Artificial Intelligence 245 (2017) 56–85

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

The virtues of idleness: A decidable fragment of resource agent logic

Natasha Alechina [a,*], Nils Bulling [b,c], Brian Logan [a], Hoang Nga Nguyen [d]

[a] School of Computer Science, University of Nottingham, UK
[b] Delft University of Technology, The Netherlands
[c] TU Clausthal, Germany
[d] Centre for Mobility and Transport, Coventry University, UK

ARTICLE INFO

Article history:
Received 26 April 2016
Received in revised form 13 December 2016
Accepted 30 December 2016
Available online 4 January 2017

Keywords:
Strategy logics
Resource constraints
Model checking

ABSTRACT

Alternating Time Temporal Logic (ATL) is widely used for the verification of multi-agent systems. We consider Resource Agent Logic (RAL), which extends ATL to allow the verification of properties of systems where agents act under resource constraints. The model checking problem for RAL with unbounded production and consumption of resources is known to be undecidable. We review existing (un)decidability results for fragments of RAL, tighten some existing undecidability results, and identify several aspects which affect decidability of model checking. One of these aspects is the availability of a 'do nothing', or *idle* action, which does not produce or consume resources. Analysis of undecidability results allows us to identify a significant new fragment of RAL for which model checking is decidable.

1. Introduction

Many problems in AI and multi-agent systems research are most naturally formulated in terms of the abilities of a group or coalition of agents. For example, a group of agents may be able to cooperate to achieve an outcome which cannot be achieved by any agent in the group acting individually. In many cases, whether the outcome can be achieved depends critically on the resources available to the agents. Money is an obvious example, but there are many kinds of resources that may be produced or consumed by the actions of agents. For example, whether a team of agents can cooperate to extinguish a fire may depend on the amount of fuel and water they have available. Several logics for reasoning about coalitional ability under resource bounds have been proposed in the literature [1–7]. These *resource logics* allow us to express properties such as: 'a coalition of agents A has a strategy (a choice of actions) requiring no more than b resources, such that whatever the actions by the agents outside the coalition, any evolution of the system generated by the strategy satisfies some temporal property'. Using model checking techniques we can then verify that a given coalition has a strategy requiring less than b resources to enforce an outcome, whatever the other agents in the system (or the environment) do. The ability to verify such properties can be useful when designing or developing a resource-constrained multi-agent system.

Unfortunately, the model checking problem for many resource logics where actions can produce resources is undecidable [2,5]. Recently, however, it was shown that some resource logics where actions can produce resources have a decidable

* Corresponding author.
E-mail addresses: nza@cs.nott.ac.uk (N. Alechina), Bulling@in.tu-clausthal.de (N. Bulling), bsl@cs.nott.ac.uk (B. Logan), hoang.nguyen@coventry.ac.uk (H.N. Nguyen).

http://dx.doi.org/10.1016/j.artint.2016.12.005

model checking problem [6–8].[1] In this paper, we investigate the reasons for the decidability or undecidability of the model checking problem for resource logics. Different syntactic and semantic choices give different variants of resource logics. Some of these choices are known to affect the decidability of the model checking problem. In particular, the decidability result in [6] was proven in the presence of two major restrictions, called, in the terminology of [2], *resource flat* and *proponent restricted*. The former assumes that agents are always re-equipped with fresh resources when they reconsider their strategies; the latter assumes that only the proponents act under resource bounds (i.e., agents outside the coalition are not resource bounded). In addition to these restrictions, another choice in the semantics is relevant for the decidability result in [6]. This choice, which is also related to the finitary and infinitary semantics of [2], stipulates that, in every model, agents always have a choice of doing nothing (executing an idle action) that produces and consumes no resources. Having an *idle* action makes model checking easier: intuitively, its availability ensures that in order to determine whether a coalition can enforce a ϕ-state after finitely many steps and within a given resource bound, we only need to find a finite strategy to enforce ϕ under the given resource bound, and after ϕ is achieved, the agents can always choose the idle action forever, which does not increase the 'cost' of the strategy. The presence of an idle action in the logic also guarantees some attractive formal properties. For example, as stated in [3], it ensures coalition monotonicity: if a coalition A can ensure a property under resource bound b, then any larger coalition can also ensure this property under the same resource bound (intuitively, the extra agents can always perform idle).

In this paper, we investigate the effects of various semantic choices, such as the availability of an idle action, on the decidability of the model checking problem for resource logics. First we show that both the resource-flat and the proponent-restricted fragments of resource agent logic remain undecidable in the presence of idle actions. We then identify and motivate a significant, non-resource-flat fragment that has a decidable model checking property in the presence of idle actions, and is not decidable otherwise. It follows that idle actions can make a difference for the decidability of model checking with respect to the semantics we consider.

The new fragment, which we call pprRAL, allows us to express statements about the existence of nested strategies for a coalition of agents given some *initial* allocation of resources. Unlike the resource-flat fragment considered in [6], where for each new strategy agents are re-equipped with a fresh set of resources, pprRAL allows us to express properties such as 'given their initial battery charge, rescue robots A can safely get to a position from which they can perform rescue while in visual contact with the base'. There are two nested strategies implicit in this property: first, the robots should be able to reach some position (not necessarily maintaining visual contact with the base), and second, from this position, the agents should be able to perform rescue while in visual contact with the base. The first strategy (getting into position) will require certain resources (in this case battery charge), and the amount of resources required will depend on the environment. Then, with *whatever resources are left*, the agents need a strategy to perform the rescue. In this example, the model checking problem essentially corresponds to finding two nested conditional resource-constrained plans, see e.g., [10]. The plans are nested because it is impossible to decouple the second plan (for rescue) from the results of the first plan (getting into position), since we do not know the resource availability for the initial state of the second plan; the resource availability in that state is determined by resource consumption of the first plan. Compared to conditional planning with resources, resource logics provide an easy way to talk not just about reachability, but also about invariants and nested goals/strategies achieved by (potentially different) coalitions.

This paper extends results presented in [7] in several respects, including: a more general definition of a decidable fragment, more elaborated intuitions regarding the (un)decidability results, detailed proofs of all theorems, and tighter undecidability results (in terms of the number of agents and resource types required for undecidability). The remainder of the paper is organised as follows. In Section 2 we briefly survey related work. In Section 3 we introduce resource agent logic, its models and the semantics. In Section 4, we review known decidability results for resource agent logic, and investigate the reasons for (un)decidability. We present new undecidability results for systems with a single resource type, and, based on these results, we motivate and introduce a new non-resource flat fragment of RAL, pprRAL. In Section 5, we present our second main technical result: a decidability result for pprRAL. We conclude in Section 6.

2. Related work

Early work on resource logics considered only the consumption of resources (i.e., no action produces resources), and initial results on the complexity of model checking were encouraging. One of the first logics capable of expressing resource requirements of agents was a version of Coalition Logic (CL)[2] called Resource-Bounded Coalition Logic (RBCL), where actions only consume (and do not produce) resources. It was introduced in [1] with the primary motivation of modelling systems of resource-bounded reasoners; however the framework is sufficiently general to model any type of action. The model checking problem for RBCL was shown to be decidable in time polynomial in the size of the transition system and of the property, and exponential in the number of resource types in [12]. A resource-bounded version of ATL, RB-ATL, where again actions only consume (and do not produce) resources was introduced in [3]. The model checking problem for this logic is also decidable in time polynomial in the size of the transition system and of the property, and exponential in the number of

[1] A preliminary version of [8] is available as a technical report [9].

[2] CL is a fragment of ATL with only the next time $\langle\langle A\rangle\rangle\mathbf{X}$ modality, introduced in [11].