

## DETECTING UNREALIZABILITY OF DISTRIBUTED FAULT-TOLERANT SYSTEMS

BERND FINKBEINER AND LEANDER TENTRUP

Saarland University, 66123 Saarbrücken  
*e-mail address:* {finkbeiner,tentrup}@cs.uni-saarland.de

**ABSTRACT.** Writing formal specifications for distributed systems is difficult. Even simple consistency requirements often turn out to be unrealizable because of the complicated information flow in the distributed system: not all information is available in every component, and information transmitted from other components may arrive with a delay or not at all, especially in the presence of faults. The problem of checking the distributed realizability of a temporal specification is, in general, undecidable. Semi-algorithms for synthesis, such as bounded synthesis, are only useful in the positive case, where they construct an implementation for a realizable specification, but not in the negative case: if the specification is unrealizable, the search for the implementation never terminates. In this paper, we introduce *counterexamples to distributed realizability* and present a method for the detection of such counterexamples for specifications given in linear-time temporal logic (LTL). A counterexample consists of a set of paths, each representing a different sequence of inputs from the environment, such that, no matter how the components are implemented, the specification is violated on *at least one* of these paths. We present a method for finding such counterexamples both for the classic distributed realizability problem and for the fault-tolerant realizability problem. Our method considers, incrementally, larger and larger sets of paths until a counterexample is found. For safety specifications in weakly ordered architectures we obtain a decision procedure, while counterexamples for full LTL and arbitrary architectures may consist of infinitely many paths. Experimental results, obtained with a QBF-based prototype implementation, show that our method finds simple errors very quickly, and even problems with high combinatorial complexity, like the Byzantine Generals’ Problem, are tractable.

### 1. INTRODUCTION

The goal of program synthesis, and systems engineering in general, is to build systems that satisfy a given specification. Sometimes, however, this goal is unattainable, because the conditions of the specification are *impossible* to satisfy in an implementation. A textbook example for such a case is the *Byzantine Generals’ Problem*, introduced in the early 1980s by Lamport et al. [LSP82]. Three generals of the Byzantine army, consisting of one commander and two lieutenants, need to agree on whether they should “attack” or “retreat.” For this purpose, the commander sends an order to the lieutenants, and all generals then exchange

---

*2012 ACM CCS:* [Theory of computation]: Logic—Modal and temporal logics; [Computer systems organization]: Dependable and fault-tolerant systems and networks.

*Key words and phrases:* Distributed Synthesis, Fault-tolerance, Coordination Logic.

messages with each other, reporting, for example, to one general which messages they have received from the other general. The problem is that one of the generals is a traitor and can therefore not be assumed to tell the truth: the tale of the Byzantine generals is, after all, just an illustration for the problem of achieving fault tolerance in distributed operating systems, where we would like to achieve consensus even if a certain subset of the nodes fail. Of course, we cannot expect the traitor to agree with the loyal generals, but we might still expect a loyal lieutenant to agree with the order issued by a loyal commander, and two loyal lieutenants to reach a consensus in case the commander is the traitor. This specification is, however, unrealizable in the setting of the three generals (and, more generally, in all settings where at least a third of the nodes are faulty).

Detecting unrealizable specifications is of great value because it avoids spending implementation effort on specifications that are impossible to satisfy. If the system consists of a single process, then unrealizable specifications can be detected with *synthesis* algorithms, which detect unrealizability as a byproduct of attempting to construct an implementation. For distributed systems, the problem is more complicated: in order to show that there is no way for the three generals to achieve consensus, we need to argue about the knowledge of each general. The key observation in the Byzantine Generals’ Problem is that the loyal generals have no way of knowing who, among the other two generals, is the traitor and who is the second loyal general. For example, the situation where the commander is the traitor and orders one lieutenant to “attack” and the other to “retreat” is *indistinguishable*, from the point of view of the loyal lieutenant who is ordered to attack, from the situation where the commander is loyal and orders both lieutenants to attack, while the traitor claims to have received a “retreat” order. Since the specification requires the lieutenant to act differently (agree with the other lieutenant vs. agree with the commander) in the two indistinguishable situations, we reach a contradiction.

Since realizability for distributed systems is in general an undecidable problem [PR90], the only available decision procedures are limited to special cases, such as pipeline and ring architectures [KV01b, FS05]. There are semi-algorithms for distributed synthesis, such as *bounded synthesis* [FS13], but the focus is on the search for implementations rather than on the search for inconsistencies: if an implementation exists, the semi-algorithm terminates with such an implementation, otherwise it runs forever. In this paper, we take the opposite approach and study *counterexamples to realizability*. Intuitively, a counterexample collects a sufficient number of scenarios such that, no matter what the implementation does, an error will occur in *at least one* of the chosen scenarios. As specifications, we consider formulas of linear-time temporal logic (LTL). It is straightforward to encode the Byzantine Generals’ Problem in LTL. Another interesting example is the famous CAP Theorem, a fundamental result in the theory of distributed computation conjectured by Brewer [Bre00]. The CAP Theorem states that it is impossible to design a distributed system that provides Consistency, Availability, and Partition tolerance (CAP) simultaneously. We assume there is a fixed number  $n$  of nodes, that every node implements the same service, and that there are direct communication links between all nodes. We use the variables  $\text{req}_i$  and  $\text{out}_i$  to denote input and output of node  $i$ , respectively. The consistency and availability requirements can then be encoded as the LTL formulas  $\bigwedge_{1 \leq i < j \leq n} (\text{out}_i \leftrightarrow \text{out}_j)$  and  $(\bigvee_{1 \leq i \leq n} \text{req}_i) \leftrightarrow (\diamond \bigvee_{1 \leq i \leq n} \text{out}_i)$ . The partition tolerance is modeled in a way that there is always at most one node partitioned from the rest of the system, i.e., we have  $n$  different fault-tolerance scenarios where in every scenario all communication links to one node are faulty.

In both examples, a finite set of input sequences suffices to force the system into violating the specification on at least one of the input sequences. In this paper, we present an efficient method for finding such counterexamples. It turns out that searching for counterexamples is much easier than the classic synthesis approach of establishing unrealizability by the non-existence of strategy trees [PR90, KV01b, FS05]. The difficulty in synthesis is to enforce the consistency condition that the strategy of a process must act the same way in all situations the process cannot distinguish. On the strategy trees, this consistency condition is not an  $\omega$ -regular (or even decidable) property. When analyzing a counterexample, on the other hand, we only check consistency on a specific set of sequences, not on a full tree. This restricted consistency condition is an  $\omega$ -regular property and can, in fact, simply be expressed in LTL as part of the temporal specification. Our QBF-based prototype implementation finds counterexamples for the Byzantine Generals' Problem and the CAP Theorem within just a few seconds.

**Related Work.** To the best of the authors' knowledge, there has been no attempt in the literature to characterize unrealizable specifications for distributed systems beyond the restricted class of architectures with decidable synthesis problems, such as pipelines and rings [KV01b, FS05]. By contrast, there is a rich literature concerning unrealizability for open systems, that is, single-process systems interacting with the environment [Chu63, ALW89, KV97]. Schuppan [Sch12] introduced the notion of unrealizable cores to identify a minimal subformula that contains the reason for unrealizability. In robotics, there have been recent attempts to analyze unrealizable specifications [RKG11]. The results are also focused on the reason for unrealizability, while our approach tries to determine if a specification is unrealizable. Moreover, they only consider the simpler non-distributed synthesis of GR(1) specifications, which is a subset of LTL. There are other approaches concerning unrealizable specifications in the non-distributed setting that also use counterexamples [CHJ08, LDS11]. There, the system specifications are assumed to be correct and the information from the counterexamples are used to modify environment assumptions in order to make the specifications realizable. The Byzantine Generals' Problem is often used as an illustration for knowledge-based reasoning in epistemic logics, see [HM84] for an early formalization. Concerning the synthesis of fault-tolerant distributed systems, there is an approach to synthesize fault-tolerant systems in the special case of strongly connected architectures [DF09]. A preliminary version of this paper appeared as [FT14]. The present paper extends our earlier work by removing restrictions to acyclic architectures and providing a completeness result for safety specifications and weakly ordered architectures, as well as a general encoding of fault-tolerant synthesis.

## 2. DISTRIBUTED REALIZABILITY

A specification is *realizable* if there exists an implementation that satisfies the specification. For distributed systems, the realizability problem is typically stated with respect to a specific system architecture. Figure 1 shows some typical example architectures: an architecture consisting of *independent* processes, a *pipeline* architecture, and a *join* architecture. The architecture describes the communication topology of the distributed system. For example, an edge from  $p_1$  to  $p_2$  labeled with  $(x, b)$  indicates that  $x$  and  $b$  are shared variables between processes  $p_1$  and  $p_2$ , where  $p_1$  writes to  $x$  and  $p_2$  reads from  $b$ . The classic *distributed realizability problem* is to decide whether there exists an implementation (or *strategy*) for

each process in the architecture, such that the joint behavior satisfies the specification. In this case, the distinction between the shared variables is unnecessary as the valuation of the variable that is read from is always equal to the valuation of the variable that is written to. In this paper, we are furthermore interested in the synthesis of fault-tolerant distributed systems, where the processes and the communication between processes may become faulty. Here, the distinction of the shared variables is used to model the different types of faults.

**LTL.** We use linear-time temporal logic (LTL) [Pnu77] to express linear-time properties, that are properties  $P \subseteq (2^\Sigma)^\omega$  for a finite alphabet  $\Sigma$ , also called *atomic propositions*. LTL consists of the temporal operators Next  $\bigcirc$  and Until  $\mathcal{U}$ . The syntax is given by the grammar

$$\varphi ::= x \mid \neg\varphi \mid \varphi \vee \psi \mid \bigcirc\varphi \mid \varphi \mathcal{U} \psi ,$$

where  $x \in \Sigma$ . We define  $\varphi \wedge \psi$  as  $\neg(\neg\varphi \vee \neg\psi)$ , the Weak Until operator  $\varphi \mathcal{W} \psi$  as  $\square\varphi \vee (\varphi \mathcal{U} \psi)$ , and the Release operator  $\varphi \mathcal{R} \psi$  as  $\neg(\neg\varphi \mathcal{U} \neg\psi)$ . We use the standard abbreviations *true*  $\equiv x \vee \neg x$  and *false*  $\equiv x \wedge \neg x$ , for some  $x \in \Sigma$ , as well as  $\varphi \rightarrow \psi \equiv \neg\varphi \vee \psi$ ,  $\varphi \leftrightarrow \psi \equiv (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$ ,  $\square\varphi \equiv \text{false} \mathcal{R} \varphi$ , and  $\diamond\varphi \equiv \text{true} \mathcal{U} \varphi$ .

For  $i \geq 0$ , the satisfaction of a path  $\sigma \in (2^\Sigma)^\omega$  on position  $i$  with respect to formula  $\varphi$ , denoted by  $\sigma, i \models_{\text{LTL}} \varphi$ , is inductively defined as

- $\sigma, i \models_{\text{LTL}} x \quad :\Leftrightarrow x \in \sigma(i)$ ,
- $\sigma, i \models_{\text{LTL}} \neg\varphi \quad :\Leftrightarrow \sigma, i \not\models_{\text{LTL}} \varphi$ ,
- $\sigma, i \models_{\text{LTL}} \varphi \vee \psi \quad :\Leftrightarrow \sigma, i \models_{\text{LTL}} \varphi \text{ or } \sigma, i \models_{\text{LTL}} \psi$ ,
- $\sigma, i \models_{\text{LTL}} \bigcirc\varphi \quad :\Leftrightarrow \sigma, i + 1 \models_{\text{LTL}} \varphi$ , and
- $\sigma, i \models_{\text{LTL}} \varphi \mathcal{U} \psi \quad :\Leftrightarrow \exists n \geq i. \sigma, n \models_{\text{LTL}} \psi \text{ and } \forall m \in \{i, \dots, n - 1\}. \sigma, m \models_{\text{LTL}} \varphi$  ,

where  $x \in \Sigma$  and  $\varphi, \psi$  are LTL formulas. We say a path  $\sigma \in (2^\Sigma)^\omega$  is satisfied by  $\varphi$ , if  $\sigma, 0 \models_{\text{LTL}} \varphi$ . The language of an LTL formula  $\llbracket \varphi \rrbracket \subseteq (2^\Sigma)^\omega$  is defined as the set of paths that satisfy  $\varphi$ . In Section 6, we use the syntactically restricted fragments *safety LTL* and *co-safety LTL* [KV01a]. These fragments are given by the grammars

$$\varphi ::= x \mid \neg x \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \bigcirc\varphi \mid \square\varphi \mid \varphi \mathcal{R} \psi, \text{ and} \quad (\text{safety LTL})$$

$$\varphi ::= x \mid \neg x \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \bigcirc\varphi \mid \diamond\varphi \mid \varphi \mathcal{U} \psi, \quad (\text{co-safety LTL})$$

respectively, where  $x \in \Sigma$ . Safety and co-safety formulas are dual with respect to negation.

**Coordination Logic.** We use Coordination logic (CL) [FS10] to give uniform and precise definitions of the various realizability problems of interest. CL is a game-based extension of LTL that makes strategies—and their observations—to first class citizens of the logic. CL divides the set of atomic propositions into strategic and coordination variables, where the latter represent observations for the strategies represented by strategic variables. The *strategy quantifier*  $\exists C \triangleright s$  introduces a strategy for  $s$  that bases its decisions only on the history of valuations of the variables in  $C$ . This allows us to pose queries on the *existence* of strategies, with specific observations, within the logic. In order to specify the realizability of distributed systems, we use the strategy quantifier  $\exists C \triangleright s$  to express the existence of an implementation for a process output  $s$  based on input variables  $C$ .




Figure 1: Example architectures

**CL Syntax.** CL formulas contain two types of variables: the set  $\mathcal{C}$  of *input* (or *coordination*) variables, and the set  $\mathcal{S}$  of *output* (or *strategy*) variables. In addition to the usual LTL operators Next  $\bigcirc$ , Until  $\mathcal{U}$ , and Release  $\mathcal{R}$ , CL has the strategy quantifier  $\exists C \triangleright s$ , which introduces an output variable  $s$  whose values must be chosen based on the history of valuations of the inputs  $C$ . The syntax<sup>1</sup> is given by the grammar

$$\varphi ::= x \mid \neg x \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \bigcirc \varphi \mid \varphi \mathcal{U} \varphi \mid \varphi \mathcal{R} \varphi \mid \exists C \triangleright s. \varphi \mid \forall C \triangleright s. \varphi ,$$

where  $x \in \mathcal{C} \cup \mathcal{S}$ ,  $C \subseteq \mathcal{C}$ , and  $s \in \mathcal{S}$ . Beside the standard abbreviations  $true \equiv x \vee \neg x$ ,  $false = x \wedge \neg x$ ,  $\diamond \varphi \equiv true \mathcal{U} \varphi$ , and  $\square \varphi \equiv false \mathcal{R} \varphi$ , we use  $\bigcirc_n \varphi$  as an abbreviation of  $n$  consecutive Next operators.

We denote by  $\mathcal{Q}$  the (possibly empty) *quantification prefix* of a formula and call the remainder the *body*. For  $Q \in \{\exists, \forall\}$ , we use  $\mathcal{Q}_Q$  if the prefix contains only  $Q$ -quantifiers. For the purposes of this paper, it suffices to consider the fragment  $CL_{\exists}$  that contains only existential quantifiers. We furthermore assume that the body is quantifier-free, i.e., that the formulas are in *prenex normal form (PNF)*.

**Examples.** We demonstrate how to express distributed realizability problems in  $CL_{\exists}$  with the example architectures from Fig. 1. The realizability of an LTL formula  $\psi_1$  in the architecture from Fig. 1(a) is expressed by the  $CL_{\exists}$  formula

$$\exists\{a\} \triangleright x. \exists\{b\} \triangleright y. \psi_1 . \quad (2.1)$$

Interprocess communication via a shared variable  $b$ , as in the pipeline architecture from Fig. 1(b), is expressed by separating the information read from  $b$  from the output written to  $b$ . In the following  $CL_{\exists}$  formula we use output variable  $x$  to denote the output written to  $b$ :

$$\exists\{a\} \triangleright x. \exists\{b\} \triangleright y. \square(b = x) \rightarrow \psi_2 \quad (2.2)$$

The LTL specification  $\psi_2$  is qualified by the input-output specification  $\square(b = x)$ , which expresses that  $\psi_2$  is required to hold under the assumption that the information written to  $b$  by process  $x$  is also the information read from  $b$  by process  $y$ . This separation between sent and received information is useful to model faults that disturb the transmission. Failing processes are specified by omitting the input-output specifications that refer to the failing processes. As an example, consider the architecture in Fig. 1(c). The  $CL_{\exists}$  formula

$$\exists\{a\} \triangleright x, y. \exists\{b, c\} \triangleright z. (\square(c = y) \rightarrow \psi_3) \wedge (\square(b = x) \rightarrow \psi_3) \quad (2.3)$$

specifies that there exists an implementation such that  $\psi_3$  is guaranteed to hold even if process  $x$  or  $y$  (but not both) fails.

<sup>1</sup>This logic is called *Extended Coordination Logic* in [FS10].




Figure 2: In (a) and (b) we sketch example strategies for  $y$  and  $x$  satisfying the  $\text{CL}_{\exists}$  formula  $\exists \emptyset \triangleright y. \exists \{a\} \triangleright x. \square (\bigcirc x \leftrightarrow a) \wedge \square (y \leftrightarrow \bigcirc \neg y)$ . In (c) we visualize the resulting computation tree on which the body (LTL) formula is evaluated.

For a formula  $\Phi$ , we differentiate two types of coordination variables, *external* and *internal*. A coordination variable  $c \in \mathcal{C}$  is internal iff the value of  $c$  is uniquely defined by the input-output specifications. In contrast, external coordination variables provide input from the environment. For example, the input  $a$  in (2.2) is external while  $b$  is internal.

**$\text{CL}_{\exists}$  Semantics.** We give a quick definition of the  $\text{CL}_{\exists}$  semantics for formulas in PNF and refer the reader to [FS10] for details and for the semantics of full CL. The semantics is based on *trees* as a representation for strategies and computations. Given a finite set of directions  $\Upsilon$  and a finite set of labels  $\Sigma$ , a (full)  $\Sigma$ -labeled  $\Upsilon$ -tree  $\mathcal{T}$  is a pair  $\langle \Upsilon^*, l \rangle$ , where  $l : \Upsilon^* \rightarrow \Sigma$  assigns each *node*  $v \in \Upsilon^*$  a label  $l(v)$ . For two trees  $\mathcal{T}$  and  $\mathcal{T}'$ , we define the joint valuation  $\mathcal{T} \oplus \mathcal{T}'$  to be the widened tree with the union of both labels. We refer to [FS10] for a formal definition. A path  $\sigma$  in a  $\Sigma$ -labeled  $\Upsilon$ -tree  $\mathcal{T}$  is an  $\omega$ -word  $\sigma_0\sigma_1\sigma_2\dots \in \Upsilon^\omega$  and the corresponding labeled path  $\sigma^{\mathcal{T}}$  is  $(l(\epsilon), \sigma_0)(l(\sigma_0), \sigma_1)(l(\sigma_0\sigma_1), \sigma_2)(l(\sigma_0\sigma_1\sigma_2), \sigma_3)\dots \in (\Sigma \times \Upsilon)^\omega$ .

For a strategy variable  $s$  that is bound by some quantifier  $QC \triangleright s. \varphi$ , we refer to  $C$  as the *scope* of  $s$ , denoted by  $\text{Scope}(s)$ . The meaning of a strategy variable  $s$  is a *strategy* or *implementation*  $f_s : (2^{\text{Scope}(s)})^* \rightarrow 2^{\{s\}}$ , i.e., a function that maps a history of valuations of input variables to a valuation of the output variable  $s$ . We represent the computation of a strategy  $f_s$  as the tree  $\langle (2^{\text{Scope}(s)})^*, f_s \rangle$  where  $f_s$  serves as the labeling function (cf. Fig. 2(a)–(b)).  $\text{CL}_{\exists}$  formulas are interpreted over *computation trees*, that are the joint valuations of the computations for strategies belonging to the strategy variables in  $\mathcal{S}$ , i.e.,  $\bigoplus_{s \in \mathcal{S}} \langle (2^{\text{Scope}(s)})^*, f_s \rangle$  (cf. Fig. 2(c)). Given a  $\text{CL}_{\exists}$  formula  $\mathcal{Q}_{\exists}. \varphi$  in prenex normal form over strategy variables  $\mathcal{S}$  and coordination variables  $\mathcal{C}$ , the formula is satisfiable if there exists a computation tree  $\mathcal{T}$  (over  $\mathcal{S}$ ), such that all paths in  $\mathcal{T}$  satisfy the LTL formula  $\varphi$ , i.e.,  $\forall \sigma \in (2^{\mathcal{C}})^\omega. \sigma^{\mathcal{T}}, 0 \models_{\text{LTL}} \varphi$ .

**From Distributed Realizability to  $\text{CL}_{\exists}$ .** We formally introduce the distributed realizability problem and show reductions from the distributed realizability problem to  $\text{CL}_{\exists}$ . Let  $V$  be a finite set of variables. An *architecture*  $\mathcal{A}$  is a tuple  $(P, p_{\text{env}}, \{I_p\}_{p \in P}, \{O_p\}_{p \in P})$ , where  $P$  is the set of processes and  $p_{\text{env}} \notin P$  is the distinct environment process.  $I_p \subseteq V$  denotes the set of input variables for process  $p$  and  $O_p \subseteq V$  denotes the set of output variables for process  $p$ . We denote by  $I = \bigcup_{p \in P} I_p$  the set of all input variables and by  $O = \bigcup_{p \in P} O_p$  the set of all output variables. The input given by the environment is  $I_{\text{env}} := I \setminus O$ , the communication variables are  $I \cap O$ . While some input may be shared across processes in

the case of broadcasting, the output variables of every pair of processes are assumed to be disjoint, i.e.,  $O_p \cap O_{p'} = \emptyset$  for all  $p \neq p' \in P$ . We represent the architecture  $\mathcal{A}$  by a directed graph  $\mathcal{G}_{\mathcal{A}} = (N, E)$ , where  $N = P \cup P_{env}$  is the set of vertices and  $E = N \times N$  the set of edges. There is an edge between two vertices  $(p, p') \in E$  if  $O_p \cap I_{p'}$ , i.e., there is a communication over shared variables between  $p$  and  $p'$ .

An *implementation* of a process  $p$  is a function  $f_p : (2^{I_p})^* \rightarrow 2^{O_p}$  which maps the history of valuation of the input variables to a subset of output variables. We say an implementation is finite state, if it can be represented by a finite transducer. An implementation for  $\mathcal{A}$  is a set of implementations for each process. The distributed realizability problem for architecture  $\mathcal{A}$  and LTL formula  $\varphi$  is to decide whether there is a finite state implementation for every process in  $\mathcal{A}$  such that the system satisfies  $\varphi$  against the environment, i.e., the joint behavior of the implementations satisfies  $\varphi$  against all input sequences given by the environment:  $\forall \sigma \in (2^I)^\omega. \sigma^T, 0 \models_{\text{LTL}} \varphi$  where  $\mathcal{T} = \bigoplus_{p \in P} \langle (2^{I_p})^*, f_p \rangle$ .

The distributed realizability problem is decidable for the class of weakly ordered architectures [FS05], which includes pipelines and rings. Weakly ordered architectures are characterized by the absence of pairs of processes, called *information forks*, that each have access to some information that is hidden from the other process. Consider a tuple  $(P', V', p, p')$ , where  $P'$  is a subset of the processes  $P \cup \{p_{env}\}$ ,  $V'$  is a subset of the variables disjoint from  $I_p \cup I_{p'}$ , and  $p, p' \in P \setminus (P' \cup \{p_{env}\})$  are two different system processes. Such a tuple is an information fork if  $P'$  together with the edges that are labeled with at least one variable from  $V'$  forms a subgraph rooted in the environment and there exist two nodes  $q, q' \in P'$  that have edges to  $p, p'$ , respectively, but are labeled with incomparable sets of variables (i.e., neither set is a subset of the other). For example, the architecture in Fig. 1(a) contains the information fork  $(\{p_{env}\}, \emptyset, p_1, p_2)$ . A *weakly ordered architecture* is an architecture that does not contain an information fork.

**Definition 2.1.** Given a  $\text{CL}_{\exists}$  formula  $\Phi = \mathcal{Q}_{\exists}. \varphi_{path} \rightarrow \varphi$ , where  $\varphi_{path} = \square \bigwedge_{(s,c) \in R} (c = s)$  for some  $R \subseteq \mathcal{C} \times \mathcal{S}$ , the induced architecture  $\mathcal{A}_{\Phi} = (P, p_{env}, \{I_p\}_{p \in P}, \{O_p\}_{p \in P})$  is defined as follows. By abuse of notation, we use  $V = \mathcal{S} \cup \mathcal{C}$  as the set of variables in the architecture and define a process  $p \in P$  as a set of strategy variables  $p \subseteq \mathcal{S}$ .

Let  $P$  be the quotient of  $\mathcal{S}$  according to the equivalence relation  $\doteq \subseteq (\mathcal{S} \times \mathcal{S})$  with  $s \doteq s'$  if  $\text{Scope}(s) = \text{Scope}(s')$ , i.e., we group the strategy variables with the same inputs together as a single process. For all  $p \in P$ , we define  $I_p$  as  $\text{Scope}(s)$  for some  $s \in p$  and  $O_p$  as the union over the defined communication variables  $\{c \mid (s, c) \in R \text{ for some } s \in p\}$  and the strategy variables that are not used for communication  $\{s \mid (s, c) \notin R \text{ for all } c \in \mathcal{C}\}$ .

**Theorem 2.1.** *The distributed realizability problem over architecture  $\mathcal{A}$  and LTL formula  $\varphi$  can be encoded as  $\text{CL}_{\exists}$  formula  $\Phi = \mathcal{Q}_{\exists}. \varphi_{path} \rightarrow \varphi$  with only prenex existential quantification.*

*Proof.* Consider an arbitrary architecture  $\mathcal{A} = (P, p_{env}, \{I_p\}_{p \in P}, \{O_p\}_{p \in P})$  and an LTL formula  $\varphi$ . We give a  $\text{CL}_{\exists}$  formula  $\Phi$  that is satisfiable if, and only if,  $\varphi$  is realizable in  $\mathcal{A}$ . For variable  $v \in V$ , we denote the coordination variable and strategy variable used in the encoding by  $c_v$  and  $s_v$ , respectively. Analogously we use  $C := \{c_v \mid v \in I\}$  as the set of coordination variables and  $S := \{s_v \mid v \in O\}$  as the set of strategy variables.

For each process  $p \in P$  we introduce a set of strategy variables  $S_p := \{s_v \mid v \in O_p\}$  with the scope  $C_p := \{c_v \mid v \in I_p\}$ . We take care of the input-output specifications by restricting the paths such that  $\varphi_{path} := \bigwedge_{v \in I \cap O} \square (c_v = s_v)$  holds, i.e., we only consider paths where the valuation of the strategy variables  $S$  representing the output variables is equal to the

valuation of the coordination variables  $I \cap O$  representing the input variables. The resulting formula  $\Phi$  for processes  $p_1, \dots, p_n$  is

$$\exists C_{p_1} \triangleright S_{p_1} \dots \exists C_{p_n} \triangleright S_{p_n} \cdot \varphi_{path} \rightarrow \varphi .$$

The correctness of the encoding follows from the semantics of  $CL_{\exists}$  and  $\mathcal{A} = \mathcal{A}_{\Phi}$ .  $\square$

**Realizability under Byzantine faults.** We model the occurrence of a Byzantine fault as a modification of the architecture, where a process reading from some shared variable no longer reads the output written by the other process, but instead some arbitrary input. An implementation is fault-tolerant if it works both in the original architecture and in the modified architecture. We encode *fault-tolerant realizability* as  $CL_{\exists}$  formulas of the general form

$$\Phi = \mathcal{Q}_{\exists} \cdot \bigwedge_{1 \leq i \leq n} (\varphi_{path_i} \rightarrow \varphi_i) , \quad (2.4)$$

where the formulas  $\varphi_{path_i}$  are obtained from  $\varphi_{path}$  by omitting constraints  $\square(c = s)$ , which indicates that there is a Byzantine fault on the shared variable  $s$ : the coordination variable  $c$  can deviate arbitrarily from  $s$ .

A disadvantage of this encoding is that we can no longer partition the coordination variables into internal and external variables, because the coordination variable  $c$  is external in the architecture where the shared variable is faulty and internal in the other architectures. In order to retain the partitioning, we use the following satisfiability preserving transformation  $partition(\Phi)$ : For each coordination variable  $c$  that is neither internal nor external, we introduce a fresh coordination variable  $c^*$  that represents the external information given in the fault architectures. Consequently, we add the condition  $\square(c = c^*)$  to all path specifications  $\varphi_{path_i}$  where  $c$  is not contained, thus, making the value  $c$  deterministic in all architectures. In the transformed formula  $partition(\Phi)$ ,  $c$  is an internal coordination variable as its value is in every architecture uniquely determined and  $c^*$  is an external coordination variable as it provides environmental input to the system.

### 3. COUNTEREXAMPLES TO DISTRIBUTED REALIZABILITY

We now introduce *counterexamples to realizability*, which are actually counterexamples to satisfiability for the  $CL_{\exists}$  formula encoding the realizability problem. The satisfiability problem for a  $CL_{\exists}$  formula in prenex form asks for an implementation for all strategy variables in the quantification prefix of the formula such that the temporal specification in the body is satisfied.

Let  $\Phi = \mathcal{Q}_{\exists} \cdot \varphi$  be a  $CL_{\exists}$  formula in prenex form over coordination variables  $\mathcal{C}$  and strategy variables  $\mathcal{S}$ , where the body of the formula is the LTL formula  $\varphi$ . A *counterexample to satisfiability* for  $\Phi$  is a (possibly infinite) set of paths  $\mathcal{P} \subseteq (2^{\mathcal{C}})^{\omega}$ , such that, no matter what strategies are chosen for the strategy variables in  $\mathcal{S}$ , there exists a path  $\sigma \in \mathcal{P}$  that violates the body  $\varphi$ . Formally,  $\mathcal{P} \subseteq (2^{\mathcal{C}})^{\omega}$  is a counterexample to satisfiability iff, for each  $s \in \mathcal{S}$  and every strategy  $f_s : (2^{Scope(s)})^* \rightarrow 2^{\{s\}}$ , it holds that there exists a path  $\sigma \in \mathcal{P}$  such that  $\sigma^{\mathcal{T}} \models_{LTL} \neg \varphi$  where  $\mathcal{T} = \bigoplus_{s \in \mathcal{S}} \langle (2^{Scope(s)})^*, f_s \rangle$ .

**Proposition 3.1.** *A  $CL_{\exists}$  formula  $\Phi$  over coordination variables  $\mathcal{C}$  and strategy variables  $\mathcal{S}$  is unsatisfiable if, and only if, there exists a counterexample to satisfiability  $\mathcal{P} \subseteq (2^{\mathcal{C}})^{\omega}$ .*

*Proof.* By the semantics of  $CL_{\exists}$  and  $\mathcal{P} = (2^{\mathcal{C}})^{\omega}$ .  $\square$



The distributed realizability problem *without* faults corresponds to  $CL_{\exists}$  formulas of the form  $\Phi = \mathcal{Q}_{\exists}.\varphi_{\text{path}} \rightarrow \varphi$ , where the  $\varphi_{\text{path}}$  defines the architecture  $\mathcal{A}_{\Phi}$ : there is an edge from one strategy variable to another if the input-output specification occurs in  $\varphi_{\text{path}}$ . A *finite* counterexample to satisfiability of  $\Phi$  is a finite set of paths  $\mathcal{P} \subseteq (2^{\mathcal{C}_{\text{ext}}})^{\omega}$  corresponding to external coordination variables  $\mathcal{C}_{\text{ext}} = \{c \mid c \text{ is an external coordination variable in } \Phi\}$ , such that for every implementation  $\mathcal{T}$  there exists a path  $\sigma \in \mathcal{P}$  such that an extension  $\sigma' \in (2^{\mathcal{C}})^{\omega}$  of  $\sigma$  violates  $\varphi$ . We say that  $\sigma' \in (2^{\mathcal{C}})^{\omega}$  is an extension of  $\sigma \in (2^{\mathcal{C}_{\text{ext}}})^{\omega}$  if for all  $i \geq 0$ , it holds that  $\sigma_i = \sigma'_i \cap \mathcal{C}_{\text{ext}}$ . Note that the extension of  $\sigma$  by the valuation of the internal coordination variables is uniquely specified by the input path  $\sigma$  and the system implementation  $\mathcal{T}$ .

**Proposition 3.2.** *A  $CL_{\exists}$  formula  $\Phi = \mathcal{Q}_{\exists}.\varphi_{\text{path}} \rightarrow \varphi$  over coordination variables  $\mathcal{C}$  and strategy variables  $\mathcal{S}$  is unsatisfiable if there exists a finite counterexample to satisfiability  $\mathcal{P} \subseteq (2^{\mathcal{C}_{\text{ext}}})^{\omega}$  of  $\Phi$ .  $\square$*

As an example consider again the  $CL_{\exists}$  formula (2.1)  $\exists\{a\} \triangleright x. \exists\{b\} \triangleright y. \psi_1$ , corresponding to the architecture from Fig. 1(a) in the previous section. Let  $\psi_1 := \square(\bigcirc y \leftrightarrow a)$ , i.e.,  $y$  must output the input  $a$  with an one-step delay. A simple counterexample for this formula consists of two paths  $\mathcal{P}_1 := \{\emptyset^{\omega}, \{a\}^{\omega}\}$  that differ in the values of  $a$ , but not in the values of  $b$ . Since process  $y$  cannot distinguish the two paths, but must produce different outputs, this leads to contradiction. Consider the same formula  $\psi_1$  for the pipeline architecture specified by (2.2)  $\exists\{a\} \triangleright x. \exists\{b\} \triangleright y. \square(b = x) \rightarrow \psi_1$ . The formula is unsatisfiable due to the delay when forwarding the input  $a$  over shared variable  $b$  (cf.  $CL_{\exists}$  semantics in Section 2, the current output of  $x$  is only visible to  $y$  in the next step).  $\mathcal{P}_1$  is a finite counterexample in this case, too: Given an implementation of  $x$  and  $y$ , we extend both paths such that the input-output specification  $\square(b = x)$  is satisfied.

The fault-tolerant distributed realizability problem corresponds to  $CL_{\exists}$  formulas of the form  $\Phi = \mathcal{Q}_{\exists}.\bigwedge_{1 \leq i \leq n} (\varphi_{\text{path}_i} \rightarrow \varphi_i)$ . If  $\varphi_i = \varphi$  for all  $i$ , the formula states that there exists an implementation such that the specification  $\varphi$  should hold in all architectures induced by the path specifications  $\varphi_{\text{path}_i}$ . Omitted specifications  $\square(b = x)$  in one of these formulas  $\varphi_{\text{path}_i}$  represent an arbitrary communication error at input  $b$  in architecture  $i$ . In this case, a counterexample identifies for each implementation one of these architectures where a contradiction occurs. A *finite* counterexample to satisfiability of  $\Phi$  are  $n$  finite sets of paths  $\mathcal{P}_i \subseteq (2^{\mathcal{C}_{\text{ext}}^i})^{\omega}$  each corresponding to external coordination variables  $\mathcal{C}_{\text{ext}}^i$  in the respective architecture  $i$ , such that for any implementation  $\mathcal{T}$  there exists an architecture  $j$  and a path  $\sigma \in \mathcal{P}_j$  such that an extension  $\sigma' \in (2^{\mathcal{C}})^{\omega}$  of  $\sigma$  violates  $\varphi_j$ .

**Proposition 3.3.** *A  $CL_{\exists}$  formula  $\Phi = \mathcal{Q}_{\exists}.\bigwedge_{1 \leq i \leq n} (\varphi_{\text{path}_i} \rightarrow \varphi_i)$  over coordination variables  $\mathcal{C}$  and strategy variables  $\mathcal{S}$  is unsatisfiable if there exists a finite counterexample to satisfiability of  $\Phi$ .  $\square$*

A counterexample for the  $CL_{\exists}$  specification (2.3)

$$\exists\{a\} \triangleright x, y. \exists\{b, c\} \triangleright z. (\square(c = y) \rightarrow \psi_3) \wedge (\square(b = x) \rightarrow \psi_3)$$

introduces paths for inputs as well as for every faulty node such that some paths model the exact input-output specification and other paths model the arbitrary node failures. The node that reads from a shared variable can, in contrast to incomplete information, react differently on the given paths, but the reaction must be consistent regarding its observations on all paths. Consider for example the specification  $\psi_3 := (\bigcirc_2 z \leftrightarrow a)$  for the CL formula

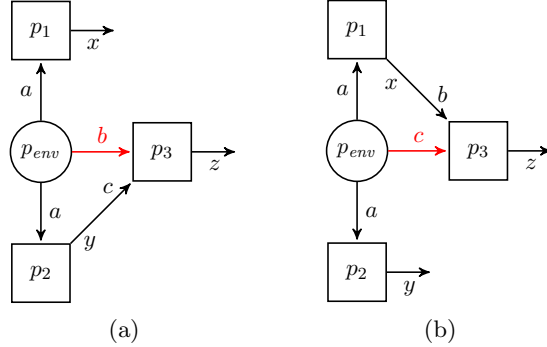


Figure 3: Visual interpretation of the fault-tolerance specification (2.3): There exists a strategy for  $x$ ,  $y$ , and  $z$  such that the specification is satisfied in both architectures.

in (2.3), that is, process  $z$  should output the input  $a$  of nodes  $x$  and  $y$ . In both architectures, depicted in Fig. 3, we introduce additional paths for the coordination variable that is omitted in the input-output specification, i.e.,  $b$  and  $c$  for the first and second conjunct, respectively. Process  $z$  cannot tell which of its inputs come from a faulty node. Since  $z$  must produce the same output on two paths it cannot distinguish, the implementation of  $z$  contradicts the specification in either architecture.

#### 4. COMPUTING FINITE COUNTEREXAMPLES

We encode the existence of finite counterexamples to realizability as a formula of *quantified propositional temporal logic (QPTL)* [KP95]. QPTL extends LTL with a *path quantifier*  $\exists p$ , where a path  $\sigma \in 2^\Sigma$  satisfies  $\exists p. \varphi$  at position  $i \geq 0$ , denoted by  $\sigma, i \models_{\text{QPTL}} \exists p. \varphi$ , if there exists a path  $\sigma' \in 2^{\Sigma \cup \{p\}}$  which coincides with  $\sigma$  except for the newly introduced atomic proposition  $p$ , such that  $\sigma', i \models_{\text{QPTL}} \varphi$ . We define the universal path quantifier  $\forall p. \varphi$  as  $\neg \exists p. \neg \varphi$ . In the following encoding, we use the path quantifier to explicitly name the paths in the counterexample.

**Distributed Realizability.** We consider first the distributed realizability problem, represented by  $\text{CL}_\exists$  formula  $\Phi = \mathcal{Q}_\exists. \varphi_{\text{path}} \rightarrow \varphi$ . Before proceeding with the general encoding, we show an example query that encodes the existence of a counterexample to satisfiability of  $\text{CL}_\exists$  formula (2.2)  $\exists\{a\} \triangleright x. \exists\{b\} \triangleright y. \Box(b = x) \rightarrow \Box(\bigcirc y \leftrightarrow a)$  by using two external paths, represented by the path variables  $a_1$  and  $a_2$ . The QPTL query is

$$\exists a_1, a_2. \forall x_1, x_2. \forall y_1, y_2. \exists b_1, b_2. \\ (x_1 = x_2 \mathcal{R} a_1 \neq a_2) \wedge (y_1 = y_2 \mathcal{R} b_1 \neq b_2) \rightarrow \quad (4.1)$$

$$\bigwedge_{i \in \{1,2\}} \Box(b_i = x_i) \wedge \bigvee_{i \in \{1,2\}} \Diamond(\bigcirc y_i \leftrightarrow a_i) . \quad (4.2)$$

This query is satisfiable, one satisfying assignment for  $a_1$  and  $a_2$  is  $\{a_1\}\emptyset^\omega$  and  $\emptyset^\omega$ , respectively. The assignments for  $b_i$  is a (alpha renamed) copy of the assignment for  $x_i$ , satisfying the condition  $\Box(b_i = x_i)$  for  $i \in \{1, 2\}$ . As already mentioned in the previous section, the value of  $b_i$  is uniquely determined by the strategies. Thus, we could optimize the query

by removing  $b_i$  altogether. However, we need the distinction between reading and writing variables later in this section.

The valuations of  $a_1$  and  $a_2$  represent two *paths* in the computation tree of the original  $\text{CL}_{\exists}$  formula. In the query above, the strategies for  $x$  and  $y$  are also evaluated on these two paths. To make the connection between the path variables  $x_1$  and  $x_2$  and the strategy variable  $x$ , we introduce *consistency* conditions between all paths considered in the query. These consistency conditions (4.1) state that valuations of universal path variables are only valid if they can be generated by the corresponding strategy. For example, the consistency condition for  $x$ ,  $(x_1 = x_2 \ \mathcal{R} \ a_1 \neq a_2)$ , states that  $a_1$  and  $a_2$  must be different in some step  $i \geq 0$  before  $x_1$  and  $x_2$  can be different in some step  $j > i$ . Hence, according to (4.1),  $x_1 = x_2$  and  $y_1 = y_2$  hold in the first step. As  $b_1 = b_2$  holds in the first step as well, the consistency condition for  $y_1 = y_2$  extends to the second step as well. Hence, we have two *paths*, given by the assignments for  $a_1$  and  $a_2$ , where  $a_1$  and  $a_2$  differ in the first position, but  $y_1 = y_2$  in the second position. Consequently, either  $(\bigcirc y_1 \leftrightarrow a_1)$  or  $(\bigcirc y_2 \leftrightarrow a_2)$  holds and we have obtained a counterexample to satisfiability of the  $\text{CL}_{\exists}$  formula.

In the following, we introduce the general encoding for computing finite counterexamples of  $\text{CL}_{\exists}$  formulas  $\Phi = \mathcal{Q}_{\exists}. \varphi_{\text{path}} \rightarrow \varphi$  by bounding the number of paths regarding the *external* coordination variables  $\mathcal{C}_{\text{ext}}$ . The bound on the number of paths is given as a function  $K : \mathcal{C} \rightarrow \mathbb{N}$  that maps each coordination variable to the number of branchings that should be considered for this variable. W.l.o.g. we can assume that  $K(c) = 0$  for every internal coordination variable  $c \in \mathcal{C}_{\text{int}}$ . For example, for coordination variables  $a$  and  $b$ , and  $K(a) = K(b) = 1$ , we encode 4 different paths, one per possible combination for the two paths for each variable. We fix an arbitrary strict order  $\prec \subseteq \mathcal{C} \times \mathcal{C}$  between the coordination variables. For a set  $C \subseteq \mathcal{C}$ , we identify  $K(C)$  by the vector in  $\mathbb{N}^{|\mathcal{C}|}$  where the position of the value  $K(c)$  for a coordination variable  $c \in C$  is determined by  $\prec$ . For our encoding in QPTL, we use the following helper functions:

- $\text{deps}(v)$  returns the set of external coordination variables that *influence* variable  $v$ . An external coordination variable  $c \in \mathcal{C}_{\text{ext}}$  influences variable  $v$  if there is a directed path from  $c$  to  $v$  in  $\mathcal{A}_{\Phi}$ . For example in the architecture of Fig. 1(c),  $x$ ,  $y$ , and  $z$  are influenced by  $a$ . A coordination variable is influenced by itself.
- $\text{branches}(C, K)$  returns the set of branches belonging to coordination variables  $C$ . A branch  $\pi$  is referenced by a vector in  $\mathbb{N}^{|\mathcal{C}|}$  and the set of branches is

$$\left\{ (n_{c_1}, \dots, n_{c_k}) \mid \{c_1 \prec \dots \prec c_k\} = C \text{ and } 1 \leq n_c \leq 2^{K(c)} \text{ for all } c \in C \right\}$$

- $\text{paths}(C, K)$  and  $\text{strategies}(S, K)$  represent the set of (path) variables in the QPTL formula that belong to the variables of the  $\text{CL}_{\exists}$  formula. For a variable  $v \in C \cup S$  it introduces for each branch  $\pi \in \text{branches}(\text{deps}(v), K)$  a separate variable  $p_{\pi}^v$  that represents the variable  $v$  belonging to this branch  $\pi$ . Formally, we define the sets

$$\begin{aligned} \text{paths}(C, K) &:= \{p_{\pi}^c \mid c \in C \wedge \pi \in \text{branches}(\text{deps}(c), K)\} \text{ and} \\ \text{strategies}(S, K) &:= \{p_{\pi}^s \mid s \in S \wedge \pi \in \text{branches}(\text{deps}(s), K)\} . \end{aligned}$$

- $\text{header}(S, K)$  creates the alternating introductions of strategies and paths according to the (partial) order given the subset relation of the strategy variables on  $\text{deps}$ . For every strategy variable  $s \in S$  we introduce all paths belonging to external coordination variables

$c \in \text{deps}(s)$  prior to  $s$  and avoid duplicate path introductions:

$$\begin{aligned} \text{header}(S, K) &:= \exists \text{paths}(\text{deps}(s_1), K) \forall \text{strategies}(\{s_1\}, K) \\ &\quad \exists \text{paths}(\text{deps}(s_2) \setminus \text{deps}(s_1), K) \forall \text{strategies}(\{s_2\}, K) \\ &\quad \dots \\ &\quad \exists \text{paths}(\text{deps}(s_n) \setminus \left( \bigcup_{i=1, \dots, n-1} \text{deps}(s_i) \right), K) \forall \text{strategies}(\{s_n\}, K) \\ &\quad \exists \text{paths}(\mathcal{C}_{\text{int}}) , \end{aligned}$$

where  $s_1, \dots, s_n$  are ordered such that for all  $i, j$  with  $i < j \leq n$ , either  $\text{deps}(s_i) \subseteq \text{deps}(s_j)$  or  $\text{deps}(s_j) \not\subseteq \text{deps}(s_i)$ , i.e., either  $\text{deps}(s_i)$  is a subset of  $\text{deps}(s_j)$  or both are incomparable.

- $\text{consistent}(S, K)$  specifies the consistency condition for the variables belonging to the strategy variables on the different branches. The variables  $p_{\pi_1}^s, \dots, p_{\pi_k}^s$  belonging to a strategy variable  $s \in S$  must be equal as long as the coordination variables in the scope of  $s$  on the branches  $\pi_1, \dots, \pi_k$  are equal. This is expressible in LTL as we only consider a finite number of branches. In detail, the formula

$$\text{consistent}(S, K) := \bigwedge_{s \in S} \bigwedge_{\substack{(\pi, \pi') \in \\ \text{branches}(\text{deps}(s), K)^2}} \left( p_{\pi}^s = p_{\pi'}^s \mathcal{R} \left( \bigvee_{c \in \text{Scope}(s)} p_{\pi}^c \neq p_{\pi'}^c \right) \right) \quad (4.3)$$

states that for every strategy variable  $s \in S$  and every pair of branches  $(\pi, \pi')$  that belongs to  $s$ , we ensure that the valuation of  $s$  on these two branches differ only if one of the variables in the scope of  $s$  was different in a prior step.

Finally, we define the QPTL encoding for  $\text{CL}_{\exists}$  formula  $\Phi$  and function  $K : \mathcal{C} \rightarrow \mathbb{N}$  as

$$\begin{aligned} \text{unsat}_{\text{dist}}(\Phi, K) &:= \text{header}(\mathcal{S}, K). \text{consistent}(\mathcal{S}, K) \rightarrow \\ &\quad \left( \bigwedge_{\pi \in \text{branches}(\mathcal{C}, K)} \varphi_{\text{path}}(\pi) \right) \wedge \left( \bigvee_{\pi \in \text{branches}(\mathcal{C}, K)} \neg \varphi(\pi) \right) , \end{aligned} \quad (4.4)$$

where  $\varphi(\pi)$  is the initialization of LTL formula  $\varphi$  on branch  $\pi$ , that is we exchange  $v$  by  $p_{\pi'}^v$  for  $v \in \mathcal{C} \cup \mathcal{S}$  where  $\pi'$  is the subvector of  $\pi$  that contains the values for coordination variables in  $\text{deps}(v)$ .

**Fault-tolerant Realizability.** In the case of possible failures, the  $\text{CL}_{\exists}$  formulas  $\Phi$  have the more general form  $\mathcal{Q}_{\exists}. \bigwedge_{1 \leq i \leq n} (\varphi_{\text{path}_i} \rightarrow \varphi_i)$ . In this setting, the set of coordination variables  $\mathcal{C}$  is in general not partitioned into external and internal coordination variables. In the first step, we apply the transformation  $\text{partition}(\Phi)$  given in Section 2 in order to get a partitioning into external and internal coordination variables. Furthermore, we generalize  $\text{deps}(v)$  to multiple architecture: An external coordination variable  $c \in \mathcal{C}_{\text{ext}}$  influences variable  $v$  if there is a directed path from  $c$  to  $v$  in one of the architectures. The QPTL query for counterexamples to  $\text{CL}_{\exists}$  formula  $\Phi$  and a bound on the number of paths, given by the functions  $K_1 \dots K_n : \mathcal{C} \rightarrow \mathbb{N}$ , is

$$\begin{aligned} \text{unsat}_{\text{fault}}(\Phi, K_1, \dots, K_n) &:= \text{header}(\mathcal{S}, K). \text{consistent}(\mathcal{S}, K) \rightarrow \\ &\quad \bigvee_{1 \leq i \leq n} \left( \bigwedge_{\pi \in \text{branches}(\mathcal{C}, K_i)} \varphi_{\text{path}_i}(\pi) \right) \wedge \left( \bigvee_{\pi \in \text{branches}(\mathcal{C}, K_i)} \neg \varphi_i(\pi) \right) , \end{aligned} \quad (4.5)$$

where  $K : \mathcal{C} \rightarrow \mathbb{N}$  is defined as  $K(c) := \max_{1 \leq i \leq n} K_i(c)$  for every  $c \in \mathcal{C}$ .




Figure 4: Visualization of the two architectures given by the fault-tolerance specification  $\exists\{a, c\} \triangleright x. \exists\{b\} \triangleright z, y. (\Box((c = y) \wedge (b = x)) \rightarrow \varphi) \wedge (\Box(c = y) \rightarrow \varphi')$ . In architecture (b), the link between shared variables  $b$  and  $x$  is faulty, hence,  $b$  is an external input.

Consider the example formula

$$\exists\{a, c\} \triangleright x. \exists\{b\} \triangleright z, y. (\Box((c = y) \wedge (b = x)) \rightarrow \varphi) \wedge (\Box(c = y) \rightarrow \varphi')$$

that models the two-way pipeline architecture, depicted in Fig. 3, where the connection between shared variables  $b$  and  $x$  may fail. We apply the transformation  $partition(\Phi)$  and get the formula

$$\exists\{a, c\} \triangleright x. \exists\{b\} \triangleright z, y. \left( \Box((c = y) \wedge (b = x)) \rightarrow \varphi \right) \wedge \left( \Box((c = y) \wedge (b = b^*)) \rightarrow \varphi' \right),$$

where  $b^*$  is a free coordination variable. The QPTL encoding  $unsat_{fault}(\Phi^*, K_1, K_2)$  with  $K_1(a) = K_2(a) = K_1(b^*) = 0$  and  $K_2(b^*) = 1$  is

$$\begin{aligned} & \exists a, b_1^*, b_2^*. \forall x_1, x_2, y_1, y_2, z_1, z_2. \exists b_1, b_2, c_1, c_2. \\ & (x_1 = x_2 \mathcal{R} c_1 \neq c_2) \wedge (y_1 = y_2 \mathcal{R} b_1 \neq b_2) \wedge (z_1 = z_2 \mathcal{R} b_1 \neq b_2) \\ & \rightarrow \Box((c_1 = y_1) \wedge (b_1 = x_1)) \wedge \neg\varphi(1) \end{aligned} \quad (4.6)$$

$$\vee \left( \bigwedge_{i \in \{1, 2\}} \Box((c_i = y_i) \wedge (b_i = b_i^*)) \right) \wedge \left( \bigvee_{i \in \{1, 2\}} \neg\varphi(i) \right). \quad (4.7)$$

In difference to distributed realizability, the encoding states that there must be a violation of the specification in some architecture. Also, not every architecture is challenged with all paths, e.g., in the query (4.6) for the first architecture, we use only one path as the path specification states that the connection between  $b$  and  $x$  is correct, while in the second architecture (4.7) we use two paths that are generated by the external inputs  $b_1^*$  and  $b_2^*$ .

**Theorem 4.1** (Correctness  $unsat_{fault}$ ). *Given a  $CL_{\exists}$  formula  $\Phi = \mathcal{Q}_{\exists}. \bigwedge_{1 \leq i \leq n} (\varphi_{path_i} \rightarrow \varphi_i)$  over coordination variables  $\mathcal{C}$  and strategy variables  $\mathcal{S}$ .  $\Phi$  is unsatisfiable if there exist functions  $K_1 \dots K_n : \mathcal{C} \rightarrow \mathbb{N}$  such that the QPTL formula  $unsat_{fault}(\Phi, K_1, \dots, K_n)$  is satisfiable.*

*Proof.* Let  $\Phi'$  be an arbitrary  $CL_{\exists}$  formula after applying the transformation  $partition(\Phi)$  given in Section 2 in order to recover the partitioning into external and internal coordination variables. Assume there exists functions  $K_1 \dots K_n : \mathcal{C} \rightarrow \mathbb{N}$  such that the QPTL formula  $unsat_{fault}(\Phi', K_1, \dots, K_n)$  is satisfiable. From the satisfiability of this QPTL formula we construct the proof that the  $CL_{\exists}$  formula  $\Phi'$  is unsatisfiable (Proposition 3.3), i.e., that for all strategies there exists a path  $\sigma$  that satisfies  $\bigvee_{1 \leq i \leq n} (\varphi_{path_i} \wedge \neg\varphi_i)$ .

We introduce the following auxiliary notation for this proof. We define  $Prev(v)$  to be the set of variables that are quantified prior to  $v$ .  $Prev_Q(v)$  denotes the set  $Prev(v)$  when

only considering the existential variables ( $Q = \exists$ ), respectively universal variables ( $Q = \forall$ ). For a variable  $v \in \mathcal{C} \cup \mathcal{S}$  and a branch  $\pi$ , we define the branching operator  $v[\pi] := p_{\pi'}^v$ , where  $\pi'$  is the subvector of  $\pi$  that contains only values for variables in  $\text{deps}(v)$ .  $v[\pi]$  is used to relate CL variables to QPTL variables, i.e., it selects the variable in QPTL formula that *belongs* to the branch  $\pi$ . Note that one QPTL variable can belong to more than one branch, for example,  $v[\pi] = v[\pi']$  if the branches  $\pi$  and  $\pi'$  differ only in the values for variables not contained in  $\text{deps}(v)$ . For  $V \subseteq \mathcal{C} \cup \mathcal{S}$ , we write  $V[\pi] := \{v[\pi] \mid v \in V\}$  for the set of QPTL variables belonging to the CL variables  $V$  on branch  $\pi$ .

Let  $s_1, \dots, s_m$  be the strategy variables  $\mathcal{S}$  ordered according to the subset relation on  $\text{deps}(s)$ . Let  $\alpha(s) : (2^{\text{Scope}(s)})^* \rightarrow 2^{\{s\}}$  be an arbitrary strategy for strategy variable  $s \in \mathcal{S}$  and fix a set of strategies  $A = \{\alpha(s) \mid s \in \mathcal{S}\}$ . By the semantics of  $\text{CL}_{\exists}$  we need to show that in the composition of these strategies  $\bigoplus_{\alpha \in A} \alpha$  there exists a path such that the LTL formula  $\bigvee_{1 \leq i \leq n} (\varphi_{\text{path}_i} \wedge \neg \varphi_i)$  is satisfied. We build the paths that satisfy the LTL formula from the branches  $\Pi = \text{branches}(\mathcal{C}, K)$  encoded in the QPTL formula. Let  $\beta(c, \pi) : (2^{\text{Prev}_{\forall}(c[\pi])})^\omega \rightarrow (2^{\{c\}})^\omega$  be the valuation of the existential path variable in the QPTL formula belonging to coordination variable  $c$  on branch  $\pi \in \Pi$ . Note that  $\beta(c, \pi)$  depends on all prior universal quantified variables, especially also on universal quantified variables (representing strategy variables) on different branches. The strategies  $\beta(c, \pi)$  for  $c \in \mathcal{C}_{\text{ext}}$  correspond to the external coordination variables, while strategies  $\beta(c, \pi)$  for  $c \in \mathcal{C}_{\text{int}}$  correspond to the extensions defined in the finite counterexamples (Section 3). We view the existential quantified variables in the formula that correspond to the external coordination variables as inputs to our system strategies  $A$  and get the canonical tuple of paths  $\mathcal{P} = (\sigma_{\pi_1}, \dots, \sigma_{\pi_k})$ , one path  $\sigma_\pi \subseteq (2^{\mathcal{C} \cup \mathcal{S}})^\omega$  for each branch  $\pi \in \Pi$ .

In the construction of the paths  $\mathcal{P}$ , we restrict a strategy  $\alpha_s \in A$  to a single branch, that is the function  $\alpha_s^\omega : (2^{\text{Scope}(s)})^\omega \rightarrow (2^{\{s\}})^\omega$  where  $\alpha_s^\omega(\sigma) = \alpha_s(\epsilon) \cdot \alpha_s(\sigma_0) \cdot \alpha_s(\sigma_0 \sigma_1) \dots$ . Let  $\gamma(s, \pi) : (2^{\text{Prev}_{\exists}(s[\pi])})^\omega \rightarrow (2^{\{s\}})^\omega$  be the strategy corresponding to the strategy variable  $s \in \mathcal{S}$  in the QPTL formula for branch  $\pi$ . We have to make sure that  $\gamma(s, \pi)$  is not more constrained than  $\alpha_s$ , i.e., the QPTL path-strategies  $\gamma(s, \pi)$  must be able to handle all behaviors of  $\alpha_s$ . As  $\text{deps}(s)[\pi] \subseteq \text{Prev}_{\exists}(s[\pi])$  the strategy  $\gamma(s, \pi)$  subsumes the strategy  $\alpha_s^\omega$ , i.e., for every branch  $\pi, \dots, \pi_n$  we can embed every strategy  $\alpha_s$  in strategies  $\gamma(s, \pi_1), \dots, \gamma(s, \pi_n)$  when only considering these branches instead of the whole computation tree. We build the paths  $\mathcal{P}$  according to the inductive definition for every  $\pi \in \Pi$  and  $i \geq 1$

$$\begin{aligned} \sigma_\pi^1 &= \bigcup_{c \in \text{deps}(s_1)} \beta(c, \pi)(\emptyset^\omega), \\ \sigma_\pi^{2i} &= \sigma_\pi^{2i-1} \cup \alpha_{s_i}^\omega(\sigma_\pi^{2i-1} \upharpoonright_{\text{Scope}(s)}), \text{ and} \\ \sigma_\pi^{2i+1} &= \sigma_\pi^{2i} \cup \bigcup_{c \in \text{deps}(s_i)} \beta(c, \pi)(\sigma_\pi^{2i}), \end{aligned}$$

where  $\sigma \upharpoonright_V$  denotes the projection of path  $\sigma$  to the variables  $V$ .

Since the  $\alpha_s^\omega$  are derived from the system strategies, it holds that they satisfy the consistency condition  $\text{consistent}(\{s\}, K)$  (4.3) in the QPTL encoding. From the satisfaction of the QPTL formula, we conclude that

$$\bigvee_{1 \leq i \leq n} \left( \bigwedge_{\pi \in \text{branches}(\mathcal{C}, K_i)} \varphi_{\text{path}_i}(\pi) \wedge \bigvee_{\pi \in \text{branches}(\mathcal{C}, K_i)} \neg \varphi_i(\pi) \right)$$

is satisfied, i.e., there exists an index  $i$  and branches  $\Pi_i = \text{branches}(\mathcal{C}, K_i)$  such that  $\bigwedge_{\pi \in \Pi_i} \varphi_{\text{path}_i}(\pi) \wedge \bigvee_{\pi \in \Pi_i} \neg \varphi_i(\pi)$  holds. From the first conjunct it follows that each path  $\sigma_\pi$  from  $\mathcal{P}$  corresponding to a branch  $\pi \in \Pi_i$  satisfies the path specification  $\varphi_{\text{path}_i}$ . We conclude with the second conjunct that one of the branches from  $\Pi$  violates the specification  $\varphi_i$ .  $\square$

**Corollary 4.2** (Correctness  $\text{unsat}_{\text{dist}}$ ). *Given a  $CL_{\exists}$  formula  $\Phi = \mathcal{Q}_{\exists} \cdot \varphi_{\text{path}} \rightarrow \varphi$  over coordination variables  $\mathcal{C}$  and strategy variables  $\mathcal{S}$ .  $\Phi$  is unsatisfiable if there exists a function  $K : \mathcal{C} \rightarrow \mathbb{N}$  such that the QPTL formula  $\text{unsat}_{\text{dist}}(\Phi, K)$  is satisfiable.*

**Remark 4.3** (Monotonicity of  $K$ ). For the distributed synthesis encoding it holds that  $K$  is monotone with respect to the satisfiability of  $\text{unsat}_{\text{dist}}(\Phi, K)$ , however, for the fault-tolerant synthesis this is not necessarily true because of the complex information dependencies between different architectures. By increasing the number of paths for a coordination variable  $c$  in an architecture  $i$  where  $c$  is not involved in a fault, i.e.,  $\square(c = s)$  is contained in  $\varphi_{\text{path}_i}$ , we restrict the choice of the environment in another architecture (where  $c$  could be involved in a fault). Hence, we have the following monotonicity condition:  $K_1 \dots K_n$  is monotone with respect to the satisfiability of  $\text{unsat}_{\text{fault}}(\Phi, K_1, \dots, K_n)$  if  $K_i(c)$  is only increased if  $c$  is not contained in the input-output specification  $\varphi_{\text{path}_i}$ .

**Example.** We consider again the Byzantine Generals' Problem with three nodes  $g_1$ ,  $g_2$ , and  $g_3$ . The first general is the commander who forwards the input  $v$  that states whether to attack the enemy or not. The encoding as  $CL_{\exists}$  formula is

$$\begin{aligned} \Phi_{\text{bgp}} := & \exists\{v\} \triangleright g_{12}, g_{13} \cdot \exists\{c_{12}\} \triangleright g_{23} \cdot \exists\{c_{13}\} \triangleright g_{32} \cdot \exists\{c_{12}, c_{32}\} \triangleright g_2 \cdot \exists\{c_{13}, c_{23}\} \triangleright g_3 \cdot \\ & (\text{operational}_{2,3} \rightarrow \text{consensus}_{2,3}) \wedge \bigwedge_{i \in \{2,3\}} (\text{operational}_{1,i} \rightarrow \text{correctval}_i) , \end{aligned} \quad (4.8)$$

where we use the following definitions

$$\begin{aligned} \text{operational}_{2,3} & := \square(c_{23} = g_{23} \wedge c_{32} = g_{32}), \\ \text{operational}_{1,3} & := \square(c_{12} = g_{12} \wedge c_{13} = g_{13} \wedge c_{32} = g_{32}), \\ \text{operational}_{1,2} & := \square(c_{12} = g_{12} \wedge c_{13} = g_{13} \wedge c_{23} = g_{23}), \\ \text{consensus}_{i,j} & := \bigcirc_3 (g_i = g_j), \text{ and} \\ \text{correctval}_i & := v \leftrightarrow \bigcirc_3 g_i . \end{aligned}$$

The quantification prefix introduces the strategies for the generals  $g_2$  and  $g_3$ , as well as the communication between the three generals as depicted in the architecture in Fig. 5(a). Note that we omit the vote of the commander  $g_1$  as it is not used in the specification. In the temporal part, we specify which failures can occur. The first conjunct, corresponding to Fig. 5(b), states that the commander is a traitor ( $\text{operational}_{2,3}$ ) which implies that the other two generals have to reach a consensus whether to attack or not ( $\text{consensus}_{2,3}$ ). The other two cases, depicted in Fig. 5(c)–(d), are symmetric and state that whenever one general is traitor the other one should agree on the decision made by the commander. The

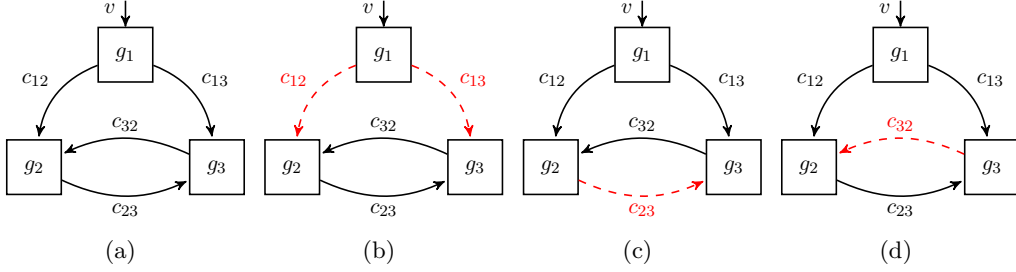


Figure 5: The Byzantine Generals' architecture. Figure (a) shows the architecture in cases all generals are loyal. Figures (b)–(d) show the possible failures, indicated by the dashed communication links.

QPTL encoding  $unsat_{fault}(\Phi_{bgp}, K_1, K_2, K_3)$  is given as

$$\begin{aligned}
& \exists paths(\{v\}, K). \forall strategies(\{g_{12}, g_{13}\}, K). \exists paths(\{c_{12}, c_{13}\}, K). \\
& \forall strategies(\{g_{23}, g_{32}\}, K). \exists paths(\{c_{23}, c_{32}\}, K). \forall strategies(\{g_2, g_3\}, K). \\
& consistent(\{g_{12}, g_{13}, g_{23}, g_{32}, g_2, g_3\}, K) \rightarrow \\
& \left( \left( \bigwedge_{\pi \in branches(\mathcal{C}, K_1)} operational_{2,3}(\pi) \wedge \bigvee_{\pi \in branches(\mathcal{C}, K_1)} \neg consensus_{2,3}(\pi) \right) \vee \right. \\
& \left( \bigwedge_{\pi \in branches(\mathcal{C}, K_2)} operational_{1,3}(\pi) \wedge \bigvee_{\pi \in branches(\mathcal{C}, K_2)} \neg correctval_3(\pi) \right) \vee \\
& \left. \left( \bigwedge_{\pi \in branches(\mathcal{C}, K_3)} operational_{1,2}(\pi) \wedge \bigvee_{\pi \in branches(\mathcal{C}, K_3)} \neg correctval_2(\pi) \right) \right) .
\end{aligned}$$

By using the functions  $K_i$  for  $1 \leq i \leq 3$  which are defined as

$$\begin{aligned}
K_1(v) &= K_2(v) = K_3(v) = 1, \\
K_1(c_{12}) &= K_1(c_{13}) = K_2(c_{23}) = K_3(c_{32}) = 1,
\end{aligned}$$

and 0 otherwise, we get a satisfying QPTL instance and, hence, proved the unsatisfiability of  $\Phi_{bgp}$ . This resembles the manual proof for the Byzantine Generals' Problem. In every situation where one of the generals is a traitor, one has to consider the possible deviations from the loyal behavior.

## 5. COMPLETENESS

In practice, finite external counterexamples are often sufficient to detect unrealizable specifications. In this section, we show that our method is in fact complete for the case of safety specifications and weakly ordered architectures. It turns out that if one of these restrictions is dropped, i.e., if one is interested in general LTL specifications, or in architectures that are not weakly ordered, then finite counterexamples may no longer exist, and the method, therefore, becomes incomplete.

A universal safety (tree) automaton is a tuple  $\mathcal{U} = (\Sigma, \Upsilon, Q, q_0, \delta)$  where  $\Sigma$  denotes a finite set of labels,  $\Upsilon$  denotes a finite set of directions,  $Q$  is a finite set of states,  $q_0 \in Q$  the designated initial state, and  $\delta$  denotes the transition function. The transition function



$\delta : Q \times \Sigma \rightarrow \mathbb{B}_\lambda^+(Q \times \Upsilon)$  maps a state and a input letter to a positive and conjunctive Boolean formula over  $Q \times \Upsilon$ . We allow the abbreviations  $\delta(q, \sigma) = \text{true}$  and  $\delta(q, \sigma) = \text{false}$ .

A  $\Sigma$ -labeled  $\Upsilon$ -transition system is a tuple  $\mathcal{T} = (T, t_0, \tau, o)$  where  $T$  is a finite set of states,  $t_0 \in T$  is the designated initial state,  $\tau : T \times \Upsilon \rightarrow T$  is the transition function, and  $o : T \rightarrow \Sigma$  is the state labeling function. A universal safety automaton accepts a  $\Sigma$ -labeled  $\Upsilon$ -transition system  $\mathcal{T}$  if  $\mathcal{T}$  has a run graph. A run graph is a directed graph  $\mathcal{G} = (G, E)$  that satisfies the following constraints:

- The vertices  $G \subseteq Q \times T$  are a subset of the product of  $Q$  and  $T$ .
- The pair of initial states  $(q_0, t_0) \in G$  is a vertex of  $\mathcal{G}$ .
- For each vertex  $(q, t) \in G$ , the set  $\{(q', v) \in Q \times \Upsilon \mid ((q, t), (q', \tau(t, v))) \in E\}$  satisfies  $\sigma(q, o(t))$ .

Since  $\mathcal{U}$  is universal, the run graph on a transition system is unique.

If we start with a temporal specification instead of a universal safety automaton, we first need the following transformations. As we only consider safety languages, we restrict the temporal specifications to *syntactically safe* [Sis94] formulas, i.e., LTL formulas where the only temporal operator is the weak until  $\varphi \mathcal{W} \psi$ . (The weak until subsumes the globally operator  $\Box \varphi$ , which is equivalent to  $\varphi \mathcal{W} \text{false}$ .)

**Proposition 5.1** ([KV01a]). *Given a syntactically safe LTL formula  $\varphi$ , we can construct a non-deterministic automaton  $\mathcal{N}_{\neg\varphi}$  on finite words that accepts the (not necessarily minimal) good prefixes of the co-safety formula  $\neg\varphi$ . The size of  $\mathcal{N}_{\neg\varphi}$  is on  $2^{\mathcal{O}(|\varphi|)}$ .*

Using this result, we build the universal safety automaton  $\mathcal{U}_\varphi$  by simulating the automaton  $\mathcal{N}_{\neg\varphi}$  along each path: if each path is not a good prefix for the negation of  $\varphi$ , then  $\varphi$  holds on every path.

**Proposition 5.2.** *Given a syntactically safe LTL formula  $\varphi$ , we can construct a universal safety automaton  $\mathcal{U}_\varphi$  with  $2^{\mathcal{O}(|\varphi|)}$  states that accepts a transition system  $\mathcal{T}$  if, and only if,  $\mathcal{T}$  satisfies  $\varphi$ .  $\square$*

We use the bound on the size of a transition system that is accepted by a universal safety automaton  $\mathcal{U}$  in order to derive a upper bound on the number of counterexample paths needed to refute the existence of an implementation.

**Proposition 5.3** ([FS13]). *If a universal safety automaton  $\mathcal{U}$  with  $n$  states accepts a transition system, then  $\mathcal{U}$  accepts a finite transition system  $\mathcal{T}$  with  $(n!)^2$  states.*

**Theorem 5.4.** *Given a syntactically safe LTL formula  $\varphi$  over inputs  $I$  and outputs  $O$ ,  $\varphi$  is unrealizable if, and only if, there exists a finite counterexample.*

*Proof.* Let  $\mathcal{U}$  be the universal safety automaton for  $\varphi$  according to Proposition 5.2. By the contraposition of Lemma 5.3 it holds that  $\mathcal{U}$  rejects all transition systems if  $\mathcal{U}$  rejects all finite transition systems with  $(n!)^2$  states. Hence, all finite transition systems with size  $(n!)^2$  have a finite counterexample path in the run of  $\mathcal{U}$ . We show that every minimal counterexample path is bounded by  $d = (n!)^2 \cdot p$ , where  $p$  is the longest loop-free path in the safety automaton  $\mathcal{U}$ .

Assume by contradiction that there exists a transition system  $\mathcal{T}$  of size  $(n!)^2$ , where the length of the minimal counterexample path  $\sigma \in (Q \times T)^*$  exceeds  $d$ . As  $d$  is the size of the product of  $\mathcal{U}$  and  $\mathcal{T}$ , there must be a repetition on  $\sigma$ , i.e., there exists  $i < j \leq |\sigma|$  such that  $\sigma[i] = \sigma[j]$ . Thus, we can shorten the counterexample to  $\sigma_{0..i-1} \cdot \sigma_{j..n-1}$  violating our minimality assumption.

As the length of the minimal counterexample is bounded by  $d$ , all minimal counterexamples are contained in the full tree of depth  $d$  that has  $(2^{|I|})^d$  paths.  $\square$

This argumentation can be generalized to weakly ordered architectures since we have an upper bound on the size of implementations, too. Given an architecture  $\mathcal{A}$  and an implementation for each process, we call the resulting system the distributed product according to architecture  $\mathcal{A}$  as the system state consist of the product of the states of the process implementations. The meaning of the distributed product is a strategy  $(2^{I_{env}}) \rightarrow 2^O$  that maps finite input sequences from the environment to the valuations of the process outputs.

**Proposition 5.5** ([FS05]). *For a weakly ordered architecture and a realizable specification  $\varphi$ , the size of the smallest implementation of every process is  $n$ -exponential, where  $n$  is the number of processes.*

**Theorem 5.6.** *Given a syntactically safe LTL formula  $\varphi$  and a weakly ordered architecture  $\mathcal{A}$ ,  $\varphi$  is unrealizable in  $\mathcal{A}$  if, and only if, there exists a finite counterexample.*

*Proof.* Let  $\mathcal{U}$  be the universal safety automaton for  $\varphi$ . By Proposition 5.5 and the unrealizability of  $\varphi$  in  $\mathcal{A}$  it holds that there does not exist  $n$ -exponential strategies for the processes. Let  $t_1, \dots, t_n$  be arbitrary  $n$ -exponential strategies for the  $n$  processes and  $\mathcal{T} : (2^{I_{env}})^* \rightarrow 2^O$  be the ( $n$ -exponential) *distributed product* according to architecture  $\mathcal{A}$ .

We choose a bound  $d$  that is greater than the size of the product of  $\mathcal{U}$  and  $\mathcal{T}$ . By the same argument as in the proof of Theorem 5.4, it follows that the length of the minimal counterexample for every such  $\mathcal{T}$  is bounded by  $d$ , hence the number of paths is bounded by  $(2^{|I_{env}|})^d$ .  $\square$

Proposition 3.1 states that the characterization of unsatisfiable formulas with counterexamples is complete. Our method, however, searches for counterexamples involving only a bounded number of external paths. In general, this is not enough, as the following propositions show.

**Proposition 5.7.** *For full LTL, the unrealizability of a specification does not imply the existence of a finite counterexample.*

*Proof.* Consider the  $\text{CL}_{\exists}$  formula  $\Phi_{\text{inf}} := \exists\{x\} \triangleright y. \varphi_{\text{inf}}$  with temporal specification  $\varphi_{\text{inf}} := \diamond(y = x)$ .  $\Phi_{\text{inf}}$  is unsatisfiable because for every strategy  $f_y : (2^{\{x\}})^* \rightarrow 2^{\{y\}}$  it holds that  $f_y$  cannot correctly guess the future value of  $x$  on every path, as the formula is evaluated over the full binary tree that branches by the valuation of  $x$  (cf.  $\text{CL}_{\exists}$  semantics in Section 2). Assume for contradiction that a finite set of paths  $P \subseteq (2^{\{x\}})^\omega$  suffices to satisfy  $\neg\varphi_{\text{inf}} = \square(y \neq x)$  against every strategy  $f_y$ . As there are only a finite set of paths  $P$ , there exists a level in the full binary tree such that every node in this level has at exactly one successor path in  $P$ . Choose the strategy  $f_y$  that assigns the nodes in this level with one successor the labeling of the next input. Such a finite state strategy exists because the nodes in each level are distinguishable by the strategy. Hence, for all paths in  $P$  it holds that  $\diamond(y = x)$  and thus no path satisfies  $\neg\varphi_{\text{inf}}$ .  $\square$

**Proposition 5.8.** *For the architecture in Fig. 1(a), the unrealizability of a safety specification does not imply the existence of a finite counterexample.*

*Proof.* Let  $D$  be a deterministic Turing machine that implements a unary counter. We use the encoding of  $D$  into the realizability problem of a safety LTL formula  $\varphi_D$  in the architecture of Fig. 1(a) [Sch14]. There does not exist a finite state implementation that

satisfies  $\varphi_D$ , but for every finite set of environment input paths, only a finite number of correct configurations of  $D$  can be asserted. Hence, for every such finite set of paths, there exists an implementation that fulfills these assertions.  $\square$

## 6. APPROXIMATIONS

Presently available QPTL solvers were unable to handle even small instances of our problem. In this section, we present two simplifications of the finite counterexample problem and their corresponding encodings. We consider the problem where we restrict the counterexample paths to a finite length by using a *weak monadic second order logic of one successor (WS1S)* encoding. Afterwards, we show a practical encoding in *quantified Boolean formulas (QBF)* that bounds the length of a counterexample path. We only consider the more general form of  $\text{CL}_{\exists}$  formulas modeling the fault-tolerance case, i.e.,  $\Phi = \mathcal{Q}_{\exists} \cdot \bigwedge_{1 \leq i \leq n} (\varphi_{\text{path}_i} \rightarrow \varphi_i)$ .

**From QPTL to S1S.** We give a short introduction to the monadic second-order theory of one successor (S1S) and show the equisatisfiable reduction from QPTL to S1S. Let  $V_1 = \{x, y, \dots\}$  and  $V_2 = \{X, Y, \dots\}$  be finite sets of first-order and second-order variables, respectively. A term  $t$  is given by the grammar

$$t ::= 0 \mid x \mid S(t) ,$$

where  $x$  is a first-order variable and  $S$  denotes the successor function of the natural numbers. We build formulas  $\varphi$  over terms by using the grammar

$$\varphi ::= t \in X \mid t = t \mid \neg \varphi \mid \varphi \vee \varphi \mid \exists x. \varphi \mid \exists X. \varphi ,$$

where  $t$  is a term,  $x$  is a first-order variable, and  $X$  is a second-order variable. We define  $\forall X. \varphi$  as  $\neg \exists X. \neg \varphi$ ,  $x \notin X$  as  $\neg(x \in X)$ , and  $x \neq y$  as  $\neg(x = y)$ . Furthermore, we use the abbreviations  $X \subseteq Y \equiv \forall z. (z \in X \rightarrow z \in Y)$  and  $X = Y \equiv X \subseteq Y \wedge Y \subseteq X$ .

The semantics of S1S is defined over first-order and second-order valuations  $\sigma_1 : V_1 \rightarrow \omega$  and  $\sigma_2 : V_2 \rightarrow 2^\omega$ , respectively. The semantics of terms is defined as

- $[0]_{\sigma_1} = 0$ ,
- $[x]_{\sigma_1} = \sigma_1(x)$ , and
- $[S(t)]_{\sigma_1} = [t]_{\sigma_1} + 1$  ,

where  $t$  is a term and  $x \in V_1$ . The semantics of formulas is defined as

- $\sigma_1, \sigma_2 \models_{\text{S1S}} t \in X \quad :\Leftrightarrow [t]_{\sigma_1} \in \sigma_2(X)$ ,
- $\sigma_1, \sigma_2 \models_{\text{S1S}} t_1 = t_2 \quad :\Leftrightarrow [t_1]_{\sigma_1} = [t_2]_{\sigma_1}$ ,
- $\sigma_1, \sigma_2 \models_{\text{S1S}} \neg \varphi \quad :\Leftrightarrow \sigma_1, \sigma_2 \not\models_{\text{S1S}} \varphi$ ,
- $\sigma_1, \sigma_2 \models_{\text{S1S}} \varphi \vee \psi \quad :\Leftrightarrow \sigma_1, \sigma_2 \models_{\text{S1S}} \varphi$  or  $\sigma_1, \sigma_2 \models_{\text{S1S}} \psi$ ,
- $\sigma_1, \sigma_2 \models_{\text{S1S}} \exists x. \varphi \quad :\Leftrightarrow \exists a \in \omega. \sigma'_1(y) = \begin{cases} \sigma_1(y) & \text{if } y \neq x \\ a & \text{otherwise} \end{cases} \wedge \sigma'_1, \sigma_2 \models_{\text{S1S}} \varphi$ , and
- $\sigma_1, \sigma_2 \models_{\text{S1S}} \exists X. \varphi \quad :\Leftrightarrow \exists A \subseteq \omega. \sigma'_2(Y) = \begin{cases} \sigma_2(Y) & \text{if } Y \neq X \\ A & \text{otherwise} \end{cases} \wedge \sigma_1, \sigma'_2 \models_{\text{S1S}} \varphi$  ,

where  $t, t_1, t_2$  are terms,  $\varphi, \psi$  are formulas,  $x, y \in V_1$ , and  $X, Y \in V_2$ . A formula  $\varphi$  is satisfiable, if there exists first-order and second-order valuations  $\sigma_1$  and  $\sigma_2$  for the free variables such that  $\sigma_1, \sigma_2 \models_{\text{S1S}} \varphi$ .

It is well known that QPTL and S1S are equally expressive [SVW85]. We use the following transformation from QPTL to S1S.

**Lemma 6.1.** *For every QPTL formula  $\varphi$  there exists an equisatisfiable S1S formula  $\psi$ .*

*Proof.* We encode the (infinite) sequences in the QPTL formula as second-order valuations in the S1S formula. For a path variable  $p$ , we denote by  $V^p$  the corresponding second-order variable. For a sequence  $\rho \in (2^{\Sigma})^\omega$  and the corresponding second-order valuation  $\sigma_2(V^p) \subseteq \omega$  we impose the invariant that  $p \in \rho_i$  if and only if  $i \in \sigma_2(V^p)$  for every  $i \geq 0$ . Given a QPTL formula  $\varphi$  over a finite set of atomic propositions  $\Sigma$  and an arbitrary S1S term  $t$  over  $V_1 = \emptyset$ . We define a S1S formula  $s1s(\varphi, t)$  over  $V_2 = \{V^p \mid p \in \Sigma\}$  such that for all  $\rho \in (2^\Sigma)^\omega$  it holds that  $\rho, [t]_{\sigma_1} \models_{\text{QPTL}} \varphi$  if and only if  $\sigma_1, \sigma_2 \models_{\text{S1S}} s1s(\varphi, t)$ , where  $\sigma_2(V^p) = \{i \in \omega \mid p \in \rho_i\}$ :

- $s1s(p, t) \quad := t \in V^p$ , for  $p \in \Sigma$ ,
- $s1s(\neg\varphi, t) \quad := \neg s1s(\varphi, t)$ ,
- $s1s(\varphi \vee \psi, t) \quad := s1s(\varphi, t) \vee s1s(\psi, t)$ ,
- $s1s(\bigcirc \varphi, t) \quad := s1s(\varphi, S(t))$ ,
- $s1s(\varphi \mathcal{U} \psi, t) \quad := \exists y. (y \geq t \wedge s1s(\psi, y) \wedge \forall z. (t \leq z < y \rightarrow s1s(\varphi, z)))$ , and
- $s1s(\exists p. \varphi, t) \quad := \exists V^p. s1s(\varphi, t)$  ,

where  $y$  and  $z$  are fresh first order variables. □

**Finite-length Counterexamples.** When we restrict the counterexamples to finite length, we lose the ability to detect violations of liveness properties. Thus, we restrict the LTL specifications  $\varphi_i$  in  $\Phi$  to syntactic *safety* properties [KV01a]. As these specifications are negated in the encoding, the resulting LTL specifications describe *co-safety* properties [KV01a]. A property  $P \subseteq (2^\Sigma)^\omega$  is co-safety if every  $\rho \in P$  has a *good prefix*, that is a prefix  $w \in (2^\Sigma)^*$  of  $\rho$  such that for all  $\rho' \in (2^\Sigma)^\omega$  it holds that  $w \cdot \rho' \in P$ .

In the weak version of S1S, the second-order quantification is restricted to finite sets. While syntactically equivalent to S1S, we change the semantics  $\models_{\text{S1S}}$  by replacing the second-order quantification rule by

$$\sigma_1, \sigma_2 \models_{\text{WS1S}} \exists X. \varphi \Leftrightarrow \exists \text{finite set } A \subseteq \omega. \sigma'_2(Y) = \begin{cases} \sigma_2(Y) & \text{if } Y \neq X \\ A & \text{otherwise} \end{cases} \wedge \sigma_1, \sigma'_2 \models_{\text{WS1S}} \varphi .$$

We define a S1S formula  $Fin(X)$  that asserts that the valuation of the second-order variable  $X$  is a finite set:

$$Fin(X) := \exists Y. (X \subseteq Y \wedge (\exists y. y \notin Y) \wedge (\forall z. (z \notin Y \rightarrow S(z) \notin Y))) .$$

We define  $ws1s(\varphi, i)$  as  $s1s(\varphi, i)$  with the difference that we interpret the result of  $ws1s(\varphi, i)$  as a WS1S formula and we only apply it to co-safety LTL formulas.

**Lemma 6.2.** *Given a QPTL formula  $\Phi = \mathcal{Q}. \varphi$  where  $\varphi$  is a co-safety LTL formula. If  $ws1s(\Phi, 0)$  is satisfiable then  $\Phi$  is satisfiable.*

*Proof.* Let  $ws1s(\Phi, 0)$  be satisfiable. By Lemma 6.1, it suffices to argue that  $s1s(\Phi, 0)$  is satisfiable. First, we can extend any existential quantification  $\exists X. \varphi$  to infinite quantification as the WS1S restriction to quantification over finite sets are special cases of the S1S semantics (these where  $\exists X. Fin(X) \wedge \varphi$  holds). Next, we show that finite universal quantification is as powerful as infinite universal quantification. Consider the co-safety LTL formula  $\varphi$ . It

holds that every satisfying assignment has a good prefix [KV01a]. Hence, the satisfaction of an infinite word  $\rho \in (2^\Sigma)^\omega$  depends only on a finite prefix. Thus, a satisfying assignment against universal quantified finite sequences implies the satisfaction against universal quantified infinite sequences.  $\square$

The WS1S encoding of  $CL_\exists$  formula  $\Phi$  and functions  $K_1 \dots K_n : \mathcal{C} \rightarrow \mathbb{N}$  is

$$\begin{aligned} \text{unsat}_{\text{fault}}^{WS1S}(\Phi, K_1, \dots, K_n) &:= \text{header}^{WS1S}(\mathcal{S}, K). \text{consistent}^{WS1S}(\mathcal{S}, K) \rightarrow \\ &\bigvee_{1 \leq i \leq n} \left( \bigwedge_{\pi \in \text{branches}(\mathcal{C}, K_i)} \varphi_{\text{path}_i}^{WS1S}(\pi) \right) \wedge \left( \bigvee_{\pi \in \text{branches}(\mathcal{C}, K_i)} \text{ws1s}(\neg\varphi_i(\pi), 0) \right), \end{aligned} \quad (6.1)$$

where  $K : \mathcal{C} \rightarrow \mathbb{N}$  is defined as  $K(c) := \max_{1 \leq i \leq n} K_i(c)$  for every  $c \in \mathcal{C}$ .  $\text{header}^{WS1S}$  introduces the second-order variables  $V^x$  that correspond to the path variables  $x$ , but is otherwise identical to the QPTL version. Furthermore,  $\varphi_{\text{path}_i}^{WS1S}$  uses second-order equivalence  $V^s = V^c$  for  $\square(c \leftrightarrow v)$ . The consistency condition  $\text{consistent}^{WS1S}(\mathcal{S}, K)$  in WS1S is defined as the conjunction of

$$\bigwedge_{(\pi, \pi') \in \text{branches}(\text{deps}(s), K)^2} \left( V_\pi^s = V_{\pi'}^s \vee \exists i \geq 0. \left( \bigvee_{c \in \text{Scope}(s)} i \in V_\pi^c \leftrightarrow i \in V_{\pi'}^c \right) \wedge \forall j \leq i. j \in V_\pi^s \leftrightarrow j \in V_{\pi'}^s \right)$$

for every  $s \in \mathcal{S}$ . This formula ensures different reactions of a universal variable on two branches is based on different prior valuations of the dependencies (cf. QPTL consistency condition (4.3)).

We give an example query based on  $CL_\exists$  formula (2.2)  $\exists\{a\} \triangleright x. \exists\{b\} \triangleright y. \square(b = x) \rightarrow \square(\bigcirc y \leftrightarrow a)$ . The WS1S query is

$$\begin{aligned} &\exists A_1, A_2. \forall X_1, X_2. \forall Y_1, Y_2. \exists B_1, B_2. \\ &\left( X_1 = X_2 \vee (\exists i \geq 0. (i \in A_1 \leftrightarrow i \in A_2) \wedge \forall j \leq i. j \in X_1 \leftrightarrow j \in X_2) \right) \wedge \\ &\left( Y_1 = Y_2 \vee (\exists i \geq 0. (i \in B_1 \leftrightarrow i \in B_2) \wedge \forall j \leq i. j \in Y_1 \leftrightarrow j \in Y_2) \right) \rightarrow \\ &\bigwedge_{i \in \{1, 2\}} (B_i = X_i) \wedge \bigvee_{i \in \{1, 2\}} \exists k \geq 0. (S(k) \in Y_i \leftrightarrow k \in A_i). \end{aligned} \quad (6.2)$$

**Theorem 6.3** (Correctness  $\text{unsat}_{\text{fault}}^{WS1S}$ ). *Given a  $CL_\exists$  formula  $\Phi = \mathcal{Q}\exists. \bigwedge_{1 \leq i \leq n} (\varphi_{\text{path}_i} \rightarrow \varphi_i)$  over coordination variables  $\mathcal{C}$  and strategy variables  $\mathcal{S}$ . Let  $\varphi_i$  be a syntactically safe LTL formula for each  $1 \leq i \leq n$ .  $\Phi$  is unsatisfiable if there exists functions  $K_1 \dots K_n : \mathcal{C} \rightarrow \mathbb{N}$  such that the WS1S formula  $\text{unsat}_{\text{fault}}^{WS1S}(\Phi, K_1, \dots, K_n)$  is satisfiable.*

*Proof.* Assume  $\text{unsat}_{\text{fault}}^{WS1S}(\Phi, K_1, \dots, K_n)$  is satisfiable. We show that this implies the satisfiability of  $\text{unsat}_{\text{fault}}(\Phi, K_1, \dots, K_n)$  and use Theorem 4.1 that states the correctness of the QPTL encoding. Lemma 6.2 states that we only have to consider finite state sequences for the LTL specifications  $\varphi_i$ . What is missing in the argumentation are the paths specifications  $\varphi_{\text{path}_i}^{WS1S}$  and the consistency condition  $\text{consistent}^{WS1S}(\Phi, K)$ . The transformed path specifications are equisatisfiable under WS1S and QPTL semantics as the (existential quantified) input variables appear after the (universal quantified) strategy variables in  $\text{header}^{WS1S}(\mathcal{S}, K)$ . Lastly, the consistency condition restricts the finite behavior of the universal quantified strategy variables and the infinite behavior does not affect the satisfiability of the co-safety properties. Hence, there exists an infinite path extension for every

existentially quantified second-order variable, and the formula remains satisfiable against any infinite path extension of the universally quantified second-order variables.  $\square$

WS1S is supported by the Mona [HJJ<sup>+</sup>95] tool. Some of our smaller instances were solved by Mona, but the Byzantine Generals' Problem failed due to memory constraints in the BDD library.

**Bounded-length Counterexamples.** Taking the simplifications even further, we not only bound the *number* of paths but also the *length* of the paths by translating the problem to the satisfiability problem of *quantified Boolean formulas (QBF)*. Quantified Boolean formulas are the extension of Boolean formulas by quantification over variables. Let  $V = \{v_1, \dots, v_n\}$  be a finite set of variables. The syntax of QBF is given by the grammar

$$\varphi ::= x \mid \neg\varphi \mid \varphi \vee \psi \mid \exists x. \varphi ,$$

where  $x \in V$ . We define the usual abbreviations  $\wedge, \forall, \text{true}, \text{false}, \rightarrow,$  and  $\leftrightarrow$ . The semantics is defined over valuations  $\sigma : V \rightarrow \{0, 1\}$ . The satisfaction of an valuation  $\sigma$  is defined as

- $\sigma \models_{\text{QBF}} x \quad :\Leftrightarrow \sigma(x) = 1,$
- $\sigma \models_{\text{QBF}} \neg\varphi \quad :\Leftrightarrow \sigma \not\models_{\text{QBF}} \varphi,$
- $\sigma \models_{\text{QBF}} \varphi \vee \psi \quad :\Leftrightarrow \sigma \models_{\text{QBF}} \varphi \text{ or } \sigma \models_{\text{QBF}} \psi, \text{ and}$
- $\sigma \models_{\text{QBF}} \exists x. \varphi \quad :\Leftrightarrow \exists a \in \{0, 1\}. \sigma'(y) = \begin{cases} \sigma(y) & \text{if } y \neq x \\ a & \text{otherwise} \end{cases} \wedge \sigma' \models_{\text{QBF}} \varphi .$

The encoding translates a QPTL variable  $x$  to Boolean variables  $x_0, \dots, x_{k-1}$ , each representing one step in the system where  $k$  is the length of the paths. We build the QBF formula by *unrolling* the QPTL formula  $\text{unsat}_{\text{fault}}(\Phi, K_1, \dots, K_n)$  for  $k$ -steps: Each variable in the quantification prefix of the QPTL formula is transformed into  $k$  Boolean variables in the QBF prefix, e.g., the 3-unrolling of  $\exists x. \forall y. \varphi$  is  $\exists x_0, x_1, x_2. \forall y_0, y_1, y_2. \varphi_{\text{unroll}}$ . The unrolling of the remaining LTL formula is given by the expansion law for Until,  $\varphi \mathcal{U} \psi \equiv \psi \vee (\varphi \wedge \bigcirc \varphi \mathcal{U} \psi)$ . After the unrolling, the QBF formula is transformed into Conjunctive Normal Form (CNF) and encoded in the QDIMACS file format, that is the standard format for QBF solvers. Already with this encoding we could solve more examples than using the WS1S approach.

Formally, we define the transformation function  $qbf(\Phi, i, k)$  that takes a QPTL formula  $\Phi$  over alphabet  $\Sigma$  and transforms it into the  $k$ -unrolling:

- $qbf(p, i, k) \quad := p_i, \text{ for } p \in \Sigma,$
- $qbf(\neg p, i, k) \quad := \neg p_i, \text{ for } p \in \Sigma,$
- $qbf(\varphi_1 \circ \varphi_2, i, k) \quad := qbf(\varphi_1, i, k) \circ qbf(\varphi_2, i, k) \text{ for } \circ \in \{\vee, \wedge\},$
- $qbf(\bigcirc \varphi, i, k) \quad := \begin{cases} qbf(\varphi, i+1, k) & \text{if } i < k \\ \text{false} & \text{otherwise} \end{cases},$
- $qbf(\varphi \mathcal{U} \psi, i, k) \quad := \begin{cases} qbf(\psi, i, k) \vee (qbf(\varphi, i, k) \wedge qbf(\varphi \mathcal{U} \psi, i+1, k)) & \text{if } i < k \\ \text{false} & \text{otherwise} \end{cases},$
- $qbf(Q p. \varphi, i, k) \quad := Q p^0, \dots, p^{k-1}. qbf(\varphi, i, k) \text{ for } Q \in \{\exists, \forall\} .$

**Corollary 6.4.** *Given a QPTL formula  $\Phi = Q. \varphi$  where  $\varphi$  is a syntactically co-safe LTL formula. If  $qbf(\Phi, 0, k)$  is satisfiable for some  $k > 0$  then  $\Phi$  is satisfiable.*

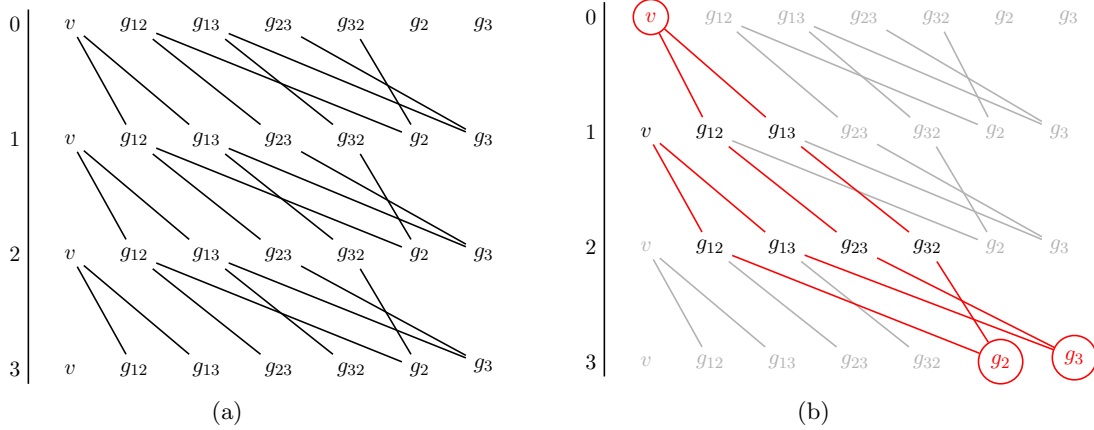


Figure 6: Example for a dependency graph of the Byzantines' Generals Problem. The graph identifies all variables that influence the variables used in the LTL specifications  $consensus_{i,j} := \bigcirc_3 (g_i = g_j)$  and  $correctval_i := v \leftrightarrow \bigcirc_3 g_i$  (variables with red circle). All variables that do not influence the LTL formula can be safely removed, e.g., the input  $v$  in step 2 and 3 never reaches the strategies  $g_2$  and  $g_3$ .

The QBF encoding of  $CL_{\exists}$  formula  $\Phi$  and functions  $K_1 \dots K_n : \mathcal{C} \rightarrow \mathbb{N}$  is

$$unsat_{fault}^{QBF}(\Phi, k, K_1, \dots, K_n) := header^{QBF}(\mathcal{S}, K, k). consistent^{QBF}(\mathcal{S}, K, k) \rightarrow \bigvee_{1 \leq i \leq n} \left( \bigwedge_{\pi \in branches(\mathcal{C}, K_i)} \varphi_{path_i}^{QBF}(\pi) \right) \wedge \left( \bigvee_{\pi \in branches(\mathcal{C}, K_i)} qbf(\neg \varphi_i(\pi), 0, k) \right), \quad (6.3)$$

where  $K : \mathcal{C} \rightarrow \mathbb{N}$  is defined as  $K(c) := \max_{1 \leq i \leq n} K_i(c)$  for every  $c \in \mathcal{C}$  and  $k > 0$ .  $header^{QBF}$  introduces the variables  $x^0, \dots, x^{k-1}$  that correspond to the path variables  $x$ , but is otherwise identical to the QPTL version. Furthermore,  $\varphi_{path}^{QBF}$  uses the equivalence  $\bigwedge_{i < k} (s^i \leftrightarrow c^i)$  for  $\square(c \leftrightarrow v)$ . The consistency condition  $consistent^{QBF}(\mathcal{S}, K, k)$  in QBF is defined as the conjunction of

$$\bigwedge_{(\pi, \pi') \in branches(deps(s), K)^2} \left( \bigwedge_{i < k} (s_{\pi}^i \leftrightarrow s_{\pi'}^i) \vee \bigvee_{j < i} \bigvee_{c \in Scope(s)} c_{\pi}^j \leftrightarrow c_{\pi'}^j \right)$$

for every  $s \in \mathcal{S}$ . This formula ensures different reactions of a strategy variable on two branches is based on different prior valuations to the dependencies (cf. QPTL consistency condition (4.3)). In this simple translation, one cause of high complexity is due to the consistency conditions between the strategy variables across different paths. However, most of these variables are not used for the counterexample itself but appear only in the consistency condition. One optimization removes these unnecessary variables from the encoding. Therefore, we collect all strategy variables and (when possible) their temporal occurrence from the LTL specification. For every used strategy variable we build the *dependency graph* that contains all variables which can influence the outcome of the strategy. In the last step, we remove all variables that are not contained in any dependency graph. This optimization is depicted in Fig. 6.

Let us consider again the pipeline example given by the  $CL_{\exists}$  formula (2.2)  $\exists\{a\} \triangleright x. \exists\{b\} \triangleright y. \Box(b = x) \rightarrow \Box(\bigcirc y \leftrightarrow a)$ . The QBF encoding with unrolling depth of 2 is

$$\begin{aligned} & \exists a_1^0, a_1^1, a_2^0, a_2^1. \forall x_1^0, x_1^1, x_2^0, x_2^1. \forall y_1^0, y_1^1, y_2^0, y_2^1. \exists b_1^0, b_1^1, b_2^0, b_2^1. \\ & (x_1^0 \leftrightarrow x_2^0) \wedge (x_1^1 \leftrightarrow x_2^1 \vee a_1^0 \leftrightarrow a_2^0) \wedge \\ & (y_1^0 \leftrightarrow y_2^0) \wedge (y_1^1 \leftrightarrow y_2^1 \vee b_1^0 \leftrightarrow b_2^0) \rightarrow \\ & \bigwedge_{i \in \{1,2\}} ((b_i^0 \leftrightarrow x_i^0) \wedge (b_i^1 \leftrightarrow x_i^1)) \wedge \bigvee_{i \in \{1,2\}} (y_i^1 \leftrightarrow a_i^0). \end{aligned} \quad (6.4)$$

**Theorem 6.5** (Correctness  $unsat_{fault}^{QBF}$ ). *Given a  $CL_{\exists}$  formula  $\Phi = \mathcal{Q}_{\exists}. \bigwedge_{1 \leq i \leq n} (\varphi_{path_i} \rightarrow \varphi_i)$  over coordination variables  $\mathcal{C}$  and strategy variables  $\mathcal{S}$ . Let  $\varphi_i$  be a syntactically safe LTL formula for each  $1 \leq i \leq n$ .  $\Phi$  is unsatisfiable if there exists functions  $K_1 \dots K_n : \mathcal{C} \rightarrow \mathbb{N}$  such that the QBF query  $unsat_{fault}^{QBF}(\Phi, k, K_1, \dots, K_n)$  is satisfiable for some  $k > 0$ .*

*Proof.* The encoding is a special case of Theorem 6.3. If the Boolean formula that is a finite  $k$ -unrolling of the LTL formula is satisfied against a finite sequence of length  $k$ , i.e., we found a counterexample within the first  $k$  steps, then it is also satisfied against any (possibly infinite) strategy.  $\square$

## 7. EXPERIMENTAL RESULTS

We have carried out our experiments on a 2.6 GHz Opteron system. For solving the QBF instances, we used a combination of the QBF preprocessor Bloqqer [BLS11] in version 031 and the QBF solver DepQBF [LB10] in version 3.0.3. For solving the WS1S instances, we used Mona [HJJ<sup>+</sup>95] in version 1.4-15.

**Byzantine Generals' Problem.** Table 1 demonstrates that the Byzantine Generals' Problem remains, despite the optimizations described above, a nontrivial combinatorial problem: we need to find a suitable set of paths for every possible combination of the strategies of the generals. The bound given in the first row reads as follows: The first component is the number of branchings for the input variable  $v$  in all three architectures. The last three components state the number of branchings for the outputs of the faulty nodes in their respective architectures. For example, bound (1, 1, 0, 0) means that we have two branches for  $v$ ,  $c_{12}$ , and  $c_{13}$ , while we have only one branch for  $c_{23}$  and  $c_{32}$ . More precisely, starting from constant zero functions  $K_1, K_2, K_3$ , the bound (1, 1, 0, 0) sets

Table 1: Result of the *Byzantine Generals' Problem* example

Bound	(0, 0, 0, 0)	(1, 0, 0, 0)	(1, 1, 0, 0)	(1, 1, 1, 0)	(1, 1, 1, 1)
Result	Unsatisfiable	Unsatisfiable	Unsatisfiable	Unsatisfiable	Satisfiable
# Clauses	57	228	2286	2904	3522
# Variables	44	143	1095	1375	1655
Time (s)	0.09	0.16	0.98	1.01	23.61

The table shows the solving time of the Byzantine Generals' Problem  $\Phi_{\text{bgp}}$  (Equation 4.8) using the QBF encoding with a fixed length of 3 unrollings. The QBF queries are solved using a combination of Bloqqer 031 and DepQBF 3.0.3.



$K_1(v) = K_2(v) = K_3(v) = K_1(c_{12}) = K_1(c_{13}) = 1$  and  $K_2(c_{23}) = K_3(c_{32}) = 0$ . To prove the unrealizability, we need one branching for the input  $v$  and one branching for every coordination variable that serves as a shared variable for a faulty node, i.e., the bound  $(1, 1, 1, 1)$ . The number of branches and thereby the formula size grows exponentially with the number of branchings for the input variables.

Table 2: Result of the *Byzantine Firing Squad Problem* example

Instance	Result	# Clauses	# Variables	Time (s)
bfsp_3	Satisfiable	803	435	0.21
bfsp_5	Satisfiable	1773	967	0.43
bfsp_10	Satisfiable	7348	3942	3.07
bfsp_20	Satisfiable	42 198	22 042	15.58
bfsp_30	Satisfiable	128 648	66 342	66.29
bfsp_40	Satisfiable	290 698	148 842	190.90
bfsp_50	Satisfiable	552 348	281 542	396.20

The table shows the solving time of the Byzantine Firing Squad Problem  $\Phi_{\text{bfsp}}$  (Equation 7.1) using the QBF encoding with a fixed length of 3 unrollings. The QBF queries are solved using a combination of Bloqqer 031 and DepQBF 3.0.3.

**Byzantine Firing Squad Problem.** We consider another classical consensus problem, the *Byzantine Firing Squad Problem* [FLM85]. In this setting, we have a set of synchronous processes, each having a distinct input and a broadcasting mechanism for communication with the other processes, depicted in Fig. 7(a). The goal of these processes is to synchronously enter a *fire* state whenever the input was initially given at one process. As before, we consider architectures with up to one faulty process, but in contrast to the Byzantine Generals' Problem, we require that this specification holds only if all processes are correct. Furthermore, we require that the non-faulty processes act *uniformly*, i.e., if a non-faulty process enters the firing state, all other non-faulty processes enter the firing state. It was shown in [FLM85] that there is no protocol that ensures correct behavior in the presence of faults. The encoding  $\Phi_{\text{bfsp}}$  for three nodes is given by the  $\text{CL}_{\exists}$  formula

$$\begin{aligned}
& \exists\{\text{req}_1\} \triangleright \text{bcast}_1. \exists\{\text{req}_1, \text{chan}_{21}, \text{chan}_{31}\} \triangleright \text{out}_1. \\
& \exists\{\text{req}_2\} \triangleright \text{bcast}_2. \exists\{\text{req}_2, \text{chan}_{12}, \text{chan}_{32}\} \triangleright \text{out}_2. \\
& \exists\{\text{req}_3\} \triangleright \text{bcast}_3. \exists\{\text{req}_3, \text{chan}_{13}, \text{chan}_{23}\} \triangleright \text{out}_3. \\
& (\Box \text{operational}_{2,3} \rightarrow \text{consistent}_{2,3}) \wedge \\
& (\Box \text{operational}_{1,3} \rightarrow \text{consistent}_{1,3}) \wedge \\
& (\Box \text{operational}_{1,2} \rightarrow \text{consistent}_{1,2}) \wedge \\
& (\Box \text{operational}_{1,2,3} \rightarrow \text{consistent}_{1,2,3} \wedge \text{uniform}_{1,2,3}) \quad , \tag{7.1}
\end{aligned}$$

where

$$\begin{aligned} \text{consistent}_N &:= \Box \bigwedge_{i < j \in N^2} (\text{out}_i = \text{out}_j) \text{ and} \\ \text{uniform}_N &:= \left( \bigvee_{i \in N} \text{req}_i \right) \rightarrow \text{O}_3 \left( \bigvee_{i \in N} \text{out}_i \right) . \end{aligned} \quad (7.2)$$

Table 2 shows that our method can detect conflicts quickly even though the  $\text{CL}_{\exists}$  encoding  $\Phi_{\text{bfsp}}$  grows quadratically in the number of nodes.

**CAP Theorem.** The CAP Theorem due to Brewer [Bre00] states that it is impossible to design a distributed system that provides Consistency, Availability, and Partition tolerance (CAP) simultaneously. For the encoding in  $\text{CL}_{\exists}$ , we assume there is a fixed number  $n$  of nodes, that every node implements the same service, and that there are direct communication links between all nodes, depicted in Fig. 7(b). We use the variables  $\text{req}_i$  and  $\text{out}_i$  to denote input and output of node  $i$ , respectively. The consistency and availability requirements are encoded as the LTL formulas  $\bigwedge_{1 \leq i < n} (\text{out}_i \leftrightarrow \text{out}_{i+1})$  and  $(\bigvee_{1 \leq i \leq n} \text{req}_i) \leftrightarrow (\text{O}_2 \bigvee_{1 \leq i \leq n} \text{out}_i)$ . The partition tolerance is modeled in a way that there is always at most one node partitioned from the rest of the system, i.e., we have  $n$  different architectures and in every architecture all communication links to one node are faulty. For two nodes, we get the  $\text{CL}_{\exists}$  formula  $\Phi_{\text{cap}}$

$$\begin{aligned} &\exists\{\text{req}_1\} \triangleright \text{com}_1. \exists\{\text{req}_1, \text{chan}_2\} \triangleright \text{out}_1. \exists\{\text{req}_2\} \triangleright \text{com}_2. \exists\{\text{req}_2, \text{chan}_1\} \triangleright \text{out}_2. \\ &(\Box(\text{chan}_1 = \text{com}_1) \rightarrow \Box((\text{out}_1 = \text{out}_2) \wedge ((\text{req}_1 \vee \text{req}_2) \leftrightarrow \text{O}_3(\text{out}_1 \vee \text{out}_2)))) \wedge \\ &(\Box(\text{chan}_2 = \text{com}_2) \rightarrow \Box((\text{out}_1 = \text{out}_2) \wedge ((\text{req}_1 \vee \text{req}_2) \leftrightarrow \text{O}_3(\text{out}_1 \vee \text{out}_2)))) . \end{aligned} \quad (7.3)$$

Table 3 shows that our method is able to find conflicts in a specification with an architecture up to 50 nodes within reasonable time. When we drop either Consistency, Availability, or Partition tolerance, the corresponding instances (AP, CP, and CA) become satisfiable. Hence, our tool does not find counterexamples in these cases.

Table 3: Result of the *CAP Theorem* example

Instance	Result	# Clauses	# Variables	Time (s)
ap_2	Unsatisfiable	1232	619	0.41
ca_2	Unsatisfiable	1408	763	1.11
cp_2	Unsatisfiable	48	42	0.09
cap_2	Satisfiable	110	84	0.11
cap_5	Satisfiable	665	426	0.28
cap_10	Satisfiable	2590	1556	1.07
cap_25	Satisfiable	15 865	9146	20.60
cap_50	Satisfiable	62 990	35 796	781.16

The table shows the solving time of the CAP Theorem  $\Phi_{\text{cap}}$  (Equation 7.3) using the QBF encoding with a fixed length of 3 unrollings. The QBF queries are solved using a combination of Bloqqer 031 and DepQBF 3.0.3.




Figure 7: Broadcasting (a) and non-broadcasting (b) communication topologies. (We omit the environment process to increase readability.)

**Discussion.** Table 4 shows a comparison of the the different encodings that we have presented in the paper. There does not exist an algorithm that decides whether a given  $\text{CL}_{\exists}$  formula is unsatisfiable. We presented an approach that bounds the number of paths and used an encoding to QPTL. The reason for incompleteness was shown in Section 5: for liveness properties one may need infinitely many paths to show unsatisfiability. Our encoding in WS1S (Mona) loses the ability to find counterexample paths of infinite length, e.g., the  $\text{CL}_{\exists}$  formula

$$\exists \emptyset \triangleright y. \diamond \square (\bigcirc y \leftrightarrow x) \quad (7.4)$$

with free coordination variable  $x$  is unsatisfiable and two paths that are infinitely often different are sufficient to prove it. The QPTL encoding is capable of finding these paths while neither the WS1S, nor the QBF, encoding is applicable. However, Mona failed to solve larger instances from Tables 1 to 3, especially it could not handle the encoding of the Byzantine Generals’ Problem. For the translation to QBF we bounded the length of the paths to  $k$  where  $k$  is an additional parameter. With this encoding we approximate the reactive behavior of our system by a finite prefix. Despite this restriction, we could prove unsatisfiability for many interesting specifications from literature. In some instances, the WS1S approximation provides stronger unrealizability guarantees. For example in the Byzantine Firing Squad Problem, the WS1S encoding is able to prove that any *finite* uniform reaction

$$\text{uniform}_N := \left( \bigvee_{i \in N} \text{req}_i \right) \rightarrow \diamond \left( \bigvee_{i \in N} \text{out}_i \right) \quad (7.5)$$

leads to an unrealizable instance (cf. Table 4), while in the QBF encoding, we have to bound the uniform reaction (7.2) to a fixed length. In practice, one would first use the QBF approximation in order to find “cheap” counterexamples. After hitting the number of paths that the QBF solver can no longer handle within reasonable time, one proceeds with more costly approximations like the WS1S encoding.

Table 4: Comparison of different encodings

	independent (2.1)	pipeline (2.2)	join (2.3)	bgp (4.8)	bfsp (7.1)	cap (7.3)	infinite (7.4)	bfsp-finite (7.1),(7.5)
QPTL	✓	(✓)	(✓)	(✓)	(✓)	(✓)	✓	(✓)
WS1S	✓	✓	✓	(✓)	✓(5)	✓(6)	-	✓(5)
QBF	✓	✓	✓	✓	✓	✓	-	-

The table compares the expressiveness of the encodings and the experimental verifiability on the examples presented in this paper. The (✓) symbol denotes that the tool failed to solve the query, the number behind the checkmark indicates up to how many nodes the query was successful. We used GOAL, Mona 1.4-15, and a combination of Bloqer 031 and DepQBF 3.0.3 for the QPTL, WS1S, and QBF queries, respectively.

## 8. BEYOND THE BYZANTINE FAULT MODEL

We formalize the fault-tolerant synthesis problem regarding multiple types of faults, their durations, observability, and the fault-tolerance requirements and provide a uniform encoding into  $CL_{\exists}$ . Faults are modeled by allowing the environment to take over the control of the faulty processes.

**Types of Faults.** We review types of faults that were considered in previous work [DS86, AAE04, DF09]. *Stuck-at faults* occur when a process stops reacting on new inputs and keep their last outputs. On *fail-stop faults* or *crashes* the process halts immediately, i.e., they stop producing any outputs. While *fail-stop faults* are detectable, *crashes* are not detectable [DS86]. *Omission faults* subsume these faults and are characterized as failures on the specified input-output specification. The most general fault type are the *Byzantine faults*, where a process can deviate arbitrarily from its specification, even in a malicious way.

The duration of a fault can be *permanent*, *intermittent*, or *transient*. Some failures may be *detectable* or *undetectable* depending on the architecture and the type of fault.

Lastly, the system usually has to fulfill different kind of specifications in the presence of faults and we use the classification of fault-tolerance requirements from [DF09]: *Masking tolerance* means that the system has to fulfill both safety and liveness specifications, *non-masking tolerance* allows the system to temporarily violate its safety properties while *fail-safe tolerance* specifies that after a fault, only the safety requirements have to be fulfilled.

**Fault-tolerance Specifications.** For simplicity, we restrict the failures to processes, but it can be easily adapted to failures of individual communication links. Consider an arbitrary architecture  $\mathcal{A} = (P, p_{env}, \{I_p\}_{p \in P}, \{O_p\}_{p \in P})$ . A *fault-tolerance scenario*  $F$  is a tuple  $(P_f, T, D, Obs, \varphi)$  where  $P_f \subseteq P$  is a subset of processes that simultaneously suffer from a fault of type  $T \in \{\textit{stuck-at}, \textit{fail-stop}, \textit{omission}, \textit{Byzantine}\}$  and duration  $D \in \{\textit{permanent}, \textit{intermittent}, \textit{transient}\}$ ,  $Obs \in \{\top, \perp\}$  denotes the observability, and  $\varphi$  is an LTL formula specifying the fault-tolerance requirements. Not all combinations of fault types and durations are possible, e.g., stuck-at faults and fail-stop faults are always permanent. A *fault-tolerance specification*  $\mathcal{F}$  is a set of fault-tolerance scenarios.

**Fault-tolerant Realizability.** The distributed fault-tolerant realizability problem is to decide, given an architecture  $\mathcal{A}$  and a fault-tolerance specification  $\mathcal{F}$ , whether there exists a finite state implementation for the processes in  $\mathcal{A}$  such that the system satisfies the fault-tolerance specification  $\mathcal{F}$ . We formalize this problem by giving a  $\text{CL}_{\exists}$  formula  $\Phi = \mathcal{Q}_{\exists} \cdot \bigwedge_{F \in \mathcal{F}} (\varphi_{\text{path}_F} \rightarrow \varphi_F)$  that models the description of the fault-tolerance scenarios given above.

Consider an arbitrary architecture  $\mathcal{A} = (P, p_{\text{env}}, \{I_p\}_{p \in P}, \{O_p\}_{p \in P})$  and a fault-tolerance specification  $\mathcal{F}$ . Using Theorem 2.1 we get the  $\text{CL}_{\exists}$  encoding  $\Psi = \mathcal{Q}_{\exists} \cdot \varphi_{\text{path}} \rightarrow \varphi$  of the distributed realizability problem for architecture  $\mathcal{A}$  and an arbitrary LTL formula  $\varphi$ . For every fault-tolerance scenario  $F \in \mathcal{F}$ , we introduce a new architecture that modifies the input-output specifications for processes affected by  $F$  according to the scenario. We denote the resulting path specification by the LTL formula  $\varphi_{\text{path}_F}$  and the fault-tolerance requirement for  $F$  as  $\varphi_F$ . The resulting  $\text{CL}_{\exists}$  formula contains a conjunct over all scenarios  $\mathcal{F}$ :

$$\Phi = \mathcal{Q}_{\exists} \cdot \bigwedge_{F \in \mathcal{F}} (\varphi_{\text{path}_F} \rightarrow \varphi_F)$$

We now define the path specifications  $\varphi_{\text{path}_F}$ . Let  $F = (P_f, T, D, \text{Obs}, \varphi_F)$  be a fault-tolerance scenario. For every process  $p \in P_f$  we introduce a fresh coordination variable  $\text{fault}_p^F$  that immediately signals whether a fault occurred at process  $p$  in the last step. If the faults are observable ( $\text{Obs} = \top$ ) we extend the scope of every strategy variable by  $\{\text{fault}_p^F \mid p \in P_f\}$ <sup>2</sup>.

Let  $\varphi_{\text{path}_F} = \varphi_{\text{path}} = \bigwedge_{v \in I \cap O} \square(c_v = s_v)$ . We remove every conjunct  $\square(c_v = s_v)$  from  $\varphi_{\text{path}_F}$  where  $v \in O_{P_f}$  is the output of a faulty process. For every  $v \in I \cap O_{P_f}$  we add the following new constraints to  $\varphi_{\text{path}_F}$  depending on the type of faults  $T$ . For readability, we replace  $\text{fault}_p^F$  by  $\text{fault}$ ,  $c_v$  by  $i$ , and  $s_v$  by  $o$ .

- $T = \text{stuck-at}$ :  $(\neg \text{fault} \wedge (i = o)) \mathcal{W} (\square \text{fault} \wedge i \leftrightarrow \bigcirc \square i)$ .
- $T = \text{fail-stop}$ :  $(\neg \text{fault} \wedge (i = o)) \mathcal{W} (\square (\text{fault} \wedge \neg i))$ .
- $T = \text{omission}$ :  $\square (\text{fault} \wedge \neg i \vee \neg \text{fault} \wedge (i = o))$ .
- $T = \text{Byzantine}$ :  $\square (\text{fault} \vee \neg \text{fault} \wedge (i = o))$ .

Lastly, we add a requirement for every  $\text{fault}_p^F$  variable regarding the specified duration  $D$ :

- $D = \text{permanent}$ :  $\neg \text{fault} \mathcal{W} (\square \text{fault})$
- $D = \text{intermittent}$ :  $\square \diamond \neg \text{fault}$ . If we want to have tighter control about the duration of faults we can use the parametric LTL operators [AELP01] to bound the length  $n$  of a fault  $\square \diamond_{\leq n} \neg \text{fault}$ , or require that length of a correct behavior is longer than  $m$  by formula  $\square (\text{fault} \rightarrow \text{fault} \mathcal{U} (\neg \text{fault} \mathcal{U}_{> m} \text{fault}))$ .
- $D = \text{transient}$ :  $\neg \text{fault} \mathcal{W} (\text{fault} \mathcal{U}_{\leq n} \square \neg \text{fault})$  where  $n$  is the maximal duration of the fault.

We simplify the undetectable permanent fault encoding by removing the  $\text{fault}$  variables, as these variables are not used in the specification. Our formulation allows for many types and durations of faults that cannot be captured with a fault-tolerance scenario. Indeed, it is possible to use arbitrary LTL formulas to specify the type and duration of a fault. Computing counterexamples for these general types of faults can be integrated into the method described in Section 4.

<sup>2</sup>These signals are called *fault-notification variables* in [DF09] as only detectable faults were considered.

## 9. CONCLUSION

We have introduced counterexamples for distributed realizability and shown how to automatically derive counterexamples from given specifications in  $CL_{\exists}$ . We used encodings in QPTL, WS1S, and QBF. In our experiments, the QBF encoding was the most efficient. Even problems with high combinatorial complexity, such as the Byzantine Generals' Problem, are handled automatically. Given that QBF solvers are likely to continue to improve in the future, even larger instances should become tractable. In future work, we plan to extend the method to a larger class of *infinite* counterexamples, which will support liveness specifications. Furthermore, we want to investigate approximation techniques for more general types of faults.

## ACKNOWLEDGMENTS

This work was partially supported by the German Research Foundation (DFG) as part of the Transregional Collaborative Research Center "Automatic Verification and Analysis of Complex Systems" (SFB/TR 14 AVACS). We thank Swen Jacobs for comments on an earlier version of this paper.

## REFERENCES

- [AAE04] Paul C. Attie, Anish Arora, and E. Allen Emerson. Synthesis of fault-tolerant concurrent programs. *ACM Trans. Program. Lang. Syst.*, 26(1):125–185, 2004.
- [AELP01] Rajeev Alur, Kousha Etessami, Salvatore La Torre, and Doron Peled. Parametric temporal logic for "model measuring". *ACM Trans. Comput. Log.*, 2(3):388–407, 2001.
- [ALW89] Martín Abadi, Leslie Lamport, and Pierre Wolper. Realizable and unrealizable specifications of reactive systems. In *ICALP*, volume 372 of *LNCS*, pages 1–17. Springer, 1989.
- [BLS11] Armin Biere, Florian Lonsing, and Martina Seidl. Blocked clause elimination for QBF. In *CADE*, volume 6803 of *LNCS*, pages 101–115. Springer, 2011.
- [Bre00] Eric A. Brewer. Towards robust distributed systems (abstract). In *PODC*, page 7. ACM, 2000.
- [CHJ08] Krishnendu Chatterjee, Thomas A. Henzinger, and Barbara Jobstmann. Environment assumptions for synthesis. In *CONCUR*, volume 5201 of *LNCS*, pages 147–161. Springer, 2008.
- [Chu63] Alonzo Church. Logic, arithmetic and automata. In *Proc. 1962 Intl. Congr. Math.*, pages 23–25, Upsala, 1963.
- [DF09] Rayna Dimitrova and Bernd Finkbeiner. Synthesis of fault-tolerant distributed systems. In *ATVA*, volume 5799 of *LNCS*, pages 321–336. Springer, 2009.
- [DS86] Danny Dolev and H. Raymond Strong. A simple model for agreement in distributed systems. In *Fault-Tolerant Distributed Computing*, volume 448 of *LNCS*, pages 42–50. Springer, 1986.
- [FLM85] Michael J. Fischer, Nancy A. Lynch, and Michael Merritt. Easy impossibility proofs for distributed consensus problems. In *PODC*, pages 59–70. ACM, 1985.
- [FS05] Bernd Finkbeiner and Sven Schewe. Uniform distributed synthesis. In *LICS*, pages 321–330. IEEE Computer Society, 2005.
- [FS10] Bernd Finkbeiner and Sven Schewe. Coordination logic. In *CSL*, volume 6247 of *LNCS*, pages 305–319. Springer, 2010.
- [FS13] Bernd Finkbeiner and Sven Schewe. Bounded synthesis. *STTT*, 15(5-6):519–539, 2013.
- [FT14] Bernd Finkbeiner and Leander Tentrup. Detecting unrealizable specifications of distributed systems. In *TACAS*, volume 8413 of *LNCS*, pages 78–92. Springer, 2014.
- [HJJ<sup>+</sup>95] Jesper G. Henriksen, Jakob L. Jensen, Michael E. Jørgensen, Nils Klarlund, Robert Paige, Theis Rauhe, and Anders Sandholm. Mona: Monadic second-order logic in practice. In *TACAS*, volume 1019 of *LNCS*, pages 89–110. Springer, 1995.
- [HM84] Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. In *PODC*, pages 50–61. ACM, 1984.

- [KP95] Yonit Kesten and Amir Pnueli. A complete proof systems for QPTL. In *LICS*, pages 2–12, 1995.
- [KV97] Orna Kupferman and Moshe Y. Vardi. Synthesis with incomplete information. In *ICTL*, 1997.
- [KV01a] Orna Kupferman and Moshe Y. Vardi. Model checking of safety properties. *Formal Methods in System Design*, 19(3):291–314, 2001.
- [KV01b] Orna Kupferman and Moshe Y. Vardi. Synthesizing distributed systems. In *LICS*, pages 389–398. IEEE Computer Society, 2001.
- [LB10] Florian Lonsing and Armin Biere. DepQBF: A dependency-aware QBF solver. *JSAT*, 7(2-3):71–76, 2010.
- [LDS11] Wenchao Li, Lili Dworkin, and Sanjit A. Seshia. Mining assumptions for synthesis. In *MEM-OCODE*, pages 43–50. IEEE, 2011.
- [LSP82] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
- [Pnu77] Amir Pnueli. The temporal logic of programs. In *FOCS*, pages 46–57, 1977.
- [PR90] Amir Pnueli and Roni Rosner. Distributed reactive systems are hard to synthesize. In *FOCS*, pages 746–757, 1990.
- [RKG11] Vasumathi Raman and Hadas Kress-Gazit. Analyzing unsynthesizable specifications for high-level robot behavior using LTLMoP. In *CAV*, volume 6806 of *LNCS*, pages 663–668. Springer, 2011.
- [Sch12] Viktor Schuppan. Towards a notion of unsatisfiable and unrealizable cores for LTL. *Sci. Comput. Program.*, 77(7-8):908–939, 2012.
- [Sch14] Sven Schewe. Distributed synthesis is simply undecidable. *Inf. Process. Lett.*, 114(4):203–207, 2014.
- [Sis94] A. Prasad Sistla. Safety, liveness and fairness in temporal logic. *Formal Asp. Comput.*, 6(5):495–512, 1994.
- [SVW85] A. Prasad Sistla, Moshe Y. Vardi, and Pierre Wolper. The complementation problem for büchi automata with applications to temporal logic (extended abstract). In *ICALP*, pages 465–474, 1985.