Minimisation of ATL* models

Serenella Cerrito[1] and Amélie David[2]

[1] IBISC, Université Evry Val d'Essonne and Paris-Saclay, France
serena.cerrito@ibisc.univ-evry.fr,
[2] Université Paris-Descartes, France
amelie.david@parisdescartes.fr

Abstract. The aim of this work is to provide a general method to minimize the size (number of states) of a model $\mathcal{M}$ of an ATL^* formula. Our approach is founded on the notion of alternating bisimulation: given a model $\mathcal{M}$, it is transformed in a stepwise manner into a new model $\mathcal{M}$' minimal with respect to bisimulation. The method has been implemented and will be integrated to the prover TATL, that constructively decides satifiability of an ATL^* formula by building a tableau from which, when open, models of the input formula can be extracted.

Keywords: Alternating-time Temporal Logic, bisimulation, model minimization, tableaux.

1 Introduction

The Alternating-time temporal logic ATL^* has been introduced in [AHK02] and proposed as a logical framework for the specification and the verification of properties of open systems, that is systems interacting with an environment whose behaviour is unknown or only partially known. The logic ATL^* can be seen as a multi-agent extension of the branching time temporal logic CTL^* where the path quantifiers are generalized to "strategic quantifiers", indexed with coalitions of agents A, ranging existentially over collective strategies of A and then universally over all paths (computations) coherent with the selected collective strategy. The language of ATL^* allows the expression of statements of the type *"Coalition A has a collective strategy to guarantee the satisfaction of the objective Ψ no matter what its opponents do"*, and can therefore model the interaction of an open system with an environment by setting the environment to be the system opponent.

The semantics of ATL^* is based on the notion of concurrent game models (CGMs), a generalisation of labelled transition systems to the multi-agent framework where an edge connecting two states is labelled by a vector describing the synchronous actions of all the agents, rather then by the action of a single agent. The aim of this work is to provide a general method to minimize the size (number of states) of a model $\mathcal{M}$ of an ATL^* formula.

Independently from the specific logic of interest, to get minimal models is useful for several tasks: hardware and software verification, fault analysis, and common sense reasoning. Several different criteria of minimality have been studied in the literature. In the case of first order classical logic, some works minimize the domain (see for instance [Hin88,Lor94], while others minimize the interpretation either of a certain set of predicates (see for instance [McC87] or of all the predicates (see for instance [BY00,Nie96,GHS01,HFK00]).

These minimality criteria can be applied to modal logics, too. Minimal model generation where certain predicates are minimal has been mostly studied in the context of non-monotonic operators and non-monotonic semantics (see for instance [GGOP08,GH09,BLW09]). In the case of modal logics, however, it is quite natural to to adopt minimality criteria founded on the notion of bisimulation. The work [PS14] presents terminating procedures for the generation of models that are minimal for a given notion of *subset-simulation* for the propositional modal logic K and all combinations of it extensions with the axioms T, B, D, 4 and 5. Roughly, what is minimized there is not the number of worlds, but the number of propositions holding at worlds.

In the specific case of temporal logics the emphasis is rather on the reduction of the size of the state space. This is crucial if the considered temporal logic has to be used to model systems whose properties need to be model-checked. In the case of CTL and CTL^*, having as semantics labelled transition systems (LTS), models are minimized with respect to bisimulation by using coarsest partition algorithms refining step by step an initial partition of the set of states of a given LTS [LIS12,KS90,PT87].

Our work is inspired by the above mentioned partition-refinement approach for LTS but treats the more complex case of ATL^* models, namely CGMs. We rewrite a CGM $\mathcal{M}$ into a bisimilar smaller model by using the definition of alternating bisimulation, that is specific to ATL^* [DGL16,ÅGJ07].

The intended application is the synthesis of ATL^* models from formal specifications by means of the software TATL, available on line via a dynamic web page [Dav]. Up to our knowledge, TATL is the only existing running system that decides the satisfiability of an ATL^* formula (and by means of a trivial preliminary rewriting also of CTL^* formulae). TATL constructively decides the satisfiability of a given ATL^* formula ϕ by exhibiting a tableau for ϕ [CDG14,Dav15]. A tableau for ϕ is built by analysing the formula and producing states of the candidate models, so as to obtain a finite graph. When the final tableau is open, it is a non-empty labelled graph representing a graph of CGMs satisfying ϕ at some initial state. The completeness proof (with respect to unsatifiability), being constructive, provides a procedure to build a model of ϕ from an open tableau [CDG14,Dav15]. Such a procedure, however, can generate a model that has an unnecessarily large number of states, because eventualities are sequentially treated to assure their realizability: eventualities that might be simultaneously realized are systematically realized one after the other. To minimize such a model is important, for instance, for the purpose of model synthesis from a formal specification written in ATL^*: CGMs that contain an unnecessary great number of states are difficult to grasp and expensive to treat (for instance to model check additional properties).

It is worthwhile observing that to rewrite a given model $\mathcal{M}$ of a formula ϕ into a model $\mathcal{M}$' that is minimal with respect to alternating bisimulation does not necessaily mean to get a model of ϕ having the minimum number of states. To illustrate this point, let us consider a very simple example.

Example 1. Let ϕ be $\langle\langle 1\rangle\rangle \bigcirc p$, stating that agent 1 can assure that p holds at a successor state. Let's assume that this agent can perform only one action at each state. Take $\mathcal{M}_1$ to have two states, 1 and 2, where 2 is the only successor of 1, 1 is the only successor of 2, and p is false at 1 but true at 2. Clearly $\mathcal{M}_1$ satisfies ϕ at state 1. Now, take $\mathcal{M}_2$ to be a model having 3 states, A, B and C, where B is the only successor of A, C is the only successor of B, A is the only successor of C and the only state not satisfying p is A. Obviously ϕ keeps true at state A. The application of our minimisation procedure to $\mathcal{M}_2$ outputs $\mathcal{M}_2$ itself, not $\mathcal{M}_1$. The reason is that any state s' of the output model must satisfy exactly the same formulae as s, where s is a bisimilar state of the input model. In $\mathcal{M}_1$, state 1 satisfies $\neg p \wedge \langle\langle 1\rangle\rangle(\bigcirc p \wedge \bigcirc\bigcirc \neg p)$ while in $\mathcal{M}_2$ state A satisfies $\neg p \wedge \langle\langle 1\rangle\rangle(\bigcirc p \wedge \bigcirc\bigcirc p)$, thus 1 and A are not bisimilar.

It is worthwhile noticing, however, that such an unnatural model of ϕ as $\mathcal{M}_2$ would not be generated by the tableau procedure for ATL^* having input $\phi = \langle\langle 1\rangle\rangle \bigcirc p$. In general, tableau construction analyses the input formula and produces tableau states (states of a candidate model) only when they are needed.

The outline of this work is the following. In Section 2 we recall some background definitions. Section 3 is the core of the paper and provides our minimisation algorithm and its foundations. Section 4 briefly discuss the implementation (ongoing work). In Section 5 we give some detailed proofs missing in the paper. Finally, we conclude and we sketch some lines of future work.

2 Preliminaries

We recall here some standard definitions about ATL^*.

Definition 1 (Concurrent Game Model). *Given a set of* atomic propositions P*, a CGM (Concurrent Game Model) is a 5-tuple*

$$\mathcal{M} = \langle \mathbb{A}, \mathbb{S}, \{\mathsf{Act}_a\}_{a\in\mathbb{A}}, \mathsf{out}, \mathsf{L}\rangle$$

such that:
- $\mathbb{A} = \{1, ..., k\}$ *is a finite non-empty set of agents;*
- $\mathbb{S}$ *is a non-empty set of states;*
- *For each* $a \in \mathbb{A}$, Act_a *is a non-empty set of actions. If* $A \subseteq \mathbb{A}$*, then* A *is a coalition of agents,* $\mathsf{Act}_A = \Pi_{a\in A}\mathsf{Act}_a$ *and* σ_A *denotes a tuple from* Act_A *(an action of the coalition* A*). If* $a \in A$, $\sigma_A(a)$ *means the action of the agent* a *in the tuple of actions* σ_A*;*
- *If* $a \in \mathbb{A}$ *and* $s \in \mathbb{S}$, $\mathsf{act}_a(s)$ *maps state* s *to a non-empty subset of* Act_a*: the set of the available actions for agent* a *at state* s*. Similarly,* $\mathsf{act}_A(s)$ *maps* s *to a non-empty subset of* Act_A *(the set of actions available to coalition* A *at state*

s). That is: $\mathsf{act}_A(s) = \Pi_{a \in A} \mathsf{act}_a(s)$*;*
• out *is a transition function, associating to each* $s \in S$ *and each* $\sigma_{\mathbb{A}} \in \mathsf{act}_{\mathbb{A}}(s)$ *a state* $\mathsf{out}(s, \sigma_{\mathbb{A}}) \in \mathbb{S}$*: the state reached when each* $a \in \mathbb{A}$ *does the action* σ_a *at* s*;*
• L *is a labelling function* $\mathsf{L} : \mathbb{S} \to \mathcal{P}(P)$*, associating to each state* s *the set of propositions holding at* s*.*

It is worthwhile observing that the above definition does not require the set $\mathbb{S}$ to be finite. In our intended application, however, where models are constructed out of open finite tableaux, it will always be finite.

The following figure shows a well known simple example of CGM, modelling the situation of a carriage pushed by two robots (the agents), one at its left and one at its right, in two opposite directions.

Fig. 1. A simple example of CGM

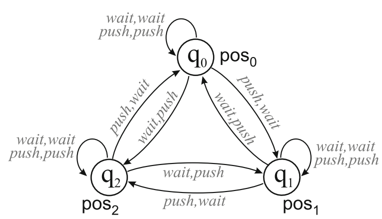

An action of a coalition A will also be called *A-move*; when A is $\mathbb{A}$ we will call it *global move*.

Below, $p \in P$ and A is a *coalition* of agents.

Definition 2 **(**ATL^* **syntax).**

State formulae : $\psi := p \mid (\neg\psi) \mid (\psi \wedge \psi) \mid (\langle\!\langle A \rangle\!\rangle \Phi)$
Path formulae : $\Phi := \psi \mid (\neg\Phi) \mid (\Phi \wedge \Phi) \mid (\bigcirc\Phi) \mid (\Box\Phi) \mid (\Phi \mathsf{U} \Phi)$

It is worthwhile observing that any ATL^* state formula is also an ATL^* path formula, while the converse is false. State formulae will always be noted by lower case Greek letters, and path formulae by upper case Greek letters. Unless explicitly stated otherwise, in the sequel by ATL^* formula we mean an ATL^* state formula.

ATL is the syntactical fragment of ATL^* obeying to the constraint that any temporal operator in a formula is prefixed by a path quantifier $\langle\!\langle A \rangle\!\rangle$, analogously to CTL w.r.t. CTL^*. Hence any ATL formula is a state formula.

The semantics for ATL^* is based on the notions of concurrent game model, *play* and *strategy*.

A play λ in a CGM $\mathcal{M}$ is an infinite sequence of elements of $\mathbb{S}$: $s_0, s_1, s_2, ...$ such that for every $i \geq 0$, there is a global move $\sigma_{\mathbb{A}} \in \mathsf{act}_{\mathbb{A}}(s_i)$ such that

$\mathsf{out}(s_i, \sigma_{\mathbb{A}}) = s_{i+1}$. Given a play λ, we denote by λ_0 its initial state, by λ_i its $(i+1)$th state, by $\lambda_{\leq i}$ the prefix $\lambda_0 ... \lambda_i$ of λ and by $\lambda_{\geq i}$ the suffix $\lambda_i \lambda_{i+1} ...$ of λ. Given a prefix $\lambda_{\leq i} : \lambda_0 ... \lambda_i$, we say that it has length $i+1$ and write $|\lambda_{\leq i}| = i+1$. An empty prefix has length 0. A (non-empty) *history* at state s is a finite prefix of a play ending with s. We denote by $\mathsf{Plays}_{\mathcal{M}}$ and $\mathsf{Hist}_{\mathcal{M}}$ respectively the set of plays and set of histories in a CGM $\mathcal{M}$.

Given a coalition $A \subseteq \mathbb{A}$ of agents, a *perfect recall strategy* F_A is a function which maps each element $\lambda = \lambda_0 ... \lambda_\ell$ of $\mathsf{Hist}_{\mathcal{M}}$ to an A-move σ_A belonging to $\mathsf{act}_A(\lambda_\ell)$ (the set of actions available to A at state λ_ℓ). Whenever F_A depends only on the state λ_ℓ the strategy is said to be *positional*. In the rest of the paper we always consider perfect recall strategies.

For any coalition A, a global action $\sigma_{\mathbb{A}}$ *extends* an A-move σ_A whenever for each agent $a \in A$, $\sigma_A(a) = \sigma_{\mathbb{A}}(a)$. Let σ_A be an A-move; the notation $\mathsf{Out}(s, \sigma_A)$ denotes the set of states $\mathsf{out}(s, \sigma_{\mathbb{A}})$ where $\sigma_{\mathbb{A}}$ is any global vector extending σ_A. Intuitively, $\mathsf{Out}(s, \sigma_A)$ denotes the set of the states that are successors of s when the coalitions A plays at s the A-move σ_A and the other agents play no matter which move.

A play $\lambda = \lambda_0, \lambda_1, ...$ is said to be *coherent with a strategy* F_A if and only if for each $j \geq 0$, $\lambda_{j+1} \in \mathsf{Out}(\lambda_j, \sigma_A)$, where σ_A is the A-move chosen by F_A at state λ_i.

The notion $\mathcal{M}$ *satisfies the formula* Φ *at state* s, noted $\mathcal{M}, s \models \Phi$, is defined by induction on ϕ as follows (omitting the obvious boolean cases):

- $\mathcal{M}, s \models p$ iff $p \in L(s)$, for any proposition $p \in \mathbb{P}$;
- $\mathcal{M}, s \models \langle\langle A \rangle\rangle \Phi$ iff there exists an A-strategy F_A such that, for all plays λ starting at s and coherent with the strategy F_A, $\mathcal{M}, \lambda \models \Phi$;
- $\mathcal{M}, \lambda \models \varphi$ iff $\mathcal{M}, \lambda_0 \models \varphi$;
- $\mathcal{M}, \lambda \models \bigcirc \Phi$ iff $\mathcal{M}, \lambda_{\geq 1} \models \Phi$;
- $\mathcal{M}, \lambda \models \Box \Phi$ iff $\mathcal{M}, \lambda_{\geq i} \models \Phi$ for all $i \geq 0$;
- $\mathcal{M}, \lambda \models \Phi \mathsf{U} \Psi$ iff there exists an $i \geq 0$ where $\mathcal{M}, \lambda_{\geq i} \models \Psi$ and for all $0 \leq j < i$, $\mathcal{M}, \lambda_{\geq j} \models \Phi$.

Given a CGM $\mathcal{M}$ and a formula ϕ, we say that $\mathcal{M}$ *satisfies* ϕ whenever there is a state s such that $\mathcal{M}, s \models \phi$; then we also say that $\mathcal{M}$ *satisfies* ϕ *at* s and that $\mathcal{M}$ *is a model of* ϕ.

The works [ÅGJ07,DGL16] define a notion of bisimulation appropriate to CGMs and analogous to the notion of bisimulation for transition systems (see, for instance, [LIS12]).

Definition 3 (Alternating Bisimulation [ÅGJ07,DGL16]). *Let* $\mathcal{M}_1 = \langle \mathbb{A}, \mathbb{S}, \{\mathsf{Act}_a\}_{a \in \mathbb{A}}, \mathsf{out}, \mathsf{L} \rangle$ *and* $\mathcal{M}_2 = \langle \mathbb{A}, \mathbb{S}', \{\mathsf{Act}'_a\}_{a \in \mathbb{A}}, \mathsf{out}', \mathsf{L}' \rangle$ *be two CGMs over the same set of atomic propositions and over the same set of agents.*

- *Let* A *be a coalition. A relation* $\beta \subseteq \mathbb{S} \times \mathbb{S}'$ *is an alternating* A*-bisimulation between* $\mathcal{M}_1$ *and* $\mathcal{M}_2$ *iff for all* $s_1 \in \mathbb{S}$ *and* $s_2 \in \mathbb{S}'$, $s_1 \beta s_2$ *implies that the following hold:*
  1 ***Local Harmony***. $\mathsf{L}(s_1) = \mathsf{L}'(s_2)$*;*

2 ***Forth***. *For any* $\alpha_A \in \mathsf{act}_A(s_1)$ *there is an* $\alpha'_A \in \mathsf{act}'_A(s_2)$ *such that for any* $t_2 \in \mathsf{Out}(s_2, \alpha'_A)$ *there exists* $t_1 \in \mathsf{Out}(s_1, \alpha_A)$ *such that* $t_1 \beta t_2$*;*

2 ***Back***. *For any* $\alpha_A \in \mathsf{act}_A(s_2)$ *there is an* $\alpha'_A \in \mathsf{act}'_A(s_1)$ *such that for any* $t_3 \in \mathsf{Out}(s_1, \alpha'_A)$ *there exists* $t_4 \in \mathsf{Out}(s_2, \alpha_A)$ *such that* $t_3 \beta t_4$*.*

- *When* β *is is an alternating* A*-bisimulation between* $\mathcal{M}_1$ *and* $\mathcal{M}_2$*, we note* $\mathcal{M}_1 \overset{\beta}{\rightleftarrows}_A \mathcal{M}_2$*;*
- *If* β *is an alternating* A*-bisimulation between* $\mathcal{M}_1$ *and* $\mathcal{M}_2$ *for every coalition* $A \subseteq \mathbb{A}$*, then* β *is a full alternating bisimulation and we note:* $\mathcal{M}_1 \overset{\beta}{\rightleftarrows} \mathcal{M}_2$*;*
- *When* β *is a full bisimulation between* $\mathcal{M}_1$ *and* $\mathcal{M}_2$*,* β *is total on* S *and its inverse is total on* S'*, then it is a global alternating bisimulation between* $\mathcal{M}_1$ *and* $\mathcal{M}_2$*. The models* $\mathcal{M}_1$ *and* $\mathcal{M}_2$ *are said to be bisimilar when such a relation* β *exists.*

Figure 2, borrowed from [DGL16], illustrates the above definition.

Fig. 2. An alternating bisimulation between two CGMs

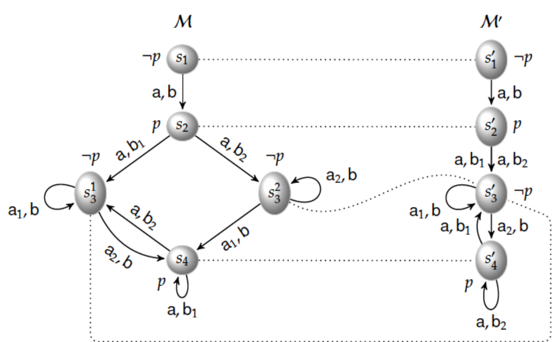

Remark 1. A full alternating bisimulation β is a fixpoint solution of the equation $X = E(X)$ where a value of X is a subset of $\mathbb{S} \times \mathbb{S}'$ such that if $\langle q, q'_2 \rangle \in X$ then $\mathsf{L}(q) = \mathsf{L}'(q')$ and, for any relation $r \subseteq \mathbb{S} \times \mathbb{S}'$, $\langle s_1, s_2 \rangle \in E(r)$ if and only if: $\langle s_1, s_2 \rangle \in r$, $\mathsf{L}(s_1) = \mathsf{L}'(s_2)$, and for every coalition A two conditions hold: (i) for any $\alpha_A \in \mathsf{act}_A(s_1)$ there is an $\alpha'_A \in \mathsf{act}'_A(s_2)$ such that for any $t_2 \in \mathsf{Out}(s_2, \alpha'_A)$ there exists $t_1 \in \mathsf{Out}(s_1, \alpha_A)$ such that $t_1 r t_2$, and (ii) for any $\alpha_A \in \mathsf{act}_A(s_2)$ there is an $\alpha'_A \in \mathsf{act}'_A(s_1)$ such that for any $t_3 \in \mathsf{Out}(s_1, \alpha'_A)$ there exists $t_4 \in \mathsf{Out}(s_2, \alpha_A)$ such that $t_3 r t_4$.

Observe also that, given a CGM having set of states $\mathbb{S}$, X may be a subset of $\mathbb{S} \times \mathbb{S}$ (*i.e* we can have $\mathbb{S} = \mathbb{S}'$).

Remark 1 will be useful in the sequel, to understand how our approach to minimization of a model constructs a maximal fixed point of the above equation in a stepwise manner.

It is also worthwhile observing that bisimilarity between CGMs is reflexive, symmetric and transitive, i.e. is an equivalence relation.

The following theorem extends to the case of ATL^* and perfect recall strategies a result presented in [ÅGJ07,DGL16] for ATL and positional strategies.

Theorem 1. *Let $\mathcal{M}$ and $\mathcal{M}'$ be two CGM.*

1. *If $\mathcal{M} \overset{\beta}{\rightleftarrows}_A \mathcal{M}'$ and $s_1 \beta s_2$, then, for any ATL^* (state) formula ϕ such that A is the only coalition occurring in ϕ, $\mathcal{M}, s_1 \models \phi$ iff $\mathcal{M}', s_2 \models \phi$.*
2. *If $\mathcal{M} \overset{\beta}{\rightleftarrows} \mathcal{M}'$ and $s_1 \beta s_2$, then, for any (state) ATL^* formula ϕ, $\mathcal{M}, s_1 \models \phi$ iff $\mathcal{M}', s_2 \models \phi$.*

The detailed proof of Theorem 1 is given in Section 5. The key idea is that if β is a full alternating bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ and F_A is a strategy for a coalition A in $\mathcal{M}$ then F_A can be simulated in $\mathcal{M}'$ by exploiting the existence of β. As a consequence of Theorem 1, if $\mathcal{M}$ and $\mathcal{M}'$ are bisimilar then, for any (state) formula ϕ, $\mathcal{M}$ is a model of ϕ if and only if $\mathcal{M}'$ is a model of ϕ.

3 Model Minimization

Our approach to the minimization of a model $\mathcal{M}$ satisfying a given formula Φ consists in rewriting it into the smallest bisimilar model in a stepwise manner. The definitions and results that follow are the foundations of our procedure.

3.1 Quotient models

Given a partition $P = \{C_1, ..., C_k\}$ of the set of the states of a CGM, we will say that each set C_i is a *cluster* of the partition P.

Definition 4 (Harmonious partition). *A harmonious partition P of a CGM $\mathcal{M}$ is a partition of the set of states of $\mathcal{M}$ such that for each cluster C of P, if $s, s' \in C$ then $\mathsf{L}(s) = \mathsf{L}(s')$.*

Given a CGM, a state s, a coalition A and a move σ_A available for A at s, we say that a state t is *reachable* from s via σ_A if $t \in \mathsf{Out}(s, \sigma_A)$, i.e. there is a global move $\sigma_{\mathbb{A}}$ extending σ_A such that $t = \mathsf{out}(s, \sigma_{\mathbb{A}})$.

Definition 5 (Behavioural equivalence of states w.r.t. a partition). *Let P be a harmonious partition of a CGM and let s and t be two states such that $\mathsf{L}(s) = \mathsf{L}(t)$.*

- *Let A be a coalition. We say that s and t are (behaviourally) A-equivalent w.r.t. P, and we note $s \equiv_{P^A} t$, when :*

  - *Given any action $\sigma_A \in \mathsf{act}_A(s)$, there is an action $\sigma'_A \in \mathsf{act}_A(t)$ such that the set of clusters of states that are reachable from t via σ'_A is a subset of the set of clusters of states that are reachable from s via σ_A.*
  - *Given any action $\sigma_A \in \mathsf{act}_A(t)$, there is an action $\sigma'_A \in \mathsf{act}_A(s)$ such that the set of clusters of states that are reachable from s via σ'_A is a subset of the set of clusters of states that are reachable from t via σ_A.*
- *We say that s and t are (behaviourally) equivalent w.r.t. P, and we note $s \equiv_P t$, when $s \equiv_{P^A} t$ for each coalition A.*

It is worthwhile observing that $\equiv_{P^A}$ (resp. $\equiv_P$) is an equivalence relation.

Remark 2. It is important to observe that given a harmonious partition P of the set of states of a CGM, the behavioural equivalence w.r.t. P of two states for a coalition does not imply their behavioural equivalence for another coalition. To see this, let us consider Example 2.

Example 2. Let $\mathcal{M}_1$ be a CGM having four states: s_1, s_2, s_3 and s_4, and three agents, 1, 2 and 3. Let p be the only boolean variable and let: $\mathsf{L}(s_1) = \mathsf{L}(s_2) = \{p\}$, $\mathsf{L}(s_3) = \mathsf{L}(s_4) = \emptyset$. Each agent can play either action 0 or action 1 at states s_1 and s_2, and only action 3 at s_3 and s_4. The transitions are: $\mathsf{out}(s_1, \langle 0,0,0\rangle) = \mathsf{out}(s_1, \langle 1,1,1\rangle) = s_3$, $\mathsf{out}(s_1, \alpha_1\rangle) = s_1$ for any other global move α_1 available at s_1, $\mathsf{out}(s_2, \langle 0,0,0\rangle) = \mathsf{out}(s_2, \langle 0,0,1\rangle) = s_4$, $\mathsf{out}(s_2, \alpha_2\rangle) = s_2$ for any other global move α_2 available at s_2, $\mathsf{out}(s_3, \langle 3,3,3\rangle) = s_3$ and $\mathsf{out}(s_4, \langle 3,3,3\rangle) = s_4$.

Let P be the harmonious partition of the set of states where s_1 and s_2 are in the cluster C_1, while s_3 and s_4 are in the cluster C_2. For $A = \{1\}$ it is easy to see that $s_1 \equiv_{P_A} s_2$. In fact, action 0 available at s_1 is simulated by action 0 available at s_2 and also action 1 available at s_1 is simulated by action 0 available at s_2. Conversely, action 0 at s_2 is simulated by action 0 (or 1) at s_1. However, for $A' = \{1,2\}$, $s_1 \not\equiv_{P'_A} s_2$. Indeed, at state s_1 the formula $\langle\!\langle 1,2\rangle\!\rangle \bigcirc \neg p$ is false, while at state s_2 is true (it suffices to play $\langle 0,0\rangle$). This shows that the equivalence of states for a coalition A does not imply the equivalence for each coalition A' such that $A \subset A'$.

The same example shows that the equivalence of states for a coalition A" does not imply the equivalence for each coalition A''' such that $A''' \subset A''$. In fact, take $A'' = \{1,2,3\}$. If the coalition plays the global move $\langle 0,0,0\rangle$ at s_1, which leads to the cluster C_1, then it can play the same move at s_2 to get the same effect; if it plays $\langle 1,1,1\rangle$ at s_1 then it can play either $\langle 0,0,0\rangle$ or $\langle 0,0,1\rangle$ at s_2 to get the same effect; finally, if it plays any move different form $\langle 0,0,0\rangle$ and $\langle 1,1,1\rangle$ at s_1, then it can play any move different form $\langle 0,0,0\rangle$ and $\langle 0,0,1\rangle$ at s_2 to get the same effect. A symmetrical reasoning on the actions that the coalition A'' can play at s_2 allow us to conclude that $s_1 \equiv_{P_{A''}} s_2$. However taking the subset coalition A''' to be $\{1,2\}$ we get, as already seen, that $s_1 \not\equiv_{P_{A'''}} s_2$.

Below we always assume that $\mathbb{S}$ is finite.

Definition 6 (Stability). *Given a partition $P = \{C_1, ..., C_n\}$ of the set of states $\mathbb{S}$ of a CGM and a relation $r \subseteq \mathbb{S} \times \mathbb{S}$, P is stable w.r.t. r when, for any $1 \leq i \leq n$, $s, t \in C_i$ implies $s\ r\ t$.*

If P is stable w.r.t. $\equiv_P$, then obviously it is stable w.r.t. $\equiv_{P^A}$ for each coalition A.

Our minimization procedure builds step by step the coarsest harmonious partition P of the set of states of the model $\mathcal{M}$ to be minimized that is stable w.r.t. $\equiv_P$. Then it builds out of P a minimal model bisimilar to $\mathcal{M}$ as a quotient of $\mathbb{S}$ with respect the equivalence $\equiv_P$.

Definition 7 (Quotient model). *Let $\mathcal{M}$ be a CGM $\langle \mathbb{A}, \mathbb{S}, \{\mathsf{Act}_a\}_{a\in\mathbb{A}}, \mathsf{out}, \mathsf{L}\rangle$. Let $P = \{C_1, ..., C_n\}$ be a harmonious partition of $\mathbb{S}$ that is stable w.r.t. $\equiv_P$. Let ρ be a function associating to each cluster C_i of P an element of C_i as representative element of C_i.*

A quotient-model $\mathcal{M}'$ of $\mathcal{M}$ w.r.t. $\equiv_P$ and ρ is defined as a

$$\mathcal{M}' = \langle \mathbb{A}, \mathbb{S}', \{\mathsf{Act}'_a\}_{a\in\mathbb{A}}, \mathsf{out}', \mathsf{L}'\rangle$$

where:

- *$\mathbb{S}'$ is the set of clusters in P: $\mathbb{S}' = \{C_1, ..., C_n\}$;*
- *C_i is connected to C_j via $\sigma_{\mathbb{A}}$ if and only if in the model $\mathcal{M}$ we have $\mathsf{out}(\rho(C_i), \sigma_{\mathbb{A}}) \in C_j$. This defines the transition function out'.*
- *The set $\{\mathsf{Act}'_a\}_{a\in\mathbb{A}}$ is constructed accordingly;*
- *For any C_i, $\mathsf{L}'(C_i) = \mathsf{L}(s)$, for any $s \in C_i$.*

Let us observe that, formally, the construction of a quotient model $\mathcal{M}'$ of $\mathcal{M}$ depends not only on the partition P, but also on the choice ρ of a representative state r_i of C_i. However, given a partition P that is stable with respect to the relation $\equiv_P$, the choice of ρ can have an effect only on labels of connecting edges in $\mathcal{M}'$ but not on the existence of a connection between two states of $\mathcal{M}'$ (that is, clusters of P). In fact, let C_i be a cluster, r_i be a state in $\mathcal{M}$ such that $\rho(C_i) = r_i$ and s be any other element of C_i. Then $s \equiv_P r_i$ by construction, therefore:

- If $\sigma_{\mathbb{A}}$ leads from r_i to a $t \in C_j$ in $\mathcal{M}$ then by definition there is some global action leading from s to some state (possibly another than t) that belongs to the same cluster C_j;
- If no global action leads from r_i to C_j in $\mathcal{M}$ then no global action leads from s to C_j in $\mathcal{M}$. In fact if some global action $\sigma_{\mathbb{A}}$ leads from s to some state in C_j then some global action $\sigma'_{\mathbb{A}}$ leads from r_i to some state in C_j, since $s \equiv_P r_i$.

Therefore a quotient model of $\mathcal{M}$ w.r.t. a harmonious partition P of $\mathcal{M}$'s states is unique modulo renaming of edge labels.

The following result states that a quotient model of $\mathcal{M}$, as defined above, is indeed bisimilar to $\mathcal{M}$.

Theorem 2. *Let $\mathcal{M}$ be a CGM $\langle \mathbb{A}, \mathbb{S}, \{\mathsf{Act}_a\}_{a\in\mathbb{A}}, \mathsf{out}, \mathsf{L}\rangle$. Let $P = \{C_1, ..., C_n\}$ be a harmonious partition of its states that is stable w.r.t. $\equiv_P$ and let ρ be a function choosing representative elements from clusters. Let $\mathcal{M}' = \langle \mathbb{A}, \mathbb{S}', \{\mathsf{Act}'_a\}_{a\in\mathbb{A}}, \mathsf{out}', \mathsf{L}'\rangle$ be a quotient-model of $\mathcal{M}$ w.r.t. $\equiv_P$ and ρ. Then the relation $\beta \subseteq \mathbb{S} \times \mathbb{S}'$ defined by: $s\beta C_i$ iff $s \in C_i$ is a global alternating bisimulation between $\mathcal{M}$ and $\mathcal{M}'$.*

The proof of Theorem 2 is given in Section 5.

As a consequence of Theorem 2 and Theorem 1 we get that if if $\mathcal{M}$ is a model, P a partition of its states that is stable w.r.t. $\equiv_P$, $\mathcal{M}'$ a corresponding quotient model, and finally, ϕ is any ATL* formula (over the given sets of propositions and agents), then $\mathcal{M}$ is a model of ϕ if and only if $\mathcal{M}'$ is a model of ϕ.

3.2 Minimization Algorithm

When the model $\mathcal{M}$ to be minimized has a finite number of states, as it is in our intended application to model minimization in TATL, a maximal bisimulation relation $\beta \subseteq \mathbb{S} \times \mathbb{S}$, hence a corresponding minimal partition P of $\mathbb{S}$ stable w.r.t. $\equiv_P$ inducing a minimal quotient model of a CGM $\mathcal{M}$, can be given a stepwise characterization and effectively constructed, analogously to the case of labelled partition systems. More precisely:

Definition 8 (Stratified bisimilarity relations).
Given a CGM $\mathcal{M} = \langle \mathbb{A}, \mathbb{S}, \{\mathsf{Act}_a\}_{a \in \mathbb{A}}, \mathsf{out}, \mathsf{L} \rangle$, the stratified alternating bisimulation relations $\beta_k \subseteq \mathbb{S} \times \mathbb{S}$ for $k \in \mathbb{N}$ are defined as follows:

- *$s_1 \beta_0 s_2$ iff $s_1, s_2 \in \mathbb{S}$ and $\mathsf{L}(s_1) = \mathsf{L}'(s_2)$;*
- *$s_1 \beta_{k+1} s_2$ iff $s_1 \beta_k s_2$, $\mathsf{L}(s_1) = \mathsf{L}'(s_2)$ and for each coalition $A \subseteq \mathbb{A}$:*
  1. ***Forth**. For any $\alpha_A \in \mathsf{act}_A(s_1)$ there is an $\alpha'_A \in \mathsf{act}'_A(s_2)$ such that for any $t_2 \in \mathsf{Out}(s_2, \alpha'_A)$ there exists $t_1 \in \mathsf{Out}(s_1, \alpha_A)$ such that $t_1 \beta_k t_2$.*
  2. ***Back**. For any $\alpha_A \in \mathsf{act}_A(s_2)$ there is an $\alpha'_A \in \mathsf{act}'_A(s_1)$ such that for any $t_3 \in \mathsf{Out}(s_1, \alpha'_A)$ there exists $t_4 \in \mathsf{Out}(s_2, \alpha_A)$ such that $t_3 \beta_k t_4$.*
- *By construction, for any k we have $\beta_{k+1} \subseteq \beta_k$. Set the relation β^* to be $\bigcap_{k \in \mathbb{N}} \beta_k$.*

When $\mid \mathbb{S} \mid$ is finite, the relation β^* can be obviously be computed in finite time since there is a j, $0 \leq j \leq\mid \mathbb{S} \mid$ such that $\beta^* = \beta_j$. By Remark 1 any full alternating bisimulation relation that is a subset of $\mathbb{S} \times \mathbb{S}$ is a fixpoint solution of the equation $X = E(X)$, where X is a subset of $S \times S$ having the property that if $\langle q, q' \rangle \in X$ then $\mathsf{L}(q) = \mathsf{L}'(q')$. We have:

Theorem 3. *The relation β^* is the maximal fixpoint solution of the equation X=E(X).*

This can be shown by arguments similar to those proving an analogous claim for labelled transition systems [HM85]. The detailed proof is given in Section 5.

Remark 3. We can observe that if P_k is the harmonious partition of $\mathbb{S}$ corresponding to a given stratified alternating bisimulation relation β_k then $s_1 \equiv_{P_k} s_2$ (as in Definition 5) if and only if $s_1 \beta_{k+1} s_2$. The two formalizations capture the same concept, but behavioural equivalence directly corresponds to the implementation of our minimization algorithm (see Section 4). Moreover, any harmonious partition P of the set of states of a model $\mathcal{M}$ is stable w.r.t. the relation $\equiv_P$ (as in Definition 6) if and only if $\equiv_P$ is a solution of the equation $X = E(X)$, although not necessarily the maximal one, corresponding to the minimal, *i.e* coarsest, partition. The partition of $\mathbb{S}$ induced by β^* is the minimal partition that is stable with respect $\equiv_P$.

Let P^* be the partition of the states $\mathbb{S}$ of a CGM $\mathcal{M}$ induced by β^*. The quotient model of $\mathcal{M}$ with respect to $\equiv_{P^*}$ is the minimization of $\mathcal{M}$ with respect to alternating bisimilarity. This yields an algorithm that minimizes $\mathcal{M}$ by computing, step by step, the partition P^* starting from an initial partition; its underlying general principle is:
Let P_0, the initial partition, be such that $s_1, s_2 \in \mathbb{S}$ belong to the same cluster if and only if $\mathsf{L}(s_1) = \mathsf{L}'(s_2)$.
For each $i > 0$ compute the i-th approximant P_i of P^* until $P_{i+1} = P_i$.
Output P_i as the value of P^*.

4 Implementation and application to TATL

We have implemented (in OCaml, the same language used for TATL) our minimization algorithm in order to add to TATL a new functionality: the minimization of the model extracted from an open tableau for an input formula ϕ by executing the procedure given by the completeness proof for ATL^* tableaux in [Dav15]. So far, TATL does not show any model, but only the tableau. The forthcoming version of TATL will allow the user to visualize the model generated by the completeness proof procedure and also its minimization. Here we give the pseudo-code of our implementation.

Algorithm 1 Main Procedure

```
P ← initial partition
change ← true
while change do
    change ← false
    for all cluster B ∈ P do
        if SPLIT(B, P) = {B1, B2} ≠ {B} then
            Refine P by replacing B by B1 and B2
            change ← true
        end if
    end for
end while
```

Algorithm 2 function SPLIT(B, P)

```
choose a state s ∈ B
B1, B2 ← ∅
for all t ∈ B do
    if EQUIVALENCE(s, t, P) then
        B1 ← B1 ∪ {t}
    else
        B2 ← B2 ∪ {t}
    end if
end for
if B2 = ∅ then
    return {B1}
else
    return {B1, B2}
end if
```

Obviously the algorithm terminates, because the number of iterations of the main loop is upper bounded by the size of the set of states, that is finite.

Algorithm 3 function EQUIVALENCE(s, t, P)

```
if s = t then
    true
else
    if L(s) = L(t) then
        clusterS ← set of successor clusters of s
        clusterT ← set of successor clusters of t
        if clusterS = clusterT then
            EQUIVALENCE_BY_COALITIONS(s, t, P)
        else
            false
        end if
    else
        false
    end if
end if
```

The core function is SPLIT that splits a cluster of the current partition P_i in two clusters whenever two states s and t in it are not behaviourally equivalent with respect to P_i; to do so it calls the function EQUIVALENCE. This last checks the behavioural equivalence of states w.r.t. the current partition for each coalition A (as in Definition 5), by means of the function EQUIVALENCE_BY_COALITIONS. For space reason, the pseudo code of this last function is not given here. This function checks if two states in a given cluster of the current partition P are behaviourally equivalent with respect to P for all coalitions or not, which inevitably makes the program to have an exponential complexity. It is necessary to check each coalition because behavioural equivalence of two states w.r.t. the current partition for a given coalition does not imply equivalence for another coalition (see Remark 2 and Example 2).

When the main procedure halts then the computed result P is the partition P^* associated to β^*. In order to prove this claim let us note P_0 the initial partition of the procedure, P_1, P_2... P_m the partitions computed in the main loop until stability, $r_0, r_1, r_2 ... r_m$ the corresponding equivalence relations, and $r = r_m$ the relation corresponding to the final result P. An easy induction on $i \in \mathbb{N}$ proves that $\beta^* \subseteq \beta_i \subseteq r_i$. Hence $\beta^* \subseteq r$. For the converse inclusion, let us observe that if P is the result of the main procedure, then P is stable w.r.t. $\equiv_P$ (see the definition of the function SPLIT). By Remark 3, r is a solution of the fixed point equation $X = E(X)$. Hence $r \subseteq \beta^*$, because, by Theorem 3, β^* is the maximal solution of such an equation. Thus $\beta^* = r$.

Also the model extraction function from a tableau (via the procedure of the completeness proof) has been implemented and partial tests of our implementation of the minimization algorithm applied to the model extracted by the tableau have been done, but a complete and representative set of test cases still needs to be constructed.

The last figure illustrates the minimization procedure via a simple example, with one agent, chosen among the tests so far done. The input formula ϕ of the tableau, as provided to the software TATL, is exhibited at the left top : it is $\langle\langle 1\rangle\rangle((\langle\langle 1\rangle\rangle\Box\langle\langle\emptyset\rangle\rangle \bigcirc \Diamond\Box a) \wedge (\bigcirc(\neg b \wedge \neg a)))$, where a and b are propositional letters and $\emptyset$ is the empty coalition. The graph on the left, having eight states, is the model of the formula produced by the completeness procedure: it satisfies ϕ at state $n1$. At the right, the minimized model, having three states and satisfying

ϕ at state $n1$. The literals holding at each state are indicated inside each state ellipse.

Fig. 3. Input (left) and output (right) of the minimization algorithm

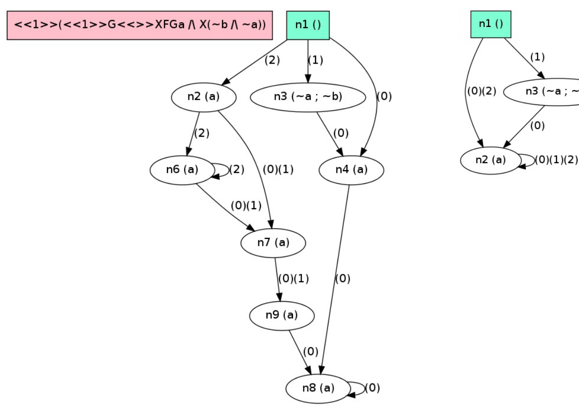

5 Proofs

5.1 Proof of Theorem 1

The proof is organized in two parts: the first one defines the preliminary notion of *simulating strategy*, the second one constitutes the core of the proof: we show that bisimilar states satisfy the same formulae.

Simulating Strategy

Let s be a state of $\mathcal{M}$ and s' be a state of $\mathcal{M}'$ such that $s\beta s'$. Below, by a *strategy* F'_A *simulating* F_A *in* $\mathcal{M}'$ we mean a strategy in $\mathcal{M}$' simulating all the plays starting at s that are coherent with F_A by plays starting at s'. The strategy F'_A is built as follows.

The notation $reach(\mathcal{M}, h, F_A)$, where h is an history in $\mathcal{M}$, denotes the set of all the states in $\mathcal{M}$ that occur in plays that stem from the last state of h and that are coherent with a given strategy F_A.

By convenience we suppose the set of actions $\{\mathsf{Act}_a\}_{a\in\mathbb{A}}$ in $\mathcal{M}$ to be enumerable, so that, for any history h, the set $reach(\mathcal{M}, h, F_A)$ is enumerable. Otherwise, we would need to to use transfinite induction (and the axiom of choice) to inductively define F'_A. Thus, let an enumeration of $reach(\mathcal{M}, h, F_A)$ be given.

The strategy F_A associates an A-move to each history h in $\mathcal{M}$. We build F'_A for histories in $\mathcal{M}'$ by first defining, by induction on $\mathbb{N}$, a partial strategy F'_A simulating F_A for histories having length 0 (the empty history), then for histories having length greater than 0. That is, we build the partial function F'_A step by step, by first simulating the A-action done by F_A at s (the history has length 1), then the actions done by F_A at states in $\mathcal{M}$ reached via such an action, and so on. Then we extend F'_A to make it total: the function F'_A so far built might not be defined for states in $\mathcal{M}'$ that are not connected by paths coherent with it to s'.

In order to define the partial function F'_A, four auxiliary notions are previously defined simultaneously by induction on $\mathbb{N}$:

1. A chain of sets $succ_0(\mathcal{M}', s') \subseteq ... \subseteq succ_n(\mathcal{M}', s')$... where $succ_n(\mathcal{M}', s')$ will be the set of states of $\mathcal{M}'$ that are reachable from s' in at most n steps if the agents in the coalition A play accordingly to the partial strategy F'_A so far constructed.
2. A chain of mappings $\zeta_0 \subseteq ...\zeta_n...$ where $\zeta_n : succ_n(\mathcal{M}', s') \to reach(\mathcal{M}, h, F_A)$, h is an history coherent with F_A starting with s (the considered state in $\mathcal{M}$), and $\zeta_n(t')\beta t'$ for every $t' \in succ_n(\mathcal{M}', s')$;
3. A chain of partial strategies for the coalition A in $\mathcal{M}'$: $F'_A(0) \subseteq ... \subseteq F'_A(n)$... where the domain of $F'_A(n)$ is $succ_{n-1}(\mathcal{M}', s')$ ($succ_{n-1}(\mathcal{M}', s') = \emptyset$ by convention);
4. An infinite sequence $S'_0, S'_1, S'_2, ...$ of subsets of states of $\mathcal{M}$'.

The inductive definition is as follows, where in the base one defines $reach_0, \zeta_0$ and $F'_A(0)$, while at step $n+1$ one defines ζ_{n+1}, $reach_{n+1}$ and $F'_A(n+1)$ and S'_n.

- Step 0. This is just an initialisation step, were we consider the empty history in $\mathcal{M}$, having length 0.
  - $succ_0(\mathcal{M}', s') = \{s'\}$;
  - $\zeta_0(s') = s$;
  - $F'_A(0) = \emptyset$.
- Step $n+1$. Here we consider histories h in $\mathcal{M}$ stemming from s and having length $n+1$, thus, in particular, also histories of length $n+1$ ending with s, since plays in $\mathcal{M}$ may contain cycles on s.
  - $S'(n) = succ_n(\mathcal{M}', s') \setminus succ_{n-1}(\mathcal{M}', s')$. Intuitively, $S'(n)$ is the set of the new states reached in $\mathcal{M}'$ from s' in n steps.
  - For every $t' \in S'_n$, note t the state $\zeta_n(t')$ of $\mathcal{M}$. Let $\sigma_A(t)$ be $F_A(h)$ where h ends with t and h as length n. Since $t\beta t'$ (by construction of ζ_n at the previous step), then we can choose an A-action at t' in $\mathcal{M}$', say ${\sigma'}_A(t')$, such that for every $q' \in \mathsf{Out}(t', {\sigma'}_A)$ there is a bisimilar state (of $\mathcal{M}$) in $\mathsf{Out}(t, \sigma_A)$, and let q be the first such state in the given enumeration of $reach(\mathcal{M}, h, F_A)$. If $q' \in reach_n(\mathcal{M}', t')$ set $\zeta_{n+1}(q') = \zeta_n(t)$ (since ζ and F'_A must be functions), otherwise set $\zeta_{n+1}(q') = q$.
  - $succ_{n+1}(\mathcal{M}', s') = succ_n(\mathcal{M}', s') \cup \bigcup_{s' \in S'_n} \mathsf{Out}(s', {\sigma'}_A)$.
  - $F'_A(n+1) = F'_A(n) \cup \bigcup_{t' \in S'_n} {\sigma'}_A(t')$.

Set F'_A (the partial strategy for $\mathcal{M}$' we are looking for) to be: $\bigcup_{n\in\mathbb{N}\ F'_A(n)}$. This defines F'_A for histories ending with a state that is reachable from s'. Then arbitrarily extend the function $F'_A(n)$ to histories ending with states of $\mathcal{M}$' where it is not yet defined.

Observe that, given s and s', the strategy F'_A simulating F_A is uniquely defined (once a choice of a σ'_A in $\mathcal{M}'$ simulating the given σ_A in $\mathcal{M}$ is made).

Proof of truth preservation by bisimilar states

The only item of Theorem 1 that needs a proof is the first one (then the second one immediately follows). In order to prove it, we first prove the following lemma:

Lemma 1. *Let A be any coalition, let β be an A-simulation between two models $\mathcal{M}$ and $\mathcal{M}'$ and let s_1 and s_2 be two states of these models such that $s_1 \beta s_2$. Suppose that all the paths λ stemming from s_1 and coherent with a given strategy F_A are such that $\mathcal{M}, \lambda \models \Psi$, where Ψ is any* ATL^* *path formula such that A is the only coalition occurring in Ψ. Then all the paths λ' stemming from s_2 and coherent with the corresponding simulating strategy F'_A are such that $\mathcal{M}', \lambda' \models \Psi$.*

The proof of the Lemma is by induction on Ψ. Recall that any state formula is also a path formula, while the converse is false.

- *Base.* Suppose that all the paths λ stemming from s_1 and coherent with a given strategy F_A are such that $\mathcal{M}, \lambda \models p$, where p is a propositional letter. This just means that p is true at s_1 and the result trivially holds by Local Harmony.
- *Inductive Step.*
  - $\Psi = \bigcirc\Psi_1$. This case is easy, given the definition of bisimulation.
  - Ψ is either $\neg\Psi_1$ or $\Psi_1 \wedge \Psi_2$. The result for these cases follow immediately from the inductive hypothesis.
  - $\Psi = \Box\Psi_1$.
    Suppose that all the paths λ in $\mathcal{M}$ stemming from s_1 and coherent with a given strategy F_A are such that $\mathcal{M}, \lambda \models \Box\Psi_1$. Therefore in $\mathcal{M}$ all the suffixes of such paths satisfy Ψ_1, that is, all the paths stemming from s_1 and coherent with F_A satisfy Ψ_1. Since by construction of the simulating strategy F'_A the states occurring in the paths stemming from s_2 and coherent with F'_A are β-images of states in in the paths λ stemming from s_1 in $\mathcal{M}$, by the inductive hypothesis all the suffixes of all the paths stemming from s_2 and coherent with F'_A satisfy Ψ_1. Hence all the paths stemming from s_2 and coherent with F'_A satisfy $\Box\Psi_1$.
  - $\Psi = \Psi_1 \mathsf{U} \psi_2$. This case is similar to the above one.
  - $\Psi = \langle\langle A\rangle\rangle\Psi_1$. Suppose that all the paths λ stemming from s_1 and coherent with a given strategy F_A are such that $\mathcal{M}, \lambda \models \langle\langle A\rangle\rangle\Psi_1$. Since $\langle\langle A\rangle\rangle\Psi_1$ is a state formula, this actually means that there is some strategy, say

G_A, such that all the paths λ stemming from s_1 and coherent with G_A are such that $\mathcal{M}, \lambda \models \Psi_1$. Thus, by inductive hypothesis, all the paths λ' stemming from s_2 and coherent with G'_A – where G'_A is the strategy simulating G_A – are such that $\mathcal{M}', \lambda' \models \Psi_1$. Therefore trivially all the paths λ' stemming from s_2 and coherent with F'_A, the strategy simulating F_A, are such that $\mathcal{M}', \lambda' \models \langle\langle A \rangle\rangle \Psi$.

Once established Lemma 1 the proof of the first item of Theorem 1 for ATL^* is almost immediate. In fact, suppose that $\mathcal{M} \overset{\beta}{\rightleftarrows}_A \mathcal{M}'$, and $s_1 \beta s_2$. If $\mathcal{M}, s_1 \models \phi$, where ϕ is an ATL^* state formula where only the coalition A occurs, then each path λ starting at s_1 satisfies ϕ, because for satisfaction of state formulae only the state $\lambda_0 = s_1$ matters. In other words, all the paths λ stemming from s_1 and coherent with the empty strategy F_A are such that $\mathcal{M}, \lambda \models \phi$. Hence, by Lemma 1, all the paths λ' stemming from s_2 and coherent with F'_A, the strategy simulating F_A –that is again empty – are such that $\mathcal{M}', \lambda' \models \phi$. But this means that $\mathcal{M}', s_2 \models \phi$. The converse, namely that if $\mathcal{M}', s_2 \models \phi$ then $\mathcal{M}, s_1 \models \phi$ also follows from Lemma 1, because the inverse relation of β is a A-bisimulation between $\mathcal{M}$' and $\mathcal{M}$.

This concludes the proof of Theorem 1.

5.2 Proof of Theorem 2

Let P, $\mathcal{M}$ and $\mathcal{M}'$ as in the statement of the theorem. By construction, β is total on S and its inverse is total on S'. What needs to be shown is that for any $\mathrm{A} \subseteq \mathbb{A}$, β is an alternating A-bisimulation between $\mathcal{M}$ and $\mathcal{M}'$.

Let A be any coalition, and suppose that $s \in C_i$. By construction local harmony between s and C_i holds.

- *Proof of the Forth Condition*
  Let us suppose that σ_A is an A-action available at s in $\mathcal{M}$. Let $r = \rho(C_i)$, that is, r is the representative element of C_i used to build the transitions in $\mathcal{M}'$ from the state C_i to its successors.
  Since the partition P of states of $\mathcal{M}$ used to build the quotient model $\mathcal{M}'$ is stable w.r.t. $\equiv_P$ (by construction), it is stable also w.r.t. $\equiv_{P^A}$. Therefore $s \equiv_{P^A} r$. By definition of $\equiv_{P^A}$, there is some A-action σ^*_A available at r in $\mathcal{M}$ such that the following property holds:
  $P1$: the set of clusters of states that are reachable from r via σ^*_A in $\mathcal{M}$ is a subset of the set of clusters of states that are reachable from s via σ_A.
  Take the A-action available at C_i that must be shown to exist for the Forth condition to hold to be σ^*_A. Thus we must show :
  If a cluster C_j belongs to $\mathsf{Out}(C_i, \sigma^*_A)$ (in $\mathcal{M}$'), then there exists $t \in \mathsf{Out}(s, \sigma_A)$ (in $\mathcal{M}$) such that $t \beta C_j$ *i.e* $t \in C_j$.
  So, suppose that in $\mathcal{M}'$: $C_j \in \mathsf{Out}(C_i, \sigma^*_A)$. By construction of $\mathcal{M}'$, this implies that in $\mathcal{M}$ there is some global action extending σ^*_A to all the agents, say $comp(\sigma^*_A)$, that leads from r to some state q belonging to C_j. Then,

by the property $P1$ above, there is also in $\mathcal{M}$ some global action $comp(\sigma_A)$ extending σ_A to all the agents that leads from s to some state t belonging to C_j, in other words there exists $t \in \mathsf{Out}(s, \sigma_A)$ such that $t \in C_j$. We are done.

– *Proof of The Back Condition*
  Let us suppose that σ_A is an A-action available at C_i in $\mathcal{M}'$. By construction, σ_A is the restriction to agents in A of a global action in $\mathcal{M}$ that leads from the representative element r of C_i to some state.
  Since the partition P of states of $\mathcal{M}$ used to build the quotient model $\mathcal{M}'$ is stable w.r.t. $\equiv_P$, it is so also w.r.t. $\equiv_{P^A}$. Therefore, since $s, r \in C_i$, $s \equiv_{P^A} r$. Therefore there is some action σ_A^* available at s in $\mathcal{M}$ such that the following property holds :
  $P2$: the set of the clusters of states that are reachable from s via σ_A^* is a subset of the set of the clusters of states that are reachable from r via σ_A.
  Take the A-action available at s that must be shown to exist for the Back condition to hold to be σ_A^*. Thus we must show :
  If a state t belongs to $\mathsf{Out}(s, \sigma_A^*)$ (in $\mathcal{M}$), then there exists a $C_j \in \mathsf{Out}(C_i, \sigma_A)$ (in $\mathcal{M}$') such that $t \in C_j$, i.e. $t\beta C_j$.
  So, suppose that a state t belongs to $\mathsf{Out}(s, \sigma_A^*)$, and let C_j be the cluster to which t belongs. By the property $P2$ the cluster C_j is reachable also from r via σ_A. Therefore, by construction of $\mathcal{M}$', $C_j \in \mathsf{Out}(C_i, \sigma_A)$. We are done.

□

5.3 Proof of Theorem 3

First, let us observe that for $i > 0$, the relation β_i of Definition 8 can be equivalently described as $E(\beta_i)$, where E is the operator defined in Remark 1, Section 2.

In order to prove the theorem, first we show that β^* is a solution of the fixpoint equation $X = E(X)$.

Since Since $\beta_0, \beta_1, \beta_2, ...$ is a descending chain (by construction) and the set $\mathbb{S}$ is finite, there must be a j such that $\forall m \geq j\ \beta_m = \beta_j$. In particular, $\beta_{j+1} = E(\beta_j)$. But β^*, defined as $\bigcup_{k\in\mathbb{N}} \beta_k{=}\beta_j = \beta_{j+1}$. Hence $\beta^* = E(\beta^*)$.

Then we show that any binary relation on $\mathbb{S}$ that is a solution to $X = E(X)$ is included in β^*, therefore that β^* is the maximal solution.

Thus, let suppose that $r \subseteq \mathbb{S}\times\mathbb{S}$ is such that $r = E(r)$. We show, by induction on $i \in \mathbb{N}$, that if $\langle s_1, s_2\rangle \in r$ then $s_1\beta_i s_2$.
Base: $i = 0$.
Since $\langle s_1, s_2\rangle \in r = E(r)$, then s_1 and s_2 have the same labels. Then obviously $s_1\beta_0 s_2$.
Inductive Step: $i > 0$.
Since $\langle s_1, s_2\rangle \in r = E(r)$, then i) $\mathsf{L}(s_1) = \mathsf{L}'(s_2)$ and (ii) for every coalition A:
(a) for any $\alpha_A \in \mathsf{act}_A(s_1)$ there is an $\alpha'_A \in \mathsf{act'}_A(s_2)$ such that for any $t_2 \in \mathsf{Out}(s_2, \alpha'_A)$ there exists $t_1 \in \mathsf{Out}(s_1, \alpha_A)$ such that such that $t_1 r t_2$, and

(b) for any $\alpha_A \in \mathsf{act}_A(s_2)$ there is an $\alpha'_A \in \mathsf{act}'_A(s_1)$ such that for any $t_3 \in \mathsf{Out}(s_1, \alpha'_A)$ there exists $t_4 \in \mathsf{Out}(s_2, \alpha_A)$ such that $t_3 r t_4$.

By inductive hypothesis $r \subseteq \beta_{i-1}$ thus $t_1 \beta_{i-1} t_2$ and $t_3 \beta_{i-1} t_4$. Therefore $s_1 \beta_i s_2$.

Therefore for each n we have $s_1 \beta_n s_2$ and we conclude that $s_1 \beta^* s_2$.
□

6 Conclusions

Up to our knowledge, the algorithm proposed in this work is the first procedure that minimizes ATL^* models with respect to alternating bisimulation.

This algorithm has a time complexity that is exponential in the size of $\mathbb{A}$, since, as observed, all the coalitions – that is all the subsets of $\mathbb{A}$– need to be checked in order to conclude that a given cluster of the current partition does not need to be split. It is interesting to compare it with the classical partition-refinement minimization algorithms for labelled transition systems, whose complexity depend only on the number n of states of the system and the number m of transitions: the algorithm in [KS90] has time complexity $O(nm)$ while the optimized algorithm in [PT87] has time complexity $m \log n$. Labelled transition systems can be seen as concurrent game structures with exactly one agent, thus it is not surprising that minimizing ATL^* models is harder, both conceptually and algorithmically, than minimizing CTL^* models. Although the problem of minimizing an ATL^* model is intrinsically exponential, it would be interesting to face issues of optimisation of our algorithm with the view of making it more efficient for practical use.

As we said, we implemented and tested our algorithm, but a large, complete and representative set of test cases is still ongoing work. When this will be finished we will add to the prover TATL the functionality of exhibiting minimized models of the input formula.

In this work we have considered only ATL^* with perfect information. Recently a definition of bisimilarity of models coping with imperfect information has been proposed [BCD+17] and it might be interesting to explore the possibility of extending our study to the minimization of models of ATL^* with imperfect information.

Aknowledgements The authors would like to thank Damien Regnault and Marta Cialdea for their careful reading of first drafts of this work and for their useful remarks. The very first ideas underlying this work rose in the context of the direction of a project of two fourth year university students at the university of Evry Val d'Essonne: Lylia Bellabiod and Théo Chelim.

References

[ÅGJ07] Thomas Ågotnes, Valentin Goranko, and Wojciech Jamroga. Alternating-time temporal logics with irrevocable strategies. In *Proceedings of the 11th*

Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2007), Brussels, Belgium, June 25-27, 2007, pages 15–24, 2007.

[AHK02] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.

[BCD+17] F. Belardinelli, R. Condurache, C. Dima, W. Jamroga, and A. V. Jones. Bisimulations for verifying strategic abilities applied to voting protocols. In *Proceedings of AAMAS17*. IFAAMAS, 2017.

[BLW09] Piero A. Bonatti, Carsten Lutz, and Frank Wolter. The complexity of circumscription in dls. *Journal of Artificial Intelligence Research*, 35, 2009.

[BY00] François Bry and Adnan Yahya. Positive unit hyperresolution tableaux and their application to minimal model generation. *Journal of Automated Reasoning*, 25(1):35–82, 2000.

[CDG14] S. Cerrito, A. David, and V. Goranko. Optimal tableau method for constructive satisfiability testing and model synthesis in the alternating-time temporal logic atl+. In *Proceedings of IJCAR 2014*, volume LNAI 8652. Springer, 2014.

[Dav] A. David. Tatl: Tableaux for atl*. `http://atila.ibisc.univ-evry.fr/tableau_ATL_star/index.php`.

[Dav15] Amélie David. Deciding ATL* satisfiability by tableaux. In *25th International Conference on Automated Deduction (CADE 2015)*, volume 9195 of *Lecture Notes in Computer Science*, pages 214–228, Berlin, Germany, August 2015.

[DGL16] Stéphane Demri, Valentin Goranko, and Martin Lange. *Temporal Logics in Computer Science*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016.

[GGOP08] Laura Giordano, Valentina Gliozzi, Nicola Olivetti, and Gian Luca Pozzato. *Reasoning about Typicality in Preferential Description Logics*, pages 192–205. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[GH09] Stephan Grimm and Pascal Hitzler. A preferential tableaux calculus for circumscriptive alco. In Axel Polleres and Terrance Swift, editors, *Web Reasoning and Rule Systems, Third International Conference, RR 2009*, volume 5837, page 40–54, Chantilly, VA, USA, 2009. Springer.

[GHS01] Lilia Georgieva, Ullrich Hustadt, and Renate A. Schmidt. Computational space efficiency and minimal model generation for guarded formulae. In *LPAR 2001 Proceedings*, volume 2250 LNAI, pages 85–99. Springer, 2001.

[HFK00] Ryuzo Hasegawa, Hiroshi Fujita, and Miyuki Koshimura. *Efficient Minimal Model Generation Using Branching Lemmas*, pages 184–199. Springer Berlin Heidelberg, Berlin, Heidelberg, 2000.

[Hin88] Jaakko Hintikka. Model minimization - an alternative to circumscription. *J. Autom. Reasoning*, 4(1):1–13, 1988.

[HM85] Matthew Hennessy and Robin Milner. Algebraic laws for nondeterminism and concurrency. *J. ACM*, 32(1):137–161, January 1985.

[KS90] Paris C. Kanellakis and Scott A. Smolka. CCS expressions, finite state processes, and three problems of equivalence. *Information and Computation*, 86(1):43–68, 1990.

[LIS12] Aceto L., Ingolfsdottir, and Jiri S. The algorithmics of bisimilarity. In Sangiorgi D. and Rutten J., editors, *Advanced topics in bisimulation and coinduction*, pages 100–171. Cambridge University Press, 2012.

[Lor94] Sven Lorenz. A tableau prover for domain minimization. *Journal of Automated Reasoning*, 13(3):375–390, 1994.

[McC87] J. McCarthy. Circumscription: A form of non-monotonic reasoning. In M. L. Ginsberg, editor, *Readings in Nonmonotonic Reasoning*, pages 145–151. Kaufmann, Los Altos, CA, 1987.

[Nie96] Ilkka Niemelä. A tableau calculus for minimal model reasoning. In *Proceedings of the Fifth Workshop on Theorem Proving with Analytic Tableaux and Related Methods*, pages 278–294, Terrasini, Italy, May 1996. Springer-Verlag.

[PS14] Fabio Papacchini and Renate A. Schmidt. *Terminating Minimal Model Generation Procedures for Propositional Modal Logics*, pages 381–395. Springer International Publishing, Cham, 2014.

[PT87] R. Paige and R.E. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987.