

# Correct-by-Construction Attack-Tolerant Systems

Robert Constable

Mark Bickford

Robbert van Renesse

Cornell University

# Definition

**Attack-tolerant** distributed systems change their protocols on-the-fly in response to apparent attacks from the environment; they substitute **functionally equivalent components** possibly more resistant to detected threats.

# Definition

A **system** is built from **components** which consist of **processes** (protocols, algorithms).

system

component

process

# Definition

A system is **correct-by-construction** if we create a correctness proof for it while creating the code. This happens if we **synthesize the program** from a constructive proof that the specification is realizable.

# Abstract

We are experimenting with libraries of attack-tolerant protocols that are **correct-by – construction** because they are synthesized from proofs of formal specifications.

We require that **variants** of tolerant protocols are automatically generated and accompanied by **machine checked proofs** that the code satisfies formal properties.

# Main Result

We have found ways to automatically produce **many provably equivalent variants of components** using formal synthesis.

Variation arises from different choices made during the **proof and code synthesis process** starting from formal specifications.

# A Discovery

In the course of this work we also discovered that it is possible to create **undefeatable attackers** for fault-tolerant consensus protocols.

Code diversity can protect against these attackers as well.

# Outline

1. The Big Picture -- 8 minutes
2. The Power of Formalization – 6 minutes
3. The Power of Adversaries – 6 minutes
4. Conclusion – 2 minutes



# The Big Picture

Specify and prove using the Logic of Events

Synthesize protocols from proofs of design

Achieve multiplicative code diversity

# Consensus is a Good Example


In modern distributed systems, e.g. the Google file system, clouds, etc., reliability against **faults** (crashes, attacks) is achieved **by replication**.



Consensus is used to coordinate write actions to keep the replicas identical. It is a **critical protocol** in modern systems used by IBM, Google, Microsoft, Amazon, EMC, etc.

# Requirements of Consensus Task

Use **asynchronous** message passing to decide on a value.



# Logical Properties of Consensus

P1: If all inputs are **unanimous** with value  $v$ , then any decision must have value  $v$ .

All  $v:T$ . ( If All  $e:E(\text{Input})$ .  $\text{Input}(e) = v$  then  
All  $e:E(\text{Decide})$ .  $\text{Decide}(e) = v$  )

**Input** and **Decide** are **event classes** that effectively partition the events and assign values to them. The **events** are points in abstract space/time at which “information flows.” More about this just below.

# Logical Properties continued

P2: All decided values are input values.

All  $e:E(\text{Decide})$ . Exists  $e':E(\text{Input})$ .

$e' < e$  &  $\text{Decide}(e) = \text{Input}(e')$

We can see that P2 will imply P1, so we take P2 as part of the requirements.

# Event Classes

If  $X$  is an **event class**, then  $E(X)$  are the events in that class. Note  $E(X)$  **effectively** partitions all events  $E$  into  $E(X)$  and  $E - E(X)$ , its complement.

Every event in  $E(X)$  has a value of some type  $T$  which is denoted  **$X(e)$** . In the case of  $E(\text{Input})$  the value is the typed input, and for  $E(\text{Decide})$  the value is the one decided.

# Events

Formally the type  $E$  of events is defined relative to the computation model which includes a definition of **processes**.

The events are the **points of space/time** at which information is exchanged. The information at an event  $e$  is **info( $e$ )**.

# Further Requirements for Consensus

The key **safety property** of consensus is that all decisions agree.

P3: Any two decisions have the same value.

This is called **agreement**.

All  $e_1, e_2: E(\text{Decide})$ .  $\text{Decide}(e_1) = \text{Decide}(e_2)$ .



# Specific Approaches to Consensus

Many consensus protocols proceed in **rounds**, **voting on values**, trying to reach agreement. We have synthesized two families of consensus protocols, the **2/3 Protocol** and the **Paxos Protocol** families.

We structure specifications around **events during the voting process**, defining **E(Vote)** whose values are pairs  $\langle n, v \rangle$ , a **ballot number**,  $n$ , and a **value**,  $v$ .

# Properties of Voting

Suppose a group  $G$  of  $n$  processes,  $P_i$ , decide by voting. If each  $P_i$  collects all  $n$  votes into a list  $L$ , and applies some **deterministic function  $f(L)$** , such as majority value or maximum value, etc., then **consensus is trivial in one step**, and the value is known at each process in the first round – possibly at very different times.

The problem is much harder because of **possible failures**.





# Fault Tolerance

Replication is used to ensure system availability in the presence of **faults**. Suppose that we assume that up to **f** processes in a group  $G$  of **n** might fail, then how do the processes reach consensus?

The **TwoThirds method** of consensus is to take  $n = 3f + 1$  and **collect only  $2f + 1$**  votes on each round, assuming that **f** processes might have failed.

# Example for $f = 1, n = 4$

Here is a sample of voting in the case  $T = \{0,1\}$ .

0	0	1	1	inputs
				
0_11	_011	001_	00_1	collected votes
1	1	0	0	next vote

---

00_1	001_	0_11	_011
0	0	1	1

where **f is majority voting**, first vote is input

# Specifying the 2/3 Method

We can specify the fault tolerant 2/3 method by introducing further event classes.

$E(\text{Vote})$ ,  $E(\text{Collect})$ ,  $E(\text{Decide})$

$E(\text{Vote})$ : the initial vote is the  $\langle 0, \text{input value} \rangle$ ,  
subsequent votes are  $\langle n, f(L) \rangle$

$E(\text{Collect})$ : collect  $2f+1$  values from  $G$  into list  $L$

$E(\text{Decide})$ : **decide**  $v$  if all collected values are  $v$

# The Hard Bits

The small example shows what can go wrong with  $2/3$ . It can **waffle forever** between 0 and 1, thus never decide.

Clearly if there is a decide event, the values agree and that unique value is an input.





Can we say anything about eventually deciding, e.g. **liveness**?

# Liveness

If  $f$  processes eventually fail, then our design will work because if  $f$  have all failed by round  $r$ , then at round  $r+1$ , all alive processes will see the same  $2f+1$  values in the list  $L$ , and thus they will all vote for  $v' = f(L)$ , so in round  $r+2$  the values will be unanimous which will trigger a decide event.

# Example for $f = 1, n = 4$

Here is a sample of voting in the case  $T = \{0,1\}$ .

0	0	1	1	inputs
				
0 01_	001_	001_	_011	collected votes
0	0	0	1	next vote

---

000_	00_1	0_01	_001
0	0	0	0

where  **$f$  is majority voting**, first vote is input, round numbers omitted.



# Safety Example

We can see in the  $f = 1$  example that once a process  $P_i$  receives  $2/3$  unanimous values, say 0, it is not possible for another process to overturn the majority decision.

Indeed this is a general property of a  $2/3$  majority, the remaining  $1/3$  cannot overturn it even if they band together on every vote.

# Safety Continued

In the general case when voting is not by majority but using  $f(L)$  and the type of values is discrete, we know that if any process  $P_i$  sees unanimous value  $v$  in  $L$ , then any other process  $P_j$  seeing a unanimous value  $v'$  will see the same value, i.e.  $v = v'$  because the two lists,  $L_i$  and  $L_j$  at round  $r$  must share a value, that is they intersect.

# Synthesizing the 2/3 Protocol from a Proof of Design

We can formally prove the safety and liveness conditions from the event logic specification given earlier.

From this **formal proof of design, pf**, we can automatically extract a protocol, first as an abstract process, then by verified compilation, a program in Java or Erlang.

# The Synthesized 2/3 Protocol

**Begin**  $r:\text{Nat}$ ,  $\text{decided}_i$ ,  $\text{vote}_i: \text{Bool}$ ,  
 $r = 0$ ,  $\text{decided}_i = \text{false}$ ,  $v_i = \text{input to } P_i$ ;  $\text{vote}_i = v_i$

**Until**  $\text{decided}_i$  **do**:

1.  $r := r+1$
2. **Broadcast**  $\langle r, \text{vote}_i \rangle$  to group  $G$
3. **Collect**  $2f+1$  round  $r$  votes in list  $L$
4.  $\text{vote}_i := \text{majority}(L)$
5. **If**  $\text{unanimous}(L)$  **then**  $\text{decided}_i := \text{true}$

**End**

# Diversity

When we prove properties of a design, there are many options at several steps, and we are able to create **multiple proofs** at low additional cost. In the process **we create new designs**.

For example, for the 2/3 protocol, Mark Bickford found a variant that is faster by varying the design proof, as mentioned in our paper – he **varies the collection method**.

# Outline

1. The Big Picture
2. The Power of Formalization
3. The Power of Adversaries
4. Conclusion

# Diversity at the Level of Proof


Multiple **formal proofs** are “simultaneously” generated. We illustrate this by viewing a proof as a tree generated top down.

# Diversity at the Level of Proof


At the root of the upside down tree is a goal of the form  $\vdash G$  which asserts that proposition  $G$  is true (provable). Interior nodes have the form  $H \vdash G$  which means that from assumptions  $H$  the goal  $G$  can be proved.



# Illustrating Multiple Proofs



# Illustrating Multiple Proofs



P1


P2

P3

P4

# Data Structure Diversity

Assuming there are four abstract protocols derived from the proof trees. For each of them it is possible to implement with different data structures, e.g. list, array, tree, set, etc.




# Programming Language Diversity

We can translate abstract programs into common programming languages such as Java, Erlang, C++, or F#. So far we use only Java and Erlang.

Combining all levels of diversity we are able to generate over 200 variants of a protocol in the best case.

# Language Diversity



4 protocols, 14 options in 4 languages,  
offers over **200 variants**

# Outline

1. The Big Picture
2. The Power of Formalization
3. The Power of Adversaries
4. Conclusion

# Role of the Environment

All distributed computing models must have a component that determines when messages between processes are delivered. We call this the **environment**. It introduces **uncertainty** into the model and determines the **schedule** of events.

# A Fundamental Theorem of about the Environment

The **Fischer/Lynch/Paterson** theorem (**FLP85**) about the computing environment says:

it is not possible to guarantee consensus among  $n$  processes when one of them might fail.

We have seen the possibility of this with the  $2/3$  Protocol which could waffle between choosing 0 or 1. The environment can act as an adversary to consensus by managing message delivery.



# The Environment as Adversary

In the setting of synthesizing protocols, I have shown that the FLP result can be made constructive (**CFLP**). This means that there is an algorithm,  $env$ , which given a potential consensus protocol  $P$  and a proof  $pf$  that it is **nonblocking** can create message ordering and a computation based on it,  $env(P, pf)$ , in which  $P$  runs forever, failing to achieve consensus.

# Perfect Attacker

The algorithm  $\text{env}(P, pf)$  is the perfect “denial of service attacker” against any consensus protocol  $P$  that is sensible (won't block).

Note, 2/3 will **block** if it waits for  $n$  replies or if it refuses to change votes as rounds progress.

# Defending Against the Perfect Attack

One way to defend against  $evn(P, pf)$  is to switch to another protocol  $P'$  if there appears to be an attack against  $P$ .

# Outline

1. The Big Picture
2. The Power of Formalization
3. The Power of Adversaries
4. Conclusion

# Conclusion

We are exploring how to build distributed systems that are **attack-tolerant by design**.

The key idea is to implement systems that respond to attacks by modifying their code in a provably safe way. We believe that the more code variants we can produce, the more resistant systems will be to attack.

# Conclusion

We initiate code generation at a very high level of abstraction by formally proving that **designs are realizable**.

By starting at such a high level, we discovered more correct options than possible by less technically advanced methods. **This discovery reveals new reasons for working formally at high levels of abstraction.**

**THE END**

# Conclusion


We are exploring how to build distributed systems that are **attack-tolerant by design**. The key idea is to implement systems that respond to attacks by modifying their code in a provably safe way. **We believe that the more code variants we can produce, the more resistant systems will be to attack.**

We have found ways to automatically produce many provably equivalent variants of components using formal synthesis. Variation arises from different choices made during synthesis.

We start at very high levels of abstraction by formally proving that designs are realizable. By starting at such a high level, we discovered more correct options than possible by less technically advanced methods. **This discovery reveals new reasons for working formally at high levels of abstraction.**



# Illustrating Multiple Proofs



P1

P2

P3

P4

# Abstract

**Attack-tolerant** distributed systems change their protocols on-the-fly in response to apparent attacks from the environment; they substitute functionally equivalent versions possibly more resistant to detected threats.

We are experimenting with libraries of attack-tolerant protocols that are **correct-by-construction** and testing them in environments that simulate specified threats, including constructive versions of the famous **FLP** imaginary adversary against fault-tolerant consensus.

We expect that all variants of tolerant protocols are automatically generated and accompanied by **machine checked proofs** that the generated code satisfies formal properties.

# Development of Event Logic

Our Event Logic is an abstract account of distributed computing inspired by the work of Winskel and Plotkin in the 80's on Petri Nets. It spans from the very abstract notion of **Event Class** down to formal models of protocols that can be compiled to **Message Automata** and from them into code in languages such as Java, ML, Erlang, F#, and so forth.

# Safety

We almost have a **proof** that our design at the level of event classes meets the requirements.

We also need to know property P2, that two decided values agree even if no processes fail.

Suppose that at some  $P_i$  the  $2f+1$  values collected in  $L$  are the same and likewise at  $P_j$  for  $j \neq i$ . Are the decided values equal?

# Diversity at the Level of Proof

Multiple **formal proofs** can be “simultaneously” generated. We illustrate this by viewing a proof as a tree generated top down.

At the root of the upside down tree is a goal of the form  $\vdash G$  which asserts that proposition  $G$  is true (provable). Interior nodes have the form  $H \vdash G$  which means that from assumptions  $H$  the goal  $G$  can be proved.

