

# The Need for Language Support for Fault-Tolerant Distributed Systems

Cezara Drăgoi<sup>1</sup>, Thomas A. Henzinger<sup>\*2</sup>, and Damien Zufferey<sup>†3</sup>

- 1 INRIA, ENS, CNRS  
Paris, France  
cezara.dragoi@inria.fr
- 2 IST Austria  
Klosterneuburg, Austria  
tah@ist.ac.at
- 3 MIT CSAIL  
Cambridge, USA  
zufferey@csail.mit.edu

---

## Abstract

Fault-tolerant distributed algorithms play an important role in many critical/high-availability applications. These algorithms are notoriously difficult to implement correctly, due to asynchronous communication and the occurrence of faults, such as the network dropping messages or computers crashing. Nonetheless there is surprisingly little language and verification support to build distributed systems based on fault-tolerant algorithms. In this paper, we present some of the challenges that a designer has to overcome to implement a fault-tolerant distributed system. Then we review different models that have been proposed to reason about distributed algorithms and sketch how such a model can form the basis for a domain-specific programming language. Adopting a high-level programming model can simplify the programmer's life and make the code amenable to automated verification, while still compiling to efficiently executable code. We conclude by summarizing the current status of an ongoing language design and implementation project that is based on this idea.

**1998 ACM Subject Classification** D.3.3 Language Constructs and Features

**Keywords and phrases** Programming language, Fault-tolerant distributed algorithms, Automated verification

**Digital Object Identifier** 10.4230/LIPIcs.SNAPL.2015.90

## 1 Introduction

Replication of data across multiple data centers allows applications to be available even in the case of failures, guaranteeing clients access to the data. For instance, state machine replication achieves fault-tolerance by cloning the application on several replicas, and ensuring that the same sequence of commands is executed on each replica using message passing. Building fault-tolerant software is challenging because designers have to face (1) asynchronous concurrency and (2) faults. Asynchrony means that replicas can interleave their actions in arbitrary ways and they communicate via message passing in a network that can delay

---

\* This research was supported in part by the European Research Council (ERC) under grant 267989 (QUAREM) and by the Austrian Science Fund (FWF) under grants S11402-N23 (RiSE) and Z211-N23 (Wittgenstein Award).

† Damien Zufferey was supported by National Science Foundation grant 1138967.



© Cezara Drăgoi, Thomas A. Henzinger, and Damien Zufferey;  
licensed under Creative Commons License CC-BY

1st Summit on Advances in Programming Languages (SNAPL'15).

Eds.: Thomas Ball, Rastislav Bodík, Shriram Krishnamurthi, Benjamin S. Lerner, and Greg Morrisett; pp. 90–102



Leibniz International Proceedings in Informatics

LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

messages for an unbounded amount of time. Hardware faults can be transient, e.g., network partition, or permanent, e.g., crashes.

In a distributed system processes only have a limited view of the global state. However, one has to enforce some global property, e.g., ensuring consistency<sup>1</sup> between replicas by making sure that they execute a similar sequence of commands. Ensuring such a property typically requires solving a consensus problem: roughly, each process has an initial value, and all non-faulty processes have to agree on a unique decision value picked from the initial ones, even in the presence of faults and uncertainty in the timing of events. Strong consistency can be implemented by using consensus iteratively. Noteworthy examples from industry are Google Megastore, which uses Paxos [37] to replicate primary user data across data centers on every write, the Chubby lock service [11], Apache Zookeeper [33], which embeds a consensus protocol called Zookeeper Atomic Broadcast [35], or Microsoft Autopilot [34] that manages a data center.

State machine replication has received a lot of attention in both academia and industry. A fundamental result on distributed algorithms [29] shows that it is impossible to reach consensus in asynchronous systems where at least one process might crash. Consequently, a large number of algorithms have been developed [24, 21, 13, 37, 57, 42, 35, 46, 48], each of them solving consensus under different assumptions on the type of faults, and the “degree of synchrony” of the system. Moreover, weaker versions of consensus that solvable under fewer network assumptions [18, 26] have been proposed.

Consensus algorithms are difficult both from a theoretical and practical perspective. Because of the impossibility of solving consensus in asynchronous networks in the presence of faults, the existing consensus algorithms not only make various assumptions on the network for which they are designed but have also a complex flow of data. For example, the widely used Paxos algorithm [37] was considered hard to understand. A few years after its first publication Lamport wrote “Paxos Made Simple” [50] to re-explain the algorithm in simpler terms. However, the challenge of understanding it still remained. As an example, 2014 saw the publication of the raft algorithm [48], which purpose is to be simpler to understand than Paxos.

From a practical perspective implementing fault-tolerant distributed algorithms is challenging because the programmer has to use clocks, timing constraints, and complex network programming using low-level constructs, e.g., sockets, event-handlers. High-level programming languages like Erlang, or in general library based approaches, offer limited the possibility to reason about faults. Moreover, these algorithms are rarely implemented as theoretically defined, but modified to fit the constraints of the system in which they are incorporated. A group of engineers working at Google wrote “Paxos Made Live” [12] to describe some of the problems encountered while building a system based on Paxos and possible solutions for them. They [12, Section 9] emphasizes the following:

- “There are significant gaps between the description of the Paxos algorithm and the needs of a real-world system. In order to build a real-world system, an expert needs to use numerous ideas scattered in the literature and make several relatively small protocol extensions. The cumulative effort will be substantial and the final system will be based on an unproven protocol.”
- “The fault-tolerance computing community has not developed the tools to make it easy to implement their algorithms.”

---

<sup>1</sup> Both strong or weak consistency are global properties.

- “The fault-tolerance computing community has not paid enough attention to testing, a key ingredient for building fault-tolerant systems.”

Despite the importance of fault-tolerant distributed algorithms, there are no automated verification techniques that can deal with the complexity of an implementation of an algorithm like Paxos. A few algorithms have been mechanically proved correct using interactive theorem provers, like Isabelle [47]. However these proofs required substantial effort and machine-checked proof of correctness for asynchronous implementations would demand even more expert time. For automated verification methods, the main difficulties are the unbounded number of replicas, the complex control structure using event-handlers, sockets, etc., and the complex data structures, e.g. unbounded buffers.

We think that the difficulty does not only come from the algorithms but from the way we think about distributed systems. Therefore, we are interested in finding an appropriate programming model for fault-tolerant distributed algorithms, that increases the confidence we have in applications based on state machine replication.

Our goal is to develop a domain-specific language for fault-tolerant distributed systems. The language must

1. be high-level to help the programmer focus on the algorithmic questions rather than spending time fiddling with low-level network and timer code;
2. compile into working, efficient systems that preserve the important properties of high-level algorithms;
3. use a programming model amenable to automated testing and verification.

Even though some of the difficulty of implementing fault-tolerant algorithms were addressed in the distributed algorithms community, there isn't a widely accepted programming language for fault-tolerant algorithms. Therefore, we think that this class of systems is relevant also for the programming language community. If we look at the past few years of programming language design and implementation, as well as software verification, i.e., POPL, PLDI, OOPSLA, and CAV, we find that only a handful of papers deal with systems where faults can occur. We believe that one should explore the advances in formal methods to increase the confidence in fault-tolerant applications. The standard formal verification approach considers the asynchronous code implemented in general programming language like C or Java, which is known for being notoriously hard. Instead we propose use alternative view over replicated systems, leading to a new programming model that eases the programmers' task and is amenable to automated verification.

In Section 2, we give an example of a fault-tolerant algorithm and explain some of the challenges in implementing and verifying it. In Section 3, we review existing work in the domain of programming languages for message-passing concurrency, models of computation for distributed systems, and the automated verification of distributed algorithms. In Section 4, we highlight several promising research directions that combine ideas across different fields to achieve the goals stated above. This section also summarizes the status of our ongoing project on this topic.

## 2 Example

Figure 1 shows a fault-tolerant algorithm [9] that solves the consensus. The algorithm is given in the formulation of [16], which is a round-based model: all processes execute the same code in lock-step manner. Each process knows initially a (potentially) different value for the variable  $x$ . The algorithm first establishes a majority of processes agreeing on the same value for  $x$ , then witnesses that majority, and finally decides that the value of  $x$  is the one


Repeat:	
Send:	send $x$ to all processes
Update:	if received more than $2n/3$ messages then
	$x :=$ the smallest most often received value
	if more than $2n/3$ received values are equal to $x$ then
	$decide(x)$

■ **Figure 1** A round based algorithm that solves Consensus.

the majority agreed on. This particular algorithm needs four replicas to tolerate one crash, i.e.,  $n > 3f$ , where  $n$  is the total number of replicas and  $f$  is the number of faulty processes.

Distributed algorithms are typically published in pseudo-code, or, in the best case, a formal but not executable specification [38]. For the implementation, the programmer has to make many decisions that influence the final system in a way that is hard to predict a priori. The first decision is to fix the fault model. The algorithm in Figure 1 tolerates only *benign* faults, i.e., messages can be lost but not modified. Allowing a *byzantine* attacker, which corrupts messages, requires using a modified algorithm [6]. After deciding on a benign faults, one still needs to decide whether to consider *crash-stop* or *crash-recovery*. Crash-stop means that a fault causes a replica to stop taking any steps. On the other hand, in a crash-recovery model faults are transient and a replica can rejoin system after a crash. Implementing an algorithm in a crash-recovery fault model requires a *stable memory*, i.e., writing to a permanent storage every change of the local state, or using a dedicated *recovery procedure* to be executed by a replica before rejoining the system after a crash. Next there is the question if the communication is reliable or if the network can *drop messages*. This will influence how long the processes wait to receive messages and whether a missing message is interpreted as a crash. All these choices influence the implementation: what kind of *failure-detector* to implement, how much additional information has to be added to messages, etc. A failure-detector is an application which keeps track of the processes in the system to know whether they have failed. In its simplest incarnation, a failure detector periodically sends heartbeat messages to the others processes in the system. Then, for performance reasons, these messages can be piggybacked on top of the protocol's messages to avoid the overhead of sending additional messages. Such small optimizations quickly add up, perturb the normal flow of the protocol, and lead to a complex, hard-to-maintain codebase. Finally, when the algorithm is implemented, there is the question of how to test or verify it. Asynchronous systems are notoriously known hard to test, debug, and prove correct.

When defining the semantics of a programming language, it is simple and natural to assume an asynchronous model, but most programmers and algorithm designers have a notion of time when they write code. The round structure of the algorithm in Fig. 1 can be seen as an abstract notion of time. From a more practical perspective, in the context of the programming language P [20], we had discussions about adding constructs to deal with time. P is a domain specific language for communication state machine which is used to program device drivers. A common coding pattern for drivers is to send a request, initialize a timer, and wait for either the response or the timer to fire. If the response comes first, then the timer needs to be canceled. This pattern introduces boilerplate that obscures the code. The proposal was to add a new type of transition, `Timeout(duration)`, that would automatically initialize/cancel the timer. However, in the end we decided against adding timeouts to keep the number of primitives in the language small.



■ **Figure 2** Paxos decomposed in four communication-closed rounds. Process  $p_1$  has the role of proposer, acceptor, and learner. Processes  $p_2$  and  $p_3$  are acceptors.

### 3 Prior Work

**Programming abstractions for message-passing systems.** Formalisms such as the  $\pi$ -calculus [44, 45] and I/O-automata [41] have mostly been used to give a formal semantics to message-passing systems and to analyze them, but not as programming model. Some of these models have been extended with time, but they are all natively asynchronous. Furthermore, faults are not part of these models. To deal with faults one has to model them on top of the provided primitives, increasing dramatically the complexity of the final system.

The Actor model [31] is probably the most successful high-level abstraction for message-passing systems. CSP [32] has also seen some success with programming languages like Occam, Ada, etc., and CSP-style rendezvous concurrency is present in many other languages. Programming languages like Erlang [5] are based on the actor model and many programming languages have libraries implementing the actor model. Erlang also has dedicated library support to help the programmers write fault-tolerant applications. Finally, at the lowest level, we find the POSIX sockets. Disturbingly enough, they are probably still the most common programming abstraction for distributed systems.

In this paper, we focus on fault-tolerance through replication, considering algorithms that make uniform and non-uniform assumptions on the replicas. The programming models designed to implement replicated computation and storage [53, 58], typically take a centralized approach, based on a trusted client that interacts with all replicas. For example, the Erlang/OTP framework [53] provides fault tolerance using supervision. An application is hierarchically structured in a tree where the leaf nodes are workers and the inner nodes are supervisors. The supervisors watch over workers and other supervisors. They are responsible for starting and shutting them down in the normal course of action and, in presence of failure, responding to a fault in their children, usually by restarting the process. However, there are implementations of Paxos in Erlang designed to integrate with the OTP library and systems like Zookeeper, that combine both approaches, a uniformly replicated core that manages a large collection of workers. We are mostly interested in replicated systems that do not rely on supervision, and make uniform assumptions on the replicas, or non-uniform assumptions that are weaker than supervision.

**Models of computation for distributed algorithms.** Distributed algorithms and systems are mature research fields. This is by no means a complete overview of these fields, but we highlight a few ideas that have not been adopted by the language and verification communities and that we believe are good candidates to form the basis of new programming abstractions. An important idea to simplify the formulation of distributed algorithms is scoping the communication. *Communication-closed rounds* [25] encapsulate all communication within rounds. Conceptually, processes operate in lock-step: in each round they send messages, receive messages, and update their local state depending on the local state at the beginning of the round and the received messages. Many distributed algorithms can be formulated naturally

in a round-based model [39, 40, 24, 16, 36, 7, 6]. Figure 2 shows how this can be done for the Paxos algorithm. Organizing the computation in rounds looks restrictive and originally was limited to synchronous systems. However, rounds can be generalized to partial synchrony [24]. Moreover, failure detectors [13] have been proposed to structure the detection and handling of failures. The combination of these ideas gives rise to round-based computational models that can *uniformly model synchronous and asynchronous systems* [30, 16, 3], e.g., Charron-Bost and Schiper [16, Table 1] showed how to do this for common fault models. In this view, an asynchronous, or faulty, system is a synchronous system with an adversarial environment that can delay, drop, or modify messages.

**Verification of distributed algorithms.** From the verification perspective, distributed algorithms are a very challenging class of systems because of several sources of unboundedness: messages come from unbounded domains, the number of processes is a parameter, and channels may also be unbounded. Indeed, most of the verification problems are undecidable for parameterized systems [4]. As alternative, model checking is done by instantiating the algorithms using a finite, usually small, number of processes [38, 54, 55, 56]. This approach works for finding bugs, but cannot prove an algorithm correct. Furthermore, the reduction to finite-state systems cannot model faithfully all types of faults, even for a finite number of processes. For instance, if we consider byzantine systems, a faulty process can send arbitrary messages that may not even be defined in the protocol.

The channels also make the verification problem hard. Unbounded FIFO channels causes undecidability even for two processes [8]. Making the channels lossy and fixing the number of processes makes the problem decidable [2], with a non-primitive recursive complexity [52]. Weaker channel models are usually at least EXPSPACE-hard for verification. A pervasive pattern in fault-tolerant systems is counting messages up to a threshold that depends on the number of processes in the system, e.g.,  $2n/3$  in Figure 1. Unfortunately, this means that these systems do not fall in the class of well-structured transition systems [27], because they do not satisfy the required monotonicity conditions. Well-structured transition systems are arguably the most general class of infinite-state systems for which verification questions are decidable [28, 1].

As a result, full formal proofs of correctness of fault-tolerant algorithms are often based on interactive proof assistants [15, 19, 17]. However, they require a substantial effort. For instance, a proof of correctness of the Paxos algorithm is about 1500 lines long [17]. Recently, we have seen the use of interactive proof assistants to program real systems [51].

## 4 Project Outline

We started to look at fault-tolerant distributed systems with the goal of verifying them [23]. We quickly realized that the situation is very grim if we look at them through classical, asynchronous language and verification models. Using the round-based models from the distributed algorithms community we can achieve a much higher degree of automation in the verification. After all, these models were developed as simplifications for manually reasoning about asynchronous systems, e.g., for expressing the FLP [29] proof “on the back of a napkin”<sup>2</sup>. Structuring the computation in communication-closed rounds is an intuitive programming model as for each round, the programmer only has to give two methods, one

<sup>2</sup> From the talk “Better late (40 years late!) than never: Monday morning quarterbacking the coordinated-attack problem” by Eli Gafni at Yale University on 2014-10-09.

to send messages, and the other to update the local state upon reception of messages. To cover asynchronous and faulty executions, one only needs to accept that not every message gets delivered. From the verification perspective, this model (1) removes the problem of interleavings, because computations look synchronous: all processes proceed in lock-step, and (2) allows removing all channels from the state of the system by looking at the system’s invariant at the boundary between rounds.

We believe that like the verification community, the programming language community has neglected to address the question of faulty systems and has been unproductively conservative in modeling them. The opportunity for the language community is to develop better abstractions for programming and reasoning about faulty distributed systems. As in the case of verification, a good place to get inspiration from are the models of the distributed algorithms community and our goal is to adapt their ideas to the needs of programming languages. However, this is not trivial: models such as [16] have been developed to be theoretically well-behaved and to simplify reasoning about faulty systems, but little has been done to compile and execute them. Also, if we adopt such a seemingly synchronous abstraction to reason about asynchronous systems, we must tell the programmer which properties of the synchronous abstraction transfer to the actual asynchronous system. Indeed, not every property is preserved. Yet we postulate that a round-based programming abstraction is still superior to the current state of affairs, where implementations need to handle faults in the asynchronous model. In particular, control-state reachability, a classical safety property for concurrent systems, is preserved by asynchrony due to its locality. Also most properties that can be checked using well-structured transition systems are preserved. In fact, it seems that most “reasonable” properties transfer from a synchronous abstraction to asynchronous executions [19, 14], and that this can be formalized.

**The heard-of model [16].** The base primitive of our language is communication-closed rounds [25] and we use the heard-of (*HO*) formalization. Conceptually, processes operate in lock-step, and distributed algorithms consist of rules that determine the new state of a process depending on the state at the beginning of the round and the messages received by the process in the current round. During a round, between the send and reception of the message, an adversarial environment can choose to drop messages. The impact of the environment is captured by the heard-of sets: for a process  $p$  its heard-of set, denoted  $HO(p)$ , contains the processes from which  $p$  may receive messages from in a given round. To model different kinds of network assumptions, the power of the environment is restricted by constraints on the *HO* set, given in linear temporal logic [49]. For instance, in a synchronous system without faults the environment respects  $\Box(\forall p, q. p \in HO(q))$ , which means  $q$  hears from every  $p$ , i.e., that every message is delivered. On the other hand, an asynchronous system with arbitrary faults does not have impose any restriction on the *HO* set. The environment can drop an arbitrary number of messages. The *HO* model can express intermediate failure models between these two extremes [16, Table 1]. For instance, a partially synchronous system with eventual reliable links and at most  $f$  crashes is  $\Diamond(\exists S. |S| > n - f \wedge \forall p. |HO(p)| = S)$ .

**The language.** We are currently working on building a complete infrastructure for implementing fault-tolerant distributed systems. The front-end consists of a domain-specific language and the corresponding specification logic. The DSL is based on the “heard-of model” [16], in which algorithms are structured in communication-closed rounds and have a synchronous semantics. The specification logic is an extension of [23] which is suitable for expressing many agreement properties such as consensus. At the back-end, we have an

automated verifier and compiler+runtime. The verifier checks the algorithms against the specification in the synchronous semantics and the compiler+runtime generates code for asynchronous systems such that a process running the algorithms cannot distinguish between the synchronous semantics and an asynchronous execution. In our system, the programmer writes the core of the application, e.g., the fault-tolerant distributed algorithm, in the high-level DSL. As the round abstraction is suitable for designing distributed algorithms, but not necessarily for the rest of the system, the algorithm interfaces with other software components like a service, e.g., a consensus service. For the verification, we require the programmer to provide an inductive invariant that captures the correctness idea behind the algorithm. The verifier generates Hoare-style verification conditions and tries to discharge them using an SMT solver. Finally, the compiler+runtime fully automatically generate asynchronous distributed code that is guaranteed to preserve the local fault-handling properties that were established for the algorithm.

**Limitations.** Round models also have their downsides. First, they impose a regular control flow. This complicates the encoding different computation paths, e.g., paths the try to speculate a potentially reliable system to reach a decision faster, or recovery paths to be taken after a crash. Protocols, with complex recovery paths should be encoded as two different algorithms: one for the normal scenario, and one for the recovery. We have started to investigating distributed systems build using only the synchronous parallel composition of the participants. Richer forms of composition could support modular composition of different algorithms, e.g., making a new algorithm by running many copies of another algorithm [22]. Finally, the runtime is dependent on the considered fault-model. A round model is a weak form of clock synchronization which requires some environment assumptions. One can guarantee the progress of an implementation of a round-based model with benign faults, by assuming partial synchrony. However, implementing a round model in the byzantine setting is roughly possible with at most  $n/3$  byzantine processes, where  $n$  is the number of replicas. Therefore, using this model does not make sense to implement an algorithm whose assumptions are weaker than those imposed by the implementation of the round model.

**Prototype implementation.** We are currently working on a prototype implementation of DSL in the SCALA programming language. Figure 3 shows the algorithm from Figure 1 in our DSL. The language is designed as a shallow embedding. The parts that require code generation, e.g. resource and state encapsulation, are handled using macros [10] which manipulate the SCALA syntax tree, and the serialization is made transparent using the pickling library [43].

To write an algorithm in this language, a programmer needs to write the program in a sequence of `Round`. Each round has a `send` and a `update` function. Intuitively, all the processes in the system, when executing a given round execute `send` followed by `update`. A runtime is responsible for delivering messages and detecting faults. Therefore, not every message might be received. Furthermore, the runtime guarantees that messages are only visible within the round in which they are sent, even if the system is asynchronous and processes progress at different speed. Late messages are discarded and late processes will try to catch up after receiving message for a “future” rounds.

To verify the algorithm, we assume that the user gives us and the specification and inductive invariants that describe the global state of the system. For our running example, the agreement property is

$$\forall i, j. \text{decided}(i) \wedge \text{decided}(j) \Rightarrow \text{decision}(i) = \text{decision}(j)$$



```

class OTR extends Algorithm[ConsensusIO] {

  val x = new LocalVariable[Int](0)
  val decision = new LocalVariable[Int](-1) //used by the verification
  val decided = new LocalVariable[Boolean](false) //used by the verification
  val callback = new LocalVariable[ConsensusIO](null)

  def process = new Process[ConsensusIO]{

    def init(io: ConsensusIO) {
      callback <- io
      x <- io.initialValue
      decided <- false
    }

    val rounds = Array[Round](
      new Round{
        type A = Int //guarantees that the types of send and receive match

        //minimal most often received value
        def mmor(mailbox: Set[(Int, ProcessID)]): Int = ...

        def send(): Set[(Int, ProcessID)] = broadcast(x)

        def update(mailbox: Set[(Int, ProcessID)]) {
          if (mailbox.size > 2*n/3) {
            x <- mmor(mailbox)
            if (mailbox.filter(msg => msg._1 == x).size > 2*n/3) {
              callback.decide(x)
              decided <- true
              decision <- x
              terminate
            }
          }
        }
      }
    )
  }
}

```

■ **Figure 3** The on-third-rule algorithm implemented in our DSL

and this property is enforced by the following invariant:

$$\forall i. \neg \text{decided}(i) \quad \vee \quad \exists v. |\{i. x(i) = v\}| > 2n/3 \wedge \forall i. \text{decided}(i) \Rightarrow \text{decision}(i) = v.$$

The invariant states that either no value has been decided, or there exists a majority of replicas formed of more than  $2n/3$  replicas, that agree on the value of  $x$ . Notice that the invariant is still simple due to the fact that we use a state invariant that holds only at the boundary between rounds. Therefore, the invariant does not need to refer to messages or the state of communication channels. Communication-closed rounds are key in enabling us to apply automated verification on this class of algorithms. An additional benefit of this model is that it enables us to reason about liveness. Any computation terminates, if there are two rounds in which all processes receive the messages send by a majority  $P$  of cardinality greater than  $2n/3$ . More precisely, in each of these two rounds the values of the HO-sets assigned by the environment satisfy the property

$$\exists P. \forall p. HO(p) = P \wedge |P| > 2n/3,$$

where  $P$  is a set of processed and  $p$  denotes any process in the system.

The verification procedure [23] is roughly based on computing the Venn regions for the sets occurring in the formula, e.g.,  $\{i. x(i) = v\}$ ,  $HO(p)$ ,  $mailbox(p)$ , reasoning about the

cardinality constraints, e.g.,  $|mailbox(p)| > 2n/3$ , and propagating the global conditions appearing in the definition of set comprehensions to constraints on the values of replica's local variables, e.g.,  $x(p) = v$  is inferred from the predicate defining the set comprehension in the invariant.

To summaries, the project aims to find a programming model that strikes a balance between program design and program analysis. More precisely, we are interested in a programming model based on high-level concepts that offer programmers freedom in reasoning about faults and asynchrony. Its formal semantics should enable the development of automated verification tools that establish the correctness of the implementations. We intend to start with one of the most studied classes of algorithms in the literature: consensus algorithms [16, 40] and their weaker forms, e.g., k-set agreement [18] or lattice agreement [26].

**Acknowledgements.** We thank Josef Widder for the fruitful discussions on this subject, Heather Miller and Philipp Haller for their help with the pickling library, and the anonymous reviewers for their comments.

## 5 Conclusion

Building correct software in the absence of faults is already a challenging task. Now, doing so when faults are integral part of the system is an even greater challenge. Addressing this will require adopting new programming models that natively include faults and/or give ways to detect and handle them. Ideally, we should aim for failure friendly programming abstractions where faults are naturally integrate without obscuring the program's logic. More realistically, we can start by identifying the features that makes those system unnecessarily complex and find better alternatives. We briefly sketched how models developed in the distributed algorithms community also solves issues that, in the PL community, make those system hard to implement and severely limit the scope of automated verification for distributed systems.

---

## References

- 1 Parosh Aziz Abdulla, Karlis Cerans, Bengt Jonsson, and Yih-Kuen Tsay. General Decidability Theorems for Infinite-State Systems. In *LICS*, pages 313–321. IEEE, 1996.
- 2 Parosh Aziz Abdulla, Aurore Collomb-Annichini, Ahmed Bouajjani, and Bengt Jonsson. Using Forward Reachability Analysis for Verification of Lossy Channel Systems. *FMSD*, 25(1):39–65, 2004.
- 3 Yehuda Afek and Eli Gafni. Asynchrony from synchrony. In Davide Frey, Michel Raynal, Saswati Sarkar, Rudrapatna K. Shyamasundar, and Prasun Sinha, editors, *Distributed Computing and Networking, 14th International Conference, ICDCN 2013, Mumbai, India, January 3-6, 2013. Proceedings*, pages 225–239, 2013.
- 4 Krzysztof R. Apt and Dexter Kozen. Limits for automatic verification of finite-state concurrent systems. *Inf. Process. Lett.*, 22(6):307–309, 1986.
- 5 Joe L. Armstrong. The development of erlang. In Simon L. Peyton Jones, Mads Tofte, and A. Michael Berman, editors, *Proceedings of the 1997 ACM SIGPLAN International Conference on Functional Programming (ICFP '97), Amsterdam, The Netherlands, June 9-11, 1997.*, pages 196–203. ACM, 1997.
- 6 Martin Biely, Bernadette Charron-Bost, Antoine Gaillard, Martin Hutle, André Schiper, and Josef Widder. Tolerating corrupted communication. In *PODC*, pages 244–253, 2007.
- 7 Fatemeh Borran, Martin Hutle, and André Schiper. Timing analysis of leader-based and decentralized Byzantine consensus algorithms. *J. Braz. Comp. Soc.*, 18(1):29–42, 2012.

- 8 Daniel Brand and Pitro Zafiropulo. On Communicating Finite-State Machines. *J. ACM*, 30(2):323–342, 1983.
- 9 Francisco Brasileiro, Fabíola Greve, Achour Mostefaoui, and Michel Raynal. Consensus in one communication step. In Victor Malyskin, editor, *Parallel Computing Technologies*, volume 2127 of *Lecture Notes in Computer Science*. Springer, 2001.
- 10 Eugene Burmako. Scala macros: Let our powers combine!: On how rich syntax and static types work with metaprogramming. In *Proceedings of the 4th Workshop on Scala, SCALA '13*, pages 3:1–3:10, New York, NY, USA, 2013. ACM.
- 11 Mike Burrows. The chubby lock service for loosely-coupled distributed systems. In *OSDI*, Berkeley, CA, USA, 2006. USENIX Association.
- 12 Tushar D. Chandra, Robert Griesemer, and Joshua Redstone. Paxos made live: An engineering perspective. In *Proceedings of the Twenty-sixth Annual ACM Symposium on Principles of Distributed Computing*, PODC '07, pages 398–407, New York, NY, USA, 2007. ACM.
- 13 Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed systems. *J. ACM*, 43(2):225–267, 1996.
- 14 Mouna Chaouch-Saad, Bernadette Charron-Bost, and Stephan Merz. A reduction theorem for the verification of round-based distributed algorithms. In *RP*, volume 5797 of *LNCS*, pages 93–106, 2009.
- 15 Bernadette Charron-Bost and Stephan Merz. Formal verification of a consensus algorithm in the heard-of model. *Int. J. Software and Informatics*, 3(2-3):273–303, 2009.
- 16 Bernadette Charron-Bost and André Schiper. The heard-of model: computing in distributed systems with benign faults. *Distributed Computing*, 22(1):49–71, 2009.
- 17 Kaustuv Chaudhuri, Damien Doligez, Leslie Lamport, and Stephan Merz. A TLA+ proof system. In *LPAR Workshops*, 2008.
- 18 S. Chaudhuri. More choices allow more faults: Set consensus problems in totally asynchronous systems. *Information and Computation*, 105(1):132 – 158, 1993.
- 19 Henri Debrat and Stephan Merz. Verifying fault-tolerant distributed algorithms in the heard-of model. *Archive of Formal Proofs*, 2012, 2012.
- 20 Ankush Desai, Vivek Gupta, Ethan K. Jackson, Shaz Qadeer, Sriram K. Rajamani, and Damien Zufferey. P: safe asynchronous event-driven programming. In Hans-Juergen Boehm and Cormac Flanagan, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '13, Seattle, WA, USA, June 16-19, 2013*, pages 321–332. ACM, 2013.
- 21 Danny Dolev, Cynthia Dwork, and Larry Stockmeyer. On the minimal synchronism needed for distributed consensus. *J. ACM*, 34:77–97, 1987. <http://doi.acm.org/10.1145/7531.7533>.
- 22 Danny Dolev and Ezra N. Hoch. On self-stabilizing synchronous actions despite byzantine attacks. In Andrzej Pelc, editor, *Distributed Computing, 21st International Symposium, DISC 2007, Lemesos, Cyprus, September 24-26, 2007, Proceedings*, volume 4731 of *Lecture Notes in Computer Science*, pages 193–207. Springer, 2007.
- 23 Cezara Dragoi, Thomas A. Henzinger, Helmut Veith, Josef Widder, and Damien Zufferey. A logic-based framework for verifying consensus algorithms. In Kenneth L. McMillan and Xavier Rival, editors, *VMCAI*, pages 161–181. Springer, 2014.
- 24 Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. Consensus in the presence of partial synchrony. *JACM*, 35(2):288–323, April 1988.
- 25 Tzilla Elrad and Nissim Francez. Decomposition of distributed programs into communication-closed layers. *Sci. Comput. Program.*, 2(3):155–173, 1982.

- 26 Jose M. Falerio, Sriram K. Rajamani, Kaushik Rajan, G. Ramalingam, and Kapil Vaswani. Generalized lattice agreement. In *ACM Symposium on Principles of Distributed Computing, PODC '12, Funchal, Madeira, Portugal, July 16-18, 2012*, pages 125–134, 2012.
- 27 Alain Finkel. A generalization of the procedure of karp and miller to well structured transition systems. In Thomas Ottmann, editor, *Automata, Languages and Programming, 14th International Colloquium, ICALP87, Karlsruhe, Germany, July 13-17, 1987, Proceedings*, volume 267 of *Lecture Notes in Computer Science*, pages 499–508. Springer, 1987.
- 28 Alain Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1-2):63–92, 2001.
- 29 Michael J. Fischer, Nancy A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985.
- 30 Eli Gafni. Round-by-round fault detectors: Unifying synchrony and asynchrony (extended abstract). In Brian A. Coan and Yehuda Afek, editors, *Proceedings of the Seventeenth Annual ACM Symposium on Principles of Distributed Computing, PODC '98, Puerto Vallarta, Mexico, June 28 - July 2, 1998*, pages 143–152, 1998.
- 31 Carl Hewitt, Peter Bishop, and Richard Steiger. A Universal Modular ACTOR Formalism for Artificial Intelligence. In *IJCAI*, pages 235–245, 1973.
- 32 C. A. R. Hoare. Communicating sequential processes. *Commun. ACM*, 21(8):666–677, 1978.
- 33 Patrick Hunt, Mahadev Konar, Flavio P. Junqueira, and Benjamin Reed. Zookeeper: wait-free coordination for internet-scale systems. In *USENIXATC*. USENIX Association, 2010.
- 34 Michael Isard. Autopilot: Automatic data center management. *Operating Systems Review*, 41(2):60–67, April 2007.
- 35 Flavio Paiva Junqueira, Benjamin C. Reed, and Marco Serafini. Zab: High-performance broadcast for primary-backup systems. In *Proceedings of the 2011 IEEE/IFIP International Conference on Dependable Systems and Networks, DSN 2011, Hong Kong, China, June 27-30 2011*, pages 245–256. IEEE, 2011.
- 36 Idit Keidar and Alexander Shraer. Timeliness, failure-detectors, and consensus performance. In *PODC*, pages 169–178, 2006.
- 37 Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 1998.
- 38 Leslie Lamport. Distributed algorithms in TLA (abstract). In *PODC*, 2000.
- 39 Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- 40 Nancy Lynch. *Distributed Algorithms*. Morgan Kaufman, 1996.
- 41 Nancy A. Lynch and Mark R. Tuttle. Hierarchical correctness proofs for distributed algorithms. In Fred B. Schneider, editor, *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing, Vancouver, British Columbia, Canada, August 10-12, 1987*, pages 137–151. ACM, 1987.
- 42 Yanhua Mao, Flavio Paiva Junqueira, and Keith Marzullo. Mencius: Building efficient replicated state machine for wans. In Richard Draves and Robbert van Renesse, editors, *8th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2008, December 8-10, 2008, San Diego, California, USA, Proceedings*, pages 369–384. USENIX Association, 2008.
- 43 Heather Miller, Philipp Haller, Eugene Burmako, and Martin Odersky. Instant pickles: generating object-oriented pickler combinators for fast and extensible serialization. In Antony L. Hosking, Patrick Th. Eugster, and Cristina V. Lopes, editors, *OOPSLA*, pages 183–202, 2013.
- 44 Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I. *Inf. Comput.*, 100(1):1–40, 1992.

- 45 Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, II. *Inf. Comput.*, 100(1):41–77, 1992.
- 46 Iulian Moraru, David G. Andersen, and Michael Kaminsky. There is more consensus in egalitarian parliaments. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles, SOSP '13*, pages 358–372, New York, NY, USA, 2013. ACM.
- 47 Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*, volume 2283 of *Lecture Notes in Computer Science*. Springer, 2002.
- 48 Diego Ongaro and John K. Ousterhout. In search of an understandable consensus algorithm. In Garth Gibson and Nikolai Zeldovich, editors, *2014 USENIX Annual Technical Conference, USENIX ATC '14, Philadelphia, PA, USA, June 19-20, 2014.*, pages 305–319, 2014.
- 49 Amir Pnueli. The temporal logic of programs. *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, 0:46–57, 1977.
- 50 Sergio Rajsbaum. Acm sigact news distributed computing column 5. *SIGACT News*, 32(4):34–58, December 2001.
- 51 Nicolas Schiper, Vincent Rahli, Robbert van Renesse, Mark Bickford, and Robert L. Constable. Shadowdb: A replicated database on a synthesized consensus core. In Michael J. Freedman and Neeraj Suri, editors, *Proceedings of the Eighth Workshop on Hot Topics in System Dependability, HotDep 2012, Hollywood, CA, USA, October 7, 2012*. USENIX Association, 2012.
- 52 Philippe Schnoebelen. Revisiting Ackermann-Hardness for Lossy Counter Machines and Reset Petri Nets. In *MFCS*, pages 616–628, 2010.
- 53 Seved Torstendahl. Open telecom platform. *Ericsson Review*, 1, 1997.
- 54 Tatsuhiro Tsuchiya and André Schiper. Model checking of consensus algorithms. In *26th IEEE Symposium on Reliable Distributed Systems (SRDS 2007), Beijing, China, October 10-12, 2007*, pages 137–148, 2007.
- 55 Tatsuhiro Tsuchiya and André Schiper. Using bounded model checking to verify consensus algorithms. In *DISC*, pages 466–480, 2008.
- 56 Tatsuhiro Tsuchiya and André Schiper. Verification of consensus algorithms using satisfiability solving. *Distributed Computing*, 23(5-6):341–358, 2011.
- 57 Josef Widder, Martin Biely, Günther Gridling, Bettina Weiss, and Jean-Paul Blaquart. Consensus in the presence of mortal Byzantine faulty processes. *Distributed Computing*, 2012.
- 58 Lantian Zheng and Andrew C. Myers. A language-based approach to secure quorum replication. In *Proceedings of the Ninth Workshop on Programming Languages and Analysis for Security, PLAS@ECOOP 2014, Uppsala, Sweden, July 29, 2014*, page 27, 2014.