(\text{snd}(\text{body}(\text{info}(\text{e}))))))) \\ & \quad \wedge (\text{m} = \text{make-Msg}(\text{``paxos preempted''}; \text{BNum}; \text{fst}(\text{snd}(\text{body}(\text{info}(\text{e})))))))))) \end{aligned}$$

The ILF is readable, and usually more informative than the EventML or pseudo-code. It is used automatically by our tactics to prove global properties of the algorithm.

Part of the Inductive Logical Form (ILF) for Paxos

$$\begin{aligned}
 & (\forall [\text{bnum}:\text{BNum}]. \forall [\text{accpts}:\text{bag}(\text{Id})]. \forall [\text{Op}, \text{Cid}:\{\text{T}:\text{Type} \mid \text{valueall-type}(\text{T})\}] . \\
 & \forall [\text{eq_Cid}:\text{EqDecider}(\text{Cid})]. \forall [\text{es}:\text{E0}']. \forall [\text{e}:\text{E}]. \forall [\text{i}:\text{Id}]. \forall [\text{m}:\text{Message}]. \\
 & \{ \langle \text{i}, \text{m} \rangle \in \text{paxos_scout_output}(\text{Cid}; \text{Op}; \text{accpts}) \text{ bnum@Loc } \circ (\text{Loc}, \\
 & \quad \text{paxos_p1b}'\text{base}(\text{Cid}; \text{Op}), \text{paxos_ScoutState}(\text{Cid}; \text{Op}; \text{accpts}; \text{eq_Cid}) \text{ bnum})(\text{e}) \\
 & \iff \boxed{ \begin{aligned} & \downarrow (\text{header}(\text{e}) = \text{``paxos p1b''}) \\ & \wedge (\text{type}(\text{info}(\text{e})) = (\text{Id} \times \text{BNum} \times ((\text{BNum} \times \mathbb{Z} \times \text{Id} \times \text{Cid} \times \text{Op}) \text{List}))) \\ & \wedge (\text{i} = \text{loc}(\text{e})) \end{aligned} } \\
 & \wedge (((\text{bnum} = (\text{fst}(\text{snd}(\text{body}(\text{info}(\text{e}))))))) \\
 & \quad \wedge (\text{bag-size}(\text{fst}(\text{State of Scout bnum at e})) < \text{paxos_threshold}(\text{accpts})) \\
 & \quad \wedge (\text{m} = \text{make-Msg}(\text{``paxos adopted''}; \\
 & \quad \quad \text{BNum} \times ((\text{BNum} \times \mathbb{Z} \times \text{Id} \times \text{Cid} \times \text{Op}) \text{List}); \\
 & \quad \quad \langle \text{bnum}, \text{snd}(\text{State of Scout (for bnum) at e}) \rangle))) \\
 & \quad \vee (((\neg(\text{bnum} = (\text{fst}(\text{snd}(\text{body}(\text{info}(\text{e}))))))) \\
 & \quad \quad \wedge (\text{m} = \text{make-Msg}(\text{``paxos preempted''}; \text{BNum}; \text{fst}(\text{snd}(\text{body}(\text{info}(\text{e}))))))))))
 \end{aligned}$$

The ILF is readable, and usually more informative than the EventML or pseudo-code. It is used automatically by our tactics to prove global properties of the algorithm.

Part of the Inductive Logical Form (ILF) for Paxos

```
( $\forall$ [bnum:BNum].  $\forall$ [accpts:bag(Id)].  $\forall$ [Op,Cid:{T:Type | valueall-type(T)} ] .  
 $\forall$ [eq_Cid:EqDecider(Cid)]. $\forall$ [es:E0'].  $\forall$ [e:E].  $\forall$ [i:Id].  $\forall$ [m:Message].  
  {<i, m>  $\in$  paxos_scout_output(Cid;Op;accpts) bnum@Loc o (Loc,  
    paxos_p1b'base(Cid;Op), paxos_ScoutState(Cid;Op;accpts;eq_Cid) bnum)(e)  
 $\iff$   $\downarrow$ (header(e) = ‘‘paxos p1b‘‘)  
   $\wedge$  (type(info(e)) = (Id  $\times$  BNum  $\times$  ((BNum  $\times$   $\mathbb{Z}$   $\times$  Id  $\times$  Cid  $\times$  Op) List)))  
   $\wedge$  (i = loc(e))  
   $\wedge$  (((bnum = (fst(snd(body(info(e)))))))  
   $\wedge$  (bag-size(fst(State of Scout bnum at e)) < paxos_threshold(accpts))  
   $\wedge$  (m = make-Msg(‘‘paxos adopted‘‘;  
    BNum  $\times$  ((BNum  $\times$   $\mathbb{Z}$   $\times$  Id  $\times$  Cid  $\times$  Op) List);  
    <bnum , snd(State of Scout (for bnum) at e)>)))  
   $\vee$  (( $\neg$ (bnum = (fst(snd(body(info(e)))))))  
     $\wedge$  (m = make-Msg(‘‘paxos preempted‘‘;BNum;fst(snd(body(info(e))))))))))}
```

The ILF is readable, and usually more informative than the EventML or pseudo-code. It is used automatically by our tactics to prove global properties of the algorithm.

Part of the Inductive Logical Form (ILF) for Paxos

$$\begin{aligned}
 & (\forall [\text{bnum}:\text{BNum}]. \forall [\text{accpts}:\text{bag}(\text{Id})]. \forall [\text{Op}, \text{Cid}:\{\text{T}:\text{Type} \mid \text{valueall-type}(\text{T})\}] . \\
 & \quad \forall [\text{eq_Cid}:\text{EqDecider}(\text{Cid})]. \forall [\text{es}:\text{E0}']. \forall [\text{e}:\text{E}]. \forall [\text{i}:\text{Id}]. \forall [\text{m}:\text{Message}]. \\
 & \quad \{ \langle \text{i}, \text{m} \rangle \in \text{paxos_scout_output}(\text{Cid}; \text{Op}; \text{accpts}) \text{ bnum} @ \text{Loc} \text{ o } (\text{Loc}, \\
 & \quad \quad \text{paxos_p1b}'\text{base}(\text{Cid}; \text{Op}), \text{paxos_ScoutState}(\text{Cid}; \text{Op}; \text{accpts}; \text{eq_Cid}) \text{ bnum}) (\text{e}) \\
 & \iff \downarrow (\text{header}(\text{e}) = \text{``paxos p1b''}) \\
 & \quad \wedge (\text{type}(\text{info}(\text{e})) = (\text{Id} \times \text{BNum} \times ((\text{BNum} \times \mathbb{Z} \times \text{Id} \times \text{Cid} \times \text{Op}) \text{List}))) \\
 & \quad \wedge (\text{i} = \text{loc}(\text{e})) \\
 & \quad \wedge (((\text{bnum} = (\text{fst}(\text{snd}(\text{body}(\text{info}(\text{e}))))))) \\
 & \quad \quad \wedge (\text{bag-size}(\text{fst}(\text{State of Scout bnum at e})) < \text{paxos_threshold}(\text{accpts})) \\
 & \quad \quad \wedge (\text{m} = \text{make-Msg}(\text{``paxos adopted''}; \\
 & \quad \quad \quad \text{BNum} \times ((\text{BNum} \times \mathbb{Z} \times \text{Id} \times \text{Cid} \times \text{Op}) \text{List}); \\
 & \quad \quad \quad \langle \text{bnum}, \text{snd}(\text{State of Scout (for bnum) at e}) \rangle))) \\
 & \quad \vee (((\neg(\text{bnum} = (\text{fst}(\text{snd}(\text{body}(\text{info}(\text{e}))))))) \\
 & \quad \quad \wedge (\text{m} = \text{make-Msg}(\text{``paxos preempted''}; \text{BNum}; \text{fst}(\text{snd}(\text{body}(\text{info}(\text{e}))))))))))
 \end{aligned}$$

The ILF is readable, and usually more informative than the EventML or pseudo-code. It is used automatically by our tactics to prove global properties of the algorithm.

Summary/ Next steps

- Synthesis of complex distributed algorithms from proofs works
 - Abstractions, automation essential
- Next steps
 - More variants of more protocols
 - Reason about capabilities/tags so that we can synthesize code that uses more CRASH technology