On a Logic for Coalitional Games with Priced-Resource Agents

Dario Della Monica[1], Margherita Napoli[1], and Mimmo Parente[3]

Dept. of Computer Science, Università di Salerno, Fisciano (Italy),
{ddellamonica|napoli|parente}@unisa.it

Alternating-time Temporal Logic (ATL) [4] and *Coalition Logic* (CL) [7, 8] are well-established logical formalisms particularly suitable to model games between dynamic coalitions of agents (like e.g. the system and the environment). In [6], Goranko has studied the relationship between the (expressive power of the) two formalisms. In particular, he has shown that CL can be embedded into ATL. Both these logics have successfully been applied to the automated verification of multi-agent systems, which is a significant topic in the recent literature of Artificial Intelligence [1]. Anyway, none of them takes into account the boundedness of the resources available for the agents. Approaches towards verification of multi-agent systems under resource constraints can be found in [2, 3, 5]. In [2], Alechina et al. introduce the logic RBCL, whose language extends the one of CL with explicit representation of resource bounds. In [3], the same authors propose an analogous extension for ATL, called RB-ATL, and give a PTIME model checking procedure mostly based on the one for ATL. In [5], Bulling and Farwer introduce the logics RAL and RAL^*. The former represents a generalization of Alechina et al.'s RB-ATL, the latter is ATL^* extended with resource bounds. The authors study several syntactic and semantic variants of RAL and RAL^* with respect to the (un)decidability of the model checking problem. In particular, while previous approaches only conceive actions consuming resources, they introduce the notion of actions producing resources. It turned out that such a new notion makes the model checking problem undecidable.

In this paper, we propose an epistemic discussion about the formalization of multi-agent systems, in which agents can cooperate to perform complex tasks and are subject to a limited availability of priced resources, which are intrinsic features of any real-world system. We highlight a certain number of problems and considerations, based on different interpretations of shortage of resources, leading to different scenarios. Our discussion hinges on existing approaches in the literature (see e.g. [2, 3, 5]) and represents an attempt to do a further step towards the formalization of such complex systems.

Formulas of the formalisms proposed in [2, 3, 5] allow one to assign an endowment of resources to the agents, by means of the so-called *team operators* (borrowed from ATL), and to state that a team of agents can perform a task. Due to the nesting of the team operators in a formula (which reflects the fact that coalitions may be dynamic, in the sense that may change in a game), during the execution of the task, the agents can be provided with a new endowment of resources to perform subtasks. This is somehow unrealistic, as it does not take into account issues related to the procurement of resources. In particular, a very significant present-day issue is that resources are available on the market (or in nature) in limited amount, and the cost for achieving them depends on such an availability (e.g., cloud computing).

First improvement. Thus, our first proposal is to introduce the notion of *price* of resources. Unlike the existing approaches, agents are equipped with an amount of money instead of an endowment of resources. They can use money for getting resources. Formulas of our logic state that a team of agents is able to perform a given task provided with a given amount of money. We also introduce a notion of *global availability* of resources on the market, the intended meaning being that, whenever an agent acquires resources from the market, the global availability is decreased, whenever it produces resources, the global availability is increased. The price of resources can be any function of the several components into play. In our approach, prices of resources depend on their global availability, the acting agent, and the physical location.

Second improvement. Another aspect that has not been fully analyzed in the literature is the problem of action producing resources. On one hand, in [2, 3], actions can only consume resources; on the other hand, in [5], the authors state that whenever actions can produce resources the model checking problem is undecidable. In this paper, we show how to realistically constrain the way in which actions can produce resources, still preserving the decidability of the model checking problem. The idea is that it is possible, at a given time, for an action to produce a resource in a quantity that is not greater of the amount that has already been consumed so far. This implies that, even if actions can produce resources, the global availability of the market will never be greater than the initial global availability, that is crucial for the model checking algorithm and realistic, as well. Indeed, such a notion makes sense as, in practical terms, it allows one to model significant real-world scenarios, such as, acquiring memory by a program, leasing a car during a travel, and, in general, any scenario in which an agent is releasing resources previously acquired.

Team and task

So far, we have talked about teams (or coalitions) of agents performing a task. But we have not clarified yet the two notions of team and task. First of all, a task is a goal that has to be reached and, for what concerns us, is represented by a logical formula that has to be satisfied. A team of agents is a subset of agents that act collectively in order to perform a task. To this end, they select a strategy that univocally determines their behavior in each possible configuration of the system. Nevertheless, the behavior of the remaining agents, that we collectively denote as the *opponent*, is undetermined.

The aim of the team is to guarantee that the task is performed independently of the opponent's behavior, that amounts to say, the task must be guaranteed for each possible strategy of the opponent.

The formalism that naturally fits our intention is the logic ATL, that allows one to fix a strategy for the agents of a team and to force a property, representing the task, to be true over all the possible executions (or outcomes) of the system. Obviously, its syntax and semantics will be extended in order to deal with resource constraints.

The special resource 'time'

One can be interested in answering questions of the kind "is it possible for the team A of agents to complete the task in x time-unit?". It is clear that the resource 'time' cannot be acquired. It is in a certain sense out of the control of the agents, as it is only possible to give time constraints a task must possibly be executed within, while it is not possible to administer it. Thus, resource 'time' will be treated in a special way with respect to other resources.

Model checking

The model checking problem consists in verifying whether a formula φ is satisfied in a location q of a game structure G, with an initial resource availability $\vec{m} \in \mathcal{M}$ ($\vec{m}$ is a vector storing the initial availability on the market of each resource).

The algorithm for model checking our logic, denoted *Priced* RB-ATL (PRB-ATL), is mostly based on the one proposed in [4] and used in [3] for model checking, respectively, ATL and its resource-bounded extension RB-ATL. Roughly speaking, it works by computing, for each sub-formula ψ of the formula φ to be model checked, the set of states in which ψ holds. The main difficulties when dealing with bounds on resources are the following. First, the set of sub-formulae must be replaced by an extended set of formulae (see [3]), that includes, for each sub-formula of the form $\langle\langle A^{\$}\rangle\rangle\psi$, all the formulae $\langle\langle A^{\$'}\rangle\rangle\psi$ for each $\$' < \$$. Second, the state does not correspond anymore to the vertices of the game structure, denoted by Q, but to configurations, that is, pairs $\langle q, \vec{m}\rangle \in Q \times \mathcal{M}$. Third, during the analysis of the computations over the game structure, the algorithm must take into account the resource availability on the market in order to guarantee that in each instant of the computation all the resources are still available, as well as to be able to compute the current prices of resources, that depend also on their availability. Finally, it must be ensured that, even if actions can produce resources, availability of each resource may not be higher than the initial availability.

Let M be the greater component appearing in the initial resource availability vector $\vec{m}$. The following theorem represents our main contribution. Full details of both the formalization and the algorithm will appear in a forthcoming paper.

Theorem 1. *The model checking problem for* PRB-ATL *is decidable in time* $O(M^r \times |\varphi|^{r+1} \times |G|)$.

Future scenarios

A further line of research in which we intend to investigate is when, given a formula in our logic, the coalitions are unknown, that is they are not specified and we may ask whether, for each nested sub-formula, there exists a team and a money endowment such that the formula is satisfied. More precisely, given a formula Ψ where $\langle\langle X_i^{\$_i}\rangle\rangle$ are the team operators occurring in it, we want to compute minimal coalitions X_i and amounts of endowment $\$_i$ such that Ψ is satisfied. Let us notice that if the minimality condition is not requested, then the problem can be trivially solved.

Another feature we are investigating is when the agents each have a price. In this scenario, in which agents are themselves resources to be acquired to perform the task, it makes sense to consider the problem of deciding which team is able to perform the task at the minimum cost.

References

[1] Thomas Ågotnes, Wiebe van der Hoek, and Michael Wooldridge. On the logic of coalitional games. In *AAMAS*, pages 153–160, 2006.

[2] Natasha Alechina, Brian Logan, Nguyen Hoang Nga, and Abdur Rakib. A logic for coalitions with bounded resources. In *Proc. of the 21st International Joint Conference on Artificial Intelligence*, IJCAI '09, pages 659–664, 2009.

[3] Natasha Alechina, Brian Logan, Nguyen Hoang Nga, and Abdur Rakib. Resource-bounded alternating-time temporal logic. In *Proc. of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*, AAMAS '10, pages 481–488, 2010.

[4] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of ACM*, 49:672–713, September 2002.

[5] Nils Bulling and Berndt Farwer. On the (un-)decidability of model checking resource-bounded agents. In *Proc. of the 19th European Conference on Artificial Intelligence*, ECAI '10, pages 567–572, 2010.

[6] Valentin Goranko. Coalition games and alternating temporal logics. In *Proc. of the 8th Conference on Theoretical Aspects of Rationality and Knowledge*, TARK '01, pages 259–272, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.

[7] Marc Pauly. A logical framework for coalitional effectivity in dynamic procedures. *Bulletin of Economic Research*, 53(4):305–324, 2001.

[8] Marc Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.