Logic-based specification and verification of homogeneous dynamic multi-agent systems

Riccardo De Masellis¹ · Valentin Goranko^{1,2}

Published online: 28 April 2020
© The Author(s) 2020

Abstract

We develop a logic-based framework for formal specification and algorithmic verification of *homogeneous* and *dynamic* concurrent multi-agent transition systems. Homogeneity means that all agents have the same available actions at any given state and the actions have the same effects regardless of which agents perform them. The state transitions are therefore determined only by the vector of numbers of agents performing each action and are specified symbolically, by means of conditions on these numbers definable in Presburger arithmetic. The agents are divided into *controllable* (by the system supervisor/controller) and *uncontrollable*, representing the environment or adversary. Dynamicity means that the numbers of controllable and uncontrollable agents may vary throughout the system evolution, possibly at every transition. As a language for formal specification we use a suitably extended version of Alternating-time Temporal Logic, where one can specify properties of the type “a coalition of (at least) n controllable agents can ensure against (at most) m uncontrollable agents that any possible evolution of the system satisfies a given objective γ ”, where γ is specified again as a formula of that language and each of n and m is either a fixed number or a variable that can be quantified over. We provide formal semantics to our logic $\mathcal{L}_{\text{HDMAS}}$ and define normal form of its formulae. We then prove that every formula in $\mathcal{L}_{\text{HDMAS}}$ is equivalent in the finite to one in a normal form and develop an algorithm for global model checking of formulae in normal form in finite HDMAS models, which invokes model checking truth of Presburger formulae. We establish worst case complexity estimates for the model checking algorithm and illustrate it on a running example.

Keywords Dynamic multi-agent systems · Logics for multi-agent systems · Logics for strategic reasoning · Model-checking

✉ Riccardo De Masellis
riccardo.demasellis@philosophy.su.se

Valentin Goranko
valentin.goranko@philosophy.su.se

¹ Stockholm University, Stockholm, Sweden

² University of Johannesburg (Visiting Professorship), Johannesburg, South Africa

1 Introduction

1.1 The framework

We consider discrete concurrent multi-agent transition systems, i.e. multi-agent systems (MAS) in which the transitions take place in a discrete succession of steps, as a result of a simultaneous (or, at least mutually independent) actions performed by all agents. Such MAS are typically modelled as *concurrent game models* (cf [1] or [6]).

Here we focus on a special type of concurrent MAS, which are *homogeneous* and *dynamic*, in a sense explained below.

The *homogeneity* means that all agents are essentially indistinguishable from each other, as their possible behaviours are determined by the same protocol. In particular, they have the same available actions at each state and the effect of these actions depends not on *which* agents perform them, but only on *how many* agents perform each action. Thus, the transitions in such systems are determined *not by the specific action profiles*, but only by *the vector of numbers of agents that perform each of the possible actions in these action profiles*. The latter can be regarded as an abstraction of the action profile. The transitions are specified symbolically, by means of conditions on these vectors, definable in Presburger arithmetic.

Typical examples of such homogeneous systems include:

- *voting procedures* where the outcome only depends on how many agents vote for each possible alternative, but not who votes for what. These also involve voting procedures where anonymity is required and the identity of agents should not be inferred by observing the system's evolution [14, 18];
- *sensor networks* of a type where protocols only depend on how many sensors send any given signal [21];
- *computer network servers*, the functioning of which only depends on how many currently connected users are performing any given action (e.g. uploading or downloading data, sending printing jobs, communicating over common channels, etc);
- *markets*, the dynamics of which only depends on how many agents are selling and how many are buying any given stock (assuming the transactions are per unit) but not exactly who does what.

The *dynamicity* of the systems that we consider means that the set (hence, the number) of agents being present (or, just acting) in the system may vary throughout the system evolution, possibly at every transition from a state to a state. All examples listed above naturally have that dynamic feature. There are different ways to interpret such dynamicity. In the extreme version, agents literally appear and disappear from the system, e.g. users joining and leaving an open network. A less radical interpretation is where the agents are in the system all the time but may become active and inactive from time to time, e.g. voters, or members of a committee, may abstain from voting in one election or decision making round, and then become active again in the next one. A more refined version is where at every state of the system performance each agent decides to act (i.e. take one of the available actions) or pass/idle, formally by performing the 'pass/idle' action. Technically, all these interpretations seem to be reducible to the latter one. However, the way we model the dynamicity here is by assuming that there is an unbounded, and possibly infinite set of 'potentially existing' agents, but that only finitely many of them are 'actually existing/

present' at each stage of the evolution of the system. Therefore, at each transition round, only finitely many currently existing agents can possibly perform an action, and each of these may also choose not to perform any action (i.e., remain inactive in that round). However, the currently inactive (or, 'non-existing') agents do not have any individual influence on the transitions. Thus, the number of currently active agents, who determine the next transition, can change from any instant to the next one, while always remaining finite. We note, however, the difference between dynamic systems, in the sense described above, and simply *parametric* systems, where the number of agents is taken as a parameter but remains fixed during the whole evolution of the system. In that sense, the present study applies both to parametric and truly dynamic systems.

In this work we develop a logic-based framework for formal specification and algorithmic verification of the behaviour of homogeneous dynamic multi-agent systems (HDMAS) of the type described above. We focus, in particular, on scenarios where the agents are divided into *controllable* (by the system supervisor or controller) and *uncontrollable*, representing the environment or an adversary. Both numbers, of controllable and uncontrollable agents, may be fixed or varying throughout the system evolution, possibly at every transition. The controllable agents are assumed to act according to a joint strategy prescribed by the supervisor/controller, with the objective to ensure the desired behaviour of the system (e.g. reaching an outcome in the voting procedure, or keeping the demand and supply of a given stock within desired bounds, or ensuring that the server will not be deadlocked by a malicious attack of adversary users, etc).

As a logical language for formal specification we introduce a suitably extended version, $\mathcal{L}_{\text{HDMAS}}$, of the alternating time temporal logic (ATL) [1]. In $\mathcal{L}_{\text{HDMAS}}$ one can specify properties of the type "A team of (at least) n controllable agents can ensure, against at most m active uncontrollable agents, that any possible evolution of the system satisfies a given objective φ ", where the objective φ is specified again as a formula of that language, and each of n and m is either a fixed number, a parameter, or a variable that can be quantified over.

To summarise the comparison: in the standard concurrent game models of MAS agents are explicitly distinguished and in the logic ATL they are explicitly referred to by their names (individually, or in coalitions). In the HDMAS framework developed here, the only distinction between the agents is whether they are controllable or not, and in the language both are referred to only by numbers.

Here is an indicative, yet generic scenario, where our framework is readily applicable for both modelling and verification.

A military fortress has k protected points of entry: $A_1, A_2 \dots A_k$, with $k > 2$. The commander of the fortress has C soldiers, hereafter called 'defenders', that can be deployed to protect these points of entry against an invading army. For each A_i , a number c_i of defenders, with $m_i \leq c_i \leq M_i$, can be deployed against n_i 'invaders'. If $c_i = M_i$, then the defenders successfully protect A_i against any number of invaders; if $c_i \neq M_i$, then entry point A_i is lost when $n_i > m_i$. Moreover, both the defender and the invading commander may receive reinforcements and re-deploy their soldiers among the entry points once a day (say, at noon), whereas the attacks can only take place at night. However, neither of them can observe the precise distribution of the soldiers of the other party, but they can observe which points of entry are currently "outpowered" by not being sufficiently protected by defenders. It is also known that the enemy must outpower more than 2 points of entry at the same time in order to successfully invade the fortress.

The framework HDMAS that we develop here will enable modelling the scenario above as well as specifying and algorithmically verifying claims of the kind: "The fortress

commander has a strategy to protect the fortress for at least d days, for a given d (or, forever) with C defenders against at most V (or, against any number of) invaders”.

1.2 Structure and content of the paper

In Sect. 2 we introduce the HDMAS framework, provide a running example, and prove some technical results needed to introduce *counting abstractions* of joint actions and strategy profiles. Using these counting abstractions, in Sect. 3 we provide formal semantics in HDMAS models for the logic $\mathcal{L}_{\text{HDMAS}}$ which we introduce there. We then define normal form of formulae of $\mathcal{L}_{\text{HDMAS}}$ and the fragment $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$, consisting of formulae in normal form. The key technical result obtained in that section is that every formula in $\mathcal{L}_{\text{HDMAS}}$ is equivalent on finite models to one in $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$. In Sect. 4 we develop an algorithm for global model checking of formulae in $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ in finite HDMAS models, which invokes model checking truth of their respective translations into Presburger formulae, and illustrate that algorithm on running examples. In Sect. 5 we establish some refined complexity estimates for the model checking algorithm, using recent complexity results obtained in [11] for fragments of Presburger arithmetic. We end with some concluding remarks on extensions and possible applications of our work in Sect. 6.

1.3 Related work

A more closely related framework to ours is Open Multi-Agent Systems (OMAS) [17]. HDMAS shares with it the characteristic ‘dynamic’ feature of agents, which can therefore leave and join the system at runtime. However HDMAS differs from OMAS in several essential aspects. First, although any finite number of agents can perform actions at each step, the evolution of OMAS depends only on the projection of those on the set of actions or, in other words, whether any action is performed by at least one agent. Thus, HDMAS makes use of the full expressivity of Presburger arithmetic. Next, the verification formalism of OMAS is a temporal epistemic logic with (universally quantified) indices spanning over agents, while ours includes strategic operators. Lastly, decidability of model-checking Open Multi-agent Systems is obtained by restricting the semantics of the models and by using cutoff techniques whereas we ultimately invoke model-checking truth of Presburger formulas.

We are aware of other threads of, more or less essentially, related work, however none of them considers formal models and verification methods for the type of homogeneous and dynamic multi-agent scenarios studied here. Therefore, we only mention them briefly as in all frameworks mentioned below, the number of agents is fixed along system executions, possibly as a parameter and the formal specification languages do not explicitly allow quantification over the number of agents.

- Counting abstraction for verification of parametric systems has been studied in [10] and [4], where techniques based on Petri nets or Vector Addition Systems with States (VASS) are used to obtain decidability of model checking.
- The work in [19] is closer to ours, as strategic reasoning is considered but only for a restricted set of properties such as reachability, coverability and deadlock avoidance. Also, assumptions on the system evolutions are made and, in particular, monotonicity with respect to a well-quasi-ordering.
- In [15] temporal epistemic properties of parametric interpreted systems are checked irrespective of the number of agents by using cutoff techniques.

- Modular Interpreted Systems [13] is a MAS framework where a decoupling between local agents and global system description is achieved, thus possibly amenable to model dynamical MAS frameworks.
- Homogeneous MAS with transitions determined by the number of acting agents have been introduced in [18].
- Population protocols [2] are parametric systems of homogeneous agents, and decidability of model checking against probabilistic linear-time specification is studied in [9].
- In [7], instead of verifying MAS with unknown number of agents, the authors propose a technique to find the minimal number of agents which, once deployed and suitably orchestrated, can carry out a manufacturing task.
- Lastly, as noted above, our logic of specification builds on the Alternating time temporal logic ATL ([1]) and extends the model checking algorithm for ATL to HDMAS .

2 Preliminaries and modelling framework

We start by introducing the basic ingredients of our framework. We assume a hereafter fixed (finite, or possibly countably infinite) **universe of potential agents** $Ag = \{ag_1, ag_2, \dots\}$, but only finitely many of them will be assumed currently present, or ‘currently existing’, at any time instant or stage of the evolution of the system. Alternatively, the universe of agents can be assumed always finite but unbounded.

Next, we consider a finite set of **action names** $Act = \{act_1, \dots, act_n\}$. We extend this set with a specific ‘idle’ action ε and define $Act^+ = Act \cup \{\varepsilon\}$. We also fix a set of distinct variables $X = \{x_1, \dots, x_n\}$ extended to $X^+ = X \cup \{x_\varepsilon\}$, called **action counters**, associated to Act and Act^+ respectively. Formally, we relate these by a mapping $\mu : Act^+ \rightarrow X^+$ such that for each $i \in \{1, \dots, n\}$, $\mu(act_i) = x_i$ and $\mu(\varepsilon) = x_\varepsilon$. Hereafter, Act , Act^+ , X , X^+ , and μ are assumed fixed, as above.

An **action profile** over a given set of actions $Act' \subseteq Act^+$ is defined as a function $p : Ag \rightarrow Act'$, assigning an action from Act' to each agent in Ag . More generally, for any subset of agents $A \subseteq Ag$, a **joint action of A** over a set of actions $Act' \subseteq Act^+$ is a function p_A assigning an action from Act' to each agent in A .

Given a function f , we will write: $\text{dom}(f)$ for the domain of f ; $f|_Z$ for the restriction of f to a domain $Z \subseteq \text{dom}(f)$; and $f[Z]$ for the image of Z under f . For technical purposes, we also consider a (unique) function $f_\emptyset$ with an empty domain.

To express relevant conditions on the number of agents performing actions in X , we make use of Presburger arithmetic (the first-order theory of natural numbers with addition and $=$). This is a fairly expressive, yet decidable theory, which makes it very natural and suitable for many computational tasks related to verification of various discrete infinite-state systems (see e.g. [12] for an introduction.)

Definition 1 (*Guards*) A **(transition) guard** g is an open (quantifier-free)¹ formula of Presburger arithmetic PrA with predicates $=$ and $<$ over variables from the set of action counters X . We denote by G the set of all guards, by $\text{Var}(g)$ the set of variables occurring

¹ The restriction to quantifier-free guards is only partly essential for the technical results, given the quantifier elimination property of Presburger arithmetic. We make that restriction mainly to keep the presentation simpler.

in a guard $g \in G$, and we use the following standard abbreviations in Presburger formulas: $n := 1 + \dots + 1$ (n times 1) and $nx := x + \dots + x$ (n times x) for any $n \in \mathbb{N}$ and $x \in X^+$.

Definition 2 An **action distribution** is any function $\mathbf{act} : X' \rightarrow \mathbb{N}$, where $X' \subseteq X^+$. The domain X' is denoted, as usual, by $\text{dom}(\mathbf{act})$. Intuitively, an action distribution assigns for every action act , through the value of the action counter $\mu(act)$, the number of agents who are assigned the action act .

Given an action distribution $\mathbf{act}$ we define:

- $\mathbf{act} \models g$, for a given guard g , if $\mathbf{act}$ satisfies g with the expected standard semantics of PrA, namely:
 $\mathbf{act} \models x_1 = x_2$ if $\mu(x_1) = \mu(x_2)$ and $\mathbf{act} \models x_1 < x_2$ if $\mu(x_1) < \mu(x_2)$;
- $\text{sum}(\mathbf{act}) := \sum_{x \in \text{dom}(\mathbf{act})} \mathbf{act}(x)$;
- $H^m := \{\mathbf{act} \mid \text{sum}(\mathbf{act}) = m\}$ is the set of action distributions where exactly m agents perform actions;
- $H := \bigcup_{m \in \mathbb{N}} H^m$ is the set of all action distributions.

We also define the mapping $\oplus : H \times H \rightarrow H$, which, given two action distributions $\mathbf{act}_1$ and $\mathbf{act}_2$, is defined if $\text{dom}(\mathbf{act}_1) = \text{dom}(\mathbf{act}_2) := Z$ and returns a new action distribution, $\mathbf{act}_1 \oplus \mathbf{act}_2$, with domain Z , defined component-wise as the sum of $\mathbf{act}_1$ and $\mathbf{act}_2$, i.e. $\mathbf{act}_1 \oplus \mathbf{act}_2(z) = \mathbf{act}_1(z) + \mathbf{act}_2(z)$ for each $z \in Z$.

Remark 1 Note that guards are defined over the set of variables X , while the domain of action distributions can also include x_e . It follows that, for any action distribution $\mathbf{act}$, the value $\mathbf{act}(x_e)$ does not have any influence on the satisfiability of a guard. More generally, for every $\mathbf{act} \in H$ and $g \in G$ we have $\mathbf{act} \models g$ iff $\mathbf{act}|_{\text{Var}(g)} \models g$.

We now relate action profiles with action distributions. Every action profile is associated with the action distribution that counts, for each action, the number of agents performing it. In that sense, action distributions are *counting abstractions* for action profiles. The formal definition follows, where we denote the set of all action profiles over Act by P and define the inverse of an action profile p as the function $p^{-1} : Act \rightarrow \wp(Ag)$ such that $p^{-1}(act) = \{ag \in Ag \mid p(ag) = act\}$.

Definition 3 The **action profile abstraction** is the function $\alpha : P \rightarrow H$ where $\alpha(p)(\mu(act)) = |p^{-1}(act)|$ for all $p \in P$ and $act \in Act^+$.

The function α partitions the set P into equivalence classes of action profiles having the same abstraction that is, two action profiles p_1 and p_2 belongs to the same equivalence class iff $\alpha(p_1) = \alpha(p_2)$.

We now introduce the abstract models of our framework.

Definition 4 A **homogeneous dynamic MAS (HDMAS)** is a structure $\mathcal{M} = \langle Ag, Act^+, S, d, \delta, AP, \lambda \rangle$ where:

- $Ag = \{ag_1, ag_2, \dots\}$ is the countable set of **agents**.
- Act^+ is the set of **action names**;

Fig. 1 The fortress example modelled as a hdmass

- S is a set of **states**;²
- $d : S \rightarrow \wp(Act^+)$ is the **action availability function**, that assigns to every state s the set of actions $d(s)$ available (to all agents) at s , and is such that $\epsilon \in d(s)$;
- $\delta : S \times S \rightarrow G$ is the **transitions guard function**, labelling possible transitions between states with guards such that:
 - $Var(\delta(s, s')) \subseteq \mu[d(s)]$ for each $s, s' \in S$ (the guards at each state only involve action counters corresponding to actions available at that state),
 - and, for each $s \in S$ and for each **act** $\in H|_{\mu[d(s)]}$, there exists a unique $s' \in S$ such that **act** $\models \delta(s, s')$ (every possible action distribution over the set of actions available at the current state determines a unique transition).
- $AP = \{p_1, p_2, \dots\}$ is a finite set of **atomic propositions**;
- $\lambda : S \rightarrow \wp(AP)$ is a **labelling function**, assigning to any state s the set of atomic propositions that are true at s .

Example 1 The fortress example presented in the introduction, with $k = 3$ entry points, can be modeled as a hdmass as follows. The set S contain two states only, displayed as circles in Fig. 1: s_1 and s_2 represents respectively the fortress being under control of the defenders or being captured. Next, we have two actions for each entry point A_i : one modelling the defensive action act_i , and the other the attacking action $\overline{act}_i$ for A_i ; therefore $Act^+ = \{act_1, \overline{act}_1, act_2, \overline{act}_2, act_3, \overline{act}_3, \epsilon\}$, with $\mu(act_i) = x_i$ and $\mu(\overline{act}_i) = \bar{x}_i$ for $i \in \{1, 2, 3\}$. All of them are allowed in s_1 and none of them in s_2 , formally: $d(s_1) = Act^+$ and $d(s_2) = \{\epsilon\}$. The guards g_1, g_2 are listed next to the picture, and an arrow is drawn from s_i to s_j and labeled with g_k iff $\delta(s_i, s_j) = g_k$. Formula g_1 guards transition from s_1 to s_2 and therefore it defines when the fortress is captured. This happens when, for each of the entry point A_i with $i \in \{1, 2, 3\}$, one of two conditions hold: 1) the number of defenders x_i is less than m_i or 2) it is less than M_i and also less than the number of attackers $\bar{x}_i$. If this is not the case, the defenders hold the fortress (loop in s_1) but once is conquered, it remains so regardless of the actions performed (g_2 is a tautology). The label of each state, as defined by the labelling function, is given next to it. We only have one atomic proposition, *captured*, false in s_1 and true in s_2 , therefore $\lambda(s_1) = \emptyset$ and $\lambda(s_2) = \{captured\}$.

Example 2 A more abstract example is given in Fig. 2, which will be used to illustrate some technical points and the model checking algorithm later. The set of actions is $Act = \{act_1, act_2, act_3\}$ and the action availability function is defined by $d(s_1) = d(s_3) = d(s_4) = Act^+$, $d(s_2) = \{act_1, act_3, \epsilon\}$, $d(s_5) = \{act_2, act_3, \epsilon\}$

² Note that S is not required in general to be finite, but some of our technical results will assume finiteness.

Fig. 2 An abstract example of a hdmass

and $d(s_6) = \{act_1, \epsilon\}$. Lastly, the labelling function is defined as: $\lambda(s_1) = \emptyset$, $\lambda(s_2) = \lambda(s_3) = \lambda(s_4) = \{p\}$ and $\lambda(s_5) = \lambda(s_6) = \{q\}$.

The restriction on δ ensures that for any number of agents and their action profile of available actions, the next state is uniquely defined. Thus, the dynamics of the system in terms of possible state transitions is fully determined symbolically by the transitions guard function δ , as defined formally below.

Definition 5 Given a hdmass $\mathcal{M}$, a **transition** in $\mathcal{M}$ is a triple (s, p, s') , where $s, s' \in S$ and $p \in P$, such that:

1. each agent ag performs an available action: $p(ag) \in d(s)$;
2. the abstraction $\alpha(p)$ satisfies the (unique) guard that labels the transition from s to s' , i.e., $\alpha(p) \models \delta(s, s')$.

Since transitions only depend on the abstractions of the action profiles, that is, on action distributions, it is immediate to see that actions profiles with the same abstraction, applied at the same state, lead to the same successor state. Formally, the following holds.

Lemma 1 Given a hdmass $\mathcal{M}$ as above, for every $s, s' \in S$, and every $p_1, p_2 \in P$, if $\alpha(p_1) = \alpha(p_2)$, then (s, p_1, s') is a transition in $\mathcal{M}$ iff (s, p_2, s') is a transition in $\mathcal{M}$.

Lemma 1 enables us to define the transition function³ of $\mathcal{M}$ directly on action distributions, rather than on action profiles.

Definition 6 Let $\mathcal{M}$ be a hdmass. The **transition function** of $\mathcal{M}$ is the partial mapping $\Delta : S \times H \rightarrow S$ defined as follows. For each $s \in S$ and $\mathbf{act} \in H$, the **outcome state** $\Delta(s, \mathbf{act})$ of $\mathbf{act}$ at s is defined and equal to $s' \in S$ iff there exists $p \in P$ such that (s, p, s') is a transition and $\alpha(p) = \mathbf{act}$; otherwise $\Delta(s, \mathbf{act})$ is undefined.

³ We remark that the assumption of determinism of hdmass is common in the study of multi-agent systems, because non-determinism can be settled easily by the actions of a fictitious new agent (Nature). Intuitively, one can transform a nondeterministic hdmass to a deterministic one by adding actions that resolve the non-determinism. Then specifications can be translated from the latter to the former by adding controllable or non-controllable agents that could execute these actions.

Infinite sequences of successor states will be called ‘plays’. Formally, a **play** is a sequence $\pi = s_0, s_1, \dots$ in S^ω , such that for every **stage** (of the play) $i \in \mathbb{N}$, there is $\mathbf{act}_i \in H$ such that $\Delta(s_i, \mathbf{act}_i) = s_{i+1}$. We denote by $\pi[i]$ the state of the i th stage of the play, for each $i \in \mathbb{N}$.

Since transitions from a given state s are defined only for action profiles that assigns to all agents only actions that are available at s , we call these **available action profiles** in s . We formally define for each state $s \in S$ the set of available action profiles in s as

$$P_s = \{p \in P \mid p(ag) \in d(s) \text{ for each } ag \in Ag\}.$$

More generally, for each set of agents $A \subseteq Ag$ we define likewise the set of **joint actions** for A available in s as

$$P_s|_A = \{p_A \in P_A \mid p_A(ag) \in d(s) \text{ for each } ag \in A\}.$$

where P_A denotes (with a mild abuse of notation) the set of all possible joint actions for A .

Next, we define a positional strategy for a given coalition of agents A as a mapping that assigns to each state s an available joint action for A .

Definition 7 Let A be a (possibly empty) set of agents and $\mathcal{M}$ be a hDMAS with a state space S . A **joint (positional) strategy for the coalition** A is a function $\sigma_A : S \rightarrow P|_A$ such that $\sigma_A(s) \in P_s|_A$ for each $s \in S$. The empty coalition has only one joint strategy $\sigma_\emptyset$, assigning the empty joint action at every state.

Hereafter we assume that at every stage of the play representing the evolution of the system, the set of all currently present agents is partitioned into two: the set of *controllable agents*, denoted by C , and the set of *uncontrollable agents*, denoted by N . Neither of these subsets (and their sizes) is fixed initially, nor during the play, but each of them can vary at each transition round.

Definition 8 Let $\mathcal{M}$ be a hDMAS, $s \in S$ be a state in it, $C, N \subseteq Ag$ be the respective current sets of controllable and uncontrollable agents, and let $p_C \in P_s|_C$. The **outcome set of p_C at s** is defined as follows:

$$\begin{aligned} out(s, p_C, N) := \{s' \in S \mid s' = \Delta(s, \alpha(p)) \text{ for some } p \in P_s|_{(C \cup N)} \\ \text{such that } p|_C = p_C\}. \end{aligned}$$

Respectively, given a joint strategy σ_C for C we define the **set of outcome plays of p_C at s (against N)** as

$$\begin{aligned} out(s, \sigma_C, N) := \{\pi = s_0, s_1, \dots \mid s_0 = s \text{ and for all } i \in \mathbb{N} \text{ there exists } p_i \in P_{s_i}|_{(C \cup N)} \\ \text{such that } p_i|_C = \sigma_C(s_i) \text{ and } \Delta(s_i, \alpha(p_i)) = s_{i+1}\}. \end{aligned}$$

The abstraction α , although defined on actions profiles, is readily extended over joint actions and naturally specifies an equivalence relation between them: two joint actions are equivalent whenever their abstraction is the same. Likewise for joint strategies, as the next definition formalizes.

Definition 9 Let $\mathcal{M}$ be a hDMAS, $C_1, C_2 \subseteq Ag$ and p_{C_1}, p_{C_2} be respective joint actions for C_1 and C_2 . We say that p_{C_1} and p_{C_2} are **equivalent**, denoted $p_{C_1} \equiv p_{C_2}$, if $\alpha(p_{C_1}) = \alpha(p_{C_2})$.

Likewise, we say that joint strategies σ_{C_1} and σ_{C_2} are **equivalent**, denoted $\sigma_{C_1} \equiv \sigma_{C_2}$ if they prescribe equivalent joint actions for C_1 and C_2 at every state.

Note that if $p_{C_1} \equiv p_{C_2}$ then $|C_1| = |C_2|$ and p_{C_1} and p_{C_2} produce the same outcome sets.

Lemma 2 *Let $\mathcal{M}$ be a HDMAS and $C_1, C_2, N_1, N_2 \subseteq Ag$ be such that, $|C_1| = |C_2|$, $|N_1| = |N_2|$, $C_1 \cap N_1 = \emptyset$, and $C_2 \cap N_2 = \emptyset$. Then:*

1. *For any $s \in S$, if p_{C_1} and p_{C_2} are two equivalent joint actions available at s , respectively for C_1 and C_2 , then $out(s, p_{C_1}, N_1) = out(s, p_{C_2}, N_2)$.*
2. *If σ_{C_1} and σ_{C_2} are two equivalent joint strategies in $\mathcal{M}$, respectively for C_1 and C_2 , then for each $s \in S$, $out(s, \sigma_{C_1}, N_1) = out(s, \sigma_{C_2}, N_2)$.*

Proof (1) Let $s' \in out(s, p_{C_1}, N_1)$. Then $s' = \Delta(s, \alpha(p_1))$ for some $p_1 \in P_s|_{(C_1 \cup N_1)}$ such that $p_1|_{C_1} = p_{C_1}$. Fix a bijection $h : C_2 \rightarrow C_1$. It can be extended to a bijection $f : Ag \rightarrow Ag$, such that $f[N_2] = N_1$.

Define $p_2 \in P_s|_{(C_2 \cup N_2)}$ so that $p_2(ag) := p_1(f(ag))$. Clearly, $\alpha(p_2) = \alpha(p_1)$. Also $p_2|_{C_2} = p_1|_{f[C_2]}$ as $f[C_2] = C_1$, hence $\alpha(p_2|_{C_2}) = \alpha(p_1|_{C_1}) = \alpha(p_{C_1}) = \alpha(p_{C_2})$ (since $p_{C_1} \equiv p_{C_2}$).

Therefore, we obtain that $s' = \Delta(s, \alpha(p_2)) \in out(s, p_{C_2}, N_2)$.

Thus, $out(s, p_{C_1}, N_1) \subseteq out(s, p_{C_2}, N_2)$.

The proof of the converse inclusion is completely symmetric.

(2) The claim follows easily by using (1). Indeed, every play $\pi = s_0, s_1, \dots$ in $out(s, \sigma_{C_1}, N_1)$ can be generated step-by-step as a play in $out(s, \sigma_{C_2}, N_2)$, by using the equivalence of σ_{C_1} and σ_{C_2} and applying (1) at every step of the construction. We leave out the routine details.

Thus, $out(s, \sigma_{C_1}, N_1) \subseteq out(s, \sigma_{C_2}, N_2)$. Again, the converse inclusion is completely symmetric. $\square$

We now prove that, as expected, the outcome sets from joint actions and strategies do not depend on the actual sets of controllable and uncontrollable agents, but only on their sizes.

Lemma 3 *Let $\mathcal{M}$ be a HDMAS, $s \in S$, with $C, N \subseteq Ag$ be the respective current sets of controllable and uncontrollable agents (hence, assumed disjoint), and let $p_C \in P_s|_C$ be an available joint action for C at s . Then for every $C' \subseteq Ag$ such that $|C'| = |C|$ there exists an available joint action $p_{C'}$ for C' at s , such that for every $N' \subseteq Ag$ where $C' \cap N' = \emptyset$, if $|N'| = |N|$, then $out(s, p_{C'}, N') = out(s, p_C, N)$.*

Proof Fix any $C' \subseteq Ag$ such that $|C'| = |C|$. Take a bijection $h : C' \rightarrow C$. It transforms canonically the joint action p_C to a joint action $p_{C'}$ available at s , defined by $p_{C'}(ag) := p_C(h(ag))$. Clearly, $\alpha(p_{C'}) = \alpha(p_C)$. Hence, by Lemma 2, $out(s, p_{C'}, N') = out(s, p_C, N)$ for every $N' \subseteq Ag$ such that $C' \cap N' = \emptyset$ and $|N'| = |N|$. $\square$

Lemma 3 easily extends to joint strategies, as follows.

Lemma 4 Let $\mathcal{M}$ be a HDMAS, $s \in S$, with $C, N \subseteq Ag$ be the respective current (disjoint) sets of controllable and uncontrollable agents, and let σ_C be a joint strategy for C . Then for every $C' \subseteq Ag$ with $|C'| = |C|$ there exists a joint strategy $\sigma_{C'}$ such that for every $N' \subseteq Ag$ where $C' \cap N' = \emptyset$, if $|N'| = |N|$, then $out(s, \sigma_{C'}, N') = out(s, \sigma_C, N)$.

Proof The argument is similar to the previous proof.

Fix any $C' \subseteq Ag$ such that $|C'| = |C|$. Take a bijection $h : C' \rightarrow C$. It transforms canonically the joint strategy σ_C to a joint strategy $\sigma_{C'}$, defined by $\sigma_{C'}(s)(ag) := \sigma_C(s)(h(ag))$.

Clearly, $\alpha(\sigma_{C'}(s)) = \alpha(\sigma_C(s))$ for every state s , hence $\sigma_C \equiv \sigma_{C'}$. Therefore, by Lemma 2, $out(s, \sigma_{C'}, N') = out(s, \sigma_C, N)$ for every $N' \subseteq Ag$ such that $C' \cap N' = \emptyset$ and $|N'| = |N|$. $\square$

Lemmas 3 and 4 essentially say that the strategic abilities in a HDMAS are determined not by the concrete sets of controllable and uncontrollable agents, but only by their respective sizes. This justifies abstracting the notions of coalitional actions and strategies in terms of action profile abstractions, to be used thereafter in our semantics and verification procedures.

Definition 10 Let $\mathcal{M}$ be a HDMAS and $C, N \in \mathbb{N}$.

1.1. An **abstract joint action for a coalition of C agents at state $s \in S$** is an action distribution $\mathbf{act}_C \in H|^C$ such that $dom(\mathbf{act}_C) = \mu[d(s)]$ (recall notation from Definition 2).

Thus, an abstract joint action for a given coalition at state s prescribes for each action available at s how many agents from the coalition take that action.

1.2. The **outcome set of states of the abstract joint action $\mathbf{act}_C$ of C controllable agents against N uncontrollable agents at s** is the set of states

$$out(s, \mathbf{act}_C, N) := \{s' \in S \mid s' = \Delta(s, \mathbf{act}_C \oplus \mathbf{act}_N) \text{ for some } \mathbf{act}_N \in H|^N \text{ such that } dom(\mathbf{act}_N) = \mu[d(s)]\}.$$

2.1. An **abstract (positional) joint strategy for a coalition of C agents** is a function $\rho_C : S \rightarrow H|^C$ such that for each $s \in S$, $\rho_C(s)$ is an abstract joint action such that $dom(\rho_C(s)) = \mu[d(s)]$.

2.2. The **outcome set of plays of an abstract joint strategy ρ_C of C controllable agents against N uncontrollable agents** is the set of plays

$$out(s, \rho_C, N) := \{\pi = s_0, s_1, \dots \mid s_0 = s \text{ and for all } i \in \mathbb{N} \text{ there is } \mathbf{act}_i \in H|^N \text{ such that } dom(\mathbf{act}_i) = \mu[d(s)] \text{ and } \Delta(s_i, \rho_C(s_i) \oplus \mathbf{act}_i) = s_{i+1}\}.$$

3 Logic for specification and verification of HDMAS

We now introduce a logic $\mathcal{L}_{\text{HDMAS}}$ for specifying and verifying properties of HDMAS, based on the Alternating-time Temporal Logic ATL. It features a strategic operator that expresses the ability of a set of controllable agents to guarantee the satisfaction a temporal objective, regardless of the actions taken by the set of uncontrollable agents. As shown in the previous section, such ability only depends on the sizes of these sets. Therefore, our strategic operator $\langle\langle *, * \rangle\rangle$ takes two arguments: the first one represent the number of controllable agents and the second—the number of uncontrollable agents currently present in the

system. Intuitively, a formula of the kind $\langle\langle C, N \rangle\rangle \chi$, with $C, N \in \mathbb{N}$ and χ being a (path) formula of $\mathcal{L}_{\text{HDMAS}}$ specifies the property:

A coalition of C controllable agents has a joint strategy to guarantee satisfaction of the objective χ against N uncontrollable agents on every play consistent with that strategy.

Each of the arguments C and N may be a concrete number, a parameter, or a variable that can be quantified over. Parameters are free variables that cannot be quantified over, which gives extra expressiveness of the language, because some syntactic restrictions will be imposed on the variables.

3.1 Formal syntax and semantics

We now fix a set of atomic propositions $\Phi = \{p_1, p_2, \dots\}$, a set of two special variables $Y = \{y_1, y_2\}$, ranging over $\mathbb{N}$, which we call **agent counters**. These will represent the numbers of controllable and uncontrollable agents respectively, and can be quantified over. We also fix a set of **agent counting parameters**⁴ $Z = \{z_1, z_2, \dots\}$, again ranging over $\mathbb{N}$, and define the set of **terms**⁵ as $T = Y \cup Z \cup \mathbb{N}$. These will be used as arguments of the strategic operators in the logical language defined below.

Definition 11 The logic $\mathcal{L}_{\text{HDMAS}}$ has two sorts of formulae, defined by mutual induction with the following grammars, where free (and bound) occurrences of variables are defined like in first-order logic (FOL):

Path formulae: $\chi ::= X\varphi \mid G\varphi \mid \psi \cup \varphi$,

where φ, ψ are state formulae.

State formulae:

$\varphi ::= \top \mid p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \langle\langle t_1, t_2 \rangle\rangle \chi \mid \forall y\varphi \mid \exists y\varphi$

where $p \in \Phi$, $t_1 \in T \setminus \{y_2\}$, $t_2 \in T \setminus \{y_1\}$, $y \in Y$, and χ is a path formula.

The cases of $\forall y\varphi$ and $\exists y\varphi$ are subject to the following syntactic constraint: all free occurrences of y in φ must have a positive polarity, viz. must be in the scope of an even number of negations.

The propositional connectives $\perp, \rightarrow, \leftrightarrow$ are defined as usual. Also, we define $F\psi := \top \cup \psi$.

Remark 2 Some remarks on the formulae in $\mathcal{L}_{\text{HDMAS}}$ are in order:

1. Note that y_1 can only occur in the first position of $\langle\langle t_1, t_2 \rangle\rangle$ and y_2 can only occur in the second position. However, the same parameter z may occur in both positions and this is one reason to allow the use of parameters, as the model checking algorithm will treat them uniformly.
2. The restriction for quantification only over positive free occurrences of variables is imposed for technical reasons. By using the duality of $\forall$ and $\exists$, that restriction can readily be relaxed to the requirement all free occurrences of the quantified variable to be of *the same* polarity (all positive, or all negative). Further relaxation, allowing both positive and negative occurrences under some restrictions, is possible, but it would complicate

⁴ The role of parameters is mostly auxiliary, just like in algebra (or in first-order logic) and will be discussed further.

⁵ To avoid cluttering the notation, we will identify here natural numbers with their numerals.

- further the syntax and the model checking algorithm, without making an essential contribution to the *useful* expressiveness of the language. Indeed, one can argue that, if a formula is to make a meaningful claim about the strategic abilities of the coalition of controllable agents which is quantified over the number of these agents, then it is natural to assume that the controllable coalition appear only in positive context in that claim⁶.
3. Some additional useful syntactic restrictions can be imposed, which (as it will be shown in the next section) do not essentially restrict the expressiveness of the language. They lead to the notion of ‘normal form’, to be introduced shortly.

Hereafter, by $\mathcal{L}_{\text{HDMAS}}$ -formulae we will mean, unless otherwise specified, state formulae of $\mathcal{L}_{\text{HDMAS}}$, whereas we will call the path formulae in $\mathcal{L}_{\text{HDMAS}}$ **temporal objectives**. In particular, for any $\mathcal{L}_{\text{HDMAS}}$ -formula ϕ of the type $\langle\langle t_1, t_2 \rangle\rangle \chi$, the path subformula χ is called the **temporal objective of ϕ** .

Some examples of $\mathcal{L}_{\text{HDMAS}}$ formulae:

- with reference to the fortress example:
 - $\langle\langle C, N_1 \rangle\rangle X \langle\langle C, N_2 \rangle\rangle X \langle\langle C, N_3 \rangle\rangle X \neg \text{captured}$, with $N_1 < N_2 < N_3$ and $N_i \in \mathbb{N}$ for $i = \{1, 2, 3\}$, says that there is a strategy for $C \in \mathbb{N}$ defenders to hold the fortress for three days against an increasing number of attackers.
 - $\exists y_1 \langle\langle y_1, N \rangle\rangle G \neg \text{captured}$ expresses that there is a number y_1 of defenders that have a strategy to hold the fortress forever against N many invaders.
 - $\forall y_2 \langle\langle C, y_2 \rangle\rangle G \neg \text{captured}$ expresses that, for any number of invaders y_2 , there is a strategy for C defenders to hold the fortress forever against y_2 invaders⁷.
 - $\forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle G \neg \text{captured}$ expresses that for any number (y_2) of invaders there is a number (y_1) of defenders who have a joint strategy to hold the fortress forever.
- lastly, an abstract example with nesting of strategic operators and quantifiers:

$$\langle\langle z_2, z_2 \rangle\rangle X p \vee \exists y_1 (\langle\langle y_1, z_1 \rangle\rangle F \langle\langle y_1, y_2 \rangle\rangle X \neg p \wedge \neg \forall y_2 \langle\langle z_1, y_2 \rangle\rangle X \neg \langle\langle y_1, z_2 \rangle\rangle p \cup q),$$
 for $z_1, z_2 \in \mathbb{Z}$.

The semantics of $\mathcal{L}_{\text{HDMAS}}$ is based on the standard, positional strategy semantics of ATL (cf [1] or [6]), applied in HDMAS models, but uses abstract joint actions and strategy profiles, rather than concrete ones. In order to evaluate formulae that contain free variables and parameters, we use a version of FOL **assignment**, here defined as a function $\theta : T \rightarrow \mathbb{N}$, where $\theta(i) = i$ for $i \in \mathbb{N}$.

Definition 12 Let $\mathcal{M}$ be a HDMAS , s be a state and θ an assignment in it. The **satisfaction relation** $\models$ is inductively defined on the structure of $\mathcal{L}_{\text{HDMAS}}$ -formulae as follows:

⁶ For instance, $\exists y_1 \langle\langle y_1, 10 \rangle\rangle X \neg \exists y_1 \langle\langle y_1, 5 \rangle\rangle p$, where y_1 appear first positively and then negatively, expresses that “there exists a coalition of controllable agents that can ensure against 10 uncontrollable agents that at the next step there is no coalition of controllable agents that can ensure against 5 uncontrollable agents the truth of p ”, which is a rather unusual specification in any practical context.

⁷ We note that such a strategy would generally depend on y_2 . However, given the monotonicity properties that our logic enjoys (proved in the next section), it turns out that the above reading of the formula is equivalent on finite models to “there exists a strategy for C defenders to hold the fortress against any number y_2 of attackers”.

1. $\mathcal{M}, s, \theta \models \top$;
2. $\mathcal{M}, s, \theta \models p$ iff $p \in \lambda(s)$;
3. $\wedge$ and $\neg$ have the standard semantics;
4. $\mathcal{M}, s, \theta \models \langle\langle t_1, t_2 \rangle\rangle \chi$ iff there exists an abstract strategy ρ_C for a coalition of $C = \theta(t_1)$ agents such that for every play π in the outcome set $out(s, \rho_C, N)$ against $N = \theta(t_2)$ uncontrollable agents the following hold:
 - (a) if $\chi = X \varphi$ then $\mathcal{M}, \pi[1], \theta \models \varphi$;
 - (b) if $\chi = G \varphi$ then $\mathcal{M}, \pi[i], \theta \models \varphi$ for every $i \in \mathbb{N}$;
 - (c) if $\chi = \varphi_1 \cup \varphi_2$ then $\mathcal{M}, \pi[i], \theta \models \varphi_2$ for some $i \geq 0$ and $\mathcal{M}, \pi[j], \theta \models \varphi_1$ for all $0 \leq j < i$;
5. $\mathcal{M}, s, \theta \models \forall y \varphi$ iff $\mathcal{M}, s, \theta[y := m] \models \varphi$ for every $m \in \mathbb{N}$, where the assignment $\theta[y := m]$ assigns m to y and agrees with θ on every other argument.
6. Likewise for $\mathcal{M}, s, \theta \models \exists y \varphi$.

The notions of **validity** and **(logical) equivalence** in $\mathcal{L}_{\text{HDMAS}}$ are defined as expected, and we will use the standard notation for them, viz. $\models \varphi$ for validity and $\varphi_1 \equiv \varphi_2$ for equivalence. We also say that two $\mathcal{L}_{\text{HDMAS}}$ -formulae, φ_1 and φ_2 are **equivalent in the finite**, denoted $\varphi_1 \equiv_{\text{fin}} \varphi_2$, if $\mathcal{M}, s, \theta \models \varphi_1$ iff $\mathcal{M}, s, \theta \models \varphi_2$ for any finite HDMAS model $\mathcal{M}$ and state s and assignment θ in $\mathcal{M}$.

Remark 3 Note the following:

1. Defining the semantics in terms of abstract joint actions and strategies in the truth definitions of the strategic operators, rather than concrete ones, is justified by Lemmas 3 and 4 which imply that the ‘concrete’ and the ‘abstract’ semantics are equivalent.
2. Just like in FOL, the truth of any $\mathcal{L}_{\text{HDMAS}}$ -formula φ only depends on the assignment of values to the parameters that occur in φ and to the variables that occur free in φ . In particular, it does not depend at all on the assignment for *closed formulae* (containing no parameters and free variables). In such cases we simply write $\mathcal{M}, s \models \varphi$.
3. Again, just like in FOL, if y has no free occurrences in φ , then $\forall y \varphi \equiv \exists y \varphi \equiv \varphi$. Thus, in order to avoid such vacuous quantification, whenever it occurs we can assume that the formula is simplified automatically according to these equivalences.

Example 3 Consider the HDMAS $\mathcal{M}$ in Example 2.

1. The closed formula $\varphi = \langle\langle 7, 5 \rangle\rangle X p$ is satisfied in state s_1 of $\mathcal{M}$. Indeed, any abstract joint strategy ρ_7 that prescribes ϵ to 3 of the controllable agents ($\rho_7(s_1)(\epsilon) = 3$) and act_3 to 4 of them ($\rho_7(s_1)(act_3) = 4$) guarantees that guard g_2 is satisfied, enforcing transition from s_1 to s_3 .
2. $\mathcal{M}, s_1 \models \neg \exists y_1 \langle\langle y_1, 11 \rangle\rangle X p$. Indeed, for any value of y_1 the abstract joint action profile for the uncontrollable agents that prescribes to all of them to perform act_3 falsifies both g_1 and g_2 , thus forces a loop to s_1 where p is false.
3. $\mathcal{M}, s_4 \models \langle\langle 7, 4 \rangle\rangle X (\forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle G p)$, as we show in Sect. 4.

3.2 Normal form and monotonicity properties

This is a technically important section, where we define the fragment $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ of normal form formulae of $\mathcal{L}_{\text{HDMAS}}$. The normal form impose essential syntactic restrictions and therefore reduce the expressiveness of the language. However, the key technical result obtained here is that every formula in $\mathcal{L}_{\text{HDMAS}}$ is equivalent *on finite models* to one in $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$. The importance of that result will be discussed further.

Definition 13 A $\mathcal{L}_{\text{HDMAS}}$ -formula ψ is **in a normal form** if:

- (NF1) There are no occurrences of $\forall y_1$ or $\exists y_2$ in ψ .
- (NF2) Every subformula $\langle\langle t_1, t_2 \rangle\rangle \chi$ of ψ where either $t_1 = y_1$ or $t_2 = y_2$ (but not both), such that that variable occurrence is bound in ψ , is immediately preceded respectively by $\exists y_1$ or $\forall y_2$.
- (NF3) Every subformula $\langle\langle y_1, y_2 \rangle\rangle \chi$, where both variable occurrences are bound in ψ , is immediately preceded either by $\forall y_2 \exists y_1$ or $\exists y_1 \forall y_2$.

Of the example formulae given after Definition 11, the first two are in normal form, while the last one is not.

We denote by $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ the fragment of $\mathcal{L}_{\text{HDMAS}}$ consisting of all formulae in normal form. We can give a more explicit definition of the formulae of $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$, by modifying the recursive definition of state formulae of $\mathcal{L}_{\text{HDMAS}}$, where the clauses $\forall y\varphi$ and $\exists y\varphi$ are replaced with the following, where χ is a temporal objective:

$$\begin{aligned} & \exists y_1 \langle\langle y_1, t_2 \rangle\rangle \chi \mid \forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle \chi \mid \\ & \forall y_2 \langle\langle t_1, y_2 \rangle\rangle \chi \mid \exists y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle \chi \end{aligned} \quad (1)$$

The same syntactic constraints as before apply. In addition, in each case above no variable quantified in the prefix of the formula may occur free in χ .

The rest of the section is devoted to prove that every formula in $\mathcal{L}_{\text{HDMAS}}$ is logically equivalent in the finite to one in $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$. That is of crucial importance, as our model checking algorithm works only on $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ formulae. Indeed, the fact that quantification in formulae in normal form does not span across multiple temporal objectives enables us to obtain fixpoint characterizations for formulae of the types listed in (1) above, presented at the end of this section, in Theorem 3. That, in turn, allows us to retain the basic structure of the recursive model checking algorithm for ATL (cf [1] or [6]).

A first important observation is that the semantics of the strategic operators in $\mathcal{L}_{\text{HDMAS}}$ is *monotonic* with respect to the number of controllable and uncontrollable agents, in a sense formalized in the following lemma.

Hereafter, for a given formula φ , term t and $k \in \mathbb{N}$, we denote by $\varphi[k/t]$ the result of uniform substitution of all *free*⁸ occurrences of t in φ by k .

Lemma 5 *For every $\mathcal{L}_{\text{HDMAS}}$ -formula φ and a term t the following monotonicity properties hold.*

⁸ The constraint to free occurrences is, of course, only relevant when t is a variable.

(C-mon) Suppose $C, C' \in \mathbb{N}$ are such that $C' > C$. Then:

(C-mon)⁺: If all free occurrences of t are positive and only in first position in strategic operators in φ then $\models \varphi[C/t] \rightarrow \varphi[C'/t]$.

(C-mon)⁻: If all free occurrences of t are negative and only in first position in strategic operators in φ then $\models \varphi[C'/t] \rightarrow \varphi[C/t]$.

(N-mon) Suppose $N, N' \in \mathbb{N}$ are such that $N' < N$. Then:

(N-mon)⁺: If all free occurrences of t are positive and only in second position in strategic operators in φ then $\models \varphi[N/t] \rightarrow \varphi[N'/t]$.

(N-mon)⁻: If all free occurrences of t are negative and only in second position in strategic operators in φ then $\models \varphi[N'/t] \rightarrow \varphi[N/t]$.

Proof (C-mon): Both claims are analogous and we prove both by simultaneous induction on the structure of φ . We will present the proof for **(C-mon)⁺**, and the claim of **(C-mon)⁻** will only be needed in the case when $\varphi = \neg\psi$, proved by using the inductive hypothesis for **(C-mon)⁻** for ψ and contraposition.

The inductive cases where the main connective of φ is $\wedge, \vee, \forall, \exists$ are easily proved by using the inductive hypothesis and the monotonicity of each of these logical connectives.

The only more essential inductive case is $\varphi = \langle\langle t, t_2 \rangle\rangle \chi$, where the inductive hypothesis is that the claim of **(C-mon)⁺** holds for the main state subformulae of χ . Note that the semantics of the strategic and temporal operators is argument-monotone, in sense that if $\models \psi \rightarrow \psi'$ then $\models \langle\langle t, t_2 \rangle\rangle X\psi \rightarrow \langle\langle t, t_2 \rangle\rangle X\psi'$ and $\models \langle\langle t, t_2 \rangle\rangle G\psi \rightarrow \langle\langle t, t_2 \rangle\rangle G\psi'$, and likewise for Until. By using that and the inductive hypothesis, we obtain that $\models \langle\langle t, t_2 \rangle\rangle \chi[C/t] \rightarrow \langle\langle t, t_2 \rangle\rangle \chi[C'/t]$. Therefore, $\models \langle\langle C, t_2 \rangle\rangle \chi[C/t] \rightarrow \langle\langle C', t_2 \rangle\rangle \chi[C'/t]$. Thus, it remains to show that $\models \langle\langle C, t_2 \rangle\rangle \chi[C/t] \rightarrow \langle\langle C', t_2 \rangle\rangle \chi[C/t]$. Let $\mathcal{M}, s, \theta \models \langle\langle C, t_2 \rangle\rangle \chi[C/t]$. Let ρ_C be an abstract strategy for C controllable agents such that every play π in the outcome set $out(s, \rho_C, \theta(t_2))$ against $\theta(t_2)$ uncontrollable agents satisfies the temporal objective $\chi[C/t]$. Then, since $C' > C$, the strategy ρ_C can be extended to strategy $\rho_{C'}$ whereby the additional $C' - C$ many agents always perform the idle action ε . Clearly, $\rho_{C'}$ ensures that $\mathcal{M}, s, \theta \models \langle\langle C', t_2 \rangle\rangle \chi[C/t]$.

(N-mon): The proof is analogous to the one for **(C-mon)**, so we only treat the inductive case of $\varphi = \langle\langle t_1, t \rangle\rangle \chi$ for the claim **(N-mon)⁺**. Similarly to the case of **(C-mon)⁺**, it boils down to proving the validity $\models \langle\langle t_1, N \rangle\rangle \chi[N/t] \rightarrow \langle\langle t_1, N' \rangle\rangle \chi[N/t]$. Let $\mathcal{M}, s, \theta \models \langle\langle t_1, N \rangle\rangle \chi[N/t]$ and let ρ_C be an abstract strategy for $C = \theta(t_1)$ controllable agents such that every play π in the outcome set $out(s, \rho_C, N)$ against N uncontrollable agents satisfies the temporal objective $\chi[N/t]$. Then the same strategy would ensure $\mathcal{M}, s, \theta \models \langle\langle t_1, N' \rangle\rangle \chi[N/t]$ for every $N' < N$, since every joint action of N' can be lifted to

a joint action of N leading to the same outcome, where the remaining $N - N'$ agents always perform the idle action ϵ . $\square$

A key consequence of the monotonicity properties is that it allows to eliminate some quantifier patterns, in the cases listed in the following lemma.

Lemma 6 *For every term t and temporal objective χ in $\mathcal{L}_{\text{HDMAS}}$, the following hold.*

1. $\forall y_1 \langle \langle y_1, t \rangle \rangle \chi \equiv \langle \langle 0, t \rangle \rangle \chi[0/y_1]$;
2. $\exists y_2 \langle \langle t, y_2 \rangle \rangle \chi \equiv \langle \langle t, 0 \rangle \rangle \chi[0/y_2]$;
3. $\forall y_1 \exists y_2 \langle \langle y_1, y_2 \rangle \rangle \chi \equiv \langle \langle 0, 0 \rangle \rangle \chi[0/y_1, 0/y_2]$;
4. $\exists y_2 \forall y_1 \langle \langle y_1, y_2 \rangle \rangle \chi \equiv \langle \langle 0, 0 \rangle \rangle \chi[0/y_1, 0/y_2]$;
5. $\forall y_2 \forall y_1 \langle \langle y_1, y_2 \rangle \rangle \chi \equiv \forall y_2 \langle \langle 0, y_2 \rangle \rangle \chi[0/y_1]$;
6. $\forall y_1 \forall y_2 \langle \langle y_1, y_2 \rangle \rangle \chi \equiv \forall y_2 \langle \langle 0, y_2 \rangle \rangle \chi[0/y_1]$;
7. $\exists y_1 \exists y_2 \langle \langle y_1, y_2 \rangle \rangle \chi \equiv \exists y_1 \langle \langle y_1, 0 \rangle \rangle \chi[0/y_2]$;
8. $\exists y_2 \exists y_1 \langle \langle y_1, y_2 \rangle \rangle \chi \equiv \exists y_1 \langle \langle y_1, 0 \rangle \rangle \chi[0/y_2]$.

Proof The logically non-trivial implications of claims 1–6 follow immediately from the polarity constraint in the definition of formulae and Lemma 5. Claims 7 and 8 follow respectively from claims 5 and 6, by commuting the quantifiers. $\square$

Lemma 6 shows that the only non-trivial cases of quantifications over formulae of the kind $\langle \langle t_1, t_2 \rangle \rangle \chi$ are those allowed in normal forms, listed in (1) (after Definition 13). We will make use of that to re-define the syntax of $\mathcal{L}_{\text{HDMAS}}$ to suit better our further technical work. First we define an **admissible quantifier prefix** Q to be a string of the form Qy_i or $Qy_iQ'y_j$ where $Q, Q' \in \{\exists, \forall\}$ and $i, j \in \{1, 2\}$, $i \neq j$. Now, we re-define the set of state formulae of $\mathcal{L}_{\text{HDMAS}}$ to be generated by the following modified grammar:

$$\varphi ::= \top \mid p \mid \neg \varphi \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid \langle \langle t_1, t_2 \rangle \rangle \chi \mid Q\varphi$$

The same positive polarity requirements as before for applying the quantifier prefixes are imposed. Clearly, this grammar is equivalent to the original grammar, i.e. it generates the same set of formulae. In the rest of the paper we adopt the new grammar above.

Next, we define recursively a **partial quantifier elimination function** PQE on path and state formulae $\xi \in \mathcal{L}_{\text{HDMAS}}$ which produces formulae $\text{PQE}(\xi)$ where all occurrences of subformulae in the left-hand sides of the equivalences in Lemma 6 are successively replaced with the corresponding right-hand sides.

```

1: Input: a quantifier prefix  $Q$  and a (state or path) formula  $\xi$  of  $\mathcal{L}_{\text{HDMAS}}$ 
2: Output: a (state or path) formula  $\xi'$  of  $\mathcal{L}_{\text{HDMAS}}$ 
3: procedure PQE( $\xi$ )
4:   case  $\xi$  of
5:      $\top \mid p$ : return  $\xi$ 
6:      $\neg\psi$ : return  $\neg\text{PQE}(\psi)$ 
7:      $\psi_1 \wedge \psi_2$ : return  $\text{PQE}(\psi_1) \wedge \text{PQE}(\psi_2)$ 
8:      $\psi_1 \vee \psi_2$ : return  $\text{PQE}(\psi_1) \vee \text{PQE}(\psi_2)$ 
9:      $\langle\langle t_1, t_2 \rangle\rangle \chi$ : return  $\langle\langle t_1, t_2 \rangle\rangle \text{PQE}(\chi)$ 
10:     $Q\varphi$ :
11:      case  $Q\varphi$  of
12:         $\forall y_1 \langle\langle y_1, t \rangle\rangle \chi$ : return  $\langle\langle 0, t \rangle\rangle \text{PQE}(\chi)[0/y_1]$ 
13:         $\exists y_2 \langle\langle t, y_2 \rangle\rangle \chi$ : return  $\langle\langle t, 0 \rangle\rangle \text{PQE}(\chi)[0/y_2]$ 
14:         $\forall y_1 \exists y_2 \langle\langle y_1, y_2 \rangle\rangle \chi$ : return  $\langle\langle 0, 0 \rangle\rangle \text{PQE}(\chi)[0/y_1, 0/y_2]$ 
15:         $\exists y_2 \forall y_1 \langle\langle y_1, y_2 \rangle\rangle \chi$ : return  $\langle\langle 0, 0 \rangle\rangle \text{PQE}(\chi)[0/y_1, 0/y_2]$ 
16:         $\forall y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle \chi$ : return  $\forall y_2 \langle\langle 0, y_2 \rangle\rangle \text{PQE}(\chi)[0/y_1]$ 
17:         $\forall y_2 \forall y_1 \langle\langle y_1, y_2 \rangle\rangle \chi$ : return  $\forall y_2 \langle\langle 0, y_2 \rangle\rangle \text{PQE}(\chi)[0/y_1]$ 
18:         $\exists y_1 \exists y_2 \langle\langle y_1, y_2 \rangle\rangle \chi$ : return  $\exists y_1 \langle\langle y_1, 0 \rangle\rangle \text{PQE}(\chi)[0/y_2]$ 
19:         $\exists y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle \chi$ : return  $\exists y_1 \langle\langle y_1, 0 \rangle\rangle \text{PQE}(\chi)[0/y_2]$ 
20:      default: return  $Q\text{PQE}(\varphi)$  ▷ Covers all other cases
21:    end case
22:     $X\psi$ : return  $X\text{PQE}(\psi)$ 
23:     $G\psi$ : return  $G\text{PQE}(\psi)$ 
24:     $\psi_1 \cup \psi_2$ : return  $\text{PQE}(\psi_1) \cup \text{PQE}(\psi_2)$ 
25:  end case
26: end procedure

```

Lemma 7 Let φ be any formula in $\mathcal{L}_{\text{HDMAS}}$ then $\text{PQE}(\varphi) \equiv \varphi$.

Proof By induction on the structure of φ (using the modified grammar), following the recursive definition of PQE. The only non-trivial cases are those in lines 12–19 and they use the equivalences in Lemma 6. For instance, let $\varphi = \forall y_1 \langle\langle y_1, t \rangle\rangle \chi$. By definition, $\text{PQE}(\forall y_1 \langle\langle y_1, t \rangle\rangle \chi) = \langle\langle 0, t \rangle\rangle \text{PQE}(\chi)[0/y_1]$ and by inductive hypothesis $\text{PQE}(\chi) \equiv \chi$, thus we get $\langle\langle 0, t \rangle\rangle \text{PQE}(\chi)[0/y_1] \equiv \langle\langle 0, t \rangle\rangle \chi[0/y_1]$. The claim now follows from case 1. in Lemma 6. All other cases are proved analogously. $\square$

Lemma 8 Let φ be any formula in $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ then $\text{PQE}(\varphi) = \varphi$.

Proof Again, induction on the structure of φ in normal form, following the recursive definition of PQE. Note, that the only cases that apply to formulae in normal form are those in lines 5–9, 20 and 22–24, which do not modify φ . $\square$

Note that after applying PQE, the resulting formula satisfies condition (NF1) in the definition of normal form.

3.3 Transformation to normal forms and fixpoint equivalences

Next, we show that quantification can always be distributed, up to equivalence in the finite, over conjunctions and disjunctions and pushed inside subformulae so that every bound variable is immediately preceded by a quantifier that binds it, which will be used further for transformations of $\mathcal{L}_{\text{HDMAS}}$ formulae to normal form.

We define by recursion a 2-argument function PUSH , applied to pairs consisting of an admissible quantifier prefix Q and a formula φ in $\mathcal{L}_{\text{HDMAS}}$, such that $\text{PUSH}(Q, \varphi)$ is a formula in $\mathcal{L}_{\text{HDMAS}}$ which satisfies conditions (NF2) and (NF3) of the definition of normal form, and which we will prove to be equivalent to $Q\varphi$. For the purpose of defining PUSH as described, we will need to define it on a wider scope, viz. applied to any state or path formula ξ , even though $Q\xi$ may not be a legitimate formula of $\mathcal{L}_{\text{HDMAS}}$.

In what follows, we denote by $\overline{Q}$ the swap of the quantifiers in the prefix Q with their duals, i.e. $\exists$ with $\forall$ and vice versa.

```

1: Input: a quantifier prefix  $Q$  and a (state or path) formula  $\xi$  of  $\mathcal{L}_{\text{HDMAS}}$ 
2: Output:  $\xi'$ , a (state or path) formula of  $\mathcal{L}_{\text{HDMAS}}$ 
3: procedure  $\text{PUSH}(Q, \xi)$ 
4:   case  $\xi$  of
5:      $\top \mid p$ : return  $\xi$ 
6:      $\neg\psi$ : return  $\neg\text{PUSH}(\overline{Q}, \psi)$ 
7:      $\psi_1 \wedge \psi_2$ : return  $\text{PUSH}(Q, \psi_1) \wedge \text{PUSH}(Q, \psi_2)$ 
8:      $\psi_1 \vee \psi_2$ : return  $\text{PUSH}(Q, \psi_1) \vee \text{PUSH}(Q, \psi_2)$ 
9:      $\langle\langle t_1, t_2 \rangle\rangle \chi$ :
10:      case  $Q$  of
11:         $Qy_i$ , where  $t_i = y_i$ , for  $i = 1$  or  $i = 2$ :
12:          or  $Qy_1Q'y_2$  or  $Q'y_2Qy_1$ , where  $t_1 = y_1, t_2 = y_2$ :
13:            return  $Q\langle\langle t_1, t_2 \rangle\rangle \text{PUSH}(Q, \chi)$ 
14:          or  $Qy_i$ , where  $t_i \neq y_i$ , for  $i = 1$  or  $i = 2$ :
15:            or  $Qy_1Q'y_2$  or  $Q'y_2Qy_1$ , where  $t_1 \neq y_1, t_2 \neq y_2$ :
16:              return  $\langle\langle t_1, t_2 \rangle\rangle \text{PUSH}(Q, \chi)$ 
17:            or  $Qy_1Q'y_2$  or  $Q'y_2Qy_1$ , where  $t_1 = y_1, t_2 \neq y_2$ :
18:              return  $Qy_1\langle\langle t_1, t_2 \rangle\rangle \text{PUSH}(Q, \chi)$ 
19:            or  $Qy_1Q'y_2$  or  $Q'y_2Qy_1$ , where  $t_1 \neq y_1, t_2 = y_2$ :
20:              return  $Q'y_2\langle\langle t_1, t_2 \rangle\rangle \text{PUSH}(Q, \chi)$ 
21:      end case
22:      $Q''y_k\psi$ :
23:      case  $Q$  of
24:         $Qy_i$ , where  $i = k$ : return  $\text{PUSH}(Q''y_k, \psi)$ 
25:         $Qy_i$ , where  $i \neq k$ : return  $\text{PUSH}(Qy_iQ''y_k, \psi)$ 
26:         $Qy_iQ'y_j$ , where  $i = k$ : return  $\text{PUSH}(Q'y_jQ''y_k, \psi)$ 
27:         $Qy_iQ'y_j$ , where  $j = k$ : return  $\text{PUSH}(Qy_iQ''y_k, \psi)$ 
28:      end case
29:      $X\psi$ : return  $X\text{PUSH}(Q, \psi)$ 
30:      $G\psi$ : return  $G\text{PUSH}(Q, \psi)$ 
31:      $\psi_1 \cup \psi_2$ : return  $\text{PUSH}(Q, \psi_1) \cup \text{PUSH}(Q, \psi_2)$ 
32:   end case
33: end procedure

```

It is quite easy to see that $\text{PUSH}(\mathcal{Q}, \varphi)$ is a formula in $\mathcal{L}_{\text{HDMAS}}$ whenever $\mathcal{Q}\varphi$ is a formula in $\mathcal{L}_{\text{HDMAS}}$. Intuitively, the function PUSH recursively pushes the quantifier prefix $\mathcal{Q}$ inside the formula by either swapping it when negation occurs or by distributing it over the others boolean connectives until it vanishes. When a strategic operator, possibly with variables as arguments that are quantified by $\mathcal{Q}$ is reached, then $\mathcal{Q}$ is placed in front of the strategic operator and is also distributed in its temporal objective, but the vacuous quantification occurring in the process is removed. Lastly, when the formula begins with another quantifier $\mathcal{Q}''y_k$, then it is prefixed by $\mathcal{Q}$, the resulting vacuous quantification, if any, is removed, and the resulting prefix is pushed inside.

Example 4 Let $\varphi \in \mathcal{L}_{\text{HDMAS}}$ be

$$\forall y_1 (\langle y_1, 5 \rangle (\forall y_2 \langle y_1, y_2 \rangle X p_1) \cup (\exists y_2 \langle y_1, y_2 \rangle F p_2) \vee \\ \exists y_1 (\forall y_2 \langle y_1, y_2 \rangle F p_3 \wedge \neg \forall y_2 \langle 3, y_2 \rangle X p_1))$$

where $p_1, p_2, p_3 \in \Phi$. Then $\text{PUSH}(\forall y_1, \varphi) =$

$$\forall y_1 \langle y_1, 5 \rangle (\forall y_1 \forall y_2 \langle y_1, y_2 \rangle X p_1) \cup (\forall y_1 \exists y_2 \langle y_1, y_2 \rangle F p_2) \vee \\ (\exists y_1 \forall y_2 \langle y_1, y_2 \rangle F p_3 \wedge \neg \forall y_2 \langle 3, y_2 \rangle X p_1)$$

Theorem 1 Let $\mathcal{Q}$ be an admissible quantifier prefix and let $\mathcal{Q}\varphi$ be a formula of $\mathcal{L}_{\text{HDMAS}}$. Then $\text{PUSH}(\mathcal{Q}, \varphi)$ is logically equivalent in the finite to $\mathcal{Q}\varphi$.

Proof We prove the claim by induction on the **nesting depth** $\text{nd}(\varphi)$ of strategic operators in the state formula φ , defined as expected.

When $\text{nd}(\varphi) = 0$ the claim is straightforward because any quantification over φ is vacuous, hence $\mathcal{Q}\varphi \equiv \varphi$ and $\text{PUSH}(\mathcal{Q}, \varphi) = \varphi$. Suppose now that $\text{nd}(\varphi) > 0$ and the claim holds for all state formulae of $\mathcal{L}_{\text{HDMAS}}$ with lower nesting depth. We will do a nested induction on the structure of φ , following the recursive definition of PUSH .

1. $\varphi = \top \mid p$. This case does not apply now, but it is, anyway, trivial for every $\mathcal{Q}$.
2. $\varphi = \neg\psi$ follows from FOL and the inductive hypothesis (IH) for ψ .
3. $\varphi = \psi_1 \wedge \psi_2$.

- (a) When $\mathcal{Q} = \forall y_i$, the claim follows immediately from the valid equivalence (proved just like in FOL) $\forall y_i(\psi_1 \wedge \psi_2) \equiv \forall y_i\psi_1 \wedge \forall y_i\psi_2$ and the IH for each of ψ_1 and ψ_2 .
- (b) When $\mathcal{Q} = \exists y_i$, it suffices to prove that $\exists y_i(\psi_1 \wedge \psi_2) \equiv_{\text{fin}} \exists y_i\psi_1 \wedge \exists y_i\psi_2$, and then use the IH for each of ψ_1 and ψ_2 . The implication from left to right is by the validity of the implication $\exists y_i(\psi_1 \wedge \psi_2) \rightarrow \exists y_i\psi_1 \wedge \exists y_i\psi_2$. To prove the converse implication, first note that, since $\exists y_i(\psi_1 \wedge \psi_2)$ is a formula of $\mathcal{L}_{\text{HDMAS}}$, all free occurrences of y_i in ψ_1 and in ψ_2 must be positive. Now, suppose first that $i = 1$ and let $\mathcal{M}, s, \theta \models \exists y_1\psi_1 \wedge \exists y_1\psi_2$ for some finite $\mathcal{M}$. Then, $\mathcal{M}, s, \theta \models \psi_1[C_1/y_1]$ and $\mathcal{M}, s, \theta \models \psi_2[C_2/y_1]$ for some $C_1, C_2 \in \mathbb{N}$. Let $C = \max(C_1, C_2)$. By the monotonicity property **(C-mon)**⁺ from Lemma 5, we obtain that $\mathcal{M}, s, \theta \models \psi_1[C/y_1]$ and $\mathcal{M}, s, \theta \models \psi_2[C/y_1]$. Therefore, $\mathcal{M}, s, \theta \models (\psi_1 \wedge \psi_2)[C/y_1]$, hence $\mathcal{M}, s, \theta \models \exists y_1(\psi_1 \wedge \psi_2)$. This proves the validity of the converse implication $(\exists y_1\psi_1 \wedge \exists y_1\psi_2) \rightarrow \exists y_1(\psi_1 \wedge \psi_2)$. The proof of the case where $i = 2$ is analogous, using the monotonicity property **(N-mon)**⁺ from Lemma 5.

- (c) Lastly, the case when $\mathcal{Q} = \mathcal{Q}_{y_i} \mathcal{Q}' y_j$ is readily reducible to the previous 2 cases, by distributing first $\mathcal{Q}' y_j$ and then $\mathcal{Q}_{y_i}$.
4. $\varphi = \psi_1 \vee \psi_2$. This case is dually analogous to the previous one.
- (a) When $\mathcal{Q} = \exists y_i$, the claim follows immediately from the valid equivalence $\exists y_i(\psi_1 \vee \psi_2) \equiv \exists y_i \psi_1 \vee \exists y_i \psi_2$ and the IH for each of ψ_1 and ψ_2 .
- (b) When $\mathcal{Q} = \forall y_i$, it suffices to prove that $\forall y_i(\psi_1 \vee \psi_2) \equiv_{\text{fin}} \forall y_i \psi_1 \vee \forall y_i \psi_2$, and then use the IH for each of ψ_1 and ψ_2 . The implication from right to left ($\forall y_i \psi_1 \vee \forall y_i \psi_2 \rightarrow \forall y_i(\psi_1 \vee \psi_2)$) is a validity, proved just like in FOL. For the converse implication, suppose first that $i = 1$ and let $\mathcal{M}, s, \theta \models \forall y_1(\psi_1 \vee \psi_2)$ for some finite $\mathcal{M}$. Then, $\mathcal{M}, s, \theta \models (\psi_1 \vee \psi_2)[0/y_1]$, hence $\mathcal{M}, s, \theta \models \psi_1[0/y_1]$ or $\mathcal{M}, s, \theta \models \psi_2[0/y_1]$. Suppose w.l.o.g. the former. Then, by the monotonicity property **(C-mon)**⁺ from Lemma 5, we obtain that $\mathcal{M}, s, \theta \models \psi_2[C/y_1]$ for any $C \in \mathbb{N}$, hence $\mathcal{M}, s, \theta \models \forall y_1 \psi_1$, so $\mathcal{M}, s, \theta \models \forall y_1 \psi_1 \vee \forall y_1 \psi_2$.
- For the case that $i = 2$, assuming that $\mathcal{M}, s, \theta \models \forall y_2(\psi_1 \vee \psi_2)$, it follows that at least one of $\mathcal{M}, s, \theta \models \psi_1[N/y_2]$ and $\mathcal{M}, s, \theta \models \psi_2[N/y_2]$ holds for infinitely many values of $N \in \mathbb{N}$. Suppose w.l.o.g. the former. Then, by the monotonicity property **(N-mon)**⁺ from Lemma 5, we obtain that $\mathcal{M}, s, \theta \models \psi_1[N/y_2]$ for any $N \in \mathbb{N}$, hence $\mathcal{M}, s, \theta \models \forall y_2 \psi_1$, so $\mathcal{M}, s, \theta \models \forall y_2 \psi_1 \vee \forall y_2 \psi_2$.
- (c) Lastly, the case when $\mathcal{Q} = \mathcal{Q}_{y_i} \mathcal{Q}' y_j$ is readily reducible to the previous 2 cases, by distributing first $\mathcal{Q}' y_j$ and then $\mathcal{Q}_{y_i}$.
5. $\varphi = \mathcal{Q}'' y_k \psi$. Again, we consider the subcases depending on $\mathcal{Q}$.
- (a) $\mathcal{Q} = \mathcal{Q}_{y_i}$, where $i = k$.
We are to show that $\mathcal{Q}_{y_i} \mathcal{Q}'' y_i \psi \equiv_{\text{fin}} \text{PUSH}(\mathcal{Q}'' y_i, \psi)$, which follows from $\mathcal{Q}_{y_i} \mathcal{Q}'' y_i \psi \equiv \mathcal{Q}'' y_i \psi$ and the IH.
- (b) $\mathcal{Q} = \mathcal{Q}_{y_i}$, where $i \neq k$.
We are to show that $\mathcal{Q}_{y_i} \mathcal{Q}'' y_k \psi \equiv_{\text{fin}} \text{PUSH}(\mathcal{Q}_{y_i} \mathcal{Q}'' y_k, \psi)$, which follows from the IH for $\mathcal{Q} = \mathcal{Q}_{y_i} \mathcal{Q}'' y_k$ and ψ .
- (c) $\mathcal{Q} = \mathcal{Q}_{y_i} \mathcal{Q}' y_j$, where $i = k$.
We are to show that $\mathcal{Q}_{y_k} \mathcal{Q}' y_j \mathcal{Q}'' y_k \psi \equiv_{\text{fin}} \text{PUSH}(\mathcal{Q}' y_j \mathcal{Q}'' y_k, \psi)$, which follows from $\mathcal{Q}_{y_k} \mathcal{Q}' y_j \mathcal{Q}'' y_k \psi \equiv \mathcal{Q}' y_j \mathcal{Q}'' y_k \psi$ and the IH.
- (c) The case $\mathcal{Q} = \mathcal{Q}_{y_i} \mathcal{Q}' y_j$, where $j = k$, is analogous.
6. $\varphi = \langle t_1, t_2 \rangle \chi$.
This inductive case—for both inductions, the external one, on **nd**(φ), and for the nested one, on the structure of φ —is the most involved case, where the finiteness of the models over which we prove the equivalence is used essentially. There are several subcases, depending on $\mathcal{Q}$ and on the main temporal connective of χ . The proof for each case is technical and some cases are longer than others, but they all use a similar approach, that essentially hinges on the finiteness of the model and the monotonicity properties from Lemma 5. These will allow us to obtain uniformly large enough values of the quantified variables, beyond which the truth values of all strategic subformulae stabilise, and thus to establish the truth of the non-trivial implications.
- We will provide a representative selection of proofs for some of the cases and will leave out the rest, which are essentially analogous, though possibly even longer.

- (a) $Q = Q_{y_i}$, where $t_i = y_i$, for $i = 1$ or $i = 2$.

We are to show that $Q_{y_i} \langle \langle t_1, t_2 \rangle \rangle \chi \equiv_{\text{fin}} Q_{y_i} \langle \langle t_1, t_2 \rangle \rangle \text{PUSH}(Q_{y_i}, \chi)$, assuming the inductive hypothesis for the main state subformulae of χ . We consider the sub-cases depending on Q , i , and the main temporal connective of χ .

Case $(\forall y_1 G)$: to prove $\forall y_1 \langle \langle y_1, t \rangle \rangle G \psi \equiv_{\text{fin}} \forall y_1 \langle \langle y_1, t \rangle \rangle G \text{PUSH}(\forall y_1, \psi)$.

By the IH for ψ , we have that $\forall y_1 \psi \equiv_{\text{fin}} \text{PUSH}(\forall y_1, \psi)$.

So, it suffices to prove that $\forall y_1 \langle \langle y_1, t \rangle \rangle G \psi \equiv_{\text{fin}} \forall y_1 \langle \langle y_1, t \rangle \rangle G \forall y_1 \psi$. By Lemma 6, $\forall y_1 \langle \langle y_1, t \rangle \rangle G \psi \equiv \langle \langle 0, t \rangle \rangle G \psi[0/y_1]$ and $\forall y_1 \langle \langle y_1, t \rangle \rangle G \forall y_1 \psi \equiv \langle \langle 0, t \rangle \rangle G \forall y_1 \psi$.

So, we have to prove that $\langle \langle 0, t \rangle \rangle G \psi[0/y_1] \equiv \langle \langle 0, t \rangle \rangle G \forall y_1 \psi$, which follows immediately, since $\forall y_1 \psi \equiv \psi[0/y_1]$, by **(C-mon)**⁺ from Lemma 5.

Case $(\exists y_1 G)$: to prove $\exists y_1 \langle \langle y_1, t \rangle \rangle G \psi \equiv_{\text{fin}} \exists y_1 \langle \langle y_1, t \rangle \rangle G \text{PUSH}(\exists y_1, \psi)$.

By the IH for ψ , we have that $\exists y_1 \psi \equiv_{\text{fin}} \text{PUSH}(\exists y_1, \psi)$. So, it suffices to prove that $\exists y_1 \langle \langle y_1, t \rangle \rangle G \psi \equiv_{\text{fin}} \exists y_1 \langle \langle y_1, t \rangle \rangle G \exists y_1 \psi$. Since $\models \psi \rightarrow \exists y_1 \psi$, we obtain validity of the implication $\exists y_1 \langle \langle y_1, t \rangle \rangle G \psi \rightarrow \exists y_1 \langle \langle y_1, t \rangle \rangle G \exists y_1 \psi$.

For the converse, suppose $\mathcal{M}, s, \theta \models \exists y_1 \langle \langle y_1, t \rangle \rangle G \exists y_1 \psi$ for some finite $\mathcal{M}$ with state space S , assignment θ and $s \in S$. Fix any $C \in \mathbb{N}$ such that $\mathcal{M}, s, \theta \models \langle \langle C, t \rangle \rangle G \exists y_1 \psi$. Since θ fixes the values of all terms, we can treat $\exists y_1 \psi$ as a closed formula. Note that, according to the syntax of $\mathcal{L}_{\text{HDMAS}}$, all occurrences of y_1 in ψ are positive. Let $W = \llbracket \exists y_1 \psi \rrbracket_{\mathcal{M}}^\theta$ be its extension in $\mathcal{M}$ (which depends on θ) and let $w \in W$. Let $f : W \rightarrow \mathbb{N}$ be a mapping assigning to every $u \in W$ a number $f(u)$ such that $\mathcal{M}, u, \theta \models \psi[f(u)/y_1]$. Now, let⁹ $f^* := \max_{u \in W} f(u)$ and $C^* := \max(f^*, C)$. Then, by **(C-mon)**⁺ from Lemma 5, we obtain that $\mathcal{M}, u, \theta \models \psi[C^*/y_1]$ for each $u \in W$, hence $\mathcal{M}, s, \theta \models \langle \langle C^*, t \rangle \rangle G \psi[C^*/y_1]$. Therefore, $\mathcal{M}, s, \theta \models \exists y_1 \langle \langle y_1, t \rangle \rangle G \varphi$. Thus, $\exists y_1 \langle \langle y_1, t \rangle \rangle G \exists y_1 \psi \rightarrow \exists y_1 \langle \langle y_1, t \rangle \rangle G \varphi$ is valid in the finite, whence the claim.

Case $(\forall y_2 G)$: to prove $\forall y_2 \langle \langle t, y_2 \rangle \rangle G \psi \equiv_{\text{fin}} \forall y_2 \langle \langle t, y_2 \rangle \rangle G \text{PUSH}(\forall y_2, \psi)$.

By the IH for ψ , we have that $\forall y_2 \psi \equiv_{\text{fin}} \text{PUSH}(\forall y_2, \psi)$. So, it suffices to prove that $\forall y_2 \langle \langle t, y_2 \rangle \rangle G \psi \equiv_{\text{fin}} \forall y_2 \langle \langle t, y_2 \rangle \rangle G \forall y_2 \psi$. The implication $\models \forall y_2 \langle \langle t, y_2 \rangle \rangle G \forall y_2 \psi \rightarrow \forall y_2 \langle \langle t, y_2 \rangle \rangle G \psi$ follows from $\models \forall y_2 \psi \rightarrow \psi$, proved just like in FOL. For the converse, suppose $\mathcal{M}, s, \theta \models \forall y_2 \langle \langle t, y_2 \rangle \rangle G \psi$ for some finite $\mathcal{M}$ with state space S , assignment θ and $s \in S$. Then, for every $N \in \mathbb{N}$, it holds that $\mathcal{M}, s, \theta \models \langle \langle t, N \rangle \rangle G \psi[N/y_2]$, i.e., there is an abstract positional joint strategy σ_N for $\theta(t)$ many controllable agents, such that ψ is true at every state on every outcome play enabled by σ_N against N uncontrollable agents. Since there are only finitely many abstract positional joint strategies for $\theta(t)$ controllable agents in $\mathcal{M}$, there is at least one such joint strategy which works for infinitely many values of N , and therefore, by **(N-mon)**, it will work for all $N \in \mathbb{N}$. Let us fix such strategy σ^c . We will show that $\mathcal{M}, s, \theta \models \forall y_2 \langle \langle t, y_2 \rangle \rangle G \forall y_2 \psi$ by proving that, for every $N \in \mathbb{N}$, if σ^c is played by $\theta(t)$ many controllable agents it ensures the truth of $\mathcal{M}, s, \theta \models \langle \langle t, N \rangle \rangle G \forall y_2 \psi$. Suppose this is not the case for some $N' \in \mathbb{N}$. Then, there is an abstract positional joint strategy σ^n for N' uncontrollable agents that guarantees reaching a state w where $\forall y_2 \psi$ fails on the unique play π generated by the pair of joint strategies (σ^c, σ^n) . Thus, $\mathcal{M}, w, \theta \not\models \forall y_2 \psi$, i.e., $\mathcal{M}, w, \theta \models \neg \forall y_2 \psi$. Therefore, $\mathcal{M}, w, \theta \models \neg \psi[N''/y_2]$ for some $N'' \in \mathbb{N}$. Let $N^* := \max(N', N'')$. Then, by **(N-mon)**⁻ from Lemma 5, we have that $\mathcal{M}, w, \theta \models \neg \psi[N^*/y_2]$. Fur-

⁹ This is where we use the finiteness of the model.

thermore, the strategy σ^n can be trivially extended to σ^{n*} for N^* uncontrollable agents (by letting the extra $N^* - N'$ uncontrollable agents idle), hence the play π is still generated by the resulting pair of joint strategies (σ^c, σ^{n*}) and the state w as above will still be reached on it. On the other hand, by the choice of σ^c , when it is played by the $\theta(t)$ many controllable agents against N^* uncontrollable agents it guarantees maintaining forever the truth of ψ , i.e., $\mathcal{M}, s, \theta \models \langle\langle t, N^* \rangle\rangle G \psi[N^*/y_2]$. In particular, that implies $\mathcal{M}, w, \theta \models \psi[N^*/y_2]$ —a contradiction. Therefore, the assumption that such N' exists is wrong, whence the claim.

Case $(\exists y_2 G)$: to prove $\exists y_2 \langle\langle t, y_2 \rangle\rangle G \psi \equiv_{\text{fin}} \exists y_2 \langle\langle t, y_2 \rangle\rangle G \text{PUSH}(\exists y_2, \psi)$.

This case is quite analogous to **Case $(\forall y_1 G)$** and is proved by using the IH for ψ , the equivalences $\exists y_2 \langle\langle t, y_2 \rangle\rangle G \psi \equiv \langle\langle t, 0 \rangle\rangle G \psi[0/y_2]$ and $\exists y_2 \langle\langle t, y_2 \rangle\rangle G \exists y_2 \psi \equiv \langle\langle t, 0 \rangle\rangle G \exists y_2 \psi$ from Lemma 6, and the monotonicity properties ($C - \text{mon}$) from Lemma 5.

Cases $(Qy_i X)$ are analogous, but a little simpler than those above.

Cases $(Qy_i U)$ are analogous, though a little longer than those above.

- (b) $Q = Qy_i$, where $t_1 \neq y_i$ and $t_2 \neq y_i$.

We are to show that $Qy_i \langle\langle t_1, t_2 \rangle\rangle \chi \equiv_{\text{fin}} \langle\langle t_1, t_2 \rangle\rangle \text{PUSH}(Qy_i, \chi)$, assuming the IH for the main state subformulae of χ . For that, it suffices to prove that the quantifier Q can be equivalently pushed inside through $\langle\langle t_1, t_2 \rangle\rangle$ and the main temporal connective of χ , e.g., that $Qy_i \langle\langle t_1, t_2 \rangle\rangle G \psi \equiv_{\text{fin}} \langle\langle t_1, t_2 \rangle\rangle Qy_i G \psi$. The non-trivial implications follow from the fact that there are only finitely many abstract positional strategies for the controllable agents in any given finite model, plus the monotonicity properties from Lemma 5. The argument for that is essentially the same as that in the proof of **Case $(\forall y_2 G)$** above.

- (c) $Q = Qy_1 Q'y_2$, where $t_1 \neq y_1$ and $t_2 \neq y_2$.

We are to show that $Qy_1 Q'y_2 \langle\langle t_1, t_2 \rangle\rangle \chi \equiv_{\text{fin}} \langle\langle t_1, t_2 \rangle\rangle \text{PUSH}(Qy_1 Q'y_2, \chi)$, assuming the IH for all state formulae of lower nesting depth, including the main state subformulae of χ .

This equivalence follows by applying case (b) twice, first for $Q = Q'y_2$ and then for $Q = Qy_1$ (the IH on the nesting of strategic operators is used here), and each time using the IH.

The case $Q = Q'y_2 Qy_1$, where $t_1 \neq y_1$ and $t_2 \neq y_2$ is completely analogous.

- (d) $Q = Qy_1 Q'y_2$, where $t_1 = y_1$ and $t_2 \neq y_2$.

We are to show that

$$Qy_1 Q'y_2 \langle\langle t_1, t_2 \rangle\rangle \chi \equiv_{\text{fin}} Qy_1 \langle\langle t_1, t_2 \rangle\rangle \text{PUSH}(Qy_1 Q'y_2, \chi),$$

assuming the IH for all state formulae of lower nesting depth, incl. the main state subformulae of χ .

E.g., when $\chi = G \psi$, we are to prove $Qy_1 Q'y_2 \langle\langle y_1, t_2 \rangle\rangle G \psi \equiv_{\text{fin}} Qy_1 \langle\langle y_1, t_2 \rangle\rangle G \text{PUSH}(Qy_1 Q'y_2, \psi)$. By the IH, $\text{PUSH}(Qy_1 Q'y_2, \psi) \equiv_{\text{fin}} Qy_1 Q'y_2 \psi$, so we are to prove that $Qy_1 Q'y_2 \langle\langle y_1, t_2 \rangle\rangle G \psi \equiv_{\text{fin}} Qy_1 \langle\langle y_1, t_2 \rangle\rangle G Qy_1 Q'y_2 \psi$.

This follows by first applying case (b) for $Q = Q'y_2$ and the IH to obtain

$Q'y_2 \langle\langle y_1, t_2 \rangle\rangle G \psi \equiv_{\text{fin}} \langle\langle y_1, t_2 \rangle\rangle G \text{PUSH}(Q'y_2, \psi) \equiv_{\text{fin}} \langle\langle y_1, t_2 \rangle\rangle G Q'y_2 \psi$, and then applying Qy_1 to both sides, then case (a) for $Q = Qy_1$, and again the IH.

- (e) The case $Q = Q'y_2 Qy_1$, where $t_1 = y_1$ and $t_2 \neq y_2$ is similar.

- (f) The cases $Q = Qy_1 Q'y_2$ and $Q = Q'y_2 Qy_1$ where $t_1 \neq y_1$ and $t_2 = y_2$ are completely analogous.

- (g) $Q = Qy_1 Q'y_2$, where $t_1 = y_1$ and $t_2 = y_2$.

We have to prove $Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle \chi \equiv_{\text{fin}} Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle \text{PUSH}(Q_{y_1} Q'_{y_2}, \chi)$, assuming the IH for all state formulae of lower nesting depth, incl. the main state subformulae of χ .

E.g., when $\chi = G\psi$, we are to prove $Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G\psi \equiv_{\text{fin}} Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G \text{PUSH}(Q_{y_1} Q'_{y_2}, \psi)$.

By the IH, $\text{PUSH}(Q_{y_1} Q'_{y_2}, \psi) \equiv_{\text{fin}} Q_{y_1} Q'_{y_2} \psi$, so we are to prove that

$$Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G\psi \equiv_{\text{fin}} Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G Q_{y_1} Q'_{y_2} \psi.$$

By case (a), we have already shown that

$$Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G\psi \equiv_{\text{fin}} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G Q'_{y_2} \psi.$$

By applying Q_{y_1} to both sides we obtain

$$Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G\psi \equiv_{\text{fin}} Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G Q'_{y_2} \psi, \text{ so it remains to prove}$$

$$Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G Q'_{y_2} \psi \equiv_{\text{fin}} Q_{y_1} Q'_{y_2} \langle\langle y_1, y_2 \rangle\rangle G Q_{y_1} Q'_{y_2} \psi.$$

For each case of Q the argument for the non-trivial implication uses the monotonicity properties from Lemma 5 and is respectively similar to that in the proof of **Case** $(\exists_{y_2} G)$ and **Case** $(\forall_{y_2} G)$ above.

The other cases for χ are similar.

- (h) The case $Q = Q'_{y_2} Q_{y_1}$, where $t_1 = y_1$ and $t_2 = y_2$ is completely analogous to the previous one.

This completes the proof for all cases in the definition of $\text{PUSH}(Q, \langle\langle t_1, t_2 \rangle\rangle \chi)$ and, therefore, the last inductive case in both inductions. □

Now we will define a recursive function NF that transforms any state or path formula ξ of $\mathcal{L}_{\text{HDMAS}}$ respectively into a state or path formula ξ^{NF} in $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$, while preserving equivalence in the finite.

- 1: **Input:** ξ , a state or path formula of $\mathcal{L}_{\text{HDMAS}}$
- 2: **Output:** ξ^{NF} , a state or path formula of $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$
- 3: **procedure** $\text{NF}(\xi)$
- 4: **case** ξ **of**
- 5: $\top \mid p$: **return** ξ
- 6: $\neg\psi$: **return** $\neg\text{NF}(\psi)$
- 7: $\psi_1 \wedge \psi_2$: **return** $\text{NF}(\psi_1) \wedge \text{NF}(\psi_2)$
- 8: $\psi_1 \vee \psi_2$: **return** $\text{NF}(\psi_1) \vee \text{NF}(\psi_2)$
- 9: $\langle\langle t_1, t_2 \rangle\rangle \chi$: **return** $\langle\langle t_1, t_2 \rangle\rangle \text{NF}(\chi)$
- 10: $Q\psi$: **return** $\text{PQE}(\text{PUSH}(Q, \text{NF}(\psi)))$
- 11: $X\psi$: **return** $X \text{NF}(\psi)$
- 12: $G\psi$: **return** $G \text{NF}(\psi)$
- 13: $\psi_1 \cup \psi_2$: **return** $\text{NF}(\psi_1) \cup \text{NF}(\psi_2)$
- 14: **end case**
- 15: **end procedure**

Intuitively, NF transforms the input formula by first applying PUSH and then PQE whenever a quantifier prefix is to be applied, thus producing a formula in a normal form.

Example 5 Let φ be as in Example 4. Then $\text{NF}(\varphi)$ is:

$$\begin{aligned} & \langle\langle 0, 5 \rangle\rangle \left((\forall y_2 \langle\langle 0, y_2 \rangle\rangle \times p_1) \cup (\forall y_2 \langle\langle 0, y_2 \rangle\rangle F p_2) \right) \vee \\ & (\exists y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle F p_3) \wedge \neg \forall y_2 \langle\langle 3, y_2 \rangle\rangle \times p_1 \end{aligned}$$

Lemma 9 Let φ a state formula of $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$. For every admissible quantifier prefix Q , if the variables occurring in Q do not occur free in φ , then $\text{PUSH}(Q, \varphi) = \varphi$.

Proof The argument is by structural induction on φ in normal form, by following the recursive definition of PUSH . The non-trivial cases are those involving quantifiers. We consider $\exists y_1 \langle\langle y_1, t_2 \rangle\rangle \chi$ and $Q = Qy_1 Q'y_2$, the other cases are proved analogously. Since $\varphi \in \mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ by hypothesis, we have that $t_2 \neq y_2$, thus by definition $\text{PUSH}(Qy_1 Q'y_2, \exists y_1 \langle\langle y_1, t_2 \rangle\rangle \chi) = \text{PUSH}(Q'y_2 \exists y_1, \langle\langle y_1, t_2 \rangle\rangle \chi) = \exists y_1 \langle\langle y_1, t_2 \rangle\rangle \text{PUSH}(Q'y_2 \exists y_1, \chi)$. By hypothesis y_2 is not free in ψ , which entails that y_2 is not free in χ and the same holds for y_1 by (NF2) in the definition of normal form. We can therefore apply the inductive hypothesis on χ to get $\exists y_1 \langle\langle y_1, t_2 \rangle\rangle \text{PUSH}(Q'y_2 \exists y_1, \chi) = \exists y_1 \langle\langle y_1, t_2 \rangle\rangle \chi$. $\square$

Lemma 10 If $\varphi \in \mathcal{L}_{\text{HDMAS}}^{\text{NF}}$, then $\text{NF}(\varphi) = \varphi$.

Proof By induction on the structure of φ in normal form, following the recursive definition of NF . All cases are straightforward, except $\varphi = Q\psi$. We consider $\varphi = \exists y_1 \langle\langle y_1, t_2 \rangle\rangle \chi$, all other cases being analogous. By the IH, $\text{NF}(\langle\langle y_1, t_2 \rangle\rangle \chi) = \langle\langle y_1, t_2 \rangle\rangle \chi$. Also, note that y_1 does not occur free in χ since $\varphi \in \mathcal{L}_{\text{HDMAS}}^{\text{NF}}$. Therefore, using Lemmas 9 and 8, we successively obtain:

$$\begin{aligned} \text{NF}(\varphi) &= \text{NF}(\exists y_1 \langle\langle y_1, t_2 \rangle\rangle \chi) = \text{PQE}(\text{PUSH}(\exists y_1, \text{NF}(\langle\langle y_1, t_2 \rangle\rangle \chi))) \\ &= \text{PQE}(\text{PUSH}(\exists y_1, \langle\langle y_1, t_2 \rangle\rangle \chi)) = \text{PQE}(\exists y_1 \langle\langle y_1, t_2 \rangle\rangle \chi) = \exists y_1 \langle\langle y_1, t_2 \rangle\rangle \chi = \varphi. \end{aligned}$$

$\square$

Theorem 2 Let φ be any formula in $\mathcal{L}_{\text{HDMAS}}$. Then:

1. $\text{NF}(\varphi) \in \mathcal{L}_{\text{HDMAS}}^{\text{NF}}$.
2. $\text{NF}(\varphi) \equiv_{\text{fin}} \varphi$.
3. $\text{NF}(\varphi)$ can be computed effectively and has length linearly bounded above by $|\varphi|$.

Proof The first claim follows by straightforward induction on the structure of φ , or just by direct inspection of the function NF .

Claim 2. is proved by induction on the structure of φ , following the cases of the recursive definition of NF . The only non-trivial case is $\varphi = Q\psi$, which follows immediately from the IH, Theorem 1, and Lemma 8.

Lastly, Claim 3. follows by direct inspection of all cases in the definitions of the functions PQE , PUSH and NF . $\square$

We conclude the section by presenting the fixpoint characterizations of formulae in (1), which provide an effective procedure for the model checking algorithm.

Theorem 3 For every terms $t, t', t'' \in T$ the following equivalences hold, where the formulae on the left are in $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$.

1. $\langle\langle t', t'' \rangle\rangle G \varphi \equiv \varphi \wedge \langle\langle t', t'' \rangle\rangle X \langle\langle t', t'' \rangle\rangle G \varphi$
2. $\langle\langle t', t'' \rangle\rangle \psi U \varphi \equiv \varphi \vee (\psi \wedge \langle\langle t', t'' \rangle\rangle X \langle\langle t', t'' \rangle\rangle \psi U \varphi)$
3. $\exists y_1 \langle\langle y_1, t \rangle\rangle G \varphi \equiv_{\text{fin}} \varphi \wedge \exists y_1 \langle\langle y_1, t \rangle\rangle X \exists y_1 \langle\langle y_1, t \rangle\rangle G \varphi$
4. $\forall y_2 \langle\langle t, y_2 \rangle\rangle G \varphi \equiv_{\text{fin}} \varphi \wedge \forall y_2 \langle\langle t, y_2 \rangle\rangle X \forall y_2 \langle\langle t, y_2 \rangle\rangle G \varphi$
5. $\exists y_1 \langle\langle y_1, t \rangle\rangle \psi U \varphi \equiv_{\text{fin}} \varphi \vee (\psi \wedge \exists y_1 \langle\langle y_1, t \rangle\rangle X \exists y_1 \langle\langle y_1, t \rangle\rangle \psi U \varphi)$
6. $\forall y_2 \langle\langle t, y_2 \rangle\rangle \psi U \varphi \equiv_{\text{fin}} \varphi \vee (\psi \wedge \forall y_2 \langle\langle t, y_2 \rangle\rangle X \forall y_2 \langle\langle t, y_2 \rangle\rangle \psi U \varphi)$
7. $\forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle G \varphi \equiv_{\text{fin}} \varphi \wedge \forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle X \forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle G \varphi$
8. $\exists y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle G \varphi \equiv_{\text{fin}} \varphi \wedge \exists y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle X \exists y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle G \varphi$
9. $\forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle \psi U \varphi \equiv_{\text{fin}} \varphi \vee (\psi \wedge \forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle X \forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle \psi U \varphi)$
10. $\exists y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle \psi U \varphi \equiv_{\text{fin}} \varphi \vee (\psi \wedge \exists y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle X \exists y_1 \forall y_2 \langle\langle y_1, y_2 \rangle\rangle \psi U \varphi)$

Proof

1. Follows directly from the semantics, just like the respective fixpoint equivalence for $\langle\langle A \rangle\rangle G$ in ATL, cf. [8].
2. Likewise, just like the respective fixpoint equivalence for $\langle\langle A \rangle\rangle U$ in ATL.
3. First, note that y_1 does not occur free in φ since $\exists y_1 \langle\langle y_1, t \rangle\rangle G \varphi \in \mathcal{L}_{\text{HDMAS}}^{\text{NF}}$.

Now, we take the equivalence 1, where $t' = y_1$, and quantify both sides with $\exists y_1$, obtaining:

$$\exists y_1 \langle\langle y_1, t \rangle\rangle G \varphi \equiv \exists y_1 (\varphi \wedge \langle\langle y_1, t \rangle\rangle X \langle\langle y_1, t \rangle\rangle G \varphi) \equiv \varphi \wedge \exists y_1 \langle\langle y_1, t \rangle\rangle X \langle\langle y_1, t \rangle\rangle G \varphi.$$

By applying Theorem 2 to both sides above and then using Lemmas 9, 10 and 8 we obtain:

$$\begin{aligned}
 & \exists y_1 \langle\langle y_1, t \rangle\rangle G \varphi \\
 & \equiv_{\text{fin}} \text{NF}(\varphi \wedge \exists y_1 \langle\langle y_1, t \rangle\rangle X \langle\langle y_1, t \rangle\rangle G \varphi) \\
 & = \text{NF}(\varphi) \wedge \text{NF}(\exists y_1 \langle\langle y_1, t \rangle\rangle X \langle\langle y_1, t \rangle\rangle G \varphi) \\
 & = \varphi \wedge \text{PQE}(\text{PUSH}(\exists y_1, \text{NF}(\langle\langle y_1, t \rangle\rangle X \langle\langle y_1, t \rangle\rangle G \varphi))) \\
 & = \varphi \wedge \text{PQE}(\text{PUSH}(\exists y_1, \langle\langle y_1, t \rangle\rangle X \langle\langle y_1, t \rangle\rangle G \text{NF}(\varphi))) \\
 & = \varphi \wedge \text{PQE}(\text{PUSH}(\exists y_1, \langle\langle y_1, t \rangle\rangle X \langle\langle y_1, t \rangle\rangle G \varphi)) \\
 & = \varphi \wedge \text{PQE}(\exists y_1 \langle\langle y_1, t \rangle\rangle \text{PUSH}(\exists y_1, X \langle\langle y_1, t \rangle\rangle G \varphi)) \\
 & = \varphi \wedge \text{PQE}(\exists y_1 \langle\langle y_1, t \rangle\rangle X \exists y_1 \langle\langle y_1, t \rangle\rangle \text{PUSH}(\exists y_1, G \varphi)) \\
 & = \varphi \wedge \text{PQE}(\exists y_1 \langle\langle y_1, t \rangle\rangle X \exists y_1 \langle\langle y_1, t \rangle\rangle G \varphi) \\
 & \equiv \varphi \wedge \exists y_1 \langle\langle y_1, t \rangle\rangle X \exists y_1 \langle\langle y_1, t \rangle\rangle G \varphi.
 \end{aligned}$$

The other cases are analogous. □

4 Model checking

In this section we develop an algorithm for model checking the fragment $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$. By virtue of Theorem 2, it will provide a model checking procedure for the whole $\mathcal{L}_{\text{HDMAS}}$.

Let φ be any state formula of $\mathcal{L}_{\text{HDMAS}}$, $\mathcal{M}$ be a HDMAS, s a state and θ an assignment in $\mathcal{M}$. The **local model checking problem** is the problem of deciding whether $\mathcal{M}, s, \theta \models \varphi$, while the **global model checking problem** is the computational problem that returns the set of

states in $\mathcal{M}$ where the input formula φ is satisfied, i.e. it is the problem of computing the **state extension** of φ in $\mathcal{M}$ given θ , formally defined as:

$$\llbracket \varphi \rrbracket_{\mathcal{M}}^{\theta} = \{s \in S \mid \mathcal{M}, s, \theta \models \varphi\}.$$

For closed formulae φ , $\llbracket \varphi \rrbracket_{\mathcal{M}}^{\theta}$ does not depend on the assignment θ , so we omit it and write $\llbracket \varphi \rrbracket_{\mathcal{M}}$.

Algorithm 4 presented here solves the global model checking problem for all $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ formulae. The core sub-procedure of the algorithm is the function PREIMG which, given a set Q of states in S and $C, N \in \mathbb{N}$, returns the set of states from which C controllable agents have a joint action, which, when played against any joint action of other N uncontrollable agents produces an outcome state in Q . We will call that set the **(C, N)-controllable pre-image of Q**. Often we will omit (C, N) , when unspecified or fixed in the context, and will write simply “the controllable pre-image of Q ”. We also extend that notion to “ (t_1, t_2) -controllable pre-image”, for any terms t_1, t_2 , the values of which are given by the assignment. When $Q = \llbracket \psi \rrbracket_{\mathcal{M}}^{\theta}$, it computes the state extension of $\langle\langle t_1, t_2 \rangle\rangle \times \psi$ which is parameterised by terms t_1, t_2 (by means of their values $\theta(t_1)$ and $\theta(t_2)$). We then extend that further to quantified extensions of $\langle\langle t_1, t_2 \rangle\rangle \times \psi$, by adding the respective quantification to the result. In all cases, we reduce the problem of computing the controllable pre-images to checking the truth of Presburger formulae.

We now proceed with some technical preparation. Recall that X^+ is the set of $n + 1$ action counters. We will also be using auxiliary integer variables $k_1, \dots, k_n, k_{\epsilon}$ and $\ell_1, \dots, \ell_n, \ell_{\epsilon}$ not contained in X^+ . Each k_i (respectively, ℓ_i) represents the number of controllable (respectively, uncontrollable) agents performing action act_i ; likewise for k_{ϵ} (resp., ℓ_{ϵ}) for the number of controllable (resp., uncontrollable) agents performing the idle action. Also, for each s in S and $i \in \{1, \dots, n\}$ we introduce an auxiliary propositional constant d_s^i which is true if and only if action act_i is available in s , i.e., $act_i \in d(s)$.

Definition 14 Given a HDMAS $\mathcal{M}$ with a finite state space S , state s in S , a subset Q of S , and terms t_1, t_2 , we define the following Presburger formulae:

$$\begin{aligned} g_Q^s(x_1, \dots, x_n) &:= \bigvee_{s' \in Q} \delta(s, s')(x_1, \dots, x_n). \\ \text{PrF}(\mathcal{M}, s, t_1, t_2, Q) &:= \exists k_1 \dots \exists k_n \exists k_{\epsilon} \left(\bigwedge_{i=1}^n (k_i \neq 0 \rightarrow d_s^i) \wedge \sum_{i=1}^n k_i + k_{\epsilon} = t_1 \wedge \right. \\ &\quad \left. \forall \ell_1 \dots \forall \ell_n \forall \ell_{\epsilon} \left(\left(\bigwedge_{i=1}^n (\ell_i \neq 0 \rightarrow d_s^i) \wedge \sum_{i=1}^n \ell_i + \ell_{\epsilon} = t_2 \right) \rightarrow g_Q^s((k_1 + \ell_1), \dots, (k_n + \ell_n)) \right) \right) \end{aligned}$$

The formula $\text{PrF}(\mathcal{M}, s, t_1, t_2, Q)$ intuitively says that there is a tuple of available actions at s such that when played by t_1 many (controllable) agents and combined with any tuple of available actions for t_2 many (uncontrollable) agents, it satisfies a guard of a transition leading to a state in Q . (The formula can be shortened somewhat, if the quantification is restricted only to k - and ℓ -variables that correspond to action counters that appear in the guard g_Q^s , which would improve the complexity estimates, as shown in Sect. 5.) That formula and its extensions with quantifiers over t_1 (when equal to y_1) and t_2 (when equal to y_2) will be used by the global model checking algorithm to compute the controllable pre-images of state extensions.

Example 6 Let us compute the state extension of the formula

$\varphi = \exists y_1 \forall y_2 \langle \langle y_1, y_2 \rangle \rangle X(p \vee q)$ in the model $\mathcal{M}$ of Example 2. First, we compute $\llbracket p \vee q \rrbracket_{\mathcal{M}} = \{s_2, s_3, s_4, s_5, s_6\}$. Then, for each state $s \in \mathcal{M}$ we check the truth of the closed Presburger formula $\exists y_1 \forall y_2 \text{PrF}(\mathcal{M}, s, y_1, y_2, \llbracket p \vee q \rrbracket_{\mathcal{M}})$ in $\mathcal{M}$.

- $\exists y_1 \forall y_2 \text{PrF}(\mathcal{M}, s_1, y_1, y_2, \llbracket p \vee q \rrbracket_{\mathcal{M}})$ is false, thus s_1 does not belong to the $\exists y_1 \forall y_2 (y_1, y_2)$ -controllable pre-image of $\llbracket p \vee q \rrbracket_{\mathcal{M}}$. Indeed 11 uncontrollable agents can force the system to stay in s_1 when they all perform act_3 ;
- $\exists y_1 \forall y_2 \text{PrF}(\mathcal{M}, s_2, y_1, y_2, \llbracket p \vee q \rrbracket_{\mathcal{M}})$ is true, hence s_2 belongs to the $\exists y_1 \forall y_2 (y_1, y_2)$ -controllable pre-image of $\llbracket p \vee q \rrbracket_{\mathcal{M}}$ trivially because all outgoing transitions from s_2 lead to states in $\llbracket p \vee q \rrbracket_{\mathcal{M}}$;
- checking all other states likewise produces the final result:
 $\llbracket \varphi \rrbracket_{\mathcal{M}} = \{s_2, s_4, s_5, s_6\}$.

Algorithm 1 Computing the controllable by t_1 agents pre-image of Q against t_2 agents (with t_1, t_2 possibly quantified).

```

1: Inputs: HDMAS  $\mathcal{M}$ ,  $t_1, t_2 \in T$ ,  $Q \subseteq S$ , assignment  $\theta$  and prefix  $\text{pfix}$ 
2: Output: the  $(t_1, t_2)$ -controllable pre-image  $Z \subseteq S$  of  $Q$ 
3: procedure PREIMG( $\mathcal{M}, t_1, t_2, Q, \theta, \text{pfix}$ )
4:   if  $t_1 \neq y_1$  does not appear in  $\text{pfix}$  then
5:      $t_1 \leftarrow \theta(t_1)$ 
6:   end if
7:   if  $t_2 \neq y_2$  does not appear in  $\text{pfix}$  then
8:      $t_2 \leftarrow \theta(t_2)$ 
9:   end if
10:   $Z \leftarrow \emptyset$ 
11:  for all  $s \in S$  do
12:    if  $\text{pfix PrF}(\mathcal{M}, s, t_1, t_2, Q)$  true then  $Z \leftarrow Z \cup \{s\}$ 
13:  end if
14:  end for
15:  return  $Z$ 
16: end procedure

```

Algorithm 2 Global model checking algorithm for closed formulae of the type $\text{pfix} \langle \langle t_1, t_2 \rangle \rangle G \psi$.

```

1: Inputs: HDMAS  $\mathcal{M}$ ,  $t_1, t_2 \in T$ , formula  $\psi$ , assignment  $\theta$  and prefix  $\text{pfix}$ 
2: Output: the set of states  $Z = \llbracket \text{pfix} \langle \langle t_1, t_2 \rangle \rangle G \psi \rrbracket_{\mathcal{M}}^{\theta}$ 
3: procedure G-FIXPOINT( $\mathcal{M}, t_1, t_2, \psi, \theta, \text{pfix}$ )
4:   $Q \leftarrow \text{GLOBALMC}(\mathcal{M}, \psi, \theta)$ 
5:   $W \leftarrow S$ 
6:   $Z \leftarrow Q$ 
7:  while  $W \not\subseteq Z$  do
8:     $W \leftarrow Z$ 
9:     $Z \leftarrow \text{PREIMG}(\mathcal{M}, t_1, t_2, \text{pfix}, W) \cap Q$ 
10:  end while
11:  return  $Z$ 
12: end procedure

```

Algorithm 3 Global model checking algorithm for closed formulae of the type $\text{pfix}\langle t_1, t_2 \rangle \psi_1 \cup \psi_2$.

```

1: Inputs: HDMAS  $\mathcal{M}$ ,  $t_1, t_2 \in T$ , formulae  $\psi_1, \psi_2$ , assignment  $\theta$  and prefix  $\text{pfix}$ 
2: Output: the set of states  $Z = \llbracket \text{pfix}\langle t_1, t_2 \rangle \psi_1 \cup \psi_2 \rrbracket_{\mathcal{M}}^{\theta}$ 
3: procedure U-FIXPOINT( $\mathcal{M}, t_1, t_2, \psi_1, \psi_2, \theta, \text{pfix}$ )
4:    $Q_1 \leftarrow \text{GLOBALMC}(\mathcal{M}, \psi_1, \theta)$ 
5:    $Q_2 \leftarrow \text{GLOBALMC}(\mathcal{M}, \psi_2, \theta)$ 
6:    $W \leftarrow \emptyset$ 
7:    $Z \leftarrow Q_2$ 
8:   while  $Z \not\subseteq W$  do
9:      $W \leftarrow Z$ 
10:     $Z \leftarrow Q_2 \cup (\text{PREIMG}(\mathcal{M}, t_1, t_2, \text{pfix}, W) \cap Q_1)$ 
11:   end while
12:   return  $Z$ 
13: end procedure

```

Algorithm 4 Global model checking algorithm for $\mathcal{L}_{\text{HDMAS}}$ -formulae.

```

1: Inputs: HDMAS  $\mathcal{M}$ , formula  $\varphi$  and assignment  $\theta$ 
2: Output: the set of states  $Z = \llbracket \varphi \rrbracket_{\mathcal{M}}^{\theta}$ 
3: procedure GLOBALMC( $\mathcal{M}, \varphi, \theta$ )
4:   case  $\varphi$  of
5:      $p$ :
6:       return  $\{s \in S \mid p \in \lambda(s)\}$ 
7:      $\neg\psi$ :
8:       return  $S \setminus \text{GLOBALMC}(\mathcal{M}, \psi, \theta)$ 
9:      $\psi_1 \wedge \psi_2$ :
10:       $Q_1 \leftarrow \text{GLOBALMC}(\mathcal{M}, \psi_1, \theta)$ 
11:       $Q_2 \leftarrow \text{GLOBALMC}(\mathcal{M}, \psi_2, \theta)$ 
12:      return  $Q_1 \cap Q_2$ 
13:      $\psi_1 \vee \psi_2$ :
14:       $Q_1 \leftarrow \text{GLOBALMC}(\mathcal{M}, \psi_1, \theta)$ 
15:       $Q_2 \leftarrow \text{GLOBALMC}(\mathcal{M}, \psi_2, \theta)$ 
16:      return  $Q_1 \cup Q_2$ 
17:      $\text{pfix}\langle t_1, t_2 \rangle X\psi$ :
18:       $Q \leftarrow \text{GLOBALMC}(\mathcal{M}, \psi, \theta)$ 
19:      return  $\text{PREIMG}(\mathcal{M}, t_1, t_2, Q, \theta, \text{pfix})$ 
20:      $\text{pfix}\langle t_1, t_2 \rangle G\psi$ :
21:      return G-FIXPOINT( $\mathcal{M}, t_1, t_2, \psi, \theta, \text{pfix}$ )
22:      $\text{pfix}\langle t_1, t_2 \rangle \psi_1 \cup \psi_2$ :
23:      return U-FIXPOINT( $\mathcal{M}, t_1, t_2, \psi_1, \psi_2, \theta, \text{pfix}$ )
24:   end case
25: end procedure

```

We now present the global model checking Algorithm 4. From here on, we denote by pfix any string from the set $\{\epsilon, \exists y_1, \forall y_2, \exists y_1 \forall y_2, \forall y_2 \exists y_1\}$, where ϵ is the empty string. In each of the cases of the algorithms, pfix is assumed to be the *longest* quantifier prefix that matches the input (sub)-formula.

1. The base case in Algorithm 4 (line 3) of φ being an atomic proposition p simply returns the set of states, the labels of which contain p .
2. The boolean cases are straightforward.
3. In the case of Nexttime formula $\text{pfix}\langle t_1, t_2 \rangle X\psi$, the algorithm first computes the state extension Q of the subformula ψ with a recursive call, and then the controllable pre-

image of Q . The computation of the respective controllable pre-image is shown in Algorithm 1. First, if any of t_1 and t_2 is not a variable that appears (i.e., is bound) in the quantifier prefix pfix , the assignment θ is applied to assign its value. Then, for each state s , if the formula $\text{pfix PrF}(\mathcal{M}, s, t_1, t_2, Q)$ is true, the algorithm adds s to the set of controllable states to be returned.

4. Algorithms 2 and 3 compute the extension of closed formulae of the type $\text{pfix}\langle\langle t_1, t_2 \rangle\rangle \chi$ with temporal objective χ starting with G and U respectively. Their structure is similar to that for global model checking of such formulae in ATL (cf. e.g. the algorithm presented in [8, Chapter 9]). They apply the iterative procedures⁴ of computing controllable pre-images that the fixpoint characterizations of the temporal operators G and U yield (ibid.). This is possible for quantified formulae as the quantifiers in formulae from $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ are propagated inside the temporal operators according to the respective fixpoint equivalences, proved in Theorem 3.

Theorem 4 Let $\mathcal{M}$ be a HDMAS , φ a $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ -formula and θ an assignment. Then

$$\llbracket \varphi \rrbracket_{\mathcal{M}}^{\theta} = \text{GLOBALMC}(\mathcal{M}, \varphi, \theta)$$

Proof By induction on the structure of $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ formulae. The boolean cases are straightforward. For nexttime formulae $\text{pfix}\langle\langle t_1, t_2 \rangle\rangle X\psi$ the claim immediately follows from the correctness of Algorithm 1, implied by the semantics of $\text{PrF}(\mathcal{M}, t_1, t_2, \text{pfix}, \llbracket \psi \rrbracket_{\mathcal{M}})$. For formulae of the type $\text{pfix}\langle\langle t_1, t_2 \rangle\rangle G\psi$ and $\text{pfix}\langle\langle t_1, t_2 \rangle\rangle \psi_1 U \psi_2$, it follows from the correctness of Algorithms 2 and 3, justified by Theorem 3. $\square$

For model checking of the full language $\mathcal{L}_{\text{HDMAS}}$, Algorithm 4 is combined with function NF , transforming constructively any $\mathcal{L}_{\text{HDMAS}}$ -formula φ to φ^{NF} in $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$, equivalent in the finite to φ by virtue of Theorem 2.

Example 7 We illustrate Algorithm 4 by sketching its application to the formula $\psi = \langle\langle 7, 4 \rangle\rangle X(\forall y_2 \exists y_1 \langle\langle y_1, y_2 \rangle\rangle Gp)$ in the HDMAS model $\mathcal{M}$ in Fig. 2. We fix any assignment θ (it does not play any role, since ψ is closed). The outer formula is a X formula, thus line 18 calls recursively the global model checking on the subformula in the temporal objective. Line 4 of $G\text{-FIXPOINT}$ initializes $Z \leftarrow \{s_2, s_3, s_4\}$, viz., states labeled with p and $W \leftarrow S = \{s_1, \dots, s_6\}$. Since $W \not\subseteq Z$, we enter the while cycle computing the fixpoint. In the numbered list below, each item i) correspond to the i th iteration cycle.

1. • $W \leftarrow \{s_2, s_3, s_4\}$;
• $\text{PREIMG}(\mathcal{M}, y_1, y_2, \{s_2, s_3, s_4\}, \theta, \forall y_2 \exists y_1) = \{s_2, s_4, s_5\}$;
• $Z \leftarrow \{s_2, s_4, s_5\} \cap \{s_2, s_3, s_4\} = \{s_2, s_4\}$.
2. • $W \leftarrow \{s_2, s_4\}$;
• $\text{PREIMG}(\mathcal{M}, y_1, y_2, \{s_2, s_4\}, \theta, \forall y_2 \exists y_1) = \{s_2, s_4, s_5\}$;
• $Z \leftarrow \{s_2, s_4, s_5\} \cap \{s_2, s_3, s_4\} = \{s_2, s_4\}$.

Now $W \leftarrow Z$ then the fixpoint is reached.

The set Z is then returned, so $\llbracket \forall y_2 \exists y_1 \langle y_1, y_2 \rangle G p \rrbracket_{\mathcal{M}} = \{s_2, s_4\}$. We now move to the outer next formula for which line 19 of GLOBALMC algorithm calls the PREIMG procedure. For each $s \in S$ the truth of formula $\text{PrF}(\mathcal{M}, s, 7, 4, \{s_2, s_4\})$ is called. The final result is $\llbracket \psi \rrbracket_{\mathcal{M}} = \{s_4, s_5\}$.

Example 8 Consider $\varphi = \langle 6, 3 \rangle X (\exists y_1 \langle y_1, 10 \rangle (\forall y_2 \exists y_1 \langle y_1, y_2 \rangle G p) U (\forall y_2 \langle 0, y_2 \rangle G q))$. We start by computing the extension of $\forall y_2 \langle 0, y_2 \rangle G q$, following Algorithm 2.

From lines 4 – 6: $Q \leftarrow \llbracket q \rrbracket_{\mathcal{M}} = \{s_5, s_6\}$; $W \leftarrow \{s_1, \dots, s_6\}$, and $Z \leftarrow \{s_5, s_6\}$.

Since $W \not\subseteq Z$, we enter the iteration cycle:

1. • $W \leftarrow \{s_5, s_6\}$;
• $\text{PREIMG}(\mathcal{M}, 0, y_2, \{s_5, s_6\}, \theta, \forall y_2) = \{s_6\}$
• $Z \leftarrow \{s_6\} \cap \{s_5, s_6\} = \{s_6\}$.
2. • $W \leftarrow \{s_6\}$;
• $\text{PREIMG}(\mathcal{M}, 0, y_2, \{s_6\}, \theta, \forall y_2) = \{s_6\}$;
• $Z \leftarrow \{s_6\} \cap \{s_5, s_6\} = \{s_6\}$.

The fixpoint is reached and $\llbracket \forall y_2 \langle 0, y_2 \rangle G q \rrbracket_{\mathcal{M}} = \{s_6\}$.

From Example 7 we get $\llbracket \forall y_2 \exists y_1 \langle y_1, y_2 \rangle G p \rrbracket_{\mathcal{M}} = \{s_2, s_4\}$. We then move to computing the extension of the until formula, following Algorithm 3. From lines 4 – 7:

$Q_1 \leftarrow \{s_2, s_4\}$; $Q_2 \leftarrow \{s_6\}$; $W \leftarrow \emptyset$ and $Z \leftarrow \{s_6\}$.

Since $Z \not\subseteq W$, we enter the iteration cycle:

1. • $W \leftarrow \{s_6\}$;
• $\text{PREIMG}(\mathcal{M}, y_1, 10, \{s_6\}, \theta, \exists y_1) = \{s_4, s_6\}$.
Indeed, from s_4 , e.g., 40 controllable agents performing act_1 guarantee that guard g_4 is satisfied.
• $Z \leftarrow \{s_6\} \cup (\{s_4, s_6\} \cap \{s_2, s_4\}) = \{s_4, s_6\}$.
2. • $W \leftarrow \{s_4, s_6\}$;
• $\text{PREIMG}(\mathcal{M}, y_1, 10, \{s_4, s_6\}, \theta, \exists y_1) = \{s_2, s_4, s_6\}$;
• $Z \leftarrow \{s_6\} \cup (\{s_2, s_4, s_6\} \cap \{s_2, s_4\}) = \{s_2, s_4, s_6\}$.
3. • $W \leftarrow \{s_2, s_4, s_6\}$;
• $\text{PREIMG}(\mathcal{M}, y_1, 10, \{s_2, s_4, s_6\}, \theta, \exists y_1) = \{s_2, s_4, s_5, s_6\}$;
• $Z \leftarrow \{s_6\} \cup (\{s_2, s_4, s_5, s_6\} \cap \{s_2, s_4\}) = \{s_2, s_4, s_6\}$.

The fixpoint is reached. Thus:

$$\llbracket \exists y_1 \langle y_1, 10 \rangle (\forall y_2 \exists y_1 \langle y_1, y_2 \rangle G p) U (\forall y_2 \langle 0, y_2 \rangle G q) \rrbracket_{\mathcal{M}} = \{s_2, s_4, s_6\}.$$

Lastly, we call $\text{PREIMG}(\mathcal{M}, 6, 3, \{s_2, s_4, s_6\}, \theta, e)$ to compute $\llbracket \varphi \rrbracket_{\mathcal{M}} = \{s_1, s_4, s_5, s_6\}$.

5 Complexity estimates

As well-known from [1], the time complexity of model checking of ATL formulae is linear in both the size of the model¹⁰ and the length of the formula. Note that in standard concurrent game models the number of agents is fixed and the transition relation is represented explicitly, by means of transitions from each state labelled with each action profile. In HDMAS models, however, the transitions are represented symbolically, in terms of the guards that determine them. An explicit representation would be infinite, in general. Thus, the question of how to measure the size of HDMAS models arises. Given a HDMAS $\mathcal{M}$, we consider the following parameters: the size $|S|$ of the state space; the size n of the action set Act , and the size $|\delta|$ of the symbolic transition guard function. The latter is defined as the sum of the length of all guards appearing in δ , where we assume a binary encoding of numbers.

Given a $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ formula φ and a HDMAS $\mathcal{M}$, the number of fixpoint computations in the global model checking algorithm is bounded by the length of $|\varphi|$. Each computation executes the while cycle at most $|S|$ times, and at each iteration, the function `PREIMAGE` is called. The pre-image algorithm cycles through all states again and invokes model checking of a PrA formula PrF each time. In the worst case $|\text{PrF}| = |\delta|$, as g_O^s could be the disjunction of almost all guards in $\mathcal{M}$. The complexity of checking the truth of a PrA-formula depends not just on its size, but more precisely on the numbers of quantifier alternations and of quantified variables in any quantifier block (cf. [12]). In our case, the maximum number of quantifier alternations is 4, while the number of variables in any quantifier block is at most $n + 1$. By applying results from [11] (cf. also [12]), these yield a worst case complexity Σ_3^{EXP} , or more precisely $\text{STA}(*, 2^{|\delta|^{O(1)}}, 3)$ when the model is not fixed, or at least n is unbounded, but it is down to $\text{STA}(*, |\delta|^{O(1)}, 3)$ when n is fixed.

Thus, the number of variables and quantifier alternation depth in PrF-formulas crucially affect the complexity of model checking of $\mathcal{L}_{\text{HDMAS}}^{\text{NF}}$ -formulae. We can distinguish the following cases of lower complexity bounds:

1. When no quantifier patterns $\exists y_1 \forall y_2$ occur, the maximal alternation depth is 3, hence the complexity is reduced to $\text{STA}(*, 2^{|\delta|^{O(1)}}, 2)$, respectively $\text{STA}(*, |\delta|^{O(1)}, 2)$.
2. If no quantification $\forall y_2$ is allowed, but the number of uncontrollable agents is a parameter, the maximal alternation depth is 2, hence the complexity is reduced to $\text{STA}(*, 2^{|\delta|^{O(1)}}, 1)$, respectively $\text{STA}(*, |\delta|^{O(1)}, 1)$.
3. In the case when the number of either controllable or uncontrollable agents is fixed or bounded, the resulting PrF-formulas become either existential or universal (by replacing the quantifiers over the actions of the bounded set of agents with conjunctions, resp. disjunctions). In these cases, the complexity drops to NP-complete if the number of actions is unbounded, resp. P-complete if that number is fixed or bounded.

6 Concluding remarks

We have proposed and explored a new, generic framework for modelling, formal specification and verification of dynamic multi-agent systems, where agents can freely join and leave during the evolution of the system. We consider indistinguishable agents and therefore the system evolution is affected only by the number of agents performing actions. As

¹⁰ The simplified algorithm presented here works in quadratic time.

neither of the currently available logics are well-suited for expressing properties of such dynamic models, we have devised a variation of the alternating time temporal logic ATL to specify strategic abilities of coalitions of controllable versus non-controllable of agents.

The framework and results presented here are amenable to various extensions, e.g. allowing any PrA-formulae as guards in HDMAS models; allowing more expressive languages, e.g. with arbitrary LTL or parity objectives, with somewhat more liberal quantification patterns in $\mathcal{L}_{\text{HDMAS}}$ (i.e., formulae of the type $\forall y \langle \langle y, y \rangle \rangle X \varphi$ and $\exists y \langle \langle y, y \rangle \rangle X \varphi$ can be added easily), adding several super-agents with controllable sets of agents, etc. The main technical challenge for some of these extensions would be to lift or extend the model checking procedure for them. Still, in particular, extending the present framework to include any finite number of different agent “types”, with each type having a different protocol, is rather straightforward, as follows. Let us fix a set of agent types $\{T_1, \dots, T_m\}$. Now each agent belong to one specific type. Definition 4 will then have $d_1, \dots, d_m$ action availability functions, one for each type, so that agents belonging to the same type have the same set of available actions in each system state, but agents belonging to different types might have different available actions. Lastly the logic will now involve m variables for the controllable agents of each type, and m other variables for the non-controllable ones of each type. The same restrictions on the use of these variables will apply in this extended logic and the notion of normal form, the technical results related to it, and the model checking algorithm for formulae in normal form, extend as expected to the multi-type case.

Of the numerous possible applications we only mention a natural link with the Colonel Blotto games [5, 20], where two players simultaneously distribute military force units across n battlefields, and in each battlefield the player (if any) that has allocated the higher number of units wins. As suggested by our fortress example, our framework can be readily applied to model and solve algorithmically multi-player and multiple-round extensions of Colonel Blotto games, which we leave to future work. More generally, dynamic resource allocation games [3] as well as verification of parameterised fault-tolerance in multi-agent systems [16] seem naturally amenable to applications of the present work.

Acknowledgements Open access funding provided by Stockholm University. This work was supported by research Grant 2015-04388 of the Swedish Research Council. We thank the reviewers for constructive comments and suggestions for improvement of the paper.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Alur, R., Henzinger, T. A., & Kupferman, O. (2002). Alternating-time temporal logic. *Journal of the ACM*, 49(5), 672–713.
2. Angluin, D., Aspnes, J., Diamadi, Z., Fischer, M. J., & Peralta, R. (2004). Computation in networks of passively mobile finite-state sensors. In *Proceedings of the twenty-third ACM PODC St. John's, Canada* (pp. 290–299).
3. Avni, G., Henzinger, T. A., & Kupferman, O. (2016). Dynamic resource allocation games. In M. Gairing & R. Savani (Eds.), *Algorithmic game theory* (pp. 153–166). Berlin: Springer.

4. Bloem, R., Jacobs, S., Khalimov, A., Konnov, I., Rubin, S., Veith, H., et al. (2016). Decidability in parameterized verification. *SIGACT News*, 47(2), 53–64.
5. Borel, E. (1953). The theory of play and integral equations with skew symmetric kernels. *Econometrica: Journal of the Econometric Society*, 21, 97–100.
6. Bulling, N., Goranko, V., & Jamroga, W. (2015). Logics for reasoning about strategic abilities in multi-player games. In J. van Benthem, S. Ghosh, & R. Verbrugge (Eds.), *Models of strategic reasoning: Logics, games, and communities* (pp. 93–136). Berlin: Springer.
7. De Giacomo, G., Vardi, M. Y., Felli, P., Alechina, N., & Logan, B. (2018). Synthesis of orchestrations of transducers for manufacturing. In *Proceedings of AAI-18, New Orleans, USA* (pp. 6161–6168).
8. Demri, S., Goranko, V., & Lange, M. (2016). *Temporal Logics in Computer Science. Cambridge Tracts in Theoretical Computer Science*. Cambridge: Cambridge University Press.
9. Esparza, J., Ganty, P., Leroux, J., & Majumdar, R. (2016). Model checking population protocols. In *36th IARCS annual conference on FSTTCS, Chennai, India* (pp. 27:1–27:14).
10. German, S. M., & Sistla, A. P. (1992). Reasoning about systems with many processes. *Journal of the ACM*, 39(3), 675–735.
11. Haase, C. (2014). Subclasses of Presburger arithmetic and the weak EXP hierarchy. In *Proceedings of CSL-LICS'14* (pp. 47:1–47:10).
12. Haase, C. (2018). A survival guide to Presburger arithmetic. *SIGLOG News*, 5(3), 67–82.
13. Jamroga, W., Ågotnes, T. (2007). Modular interpreted systems. In *Proceedings of AAMAS* (pp. 131:1–131:8). ACM
14. Jamroga, W., Knapik, M., & Kurpiewski, D. (2018). Model checking the SELENE e-voting protocol in multi-agent logics. In *Proceedings of the third international joint conference on electronic voting, E-Vote-ID 2018. Lecture notes in computer science* (Vol. 11143, pp. 100–116). Springer.
15. Kouvaros, P., & Lomuscio, A. (2016). Parameterised verification for multi-agent systems. *Artificial Intelligence*, 234, 152–189.
16. Kouvaros, P., & Lomuscio, A. (2017). Verifying fault-tolerance in parameterised multi-agent systems. *Proceedings of IJCAI, 2017*, 288–294.
17. Kouvaros, P., Lomuscio, A., Pirovano, E., & Punchihewa, H. (2019). Formal verification of open multi-agent systems. In *Proceedings of AAMAS'19* (pp. 179–187).
18. Pedersen, T., & Dyrkolbotn, S. K. (2013). Agents homogeneous: A procedurally anonymous semantics characterizing the homogeneous fragment of ATL. In *Proceedings of PRIMA* (Vol. 2013, pp. 245–259).
19. Raskin, J., Samuelides, M., & Begin, L. V. (2005). Games for counting abstractions. *Electronic Notes in Theoretical Computer Science*, 128(6), 69–85.
20. Roberson, B. (2006). The Colonel Blotto game. *Economic Theory*, 29(1), 1–24.
21. Vinyals, M., Rodriguez-Aguilar, J. A., & Cerquides, J. (2011). A survey on sensor networks from a multiagent perspective. *The Computer Journal*, 54(3), 455–470. <https://doi.org/10.1093/comjnl/bxq018>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.