

On the Complexity of Consistency and Complete State Coding for Signal Transition Graphs

Javier Esparza*

Institut fuer Informatik
Technische Universitaet Muenchen, Germany
esparza@in.tum.de

Petr Jančar†

Center of Applied Cybernetics
Dept of CS, TU Ostrava, Czechia
Petr.Jancar@vsb.cz

Alexander Miller

Institute for Formal Methods in Computer Science
Univ. Stuttgart, Germany
Alexander.Miller@informatik.uni-stuttgart.de

Abstract. Signal Transition Graphs (STGs) are a popular formalism for the specification of asynchronous circuits. A necessary condition for the implementability of an STG is the existence of a consistent and complete state encoding. For an important subclass of STGs, the marked graph STGs, we show that checking consistency is polynomial, but checking the existence of a complete state coding is co-NP-complete. In fact, co-NP-completeness already holds for acyclic and 1-bounded marked graph STGs and for live and 1-bounded marked graph STGs. We add some relevant results for free-choice, bounded, and general STGs.

Keywords: signal transition graphs, consistency, complete state coding, Petri nets, complexity

1. Introduction

Signal transition graphs (STGs) are a popular formalism for specifying asynchronous circuits [3, 12]. They are Petri nets in which the firing of a transition is interpreted as rising or falling of a signal in the circuit. Not every STG can be implemented as a physical circuit. A central question related to implementability of an STG is whether it admits a so-called *consistent* and *complete state coding*. Most

*This author was affiliated with the University of Stuttgart when this work was done.

†This author is supported by the Czech Ministry of Education, Grant No. 1M0567

papers in the literature consider only the completeness part, assuming that the STG is already consistent, and call the existence of a complete state coding the *CSC property*. This property, and the stronger unique state coding property (*USC property* for short) have been studied in many papers (see e.g. [1, 9, 10, 11, 14, 16, 17]).

In this paper we reason about the computational complexity of deciding if a given STG has a consistent and complete state coding, viewing the consistency problem separately. We obtain new results for STGs whose underlying nets are marked graphs and free-choice nets; for completeness, we also sketch some straightforward results for STGs with more general underlying nets (bounded or even arbitrary).

We first explore the consistency problem for marked graph STGs. In [7] a polynomial algorithm was given to check consistency of live, bounded, and cyclic free-choice STGs, which include live and bounded marked graph STGs as a subclass. Here we show that consistency is polynomial for arbitrary marked graph STGs by means of a new algorithm based on linear programming.

A natural question is whether these polynomiality results also hold for the CSC or USC problems (i.e., the problems of checking the CSC or USC properties), at least for the class of live and 1-bounded marked graph STGs. Our main result shows that both problems are co-NP-complete, and so that polynomial algorithms are unlikely. This result explains why the algorithms of [1, 9, 10, 11, 14, 16, 17] have exponential runtime or can only decide some necessary or sufficient conditions for the CSC or USC properties to hold. These algorithms are discussed in detail in the final section.

Our co-NP-completeness result is rather robust. We prove that the CSC and USC problems remain co-NP-hard for 1-bounded and acyclic marked graph STGs, and that they remain in co-NP for arbitrary marked graph STGs and for live and bounded free-choice STGs.

Moving to more general classes, we show that the consistency, CSC and USC problems are PSPACE-complete for 1-bounded STGs, and that the consistency problem remains PSPACE-hard in the free-choice case. Finally, we clarify the relation between the consistency, USC and CSC problems for general STGs, and the fireability and reachability problems for general Petri nets.

The paper is structured as follows. Section 2 presents basic definitions and a characterization of consistency. Section 3 presents the results about marked-graph STGs; it is the core of the paper. Section 4 deals with free-choice and Section 5 with general STGs. Section 6 contains conclusions and discusses related work.

Remark. A preliminary version of this paper appeared as the conference paper [8].

2. Basic definitions

A *net* is a triple (P, T, F) , where P and T are disjoint sets of *places* and *transitions*, respectively, and F is a function $(P \times T) \cup (T \times P) \rightarrow \{0, 1\}$. Places and transitions are generically called *nodes*; we also note that a net can be viewed as a (bipartite) graph. Places are graphically represented as circles; transitions are usually drawn like boxes, but we just use their labels in the figures. If $F(x, y) = 1$ then we say that there is an *arc* from x to y . The *preset* of a node x , denoted by $\bullet x$, is the set of its *input nodes*, i.e., the set $\{y \in P \cup T \mid F(y, x) = 1\}$. The *postset* of x , denoted by x^\bullet , contains its *output nodes*, i.e., the set $\{y \in P \cup T \mid F(x, y) = 1\}$.

A *marking* M of a net (P, T, F) is a mapping $P \rightarrow \mathbb{N}$ (where \mathbb{N} denotes the set of natural numbers including 0). Graphically, a marking is represented by drawing $M(p)$ tokens on the circle representing the place p . A marking M *enables* a transition t if it puts at least one token on each place $p \in \bullet t$, i.e.,

if $M(p) \geq 1$ for each $p \in \bullet t$. If t is enabled at M , then it can *fire* (or *occur*) and its firing (occurrence) *leads to* a new marking M' , obtained by removing a token from each place in the preset of t , and adding a token to each place in its postset; formally, $M'(p) = M(p) + F(t, p) - F(p, t)$ for every place p . $M \xrightarrow{t}$ denotes that t is enabled at M , and $M \xrightarrow{t} M'$ moreover denotes that firing t leads to M' .

The notation $M \xrightarrow{\sigma}$, $M \xrightarrow{\sigma} M'$ is extended to finite sequences $\sigma \in T^*$ in the natural way. When $M \xrightarrow{\sigma} M'$, for $\sigma = t_1 t_2 \dots t_n$, we speak of an *occurrence sequence from M to M'* , meaning the sequence

$$M \xrightarrow{t_1} M_1 \xrightarrow{t_2} \dots M_{n-1} \xrightarrow{t_n} M'$$

By the *Parikh vector* of $\sigma \in T^*$, denoted by $\vec{\sigma}$ or $P(\sigma)$, we mean the mapping (or the corresponding vector) $T \rightarrow \mathbb{N}$ such that $\vec{\sigma}(t)$ is the number of occurrences of t in σ .

The *incidence matrix* of N is the matrix $C_N: P \times T \rightarrow \{-1, 0, +1\}$ given by $C_N(p, t) = F(t, p) - F(p, t)$. We note that if $M \xrightarrow{\sigma} M'$ then $M + C_N \cdot \vec{\sigma} = M'$.

A *Petri net* is a pair (N, M_0) where N is a net and M_0 is a marking of N , called the *initial marking*. A marking M is called *reachable* if there exists an occurrence sequence from M_0 to M ; we also denote this by $M_0 \rightarrow^* M$. We call

$$M_0 + C_N \cdot X \geq 0$$

the *marking inequation*. We note that $M_0 \xrightarrow{\sigma} M$ implies $M_0 + C_N \cdot \vec{\sigma} = M$; $\vec{\sigma}$ is thus a (nonnegative integer) solution of the marking inequation.

A marking M of a net N is *n-bounded* if $M(p) \leq n$ for every place p . A Petri net (N, M_0) is *n-bounded* if all its reachable markings are *n-bounded*.

A transition t is *fireable* in (N, M_0) if there is σ such that $M_0 \xrightarrow{\sigma} M$ and $M \xrightarrow{t}$. A Petri net (N, M_0) is *live* if each transition t is fireable in (N, M) for each M reachable from M_0 . A transition is *dead* at a marking M if t is not fireable in (N, M) .

A net N is called a *marked graph* if every place has at most one input and at most one output transition. $N = (P, T, F)$ is a *free-choice* net if: for each place p and every transition t , if $F(p, t) = 1$ then $F(p', t') = 1$ for every $p' \in \bullet t$, $t' \in p^\bullet$. In a free-choice net, if some output transition of a place is enabled at a marking, then all its output transitions are enabled, and it is possible to “freely” choose among them.

Signal transition graphs. Let $A = \{a_1, \dots, a_n\}$ be a set (alphabet) of *signals* partitioned into *input* and *output* signals. Rising and falling of a signal a is denoted by a^+ and a^- , respectively. (In some proofs we also use the notation $+a$ and $-a$, which is more convenient for using sub- and superscripts.) We call an element of $\mathcal{L} = A \times \{+, -\}$ a *label*. A *signal transition graph* (STG) is a triple $S = (N, M_0, \ell)$, where (N, M_0) is a Petri net and ℓ is a *labelling function* that assigns to each transition of N a label in \mathcal{L} .

A signal transition graph is a specification of the behaviour of the circuit under some assumptions on the environment. An STG S is implementable if there exists a *state coding mapping* λ (we also use the term *binary encoding*) that associates with each reachable marking M a vector of *signal values* $\lambda(M) \in \{0, 1\}^n$ satisfying the following two properties:

- (1) *Consistency.* If $M \xrightarrow{t} M'$ and t is labelled by a_i^+ , then the i -th components of $\lambda(M)$ and $\lambda(M')$ are 0 and 1, respectively, and all other components have the same value in $\lambda(M)$ and $\lambda(M')$. If t

is labelled by a_i^- , then the i -th components of $\lambda(M)$ and $\lambda(M')$ are 1 and 0, respectively, and all other components have the same value in $\lambda(M)$ and $\lambda(M')$.

- (2) *Completeness*: if two different reachable markings M, M' satisfy $\lambda(M) = \lambda(M')$, then they enable exactly the same output labels.

Consistency is obviously necessary for implementability. Completeness is necessary because the state of an implementation is completely determined by the signal values of all signals. Therefore, if some output signal is enabled at M but not at M' , then M and M' must correspond to different states of the implementation, and so they must differ in the value of at least one signal.

We define the *consistency problem* as the problem of deciding if a given STG is consistent, i.e., if it admits a binary encoding λ satisfying (1). The Complete State Coding problem, *CSC problem* for short, is the problem to decide if a given STG (usually already assumed consistent) has the *CSC property*, i.e., admits a binary encoding satisfying (1) and (2). A stricter version is the *USC problem* (unique state coding) where we ask if a given STG has the *USC property*, i.e., admits an *injective* binary encoding λ satisfying (1) (thus $\lambda(M) \neq \lambda(M')$ for any two different reachable M, M').

STGs naturally inherit many notions from their underlying (Petri) nets. We already used this when speaking about ‘enabling a label’, e.g. $M \xrightarrow{a^+}$ (meaning that M enables a transition with label a^+). Thus we will freely speak about n -bounded, live, marked graph, or free-choice STGs, etc. We can also use notions like a is dead at M (meaning that each transition with label a^+ or a^- is dead at M).

We also freely use notation like $M \xrightarrow{u} M'$ for sequences of labels (meaning that there is a transition sequence $\sigma = t_1 t_2 \cdots t_m$ such that $M \xrightarrow{\sigma} M'$ and $u = \ell(t_1)\ell(t_2) \cdots \ell(t_m)$). We can occasionally even mix, and consider u as a sequence of transitions and labels, when this should not cause confusion. We also use expressions like u is a -free, meaning that there is no a^+ nor a^- in u ; and if u contains transitions, we mean that those transitions do not have labels a^+, a^- . Recall that $P(u)$ denotes the Parikh vector of u ; We denote by $P(u)(a^+)$ the number of transitions with label a^+ in u .

Finally we note that since the circuit implementation of an STG can be seen as a finite object with at most 2^n states, where n is the number of signals, STGs used in practice are bounded, most of them are even 1-bounded; but in principle unbounded STGs can make sense.

We finish the section by a characterization of consistency, i.e., we look in more detail on when an STG is inconsistent.

Proposition 2.1. An STG $S = (N, M_0, \ell)$ is inconsistent (i.e., it admits no consistent binary encoding) iff there is

a pair (M, a) where $M_0 \xrightarrow{*} M$ and a is a signal

such that one of the following conditions holds:

- (1) M enables ua^+ and va^-
for some a -free sequences u, v ,
- (2) M enables a^+ua^+ or a^-ua^-
for some a -free sequence u ,
- (3) M is reachable by w_1a^+u and by w_2a^-v
for some a -free sequences u, v (and some w_1, w_2).

Proof:

If there is a pair (M, a) such that one of the conditions holds then S is obviously inconsistent.

If there is no such pair then we can (soundly) define the following (partial) encoding λ :
For each reachable M and signal a we put

- $\lambda(M)(a) = 0$ if M enables ua^+ for an a -free sequence u ,
- $\lambda(M)(a) = 1$ if M enables ua^- for an a -free sequence u .

We note that if $\lambda(M)(a)$ is (so far) undefined then M is a -dead; we then put

- $\lambda(M)(a) = 1$ when M can be reached by wa^+u for some a -free sequence u ,
- $\lambda(M)(a) = 0$ when M can be reached by wa^-u for some a -free sequence u ,
- $\lambda(M)(a) = \lambda(M_0)(a)$ otherwise (i.e., when M is reachable only by a -free sequences); if $M = M_0$ we define $\lambda(M_0)(a) = 0$ (we could use $\lambda(M_0)(a) = 1$ as well here).

One can easily check that λ is a consistent binary encoding. □

3. Marked graphs

In this section we show that consistency can be decided in polynomial time for all marked graph STGs and that both the CSC problem and the USC problem are co-NP-complete for them, even in the case of 1-bounded acyclic marked graphs and in the case of live 1-bounded marked graphs.

3.1. Consistency

In [7] it is shown that consistency of live, bounded, and cyclic free-choice STGs can be decided in polynomial time. (A Petri net is *cyclic* if the initial marking is reachable from every reachable marking, i.e., if it is always possible to return to the initial marking). Since live and bounded marked graphs are always cyclic (see for instance [4]), and marked graphs are a special case of free-choice nets, [7] provides a polynomial algorithm deciding consistency of live and bounded marked graph STGs. We now show a polynomial algorithm for all marked graph STGs.

We start by recalling some simple properties of marked graphs and derive a simpler variant of Proposition 2.1, valid for marked graphs. One such property is that if M enables a sequence with n occurrences of t and $M \xrightarrow{t'} M'$ for $t' \neq t$ then M' enables a sequence with n occurrences of t as well; if $t' = t$ then M' enables a sequence with $n-1$ occurrences of t .

By $P(u)(t)$ we denote the number of occurrences of t in a transition sequence u (P stands for the Parikh vector).

Claim 3.1. Let M be a marking of a marked graph. If $M \xrightarrow{u} M_1$ and $M \xrightarrow{v} M_2$ then $M \xrightarrow{w} M'$ for some w and M' such that

$$\forall t : P(w)(t) = \max\{P(u)(t), P(v)(t)\}.$$

Moreover, if $M_1 \xrightarrow{t}$ and $P(v)(t) \leq P(u)(t)$ then $M' \xrightarrow{t}$.

Proof:

We can perform u and then a sequence consisting of $P(v)(t) - P(u)(t)$ occurrences of each t for which $P(u)(t) < P(v)(t)$. \square

Slightly abusing notation, by $\max(u, v)$ we will denote the w guaranteed by the claim.

Proposition 3.1. A marked graph STG $S = (N, M_0, \ell)$ is inconsistent iff one of the following conditions holds:

(1') there is a reachable M ($M_0 \rightarrow^* M$) such that

$$M \xrightarrow{a^+} \text{ and } M \xrightarrow{a^-}$$

(2') there is a reachable M such that

$$M \xrightarrow{a^+ua^+} \text{ or } M \xrightarrow{a^-ua^-}$$

for some signal a and some a -free sequence u .

Proof:

If (1') or (2') holds then S is obviously inconsistent.

Now assume that S is inconsistent. Then we know that there is a reachable M and a signal a such that one of the conditions (1), (2), (3) of Proposition 2.1 holds. It is sufficient to show that this implies (1') or (2').

If (M, a) satisfies (2) then (2') holds. If (M, a) satisfies (1), i.e. $M \xrightarrow{ua^+}$ and $M \xrightarrow{va^-}$ for a -free sequences u, v , then $M \xrightarrow{\max(u,v)}$ M' and $M' \xrightarrow{a^+}$, $M' \xrightarrow{a^-}$ (recall Claim 3.1); thus (1') holds.

We finish by deriving a contradiction from the assumption that the inconsistency of S can not be shown by using (1) nor (2) while we have (M, a) satisfying (3). Hence

$$M_0 \xrightarrow{w_1 t_1 u} M \text{ and } M_0 \xrightarrow{w_2 t_2 v} M,$$

where $\ell(t_1) = a^+$, $\ell(t_2) = a^-$, and u, v are a -free.

Necessarily, all transitions labelled by a^+ or a^- are dead in M —otherwise there would exist (M', a) satisfying (2). Thus both t_1, t_2 are dead in M , which means that w_1 contains the maximal possible number of occurrences of t_2 , while w_2 contains the maximal possible number of occurrences of t_1 .

Let $w_1 = u_1 t_2 v_1$ where v_1 is t_2 -free. Similarly $w_2 = u_2 t_1 v_2$ where v_2 is t_1 -free. We note that $P(u_2)(t_2) \leq P(u_1)(t_2)$, and $P(u_1)(t_1) \leq P(u_2)(t_1)$.

Hence $M_0 \xrightarrow{\max(u_1, u_2)}$ M' where M' enables both t_1 and t_2 , so we have both $M' \xrightarrow{a^+}$ and $M' \xrightarrow{a^-}$. Thus (M', a) satisfies (1)—a contradiction. \square

It is now sufficient to show that conditions (1'), (2') of Proposition 3.1 can be checked in polynomial time.

To this aim, we recall further useful observations about marked graphs. We note that, given a marked graph STG $S = (N, M_0, \ell)$, we can check in polynomial time if there is a cycle in N which is not marked at M_0 (i.e., its places have no tokens in M_0). The places of such a cycle can be safely removed, since no transition in the cycle can ever occur.

We call a marked graph (N, M_0) *normalized* if every cycle in N is marked at M_0 .

Claim 3.2. Let (N, M_0) be a normalized marked graph, and consider the inequation $M_0 + C_N \cdot X \geq 0$, where C_N is the incidence matrix of N . An integer vector $X_0 \geq 0$ is a solution of this inequation if and only if $M_0 \xrightarrow{\sigma}$ for a transition sequence σ whose Parikh vector is X_0 . Moreover, if $M_0 \xrightarrow{\sigma} M$ then $M_0 + C_N \cdot X_0 = M$.

Proof:

The only nonobvious claim is that a solution X_0 implies the existence of an appropriate σ . But this can be done easily by induction on $|X_0|$: among the transitions t with $X_0(t) \geq 1$, some must be enabled at M_0 —otherwise we would find a cycle unmarked in M_0 ; we can fire such a transition and use the induction hypothesis. \square

Now we come to the polynomiality claims, which can be quickly established by using linear programming (which is a well-known polynomial problem).

Proposition 3.2. For normalized marked graph STGs, checking (1') of Proposition 3.1 can be done in polynomial time.

Proof:

Let $S = (N, M_0, \ell)$ be a normalized marked graph STG. If S satisfies (1') then N contains transitions t_1, t_2 with $\ell(t_1) = a^+$ and $\ell(t_2) = a^-$ such that there is σ for which

$$M_0 \xrightarrow{\sigma} M, M \xrightarrow{t_1}, M \xrightarrow{t_2}.$$

We note that $M \geq M_{t_1} + M_{t_2}$ where M_t denotes the marking that puts one token in each input place of t and no tokens elsewhere. The Parikh vector of σ is thus a solution of the linear inequation

$$M_0 + C_N \cdot X \geq M_{t_1} + M_{t_2}.$$

On the other hand, if the inequation has a nonnegative, rational solution X_0 then the integer vector $\lfloor X_0 \rfloor$ is also a solution, as one can easily check. Claim 3.2 then guarantees the existence of an appropriate σ , meaning that S satisfies (1').

Thus checking (1') can be done by solving the inequations for all appropriate pairs t_1, t_2 . \square

Proposition 3.3. For normalized marked graph STGs not satisfying (1'), checking (2') can be done in polynomial time.

Proof:

Let $S = (N, M_0, \ell)$ be a normalized marked graph STG which does not satisfy (1'); i.e., no reachable M can enable both a^+ and a^- . From this we can derive that (M_0, a) does not satisfy (1) of Proposition 2.1. Therefore, in every occurrence sequence containing occurrences of the signal a , the first occurrence of a always has the same sign. Which sign this is, $+$ or $-$, can be determined very efficiently, e.g. by firing any maximal transition sequence in which each transition of S occurs at most once (such a sequence contains all transitions that can ever be enabled).

Consider signal a , and assume we have found that a^+ is fireable as the first of a^+, a^- . (The case with a^- being the first is similar.)

Let us now solve the linear programming problems

$$\begin{array}{ll}
\text{maximize} & f(X) \\
\text{subject to} & X \geq 0, M_0 + C_N \cdot X \geq 0, \\
\text{minimize} & f(Y) \\
\text{subject to} & Y \geq 0, M_0 + C_N \cdot Y \geq 0,
\end{array}$$

where

$$f(X) = \sum_{t \in \ell^{-1}(a^+)} X(t) - \sum_{t \in \ell^{-1}(a^-)} X(t).$$

If we find that it is NOT the case that both problems have optimal solutions X_{op}, Y_{op} with $f(X_{op}) = 1$ and $f(Y_{op}) = 0$ then we claim ‘(2’) holds’.

To check (2’), we run the above procedure for each signal a separately, and claim that (2’) holds when one signal gives rise to this claim, otherwise we claim that (2’) does not hold. The overall time of this algorithm is surely polynomial; it remains to show its correctness.

Let us again consider signal a where a^+ is fireable as the first (of a^+, a^-). It is obvious that condition (2’) holds for signal a iff there is a transition sequence σ , with Parikh vector X , such that $M_0 \xrightarrow{\sigma}$ and $f(X) \geq 2$ (there are two a^+ ’s without any a^- in-between) or $f(X) \leq -1$ (two a^- ’s without any a^+ in-between).

So if (2’) holds for signal a then the procedure for a surely gives rise to the claim ‘(2’) holds’.

If (2’) does not hold for a then we have $f(X) \in \{0, 1\}$ for each *integer* admissible solution X (due to Claim 3.2). We want to show that the procedure for a finds some optimal solutions X_{op}, Y_{op} with $f(X_{op}) = 1$ and $f(Y_{op}) = 0$ (and thus does not give rise to the claim ‘(2’) holds’).

To see this, we recall that all solutions of $X \geq 0, M_0 + C_N \cdot X \geq 0$ constitute a polyhedron. The optimal solutions X_{op}, Y_{op} exist if and only if $f(X)$ is bounded from above and from below on the polyhedron, and then such solutions can be found in the extremal points. The fact that $f(X) \in \{0, 1\}$ for all integer X easily implies that $f(X)$ is bounded for all (admissible) X ; thus the optimal solutions exist.

We now note that every row of C_N contains at most one +1 and at most one -1, which means that matrix C_N is (totally) unimodular. Hence the extremal points of the polyhedron are integer vectors (cf. e.g. [13]). Thus the procedure for a indeed finds some optimal solutions X_{op}, Y_{op} with $f(X_{op}) = 1$ and $f(Y_{op}) = 0$. \square

Theorem 3.1. Consistency of marked graph STGs can be decided in polynomial time.

Proof:

The polynomial algorithm first normalizes the STG and then uses the algorithms guaranteed by Propositions 3.2 and 3.3 to check if one of the conditions (1’), (2’) of Proposition 3.1 holds. \square

3.2. Complete state coding

In this subsection we show the announced co-NP-completeness results for the CSC problem and the USC problem on (consistent) marked graph STGs.

The next lemma is the main technical result of the paper. We say that an occurrence sequence is *balanced* if for every signal a the sequence contains the same number of occurrences of transitions labelled by a^+ and of transitions labelled by a^- .

Lemma 3.1. The following problem is NP-complete:

Instance: a (consistent) STG $S = (N, M_0, \ell)$ such that (N, M_0) is a 1-bounded, acyclic marked graph.

Question: is there an occurrence sequence $M_0 \xrightarrow{\sigma} M_1 \xrightarrow{\tau} M_2$ of S such that τ is nonempty and balanced?

Proof:

Membership in NP is clear: In any net (N, M_0) which is 1-bounded and acyclic, each transition can appear at most once in any occurrence sequence. So a nondeterministic algorithm can just guess a sequence $\sigma\tau$ of pairwise distinct transitions and verify that it is performable from M_0 and that τ is nonempty and balanced.

The main point is NP-hardness, which we show by a reduction from CNF-SAT. Let φ be a boolean formula in conjunctive normal form

- with m clauses c_1, \dots, c_m ,
- and n variables x_1, \dots, x_n .

(E.g., formula $(x_1 \vee \overline{x_2} \vee x_3) \wedge (x_2 \vee \overline{x_4})$ has 2 clauses and 4 variables.)

Our aim is to show a polynomial construction of a certain STG $S_\varphi = (N, M_0, \ell)$, with (N, M_0) being a 1-bounded acyclic marked graph, so that φ is satisfiable iff S_φ admits $M_0 \xrightarrow{\sigma} M_1 \xrightarrow{\tau} M_2$ for some sequence σ and some nonempty balanced sequence τ .

The construction is based on the fact that there is a truth assignment

$$\mathcal{A} : \{x_1, x_2, \dots, x_n\} \rightarrow \{0, 1\}$$

satisfying φ if and only if there is a *consistent choice of literals*, by which we mean a mapping

$$l : \{c_1, c_2, \dots, c_m\} \rightarrow \{x_1, \overline{x_1}, x_2, \overline{x_2}, \dots, x_n, \overline{x_n}\}$$

attaching to each clause c_i one of its literals, denoted $l(c_i)$, in such a way that $l(c_i) \neq \overline{l(c_j)}$ for all i, j (i.e., it is forbidden that one clause ‘chooses’ x while another clause ‘chooses’ \overline{x}).


We can easily observe that any consistent choice of literals l naturally provides a satisfying truth assignment \mathcal{A} (which can be specified arbitrarily for variables not appearing in the range of l); and any satisfying truth assignment enables to define (maybe several) consistent choices of literals.

We now describe the STG S_φ , providing also informal comments which will ease the later correctness proof. Figure 1 shows the overall structure of S_φ .

We need a few remarks about the notation. We construct $S_\varphi = (N, M_0, \ell)$ where N is an acyclic marked graph. All the minimal elements with respect to the flow relation will be places, and precisely these places will be initially marked (i.e., each will carry one token). We say that there is an *arc from transition t_1 to transition t_2* when there is an (intermediate) place p (initially unmarked) and arcs $t_1 \rightarrow p$, $p \rightarrow t_2$. (This is, in fact, a usual convention which we also use for drawing marked graphs.)

Each symbol of Figure 1 (i.e., each V_T^1, \dots, C_m^P) stands for an acyclic marked graph. The arrow $V_T^1 \rightarrow N_\S$ has the following meaning: V_T^1 has a transition t which is the unique maximal element in V_T^1 (w.r.t. the order induced by the flow relation), N_\S has a transition u which is the unique minimal element in N_\S , and the (overall) net N contains an arc leading from t to u (with an intermediate place—using our convention). The meaning of the other arrows in the structure is analogous.

It will be clear (after we finish the construction) that any complete behaviour of S_φ can be divided into three phases:

Figure 1. The overall structure of S_φ

- I. first, all transitions in $V_T^1, \dots, V_T^n, C_1^N, \dots, C_m^N$ occur,
- II. then all transitions of $N_\$$ follow,
- III. and finally all transitions in $V_F^1, \dots, V_F^n, C_1^P, \dots, C_m^P$ occur.

The complete behaviours of S_φ differ only in the order in which transitions occur in the phases I and III. We proceed to describe the marked graphs corresponding to $N_\$, V_T^1, \dots, V_T^n, C_1^N, \dots, C_m^N$. Since we need to use both sub- and superscripts, we change the notation and write $+a$ and $-a$ instead of a^+ and a^- . The net $N_\$,$ enabled after the whole phase I is finished, has one single (complete) behaviour, shown in Figure 2.

$$+\$ -x^1 -x^2 \cdots -x^n -c_1 -c_2 \cdots -c_m -\$$$

Figure 2. (Linear) behaviour of $N_\$$

This means that the signal set of S_φ contains (among others):

- a signal c_i for every clause ($1 \leq i \leq m$);
- a signal x^j for every variable ($1 \leq j \leq n$);
- a (special) signal $\$$.

Signal $\$$ will not appear anywhere else but in $N_\$$. It will be the case that any nonempty balanced sequence must include all transitions of $N_\$,$ and so such a sequence will necessarily contain the whole phase II.

For the rest of the proof let bal denote any non-empty and balanced sequence such that $M_0 \xrightarrow{\sigma} M_1 \xrightarrow{bal} M_2$. In bal , each falling $-x^j$ ($1 \leq j \leq n$) must be compensated by a raising $+x^j$; the label $+x^j$ will appear just on the maximal (i.e., the last) transition of V_T^j (cf. Figure 3) and on the minimal (i.e., the first) transition of V_F^j (cf. Figure 4). So precisely one of the subnets V_T^j, V_F^j will contribute to bal . We interpret this as ‘choosing’ a truth assignment \mathcal{A} .

Similarly, each falling $-c_i$ ($1 \leq i \leq m$) must be compensated by a raising $+c_i$; the label $+c_i$ will appear just once in C_i^N and once in C_i^P , now ‘almost’ as the last transition and ‘almost’ as the first transition, respectively. Again, exactly one of the subnets C_i^N, C_i^P will contribute to bal .

Now we continue with the details of our construction. We extend the signal set used so far by

- a signal p_i^j for each pair i, j ($1 \leq i \leq m, 1 \leq j \leq n$) such that clause c_i contains literal x_j (p stands for ‘positive’);
- a signal n_i^j for each pair i, j ($1 \leq i \leq m, 1 \leq j \leq n$) such that clause c_i contains literal \bar{x}_j (n stands for ‘negative’).

(As usual, we can assume that no clause c_i of formula φ contains a complementary pair of literals.)

Given j ($1 \leq j \leq n$), let $\{c_{i_1}, c_{i_2}, \dots, c_{i_a}\}$ be the set of clauses containing (the positive) literal x_j . The (sub)net V_T^j (representing setting x_j to ‘true’) is depicted in Figure 3. Thus V_T^j ‘emits’ labels




Figure 3. The net V_T^j

$+p_{i_1}^j, +p_{i_2}^j, \dots, +p_{i_a}^j$ in any order, and then finishes by $+x^j$.

Now let $\{c_{k_1}, c_{k_2}, \dots, c_{k_b}\}$ be the set of clauses containing (the negative) literal \bar{x}_j . The (sub)net V_F^j (representing setting x_j to ‘false’) is depicted in Figure 4. Thus, after the label $-\$$ of $N_\$$ occurs, V_F^j




Figure 4. The net V_F^j

‘emits’ label $+x^j$ and then labels $-n_{k_1}^j, -n_{k_2}^j, \dots, -n_{k_b}^j$ in any order.

We now define the subnets C_i^N, C_i^P . Recall that the sequence bal will contain either transitions of C_i^N or C_i^P , but not of both. This corresponds to ‘choosing’ either a positive or a negative literal $l(c_i)$ from c_i . Which literal is chosen will depend on which transitions of the corresponding net occur in bal , and is explained later.

The nets C_i^N and C_i^P have no concurrency. They use additional ‘parenthetical’ signals. More precisely, we enhance the signal set by

- a signal \square_i^j for each pair $i, j, 1 \leq i \leq m$ and $1 \leq j \leq n_i$, where n_i is the number of negative literals in c_i ;

- a signal Δ_i^j for each pair i, j , $1 \leq i \leq m$ and $1 \leq j \leq p_i$, where p_i is the number of positive literals in c_i .

Given i ($1 \leq i \leq m$), let $\{\overline{x_{j_1}}, \overline{x_{j_2}}, \dots, \overline{x_{j_a}}\}$ be the set of negative literals of the clause c_i .

The (sub)net C_i^N has a (marked) place as the least element (w.r.t. the flow relation). And the only (complete) behaviour of C_i^N is the sequence of labels shown in Figure 5. The key observation is that if

$$\begin{aligned} & +\square_i^1 + n_i^{j_1} - n_i^{j_2} - \square_i^1 + \square_i^2 + n_i^{j_2} - n_i^{j_3} - \square_i^2 \dots \\ \dots & + \square_i^{a-1} + n_i^{j_{a-1}} - n_i^{j_a} - \square_i^{a-1} + \square_i^a + n_i^{j_a} + c_i - \square_i^a \end{aligned}$$

Figure 5. (Linear) behaviour of C_i^N

the label $+c_i$ of C_i^N belongs to the balanced sequence bal , then bal must also contain $-\square_i^a$, and thus, by balancedness, also $+\square_i^a$. But then bal also contains $+n_i^{j_a}$, and so it must also contain $-n_i^{j_a}$. If we add the label $-n_i^{j_a}$ of C_i^N to bal , then we are forced to add $-\square_i^{a-1}$ as well, and thus also $+\square_i^{a-1}$ and $+n_i^{j_{a-1}}$; etc. So if labels of C_i^N occur in bal , then bal contains an occurrence of some $+n_i^j$, where $\overline{x_j}$ is a literal of c_i , such that the ‘balancing’ occurrence of $-n_i^j$ does not come from C_i^N , and so it must come from V_F^j . We interpret this as ‘choosing’ the literal $\overline{x_j}$ of c_i , i.e., as setting $l(c_i) = \overline{x_j}$.

The (sub)net C_i^P is similar. We let $\{x_{k_1}, x_{k_2}, \dots, x_{k_b}\}$ be the set of positive literals of the clause c_i . The least element of C_i^P (w.r.t. the flow relation) is a transition labelled by Δ_i^1 ; it follows from the overall structure that this transition is enabled after $-\$$ occurs. The only (complete) behaviour of C_i^P (after being enabled) is the sequence of labels shown in Figure 6. And we reason similarly as above. If

$$+\Delta_i^1 + c_i - p_i^{k_1} - \Delta_i^1 + \Delta_i^2 + p_i^{k_1} - p_i^{k_2} - \Delta_i^2 + \Delta_i^b + p_i^{k_{b-1}} - p_i^{k_b} - \Delta_i^b$$

Figure 6. (Linear) behaviour of C_i^P

the label $+c_i$ from C_i^P belongs to bal , then bal must also contain $+\Delta_i^1$, and thus also $-\Delta_i^1$ etc. So if labels of C_i^P occur in bal , then bal contains an occurrence of some $-p_i^j$, where x_j is a literal of c_i , such that the ‘balancing’ occurrence $+p_i^j$ does not come from C_i^N , and so it must come from V_T^j . We interpret this as ‘choosing’ the literal x_j of c_i , i.e., as setting $l(c_i) = x_j$.

For illustration, Figure 7 shows the STG S_φ for a simple formula φ . In fact, it is a slightly different variant of S_φ , using more concurrency and additional signals T^j, F^j which stress the correspondence with setting variable x^j to true or false. (The shaded region shows a balanced sequence, which will be also discussed later.)

We have thus completed the (obviously polynomial) construction of S_φ , and we can easily check that S_φ is a consistent 1-bounded acyclic marked graph. We have also provided some intuition why the reduction works, i.e., why

$$\varphi \text{ is satisfiable} \iff S_\varphi \text{ admits } M_0 \xrightarrow{\sigma} M_1 \xrightarrow{\tau} M_2 \text{ for some nonempty balanced } \tau.$$

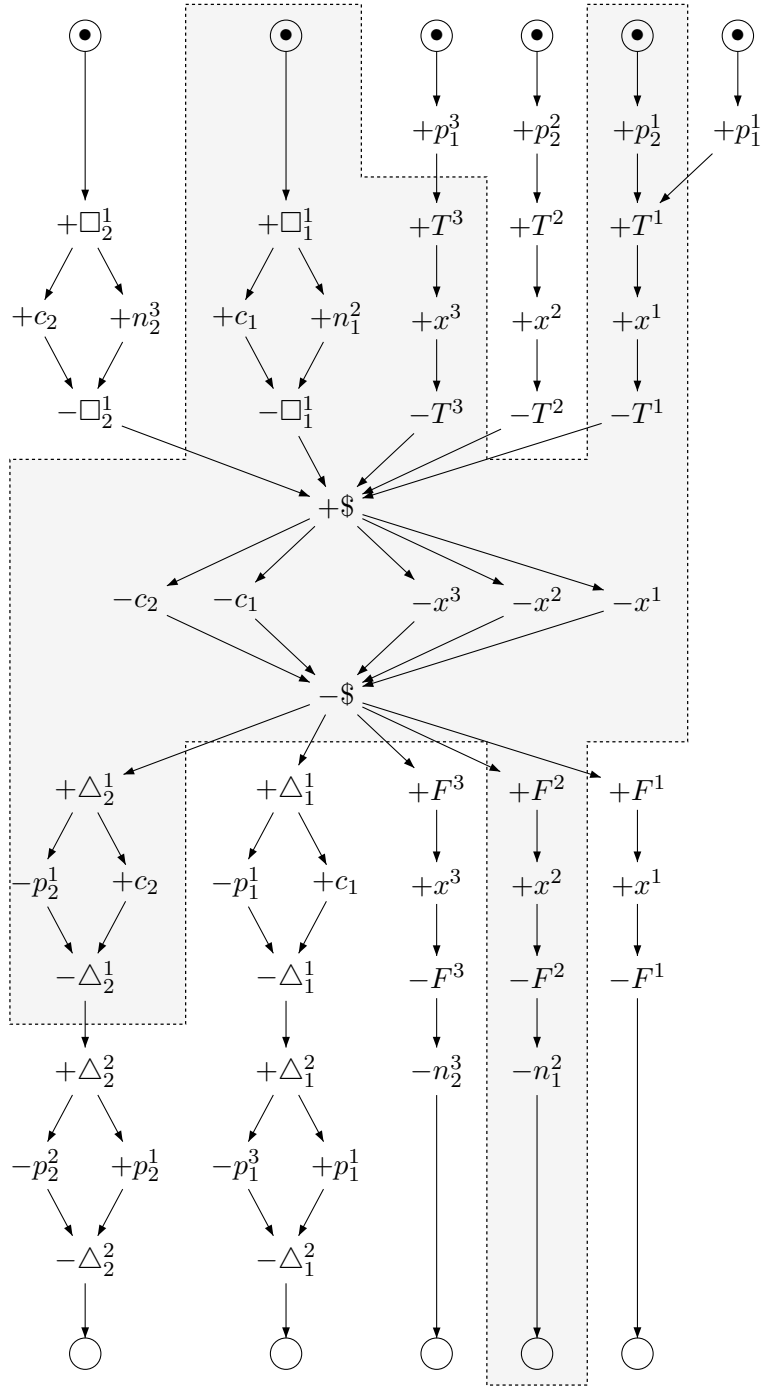


Figure 7. STG S_φ for $\varphi \equiv (x_1 \vee \bar{x}_2 \vee x_3) \wedge (x_1 \vee x_2 \vee \bar{x}_3)$

Now we summarize the correctness arguments, first informally and then in more detail.

If φ is satisfiable, then we ‘choose’ a satisfying truth assignment \mathcal{A} and for each clause c_i we ‘choose’ a literal l_i such that \mathcal{A} makes l_i true, where ‘choose’ has the meaning described above. This leads to a balanced sequence bal . On the other hand, if a balanced sequence bal can be found, then the corresponding ‘choice’ of literals must be consistent (and so φ is satisfiable): if both x_j and \bar{x}_j are ‘chosen’, then both $+n_i^j$ and $-p_k^j$ appear in bal , and both V_T^j and V_F^j must contribute to bal , which, as we have seen, is not possible. A more detailed formulation of these arguments follows.

(“ \implies ”) Suppose that φ is satisfiable; we shall show that there is an occurrence sequence $M_0 \xrightarrow{\sigma} M_1 \xrightarrow{\tau} M_2$ in S_φ such that τ is nonempty and balanced. Let

$$l : \{c_1, c_2, \dots, c_m\} \rightarrow \{x_1, \bar{x}_1, x_2, \bar{x}_2, \dots, x_n, \bar{x}_n\}$$

be a consistent choice of literals (which must exist since φ is satisfiable).

We now define certain sequences $\sigma(X)$, $\tau(X)$ for the appropriate subnets X . The required sequence σ will then be defined as

$$\sigma = \sigma(V_T^1) \cdots \sigma(V_T^n) \sigma(C_1^N) \cdots \sigma(C_n^N),$$

and τ will be defined as $\tau = \tau_1 \tau_2 \tau_3$, where

$$\tau_1 = \tau(V_T^1) \cdots \tau(V_T^n) \tau(C_1^N) \cdots \tau(C_n^N),$$

$$\tau_2 = \tau(N_\S),$$

$$\tau_3 = \tau(V_F^1) \cdots \tau(V_F^n) \tau(C_1^P) \cdots \tau(C_n^P).$$

As expected, we define

$$\tau(N_\S) = +\$ -x^1 \cdots -x^n -c_1 \cdots -c_m -\$.$$

For each j , $1 \leq j \leq n$, we proceed as follows. If x_j is in the range of l (x_j was chosen by at least one clause), we define $\tau(V_T^j)$ to be a sequence finishing by $+x^j$ and otherwise consisting of precisely those $+p_i^j$ for which $x_j = l(c_i)$; $\sigma(V_T^j)$ contains all others $+p_i^j$ appearing in V_T^j . And we define $\tau(V_F^j)$ as the empty sequence.

If \bar{x}_j is in the range of l , we define $\sigma(V_T^j)$ to be a sequence containing all transitions (i.e., labels) of V_T^j , and we define $\tau(V_T^j)$ as empty. But now $\tau(V_F^j)$ is nonempty; it starts with $+x_j$ and otherwise contains precisely those $-n_i^j$ for which $\bar{x}_j = l(c_i)$.

The above definition is sound since both x_j, \bar{x}_j can not belong to the range of l (l is consistent). If none of x_j, \bar{x}_j belongs to the range of l , we can define $\sigma(V_T^j)$ as a sequence consisting of all $+p_i^j$ in V_T^j , and we put $\tau(V_T^j) = +x_j$ and $\tau(V_F^j)$ empty.

Now for each i , $1 \leq j \leq m$, we proceed as follows. If $l(c_i) = x_j$ (thus $\tau(V_T^j)$ is nonempty), we define $\tau(C_i^P)$ as the sequence

$$+\Delta_i^1 +c_i \cdots -p_i^{k_d} -\Delta_i^d,$$

where $k_d = j$ (recall Figure 6); thus $+c_i$ and $-p_i^j$ is left to be compensated in $\tau(C_i^P)$. And we define $\sigma(C_i^N)$ as the whole behaviour of C_i^N , and $\tau(C_i^N)$ as empty.

If $l(c_i) = \bar{x}_j$ (thus $\tau(V_F^j)$ is nonempty), we define

$$\sigma(C_i^N) = +\square_i^1 +n_i^{j_1} \cdots -n_i^{j_d} -\square_i^{d-1}$$

$$\tau(C_i^N) = +\square_i^d +n_i^{j_d} \cdots +c_i -\square_i^a,$$

where $j_d = j$ (recall Figure 5); thus $+c_i$ and $+n_i^j$ is left to be compensated in $\tau(C_i^P)$. And we define $\tau(C_i^P)$ as empty.

It is now straightforward to check that the defined sequence $\sigma\tau$ can really occur from M_0 , and that τ is nonempty and balanced. (For each j , $1 \leq i \leq m$, either $\tau(V_T^j)$ or $\tau(V_F^j)$ is nonempty; thus the occurrence $-x^j$ of $\tau(N_\S)$ is compensated. For each i , $1 \leq i \leq m$, either $\tau(C_i^N)$ or $\tau(C_i^P)$ is nonempty. Suppose that it is $\tau(C_i^N)$, the other case being similar. Then $\tau(C_i^N)$ is balanced except of the two ‘superfluous’ occurrences: $+c_i$ and $+n_i^j$, where $l(c_i) = \bar{x}_j$. But $+c_i$ is compensated by $-c_i$ of $\tau(N_\S)$, and $+n_i^j$ is compensated by $-n_i^j$ of $\tau(V_F^j)$.)

Remark. The formula φ of Figure 7 is satisfiable. The shaded region corresponds to the balanced sequence for the choice $l(c_1) = \bar{x}_2$, $l(c_2) = x_1$ (and setting x_3 to true).

(“ \Leftarrow ”) We now assume that there is an occurrence sequence $M_0 \xrightarrow{\sigma} M_1 \xrightarrow{\tau} M_2$ in S_φ such that τ is nonempty and balanced; we shall show that φ is satisfiable. We can easily check that τ must contain at least one transition of N_\S ; but this obviously implies that τ contains all transitions of N_\S .

Hence τ can be written

$$\tau = \tau'_1 \tau'_2 \tau'_3,$$

where

$$\tau'_2 = \tau(N_\S) = +\$ -x^1 \cdots -x^n -c_1 \cdots -c_m -\$.$$

For each j , $1 \leq j \leq n$, the occurrence $-x^j$ in τ'_2 must be compensated by $+x^j$ in either τ'_1 or τ'_3 ; so precisely one of the nets V_T^j and V_F^j contributes to τ . This naturally corresponds to a truth assignment \mathcal{A} .

For each i , $1 \leq i \leq m$, the occurrence $-c_i$ in τ'_2 must be compensated by $+c_i$ in either τ'_1 or τ'_3 . It is thus clear that precisely one of the nets C_i^N and C_i^P contributes to τ ; in this contribution, precisely one $+n_i^j$ or $-p_i^j$, respectively, is left to be compensated. This naturally defines a choice of literals l ; we will show that this choice l is consistent.

The only possibility how such a ‘superfluous’ $+n_i^j$ ($l(c_i) = \bar{x}_j$) is compensated in τ is by $-n_i^j$ in V_F^j , which means that $+x_j$ of V_F^j appears in τ (i.e., $\mathcal{A}(x_j) = 0$). If it is $-p_i^j$ which is to be compensated ($l(c_i) = x_j$), τ must contain $+p_i^j$ of V_T^j , which means that $+x_j$ of V_T^j appears in τ (i.e., $\mathcal{A}(x_j) = 1$).

This implies that the above mentioned choice of literals l is indeed consistent, which means that φ is satisfiable. \square

The previous lemma is now used to derive the desired co-NP-hardness results.

Proposition 3.4. Both the CSC problem and the USC problem are co-NP-hard for (consistent) STGs whose underlying nets are 1-bounded acyclic marked graphs.

Proof:

We use the STG S_φ constructed in the proof of Lemma 3.1, recalling that it is a consistent 1-bounded acyclic marked graph; let us denote its (unique) consistent binary encoding by b .

Assume now that S_φ does not have the USC property. This means that there are occurrence sequences

$$M_0 \xrightarrow{\sigma_1} M_1, \quad M_0 \xrightarrow{\sigma_2} M_2,$$

such that

- $M_1 \neq M_2$ (i.e., σ_1 and σ_2 do not contain the same transitions),

- $b(M_1) = b(M_2)$.

We explore the following cases C1), C2), C3), covering all possibilities.

C1) One of σ_1, σ_2 contains +\$ but not -\$ (turns out impossible):

In this case, both σ_1, σ_2 must contain +\$ and not -\$ (since $b(M_1) = b(M_2)$). (I.e., both M_1, M_2 are in the ‘middle segment’.) But this is impossible, since we obviously cannot have both $M_1 \neq M_2$ and $b(M_1) = b(M_2)$.

C2) Both σ_1, σ_2 do not contain +\$ or both contain -\$ (turns out impossible):

Either both M_1, M_2 are in the ‘left segment’ or both are in the ‘right segment’. Let us assume the subcase where both σ_1, σ_2 do not contain +\$ (i.e., both are in the ‘left segment’). We observe that σ_1 and σ_2 must contain the same transitions of V_T^j , for each j (since $b(M_1) = b(M_2)$).

So there must be a transition of C_i^N (for some i), which is, say, in σ_2 but not in σ_1 ; i.e., σ_2 contains a longer prefix of the behaviour of C_i^N (cf. Figure 5) than σ_1 . But then the difference between the two prefixes must obviously be balanced, which can be easily checked to be impossible.

The subcase where both σ_1, σ_2 contain -\$ can similarly be shown as impossible.

C3) One of σ_1, σ_2 does not contain +\$, and the other contains -\$:

We can assume that σ_1 does not contain +\$ and σ_2 contains -\$. (I.e., M_1 is in the ‘left segment’, M_2 is in the ‘right segment’.) This implies that there is (a nonempty) τ such that $M_0 \xrightarrow{\sigma_1} M_1 \xrightarrow{\tau} M_2$; necessarily, τ is balanced. And from the proof of Lemma 3.1 we know that this is possible if and only if φ is satisfiable.

Thus cases C1), C2) turn out to be impossible, and C3) is possible if and only if there is a (nonempty) τ such that $M_0 \rightarrow^* M_1 \xrightarrow{\tau} M_2$; necessarily, τ is balanced. Moreover, such M_1, M_2 (with $b(M_1) = b(M_2)$) enable different sets of signals, so the CSC property is violated—when viewing all signals as output signals. Therefore we can apply Lemma 3.1 to finish the proof. \square

Proposition 3.5. Both the CSC problem and the USC problem are co-NP-hard for live 1-bounded marked graph STGs.

Proof:

Consider the USC problem. We reuse the Petri net S_φ from the proof of Lemma 3.1. We note that the behaviour obtained by firing all transitions of S_φ is not balanced; i.e., $b(M_0)$ and $b(M_f)$, where b is the consistent boolean encoding and M_f is the final marking, differ on some signals.

Remark. For concreteness, these unbalanced signals are $x^j, c_i, n_i^{j_d}$ (for $d = 2, 3, \dots, a$), and $p_i^{k_d}$ (for $d = 1, 2, \dots, b-1$).

We define a new STG S'_φ by adding a ‘final segment’ to S_φ : we add a fresh signal f and construct a ‘linear’ net N_f with the behaviour

$$+f \ell_1 \ell_2 \dots \ell_k -f,$$

where ℓ_i are the labels compensating the unbalance of S_φ ; they include $-x^j$, $-c_i$, $+n_i^{j2}$, etc.; we note that each nonempty sequence of transitions of N_f is unbalanced. The net N_f will be prompted in S'_φ after all transitions of S_φ occur; the final transition of N_f will then restore the initial marking M_0 .

Hence S'_φ is an STG whose underlying net is a live and 1-bounded marked graph. It is easy to see that any sequence containing precisely one occurrence of each transition of S'_φ is balanced. Let b' be the unique consistent boolean encoding of S'_φ .

We show that S_φ has the USC property iff S'_φ has the USC property, which proves the second part of the proposition.

It is trivial that if S_φ does not have the USC property, then S'_φ does not have it either. For the other direction, assume that S'_φ does not have the USC property. Then there is a *witness of the USC-violation*, i.e., two occurrence sequences

$$M_0 \xrightarrow{\sigma_1} M_1, \quad M_0 \xrightarrow{\sigma_2} M_2$$

as in the proof of Proposition 3.4.

Let us assume that the witness is *minimal* in the sense that neither σ_1 nor σ_2 can be shortened. We prove that this minimal witness also corresponds to a USC-violation in the Petri net S_φ . It suffices to show that neither σ_1 nor σ_2 contains a transition labelled by the signal f .

Assume that one of σ_1 and σ_2 , say σ_2 , contains an occurrence of the signal f . Since $b'(M_1) = b'(M_2)$, we can easily check that the assumption $b'(M_2)(f) = 1$ would force $M_1 = M_2$, a contradiction. So $b'(M_2)(f) = 0$, which means that the last occurrence of f in σ_2 is $-f$. But then σ_2 can be (rearranged and) written as $\sigma_2 = \sigma'_2 \sigma$ where σ contains precisely one occurrence of each transition of S'_φ . This implies $M_0 \xrightarrow{\sigma'_2} M_2$, which contradicts our minimality assumption.

Consider now the CSC property. Assume that all signals are output signals. We show that S_φ has the CSC property iff S'_φ has the CSC property. As in the USC case, it is trivial that if S_φ does not have the CSC property, then S'_φ does not have it either. For the other direction, assume S_φ has the CSC property. We have shown in Lemma 3.1 that in this case S_φ has the USC property as well. So, by the first part of this proof concerning the USC property, S'_φ has the USC property. Since USC implies CSC, S'_φ has the CSC property, and we are done. \square

We now show the upper bound, a lemma which was already (implicitly) proved in [1].

Lemma 3.2. Both the CSC problem and the USC problem are in co-NP for (bounded or unbounded) marked graph STGs.

Proof:

Let $S = (N, M_0, \ell)$ be a normalized and consistent marked graph STG. (We recall that consistency of S can be checked in polynomial time.) It is sufficient to deal with the CSC problem; the claim for the USC problem will follow easily.

We observe that S does not have the CSC property if and only if there are sequences u_1, u_2 such that

- $M_0 \xrightarrow{u_1} M_1, M_0 \xrightarrow{u_2} M_2,$
- $M_1 \neq M_2,$
- for each signal a :

$$P(u_1)(a^+) - P(u_1)(a^-) = P(u_2)(a^+) - P(u_2)(a^-),$$

- M_1, M_2 enable different output signals.

To check that there is such a ‘CSC-violation’, a nondeterministic (polynomial) algorithm guesses a place p such that $M_1(p) \neq M_2(p)$, and guesses further whether $M_1(p) > M_2(p)$ or $M_2(p) > M_1(p)$ holds. The algorithm proceeds to guess an output signal a , and which of M_1, M_2 enables a . Assume w.l.o.g. the guess is that M_1 enables a and M_2 does not. The algorithm guesses which places of M_1 carry at least one token (including all the input places of some transition labelled by a) and which places of M_2 carry no token (including at least one input place of each transition labelled by a). The algorithm translates all these guesses into a system of linear inequalities, guesses an integer solution of polynomial size, and checks in polynomial time that it is indeed a solution. (Variables for transition sequences are replaced by variables for their Parikh vectors, and Claim 3.2 is used.) \square

Putting together Propositions 3.4 and 3.5 and Lemma 3.2 we obtain:

Theorem 3.2. The CSC problem and the USC problem are co-NP-complete for marked graph STGs, and stay co-NP-hard for live and 1-bounded marked graph STGs as well as for 1-bounded acyclic marked graph STGs.

Remark. Notice that in the marked graphs produced by the reduction from the proof of Lemma 3.1 there are different transitions carrying the same label. The case with injective labelling (each transition has its unique label) might well admit a polynomial algorithm but we leave this problem open here.

4. Live and bounded free-choice nets

As already mentioned, [7] shows that consistency can be decided in polynomial time for live and bounded free-choice STGs that are moreover cyclic, meaning that the initial marking is reachable from every reachable marking. It is not known whether the polynomiality result still holds if the cyclicity condition is removed, and we leave this problem open.

We now show co-NP-completeness of the CSC problem and of the USC problem for live and bounded free-choice STGs. Since live and bounded marked graphs are cyclic, Theorem 3.2 gives co-NP-hardness even for *cyclic* live and bounded free-choice STGs. So we just need to show that the complementary problem is in NP. We proceed similarly as in the marked graph case, first recalling a known result analogous to Claim 3.2; for this we use the following notation:

For a net $N = (P, T, F)$ and $X : T \rightarrow \mathbb{N}$, we denote by $N_X = (P_X, T_X, F_X)$ the subnet of N defined as follows: T_X is the set of transitions of T for which $X(t) \geq 1$, $P_X = \bullet T_X \cup T_X^\bullet$, and F_X is the projection of F on $(P_X \times T_X) \cup (T_X \times P_X)$. We also recall that $Q \subseteq P$ is a *trap* in $N = (P, T, F)$ if $Q^\bullet \subseteq \bullet Q$. (If a trap is marked, i.e., has at least one token, it cannot be unmarked). Here we consider only nonempty traps $Q \neq \emptyset$.

Lemma 4.1. ([15])

Let (N, M_0) be a live and bounded free-choice Petri net, and let C_N be its incidence matrix. An integer vector $X_0 \geq 0$ is the Parikh vector of a transition sequence enabled at M_0 if and only if

1. $M_0 + C_N \cdot X_0 \geq 0$, and
2. $M = M_0 + C_N \cdot X_0$ marks all traps of N_{X_0} .

Theorem 4.1. The CSC problem and the USC problem are co-NP-complete for live and bounded free-choice STGs.

Proof:

As mentioned above, co-NP-hardness follows from Theorem 3.2 (even when the Petri nets are also cyclic).

A nondeterministic polynomial algorithm for showing that a given (consistent) live and bounded free-choice STG does not have the CSC property (or the USC property) can be constructed as in the proof of Lemma 3.2, using Lemma 4.1 instead of Claim 3.2.

A little difficulty is the fact that a (nonnegative integer) solution of $M_0 + C_N \cdot X \geq 0$ may not be the Parikh vector of an occurrence sequence. The algorithm handles this problem by guessing (and requiring in the system of inequalities) which components of X are positive and which are zero; then it guesses a subset P' of places of N_X , verifies that P' does not contain a trap in N_X (which can be easily done in polynomial time) and requires (in the constructed system of inequalities) that $M_0 + C_N \cdot X$ is positive for all places of N_X outside P' . \square

In the next section we show the importance of the assumption of liveness.

5. More general nets

Here we study the complexity of the consistency, CSC, and USC problems for more general classes of STGs.

By a straightforward use of standard techniques of Petri net theory (using the reachability problem for k -bounded nets) we can show:

Proposition 5.1. The consistency problem, the CSC problem and the USC problem are PSPACE-complete for k -bounded nets (for any fixed k).

The relevant proofs can be found in the appendix.

It is worth noting that free-choice does not make this simpler:

Proposition 5.2. The consistency problem for 1-bounded free-choice STGs (not necessarily live) is PSPACE-complete.

Proof:

An arbitrary 1-bounded STG can be transformed into a 1-bounded free-choice STG by means of the operation illustrated in Figure 8 while preserving consistency. (This operation is closely related to the “releasing arcs”-technique, see e.g. [4].)

In more detail: For changing an arbitrary 1-bounded STG into a 1-bounded free-choice STG, we use transformations of the following type:

- inside an original arc from a place p to a transition t , insert a fresh label f^+ (i.e., add a new freshly labelled transition and an additional place), and
- split t into t_1, t_2 , adding a new place p' and arcs $t_1 \rightarrow p', p' \rightarrow t_2$, where t_1 (t_2) inherits the current input (output) places of t ; put $\ell(t_1) = \ell(t)$ and $\ell(t_2) = f^-$.




Figure 8. Transforming a 1-bounded STG into a 1-bounded free-choice STG

Figure 8 shows the result of two such transformations, where f_2 was inserted before f_1 (in fact, the order of the f_i^- labels is irrelevant); f_2^- now has the original output places of l_2 .

Using sufficiently many such transformations, we can obviously transform each 1-bounded STG $S = (N, M_0, \ell)$ into $\bar{S} = (\bar{N}, \bar{M}_0, \bar{\ell})$ which is 1-bounded and free-choice; \bar{M}_0 coincides with M_0 on places from N and is 0 elsewhere.

An important observation is that S is consistent iff \bar{S} is consistent. To show this, it suffices to show that using the described transformation, i.e., the step performed for one arc and changing an intermediate S' to S'' , keeps the (in)consistency untouched. We can easily note that if S' is inconsistent then so is S'' . And 1-boundedness guarantees that S'' can not become inconsistent due to the added f^+ , f^- ; thus an inconsistency witness in S'' naturally ‘translates’ into S' . \square

Using reductions from and to the reachability problem of general Petri nets, we can show

Proposition 5.3. The consistency problem and the CSC problem for general STGs are decidable but EXPSPACE-hard.

The proofs are also in the appendix.

6. Conclusions and related work

We have explored the complexity of the consistency and the CSC problem for several classes of STGs. The main result shows that deciding the CSC property is co-NP-complete even for 1-bounded and acyclic marked graph STGs and for 1-bounded and live marked graph STGs. The same result holds for the USC property. This result explains why none of the existing approaches for checking the USC or the CSC property in marked graph STGs is polynomial and complete.

In [14] the USC property was studied for live and 1-bounded marked graph STGs with injective labelling (i.e., one up-transition and one down-transition per signal). A sufficient condition for the USC property to hold is presented, and it is shown that it can be checked in polynomial time. The condition is

conjectured to be also necessary, which would imply that checking the USC property is polynomial. The reduction used in our NP-completeness result transforms a formula into an STG in which several signals have two up- and two down-transitions, and so it does not apply to this case. The complexity of the USC property for this particular case is left for future research.

In [16] the result of [14] is extended to the case where the STG may have several up- and down-transitions per signal. The paper presents a generalization of the sufficient condition of [14]. Our NP-completeness result shows that if $P \neq NP$, then the condition is not necessary, or it cannot be checked in polynomial time, or both. In fact, we conjecture that the condition is neither necessary, nor can be checked in polynomial time (it requires to establish a property for a potentially exponential number of objects).

In [17] it is shown that a live and 1-bounded marked-graph STG violates the USC property iff the STG has a so-called complementary path. The paper proposes an algorithm that searches for such paths. The worst-case complexity of the algorithm is exponential, and by our result this is unavoidable unless $P=NP$.

In [11] a polynomial algorithm is presented that detects all violations of the CSC property in a live and bounded free-choice STG. However, the algorithm may also give false positives, i.e., it may detect false violations. Our result shows that if $P \neq NP$ then every polynomial algorithm must produce false positives or false negatives.

In [1] a procedure is described that, given a marked-graph STG, constructs in polynomial time an Integer Linear Programming (ILP) problem such that the STG violates the CSC property if and only if the problem has a solution. Our result shows that, unless $P=NP$, ILP is necessary, and cannot be replaced by ordinary Linear Programming (recall that Linear Programming problems can be solved in polynomial time).

In [9, 10] it is shown how to check the CSC property for arbitrary bounded STGs using net unfoldings and ILP-solvers or SAT-solvers. Given a bounded STG S , an object is constructed called the unfolding of S . This unfolding is used to generate an ILP problem (a boolean formula) such that S violates the CSC property iff the ILP problem has a solution (iff the formula is satisfiable). If S is a live and 1-bounded marked graph, then the unfolding of S has polynomial size in S ([5], Theorem 4.14). This shows that, even for marked graphs, ILP-solvers or SAT-solvers are unlikely to be replaceable by other tools with polynomial running time: if $P \neq NP$, then no polynomial algorithm taking the unfolding of S as input can decide the CSC or the USC property.

Finally, it could be argued that the important problem in practice is not to decide whether a given STG satisfies the CSC property, but to transform an STG that does not satisfy the CSC property into another one that does. In [2] an automatic, very efficient procedure for such a transformation is presented. Unfortunately, the procedure adds many additional signals (one per place of the STG), and so in most cases its output is only useful as a first approximation to the design. The optimization of this first approximation has to be carried out by a (possibly automatic) trial and error procedure in which a candidate for an optimized STG is guessed. The candidate must be checked for the CSC property, which brings us back to the problem discussed in this paper.

Acknowledgments. The first author thanks Jordi Cortadella and José Carmona for helpful discussions. He also thanks the Alexander von Humboldt Foundation for supporting the cooperation between universities in Stuttgart and Brno.

References

- [1] Carmona, J., Cortadella, J.: ILP Models for the Synthesis of Asynchronous Control Circuits, *2003 International Conference on Computer-Aided Design (ICCAD'03), November 9-13, 2003, San Jose, CA, USA*, IEEE Computer Society / ACM, 2003, ISBN 1-58113-762-1.
- [2] Carmona, J., Cortadella, J., Pastor, E.: A structural encoding technique for the synthesis of asynchronous circuits, *Proc. Int. Conf. on Application of Concurrency Theory to System Design*, IEEE Computer Society, 2001.
- [3] Chu, T.-A.: *Synthesis of Self-Timed VLSI Circuits from Graph-theoretic Specifications*, Ph.D. Thesis, MIT, 1987.
- [4] Desel, J., Esparza, J.: *Free Choice Petri Nets*, vol. 40 of *Cambridge Tracts in Theoretical Computer Science*, Cambridge University Press, 1995.
- [5] Esparza, J.: Model checking using net unfoldings, *Science of Computer Programming*, **23**, 1994, 151–195.
- [6] Esparza, J.: Decidability and Complexity of Petri Net Problems – an Introduction, *Lectures on Petri Nets I: Basic Models. Advances in Petri Nets* (G. Rozenberg, W. Reisig, Eds.), number 1491 in *Lecture Notes in Computer Science*, 1998.
- [7] Esparza, J.: A Polynomial-Time Algorithm for Checking Consistency of Free-Choice Signal Transition Graphs, *Fundamenta Informaticae*, **62**(2), 2004, 197–220.
- [8] Esparza, J., Jančar, P., Miller, A.: On the Complexity of Consistency and Complete State Coding for Signal Transition Graphs, *Proceedings of the 6th International Conference on Application of Concurrency to System Design (ACSD 2006)*, IEEE Computer Society, Turku, Finland, June 2006.
- [9] Khomenko, V., Koutny, M., Yakovlev, A.: Detecting State Coding Conflicts in STGs Using Integer Programming, *Proc. of the Design, Automation and Test in Europe Conference and Exhibition*, IEEE Computer Society, 2002.
- [10] Khomenko, V., Koutny, M., Yakovlev, A.: Detecting State Coding Conflicts in STG Unfoldings using SAT, *Proc. of the 4th Int. Conf. on Application of Concurrency to System Design*, IEEE Computer Society, 2004.
- [11] Pastor, E., Cortadella, J.: Polynomial algorithms for the synthesis for hazard-free circuits from signal transition graphs, *1993 International Conference on Computer-Aided Design (ICCAD'93), Santa Clara, CA, USA*, IEEE Computer Society / ACM, 1993.
- [12] Rosenblum, L., Yakovlev, A.: Signal Graphs: from self-timed to timed ones, *Proc. Int. Workshop on Timed Petri nets*, IEEE Computer Society, 1985.
- [13] Schrijver, A.: *Theory of Linear and Integer Programming*, Wiley, 1986.
- [14] Vanbekbergen, P., Catthoor, F., Goossens, G., Man, H. D.: Optimized synthesis of asynchronous control circuits from graph-theoretic specifications, *1990 International Conference on Computer-Aided Design (ICCAD'90)*, IEEE Computer Society, 1990.
- [15] Yamasaki, H., Huang, J., Murata, T.: Reachability Analysis of Petri Nets via Structural and behavioral Classifications of Transitions, *Petri Net Newsletter*, (60), 2001, 5–21.
- [16] Ykman-Couvreur, C., Lin, B., Goossens, G., Man, H. D.: Synthesis and optimization of asynchronous controllers based on extended lock graph theory, *4th European Conference on Design Automation, Paris, France*, IEEE Computer Society, 1993.
- [17] Yu, M., Subrahmanyam, P.: A new approach for checking the unique state coding property of signal transition graphs, *Proc. 3rd Int. European Conference on Design Automation*, IEEE Computer Society, 1992.

7. Appendix (Bounded nets and general nets)

This appendix provides proofs for the results mentioned in Section 5. I.e., it demonstrates the hardness of the consistency problem and the CSC problem in the general case, where the underlying net of an STG can be an arbitrary Petri net, and in the subcase when a bound on the number of tokens in each place is given (which can be viewed as ‘capacity’). This is done by showing suitable polynomial reductions from the reachability problem to the consistency problem and to the CSC problem, respectively.

In fact, the reachability problem can also provide an upper bound, so the consistency and CSC problems can be roughly viewed as equivalent with the reachability (w.r.t. the computational complexity). We also clarify the case of consistency in more detail, by showing a relation to the fireability problem (which is straightforwardly equivalent to the coverability problem).

The proof ideas use the usual techniques, so we do not describe them very formally nor in great detail. For completeness, we start by recalling definitions and the known complexity results for reachability and fireability. (Precise references can be found, e.g., in [6].)

The *reachability problem* (RP)

Instance: a Petri net (N, M_0) and a marking M .

Question: Is $M_0 \longrightarrow^* M$?

Theorem 7.1. For (general) Petri nets, RP is decidable and EXPSpace-hard. For k -bounded nets (for any fixed k), RP is PSPACE-complete.

The *fireability problem* (FP)

Instance: a Petri net (N, M_0) and a transition t .

Question: Is t fireable (i.e., is there some M such that $M_0 \longrightarrow^* M \xrightarrow{t}$)?

Theorem 7.2. For (general) Petri nets, FP is EXPSpace-complete. For k -bounded nets (for any fixed k), FP is PSPACE-complete.

7.1. Nonreachability reduces to consistency

We show how an instance $(N, M_0), M_1$ of the reachability problem in general nets can be transformed into an STG which is consistent if and only if M_1 is not reachable in (N, M_0) . Moreover, if (N, M_0) is 1-bounded then the constructed STG is also 1-bounded. In addition, we clarify the difference between conditions (1) and (2) of Proposition 2.1 on one hand and condition (3) on the other hand. Conditions (1), (2) turn out to be equivalent to the fireability problem; it is condition (3) which is as difficult as reachability.

We start with a simple construction that we use several times:

Construction 1: Given a Petri net (N, M_0) , we denote by $S_{(N, M_0)}$ the STG obtained as follows:

- As the set of signals of $S_{(N, M_0)}$ we take the set of transitions of N .
- In N , we replace every transition t by a place p_t and two transitions t_1, t_2 , labelled by t^+, t^- , respectively. Transition t_1 inherits the input places of t , and has p_t as the unique output place; transition t_2 has p_t as the unique input place and inherits the output places of t .

- We add a (run-)place r , and an arc $r \rightarrow t_1$ and $t_2 \rightarrow r$, for every transition t of N .
- The initial marking of $S_{(N, M_0)}$ coincides with M_0 on the places inherited from N ; moreover, r carries 1 token, and places p_t are empty.

Observation 7.1. The STG $S_{(N, M_0)}$ is consistent, and it tightly simulates the behaviour of (N, M_0) .

Proposition 7.1. The reachability problem for general Petri nets (for k -bounded Petri nets) is polynomially reducible to the inconsistency problem for general STGs (for k -bounded STGs).

Proof:

Assume an instance of the reachability problem: $(N, M_0), M_1$. We consider the following construction of an STG S :

- Start with $S_{(N, M_0)}$ as described in Construction 1.
- Add a transition t_f labelled by a fresh label f^+ , and an arc from (the run-)place r to t_f . (Thus the added t_f can fire at most once, by which a dead marking is reached, corresponding to a reachable marking of N .)
- Add a new (starting) place s ; the initial marking of S will put 1 token in s ; all other places (including r) will be initially empty.
- Add a transition z_1 , labelled by a fresh a^+ . It takes the token from s and installs M_1 in the places inherited from N . (Marking M_1 thus becomes ‘frozen’.)
- Finally add transitions z_2, z_3 , labelled with a^+ and a^- , a place p_z , and the arcs

$$s \rightarrow z_2(a^+) \rightarrow p_z \rightarrow z_3(a^-) \rightarrow r$$

as well as additional arcs from z_3 which install M_0 in the places inherited from N .

We observe that the constructed S can start either with firing z_1 (label a^+), reaching the ‘frozen’ M_1 , or with firing $z_2 z_3$ ($a^+ a^-$) after which it behaves like $S_{(N, M_0)}$, with a possibility to ‘freeze’ any reachable marking of (N, M_0) . We also note that if (N, M_0) is k -bounded then S is also k -bounded.

It is clear that S can not provide any inconsistency witness (M, a) of the form (1) and (2) of Proposition 2.1; there might be a witness satisfying (3) but this happens if and only if M_1 is reachable in (N, M_0) . \square

The previous reduction was based on condition (3). For completeness, we show that the existence of an inconsistency witness of form (1) or (2) is ‘easier’, namely polynomially equivalent to the fireability problem.

Claim 7.1. The problem of deciding, given an STG S , if there is a pair (M, a) satisfying conditions (1) or (2) of Proposition 2.1 is polynomially equivalent to the fireability problem.

Proof:

We first show that the problem if S provides a pair (M, a) satisfying (1) can be reduced to the fireability problem.

Let us fix a (a signal), and define:

- S_a^+ is a ‘copy’ of S from which we remove all transitions labelled by a^- together with their adjacent arcs, and we add a (run-)place r_a^+ with 1 token. For every transition t of S_a^+ which is not labelled by a^+ we add arcs $r_a^+ \rightarrow t, t \rightarrow r_a^+$; in the case of t labelled a^+ we only add $r_a^+ \rightarrow t$.

We observe that S_a^+ behaves like S until a first occurrence of an a -label; this a -label must be a^+ , and the computation of S_a^+ is thus finished.

Similarly we proceed for a^- :

- S_a^- is a ‘copy’ of S from which we remove all transitions labelled by a^+ together with their adjacent arcs, and we add a (run-)place r_a^- with 1 token. For every transition t of S_a^- which is not labelled by a^- we add arcs $r_a^- \rightarrow t, t \rightarrow r_a^-$; in the case of t labelled a^- we only add $r_a^- \rightarrow t$.

Now, we put the STGs S_a^+ and S_a^- side by side. We add new places p_a^+ and p_a^- , which are initially empty, and a new transition t_f^a with the arcs $p_a^+ \rightarrow t_f^a$ and $p_a^- \rightarrow t_f^a$. Moreover, for each t in S_a^+ labelled by a^+ we add $t \rightarrow p_a^+$, and for each t in S_a^- labelled by a^- we add $t \rightarrow p_a^-$.

We have thus got a Petri net (N, M_0) where t_f^a is fireable if and only if the initial marking M of S satisfies (1).

To reach our goal, we further modify the net (N, M_0) :

- Add a new (run-)place r , initially with a token, and let r_a^+, r_a^- be initially empty.
- For each transition t of S , add an additional copy of t (to (N, M_0)), with the arcs $r \rightarrow t, t \rightarrow r$. For each arc $p \rightarrow t$ in S , add arcs $p_1 \rightarrow t, p_2 \rightarrow t$, where p_1 and p_2 are the copies of p in S_a^+ and S_a^- , respectively. Similarly for the output arcs $t \rightarrow p$.
- Finally add a transition t_a , with the arcs $r \rightarrow t_a, t_a \rightarrow r_a^+, t_a \rightarrow r_a^-$.

We observe that the arisen net (N', M'_0) in the first phase simulates S synchronously on both the places in S_a^+ and the places in S_a^- . To enable t_f^a , this first phase must stop by firing t_a (which unmarks r and marks r_a^+, r_a^-). Transition t_f^a can then indeed be enabled if and only if the corresponding M , reachable in S and copied in both S_a^+ and S_a^- , satisfies (1).

Thus we have described a polynomial algorithm which, given an STG S and signal a , constructs (N', M'_0) so that S has an inconsistency witness (M, a) satisfying (1) iff t_f^a is fireable in (N', M'_0) . The construction can be completed by subnets S_b^+ and S_b^- for all signals b , and some straightforward modifications, one of them ensuring that firing any t_f^b will enable an additional distinguished transition t_f .

Using similar techniques, we can extend the overall construction to show that

there is a polynomial algorithm which, given an STG S , constructs a net (N_S, M_S) , with a distinguished transition t_f , so that S has an inconsistency witness (M, a) satisfying (1) or (2) iff t_f is fireable in (N_S, M_S) .

For the other direction, assume an instance (N, M_0) , t of the fireability problem. We start with constructing $S_{(N, M_0)}$ by Construction 1; then we add a new transition z , whose only input place is p_t and whose label is t^+ . It is clear that the so constructed STG S has some (M, a) satisfying (1), or (2), iff t is fireable in (N, M_0) . \square

7.2. Nonreachability reduces to CSC

Proposition 7.2. The reachability problem for general Petri nets (for k -bounded Petri nets) is polynomially reducible to the negation of the CSC problem for general STGs (for k -bounded STGs).

Proof:

We use the single-place-zero reachability problem. So an instance is (N, M_0) and a place p_0 , and the question is if there is a reachable M with $M(p_0) = 0$. Given such an instance, we construct a (consistent) STG S as follows:

- We start with the (consistent) STG $S_{(N, M_0)}$ from Construction 1; recall that it has a run-place r .
- We now add two transitions t_1 and t_2 , labelled with fresh signals a^+ and a^- , respectively, two new places p_1, r_1 , and the following arcs:

$$r \rightarrow t_1(a^+) \rightarrow p_1 \rightarrow t_2(a^-) \rightarrow r_1.$$

- We add another two transitions t_3, t_4 , labelled with fresh b^+, b^- , a place p_2 , and the depicted arcs:

$$r \rightarrow t_3(b^+) \rightarrow p_2 \rightarrow t_4(b^-), t_4 \rightarrow r_1, t_4 \rightarrow p_0$$

- Finally we add a new transition t_5 labelled by o^+ , where o is defined as the only output signal; and we add the arcs $r_1 \rightarrow t_5$ and $p_0 \rightarrow t_5$.

We observe that the constructed S is still consistent, we denote the consistent boolean encoding by b . Now assume that (N, M_0) can reach M such that $M(p_0) = 0$. Then in $S = (N', M'_0, \ell)$ we have

- $M'_0 \xrightarrow{*} M'$, where $M'(p_0) = 0$,
- $M' \xrightarrow{a^+ a^-} M_1$, where $M_1(p_0) = 0$,
- $M' \xrightarrow{b^+ b^-} M_2$, where $M_2(p_0) = 1$,

So $M_1 \neq M_2$ and necessarily $b(M_1) = b(M_2)$; but M_1 does not enable o^+ and M_2 does. Hence S does not have the CSC property.

On the other hand, if all reachable M in (N, M_0) satisfy $M(p) \geq 1$, then S obviously has the CSC property.

Finally we note that if (N, M_0) is k -bounded then the STG S is ‘almost’ k -bounded. The only problem (increasing to $k+1$) can be caused by the arc $t_4 \rightarrow p_0$. But we can replace it by an arc $t_4 \rightarrow p'$ for a new place p' , and add a further transition t'_5 labelled by o^+ , with the arcs $r_1 \rightarrow t'_5, p' \rightarrow t'_5$. \square

7.3. PSPACE-completeness for bounded nets

Let us consider k -bounded STGs. Propositions 7.1 and 7.2 show that both the consistency problem and the CSC problem are PSPACE-hard. For showing that these problems are in PSPACE, it is sufficient to consider nondeterministic algorithms for the complementary problems (recall that $\text{PSPACE}=\text{NPSPACE}$). But such algorithms are obvious; so we have:

Proposition 7.3. Both the consistency problem and the CSC problem are PSPACE-complete for (explicitly) bounded nets.

7.4. Reduction to reachability

To provide some upper bound on the complexity of the consistency and CSC problems in the general case, we show reductions to the reachability problem. By a reduction we mean showing how an instance (of the consistency or CSC problem) can be answered by solving possibly several instances of the reachability problem, all of them being constructed in polynomial time.

Proposition 7.4. The consistency problem (for general STGs) is reducible to the reachability problem.

Proof:

Due to Claim 7.1, it is sufficient to handle condition (3) of Proposition 2.1. Given an STG, to decide if there is a marking M reachable by $w_1 a^+ u$ and by $w_2 a^- v$ for some a -free sequences u, v , we can let run two copies S', S'' of S independently (each having its own run-place). S' has the possibility to ‘freeze’ a marking reached by a sequence where a^+ was the last a -transition, while S'' has the possibility to ‘freeze’ a marking reached by a sequence where a^- was the last a -transition. In the final phase, the markings in both copies S', S'' will be ‘compared’: for each place of S , a transition which takes a token from both p' and p'' , which are the copies of p in S' and S'' , respectively, will fire as long as possible. The everywhere-zero marking can thus be reached if and only if the markings reached in S' and S'' coincide. \square

Proposition 7.5. The CSC problem (for general STGs) is reducible to the reachability problem.

Proof:

Let $S = (N, M_0, \ell)$ be an STG (which can be supposed to be consistent). For each signal a , we add places p_{a+} and p_{a-} , and we check which of the two signals a^+ and a^- can be enabled first. In the first case, we put 1 token in p_{a-} and 0 tokens in p_{a+} , and in the second case vice versa. We add further arcs, such that each a^+ -transition takes a token from p_{a-} and puts a token in p_{a+} , and each a^- -transition takes one from p_{a+} and puts it in p_{a-} .

So the modified net faithfully simulates the original S ; moreover, each reachable marking contains explicit information about the current (consistent) binary encoding.

Again, we can use two copies of S , use run-places for distinct phases of computation etc., such that this allows to choose any two reachable markings M_1, M_2 , after which it will be guaranteed that a specified (sub)marking will be reachable if and only if M_1 and M_2 have the same binary encoding but one of these markings enables a certain output signal while the other does not. (We can solve this for each output signal separately.) \square