An Outline of Parameterised Resource-Bounded ATL

Natasha Alechina
Department of Information and Computing Sciences,
University of Utrecht
n.a.alechina@uu.nl

Stéphane Demri
LSV, ENS de Cachan,
CNRS
demri@lsv.ens-cachan.fr

Brian Logan
School of Computer Science,
University of Nottingham
brian.logan@nottingham.ac.uk

ABSTRACT

It is often advantageous to be able to extract resource requirements in resource logics of strategic ability, rather than to verify whether a fixed resource requirement is sufficient for achieving a goal. We study Parameterised Resource-Bounded Alternating Time Temporal Logic where parameter extraction is possible. We give a parameter extraction algorithm and show that parameter extraction can be done in 2EXPTIME, and this bound is tight. A longer version of this note appeared in AAAI 2020 [2].

1 INTRODUCTION

We show how to compute the resources required by a successful strategy for a coalition of agents where the goal of the coalition is specified as an achievement or non-termination property in temporal logic. For positive formulas (where no modal operator is in the scope of a negation), we compute the set of Pareto optimal resource values, as in games on multi-weighted graphs [5, 6].

2 SYNTAX AND SEMANTICS

The syntax of ParRB±ATL(n,r) is defined relative to the following three sets:

- $Agt = \{a_1, \ldots, a_n\}$ is a set of $n \geq 1$ agents
- $Res = \{res_1, \ldots, res_r\}$ is a set of $r \geq 1$ resource types
- Π denotes a set of propositions, and Var a countably infinite set of variables.

We write ParRB±ATL(n,r) for the logic ParRB±ATL with n agents and r resources, interpreted over models with the same numbers of agents and resources. When talking about all possible n and r, we write ParRB±ATL.

Formulas of ParRB±ATL(n,r) are defined by the following syntax

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

where $p \in \Pi$ is a proposition, $A \subseteq Agt$, and $\mathbf{x}$ is a sequence of variables $x_1, \ldots, x_r$ from Var of length r, intuitively placeholders for amounts of resource 1 to resource r.

The models of the logic ParRB±ATL are resource-bounded concurrent game structures (RB-CGS) introduced in [3].

Definition 2.1. A resource-bounded concurrent game structure is a tuple $M = (Agt, Res, \Pi, S, \pi, Act, d, c, \delta)$ where:

- Agt, Res, Π as above,
- set of states S, assignment $\pi : \Pi \to \wp(S)$, partial transition function are as in CGS;
- Act is a non-empty set of actions which includes $idle$, and $d : S \times Agt \to \wp(Act)$ is a function which assigns to each $s \in S$ a set of actions available to each agent $a \in Agt$. For every $s \in S$ and $a \in Agt$, $idle \in d(s, a)$.
- $c : S \times Agt \times Act \to \mathbb{Z}^r$ is a partial function which maps a state s, and agent a and an action $\alpha \in d(s, a)$ to a vector of integers where the integer in position i indicates consumption or production of resource res_i by the action (positive value for consumption and negative value for production). We stipulate that $c(s, a, idle) = \mathbf{0}$ for all $s \in S$ and $a \in Agt$ where $\mathbf{0} = (0, \ldots, 0) \in \mathbb{N}^r$.

The notational conventions and definitions of joint actions, outcomes, perfect recall strategies, and infinite computations consistent with a strategy are standard for ATL; we refer the reader to [2].

The cost of a joint action σ by agents in coalition A is defined as:

$$cost(s, \sigma) = \sum_{a \in A} c(s, a, \sigma_a)$$

(component-wise sum of vectors of costs for all individual actions in σ).

Given a bound $\mathbf{b} \in \mathbb{N}^r$, a computation $\lambda \in out(s, F_A)$ (a computation starting in state s which is consistent with A's strategy F_A) is $\mathbf{b}$-consistent with F_A iff, for every $i \geq 0$,

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

(the cost of any prefix of the computation is (component-wise, for each resource value) below the bound $\mathbf{b}$). Note that this definition implies that if agents start with resource allocation $\mathbf{b}$, then on every prefix of the computation the amount of resources they have is never below $\mathbf{0}$. Furthermore, the constraint on $\mathbf{b}$-consistent computations involves only action costs for agents from the coalition A (and not on the agents outside A).

F_A is a $\mathbf{b}$-strategy iff for any state s, all computations consistent with this strategy starting in s are $\mathbf{b}$-consistent.

Given an RB-CGS M, a state s of M, the truth of a ParRB±ATL formula φ with respect to M, s and assignment v from the set of variables to $\mathbb{N}$ is defined inductively on the structure of φ as follows:

- $M, s, v \models p$ iff $s \in \pi(p)$;
- $M, s, v \models \neg\phi$ iff $M, s, v \not\models \phi$;
- $M, s, v \models \phi \vee \psi$ iff $M, s, v \models \phi$ or $M, s, v \models \psi$;
- $M, s, v \models \langle\!\langle A^{\mathbf{x}} \rangle\!\rangle \bigcirc \phi$ iff there exists a $v(\mathbf{x})$-strategy F_A such that for all $\lambda \in out(s, F_A)$: $M, \lambda[1], v \models \phi$;
- $M, s, v \models \langle\!\langle A^{\mathbf{x}} \rangle\!\rangle \Box \phi$ iff there exists a $v(\mathbf{x})$-strategy F_A such that for all $\lambda \in out(s, F_A)$ and $i \geq 0$: $M, \lambda[i], v \models \phi$;
- $M, s, v \models \langle\!\langle A^{\mathbf{x}} \rangle\!\rangle \phi \,\mathcal{U}\, \psi$ iff there exists a $v(\mathbf{x})$-strategy F_A such that for all $\lambda \in out(s, F_A)$, there exists $i \geq 0$ such that $M, \lambda[i], v \models \psi$ and $M, \lambda[j], v \models \phi$ for all $j \in \{0, \ldots, i-1\}$.

Submitted to LAMAS 2020, 2020

https://doi.org/doi

3 PROBLEM AND COMPLEXITY RESULTS

We show how to characterise the resource values for which a formula of ParRB±ATL holds. For positive formulas (possibly with nested modalities) this can be used to extract the Pareto optimal values of resources required to satisfy the formula.

Definition 3.1. The following problem is the *function form of the model-checking problem* for ParRB±ATL(n,r).

Input: $n, r \geq 1$ (in unary), a ParRB±ATL(n,r) formula φ, a finite model M, and a state s.

Output: Compute a formula γ describing the set of assignments v such that $M, s, v \models \varphi$, where γ is of the following form (for $b \in \mathbb{N}$):

$$\gamma := x \geq b \mid \top \mid \bot \mid \neg\gamma \mid \gamma \wedge \gamma \mid \gamma \vee \gamma$$

Note that such a formula always exists. For each subformula starting with a coalition modality of the form $\langle\langle A^{\mathbf{x}} \rangle\rangle \psi$, the set of assignments satisfying it can be described by a finite disjunction of the form $\mathbf{x} \geq \mathbf{b}_1 \vee \ldots \vee \mathbf{x} \geq \mathbf{b}_k$. This is because there are finitely many Pareto optimal values $\mathbf{b}$ for $\mathbf{x}$ (such that a $\mathbf{b}$-consistent strategy satisfying ψ exists, and there is no $\mathbf{b}' < \mathbf{b}$ such that a $\mathbf{b}'$-consistent strategy exists). If there is a $\mathbf{b}$-consistent strategy satisfying ψ, then for every resource allocation $\mathbf{b}' \geq \mathbf{b}$, the same strategy is also $\mathbf{b}'$-consistent, hence all assignments satisfying $\langle\langle A^{\mathbf{x}} \rangle\rangle \psi$ can be described by a disjunction of the form $\mathbf{x} \geq \mathbf{b}_1 \vee \ldots \vee \mathbf{x} \geq \mathbf{b}_k$, where $\mathbf{b}_1, \ldots, \mathbf{b}_k$ are Pareto-optimal values for $\mathbf{x}$. Nestings of coalition modalities give rise to conjunctions of constraints, and boolean combinations to corresponding boolean combinations of constraints (see Algorithm 1).

THEOREM 3.2. *The function form of the model-checking problem for ParRB±ATL can be solved in 2EXPTIME.*

2EXPTIME lower bound is an easy consequence of 2EXPTIME-hardness of the model-checking problem for RB±ATL (where instead of variables $\mathbf{x}$ all modalities have concrete values $\mathbf{b}$), the result shown in [1]. The model-checking problem for RB±ATL can be solved using the function form of the model-checking problem for ParRB±ATL as follows: compute the constraint formula γ, and then check in linear time whether the concrete bounds in the RB±ATL formula satisfy γ.

To prove the upper bound, we need the notion of a witness for a formula starting with a strategic modality $\langle\langle A \rangle\rangle$. It is a partial strategy which can be represented as a finite tree where nodes are pairs of a state and a resource amount coalition A has in that state. The leaves of the tree represent either achivement of the goal (the second argument of Until modality), or a state on a non-consuming loop for a Box modality. Note that for Until (achievement) formulas such a partial strategy can be expanded to a complete strategy for the same initial resource allocation by choosing the *idle* action after the goal formula is achieved.

A *witness* w for $M, s, \mathbf{x} \mapsto \mathbf{b} \models \langle\langle A^{\mathbf{x}} \rangle\rangle \psi_1 \,\mathcal{U}\, \psi_2$ ($\langle\langle A^{\mathbf{x}} \rangle\rangle \Box \phi$) is a finite tree with root $(s, \mathbf{b})$ where

- all nodes of the tree are of the form $(s', \mathbf{b}')$ where s' is a state (resulting from executing the joint action in the state of the parent node) and $\mathbf{b}'$ the vector of resource amounts A have in s' as a result, and $\mathbf{b}' \geq 0$,
- in all leaf nodes $(s', \mathbf{b}')$ of the tree, s' satisfies ψ_2
- in all non-leaf nodes $(s', \mathbf{b}')$ of the tree, s' satisfies ψ_1
- (in all nodes, s' satisfies ϕ)

$\|w\|$ is the norm of the witness (the maximal value in any $\mathbf{b}'$). max_M is the largest absolute value of an action cost in M.

The crux of establishing a 2EXPTIME upper bound for the ParRB±ATL model-checking problem is showing that there is a maximal value of the amount of resource needed to reach a particular state or to enter and execute a non-resource-consuming loop, and this bound depends only on the model and not on the property to be checked. The proof of theorem builds on the results of [6] to show that if a witness for $\langle\langle A^{\mathbf{x}} \rangle\rangle \psi_1 \,\mathcal{U}\, \psi_2$ ($\langle\langle A^{\mathbf{x}} \rangle\rangle \Box \phi$) being true in M, s exists, then there exists a witness w with $\|w\| \leq B_M$ where $B_M = (4(|S| + 1) \cdot max_M)^{2(r+3)^3}$.

4 MODEL-CHECKING ALGORITHM

Algorithm 1, given a model M and a ParRB±ATL formula φ with variables $x_1, \ldots, x_n$, returns a set of pairs $\|\phi\|_M$ of the form (s, γ) where s is a state and γ a constraint formula defining the set of assignments v to $x_1, \ldots, x_n$ such that $M, s, v \models \phi$. This set is finite because the set of states is finite, and γ is a finite formula (as argued after the definition of the model-checking problem; essentially this is due to the finite number of Pareto values for variables in each subformula starting with a coalition modality, and this in turn is due to the fixed number of resource types, and monotonicity of coalition modalities).

Algorithm 1 Computing $\|\phi_0\|$

function PARRB±ATL-MC(M, ϕ_0)
  for $\phi' \in Sub(\phi_0)$ **do**
    case $\phi' = p$
      $\|p\|_M \leftarrow \{(s, \top) : s \in \pi(p)\} \cup \{(s, \bot) : s \notin \pi(p)\}$
    case $\phi' = \neg\psi$
      $\|\neg\psi\|_M \leftarrow \{(s, \neg\gamma) : (s, \gamma) \in \|\psi\|_M\}$
    case $\phi' = \psi_1 \wedge \psi_2$
      $\|\psi_1 \wedge \psi_2\|_M \leftarrow \{(s, \gamma_1 \wedge \gamma_2) : (s, \gamma_1) \in \|\psi_1\|_M \wedge (s, \gamma_2) \in \|\psi_2\|_M\}$
    case $\phi' = \psi_1 \vee \psi_2$
      $\|\psi_1 \vee \psi_2\|_M \leftarrow \{(s, \gamma_1 \vee \gamma_2) : (s, \gamma_1) \in \|\psi_1\|_M \wedge (s, \gamma_2) \in \|\psi_2\|_M\}$
    case $\phi' = \langle\langle A^{\mathbf{x}} \rangle\rangle \bigcirc \psi$
      $\|\langle\langle A^{\mathbf{x}} \rangle\rangle \bigcirc \psi\|_M \leftarrow \{(s, \gamma) : \gamma = \bigvee_{\sigma \in D_A(s) \wedge out(s,\sigma) \subseteq \|\psi\|} \mathbf{x} \geq cost(\sigma, s)\}$
    case $\phi' = \langle\langle A^{\mathbf{x}} \rangle\rangle \psi_1 \,\mathcal{U}\, \psi_2$
      $\|\langle\langle A^{\mathbf{x}} \rangle\rangle \psi_1 \,\mathcal{U}\, \psi_2\| \leftarrow \{(s, \gamma) : \gamma = \bigvee_{w \in W(M, s, \langle\langle A^{\mathbf{x}} \rangle\rangle \psi_1 \mathcal{U} \psi_2)} \gamma(w)\}$
    case $\phi' = \langle\langle A^{\mathbf{x}} \rangle\rangle \Box \psi$
      $\|\langle\langle A^{\mathbf{x}} \rangle\rangle \Box \psi\| \leftarrow \{(s, \gamma) : \gamma = \bigvee_{w \in W(M, s, \langle\langle A^{\mathbf{x}} \rangle\rangle \Box \psi)} \gamma(w)\}$
  return $\|\phi_0\|_M$

In the algorithm, $W(M, s, \langle\langle A^{\mathbf{x}} \rangle\rangle \psi_1 \,\mathcal{U}\, \psi_2)$ ($W(M, s, \langle\langle A^{\mathbf{x}} \rangle\rangle \Box \psi)$) is the set of all possible non-dominated (minimal) witnesses w for $\langle\langle A^{\mathbf{x}} \rangle\rangle \psi_1 \,\mathcal{U}\, \psi_2$ ($W(M, s, \langle\langle A^{\mathbf{x}} \rangle\rangle \Box \psi)$) being true in M, s with $\|w\| \leq B_M$.[1] The definition of a witness is modified for the general case,

[1] A witness is non-dominated if there is no other witness where the largest resource vector anywhere in the tree (the norm) is strictly less, and no other norm-minimal witness with a strictly smaller value at the root exists.

replacing the requirement that ψ_1, ψ_2 or ϕ is true in s' with a requirement that s' satisfies ψ_1, ψ_2 or ϕ for some assignment to $\mathbf{x}$ (that s' is associated with a satisfiable constraint formula). Intuitively, if s' satisfies a propositional variable p, then the corresponding constraint is $\top$, and if s' satisfies $\neg\langle\langle A^{\mathbf{x}}\rangle\rangle \psi_1 \mathcal{U} \psi_2$ for values of $\mathbf{x}$ less than $\mathbf{b}$, the corresponding constraint is $\mathbf{x} < \mathbf{b}$. $\gamma(w)$ is a formula describing constraints on variables in the witness w.

The complexity result follows because the set of non-dominated witnesses $W(M, s, \phi')$ for each subformula ϕ' of ϕ_0 can be computed in AEXSPACE (by an alternating Turing machine [4] using exponential space). Hence the whole algorithm can be executed in AEXSPACE which is equivalent to 2EXPTIME.

If ϕ is a positive formula then the corresponding constraint formula γ only has subformulas of the form $\mathbf{x} \geq \mathbf{b}$. Since the constraint formula describes non-dominated witnesses, it specifies the *optimal* values of resource parameters $\mathbf{b}$ (the Pareto frontier).

5 CONCLUSION

Algorithm 1 returns a set of states satisfying a ParRB±ATL formula together with a constraint formula describing the set of assignments to resource variables for which the ParRB±ATL formula is true in that state. The algorithm can be easily modified to return the witness (finite representation of the strategy). The complexity of the algorithm is very high (2EXPTIME). However, only the number of resource types and the *log* of the largest cost in the model ($log\ max_M$) are in the exponent; the running time is polynomial in the number of states and transitions in the model.

In future work we plan to implement the algorithm.

Acknowledgements. We thank LAMAS 2020 reviewers for their thoughtful, detailed and constructive reviews.

REFERENCES

[1] Natasha Alechina, Nils Bulling, Stéphane Demri, and Brian Logan. 2018. On the complexity of resource-bounded logics. *Theoretical Computer Science* 750 (2018), 69–100. https://doi.org/10.1016/j.tcs.2018.01.019

[2] Natasha Alechina, Stéphane Demri, and Brian Logan. 2020. Parameterised Resource-Bounded ATL. In *Proceedings of the Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI-20)*. AAAI, AAAI Press, New York, USA.

[3] Natasha Alechina, Brian Logan, Nguyen Hoang Nga, and Abdur Rakib. 2010. Resource-bounded alternating-time temporal logic. In *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, Wiebe van der Hoek, Gal Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen (Eds.). IFAAMAS, IFAAMAS, 481–488.

[4] Ashok K. Chandra and Larry J. Stockmeyer. 1976. Alternation. In *17th Annual Symposium on Foundations of Computer Science*. 98–108.

[5] J.B. Courtois and S. Schmitz. 2014. Alternating Vector Addition Systems with States. In *39th International Symposium on Mathematical Foundations of Computer Science (MFCS'14) (Lecture Notes in Computer Science)*, Vol. 8634. Springer, 220–231.

[6] M. Jurdzinski, R. Lazić, and S. Schmitz. 2015. Fixed-Dimensional Energy Games are in Pseudo-Polynomial Time. In *Proceedings of ICALP 2015*. 260–272.