

Tractable Temporal Reasoning

Keywords: theorem proving; temporal reasoning; mathematical foundations

Abstract

Temporal reasoning is widely used within both Computer Science and Artificial Intelligence. The basic idea of proof, within propositional, discrete temporal logics, is both intuitive and appealing. However, the underlying complexity of such discrete temporal logics has led to use of either simplified formalisms, such as temporal interval algebras, or simplified techniques, such as model checking.

In this paper we show that tractable sub-classes of propositional linear temporal logic can be developed, based on the use of XOR fragments of the logic. We not only show that such fragments can be decided, tractably, via clausal temporal resolution, but also show the benefits of combining multiple XOR fragments. For such combinations we establish completeness and complexity (of the resolution method), and also describe how such a temporal language might be used in a number of application areas, for example the verification of multi-agent systems. This new approach to temporal reasoning provides a framework in which tractable temporal logics can be engineered by intelligently combining appropriate XOR fragments.

1 Introduction

Temporal logics have been used to describe a wide variety of systems, from both Computer Science and Artificial Intelligence. The basic idea of proof, within propositional, discrete temporal logics, is also both intuitive and appealing. However the complexity of satisfiability for such logics is high. For example, the complexity of satisfiability for propositional linear time temporal logic (PTL) is PSPACE-complete [Sistla and Clarke, 1985]. Consequently, model checking [Clarke *et al.*, 1999] has received much attention it also allows users to check that a temporal property holds for some underlying model of the system.

Often temporal problems involve an underlying structure, such as an automaton, where a key property is that the automaton can be in exactly one state at each moment. Such problems frequently involve several process or agents, each with underlying automaton-like structures, and we are interested in properties relating to how the agents progress under

particular models of concurrency such as synchrony, asynchrony etc., or particular coordination or cooperation actions.

In this paper we consider a new fragment of PTL that incorporates the use of XOR operators, denoted $(q_1 \oplus q_2 \oplus \dots \oplus q_n)$ meaning that exactly one q_i holds for $1 \leq i \leq n$. Since the complexity of unsatisfiability for XOR clauses in classical propositional logic is low [Schaefer, 1978], there is the potential to carry much of this over to the temporal case.

Thus, in this paper we provide several results. First, we introduce the PTL fragment to be considered, called TLX, and show a complete clausal resolution system for this. The fragment allows us to split the underlying set of propositions into distinct subsets such that each subset (except one) represents a set of propositions where exactly one proposition can hold (termed *XOR sets*); the remaining set has no such constraints. Then we show that deciding unsatisfiability of specifications in such a logic is, indeed, tractable.

2 XOR Temporal Logic

The logic we consider is called “TLX”, and its syntax and semantics essentially follow that of PTL [Gabbay *et al.*, 1980], with models (isomorphic to the Natural Numbers, \mathbb{N}) of the form: $\sigma = t_0, t_1, t_2, t_3, \dots$ where each state, t_i , is a set of proposition symbols, representing those proposition symbols which are satisfied in the i^{th} moment in time. The notation $(\sigma, i) \models A$ denotes the truth (or otherwise) of formula A in the model σ at state index $i \in \mathbb{N}$. This leads to semantic rules:

$$\begin{aligned} (\sigma, i) \models \bigcirc A & \text{ iff } (\sigma, i+1) \models A \\ (\sigma, i) \models \bigtriangleleft A & \text{ iff } \exists k \in \mathbb{N}. (k \geq i) \text{ and } (\sigma, k) \models A \\ (\sigma, i) \models \square A & \text{ iff } \forall j \in \mathbb{N}. \text{ if } (j \geq i) \text{ then } (\sigma, j) \models A \end{aligned}$$

For any formula A , model σ , and state index $i \in \mathbb{N}$, then either $(\sigma, i) \models A$ holds or $(\sigma, i) \models A$ does not hold, denoted by $(\sigma, i) \not\models A$. If there is some σ such that $(\sigma, 0) \models A$, then A is said to be *satisfiable*. If $(\sigma, 0) \models A$ for all models, σ , then A is said to be *valid* and is written $\models A$.

The main novelty in TLX is that it is parameterised by XOR-sets $\mathcal{P}_1, \mathcal{P}_2, \dots$, and the formulae of $\text{TLX}(\mathcal{P}_1, \mathcal{P}_2, \dots)$ are constructed under the restrictions that *exactly* one proposition from every set \mathcal{P}_i is true in every state. For example, if we consider just one set of propositions \mathcal{P} , we have

$$\square(p_1 \oplus p_2 \oplus \dots \oplus p_n).$$

Furthermore, we assume that there exists a set of propositions in addition to those defined by the parameters, and that

these propositions are unconstrained as normal. Thus, TLX() is essentially a standard propositional, linear temporal logic, while $\text{TLX}(\mathcal{P}, \mathcal{Q}, \mathcal{R})$ is a temporal logic containing at least the propositions $\mathcal{P} \cup \mathcal{Q} \cup \mathcal{R}$, where $\mathcal{P} = \{p_1, p_2, \dots, p_l\}$, $\mathcal{Q} = \{q_1, q_2, \dots, q_m\}$, and $\mathcal{R} = \{r_1, r_2, \dots, r_n\}$, but also satisfying

$$\square[(p_1 \oplus p_2 \oplus \dots \oplus p_l) \wedge (q_1 \oplus q_2 \oplus \dots \oplus q_m) \wedge (r_1 \oplus r_2 \oplus \dots \oplus r_n)]$$

2.1 Normal Form

Assume we have n sets of XOR propositions $P_1 = \{p_{11}, \dots, p_{1N_1}\}, \dots, P_n = \{p_{n1}, \dots, p_{nN_n}\}$ and a set of additional propositions $A = \{a_1, \dots, a_{N_a}\}$. In the following:

- \hat{P}_{ij}^- denotes a conjunction of negated XOR propositions from the set P_i ;
- \check{P}_{ij}^+ denotes a disjunction of (positive) XOR propositions from the set P_i ;
- \hat{A}_i denotes a conjunction of non-XOR literals;
- \check{A}_i denotes a disjunction of non-XOR literals.

A normal form for TLX is of the form $\square \bigwedge_i C_i$ where each C_i is an *initial*, *step* or *sometime* clause (respectively) as follows:

$$\begin{aligned} \text{start} &\Rightarrow \check{P}_{1i}^+ \vee \dots \vee \check{P}_{ni}^+ \vee \check{A}_i \\ \hat{P}_{1j}^- \wedge \dots \wedge \hat{P}_{nj}^- \wedge \hat{A}_j &\Rightarrow \bigcirc(\check{P}_{1j}^+ \vee \dots \vee \check{P}_{nj}^+ \vee \check{A}_j) \\ \text{true} &\Rightarrow \diamond(\check{P}_{1k}^+ \vee \dots \vee \check{P}_{nk}^+ \vee \check{A}_k). \end{aligned}$$

Note that due to the semantics of the XOR clauses, if $i \neq k$

$$\begin{aligned} p_{ji} \wedge p_{jk} &\equiv \text{false} & \neg p_{ji} \vee \neg p_{jk} &\equiv \text{true} \\ \text{and} & \bigwedge_{i=1}^{N_j} \neg p_{ji} &\equiv \text{false} & \bigvee_{i=1}^{N_j} p_{ji} &\equiv \text{true}. \\ \text{Also } p_{ji} &\equiv \bigwedge_{p_{jk} \in P_j, k \neq i} \neg p_{jk} & \neg p_{ji} &\equiv \bigvee_{p_{jk} \in P_j, k \neq i} \neg p_{jk} \end{aligned}$$

allow us to maintain positive XOR propositions on the right hand sides of clauses and negated XOR propositions on the left hand side of clauses.

2.2 Resolution Rules

We decide the validity of formulae in TLX using a form of clausal temporal resolution [Fisher *et al.*, 2001]. The resolution rules are split into three types: *initial resolution*, *step resolution* and *temporal resolution*. These are presented in Fig. 1. Initial resolution resolves constraints holding in the initial moment in time. Step resolution involves resolving two step clauses or deriving additional constraints when a contradiction in the next moment is derived. Temporal resolution resolves a sometime clause with a constraint that ensures that the right hand side of this clause cannot occur.

In the conclusion of the following resolution rules $\text{com}(\check{P}_{ij}^+, \check{P}_{ik}^+)$ denotes the disjunction of the propositions in both \check{P}_{ij}^+ and \check{P}_{ik}^+ or **false** if there are no propositions common to both. For example, $\text{com}(p_1 \vee p_2, p_2 \vee p_3) = p_2$.

Observe that IRES_A and SRES_A apply classical resolution to the right hand side of the parent clauses whereas IRES_{P_k}

and SRES_{P_k} involve XOR resolution. Note we can only apply IRES_A and SRES_A between clauses with complementary (non-XOR) literals on the right hand side. We can also apply the IRES_{P_k} and SRES_{P_k} rules to these clauses but the disjunct $\check{A}_1 \vee \check{A}_2$ on the right hand side of the conclusion will be equivalent to **true**.

3 Soundness and Completeness

Similarly to [Fisher *et al.*, 2001; Degtyarev *et al.*, 2006], one can show that whenever the parent clauses are satisfiable then so is the resolvent. Since all the rules of *initial*, and *step* resolution follow the same pattern, we first prove the classical propositional counterpart of the completeness theorem, and then use it to prove the completeness of temporal resolution. Consider the following *classical* set of resolution rules consisting of the rule RES_A :

$$\frac{(\check{P}_{11}^+ \vee \dots \vee \check{P}_{n1}^+ \vee \check{A}_1 \vee a); \quad (\check{P}_{12}^+ \vee \dots \vee \check{P}_{n2}^+ \vee \check{A}_2 \vee \neg a)}{(\check{P}_{11}^+ \vee \dots \vee \check{P}_{n1}^+ \vee \check{P}_{12}^+ \vee \dots \vee \check{P}_{n2}^+ \vee \check{A}_1 \vee \check{A}_2)}$$

and, for every $k \in \{1, \dots, n\}$, the rule RES_{P_k} :

$$\frac{(\check{P}_{11}^+ \vee \dots \vee \check{P}_{k1}^+ \vee \dots \vee \check{P}_{n1}^+ \vee A_1) \quad (\check{P}_{12}^+ \vee \dots \vee \check{P}_{k2}^+ \vee \dots \vee \check{P}_{n2}^+ \vee A_2)}{(\check{P}_{11}^+ \vee \check{P}_{12}^+ \vee \dots \vee \text{com}(\check{P}_{k1}^+, \check{P}_{k2}^+) \vee \dots \vee \check{P}_{n1}^+ \vee \check{P}_{n2}^+ \vee \check{A}_1 \vee \check{A}_2)}$$

Lemma 1 *If a set of classical propositional clauses is unsatisfiable then its unsatisfiability can be established by the rules RES_A and RES_{P_k} in $O(N_1 \times N_2 \times \dots \times N_n \times 2^{N_a})$ time.*

Proof: First we show that if an unsatisfiable set of clauses \mathcal{C} does not contain non-XOR literals, then its unsatisfiability can be established by rules RES_{P_k} . Note that any such set of clauses \mathcal{C} is unsatisfiable if, and only if, for every l , $0 < l \leq n$, and every set of propositions p_1, p_2, \dots, p_l , where $p_i \in P_i$, the set $\mathcal{C}_{p_1, \dots, p_l}$ of clauses from \mathcal{C} , which contain *none* of p_1, \dots, p_l , is nonempty. Indeed, otherwise every clause from \mathcal{C} contains at least one of the propositions p_1, \dots, p_l , so making p_1, \dots, p_l true satisfies \mathcal{C} .

Assume all clauses from \mathcal{C} consist of propositions from P_1, \dots, P_k only (originally, $k = n$) and show that with the rule RES_{P_k} one can obtain an unsatisfiable set of clauses \mathcal{C}' in which all clauses consist of propositions from P_1, \dots, P_{k-1} only.

Take arbitrary propositions $p_1 \in P_1, p_2 \in P_2, \dots, p_{k-1} \in P_{k-1}$ and take arbitrary clauses $C_1 \in \mathcal{C}_{p_1, p_2, \dots, p_{k-1}, p_{k,1}}, C_2 \in \mathcal{C}_{p_1, p_2, \dots, p_{k-1}, p_{k,2}}, \dots, C_{N_k} \in \mathcal{C}_{p_1, p_2, \dots, p_{k-1}, p_{k, N_k}}$. Applying rule RES_{P_k} to C_1, \dots, C_{N_k} one can obtain a clause \mathcal{C}' consisting of propositions from P_1, \dots, P_{k-1} only such that \mathcal{C}' does not contain any of p_1, \dots, p_{k-1} . The set \mathcal{C}' is formed from such clauses \mathcal{C}' for all possible combinations of $p_1 \in P_1, p_2 \in P_2, \dots, p_{k-1} \in P_{k-1}$. Clearly, for every l , $0 < l \leq n$, and every set of propositions p_1, p_2, \dots, p_l , where $p_i \in P_i$, the set $\mathcal{C}'_{p_1, \dots, p_l}$ is nonempty, hence, \mathcal{C}' is unsatisfiable. Applying this reasoning at most n times, one can obtain an empty clause.

Consider now a set of clauses \mathcal{C} , which may contain non-XOR literals. For arbitrary $p_1 \in P_1, \dots, p_n \in P_n$ consider $\mathcal{C}_{p_1, \dots, p_n}$. Similarly to the previous case, every such

Initial Resolution:

$$\boxed{\text{IRES}_A} \quad \begin{array}{l} \text{start} \Rightarrow (\check{P}_{11}^+ \vee \dots \vee \check{P}_{n1}^+ \vee \check{A}_1 \vee a) \\ \text{start} \Rightarrow (\check{P}_{12}^+ \vee \dots \vee \check{P}_{n2}^+ \vee \check{A}_2 \vee \neg a) \\ \hline \text{start} \Rightarrow (\check{P}_{11}^+ \vee \dots \vee \check{P}_{n1}^+ \vee \check{P}_{12}^+ \vee \dots \vee \check{P}_{n2}^+ \vee \check{A}_1 \vee \check{A}_2) \end{array}$$

For every $k \in \{1, \dots, n\}$ we have the rule.

$$\boxed{\text{IRES}_{P_k}} \quad \begin{array}{l} \text{start} \Rightarrow (\check{P}_{11}^+ \vee \dots \vee \check{P}_{k1}^+ \vee \dots \vee \check{P}_{n1}^+ \vee \check{A}_1) \\ \text{start} \Rightarrow (\check{P}_{12}^+ \vee \dots \vee \check{P}_{k2}^+ \vee \dots \vee \check{P}_{n2}^+ \vee \check{A}_2) \\ \hline \text{start} \Rightarrow (\check{P}_{11}^+ \vee \check{P}_{12}^+ \vee \dots \vee \text{com}(\check{P}_{k1}^+, \check{P}_{k2}^+) \vee \dots \vee \check{P}_{n1}^+ \vee \check{P}_{n2}^+ \vee \check{A}_1 \vee \check{A}_2) \end{array}$$

Step Resolution:

$$\boxed{\text{SRES}_A} \quad \begin{array}{l} \hat{A}_1 \wedge \hat{P}_{11}^- \wedge \dots \wedge \hat{P}_{n1}^- \Rightarrow \bigcirc(\check{P}_{11}^+ \vee \dots \vee \check{P}_{n1}^+ \vee \check{A}_1 \vee a) \\ \hat{A}_2 \wedge \hat{P}_{12}^- \wedge \dots \wedge \hat{P}_{n2}^- \Rightarrow \bigcirc(\check{P}_{12}^+ \vee \dots \vee \check{P}_{n2}^+ \vee \check{A}_2 \vee \neg a) \\ \hline \hat{A}_1 \wedge \hat{A}_2 \wedge \hat{P}_{11}^- \wedge \dots \wedge \hat{P}_{n1}^- \wedge \hat{P}_{12}^- \wedge \dots \wedge \hat{P}_{n2}^- \Rightarrow \bigcirc(\check{P}_{11}^+ \vee \dots \vee \check{P}_{n1}^+ \vee \check{P}_{12}^+ \vee \dots \vee \check{P}_{n2}^+ \vee \check{A}_1 \vee \check{A}_2) \end{array}$$

For every $k \in \{1, \dots, n\}$ we have the rule

$$\boxed{\text{SRES}_{P_k}} \quad \begin{array}{l} \hat{A}_1 \wedge \hat{P}_{11}^- \wedge \dots \wedge \hat{P}_{n1}^- \Rightarrow \bigcirc(\check{P}_{11}^+ \vee \dots \vee \check{P}_{k1}^+ \vee \dots \vee \check{P}_{n1}^+ \vee A_1) \\ \hat{A}_2 \wedge \hat{P}_{12}^- \wedge \dots \wedge \hat{P}_{n2}^- \Rightarrow \bigcirc(\check{P}_{12}^+ \vee \dots \vee \check{P}_{k2}^+ \vee \dots \vee \check{P}_{n2}^+ \vee A_2) \\ \hline \hat{A}_1 \wedge \hat{A}_2 \wedge \hat{P}_{11}^- \wedge \dots \wedge \hat{P}_{n1}^- \wedge \hat{P}_{12}^- \wedge \dots \wedge \hat{P}_{n2}^- \Rightarrow \bigcirc(\check{P}_{11}^+ \vee \check{P}_{12}^+ \vee \dots \vee \text{com}(\check{P}_{k1}^+, \check{P}_{k2}^+) \vee \dots \vee \check{P}_{n1}^+ \vee \check{P}_{n2}^+ \vee A_1 \vee A_2) \end{array}$$

$$\boxed{\text{CONV}} \quad \begin{array}{l} \hat{A}_1 \wedge \hat{P}_{11}^- \wedge \dots \wedge \hat{P}_{n1}^- \Rightarrow \bigcirc \text{false} \\ \hline \text{start} \Rightarrow (\neg \hat{A}_1^- \vee \neg \hat{P}_{11}^- \vee \dots \vee \neg \hat{P}_{n1}^-); \quad \text{true} \Rightarrow \bigcirc(\neg \hat{A}_1^- \vee \neg \hat{P}_{11}^- \vee \dots \vee \neg \hat{P}_{n1}^-) \end{array}$$

Temporal Resolution:

$$\boxed{\text{TRES}} \quad \begin{array}{l} L \Rightarrow \square(\neg \check{P}_{11}^+ \wedge \dots \wedge \neg \check{P}_{n1}^+ \wedge \neg A_1) \\ \text{true} \Rightarrow \diamond(\check{P}_{11}^+ \vee \dots \vee \check{P}_{n1}^+ \vee A_1) \\ \hline \text{start} \Rightarrow \neg L \quad \text{true} \Rightarrow \bigcirc \neg L \end{array}$$

Figure 1: Resolution Rules for the XOR Fragment

$\mathcal{C}_{p_1, \dots, p_n}$ should be nonempty. Consider the set $\tilde{\mathcal{C}}_{p_1, \dots, p_n}$ of clauses obtained by deleting all XOR-propositions from clauses of $\mathcal{C}_{p_1, \dots, p_n}$. Every $\tilde{\mathcal{C}}_{p_1, \dots, p_n}$ must be unsatisfiable (otherwise, extending the satisfying assignment for $\tilde{\mathcal{C}}_{p_1, \dots, p_n}$ with p_1, \dots, p_n we satisfy all the clauses in \mathcal{C}). Then classical binary resolution will be able to prove unsatisfiability of $\tilde{\mathcal{C}}_{p_1, \dots, p_n}$. Applying RES_A “in the same way”, one can obtain a clause C' , which does not contain neither non-XOR literals, nor p_1, \dots, p_n . The set \mathcal{C}' , formed from such clauses C' for all possible combinations of $p_1 \in P_1, p_2 \in P_2, \dots, p_{k-1} \in P_{k-1}$, is an unsatisfiable set of clauses not containing non-XOR literals.

Finally, one can see that it is possible to implement the described procedure in $O(N_1 \times N_2 \times \dots \times N_n \times 2^{N_a})$ time. \square

Next we sketch the proof of completeness of temporal resolution, which is obtained combining the ideas of [Fisher *et al.*, 2001; Degtyarev *et al.*, 2002] and Lemma 1.

Definition 1 (Behaviour Graph) We split the set of temporal clauses into three groups. Let \mathcal{I} denote the initial clauses;

\mathcal{T} be the set of all step clauses; and \mathcal{E} be the sometime clauses.

Given a set of clauses over a set of propositional symbols \mathcal{P} , we construct a finite directed graph G as follows. The nodes of G are interpretations of the set of propositions, that satisfy the XOR constraints over the XOR subsets. Notice that, because of the XOR-constraints, exactly one proposition from each set of XOR propositions P_i and any subset of propositions in A are true in I . This means that there at most $N_1 \times N_2 \times \dots \times N_n \times 2^{N_a}$ nodes in the behaviour graph.

For each node, I , we construct an edge in G to a node I' if, and only if, the following condition is satisfied:

- For every step clause $(P \Rightarrow \bigcirc Q) \in \mathcal{T}$, if $I \models P$ then $I' \models Q$.

A node, I , is designated an initial node of G if $I \models \mathcal{I}$. The behaviour graph G of the set of clauses is the maximal sub-graph of G given by the set of all nodes reachable from initial nodes.

If G is empty then the set \mathcal{I} is unsatisfiable. In this case there must exist a derivation by IRES_A and IRES_{P_k} as described in Lemma 1 (and in $O(N_1 \times N_2 \times \dots \times N_n \times 2^{N_a})$ time).

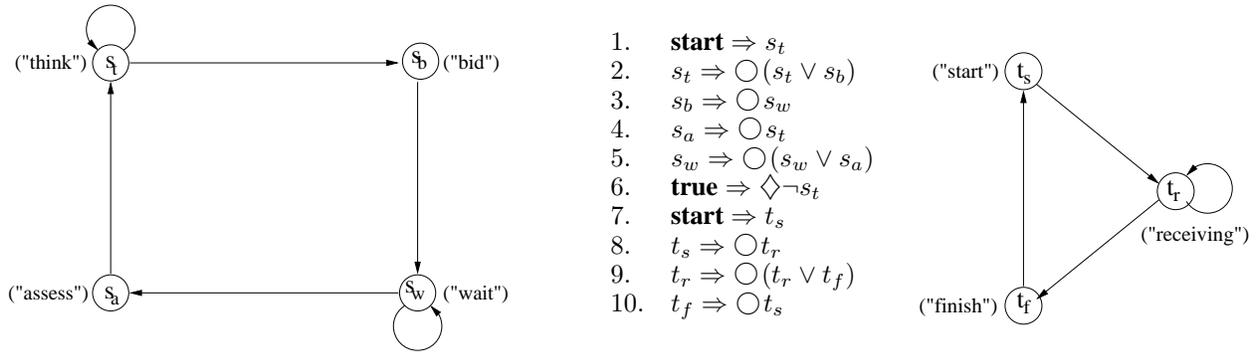


Figure 2: Automata for agents S and T , together with corresponding clauses in normal form.

Now suppose G is not empty. Let I be a node of G which has no successors. Let $\{(P_i \Rightarrow \bigcirc Q_i)\}$ be the set of all step clauses such that $I \models P_i$, then $\bigwedge Q_i$ is unsatisfiable. Using Lemma 1, one can show that step resolution proves $\bigwedge P_i \Rightarrow \bigcirc \text{false}$. After the set of clauses is extended by the conclusion of the CONV rule, $\bigvee \neg P_i$, the node I is deleted from the graph.

In the case when all nodes of G have a successor, a contradiction can be derived with the help of the temporal resolution rule TRES. Note that we impose no restriction on this rule (it coincides with the temporal resolution rule for the general calculi presented in [Fisher *et al.*, 2001; Degtyarev *et al.*, 2002]) and the proof of completeness is no different from what is already published [Fisher *et al.*, 2001; Degtyarev *et al.*, 2002].

4 Complexity

Again, we consider initial and step resolution first.

Lemma 2 *Using the rules of initial and step resolution, it is possible to reduce a set of temporal clauses to one whose behaviour graph does not have nodes without successors in $O\left((N_1 \times N_2 \times \dots \times N_n \times 2^{N_a})^3\right)$ time.*

Proof: Consider the following resolution strategy. For every set of propositions $p_1 \in P_1, \dots, p_n \in P_n$ and $a \in A$, consider the set of all step-clauses

$$\hat{A}_1 \wedge \hat{P}_{11}^- \wedge \dots \wedge \hat{P}_{n1}^- \Rightarrow \bigcirc (\check{P}_{11}^+ \vee \dots \vee \check{P}_{n1}^+ \vee \check{A}_1)$$

such that $\hat{A}_1, \hat{P}_{11}^-, \dots, \hat{P}_{n1}^-$ do not contain any of a, p_1, \dots, p_n (there are at most $N_1 \times N_2 \times \dots \times N_n \times 2^{N_a}$ such sets of clauses), and try establishing the unsatisfiability of the conjunction of the right-hand sides together with the universal clauses by step resolution (as Lemma 1 shows, this can be done in $O(N_1 \times N_2 \times \dots \times N_n \times 2^{N_a})$ time. Then, all nodes without successors will be deleted from the behaviour graph (but some new such nodes may emerge). After $N_1 \times N_2 \times \dots \times N_n \times 2^{N_a}$ repetitions, we obtain a graph in which every node has a successor. \square

Lemma 3 *Given a set of temporal clauses, it is possible to find L such that $L \Rightarrow \square \bigwedge_k \neg q_k$, as required for the TRES*

rule, in time polynomial in $N_1 \times N_2 \times \dots \times N_n \times 2^{N_a}$.

Proof: To find such L , it suffices to find a strongly connected component in the behaviour graph of the set of clauses, such that for every node I of this component, $I \models \bigwedge_k \neg q_k$. The simplest brute-force algorithm would analyse all pairs of nodes (and there are $(N_1 \times N_2 \times \dots \times N_n \times 2^{N_a})^2$ such pairs), and this can be done more efficiently with step resolution as in [Degtyarev *et al.*, 2006]. \square

Theorem 4 *If a set of temporal clauses is unsatisfiable, temporal resolution will deduce a contradiction in time polynomial in $N_1 \times N_2 \times \dots \times N_n \times 2^{N_a}$.*

5 Example

Having described the underlying approach, we will now consider an example that makes use of some of these aspects. In particular, we will have multiple XOR fragments, together with standard propositions (unconstrained by XOR clauses). The example we will use is a simplification and abstraction of agent negotiation protocols; see, for example [Ballarini *et al.*, 2006]. Here, several (in our case, two) agents exchange information in order to come to some agreement. Each agent essentially has a simple control cycle, which can be represented as a finite state machine. In fact, we have simplified these still further, and sample basic control cycles are given in Fig. 2 (for both agents S and T).

Thus, we aim to use these automata as models of the agents, then formalise these within our logic. Importantly, we will add additional clauses (and propositions) characterising agreements or concurrency and, finally, we will show how our resolution method can be used to carry out verification.

We begin by characterising each agent separately as a set of clauses within our logic. To achieve this, we use a set of propositions for each agent. Thus, the automaton describing agent S is characterised through propositions of the form s_a, s_b , etc., while the automaton describing agent T is characterised using propositions such as t_r, t_s , etc. Both these sets are XOR sets. Thus, exactly one of s_a, s_b, \dots , and exactly one of t_r, t_s, \dots , must be true at any moment in time.

Now, the set of clauses characterising the two automata are given in Fig. 2. Regarding automaton S 's description, note that clause 6 ensures that the automaton is infinitely often in

1.	start	$\Rightarrow s_t$			
2.	$\neg s_b \wedge \neg s_w \wedge \neg s_a$	$\Rightarrow \bigcirc(s_t \vee s_b)$	10.	$\neg t_s \wedge \neg t_r$	$\Rightarrow \bigcirc t_s$
3.	$\neg s_t \wedge \neg s_w \wedge \neg s_a$	$\Rightarrow \bigcirc s_w$	11.	true	$\Rightarrow \diamond agree$
4.	$\neg s_t \wedge \neg s_b \wedge \neg s_w$	$\Rightarrow \bigcirc s_t$	12.	$(agree \wedge \neg s_t \wedge \neg s_b \wedge \neg s_a \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc s_a$
5.	$\neg s_t \wedge \neg s_b \wedge \neg s_a$	$\Rightarrow \bigcirc(s_w \vee s_a)$	13.	$(agree \wedge \neg s_t \wedge \neg s_b \wedge \neg s_a \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc t_f$
6.	true	$\Rightarrow \diamond(s_b \vee s_w \vee s_a)$	14.	$(\neg agree \wedge \neg s_t \wedge \neg s_b \wedge \neg s_a)$	$\Rightarrow \bigcirc s_w$
7.	start	$\Rightarrow t_s$	15.	$(\neg agree \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc t_r$
8.	$\neg t_r \wedge \neg t_f$	$\Rightarrow \bigcirc t_r$	16.	$(agree \wedge \neg s_t \wedge \neg s_b \wedge \neg s_a \wedge \neg t_r)$	$\Rightarrow \bigcirc s_w$
9.	$\neg t_s \wedge \neg t_f$	$\Rightarrow \bigcirc(t_r \vee t_f)$	17.	$(agree \wedge \neg s_w \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc t_r$
18.	true	$\Rightarrow \bigcirc(s_b \vee s_w \vee s_a \vee t_r \vee t_f)$			
19.	$(\neg s_t \wedge \neg s_b \wedge \neg s_w \wedge \neg t_s \wedge \neg t_r)$	$\Rightarrow \bigcirc \text{false}$			[18, 10, 4 SRES _{P_k}]
20.	true	$\Rightarrow \bigcirc(s_t \vee s_b \vee s_w \vee t_s \vee t_r)$			[19 CONV]
21.	$(agree \wedge \neg s_t \wedge \neg s_b \wedge \neg s_a \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc \text{false}$			[20, 12, 13 SRES _{P_k}]
22.	true	$\Rightarrow \bigcirc(\neg agree \vee s_t \vee s_b \vee s_a \vee t_s \vee t_f)$			[21 CONV]
23.	$(\neg agree \wedge \neg s_t \wedge \neg s_b \wedge \neg s_a \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc \neg agree$			[22, 14, 15 SRES _{P_k}]
24.	true	$\Rightarrow \bigcirc(agree \vee s_t \vee s_b \vee s_a \vee t_s \vee t_f)$			[23, 15, 14, 11 TRES]
25.	true	$\Rightarrow \bigcirc(s_t \vee s_b \vee s_a \vee t_s \vee t_f)$			[24, 22 SRES _A]
26.	$(\neg s_t \wedge \neg s_w \wedge \neg s_a)$	$\Rightarrow \bigcirc(t_s \vee t_f)$			[25, 3 SRES _{P_k}]
27.	$(\neg agree \wedge \neg s_t \wedge \neg s_w \wedge \neg s_a \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc \text{false}$			[26, 15 SRES _{P_k}]
28.	true	$\Rightarrow \bigcirc(agree \vee s_t \vee s_w \vee s_a \vee t_s \vee t_f)$			[27 CONV]
29.	$(agree \wedge \neg s_t \wedge \neg s_w \wedge \neg s_a \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc \text{false}$			[26, 17 SRES _{P_k}]
30.	true	$\Rightarrow \bigcirc(\neg agree \vee s_t \vee s_w \vee s_a \vee t_s \vee t_f)$			[29 CONV]
31.	true	$\Rightarrow \bigcirc(s_t \vee s_w \vee s_a \vee t_s \vee t_f)$			[28, 30 SRES _A]
32.	$\neg s_b \wedge \neg s_w \wedge \neg s_a$	$\Rightarrow \bigcirc(s_t \vee t_s \vee t_f)$			[31, 2 SRES _{P_k}]
33.	$(\neg agree \wedge \neg s_b \wedge \neg s_w \wedge \neg s_a \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc s_t$			[32, 15 SRES _{P_k}]
34.	$(agree \wedge \neg s_b \wedge \neg s_w \wedge \neg s_a \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc s_t$			[32, 17 SRES _{P_k}]
35.	true	$\Rightarrow \bigcirc(s_b \vee s_w \vee s_a \vee t_s \vee t_f)$			[33, 15, 34, 17, 6 TRES]
36.	$(\neg agree \wedge \neg s_t \wedge \neg s_b \wedge \neg s_w \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc \text{false}$			[35, 15, 4 SRES _{P_k}]
37.	true	$\Rightarrow \bigcirc(agree \vee s_t \vee s_b \vee s_w \vee t_s \vee t_f)$			[36 CONV]
38.	$(agree \wedge \neg s_t \wedge \neg s_b \wedge \neg s_w \wedge \neg t_s \wedge \neg t_f)$	$\Rightarrow \bigcirc \text{false}$			[35, 17, 4 SRES _{P_k}]
39.	true	$\Rightarrow \bigcirc(\neg agree \vee s_t \vee s_b \vee s_w \vee t_s \vee t_f)$			[38 CONV]
40.	true	$\Rightarrow \bigcirc(s_t \vee s_b \vee s_w \vee t_s \vee t_f)$			[37, 39 SRES _A]
41.	true	$\Rightarrow \bigcirc(t_s \vee t_f)$			[40, 35, 31, 25 SRES _{P_k}]
42.	$\neg t_r \wedge \neg t_f$	$\Rightarrow \bigcirc \text{false}$			[41, 8 SRES _{P_k}]
43.	start	$\Rightarrow t_r \vee t_f$			[42 CONV]
44.	start	$\Rightarrow \text{false}$			[43, 7 IRES _{P_k}]

Figure 3: Resolution Proof for Automata Agents Example.

a state other than s_t , ensuring that the automaton can not remain in state s_t forever.

We can also characterise how the computations within each automaton relate. To begin with, we assume a simple, synchronous, concurrent model where both automata make a transition at the same time (see Section 5 for variations on this). Next we add a key aspect in negotiation protocols, namely a description of what happens when an *agreement* is reached. In our example, this is characterised as a synchronised communication act. Logically, we use the proposition *agree* to denote this, and add the following clauses.

11. **true** $\Rightarrow \diamond agree$
12. $(agree \wedge s_w \wedge t_r) \Rightarrow \bigcirc s_a$
13. $(agree \wedge s_w \wedge t_r) \Rightarrow \bigcirc t_f$
14. $(\neg agree \wedge s_w) \Rightarrow \bigcirc s_w$
15. $(\neg agree \wedge t_r) \Rightarrow \bigcirc t_r$
16. $(s_w \wedge agree \wedge \neg t_r) \Rightarrow \bigcirc s_w$
17. $(\neg s_w \wedge agree \wedge t_r) \Rightarrow \bigcirc t_r$

Here, we say that agreements *will* occur infinitely often in the future (clause 11). Clauses 12 and 13 capture the exact synchronisation. If an agreement occurs while automaton S is in state s_w and automaton T is in t_r , then the automata make transitions forward to states s_a and t_f respectively. Finally, clauses 14–17 ensure that, if no synchronised agreement is possible, then the automata remain in their relevant states.

The clauses above represent the specification of a simple system. As an example of how resolution can be used, we also wish to verify that the system is *simultaneously* in states s_t and t_s infinitely often. To verify this, we add the negation of this property, as characterised by clause 18:

18. **true** $\Rightarrow \bigcirc(\neg s_t \vee \neg t_s)$

Thus, if we can derive a contradiction from clauses 1–18 then we know the negated property is valid for this specification. We first rewrite clauses 1–18 in the correct format for the normal form. The refutation is given in Figure 3.

The example above essentially captures activity within a

synchronous, truly concurrent, system. If we wish to move to more complex models of computation, we can do so, essentially by introducing the notion of a *turn*. Thus, when it is automaton S 's turn to move, $turn_s$ is true; when it is automaton T 's turn to move, $turn_t$ is true. Then, each clause describing an automaton transition, for example, $3. s_b \Rightarrow \bigcirc s_w$ is replaced by two clauses

$$\begin{aligned} 3a. \quad & (s_b \wedge turn_s) \Rightarrow \bigcirc s_w \\ 3b. \quad & (s_b \wedge \neg turn_s) \Rightarrow \bigcirc s_b. \end{aligned}$$

In the example above, $turn_s$ and $turn_t$ are effectively both true together (and forever). However, we can modify the synchronisation clauses and model a different form of concurrency. For example, if we were to introduce *interleaving* concurrency, we might use the following clauses¹:

$$\mathbf{start} \Rightarrow turn_s \quad turn_s \Rightarrow \bigcirc turn_t \quad turn_t \Rightarrow \bigcirc turn_s$$

If we go further still, and introduce an asynchronous model of concurrency, then we might get

$$\mathbf{true} \Rightarrow \diamond turn_s \quad \mathbf{true} \Rightarrow \diamond turn_t$$

In both the above cases if we want to ensure that exactly one of $turn_s$ and $turn_t$ hold at each moment we implicitly have $\square(turn_s \oplus turn_t)$ and so we are effectively using $TLX(\mathcal{S}, \mathcal{T}, \{turn_s, turn_t\})$.

6 Concluding Remarks and Related Work

In this paper we have developed a tractable sub-class of temporal logic, based on the central use of XOR operators. This logic can be decided, tractably, via clausal temporal resolution. Importantly, multiple XOR fragments can be combined. This new approach to temporal reasoning provides a framework in which tractable temporal logics can be engineered by intelligently combining appropriate XOR fragments.

The complexity result means that TLX is more amenable to efficient implementation than other similar temporal logics. Moreover, since no two propositions from the same XOR set can occur in the right- (or left-) hand side of any temporal clause, one can efficiently represent disjunctions of (positive) propositions (and conjunctions of negated propositions) as bit vectors and the rules of temporal resolution as bit-wise operations on such bit vectors. Thus, temporal reasoning in TLX can be efficient not only in theory, but also in practice.

Demri and Schnoebelen [2002] consider sub-fragments of PTL, particularly those restricting the number of propositions, the temporal operators allowed, and the depth of temporal nesting in formulae. Demri and Schnoebelen show that, since the formulae tackled in practical model checking often fall within such fragments, then this provides a natural explanation for the viability of model checking in PTL.

Recent results relating to a clausal resolution calculus for propositional temporal logics can be found in [Fisher *et al.*, 2001; Hustadt and Konev, 2003; Hustadt *et al.*, 2004]. Since deciding unsatisfiability of PTL is also PSPACE-complete, then deductive verification of PTL formulae would seem to be an impractical way to proceed. However, just as Demri and

Schnoebelen showed how PTL model checking can be seen as being tractable when we consider fragments of PTL, so we have been examining fragments of PTL that allow clausal resolution to be tractable.

Related to the fragment presented in this paper is a more restricted case in [Dixon *et al.*, 2006]. In that paper, a particular fragment allowing two XOR sets of propositions but where the allowable clauses were further restricted is considered and corresponding resolution calculus given. This restricted fragment is used to represent Büchi Automata so that an the set of clauses representing a Büchi Automaton is satisfiable if and only if the Büchi Automaton is non-empty. The complexity of applying the resolution calculus is polynomial. Further, it is easy to show that every resolvent within the calculus of [Dixon *et al.*, 2006] can be derived by applying resolution rules (possibly several times) from the resolution calculus proposed in this paper restricted to two XOR sets and no additional propositions.

References

- [Ballarini *et al.*, 2006] P. Ballarini, M. Fisher, and M. Wooldridge. Automated Game Analysis via Probabilistic Model Checking: a case study. *Electronic Notes in Theoretical Computer Science*, 149(2):125–137, 2006.
- [Clarke *et al.*, 1999] E.M. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, December 1999.
- [Degtyarev *et al.*, 2002] A. Degtyarev, M. Fisher, and B. Konev. A Simplified Clausal Resolution Procedure for Propositional Linear-Time Temporal Logic. In *Proc. TABLEAUX-02, LNCS* vol. 2381, pages 85–99. Springer-Verlag, 2002.
- [Degtyarev *et al.*, 2006] A. Degtyarev, M. Fisher, and B. Konev. Monodic Temporal Resolution. *ACM Transactions on Computational Logic*, 7(1), January 2006.
- [Demri and Schnoebelen, 2002] S. Demri and P. Schnoebelen. The Complexity of Propositional Linear Temporal Logic in Simple Cases. *Information and Computation*, 174(1):84–103, 2002.
- [Dixon *et al.*, 2006] C. Dixon, M. Fisher, and B. Konev. Is There a Future for Deductive Temporal Verification? In *Proc. TIME-06*. IEEE Computer Society Press, 2006.
- [Fisher *et al.*, 2001] M. Fisher, C. Dixon, and M. Peim. Clausal Temporal Resolution. *ACM Transactions on Computational Logic*, 2(1):12–56, January 2001.
- [Gabbay *et al.*, 1980] D. Gabbay, A. Pnueli, S. Shelah, and J. Stavi. The Temporal Analysis of Fairness. In *Proc. POPL-80*, pages 163–173, January 1980.
- [Hustadt and Konev, 2003] U. Hustadt and B. Konev. TRP++ 2.0: A Temporal Resolution Prover. In *Proc. CADE-19, LNAI* vol. 2741, pages 274–278. Springer, 2003.
- [Hustadt *et al.*, 2004] U. Hustadt, B. Konev, A. Riazanov, and A. Voronkov. TeMP: A Temporal Monodic Prover. In *Proc. IJCAR-04, LNAI* vol. 3097, pages 326–330. Springer, 2004.
- [Schaefer, 1978] T. J. Schaefer. The Complexity of Satisfiability Problems. In *Proc. STOC-78*, pages 216–226, 1978.
- [Sistla and Clarke, 1985] A. P. Sistla and E. M. Clarke. Complexity of Propositional Linear Temporal Logics. *Journal of the ACM*, 32(3):733–749, July 1985.

¹Note that a different model of concurrency might also require modification in the *agreement* clauses.