

Efficient First-Order Temporal Logic for Infinite-State Systems

Clare Dixon Michael Fisher Boris Konev Alexei Lisitsa

Department of Computer Science, University of Liverpool
Liverpool L69 3BX, United Kingdom

{C.Dixon, M.Fisher, B.Konev, A.Lisitsa}@csc.liv.ac.uk

Abstract


In this paper we consider the specification and verification of infinite-state systems using temporal logic. In particular, we describe parameterised systems using a new variety of first-order temporal logic that is both powerful enough for this form of specification and tractable enough for practical deductive verification. Importantly, the power of the temporal language allows us to describe (and verify) asynchronous systems, communication delays and more complex properties such as liveness and fairness properties. These aspects appear difficult for many other approaches to infinite-state verification.

1 Introduction

First-order temporal logic (FOTL) has been shown to be a powerful formalism for expressing sophisticated dynamic properties. Unfortunately, this power also leads to strong intractability. Recently, however, a fragment of FOTL, called *monodic* FOTL, has been investigated, both in terms of its theoretical [24, 22] and practical [7, 26, 25] properties. Essentially, monodicity allows for one free variable in every temporal formula. Although clearly restrictive, this fragment has been shown to be useful in expressive description logics, infinite-state verification, and spatio-temporal logics [3, 31, 27, 20, 19].

We here develop a new temporal logic, combining decidable fragments of monodic FOTL [24] with recent developments in XOR temporal logics [14], and apply this to the verification of parameterised systems. We use a communicating finite state machine model of computation, and can specify not only basic synchronous, parameterised systems with instantaneous broadcast communication [17], but the powerful temporal language allows us also to specify asynchronously executing machines and more sophisticated communication properties, such as delayed delivery of messages. In addition, and in contrast to many other approaches [29, 10, 2], not only safety, but also liveness and fairness properties, can be verified through automatic deductive verification. Finally, in contrast to work on regular model checking [1] and constraint based verification using counting abstraction [17], the logical approach is both complete and decidable.

The verification of concurrent systems often comes down to the analysis of multiple finite-state automata, for example of the following form.



In describing such automata, both automata-theoretic and logical approaches may be used. While *temporal logic* [16] provides a clear, concise and intuitive description of the system, automata-theoretic techniques such as *model checking* [6] have been shown to be more useful in practice. Recently, however, a propositional, linear-time temporal logic with improved deductive properties has been introduced [13, 14], providing the possibility of practical deductive verification in the future. The essence of this approach is to provide an XOR constraint between key propositions. These constraints state that exactly one proposition from a XOR set can be true at any moment in time. Thus, the automaton above can be described by the following clauses which are implicitly in the scope of a ‘ \square ’ (‘always in the future’) operator.

1. **start** $\Rightarrow s_t$
2. $s_t \Rightarrow \bigcirc(s_t \vee s_a)$
3. $s_b \Rightarrow \bigcirc s_t$
4. $s_a \Rightarrow \bigcirc s_w$
5. $s_w \Rightarrow \bigcirc(s_w \vee s_b)$

Here ‘ \bigcirc ’ is a temporal operator denoting ‘at the next moment’ and ‘**start**’ is a temporal operator which holds only at the initial moment in time. The inherent assumption that at any moment in time exactly one of s_a, s_b, s_t or s_w holds, is denoted by the following.

$$\square(s_a \oplus s_b \oplus s_t \oplus s_w)$$

With the complexity of the decision problem (regarding s_a, s_b , etc) being polynomial, then the properties of any finite collection of such automata can be tractably verified using this *propositional XOR* temporal logic.

However, one might argue that this deductive approach, although elegant and concise, is still no better than a model checking approach, since it targets just *finite* collections of (finite) state machines. Thus, this naturally leads to the question of whether the XOR temporal approach can be extended to *first-order temporal logics* and, if so, whether a form of tractability still applies. In such an approach, we can consider *infinite* numbers of finite-state automata (initially, all of the same structure). Previously, we have shown that FOTL can be used to elegantly specify such a system, simply by assuming the argument to each predicate represents a particular automaton [19]. Thus, in the following $s_a(X)$ is true if automaton X is in state s_a :

1. **start** $\Rightarrow \exists x. s_t(x)$
2. $\forall x. (s_t(x) \Rightarrow \bigcirc(s_t(x) \vee s_a(x)))$
3. $\forall x. (s_b(x) \Rightarrow \bigcirc s_t(x))$
4. $\forall x. (s_a(x) \Rightarrow \bigcirc s_w(x))$
5. $\forall x. (s_w(x) \Rightarrow \bigcirc(s_w(x) \vee s_b(x)))$

Thus, FOTL can be used to specify and verify broadcast protocols between synchronous components [17]. In this paper we define a logic, FOTLX, which allows us to not only to specify and verify systems of the above form, but also to specify and verify more sophisticated asynchronous systems, and to carry out verification with a reasonable complexity.

2 FOTLX

2.1 First-Order Temporal Logic

First-Order (discrete, linear time) Temporal Logic, FOTL, is an extension of classical first-order logic with operators that deal with a discrete and linear model of time (isomorphic to the Natural Numbers, \mathbb{N}).

Syntax. The symbols used in FOTL are

- *Predicate symbols:* P_0, P_1, \dots each of which is of a fixed arity (null-ary predicate symbols are *propositions*);
- *Variables:* x_0, x_1, \dots ;
- *Constants:* c_0, c_1, \dots ;
- *Boolean operators:* $\wedge, \neg, \vee, \Rightarrow, \equiv$, **true** ('true'), **false** ('false');
- *First-order Quantifiers:* \forall ('for all') and \exists ('there exists'); and
- *Temporal operators:* \square ('always in the future'), \diamond ('sometime in the future'), \bigcirc ('at the next moment'), U (until), W (weak until), and **start** (at the first moment in time).

Although the language contains constants, neither equality nor function symbols are allowed.

The set of well-formed FOTL-formulae is defined in the standard way [24, 7]:

- Booleans **true** and **false** are atomic FOTL-formulae;
- if P is an n -ary predicate symbol and $t_i, 1 \leq i \leq n$, are variables or constants, then $P(t_1, \dots, t_n)$ is an atomic FOTL-formula;
- if ϕ and ψ are FOTL-formulae, so are $\neg\phi, \phi \wedge \psi, \phi \vee \psi, \phi \Rightarrow \psi$, and $\phi \equiv \psi$;
- if ϕ is an FOTL-formula and x is a variable, then $\forall x\phi$ and $\exists x\phi$ are FOTL-formulae;
- if ϕ and ψ are FOTL-formulae, then so are $\square\phi, \diamond\phi, \bigcirc\phi, \phi U \psi, \phi W \psi$, and **start**.

A *literal* is an atomic FOTL-formula or its negation.

$\mathfrak{M}_n \models^a \mathbf{true}$		$\mathfrak{M}_n \not\models^a \mathbf{false}$
$\mathfrak{M}_n \models^a \mathbf{start}$	iff	$n = 0$
$\mathfrak{M}_n \models^a P(t_1, \dots, t_m)$	iff	$\langle I_n^a(t_1), \dots, I_n^a(t_m) \rangle \in I_n(P)$, where $I_n^a(t_i) = I_n(t_i)$, if t_i is a constant, and $I_n^a(t_i) = \mathbf{a}(t_i)$, if t_i is a variable
$\mathfrak{M}_n \models^a \neg\phi$	iff	$\mathfrak{M}_n \not\models^a \phi$
$\mathfrak{M}_n \models^a \phi \wedge \psi$	iff	$\mathfrak{M}_n \models^a \phi$ and $\mathfrak{M}_n \models^a \psi$
$\mathfrak{M}_n \models^a \phi \vee \psi$	iff	$\mathfrak{M}_n \models^a \phi$ or $\mathfrak{M}_n \models^a \psi$
$\mathfrak{M}_n \models^a \phi \Rightarrow \psi$	iff	$\mathfrak{M}_n \models^a (\neg\phi \vee \psi)$
$\mathfrak{M}_n \models^a \phi \equiv \psi$	iff	$\mathfrak{M}_n \models^a ((\phi \Rightarrow \psi) \wedge (\psi \Rightarrow \phi))$
$\mathfrak{M}_n \models^a \forall x\phi$	iff	$\mathfrak{M}_n \models^b \phi$ for every assignment \mathbf{b} that may differ from \mathbf{a} only in x and such that $\mathbf{b}(x) \in D_n$
$\mathfrak{M}_n \models^a \exists x\phi$	iff	$\mathfrak{M}_n \models^b \phi$ for some assignment \mathbf{b} that may differ from \mathbf{a} only in x and such that $\mathbf{b}(x) \in D_n$
$\mathfrak{M}_n \models^a \bigcirc\phi$	iff	$\mathfrak{M}_{n+1} \models^a \phi$;
$\mathfrak{M}_n \models^a \diamond\phi$	iff	there exists $m \geq n$ such that $\mathfrak{M}_m \models^a \phi$;
$\mathfrak{M}_n \models^a \square\phi$	iff	for all $m \geq n$, $\mathfrak{M}_m \models^a \phi$;
$\mathfrak{M}_n \models^a (\phi \cup \psi)$	iff	there exists $m \geq n$, such that $\mathfrak{M}_m \models^a \psi$ and, for all $i \in \mathbb{N}$, $n \leq i < m$ implies $\mathfrak{M}_i \models^a \phi$;
$\mathfrak{M}_n \models^a (\phi \mathcal{W} \psi)$	iff	$\mathfrak{M}_n \models^a (\phi \cup \psi)$ or $\mathfrak{M}_n \models^a \square\phi$.

Figure 1: Semantics of FOTL.

Semantics, Intuitively, FOTL formulae are interpreted in *first-order temporal structures* which are sequences \mathfrak{M} of *worlds*, $\mathfrak{M} = \mathfrak{M}_0, \mathfrak{M}_1, \dots$ with truth values in different worlds being connected via temporal operators.

More formally, for every moment of time $n \geq 0$, there is a corresponding *first-order* structure, $\mathfrak{M}_n = \langle D_n, I_n \rangle$, where every D_n is a non-empty set such that whenever $n < m$, $D_n \subseteq D_m$, and I_n is an interpretation of predicate and constant symbols over D_n . We require that the interpretation of constants is *rigid*. Thus, for every constant c and all moments of time $i, j \geq 0$, we have $I_i(c) = I_j(c)$.

A (*variable*) *assignment* \mathbf{a} is a function from the set of individual variables to $\cup_{n \in \mathbb{N}} D_n$. We denote the set of all assignments by \mathfrak{A} . The set of variable assignments \mathfrak{A}_n corresponding to \mathfrak{M}_n is a subset of the set of all assignments, $\mathfrak{A}_n = \{\mathbf{a} \in \mathfrak{A} \mid \mathbf{a}(x) \in D_n \text{ for every variable } x\}$; clearly, $\mathfrak{A}_n \subseteq \mathfrak{A}_m$ if $n < m$.

The *truth* relation $\mathfrak{M}_n \models^a \phi$ in a structure \mathfrak{M} , is defined inductively on the construction of ϕ *only for those assignments \mathbf{a} that satisfy the condition $\mathbf{a} \in \mathfrak{A}_n$* . See Fig. 1 for details. \mathfrak{M} is a *model* for a formula ϕ (or ϕ is *true* in \mathfrak{M}) if, and only if, there exists an assignment \mathbf{a} in D_0 such that $\mathfrak{M}_0 \models^a \phi$. A formula is *satisfiable* if, and only if, it has a model. A formula is *valid* if, and only if, it is true in any temporal structure \mathfrak{M} under any assignment \mathbf{a} in D_0 .

The models introduced above are known as *models with expanding domains* since $D_n \subseteq D_{n+1}$. Another important class of models consists of *models with constant domains* in which the class of first-order temporal structures, where FOTL formulae are interpreted, is restricted to structures $\mathfrak{M} = \langle D_n, I_n \rangle$, $n \in \mathbb{N}$, such that $D_i = D_j$ for all $i, j \in \mathbb{N}$. The notions of truth and validity are defined similarly to the expanding domain case. It is known [32] that satisfiability over expanding domains can be reduced to satisfiability over constant domains with only a polynomial increase in the size of formulae.

2.2 Monodicity and Monadicity

The set of valid formulae of FOTL is not recursively enumerable. Furthermore, it is known that even “small” fragments of FOTL, such as the *two-variable monadic* fragment (where all predicates are unary), are not recursively enumerable [30, 24]. However, the set of valid *monodic* formulae is known to be finitely axiomatisable [33].

Definition 1 An FOTL-formula ϕ is called *monodic* if, and only if, any subformula of the form $T\psi$, where T is one of \bigcirc , \square , \diamond (or $\psi_1 T \psi_2$, where T is one of \bigcup , \bigcap), contains at most one free variable.

We note that the addition of either equality or function symbols to the monodic fragment generally leads to the loss of recursive enumerability [33, 8, 22]. Thus, monodic FOTL is expressive, yet even small extensions lead to serious problems. Further, even with its recursive enumerability, monodic FOTL is generally undecidable. To recover decidability, the easiest route is to restrict the first order part to some decidable fragment of first-order logic, such as the guarded, two-variable or monadic fragments. We here choose the latter, since monadic predicates fit well with our intended application to parameterised systems. Recall that monodicity requires that all predicates have arity of at most ‘1’. Thus, we use monadic, monodic FOTL [7].

A practical approach to proving monodic temporal formulae is to use *fine-grained temporal resolution* [26], which has been implemented in the theorem prover TeMP [25]. In the past, TeMP has been successfully applied to problems from several domains [21], in particular, to examples specified in the temporal logics of knowledge (the fusion of propositional linear-time temporal logic with multi-modal S5) [15, 11, 13]. From this work it is clear that monodic first-order temporal logic is an important tool for specifying complex systems. However, it is also clear that the complexity, even of *monadic* monodic first-order temporal logic, makes this approach difficult to use for larger applications [21, 19].

2.3 XOR Restrictions

An additional restriction we make to the above logic involves implicit XOR constraints over predicates. Such restrictions were introduced into temporal logics in [13], where the correspondence with Büchi automata was described, and generalised in [14]. In both cases, the decision problem is of much better (generally, polynomial) complexity than that for the standard, unconstrained, logic. However, in these papers only *propositional* temporal logic was considered. We now add such an XOR constraint to FOTLX.

The set of predicate symbols $\Pi = \{P_0, P_1, \dots\}$, is now partitioned into a set of XOR-sets, X_1, X_2, \dots, X_n , with one *non-XOR* set N such that

1. all X_i are disjoint with each other,
2. N is disjoint with every X_i ,
3. $\Pi = \bigcup_{j=0}^n X_j \cup N$, and
4. for each X_i , exactly *one* predicate within X_i is satisfied (for any element of the domain) at any moment in time.

Example 1 Consider the formula

$$\forall x. ((P_1(x) \vee P_2(x)) \wedge (P_4(x) \vee P_7(x) \vee P_8(x)))$$

where $\{P_1, P_2\} \subseteq X_1$ and $\{P_4, P_7, P_8\} \subseteq X_2$. The above formula states that, for any element of the domain, a , then one of $P_1(a)$ or $P_2(a)$ must be satisfied and one of $P_4(a)$, $P_7(a)$ or $P_8(a)$ must be satisfied.

2.4 Normal Form

To simplify our description, we will define a *normal form* into which FOTLX formulae can be translated. In the following:

- $\hat{X}_{ij}^-(x)$ denotes a conjunction of negated XOR predicates from the set X_i ;
- $\check{X}_{ij}^+(x)$ denotes a disjunction of (positive) XOR predicates from the set X_i ;
- $\hat{N}_i(x)$ denotes a conjunction of non-XOR literals;
- $\check{N}_i(x)$ denotes a disjunction of non-XOR literals.

A *step* clause is defined as follows:

$$\hat{X}_{1j}^-(x) \wedge \dots \wedge \hat{X}_{nj}^-(x) \wedge \hat{N}_j(x) \Rightarrow \bigcirc (\check{X}_{1j}^+(x) \vee \dots \vee \check{X}_{nj}^+(x) \vee \check{N}_j(x))$$

A *monodic temporal problem in Divided Separated Normal Form (DSNF)* [7] is a quadruple $\langle \mathcal{U}, \mathcal{I}, \mathcal{S}, \mathcal{E} \rangle$, where:

1. the universal part, \mathcal{U} , is a finite set of arbitrary closed first-order formulae;
2. the initial part, \mathcal{I} , is, again, a finite set of arbitrary closed first-order formulae;
3. the step part, \mathcal{S} , is a finite set of step clauses; and
4. the eventuality part, \mathcal{E} , is a finite set of eventuality clauses of the form $\diamond L(x)$, where $L(x)$ is a unary literal.

In what follows, we will not distinguish between a finite set of formulae \mathcal{X} and the conjunction $\bigwedge \mathcal{X}$ of formulae within the set. With each monodic temporal problem, we associate the formula

$$\mathcal{I} \wedge \square \mathcal{U} \wedge \square \forall x \mathcal{S} \wedge \square \forall x \mathcal{E}.$$

Now, when we talk about particular properties of a temporal problem (e.g., satisfiability, validity, logical consequences etc) we mean properties of the associated formula.

Every monodic FOTLX formula can be translated to the normal form in satisfiability preserving way using a renaming and unwinding technique which substitutes non-atomic subformulae and replaces temporal operators by their fixed point definitions as described, for example, in [18]. A step in this transformation is the following: We recursively rename each innermost open subformula $\xi(x)$, whose main connective is a temporal operator, by $P_\xi(x)$, where $P_\xi(x)$ is a new unary predicate, and rename each innermost closed subformula ζ , whose main connective is a temporal operator, by p_ζ , where p_ζ is a new propositional variable. While renaming introduces new, non-XOR predicates and propositions, practical problems stemming from verification are nearly in the normal form, see Section 3.

2.5 Complexity

First-order temporal logics are notorious for being of a high complexity. Even decidable sub-fragments of monodic first-order temporal logic can be too complex for practical use. For example, satisfiability of monodic monadic FOTL logic is known to be EXPSPACE-complete [23]. However, imposing XOR restrictions we obtain better complexity bounds.

Theorem 1 *Satisfiability of monodic monadic FOTLX formulae (in the normal form) can be decided in $2^{O(N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a})}$ time, where N_1, \dots, N_n are cardinalities of the sets of XOR predicates, and N_a is the cardinality of the set of non-XOR predicates.*

Before we sketch the proof of this result, we show how the XOR restrictions influence the complexity of the satisfiability problem for monadic first-order (non-temporal) logic.

Lemma 2 *Satisfiability of monadic first-order formulae can be decided in $\text{NTime}(O(n \cdot N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a}))$, where n is the length of the formula, and N_1, \dots, N_n, N_a are as in Theorem 1.*

Proof As in [4], Proposition 6.2.9, the non-deterministic decision procedure first guesses a structure and then verifies that the structure is a model for the given formula. It was shown, [4], Proposition 6.2.1, Exercise 6.2.3, that if a monadic first-order formula has a model, it also has a model, whose domain is the set of all *predicate colours*. A predicate colour, γ , is a set of unary literals such that for every predicate $P(x)$ from the set of all predicates $X_1 \cup \dots, X_n \cup N$, either $P(x)$ or $\neg P(x)$ belongs to γ . Notice that under the conditions of the lemma, there are at most $N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a}$ different predicate colours. Hence, the structure to guess is of $O(N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a})$ size.

It should be clear that one can evaluate a monadic formula of the size n in a structure of the size $O(N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a})$ in deterministic $O(n \cdot N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a})$ time. Therefore, the overall complexity of the decision procedure is $\text{NTime}(O(n \cdot N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a}))$. \square

Proof [of Theorem 1, Sketch] For simplicity of presentation, we assume the formula contains no propositions. Satisfiability of a monodic FOTL formula is equivalent to a property of the *behaviour graph* for the formula, checkable in time polynomial in the product of the number of different predicate colours and the size of the graph, see [7], Theorem 5.15. For unrestricted FOTL formulae, the size of the behaviour graph is double exponential in the number of predicates. We estimate now the size of the behaviour graph and time needed for its construction for FOTLX formulae.

Let Γ be a set of predicate colours and ρ be a map from the set of constants, $\text{const}(\mathbf{P})$, to Γ . A couple $\langle \Gamma, \rho \rangle$ is called a *colour scheme*. Nodes of the behaviour graph are colour schemes. Clearly, there are no more than $2^{O(N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a})}$ different colour schemes. However, not every colour scheme is a node of the behaviour graph: a colour scheme \mathcal{C} is a node if, and only if, a monadic formula of first-order (non-temporal) logic, constructed from the given FOTLX formula and the colour scheme itself, is satisfiable (for details see [7]). A similar first-order monadic condition determines which nodes are connected with edges. It can be seen that the size of the formula is polynomial in both cases. By Lemma 2, satisfiability of monadic first-order formulae can be decided in deterministic $2^{O(N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a})}$ time.

Overall, the behaviour graph, representing all possible models, for an FOTLX formula can be constructed in $2^{O(N_1 \cdot N_2 \cdot \dots \cdot N_n \cdot 2^{N_a})}$ time. \square

3 Infinite-State Systems

In previous work, notably [17, 9] a parameterised finite state machine based model, suitable for the specification and verification of protocols over arbitrary numbers of processes was defined. Essentially, this uses a family of identical, and synchronously executing, finite state automata with a rudimentary form of communication: if one automaton makes a transition (an action) a , then it is required that *all* other automata simultaneously make a complementary transition (reaction) \bar{a} . In [19] we translated this automata model

into monodic FOTL and used automated theorem proving in that logic to verify parameterised cache coherence protocols [10]. The model assumed not only synchronous behaviour of the communicating automata, but instantaneous broadcast.

Here we present a more general model suitable for specification of both synchronous and asynchronous systems (protocols) with (possibly) delayed broadcast and give its faithful translation into FOTLX. This not only exhibits the power of the logic but, with the improved complexity results of the previous section, provides a route towards the practical verification of temporal properties of infinite state systems.

3.1 Process Model

We begin with a description of both the asynchronous model, and the delayed broadcast approach.

Definition 2 (Protocol) *A protocol, P is a tuple $\langle Q, I, \Sigma, \tau \rangle$, where*

- Q is a finite set of states;
- $I \subseteq Q$ is a set of initial states;
- $\Sigma = \Sigma_L \cup \Sigma_M \cup \bar{\Sigma}_M$, where
 - Σ_L is a finite set of local actions;
 - Σ_M is a finite set of broadcast actions, i.e. “send a message”;
 - $\bar{\Sigma}_M = \{\bar{\sigma} \mid \sigma \in \Sigma_M\}$ is the set of broadcast reactions, i.e. “receive a message”;
- $\tau \subseteq Q \times \Sigma \times Q$ is a transition relation that satisfies the following property

$$\forall \sigma \in \Sigma_M. \forall q \in Q. \exists q' \in Q. \langle q, \bar{\sigma}, q' \rangle \in \tau$$

i.e., “readiness to receive a message in any state”.

Further, we define a notion of global machine, which is a set of n finite automata, where n is a parameter, each following the protocol and able to communicate with others via (possibly delayed) broadcast. To model asynchrony, we introduce a special automaton action, $idle \notin \Sigma$, meaning the automaton is not active and so its state does not change. At any moment an arbitrary group of automata may be idle and all non-idle automata perform their actions in accordance with the transition function τ ; different automata may perform different actions.

Definition 3 (Asynchronous Global Machine) *Given a protocol, $\mathcal{P} = \langle Q, I, \Sigma, \tau \rangle$, the global machine \mathcal{M}_G of dimension n is the tuple $\langle Q_{\mathcal{M}_G}, I_{\mathcal{M}_G}, \tau_{\mathcal{M}_G}, \mathcal{E} \rangle$, where*

- $Q_{\mathcal{M}_G} = Q^n$
- $I_{\mathcal{M}_G} = I^n$
- $\tau_{\mathcal{M}_G} \subseteq Q_{\mathcal{M}_G} \times (\Sigma \cup \{idle\})^n \times Q_{\mathcal{M}_G}$ is a transition relation that satisfies the following property

$$\begin{aligned} & \langle \langle s_1, \dots, s_n \rangle, \langle \sigma_1, \dots, \sigma_n \rangle, \langle s'_1, \dots, s'_n \rangle \rangle \in \tau_{\mathcal{M}_G} \\ & \text{iff} \\ & \forall 1 \leq i \leq n. [(\sigma_i \neq idle \Rightarrow \langle s_i, \sigma_i, s'_i \rangle \in \tau) \\ & \quad \wedge (\sigma_i = idle \Rightarrow s_i = s'_i)]. \end{aligned}$$

- $\mathcal{E} = 2^{\Sigma_M}$ is a communication environment, that is a set of possible sets of messages in transition.

An element $G \in Q_{\mathcal{M}_G} \times (\Sigma \cup \{\text{idle}\})^n \times \mathcal{E}$ is said to be a global configuration of the machine.

A run of a global machine \mathcal{M}_G is a possibly infinite sequence $\langle s^1, \sigma^1, E_1 \rangle \dots \langle s^i, \sigma^i, E_i \rangle \dots$ of global configurations of \mathcal{M}_G satisfying the properties (1)–(6) listed below. In this formulation we assume $s^i = \langle s_1^i, \dots, s_n^i \rangle$ and $\sigma^i = \langle \sigma_1^i, \dots, \sigma_n^i \rangle$.

1. $s^1 \in I^n$
 (“initially all automata are in initial states”);
2. $E_1 = \emptyset$
 (“initially there are no messages in transition”);
3. $\forall i. \langle s^i, \sigma^i, s^{i+1} \rangle \in \tau_{\mathcal{M}_G}$
 (“an arbitrary part of the automata can fire”);
4. $\forall a \in \Sigma_M. \forall i. \forall j. ((\sigma_j^i = a) \Rightarrow \forall k. \exists l \geq i. (\sigma_k^l = \bar{a}))$
 (“delivery to all participants is guaranteed”);
5. $\forall a \in \Sigma_M. \forall i. \forall j. [(\sigma_j^i = \bar{a}) \Rightarrow (a \in E_i) \vee \exists k. \sigma_k^i = a]$ (“one can receive only messages kept by the environment, or sent at the same moment of time ”)

In order to formulate further requirements we introduce the following notation:

$$Sent_i = \{a \in \Sigma_M \mid \exists j. \sigma_j^i = a\}$$

$Delivered_k =$

$$\left\{ a \in \Sigma_M \left| \begin{array}{l} \exists i \leq k. (a \in Sent_i) \wedge \\ (\forall l. (i < l < k) \rightarrow a \notin Sent_l) \wedge \\ (\forall j. \exists l. (i \leq l \leq k) \wedge (\sigma_j^l = \bar{a})) \end{array} \right. \right\}$$

Then, the last requirement the run should satisfy is

6. $\forall i. E_{i+1} = (E_i \cup Sent_i) - Delivered_i$

Example: Asynchronous Floodset Protocol. We illustrate the use of the above model by presenting the specification of an asynchronous FloodSet protocol in our model. This is a variant of the *FloodSet algorithm with alternative decision rule* (in terms of [28], p.105) designed for solution of the Consensus problem.

The setting is as follows. There are n processes, each having an *input bit* and an *output bit*. The processes work asynchronously, run the same algorithm and use *broadcast* for communication. The broadcasted messages are guaranteed to be delivered, though possibly with arbitrary delays. (The process is described graphically in Fig. 2.)

The goal of the algorithm is to eventually reach an agreement, i.e. to produce an output bit, which would be the same for all processes. It is required also that if all processes have the same input bit, that bit should be produced as an output bit.

The asynchronous FloodSet protocol we consider here is adapted from [28]. Main differences with original protocol are:

- the original protocol was synchronous, while our variant is asynchronous;




Figure 2: Asynchronous FloodSet Protocol Process.

- the original protocol assumed instantaneous message delivery, while we allow arbitrary delays in delivery; and
- although the original protocol was designed to work in the presence of crash (or fail-stop) failures, we assume, for simplicity, that there are no failures.

Because of the absence of failures the protocol is very simple and unlike the original one does not require “retransmission” of any value. We will show later (in Section 3.3) how to include the case of crash failures in the specification (and verification). Thus, the asynchronous FloodSet protocol is defined, informally, as follows.

- At the first round of computations, every process broadcasts its input bit.
- At every round the (tentative) output bit is set to the minimum value ever seen so far.

The correctness criterion for this protocol is that, eventually, the output bits of all processes will be the same.

Now we can specify the asynchronous FloodSet as a protocol $\langle Q, I, \Sigma, \tau \rangle$, where $Q = \{i_0, i_1, o_0, o_1\}$; $I = \{i_0, i_1\}$; $\Sigma = \Sigma_m \cup \bar{\Sigma}_m \cup \Sigma_L$ with $\Sigma_m = \{0, 1\}$, $\bar{\Sigma}_m = \{\bar{0}, \bar{1}\}$, $\Sigma_L = \emptyset$. The transition relation $\tau = \{\langle i_0, 0, o_0 \rangle, \langle o_0, \bar{0}, o_0 \rangle, \langle o_0, \bar{1}, o_0 \rangle, \langle i_1, 1, o_1 \rangle, \langle o_1, \bar{0}, o_1 \rangle, \langle o_1, \bar{1}, o_1 \rangle\}$.

3.2 Temporal Translation

Given a protocol $\mathcal{P} = \langle Q, I, \Sigma, \tau \rangle$, we define its translation to FOTLX as follows.

For each $q \in Q$, introduce a monadic predicate symbol P_q and for each $\sigma \in \Sigma \cup \{idle\}$ introduce a monadic predicate symbol A_σ . For each $\sigma \in \Sigma_M$ we introduce also a propositional symbol m_σ .

Intuitively, elements of the domain in the temporal representation will represent exemplars of finite automata, and the formula $P_q(x)$ is intended to represent “automaton x is in state q ”. The formula $A_\sigma(x)$ is going to represent “automaton x performs action σ ”. Proposition m_σ will denote the fact “message σ is in transition” (i.e. it has been sent but not all participants have received it.)

Because of intended meaning we define two XOR-sets: $X_1 = \{P_q \mid q \in Q\}$ and $X_2 = \{A_\sigma \mid \sigma \in \Sigma \cup \{idle\}\}$. All other predicates belong to the set of non-XOR predicates.

- I.** Each automaton either performs one of the actions available in its state, or is idle:
 $\Box[\forall x. P_q(x) \rightarrow A_{\sigma_1}(x) \vee \dots \vee A_{\sigma_k}(x) \vee A_{idle}(x)]$, where $\{\sigma_1, \dots, \sigma_k\} = \{\sigma \in \Sigma \mid \exists r \langle q, \sigma, r \rangle \in \tau\}$.
- II.** Action effects (non-deterministic actions): $\Box[\forall x P_q(x) \wedge A_\sigma(x) \rightarrow \bigcirc \bigvee_{\langle q, \sigma, r \rangle \in \tau} P_r(x)]$ for all $q \in S$ and $\sigma \in \Sigma$.
- III.** Effect of being idle: $\Box[\forall x P_q(x) \wedge A_{idle}(x) \rightarrow \bigcirc P_q(x)]$, for all $q \in S$
- IV.** Initially there are no messages in the transition and all automata are in initial states: **start** $\rightarrow \neg m_\sigma$ for all $\sigma \in \Sigma_m$ and **start** $\rightarrow \forall x \bigvee_{q \in I} P_q(x)$.
- V.** All messages are eventually received (Guarantee of Delivery): $\Box[\exists y A_\sigma(y) \rightarrow \forall x \diamond A_{\bar{\sigma}}(x)]$, for all $\sigma \in \Sigma_m$.
- VI.** Only messages kept in the environment (are in transition), or sent at the same moment of time can be received: $\Box[\forall x A_{\bar{\sigma}}(x) \rightarrow m_\sigma \vee \exists y A_\sigma(y)]$ for all $\sigma \in \Sigma_m$.
- VII.** Finally, for all $\sigma \in \Sigma_m$, we have the conjunction of the following formulae:
1. **start** $\rightarrow \forall x. \neg Received_\sigma(x)$
 2. $\Box[\forall x. (A_{\bar{\sigma}}(x) \wedge \neg \forall y. Received_\sigma(y)) \rightarrow \bigcirc Received_\sigma(x)]$
 3. $\Box[\forall x. (Received_\sigma(x) \wedge \neg \forall y. Received_\sigma(y)) \rightarrow \bigcirc Received_\sigma(x)]$
 4. $\Box[\forall x. (\neg(A_{\bar{\sigma}}(x) \vee Received_\sigma(x)) \wedge \neg \forall y. Received_\sigma(y)) \rightarrow \bigcirc \neg Received_\sigma(x)]$
 5. $\Box[\forall x. Received_\sigma \rightarrow \bigcirc \neg m_\sigma]$
 6. $\Box[\exists x. A_\sigma(x) \wedge \neg \forall y. Received_\sigma(y) \rightarrow \bigcirc m_\sigma]$
 7. $\Box[\neg \exists x. A_\sigma(x) \wedge \neg \forall y. Received_\sigma(y) \rightarrow (m_\sigma \leftrightarrow \bigcirc m_\sigma)]$

Figure 3: Temporal Specification of Abstract Protocol Structure.

We define the temporal translation of \mathcal{P} , called $T_{\mathcal{P}}$, as a conjunction of the formulae in Fig. 3. Note that, in order to define the temporal translation of requirement (6) above, (on the dynamics of environment updates) we introduce the unary predicate symbol $Received_\sigma$ for every $\sigma \in \Sigma_m$.

We now consider the correctness of the temporal translation. This translation of protocol \mathcal{P} is faithful in the following sense.

Proposition 1 *Given a protocol, \mathcal{P} , and a global machine, \mathcal{M}_G , of dimension n , then any temporal model M_1, M_2, \dots of $T_{\mathcal{P}}$ with the finite domain c_1, \dots, c_n of size n represents some run $\langle s^1, \sigma^1, E_1 \rangle \dots \langle s^i, \sigma^i, E_i \rangle \dots$ of \mathcal{M}_G as follows:*

$\langle \langle s_1, \dots, s_n \rangle, \langle \sigma_1, \dots, \sigma_n \rangle, E \rangle$ is i -th configuration of the run iff $M_i \models P_{q_1}(c_1) \wedge \dots \wedge P_{q_n}(c_n)$, $M_i \models A_{\sigma_1}(c_1) \wedge \dots \wedge A_{\sigma_n}(c_n)$ and $E = \{\sigma \in \Sigma_m \mid M_i \models m_\sigma\}$

Dually, for any run of \mathcal{M}_G there is a temporal model of $T_{\mathcal{P}}$ with a domain of size n representing this run.

Proof By routine inspection of the definitions of runs, temporal models and the translation. \square

3.3 Variations of the model

The above model allows various modifications and corresponding version of Proposition 1 still holds.

Determinism. The basic model allows non-deterministic actions. To specify the case of deterministic actions only, one should replace the “Action Effects” axiom in Fig. 3 by the following variant:

$$\Box[\forall x. P_q(x) \wedge A_\sigma(x) \rightarrow \bigcirc P_r(x)]$$

for all $\langle q, \sigma, r \rangle \in \tau$

Explicit bounds on delivery. In the basic mode, no explicit bounds on delivery time are given. To introduce bounds one has to replace the “Guarantee of Delivery” axiom with the following one:

$$\Box[\exists y. A_\sigma(y) \rightarrow \forall x. \bigcirc A_{\bar{\sigma}}(x) \vee \bigcirc A_{\bar{\sigma}}(x) \vee \dots \vee \bigcirc^n A_{\bar{\sigma}}(x)]$$

for all $\sigma \in \Sigma_m$ and some n (representing the maximal delay).

Finite bounds on delivery. One may replace the “Guarantee of Delivery” axiom with the following one

$$\Box[\exists y. A_\sigma(y) \rightarrow \Diamond \forall x. Received_{\bar{\sigma}}(x)]$$

for all $\sigma \in \Sigma_m$.

Crashes. One may replace the “Guarantee of Delivery” axiom by an axiom stating that only the messages sent by normal (non-crashed) participants will be delivered to all participants. (See [19] for examples of such specifications in a FOTL context.)

Guarded actions. One can also extend the model with guarded actions, where action can be performed depending on global conditions in global configurations.

Returning to the FloodSet protocol, one may consider a variation of the asynchronous protocol suitable for resolving the Consensus problem in the presence of *crash failures*. We can modify the above setting as follows. Now, processes may fail and, from that point onward, such processes send no further messages. Note, however, that the messages sent by a process *in the moment of failure* may be delivered to *an arbitrary subset* of the non-faulty processes.

The goal of the algorithm also has to be modified, so only *non-faulty* processes are required to eventually reach an agreement. Thus, the FloodSet protocol considered above is modified by adding the following rule:

- At every round (later than the first), a process broadcasts any value *the first time it sees it*.

Now, in order to specify this protocol the variation of the model with crashes should be used. The above rule can be easily encoded in the model and we leave it as an exercise for the reader.

An interesting point here is that the protocol is actually correct under the assumption that *only finitely many processes may fail*. This assumption is automatically satisfied in our automata model, but not in its temporal translation. Instead, one may use the above *Finite bounds on delivery* axiom to prove the correctness of this variation of the algorithm.

3.4 Verification

Now we have all the ingredients to perform the verification of parameterised protocols. Given a protocol \mathcal{P} , we can translate it into a temporal formula $T_{\mathcal{P}}$. For the temporal representation, χ of a required correctness condition, we then check whether $T_{\mathcal{P}} \rightarrow \chi$ is valid temporal formula. If it is valid, then the protocol is correct for all possible values of the parameter (sizes).

Correctness conditions can, of course, be described using any legal FOTLX formula. For example, for the above FloodSet protocol(s) we have a liveness condition to verify:

$$\diamond(\forall x. o_0(x) \vee \forall x. o_1(x))$$

or, alternatively

$$\diamond \left[\begin{array}{l} (\forall x. \text{Non-faulty}(x) \rightarrow o_0(x)) \vee \\ (\forall x. \text{Non-faulty}(x) \rightarrow o_1(x)) \end{array} \right]$$

in the case of a protocol working in presence of processor crashes.

While space precludes describing many further conditions, we just note that, in [19], we have demonstrated how this approach can be used to verify safety properties, i.e with $\chi = \square\phi$. Since we have the power of FOTLX, but with decidability results, we can also automatically verify fairness formulae of the form $\chi = \square\diamond\phi$.

4 Concluding Remarks

In the propositional case, the incorporation of XOR constraints within temporal logics has been shown to be advantageous, not only because of the reduced complexity of the decision procedure (essentially, polynomial rather than exponential; [14]), but also because of the strong fit between the scenarios to be modelled (for example, finite-state verification) and the XOR logic [13]). The XOR constraints essentially allow us to select a set of names/propositions that must occur exclusively. In the case of verification for finite state automata, we typically consider the automaton states, or the input symbols, as being represented by such sets. Modelling a scenario thus becomes a problem of engineering suitable (combinations of) XOR sets.

In this paper, we have developed an XOR version of FOTL, providing: its syntax and semantics; conditions for decidability; and detailed complexity of the decision procedure. As well as being an extension and combination of the work reported in both [7] and [14], this work forms the basis for tractable temporal reasoning over infinite state problems. In order to motivate this further, we considered a general model concerning the verification of infinite numbers of identical processes. We provide an extension of the work in [19] and [1, 2], tackling liveness properties of infinite-state systems, verification of asynchronous infinite-state systems, and varieties of communication within infinite-state systems. In particular, we are able to capture some of the more complex aspects of *asynchrony* and *communication*, together with the verification of more sophisticated *liveness* and *fairness* properties.

The work in [19] on basic temporal specification such as the above have indeed shown that deductive verification can here be attempted but is expensive — the incorporation of XOR provides significant improvements in complexity.

4.1 Related Work

The properties of first-order temporal logics have been studied, for example, in [24, 23]. Proof methods for the monodic fragment of first order-temporal logics, based on resolution or tableaux have been proposed in [7, 26, 27].

Model checking for parameterised and infinite state-systems is considered in [1]. Formulae are translated into to a Büchi transducer with regular accepting states. Techniques from regular model checking are then used to search for models. This approach has been applied to several algorithms verifying safety properties and some liveness properties.

Constraint based verification using counting abstractions [9, 10, 17], provides complete procedures for checking safety properties of broadcast protocols. However, such approaches

- have theoretically non-primitive recursive upper bounds for decision procedures (although they work well for small, interesting, examples) — in our case the upper bounds are definitely primitive-recursive;
- are not suitable (or, have not been used) for asynchronous systems with delayed broadcast — it is not clear how to adapt these methods for such systems; and
- typically lead to undecidable problems if applied to liveness properties.

4.2 Future Work

Future work involves exploring further the framework described in this paper in particular the development of an implementation to prove properties of protocols in practice. Further, we would like to see if we can extend the range of systems we can tackle beyond the monodic fragment.

We also note that some of the variations we might desire to include in Section 3.3 can lead to undecidable fragments. However, for some of these variations, we have correct although (inevitably) incomplete methods, see [19]. We wish to explore these boundaries further.

References

- [1] P. A. Abdulla, B. Jonsson, M. Nilsson, J. d’Orso, and M. Saksena. Regular Model Checking for LTL(MSO). In *Proc. 16th International Conference on Computer Aided Verification (CAV)*, volume 3114 of *LNCS*, pages 348–360. Springer, 2004.
- [2] P. A. Abdulla, B. Jonsson, A. Rezine, and M. Saksena. Proving Liveness by Backwards Reachability. In *Proc. 17th International Conference on Concurrency Theory (CONCUR)*, volume 4137 of *LNCS*, pages 95–109. Springer, 2006.
- [3] A. Artale, E. Franconi, F. Wolter, and M. Zakharyashev. A Temporal Description Logic for Reasoning over Conceptual Schemas and Queries. In *Proc. European Conference on Logics in Artificial Intelligence (JELIA)*, volume 2424 of *LNCS*, pages 98–110. Springer, 2002.
- [4] E. Börger, E. Grädel, and Yu. Gurevich. *The Classical Decision Problem*. Springer, 1997.
- [5] J. Brotherston, A. Degtyarev, M. Fisher, and A. Lisitsa. Implementing Invariant Search via Temporal Resolution. In *Proc. International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, volume 2514 of *LNCS*, pages 86–101. Springer Verlag, 2002.
- [6] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, Dec. 1999.
- [7] A. Degtyarev, M. Fisher, and B. Konev. Monodic Temporal Resolution. *ACM Transactions on Computational Logic*, 7(1):108–150, January 2006. ([arXiv:cs.LO/0306041](https://arxiv.org/abs/cs.LO/0306041))

- [8] A. Degtyarev, M. Fisher, and A. Lisitsa. Equality and Monodic First-Order Temporal Logic. *Studia Logica*, 72(2):147–156, Nov. 2002.
- [9] G. Delzanno. Automatic Verification of Parameterized Cache Coherence Protocols. In *Proc. 12th International Conference on Computer Aided Verification (CAV)*, volume 1855 of *LNCS*, pages 53–68, 2000.
- [10] G. Delzanno. Constraint-based verification of parametrized cache coherence protocols. *Formal Methods in System Design*, 23(3):257–301, 2003.
- [11] C. Dixon. Using Temporal Logics of Knowledge for Specification and Verification—a Case Study. *Journal of Applied Logic*, 4(1): 50–78, 2006.
- [12] C. Dixon, M.C. Fernández-Gago, M. Fisher, and W. van der Hoek. Using Temporal Logics of Knowledge in the Formal Verification of Security Protocols. In *Proc. International Symposium on Temporal Representation and Reasoning (TIME)*, pages 148–151, 2004. IEEE CS Press,
- [13] C. Dixon, M. Fisher, and B. Konev. Is There a Future for Deductive Temporal Verification? In *Proc. International Symposium on Temporal Representation and Reasoning (TIME)*, pages 11–18, 2006. IEEE CS Press.
- [14] C. Dixon, M. Fisher, and B. Konev. Tractable Temporal Reasoning. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, 2007.
- [15] C. Dixon, M. Fisher, and M. Wooldridge. Resolution for Temporal Logics of Knowledge. *Journal of Logic and Computation*, 8(3):345–372, 1998.
- [16] E. A. Emerson. Temporal and Modal Logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, pages 996–1072. Elsevier, 1990.
- [17] J. Esparza, A. Finkel, and R. Mayr. On the Verification of Broadcast Protocols. In *Proc. 14th IEEE Symposium on Logic in Computer Science (LICS)*, pages 352–359. IEEE CS Press, 1999.
- [18] M. Fisher, C. Dixon, and M. Peim. Clausal Temporal Resolution. *ACM Transactions on Computational Logic*, 2(1):12–56, Jan. 2001. ([arXiv:cs.LO/9907032](https://arxiv.org/abs/cs.LO/9907032))
- [19] M. Fisher, B. Konev, and A. Lisitsa. Practical Infinite-state Verification with Temporal Reasoning. In *Verification of Infinite State Systems and Security*. IOS Press, January 2006.
- [20] D. Gabelaia, R. Kontchakov, A. Kurucz, F. Wolter, and M. Zakharyashev. On the Computational Complexity of Spatio-Temporal Logics. In *Proc. 16th International Florida Artificial Intelligence Research Society Conference (FLAIRS)*, pages 460–464. AAAI Press, 2003.
- [21] M.-C. F. Gago, U. Hustadt, C. Dixon, M. Fisher, and B. Konev. First-Order Temporal Verification in Practice. *Journal of Automated Reasoning*, 34(3):295–321, 2005.
- [22] I. Hodkinson. Monodic Packed Fragment with Equality is Decidable. *Studia Logica*, 72(2):185–197, 2002.
- [23] I. Hodkinson, R. Kontchakov, A. Kurucz, F. Wolter, and M. Zakharyashev. On the Computational Complexity of Decidable Fragments of First-Order Linear Temporal Logics. In *Proc. International Symposium on Temporal Representation and Reasoning (TIME)*, pages 91–98. IEEE CS Press, 2003.
- [24] I. Hodkinson, F. Wolter, and M. Zakharyashev. Decidable Fragments of First-Order Temporal Logics. *Annals of Pure and Applied Logic*, 2000.
- [25] U. Hustadt, B. Konev, A. Riazanov, and A. Voronkov. TeMP: A Temporal Monodic Prover. In *Proc. 2nd International Joint Conference on Automated Reasoning (IJCAR)*, volume 3097 of *LNAI*, pages 326–330. Springer, 2004.
- [26] B. Konev, A. Degtyarev, C. Dixon, M. Fisher, and U. Hustadt. Mechanising First-Order Temporal Resolution. *Information and Computation*, 199(1-2):55–86, 2005.
- [27] R. Kontchakov, C. Lutz, F. Wolter, and M. Zakharyashev. Temporalising Tableaux. *Studia Logica*, 76(1):91–134, 2004.

- [28] N. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers, San Mateo, CA, 1996.
- [29] M. Maidl. A Unifying Model Checking Approach for Safety Properties of Parameterized Systems. *LNCS*, 2102:311–323, 2001.
- [30] S. Merz. Decidability and Incompleteness Results for First-Order Temporal Logic of Linear Time. *Journal of Applied Non-Classical Logics*, 2:139–156, 1992.
- [31] H. Sturm and F. Wolter. A Tableau Calculus for Temporal Description Logic: the Expanding Domain Case. *Journal of Logic and Computation*, 12(5):809–838, 2002.
- [32] F. Wolter and M. Zakharyashev. Decidable Fragments of First-Order Modal Logics. *Journal of Symbolic Logic*, 66:1415–1438, 2001.
- [33] F. Wolter and M. Zakharyashev. Axiomatizing the Monodic Fragment of First-Order Temporal Logic. *Annals of Pure and Applied Logic*, 118(1-2):133–145, 2002.