Playing with Norms: Tractability of Normative Systems for Homogeneous Game Structures

Sjur Dyrkolbotn
Durham School of Law, Durham University
Palatine Centre, Stockton Road
Durham DH1 3LE, UK
s.k.dyrkolbotn@durham.ac.uk

Piotr Kaźmierczak
Dept. of Computing, Mathematics and Physics
Bergen University College
P.O. Box 7030, 5020 Bergen, Norway
phk@hib.no

ABSTRACT

In this paper we study the connection between anonymous (normal form) games and normative systems for homogeneous concurrent game structures. We present a variant of the popular strategic logic ATL that allows for reasoning about norm compliance and the strategic ability of coalitions of agents. We broaden the notion of a normative system compared to earlier work and introduce two categories: individual and collective norms. Then we establish a technical link between these two notions, and between the notion of an individual norm and an anonymous game, as studied in recent work on algorithmic game theory. These connections enable us to show that model checking on homogeneous concurrent game structures with norms is tractable in the number of agents.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Theory, Economics

Keywords

Strategic logics; Anonymous games; Normative systems; Tractability; Computational complexity; Concurrent Game Structures

1. INTRODUCTION

Normative systems (or *social laws*) have received a lot of attention in the multi-agent systems community in recent years, and have proved to be a very useful framework for agent coordination.[1] The idea is simple: we put *behavioral constraints* on agents and check whether they can achieve given *objectives* while complying to these constraints. While initially proposed and analyzed by Shoham & Tennenholtz [13, 14], normative systems were further studied using modal logic by Ågotnes *et al.* [1, 3, 2, 4] with *Computation Tree Logic* (CTL) as a language for expressing objectives and agent-labeled Kripke structures as system models, and by van der Hoek *et al.* [9] who used *Alternating-time temporal logic* (ATL) to express objective and *Concurrent Game Structures* (CGSs) as models.

The normative systems used in these papers suffer from two major shortcomings:

1. They are not expressive enough: A normative system is understood as a list of forbidden actions, one list per agent, per state of the system. That does not allow us to model *normative interaction*. There is no coordination involved in norm compliance – the normative constraints imposed on an agent's behavior do not in any way depend on the actions performed by the other agents.

2. They are not tractable: Checking whether or not a given normative system can ensure that a given objective is met is typically (co)-NP-complete, and exact algorithms tend to use time exponential in the number of agents.

In this paper we address both of these shortcomings by introducing a logic for reasoning about a broader notion of norms, for which we can also show that model checking can be done in time polynomial in the number of agents in the system. This means that the logic is tractable for analyzing and verifying properties of large scale systems containing a great number of agents, as long as it involves a constant number of other elements, such as states and actions.

Tractability follows from imposing a restriction on the underlying system, formulated with respect to CGS's. We require, in particular, that the system is *homogeneous*, meaning that it allows a branching-time future that depends solely on the *number* of agents performing various actions, and does not differentiate based on agents' identity. This restriction stems from [11] where an equivalent role-based semantics for ATL was given to facilitate fast model checking of models with few roles. Homogeneity arises from restricting attention to models with a single role, and the corresponding class of structures was axiomatized in [12].

The norms we consider involve no corresponding homogeneity restriction and they can prescribe different sets of

[1]Throughout the paper we use 'normative systems' or 'norms' whenever we refer to behavioral constraints on our agents. It is similar to the notion of a 'social law' as defined by Shoham & Tennenholtz [13, 14].

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.*

legal actions to different agents, depending on the state of the system. Furthermore, we introduce norms that allow us to model coordination, such that the set of legal actions for an agent depends on the actions performed by other agents, up to equivalence with respect to the underlying structure. We consider two ways of doing this:

- Individually: Each agent has a list of forbidden actions, but the list depends on what the other agents do.

- Collectively: The normative system forbids tuples of actions directly.

We demonstrate that these two notions are equivalent as long as we assume full norm compliance. Then we discuss partial compliance and argue that individual norms raise some conceptual questions that we can only answer successfully if we recognize that individual and collective norms are genuinely different, and require different notions of partial compliance in order to do justice to fundamental intuitions and modeling aims.

Adding normative systems to homogeneous structures results in a logic that no longer exhibits homogeneity with respect to the strategic ability of agents under norm compliance. Hence norms serve to *reintroduce* individuality to the system. However, we show that model checking remains tractable. This follows from a link with recent work in algorithmic game theory, namely showing how normative systems can be efficiently *incentivized* by a normal form game which rewards agents for following the norm. We observe, in particular, that the set of Nash Equilibria for this game corresponds exactly to the legal actions for the agents under the corresponding normative system. Furthermore, we show that all such normative games belong to a class of games which admit compact representation facilitating polynomial time computation of the set of Nash Equilibria.

In Section 2 we introduce basic notation and present a terse summary of the main technical concepts we rely on. Then, in Section 3 we introduce formal definitions of collective and individual norms for homogeneous concurrent game structures and present a technical elaboration of these concepts that culminates in the definition of truth for a language of norm compliance on homogeneous structures with norms. In Section 4 we present our main result, encoding normative systems as anonymous games, and we use this correspondence to establish tractability of model checking. We conclude and discuss directions for future work in Section 5.

2. PRELIMINARIES

The technical results presented in this paper arise from a novel combination of results on three distinct formal notions: homogeneous concurrent game structures, anonymous normal form games, and normative systems. We now proceed with detailed, but somewhat terse, definitions of all the necessary formal background regarding the first two concepts, and we devote a separate section to normative systems, since our definitions significantly broaden the notion of a norm previously considered in related work on strategic logics for multi-agent systems.

Notation and basic concepts

We start by introducing some notation and basic definitions that we will use in the remainder of the paper. Throughout we assume given a set $\Sigma = \{1, \ldots, n\}$ of agents and a set $\mathbb{A} = \{p_1, \ldots, p_m\}$ of actions.[2] We say that given any function $s : X \to Y$ we let $dom(s) = X$ denote its domain and $targ(s) = \{y \in Y \mid \exists x \in X : s(x) = y\}$ denote the subset of the target that s is onto. We let s_{-i} denote the function $s \upharpoonright dom(s) \setminus \{i\}$ (where $\upharpoonright$ is used to signify partial function application) and we let (s_{-i}, p) denote the function s' defined by:

$$s'(j) = \begin{cases} s(j) \text{ if } j \neq i \\ p \text{ otherwise.} \end{cases}$$

Also, for an arbitrary function $s : X \to Y$ we let $s^- : Y \to 2^X$ denote the function defined by:

$$s^-(y) = \{x \mid s(x) = y\}$$

for all $y \in Y$. We will often use the notation $\#(y, s) = |s^-(y)| = |\{x \in X \mid s(x) = y\}|$ and $(\#(y, s))_{y \in Y}$, which we will refer to as the *profile induced by s*. It is a vector which records, for each $y \in Y$, the number of elements of X that "choose" y. Given a coalition $C \subseteq \Sigma$, an action-tuple for C is a function $s \subseteq \mathbb{A}^C$, where $s(i)$ is player i's action. Hence the function-space $\mathbb{A}^C$ contains all possible action-tuples for the coalition C. Using this notation, we define the set $P^C(\mathbb{A}) = \{(\#(y, s))_{y \in \mathbb{A}} \mid s \in \mathbb{A}^C\}$. We refer to elements $F \in P^C(\mathbb{A})$ as *C-profiles* or just partial profiles if C is not specified. Given a profile $F \in P^C(\mathbb{A})$ and an action $p_1 \in \mathbb{A}$, we use (F, p_1) to denote the profile F' defined by

$$F'(p_2) = \begin{cases} F(p_2) + 1 \text{ if } p_2 = p_1 \\ F(p_2) \text{ otherwise.} \end{cases}$$

We let F_{-p_1} denote the profile F' defined by

$$F'(p_2) = \begin{cases} F(p_1) - 1 \text{ if } p_2 = p_1 \\ F(p_2) \text{ otherwise.} \end{cases}$$

Moreover, for two profiles F, F' the profile $F + F'$ is defined for all $p \in \mathbb{A}$ as follows:

$$F + F'(p) = F(p) + F'(p)$$

For all $F \in P^C(\mathbb{A})$, we also use $ext(F) = \{F' \in P^\Sigma(\mathbb{A}) \mid F' \upharpoonright C = F\}$. We notice that we can generate the set $P^C(\mathbb{A})$ without making use of the set of action-profiles for C, hence we avoid traversing a set that has exponential size in the number of agents. In particular, given a coalition C and a vector $F \in \mathbb{N}^\mathbb{A}$ such that $\sum_{p \in \mathbb{A}} F(p) = |C|$, we observe that there is at least one action-tuple $s_F \in \mathbb{A}^C$ such that $(\#(y, s_F))_{y \in \mathbb{A}} = F$. Conversely, we have, for all $s \in \mathbb{A}^C$, $\sum_{p \in \mathbb{A}} \#(p, s) = |C|$. It follows that $P^C(\mathbb{A}) = \{F \in \mathbb{N}^\mathbb{A} \mid \sum_{p \in \mathbb{A}} F(p) = |C|\}$ is a compact characterization of the set of profiles for a coalition.

NCHATL and Concurrent Game Structures

The logical language we use, $\mathcal{L}_{\text{NCHATL}}$, is based on ATL [5], extended with one extra operator that we use to express

[2]We will only consider semantic structures for which the actions are shared among the agents.

norm compliance. Formally, the language is generated by the following BNF:

$$\varphi ::= \top \mid p \mid \neg\varphi \mid \varphi\vee\varphi \mid \langle\langle C\rangle\rangle\bigcirc\varphi \mid \langle\langle C\rangle\rangle\Box\varphi \mid \langle\langle C\rangle\rangle\varphi\mathcal{U}\varphi \mid \langle C\rangle\varphi$$

where p is a propositional symbol, and $C \subseteq \Sigma$ is a coalition of agents.[3]

The language of NCHATL contains three types of modalities for talking about pairs (H, χ) where H is a semantic structure and χ is a normative system.

- $\bigcirc$, $\Box$ and $\mathcal{U}$ are standard *temporal* operators known from many temporal logics, and stand for "next state", "some future state" and "until", respectively;
- $\langle\langle C\rangle\rangle$ is a *strategic ability* operator, and its intuitive meaning is that the coalition $\langle\langle C\rangle\rangle \bigcirc \varphi$ has a *joint strategy* for enforcing a formula φ in the next state of H;
- finally $\langle C\rangle$ is the *norm compliance* operator, which has the intuitive reading that the agents in C are willing to comply to χ.

For example, if we have some objective φ and we wonder if C complying to χ enables D to ensure φ in the next state of H, we would test if the formula $\langle C\rangle\langle\langle D\rangle\rangle \bigcirc \varphi$ is true on (H, χ). In this way, properties of H that depend on the choices of the agents can be specified with great flexibility using χ, without altering the specification of the underlying system H. In particular, rather than hard-coding predictions and intentions concerning agent behavior in the system design, we give this aspect a separate semantic characterization using χ, enabling us also to address and study interactions between the distinct modalities of norm compliance and strategic ability. Further motivation for using normative systems in verification and specification of multi-agent systems can be found in [4].

In the following, we specify H and χ compactly, and in a way that allows for great freedom in the choice of χ. To this end, we first recap the definition of a *concurrent game structure* [5], which provides the backbone for our definition of H.

Definition 1. A CGS is a tuple $S = \langle \Sigma, Q, \Pi, (\mathbb{A}_q)_{q\in Q}, \pi, \delta\rangle$ where:

- Q is the non-empty set of states.
- Π is a set of propositional letters and $\pi : Q \to 2^{\Pi}$ maps each state to the set of propositions true in it.
- $\mathbb{A}_q \subseteq \mathbb{A}$ is the set of actions available at q.
- δ is the transition function. For each $q \in Q$ and any action-tuple $s \in \mathbb{A}_q^{\Sigma}$, it returns a state $q' = \delta(q, s) \in Q$, referred to as a *successor* of q.

The crucial condition which allows for tractable model checking of $\mathcal{L}_{\text{NCHATL}}$ is *homogeneity*, which was first introduced (implicitly) in [11], and subsequently axiomatized in [12]. It is formalized as follows, such that S is said to be homogeneous if the following hold, for all $q \in Q$ and all action-tuples $s, s' \in \mathbb{A}_q^{\Sigma}$:

$$(\#(y, s))_{y\in\mathbb{A}_q} = (\#(y, s'))_{y\in\mathbb{A}_q} \Rightarrow \delta(q, s) = \delta(q, s') \qquad (1)$$

[3]This language closely resembles the language of *Norm Compliance* CTL presented in [4].

Alternatively, the axiomatic characterization in terms of ATL (using the language of $\mathcal{L}_{\text{NCHATL}}$ without the norm-compliance operator) is also very intuitive, saying that a CGS is homogeneous as long as any two coalitions of the same size have the same strategic ability in the next-time step:

$$\langle\langle C\rangle\rangle \bigcirc \varphi \leftrightarrow \langle\langle D\rangle\rangle \bigcirc \varphi \text{ if } |C| = |D| \qquad (2)$$

For space reasons we omit the formal definition of truth on CGSs, and only remark that the crucial property of homogeneous CGSs is that they permit testing the truth of ATL-formulas with only a polynomial time dependence on the number of agents.[4] This result follows from compact representation of such structures, obtained by replacing explicit action-tuples by profiles. In particular, a homogeneous CGS can be equivalently represented by a structure of the following kind [11, 12]:

Definition 2. An HCGS is a tuple $H = \langle \Sigma, Q, \Pi, \pi, (\mathbb{A}_q)_{q\in Q}, \delta\rangle$ where

- Σ, Q, Π, π and $(\mathbb{A}_q)_{q\in Q}$ are defined as in Definition 1,
- For every state $q \in Q$, and every $F \in P(q) = P^{\Sigma}(\mathbb{A}_q)$, δ maps F to a successor, $\delta(q, F) = q' \in Q$.

The profiles $F \in P(q)$ assign a natural number to each action such that the sum of these numbers (over all actions) sums up to n (the number of agents). The intended meaning is that the profile describes, for all actions, how many agents perform that action. We also have C-profiles at $q \in Q$, for all $q \in Q, C \subseteq \Sigma$

$$P^C(q) = P^C(\mathbb{A}_q) = \left\{ F \in \mathbb{N}^{\mathbb{A}_q} \;\middle|\; \sum_{p\in\mathbb{A}_q} F(p) = |C| \right\} \qquad (3)$$

Example 1. Consider an example of an HCGS that models a jazz trio.

- The trio consists of three players: Keith, Gary and Jack. $\Sigma = \{k, g, j\}$.
- A scenario that we are modeling is a concert, which makes for two states: one where the players are idle (pre-performance), and one in which they are playing. $Q = \{q_0, q_1\}$, $\Pi = \{idle, play\}$, $\pi(q_0) = \{idle\}, \pi(q_1) = \{play\}$.
- Agents have four actions available in the initial state: to remain idle, to play the piano, to play the double-bass, and to play the drums, which are actions p_1, p_2, p_3 and p_4, respectively. In state q_1, agents can perform action *done*, which is action p_5. $\mathbb{A}_{q_0} = \{p_1, p_2, p_3, p_4\}$, $\mathbb{A}_{q_1} = \{p_5\}$.
- Every agent is allowed to play, and to become 'idle' once the concert is over. We only list transitions that lead from one state to a new one, all other possible transitions are reflexive arrows. In particular, we have $\delta(q_0, 0, 1, 1, 1) = q_1$ and $\delta(q_1, 3) = q_0$.

[4]This is not true for CGSs in general since their size is typically exponential in the number of agents. There are, in particular, exponentially many distinct action-tuples at every state. This is a major argument against the usefulness of ATL in practical multi-agent modeling scenarios, as pointed out, for instance, in [10].

We return to this scenario in Example 2 when we introduce a normative system to help ensure that the concert becomes a success.

We postpone a formal definition of truth on HCGSs until Definition 8, when we will define truth for (H, χ), where H is an HCGS and χ is a normative system. Before we turn to a formalization of norms, we present the second basic building block of our approach, namely *anonymity* in the sense of algorithmic game theory [6, 7].

Anonymous games

A normative system describes constraints on behavior, but does not provide any mechanism for making sure that agents conform to these constraints. To do this, one may associate a mechanism of reward and/or punishment with the normative system, to provide an incentive for agents to follow the norm. We will do this in Section 4, partly to show how our norms can be backed up by an incentive mechanism, but also to show how recent work in algorithmic game theory can be used to ensure efficient model checking. In the following we provide the necessary background on the game theoretic notions we rely on.

Definition 3. A (normal form) *game* (with shared actions) is a tuple $\Gamma = (\Sigma, \mathbb{A}, (\mu_i)_{i\in\Sigma})$ such that for all action-tuples $s \in \mathbb{A}^{\Sigma}$, $\mu_i(s) \in \mathbb{R}$ is the *payoff* for agent i. We also define the following notions:

- For any action-tuple s, agent i's best response to s is $\mathsf{b}_i(s) = \{p \in \mathbb{A} \mid \forall p' \in \mathbb{A} : \mu_i(s_{-i}, p) \geq \mu_i(s_{-i}, p')\}$.
- An action-tuple s is a Nash Equilibrium in pure strategies if for all $i \in \Sigma$ we have $s(i) \in \mathsf{b}_i(s)$. That is, no player can do better by changing his action, assuming all other actions remain fixed. We let $\mathsf{NE}(\Gamma)$ denote the set of all Nash Equilibria in Γ.

We notice that for any game with shared actions the number of distinct Σ-profiles is polynomial in Σ. This motivates studying classes of games for which the payoffs of players remain invariant under (classes of) action-tuples that induce the same profile. The following two notions from recent work in game-theory both rely on this idea [6, 7].

Definition 4. A game is said to be:

- *Homogeneous* if for all agents $i \in \Sigma$ and all action-tuples $s, s' \in \mathbb{A}^{\Sigma}$, we have
$$(\#(y, s))_{y\in\mathbb{A}} = (\#(y, s'))_{y\in\mathbb{A}} \Rightarrow \mu_i(s) = \mu_i(s')$$
- *Anonymous* if for all agents $i \in \Sigma$ and all action tuples $s, s' \in \mathbb{A}^{\Sigma}$, we have
$$s(i) = s'(i) \ \& \ ((\#(y, s_{-i}))_{y\in\mathbb{A}} = (\#(y, s'_{-i}))_{y\in\mathbb{A}} \Rightarrow \mu_i(s) = \mu_i(s')$$

In short, the intuition behind homogeneous games is that outcomes for the players remain invariant under permutations of actions over the players.[5] In anonymous games, agents may differentiate between different tuples that induce the same profile, but only with respect to their own action; they do not care about the identity of the agents performing the remaining actions that induces the profile.

For homogeneous and anonymous games, we can express the payoff for each player as a function of profiles rather than explicit action-tuples. To this end we first define, for all $C \subseteq \Sigma, F \in P^C(\mathbb{A})$ and all $1 \leq k \leq m$ the following number, which will allow us to define a canonical action-tuple s_F that induces F for any $F \in P^C(\mathbb{A})$. Keep in mind here that $\mathbb{A} = \{p_1, \dots, p_m\}$, since we will rely on this ordering below.

$$\begin{array}{l} \rho(0, F) = 1 \ \& \ \rho(1, F) = F(p_1) \\ \rho(i, F) = \rho(i-1, F) + F(p_i) \text{ for all } 2 \leq i \leq m \end{array} \quad (4)$$

Then we can define $s_F \in \mathbb{A}^C$ as follows, for all $F \in P^C(\mathbb{A}), i \in \Sigma$:

$$s_F(i) = p_j \text{ for } j \text{ such that } \rho(j-1, F) \leq i \leq \rho(j, F) \quad (5)$$

Clearly, we have $(\#(y, s_F))_{y\in\mathbb{A}} = F$, and this representation (easily computable in polynomial time) can now be used to define a compact version of the payoff function for homogeneous and anonymous games.

$$\mu_i(F, p) ::= \mu_i(s_F, p) \text{ for all } F \in P^{\Sigma\setminus\{i\}}(\mathbb{A}), p \in \mathbb{A}$$

The action p in $\mu_i(F, p)$ is to be thought of as the action performed by agent i. Notice that we need to keep track of this p explicitly in order to also cover the case of anonymous games. We may now lift the best response function accordingly, for all agents $i \in \Sigma$.

$$\forall F \in P^{\Sigma\setminus\{i\}}(\mathbb{A}) : \mathsf{b}_i(F) = \{p \mid \forall r \in \mathbb{A} : \mu_i(F, r) \leq \mu_i(F, p)\} \quad (6)$$

From the definition of anonymity and homogeneity, it immediately follows that this representation is faithful, in the following sense, for all $s \in \mathbb{A}^{\Sigma}$:

$$\mathsf{b}_i(s) = \mathsf{b}_i((\#(y, s_{-i}))_{y\in\mathbb{A}}) \quad (7)$$

The crucial property of homogeneous and anonymous games is that the set of Nash Equilibria can be computed by looking only at the representation that relies on profiles rather than explicit action-tuples. In particular, as first observed in [6] and elaborated upon in [7], it follows from Hall's marriage theorem, a fundamental result from combinatorics [8], that a profile $F \in P^{\Sigma}(\mathbb{A})$ corresponds to a Nash Equilibrium s of Γ with $(\#(y, s))_{y\in\mathbb{A}} = F$ if, and only if, the following condition holds:

$$\forall P \subseteq \mathbb{A} : |\{i \in \Sigma \mid \exists p \in P : p \in \mathsf{b}_i(s_F)\}| \geq \sum_{p\in P} F(p) \quad (8)$$

From this characterization and Equation 7 it is easy to see that computing the set of profiles corresponding to $\mathsf{NE}(\Gamma)$ can be done in polynomial time in the number of agents. In fact, as shown in [7], the problem is in the complexity class TC_0. We will make use of this fact later to show that the normative systems for homogeneous structures admit tractable model checking procedures.

3. NORMATIVE SYSTEMS AND TRUTH

Unlike previous work on formal logical representation of normative systems, we do not restrict attention to simple lists of forbidden actions. Rather, we want to consider norms that reflect the way in which norms are inherently *social*, such that what an agent should do can depend on what other

[5]This is sometimes called *self-anonymous* in the game theory literature.

agents are doing. Compliance to such norms might then require coordination – they can be impossible to comply to unless the agents jointly reach a decision about what to do, so that each agent knows what actions the other agents intend to perform. In light of this, one may also consider norms that do not consist in specifying illegal actions for individual agents, but rather specifies directly the joint actions that are not allowed. In particular, one may consider both *individual* and *collective* normative systems.

Below we provide definitions of both these kinds of normative systems for homogeneous concurrent game structures. Both of them potentially involve coordination, and we observe that under the assumption of full norm compliance the two notions are the equivalent. We argue, however, that for partial compliance, a non-obvious notion is needed to do justice to individual norms, setting them apart from collective ones. The basic definition is as follows:

Definition 5. Given an HCGS H,

- An individual norm is a collection of functions $\eta = \{\eta_i \mid i \in \Sigma\}$, such that for all $q \in Q, i \in \Sigma, F \in P^{\Sigma\setminus\{i\}}(q)$ we have
$$\eta_i(q,F) \subset \mathbb{A}_q$$
- A collective norm is function κ such that for all $q \in Q$
$$\kappa(q) \subset P(q)$$

The normative systems considered in previous work on strategic logics correspond to a restricted class of individual norms, defined below.

Definition 6. Given an HCGS H and an individual normative system η, we say that η is *simple* if for all $q \in Q, i \in \Sigma$ and all $F, F' \in P^{\Sigma\setminus\{i\}}(q)$
$$\eta_i(q,F) = \eta_i(q,F')$$

Simple norms capture normative systems/social laws for which the forbidden actions only depend on the state, as in [4, 9]. Given an individual norm η we define:

$$\mathsf{L}_\eta(q) = \{s \in \mathbb{A}^\Sigma \mid \forall i \in \Sigma : s(i) \in \mathbb{A}_q \setminus \eta_i(q, (\#(y, s_{-i}))_{y\in\mathbb{A}_q})\} \tag{9}$$

These are the *legal* action-tuples at q, when we assume full compliance to the normative system η; everyone chooses an action that is permitted in light of what all the other agents choose to do. We say that an individual normative system is *consistent* if for all $q \in Q$ we have $\mathsf{L}_\eta(q) \neq \emptyset$. Intuitively, consistency of a normative system means that it is possible to comply with it, and in the following we assume that all normative systems are consistent.[6]

We also define the legal action tuples at q when we assume full compliance to a collective norm κ:

$$\mathsf{L}_\kappa(q) = \{s \in \mathbb{A}^\Sigma \mid (\#(y,s))_{y\in\mathbb{A}_q} \in P(q) \setminus \kappa(q)\} \tag{10}$$

[6]It will follow from Theorem 1 that consistency can be checked in polynomial time. The requirement is made here to ensure that our models remain serial after implementation of a norm. Notice that all simple norms are consistent, since all agents have at least one legal action at every state. We mention that this is a standard assumption from the literature, see e.g., [4].

Let us now define:

$$P_\eta(q) = \{(\#(y,s))_{y\in\mathbb{A}_q} \mid s \in \mathsf{L}_\eta(q)\} \tag{11}$$

These are the profiles induced by the legal action tuples at q, and they provide us with a translation of individual norms into collective norms. In particular, given an individual norm η we obtain the collective norm ${}_\eta\kappa$ defined at all $q \in Q$ by

$${}_\eta\kappa(q) = P(q) \setminus P_\eta(q) \tag{12}$$

Let us now go the other way; from collective norms to individual norms. Assume we have given a collective norm κ. Then we first choose some profiles $f = \{f_q \in \mathbb{N}^\mathbb{A} \mid q \in Q\}$ such that for all $q \in Q$ we have $f_q \in P(q) \setminus \kappa(q)$. This allows us to represent the norm κ by the individual norm ${}^f_\kappa\eta$ defined as follows, for all $q \in Q, i \in \Sigma$ and $F \in P^{\Sigma\setminus\{i\}}(q)$:

$${}^f_\kappa\eta_i(q,F) = \{p \in \mathbb{A}_q \mid p \neq s_{f_q}(i) \text{ and } (F,p) \in \kappa(q)\} \tag{13}$$

Here s_{f_q} is the canonical action-tuple inducing f_q, as defined in Equation 5. Adequacy of the translation is now easy to establish, and we omit the proof for space reasons.[7]

PROPOSITION 1. *For any* HCGS H *and any normative system* κ *on* H*, we have* $P_{{}^f_\kappa\eta}(q) = \kappa(q)$ *for all* $q \in Q$*.*

We now turn to partial compliance, first for collective norms. Given a coalition C, we define the set of C-profiles that are compliant to κ as follows:

$$P^C_\kappa(q) = \{F \in P^C(q) \mid ext(F) \cap (P(q) \setminus \kappa(q)) \neq \emptyset\} \tag{14}$$

So a C-profile is compliant to κ if it can be extended to a complete profile that is legal.[8]

Let us now turn to partial compliance for individual norms. It is tempting to say that partial compliance by $C \subseteq \Sigma$ to η can be defined using the translation to collective norms directly, as follows:

$$\{F \in P^C(q) \mid ext(F) \cap P_\eta(q) \neq \emptyset\} \tag{15}$$

However, this is not adequate since it admits F as a partially compliant C-profile whenever F can be induced by a choice of actions for *some* coalition of size $|C|$ that adhere to η. It does not ensure that F can be induced by a choice of actions that is permissible for C. Hence for individual norms, the notion defined below is more appropriate. It refers to the set of explicit action-tuples to ensure that a partially compliant profile for C can actually be induced by actions that are allowed for C. In particular, we define the set of C-profiles that are partially compliant to an individual norm η as follows, for all $q \in Q$:

$$P^C_\eta(q) = \{(\#(y,s))_{y\in\mathbb{A}_q} \mid s \in \mathbb{A}^C \ \&\ \exists s' \in \mathsf{L}_\eta(q) : s = s' \restriction C\} \tag{16}$$

Hence we require that the C-profile can be induced by a C-action-tuple which can be extended to a legal action-tuple. This clearly implies that the profile can be extended to a

[7]In particular, it is straightforward to see that we allow all profiles that are permitted under κ. However, to ensure that every agent always has some legal action to perform, we allow by default the canonical actions that induces the legal profiles f_q, for all $q \in Q$.

[8]For future work we would like to consider more subtle notions of compliance, but this seems like the obvious place to start, a minimum requirement that is hard to dispute.

legal profile, but it is stronger; it also requires the agents in C to act in such a way as *they* are permitted to act under η. It is not enough that *some* collection of $|C|$ agents could have acted in this way.

Truth for NCHATL *on* (H, χ)

We are almost ready to define truth of $\mathcal{L}_{\text{NCHATL}}$ formulas on (H, χ) where H is an HCGS and χ is an individual or a collective norm. In order to do so, we first need to define the notion of a *compliant strategy* with respect to a norm $\chi \in \{\eta, \kappa\}$ which is either collective or individual. This, in turn, requires us to define the set of legal D-profiles given C-compliance to χ:

$$P^C_\chi(q, D) = \{F \in P^D(q) \mid \exists F_1 \in P^{C\cap D}_\chi(q) : \exists F_2 \in P^{D\setminus C}(q) : F = F_1 + F_2\} \tag{17}$$

Notice that this definition requires $C \cap D$ to comply with the norm *independently* from the remaining agents in C. For collective and simple norms, this makes no difference, but for individual norms involving coordination this means that we require something stronger than norm-compliance of C as a coalition. Rather, we interpret D as the coalition of agents which coordinate their actions, so the fact that C follow the norm is taken to mean that when D is acting as a group, without necessarily including all members of C, then $C \cap D$ also follow the norm as a sub-group of C. However, we also require, when defining the extensions of D's strategy below, that C follow the norm as a whole. Hence in effect we model the situation when $C \cap D$ and $C \cap (\Sigma \setminus D)$ follow the norm independently of one another.

We could define this differently by assuming only that C follow the norm as a whole and do not necessarily coordinate their actions inside D to meet the requirements of η. Nothing hinges on this for the results that are about to follow, but our solution seems more natural with respect to the intended readings of $\langle\langle D\rangle\rangle$ and $\langle C\rangle$.

We say that a *computation* is an infinite sequence $\lambda = q_0q_1\ldots$ of states such that for all positions $i \geq 0$, q_{i+1} is a successor of q_i and we follow standard abbreviations. Hence a q-computation denotes a computation starting at q, and $\lambda[i]$, $\lambda[0, i]$ and $\lambda[i, \infty]$ denote the i-th state, the finite prefix $q_0q_1\ldots q_i$ and the infinite suffix $q_iq_{i+1}\ldots$ of λ for any computation λ and its position $i \geq 0$, respectively.

This leads us to the following notion of a *strategy* for a coalition D assuming that C complies with the normative system.

Definition 7. A $\chi \restriction C$-compatible D-strategy is a map $s_D : Q \to \bigcup_{q\in Q} P^C_\chi(q, D)$ such that

$$s_D(q) \in P^C_\chi(q, D) \text{ for each } q \in Q$$

We denote the set of all such strategies by $strat^\chi_C(D)$.

Notice that if $s \in strat^\chi_C(\Sigma)$ for some $C \subseteq \Sigma$, then if we apply $\delta(q)$ to $s(q)$ we obtain a unique new state $q' = \delta(q, s(q))$. Iterating, we get the *induced* computation $\lambda_{s,q} = q_0q_1\ldots$ such that $q = q_0$ and $\forall i \geq 0 : \delta(q_i, (s(q_i))) = q_{i+1}$. Given $s_D \in strat^\chi_C(D)$ and a state q we get an associated *set* of computations $out(s_D, q)$. This is the set of all computations that can result when at any state, D is acting in the way specified by s_D. That is

$$out(s_D, q) := \{\lambda_{s,q} \mid s \in strat^\chi_C(\Sigma) \text{ and } s_D \leq s\} \tag{18}$$

We are now ready for the main definition of this section.[9]

Definition 8. Given a normative HCGS (H, χ) with $\chi \in \{\eta, \kappa\}$ being either a collective or individual norm, a state q and a coalition $C \subseteq \Sigma$, truth of φ on (H, χ) under C-compliance is defined inductively.

- $H, \chi, C, q \models p$ iff $q \in \pi(p)$
- $H, \chi, C, q \models \neg\varphi$ iff $H, \chi, C, q \not\models \varphi$
- $H, \chi, C, q \models \varphi \vee \psi$ iff $H, \chi, C, q \models \varphi$ or $H, \chi, C, q \models \psi$
- $H, \chi, C, q \models \langle\langle D\rangle\rangle \bigcirc \varphi$ iff $\exists s_D \in strat^\chi_C(D) : \forall \lambda \in out(s_D, q) : \lambda[1] \models \varphi$
- $H, \chi, C, q \models \langle\langle D\rangle\rangle \Box \varphi$ iff $\exists s_D \in strat^\chi_C(D) : \forall \lambda \in out(s_D, q) : \forall i \geq 0 : \lambda[i] \models \varphi$
- $H, \chi, C, q \models \langle\langle D\rangle\rangle \varphi \mathcal{U} \psi$ iff $\exists s_D \in strat^\chi_C(D) : \forall \lambda \in out(s_D, q) : \exists i \geq 0 : (\lambda[i] \models \psi \wedge \forall j \in [i] : \lambda[j] \models \varphi)$
- $H, \chi, C, q \models \langle D\rangle\varphi$ iff $H, \chi, D, q \models \varphi$

Example 2. We continue our running jazz trio example. The trio has one objective – to play the concert – which we can express by the following NCHATL formula: $\langle\langle\{k, g, j\}\rangle\rangle \bigcirc$ *play*. This formula is true on the model, but it does not *guarantee* a concert. For that we need to ensure this one: $[\![\{k, g, j\}]\!] \bigcirc$ *play*. We can design a set of individual norms to make this formula true, assuming that all agents comply:

- $\eta_k(q_0, \langle x, y, z, w\rangle) = \eta_g(q_0, \langle x, y, z, w\rangle) = \eta_j(q_0, \langle x, y, z, w\rangle) = \{p_2, p_3, p_4\}$ if $x \neq 0$
- $\eta_k(q_0, \langle 0, 0, 1, 1\rangle) = \{p_1, p_3, p_4\}$;
- $\eta_g(q_0, \langle 0, 1, 0, 1\rangle) = \{p_1, p_2, p_4\}$;
- $\eta_j(q_0, \langle 0, 1, 1, 0\rangle) = \{p_1, p_2, p_3\}$;

The first point ensures that an agent is allowed to play *only* when all the other agents also play an instrument. So if we add a new state to the model, for the case that only some of the agents play, we would now be able to prevent this state from coming about by implementing this normative system. This illustrates the increased expressive power we get from considering coordination. Also note that under full compliance we are sure that the agents play the instruments that they *know* how to play: Keith plays piano, Gary plays the double-bass and Jack plays the drums.

But at this point we also see the limit of our approach, since with an underlying homogeneous structure, we cannot distinguish this from a situation where everyone plays but plays a different instrument than his own. For instance, even if Jack and Gary switch their instruments, our norm still requires Keith to play, even if the concert should probably not go ahead in this case. However, we can capture (parts of) this distinction, since it corresponds to a scenario when the norm is violated by Jack and Gary. A direction for future research that we think will be fruitful is to attempt to use the norm structure more actively, by developing extensions of the logic for describing how bringing about a concert in the wrong way, in violation of the norm, could influence the

[9] Notice that our definition requires also C to follow the norm as a group, when we compute the set of possible extensions of D's strategy.

future development of the system. This could perhaps be modeled by letting the violation itself directly influence the agents' beliefs, desires and intentions. A concert is a concert, one might say, but the audience might not want to come to the next one if Jack the drummer suddenly starts playing the double-bass.

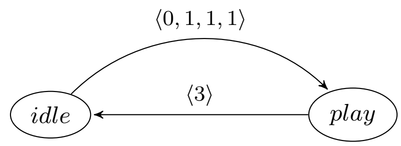

Figure 1: Jazz trio example (reflexive arrows omitted).

4. PARTIAL COMPLIANCE IS TRACTABLE

In this section we will implement norms using anonymous games, and use this to establish the main result, namely that model checking $\mathcal{L}_{\text{NCHATL}}$ against Definition 8 takes polynomial time in the number of agents. We only show this for individual norms η, but due to the representation given in Equation 13, allowing us to represent any collective norm as an individual norm, the result for collective norms follows as an immediate corollary.

Now, for an individual normative system η and every $q \in Q$, we define the corresponding anonymous game $\mathsf{G}_\eta(q) = (\Sigma, \mathbb{A}_q, (\mu_i^q)_{i\in\Sigma})$ where for all $i \in \Sigma$ and every $s \in \mathbb{A}_q^\Sigma$ we take

$$\mu_i^q(s) = \begin{cases} 1 \text{ if } s(i) \in \mathbb{A}_q \setminus \eta_i(q, (\#(y, s_{-i}))_{y\in\mathbb{A}}) \\ 0 \text{ otherwise} \end{cases} \tag{19}$$

The following proposition follows easily, and we omit the proof.

Proposition 2. *Given an* HCGS H *and a state* $q \in Q$, $s \in \mathbb{A}_q^\Sigma$ *is a Nash equilibrium for* $\mathsf{G}_\eta(q)$ *if, and only if,* $s \in \mathsf{L}_\eta(q)$, *where* $\mathsf{L}_\eta(q)$ *is as defined in Equation 10.*

In light of this fact, the following becomes a corollary of the fact that computing Nash Equilibria in anonymous games is tractable.

Theorem 1. *Computing* $P_\eta(q)$ *takes polynomial time (actually, is in* TC_0*).*

Proof. It follows from Proposition 2 and Equation 11 that the set $P_\eta(q)$ is the set of profiles that can be induced by a Nash Equilibrium of $\mathsf{G}_\eta(q)$. Since this set can be computed in polynomial time using the characterization provided in Equation 8, the result follows. (For the claim in the parentheses, we refer to [7] who show that computing the set of profiles for Nash Equilibrium of an anonymous game is in the complexity class TC_0). □

We now demonstrate that partial compliance can be reasoned with in the same efficient manner as full compliance. That is, we must be able to compute the set $P_\eta^C(q)$ defined in Equation 16 in polynomial time for all $q \in Q$, $C \subseteq \Sigma$. Given the correspondence to anonymous games, we can do this if we can answer the following question in polynomial time: *given an anonymous game and a* C*-profile* F*, can* F *be induced by a partial action-tuple for* C *players,* s*, such that* s *can be extended to a Nash Equilibrium?.*

Notice that a C-profile need not have this property even if it can be extended to a profile that is induced by a Nash Equilibrium. If the game is homogeneous this holds trivially, but for anonymous games we cannot arbitrarily permute actions in the Nash Equilibrium to ensure that the desired C-profile is induced by explicit actions for members of C, and not just actions for members of some other coalition of the same size. Notice that this corresponds exactly to our observation about the special care that was needed in defining partial compliance to individual norms.

For all anonymous games $\Gamma = (\Sigma, \mathbb{A}, (\mu_i)_{i\in\Sigma})$ and all $C \subseteq \Sigma$ we say that $F \in P^C(\mathbb{A})$ is a partial Nash profile for Γ whenever

$$\exists s \in \mathsf{NE}(\Gamma) : (\#(y, s \restriction C))_{y\in\mathbb{A}} = F \tag{20}$$

Then the problem we must solve is the problem of deciding whether a given $F \in P^C(\mathbb{A})$ satisfies Equation 20. An answer to this suffices to establish our main result since the notion of a partial Nash profile corresponds to partial compliance in the obvious way, c.f., Proposition 2. In particular, we have the following.

Lemma 1. *For any* HCGS H*, any individual norm* η *and any* $q \in Q$*, we have* $F \in P_\eta^q(C)$ *if, and only if,* $F \in P^q(C)$ *is a partial Nash profile for* $\mathsf{G}_\eta(q)$.

Testing by brute force whether Equation 20 holds for some $F \in P^C(\mathbb{A})$ requires us to consider the potentially exponential set of all Nash Equilibria for the game Γ. However, we can simplify this by making another application of Hall's Lemma, analogous to that used to establish Equation 8. In particular, it is not hard to show that $F \in P^C(\mathbb{A})$ is a partial Nash profile for Γ if, and only if, there exists some $F' \in P^{\Sigma\setminus C}(\mathbb{A})$ such that the following two conditions hold:

$$\begin{aligned} (1):\ & \forall P \subseteq \mathbb{A}: \\ & |\{i \in C \mid \exists p \in P : p \in \mathsf{b}_i(s_{F+F'})\}| \geq \sum_{p\in P} F(p) \\ (2):\ & \forall P \subseteq \mathbb{A}: \\ & |\{i \in \Sigma \setminus C \mid \exists p \in P : p \in \mathsf{b}_i(s_{F+F'})\}| \geq \sum_{p\in P} F'(p) \end{aligned} \tag{21}$$

Using Equation 7 these conditions can be checked in polynomial time in the number of agents. Moreover, remember that running through different C-profiles, functions from $\mathbb{A}$ to $|C|$, is tractable by brute force when the number of actions remains fixed. Hence our main technical result follows.

Theorem 2. *For any* HCGS H*, any individual norm* η *and any* $q \in Q$*, we can compute* $P_\eta^C(q)$ *in polynomial time in the number of agents.*

From this it follows easily that model checking is tractable in the agents. In particular, a simple adaptation of the model checking algorithm presented in [11], where we quantify over $P_\chi^C(q, D)$ in place of $P^C(q)$, witnesses to the truth of the following Theorem.

Theorem 3. *Given an* HCGS H*, a norm* $\chi \in \{\eta, \kappa\}$*, a* $\mathcal{L}_{\text{NCHATL}}$*-formula* φ *a state* $q \in Q$ *and a coalition* $C \subseteq \Sigma$*, checking* $H, \chi, C, q \models \varphi$ *takes polynomial time in the number of agents.*

5. CONCLUSIONS AND FUTURE WORK

We studied normative systems for concurrent game structures, exploiting recent ideas from algorithmic game theory to ensure both tractability and greater expressiveness of the norms we consider. The tractability result was made possible by introducing homogeneity on the level of the system description, but as we showed, heterogeneous properties of interacting agents may be regained in a structured way using norms. Hence we provide a new perspective on where it is appropriate to encode heterogeneous properties of agents. For complexity reasons, introducing heterogeneity as a non-primitive notion, allowing it to arise in a limited way from mechanisms that aim to describe certain types of interactions between autonomous agents, might be more appropriate than embracing heterogeneity as an irreducible modeling assumption.

This perspective on heterogeneity also touches on points that may have conceptual significance above and beyond issues to do with complexity. In many cases homogeneity and/or anonymity arise naturally, for reasons to do with privacy, fairness, or lack of information. We point to [12] for a further discussion on the conceptual significance of the notion of homogeneity and anonymity in logics for multi-agent systems.

An interesting direction for future research is to investigate connections with cooperative game theory and social choice theory. In this paper we used normal form games to characterize and provide an incentive for a very simple *binary* notion of compliance for both individuals and coalitions. Intuitively, the notion of compliance we formalized was built on the idea that a coalition can be said to comply if, and only if, it acts in such a way that it is possible for all the remaining agents to act in accordance with the norm. A very natural next step is to consider *degrees* of compliance, where coalitions are rewarded on a gliding scale depending on how *easy* it becomes for the rest of the agents to fulfill the norm. This, in turn, leads us to cooperative game theory, where the norms themselves can give agents an *incentive* to form coalitions, to ensure higher rewards from better quality of compliance.

On the technical side, we hope to generalize $\mathcal{L}_{\text{NCHATL}}$ to allow quantification over coalitions, following the approach of [4]. For the structures considered there, the decision problems that arise after introducing such quantification are generally not tractable. For homogeneous structures, on the other hand, quantifying over coalitions should be possible, even in the context of broader norms, by relying on a compact representation. Essentially, a partitioning of coalitions into a limited number of equivalence classes would ensure efficient model-checking procedures also when quantification is involved.

We also hope to investigate more closely the link with non-homogeneous structures. In particular, it seems that we may in many cases be able to recognize different degrees of homogeneity within systems that are already formulated using heterogeneous means. Transforming such models to homogeneous models, using norms to model heterogeneous properties, is a challenge that will be considered.

In conclusion, we think broad norms on homogeneous structures provide a logical formalism for modeling multi-agent systems that have many attractive features, making it a suitable template for continued formal work on norms and strategic interaction.

Acknowledgments. Piotr Kaźmierczak's research was supported by the Research Council of Norway project number 194521 (FORMGRID).

6. REFERENCES

[1] T. Ågotnes, W. van Der Hoek, J. A. Rodríguez-Aguilar, C. Sierra, and M. Wooldridge. On the Logic of Normative Systems. In *Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 07)*, pages 1175–1180, 2007.

[2] T. Ågotnes, W. van der Hoek, M. Tennenholtz, and M. Wooldridge. Power in normative systems. In *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, pages 145–152, 2009.

[3] T. Ågotnes, W. van der Hoek, and M. Wooldridge. Normative system games. In *Proc. of 6th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, pages 1–8, 2007.

[4] T. Ågotnes, W. van der Hoek, and M. Wooldridge. Robust normative systems and a logic of norm compliance. *Logic Journal of the IGPL*, 18(1):4–30, 2009.

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

[6] M. Blonski. Characterization of pure strategy equilibria in finite anonymous games. *Journal of Mathematical Economics*, 34:225–233, 2000.

[7] F. Brandt, F. Fischer, and M. Holzer. Symmetries and the complexity of pure nash equilibrium. *Journal of Computer and System Sciences*, 75(3):163–177, 2009.

[8] P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, s1-10(1):26–30, 1935.

[9] W. Hoek, M. Roberts, and M. Wooldridge. Social laws in alternating time: effectiveness, feasibility, and synthesis. *Synthese*, 156(1):1–19, 2007.

[10] W. Jamroga and J. Dix. Do agents make model checking explode (computationally)? In M. Pechoucek, P. Petta, and L. Z. Varga, eds., *Multi-Agent Systems and Applications IV (LNAI Volume 3690)*, 2005.

[11] T. Pedersen, S. Dyrkolbotn, P. Kaźmierczak, and E. Parmann. Concurrent game structures with roles. In *Proc. of the 1st Int. Workshop on Strategic Reasoning*, volume 112 of *EPTCS*, pages 61–69. Open Publishing Association, 2013.

[12] T. Pedersen and S. K. Dyrkolbotn. Agents homogeneous: A procedurally anonymous semantics characterizing the homogeneous fragment of ATL. In *Proc. of PRIMA 2013: Principles and Practice of Multi-Agent Systems*, volume 8291 of *LNCS*, pages 245–259. Springer Berlin Heidelberg, 2013.

[13] Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies. In *Proc. of the 10th National Conf. on Artificial Intelligence*, AAAI'92, pages 276–281. AAAI Press, 1992.

[14] Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: Off-line design. *Artificial Intelligence*, 73:231–252, 1995.