Science of Computer Programming 151 (2018) 48–60

Contents lists available at ScienceDirect

Science of Computer Programming

www.elsevier.com/locate/scico

Reversible computation vs. reversibility in Petri nets

Kamila Barylska [a,*], Maciej Koutny [b], Łukasz Mikulski [a], Marcin Piątkowski [a]

[a] Faculty of Mathematics and Computer Science, Nicolaus Copernicus University, Chopina 12/18, 87-100 Toruń, Poland
[b] School of Computing Science, Newcastle University, Newcastle upon Tyne, NE1 7RU, United Kingdom

ARTICLE INFO

Article history:
Received 7 December 2016
Received in revised form 16 October 2017
Accepted 17 October 2017
Available online 2 November 2017

Keywords:
Petri net
Reversibility
Reversible computation
Decidability

ABSTRACT

Petri nets are a general formal model of concurrent systems which supports both action-based and state-based modelling and reasoning. One of important behavioural properties investigated in the context of Petri nets has been reversibility, understood as the possibility of returning to the initial marking from any reachable net marking. Thus reversibility in Petri nets is a global property. Reversible computation, on the other hand, is typically a local mechanism by which a system can undo some of the executed actions.
This paper is concerned with the modelling of reversible computation within Petri nets. A key idea behind the proposed construction is to add 'reverses' of selected transitions, and the paper discusses its different implementations.
Adding reverses can severely impact on the behaviour of a Petri net. Therefore it is important, in particular, to be able to determine whether the modified net has a similar set of states as the initial one. We first prove that the problem of establishing whether the initial and modified nets have the same reachable markings is undecidable, even in the restricted case considered in this paper. We then show that the problem of checking whether the reachability sets of the two nets cover the same markings is decidable.

1. Introduction

Petri nets are a general formal model of concurrent systems which supports both action-based and state-based modelling and reasoning. One of important behavioural properties investigated in the context of Petri nets has been reversibility. At first, reversible Petri nets were defined as nets in which every transition a can be directly undone (i.e. in the net there exists a transition b with the exactly opposite effect) [10,23]. Later on, the notion of reversibility in Petri nets started to be understood as the possibility of returning to the initial marking (a global state) from any reachable marking. However, it is not required that any specific transitions (global states) are used to bring the net back to the initial marking.

Reversibility in the later formulation has been investigated for years, for example, in the context of enforcing controllability in discrete event systems [29,31,40]. Intuitively, it is a global property which is related to the existence of home states [7,27], i.e., those markings which can be reached from all forward reachable markings. The former notion, however, is still occasionally in use [11,26], but – to avoid ambiguity – recently renamed to *symmetric* Petri nets [19].

Unlike currently considered Petri net reversibility, reversible computation typically refers to a local mechanism by which a system can undo (the effect of) some of the already executed actions. Such an approach has been applied, in particular, to

* Corresponding author.
E-mail addresses: kamila.barylska@mat.umk.pl (K. Barylska), maciej.koutny@newcastle.ac.uk (M. Koutny), lukasz.mikulski@mat.umk.pl (Ł. Mikulski), marcin.piatkowski@mat.umk.pl (M. Piątkowski).

https://doi.org/10.1016/j.scico.2017.10.008
0167-6423/

various kinds of process calculi and event structures (see, e.g., [4,9,12,13,30,32–34]). A category theory based rendering of reversible computation with an application to Petri nets has been proposed in [14].

1.1. Previous work

A good deal of decision problems related to reversibility as well as home states and home spaces has been investigated over the past decades. These problems were usually considered within the domain of potentially infinite-state Place/Transition-net (p/t-nets) and their subclasses, as most problems become trivial for finite-state models. Typically, these problems are of one of two kinds.

In the case of the first kind of problems, one wants to establish whether a given marking (or a set of markings) satisfies a desirable property. For instance, the fundamental home state problem is concerned with establishing whether a given marking of a given p/t-net is a home state. The problem was shown in [1] to be decidable, as well as its restricted version consisting in deciding whether the initial marking of a p/t-net is a home state. Another example problem is that of establishing whether a linear set of markings is a home space of a given p/t-net, and [15] demonstrated that such a problem is decidable. Problems of the second kind put the emphasis on the existence of a marking (or set of markings) satisfying a desirable property. For example, the fundamental home state existence problem, shown to be decidable in [6], is to establish whether there exists a home state for a given p/t-net.

Although there are several positive decidability results related to reversibility, in general, the complexity of potential solutions appears to be high or difficult to establish. For example, the problem of the reversibility property is decidable but its complexity is still unknown [27], and [6] demonstrated that the problem of home state existence is at least as hard as the reachability problem [25]. These, rather pessimistic results, meant that the quest for effective algorithms, and indeed decidable problems, has for many years been carried out within special subclasses of p/t-nets. Such subclasses are often defined by imposing restrictions on the structure of a net, or by assuming boundedness, with the resulting submodels of p/t-nets being still relevant for a wide range of practical applications.

For example, it was shown in [8] that all live and bounded free-choice nets have home states, and the free-choice assumption cannot be changed to asymmetric choice. The home space problem is polynomial for live and bounded free-choice Petri nets [5,16], and they also were shown to have home states [39]. Other, progressively less restricted, net classes were considered in [5,37,27,35,38].

1.2. Our contribution

This paper is concerned with the modelling of reversible computation in Petri nets. A key idea is to add *reversed* versions of selected net transitions, each such reversed transition being obtained by simply changing the directions of adjacent arcs. The resulting reversible computations implement in a direct way what can be seen as the *undoing* of an executed action, and the simple form of such an undoing is possible thanks to the local nature of marking changes effected by net transitions.

Adding reversed transitions can greatly impact on the behaviour of the system. As the examples provided in the paper show, even adding a single reverse may change the set of reachable markings. It is therefore crucial to be able to determine whether the modified net has similar set of states as the initial one. In this paper we present two key results. First, we prove that the problem of establishing whether the initial net and that resulting from adding reverse transitions have the same reachable markings is undecidable even in the case of adding a single reverse. This is a strong result indicating that unless reversing of transitions is applied to restricted classes of Petri nets, such as bounded nets, controlling reversibility (so that the state space of a system does not grow) is too hard a task. We then turn to more relaxed requirement on the state space of the 'reversed' net by stipulating that what one requires is that the two nets 'cover' the same sets of markings. We then demonstrate that the problem of checking whether the reachability sets of the two nets are equivalent w.r.t. coverability is decidable.

It should be noted that focusing on coverability still has a significant application potential. For example, if all the markings covered by the initial Petri net are safe on a given subset of places, then all the reachable markings of the 'reversed' net are guaranteed to be safe on this subset of places as well, provided that the nets cover the same sets of markings.

We also focus on higher level of abstraction. Reversible computation typically refers to a local mechanism by which a system can undo some of the already executed actions. Moreover, one might consider different variants of reversing. The first approach we consider insists on reversing only transition actually executed. The other two allow to reverse an action a as long as in has not been prohibited by another one. Here we distinguish two cases: a 'stopper' transition (after the execution of the stopper, reversing of a is suspended until the next execution of a), and a 'blocker' transition (after the execution of the blocker, reversing of a is disallowed forever).

1.3. Organisation of this paper

The paper is organised as follows. In Section 2, we recall some basic definitions concerning Petri nets and their behavioural properties. Section 3 contains examples motivating our work and facilitating the understanding of the proposed approach. In Section 4, we provide the proof of undecidability of the problem of establishing whether two given nets have the same sets of reachable markings. In the Section 5, we prove that the problem of checking whether the reachability