

A Logic for Coalitions with Bounded Resources*

Natasha Alechina

Brian Logan

Nguyen Hoang Nga

Abdur Rakib

School of Computer Science

University of Nottingham

Nottingham, UK

{nza,bsl,hnn,rza}@cs.nott.ac.uk

Abstract

Recent work on Alternating-Time Temporal Logic and Coalition Logic 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. This paper presents a Resource-Bounded Coalition Logic (RBCL) which has explicit representation of resource bounds in the language, and gives a complete and sound axiomatisation of RBCL.

1 Introduction

Recent work on Alternating-Time Temporal Logic ATL and Coalition Logic CL (for example, [Pauly, 2001; Goranko, 2001; Pauly, 2002; Alur *et al.*, 2002; Wooldridge *et al.*, 2007]) 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.

The motivation for this work is in verifying problems of the following form: can a set of agents C achieve a state of the world satisfying ϕ under the given resource bound b . Essentially, this is the *successful coalition under resource bound* problem from [Wooldridge and Dunne, 2006]. Unlike Wooldridge and Dunne [2006], we consider multi-shot games where the agents need to perform a *sequence* of actions to achieve the goal. As a motivating example, consider a system of distributed reasoning agents as described in [Alechina *et al.*, 2008], whose actions involve inferring new clauses from their knowledge bases and communicating results to other agents. Clearly these activities require resources such as time (which can be identified with the number of inference steps), memory (the space required to store premises in reasoning, and any intermediate lemmata, which can be measured as the number of clauses in the agents memory at any one time) and communication bandwidth (which can be measured as the number of communicated clauses). Properties of interest for such systems include ‘the set of reasoners A can derive the clause $[p]$ under the resource bound 10 for time, 3 for memory and 2 for communication’. In general, we would like to be able to express any properties of systems

where abilities of individual agents and coalitions of agents are constrained by available resources in a non-trivial way, and the properties relating to resource bounds are important.

In this paper, we propose a logic in which we can describe systems of agents specified in terms of the resources required to perform actions. First we introduce a simple formalism $RBCL_1$ which describes single-step strategies and is based on Coalition Logic [Pauly, 2002] extended with resource bounds. We then motivate multi-step strategies and introduce a more complex logic $RBCL$ which can express multi-step strategies (in a sense corresponding to the Extended Coalition Logic of [Pauly, 2001]). We give a sound and complete axiomatisation of $RBCL$.

2 Formalising Single Step Strategies

We assume a set of agents $A = \{1, \dots, n\}$ and a set of resources $R = \{1, \dots, r\}$. Agents can perform actions from a set $\Sigma = \cup_{i \in A} \Sigma_i$, where Σ_i is the set of actions agent i can perform. Each action $a \in \Sigma$ has an associated cost $Res(a)$, which is a vector of costs (assumed to be natural numbers) for each resource in R . A coalition $C \subseteq A$ can execute a joint action a_C , where a_C is a tuple of actions (a_1, \dots, a_k) (we assume for simplicity unless otherwise stated that $C = \{1, \dots, k\}$ for some $k \leq n$). For the moment, let us stipulate that the cost of a joint action a_C corresponds to the vector sum of costs of actions in a_C (we will generalise the way to combine costs for different resources later). We compare vectors of resources using pointwise vector comparison \leq , e.g., for $b = (b_1, \dots, b_r)$ and $d = (d_1, \dots, d_r)$, $b \leq d$ iff for each $j \leq r$, $b_j \leq d_j$.

The language of $RBCL_1$ is defined relative to the sets A and R and a set of propositional variables $Prop$. A formula is defined as follows:

$$p \mid \neg\phi \mid \phi \wedge \psi \mid [C^b]\phi$$

where $p \in Prop$, $C \subseteq A$, and $b \in \mathbb{N}^r$. The intuitive meaning of $[C^b]\phi$ for $C \neq \emptyset$ is that coalition C can force the outcome ψ under resource bound b , or in other words C has a strategy costing at most b which enables them to achieve a ϕ -state no matter what the other agents $\bar{C} = A \setminus C$ do. For the empty coalition, $[\emptyset^b]\phi$ means that if the grand coalition A executes any joint action which together costs at most b , then the system will end up in a ϕ state; that is, ϕ is unavoidable if A acts within resource bound b .

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

We define models of $RBCL_1$ as transition systems, where in each state agents execute some actions in parallel which determines the next state. These are essentially the same as the models for coalition logic, with the addition of costs of actions. First we define action frames which underlie the models:

Definition 1. A resource-bounded action (RBA) frame F is a tuple $(A, R, \Sigma = \cup_{i \in A} \Sigma_i, S, T, o, Res)$ where:

A is a non-empty set of agents,

R is a non-empty set of resources,

Σ is the set of actions agents can perform,

S is a non-empty set of states,

$T : S \times A \rightarrow \wp(\Sigma_i)$ assigns to each state the set of actions available to agent i in this state; this set is always non-empty as it contains an action *noop* with $Res(\text{noop}) = \vec{0} = (0, \dots, 0)$,

o is the outcome function which takes a state s and a joint action a_A and returns the state resulting from the execution of a_A by the agents in s ,

$Res : \Sigma \rightarrow \mathbb{N}^r$ is the resource requirement function.

In the case of joint actions, we generalise the function T as follows: a joint action $a_C \in T(s, C)$ iff $a_i \in T(s, i)$ for all $i \in C$. By $Res(a_C)$ we denote the vector sum of $Res(a_i)$ for $i \in C$.

Definition 2. A single-step resource-bounded action model M is a pair (F, V) where F is an RBA frame, and $V : S \rightarrow \wp(Prop)$ is an assignment function.

The truth definition for single-step RBA models is as follows:

- $M, s \models p$ iff $p \in V(s)$
- $M, s \models \neg\phi$ iff $M, s \not\models \phi$
- $M, s \models \phi \wedge \psi$ iff $M, s \models \phi$ and $M, s \models \psi$
- $M, s \models [C^b]\phi$ for $C \neq \emptyset$ iff there is $a_C \in T(s, C)$ with $Res(a_C) \leq b$ such that for every joint action $a_{\bar{C}} \in T(s, \bar{C})$ by the agents not in C , the outcome of the resulting tuple of actions executed in s satisfies ϕ : $M, o(s, (a_C, a_{\bar{C}})) \models \phi$
- $M, s \models [\emptyset^b]\phi$ iff the outcome of any joint action $a_A \in T(s, A)$ with $Res(a_A) \leq b$ executed in s satisfies ϕ : $M, o(s, a_A) \models \phi$.

The notions of satisfiability and validity are standard. Let us call the set of all formulas valid in single-step RBA models $RBCL_1$ (where 1 refers to considering only one-step strategies, as in Coalition Logic).

Theorem 1. $RBCL_1$ is completely axiomatised by the following set of axiom schemas and inference rules:

A0 All propositional tautologies

A1 $[C^b]\top$

A2 $\neg[C^b]\perp$

A3 $\neg[\emptyset^b]\varphi \rightarrow [A^b]\neg\varphi$

A4 $[C^b](\varphi \wedge \psi) \rightarrow [C^b]\varphi$

A5 $[C^b]\varphi \rightarrow [C^d]\varphi$ where $d \geq b$ if $C \neq \emptyset$ or $d \leq b$ if $C = \emptyset$

A6 $[C^b]\varphi \wedge [D^d]\psi \rightarrow [(C \cup D)^{b+d}](\varphi \wedge \psi)$

MP $\vdash \varphi, \vdash \varphi \rightarrow \psi \Rightarrow \vdash \psi$

Equivalence $\vdash \varphi \leftrightarrow \psi \Rightarrow \vdash [C^b]\varphi \leftrightarrow [C^b]\psi$

The notions of derivability and consistency are standard. Note that if we erase the resource superscript in the axiomatisation above, we get the complete axiomatisation of Coalition Logic as given in [Pauly, 2002], and a trivial formula resulting from **A5**. The rule of monotonicity (RM) is derivable, as in Coalition Logic, that is if $\vdash \varphi \rightarrow \psi$, then $\vdash [C^b]\varphi \rightarrow [C^b]\psi$.

We omit the completeness proof here as it is a special case of completeness proof of RBCL given below.

As an illustration, we state a property of the first example given in [Wooldridge and Dunne, 2006], that the coalition of agents 1 and 2 can achieve g_1 within the resource bound $(3, 2)$. This can be expressed in $RBCL_1$ as $[\{1, 2\}^{(3,2)}]g_1$.

3 Formalising Multi-step Strategies and Arbitrary Resource Combinators

In what follows, we generalise the logic described above in two ways. First, we consider multi-step strategies, as in Extended Coalition Logic with the $[C^*]$ operator [Pauly, 2001], or as in ATL. The reason for this is as follows. We are interested in the costs of strategies which involve multiple steps. For example, suppose a coalition C can enforce their goal ϕ in three steps: $[C^{b_1}][C^{b_2}][C^{b_3}]\phi$. We can deduce from this that the agents have a strategy to achieve ϕ which costs at most $b_1 + b_2 + b_3$. However expressing the fact in this way is rather clumsy. Even worse, to say that ‘ C has some strategy which achieves ϕ in three steps which costs at most b ’, we have to use a disjunction over all possible vectors of natural numbers b_1, b_2, b_3 which sum up to b : $\bigvee_{b_1+b_2+b_3=b} [C^{b_1}][C^{b_2}][C^{b_3}]\phi$. Hence we extend the set of actions, or strategies, with sequential compositions of actions.

The second direction of generalisation involves the way we calculate the resource requirements of complex actions. We argue that not all resource costs should be combined using simple addition. For example, if one of resources is time and the agents execute their actions concurrently, then if each individual action costs one unit of time, the parallel combination of those actions also costs one unit of time. If one of the resources is memory, one can argue that if action a_1 requires n units of memory and action a_2 requires m units of memory, then executing actions a_1 and a_2 sequentially requires $\max(n, m)$ units of memory. For generality, we introduce two cost combinators to express how amounts of resources are combined in parallel and sequentially. If two actions a_1 and a_2 are performed in parallel, then the cost of executing them is $Res(a_1) \oplus Res(a_2)$ and the cost of executing them sequentially is $Res(a_1) \otimes Res(a_2)$ (where \oplus and \otimes may be sum for some resources, and \max or some other combinator for others).

In the rest of the paper, we assume that the set of resources R always includes time, that every action costs exactly one unit of time, and that the time cost is the last component of every cost vector. The cost of the *noop* action is redefined as $(0, \dots, 0, 1)$. We denote by $t(b)$ the time component of cost

vector b . In particular, $t(Res(a)) = 1$ for any $a \in \Sigma$. In the language, only operators $[C^b]$ with $t(b) \geq 1$ are allowed.

3.1 Strategies and Multi-step RBA Models

Given an RBA frame $F = (A, R, \Sigma, S, T, o, Res)$, a *strategy* for an agent $i \in A$ is a function $f_i : S^+ \rightarrow \Sigma_i$ from finite non-empty sequences of states to actions, such that $f_i(\lambda s) = a \in T(s, i)$, where λs is a sequence of states ending in state s . Intuitively, f_i says what action agent i should perform in state s given the previous history of the system. A strategy for a coalition C is a set $F_C = \{f_1, \dots, f_k\}$ of strategies for each agent.

For a sequence $\lambda = s_0 s_1 \dots \in S^\omega$, we denote $\lambda[i] = s_i$ and $\lambda[i, j] = s_i \dots s_j$.

The set of possible computations generated by a strategy F_C from a state s_0 , $out(s_0, F_C)$, is

$$\{\lambda \mid \lambda[0] = s_0 \wedge \forall j > 0 : \lambda[j+1] \in o^*(\lambda[j], (f_i(\lambda[0, j]))_{i \in C})\}$$

where $o^*(s, a_C) = \{o(s, (a_C, a_{\bar{C}})) \mid a_{\bar{C}} \in T(s, \bar{C})\}$. Now we define the cost of a multi-step strategy. Let $\lambda \in out(s_0, F_C)$. The cost of F_C over a prefix $\lambda[0, m]$ where $m > 0$ is defined inductively as follows:

$$Res(\lambda[0, 1], F_C) = \bigoplus_{i \in C} Res(f_i(\lambda[0])) \quad (\text{where } Res(f_i(\lambda[0])) \text{ is the cost of the action by agent } i \text{ in } \lambda[0], \text{ and } \bigoplus_{i \in C} \text{ is the operator for combining the costs of actions executed in parallel by the agents in } C)$$

$$Res(\lambda[0, m], F_C) = Res(\lambda[0, m-1], F_C) \otimes (\bigoplus_{i \in C} Res(f_i(\lambda[0, m-1]))) \quad \text{for } m > 1 \text{ (this is the cost of the previous } m-1 \text{ steps in the strategy combined sequentially with the cost of the } m \text{th step).}$$

Definition 3. A *multi-step resource-bounded action model* M is a pair (F, V) where F is an RBA frame, and $V : S \rightarrow \wp(Prop)$ is an assignment function, and the truth definition for the $[C^b]$ modality is

- $M, s \models [C^b]\phi$ for $C \neq \emptyset$ iff there is a strategy F_C such that for all $\lambda \in out(s, F_C)$, there exists $m > 0$ such that $Res(\lambda[0, m], F_C) \leq b$ and $M, \lambda[m] \models \phi$,
- $M, s \models [\emptyset^b]\phi$ iff for all strategies F_A , computations $\lambda \in out(s, F_A)$, and $m > 0$ such that $Res(\lambda[0, m], F_A) \leq b$, $M, \lambda[m] \models \phi$.

The set of all formulas valid in multi-step RBA models will be denoted by RBCL.

3.2 Effectivity Structures

For proving completeness of RBCL, it is easier to work with an alternative semantics, given not in terms of multi-step RBA models, but in terms of effectivity structures. These are closely related to RBA models, and we will show that effectivity structures satisfying some natural properties give rise to an alternative semantics for RBCL.

Let $\wp(A)^{\mathbb{B}} = \{C^b \mid C \subseteq A, b \in \mathbb{N}^r, t(b) \geq 1\}$. Intuitively, this is the set of all possible coalitions with all possible resource allocations. An *effectivity structure* is a function $E : S \rightarrow (\wp(A)^{\mathbb{B}} \rightarrow \wp(\wp(S)))$ which describes, for each state in S , which subsets of S a coalition C can force under resource bound b .

Given an RBA frame F , the *effectivity structure corresponding to F* is defined as follows:

- For $C \neq \emptyset$, $X \in E(s)(C^b)$ iff there exists a strategy F_C such that for all $\lambda \in out(s, F_C)$, there exists $m > 0$ such that $Res(\lambda[0, m], F_C) \leq b$ and $\lambda[m] \in X$.
- $X \in E(s)(\emptyset^b)$ iff for all strategies F_A , computations $\lambda \in out(s, F_A)$, and $m > 0$ such that $Res(\lambda[0, m], F_A) \leq b$, we have $\lambda[m] \in X$.

In other words, $X \in E(s)(C^b)$, where C is not the empty coalition, means that the coalition C has a strategy to bring about X within the bound b . $X \in E(s)(\emptyset^b)$ means that whatever the grand coalition does under the cost b , the system always goes to a state in X , i.e. it cannot avoid X .

3.3 Characterising Effectivity in RBA Frames

Every RBA frame gives rise to an effectivity structure, but the reverse does not hold. In this section, we characterise properties which an effectivity structure should satisfy to be an effectivity structure corresponding to an RBA frame. Following Pauly in [Pauly, 2002], we call such effectivity structures *playable* (RB-playable, where RB stands for resource-bounded).

Below we state some useful properties of RB-playable effectivity structures. These are very similar (apart from the resource bound) to the properties of playable effectivity structures listed in [Pauly, 2002] and are given the same names:

An effectivity structure E is *outcome monotonic* iff
 $X \in E(s)(C^b) \Rightarrow X' \in E(s)(C^b)$ for all $X' \supseteq X$

An effectivity structure E is *coalition monotonic* iff
 $X \in E(s)(C^b) \Rightarrow X \in E(s)(D^b)$ for all $D \supseteq C$

An effectivity structure E is *A-maximal* iff
 $X \notin E(s)(\emptyset^b) \Rightarrow \bar{X} \in E(s)(A^b)$

An effectivity structure E is *A-minimal* iff
 $X \in E(s)(A^b) \wedge Y \notin E(s)(A^b) \Rightarrow X \setminus Y \in E(s)(A^b)$
 (this property is not listed in [Pauly, 2002], but its analogue is derivable there)

An effectivity structure E is *regular* iff
 for $C = \emptyset$ or $C = A$, $X \in E(s)(C^b) \Rightarrow \bar{X} \notin E(s)(\bar{C}^b)$
 for all other C , $X \in E(s)(C^b) \Rightarrow \bar{X} \notin E(s)(\bar{C}^{b'})$ for all $t(b) = t(b') = 1$.

An effectivity structure E is *super-additive* iff the following holds, for all $t(b) = t(d) = 1$, $C \cap D = \emptyset$:
 - if $C \neq \emptyset$ and $D \neq \emptyset$, $X_1 \in E(s)(C^b)$ and $X_2 \in E(s)(D^d) \Rightarrow X_1 \cap X_2 \in E(s)((C \cup D)^{b \oplus d})$.
 - if $C = \emptyset$, $X_1 \in E(s)(\emptyset^d)$ and $X_2 \in E(s)(D^d) \Rightarrow X_1 \cap X_2 \in E(s)(D^d)$.

Note that super-additivity applies only to effectivity functions corresponding to single-step strategies (with the time component of the bound equal to 1). Since effectivity functions correspond to multi-step strategies, we need to extend super-additivity to more than single-step strategies, that is to the case when the time component is greater than 1:

An effectivity structure E is *general super-additive* iff it is super-additive and $X_1 \in E(s)(\emptyset^b)$ and $X_2 \in E(s)(C^b) \Rightarrow X_1 \cap X_2 \in E(s)(C^b)$.

We also need properties corresponding to sequential composition of strategies:

An effectivity structure E is *super-transitive* iff the following holds for all $C \neq \emptyset$: $\{s' \in S \mid X \in E(s')(C^{b_2})\} \in E(s)(C^{b_1}) \Rightarrow X \in E(s)(C^{b_1 \otimes b_2})$ (if a set of states where X is obtainable under b_2 can be enforced under b_1 , then X can be enforced by the combined strategy under $b_1 \otimes b_2$).

An effectivity structure E is *transitive* iff for any b with $t(b) > 1$ and $C \neq \emptyset$: $X \in E(s)(C^b) \Rightarrow \exists b' < b : X \in E(s)(C^{b'})$ (X can be achieved under a tighter bound b') or $\exists b_1 \otimes b_2 = b : \{s' \in S \mid X \in E(s')(C^{b_2})\} \in E(s)(C^{b_1})$ (X can be achieved by combining two strategies costing b_1 and b_2 such that $b_1 \otimes b_2 = b$).

Finally, the following property is specific to resource bounds:

An effectivity structure E is *bound-monotonic* iff $X \in E(s)(C^b) \Rightarrow X \in E(s)(C^d)$ for all $d \geq b$ if $C \neq \emptyset$ or $d \leq b$ if $C = \emptyset$.

Bound-monotonicity is a very natural property: if a non-empty coalition can achieve something under the bound b , then it can achieve it with a more generous resource allowance. For $C = \emptyset$, this property means that if an outcome cannot be avoided when the grand coalition is restricted to strategies which cost at most b , then it cannot be avoided if A uses fewer resources (hence has fewer strategies available).

It is easy to prove that the properties above are true for any effectivity structure obtained from a RBA frame. Conversely, RB-playable effectivity structures defined below are effectivity structures of an RBA frame.

Definition 4. An effectivity structure $E : S \rightarrow (\wp(A)^{\mathbb{B}} \rightarrow \wp(\wp(S)))$ is *RB-playable* iff, for every $s \in S$, E has the following properties:

1. For all $C^b \in \wp(A)^{\mathbb{B}}$, $S \in E(s)(C^b)$
2. For all $C^b \in \wp(A)^{\mathbb{B}}$, $\emptyset \notin E(s)(C^b)$
3. Outcome-monotonicity
4. A-maximality
5. Super-additivity
6. Super-transitivity
7. Transitivity
8. Bound-monotonicity

It can be shown that RB-playability implies the other properties listed above.

Lemma 1. Let E be a RB-playable effectivity structure, then E has the following properties:

1. Coalition monotonicity
2. A-minimality
3. Regularity
4. General super-additivity

The proof is omitted due to lack of space. First, general super-additivity is proved by induction on resource bounds using super-additivity, and the proof for other properties uses general super-additivity.

Theorem 2. An effectivity structure is RB-playable iff it is the effectivity structure of some RBA frame.

Proof. It is easy to check that effectivity structures obtained from RBA frames satisfy all properties of RB-playability. For the other direction, given an RB-playable effectivity structure E , we will construct a RBA frame such that its effectivity structure is E .

Let E be RB-playable. The construction of the RBA frame is similar to that in Coalition Logic extended with costs for actions. First, we define the set of possible actions for each agent at each state $s \in S$ with their associated costs Res . Then, the construction is completed by defining the outcome function o .

For $i \in A$ and b such that $t(b) = 1$, we denote $\mathcal{C}_i^b = \{C^d \mid i \in C \wedge t(d) = 1 \wedge d \geq b\}$ the set of all coalitions that i may participate in while contributing b amount of resources. Note that for all actions $t(b)$ is always 1.

For $s \in S$, we define

$$\Gamma(s, i) = \{g_i^b(s) : \mathcal{C}_i^b \rightarrow \wp(S) \mid g_i^b(s)(C^d) \in E(s)(C^d)\}$$

$\Gamma(s, i)$ is the set of option functions for an agent i at state s . Each option function in $\Gamma(s, i)$ is a mapping $g_i^b(s)$ which determines the outcome if agent i agrees to participate in a coalition. How an agent agrees to participate in a coalition will be specified later when we define the outcome function.

Let $H = \{h : \wp(S) \rightarrow S \mid h(X) \in X\}$ be the set of choice functions, that is if an agent has the power to decide the outcome, it will use some h function to do that. Then, we define the set of available actions for an agent i at a state s as follows:

$$T(s, i) = \Gamma(s, i) \times \mathbb{N} \times H$$

Each action is a triple $(g^b(s), t, h)$ consisting of an option function g^b , an index t (a natural number) and a choice function h . Informally, option functions will determine how the agents group together to form coalitions and then which outcome options they will choose. The index determines which agent has the power to decide the outcome based on its associated h function. For now, we assign that $Res((g^b(s), t, h)) = b$.

Let $\Sigma_i = \bigcup_{s \in S} T(s, i)$. We now define the outcome of a joint action $\sigma \in \Sigma_A$ at a state s . Assume that $\sigma = \{(g_i^{b_i}(s), t_i, h_i) \mid i = 1, \dots, n\}$ in which $t(b_i) = 1$ for all $i \in A$. Let $b_C = \oplus_{i \in C} b_i$, $g = (g_i^{b_i}(s))_{i \in A}$. We denote $P(f, C)$ the coarsest partition $\langle C_1, \dots, C_m \rangle$ of C such that:

$$\forall l \leq m \forall i, j \in C_l : g_i^{b_i}(C_l^{b_{C_l}}) = g_j^{b_j}(C_l^{b_{C_l}})$$

We define how coalitions are formed based on g as follows:

$$\begin{aligned} P_0(g) &= \langle A \rangle \\ P_1(g) &= \langle P(g, A) \rangle = \langle C_{1,1}, \dots, C_{1,k_1} \rangle \\ P_2(g) &= \langle P(g, C_{1,1}), \dots, P(g, C_{1,k_1}) \rangle \\ &= \langle C_{2,1}, \dots, C_{2,k_2} \rangle \\ &\vdots \\ P_\eta(g) &= \langle C_{\eta,1}, \dots, C_{\eta,k_\eta} \rangle \end{aligned}$$

As A is finite, the above computation reaches some η such that $P_\eta(g) = P_{\eta+1}(g)$. Let $P(g) = P_\eta(g)$ which shows how agents are grouped into coalitions.

For technical convenience, let $E^o(s)(A^b)$ denote the collection of minimal sets in $E(s)(A^b)$. By A-minimality, it

is easy to show that $E^o(s)(A^b)$ contains only singletons. In other words, by outcome-monotonicity, we have $X \in E(s)(A^b)$ if and only if $X \supseteq X^o$ for some $X^o \in E^o(s)(A^b)$. By regularity, we have $X \in E(s)(\emptyset^b)$ if and only if $X \supseteq \cup E^o(s)(A^b)$.

Assume that $P(g) = \langle C_1, \dots, C_m \rangle$. For convenience, let $g(C_l) = g_i^{b_i}(C_l^{b_{C_l}})$ for some $i \in C_l$ where $l \leq m$.

We define $G(g) = \bigcap_{l \leq m} g(C_l) \cap (\cup E^o(s)(A^{b_A}))$. By super-additivity and the fact that $\emptyset \notin E(s)(A^{b_A})$ as E is RB-playable, it is straightforward to show that $G(g) \neq \emptyset$.

Let $t_0 = (\sum_{i \in A} t_i \bmod n) + 1$. The outcome function is defined as follows: $o(s, \sigma) = h_{t_0}(G(g))$. Let E_F be the effectivity structure of the frame constructed above. We claim that $E = E_F$.

Firstly, we show the left-to-right inclusion by induction on bounds. In the base case, assume $X \in E(s)(C^b)$ in which $t(b) = 1$. Choose the actions for agents in C as follows,

$$\begin{aligned} a_1 &= (g_1^b, t_1, h_1) \\ a_2 &= (g_2^{\bar{0}}, t_2, h_2) \\ &\vdots \\ a_k &= (g_k^{\bar{0}}, t_k, h_k) \end{aligned}$$

where $g_1^b(D^d) = g_i^{\bar{0}}(D^d) = X$ for all $i = 2, \dots, k$, $D \supseteq C$, $d \geq b$. Notice that the choices of $g_1^b, g_2^{\bar{0}}, \dots, g_k^{\bar{0}}$ must exist because of bound-monotonicity and coalition-monotonicity. Moreover, the choices of t_i and h_i , where $i = 1, \dots, k$, are arbitrary. Let $\sigma_C = \{(g_1^b, t_1, h_1), (g_2^{\bar{0}}, t_2, h_2), \dots, (g_k^{\bar{0}}, t_k, h_k)\}$.

Let $\sigma_{\bar{C}}$ be an arbitrary joint action for \bar{C} . Let $\sigma = (\sigma_C, \sigma_{\bar{C}})$ and let g be the set of the option functions from σ . By the choice of σ_C , C must be a subset of a partition C_l in $P(g)$. Then, we have

$$o(s, \sigma) = h_{t_0}(G(g)) \in G(g) \subseteq g(C_l) = X$$

Hence, $X \in E_F(s)(C^b)$.

For the induction step, let $X \in E(s)(C^b)$ in which $t(b) > 1$. If $X \in E(s)(C^{b'})$, by the induction hypothesis, we have $X \in E_F(s)(C^{b'})$. Therefore, bound-monotonicity implies that $X \in E_F(s)(C^b)$.

If $X \notin E(s)(C^{b'})$ for any $b' < b$, by transitivity there are $b_1 \otimes b_2 = b$ such that

$$\{s' \mid X \in E(s')(C^{b_2})\} \in E(s)(C^{b_1})$$

By the induction hypothesis, we have

$$\{s' \mid X \in E(s')(C^{b_2})\} \in E_F(s)(C^{b_1})$$

and

$$\begin{aligned} &\{s' \mid X \in E(s')(C^{b_2})\} \\ &\subseteq \{s' \mid X \in E_F(s')(C^{b_2})\} \end{aligned}$$

By outcome-monotonicity, we have

$$\{s' \mid X \in E_F(s')(C^{b_2})\} \in E_F(s)(C^{b_1})$$

Hence, by super-transitivity $X \in E_F(s)(C^b)$.

For the other direction, we consider two cases in which $C = A$ and $C \subset A$.

Assume that $X \notin E(s)(A^b)$. By A -maximality, we obtain $\bar{X} \in E(s)(\emptyset^b)$. However, the previous proof implies that $\bar{X} \in E_F(s)(\emptyset^b)$. As E_F is RB-playable, by regularity we have $X \in E_F(s)(A^b)$.

For the case of $C \subset A$, the proof is done by induction on bounds. Assume that $X \notin E(s)(C^b)$ in which $t(b) = 1$ and $C \subset A$, i.e. there is $i_0 \in A \setminus C$. Let $\sigma_C = \{(g_i^{b_i}(s), t_i, h_i) \mid i \in C\}$ be an joint action for C such that $Res(\sigma_C) \leq b$. We choose a strategy for \bar{C} $\sigma_{\bar{C}} = \{(g_i^{b_i}(s), t_i, h_i) \mid i \in \bar{C}\}$ such that:

- $b_i = \bar{0}$ for all $i > k$
- $g_i^{b_i}(s)(D^d) = S$ for all $i \in \bar{C}$, $D \supseteq \bar{C}$, $d \geq b_i$
- $(\sum_{i \in A} t_i \bmod n) + 1 = i_0$
- h_i for $i \neq i_0$ is arbitrary, we will select h_{i_0} shortly

As before, let $\sigma = (\sigma_C, \sigma_{\bar{C}})$ and g the collection of option functions in σ . We use notation $b_D = \oplus_{i \in D} b_i$ for any $D \subseteq A$.

By the choice of option functions in $\sigma_{\bar{C}}$, it follows that \bar{C} is the subset of some partition C_l of $P(g)$. For other partitions, super-additivity shows that $G(g) \in E(s)(\bar{C}_l^{b_{\bar{C}_l}})$. By coalition-monotonicity and bound-monotonicity, we have that $G(g) \in E(s)(C^b)$. As $X \notin E(s)(C^b)$, it follows that $G(g) \not\subseteq X$ by outcome-monotonicity, i.e. there is some $s_0 \in G(g) \setminus X$. Select h_{i_0} such that $h_{i_0}(G(g)) = s_0$, then

$$o(s, \sigma) = h_{i_0}(G(g)) = s_0 \notin X$$

Hence, $X \notin E_F(s)(C^b)$.

In the induction step, assume that $X \notin E(s)(C^b)$ where $t(b) > 1$. Bound-monotonicity shows that for all $b' \leq b$, $X \notin E(s)(C^{b'})$ and super-transitivity implies that for all $b_1 \otimes b_2 = b$,

$$\{s' \mid X \in E(s')(C^{b_2})\} \notin E(s)(C^{b_1})$$

By the induction hypothesis, we have that for all $b' < b$, $X \notin E_F(s)(C^{b'})$ and for all $b_1 \otimes b_2 = b$,

$$\{s' \mid X \in E(s')(C^{b_2})\} \notin E_F(s)(C^{b_1})$$

and $\{s' \mid X \in E(s')(C^{b_2})\} = \{s' \mid X \in E_F(s')(C^{b_2})\}$. Then, $\{s' \mid X \in E_F(s')(C^{b_2})\} \notin E_F(s)(C^{b_1})$. Therefore, transitivity implies that $X \notin E_F(s)(C^b)$. \square

4 Axiomatisation of RBCL

In this section we define models based on playable effectivity structures, and give a complete axiomatisation for the set of validities in those models.

Definition 5. A resource-bounded effectivity model $M = (S, E, V)$ is a triple consisting of a non-empty set of states, a RB-playable effectivity structure and a valuation function $V : Prop \rightarrow \wp(S)$. The truth definition for $[C^b]$ modalities is as follows:

- $M, s \models [C^b]\varphi$ iff $\varphi^M \in E(s)(C^b)$ where $\varphi^M = \{s' \mid M, s' \models \varphi\}$

Theorem 3. *The sets of formulas valid in multi-step RBA models and in resource-bounded effectivity models are equal.*

This follows from the correspondence between RBA frames and RB-playable effectivity structures, and the correspondence between the two truth definitions. Therefore the next result also provides an axiomatisation for RBCL.

Theorem 4. *The following set of axiom schemas and inference rules provides a sound and complete axiomatisation of the set of validities over all resource-bounded effectivity models:*

A0-A5, MP and Equivalence given above

A6 $[C^b]\varphi \wedge [D^d]\psi \rightarrow [(C \cup D)^{b \oplus d}](\varphi \wedge \psi)$ where $C \cap D = \emptyset$, $t(b) = t(d) = 1$ and $b = d$ if either C or D is empty

A7 $[C^{b_1}][C^{b_2}]\varphi \rightarrow [C^{b_1 \otimes b_2}]\varphi$ for $C \neq \emptyset$

A8 $[C^b]\varphi \rightarrow \bigvee_{b' < b} [C^{b'}]\varphi \vee \bigvee_{b_1 \otimes b_2 = b} [C^{b_1}][C^{b_2}]\varphi$ for all $C \neq \emptyset$

Proof. The proof of soundness is straightforward. We concentrate on proving completeness. We prove completeness by constructing a canonical model. Let us denote by \vdash_{Λ} derivability in the axiom system above. Let S^{Λ} be the set of all Λ -maximally consistent sets. For any formula φ , we denote $\tilde{\varphi} = \{s \in S^{\Lambda} \mid \varphi \in s\}$. Then, we define the canonical valuation function $V^{\Lambda}(p) = \tilde{p}$.

We define the canonical effectivity structure E^{Λ} by induction on b as follows:

- For all b such that $t(b) = 1$ and $C \neq A$, $X \in E^{\Lambda}(s)(C^b)$ iff $\exists \tilde{\varphi} \subseteq X : [C^b]\varphi \in s$. $X \in E^{\Lambda}(s)(A^b)$ iff $\bar{X} \notin E^{\Lambda}(s)(\emptyset^b)$.
- For all b such that $t(b) > 1$ and $C \neq \emptyset$, $X \in E^{\Lambda}(s)(C^b)$ iff $X \in E^{\Lambda}(s)(C^{b'})$ for some $b' < b$ or there are $b_1 \otimes b_2 = b$ such that $\{s' \mid X \in E^{\Lambda}(s')(C^{b_2})\} \in E^{\Lambda}(s)(C^{b_1})$. $X \in E^{\Lambda}(s)(\emptyset^b)$ iff $\bar{X} \notin E^{\Lambda}(s)(A^b)$.

The following property (*) is crucial for the proof:

$$(*) \quad \tilde{\varphi} \in E^{\Lambda}(s)(C^b) \text{ iff } [C^b]\varphi \in s$$

We prove it by induction on the bounds. In the base case, assume that $\tilde{\varphi} \in E^{\Lambda}(s)(C^b)$ for some $t(b) = 1$. For $C \neq A$, $\tilde{\varphi} \in E^{\Lambda}(s)(C^b)$ iff $\exists \tilde{\psi} \subseteq \tilde{\varphi} : [C^b]\psi \in s$ iff $[C^b]\varphi \in s$ (by $\vdash_{\Lambda} \psi \rightarrow \varphi$ and RM). If $C = A$, we have $\tilde{\varphi} \in E^{\Lambda}(s)(A^b)$ iff $\neg\tilde{\varphi} \notin E^{\Lambda}(s)(\emptyset^b)$ iff $\neg[\emptyset^b]\neg\varphi \in s$ (as just proved) iff $[A^b]\varphi \in s$ (by axiom A3).

For the induction step, assume that $\tilde{\varphi} \in E^{\Lambda}(s)(C^b)$ where $t(b) > 1$. For $C \neq \emptyset$, there are two cases to consider. (1) $\tilde{\varphi} \in E^{\Lambda}(s)(C^{b'})$ for some $b' < b$. By the induction hypothesis, we have $[C^{b'}]\varphi \in s$. Then, axiom A5 implies that $[C^b]\varphi \in s$. (2) There are $b_1 \otimes b_2 = b$ such that

$$\{s' \mid \tilde{\varphi} \in E^{\Lambda}(s')(C^{b_2})\} \in E^{\Lambda}(s)(C^{b_1}).$$

By the induction hypothesis, we have $[C^{b_2}]\varphi \in E^{\Lambda}(s)(C^{b_1})$, so $[C^{b_1}][C^{b_2}]\varphi \in s$. Therefore, by axiom A7 $[C^b]\varphi \in s$.

If $C = \emptyset$, we have $\tilde{\varphi} \in E^{\Lambda}(s)(\emptyset^b)$ iff $\neg\tilde{\varphi} \notin E^{\Lambda}(s)(A^b)$ iff $\neg[A^b]\neg\varphi \in s$ (as just proved) iff $[\emptyset^b]\varphi \in s$ (by axiom A3).

The proof that E^{Λ} is RB-playable is straightforward given the property (*), definition of E^{Λ} and the axioms of Λ .

Since we have already shown (*), the truth lemma $(M^{\Lambda}, s \models \varphi \text{ iff } \varphi \in s)$ is also straightforward. \square

As an illustration, we can state a property from [Alechina *et al.*, 2008] that two reasoners can derive an empty clause within resource bounds 4 for memory, 1 for communication, and 5 for time: $\{[1, 2]^{(4,1,5)}\} B_2 \perp$.

5 Conclusions and Further Work

We have proposed a complete and sound logic *RBCL* where we can express costs of (multi-step) strategies and hence coalitional ability under resource bounds in multi-shot games. The logic is expressive enough to formalise decision problems of Coalitional Resource Games [Wooldridge and Dunne, 2006] and properties of resource-bounded communicating reasoners from [Alechina *et al.*, 2008]. *RBCL* is related to both Coalition Logic and ATL. The $[C^b]$ operators in *RBCL* correspond to the $[C^*]$ operator in Extended Coalition Logic which stands for a finite iteration of $[C]$ modalities [Pauly, 2001], or to the $\langle\langle C \rangle\rangle F$ operator of ATL [Alur *et al.*, 2002].

In future work, we plan to investigate automated verification of properties of RBCL using model-checking. The model-checking algorithm for ATL given in [Alur *et al.*, 2002] can be modified to produce a model-checking algorithm for RBCL without increase in complexity. We anticipate that we may also be able to exploit resource bounds to optimise the model-checking algorithm or to use bounded model-checking for some properties.

References

- [Alechina *et al.*, 2008] N. Alechina, B. Logan, Nguyen H. N., and A. Rakib. Verifying time, memory and communication bounds in systems of reasoning agents. In *Proc. 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 736–743. IFAAMAS, 2008.
- [Alur *et al.*, 2002] R. Alur, T. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. of the ACM*, 49(5):672–713, 2002.
- [Goranko, 2001] V. Goranko. Coalition games and alternating temporal logics. In *Proc. 8th Conf. on Theoretical Aspects of Rationality and Knowledge (TARK VIII)*, pages 259–272. Morgan Kaufmann, 2001.
- [Pauly, 2001] M. Pauly. *Logic for Social Software*. Ph.D. thesis, ILLC, University of Amsterdam, 2001.
- [Pauly, 2002] M. Pauly. A modal logic for coalitional power in games. *J. Log. and Comp.*, 12(1):149–166, 2002.
- [Wooldridge and Dunne, 2006] M. Wooldridge and P. E. Dunne. On the computational complexity of coalitional resource games. *Artif. Intel.*, 170(10):835–871, 2006.
- [Wooldridge *et al.*, 2007] M. Wooldridge, T. Ågotnes, P. E. Dunne, and W. van der Hoek. Logic for Automated Mechanism Design – A Progress Report. In *Proc. 22nd AAAI Conf. on Artificial Intelligence*, pages 9–16. AAAI Press, 2007.