Artificial Intelligence 234 (2016) 152–189

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Parameterised verification for multi-agent systems

Panagiotis Kouvaros *, Alessio Lomuscio *

Department of Computing, Imperial College London, UK

ARTICLE INFO

Article history:
Received 28 October 2014
Received in revised form 5 January 2016
Accepted 13 January 2016
Available online 18 January 2016

Keywords:
Multi-agent systems
Validation
Parameterised verification
Cutoffs

ABSTRACT

We study the problem of verifying role-based multi-agent systems, where the number of components cannot be determined at design time. We give a semantics that captures parameterised, generic multi-agent systems and identify three notable classes that represent different ways in which the agents may interact among themselves and with the environment. While the verification problem is undecidable in general we put forward cutoff procedures for the classes identified. The methodology is based on the existence of a notion of simulation between the templates for the agents and the template for the environment in the system. We show that the cutoff identification procedures as well as the general algorithms that we propose are sound; for one class we show the decidability of the verification problem and present a complete cutoff procedure. We report experimental results obtained on MCMAS-P, a novel model checker implementing the parameterised model checking methodologies here devised.

1. Introduction

With the development and deployment of autonomous agents and multi-agent systems (MAS) in diverse applications such as robot-based search-and-rescue [1], web-services [2], personal negotiation assistants [3], a growing need has emerged to develop powerful and versatile methodologies for the validation and verification of MAS. Model checking [4] is a leading logic-based technique for the verification of systems that has emerged in the past twenty years. Model checking enables us to check whether a model M_S representing a system S, satisfies a formula ϕ_P encoding a specification P.

While plain reactive systems [5] are typically specified by means of reachability or purely temporal statements, autonomous agents are typically specified by means of high level properties inspired from AI. As a consequence, in the case of MAS the specification ϕ_P is typically given in agent-based logics, such as epistemic logic [6], BDI [7], Desires-Goal-Intention [8], and ATL [9]. Over the past ten years a number of techniques have been put forward for the efficient model checking of MAS against agent-based specifications including binary decision diagrams [10,11], abstraction [12], partial order reduction [13], bounded model checking [14], parallel model checking [15], thereby making it possible to verify systems with large state spaces. Yet, since the number of states is exponential in the number of agents in the system, systems of many agents typically remain intractable.

A further difficulty consists on the fact that some agent-based protocols, such as auctions, do not specify how many agents may be present at runtime. By model checking we may be able to verify a system for a *given number of agents*. But this does not enable us to draw any conclusion as to whether the specification would still hold should more agents be present. Intuitively, additional agents may possibly interfere with the system in unpredicted ways resulting in the specification to be

* Corresponding authors.
E-mail addresses: p.kouvaros@imperial.ac.uk (P. Kouvaros), a.lomuscio@imperial.ac.uk (A. Lomuscio).

http://dx.doi.org/10.1016/j.artint.2016.01.008

violated. Yet, our practical experience (e.g., networking and security protocols) tells us that some, albeit not all, protocols are correct irrespective of the number of components. Any technique that enables us to verify specifications independently of the number of agents present would clearly be beneficial in validating a wide range of MAS.

Cutoffs have been studied in the formal analysis of systems to try to address this, often in the context of networking protocols [16,17]. A cutoff for a specification is the number of components that need to be analysed to be able to draw general conclusions that hold irrespective of the number of components in a system. Since the problem in its generality is undecidable [18], sound but incomplete methods have been put forward [17,19,20] that impose restrictions on the systems and the properties to be studied. However, as we discuss below the current literature does not address the needs of MAS, or AI systems in general, as they are tailored to temporal specifications only and they often rely on specific semantics that abstract from the particular way in which agents may interact.

The aim of this paper is to present a technique for the automatic verification of MAS populated by arbitrarily many agents adhering to different roles. In particular we isolate three classes of MAS for which we show that cutoffs can be given when certain sufficient conditions are met. We illustrate the semantic classes correspond to different ways in which the agents may interact among themselves and with the environment. In addition to exploring the theoretical side of the problem we also present an implementation based on ideas here presented and discuss the experimental results obtained.

1.1. Parameterised model checking

The traditional model checking problem [4] concerns establishing whether a specification ϕ_P representing a property P holds on a finite model M_S built from a finite number of components implementing the system S, or $M_S \models \phi_P$. In the traditional approach the behaviours of all the components are specified beforehand; the model M_S resulting from their synchronisation is then constructed and the property ϕ_P is then checked.

While the traditional model checking problem establishes whether a particular system satisfies a given specification, the parameterised model checking problem (PMCP) is concerned with establishing whether any system composed of any number of agents following a certain behavioural template satisfies a given specification. Clearly any attempt to reduce the parameterised model checking problem to the standard model checking problem would entail checking an infinite number of models, i.e., all possible systems built from any number of agents. Given the number of agents is not bounded it would also imply checking models of unbounded size.

In traditional computer science the PMCP can potentially be used to verify specific networking protocols and a wide range of distributed algorithms. In MAS and AI in general, techniques for the PMCP could in principle be used to establish properties of a wide and diverse range of systems ranging from robotic swarms to e-commerce applications where the number of agents is not known at design time.

In the general setting the PMCP is undecidable [18]. However, given its importance, it is of interest to develop sound but incomplete techniques to solve it. The PMCP is typically formulated in a finitary, abstract way by giving a template for the agents in the system, a template for the environment, and the formula to be verified. By providing the parameter n specifying the actual number of agents in the system, we can then construct a concrete system upon which the standard model checking problem can be solved. A way to limit the generality of the problem is to restrict the systems considered. For example, we may consider a specific topology, e.g., rings, when analysing network protocols for an unbounded number of hosts. In this paper we follow a different approach. We do not impose many constraints in terms of how the agents may behave, but we are constrain their interaction.

1.2. Related work

In the past 10 years several methods have been put forward for verifying MAS by means of symbolic model checking. Most techniques support epistemic specifications [13,14,21–24]; others target deontic specifications [25,26], or specifications expressing strategic abilities [27,28]. The resulting performance differs depending on a number of assumptions; symbolic checkers such as MCK [10], MCMAS [11] and VerICS [29] are all capable of handling state-spaces of the region of 10^{15} and beyond.

While these techniques have received considerable attention, they all suffer from a key limitation in that they only deal with closed MAS where the number of components is known at design time. This makes it impossible to verify MAS where the number of agents is not known at design time.

Verification of systems with an arbitrarily large number of components has been investigated, however, in the context of reactive systems where the problem has been shown to be undecidable in general [18]. The techniques put forward typically assume a number of restrictions either on the systems or in the specifications considered so that either soundness or decidability can be retained. The approaches can be classified into *abstraction techniques, network invariant techniques, regular model checking*, and *cutoff techniques*.

Abstraction techniques [30–37] rely on the analysis of a single finite state *abstract system* encoding all possible concrete systems. Typically these methods require manual guidance for obtaining the abstract mapping. Further, they are often incomplete: if a certain specification is falsified in the abstract model, then it does not necessarily follow that there is a concrete system falsifying the specification. Among these techniques we identify *counter abstraction* and *environment abstraction*.