On Stubborn Sets in the Verification of Linear Time Temporal Properties

Kimmo Varpaaniemi
Helsinki University of Technology, Digital Systems Laboratory
P.O. Box 1100, FIN-02015 HUT, Finland
E-mail: Kimmo.Varpaaniemi@hut.fi

Abstract. The stubborn set method is one of the methods that try to relieve the state space explosion problem that occurs in state space generation. This paper is concentrated on the verification of nexttime-less LTL (linear time temporal logic) formulas with the aid of the stubborn set method. The contribution of the paper is a theorem that gives us a way to utilize the structure of the formula when the stubborn set method is used and there is no fairness assumption. Connections to already known results are drawn by modifying the theorem to concern verification under fairness assumptions.

1 Introduction

Reachability analysis, also known as *exhaustive simulation* or *state space generation*, is a powerful formal method for detecting errors in concurrent and distributed finite state systems. Strictly speaking, infinite state systems can be analyzed, too, but reachability analysis methods are typically such that they cannot process more than a finite set of states. Nevertheless, we can quite well try to find errors even in cases where we do not know if or not the complete state space of the system is finite.

Anyway, reachability analysis suffers from the so called *state space explosion problem*, i.e. the complete state space of a system can be far too large w.r.t. the resources needed to inspect all states in the state space. Fortunately, in a variety of cases we do not have to inspect all reachable states of the system in order to get to know if or not errors of a specified kind exist.

The *stubborn set method* [22–26], and the *sleep set method* [8, 14, 16] are state search techniques that are based on the idea that when two executions of action sequences are sufficiently similar to each other, it is not necessary to investigate both of the executions. *Persistent sets* [8, 9] and *ample sets* [16–18] are strikingly similar to stubborn sets, at least if we consider the actual construction algorithms that have been suggested for stubborn, persistent and ample sets. This similarity is made explicit in [13] where a set is said to be a *stamper set* whenever the set is stubborn or ample or persistent in some way. Other closely related techniques have been presented in e.g. [1, 6, 10, 12, 15, 19, 28, 29]. This paper is concentrated on the theory of the stubborn set method.

Place/transition nets [21] are the formalism to which the stubborn set method is applied in this paper. The main reason for this choice is that there is hardly no simple and well-known formalism where the whole theory of the stubborn set method could be put into practice in a more fine-grained way. (For example, the difference between (general) *dynamic stubbornness* and *strong dynamic stubbornness* [27] is significant in place/transition nets but does not seem to have any useful analogy in the theory of stubborn sets for process algebras [25].)

For historical reasons, "stubbornness" without any preceding attribute is defined in a way that directly indicates how such sets can be computed. When one wants to show results concerning the theoretical properties of the stubborn set method, dynamic stubbornness is a more appropriate notion. When definitions are as they should be, stubbornness implies dynamic stubbornness but not vice versa.

Linear time temporal logics [4] give us a straightforward though of course a limited way to express what should or should not happen in a concurrent or distributed system. Depending on the context, the abbreviation *LTL* refers either to a specific linear time temporal logic or to "a linear time temporal logic in general". In LTL, the satisfaction of a formula is measured w.r.t. an infinite or deadlock-ended execution. A formula is *valid at a state* iff the formula is satisfied by all those infinite and deadlock-ended executions that start from the state. Verifying a formula typically means showing that the formula is valid at the initial state of the system that is under analysis. Validity is sometimes redefined in such a way that the requirement of satisfaction is restricted to paths of a certain kind. *Fairness assumptions* [5] are one form of such a restriction. A definition of fairness expresses some kind of progress that is expected in situations of a certain kind.

On-the-fly verification of a property means that the property is verified during state space generation, in contrary to the traditional approach where properties are verified after state space generation. As soon as it is known whether the property holds, the generation of the state space can be stopped. Since an erroneous system can have much more states than the intended correct system, it is important to find errors as soon as possible. On the other hand, even in the case that all states become generated, the overhead caused by on-the-fly verification, compared to non-on-the-fly verification, is often negligible.

An LTL formula can be verified on-the-fly by means of a *Büchi automaton* [7]. A Büchi automaton that accepts sequences satisfying the negation of the formula can be constructed automatically and intersected with the state space of the modelled system during the construction of the latter. The state space of the system can easily be thought of as a Büchi automaton. The formula is valid in the state space of the system iff the intersection to be computed, also a Büchi automaton, accepts no sequence.

In the fundamental presentation of stubborn sets in the verification of nexttime-less LTL-formulas [23], the computation of stubborn sets is directed by atomic formulas only, and the reduced state space can be used for verifying any nexttime-less LTL-formula that is constructible from those atomic formulas.

Unfortunately, the state space generation algorithm in [23] tends to generate the complete state space when verification is done under some of the most typical fairness assumptions. (In [23], all reduction is gained by utilizing transitions that are "sufficiently uninteresting". A typical fairness assumption makes all transitions "too interesting" in this sense.) The approaches [16, 17] improve the approach of [23] by utilizing the structure of the formula and by allowing a fairness assumption. A weakness in [16, 17] is that the structure of the formula is utilized only in cases when fairness is assumed or the formula expresses a safety property. This paper improves the method by utilizing the structure of the formula when fairness is not assumed and the formula is arbitrary. (The expression "fairness is not assumed" should be read to mean "no kind of fairness is assumed" though the latter may sound like "unfairness is assumed".) Though the recently published alternative solution [13] can be considered more goal-oriented, it does not cover our approach.

We also consider the verification of nexttime-less LTL-formulas when fairness is assumed. For convenience, we concentrate on *operation fairness* [16], though we could in principle handle some of the weaker fairness assumptions mentioned by [16] in the same way. The LTL verification approach in [23] can systematically be modified to handle fairness assumptions efficiently, and our approach can be modified quite similarly. It is by no means surprising that we essentially end up in an approach similar to those in [16, 17].

The rest of this paper has been organized as follows. Section 2 presents basic definitions related to place/transition nets. Our version of a linear time temporal logic is presented in Section 3. Section 4 defines *dynamic stubbornness*. Section 5 is devoted to the main preservation theorem of this paper, concerning verification without fairness assumptions. Section 6 extends the results of Section 5 to concern verification with fairness assumptions. Conclusions are then drawn in Section 7.

2 Place/transition nets

This section presents basic definitions related to *place/transition nets with infinite capacities* [21]. (Capacities do not increase expression power and are typically eliminated anyway, so we do not include them in the definitions.) We shall use N to denote the set of non-negative integer numbers, 2^X to denote the set of subsets of the set X, X^* (respectively, X^∞) to denote the set of finite (respectively, infinite) words over the alphabet X, and ε to denote the empty word. For any alphabet X and for any $\rho \in X^\infty$, ρ is thought of as a function from N to X in such a way that $\rho = \rho(0)\rho(1)\rho(2)\ldots$.

Definition 2.1 A *place/transition net* is a quadruple $\langle S, T, W, M_0 \rangle$ such that S is the set of *places*, T is the set of *transitions*, $S \cap T = \emptyset$, W is a function from $(S \times T) \cup (T \times S)$ to N, and M_0 is the *initial marking* (*initial state*), $M_0 \in \mathcal{M}$ where $\mathcal{M}$ is the set of *markings* (*states*), i.e. functions from S to N. The net is *finite* iff $S \cup T$ is finite. If $x \in S \cup T$, then the set of *input elements* of x is

${}^\bullet x = \{y \mid W(y,x) > 0\}$, the set of *output elements* of x is $x^\bullet = \{y \mid W(x,y) > 0\}$, and the set of *adjacent elements* of x is $x^\bullet \cup {}^\bullet x$. A transition t *leads (can be fired) from a marking M to a marking M'* ($M[t\rangle M'$ for short) iff

$$\forall s \in S \; M(s) \geq W(s,t) \wedge M'(s) = M(s) - W(s,t) + W(t,s).$$

A transition t is *enabled at a marking M* iff t leads from M to some marking. A marking M is *terminal* iff no transition is enabled at M. □

In our figures, places are circles, transitions are rectangles, and the initial marking is shown by the distribution of tokens, black dots, onto places. A directed arc, i.e. an arrow, is drawn from an element x to an element y iff x is an input element of y. Then $W(x,y)$ is called the *weight* of the arc. As usual, the weight is shown iff it is not equal to 1.

Definition 2.2 Let $\langle S, T, W, M_0 \rangle$ be a place/transition net. The set T^* (respectively, T^∞) is called the *set of finite* (respectively, *infinite*) *transition sequences of the net.* Let f be a function from $\mathcal{M}$ to 2^T. A finite transition sequence σ *f-leads (can be f-fired) from a marking M to a marking M'* iff $M[\sigma\rangle_f M'$, where

$$\forall M \in \mathcal{M} \; M[\varepsilon\rangle_f M, \text{ and}$$

$$\begin{aligned} &\forall M \in \mathcal{M} \; \forall M' \in \mathcal{M} \; \forall \delta \in T^* \; \forall t \in T \\ &M[\delta t\rangle_f M' \Leftrightarrow (\exists M'' \in \mathcal{M} \; M[\delta\rangle_f M'' \wedge t \in f(M'') \wedge M''[t\rangle M'). \end{aligned}$$

A finite transition sequence σ is *f-enabled at a marking M* ($M[\sigma\rangle_f$ for short) iff σ f-leads from M to some marking. An infinite transition sequence σ is f-enabled at a marking M ($M[\sigma\rangle_f$ for short) iff all finite prefixes of σ are f-enabled at M. A marking M' is *f-reachable from a marking M* iff some finite transition sequence f-leads from M to M'. A marking M' is an *f-reachable marking* iff M' is f-reachable from M_0. The *f-reachability graph* of the net is the pair $\langle V, A \rangle$ such that the set of vertices V is the set of f-reachable markings, and the set of edges A is $\{\langle M, t, M' \rangle \mid M \in V \wedge M' \in V \wedge t \in f(M) \wedge M[t\rangle M'\}$. □

Let Ψ be the function from $\mathcal{M}$ to 2^T such that for each marking M, $\Psi(M) = T$. From now on in this paper, we use a plain "$\rangle$" instead of "$\rangle_\Psi$", and as far as the notions of Definition 2.2 are concerned, we replace "Ψ-xxx" by "xxx" (where xxx is any word), with the exception that the Ψ-reachability graph of the net is called the *full reachability graph* of the net. When f is clear from the context or is implicitly assumed to exist and be of a kind that is clear from the context, then the f-reachability graph of the net is called the *reduced reachability graph* of the net.

Definition 2.3 Let $\langle S, T, W, M_0 \rangle$ be a place/transition net. Let f be a function from $\mathcal{M}$ to 2^T and let G be the f-reachability graph of the net. For any edge $\langle M, t, M' \rangle$ of G, t is called the *label* of the edge. (The labelling of the paths of G then follows by a natural extension.) A path of G is called a *terminal path* iff the path is finite and no nonempty transition sequence is f-enabled at the last vertex of the path. □

Definition 2.4 Let $\langle S, T, W, M_0\rangle$ be a place/transition net. Let $T_s \subseteq T$. A finite transition sequence, δ, T_s*-exhausts* a finite transition sequence σ iff for each $t \in T_s$, the number of t's in δ is greater than or equal to the number of t's in σ. The function $\Re$ from $(T^* \cup T^\infty) \times 2^T$ to $T^* \cup T^\infty$ is defined by requiring that for each $Y \in 2^T$, $\Re(\varepsilon, Y) = \varepsilon$, and for each $t_1 \in Y$, for each $t_2 \in T \setminus Y$, for each $\delta \in T^*$ and for each $\rho \in T^\infty$, $\Re(t_1\delta, Y) = t_1\Re(\delta, Y)$, $\Re(t_2\delta, Y) = \Re(\delta, Y)$, and $\Re(\rho, Y) = \Re(\rho(0), Y)\Re(\rho(1), Y)\Re(\rho(2), Y)\ldots$. For any $Y \in 2^T$ and for any $\sigma \in T^* \cup T^\infty$, $\Re(\sigma, Y)$ is called the Y*-restriction* of σ. Let $\Upsilon \subseteq 2^T$. A finite or an infinite transition sequence δ is Υ*-equivalent* to a finite or an infinite transition sequence σ iff for each $Y \in \Upsilon$, $\Re(\delta, Y) = \Re(\sigma, Y)$. Let $\mathcal{T} = \{\{t\} \mid t \in T\}$. A finite or an infinite transition sequence δ is a *permutation of a finite or an infinite transition sequence* σ iff δ is $\mathcal{T}$-equivalent to σ. $\square$

The above Υ can be considered as a set of views to the behaviour of the net. If $T = \{a, b, c, d, e, f, g\}$ and $\Upsilon = \{\{a, b\}, \{c, d\}, \{d, f\}\}$ then $gbdcefa$ is Υ-equivalent to $badfc$ since both of these sequences have the $\{a, b\}$-restriction ba, the $\{c, d\}$-restriction dc, and the $\{d, f\}$-restriction df.

Note that in the case of infinite sequences, the above definition of a permutation does not pay any attention to the possible repeated patterns in the sequences. So, for example the sequence obtained by repeating $bbba$ infinitely many times is a permutation of the sequence obtained by repeating ab infinitely many times.

3 An LTL

This section presents one version of a linear time temporal logic. The presentation assumes that the system to be analyzed has a place/transition net model. Our LTL has effectively the same syntax as the Propositional Linear Temporal Logic (PLTL) in [4]. The semantics are also effectively the same, with the exception that we consider finite executions, too. We make this difference because deadlock-ended executions are important to us whereas the semantic definitions for PLTL assume that every state has a successor.

A formula in our LTL is either atomic or of the form $\perp$, $(A) \Rightarrow (B)$, $\bigcirc(A)$ or $(A)\mathcal{U}(B)$ where A and B are formulas. The following are syntactic abbreviations: $\neg(A)$ means $(A) \Rightarrow (\perp)$, $\top$ means $\neg(\perp)$, $(A)\vee(B)$ means $(\neg(A)) \Rightarrow (B)$, $(A)\wedge(B)$ means $\neg((\neg(A))\vee(\neg(B)))$, $\Diamond(A)$ means $(\top)\mathcal{U}(A)$, and $\Box(A)$ means $\neg(\Diamond(\neg(A)))$. An atomic formula is a subset of markings of the net, i.e. a subset of $\mathcal{M}$. In our examples, all atomic formulas are of the form "$M(s)$ op k" where $k \in N$, s is a place in the net, op is a comparison operator and the actual meaning of the formula "$M(s)$ op k" is $\{M \in \mathcal{M} \mid M(s) \text{ op } k\}$.

The operators $\perp$, $\Rightarrow$, $\neg$, $\top$, $\wedge$ and $\vee$ are called *propositional.* The other operators are then called non-propositional or *temporal.* Non-propositional operators have the following names: $\bigcirc$ is "nexttime", $\mathcal{U}$ is "until", $\Diamond$ is "eventually" and $\Box$ is "henceforth". A formula is *nexttime-less* iff the formula does not contain any $\bigcirc$. By a *Boolean combination* of formulas from a collection we mean a formula that can be constructed from the formulas of the collection by using propositional operators only. (A single formula can be used several times in the combination whereas it is not necessary to use all formulas of the collection.)

The rules of satisfaction of a formula are given w.r.t. finite and infinite paths in a (reduced or full) reachability graph of the net and are as follows. (We assume that a path always contains at least one vertex and starts with a vertex. Moreover, each finite path ends with a vertex. Also, paths x and y can be concatenated into a path xy iff x is finite and the last vertex of x is the first vertex of y. The path xy is then the path "x continued by y.")

- A path satisfies an atomic formula p iff the first vertex of the path is in p.
- No path satisfies $\perp$.
- A path satisfies $(A) \Rightarrow (B)$ iff the path satisfies B or does not satisfy A.
- A path x satisfies $\bigcirc(A)$ iff there is at least one edge in the path and A is satisfied by the path obtained from x by removing the first vertex and the first edge.
- A path x satisfies $(A)\mathcal{U}(B)$ iff there is a path z and a finite path y such that $x = yz$, z satisfies B, and for any finite paths v and u, $y = uv \neq u$ implies that vz satisfies A.

A formula is *valid at a marking in the graph* iff the formula is satisfied by all those infinite and terminal paths of the graph that start from the marking. (So, a formula $(A) \wedge (B)$ is valid at a marking iff both of A and B are valid at the marking. On the other hand, $(A) \vee (B)$ can be valid at a marking even in the case that neither A nor B is valid at the marking.) *Verifying a formula* means showing that the formula is valid at the initial marking in the full reachability graph of the net.

For convenience, validity is sometimes redefined in such a way that the requirement of satisfaction is restricted to paths of a certain kind. The restriction may or may not be expressible in LTL. *Fairness assumptions* [5] are one form of such a restriction. Fairness is basically an informal concept, and the choice of a formal definition depends much on the context. Anyway, a definition of fairness expresses some kind of progress that is expected in situations of a certain kind. Also, some definitions of fairness have turned out to be of general interest. To this paper, we have chosen one of such definitions, *operation fairness* [16] that is a certain type of *strong fairness* [5].

Definition 3.1 Let $\langle S, T, W, M_0 \rangle$ be a place/transition net. A path in the full reachability graph of the net is *operation fair* iff the following holds for each transition t : if t is enabled infinitely many times on the path, then the path contains infinitely many occurrences of t. (Note that all finite paths are thus operation fair.) □

Operation fairness cannot be expressed in our LTL because our version of LTL has no general way to describe the occurrence of a transition in such a way that the description would match only that transition. On the other hand, as can be seen from [11, 16], operation fairness is easily expressible in action-oriented versions of LTL, at least if the net does not have infinitely many transitions. Though we cannot express operation fairness in our LTL, we can still handle it formally without difficulties, as we shall see in Section 6.

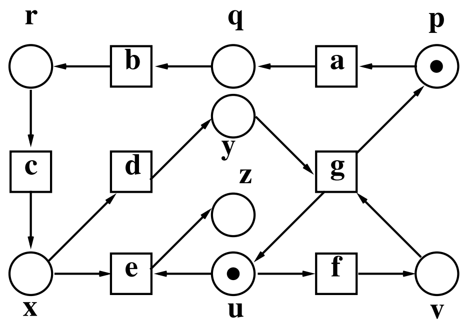

Fig. 1. In the full reachability graph of this net, $(abfcdg)(abfcdg)(abfcdg)\ldots$ labels an operation fair path while $(abcfdg)(abcfdg)(abcfdg)\ldots$ does not.

Operation fairness is not guaranteed to be preserved when the order of firing of transitions is changed in such a way that the resulting path has no suffix that would be a suffix of the original path. In the net in Figure 1, the path starting from the initial marking M_0 and being labelled by $(abcfdg)(abcfdg)(abcfdg)\ldots$ is not operation fair though the path starting from M_0 and being labelled by $(abfcdg)(abfcdg)(abfcdg)\ldots$ is operation fair.

4 Dynamic stubbornness

When one wants to show results concerning the theoretical properties of the stubborn set method, it is often best to use a dynamic definition of stubbornness. The below principles D1 and D2 are the principles 1^* and 2^* of [20], respectively. Dynamic stubbornness has also been handled in e.g. [24, 27].

Definition 4.1 Let $\langle S, T, W, M_0 \rangle$ be a place/transition net. Let M be a marking of the net. A set $T_s \subseteq T$ *fulfils the first principle of dynamic stubbornness* (*D1* for short) *at* M iff $\forall \sigma \in (T \setminus T_s)^* \ \forall t \in T_s \ M[\sigma t\rangle \Rightarrow M[t\sigma\rangle$. A transition t is a *dynamic key transition of a set* $T_s \subseteq T$ *at* M iff $t \in T_s$ and $\forall \sigma \in (T \setminus T_s)^* \ M[\sigma\rangle \Rightarrow M[\sigma t\rangle$. A set $T_s \subseteq T$ *fulfils the second principle of dynamic stubbornness* (*D2* for short) *at* M iff T_s has a dynamic key transition at M. A set $T_s \subseteq T$ is *dynamically stubborn at* M iff T_s fulfils D1 and D2 at M. A function f from $\mathcal{M}$ to 2^T is *a dynamically stubborn function* iff for each marking M, either $f(M)$ is dynamically stubborn at M or no transition is enabled at M. $\square$

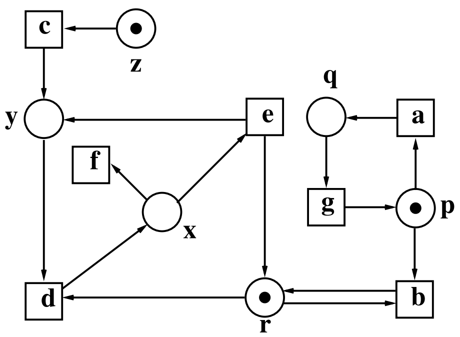

Fig. 2. A net demonstrating dynamic stubbornness.

An interesting thing in Definition 4.1 is that it does not require any true dependency relation between transitions. For example, consider the net in Figure 2. The transition sequences *cdebdf*, *cbdedf* and *bcdedf* all lead from the initial marking M_0 to the same terminal marking, and the only essential difference between the sequences is the position of b. Though we can well consider b independent of c, it is difficult to imagine even any "flexible" dependency relation that would make b independent of d at all "important" markings. The set $\{a, b\}$ is dynamically stubborn at M_0. If h is a dynamically stubborn function and $h(M_0) = \{a, b\}$, the h-reachability graph has no path where c would be fired at M_0.

5 A preservation theorem

Let us call a formula *directly temporal* iff the outermost operator of the formula is a non-propositional operator. A nexttime-less LTL-formula can be transformed into a nexttime-less LTL-formula where directly temporal subformulas are as short as possible [16]. Then a suitable reduced reachability graph can be generated by using the stubborn set method, provided that the conditions in Proposition 5.1 are satisfied. Note that any formula can be seen as a Boolean combination of directly temporal subformulas. The $\Box(\bigcirc(\top))$ formula occurring in Proposition 5.1 is satisfied by every infinite path whereas no terminal path satisfies it.

Proposition 5.1 *Assumptions:*

(P1) $\langle S, T, W, M_0 \rangle$ *is a place/transition net. (The net and the full reachability graph of the net can be finite or infinite.)*
(P2) Φ *is a collection of nexttime-less LTL-formulas. (Φ can be finite or infinite.)*
(P3) Π *is a function from* $2^{\mathcal{M}}$ *to* 2^S *in such a way that whenever we have a subset* p *of* $\mathcal{M}$*) and markings* M *and* M' *for which* $M \in p$ *and* $M' \notin p$*, there exists* $s \in \Pi(p)$ *for which* $M(s) \neq M'(s)$*.*

(P4) Ξ *is a function from* Φ *to* 2^T *in such a way that for each* $\phi \in \Phi$ *and for each atomic subformula* p *of* ϕ, $\{t \in T \mid \exists s \in \Pi(p) \;\; W(s,t) \neq W(t,s)\} \subseteq \Xi(\phi)$.

(P5) Υ *is a (finite or an infinite) subset of* 2^T *such that* $\{\Xi(\phi) \mid \phi \in \Phi\} \subseteq \Upsilon$.

(P6) f *is a function from* $\mathcal{M}$ *to* 2^T *in such a way that every terminal path in the* f*-reachability graph of the net is a terminal path of the full reachability graph of the net. (The* f*-reachability graph of the net can be finite or infinite.)*

(P7) For each terminal path starting from M_0 *in the full reachability graph, there exists a terminal path starting from* M_0 *in the* f*-reachability graph in such a way that the labels of the paths are* Υ*-equivalent.*

(P8) For each infinite path starting from M_0 *in the full reachability graph, there exists an infinite path starting from* M_0 *in the* f*-reachability graph in such a way that the labels of the paths are* Υ*-equivalent.*

Claim: For any boolean combination ϕ *of the formulas in* $\Phi \cup \{\Box(\bigcirc(\top))\}$, ϕ *is valid at* M_0 *in the full reachability graph of the net iff* ϕ *is valid at* M_0 *in the* f*-reachability graph of the net.*

Proof. The "only if" -part of the claim is obvious. The "if" -part can be shown by using a transformation from a path into a *propositional sequence* [13, 17] and by utilizing *equivalence up to stuttering* [13, 17]. □

There is actually nothing new or amazing in Proposition 5.1, and its only purpose is to serve as an interface to Theorem 5.7, i.e. instead of talking about formulas we can talk about Υ-equivalence. Claims of Theorem 5.7 occur as assumptions in Proposition 5.1.

Theorem 5.7, the goal of this section, is a refinement of Theorem 2 of [23] and gives us better chances for reduction. The refinement is strongly inspired by [16, 17]. The new aspect in Theorem 5.7 is that we do not preserve all orders of *visible* transitions. A transition is visible iff at least one member of the above defined Υ contains the transition. Roughly speaking, visible transitions are those transitions that determine the satisfaction of the atomic subformulas of the interesting formulas. In a verification task, if the original formula to be verified is ϕ_0 and an equivalent formula obtained by transformation is ϕ_1, then the collection of interesting formulas consists of directly temporal formulas such that ϕ_1 is a Boolean combination of the formulas in the collection. (If ϕ_1 itself is directly temporal, then the collection is simply $\{\phi_1\}$.)

Let us look at the net in Figure 3. Clearly, the full reachability graph of the net has no terminal path but has exactly two infinite paths that start from the initial marking. Among these two paths, the path labelled by $accc\ldots$ satisfies the formula $\Diamond(M(q) = 0)$ while the path labelled by $baccc\ldots$ does not. However, $\mathcal{M}$ has no markings M_1 and M_2 for which it would be that $M_1[b\rangle M_2$ and either $M_1(q) = 0 \neq M_2(q)$ or $M_2(q) = 0 \neq M_1(q)$. We still consider b as a visible transition w.r.t. the atomic formula "$M(q) = 0$", since in the sequel, transitions like this would anyway be treated like the "pedantically visible" transitions.

The assumptions of Theorem 5.7 are the following.

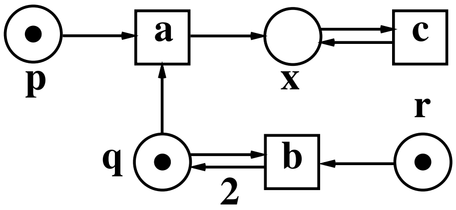

Fig. 3. Both of a and b are visible w.r.t. the atomic formula "$M(q) = 0$".

(A1) $\langle S, T, W, M_0 \rangle$ is a place/transition net, $\Upsilon \subseteq 2^T$, and $J = T \setminus \cup_{Y \in \Upsilon} Y$. (The net, Υ and the full reachability graph of the net can be finite or infinite.)
(A2) f is a dynamically stubborn function from $\mathcal{M}$ to 2^T. (The f-reachability graph of the net can be finite or infinite.)
(A3) For any $Y \in \Upsilon$ and for any marking M, $Y \subseteq f(M)$ or $\{t \in Y \cap f(M) \mid M[t\rangle\} = \emptyset$ (or both).
(A4) For any marking M, if $f(M)$ does not contain all those transitions that are enabled at M, then some transition in J is a dynamic key transition of $f(M)$ at M.
(A5) For any $t \in T \setminus J$, every infinite path (starting from a marking whatsoever) in the f-reachability graph of the net contains at least one marking M such that $t \in f(M)$.

Coarsely speaking, A3 prevents us from changing the order of transitions that are visible w.r.t. a single member of Φ while A4 and A5 prevent us from ignoring any member of Φ. The transitions in J are invisible w.r.t. all members of Φ. There is a following correspondence between A3 – A5 and the assumptions 2 – 4 of Theorem 2 of [23]: if $|\Upsilon| = 1$, n is between 3 and 5 and the f-reachability graph is finite, An becomes assumption $n - 1$ of [23].

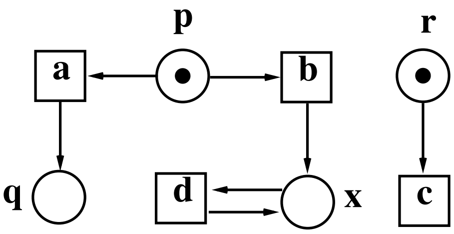

Fig. 4. An example net for verification.

Let us consider an example where we try to verify the formula

$$(\Diamond(M(q) = 1)) \vee (\Diamond(M(r) = 0))$$

about the net in Figure 4. We can let $\Upsilon = \{\{a\}, \{c\}\}$ ($|\Upsilon| = 2$) since the satisfaction of $M(q) = 1$ can be affected by a only whereas the satisfaction of $M(r) = 0$ can be affected by c only. Let us choose $\{a, b\}$ for the dynamically stubborn set at the initial marking. This choice respects all of A1 – A5. (Note that [23] would not accept such a choice but would require us to take all enabled transitions into the set. We have thus gained reduction w.r.t. [23].) At any other encountered nonterminal marking, we let the dynamically stubborn set contain all enabled transitions since A1 – A5 would otherwise be violated. (The same would have to be done if the conditions in [23] would have to be satisfied instead.) The reduced reachability graph has exactly one terminal path that starts from the initial marking, and the label of that path is ac. The labels of the infinite paths starting from the initial marking in the reduced reachability graph are $bddd\ldots$, $bcddd\ldots$, $bdcddd\ldots$, $bddcddd\ldots$, etc. From these paths the path labelled by $bddd\ldots$ invalidates the formula.

Let us then verify the formula

$$(\Diamond(M(x) = 1)) \vee (\Diamond(M(r) = 0)).$$

Using similar reasoning as above, we can let $\Upsilon = \{\{b\}, \{c\}\}$ ($|\Upsilon| = 2$). (Though d is connected to x, d cannot affect the satisfaction of $M(x) = 1$.) Proceeding as above, we actually get exactly the same reduced reachability graph, but that is merely a coincidence. Since there is no counterexample to the formula, we conclude that the formula is valid at the initial marking. (This is indeed an example of a disjunction that is valid despite of the fact that none of the disjuncts is valid.)

Let us also look what would be the consequences if some of A3 – A5 were dropped. Dropping A3 could make us draw a wrong conclusion about

$$\Box(((M(r) = 1) \vee (M(q) = 1)) \vee (\Box(M(q) = 0))).$$

When $\Upsilon = \{\{a, c\}\}$ ($|\Upsilon| = 1$), we could choose $\{a, b\}$ for the dynamically stubborn set at the initial marking. The only counterexample to the formula, i.e. the path starting from the initial marking and being labelled by ca, would then be lost.

Dropping A4 could make us draw a wrong conclusion about $\Diamond(M(r) = 0)$. When $\Upsilon = \{\{c\}\}$, we could choose $\{c\}$ for the dynamically stubborn set at the initial marking. The only counterexample to the formula, i.e. the path starting from the initial marking and being labelled by $bddd\ldots$, would then be lost.

Dropping A5 could make us draw a wrong conclusion about

$$\Box(((M(x) = 0) \vee (M(r) = 0)) \vee (\Box(M(r) = 1))).$$

When $\Upsilon = \{\{b, c\}\}$ ($|\Upsilon| = 1$), we could let $\{a, b, c\}$ be the dynamically stubborn set at the initial marking and choose $\{d\}$ to be the dynamically stubborn set at the marking to which b leads from the initial marking. All the counterexamples to the formula, i.e. the paths where c occurs after b, would then be lost.

In the net in Figure 5, omitting the attribute "dynamic key" in A4 could make us draw a wrong conclusion about

$$(\Diamond(M(q) = 1)) \vee (\Diamond(M(y) = 1)).$$

When $\Upsilon = \{\{a, d\}\}$ ($|\Upsilon| = 1$), we could choose $\{a, b, d\}$ for the dynamically stubborn set at the initial marking. The only counterexample to the formula, i.e. the path starting from the initial marking and being labelled by $ceee\ldots$, would then be lost.

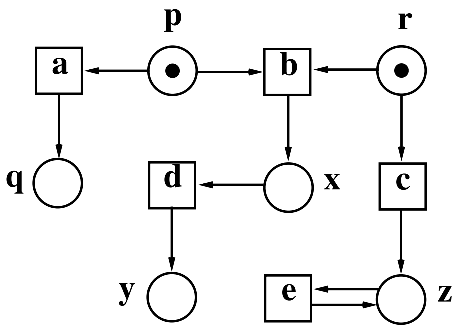

Fig. 5. A net that motivates the assumption A4.

We now start working towards Theorem 5.7. Lemma 5.2 tells us that the used transition selection function respects the important orderings of transitions.

Lemma 5.2 *Assumptions: A1, A2 and A3.*
Claim: For each nonterminal marking M, for each t in $f(M)$ and for each σ in $(T \setminus f(M))^*$, if $M[\sigma t\rangle$, then $M[t\sigma\rangle$, and $t\sigma$ is Υ-equivalent to σt.

Proof. Let M be a nonterminal marking, $t \in f(M)$ and $\sigma \in (T \setminus f(M))^*$ in such a way that $M[\sigma t\rangle$. From D1 (and, as goes without saying, from A2) it follows that $M[t\sigma\rangle$. Let $Y \in \Upsilon$. If $Y \subseteq f(M)$, then $\sigma \in (T \setminus Y)^*$ and thus $\Re(t\sigma, Y) = t = \Re(\sigma t, Y)$. If $Y \not\subseteq f(M)$, then A3 has the effect that $t \notin Y$, so $\Re(t\sigma, Y) = \Re(\sigma, Y) = \Re(\sigma t, Y)$. □

Lemma 5.3 guarantees that the possible terminal paths of the full reachability graph are sufficiently represented in the reduced reachability graph.

Lemma 5.3 *Assumptions: A1, A2 and A3.*
Claim: For each finite transition sequence σ'' and for each marking M'', if σ'' leads from M'' to a terminal marking M_d, then there exists a permutation δ'' of σ'' in such a way that $M''[\delta''\rangle_f M_d$ and δ'' is Υ-equivalent to σ''.

Proof. We use induction on the length of σ''. The claim holds trivially when restricted to $\sigma'' = \varepsilon$. Our induction hypothesis is that the claim holds when restricted to any σ'' of length $n \geq 0$.

Let a finite transition sequence σ of length $n+1$ lead from a marking M to a terminal marking M_d. From D2 it follows that there exist $t \in f(M)$, $\delta \in (T \setminus f(M))^*$ and $\delta' \in T^*$ in such a way that $\sigma = \delta t \delta'$. From Lemma 5.2 it follows that there exists a marking M' in such a way that $M[t\rangle_f M'$, $M'[\delta\delta'\rangle M_d$, and $t\delta\delta'$ is Υ-equivalent to σ.

By the induction hypothesis, there exists a permutation δ'' of $\delta\delta'$ in such a way that $M'[\delta''\rangle_f M_d$ and δ'' is Υ-equivalent to $\delta\delta'$. So, $t\delta''$ is a permutation of σ in such a way that $M[t\delta''\rangle_f M_d$ and $t\delta''$ is Υ-equivalent to σ. □

Lemma 5.4 guarantees that the possible infinite invisible transition sequences are sufficiently represented in the reduced reachability graph.

Lemma 5.4 *Assumptions: A1, A2 and A4.*

Claim: For each $\sigma'' \in J^\infty$ and for each $M'' \in \mathcal{M}$, if $M''[\sigma''\rangle$ then there exists $\delta'' \in J^\infty$ in such a way that $M''[\delta''\rangle_f$.

Proof. Let $\sigma \in J^\infty$ and $M \in \mathcal{M}$ such that $M[\sigma\rangle$. A4 and D1 guarantee that we can define a function τ from N to T, a function μ from N to $\mathcal{M}$ and a function θ from N to J^∞ as follows.

Firstly, $\mu(0) = M$ and $\theta(0) = \sigma$. Let then $k \in N$. If $\theta(k)$ contains a transition from $f(\mu(k))$, we let $t_k \in f(\mu(k))$, $\gamma_k \in (J \setminus f(\mu(k))^*$ and $\zeta_k \in J^\infty$ be such that $\gamma_k t_k \zeta_k = \theta(k)$ and require that $\tau(k) = t_k$, $\mu(k)[\tau(k)\rangle_f \mu(k+1)$ and $\theta(k+1) = \gamma_k \zeta_k$.

In the remaining case, we choose a transition t_k from $f(\mu(k)) \cap J$ in such a way that $\mu(k)[t_k \sigma \theta(k)\rangle$ and require that $\tau(k) = t_k$, $\mu(k)[\tau(k)\rangle_f \mu(k+1)$ and $\theta(k+1) = \theta(k)$.

The function τ represents an infinite transition sequence that is f-enabled at M. □

Lemma 5.5 states that if we have an infinite or a finite sequence in the full reachability graph, we can choose an arbitrary finite prefix of the sequence in such a way that there is a sequence that is Υ-equivalent to the original sequence and has a finite prefix that is f-enabled and covers the prefix we chose. (To remember the meaning of "exhausting", see Definition 2.4.)

Lemma 5.5 *Assumptions: A1, A2, A3, A4 and A5.*

Claim: For each $\sigma'' \in T^*$, for each $\rho' \in T^* \cup T^\infty$ and for each $M' \in \mathcal{M}$, if $M'[\sigma''\rho'\rangle$ and M' is f-reachable from M_0, then there exist $\gamma'' \in T^*$, $\delta_1 \in T^*$, $\delta_2 \in T^*$, $\rho'' \in T^* \cup T^\infty$ and $M'' \in \mathcal{M}$ in such a way that $\gamma''\rho'' = \rho'$, $M'[\delta_1\rangle_f M''$, $M''[\delta_2\rho''\rangle$, δ_1 $(T \setminus J)$-exhausts σ'', and $\delta_1\delta_2$ is Υ-equivalent to $\sigma''\gamma''$.

Proof. We use induction on the length of σ''. The claim holds trivially when restricted to $\sigma'' = \varepsilon$. Our induction hypothesis is that the claim holds when restricted to any σ'' of length $n \geq 0$. The claim holds trivially when restricted to any $\sigma'' \in J^*$ since in that case, ε is Υ-equivalent to σ'' and $(T \setminus J)$-exhausts σ'', so $\gamma'' = \delta_1 = \delta_2 = \varepsilon$, $\rho'' = \rho'$ and $M'' = M'$ are suitable choices for that case. Let then $\sigma \in T^* \setminus J^*$, $\rho \in T^* \cup T^\infty$ and $M \in \mathcal{M}$ be such that $M[\sigma\rho\rangle$, M is f-reachable from M_0 and the length of σ is $n+1$.

Let L be the set of those transitions that occur in σ. A4 and D1 guarantee that we can define functions ξ, β and η from N to T^*, a function μ from N to $\mathcal{M}$ and a function θ from N to $T^* \cup T^\infty$ as follows. Firstly, $\xi(0) = \varepsilon$, $\beta(0) = \varepsilon$, $\eta(0) = \varepsilon$, $\mu(0) = M$ and $\theta(0) = \rho$. Let then $k \in N$. If there exists $\tau \in L$ such that $\mu(k)[\tau\rangle_f$, then $\xi(k+1) = \xi(k)$, $\beta(k+1) = \beta(k)$, $\eta(k+1) = \eta(k)$, $\mu(k+1) = \mu(k)$ and $\theta(k+1) = \theta(k)$. Otherwise, if $\eta(k)$ contains a transition from $f(\mu(k))$, we let $\tau_k \in f(\mu(k))$, $\gamma_k \in (T \setminus f(\mu(k)))^*$ and $\zeta_k \in T^*$ be such that $\gamma_k\tau_k\zeta_k = \eta(k)$ and require that $\xi(k+1) = \xi(k)$, $\beta(k+1) = \beta(k)\tau_k$, $\eta(k+1) = \gamma_k\zeta_k$, $\mu(k)[\tau_k\rangle_f\mu(k+1)$ and $\theta(k+1) = \theta(k)$. In the remaining case, if $\theta(k)$ contains a transition from $f(\mu(k))$, we let $\tau_k \in f(\mu(k))$, $\gamma_k \in (T \setminus f(\mu(k))^*$ and $\zeta_k \in T^* \cup T^\infty$ be such that $\gamma_k\tau_k\zeta_k = \theta(k)$ and require that $\xi(k+1) = \xi(k)\gamma_k\tau_k$, $\beta(k+1) = \beta(k)\tau_k$, $\eta(k+1) = \eta(k)\gamma_k$, $\mu(k)[\tau_k\rangle_f\mu(k+1)$ and $\theta(k+1) = \zeta_k$. In the ultimate remaining case, we choose a transition τ_k from $f(\mu(k)) \cap J$ in such a way that $\mu(k)[\tau_k\sigma\eta(k)\theta(k)\rangle$ and require that $\xi(k+1) = \xi(k)$, $\beta(k+1) = \beta(k)\tau_k$, $\eta(k+1) = \eta(k)$, $\mu(k)[\tau_k\rangle_f\mu(k+1)$ and $\theta(k+1) = \theta(k)$.

Clearly, for each $k \in N$, $\xi(k)\theta(k) = \rho$ and $M[\beta(k)\rangle_f\mu(k)$. From Lemma 5.2 it follows that for each $k \in N$, $\mu(k)[\sigma\eta(k)\theta(k)\rangle$, and $\beta(k)\sigma\eta(k)$ is Υ-equivalent to $\sigma\xi(k)$.

Let us first assume that there are no $k' \in N$ and $\tau' \in L$ that would satisfy $\mu(k')[\tau'\rangle_f$. Let us call this assumption B. Since $\sigma \in T^* \setminus J^*$ and L is the set of those transitions that occur in σ, the set $L \setminus J$ is not empty. Let t be any transition in $L \setminus J$. From B and A5 it follows that there exists $k'' \in N$ such that $t \in f(\mu(k''))$. Consequently, there must be some $k_1 \leq k''$, $t' \in L \cap f(\mu(k_1))$, $\gamma \in (L \setminus f(\mu(k_1)))^*$ and $\gamma' \in L^*$ such that $\sigma = \gamma t'\gamma'$. Since $\mu(k_1)[\sigma\rangle$, from D1 it follows that $\mu(k_1)[t'\rangle_f$. We have thus reached a contradiction with B.

So, we can choose $k' \in N$ and $\tau' \in L$ such that $\mu(k')[\tau'\rangle_f$. Since $\mu(k')[\sigma\rangle$, there are some $t_1 \in f(\mu(k'))$, $\delta \in (T \setminus f(\mu(k')))^*$ and $\delta' \in T^*$ such that $\sigma = \delta t_1\delta'$. Since $\mu(k')[\sigma\eta(k')\theta(k')\rangle$, from Lemma 5.2 it follows that there exists a marking M_1 such that $\mu(k')[t_1\rangle_f M_1$, $M_1[\delta\delta'\eta(k')\theta(k')\rangle$, and $t_1\delta\delta'$ is Υ-equivalent to σ. So, $\beta(k')t_1\delta\delta'\eta(k')$ is Υ-equivalent to $\sigma\xi(k')$ since $\beta(k')\sigma\eta(k')$ is Υ-equivalent to $\sigma\xi(k')$.

By the induction hypothesis, there exists $\gamma'' \in T^*$, $\delta_1 \in T^*$, $\delta_2 \in T^*$, $\rho'' \in T^* \cup T^\infty$ and $M_2 \in \mathcal{M}$ in such a way that $\gamma''\rho'' = \eta(k')\theta(k')$, $M_1[\delta_1\rangle_f M_2$, $M_2[\delta_2\rho''\rangle$, δ_1 $(T \setminus J)$-exhausts $\delta\delta'$, and $\delta_1\delta_2$ is Υ-equivalent to $\delta\delta'\gamma''$. Then $M[\beta(k')t_1\delta_1\rangle_f M_2$ and $\beta(k')t_1\delta_1$ $(T \setminus J)$-exhausts $\delta t_1\delta' = \sigma$.

Let us first consider the case that γ'' is shorter than $\eta(k')$. Let $\delta_3 \in T^*$ be such that $\gamma''\delta_3 = \eta(k')$. Then $\delta_3\theta(k') = \rho''$. We thus have that $M_2[\delta_2\delta_3\theta(k')\rangle$. On the other hand, $\xi(k')\theta(k') = \rho$. Moreover, $\beta(k')t_1\delta_1\delta_2\delta_3$ is Υ-equivalent to $\sigma\xi(k')$ since $\delta_1\delta_2\delta_3$ is Υ-equivalent to $\delta\delta'\gamma''\delta_3 = \delta\delta'\eta(k')$ whereas $\beta(k')t_1\delta\delta'\eta(k')$ is Υ-equivalent to $\sigma\xi(k')$.

Let us then consider the case that γ'' is at least as long as $\eta(k')$. Let $\delta_4 \in T^*$ be such that $\eta(k')\delta_4 = \gamma''$. Then $\delta_4\rho'' = \theta(k')$. We thus have that $\xi(k')\delta_4\rho'' = \xi(k')\theta(k') = \rho$. On the other hand, $M_2[\delta_1\delta_2\rho''\rangle$. Moreover, $\beta(k')t_1\delta_1\delta_2$ is Υ-equivalent to $\sigma\xi(k')\delta_4$ since $\delta_1\delta_2$ is Υ-equivalent to $\delta\delta'\gamma'' = \delta\delta'\eta(k')\delta_4$ whereas $\beta(k')t_1\delta\delta'\eta(k')$ is Υ-equivalent to $\sigma\xi(k')$. $\square$

If we have a series of finite prefixes of an infinite sequence in the full reachability graph, Lemma 5.5 gives us a series of finite sequences in the reduced reachability graph, but the series is not necessarily a series of prefixes of any single infinite sequence. There is still one thing we can do: we can move along a path in the reduced reachability graph and apply Lemma 5.5 to any of the markings on the path. Guaranteeing Υ-equivalence is then not difficult at all since for any infinite sequence, Lemma 5.5 just modifies some finite prefix of the sequence and leaves the rest of the infinite sequence untouched. We then just have to make sure that we choose a prefix that includes some of the so far untouched part of the original infinite sequence. This is the idea of the proof of the below Lemma 5.6.

Lemma 5.6 *Assumptions: A1, A2, A3, A4 and A5.*

Claim: For each $\sigma'' \in T^\infty$ for each $M'' \in \mathcal{M}$, if $M''[\sigma''\rangle$, M'' is f-reachable from M_0 and σ'' contains infinitely many occurrences of transitions from $T \setminus J$, then there exists $\delta'' \in T^\infty$ in such a way that $M''[\delta''\rangle_f$ and δ'' is Υ-equivalent to σ''.

Proof. Let $\sigma \in T^\infty$ and $M \in \mathcal{M}$ be such that $M[\sigma\rangle$, M is f-reachable from M_0 and σ contains infinitely many occurrences of transitions from $T \setminus J$. By Lemma 5.5 we can define functions β, γ, δ, λ, ξ, η and ζ from N to T^*, a function μ from N to $\mathcal{M}$ and functions θ and ρ from N to T^∞ as follows. Firstly, $\beta(0) = \varepsilon$, $\mu(0) = M$, $\gamma(0) = \varepsilon$, $\delta(0) = \varepsilon$, $\lambda(0) = \varepsilon$, $\xi(0) = \varepsilon$, $\eta(0) = \varepsilon$, $\zeta(0) = \varepsilon$, $\theta(0) = \sigma$ and $\rho(0) = \sigma$.

Let then $k \in N$. We choose $\beta(k+1)$, $\mu(k+1)$, $\gamma(k+1)$, $\delta(k+1)$, $\lambda(k+1)$, $\xi(k+1)$, $\eta(k+1)$, $\zeta(k+1)$, $\theta(k+1)$ and $\rho(k+1)$ in such a way that $\xi(k+1)\rho(k+1) = \theta(k)$, $\xi(k+1) \in T^* \setminus J^*$, $\gamma(k+1)\theta(k+1) = \rho(k+1)$, $\beta(k+1) = \beta(k)\delta(k+1)$, $\mu(k)[\delta(k+1)\rangle_f \mu(k+1)$, $\eta(k+1) = \eta(k)\xi(k+1)\gamma(k+1)$, $\zeta(k+1) = \eta(k)\xi(k+1)$, $\mu(k+1)[\lambda(k+1)\theta(k+1)\rangle$, $\delta(k+1)$ $(T \setminus J)$-exhausts $\lambda(k)\xi(k+1)$, and $\delta(k+1)\lambda(k+1)$ is Υ-equivalent to $\lambda(k)\xi(k+1)\gamma(k+1)$.

From this definition it follows that for each $k \in N$, $\eta(k+1)\theta(k+1) = \eta(k)\xi(k+1)\gamma(k+1)\theta(k+1) = \eta(k)\xi(k+1)\rho(k+1) = \eta(k)\theta(k)$, $\beta(k+1) = \beta(k)\delta(k+1)$ $(T \setminus J)$-exhausts $\beta(k)\lambda(k)\xi(k+1)$, and $\beta(k+1)\lambda(k+1) = \beta(k)\delta(k+1)\lambda(k+1)$ is Υ-equivalent to $\beta(k)\lambda(k)\xi(k+1)\gamma(k+1)$. So, if $\beta(k)\lambda(k)$ is Υ-equivalent to $\eta(k)$, then $\beta(k+1)$ $(T \setminus J)$-exhausts $\eta(k)\xi(k+1) = \zeta(k+1)$ and $\beta(k+1)\lambda(k+1)$ is Υ-equivalent to $\eta(k)\xi(k+1)\gamma(k+1) = \eta(k+1)$.

By induction we get that for each $k \in N$, $\eta(k)\theta(k) = \sigma$, $\beta(k)$ $(T \setminus J)$-exhausts $\zeta(k)$ and $\beta(k)\lambda(k)$ is Υ-equivalent to $\eta(k)$. On the other hand, $\eta(k) = \zeta(k)\gamma(k)$, $M[\beta(k)\rangle_f \mu(k)$, $\mu(k)[\lambda(k)\theta(k)\rangle$, and $\zeta(k+1)$ contains more occurrences of transitions from $T \setminus J$ than $\zeta(k)$ contains.

Since for any $k \in N$, $\beta(k+1) = \beta(k)\delta(k+1)$, the function β represents an infinite transition sequence that is f-enabled at M. Let ω be this infinite sequence. From above it follows that for any $Y \in \Upsilon$, every finite prefix of the Y-restriction of ω is a finite prefix of the Y-restriction of σ, and every finite prefix of the Y-restriction of σ is a finite prefix of the Y-restriction of ω. The infinite sequence ω is thus Υ-equivalent to the infinite sequence σ. $\square$

We are now ready to collect together the results we have obtained and prove the desired theorem. The task is simple since all the hard work has been done in proving the lemmas. Note that according to A1 – A5, "everything is possibly infinite".

Theorem 5.7 *Assumptions: A1, A2, A3, A4 and A5.*
Claims:

(C1) Every terminal path in the f-reachability graph of the net is a terminal path of the full reachability graph of the net.
(C2) For each terminal path starting from M_0 in the full reachability graph, there exists a terminal path starting from M_0 in the f-reachability graph in such a way that the labels of the paths are Υ-equivalent.
(C3) For each finite path starting from M_0 in the full reachability graph, there exists a finite path starting from M_0 in the f-reachability graph in such a way that the labels of the paths are Υ-equivalent.
(C4) For each infinite path starting from M_0 in the full reachability graph, there exists an infinite path starting from M_0 in the f-reachability graph in such a way that the labels of the paths are Υ-equivalent.

Proof. C1 follows trivially from D2. C2 is an immediate consequence of Lemma 5.3. C3 follows directly from Lemma 5.5, by letting $\rho' = \varepsilon$.

From Lemma 5.5, by letting $\rho' \in J^\infty$, and from Lemma 5.4 it directly follows that C4 holds when restricted to a path where some suffix of the label of the path is in J^∞. From Lemma 5.6 it immediately follows that C4 holds when restricted to a path where no suffix of the label is in J^∞. □

As we see from Proposition 5.1, C3 is actually not needed in our LTL verification problem. However, C3 is interesting by its own virtue, at least if Υ-equivalence is thought of as a behavioural equivalence.

6 Treating operation fairness

We now consider verification under the assumption of operation fairness. In order to guarantee that operation fair paths are sufficiently retained in a reduction, we extend the assumptions A1–A5 by the following assumption A6 and then drop assumption A4 since A4 and A6 together would simply force us to generate the full reachability graph.

(A6) Let ϖ be a function from T to 2^T such that for each $t \in T$, $\{t' \in T \mid \exists s \in {}^{\bullet}t \;\; W(s,t') \neq W(t',s)\} \subseteq \varpi(t)$. Then $\mathcal{T} \cup \mathcal{Y} \subseteq \Upsilon$ where $\mathcal{T} = \{\{t\} \mid t \in T\}$ and $\mathcal{Y} = \{\varpi(t) \mid t \in T\}$.

From A6 it follows that all transitions are visible. We observe that if two infinite paths in the full reachability graph start from the same marking and have $(\mathcal{T} \cup \mathcal{Y})$-equivalent labels, then both of the paths are operation fair or neither of them is operation fair. The set $\varpi(t)$ contains at least all those transition that are "visible w.r.t. the enabledness of t", where visibility is understood in the same way as in the discussion after Proposition 5.1. The separation of $\mathcal{T}$ and $\mathcal{Y}$ reflects the fact that the definition of operation fairness does not say anything about what should happen if a transition is enabled at most finitely many times. If we had defined an action-oriented version of LTL [11, 16], operation fairness could have been expressed as an ordinary formula (except possibly in the case that the set of transitions is infinite), and A6 would have been obtained as a side effect of the ordinary construction principles of Υ.

If we return to the example concerning the net in Figure 1, we see that the sequences $(abcfdg)(abcfdg)(abcfdg)\ldots$ and $(abfcdg)(abfcdg)(abfcdg)\ldots$ are not $\mathcal{Y}$-equivalent since both of c and f must be in $\varpi(e)$.

Note that $\mathcal{Y}$-equivalence does not imply $\mathcal{T}$-equivalence. If a net has transitions but no place, we can let $\varpi(t) = \emptyset$ for each transition t, with the consequence that any two transition sequences are $\mathcal{Y}$-equivalent.

Lemma 6.1 is much like Lemma 5.5. The difference is that assumption A4 has been replaced by assumption A6, the sequence to be transformed is definitely infinite and a label of an operation fair path, the result of the transformation is a permutation of the original sequence, and the prefix covering condition has been fixed according to the fact that $J = \emptyset$.

Lemma 6.1 *Assumptions: A1, A2, A3, A5 and A6.*

Claim: For each $\sigma'' \in T^*$, for each $\rho' \in T^\infty$ and for each $M' \in \mathcal{M}$, if $M'[\sigma''\rho'\rangle$, M' is f-reachable from M_0 and the path starting from M' and being labelled by $\sigma''\rho'$ in the full reachability graph is operation fair, then there exist $\gamma'' \in T^*$, $\delta_1 \in T^*$, $\delta_2 \in T^*$, $\rho'' \in T^* \cup T^\infty$ and $M'' \in \mathcal{M}$ in such a way that $\gamma''\rho'' = \rho'$, $M'[\delta_1\rangle_f M''$, $M''[\delta_2\rho''\rangle$, δ_1 T-exhausts σ'', and $\delta_1\delta_2$ is Υ-equivalent to $\sigma''\gamma''$.

Proof. Taking into account that $J = \emptyset$, it suffices to repeat the proof of Lemma 5.5 literally, with the following modifications: the induction hypothesis must refer to the claim of the lemma being proved, a trivial observation concerning operation fairness is needed in the induction step, and the reference to A4 has to be replaced by a reference to D2 and operation fairness. (D2 and operation fairness together guarantee that we never get into a situation where a transition from J would be needed.) □

Lemma 6.2 is a similar continuation to Lemma 6.1 as Lemma 5.6 is to Lemma 5.5.

Lemma 6.2 *Assumptions: A1, A2, A3, A5 and A6.*

Claim: For each $\sigma'' \in T^\infty$ for each $M'' \in \mathcal{M}$, if $M''[\sigma''\rangle$, M'' is f-reachable from M_0, and the path starting from M'' and being labelled by σ'' in the full reachability graph is operation fair, then there exists $\delta'' \in T^\infty$ in such a way that $M''[\delta''\rangle_f$ and δ'' is Υ-equivalent to σ''.

Proof. Taking into account that $J = \emptyset$, it suffices to repeat the proof of Lemma 5.6 literally, with the following modifications: a trivial observation concerning operation fairness is needed before applying any inductive argument, and the reference to Lemma 5.5 must be replaced by a reference to Lemma 6.1. □

We are now ready to present a preservation theorem for operation fair paths.

Theorem 6.3 *Assumptions: A1, A2, A3, A5 and A6.*
Claims: C1 and C2 of Theorem 5.7 and

(C5) For each operation fair infinite path starting from M_0 in the full reachability graph, there exists an operation fair infinite path starting from M_0 in the f-reachability graph in such a way that the labels of the paths are Υ-equivalent.

Proof. Again, C1 follows trivially from D2 and C2 is an immediate consequence of Lemma 5.3. C5 in turn follows directly from A6, Lemma 6.2 and the definition of operation fairness. □

Though $(\mathcal{T} \cup \mathcal{Y})$-equivalence preserves operation fairness and its negation, Theorem 6.3 does not promise anything that would concern the paths that are not operation fair. Let Υ' be an arbitrary subset of 2^T, thus not required to satisfy A6. By substituting Υ' for Υ in Theorem 5.7 and $\Upsilon' \cup \mathcal{T} \cup \mathcal{Y}$ for Υ in Theorem 6.3, one could present a corollary to be applied when operation fair counterexamples are expected but a total absence of operation fair counterexamples makes any counterexample acceptable. However, nothing prevents us from simply verifying a formula first under fairness assumptions and then without fairness assumptions. Such a simple approach is even recommendable since retaining several less than strictly related things during a single state space construction is one of the most typical ways to promote state space explosion.

Theorem 6 is effectively so close to the corresponding theorems in [16, 17] that we have not essentially improved the stubborn set method in verification under fairness assumptions.

7 Conclusions

This paper has considered relieving of the state space explosion problem that occurs in the analysis of concurrent and distributed systems. We have concentrated on one method for that purpose: the stubborn set method. We are fully aware of the fact that the stubborn set method has no special position among verification heuristics. It is also clear that in industrial-size cases, one method alone is typically almost useless. Our motivation is that whenever a method is used, it should be used reasonably.

The contribution of this paper is Theorem 5.7 that gives us a way to utilize the structure of the formula when the stubborn set method is used but fairness is not assumed. Algorithmic implementations can be derived from this theorem in the same way as in [23].

The tester approach in [26] can be considered more goal-oriented than our approach, but so far we have not found any automatic way to construct a useful tester for an arbitrary formula. In [13], a visibility relaxation heuristic for improving the tester technique is presented and the heuristic is shown to apply very well to automatically constructible Büchi automata, too. However, this relaxation technique does not cover our approach. Let us consider a verification task where we need a Büchi automaton that accepts exactly the sequences that satisfy an $n$ary conjunction of formulas. (The formula to be verified then corresponds to an $n$ary disjunction. If an $n$ary conjunction were to be verified, we could verify it simply by verifying its conjuncts separately.) As can be seen from the construction description [7, 13] and from Lemma 6 of [13], all conjuncts become represented in every state of the automaton. Consequently, the visibility relaxation heuristic in [13] does not take any obvious advantage of the fact that the $n$ary conjunction in question is a Boolean combination.

The use of stubborn sets in various formalisms and logics is a fruitful area of future research. On the other hand, we should, by means of large case studies, try to find out what the central problems in the application of the method are and how these problems could be alleviated.

Acknowledgements

This work has been supported by Helsinki Graduate School in Computer Science and Engineering and The Academy of Finland. Discussions with Associate Professor Antti Valmari from Tampere University of Technology have affected the paper in the following things: the assumption A4 (in Section 5) and the idea of intersecting an automatically constructed Büchi automaton with a reduced state space during the construction of the latter.

References

1. Bause, F.: *Analysis of Petri Nets with a Dynamic Priority Method.* Azéma, P., and Balbo, G. (Eds.), Proceedings of PN '97, Toulouse. LNCS 1248, Springer-Verlag 1997, pp. 215–234.
2. Courcoubetis, C. (Ed.): Proceedings of CAV '93, Elounda, Greece. LNCS 697, Springer-Verlag 1993, 504 p.
3. Dembiński, P., and Średniawa, M. (Eds.): Proceedings of PSTV '95, Warsaw. Chapman & Hall 1996, 453 p.
4. Emerson, E.A.: *Temporal and Modal Logic.* van Leeuwen, J. (Ed.), Handbook of Theoretical Computer Science, Vol. B. Elsevier 1990, pp. 995–1072.
5. Francez, N.: *Fairness.* Springer-Verlag 1986, 295 p.
6. Gerth, R., Kuiper, R., Peled, D., and Penczek, W.: *A Partial Order Approach to Branching Time Logic Model Checking.* Proceedings of the 3rd Israel Symposium on Theory of Computing and Systems, Tel Aviv 1995. IEEE Computer Society Press, Los Alamitos CA 1995, pp. 130–140.
7. Gerth, R., Peled, D., Vardi, M.Y., and Wolper, P.: *Simple On-the-Fly Automatic Verification of Linear Temporal Logic.* In [3], pp. 3–18.
8. Godefroid, P.: *Partial-Order Methods for the Verification of Concurrent Systems — An Approach to the State-Explosion Problem.* LNCS 1032, Springer-Verlag 1996, 143 p.

9. Godefroid, P., and Pirottin, D.: *Refining Dependencies Improves Partial-Order Verification Methods.* In [2], pp. 438–449.
10. Janicki, R., and Koutny, M.: *Using Optimal Simulations to Reduce Reachability Graphs.* Clarke, E.M., and Kurshan, R.P. (Eds.), Proceedings of CAV '90, New Brunswick NJ. LNCS 531, Springer-Verlag 1991, pp. 166–175.
11. Kaivola, R.: *Equivalence, Preorders and Compositional Verification for Linear Time Temporal Logic and Concurrent Systems.* Doctoral thesis, University of Helsinki, Department of Computer Science, Report A-1996-1, 1996, 185 p.
12. Katz, S., and Peled, D.: *Verification of Distributed Programs Using Representative Interleaving Sequences.* Distributed Computing 6 (1992) 2, pp. 107–120.
13. Kokkarinen, I., Peled, D., and Valmari, A.: *Relaxed Visibility Enhances Partial Order Reduction.* Grumberg, O. (Ed.), Proceedings of CAV '97, Haifa. LNCS 1254, Springer-Verlag 1997, pp. 328–339.
14. Koutny, M., and Pietkiewicz-Koutny, M.: *On the Sleep Sets Method for Partial Order Verification of Concurrent Systems.* University of Newcastle upon Tyne, Department of Computing Science, Technical Report 495, Newcastle upon Tyne 1994.
15. Overman, W.T.: *Verification of Concurrent Systems: Function and Timing.* PhD thesis, University of California at Los Angeles 1981, 174 p.
16. Peled, D.: *All from One, One for All: on Model Checking Using Representatives.* In [2], pp. 409–423.
17. Peled, D.: *Combining Partial Order Reductions with On-the-Fly Model-Checking.* Formal Methods in System Design 8 (1996) 1, pp. 39–64.
18. Peled, D., and Penczek, W.: *Using Asynchronous Büchi Automata for Efficient Automatic Verification of Concurrent Systems.* In [3], pp. 115–130.
19. Ramakrishna, Y.S., and Smolka, S.A.: *Partial-Order Reduction in the Weak Modal Mu-Calculus.* Mazurkiewicz, A., and Winkowski, J. (Eds.), Proceedings of CONCUR '97, Warsaw. LNCS 1243, Springer-Verlag 1997, pp. 5–24.
20. Rauhamaa, M.: *A Comparative Study of Methods for Efficient Reachability Analysis.* Helsinki University of Technology, Digital Systems Laboratory Report A 14, 1990, 61 p.
21. Reisig, W.: *Petri Nets: An Introduction.* EATCS Monographs on Theoretical Computer Science 4, Springer-Verlag 1985, 161 p.
22. Sloan, R.H., and Buy, U.: *Stubborn Sets for Real-Time Petri Nets.* Formal Methods in System Design 11 (1997) 1, pp. 23–40.
23. Valmari, A.: *A Stubborn Attack on State Explosion.* Formal Methods in System Design 1 (1992) 4, pp. 297–322.
24. Valmari, A.: *Stubborn Sets of Coloured Petri Nets.* Proceedings of PN '91, Gjern, Denmark, pp. 102–121.
25. Valmari, A.: *Alleviating State Explosion during Verification of Behavioural Equivalence.* University of Helsinki, Department of Computer Science, Report A-1992-4, 1992, 57 p.
26. Valmari, A.: *On-the-Fly Verification with Stubborn Sets.* In [2], pp. 397–408.
27. Varpaaniemi, K.: *On Combining the Stubborn Set Method with the Sleep Set Method.* Valette, R. (Ed.), Proceedings of PN '94, Zaragoza. LNCS 815, Springer-Verlag 1994, pp. 548–567.
28. Vernadat, F., Azéma, P., and Michel, F.: *Covering Step Graph.* Billington, J., and Reisig, W. (Eds.): Proceedings of PN '96, Osaka. LNCS 1091, Springer-Verlag 1996, pp. 516–535.
29. Yoneda, T., Shibayama, A., Schlingloff, B.-H., and Clarke, E.M.: *Efficient Verification of Parallel Real-Time Systems.* In [2], pp. 321–332.