Theoretical Computer Science 546 (2014) 225–236

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Rabin's theorem in the concurrency setting: A conjecture

P.S. Thiagarajan [a,*], Shaofa Yang [b]

[a] School of Computing, National University of Singapore, Singapore
[b] State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences, China

ARTICLE INFO

Available online 6 March 2014

Keywords:
Rabin's theorem
Monadic second order logic
Event structure
Petri net
Grid-free

ABSTRACT

Rabin's theorem says that the monadic second order theory of the infinite binary tree is decidable. This result has had a far reaching influence in the theory of branching time temporal logics. A simple consequence of Rabin's theorem is that for every finite state transition system, the monadic second order theory of its computation tree is decidable. Concurrency theory strongly suggests that finite 1-safe Petri nets (or simply, *net systems*) are a natural generalization of the notion of a finite state transition system while labelled event structures arising as the unfoldings of net systems are the proper counterparts to the computation trees obtained by unwinding finite (sequential) transition systems. It is easy to define the monadic second order theory of such event structures. It turns out that unlike the sequential case, not every net system (i.e. its event structure unfolding) has a decidable monadic second order theory. This gives rise to the question: Which net systems admit a decidable monadic second order theory? Here we present a conjecture based on a property called grid-freeness. Our conjecture is that a net system has a decidable monadic second order theory iff its event structure unfolding is grid-free. We show that it is decidable whether a net system has this property. We also prove that the monadic second order theory of a net system is undecidable if its event structure unfolding is not grid-free. In addition we show that our conjecture can be effectively reduced to the sub-class of *free choice* net systems. Finally we point out how the positive resolution of our conjecture will settle the decidability of a range of distributed controller synthesis problems.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

Rabin's theorem asserts the decidability of the monadic second order (MSO) theory of the infinite binary tree [1]. This fundamental result has played a dominant role in the field of branching time temporal logics [2]. A simple but important consequence of Rabin's result is that the MSO theory of every finite state transition systems is decidable. More precisely, the MSO theory of the computation tree that arises as the unwinding of a finite state transition system is decidable.

A natural question that arises is: What is a sound generalization of Rabin's theorem in the concurrency setting? Here we offer a conjecture regarding this question. This conjecture was formulated a number of years ago but has so far resisted a proof—or a counterexample—despite sustained efforts by the present authors (and—perhaps less sustained but—keen attempts by a number of other researchers). Our main purpose in presenting it here is to bring it to the attention of a larger community with the hope that it might get resolved.

* Corresponding author.
E-mail addresses: thiagu@comp.nus.edu.sg (P.S. Thiagarajan), yangsf@ios.ac.cn (S. Yang).

http://dx.doi.org/10.1016/j.tcs.2014.03.010
0304-3975/© 2014 Elsevier B.V. All rights reserved.

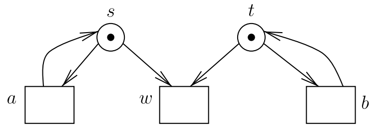

Fig. 1. The net system N_1.

In what follows we will present the material in terms of finite 1-safe Petri nets and related notions. This is mainly for convenience. Our conjecture and the results we present can easily be phrased in terms of any one of the related formalisms such as a network of transition systems that synchronize on common actions and various other types of asynchronous transition systems [3]. For brevity we will refer to finite 1-safe Petri nets as net systems.

The behaviors of finite state transition systems can be specified and verified using a rich set of linear time and branching time logics. One can view the monadic second order (MSO) logic of 1-successor interpreted over infinite runs of finite state transition systems as the canonical linear time logic and the MSO logic of n-successors interpreted over infinite tree unwindings of finite state transition systems as the canonical branching time logic [2]. The decidability of the former was settled by Büchi [4] while the decidability of the latter constitutes Rabin's theorem. Linear time and branching temporal logics such as *LTL*, *CTL*, *CTL** and the modal μ-calculus can be viewed as sub-logics that have a more convenient syntax and in which expressive power is often traded for a more efficient verification procedure.

In the case of concurrent systems the situation is similar when it comes to linear time behaviors. Mazurkiewicz traces—viewed as restricted labelled partial orders—constitute a natural generalization of sequences and the MSO logic of sequences can be smoothly extended to Mazurkiewicz traces [5]. In the branching time setting, labelled event structures [6] serve as an appropriate extension of computation trees. Further, just as a transition system can be unwound into a computation tree, a net system can be unwound into a labelled event structure [6]. One can also define an MSO logic for labelled event structures that can be viewed as a natural and conservative extension of the MSO logic over computation trees. At this stage however the correspondence between the sequential and concurrent settings breaks down.

Suppose we say that the MSO theory of a transition system is the MSO theory of the (labelled) computation tree obtained as its unwinding. As mentioned above, it follows from Rabin's theorem, that the MSO theory of (the tree unwinding of) every finite state transition system is decidable. In the concurrent setting, one can correspondingly say that the MSO theory of a net system is the MSO theory of the event structure obtained as its (event structure) unfolding. It turns out it is *not* the case that the MSO theory of every net system is decidable. Fig. 1 shows a simple net system (suggested by Igor Walukiewicz). The fact that its MSO theory is undecidable will follow at once from the main result of Section 4. At the same time sub-classes of net systems do have decidable MSO theories. For instance, if a net system does not exhibit any concurrency then its MSO theory is decidable thanks to Rabin's theorem. On the other hand if a net system does not exhibit any conflicts then too its MSO theory is decidable. These observations will follow easily from the material presented in Section 2. Hence the interesting question now is: What is the exact sub-class of net systems that have a decidable MSO theory? Our conjecture is:

The MSO theory of a net system is decidable iff its event structure unfolding is *grid-free*.

At the level of event structures grid-freeness demands the absence of the two dimensional grid as a conflict-free sub-structure (an exact definition is provided in the next section). This is a natural restriction since a variety of tiling and coloring problems for the two dimensional grid are undecidable [7]. As we show in Section 4 we can easily encode such problems in the MSO logic of event structures. At the system level, grid-freeness boils down to a restriction on the allowed communication patterns. It says that if two disjoint groups of processes in a distributed system have both *independently* executed sufficiently many steps then they will never hear from each other again. Using this insight we show that it is decidable whether the event structure unfolding of a net system is grid-free.

Finally, we also show that it involves no loss of generality to assume that the net system under study is *free choice*. The free choice property induces an important behavioral property called *confusion-freeness* which in turn adds considerable additional structure to the event structure unfolding [8,9]. Further the sub-class of free choice net systems have a pleasing and rich theory [10]. Thus our reduction shows that the full conjecture can be tackled in the more restricted and familiar setting provided by the free choice property.

We discuss in the final section the major barrier that must be overcome to settle the difficult open half of the conjecture. We also point to the rich existing literature on graph classes with decidable MSO theories [11–15] with the hope that a versatile reader may be able to exploit these results to settle our conjecture. In addition, we sketch how the positive resolution of our conjecture will have a significant impact on a variety of distributed controller synthesis problems.

In the literature there are two related results. The first one basically [16] says that the MSO theory of *every* net system is decidable provided quantification over sets is restricted to *conflict-free* subsets of events in the logic. However it is easy to find net systems that exhibit conflict and concurrency and yet admit a decidable MSO theory. Hence this result—valuable though it is—does not quite generalize Rabin's theorem. Further, as we point out in the final section, the restriction to quantification over conflict-free sets reduces its applicability in the context of distributed controller synthesis problems.

The second result [17] is phrased in terms of networks of transition systems that communicate by synchronizing on common actions [5]. The result says that the MSO theory of *connectedly* communicating processes (CCP) is decidable. As the name suggests, in a CCP, processes are required to communicate with each other frequently. More precisely, there is a bound K that depends only on the finite presentation of the system such that if process p executes K steps without hearing from process q either directly or indirectly and reaches a global state s, then starting from s it will never hear from q again, directly or indirectly. The corresponding class of net systems can exhibit both concurrency and conflict and yet their—unrestricted—MSO theory will be decidable.

Our motivation for studying branching time temporal logics in a concurrent setting mainly has to do with distributed controller synthesis. Specifically, for distributed systems one is often interested in constructing local finite memory strategies that do not depend on global information and hence can be synthesized as a distributed controller. For tackling these problems it is natural to look at partial order based branching time behaviors. Indeed, it is easy to show that a rich set of distributed controller synthesis problems can be effectively solved for net systems—and communicating networks of finite state systems—whose MSO theories are decidable. We address this application again in the concluding section.

In summary, we present here a conjecture which states that a natural extension of Rabin's theorem goes through in a partial order based branching time setting precisely when the grid-free property holds. We use net systems as the canonical system model and their event structure unfoldings as the models for the MSO logic to formulate the conjecture. We then prove it is decidable whether the event structure unfolding of a net system is grid-free. We also prove that if a net system is not grid-free then its MSO theory is undecidable. Finally we show that our conjecture can be reduced to the sub-class of *free choice* net systems.

1.1. Plan of the paper

In the next section we present the basic notions and related terminology and formulate our conjecture. In Section 3 we show that grid-freeness is a decidable property and in Section 4 we settle the easy half of the conjecture. In the subsequent section we establish the reduction to the free choice case and then conclude with a discussion.

2. Net systems and their event structure unfoldings

2.1. Net systems

A net system is a structure $N = (S, \Sigma, F, M_{in})$ where S and Σ are disjoint finite sets of places and actions, $F \subseteq (S \times \Sigma) \cup (\Sigma \times S)$ is the flow relation and $M_{in} \subseteq S$ is the initial marking. We have used "actions" instead of the more traditional "events" or "transitions" to avoid clashes with the terminology to be deployed below. We let a, b with or without subscripts to range over Σ.

As usual, for $v \in S \cup \Sigma$, we set ${}^\bullet v = \{u \mid (u, v) \in F\}$ and $v^\bullet = \{u \mid (v, u) \in F\}$. A *marking* of N is a subset of S. The transition relation $\longrightarrow_N \subseteq 2^S \times \Sigma \times 2^S$ is given by: $M \xrightarrow{a}_N M'$ iff ${}^\bullet a \subseteq M$, $(a^\bullet - {}^\bullet a) \cap M = \emptyset$ and $M' = (M - {}^\bullet a) \cup a^\bullet$. This is the usual firing rule for net systems stated here in terms of a transition relation. Where N is clear from the context we will write $\longrightarrow$ instead of $\longrightarrow_N$. Indeed we will follow these conventions for other entities associated with N to be introduced below. The transition relation $\longrightarrow$ is extended to sequences of actions as follows. For convenience this extension is also denoted as $\longrightarrow$ and is given by:

- $M \xrightarrow{\varepsilon} M$ for every marking M.
- Suppose $M \xrightarrow{\sigma} M'$ for $\sigma \in \Sigma^\star$ and $M' \xrightarrow{a} M''$ then $M \xrightarrow{\sigma a} M''$.

We will say that σ is a firing sequence at M if there exists a marking M' such that $M \xrightarrow{\sigma} M'$. We let *FS* denote the set of firing sequences at M_{in}. Finally we will say that the marking M is reachable iff there exists a firing sequence σ in *FS* such that $M_{in} \xrightarrow{\sigma} M$.

Suppose M is a reachable marking and $a \in \Sigma$ such that ${}^\bullet a \subseteq M$ but $a^\bullet \cap M \neq \emptyset$. Then according to our firing rule, a is *not* enabled at M. One usually calls this a *contact* situation. N is said to be contact-free if for every reachable marking M and every $a \in \Sigma$, if ${}^\bullet a \subseteq M$ then $a^\bullet \cap M = \emptyset$. It involves no loss of generality to assume that N is contact-free [18]. And we shall do so through the rest of the paper. (Basically we are then dealing with 1-safe Petri nets.)

2.2. Event structures

Given a net system N, there is a canonical way of associating a Σ-labelled event structure with N.

First recall that an event structure (often called a prime event structure) is a triple $ES = (E, \leqslant, \#)$ where (i) $(E, \leqslant)$ is a countable poset and (ii) for every $e \in E$, $\downarrow e = \{e' \in E \mid e' \leqslant e\}$ is a finite set and (iii) $\# \subseteq E \times E$ is an irreflexive and symmetric relation such that, for every e_1, e_2 and e_3, if $e_1 \# e_2$ and $e_2 \leqslant e_3$, then $e_1 \# e_3$.

E is the set of events, $\leqslant$ the causality relation and $\#$ the conflict relation. We define the concurrency relation $co \subseteq E \times E$ via: $e \; co \; e'$ iff $e \nleq e'$ and $e' \nleq e$ and $(e, e') \notin \#$. For $E' \subseteq E$ we define $\downarrow E' = \bigcup_{e \in E'} \downarrow e$. The states of an event structure are usually called configurations. A configuration is a set of events that have occurred so far while respecting the causality relation and the mutual exclusiveness implied by the conflict relation. More precisely, $c \subseteq E$ is a configuration iff $\downarrow c = c$ and $(c \times c) \cap \# = \emptyset$.

A Σ-labelled event structure is a structure $(E, \leqslant, \#, \lambda)$ where $(E, \leqslant, \#)$ is an event structure and $\lambda : E \to \Sigma$ a labelling function.

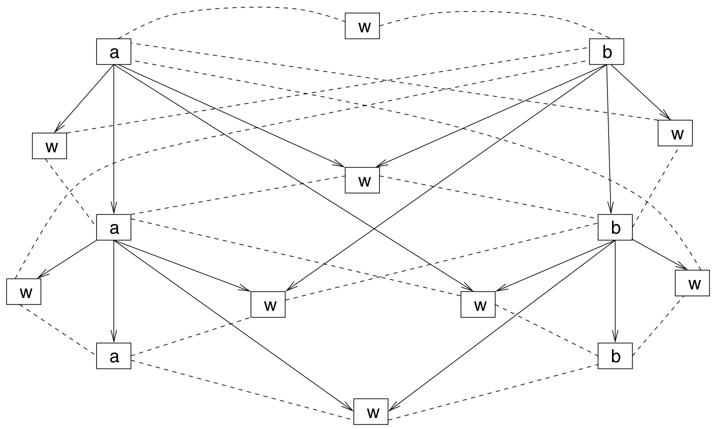

Fig. 2. The event structure of N_1.

2.3. The event structure unfolding of a net system

Next recall that the (Mazurkiewicz) trace alphabet associated with N is (Σ, I) where $a\ I\ b$ iff $(a^\bullet \cup {}^\bullet a) \cap (b^\bullet \cup {}^\bullet b) = \emptyset$. $D = (\Sigma \times \Sigma) - I$ is the dependence relation. The independence relation I induces in the usual way the equivalence relation $\sim_I \subseteq \Sigma^\star \times \Sigma^\star$. It is the least equivalence relation that satisfies: $\sigma_1 ab\sigma_2 \sim_I \sigma_1 ba\sigma_2$ whenever $a\ I\ b$. The $\sim_I$-equivalence class containing the string σ, called a *(Mazurkiewicz) trace*, will be denoted as $\langle\sigma\rangle$. If σ is non-null then $last(\sigma)$ will stand for the last letter of σ. The trace $\langle\sigma\rangle$ is *prime* iff σ is non-null and for every $\sigma' \in \langle\sigma\rangle$, $last(\sigma) = last(\sigma')$. Finally, the partial ordering relation $\sqsubseteq$ over traces is given by: $\langle\sigma\rangle \sqsubseteq \langle\tau\rangle$ iff there exists $\sigma' \in \langle\sigma\rangle$, $\tau' \in \langle\tau\rangle$ such that σ' is a prefix of τ'.

The event structure unfolding of N is $ES = (E, \leqslant, \#, \lambda)$ where

- E is the set of traces given by: $\langle\sigma\rangle \in E$ iff $\sigma \in FS$ and $\langle\sigma\rangle$ is a prime trace.
- $\leqslant$ is $\sqsubseteq$, restricted to $E \times E$.
- Let $e, e' \in E$. Then $e\ \#\ e'$ iff there does *not* exist a firing sequence σ in FS such that $e \sqsubseteq \langle\sigma\rangle$ and $e' \sqsubseteq \langle\sigma\rangle$.
- $\lambda : E \to \Sigma$ is given by: $\lambda(\langle\sigma\rangle) = last(\sigma)$.

In Fig. 2 we show an initial fragment of the event structure unfolding of the net system shown in Fig. 1. The dashed edges represent the "minimal" conflict relation $\#_{\min}$ where $e\ \#_{\min}\ e'$ iff $e\ \#\ e'$ and $(\downarrow e \times \downarrow e') \cap \# = \{(e, e')\}$. The directed edges represent the "immediate" causality relation $\lessdot$ where $e \lessdot e'$ iff $e < e'$ and $e \leqslant e'' \leqslant e'$ implies $e = e''$ or $e'' = e'$.

3. The conjecture

Through the remaining parts of the paper we fix a net system N and its event structure unfolding $ES = (E, \leqslant, \#, \lambda)$ as defined above.

3.1. The MSO logic

The syntax of the MSO logic over $ES = (E, \leqslant, \#, \lambda)$ is:

$$MSO(ES) ::= R_a(x) \mid x \leqslant y \mid x \in X \mid \exists x(\varphi) \mid \exists X(\varphi) \mid \sim\varphi \mid \varphi_1 \vee \varphi_2.$$

where $a \in \Sigma$, $x, y, \ldots$ are individual variables and $X, Y, \ldots$ are set variables. For convenience, we have used $\leqslant$ as a non-logical symbol while it also stands for the causality relation of ES. We will do this for other derived relations as well. This should cause no confusion.

An interpretation $\mathcal{I}$ assigns to every individual variable an event in E and every set variable, a subset of E. The notion of ES satisfying a formula φ under an interpretation $\mathcal{I}$, denoted $ES \models_{\mathcal{I}} \varphi$, is defined as follows.

- $ES \models_{\mathcal{I}} R_a(x)$ iff $\lambda(\mathcal{I}(x)) = a$.
- $ES \models_{\mathcal{I}} x \leqslant y$ iff $\mathcal{I}(x) \leqslant \mathcal{I}(y)$.
- $ES \models_{\mathcal{I}} x \in X$ iff $\mathcal{I}(x) \in \mathcal{I}(X)$.
- $ES \models_{\mathcal{I}} \exists x(\varphi)$ iff there exists $e \in E$ and an interpretation $\mathcal{I}'$ such that $ES \models_{\mathcal{I}'} \varphi$ where $\mathcal{I}'$ satisfies: $\mathcal{I}'(x) = e$ and $\mathcal{I}'(y) = \mathcal{I}(y)$ if y is an individual variable other than x. Further, $\mathcal{I}'(X) = \mathcal{I}(X)$ for every set variable X.
- $ES \models_{\mathcal{I}} \exists X(\varphi)$ iff there exists $E' \subseteq E$ and an interpretation $\mathcal{I}'$ such that $ES \models_{\mathcal{I}'} \varphi$ where $\mathcal{I}'$ satisfies: $\mathcal{I}'(x) = \mathcal{I}(x)$ for every individual variable x and $\mathcal{I}'(X) = E'$. Further, $\mathcal{I}'(Y) = \mathcal{I}(Y)$ if Y is a set variable other than X.

- $ES \models_{\mathcal{I}} \sim\varphi$ and $ES \models_{\mathcal{I}} \varphi_1 \vee \varphi_2$ are defined in the standard way.

The notion of a free variable and bound variables in a formula are defined as usual and a sentence is a formula which has no free variables. $ES \models \varphi$ will denote that ES is a model of the sentence φ. The MSO theory of ES is the set of sentences it satisfies. By the MSO theory of N we mean the MSO theory of its event structure unfolding. The MSO theory of N is said to be decidable iff there exists an effective procedure using which one can determine for each sentence φ in $MSO(ES)$, whether $ES \models \varphi$.

The MSO logic is quite expressive. To start with, we note that the other connectives of propositional logic such as $\wedge$, $\Rightarrow$ (implies) and $\equiv$ (if and only if), universal quantification over individual and set variables ($\forall x(\varphi)$, $\forall X(\varphi)$) as well as the set inclusion relation $\subseteq$ ($X \subseteq Y$) can all be defined easily.

Next we observe that the conflict relation of ES can be defined via:

- $x \mathrel{\widehat{\#}} y \triangleq \sim(x \leqslant y) \wedge \sim(y \leqslant x) \wedge \bigvee_{(a,b)\in D}(R_a(x) \wedge R_b(y))$. (Recall that D is the dependency relation over Σ.)
- $x \# y \triangleq \exists x' \exists y'(x' \leqslant x \wedge y' \leqslant y \wedge x' \mathrel{\widehat{\#}} y')$.

We can now define the concurrency relation *co* via: $x \; co \; y \triangleq \sim(x \leqslant y) \wedge \sim(y \leqslant x) \wedge \sim(x \# y)$. Recall that a set of events c is a *configuration* iff it is conflict-free (if $e, e' \in c$ then it is not the case that $e \# e'$) and downward closed ($e \in c$ and $e' \leqslant e$ implies $e' \in c$). Clearly the notion of a configuration can also be defined in the MSO logic. Consequently a rich variety of dynamical properties of the event structure unfolding can be defined in the logic.

3.2. Net systems with decidable MSO theories

The net system $N = (S, \Sigma, F, M_{in})$ is said to be *sequential* iff $|{}^{\bullet}a| = |a^{\bullet}| = 1$ for every a in Σ. Furthermore $|M_{in}| = 1$. Let $M_{in} = \{s_{in}\}$. Clearly $TS = (S, s_{in}, \Sigma, \longrightarrow)$ is finite state transition system and it is easy to see that the event structure unfolding of N will be order-isomorphic to the standard unwinding of TS as a computation tree rooted at s_{in}. Hence by Rabin's theorem the MSO theory of every sequential net system is decidable.

Next we define the net system $N = (S, \Sigma, F, M_{in})$ to be *(structurally) conflict-free* iff $|{}^{\bullet}s| = |s^{\bullet}| = 1$ for every s in S. Let $ES = (E, \leqslant, \#, \lambda)$ be the event structure unfolding of the conflict-free net system N. Then it is easy to verify that $\# = \emptyset$ and that ES is a possibly infinite (Mazurkiewicz) trace over the trace alphabet (Σ, I) in its standard representation as a Σ-labeled partially ordered set [5].

To see that the MSO theory of ES is decidable we first recall the MSO logic of (finite and infinite) traces over (Σ, I) [5]. Let us denote this logic as $MSO(\Sigma, I)$. Its syntax will be the same as the syntax of the MSO logic over event structures that we have defined but it will be interpreted over finite and infinite traces generated by the trace alphabet (Σ, I). In this setting a trace is represented as a Σ-labeled partial order $(\widehat{E}, \widehat{\leqslant}, \hat{\lambda})$ that satisfies:

- $(\widehat{E}, \widehat{\leqslant})$ is a poset.
- For each e, $\downarrow e$ is a finite set.
- If $e \mathrel{\widehat{\lessdot}} e'$ then $\hat{\lambda}(e) \; D \; \hat{\lambda}(e')$.
- For any e, e', if $\hat{\lambda}(e) \; D \; \hat{\lambda}(e')$ then $e \mathrel{\widehat{\leqslant}} e'$ or $e' \mathrel{\widehat{\leqslant}} e$.

The satisfiability problem for $MSO(\Sigma, I)$ is decidable [5]. More precisely, given a sentence φ in $MSO(\Sigma, I)$, one can effectively determine whether there exists a trace $(\widehat{E}, \widehat{\leqslant}, \hat{\lambda})$ ($\widehat{E}$ may be finite or infinite) over (Σ, I) such that $(\widehat{E}, \widehat{\leqslant}, \hat{\lambda}) \models \varphi$. In what follows, we construct the sentence φ_{ES} such that a trace is a model for φ_{ES} iff it is the event structure unfolding ES of N. It will then follow that for a sentence φ in the MSO logic of ES, $ES \models \varphi$ iff there exists a trace over (Σ, I) which is a model for $\varphi_{ES} \wedge \varphi$.

Let $config(c)$ be the predicate defined in $MSO(\Sigma, I)$ asserting that c (treated as a set variable) is a finite configuration. We next define the predicate $en(c, x) \triangleq config(c) \wedge (x \notin c) \wedge config(c \cup \{x\})$. We are being informal here for convenience. One can easily convert these semi-formal descriptions to formulas. We will allow ourselves this freedom through the rest of the paper to avoid notational clutter.

Next we define the predicate $min(x) \triangleq \forall y(y \leqslant x \Rightarrow y = x)$. Let $A_0, A_1 \subseteq \Sigma \times \Sigma$ be given by: $(b, a) \in A_0$ iff there exists $s \in S$ with ${}^{\bullet}s = \{b\}$ and $s^{\bullet} = \{a\}$ and $s \notin M_{in}$; $(b, a) \in A_1$ iff there exists $s \in S$ with ${}^{\bullet}s = \{b\}$ and $s^{\bullet} = \{a\}$ and $s \in M_{in}$. Now the formula φ_{ES} is $\varphi_1 \wedge \varphi_2$ where:

- $\varphi_1 \triangleq \forall x(min(x) \equiv \bigvee_{{}^{\bullet}a \subseteq M_{in}} R_a(x))$.
- φ_2 will say that if c is a configuration then there exists x satisfying $en(c, x)$ and $R_a(x)$ iff (i) in case $(b, a) \in A_1$, then $y \in c$ and $R_a(y)$ implies there exists $y' \in c$ such that $y < y'$ and $R_b(y')$; (ii) in case $(b, a) \in A_0$, there exists $y \in c$ such that $R_b(y)$ and if $y' \in c$ and $R_a(y')$ it is not the case that $y < y'$.

Now it is easy to verify that for any sentence φ in the MSO logic of ES, $ES \models \varphi$ iff there exists a trace which is a model for $\varphi_{ES} \wedge \varphi$. Hence the MSO theory of every conflict-free net system is decidable.

The class of systems identified in [17] can also be represented as net systems (accompanied by a labeling of the actions). This class will *properly* include the class of sequential and conflict-free net systems. From the main result of [17] it follows that the MSO theory of every net system in this class is also decidable.

3.3. Grid-freenes

It turns out however that if the event structure unfolding of a net system is not *grid-free* then its MSO theory is undecidable.

Definition 1. The event structure $ES = (E, \leqslant, \#)$ is *grid-free* iff there does *not* exist pairwise disjoint subsets X, Y, Z of E satisfying the following conditions.

- $X = \{x_0, x_1, x_2, \ldots\}$ is an infinite set with $x_0 < x_1 < x_2 \cdots$.
- $Y = \{y_0, y_1, y_2, \ldots\}$ is an infinite set with $y_0 < y_1 < y_2 \cdots$.
- $X \times Y \subseteq co$.
- There exists an injective mapping $g : X \times Y \to Z$ satisfying: If $g(x_i, y_j) = z$ then $x_i < z$ and $y_j < z$. Furthermore, if $i' > i$ then $x_{i'} \not< z$ and if $j' > j$ then $y_{j'} \not< z$.

The Σ-labelled event structure $(E, \leqslant, \#, \lambda)$ is said to be grid-free iff $(E, \leqslant, \#)$ is grid-free. We shall also say that N is grid-free just in case ES_N is grid-free. This leads to:

Conjecture *(Thiagarajan). The MSO theory of a net system is decidable iff it is grid-free.*

3.4. Grid-freeness is a recursive property

Before examining the conjecture in more detail, we wish to show that one can decide if a net system is grid-free. Informally, grid-freeness of N can be understood as follows. First we note that one can view N as the parallel composition of a network of finite transition systems that synchronize on common actions by performing a transformation called place complementation [19]. In this light, the net system N is grid-free iff there exists a K which depends only on the finite presentation of N such that starting from a reachable marking M, if N can reach a marking M' via a firing sequence σ at M, then the following is satisfied for any two agents p and q.

(i) If p executes K transitions along σ without hearing from q, directly or indirectly;
(ii) and q executes K transitions along σ without hearing from p, directly or indirectly;
(iii) then starting from M' the agents p and q will never hear from each other again, directly or indirectly.

Here is a more precise characterization of grid-freeness.

Theorem 2. *N is grid-free iff there does* not *exist a reachable marking M and firing sequences $\alpha, \beta, \gamma \in \Sigma^+$ at M such that*

- *α, β are "cycles". That is, $M \xrightarrow{\alpha} M$ and $M \xrightarrow{\beta} M$.*
- *α, β are independent. In other words, if a occurs in α and b occurs in β then a I b.*
- *$\langle\gamma\rangle$ is a* prime *trace and there exists a non-null prefix α' of α and a non-null prefix β' of β such that $\langle\alpha'\rangle \sqsubseteq \langle\gamma\rangle$ and $\langle\beta'\rangle \sqsubseteq \langle\gamma\rangle$.*

To prove this result we need some additional notations and preliminary observations. Suppose A is a non-null finite subset of E with $ES = (E, \leqslant, \#, \lambda)$. Then a linearization of A is a sequence of events $\tau = e_1 e_2 \ldots e_k$ such that each e in A appears exactly once in τ and if $e_i < e_j$ in ES then $i < j$. We let $lin(A)$ denote the set of linearizations of A. Next we denote the finite configurations of ES as C_{ES}. We will often write C when ES is clear from the context.

Let TR_N be the set of finite traces corresponding to the firing sequences of N. In other words, $TR_N = \{\langle\sigma\rangle \mid \sigma \in FS\}$. We now define the map $ct : C \to TR_N$ via: $ct(c) = \langle\sigma\rangle$ iff there exists τ in $lin(c)$ such that $\tau = e_1 e_2 \ldots e_k$ and $\sigma = a_1 a_2 \ldots a_k$ with $a_i = \lambda(e_i)$ for $1 \leqslant i \leqslant k$. In case $c = \emptyset$ we set $ct(c) = \langle\varepsilon\rangle$. It is easy to verify that ct is a well-defined map. Conversely we define the map $tc : TR_N \to C$ via: $tc(\langle\sigma\rangle) = \{e \mid e$ is prime trace and $e \sqsubseteq \langle\sigma\rangle\}$. Again it follows from the definitions that tc is well-defined. Next the map $mark : C \to 2^S$ is defined as: $mark(c) = M$ iff there exists a firing sequence σ in FS such that $ct(c) = \langle\sigma\rangle$ and $M_{in} \xrightarrow{\sigma} M$. Thus $mark(c)$ is a reachable marking and the fact that $mark$ is well-defined is also easy to check. The following facts too follow easily from the literature [5] and hence we will state them without a proof. In the statement of the result, $Alph(\sigma)$ will denote the set of letters that appear in the sequence σ one or more times.

Proposition 3. *Recall that $ES = (E, \leqslant, \#, \lambda)$. Let e, e' range over E. Then the following hold:*

(i) *For each* e, $\downarrow e \in C$ *and* $ct(\downarrow e)$ *is a prime trace.*
(ii) *If* e *co* e', *then* $\lambda(e)\ I\ \lambda(e')$.
(iii) *If* $\lambda(e)\ D\ \lambda(e')$, *then one of* $e \leqslant e'$, $e' \leqslant e$, $e \# e'$ *must hold.*
(iv) *If* e *co* e', *then* $\downarrow e \cap \downarrow e'$ *and* $\downarrow e \cup \downarrow e'$ *are both configurations.*
(v) $e \# e'$ *iff there does not exist a configuration* c *in* C *such that* $e, e' \in c$.
(vi) *Suppose* $c, c' \in C$ *where* $c \subset c'$. *Let* $M = mark(c)$. *If* $ct(c) = \langle\sigma\rangle$ *for some firing sequence* σ *in FS then there exists a firing sequence* α *at* M *such that* $ct(c') = \langle\sigma\alpha\rangle$.
(vii) *If* $e < e'$ *then there exists a sequence of events* $d_1 d_2 \ldots d_m$ *such that* $d_1 = e$, $d_m = e'$ *and for* $1 \leqslant j < m$, $d_j < d_{j+1}$ *and* $\lambda(d_j)\ D\ \lambda(d_{j+1})$. *Further,* $\lambda(d_i) \neq \lambda(d_j)$ *for* $1 \leqslant i < j \leqslant m$ *and hence* $m \leqslant |\Sigma|$.

With these preliminaries out of the way, we can prove Theorem 2.

Proof of Theorem 2. Suppose there exist a reachable marking M and firing sequences $\alpha, \beta, \gamma, \alpha', \beta'$ satisfying the asserted conditions. Fix a firing sequence σ at the initial marking M_{in} such that $M_{in} \xrightarrow{\sigma} M$.

Construct the set $X = \{x_0, x_1, \ldots\}$ of events using σ, α, α' as follows. Fix event x_0 to be some event in $tc(\langle\sigma\alpha'\rangle) - tc(\langle\sigma\rangle)$. For each $i > 0$, choose x_i to be some event in $tc(\langle\sigma\alpha^i\alpha'\rangle) - tc(\langle\sigma\alpha^i\rangle)$ such that $\lambda(x_i) = \lambda(x_0)$. The choice of x_i can be easily made unique as N is finite. Define the set $Y = \{y_0, y_1, \ldots\}$ of events using σ, β, β' similarly. Specifically, the event y_0 is fixed to be some event in $tc(\langle\sigma\beta'\rangle) - tc(\langle\sigma\rangle)$. For each $i > 0$, y_i is some event in $tc(\langle\sigma\beta^i\beta'\rangle) - tc(\langle\sigma\beta^i\rangle)$ such that $\lambda(y_i) = \lambda(y_0)$.

Let $Z = \{z_{ij} \mid i, j \in \mathbb{N}\}$ where z_{ij} is the event $\langle\sigma\alpha^i\beta^j\gamma\rangle$. Define the map $g : X \times Y \to Z$ as $g(x_i, y_j) = z_{ij}$ for $i, j \in \mathbb{N}$. It follows easily that X, Y, Z fulfill the conditions asserted in Definition 1. Hence N is not grid-free.

To prove the second half, suppose N is not grid-free, that is, there exist X, Y, Z fulfilling the conditions in Definition 1 with the associated injective mapping g from $X \times Y$ to Z. Let $X = \{x_0, x_1, \ldots\}$, $Y = \{y_0, y_1, \ldots\}$. The proof will be based on the following claim.

Claim 4. *There exist* x, x' *in* X *and* y, y' *in* Y *which satisfy the conditions:*

(i) $x < x'$, $y < y'$.
(ii) $mark(\downarrow x) = mark(\downarrow x')$, *and* $mark(\downarrow y) = mark(\downarrow y')$.
(iii) $(\downarrow x' - \downarrow x) \cap \downarrow y' = \emptyset$. *And* $(\downarrow y' - \downarrow y) \cap \downarrow x' = \emptyset$.
(iv) $Alph(\downarrow x' - \downarrow x) \times Alph(\downarrow y' - \downarrow y) \subseteq I$.

Proof. Clearly if X' is an infinite subset of X and Y' an infinite subset of Y, then X', Y', Z also fulfills the conditions in Definition 1. Noting this and that N is finite, it involves no loss of generality to further assume that X, Y satisfies:

- $mark(\downarrow x_0) = mark(\downarrow x_1) = \cdots$ and $\lambda(x_0) = \lambda(x_1) = \cdots$.
- $mark(\downarrow y_0) = mark(\downarrow y_1) = \cdots$ and $\lambda(y_0) = \lambda(y_1) = \cdots$.

For each $i > 0$, let $E_i = \downarrow x_i - \downarrow x_{i-1}$ and $AL_i = \{\lambda(e) \mid e \in E_i\}$. Thus one can find indices $h_1, h_2, \ldots, h_k$ such that:

- $k = |\Sigma| + 1$. And $0 < h_1 < h_2 < \cdots < h_k$.
- $AL_{h_1} = AL_{h_2} = \cdots = AL_{h_k}$.

We now fix $x = x_{h_k - 1}$, $x' = x_{h_k}$. We then fix y, y' in Y in the same manner.

It follows that x, x', y, y' satisfy the first two conditions stated in Claim 4. To show that x, x', y, y' satisfy the third condition, we first argue that:

($\star$) If $e \in \downarrow x_{h_k} - \downarrow x_{h_k - 1}$ then $x_{h_1} < e$.

Let $e \in \downarrow x_{h_k} - \downarrow x_{h_k - 1}$. Since $e \leqslant x_{h_k}$, it follows from Proposition 3 that there exist events $e_1, \ldots, e_u$ in $\downarrow x_{h_k} - \downarrow x_{h_k - 1}$ such that $u \leqslant |\Sigma|$, $e = e_1$, $e_u = x_{h_k}$, $e_1 < \cdots < e_u$, and $\lambda(e_j)\ D\ \lambda(e_{j+1})$ for $1 \leqslant j < u$. As $AL_{h_1} = AL_{h_2} = \cdots = AL_{h_k}$, and also $\lambda(x_{h_1}) = \lambda(x_{h_2}) = \cdots = \lambda(x_{h_k})$, one can pick events $f_1, f_2, \ldots, f_u$, such that $f_1 = e$, $f_u = x_{h_1}$, and for $1 < j < u$, $f_j \in E_{h_{k-j+1}}$ with $\lambda(f_j) = \lambda(e_j)$ (note that $u \leqslant |\Sigma|$ and $k = |\Sigma| + 1|$).

We argue that $f_2 < f_1$. Since $\lambda(f_1)\ D\ \lambda(f_2)$, one of $f_1 \leqslant f_2$, $f_2 \leqslant f_1$, $f_1 \# f_2$ must hold (cf. Proposition 3). The case $f_1 \# f_2$ is ruled out by that $f_1 \in \downarrow x_{h_k}$, $f_2 \in \downarrow x_{h_{k-1}}$, and $x_{h_{k-1}} \leqslant x_{h_k}$. The case $f_1 \leqslant f_2$ leads to a contradiction, due to that $f_1 \notin \downarrow x_{h_k - 1}$, $f_2 \leqslant x_{h_{k-1}}$, and $x_{h_{k-1}} \leqslant x_{h_k - 1}$ (as $h_{k-1} < h_k$). Thus, it can only be that $f_2 < f_1$.

By extending the above reasoning, one conclude that $f_u < \cdots < f_1$, and in particular, $x_{h_1} < e$.

Having shown ($\star$) above, it follows that $(\downarrow x_{h_k} - \downarrow x_{h_k - 1}) \cap \downarrow y' = \emptyset$. For otherwise, if $e \in (\downarrow x_{h_k} - \downarrow x_{h_k - 1}) \cap \downarrow y'$, then $x_{h_1} < e < y'$ contradicting that $X \times Y \subseteq co$. That $(\downarrow y' - \downarrow y) \cap \downarrow x' = \emptyset$ can be similarly shown, and thus x, x', y, y' fulfills condition (iii) in Claim 4.

Lastly, to see that x, x', y, y' fulfills condition (iv), note that for any $e \in \downarrow x' - \downarrow x$, $f \in \downarrow y' - \downarrow y$, we have $e\ co\ f$ and thus $\lambda(e)\ I\ \lambda(f)$ by Proposition 3. For, $e \leqslant f$ would lead to $e \leqslant y'$, contradicting that $(\downarrow x' - \downarrow x) \cap \downarrow y' = \emptyset$. Similarly, $f \leqslant e$ cannot hold. And $e \# f$ contradicts that there exists z in Z with $x' < z$, $y' < z$.

—End of the proof of Claim 4 □

Having established Claim 4, let $M = mark(\downarrow x \cup \downarrow y)$ and σ be such that $ct(\downarrow x \cup \downarrow y) = \langle\sigma\rangle$. The firing sequences α, β at M are chosen such that: $ct(\downarrow x' \cup \downarrow y) = \langle\sigma\alpha\rangle$ and $ct(\downarrow y' \cup \downarrow x) = \langle\sigma\beta\rangle$. Note that $Alph(\alpha) = \{\lambda(e) \mid e \in \downarrow x' - (\downarrow x \cup \downarrow y)\}$. But then $\downarrow x' - (\downarrow x \cup \downarrow y) = \downarrow x' - \downarrow x$ due to condition (iii) in Claim 4. Similarly, $Alph(\beta) = \{\lambda(e) \mid e \in \downarrow y' - \downarrow y\}$. Thus $Alph(\alpha) \times Alph(\beta) \subseteq I$.

Let z be in Z such that $g(x', y') = z$. Since $x' < z$ and $y' < z$, we know that $\downarrow x' \cup \downarrow y' \subset \downarrow z$. Since $Alph(\alpha) \times Alph(\beta) \subseteq I$, we are also assured that $ct(\downarrow x' \cup \downarrow y') = \langle\sigma\alpha\beta\rangle$. We now fix γ' to be such that $ct(\downarrow z) = \langle\sigma\alpha\beta\gamma'\rangle$ and set $\gamma = \alpha\beta\gamma'$. It follows that γ is a prime trace. We note that $\langle\sigma\alpha\beta\rangle = \langle\sigma\beta\alpha\rangle$ and $\langle\sigma\alpha\rangle \sqsubseteq \langle\sigma\gamma\rangle$, $\langle\sigma\beta\rangle \sqsubseteq \langle\sigma\gamma\rangle$. By taking $\alpha' = \alpha$, $\beta' = \beta$, it is now easy to verify that M, α, β, γ fulfill the conditions in Theorem 2.

—End of the proof of Theorem 2 □

We can now show one can effectively determine if a net system N is grid-free. One can do so by constructing a finite automaton running over strings in $\Sigma^\star$ which checks for existence of M, α, β, γ. The construction we have in mind in fact makes use of the following characterization of grid-freeness which follows easily from the proof of Theorem 2 and Proposition 3.

Corollary 5. *N is grid-free iff there does* not *exist a reachable marking M and firing sequences α, β, γ at M such that:*

- $M \xrightarrow{\alpha} M$, $M \xrightarrow{\beta} M$.
- $Alph(\alpha) \times Alph(\beta) \subseteq I$.
- *Each of* $\langle\alpha\rangle, \langle\beta\rangle, \langle\gamma\rangle$ *is a prime trace.*
- $\langle\alpha\rangle \sqsubseteq \langle\gamma\rangle$, $\langle\beta\rangle \sqsubseteq \langle\gamma\rangle$.

One can easily further demand that α, β, γ are of length at most K for some constant K that depends only on N. Further, trace theory [5] suggests that the last condition can be phrased in terms of sequences. That is, there exists a sequence γ' in Σ^+ which satisfies the following three conditions: Firstly, $\langle\gamma\rangle = \langle\alpha\beta\gamma'\rangle = \langle\beta\alpha\gamma'\rangle$. Secondly, there exists a subsequence $a_1a_2\ldots a_u$ of γ' such that $u \leqslant |\Sigma|$, $last(\alpha)\ D\ a_1$, $a_i\ D\ a_{i+1}$ for $1 \leqslant i < u$, and $a_u = last(\gamma')$. Lastly, there exists a subsequence $b_1b_2\ldots b_v$ of γ' such that $v \leqslant |\Sigma|$, $last(\beta)\ D\ b_1$, $b_i\ D\ b_{i+1}$ for $1 \leqslant i < v$, and $b_v = last(\gamma')$. It is also well-known [5] that one can effectively construct a finite state automaton which accepts a sequence τ iff $\langle\tau\rangle$ is a prime trace. With these observations in place, it is easy to construct a finite state automaton $\mathcal{A}$ which first guesses a finite firing sequence σ in FS to reach a marking M. It then guesses sequences of bounded lengths α, β, γ' and checks whether $\alpha\beta\gamma'$ is a firing sequence at M and whether α, β and γ' fulfill the conditions stated above. It then follows that N is grid-free iff the language recognized by $\mathcal{A}$ is empty.

4. Undecidability

Here we shall prove the relatively easy half of the conjecture, namely, if a net system is not grid-free then its MSO theory is undecidable.

Assume that N is not grid-free. Then we shall reduce each instance CP of the following coloring problem to a sentence φ in $MSO(ES)$ such that CP has a solution iff $ES \models \varphi$. It follows easily from [7,20] that this coloring problem is undecidable. Hence we will be able to conclude that the MSO theory of N is undecidable. An instance CP of the coloring problem consists of a finite set of colors $Col = \{\kappa_0, \kappa_1, \ldots, \kappa_J\}$ and two functions $R : Col \to 2^{Col} - \{\emptyset\}$ and $U : Col \to 2^{Col} - \{\emptyset\}$. The problem is to find a coloring function $f : \mathbb{N} \times \mathbb{N} \to Col$ such that $f(0,0) = \kappa_0$ and for each $i, j \in \mathbb{N}$ it is the case that $f(i+1, j) \in R(f(i,j))$ and $f(i, j+1) \in U(f(i,j))$.

We will construct a sentence of the form $\varphi_{grid} \wedge \varphi_{CP}$ such that $ES \models \varphi_{grid} \wedge \varphi_{CP}$ iff CP admits a solution. To aid readability we will describe φ_{grid} and φ_{CP} in an informal fashion. It will be clear how this description can be converted into a sentence in the MSO logic of N.

First note that the equality predicate is definable in the logic via: $x = y \triangleq x \leqslant y \wedge y \leqslant x$. The sentence φ_{grid} will basically say that there exist sets X, Y and Z such that they together constitute the two dimensional grid. With this in mind, we first define the predicate $grid(x, y, z)$ as:

$$\begin{aligned}
&(x \in X \wedge y \in Y \wedge z \in Z) \wedge (x < z \wedge y < z) \\
&\quad \wedge \big(\forall x'\big(x' \in X \wedge x < x' \Rightarrow \sim\!\big(x' < z\big)\big) \\
&\quad \wedge \big(\forall y'\big(y' \in Y \wedge y < y' \Rightarrow \sim\!\big(y' < z\big)\big).
\end{aligned}$$

The sentence φ_{grid} will then consist of the conjunction of the following statements.

- There exist three sets X, Y and Z such that they are pairwise disjoint and x *co* y for every x in X and y in Y.
- X, Y and Z are infinite sets. This is easy to express. For instance one can say X is an infinite set via: $\exists x(x \in X) \wedge \forall x(x \in X \Rightarrow \exists x'(x' \in X \wedge x < x'))$.
- $\leqslant$ restricted to X and Y are linear orders.
- $\forall x \forall y(x \in X \wedge y \in Y \Rightarrow \exists z(grid(x, y, z)))$.
- $\forall x \forall y \forall z \forall z'(grid(x, y, z) \wedge grid(x, y, z') \Rightarrow z = z')$.

To construct φ_{CP} we will make use of the predicates $Right(x, x')$ and $Up(y, y')$ defined as:

$$Right(x, x') \triangleq x < x' \wedge \forall x''\big((x'' \in X \wedge x \leqslant x'' \leqslant x') \Rightarrow x = x'' \vee x'' = x'\big),$$
$$Up(y, y') \triangleq y < y' \wedge \forall y''\big((y'' \in Y \wedge y \leqslant y'' \leqslant y') \Rightarrow y = y'' \vee y'' = y'\big).$$

We can now construct φ_{CP} as the conjunction of following statements.

- There exist $Z_0, Z_1, \ldots, Z_J$ such that for $1 \leqslant j \leqslant J$, each Z_j is a subset of Z. And each z in Z belongs to exactly one Z_j in the family $\{Z_0, Z_1, \ldots, Z_J\}$.
- If x_0 is the least element of X and y_0 is the least element of Y and $grid(x_0, y_0, z)$, then $z \in Z_0$.
- If $grid(x, y, z)$, $z \in Z_j$, $Right(x, x')$, $grid(x', y, z')$ and $z' \in Z_\ell$ then $\kappa_\ell \in R(\kappa_j)$.
- If $grid(x, y, z)$, $z \in Z_j$, $Up(y, y')$, $grid(x, y', z')$ and $z' \in Z_\ell$ then $\kappa_\ell \in U(\kappa_j)$.

It is now easy to argue that $ES \models \varphi_{grid} \wedge \varphi_{CP}$ iff CP has a solution. The crucial point being of course $ES \models \varphi_{grid}$ iff ES is not grid-free. This leads to:

Theorem 6. *If a net system is not grid-free then its MSO theory is undecidable.*

Obviously the net system shown in Fig. 1 is not grid-free. Hence by the above theorem its MSO theory is undecidable.

5. Reduction to the free choice case

Here we wish to show that it suffices to prove the decidability half of the conjecture for the restricted class of free choice net systems. The net system $N = (S, \Sigma, F, M_{in})$ is *free choice* iff for every (s, a) in $F \cap (S \times \Sigma)$ it is the case $s^\bullet = \{a\}$ or ${}^\bullet a = \{s\}$. The net system shown in Fig. 1 is not free choice since $(s, w) \in F$ but $s^\bullet = \{a, w\}$ while ${}^\bullet w = \{s, t\}$. As mentioned earlier, free choice net systems have a rich theory and their event structure unfoldings are endowed with considerable additional structure due the behavioral property known as confusion-freeness [10,8].

We wish to show the following: Given a net system N one can effectively construct a free choice net system N' such that, for any sentence φ in the MSO theory of N one can effectively determine a sentence φ' in the MSO theory of N' with the property that $ES \models \varphi$ iff $ES' \models \varphi'$. Here ES is the event structure unfolding of N and ES' is the event structure unfolding of N'.

Let $N = (S, \Sigma, F, M_{in})$ be a net system which is not free choice. First let $F_{nfc} \subseteq F$ be given by:

$$(s, a) \in F_{nfc} \text{ iff } s \in S,\ a \in \Sigma \text{ and } s^\bullet \neq \{a\} \text{ and } {}^\bullet a \neq \{s\}.$$

Then $N' = (S', F', \Sigma', M'_{in})$ is defined via:

- $S' = S \cup S_{fc}$ where $S_{fc} = \{s_a \mid (s, a) \in F_{nfc}\}$.
- $\Sigma' = \Sigma \cup \Sigma_{fc}$ where $\Sigma_{fc} = \{a_s \mid (s, a) \in F_{nfc}\}$.
- $F' = (F - F_{nfc}) \cup \{(s, a_s), (a_s, s_a), (s_a, a) \mid (s, a) \in F_{nfc}\}$.
- $M'_{in} = M_{in}$.

It is easy to verify that N' is a 1-safe free choice net system.

In Fig. 3 we show the resulting free choice net system when we apply the above construction to the net system shown in Fig. 1.

In order to establish the required relationship between N and N' we need some additional notations. For a in Σ, let $pre_{nfc}(a) = \{s \mid (s, a) \in F_{nfc}\}$. We now define $pred(a) = a_{s_1} a_{s_2} \ldots a_{s_m}$ if $pre_{nfc}(a) = \{s_1, s_2, \ldots, s_m\}$. (The indexing of this set can be made canonical by fixing a linear order over S). In case $pre_{nfc}(a) = \emptyset$ we set $pred(a) = \varepsilon$.

FS' is the set of firing sequences of N' and (Σ', I') be trace alphabet of N' induced by F' in the usual way. We denote by $\sim_{I'}$ the associated equivalence relation $\sim_{I'} \subseteq (\Sigma')^\star \times (\Sigma')^\star$. Now suppose $\sigma' \in FS'$. We will say that σ' is in *standard* form if it can be written as:

$$\sigma' = pred(a_1)a_1 pred(a_2)a_2 \ldots pred(a_n)a_n \tau' \quad \text{where } a_i \in \Sigma \text{ for } 1 \leqslant i \leqslant n \text{ and } \tau' \in (\Sigma_{fc})^\star.$$

We will refer to τ' as the *tail* of σ' if it is in standard form. Further σ' will be said to be *complete* if it is in standard from and its tail is the null sequence. It is easy to use these definitions to establish the next result.

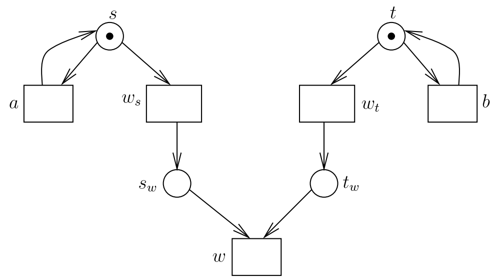

Fig. 3. The free choice net system N_1'.

Lemma 7. *If $\sigma' \in FS'$ then there exists $\sigma'' \in FS'$ such that $\sigma' \sim_{I'} \sigma''$ and σ'' is in standard form. Moreover $|\tau'| \leqslant |F_{nfc}|$ where τ' is the tail of σ''.*

Next let $ES = (E, \leqslant, \#, \lambda)$ be the event structure unfolding of N while $ES' = (E', \leqslant', \#', \lambda')$ is the event structure unfolding of N'. Then N and N' are strongly related to each other in the following sense.

Lemma 8. *The following statements hold:*

(i) *$FS = Proj_\Sigma(FS')$ where $Proj_\Sigma(\sigma)$ is the sequence obtained by erasing from σ all appearances of letters that are not in Σ. (Recall that FS is the set of firing sequences of N).*
(ii) *$I = I' \cap (\Sigma \times \Sigma)$.*
(iii) *Let $ES'' = (E'', \leqslant'', \#'', \lambda'')$ be given by: $E'' = \{e'' \mid e'' \in E', \lambda(e'') \in \Sigma\}$. Further, $\leqslant''$ is $\leqslant'$ restricted to $E'' \times E''$, $\#''$ is $\#'$ restricted to $E'' \times E''$ and λ'' is λ' restricted to E''. Then ES'' is an event structure. Moreover ES and ES'' are isomorphic in the sense that there is a bijection $H : E \to E''$ which satisfies: for every $e_1, e_2 \in E$, $e_1 \leqslant e_2$ iff $H(e_1) \leqslant'' H(e_2)$, and $e_1 \# e_2$ iff $H(e_1) \#'' H(e_2)$. Moreover $\lambda(e) = \lambda''(H(e))$ for every $e \in E$.*

Proof. To prove the first part of the lemma, let $\sigma \in FS$. If $\sigma = \varepsilon$ then $\sigma \in Proj_\Sigma(FS')$ since $\varepsilon \in FS'$. So suppose $\sigma = a_1a_2 \dots a_n$. Then from the construction of N' it follows that $\sigma' = pred(a_1)a_1pred(a_2)a_2 \dots pred(a_n)a_n \in FS'$. Moreover $Proj_\Sigma(\sigma') = \sigma$. Hence $FS \subseteq Proj_\Sigma(FS')$.

Next consider $\sigma' \in FS'$. If $\sigma' \sim_{I'} \sigma''$ then $Proj_\Sigma(\sigma') = Proj_\Sigma(\sigma'')$. Hence by the previous lemma it involves no loss of generality to assume that σ' is in standard form. Suppose $\sigma' = pred(a_1)a_1pred(a_2)a_2 \dots pred(a_n)a_n\tau'$. Then clearly $\sigma'' = pred(a_1)a_1pred(a_2)a_2 \dots pred(a_n)a_n$ is complete. An easy argument shows that $a_1a_2 \dots a_n \in FS$. But then $Proj_\Sigma(\sigma') = Proj_\Sigma(\sigma'')$ and hence $Proj_\Sigma(\sigma') \in FS$. Thus $Proj_\Sigma(FS') \subseteq FS$.

The second part follows from the definitions. The proof of the last part of the lemma is tedious but straightforward and hence we shall omit it. □

Now let φ be a sentence in the MSO theory of N and φ' be the sentence in the MSO theory of N' where φ' is defined as:

$$\varphi' = \exists E \left(\forall x \left(x \in E \equiv \bigvee_{a \in \Sigma} R_a(x) \right) \wedge \|\varphi\| \right)$$

In the above sentence E is being used—by abuse of notation—as a set variable. Its intended interpretation is the set of events of ES and this is captured by $\forall x(x \in E \equiv \bigvee_{a\in\Sigma} R_a(x))$. The formula $\|\varphi\|$ is defined via structural induction as:

- $\|R_a(x)\| \triangleq x \in E \wedge R_a(x)$.
- $\|x \leqslant y\| \triangleq x \in E \wedge y \in E \wedge x \leqslant y$.
- $\|x \in X\| \triangleq x \in E \wedge \forall y(y \in X \Rightarrow y \in E) \wedge x \in X$.
- $\|\exists x(\varphi)\| \triangleq \exists x(x \in E \wedge \|\varphi\|)$.
- $\|\exists X(\varphi)\| \triangleq \exists X(\forall x(x \in X \Rightarrow x \in E) \wedge \|\varphi\|)$.
- $\|{\sim}\varphi\| \triangleq {\sim}\|\varphi\|$.
- $\|\varphi_1 \vee \varphi_2\| \triangleq \|\varphi_1\| \vee \|\varphi_2\|$.

It is now easy to show:

Theorem 9. *Let φ be a sentence in the MSO theory of N and φ' be the sentence in the MSO theory of N' defined as above. Then $ES \models \varphi$ iff $ES' \models \varphi'$ where ES is the event structure unfolding of N and ES' is the event structure unfolding of N'.*

6. Discussion

An obvious line of attack to settle the conjecture is to reduce it to Rabin's theorem. In fact this seems to be the most viable approach as well since the MSO theory of regular trees is very powerful (see e.g. [1,21]). The idea is to construct a regular tree to represent the events of ES_N where N is a grid-free net system. The hard part is to find an encoding such that the fact that it is representing the events of ES can be expressed in the MSO logic over trees. This is in fact the route followed in [17]. The key difficulty in the current setting is to find an encoding in which one can describe the causality relation $<$. It could well be that with $e < e''$, $e' < e''$ and $e\ co\ e'$, the node representing e or the node representing e' does not appear as an ancestor of the node representing e''. One must still somehow describe the fact that e (or e') is earlier than e'' (in the causal order of the event structure). The problem gets aggravated by the fact that $\downarrow e'' - \downarrow e$ or $\downarrow e'' - \downarrow e'$ may be arbitrarily large.

There are of course other important techniques for proving decidability of the MSO theory of an infinite (directed or undirected) graph. For instance, it has been shown in [12–14] that if the MSO theory of a graph G is decidable then the unfolding of G from a definable vertex and also the "tree iteration" of G both have decidable MSO theories. Building on Rabin's theorem, these techniques have led to identification of numerous infinite graphs with decidable MSO theories [22,15]. In particular, by extending these techniques, a hierarchy of classes of infinite graphs with decidable MSO theories has been constructed [23]. It is not clear whether these results could be exploited in our setting.

Interestingly there is an intimate relation between exclusion of *grids* as *minors* and bounded *tree-width* in the theory of finite undirected graphs [24] (see also [25, Chapter 12]). This connection between tree-width and exclusion of grids may also yield some new ideas for settling our conjecture.

Another related theme is a conjecture of Seese [11] which states if the MSO theory of a class of graphs is decidable, then this class of graphs is MSO interpretable: it can be suitably encoded by MSO formulas interpreted over some class of trees (see e.g. [26] for a precise formulation). Partial results concerning Seese's conjecture can be found in [27,26,28,15]. It is again not clear if this line of work can be exploited for settling our conjecture.

Turning now to our main motivation for studying the MSO theory of net systems, namely distributed controller synthesis, the positive result we have for connectedly communicating processes [17] is illustrative. It shows that in the MSO theory of the associated event structure one can easily assert: (i) the existence of a controlled set of events that result from local strategies; (ii) the fact that each local strategy depends only on the sequence of actions it has performed and the actions of other agents it comes to know of through synchronizations; and (iii) the controlled behavior satisfies a global specification. Consequently for systems that have a decidable MSO theory one can at once solve such distributed controller synthesis problems for a variety of local strategies and a variety of linear time or branching time global specifications. It is worth pointing out here that this approach will not go through if quantification over sets is restricted to conflict-free subsets of events as required in [16]. Finally, an additional benefit of settling our conjecture would be that the natural counterparts to branching time logics such as *CTL*, *CTL** in the concurrency setting will have effective procedures for checking satisfiability in the presence of grid-freeness.

References

[1] M. Rabin, Decidability of second order theories and automata on infinite trees, Trans. AMS 141 (1969) 1–35.
[2] W. Thomas, Automata on Infinite Objects, Handbook Theoret. Comput. Sci., vol. B, Elsevier, 1990.
[3] G. Winskel, M. Nielsen, Models for Concurrency, Handbook of Logic in Comput. Sci., vol. 3, Oxford Univ. Press, 1994.
[4] J.R. Büchi, On a decision method in restricted second-order arithmetic, in: 1960 Int. Congr. for Logic, Methodology and Philosophy of Science, Stanford, Stanford Univ. Press, 1962, pp. 1–11.
[5] V. Diekert, G. Rozenberg (Eds.), The Book of Traces, World Scientific, Singapore, 1995.
[6] M. Nielsen, G. Plotkin, G. Winskel, Petri nets, event structures and domains, Part 1, Theoret. Comput. Sci. 13 (1981) 85–108.
[7] D. Harel, Recurring dominoes: Making the highly undecidable highly understandable, Annals Discrete Math. 24 (1985) 51–72.
[8] G. Kahn, G. Plotkin, Concrete domains, Theoret. Comput. Sci. 121 (1993) 187–277.
[9] G. Winskel, Event structures, in: Petri Nets: Applications and Relationships to Other Models of Concurrency, in: Lecture Notes in Comput. Sci., vol. 255, Springer, 1987, pp. 325–392.
[10] J. Desel, J. Esparza, Free Choice Petri Nets, Cambridge Tracts Theoret. Comput. Sci., vol. 40, Cambridge Univ. Press, 2005.
[11] D. Seese, The structure of the models of decidable monadic theories of graphs, Ann. Pure Appl. Logic 53 (1991) 169–195.
[12] A.L. Semenov, Decidability of monadic theories, in: MFCS 1984, in: Lecture Notes in Comput. Sci., vol. 176, Springer, 1984, pp. 162–175.
[13] B. Courcelle, I. Walukiewicz, Monadic second-order logic, graph coverings and unfoldings of transition systems, Ann. Pure Appl. Logic 92 (1998) 95–208.
[14] I. Walukiewicz, Monadic second-order logic on tree-like structures, Theoret. Comput. Sci. 275 (2002) 311–346.
[15] B. Courcelle, J. Engelfriet, Graph Structure and Monadic Second-Order Logic: A Language-Theoretic Approach, Encyclopedia Math. Appl., vol. 138, Cambridge Univ. Press, 2012.
[16] P. Madhusudan, Model-checking trace event structures, in: LICS 2003, IEEE Press, 2003, pp. 371–380.
[17] P. Madhusudan, P. Thiagarajan, S. Yang, The MSO theory of connectedly communicating processes, in: FSTTCS 2005, in: Lecture Notes in Comput. Sci., vol. 3821, Springer, 2006, pp. 201–212.
[18] P.S. Thiagarajan, Elementary net systems, in: Petri Nets: Central Models and Their Properties, Advances in Petri Nets 1986, Part I, in: Lecture Notes in Comput. Sci., vol. 254, Springer, 1987, pp. 26–59.
[19] W. Reisig, Petri Nets: An Introduction, Monogr. Theoret. Comput. Sci. EATCS Ser., Springer, 1985.
[20] R. Parikh, Decidability and undecidability in distributed transition systems, in: A Perspective in Theoretical Comput. Sci. — Commemorative Volume for Gift Siromoney, World Scientific, Singapore, 1988, pp. 199–209.
[21] Y. Gurevich, Monadic second-order theories, in: Model-Theoretical Logics, Perspective in Mathematical Logic, Springer, 1985, pp. 479–506, Ch. XIII.
[22] W. Thomas, Constructing infinite graphs with a decidable MSO-theory, in: MFCS 2003, in: Lecture Notes in Comput. Sci., vol. 2747, Springer, 2003, pp. 113–124.

[23] D. Caucal, On infinite transition graphs having a decidable monadic theory, Theoret. Comput. Sci. 290 (2003) 79–115.
[24] N. Robertson, P. Seymour, Graph minors V: Excluding a planar graph, J. Combin. Theory Ser. B 41 (1986) 92–114.
[25] R. Diestel, Graph Theory, third edition, Springer, 2006.
[26] B. Courcelle, S. Oum, Vertex-minors, monadic second-order logic and a conjecture by Seese, J. Combin. Theory Ser. B 97 (2007) 91–126.
[27] B. Courcelle, The monadic second-order logic of graphs XV: On a conjecture by D. Seese, J. Appl. Log. 4 (2006) 79–114.
[28] P. Hliněný, D. Seese, Trees, grids and MSO decidability: From graphs to matroids, Theoret. Comput. Sci. 351 (2006) 372–393.