Characterising Finite Domains in Monodic First-Order Temporal Logic

Boris Konev
Department of Computer Science
Liverpool University
Liverpool, UK
Konev@liverpool.ac.uk

Anatoli Degtyarev
Department of Computer Science
Kings College London
London, UK
Anatoli.Degtiarev@kcl.ac.uk

Michael Fisher
Department of Computer Science
Liverpool University
Liverpool, UK
MFisher@liverpool.ac.uk

Alexei Lisitsa
Department of Computer Science
Liverpool University
Liverpool, UK
A.Lisitsa@liverpool.ac.uk

Abstract

In this paper we introduce a syntactic characterisation of finite domains in first-order temporal logics. In addition to showing that this characterisation is complete with respect to finite domains, we show that the formulae are still within a decidable fragment of first-order temporal logic. This allows us to automatically verify certain finite properties of temporal specifications simply by adding the characteristic formula to the specification and carrying out verification as usual.

1 Introduction

First-order languages allow statements about a particular domain of reference. Since these domains can be arbitrary, such languages are widely used in the description and analysis of complex situations. However, the arbitrary nature of the domain often brings problems. To avoid this, we might fix our domain of discourse — but this leads to inflexibility. Thus, we would like to constrain our domain but maintain its flexibility.

A popular (and viable) solution to this problem is to use modal operators to describe axioms which effectively limit the domain of discourse. The best know of these are the Barcan Formula (BF)

$$\forall x\ \Box p(x)\ \Rightarrow\ \Box\forall x\ p(x)$$

and the Converse Barcan Formula (CBF) [1]

$$\Box\forall x\ p(x)\ \Rightarrow\ \forall x\ \Box p(x)\,.$$

These axioms effectively provide some limited second-order expressive power constraining the domains to be either (non-strictly) decreasing in size (BF) or (non-strictly) increasing in size (CBF).

Discrete temporal logic allows us to go further [13, 12]. As can be seen from the fact that a form of arithmetic induction can be encoded in first-order discrete, linear temporal logic, i.e.

$$[(\exists x\ q(x))\ \wedge\ \Box(\forall y\ q(y)\Rightarrow\bigcirc q(y))]\ \Rightarrow\ [\exists z\ \Box q(z)]$$

then the combination of '$\bigcirc$' and '$\Box$' gives us quite significant expressive power.

In this paper we aim to use first-order discrete, linear temporal logic (FOTL) in order to characterise a further important property of domains in first-order languages. Namely, we wish to characterise the statement *"the domain, D, is of finite size"*. The full FOTL language is very powerful (the set of valid statements is not recursively enumerable) and so one can easily define a form of finiteness inside of this logic. Let ψ be the conjunction of the following formulae[1]:

- $\forall x.\ \neg F(x)$
- $\Box(\forall x.\ F(x)\rightarrow\bigcirc F(x))$
- $\Box(\forall x,y.\ ((\neg F(x)\wedge\neg F(y)\wedge\bigcirc F(x)\wedge\bigcirc\ \rightarrow x=y)$
- $\Diamond\Box(\forall y.\ \bigcirc F(y)\rightarrow F(y))$

It is easy to see that for any temporal model of ψ, there exists a bound N such that the cardinality of the interpretation $F^{\mathcal{I}_n}$ of the predicate F in the state (moment of time) n is

[1] These formulae are taken almost literally from [15], where these have been used to show non-axiomatisability of fragments of FOTL.

bounded by N. Now, given any FOTL formula ϕ, one can *relativise* quantifiers to F in such a way that ϕ is finitely satisfiable if, and only if, its relativisation if satisfiable in general. The applicability of this reduction is limited. It reduces a non-enumerable logic to a non-enumerable logic, and so no complete deductive verification method can be developed using this approach.

FOTL can be considered as a combination of classical first-order logic and propositional temporal logic, in other words classical first order logic is enriched by temporal operators. We can classify forms of the combination depending on which variety of first-order formulae are placed under temporal operators. The simplest combination is the case when only closed first-order formulae under temporal operators are allowed. We call such temporal formulae grounded. The monodic fragment [8, 15] is the case when formulae with at most one free variable are allowed under temporal operators.

We intend to provide a logical characterisation of the finiteness of a domain, but we aim to do so in a complete monodic fragment of first-order temporal logic. Notice that the reduction above destroys monodicity (and, hence, completeness) and cannot be used. We demonstrate that the addition of *finiteness conditions* leads to a complete characterisation of finiteness within monodic first-order temporal logic. Moreover, we show that w.r.t. finite satisfiability the monodic fragment is reduced to the grounded fragment. This will give us a way to automatically assess the truth of formulae when their domain of discourse is restricted to be finite using proof tools for the monodic fragment [9]. This then has many applications, typically in formal verification where the correctness of a system can be assessed in terms of finite data, finite storage, or finite processes. A particular example we have considered previously in [7] uses elements of the domain to represent (distributed) processes. The important requirement of fault-tolerance is that some behaviour can be shown to remain correct even under a *finite* number of processor failures. It was in this earlier paper that we first suggested that first-order temporal logic might be an appropriate tool for describing such finiteness. However, in [7], we provided only suggestions and included no completeness or complexity proofs. In the current paper we prove completeness of this deductive approach to verification.

Although decidability of monodic fragments holds also for the case of semantics where only temporal structures over *finite domains* are allowed [8], the proof is model-theoretic and no practical procedure has been developed.

2 Preliminaries

The language of FOTL is an extension of classical first-order logic by temporal operators for a discrete linear model of time (isomorphic to $\mathbb{N}$, that is, the most commonly used model of time). The signature of FOTL (without equality and function symbols) consists of a countably infinite set of *variables* x_0, x_1, ..., a countably infinite set of *constants* $c_0, c_1, \ldots$, a non-empty set of *predicate symbols* $P, P_0, \ldots$, each with a fixed arity ≥ 0, the *propositional operators* $\top$, $\neg$, $\vee$, the *quantifiers* $\exists x_i$ and $\forall x_i$, and the *temporal operators* $\Box$ ('always in the future'), $\Diamond$ ('eventually in the future'), $\bigcirc$ ('at the next moment'), and $\mathcal{U}$ ('until'). The set of formulae of FOTL is defined as follows: $\top$ is a FOTL formula; if P is an n-ary predicate symbol and t_1, ..., t_n are variables or constants, then $P(t_1, \ldots, t_n)$ is an *atomic* FOTL formula; if φ and ψ are FOTL formulae, then so are $\neg\varphi$, $\varphi \vee \psi$, $\exists x\varphi$, $\forall x\varphi$, $\Box\varphi$, $\Diamond\varphi$, $\bigcirc\varphi$, and $\varphi\,\mathcal{U}\,\psi$. We also use $\bot$, $\wedge$, and $\Rightarrow$ as additional operators, defined using $\top$, $\neg$, and $\vee$. Free and bound variables of a formula are defined in the standard way, as well as the notions of open and closed formulae. Given a formula φ, we write $\varphi(x_1, \ldots, x_n)$ to indicate that all the free variables of φ are among x_1, ..., x_n. As usual, a *literal* is either an atomic formula or its negation.

Formulae of this logic are interpreted over structures $\mathfrak{M} = (D, I_n)_{n\in\mathbb{N}}$ that associate with each element n of $\mathbb{N}$, representing a moment in time, a first-order structure $\mathfrak{M}_n = (D, I_n)$, where D is a non-empty domain and I_n is an interpretation. An *assignment* $\mathfrak{a}$ is a function from the set of variables to D. The application of an assignment to terms is defined in the standard way, in particular, $\mathfrak{a}(c) = c$ for every constant c. The *truth relation* $\mathfrak{M}_n \models^{\mathfrak{a}} \varphi$ is defined in Fig. 2.

In this paper we assume that the interpretation of constants is *rigid*, that is, $I_n(c) = I_m(c)$ for all $n, m \in \mathbb{N}$.

The set of valid formulae of this logic is not recursively enumerable. However, the set of valid *monodic* formulae is known to be finitely axiomatisable [15]. A formula φ of FOTL is called *monodic* if any subformula of φ of the form $\bigcirc\psi$, $\Box\psi$, $\Diamond\psi$, or $\psi_1\,\mathcal{U}\,\psi_2$ contains at most one free variable. For example, the formulae $\forall x\,\Box\exists y P(x, y)$ and $\forall x\,\Box P(x, c)$ are monodic, while $\forall x\forall y(P(x, y) \Rightarrow \Box P(x, y))$ is not monodic.

Every monodic temporal formula can be transformed into an equi-satisfiable normal form, called *divided separated normal form (DSNF)* [10].

Definition 1 *A* monodic temporal problem P in DSNF *is a quadruple* $\langle \mathcal{U}, \mathcal{I}, \mathcal{S}, \mathcal{E}\rangle$*, where*

1. *the step part* $\mathcal{S}$ *is a finite set of* ground *step clauses of the form* $p \Rightarrow \bigcirc q$*, where* p *and* q *are propositions, and* non-ground *step clauses* $P(x) \Rightarrow \bigcirc Q(x)$*, where* P *and* Q *are unary predicate symbols and* x *is a variable;*
2. *the universal part* $\mathcal{U}$ *and the initial part* $\mathcal{I}$ *are finite sets of first-order formulae; and*

$\mathfrak{M}_n \models^{\mathfrak{a}} \top$
$\mathfrak{M}_n \models^{\mathfrak{a}} P(t_1, \ldots, t_n)$ iff $(I_n(\mathfrak{a}(t_1)), \ldots, I_n(\mathfrak{a}(t_n))) \in I_n(P)$
$\mathfrak{M}_n \models^{\mathfrak{a}} \neg\varphi$ iff not $\mathfrak{M}_n \models^{\mathfrak{a}} \varphi$
$\mathfrak{M}_n \models^{\mathfrak{a}} \varphi \vee \psi$ iff $\mathfrak{M}_n \models^{\mathfrak{a}} \varphi$ or $\mathfrak{M}_n \models^{\mathfrak{a}} \psi$
$\mathfrak{M}_n \models^{\mathfrak{a}} \exists x\varphi$ iff $\mathfrak{M}_n \models^{\mathfrak{b}} \varphi$ for some assignment $\mathfrak{b}$ that may differ from $\mathfrak{a}$ only in x
$\mathfrak{M}_n \models^{\mathfrak{a}} \forall x\varphi$ iff $\mathfrak{M}_n \models^{\mathfrak{b}} \varphi$ for every assignment $\mathfrak{b}$ that may differ from $\mathfrak{a}$ only in x
$\mathfrak{M}_n \models^{\mathfrak{a}} \bigcirc\varphi$ iff $\mathfrak{M}_{n+1} \models^{\mathfrak{a}} \varphi$
$\mathfrak{M}_n \models^{\mathfrak{a}} \Diamond\varphi$ iff there exists $m \geq n$ such that $\mathfrak{M}_m \models^{\mathfrak{a}} \varphi$
$\mathfrak{M}_n \models^{\mathfrak{a}} \Box\varphi$ iff for all $m \geq n$, $\mathfrak{M}_m \models^{\mathfrak{a}} \varphi$
$\mathfrak{M}_n \models^{\mathfrak{a}} \varphi \mathcal{U} \psi$ iff there exists $m \geq n$ such that $\mathfrak{M}_m \models^{\mathfrak{a}} \psi$ and $\mathfrak{M}_i \models^{\mathfrak{a}} \varphi$ for every $i, n \leq i < m$

.

Figure 1. Definition of the truth relation $\mathfrak{M}_n \models^{\mathfrak{a}} \varphi$.

3. *the eventuality part $\mathcal{E}$ is a finite set of formulae of the form $\Diamond L(x)$ (a* non-ground *eventuality clause) and $\Diamond l$ (a* ground eventuality *clause), where l is a propositional literal and $L(x)$ is a unary non-ground literal with variable x as its only argument.*

With each monodic temporal problem $\langle \mathcal{U}, \mathcal{I}, \mathcal{S}, \mathcal{E} \rangle$ we associate the FOTL formula $\mathcal{I} \wedge \Box\mathcal{U} \wedge \Box\forall x\mathcal{S} \wedge \Box\forall x\mathcal{E}$. When we talk about particular properties of a temporal problem (e.g., satisfiability, validity, logical consequences, etc) we refer to properties of this associated formula.

Theorem 1 (see [3], Theorem 1) *Any monodic first-order temporal formula can be transformed into an equisatisfiable monodic temporal problem in DSNF with at most a linear increase in the size of the problem. The transformation also preserves satisfiability over finite domains.*

Definition 2 *Let $\mathsf{P} = \langle \mathcal{I}, \mathcal{U}, \mathcal{S}, \mathcal{E} \rangle$ be a monodic temporal problem. By $\mathsf{Pred}(\mathsf{P})$ we denote the set of all (unary) predicates occurring to $\mathcal{S} \cup \mathcal{E}$. By $\mathsf{Prop}(\mathsf{P})$ we denote the set of all propositional symbols occurring to $\mathcal{S} \cup \mathcal{E}$.*

3 Characterisation of Finite Satisfiability

In this section we explore properties of temporal problems and prove that satisfiability over finite domains can be reduced to general satisfiability. We start with separating two possible causes of the domain being infinite: the first-order and inductive. Simply putting a first-order formula without a finite domain model into the universal part $\mathcal{U}$ will, clearly, make any temporal problem unsatisfiable over finite domains. On the other hand, the temporal formula $\Box\exists x P(x) \wedge \Box\forall x(P(x) \Rightarrow \bigcirc\Box\neg P(x))$ does not have a finite domain model because of temporal restrictions. We show that monodic satisfiability over finite domains can be reduced to monodic satisfiability over arbitrary domains provided that the first-order component of a monodic temporal problem lies in a class of formulae having the finite model property.

Definition 3 (Finiteness condition) *Let P be a monodic temporal problem. Let $\mathcal{C}$ be the set of constants occurring in P. Let $\mathcal{F}(x)$ be a first-order formula containing at most one free variable x such that $\mathcal{F}(x)$ is built using propositions from $\mathsf{Prop}(\mathsf{P})$, predicate symbols from $\mathsf{Pred}(\mathsf{P})$, and constants from $\mathcal{C}$. The* finiteness condition *for $\mathcal{F}(x)$ is the following formula*

$$\mathsf{Fin}_{\mathcal{F}} = \Box\,(\forall x(\mathcal{F}(x) \Rightarrow \bigcirc\Box\neg\mathcal{F}(x))) \Rightarrow \Diamond\Box\forall x\neg(\mathcal{F}(x)).$$

Intuitively, a finiteness condition states that if the property $\mathcal{F}(x)$ is such that when it holds on a domain element, it will never hold on the same element in the future, and the domain is finite, eventually, there will be no element with property $\mathcal{F}$.

Definition 4 *We say that a class of* first-order *formulae has the property $\mathfrak{F}$ if it has the finite model property, is closed under conjunctions, and contains all monadic formulae. A class of first-order* temporal *monodic formulae is said to have the property $\mathfrak{T}\mathfrak{F}$ if for any its formula φ and its DSNF translation $\mathsf{P}_\varphi = \langle \mathcal{I}, \mathcal{U}, \mathcal{S}, \mathcal{E} \rangle$, $\mathcal{I} \cup \mathcal{U}$ belongs to a class of first-order formulae having the property $\mathfrak{F}$.*

Theorem 2 *Let P be a monodic temporal problem such that $\mathcal{U} \cup \mathcal{I}$ belong to a class of first-order formulae, which has the property $\mathfrak{F}$. Then*

1. *If P is finitely satisfiable, for any conjunction of finiteness conditions $\bigwedge_{\mathcal{F}_i} \mathsf{Fin}_{\mathcal{F}_i}$ the formula $(\mathsf{P} \wedge \bigwedge_{\mathcal{F}_i} \mathsf{Fin}_{\mathcal{F}_i})$ is finitely satisfiable.*
2. *If P does not have a model with a finite domain then there exist $\mathcal{F}_1, \ldots, \mathcal{F}_M$ such that $(\mathsf{P} \wedge \bigwedge_{i=1}^{M} \mathsf{Fin}_{\mathcal{F}_i})$ is (not necessarily finitely) unsatisfiable.*

Proof Proved in Section 7. ❑

It follows from the proof of Theorem 2 that it suffices to consider finiteness conditions of a specific syntactic form. It is easy to see that it is possible to enumerate all finiteness conditions of this syntactic form.

Corollary 3 *Let P be a monodic temporal problem satisfying the conditions of Theorem 2. Let $\mathsf{Fin}_{\mathcal{F}_1}, \ldots, \mathsf{Fin}_{\mathcal{F}_N}$ be the set of* all *finiteness conditions of the syntactic form from the proof of Theorem 2. Then the following holds: P does not have a model with a finite domain if, and only if, $(\mathsf{P} \wedge \bigwedge_{j=1}^{N} \mathsf{Fin}_{\mathcal{F}_j})$ is (not necessarily finitely) unsatisfiable.*

In many practical cases, it suffices, however, to consider a small number of finiteness conditions.

Example 1 *Consider the following temporal problem*

$$\mathcal{I} = \{a\} \qquad \mathcal{E} = \emptyset$$
$$\mathcal{U} = \left\{ \begin{array}{l} a \Rightarrow \exists x P(x) \\ \forall x \neg(P(x) \wedge Q(x)) \end{array} \right\}$$
$$\mathcal{S} = \left\{ \begin{array}{rcl} (P(x) \wedge a) & \Rightarrow & \bigcirc (P(x) \wedge \neg a) \\ (Q(x) \wedge a) & \Rightarrow & \bigcirc (P(x) \wedge \neg a) \\ (P(x) \wedge \neg a) & \Rightarrow & \bigcirc (Q(x) \wedge a) \end{array} \right\}$$

It can be seen that P has only infinite models. Consider

$$\mathsf{Fin} = \big(\Box(\forall x((P(x) \wedge a) \Rightarrow \bigcirc \Box \neg(P(x) \wedge a))) \Rightarrow \Diamond \Box (\forall x \neg(P(x) \wedge a))\big).$$

It can be seen that $\mathsf{P} \wedge \mathsf{Fin}$ is (generally) unsatisfiable.

We formulate Theorem 2 for monodic temporal problem since it is easier to impose restrictions on the first-order component of the problem and because it is possible to explicitly enumerate all possible finiteness conditions for a problem. We can extend this result to formulae not in DSNF using *general finiteness condition*. Let $\mathcal{P}(x)$ be any monodic temporal formula with at most one free variable x. Then

$$\Box (\forall x(\mathcal{P}(x) \Rightarrow \bigcirc \Box \neg \mathcal{P}(x))) \Rightarrow \Diamond \Box \forall x \neg(\mathcal{P}(x))$$

is a general finiteness condition for $\mathcal{P}(x)$. Notice that $\mathcal{P}(x)$ may include temporal operators.

Corollary 4 *Let φ belongs to a class of temporal monodic formulae which has the property $\mathfrak{TF}$. Then if φ does not have a model with a finite domain then there exist general finiteness conditions $G_1, \ldots, G_M$ such that $(\varphi \wedge \bigwedge_{i=1}^{M} G_i)$ is (not necessarily finitely) unsatisfiable.*

Proof (Sketch) Let P_φ be the DSNF translation of φ, and $\mathcal{F}_1, \ldots, \mathcal{F}_M$ such that $(\mathsf{P}_\varphi \wedge \bigwedge_{i=1}^{M} \mathsf{Fin}_{\mathcal{F}_i})$ is unsatisfiable. We show that there exist $\mathcal{P}_1, \ldots, \mathcal{P}_M$ such that $(\varphi \wedge \bigwedge_{i=1}^{M} G_i)$ is also unsatisfiable, where G_i is the general finiteness condition for $\mathcal{P}_i(x)$. The proof relies on properties of translation of monodic formulae to DSNF [10]. The reduction is based on using a renaming technique to substitute non-atomic subformulae and replacing temporal operators by their fixed point definitions described e.g. in [5]. Thus, whenever $\mathcal{F}_i(x)$ contains a predicate or proposition symbol used to rename ψ, a subformula of φ, in the same position $\mathcal{P}(x)$ will have ψ itself. Predicate and propositional symbols introduced in fixed point definitions also correspond to combinations of temporal subformulae of φ. We only consider here the case of unwinding the eventuality; it can be seen that all other cases from [10] also go through. Suppose a subformula of the form $\Diamond L(x)$ was renamed and a new predicate name $P(x)$ was introduced. Then the following definition is added to the formula

$$\Box \forall x(P(x) \Rightarrow \Diamond L(x))$$

This expression is translated to DSNF with the help of a new predicate symbol $waitforL(x)$.

$$\Box \forall x \big(((P(x) \wedge \neg L(x)) \implies waitforL(x)) \wedge ((waitforL(x) \wedge \bigcirc \neg L(x)) \implies \bigcirc waitforL(x)) \wedge \Diamond \neg waitforL(x) \big)$$

It can be seen that in any model of P_φ we have $waitforL(x)$ is true on an element x if, and only if, $\neg L(x) \wedge \Diamond L(x)$ is true on x. Therefore, if $\mathcal{F}_i(x)$ contains $waitforL(x)$, in the same position $\mathcal{P}(x)$ will have $\neg L(x) \wedge \Diamond L(x)$. ❑

4 Axiomatisation of monodic logic over finite domains

Theorem 2 and Corollary 4 show how to reduce satisfiability of monodic formulae over finite domains to satisfiability over all domains. From that one can derive results on axiomatisation of (fragments of) monodic temporal logic over finite domains. For general semantics, i.e of not necessarily finite domains the finite axiomatisation of that logic is given in [15]. Notice that due to the fact that first-order predicate logic over finite domains is not recursively axiomatisable [14], there can be no recursive axiomatisation of monodic temporal logic over finite domains. We show here that by restricting the language to a class of formulae having the property $\mathfrak{TF}$, and adding to the axiomatisation from [15] the finiteness principles as an axiom scheme we get a complete axiomatisation over finite domains.

Hilbert-style axiomatisation Ax_{mon} of monodic temporal logic [15] includes the usual axiom schemata and inference rules for classical first-order logic extended by two temporal inference rules:

$$\frac{\varphi}{\bigcirc \varphi} \qquad \text{and} \qquad \frac{\chi \rightarrow \neg\psi \wedge \bigcirc \chi}{\chi \rightarrow \neg(\varphi \,\mathcal{U}\, \psi)}.$$

Theorem 5 *Let Ax^{fin}_{mon} be the above axiomatisation Ax_{mon} extended by the finiteness axiom scheme*

$$\Box (\forall x(\varphi(x) \Rightarrow \bigcirc \Box \neg \varphi(x))) \Rightarrow \Diamond \Box \forall x \neg(\varphi(x)).$$

For any formula φ from a class of monodic temporal formulae, which has the property $\mathfrak{TF}$ we have that φ is valid over finite domains if and only if φ is derivable in Ax^{fin}_{mon}.

Proof Soundness is straightforward. To prove completeness assume that a formula φ is valid over finite domains. Then $\neg\varphi$ is unsatisfiable over finite domains. By Corollary 4 there exists general finiteness conditions $G_1,\ldots,G_M$ such that $(\neg\varphi \wedge \bigwedge_{i=1}^{M} G_i)$ is unsatisfiable (over general domains). It follows that $\bigwedge_{i=1}^{M} G_i \rightarrow \varphi$ is valid over all domains. By completeness of Ax_{mon} the last formula is derivable in Ax_{mon}. Using finiteness axiom scheme and modus ponens the formula φ is derivable in Ax^{fin}_{mon}. ❑

Note that Theorem 2 can be also seen as a variant of Theorem 5 that specifies explicitly which instances of the finite axiom scheme are needed for proving (see also Corollary 3).

5 Ground Temporal Problems

We show now that monodic temporal problem satisfying conditions of Theorem 2 can be further reduced, preserving satisfiability over finite domains, to a monodic temporal problem in which the step and eventuality parts only contain propositions.

Definition 5 *A temporal problem P is called* ground *if all the step clauses and the eventuality clauses of P are ground. A temporal problem P is called a* ground eventuality *problem if all the eventualities of P are ground.*

It is known that if the eventuality part of a temporal problem is ground, the problem can be reduced to a ground problem.

Theorem 6 ([3]) *Every ground eventuality monodic temporal problem can be reduced to a satisfiability equivalent ground monodic problem with an exponential growth in size of the given problem.*

We now consider finite satisfiability.

Theorem 7 (grounding eventualities over finite domains) *Any monodic temporal problem P can be reduced to a ground-eventuality monodic temporal problem P' such that P is finitely satisfiable if, and only if, P' is finitely satisfiable. The size of P' is is linear in the size of P.*

Proof [Theorem 7] The reduction is an immediate consequence of the two following lemmas. We use the (past-time) operator $\mathcal{S}$ whose semantics is defined as follows

$$\begin{array}{l}\mathfrak{M}_n \models^{\mathfrak{a}} \phi\,\mathcal{S}\,\psi \text{ iff there exists a } k \in \mathbb{N},\\ \qquad \text{such that } 0 \leqslant k < n \text{ and } \mathfrak{M}_k \models^{\mathfrak{a}} \psi\\ \qquad \text{and, for all } j \in \mathbb{N},\ \text{if } k \leqslant j < n \text{ then } \mathfrak{M}_j \models^{\mathfrak{a}} \phi\end{array}$$

and which is then eliminated by Lemma 9, following the approach in [4]. ❑

Lemma 8 *The formula $\Phi \wedge \Box\forall x \Diamond L(x)$, where Φ is an arbitrary temporal formula, is satisfiability equivalent* over finite domains *to*

$$\Phi \wedge l \wedge \Box(l \Rightarrow \Diamond\bigcirc(\forall x(\neg l\,\mathcal{S}\,L(x)) \wedge l)),$$

where l is a new propositional symbol.

See Appendix for proof of Lemma 8.

Lemma 9 *Formula $\Box\forall x(A(x) \Rightarrow B(x)\,\mathcal{S}\,C(x))$ is satisfiability equivalent to*

$$\left\{\begin{array}{c}\forall x \neg P(x) \quad \wedge\\ \Box\forall x(A(x) \Rightarrow P(x)) \quad \wedge\\ \Box\forall x(C(x) \vee (B(x) \wedge P(x))\end{array}\right\} \equiv \bigcirc P(x)),$$

where P is a new predicate symbol.

Proof [Lemma 9] Follows straightforwardly from consideration of possible models [4]. ❑

Theorem 10 *Every monodic temporal problem $\mathsf{P} = \langle \mathcal{I}, \mathcal{U}, \mathcal{S}, \mathcal{E}\rangle$ such that $\mathcal{I} \cup \mathcal{U}$ belongs to a class of first-order formulae, which has the property $\mathfrak{F}$ can be reduced to a ground monodic problem $\mathsf{P}^{\mathsf{G,Fin}}$ such that P is finitely satisfiable if, and only if, $\mathsf{P}^{\mathsf{G,Fin}}$ is satisfiable.*

Proof Let $\mathsf{P}^{\mathsf{Fin}}$ be the monodic temporal problem guaranteed by Theorem 2 for P and $\mathsf{P}^{\mathsf{G,Fin}}$ the ground monodic temporal problem guaranteed by Theorems 7 and 6 for $\mathsf{P}^{\mathsf{Fin}}$. Suppose P is finitely satisfiable. Then, by Theorem 2, $\mathsf{P}^{\mathsf{Fin}}$ is finitely satisfiable, and so $\mathsf{P}^{\mathsf{G,Fin}}$ is finitely satisfiable. Conversely, suppose $\mathsf{P}^{\mathsf{G,Fin}}$ is satisfiable. Then, it can be seen that $\mathsf{P}^{\mathsf{Fin}}$ is satisfiable and, by Theorem 2, P is finitely satisfiable. ❑

6 Applications to Formal Verification

In [7] the notion of a general finiteness condition (Definition 3) has been introduced under the name of the finite clock axiom which is, in fact, a scheme of axioms. The finite clock axiom was used as an additional proof principle in deductive verification of *parametrised infinite state systems* comprising *arbitrary* numbers of identical processes. Such systems has become increasingly important [2]. Practical problems of an open, distributed nature often fit into this model, for example robot swarms of arbitrary sizes.

When modelling parametrised systems in temporal logic, informally, elements of the domain correspond to processes, and predicates to states of such processes [6]. For example $idle(x)$ means that a process x is in the idle state, $\Diamond\forall y.\ agreement(y)$ means that, eventually, all processes will be in agreement, while $\exists z.\ \Box inactive(z)$ means that there is at least one process that is always inactive. (See [6]

for further details.) For many protocols, especially when fault tolerance is concerned, it is essential that the number of processes is finite. The straightforward generalisation to infinite numbers of processes makes many protocols incorrect.

It was noticed in [7] that every instance of the finite clock axiom scheme holds in every temporal model with a finite domain, and that in many cases properties of parametrised infinite systems can be proved with the help of the axiom. This was demonstrated in a case study of the FloodSet algorithm for the Consensus problem.

The setting is as follows: There are n processes, each having an *input bit* and an *output bit*. The processes work synchronously, run the same algorithm and use *broadcasting* for communication. Any message sent by a non-faulty process is instantaneously delivered to all other processes. Some processes may fail and, from that point onward, such processes do not send any 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 processes. 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 non-faulty processes. It is required also that if all processes have the same input bit, that bit should be produced as an output bit.

In [7] it has been shown that a variant of the *FloodSet algorithm with alternative decision rule* [11], designed to solve the Consensus problem in presence of crash failures, can be specified (naturally) within monodic monadic temporal logic without equality. This variant of the FloodSet algorithm operates as follows.

- In the first round of computations, every process broadcasts its input bit.
- In every later round, a process broadcasts any value *the first time it sees it*.
- In every round the (tentative) output bit is set to the minimum value seen so far.

The correctness criterion for this algorithm is that, eventually, the output bits of all non-faulty processes will be the same. It is crucial for the correctness of the algorithm that only a finite number of processes can fail. This can be captured by an instance of the finite clock axiom,

$$\Box\,(\forall x(Fail(x) \Rightarrow \bigcirc\,\Box\neg Fail(x))) \Rightarrow \Diamond\,\Box\forall x\neg(Fail(x)),$$

where the predicate $Fail(x)$ is true if, and only if, the process x fails. [7] gives a formal proof of the correctness of the FloodSet algorithm.

It was an empirical observation in [7] that finiteness conditions help in proving properties of parametrised infinite state systems. Theorem 5 stipulates completeness of the deductive approach to verification given in [7].

In [7] we have also considered a family of protocols which terminate after a certain (but unknown) number of steps. An example of such a protocol is any protocol, where every process sends only a finite number of messages. It was shown that the monodic monadic specification of an eventually stable protocol is satisfiable over finite domains if, and only if, it is satisfiable over arbitrary domains.

The *stabilisation principle* for a temporal problem P is the formula:

$$\mathsf{Stab} = \Diamond\,\Box(\forall x \bigwedge_{P\in\mathsf{Pred}(\mathsf{P})} [P(x) \equiv \bigcirc P(x)]).$$

Proposition 11 ([7]) *Let P be a monodic monadic temporal problem. Then $\mathsf{P} \wedge \mathsf{Stab}$ is satisfiable in a model with a finite domain if, and only if, $\mathsf{P} \wedge \mathsf{Stab}$ is satisfiable in a model with an arbitrary domain.*

We give a new proof of the proposition now. Let $\mathcal{F}(x)$ be any first-order formula containing at most one free variable x such that $\mathcal{F}(x)$ is built using propositions from $\mathsf{Prop}(\mathsf{P})$, predicate symbols from $\mathsf{Pred}(\mathsf{P})$, and constants from $\mathcal{C}$. Notice that whenever $\mathfrak{M} \models \mathsf{Stab}$ we have $\mathfrak{M} \not\models \Box\forall x(\mathcal{F}(x) \Rightarrow \bigcirc\,\Box\neg\mathcal{F}(x))$, and thus $\mathfrak{M} \models \mathsf{Fin}_{\mathcal{F}}$. Therefore, the stabilisation principle implies any finiteness condition, and Proposition 11 is a trivial consequence of Theorem 2.

7 Proof of Theorem 2

Theorem 2 follows from Lemmas 15 and 16 below. First, we introduce additional concepts. Let $\mathsf{P} = \langle \mathcal{U}, \mathcal{I}, \mathcal{S}, \mathcal{E}\rangle$ be a monodic temporal problem. A *predicate colour* γ is a set of unary literals such that for every $P(x) \in \mathsf{Pred}(\mathsf{P})$, either $P(x)$ or $\neg P(x)$ belongs to γ. A *propositional colour* θ is a set of propositional literals such that for every $p \in \mathsf{Prop}(\mathsf{P})$, either p or $\neg p$ belongs to θ. Let Γ be a set of predicate colours, θ be a propositional colour, and ρ be a map from the set of constants, $\mathrm{const}(\mathsf{P})$, to Γ. A triple $\langle \Gamma, \theta, \rho\rangle$ is called a *colour scheme*, and ρ is called a *constant distribution*. We write sometime $\gamma \in \mathcal{C}$ when $\gamma \in \Gamma$ and $\mathcal{C} = \langle \Gamma, \theta, \rho\rangle$.

For every colour scheme $\mathcal{C} = \langle \Gamma, \theta, \rho\rangle$ let us construct the formulae $\mathcal{F}_{\mathcal{C}}$, $\mathcal{A}_{\mathcal{C}}$, $\mathcal{B}_{\mathcal{C}}$ in the following way. For every $\gamma \in \Gamma$ and for every θ, introduce the conjunctions:

$$F_\gamma(x) = \bigwedge_{L(x)\in\gamma} L(x); \qquad F_\theta = \bigwedge_{l\in\theta} l.$$

Let

$$\begin{array}{lcl} A_\gamma(x) & = & \bigwedge\{L(x) \mid L(x) \Rightarrow \bigcirc M(x) \in \mathcal{S},\ L(x) \in \gamma\},\\ B_\gamma(x) & = & \bigwedge\{M(x) \mid L(x) \Rightarrow \bigcirc M(x) \in \mathcal{S},\ L(x) \in \gamma\},\\ A_\theta & = & \bigwedge\{l \mid l \Rightarrow \bigcirc m \in \mathcal{S},\ l \in \theta\},\\ B_\theta & = & \bigwedge\{m \mid l \Rightarrow \bigcirc m \in \mathcal{S},\ l \in \theta\}. \end{array}$$

(Recall that there are no two different step clauses with the same left-hand side.) Now $\mathcal{F}_\mathcal{C}$, $\mathcal{A}_\mathcal{C}$, $\mathcal{B}_\mathcal{C}$ are of the following forms:

$$\mathcal{F}_\mathcal{C} = \bigwedge_{\gamma\in\Gamma} \exists x F_\gamma(x) \wedge F_\theta \wedge \bigwedge_{c\in \mathrm{const}(P)} F_{\rho(c)}(c) \wedge \forall x \bigvee_{\gamma\in\Gamma} F_\gamma(x),$$
$$\mathcal{A}_\mathcal{C} = \bigwedge_{\gamma\in\Gamma} \exists x A_\gamma(x) \wedge A_\theta \wedge \bigwedge_{c\in \mathrm{const}(P)} A_{\rho(c)}(c) \wedge \forall x \bigvee_{\gamma\in\Gamma} A_\gamma(x),$$
$$\mathcal{B}_\mathcal{C} = \bigwedge_{\gamma\in\Gamma} \exists x B_\gamma(x) \wedge B_\theta \wedge \bigwedge_{c\in \mathrm{const}(P)} B_{\rho(c)}(c) \wedge \forall x \bigvee_{\gamma\in\Gamma} B_\gamma(x).$$

Definition 6 (Behaviour Graph) *Now, given a temporal problem $\mathsf{P} = \langle \mathcal{U}, \mathcal{I}, \mathcal{S}, \mathcal{E}\rangle$ we define a finite directed graph G as follows. Every vertex of G is a colour scheme $\mathcal{C}$ for P such that $\mathcal{U} \cup \mathcal{F}_\mathcal{C}$ is satisfiable. For each vertex $\mathcal{C} = \langle \Gamma, \theta, \rho\rangle$, there is an edge in G to $\mathcal{C}' = \langle \Gamma', \theta', \rho'\rangle$, if $\mathcal{U} \wedge \mathcal{F}_{\mathcal{C}'} \wedge \mathcal{B}_\mathcal{C}$ is satisfiable. They are the only edges originating from $\mathcal{C}$. A vertex $\mathcal{C}$ is designated as an* initial *vertex of G if $\mathcal{I} \wedge \mathcal{U} \wedge \mathcal{F}_\mathcal{C}$ is satisfiable. The* behaviour graph *H of P is the sub-graph of G induced by the set of all vertices reachable from the initial vertices.*

Definition 7 (Path; Path Segment) *A* path*, π, through a behaviour graph, H, is a function from $\mathbb{N}$ to the vertices of the graph such that for any $i \geq 0$ there is an edge $\langle \pi(i), \pi(i+1)\rangle$ in H. In the similar way, we define a* path segment *as a function from $[m, n]$, $m < n$, to the vertices of H with the same property.*

Lemma 12 ([3]) *Let $\mathsf{P}_1 = \langle \mathcal{U}_1, \mathcal{I}, \mathcal{S}, \mathcal{E}\rangle$ and $\mathsf{P}_2 = \langle \mathcal{U}_2, \mathcal{I}, \mathcal{S}, \mathcal{E}\rangle$ be two problems such that $\mathcal{U}_1 \subseteq \mathcal{U}_2$. Then the behaviour graph of P_2 is a sub-graph of the behaviour graph of P_1.*

Definition 8 (Suitability) *For $\mathcal{C} = (\Gamma, \theta, \rho)$ and $\mathcal{C}' = (\Gamma', \theta', \rho')$, let $(\mathcal{C}, \mathcal{C}')$ be an ordered pair of colour schemes for a temporal problem P. An ordered pair of predicate colours (γ, γ') where $\gamma \in \Gamma$, $\gamma' \in \Gamma'$ is called* suitable *if the formula $\mathcal{U} \wedge \exists x(F_{\gamma'}(x) \wedge B_\gamma(x))$ is satisfiable. Similarly, an ordered pair of propositional colours (θ, θ') is suitable if $\mathcal{U} \wedge F_{\theta'} \wedge B_\theta$ is satisfiable, and an ordered pair of constant distributions (ρ, ρ') is suitable if, for every $c \in C$, the pair $(\rho(c), \rho'(c))$ is suitable.*

Note that the satisfiability of $\exists x(F_{\gamma'}(x) \wedge B_\gamma(x))$ implies $\models \forall x(F_{\gamma'}(x) \Rightarrow B_\gamma(x))$ as the conjunction $F_{\gamma'}(x)$ contains a valuation at x of *all* predicates occurring in $B_\gamma(x)$.

Lemma 13 ([3]) *Let H be the behaviour graph for the problem $\mathsf{P} = \langle \mathcal{U}, \mathcal{I}, \mathcal{S}, \mathcal{E}\rangle$ with an edge from a vertex $\mathcal{C} = \langle \Gamma, \theta, \rho\rangle$ to a vertex $\mathcal{C}' = \langle \Gamma', \theta', \rho'\rangle$. Then*

1. *for every $\gamma \in \Gamma$ there exists a $\gamma' \in \Gamma'$ such that the pair (γ, γ') is suitable;*
2. *for every $\gamma' \in \Gamma'$ there exists a $\gamma \in \Gamma$ such that the pair (γ, γ') is suitable;*
3. *the pair of propositional colours (θ, θ') is suitable;*
4. *the pair of constant distributions (ρ, ρ') is suitable.*

Definition 9 (Run/E-Run) *Let π be a path through a behaviour graph H of a temporal problem P, and $\pi(i) = (\Gamma_i, \theta_i, \rho_i)$. By a* run *in π we mean a function $r(n)$ from $\mathbb{N}$ to $\bigcup_{i\in\mathbb{N}} \Gamma_i$ such that for every $n \in \mathbb{N}$, $r(n) \in \Gamma_n$ and the pair $(r(n), r(n+1))$ is suitable. In the similar way, we define a* run segment *as a function from $[m, n]$, $m < n$, to $\bigcup_{i\in\mathbb{N}} \Gamma_i$ with the same property. A run r is called an* e-run *if $\forall i \geq 0 \forall \Diamond L(x) \in \mathcal{E} \exists j > i(L(x) \in r(j))$*[2].
Let π be a path, the set of all runs in π is denoted by $\mathcal{R}(\pi)$, and the set of all e-runs in π is denoted by $\mathcal{R}_e(\pi)$. If π is clear, we may omit it.

Let H be the behaviour graph for a temporal problem $\mathsf{P} = \langle \mathcal{U}, \mathcal{I}, \mathcal{S}, \mathcal{E}\rangle$ and $\pi = \mathcal{C}_0, \ldots, \mathcal{C}_n, \ldots$ be a path in H where $\mathcal{C}_i = \langle \Gamma_i, \theta_i, \rho_i\rangle$. Let $\mathcal{G}_0 = \mathcal{I} \cup \{\mathcal{F}_{\mathcal{C}_o}\}$ and $\mathcal{G}_n = \mathcal{F}_{\mathcal{C}_n} \wedge \mathcal{B}_{\mathcal{C}_{n-1}}$ for $n \geq 1$. According to the definition of a behaviour graph, the set $\mathcal{U} \cup \{\mathcal{G}_n\}$ is satisfiable for every $n \geq 0$.

The following lemma can be proved similarly to Lemma 27 in [8].

Lemma 14 *Let $\mathcal{U} \cup \mathcal{I}$ belong to a class of first-order formulae, which has the property $\mathfrak{F}$. Then there exists a number N such that for every $k > N$ and every $n \geq 0$, if the set $\mathcal{U} \cup \{\mathcal{G}_n\}$ is satisfiable then there exists a model $\mathfrak{M}_n = \langle D, I_n\rangle$ of $\mathcal{U} \cup \{\mathcal{G}_n\}$ such that for every $\gamma \in \Gamma_n$ the set $D_{(n,\gamma)} = \{a \in D \mid \mathfrak{M}_n \models F_\gamma(a)\}$ is of cardinality k.*

Let π be a path through H consisting of an initial segment and a looping part, $\pi = \pi_1(\pi_2)^*$, for some π_1 and π_2. By $\mathcal{R}_e^f(\pi)$ we denote the set of e-runs such that for every $r \in \mathcal{R}_e^f(\pi)$ and every $i \geq 0$ we have $r(l_1 + i) = r(l_1 + l_2 + i)$, where l_1 and l_2 are lengths of π_1 and π_2, resp.

Lemma 15 *Let P be a monodic temporal problem such that $\mathcal{U} \cup \mathcal{I}$ belong to a class of first-order formulae, which has the property $\mathfrak{F}$. Let H be the behaviour graph for P. Then P is finitely satisfiable if, and only if, there exists a path $\pi = \pi_1(\pi_2)^*$ through H such that the following conditions hold.*

(a) $\pi(0)$ is an initial vertex of H;

(b) for every colour scheme $\mathcal{C} = \pi(i)$, $i \geq 0$, and every ground eventuality literal $\Diamond l \in \mathcal{E}$ there exists a colour scheme $\mathcal{C}' = \pi(j)$, $j > i$, such that $l \in \theta'$;

(c) for every colour scheme $\mathcal{C} = \pi(i)$, $i \geq 0$ and every predicate colour γ from the colour scheme there exists an e-run $r \in \mathcal{R}_e^f(\pi)$ such that $r(i) = \gamma$;

[2] To make the presentation compact, we abuse the notation by allowing the use of logical symbols at meta-level.

(d) for every constant $c \in \mathcal{L}$, the function $r_c(n)$ defined by $r_c(n) = \rho_n(c)$, where ρ_n is the constant distribution from $\pi(n)$, is an e-run in π; and

See Appendix for proof of Lemma 15.

Example 2 *Consider the following temporal problem.*

$$\mathcal{I} = \{\exists x P(x) \wedge \exists y \neg P(y)\}, \quad \mathcal{U} = \emptyset, \quad \mathcal{E} = \{\Diamond P(x)\}$$

$$\mathcal{S} = \left\{ \begin{array}{rcl} (P(x) \wedge a) & \Rightarrow & \bigcirc(\neg P(x) \wedge \neg a) \\ (\neg P(x) \wedge a) & \Rightarrow & \bigcirc(P(x) \wedge \neg a) \\ (P(x) \wedge \neg a) & \Rightarrow & \bigcirc(P(x) \wedge a) \\ (\neg P(x) \wedge \neg a) & \Rightarrow & \bigcirc(\neg P(x) \wedge a) \end{array} \right\}$$

In this example, the behaviour graph contains two nodes, $C_1 = \langle (P, \neg P), a \rangle \qquad C_2 = \langle (P, \neg P), \neg a \rangle$.

Clearly, in this example we have a model with a finite domain, but we have to consider paths segments of the length 4 when cycling $(C_1, C_2, C_1, C_2)^*$ in order to extract a finite domain model.

Lemma 16 *Let P be a monodic temporal problem such that $\mathcal{U} \cup \mathcal{I}$ belong to a class of first-order formulae, which has the property $\mathfrak{F}$. Let $\mathcal{CS}$ be the set of colour schemes, and let* Fin *be the following formula.*

$$\bigwedge_{\mathcal{C} \in \mathcal{CS}} \Big(\Big(\Box\,(\forall x((\mathcal{F}_{\mathcal{C}} \wedge F_\gamma(x)) \Rightarrow \bigcirc \Box \neg(\mathcal{F}_{\mathcal{C}} \wedge F_\gamma(x)))) \Rightarrow \Diamond \Box \forall x \neg(\mathcal{F}_{\mathcal{C}} \wedge F_\gamma(x)) \Big) \Big).$$

Let $(\mathsf{P} \wedge \mathsf{Fin})$ be satisfiable. Then there exists a path through the behaviour graph for P satisfying the conditions of Lemma 15.

See Appendix for proof of Lemma 16.

8 Concluding Remarks

In this paper we have characterised reasoning over finite domains in monodic fragments of first-order temporal logic. It has turned out that intuitive finiteness principles are sufficient to get a complete characterisation. On the one hand, presented results continue the modal logic tradition on capturing natural properties of the domains by appropriate axioms/proof principles, on the other hand they provide a foundation for deductive verification of parametrised systems with finite resources. The future work includes automation of proof search for monodic temporal logic with finiteness principles. The main issue to be addressed here is how to search efficiently for appropriate instances of the finiteness condition. The development of specialised proof systems, e.g the extensions of resolution-based calculi [3] looks as a promising next step. Axiomatisation of reasoning in finite for other fragments of FOTL is also of interest. For non-recursively axiomatisable fragments, including full FOTL one may ask whether suitable finiteness principles would be enough to establish the relative completeness.

References

[1] R. C. Barcan (Marcus). A functional calculus of rst order based on strict implication. *Journal of Symbolic Logic*, 11:1–16, 1946.

[2] M. Calder and A. Miller. An automatic abstraction technique for verifying featured, parameterised systems. *Theoretical Computer Science*, 404(3):235–255, 2008.

[3] A. Degtyarev, M. Fisher, and B. Konev. Monodic temporal resolution. *ACM Transactions on Computational Logic*, 7(1):108–150, 2006.

[4] M. Fisher. A normal form for temporal logics and its applications in theorem proving and execution. *Journal of Logic and Computation*, 7(4):429–456, 1997.

[5] M. Fisher, C. Dixon, and M. Peim. Clausal temporal resolution. *ACM Transactions on Computational Logic*, 2(1):12–56, 2001.

[6] M. Fisher, B. Konev, and A. Lisitsa. Practical infinite-state verification with temporal reasoning. In *Verification of Infinite State Systems and Security*. IOS Press, 2006.

[7] M. Fisher, B. Konev, and A. Lisitsa. Temporal verification of fault-tolerant protocols. In *Methods Models and Tools for Fault Tolerance*, volume 5454, 2009.

[8] I. Hodkinson, F. Wolter, and M. Zakharyaschev. Decidable fragments of first-order temporal logics. *Annals of Pure and Applied Logic*, 106:85–134, 2000.

[9] U. Hustadt, B. Konev, A. Riazanov, and A. Voronkov. **TeMP**: A temporal monodic prover. In *Proceedings IJCAR 2004*, volume 3097 of *LNAI*, pages 326–330. Springer, 2004.

[10] 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.

[11] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufman, 1996.

[12] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems: Specification*. Springer, 1992.

[13] A. Pnueli. The Temporal Logic of Programs. In *Proceedings of the Eighteenth Symposium on the Foundations of Computer Science (FOCS)*, pages 46–57, 1977.

[14] B. Trakhtenbrot. The impossibility of an algorithm for the decision problem for finite models. *Dokl. Akad. Nauk SSSR*, 70:596–572, 1950. English translation in: AMS Transl. Ser. 2, vol.23 (1063), 1–6.

[15] F. Wolter and M. Zakharyaschev. Axiomatizing the monodic fragment of first-order temporal logic. *Annals of Pure and Applied logic*, 118:133–145, 2002.

A Proofs of Lemmas 8, 15, and 16

Proof [Lemma 8] Let us reformulate the given problem in a two-sorted temporal language with variables over $\mathbb{N}$ for the temporal sort, i.e. $\Phi \wedge \forall x \forall n \exists m(n \leq m \wedge L(m,x))$, meaning that for all x, $L(x)$ is satisfied infinitely often. This problem is satisfiability equivalent *over finite domains* to the following (this can be easily checked by considering possible models):

$$\Phi \wedge \forall n \exists m(m > n \wedge \forall x \exists k(n \leq k < m \wedge L(k,x))) \quad (1)$$

which states, informally, that for each moment of time, n, there is a moment $m > n$, such that for all x the eventuality $\Diamond L(x)$ is satisfied "after n and before m" (there are only finitely many elements in a domain).

We prove, that given a model for (1), it is possible to find a model for

$$\Phi \wedge l \wedge \Box(l \Rightarrow \Diamond \bigcirc (\forall x(\neg l \,\mathcal{S}\, L(x)) \wedge l)), \qquad (2)$$

and vice versa. First, consider a model $\mathfrak{M} = \langle D, I\rangle$ for (1). We construct a model $\mathfrak{M}' = \langle D, I'\rangle$ for (2) by extending $\mathfrak{M}$ with a new proposition l and defining its value as follows. Formula (1) states that for each moment of time n there exists a future moment m when a certain property holds, defining thus a moment m for each n. Let us construct a sequence of times defined by (1) starting from 0, i.e. $m_0 = 0$, $m_1 = m(0), \ldots, m_{i+1} = m(m_i)$; and let us define l in $\mathfrak{M}'$ as **True** at those times and as **False** everywhere else. Note that for every element d of the domain D and for all $i \geq 0$, there exists a moment $k : m_i \leq k < m_{i+1}$ such that $L(d,k)$. Therefore, $\mathfrak{M}_{m_{i+1}} \models \forall x(\neg l \,\mathcal{S}\, L(x))$; hence, $\mathfrak{M}_{m_i} \models \Diamond \bigcirc (\forall x(\neg l \,\mathcal{S}\, L(x)) \wedge l)$, making (2) **True** in $\mathfrak{M}'$.

We show that any model for (2), $\mathfrak{M}$, is also a model for (1). It is enough to show that for infinitely many n's there exists an m such that $m > n \wedge \forall x \exists k(n \leq k < m \wedge L(x,k))$ holds. Let m_i be the sequence of all moments when l is **True** (there are infinitely many of these). We show that for all $i \geq 0$, $n = m_i$, and $m = m_{i+1}$, the formula $\forall x \exists k(n \leq k < m \wedge L(x,k))$ is true in $\mathfrak{M}$. Indeed, $\mathfrak{M}_n \models l$, $\mathfrak{M}_m \models l$; by semantics of the operator "$\mathcal{S}$", $\mathfrak{M}_m \models \forall x(\neg l \,\mathcal{S}\, L(x))$ means that $\mathfrak{M}_m \models \forall x \exists k(n \leq k < m \wedge L(x,k))$. ❑

Proof [Lemma 15]

Following [8] let N be the number supplied by Lemma 14. Let us define a domain $D = \{\langle r,k\rangle \mid r \in \mathcal{R}_e^f, k < N\}$. Then for every $n \in \mathbb{N}$ we have $D = \bigcup_{\gamma \in \Gamma_n} D_{(n,\gamma)}$, where $D_{(n,\gamma)} = \{\langle r,k\rangle \mid r(n) = \gamma\}$ and $\left|D_{(n,\gamma)}\right| = N$.

Hence, by Lemma 14, for every $n \in \mathbb{N}$ there exists an $\mathcal{L}$-structure $\mathfrak{M}_n = \langle D, I_n\rangle$ satisfying $\mathcal{U} \cup \{\mathcal{G}_n\}$ such that $D_{(n,\gamma)} = \{\langle r,k\rangle \in D \mid \mathfrak{M}_n \models F_\gamma(\langle r,k\rangle)\}$. Moreover, we can suppose that $c^{I_n} = \langle r_c, 0\rangle$ for every constant $c \in \mathrm{const}(\mathsf{P})$. A potential first order temporal model is $\mathfrak{M} = \langle D, I\rangle$, where $I(n) = I_n$ for all $n \in \mathbb{N}$. To be convinced of this we have to check validity of step and eventuality clauses. (Recall that satisfiability of $\mathcal{I}$ and $\mathcal{U}$ in $\mathfrak{M}_0$ is implied by satisfiability of $\mathcal{G}_0$ in $\mathfrak{M}_0$ and definition of a behaviour graph.)

Let $\Box\forall x(P_i(x) \Rightarrow \bigcirc R_i(x))$ be an arbitrary step clause; we show that it is true in $\mathfrak{M}$. Namely, we show that for every $n \geq 0$ and every $\langle r,k\rangle \in D$, if $\mathfrak{M}_n \models P_i(\langle r,k\rangle)$ then $\mathfrak{M}_{n+1} \models R_i(\langle r,k\rangle)$. Suppose $r(n) = \gamma \in \Gamma_n$ and $r(n+1) = \gamma' \in \Gamma'$, where (γ, γ') is a suitable pair in accordance with the definition of a run. It follows that $\langle r,k\rangle \in D_{(n,\gamma)}$ and $\langle r,k\rangle \in D_{(n+1,\gamma')}$, in other words $\mathfrak{M}_n \models F_\gamma(\langle r,k\rangle)$ and $\mathfrak{M}_{n+1} \models F_{\gamma'}(\langle r,k\rangle)$. Since $\mathfrak{M}_n \models P_i(\langle r,k\rangle)$ then $P_i(x) \in \gamma$. It follows that $R_i(x)$ is a conjunctive member of $B_\gamma(x)$. Since the pair (γ, γ') is suitable, it follows that the conjunction $\exists x(F_{\gamma'}(x) \wedge B_\gamma(x))$ is satisfiable and, moreover, $\models \forall x(F_{\gamma'}(x) \Rightarrow B_\gamma(x))$. Together with $\mathfrak{M}_{n+1} \models F_{\gamma'}(\langle r,k\rangle)$ this implies that $\mathfrak{M}_{n+1} \models R_i(\langle r,k\rangle)$. Propositional step clauses are treated in a similar way.

Let $(\Box\forall x)\Diamond L(x)$ be an arbitrary eventuality clause. We show that for every $n \geq 0$ and every $\langle r,k\rangle \in D$, $r \in \mathcal{R}_e, k < Na$, there exists $m > n$ such that $\mathfrak{M}_m \models L(\langle r,k\rangle)$. Since r is an e-run, there exists $\mathcal{C}' = \pi(m)$ for some $m > n$ such that $r(m) = \gamma' \in \Gamma'$ and $L(x) \in \gamma'$. It follows that $\langle r,k\rangle \in D_{(m,\gamma')}$, that is $\mathfrak{M}_m \models F_{\gamma'}(\langle r,k\rangle)$. In particular, $\mathfrak{M}_m \models L(\langle r,k\rangle)$. Propositional eventuality clauses are considered in a similar way. ❑

Proof [Lemma 16] Let temporal problem $\mathsf{P_F} = \langle \mathcal{I}_\mathsf{F}, \mathcal{U}_\mathsf{F}, \mathcal{S}_\mathsf{F}, \mathcal{E}_\mathsf{F}\rangle$ be the result of the normal form transformation for $(\mathsf{P} \wedge \mathsf{Fin})$. One can prove that there exists a path $\pi_\mathsf{F} = \pi_1(\pi_2)^*$ through H_F, the behaviour graph for $\mathsf{P_F}$, such that the following conditions hold.

(a) $\pi_\mathsf{F}(0)$ is an initial vertex of H_F;

(b) for every colour scheme $\mathcal{C} = \pi_\mathsf{F}(i)$, $i \geq 0$, and every ground eventuality literal $\Diamond l \in \mathcal{E}_\mathsf{F}$ there exists a colour scheme $\mathcal{C}' = \pi_\mathsf{F}(j)$, $j > i$, such that $l \in \theta'$;

(c) for every colour scheme $\mathcal{C} = \pi_\mathsf{F}(i)$, $i \geq 0$ and every predicate colour γ from the colour scheme there exists an e-run $r \in \mathcal{R}_e(\pi_\mathsf{F})$ such that $r(i) = \gamma$;

(d) for every constant $c \in \mathcal{L}$, the function $r_c(n)$ defined by $r_c(n) = \rho_n(c)$, where ρ_n is the constant distribution from $\pi_\mathsf{F}(n)$, is an e-run in π_F.

By Lemma ? from [3], there exists a model $\mathfrak{M}$ for $\mathsf{P_F}$ with a domain D such that $\pi_\mathsf{F}(i) = (\Gamma_i^M, \theta_i^M, \rho_i^M)$, where $\Gamma_i^M = \{\gamma^i_{(a)} \mid a \in D\}$, $\rho_i^M(c) = \gamma^i_{(c^I)}$,

$$\gamma^i_{(a)} = \{P \mid P \in \mathsf{Pred(P)},\ \mathfrak{M} \models P(a)\} \cup \{\neg P \mid P \in \mathsf{Pred(P)},\ \mathfrak{M} \not\models P(a)\}$$

and

$$\theta_i^M = \{p \mid p \in \mathsf{Prop}(\mathsf{P}),\ \mathfrak{M} \models p\} \cup \{\neg p \mid p \in \mathsf{Prop}(\mathsf{P}),\ \mathfrak{M} \not\models p\}.$$

Furthermore, for every $a \in D$, $r_a(i) = \gamma^i_{(a)}$ is a e-run in π_{F}; and for every e-run r' in π_{F} there exists $b \in D$ such that $r'(i) = \gamma^i_{(b)}$.

Let γ be a predicate colour for P_{F}. By $\gamma|_{\mathsf{P}}$ we denote the predicate colour for P obtained from γ by removing any literals P, $\neg P$ such that $P \in \mathsf{Pred}(\mathsf{P}_{\mathsf{F}})$ but $P \notin \mathsf{Pred}(\mathsf{P})$ (that is, removing symbols originating from the translation of Fin). Similarly, $\theta|_{\mathsf{P}}$ is obtained from a propositional colour for P_{F} by removing all propositional symbols originating from Fin. For a colour scheme $\mathcal{C}$ for P_{F}, by $\mathcal{C}|_{\mathsf{P}}$ we denote the colour scheme for P obtained by removing symbols originating from Fin. Clearly, $\pi_{\mathsf{F}}|_{\mathsf{P}}$ is a path in H, the behaviour graph for P, and all the conditions of Lemma 15, except possibly (c), are satisfied. We show that condition (c) is also satisfied.

Consider the set of runs $\mathcal{R} = \{r \in \mathcal{R}_e(\pi_{\mathsf{F}})$ such that $r(l_1 + i)|_{\mathsf{P}} = r(l_1 + l_2 + i)|_{\mathsf{P}}$, where l_1 and l_2 are lengths of π_1 and π_2, resp$\}$. It can be seen that $\mathcal{R}|_{\mathsf{P}} = \mathcal{R}^f_e(\pi_{\mathsf{F}}|_{\mathsf{P}})$. Notice that if for every $i \geq l_1$ and every predicate colour $\gamma \in \pi(i)$ there exists an e-run $r \in \mathcal{R}$ such that $r(i)|_{\mathsf{P}} = \gamma|_{\mathsf{P}}$, (that is, property (c) holds for the looping part of $\pi_{\mathsf{F}}|_{\mathsf{P}}$), there exists an e-run $r \in \mathcal{R}$ such that $(r(i))|_{\mathsf{P}} = \gamma|_{\mathsf{P}}$ for all $i < l_1$ and $\gamma|_{\mathsf{P}} \in (\pi_{\mathsf{F}}(i))|_{\mathsf{P}}$ (that is, property (c) holds for the initial segment of $\pi_{\mathsf{F}}|_{\mathsf{P}}$).

Assume now that for some $i \geq l_1$ and some $\gamma \in \pi_{\mathsf{F}}(i)$ there exists no e-run $r \in \mathcal{R}$ such that $(r(i))|_{\mathsf{P}} = \gamma|_{\mathsf{P}}$. Then for any run $r \in \mathcal{R}_e(\pi_{\mathsf{F}})$ such that $(r(i))|_{\mathsf{P}} = \gamma|_{\mathsf{P}}$ we have $r(i+l_2) \neq r(i)$. But then there exist $a \in D$ such that $\mathfrak{M}_i \models F_\gamma(a)$ and for all $j > i$, $\mathfrak{M}_j \not\models \mathcal{F} \wedge F_\gamma(a)$. Since $\mathfrak{M}$ is a model for P_{F}, $\mathfrak{M}_i \models \Diamond\Box\forall x\neg(\mathcal{F} \wedge F_\gamma(x))$ contradicting $\mathfrak{M}_{i+k\cdot l_2} \models \mathcal{F}$ for every $k \geq 0$. ❑