Resource-bounded alternating-time temporal logic

Natasha Alechina
University of Nottingham
Nottingham, UK
nza@cs.nott.ac.uk

Brian Logan
University of Nottingham
Nottingham, UK
bsl@cs.nott.ac.uk

Nguyen Hoang Nga
University of Nottingham
Nottingham, UK
hnn@cs.nott.ac.uk

Abdur Rakib
University of Nottingham
Nottingham, UK
rza@cs.nott.ac.uk

ABSTRACT

Many problems in AI and multi-agent systems research are most naturally formulated in terms of the abilities of a coalition of agents. There exist several excellent logical tools for reasoning about coalitional ability. However, coalitional ability can be affected by the availability of resources, and there is no straightforward way of reasoning about resource requirements in logics such as Coalition Logic (CL) and Alternating-time Temporal Logic (ATL). In this paper, we propose a logic for reasoning about coalitional ability under resource constraints. We extend ATL with costs of actions and hence of strategies. We give a complete and sound axiomatisation of the resulting logic Resource-Bounded ATL (RB-ATL) and an efficient model-checking algorithm for it.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: [multiagent systems]

General Terms

Theory, Verification

Keywords

Logics for agency, Verification of MAS

1. INTRODUCTION

In many situations a group of agents can cooperate to achieve an outcome which cannot be achieved by any agent in the group acting individually. For example, in the prisoners dilemma, a single prisoner cannot ensure the optimal outcome, while a coalition of two prisoners can. Similarly, it may be possible for a set of cooperating agents to solve a difficult computational problem by distributing it, while a single agent may not have sufficient memory or processor power to solve it. In the latter case, there is an interaction between the amount of resources available to the agents (or the amount of resources which they are willing to contribute), and their ability to achieve their goal.

In this paper we propose a logic, RB-ATL, for reasoning about coalitional ability under resource constraints. RB-ATL allows us to express and verify properties such as

Cite as: Resource-bounded alternating-time temporal logic, Natasha Alechina, Brian Logan, Hoang Nga Nguyen and Abdur Rakib, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 481-488

(1) 'a coalition of agents A has a strategy for achieving a property ϕ provided they have resources b, but they cannot enforce ϕ under a tighter resource bound b_1',

(2) 'A has a strategy to maintain the property ϕ with resources b',

(3) 'A can maintain ϕ until ψ becomes true provided A has resources b'.

There exists work on introducing resource bounds in coalition logic [3] and temporal logic [5]. We believe that our contribution presents a significant advance on this work. Specifically, the logic RBCL defined in [3] can express properties of the form (1) but not of the form (2) and (3) (since it generalises Coalition Logic and does not have the full set of temporal operators). Also, [3] does not analyse model-checking complexity of RBCL. In [5], a logic RTL* is introduced, which is CTL* extended with quantifiers representing the cost of paths. Using CTL* as a starting point means that only single-agent systems can be analysed in RTL*. The setting of [5] is also different from the one presented in this paper, in that the actions not only consume but also produce resources. The model-checking problem for RCTL* is quite complex; only partial solutions (e.g. for RCTL rather than RCTL* where actions only consume resources) are presented in [5]. No axiomatisation of RCTL* is given.

The rest of this paper is organised as follows. In section 2, we present the syntax and semantics of RB-ATL. In section 3 we provide a sound and complete axiomatisation of RB-ATL. In section 4, we give a model-checking algorithm for RB-ATL. Finally, we survey related work and conclude. The Appendix contains some of the proofs.

2. SYNTAX AND SEMANTICS OF RB-ATL

Consider a system of agents which can perform actions to change the state (we assume concurrent execution of actions by all agents). We denote the set of agents by N. In order to reason about resources, we assume that actions have costs. Let R be a set of resources (such as money, energy, or anything else which may be required by an agent for performing an action). We assume that a cost of an action, for each of the resources, is a natural number. The set of resource bounds $\mathbb{B}$ over R is defined as $\mathbb{B} = \mathbb{N}^r$, where $r = |R|$.

2.1 Syntax of RB-ATL

The syntax of RB-ATL is defined as follows, where A is a nonempty subset of N and $b \in \mathbb{B}$.

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \psi \mid \langle\langle A^b \rangle\rangle \bigcirc \varphi \mid \langle\langle A^b \rangle\rangle \Box \varphi \mid \langle\langle A^b \rangle\rangle \varphi \mathcal{U} \psi$$

Here, $\langle\langle A^b \rangle\rangle \bigcirc \varphi$ means that a coalition A can ensure that the next state satisfies φ under resource bound b. $\langle\langle A^b \rangle\rangle \Box \varphi$ means that A has a strategy to make sure that φ is always true, and the cost of this strategy is at most b. Similarly, $\langle\langle A^b \rangle\rangle \varphi \mathcal{U} \psi$ means that A has a strategy to enforce ψ while maintaining the truth of φ, and the cost of this strategy is at most b.

The corresponding operators for the empty coalition are defined as follows:

$$\langle\langle \emptyset^b \rangle\rangle \bigcirc \varphi =_{df} \neg \langle\langle N^b \rangle\rangle \bigcirc \neg \varphi$$

$$\langle\langle \emptyset^b \rangle\rangle \Box \varphi =_{df} \neg \langle\langle N^b \rangle\rangle \top \mathcal{U} \neg \varphi$$

$$\langle\langle \emptyset^b \rangle\rangle \varphi \mathcal{U} \psi =_{df} \neg (\langle\langle N^b \rangle\rangle \neg \psi \mathcal{U} \neg (\varphi \vee \psi) \vee \langle\langle N^b \rangle\rangle \Box \neg \psi)$$

We explain why we make the $\langle\langle \emptyset^b \rangle\rangle$ modalities definable rather than primitive symbols after the truth definition for RB-ATL.

2.2 Semantics of RB-ATL

To interpret this language, we extend the definition of concurrent game structures [4] with resource requirements for executing actions. For consistency with [4], in what follows we refer to agents as 'players' and actions as 'moves'.

DEFINITION 1. *A Resource-bounded Concurrent Game Structure (RB-CGS) is a tuple $S = (n, r, Q, \Pi, \pi, d, c, \delta)$ where:*

- $n \geq 1$ *is the number of players (agents), we denote the set of players $\{1, \ldots, n\}$ by N*
- r *is the number of resources*
- Q *is a non-empty set of states*
- Π *is a finite set of propositional variables*
- $\pi : Q \to \wp(\Pi)$ *is a function which assigns each state in Q a subset of propositional variables*
- $d : Q \times N \to \mathbb{N}$ *is a function which indicates the number of available moves (actions) for each player $a \in N$ at a state $q \in Q$ such that $d(q, a) \geq 1$. At each state $q \in Q$, we denote the set of joint moves available for all players in N by $D(q)$. That is*
$$D(q) = \{1, \ldots, d(q, 1)\} \times \ldots \times \{1, \ldots, d(q, n)\}$$
- $c : Q \times N \times \mathbb{N} \to \mathbb{B}$ *is a partial function which indicates the minimal amount of resources required by each move available to each agent at a specific state.*
- $\delta : Q \times \mathbb{N}^{|N|} \to Q$ *is a partial function where $\delta(q, m)$ is the next state from q if the players execute the move $m \in D(q)$.*

We denote by $\bar{0}$ the smallest resource bound $(0, \ldots, 0)$. We assume that each agent in each state has an available action with $\bar{0}$ cost (intuitively, it has the option of doing nothing).

Given a RB-CGS S, we denote the set of infinite sequences of states (computations) by Q^ω. For a computation $\lambda = q_0 q_1 \ldots \in Q^\omega$, we use the notation $\lambda[i] = q_i$ and $\lambda[i, j] = q_i \ldots q_j$. We denote the set of finite non-empty sequences of states by Q^+.

DEFINITION 2. *Given a RB-CGS S and a state $q \in Q$, a move (or a joint action) for a coalition $A \subseteq N$ is a tuple $\sigma_A = (\sigma_a)_{a \in A}$ such that $1 \leq \sigma_a \leq d(q, a)$.*

By $D_A(q)$ we denote the set of all moves for A at state q. Given a move $m \in D(q)$, we denote by m_A the actions executed by A, $m_A = (m_a)_{a \in A}$. We define the set of all possible outcomes of a move $\sigma_A \in D_A(q)$ at state q as follows:

$$out(q, \sigma_A) = \{q' \in Q \mid \exists m \in D(q) : m_A = \sigma_A \wedge q' = \delta(q, m)\}$$

The cost of a move $\sigma_A \in D_A(q)$ is defined as $cost(q, \sigma_A) = \Sigma_{a \in A} c(q, a, \sigma_a)$. (Note that we use c for the cost of single actions and $cost$ for the cost of joint actions).

DEFINITION 3. *Given a RB-CGS S, a strategy for a subset of players $A \subseteq N$ is a mapping F_A which associates each sequence $\lambda q \in Q^+$ to a move in $D_A(q)$.*

A computation $\lambda \in Q^\omega$ is consistent with F_A iff for all $i \geq 0$, $\lambda[i+1] \in out(\lambda[i], F_A(\lambda[0, i]))$. We denote by $out(q, F_A)$ the set of all such sequences λ starting from q, i.e. where $\lambda[0] = q$.

To compare costs and resource bounds, we use the usual pointwise vector comparison, that is, $(b_1, \ldots, b_r) \leq (d_1, \ldots, d_r)$ iff $b_i \leq d_i$ for $i \in \{1, \ldots, r\}$. We also use pointwise vector addition: $(b_1, \ldots, b_r) + (d_1, \ldots, d_r) = (b_1 + d_1, \ldots, b_r + d_r)$.

DEFINITION 4. *Given a bound b, a computation $\lambda \in out(q, F_A)$ is b-consistent with F_A iff, for every $i \geq 0$,*

$$\sum_{j=0}^{i} cost(\lambda[i], F_A(\lambda[0, i])) \leq b.$$

We denote by $out(q_0, F_A, b)$ the set of all b-consistent computations. A strategy F_A is a b-strategy iff $out(q, F_A) = out(q, F_A, b)$ for any $q \in Q$.

In other words, all executions of a b-strategy cost at most b resources. Note that this means that each computation of such a strategy starts with a finite prefix where some non-$\bar{0}$ cost actions are executed, and continues with an infinite sequence of $\bar{0}$-cost actions.

2.3 Truth definition for RB-ATL

Given a RB-CGS $S = (n, r, Q, \Pi, \pi, d, c, \delta)$, the truth definition for RB-ATL is given inductively as follows:

- $S, q \models p$ iff $p \in \pi(q)$
- $S, q \models \neg \varphi$ iff $S, q \not\models \varphi$
- $S, q \models \varphi \vee \psi$ iff $S, q \models \varphi$ or $S, q \models \psi$
- $S, q \models \langle\langle A^b \rangle\rangle \bigcirc \varphi$ iff there exists a b-strategy F_A such that for all $\lambda \in out(q, F_A)$, $S, \lambda[1] \models \varphi$ iff there is a move $\sigma_A \in D_A(q)$ such that for all $q' \in out(\sigma_A)$, $S, q' \models \varphi$
- $S, q \models \langle\langle A^b \rangle\rangle \Box \varphi$ iff there exists a b-strategy F_A for any $\lambda \in out(q, F_A)$, $S, \lambda[i] \models \varphi$ for all $i \geq 0$
- $S, q \models \langle\langle A^b \rangle\rangle \varphi \mathcal{U} \psi$ iff there exists a b-strategy F_A such that for all $\lambda \in out(q, F_A)$, there exists $i \geq 0$ such that $S, \lambda[i] \models \psi$ and $S, \lambda[j] \models \psi$ for all $j \in \{0, \ldots, i-1\}$

Note that given the definition of $\langle\langle \emptyset^b \rangle\rangle \bigcirc \varphi$ as $\neg \langle\langle N^b \rangle\rangle \bigcirc \neg \varphi$, its truth definition is

- $S, q \models \langle\langle \emptyset^b \rangle\rangle \bigcirc \varphi$ iff for every b-strategy F_N for all $\lambda \in out(q, F_A)$, $S, \lambda[1] \models \varphi$

Intuitively, $\langle\langle \emptyset^b \rangle\rangle \bigcirc \varphi$ means that φ is inevitably true in the next state *provided that the grand coalition N executes a move which costs no more than b*. Observe that if we used the same truth definition for empty and non-empty coalition modalities, then all $\langle\langle \emptyset^b \rangle\rangle \bigcirc \varphi$ formulas (for any b) would be equivalent and mean simply that φ is inevitable in the next state. A similar collapse would happen for $\langle\langle \emptyset^b \rangle\rangle \Box$ and $\langle\langle \emptyset^b \rangle\rangle\, \mathcal{U}$ temporal operators as well. For this reason, we have chosen to make empty coalition modalities a special (definable) case.

3. AXIOMATISATION

In this section we present the axiomatic system for RB-ATL. To make the formulas below more readable, we define the following abbreviations:

$$\langle\langle A^b \rangle\rangle \bigcirc \Box\varphi = \bigvee_{b_1+b_2=b, b_1\neq\bar{0}} \langle\langle A^{b_1} \rangle\rangle \bigcirc \langle\langle A^{b_2} \rangle\rangle \Box\varphi$$
$$\langle\langle A^b \rangle\rangle \bigcirc \varphi\mathcal{U}\psi = \bigvee_{b_1+b_2=b, b_1\neq\bar{0}} \langle\langle A^{b_1} \rangle\rangle \bigcirc \langle\langle A^{b_2} \rangle\rangle \varphi\mathcal{U}\psi$$

The axiomatic system consists of the following axioms and rules of inference, where A, A_1 and A_2 are non-empty subsets of N, and b, $b_1, b_2 \in \mathbb{B}$.

Axioms.

(PL) Tautologies of Propositional Logic

($\bot$) $\neg\langle\langle A^b \rangle\rangle \bigcirc \bot$

($\top$) $\langle\langle A^b \rangle\rangle \bigcirc \top$

(B) $\langle\langle A^{b_1} \rangle\rangle \bigcirc \varphi \rightarrow \langle\langle A^{b_2} \rangle\rangle \bigcirc \varphi$
where $b_1 \leq b_2$

(S) $\langle\langle A_1^{b_1} \rangle\rangle \bigcirc \varphi \wedge \langle\langle A_2^{b_2} \rangle\rangle \bigcirc \psi \rightarrow \langle\langle (A_1 \cup A_2)^{b_1+b_2} \rangle\rangle \bigcirc (\varphi \wedge \psi)$
where $A_1 \cap A_2 = \emptyset$

($S_\emptyset$) $\langle\langle \emptyset^{b_1} \rangle\rangle \bigcirc \varphi \wedge \langle\langle \emptyset^{b_2} \rangle\rangle \bigcirc \psi \rightarrow \langle\langle \emptyset^{b_1} \rangle\rangle \bigcirc (\varphi \wedge \psi)$
where $b_1 \leq b_2$

(S_N) $\langle\langle N^{b_1} \rangle\rangle \bigcirc \varphi \wedge \langle\langle \emptyset^{b_2} \rangle\rangle \bigcirc \psi \rightarrow \langle\langle N^{b_1} \rangle\rangle \bigcirc (\varphi \wedge \psi)$
where $b_1 \leq b_2$

($FP_\Box$) $\langle\langle A^b \rangle\rangle \Box\varphi \leftrightarrow \varphi \wedge (\langle\langle A^b \rangle\rangle \bigcirc \Box\varphi \vee \langle\langle A^{\bar{0}} \rangle\rangle \bigcirc (\langle\langle A^b \rangle\rangle \Box\varphi)$

($FP_\mathcal{U}$) $\langle\langle A^b \rangle\rangle \varphi\mathcal{U}\psi \leftrightarrow \psi$
$\vee(\varphi \wedge (\langle\langle A^b \rangle\rangle \bigcirc \varphi\mathcal{U}\psi \vee \langle\langle A^{\bar{0}} \rangle\rangle \bigcirc \langle\langle A^b \rangle\rangle \varphi\mathcal{U}\psi))$

($N_\bigcirc$) $\langle\langle \emptyset^b \rangle\rangle \bigcirc \varphi \leftrightarrow \neg\langle\langle N^b \rangle\rangle \bigcirc (\neg\varphi)$

($N_\Box$) $\langle\langle \emptyset^b \rangle\rangle \Box\varphi \leftrightarrow \varphi \wedge \neg\langle\langle N^b \rangle\rangle \top\mathcal{U}\neg\varphi$

($N_\mathcal{U}$) $\langle\langle \emptyset^b \rangle\rangle \varphi\mathcal{U}\psi \leftrightarrow \neg(\langle\langle N^b \rangle\rangle \neg\psi\mathcal{U}\neg(\varphi \vee \psi) \vee \langle\langle N^b \rangle\rangle \Box\neg\psi)$

Inference rules.

(MP) $\dfrac{\varphi, \varphi \rightarrow \psi}{\psi}$

($\langle\langle A^b \rangle\rangle\bigcirc$-Monotonicity) $\dfrac{\varphi \rightarrow \psi}{\langle\langle A^b \rangle\rangle \bigcirc \varphi \rightarrow \langle\langle A^b \rangle\rangle \bigcirc \psi}$

($\langle\langle \emptyset^b \rangle\rangle\Box$-Necessitation) $\dfrac{\varphi}{\langle\langle \emptyset^b \rangle\rangle \Box\varphi}$

($\langle\langle A^b \rangle\rangle\Box$-Induction)

$$\frac{\theta \rightarrow (\varphi \wedge (\langle\langle A^b \rangle\rangle \bigcirc \Box\varphi \vee \langle\langle A^{\bar{0}} \rangle\rangle \bigcirc \theta))}{\theta \rightarrow \langle\langle A^b \rangle\rangle \Box\varphi}$$

($\langle\langle A^b \rangle\rangle\mathcal{U}$-Induction)

$$\frac{(\psi \vee (\varphi \wedge (\langle\langle A^b \rangle\rangle \bigcirc \varphi\mathcal{U}\psi \vee \langle\langle A^{\bar{0}} \rangle\rangle \bigcirc \theta))) \rightarrow \theta}{\langle\langle A^b \rangle\rangle \varphi\mathcal{U}\psi \rightarrow \theta}$$

Before proving soundness and completeness, we give an intuitive explanation of the axioms and compare them with the axiomatic system for ATL given in [7].

First of all, observe that with the resource bounds removed, the axioms ($\bot$), ($\top$), (**S**), ($N_\bigcirc$) and the inference rules ($\langle\langle A^b \rangle\rangle\bigcirc$-Monotonicity) and ($\langle\langle \emptyset^b \rangle\rangle\Box$-Necessitation) are identical to their ATL counterparts. The axiom (B) says that if A can enforce φ under a resource bound b_1, then it can also enforce φ if it has more than b_1 resources. Note that for the empty coalition, the relationship between the bounds is reversed: an outcome which is inevitable when the grand coalition can only use b resources, is also inevitable when the grand coalition can use fewer resources:

$$\langle\langle \emptyset^b \rangle\rangle \bigcirc \varphi \rightarrow \langle\langle \emptyset^d \rangle\rangle \bigcirc \varphi \quad \text{for } d \leq b$$

With resource bounds removed, this axiom obviously becomes trivial. The axiom ($FP_\Box$) is similar to its ATL counterpart. However, unlike in ATL, there are two ways to 'unwind' $\langle\langle A^b \rangle\rangle \Box\varphi$ in RB-ATL: one way is to make a move which costs a non-trivial amount of resources b_1, and then maintain φ with $b - b_1$ resources; the second way is to make a trivial $\bar{0}$-cost move, and then maintain φ with b resources. Similarly for ($FP_\mathcal{U}$). Finally, the rules ($\langle\langle A^b \rangle\rangle\Box$-Induction) and ($\langle\langle A^b \rangle\rangle\mathcal{U}$-Induction) correspond to the ATL axioms ($GFP_\Box$) and ($LFP_\mathcal{U}$); the first one says that $\Box$ corresponds to the greatest fixed point and the second that $\mathcal{U}$ corresponds to the least fixed point. This will be made more precise after we give fixed point characterisations of the temporal operators.

3.1 Fixed point characterisations of temporal operators

Consider an operation $[\langle\langle A^b \rangle\rangle\bigcirc]$ which given a set of states X, returns the set of states from where A can enforce an outcome to be in X under resource bound b (this is the same as $Pre(A, X, b)$ defined in Section 4, which is in turn similar to Pre from [4]):

DEFINITION 5. $[\langle\langle A^b \rangle\rangle\bigcirc] : \wp(Q) \rightarrow \wp(Q)$ *is defined as follows: given a set $X \subseteq Q$, $[\langle\langle A^b \rangle\rangle\bigcirc](X)$ is the set*

$$\{q \mid \exists\sigma \in D_A(q) : cost(\sigma) \leq b \wedge out(q, \sigma) \subseteq X\}$$

Let us define $\|\varphi\| = \{q \in Q \mid S, q \models \varphi\}$. It is straightforward that:

$$\|\langle\langle A^b \rangle\rangle \bigcirc \varphi\| = [\langle\langle A^b \rangle\rangle\bigcirc](\|\varphi\|)$$

Recall that if f is a monotone operator $2^Q \longrightarrow 2^Q$ (that is, $X \subseteq Y$ implies $f(X) \subseteq f(Y)$), then X is a fixed point of f if $F(X) = X$. By the Knaster-Tarski theorem, f has the least and the greatest fixed point. The least fixed point of f is denoted by $\mu X.f(X)$ and the greatest fixed point by $\nu X.f(X)$. We are going to show that the meanings of $\Box$ and $\mathcal{U}$ correspond to the greatest and the least fixed points of certain operations on sets of states.

LEMMA 1. *For all $q \in Q$, the following fixed point characterisations hold:*

1. *$q \in \|\langle\langle A^b \rangle\rangle \Box\varphi\|$ iff $q \in \nu X.\|\varphi\| \cap (\|\langle\langle A^b \rangle\rangle \bigcirc \Box\varphi\| \cup [\langle\langle A^{\bar{0}} \rangle\rangle\bigcirc](X))$ iff there is a b-strategy F_A for A such that for all $\lambda \in out(q, F_A)$, $\lambda[i] \in \|\varphi\|$ for all $i \geq 0$*
2. *$q \in \|\langle\langle A^b \rangle\rangle \varphi\,\mathcal{U}\psi\|$ iff $q \in \mu X.\|\psi\| \cup (\|\varphi\| \cap (\|\langle\langle A^b \rangle\rangle \bigcirc \varphi\,\mathcal{U}\psi\| \cup [\langle\langle A^{\bar{0}} \rangle\rangle\bigcirc](X)))$ iff there is a b-strategy F_A for A such that for all $\lambda \in out(q, F_A)$, there exists $i \geq 0$ such that $\lambda[i] \in \|\psi\|$ and $\lambda[j] \in \|\varphi\|$ for all $j \leq i$*

A proof of Lemma 1(1) is given in the Appendix. The second part of the lemma can be proved in a similar way.

3.2 Soundness of RB-ATL

We prove that the axioms of RB-ATL are valid.

($\bot$) is valid because there is no b-strategy F_A such that for all $\lambda \in out(q, F_A)$, $\lambda[1]$ makes $\bot$ true.

($\top$) is valid because A has a $\overline{0}$-strategy F_A such that for all $\lambda \in out(q, F_A)$, $\lambda[1]$ makes $\top$ true.

(B) is valid because if there is a b_1-strategy F_A such that for all $\lambda \in out(q, F_A)$, $\lambda[1]$ makes φ true, then the same F_A is also a b_2-strategy which has the same property.

(S) is valid because if there exists a strategy F_{A_1} to enforce φ and a strategy F_{A_2} to enforce ψ, then there exists a joint strategy $F_{A_1 \cup A_2}$ (with the same moves for A_1 and A_2 as F_{A_1} and F_{A_2}, respectively) to enforce both φ and ψ.

(S_N) is valid because if there exists a b_1-strategy F_N to enforce φ, and for all strategies of N which cost at most b_2 ψ is inevitable, then $\varphi \wedge \psi$ can be enforced in by F_N. ($S_\emptyset$) is obvious; ($N_\bigcirc$), ($N_\Box$) and ($N_\mathcal{U}$) correspond to definitions.

($\mathbf{FP}_\Box$) is valid by Lemma 1(1) and ($\mathbf{FP}_\mathcal{U}$) by by Lemma 1(2).

($\langle\langle A^b \rangle\rangle \bigcirc$-Monotonicity), ($\langle\langle A^b \rangle\rangle \Box$-Monotonicity) and ($\langle\langle A^b \rangle\rangle\, \mathcal{U}$-Monotonicity) clearly preserve validity, since if $\|\varphi\| \subseteq \|\psi\|$ and an outcome in $\|\varphi\|$ can be enforced, then an outcome in $\|\psi\|$ can also be enforced by the same strategy.

($\langle\langle \emptyset^b \rangle\rangle \Box$-Necessitation) is valid since if φ is logically true, then it is inevitable in perpetuity.

($\langle\langle A^b \rangle\rangle \Box$-Induction) and ($\langle\langle A^b \rangle\rangle \mathcal{U}$-Induction) preserve validity by Lemma 1.

3.3 Completeness of RB-ATL

The proof of completeness is based on [7]. We construct a satisfying model for a formula φ_0 which is consistent with the axiomatic system for RB-ATL.

In the proof, we assume when convenient that all formulas are in negation normal form of RB-ATL; the precise definitions of normal form RB-ATL and semantics for $\neg\langle\langle A^b \rangle\rangle \bigcirc \varphi$, $\neg\langle\langle A^b \rangle\rangle \Box \varphi$ and $\neg\langle\langle A^b \rangle\rangle \varphi\, \mathcal{U} \psi$ are similar to that for ATL and omitted due to the lack of space. We also treat $\langle\langle \emptyset^b \rangle\rangle \bigcirc$ as a primitive symbol in normal form RB-ATL, with the truth definition given earlier in the paper.

The model is constructed in a way very similar to the construction in [7]. It is assembled from finite trees where nodes are labelled by sets of formulas. First we define the set of formulas used in the labelling.

Definition 6. *The closure $cl(\varphi_0)$ is the smallest set of formulas satisfying the following closure conditions:*

- *all sub-formulas of φ_0 including φ_0 itself are in $cl(\varphi_0)$*
- *if $\langle\langle A^b \rangle\rangle \Box \varphi$ is in $cl(\varphi_0)$, then so are $\langle\langle A^{b_1} \rangle\rangle \bigcirc \langle\langle A^{b_2} \rangle\rangle \Box \varphi$ for all $b_1 + b_2 = b$ and also $\langle\langle A^{\overline{0}} \rangle\rangle \bigcirc \langle\langle A^b \rangle\rangle \Box \varphi$*
- *if $\langle\langle A^b \rangle\rangle \varphi \mathcal{U} \psi$ is in $cl(\varphi_0)$, then so are $\langle\langle A^{b_1} \rangle\rangle \bigcirc \langle\langle A^{b_2} \rangle\rangle \varphi \mathcal{U} \psi$ for all $b_1 + b_2 = b$ and also $\langle\langle A^{\overline{0}} \rangle\rangle \bigcirc \langle\langle A^b \rangle\rangle \varphi \mathcal{U} \psi$*
- *if φ is in $cl(\varphi_0)$, then so is $\sim\varphi$ (normal form negation of φ)*
- *$cl(\varphi_0)$ is also closed under finite positive boolean operators ($\vee$ and $\wedge$) up to tautology equivalence.*

Note that $cl(\varphi_0)$ is finite. Let Γ be a set of maximal consistent subsets of $cl(\varphi_0)$. We define trees (T, V) over Γ in exactly the same way as in [7] (T is the set of nodes and $V : T \longrightarrow \Gamma$). Intuitively, nodes in a tree are identified with finite words corresponding to the sequence of joint actions by the grand coalition which leads to that node. The root is the empty word ϵ and each node c corresponds to a finite computation the last state of which is c. A formula is in $V(c)$ intuitively means that the formula is true in c. As in [7], the construction proceeds in three stages. The first stage is producing locally consistent trees, namely trees where the labelling satisfies conditions on successor nodes which makes it possible to prove a truth lemma for the next step modalities. The second stage is proving the existence of trees which realise eventualities (essentially, make the labelling consistent with the truth conditions for the $\Box$ and $\mathcal{U}$ modalities). Finally, the finite trees realising eventualities are combined into one infinite tree model.

Definition 7. *A tree (T, V) is locally consistent if and only if for any interior node $t \in T$:*

1. *if $\langle\langle A^b \rangle\rangle \bigcirc \varphi$ in $V(t)$, then there is a move σ_A such that $cost(\sigma_A) \leq b$ and for any $c \in out(\sigma_A)$ we have $\varphi \in V(c)$*
2. *if $\neg\langle\langle A^b \rangle\rangle \bigcirc \varphi$ in $V(t)$, then for any move σ_A with $cost(\sigma_A) \leq b$, there exists $c \in out(\sigma_A)$ where $\neg\varphi \in V(c)$*

Two following lemmas are used as a crucial step in the local consistency proof.

Lemma 2. *Let $\Phi = \{\langle\langle A_1^{b_1} \rangle\rangle \bigcirc \varphi_1, \ldots, \langle\langle A_k^{b_k} \rangle\rangle \bigcirc \varphi_k, \neg\langle\langle A^b \rangle\rangle \bigcirc \varphi\}$ be a consistent set of formulas in which:*

- *all A_i are both non-empty and pairwise disjoint*
- *$\bigcup_i A_i \subseteq A$*
- *$\Sigma_i b_i \leq b$*

We have $\Psi = \{\varphi_1, \ldots, \varphi_k, \neg\varphi\}$ is also consistent.

Lemma 3. *Let $\Phi = \{\langle\langle A_1^{b_1} \rangle\rangle \bigcirc \varphi_1, \ldots, \langle\langle A_k^{b_k} \rangle\rangle \bigcirc \varphi_k, \langle\langle \emptyset^{e_1} \rangle\rangle \bigcirc \chi_1, \ldots, \langle\langle \emptyset^{e_m} \rangle\rangle \bigcirc \chi_m\}$ be a consistent set of formulas in which:*

- *all A_i are both non-empty and pairwise disjoint*
- *$\Sigma_i b_i \leq e_j$ for all j*

We have $\Psi = \{\varphi_1, \ldots, \varphi_k, \chi_1, \ldots, \chi_m\}$ is also consistent.

The proofs of the lemmas use axioms (S), (S_N), $(S_\emptyset)$ and (B).

Lemma 4. *Let Φ be a finite consistent set of formulas. Let $\Phi_\bigcirc$ be the subset of Φ which contains all formulas of the form $\langle\langle A^b \rangle\rangle \bigcirc \varphi$ or their negations. If $|\Phi_\bigcirc| < k$, then there is a tree (T, V) which is of height one, branching factor k^n, has $V(\epsilon) = \Phi$ and is locally consistent.*

The proof is given in the Appendix.

The next stage of the proof is to consider what conditions on tree labelling we need in order to be able to prove the truth lemma for other temporal modalities. The definition of what it means to 'realise' formulas of the form $\langle\langle A^b \rangle\rangle \varphi \mathcal{U} \psi$, $\neg\langle\langle A^b \rangle\rangle \Box \varphi$, $\langle\langle A^b \rangle\rangle \Box \varphi$, $\neg\langle\langle A^b \rangle\rangle \varphi \mathcal{U} \psi$ is similar to the one in [7] (essentially the truth conditions for the formulas with 'satisfied' replaced by 'in the labelling of'). In what follows, $\Psi_\bigcirc$ is the set of formulas of the form $\langle\langle A^b \rangle\rangle \bigcirc \varphi$ or $\neg\langle\langle A^b \rangle\rangle \bigcirc \varphi$ from $cl(\varphi_0)$.

LEMMA 5. *For each formula $\langle\langle A^b\rangle\rangle\varphi\mathcal{U}\psi$ and $x \in \Gamma$, there is finite tree (T, V) over Γ such that:*

- *(T, V) is of fixed branching degree k^n where $k = |\Psi_{\bigcirc}| + 1$*
- *(T, V) is locally consistent*
- *$V(\epsilon) = x$*
- *if $\langle\langle A^b\rangle\rangle\varphi\mathcal{U}\psi \in x$ then (T, V) realises $\langle\langle A^b\rangle\rangle\varphi\mathcal{U}\psi$ from ϵ*

LEMMA 6. *For each formula $\neg\langle\langle A^b\rangle\rangle\Box\varphi$ and $x \in \Gamma$, there is finite tree (T, V) over Γ such that:*

- *(T, V) is of fixed branching degree k^n where $k = |\Psi_{\bigcirc}| + 1$*
- *(T, V) is locally consistent*
- *$V(\epsilon) = x$*
- *if $\neg\langle\langle A^b\rangle\rangle\Box\varphi \in x$ then (T, V) realises $\neg\langle\langle A^b\rangle\rangle\Box\varphi$ from ϵ*

The proofs of these lemmas are similar to the corresponding proofs in [7], but also use induction on the bound b.

Now we have almost all the ingredients for constructing the model for φ_0. For each consistent set x in Γ and an eventuality φ of $cl(\varphi_0)$, we have a finite tree $(T_{x,\varphi}, V_{x,\varphi})$ with the root having label x which realises φ. Let the eventualities in $cl(\varphi_0)$ be listed as $\varphi_0^e, \ldots, \varphi_m^e$. In the following, we have the definition of the final tree.

DEFINITION 8. *The final tree $(T_{\varphi_0}, V_{\varphi_0})$ is constructed inductively as follows.*

- *Initially, select an arbitrary $x \in \Gamma$ such that $\varphi_0 \in x$. As that formula is consistent, such a set exists. Let $(T_{x,\varphi_0^e}, V_{x,\varphi_0^e})$ be the initial tree.*
- *Given the tree constructed so far and the last used eventuality φ_i^e. Then, for every leaf labelled by $y \in \Gamma$ of the currently constructed tree, we replace it with the tree $(T_{y,\varphi_j^e}, V_{y,\varphi_j^e})$ in which $j = i + 1$ if $i < m$ or $j = 0$ if otherwise.*

Let S_{φ_0} be the model which is based on $(T_{\varphi_0}, V_{\varphi_0})$.

LEMMA 7. *If $\langle\langle A^b\rangle\rangle\varphi\mathcal{U}\psi$ or $\neg\langle\langle A^b\rangle\rangle\Box\varphi$ is in the label of some node t of $(T_{\varphi_0}, V_{\varphi_0})$, it is realised from t.*

LEMMA 8. *If $\neg\langle\langle A^b\rangle\rangle\varphi\mathcal{U}\psi$ or $\langle\langle A^b\rangle\rangle\Box\varphi$ is in the label of some node t of $(T_{\varphi_0}, V_{\varphi_0})$, it is realised from t.*

Finally, we show the following truth lemma.

LEMMA 9. *For every node t of $(T_{\varphi_0}, V_{\varphi_0})$ and every formula $\varphi \in cl(\varphi_0)$, if $\varphi \in V_{\varphi_0}(t)$ then $S_{\varphi_0}, t \models \varphi$.*

PROOF. The proof is done by induction on the structure of φ.

- For the cases of propositions, negative proposition and disjunction, the proofs are trivial.
- Assume $\varphi = \langle\langle A^b\rangle\rangle \bigcirc \psi$, Lemma 4 ensures that there is a move $\sigma \in \Delta_A$ of cost at most b such that for all $c \in out(t, \sigma)$, we have $\psi \in V(tc)$. Then by the induction hypothesis, we have that $S_{\varphi_0}, tc \models \psi$. Then, $S_{\varphi_0}, t \models \langle\langle A^b\rangle\rangle\bigcirc \psi$
- Assume $\varphi = \neg\langle\langle A^b\rangle\rangle \bigcirc \psi$, Lemma 4 ensures that there is a co-move $\sigma \in \Delta_A$ such that for all $c \in out(t, \sigma, b)$, we have $\sim \psi \in V(tc)$. Then by the induction hypothesis, we have that $S_{\varphi_0}, tc \models\sim \psi$. Then, $S_{\varphi_0}, t \models \neg\langle\langle A^b\rangle\rangle \bigcirc \psi$
- For the cases of $\langle\langle A^b\rangle\rangle\varphi\mathcal{U}\psi$, $\neg\langle\langle A^b\rangle\rangle\Box\varphi$, $\neg\langle\langle A^b\rangle\rangle\varphi\mathcal{U}\psi$ and $\langle\langle A^b\rangle\rangle\Box\varphi$, the proofs are trivial due to the two previous lemmas.

□

Finally, we have the following theorem.

THEOREM 1. *The axiom system for RB-ATL is sound and complete.*

4. MODEL-CHECKING RB-ATL

In this section we describe a model-checking algorithm for RB-ATL which runs in time polynomial in the size of the formula and the structure (if we treat the number of resources as a constant). The algorithm is similar to the model-checking algorithm for ATL given in [4]. The main differences from the algorithm for ATL are that we need to take the costs of strategies into account, and, instead of working with a straightforward set of subformulas $Sub(\phi)$ of a given formula ϕ, we work with an extended set of subformulas $Sub^+(\phi)$. $Sub^+(\phi)$ includes $Sub(\phi)$, and in addition:

- if $\langle\langle A^b\rangle\rangle\Box\psi \in Sub(\phi)$, then $\langle\langle A^{b'}\rangle\rangle\Box\psi \in Sub^+(\phi)$ for all $b' < b$
- if $\langle\langle A^b\rangle\rangle\varphi\mathcal{U}\psi \in Sub(\phi)$, then $\langle\langle A^{b'}\rangle\rangle\varphi\mathcal{U}\psi \in Sub^+(\phi)$ for all $b' < b$

THEOREM 2. *Given a structure $S = (n, r, Q, \Pi, \pi, d, c, \delta)$ and a formula ϕ, there is an algorithm which returns the set of states $[\phi]_S$ satisfying ϕ: $[\phi]_S = \{q \mid S, q \models \phi\}$, which runs in time $O(|\phi|^{2r+1} \times |S|)$.*

PROOF. Consider the following model-checking algorithm:

for every ϕ' in $Sub^+(\phi)$:

case $\phi' == p$: $[\phi']_S = \{q \mid p \in \pi(q)\}$

case $\phi' == \neg\psi$: $[\phi']_S = Q \setminus [\psi]_S$

case $\phi' == \psi_1 \wedge \psi_2$: $[\phi']_S = [\psi_1]_S \cap [\psi_2]_S$

case $\phi' == \langle\langle A^b\rangle\rangle \bigcirc \psi$: $[\phi']_S = Pre(A, [\psi]_S, b)$

case $\phi' == \langle\langle A^{\bar{0}}\rangle\rangle\Box\psi$:
  $\rho := [true]; \tau := [\psi]_S$;
  while $\rho \not\subseteq \tau$ **do** $\rho := \tau; \tau := Pre(A, \rho, \bar{0}) \cap [\psi]_S$ **od**;
  $[\phi']_S := \rho$

case $\phi' == \langle\langle A^b\rangle\rangle\Box\psi$ for $b > \bar{0}$:
  $\rho := [false]; \tau := [false]$;
  foreach $b' < b$ **do**
    $\tau := Pre(A, [\langle\langle A^{b'}\rangle\rangle\Box\psi]_S, b - b') \cap [\psi]_S$
    while $\tau \not\subseteq \rho$ **do**
      $\rho := \rho \cup \tau; \tau := Pre(A, \rho, \bar{0}) \cap [\psi]_S$
    od
  od;
  $[\phi']_S := \rho$

case $\phi' == \langle\langle A^{\bar{0}}\rangle\rangle\psi_1\mathcal{U}\psi_2$:
  $\rho := [false]; \tau := [\psi_2]_S$;
  while $\tau \not\subseteq \rho$ **do** $\rho := \rho \cup \tau; \tau := Pre(A, \rho, \bar{0}) \cap [\psi_1]_S$
  od;
  $[\phi']_S := \rho$

case $\phi' == \langle\langle A^b\rangle\rangle \psi_1 \mathcal{U} \psi_2$ for $b > \overline{0}$:
$\rho := [false]; \tau := [false]$;
foreach $b' < b$ **do**
$\tau := Pre(A, [\langle\langle A^{b'}\rangle\rangle \psi_1 \mathcal{U} \psi_2]_S, b - b') \cap [\psi_1]_S$
while $\tau \not\subseteq \rho$ **do**
$\rho := \rho \cup \tau; \tau := Pre(A, \rho, \overline{0}) \cap [\psi_1]_S$
od
od;
$[\phi']_S := \rho$

Pre is a function which given a coalition A, a set $\rho \subseteq Q$ and a bound b returns a set of states q in which A has a move σ_A with cost $cost(q, \sigma_A) \leq b$ such that $out(q, \sigma_A) \subseteq \rho$.

The cases for propositional variables, negation, conjunction and $\langle\langle A^b\rangle\rangle \bigcirc \psi$ are straightforward. The cases where the resource bound is $\overline{0}$ are also similar to [4]. However the cases for $\langle\langle A^b\rangle\rangle \Box \psi$ and $\langle\langle A^b\rangle\rangle \psi_1 \mathcal{U} \psi_2$ for $b > \overline{0}$ have no counterpart in the ATL algorithm, and we explain these in some detail. First, note that the cases for $\langle\langle A^{\overline{0}}\rangle\rangle \Box \psi$ and $\langle\langle A^{\overline{0}}\rangle\rangle \psi_1 \mathcal{U} \psi_2$ are the standard greatest and least fixed point computations respectively, which consider only $\overline{0}$-cost moves. In particular, $[\langle\langle A^{\overline{0}}\rangle\rangle \Box \psi]_S = Pre(A, [\langle\langle A^{\overline{0}}\rangle\rangle \Box \psi]_S, \overline{0}) \cap [\psi]_S$ and $[\langle\langle A^{\overline{0}}\rangle\rangle \psi_1 \mathcal{U} \psi_2]_S = Pre(A, [\langle\langle A^{\overline{0}}\rangle\rangle \psi_1 \mathcal{U} \psi_2]_S, \overline{0}) \cap [\psi_2]_S$. $[\langle\langle A^{\overline{0}}\rangle\rangle \Box \psi]_S$ contain all states where A has a $\overline{0}$-cost strategy to maintain ψ forever. Note that A has a b-cost strategy to maintain ψ forever if and only if it has a b-cost strategy to force the system into one of the $[\langle\langle A^{\overline{0}}\rangle\rangle \Box \psi]_S$ states, while maintaining ψ. In other words, in order to compute $\langle\langle A^b\rangle\rangle \Box \psi$ for $b > \overline{0}$, we need to compute $\langle\langle A^b\rangle\rangle \psi \mathcal{U} \langle\langle A^{\overline{0}}\rangle\rangle \Box \psi$. This explains the similarity between the cases of $\langle\langle A^b\rangle\rangle \Box \psi$ and $\langle\langle A^b\rangle\rangle \psi_1 \mathcal{U} \psi_2$ for $b > \overline{0}$. In the case of $\langle\langle A^b\rangle\rangle \Box \psi$, in the first execution of the **foreach** $b' < b$ loop, we have $b' = \overline{0}$ and $\tau = Pre(A, [\langle\langle A^{\overline{0}}\rangle\rangle \Box \psi]_S, b) \cap [\psi]_S$, which includes $Pre(A, [\langle\langle A^{\overline{0}}\rangle\rangle \Box \psi]_S, \overline{0}) \cap [\psi]_S$, hence it also includes $[\langle\langle A^{\overline{0}}\rangle\rangle \Box \psi]_S$. In the nested **while** loop, ρ accumulates the results and τ adds the ψ-states from where A has a $\overline{0}$ strategy to enforce the outcome to be in ρ. In the outer loop, b' bounds are used in some order consistent with $<$, namely satisfying the condition that if $b_i < b_j$ then b_i is used before b_j.

In the case for $\langle\langle A^b\rangle\rangle \psi_1 \mathcal{U} \psi_2$ where $b > \overline{0}$, after the first iteration of the **foreach** $b' < b$ loop, τ is $[\langle\langle A^{\overline{0}}\rangle\rangle \psi_1 \mathcal{U} \psi_2]_S$ which includes $[\psi_2]_S$. The rest is very similar to the case for $\langle\langle A^b\rangle\rangle \Box \psi$ where $b > \overline{0}$.

Note that $|\{b' \mid b' \leq b\}| = b^r$. If ϕ contains operators with non-$\overline{0}$ bounds, $|Sub^+(\phi)| = |\phi| \times |\phi|^r$. In the $\langle\langle A^b\rangle\rangle \Box \psi$ and $\langle\langle A^b\rangle\rangle \psi_1 \mathcal{U} \psi_2$ cases, the outer loop is executed $|\phi|^r$ times and the inner loop is executed in total at most $|S|$ times. This gives us complexity $O(|\phi| \times |\phi|^r \times |\phi|^r \times |S|)$, or $O(|\phi|^{2r+1} \times |S|)$. If r is treated as a constant factor,[1] we get complexity polynomial in $|\phi|$ and $|S|$. □

5. RELATED WORK

Recent work on Alternating-Time Temporal Logic and Coalition Logic (for example, [9, 6, 10, 4, 7, 1]) has allowed the expression of many interesting properties of coalitions and strategies. However, there is no natural way of expressing resource requirements in these logics.

Resources were considered in Coalitional Resource Games in [12] from the point of view of decision complexity of various properties of games. A logic for describing such games and a model-checking procedure for it were proposed in [2]; however the only modality that logic has is $\langle\langle A^b\rangle\rangle \bigcirc$ (only one step games were considered).

As mentioned in the introduction, resource bounds were added to extended coalition logic in [3]; essentially the logic RBCL described there has modalities $\langle\langle A^b\rangle\rangle \top \mathcal{U} \varphi$. A sound and complete axiomatisation for the logic is given, but no model-checking procedure is described.

The model-checking problem for CTL^* extended with resource path quantifiers ($RCTL^*$) was studied in [5]. In [5], actions not only consume but also produce resources. This potentially allows for modelling and verifying a larger class of (single-agent) systems. However it also significantly increases the complexity of model-checking; in fact no decidability result or model-checking algorithm are given for the general case of $RCTL^*$ or $RCTL$, only for fragments of these logics.

6. CONCLUSIONS

We have provided a complete and sound axiomatisation of RB-ATL, a logic which extends ATL with resource bounds. The resulting logic can express interesting properties of coalitions of agents involving resource limitations. For example, it can express that a coalition can maintain the system in a ϕ-state indefinitely given a finite amount of resources (this essentially means that after a while ϕ can be maintained for free). We have also presented a polynomial-time model-checking algorithm for RB-ATL.

The semantics for RB-ATL presented in this paper, in particular the assumption that actions only consume but never produce resources, is motivated by verifying resource requirements for systems of agents where resources of interest are time, memory, bandwidth etc., which cannot be generated by agents. It would be interesting to also consider semantics where actions can have a negative cost, such as in [5].

7. ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council [grant EP/E031226/1].

8. REFERENCES

[1] T. Ågotnes, W. van der Hoek, and M. Wooldridge. Reasoning about coalitional games. *Artif. Intell.*, 173(1):45–79, 2009.

[2] N. Alechina, B. Logan, N. H. Nga, and A. Rakib. Expressing properties of coalitional ability under resource bounds. In X. He, J. F. Horty, and E. Pacuit, editors, *Logic, Rationality, and Interaction, Second International Workshop, LORI 2009, Proceedings*, volume 5834 of *Lecture Notes in Computer Science*, pages 1–14. Springer, 2009.

[3] N. Alechina, B. Logan, H. N. Nguyen, and A. Rakib. A logic for coalitions with bounded resources. In C. Boutilier, editor, *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 659–664. AAAI Press, 2009.

[4] R. Alur, T. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.

[5] N. Bulling and B. Farwer. RTL and RTL*: Expressing abilities of resource-bounded agents. In J. Dix, M. Fisher, and P. Novák, editors, *Proceedings of the 10th International Workshop on Computational Logic in Multi-Agent Systems*, pages 2–19, 2009.

[6] V. Goranko. Coalition games and alternating temporal logics. In *Proceeding of the Eighth Conference on Theoretical*

[1] We also treat the number of agents as a constant factor; the complexity of ATL model-checking without this assumption was shown to be exponential in the number of agents in [8]. Other assumptions implicit in the formulation of the problem, such as the set of states being given explicitly, are discussed in [11].

Aspects of R ationality and Knowledge (TARK VIII, pages 259–272. Morgan Kaufmann, 2001.

[7] V. Goranko and G. van Drimmelen. Complete axiomatization and decidability of alternating-time temporal logic. *Theor. Comput. Sci.*, 353(1-3):93–117, 2006.

[8] W. Jamroga and J. Dix. Do agents make model checking explode (computationally)? In M. Pechoucek, P. Petta, and L. Z. Varga, editors, *Proc. 4th International Central and Eastern European Conference on Multi-Agent Systems (CEEMAS 2005)*, volume 3690 of *Lecture Notes in Computer Science*, pages 398–407. Springer, 2005.

[9] M. Pauly. *Logic for Social Software*. Ph.D. thesis, ILLC, University of Amsterdam, 2001.

[10] M. Pauly. A modal logic for coalitional power in games. *J. Log. Comput.*, 12(1):149–166, 2002.

[11] W. van der Hoek, A. Lomuscio, and M. Wooldridge. On the complexity of practical ATL model checking. In H. Nakashima, M. P. Wellman, G. Weiss, and P. Stone, editors, *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006)*, pages 201–208. ACM, 2006.

[12] M. Wooldridge and P. E. Dunne. On the computational complexity of coalitional resource games. *Artif. Intell.*, 170(10):835–871, 2006.

APPENDIX

Proof of Lemma 1(1)

For convenience, let us denote $f(X) = \|\varphi\| \cap (\|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\| \cup [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](X))$. First we show that $f(X)$ is monotone. Let $X_1 \subseteq X_2 \subseteq Q$. Let $q \in f(X_1)$, then $q \in \|\varphi\|$ and either $q \in \|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\|$ or $q \in [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](X_1)$. By the definition of $[\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc]()$, it is easy to see that $q \in [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](X_2)$ if $q \in [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](X_1)$, hence $q \in f(X_2)$.

Therefore, $f(X)$ is monotone and there is a greatest fixed point $\nu X.f(X)$. We now show that $Y = \|\langle\langle A^b\rangle\rangle\Box\varphi\|$ is a post-fixed point of $f(X)$, i.e. $f(Y) \subseteq Y$. Let $q \in Y$. By the truth definition, we have that there is a b-strategy F_A such that for any $\lambda \in out(q, F_A)$, $\lambda[i] \in \|\varphi\|$ for all $i \geq 0$. Then, $q = \lambda[0] \in \|\varphi\|$. Let $b' = cost(q, F_A(q))$. For every $q' \in out(q, F_A(q))$ we define a $(b - b')$-strategy F'_A which is the remainder of F_A from q' as follows, $F'_A(\kappa) = F_A(q\kappa)$ for all $\kappa \in Q^\omega$ starting at q'. Then, for all $\kappa \in out(q', F'_A)$, we have that $q\kappa \in out(q, F_A)$. It is straightforward that any computation in $out(q', F'_A)$ costs at most $b - b'$. Then, for all $i \geq 0$, we have that $\kappa[i] \in \|\varphi\|$, hence $q' \in \|\langle\langle A^{b-b'}\rangle\rangle\Box\varphi\|$. Thus, $q \in [\langle\langle A^{b'}\rangle\rangle\bigcirc](\|\langle\langle A^{b-b'}\rangle\rangle\Box\varphi\|)$. If $b' \neq 0$, we have that $q \in \|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\|$, otherwise $q \in [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](\|\langle\langle A^b\rangle\rangle\Box\varphi\|)$. This means that $q \in f(\|\langle\langle A^b\varphi\rangle\rangle\Box\|)$. In order to show that $Y = \|\langle\langle A^b\rangle\rangle\Box\varphi\|$ is, in fact, the greatest fixed point of $f(X)$, we show that for every post-fixed point Z, $Z \subseteq Y$. We show the inclusion by induction on the bound b.

In the base case, $b = 0$, we have $f(X) = \|\varphi\| \cap [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](X)$. Assume $q \in Z$, then $q \in \|\varphi\| \cap [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](Z)$ as Z is a post fixed point of $f(X)$. We now define a 0-strategy F_A which will maintain φ for any consistent computation. The definition will be done by induction on the length of inputs for F_A. Moreover, we only define F_A for inputs which will be used later for the coalition to determine which joint action to perform in order to maintain φ. Let $input^n$ denote the set of such inputs of length n. Initially, $input^1 = \{q\}$. We will define F_A and $input^{i+1}$ inductively on i so that the last element of any member of $input^{i+1}$ is always in Z.

- When $i = 1$, we have that $q \in \|\varphi\|$ and there is a move $\sigma_A \in D_A(q)$ with $cost(q, \sigma_A) = 0$ such that $out(q, \sigma_A) \subseteq Z$. Let $F_A(q) = \sigma_A$ and $input^2 = \{qq' \mid q' \in out(q, F_A(q))\}$. For any such q', we have $q' \in Z \subseteq f(Z)$.
- When $i > 1$, for any $\lambda \in input^i$, we have that $\lambda[i-1] \in Z \subseteq f(Z)$ by the induction hypothesis. We have that $\lambda[i-1] \in \|\varphi\|$ and there is a move $\sigma_A \in D_A(\lambda[i-1])$ with $cost(\lambda[i-1], \sigma_A) = 0$ such that $out(\lambda[i-1], \sigma_A) \subseteq Z$. Let $F_A(\lambda) = \sigma_A$ and $input^{i+1}(\lambda) = \{\lambda q' \mid q' \in out(\lambda[i-1], F_A(\lambda))\}$.

  Finally, we define $input^{i+1} = \bigcup_{\lambda \in input^i} input^{i+1}(\lambda)$ By definition of $input^{i+1}(\lambda)$, it is easy to see that for any $\lambda' \in input^{i+1}$, $\lambda'[i] \in Z \subseteq f(Z)$.

After defining F_A, we have that for any $\lambda \in out(q, F_A)$ and $i \geq 0$, $\lambda[0, i] \in input^{i+1}$, hence $\lambda[i] \in Z \subseteq f(Z)$. Therefore, $\lambda[i] \in \|\varphi\|$. This shows that $q \in Y$.

In the induction step, $b > 0$, we have $f(X) = \|\varphi\| \cap (\|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\| \cup [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](Z))$. Assume $q \in Z$, then $q \in \|\varphi\|$ and either $q \in \|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\|$ or $q \in [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](Z)$. Similarly to the base case, we also define a b-strategy F_A which will maintain φ for any consistent computation. The definition will be also be done by induction on the length of inputs for F_A. Moreover, we only define F_A for inputs which will be used later for the coalition to determine which joint action to perform in order to maintain φ. Let $input^n$ denote the set of such inputs of length n. Initially, $input^1 = \{q\}$. We will define F_A and $input^{i+1}$ inductively on i such that the last element of any member of $input^{i+1}$ is always either in $\|\langle\langle A^{b_2}\rangle\rangle\Box\varphi\|$ if the accumulated cost along that member is no more than b_1 for some $b_1 + b_2 = b$ or in Z if the same cost is zero.

- When $i = 1$, we have that $q \in \|\varphi\|$ and either $q \in \|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\|$ or $q \in [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](Z)$. If $q \in \|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\|$, there is $b_1 + b_2 = b$ such that $q \in [\langle\langle A^{b_1}\rangle\rangle\bigcirc](\|\langle\langle A^{b_2}\rangle\rangle\Box\varphi\|)$. Then, there is a move $\sigma_A \in D_A(q)$ with $cost(q, \sigma_A) \leq b_1$ such that $out(q, \sigma_A) \subseteq \|\langle\langle A^{b_2}\rangle\rangle\Box\varphi\|$. By the induction hypothesis, for any $q' \in out(q, \sigma_A)$, there is another b_2-strategy from q to maintain φ, we define $F_A(qq'\lambda) = F'_A(q'\lambda)$ for all $\lambda \in Q^*$. Let $F_A(q) = \sigma_A$ and $input^2 = \{qq' \mid q' \in out(q, \sigma_A)\}$. It is obvious that all such $q' \in \|\langle\langle A^{b_2}\rangle\rangle\Box\varphi\|$ and the cost along qq' is at most b_1.

  If $q \in [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](Z))$, there is a move $\sigma_A \in D_A(q)$ with $cost(q, \sigma_A) \leq 0$ such that $out(q, \sigma_A) \subseteq Z$. Let $F_A(q) = \sigma_A$ and $input^2 = \{qq' \mid q' \in out(q, \sigma_A)\}$. It is obvious that all such $q' \in Z$ and the cost along qq' is zero.
- When $i > 1$, for any $\lambda \in input^i$, we have that either **(i)** $\lambda[i-1] \in \|\langle\langle A^{b_2}\rangle\rangle\Box\varphi\|$ if $\sum_{j<i-1} cost(\lambda[j], F_A(\lambda[0,j])) \leq b_1$ for some $b_1 + b_2 = b$ or **(ii)** $\lambda[i-1] \in Z \subseteq f(Z)$ if $\sum_{j<i-1} cost(\lambda[j], F_A(\lambda[0,j])) = 0$ by the induction hypothesis.

  (i) If $\lambda[i-1] \in \|\langle\langle A^{b_2}\rangle\rangle\Box\varphi\|$, then F_A has been defined. Let $input^{i+1}(\lambda) = \{\lambda q' \mid q' \in out(\lambda[i-1], F_A(\lambda[0, i-1]))\}$. Let $b' = cost(\lambda[i-1], F_A(\lambda[0, i-1]))$. By the induction hypothesis, as $b_2 < b$, we have that all such $q' \in \|\langle\langle A^{b_2-b'}\rangle\rangle\Box\varphi\|$ and $\sum_{j<i} cost(\lambda[j], F_A(\lambda[0,j])) \leq b_1 + b'$.

  (ii) If $\lambda[i-1] \in Z \subseteq f(Z)$, then $\lambda[i-1] \in \|\varphi\|$ and either **(ii-1)** $\lambda[i-1] \in \|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\|$ or **(ii-2)** $\lambda[i-1] \in [\langle\langle A^{\bar{0}}\rangle\rangle\bigcirc](Z)$.

  (ii-1) If $\lambda[i-1] \in \|\langle\langle A^b\rangle\rangle \bigcirc \Box\varphi\|$, there is $b_1 + b_2 = b$ such that $\lambda[i-1] \in [\langle\langle A^{b_1}\rangle\rangle\bigcirc](\|\langle\langle A^{b_2}\rangle\rangle\Box\varphi\|)$. Then, there is a move $\sigma_A \in D_A(\lambda[i-1])$ with $cost(\lambda[i-1], \sigma_A) \leq b_1$ such that $out(\lambda[i-1], \sigma_A) \subseteq \|\langle\langle A^{b_2}\rangle\rangle \Box\varphi\|$. By the induction hypothesis, for any $q' \in out(\lambda[i], \sigma_A)$, there is another b_2-strategy from q to maintain φ, we define $F_A(\lambda q'\kappa) = F'_A(q'\kappa)$ for all $\kappa \in Q^*$. Let $F_A(\lambda) = \sigma_A$ and $input^{i+1}(\lambda) = \{\lambda q' \mid q' \in out(\lambda[i-1], \sigma_A)\}$. Then, for all such q' we have $q' \in \|\langle\langle A^{b_2}\rangle\rangle \bigcirc \Box\varphi\|$ and $\sum_{j<i} cost(\lambda[j], F_A(\lambda[0,j])) \leq b_1$.

(ii-2) If $\lambda[i-1] \in [\langle\langle A^{\overline{0}}\rangle\rangle\bigcirc](Z))$, there is a move $\sigma_A \in D_A(\lambda[i-1])$ with $cost(\lambda[i-1], \sigma_A) = 0$ such that $out(\lambda[i-1], \sigma_A) \subseteq Z$. Let $F_A(\lambda) = \sigma_A$ and $input^{i+1}(\lambda) = \{\lambda q' \mid q' \in out(\lambda[i-1], \sigma_A)\}$. Then, for all such q', we have that $q' \in Z$ and $\sum_{j<i} cost(\lambda[j], F_A(\lambda[0,j])) \leq 0$.

So, $input^{i+1} = \bigcup_{\lambda \in input^i} input^{i+1}(\lambda)$. After defining F_A, we have that for any $\lambda \in out(q, F_A)$ and $i \geq 0$, $\lambda[0,i] \in input^{i+1}$, hence $\lambda[i] \in Z \subseteq f(Z)$. Therefore, $\lambda[i] \in \|\varphi\|$. This shows that $q \in Y$.

So, Y is the greatest post-fixed point of $f(X)$, hence also the greatest fixed point of $f(X)$.

Proof of Lemma 4

Firstly, we have $\neg\langle\langle N^b\rangle\rangle \bigcirc \varphi$ and $\neg\langle\langle \emptyset^b\rangle\rangle \bigcirc \varphi$ are equivalent to $\langle\langle \emptyset^b\rangle\rangle \bigcirc \neg\varphi$ and $\langle\langle N^b\rangle\rangle \bigcirc \neg\varphi$, respectively. Therefore, we only consider the case when $\Phi_\bigcirc$ does not contain formulas of the form $\neg\langle\langle N^b\rangle\rangle \bigcirc \varphi$ and $\neg\langle\langle \emptyset^b\rangle\rangle \bigcirc \varphi$.

Let us assume that

$$\begin{aligned}\Phi_\bigcirc \quad = \quad & \{\langle\langle A_0^{b_0}\rangle\rangle \bigcirc \varphi_0, \ldots, \langle\langle A_{m-1}^{b_{m-1}}\rangle\rangle \bigcirc \varphi_{m-1}\} \cup \\ & \{\neg\langle\langle B_0^{d_0}\rangle\rangle \bigcirc \psi_0, \ldots, \neg\langle\langle B_{l-1}^{d_{l-1}}\rangle\rangle \bigcirc \psi_{l-1}\} \cup \\ & \{\langle\langle \emptyset^{e_0}\rangle\rangle \bigcirc \chi_0, \ldots, \langle\langle \emptyset^{e_{h-1}}\rangle\rangle \bigcirc \chi_{h-1}\}\end{aligned}$$

where all A_i are non-empty, all B_i are both non-empty and not equal to the grand coalition N.

Let e be a bound of resources such that $e > e_i$ for all $i \in \{0, \ldots, h-1\}$. We construct a tree with a root labelled by Φ and n^k children, each is denoted by $c = (a_1, \ldots, a_n)$ in which $a_i \in \{0, \ldots, k-1\}$. Intuitively, we allow each agent i to perform k different actions denoted by numbers from 0 to $k-1$, where the special action $k-1$ for all agents will be considered as the zero-cost idle action. We shall denote $c(i) = a_i$ for the action performed by agent i with the corresponding outcome c. Now we define the label $V(c)$ for each such node c. For each $\langle\langle A_p^{b_p}\rangle\rangle \bigcirc \varphi_p \in \Phi_\bigcirc$ wherein $A_p \neq \emptyset$, φ_p is added to $V(c)$ whenever $c(i) = p$ for all $p \in A_p$. Let $\min_{A_P}$ be the smallest number in A_P, we assign the cost of action p performed by $\min_{A_P}$ be b_p, i.e. $cost(p, \min_{A_p}) = b_p$. For all other $j \in A_p$ which is not $\min_{A_P}$, we set $cost(p, j) = \overline{0}$.

After considering all such $\langle\langle A_p^{b_p}\rangle\rangle \bigcirc \varphi_p \in \Phi_\bigcirc$, for all other unsigned-cost actions, i.e. actions $\geq m-1$ but $< k-1$ for all agents, we simply set their costs to be e. The action $k-1$ performed by all agents is defined to associate with the cost $\overline{0}$. We denote $cost(c) = \Sigma_i cost(c(i), i)$. Then, for each $\langle\langle \emptyset_p^{e_p}\rangle\rangle \bigcirc \chi_p \in \Phi_\bigcirc$, χ_p is added to $V(c)$ whenever $cost(c) \leq b_p$.

Finally, we will add at most one formula from the negation formulas of $\Phi_\bigcirc$ to $V(c)$. We denote $cost(c, A) = \Sigma_{i\in A} cost(c(i), i)$. For each c, let $\Phi_\bigcirc^-(c) = \{\neg\langle\langle B^d\rangle\rangle \bigcirc \psi \in \Phi_\bigcirc \mid cost(c, B) \leq d\} = \{\neg\langle\langle B_{i_0}^{d_{i_0}}\rangle\rangle \bigcirc \psi_{i_0}, \ldots, \neg\langle\langle B_{i_{l_c-1}}^{d_{l_c-1}}\rangle\rangle \bigcirc \psi_{l_c-1}\}$ in which $i_0 < i_1 < \ldots < i_{l_c-1}$. Let $I_c = \{i \mid m \leq c(i) < m + l_c\}$ and $j = \Sigma_{i\in I_c}(c(i) - m) \mod l_c$. Consider $\neg\langle\langle B_{i_j}^{d_j}\rangle\rangle \bigcirc \psi_j$: if $N \setminus B_{i_j} \subseteq I_c$, then $\neg\psi_{i_j}$ is added into $V(c)$.

We now need to show that our simple tree is locally consistent. In the first step, we show that all labels are consistent. It is obvious that $V(\epsilon) = \Phi$ is consistent.

Let us firstly consider every child c of the root in which $\neg\psi_q \in V(c)$ from some negation formula in $\Phi_\bigcirc$. This will imply that there will be no $\chi \in V(c)$ from the formulas of the form $\langle\langle \emptyset^b\rangle\rangle \bigcirc \chi$ in $\Phi_\bigcirc$. The reason is that as some $\neg\psi_q \in V(c)$, there will be some agent performing an action $i \geq m$. As the cost of this action is e, $cost(c) \geq e$, hence, no χ will be added into $V(c)$.

When there is no $\varphi \in V(c)$ from the formulas of the form $\langle\langle A^b\rangle\rangle \bigcirc \varphi$ in $\Phi_\bigcirc$, the proof is trivial as there is only one $\neg\psi_q \in V(c)$. If there are some $\varphi_p \in V(c)$ where $\langle\langle A_p^{b_p}\rangle\rangle \bigcirc \varphi_p \in \Phi_\bigcirc$, then for each p, $c(i) = p < m$ for all $i \in A_p$. Hence, all A_p are pairwise disjoint. Moreover, we have that $B_q \supseteq N \setminus I_c$ and all agents in $\bigcup_{\varphi_p \in V(c)} A_p$ perform some action $< m$, hence $B_q \supseteq \bigcup_{\varphi_p \in V(c)} A_p$. From the selection of ψ_q, we also have that $d_q \geq \Sigma_{i\in B_q} cost(c(i), i) \geq \Sigma_{i \in \bigcup_{\varphi_p \in V(c)} A_p} cost(c(i), i) \geq \Sigma_{\varphi_p \in V(c)} b_p$. This simply shows that the set of $\langle\langle A_p^{b_p}\rangle\rangle \bigcirc \varphi_p \in \Phi_\bigcirc$ where $\varphi_p \in V(c)$ and $\neg\langle\langle B_q^{d_q}\rangle\rangle \bigcirc \psi_q$ satisfies the conditions of Lemma 2. Therefore, $V(c)$ is consistent.

Now, we consider every child c of the root in which there is no $\neg\psi \in V(c)$ from some negation formula in $\Phi_\bigcirc$.

When there is no $\varphi \in V(c)$ from the formulas of the form $\langle\langle A^b\rangle\rangle \bigcirc \varphi$ in $\Phi_\bigcirc$, the proof is trivial as there are at most only some $\chi_q \in V(c)$. If there are some $\varphi_p \in V(c)$ where $\langle\langle A_p^{b_p}\rangle\rangle \bigcirc \varphi_p \in \Phi_\bigcirc$ and $A_p \neq \emptyset$, then for each p, $c(i) = p < m$ for all $i \in A_p$. Hence, all A_p are pairwise disjoint. For any $\chi_q \in V(c)$ by some $\langle\langle \emptyset^{e_q}\rangle\rangle \bigcirc \chi_q \in \Phi_\bigcirc$, we have that $e_q \geq cost(c) \geq \Sigma_p b_p$. This simply shows that the set of $\langle\langle A_p^{b_p}\rangle\rangle \bigcirc \varphi_p \in \Phi_\bigcirc$ where $\varphi_p \in V(c)$ and $\langle\langle \emptyset_q^{e_q}\rangle\rangle \bigcirc \chi_q$ satisfies the conditions of Lemma 3. Therefore, $V(c)$ is consistent.

Let us now check the conditions of local consistency on the newly built tree. For $\langle\langle A_p^{b_p}\rangle\rangle \bigcirc \varphi_p \in \Phi_\bigcirc$, it is straightforward that the move σ_{A_p} where all agents in A_p perform action $p < m$ of cost b_p and for any $c \in out(\sigma_{A_p})$, $\varphi_p \in V(c)$. For $\neg\langle\langle B_p^{d_p}\rangle\rangle \bigcirc \psi_p \in \Phi_\bigcirc$ and σ being an arbitrary move of agents in B_p with cost at most equal to d_p, we will point out an output $c \in out(\sigma_{B_p})$ such that $\neg\psi_p \in V(c)$. That is, we need to select suitable actions for agents outside B_p such that the output's label will contain $\neg\psi_p$. We will select only actions $\geq m$ for those agents, then their costs are fixed at e. That is to say even without knowing which actions are performed by agents out of B_p, we know the cost of each action as well as the total cost of c. Hence, l_c is determined as well as the set $\Phi_\bigcirc^-(c) = \{\neg\langle\langle B_{i_0}^{d_{i_0}}\rangle\rangle \bigcirc \psi_{i_0}, \ldots, \neg\langle\langle B_{i_{l_c-1}}^{d_{i_{l_c-1}}}\rangle\rangle \bigcirc \psi_{i_{l_c-1}}\}$ where $p = i_r$ for some $r < l_c$. Let σ_i be the action performed by agent i in B_p, we have $c(i) = \sigma_i$ for all $i \in B_p$. Let $I' = \{i \in B_q \mid m \leq c(i) < m + l_c\}$ and $j' = \Sigma_{i\in I'}(c(i) - m) \mod l_c$. We select an arbitrary $i \notin B_p$ and set $c(i) = m + (r - j') \mod l_c$. For all other $i \notin B_p$, let $c(i) = m$. Then, we have $I_c = \{i \mid m \leq c(i) < m + l_c\}$ and $\Sigma_{i\in I}(c(i) - m) \mod l_c = r$, hence $\neg\psi_p \in V(c)$. Notice that an action $\geq m$ and $< k-1$ always costs e.