

Correct-by-Construction Methodology

CRASH Methodology for Correct-by-construction Attack-tolerant Systems

Also see the: [CRASH Project Home](#) | [About CRASH](#) | [Software](#) | [People](#) | [Publications](#)

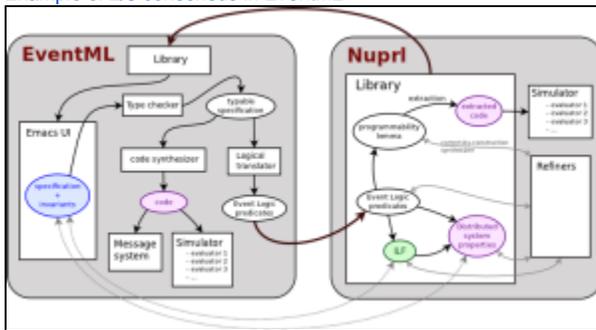
Background

Our method uses formal proofs of high-level system requirements then to synthesize code from the proven specification. Within our work on distributed systems we use the following concepts:

- Logic of Events: a simple formal theory of mathematical structure corresponding to message sequence diagrams
 - [Logic of Events presentation at LADA 2012](#) by Robert Constable, Mark Bickford, and Vincent Rahli. 2012
 - [Automated Proof of Authentication Protocols in a Logic of Events](#) by Mark Bickford. 2010
 - [A Causal Logic of Events in Formalized Computational Type Theory](#) by Mark Bickford, Robert L. Constable. 2005
- Event Classes: abstract description of processes
 - [Investigating Correct-by-Construction Attack-Tolerant Systems](#) by Robert Constable, Mark Bickford, Robbert Van Renesse. 2010
 - [Generating Event Logics with Higher-Order Processes as Realizers](#) by Mark Bickford, Robert L. Constable, David Guaspari. 2010
 - [Component Specification Using Event Classes](#) by Mark Bickford. 2009

The main tool we use is [EventML](#) which is a programming and specification language. EventML, built by [Vincent Rahli](#), cooperates with the Nuprl interactive theorem prover at every stage of program development to help programmers ensure correctness, document the code, and support modifications and improvements. It generates an *Inductive Logical Form* that proves the specification and can also automatically synthesize code.

- Resources for EventML
 - EventML documentation and downloads available at <http://www.nuprl.org/software/>
 - [EventML Tutorial](#)
 - [Example of 2/3 consensus in EventML](#)



Code *diversity* is created during the process. We can introduce variants at the EventML specification and code synthesis.

Methodology

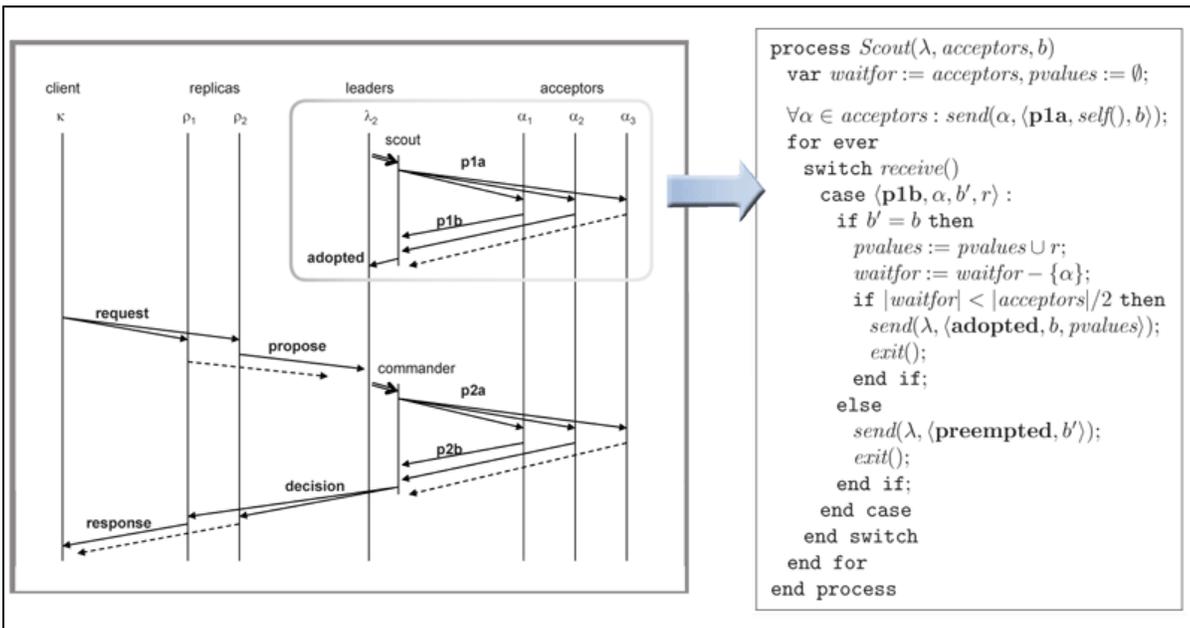
To create correct-by-construction code we:

1. Write the specification in [EventML](#)
2. Automatically generate and prove an Inductive Logical Form of the specification
3. Synthesize code
4. Diversify and deploy code

For more details about the example below see Mark's [presentation at the May 2012 CRASH meeting](#).

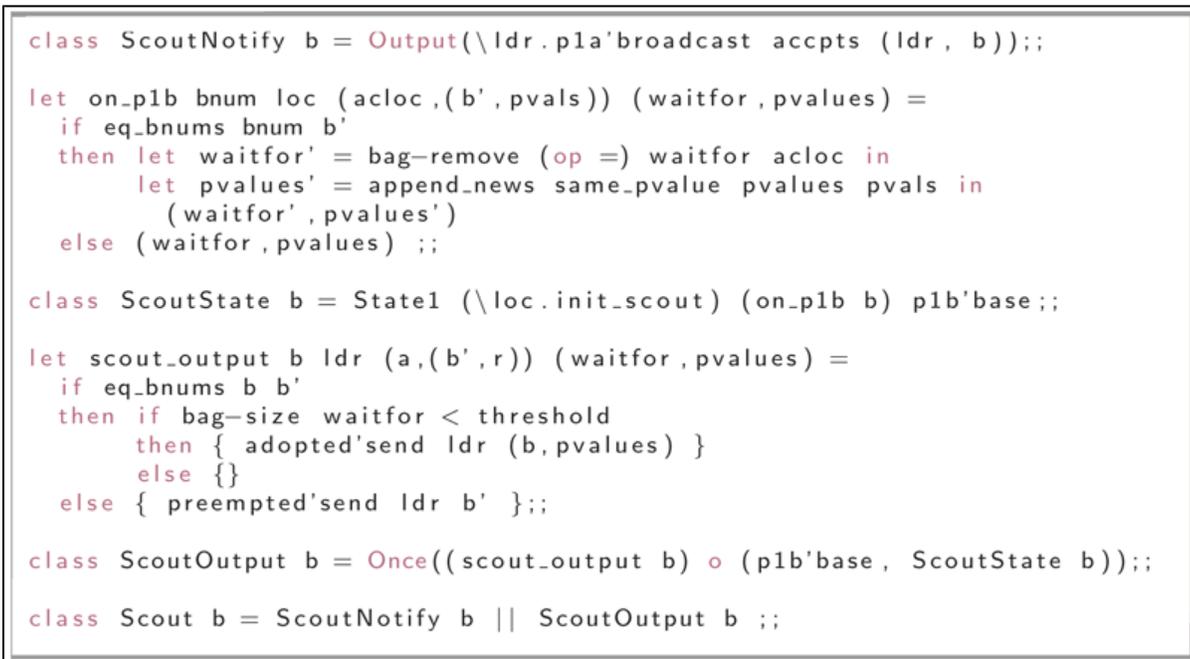
Example with Consensus

In this example we will look at Paxos consensus, focusing on specifying the Scout protocol.



Step 1 - We write the Scout specification in EventML

EventML uses classes from the Formal Digital Library to describe events and protocols.



Step 2 - Generate the Inductive Logical Form (ILF) from the Scout specification

Here EventML interfaces with Nuprl's formal library and theorem prover. Using the distributed prover, we generate a proof of the protocol which results in a readable Inductive Logical Form. If there are any issues with the proof we revise the specification and reiterate the process with the ILF.

Inductive Logical Form proof of specification

```

 $\forall [bnum:BNum]. \forall [acpts:bag(Id)]. \forall [Op,Cid:\{T:Type \mid valueall-type(T)\}].$ 
 $\forall [eq\_Cid:EqDecider(Cid)]. \forall [es:EO']. \forall [e:E]. \forall [i:Id]. \forall [m:Message].$ 
 $\{ \langle i, m \rangle \in paxos\_scout\_output(Cid;Op;acpts) \text{ bnum@Loc } o (Loc,$ 
 $\quad paxos\_p1b'base(Cid;Op), paxos\_ScoutState(Cid;Op;acpts;eq\_Cid) \text{ bnum})(e)$ 
 $\iff \downarrow (\text{header}(e) = \text{'paxos p1b'})$ 
 $\quad \wedge (\text{type}(\text{info}(e)) = (Id \times BNum \times ((BNum \times \mathbb{Z} \times Id \times Cid \times Op) \text{ List})))$ 
 $\quad \wedge (i = \text{loc}(e))$ 
 $\quad \wedge (((bnum = \text{fst}(\text{snd}(\text{body}(\text{info}(e))))))$ 
 $\quad \wedge (\text{bag-size}(\text{fst}(\text{State of Scout } bnum \text{ at } e)) < \text{paxos\_threshold}(\text{acpts}))$ 
 $\quad \wedge (m = \text{make-Msg}(\text{'paxos adopted'},$ 
 $\quad \quad BNum \times ((BNum \times \mathbb{Z} \times Id \times Cid \times Op) \text{ List});$ 
 $\quad \quad \langle bnum, \text{snd}(\text{State of Scout (for } bnum \text{ at } e)\rangle))$ 
 $\quad \vee ((\neg(bnum = \text{fst}(\text{snd}(\text{body}(\text{info}(e))))))$ 
 $\quad \wedge (m = \text{make-Msg}(\text{'paxos preempted'}, BNum; \text{fst}(\text{snd}(\text{body}(\text{info}(e))))))$ 

```

A scout is a subprocess that tries to get a ballot number adopted by a majority of acceptors. A scout, working on ballot number $bnum$, sends a message m to location i if and only if the following happens:

1. The scout reacts to the receipt of $p1b$ message (from an acceptor).
2. The $p1b$ message contains a location, ballot number, and proposal list.
3. The scout sends message m to itself.
4. If a majority of acceptors have adopted the ballot number $bnum$, then message m is an *adopted* message.
5. However, if the scout receives a $p1b$ message that contains a ballot number different from the one it is currently working on, the scout sends a *preempted* message.

Step 3 - Synthesize Consensus Code

Using EventML's correct-by-construction synthesizer we generate code from the specification.

Step 4 - Diversify & Deploy

Engaging our diversity in classes and state machines we generate multiple verified versions of the code. Then the final step before deployment is to test the code in the EventML simulator. The result is a correct-by-construction synthesized version of Paxos with multiple code variants.

Example Deployment: ShadowDB

ShadowDB is a replicated database, created by [Nicolas Schiper](#), on top of a synthesized consensus core. The primary defines the order in which transaction updates are applied. When crashes occur, consensus is used to reconfigure the set of replicas and agree on a prefix of executed transactions.

