Reasoning about Temporary Coalitions and LTL-definable Ordered Objectives in Infinite Concurrent Multiplayer Games

Extended Abstract

Dimitar P. Guelev
Institute of Mathematics and Informatics
Bulgarian Academy of Sciences, Sofia, Bulgaria
gelevdp@math.bas.bg

We propose enhancing the use of propositions for denoting decisions and strategies as established in temporal languages such as CTL*, if interpreted on concurrent game models. The enhancement enables specifying varying coalition structure. In quantified CTL* this technique also enables quantifying over coalition structure, and we use it to quantify over an extended form of strategy profiles which capture temporary coalitions. We also extend CTL* by a temporal form of a binary preference operator that can be traced back to the work of Von Wright. The resulting extension of quantified CTL* can be used to spell out conditions on the rationality of behaviour in concurrent multiplayer games such as what appear in solution concepts, with players having multiple individual objectives and preferences on them, and with the possibility to form temporary coalitions taken in account. We propose complete axiomatisations for the extension of CTL* by the temporal preference operator. The decidability of the logic is not affected by that extension.

keywords: strategic ability, temporary coalitions, rational synthesis, preference, concurrent multiplayer games.

Introduction

In Alternating-time Temporal Logic (ATL, [AHK97, AHK02]), $\langle\langle\Gamma\rangle\rangle\varphi$ means that a joint strategy for the players from Γ exists which, if implemented, would cause the (rest of the) play to satisfy temporal condition φ. Willfully carrying out a strategy to enforce a temporal property which refers to the entire future requires Γ to act as a *permanent* coalition with that property as the *shared* objective and the joint strategy as the member players' lifetime agenda. This means that ATL's game-theoretic operator, if regarded as a pattern for statements, is an immediate match for statements on the strategic ability of permanent coalitions wrt shared long term objectives. Strategy Logic (SL, [MMV10]) is built around a construct which expresses that implementing a given global strategy profile would enforce a temporal property that similarly refers to the entire future of the play. This is not to say that permanent coalitions are the limit of expressive power for ATL and SL in some ultimate sense. The skillful use of logical connectives, superposition and quantification can certainly be relied on to achieve a variety of other statements about strategic ability.

In this paper we consider a setting where Concurrent Game Models (CGMs) and the temporal sublanguage take their established roles much like in ATL, but aiming to formulate statements about strategic ability and rational behaviour on behalf of players who are prepared to enter *temporary* coalitions while in pursuit of their *individual* long term objectives, without having to combine those objectives into long

Submitted to:
SR 2020

term common ones such as what the plain use of the constructs offered in ATL and SL appears to be best suited for. To this end we fall back on propositionally quantified CTL* (henceforth QCTL*) as the temporal logic language while keeping CGMs as the semantics. We extend the established way of using propositional variables for naming strategies and expressing quantification over strategies in QCTL* to provide notation for speculating about the varying coalition structure and the 'short term' agendas of temporary coalitions combined.

To enable judgements which link the viability of temporary coalitions and their agendas to the players' individual long term objectives, we assume that the preferences of each player define a partial order on plays, and extend CTL* by a temporal binary preference operator: $\varphi <_i \psi$ means that i prefers plays which satisfy φ to plays satisfying ψ. The operator can be viewed a temporal extension of the abstract binary preference operator which can be traced back to the work of Von Wright [vW63]. Interestingly, the possibility to specify step-by-step progress towards objectives, and the evolution of objectives in the course of their partial fulfillment or forfeiture admit a rather straightforward expression using this operator as an addition to CTL* in combinations with LTL's guarded normal form and the separated normal form of LTL formulas. The interaction is encoded in some axioms for $<_i$ we propose as part of the second of two complete axiomatisations for it. The first axiomatisation enables the elimination of $<_i$, in case the preferences of every player are suitably expressed by some finite list of comparisons between LTL-definable sets of plays. The second axiomatisation makes no such assumption; indeed its completeness proof, which is not included in this extended abstract, entails that a formula in CTL* with $<_i$ is satisfiable only if so wrt the preferences of the players being expressible this way. Since the possibility to work towards objectives in multiple steps is specific to those objectives being temporal, we believe that this temporal species of a binary preference operator is of some interest on its own too.

Meta-theoretical results on *temporalization* were first obtained in [FG92]. More general ways for combining logics were later proposed by Gabbay and others. Our extension of ATL* with $<$ is broadly compatible with that framework. In the sense of [GS98], it can be viewed as a kind of *fusion* of ATL* and Preference Logic as in [vBvOR05], except that the state space we interpret Preference Logic in CGMs on is not the same as that for ATL* but is the set of the infinite plays in the considered CGM.

Finding strategy profiles which meet the requirements of solution concepts in general is referred to as *rational synthesis*. For *temporal* objectives, rational synthesis was proposed in [FKL10] and further investigated in [KPV16] with the focus on rational behaviour in the sense of Nash equilibrium wrt the objectives of some finite number of players, with the additional condition that the objective of a distinguished player, the *system*, is achieved. One key observation about solution concepts with coalitions is that players can use their knowledge of the actions to be taken by their allies when deciding whether to enter a coalition and embrace its agenda. Non-members are just assumed to be acting rationally wrt their respective preferences and the adopted solution concept. With temporary coalitions, this includes perpetually revising the coalition structure. In this paper, we formulate temporary coalition variants of the concepts of the *game core* and *dominant strategies* to demonstrate the use of the notation we are proposing. We do not claim these variants to be the unique options for extending these solution concepts. Indeed, solution concepts being so varied is what justifies studying logical languages which are both versatile and amenable to algorithmic methods. Results on rational synthesis after [FKL10] in the literature are now commonly derived just from the possibility to express the solution concepts in a decidable subset of SL too. Of course, optimality requires looking at algorithms which are taylored for specific concepts; cf. e.g., [BBMU15, BBMU12] on Nash equilibrium.

1 Preliminaries

Concurrent Game Models We consider CGMs $M = \langle W, w_I, \langle Act_i : i \in \Sigma\rangle, o, V\rangle$ for some given sets of *players* $\Sigma = \{1,\dots,N\}$ and atomic propositions AP. W is the set of *states*, $w_I \in W$ being the *initial* state, Act_i is the set of *actions* of player $i \in \Sigma$, $Act_\Gamma \hat{=} \prod\limits_{i\in\Gamma} Act_i$, $o : W \times Act_\Sigma \to W$ is the *outcome* function, and $V \subseteq W \times AP$ is the *valuation* relation. Given a $\mathbf{w} \in W^+$, we put

$$R^{\mathrm{inf}}_M(\mathbf{w}) \hat{=} \{\mathbf{v} \in W^\omega : \mathbf{v}^0 \dots \mathbf{v}^{|\mathbf{w}|-1} = \mathbf{w}, (\forall k < \omega)(\exists a \in Act_\Sigma)(\mathbf{v}^{k+1} = o(\mathbf{v}^k, a))\}.$$

for the set of all the infinite plays in M which are continuations of $\mathbf{w}$. The set $R^{\mathrm{fin}}_M(\mathbf{w}) \subseteq W^+$ of the *finite* continuations of $\mathbf{w}$ is defined similarly.

Given a $p \in AP$ and an $X \subseteq W$, M^X_p denotes the CGM $\langle W, w_I, \langle Act_i : i \in \Sigma\rangle, o, V^X_p\rangle$ where $V^X_p(w,p) \leftrightarrow w \in X$ and $V^X_q(w,q) \leftrightarrow V(w,q)$ for $q \in AP \setminus \{p\}$.

In the sequel we tacitly assume an arbitrary fixed finite CGM M with its components named as above.

Strategies A (total deterministic) *strategy* for player $i \in \Sigma$ is a function of type $W^+ \to Act_i$. A *strategy profile for* $\Gamma \subseteq \Sigma$ is a tuple of strategies $\mathbf{s} = \langle \mathbf{s}_i : i \in \Gamma\rangle$, one for every member i of Γ. We denote the set of all strategy profiles for Γ by S_Γ. Strategies as above apply to permanent coalitions.

Alternating-time Temporal Logic (ATL**, [AHK97, AHK02])** This is now a group of logics with the common defining feature being the game-theoretic modalities $\langle\langle\Gamma\rangle\rangle$ where $\Gamma \subseteq \Sigma$. ATL* syntax includes *state* formulas φ and *path* formulas ψ. Path formulas are PLTL (LTL with past) formulas with ATL* state formulas allowed to appear along with the atomic propositions. We use ATL* with linear past. A recent study on ATLs with linear past, with comprehensive references on the overall topic can be found in [BMS20]. The syntax of ATL* with past can be defined by the BNFs:

$$\varphi ::= \bot \mid p \mid \varphi \Rightarrow \varphi \mid \langle\langle\Gamma\rangle\rangle\psi \qquad \psi ::= \varphi \mid \psi \Rightarrow \psi \mid \bigcirc\psi \mid (\psi \mathsf{U} \psi) \mid \ominus\psi \mid (\psi \mathsf{S} \psi)$$

As it becomes clear below, the past operators facilitate reference to the players' objectives as formulated wrt the beginning of a play. ATL* semantics on CGMs comes in several variants; we choose the one in which a CGM M and a finite play $\mathbf{w} \in R^{\mathrm{fin}}_M(w_I)$ need to be specified to define $\models$ on state formulas, and an infinite play, and a position $k < \omega$ in it are needed for path formulas. Given a $\Gamma \subseteq \Sigma$, a $\mathbf{w} \in R^{\mathrm{fin}}_M(w_I)$ and an $\mathbf{s} \in S_\Gamma$, the set of the infinite continuations of $\mathbf{w}$ in which the players from Γ act according to $\mathbf{s}$ is:

$$\mathsf{out}(\mathbf{w},\mathbf{s}) \hat{=} \{\mathbf{v} \in R^{\mathrm{inf}}_M(\mathbf{w}) : (\forall k \geq |\mathbf{w}|)(\exists \mathbf{b} \in Act_{\Sigma\setminus\mathrm{dom}\mathbf{s}})(\mathbf{v}^k = o(\mathbf{v}^{k-1}, \mathbf{s}(\mathbf{v}^0 \dots \mathbf{v}^{k-1}) \cup \mathbf{b}))\}.$$

The definining clauses for $\models$ are as follows:

$M,\mathbf{w} \models p$	iff	$V(\mathbf{w}^{\lvert\mathbf{w}\rvert-1}, p)$ for $p \in AP$
$M,\mathbf{w} \models \langle\langle\Gamma\rangle\rangle\varphi$	iff	there exists an $\mathbf{s} \in S_\Gamma$ s. t. $M,\mathbf{v},\lvert\mathbf{w}\rvert-1 \models \varphi$ for all $\mathbf{v} \in \mathsf{out}(\mathbf{w},\mathbf{s})$
$M,\mathbf{w},k \models \bigcirc\varphi$	iff	$M,\mathbf{w},k+1 \models \varphi$
$M,\mathbf{w},k \models \ominus\varphi$	iff	$k > 0$ and $M,\mathbf{w},k-1 \models \varphi$
$M,\mathbf{w},k \models (\varphi \mathsf{U} \psi)$	iff	for some $n < \omega$, $M,\mathbf{w},k+n \models \psi$ and $M,\mathbf{w},k+m \models \varphi$ for all $m < n$.
$M,\mathbf{w},k \models (\varphi \mathsf{S} \psi)$	iff	for some $n \leq k$, $M,\mathbf{w},k-n \models \psi$ and $M,\mathbf{w},k-m \models \varphi$ for all $m < n$.

The clauses about $\bot$, $\top$ and $\Rightarrow$ are as usual; $\vee$, $\wedge$ and $\Leftrightarrow$ and the derived PLTL operators $\Diamond$, $\Box$, $\boxminus$, $\diamondminus$, etc. are defined as usual too. The beginning of time is indicated by $\mathsf{I} \hat{=} \neg\ominus\top$. The dual $\neg\langle\langle\Gamma\rangle\rangle\neg$ of $\langle\langle\Gamma\rangle\rangle$ is written $[[\Gamma]]$. ATL's $\langle\langle\emptyset\rangle\rangle$ and $[[\emptyset]]$ are equivalent to CTL's $\forall$ and $\exists$, respectively, the genetic link between

ATL and CTL. In the sequel, for the sake of brevity, we write $\forall$ and $\exists$ for $\langle\langle\emptyset\rangle\rangle$ and $[[\emptyset]]$, respectively. For $\Gamma \subseteq \Sigma$, $-\Gamma \hat{=} \Sigma \setminus \Gamma$.

For future state formulas φ, we put $[[\varphi]]_M \hat{=} \{\mathbf{v}^{|\mathbf{v}|-1} : M, \mathbf{v} \models \varphi\}$. For future state φ, $M, \mathbf{v} \models \varphi$ does not impose conditions on $\mathbf{v}^0 \dots \mathbf{v}^{|\mathbf{v}|-2}$. Therefore $M, \mathbf{v} \models \varphi$ is equivalent to $\mathbf{v}^{|\mathbf{v}|-1} \in [[\varphi]]_M$ for such formulas.

We write $\mathrm{Var}(\varphi)$ and $\mathrm{Subf}(\varphi)$ for the set of the atomic propositions which occur in formula φ and the set of φ's subformulas, including φ itself, respectively.

Unwinding CGMs Given a CGM M as above, the *unwinding* $M^T \hat{=} \langle W^T, w_I^T, \langle Act_i : i \in \Sigma\rangle, o^T, V^T\rangle$ of M is defined as follows:

$$\begin{array}{lcllcl} W^T & \hat{=} & W(Act_\Sigma\, W)^* & o^T(w^0\mathbf{a}^1 \cdots \mathbf{a}^n w^n, \mathbf{b}) & \hat{=} & w^0\mathbf{a}^1 \cdots \mathbf{a}^n w^n \mathbf{b} o(w^n, \mathbf{b}) \\ w_I^T & \hat{=} & w_I & V^T(w^0\mathbf{a}^1 \cdots \mathbf{a}^n w^n, p) & \hat{=} & V(w^n, p) \end{array}$$

M^T and M are bisimilar and $R_M^{\inf}(w_I)$ and $R_{M^T}^{\inf}(w_I^T)$ are isomorphic. Importantly, o^T is invertible.

QCTL* The use of $\exists$ and $\forall$ as abbreviations for ATL's $[[\emptyset]]$ and $\langle\langle\emptyset\rangle\rangle$, respectively, gives rise to a sublanguage of ATL* (with past) which coincides with the language of CTL*, for the same vocabulary. The interpretation of $\exists$ and $\forall$ which is entailed by their use as abbreviations in CGMs M is consistent with the same mode of referencing the past in the CTL* definitions of their *structure* semantics wrt models $\langle W, w_I, R, V\rangle$ where W, w_I, and V are as in the CGM M, and the transition relation R is defined using the CGM's outcome function o by the equivalence $R(w, v) \leftrightarrow (\exists \mathbf{a} \in Act_\Sigma)(v = o(w, \mathbf{a}))$.

Along with its structure semantics, QCTL* admits also a *tree* semantics, which can be spelled out using the same defining clauses, except that M^T, the unwinding of the given model M, appears on the LHS of $\models$. The tree semantics differs from the structure one only wrt propositional quantification. In the structure semantics,

$$M, \mathbf{w} \models \exists p\varphi \quad \text{iff} \quad \text{there exists an } X \subseteq W \text{ s. t. } M_p^X \models \varphi.$$

With the values of quantified variables being varied on the states of the original $\langle W, w_I, R, V\rangle$, the recurrence of states along paths entails corresponding repeats of the variables' values. Unwindings display no such dependency as repeated occurrences of the original model's states along paths are replaced by distinct states, and the condition $X \subseteq W$ becomes replaced by $X \subseteq W^T$. In this work we refer mostly to unwindings, which can be regarded as using the tree semantics of arbitrary CGMs too. Model-checking and validity in QCTL* are decidable on the class of tree-based Kripke models [Fre01, Fre06] such as what underly unwindings M^T.

2 Strategy Profiles with Temporary Coalitioning and a Vocabulary for Temporary Coalitions

In this section we introduce our propositional vocabulary for making statements about strategic ability with the co-existence of temporary coalitions taken in account. In doing so, we build on a technique which enables the use of propositional variables for the naming of sets of transitions in modal languages. In QCTL*, along with the naming of strategies, this technique enables the expression of quantification over strategies on CGMs. In our approach, dedicated collections of propositional variables specify both sets of decisions in the established way and the coalition structure in place upon every transition. We upgrade the underlying semantic notion of strategy profile to incorporate coalition structure first. Let $\mathsf{part}(\Gamma)$ stand for the set of the exhaustive partitions of Γ into disjoint nonempty subsets.

Definition 1 Assuming a CGM M as usual, and a *strategy profile with temporary coalitions* (SPTC) is a function of type $W^+ \to Act_\Sigma \times \mathsf{part}(\Sigma)$.

Informally, given a finite play, a SPTC specifies both a decision on how to continue the play and a partitioning of Σ into coalitions, the competing collective authors of that decision. Now let us explain how our vocabulary for SPTC builds on the conventions for naming decisions in a propositional temporal language.

Naming decisions and strategies by propositions in QCTL* works as follows. In an arbitrary CGM M with its components named as above, propositions p define the sets of states $[\![p]\!] \hat{=} \{w \in W : V(w,p)\}$. In the unwinding M^T of a CGM M, the invertibility of o^T entails that any subset S of $W^T \times Act_\Sigma$ can be recovered from the corresponding $\{o(w,\mathbf{a}) : \langle w,\mathbf{a}\rangle \in S\}$, which is a set of states, and therefore can be denoted by a proposition. Finite plays can be determined from their last states in M^T, which renders memorylessness of strategies vacuous. Hence any strategy can be determined unambiguously from the set of the target states of the transitions it generates. A proposition which is true in states that are reachable by transitions that are consistent with the strategy and false elsewhere can serve to name the strategy. Hence $\exists s(\delta_\Gamma(s) \wedge \forall\bigcirc(s \Rightarrow \varphi))$, a QCTL* formula, in which $\delta_\Gamma(s)$ presumedly restricts s to range over those $S \subseteq W^T \times Act_\Sigma$ which fit the description of strategies for Γ, expresses that Γ can enforce φ by implementing s for one step.

This observation combines beneficially with the facts that QCTL* is decidable on tree-based Kripke models (aka execution trees) which are the unwindings of finite Kripke models [**?**], and, unsurprisingly, finite CGMs reduce to finite Kripke models, if the transition relation is defined by putting $R(w,v) \leftrightarrow (\exists \mathbf{a} \in Act_\Sigma)(v = o(w,\mathbf{a}))$, as mentioned in the Preliminaries section.

Specifying evolving coalition structure We specify evolving coalition structure by enhancing the above use of propositions. The propositional variables which we introduce serve to both name decisions and specify the coalition structure in place upon carrying out the decisions. Specifying a SPTC s takes a collection $\mathbf{s} \hat{=} \langle \mathbf{s}_\Gamma : \Gamma \subseteq \Sigma\rangle$ of variables. The extension $(M^T)^{\langle X_\Gamma(s):\Gamma\subseteq\Sigma\rangle}_{\langle \mathbf{s}_\Gamma:\Gamma\subseteq\Sigma\rangle}$ of M^T in which the variables from $\mathbf{s}$ specify s is defined by putting

$$X_\Gamma(s) \hat{=} \{o^T(\mathbf{w}^{|\mathbf{w}|-1}, \mathbf{a}|_\Gamma \cup \mathbf{b}) : \mathbf{w} \in R^{\mathrm{fin}}_{M^T}(w_I^T), \langle \mathbf{a}, C\rangle = s(\mathbf{w}), \Gamma \in C, \mathbf{b} \in Act_{-\Gamma}\}. \quad (1)$$

This means that $(M^T)^{\langle X_\Gamma(s):\Gamma\subseteq\Sigma\rangle}_{\langle \mathbf{s}_\Gamma:\Gamma\subseteq\Sigma\rangle}, \mathbf{v} \models \mathbf{s}_\Gamma$ indicates that $\mathbf{v}^{|\mathbf{v}|-1}$ can be reached from $\mathbf{v}^{|\mathbf{v}|-2}$ by Γ's part of the c-decision for $\mathbf{v}^0 \cdots \mathbf{v}^{|\mathbf{v}|-2}$, with the express condition that Γ's part of this decision is a decision of Γ *as a coalition*, and not a coincidental collection of unrelated decisions by the players from Γ.

By this convention we can write, e.g., $\exists z(\delta_\Gamma(z) \wedge \forall\bigcirc(z \wedge \mathbf{s}_{-\Gamma} \Rightarrow \varphi))$ to express that the members of Γ can enforce φ provided that the rest of the players make a $\mathbf{s}_{-\Gamma}$-move *as a coalition*. This condition on $-\Gamma$'s concurrent move may be of some consequence, if, e.g., the SPTC denoted by $\mathbf{s}$ meets the requirements a solution concept may be imposing, including requirements on the viability of coalitions. Then the $-\Gamma$ decisions which lead to $\mathbf{s}_{-\Gamma}$-states would range over those single step coalition agendas which, according to the adopted solution concept, appear to be sufficiently attractive to unite $-\Gamma$.

In the rest of the paper $\mathbf{X}(s)$ abbreviates $\langle X_\Gamma(s) : \Gamma \subseteq \Sigma\rangle$; hence $(M^T)^{\mathbf{X}(s)}_{\mathbf{s}}$ stands for $(M^T)^{\langle X_\Gamma(s):\Gamma\subseteq\Sigma\rangle}_{\langle \mathbf{s}_\Gamma:\Gamma\subseteq\Sigma\rangle}$. As it becomes clear below, expressing comparisons between alternative SPTC may take CTL* formulas which refer to more than one such system of variables of the form $\langle \mathbf{s}_\Gamma : \Gamma \subseteq \Sigma\rangle$.

The apparent lack of semantics for coalition structure Note that, unlike decisions and strategies, coalition structure cannot be observed by examining the (global) transitions themselves in CGMs. Therefore it is crucial to notice that the intended meaning of these dedicated variables is achieved only in extensions of M^T by dedicated valuations for these variables, or by their *bound occurrences*, with the quantification also indicating that the coalitioning in question is *hypothetical*. This convention leads to averting the perceived necessity to upgrade CGMs for registering coalition structure by anything more than the valuation for the dedicated propositional variables introduced above.

Constraining the vocabulary to express well-formed coalition structure For a system $\mathbf{s} \hat{=} \langle \mathbf{s}_\Gamma : \Gamma \subseteq \Sigma \rangle$ of variables to really specify how Σ is partitioned into disjoint coalitions, $\mathbf{s}$ must assign every player to a unique coalition, with the decisions of unallied players i being denoted by the respective $\mathbf{s}_{\{i\}}$. Along with expressing that the reference state is the target of a transition in which all the decisions were $\mathbf{s}$-decisions, the formulas $\widetilde{\mathbf{s}}_\Gamma$ below express that each of the players from Γ belongs to a unique coalition within Γ and, furthermore, no coalitions include players from both inside and outside Γ. The following equivalence defines $\widetilde{\mathbf{s}}_\Gamma$, $\Gamma \subseteq \Sigma$, recursively:

$$\widetilde{\mathbf{s}}_\Gamma \Leftrightarrow \bigwedge_{\Gamma' \subseteq \Sigma, \Gamma' \cap \Gamma \neq \emptyset, \Gamma' \setminus \Gamma \neq \emptyset} \neg \mathbf{s}_{\Gamma'} \wedge \left(\mathbf{s}_\Gamma \wedge \bigwedge_{\emptyset \subsetneq \Delta \subsetneq \Gamma} \neg \mathbf{s}_\Delta \vee \bigvee_{\emptyset \subsetneq \Delta \subsetneq \Gamma} (\widetilde{\mathbf{s}}_\Delta \wedge \widetilde{\mathbf{s}}_{\Gamma \setminus \Delta}) \right). \tag{2}$$

Since $-\Sigma = \emptyset$, $\widetilde{\mathbf{s}}_\Sigma$ only states that $\mathbf{s}$ is assigning every player to a unique coalition.[1] Next we spell out some more restrictions on $\mathbf{s}$ which further entail that a state which satisfies $\widetilde{\mathbf{s}}_\Sigma$ ought to be the target state of a unique global decision, with this coalition structure in place. Namely, we restrict $\mathbf{s}_\Gamma$ to evaluate to the set of the target states of transitions with some specific $\mathbf{a} \in Act_\Gamma$ as Γ's projection of the involved global decision, in case Γ appears in the current coalition structure. This is achieved by including in AP names for the sets $\{o^T(w, \mathbf{a}) : \mathbf{a}_i = a\}$ all $a \in Act_i$, $i \in \Sigma$. To enable this, assume that Act_i, $i \in \Sigma$, are pairwise disjoint subsets of AP, and V^T is defined on them by the equivalences:

$$V^T(w_I^T, a) \leftrightarrow \bot, \qquad V^T(w^0 \mathbf{a}^1 \cdots \mathbf{a}^{n+1} w^{n+1}, a) \leftrightarrow a = \mathbf{a}_i^{n+1} \text{ for all } a \in Act_i, i \in \Sigma.$$

In arbitrary concurrent game structures, the target sets of distinct actions by the same player need not be disjoint like in unwindings M^T and recovering previous actions from their target states like above may be impossible. However, given an arbitrary finite CGM M, an expansion $\bar{M}$ of M can be defined, which is still finite, has an outcome function which is invertible wrt decisions, and has unwinding $\bar{M}^T$ which is isomorphic to M^T, the unwinding of the original M. $\bar{M} \hat{=} \langle \bar{W}, \bar{w}_I, \langle Act_i : i \in \Sigma \rangle, \bar{o}, \bar{V} \rangle$ can be viewed as extending M to store (only) the latest global decisions in their target states, as opposed to M^T, which stores whole histories:

$$\begin{array}{lll lll}
\bar{W} & \hat{=} & W \times (Act_\Sigma \cup \{*\}) & \bar{V}(\langle w, \mathbf{a} \rangle, p) & \hat{=} & V(w,p) \text{ for } p \in AP \setminus \bigcup_{i \in \Sigma} Act_i \\
\bar{w}_I & \hat{=} & \langle w_I, * \rangle & \bar{V}(\langle w, \mathbf{a} \rangle, a) & \hat{=} & a = \mathbf{a}_i \text{ for } a \in Act_i, i \in \Sigma \\
\bar{o}(\langle w, \mathbf{b} \rangle, \mathbf{a}) & \hat{=} & \langle o(w, \mathbf{a}), \mathbf{a} \rangle & & &
\end{array}$$

That is why M^T, with the valuations of the action-naming atomic propositions as defined, for finite M, is still the unwinding of a finite CGM. We highlight this fact because it is relevant for the decidability of satisfaction of formulas in it.

Given an $\mathbf{a} \in Act_\Gamma$ for some $\Gamma \subseteq \Sigma$, let $\hat{\mathbf{a}} \hat{=} \bigwedge_{i \in \Gamma} \mathbf{a}_i$. Let

$$\tilde{\delta}(\mathbf{s}) \hat{=} \forall \Box \exists \bigcirc \widetilde{\mathbf{s}}_\Sigma \wedge \bigwedge_{\Gamma \subseteq \Sigma} \forall \Box \left(\forall \bigcirc \neg \mathbf{s}_\Gamma \vee \bigvee_{\mathbf{a} \in Act_\Gamma} \forall \bigcirc (\mathbf{s}_\Gamma \Leftrightarrow \hat{\mathbf{a}}) \right)$$

[1] We prefer the above recursive definition of $\widetilde{\mathbf{s}}_\Gamma$ to a 'flat' encoding of the conditions on $\mathbf{s}$ as the recursive definition gives easy access to a useful corollary of the property, namely that either the whole of Γ is a coalition, or some proper subset of its, e.g. some subset of a Δ from the RHS disjunct in the parentheses, is a coalition. Note that these conditions are not perfectly 'modular', as $\widetilde{\mathbf{s}}_\Gamma$ rules out coalitions which only overlap with Γ and can extend among $\Sigma \setminus \Gamma$ as well.

For any collection $\mathbf{Y} \hat{=} \langle Y_\Gamma : \Gamma \subseteq \Sigma \rangle$ of subsets of $W^T \setminus \{w_I^T\}$, it can be shown that $(M^T)^{\mathbf{Y}}_{\mathbf{s}}, w_I^T \models \tilde{\delta}(\mathbf{s})$ iff there exists a SPTC s for M^T such that $\mathbf{Y} = \mathbf{X}(s)$ where $\mathbf{X}(s)$ is as in (1).

Some basic expressions in SPTC **vocabularies** Given a SPTC s, $(M^T)^{\mathbf{X}(s)}_{\mathbf{s}}, \mathbf{w}, 0 \models \Box \bigcirc \widetilde{\mathbf{s}}_\Sigma$ means that play $\mathbf{w}$ in M^T is consistent with s. SPTC s enforces φ, if

$$(M^T)^{\mathbf{X}(s)}_{\mathbf{s}}, w_I^T \models \forall \left(\Box \bigcirc \widetilde{\mathbf{s}}_\Sigma \Rightarrow \varphi \right). \tag{3}$$

SPTC s enables $i \in \Sigma$ to achieve φ, if

$$(M^T)^{\mathbf{X}(s)}_{\mathbf{s}}, w_I^T \models \forall \left(\Box \bigcirc \bigvee_{i \in \Gamma \subseteq \Sigma} \mathbf{s}_\Gamma \Rightarrow \varphi \right). \tag{4}$$

In words, if every step of a play is consistent with the agenda of the s-coalition where i belongs, then that play is bound to satisfy φ, regardless of the agendas of the co-existing temporary coalitions. In plays which satisfy $\Box \bigcirc \widetilde{\mathbf{s}}_\Sigma$, i belongs to exactly one coalition at any time. Importantly, it is not guaranteed that the hypothetical allies of i would be really interested in participating in the designated coalitions.

Strategy contexts with temporary coalitions Given our conventions about the use of systems $\mathbf{s} = \langle \mathbf{s}_\Gamma : \Gamma \subseteq \Sigma \rangle$ of propositional variables for naming SPTC, we can define strategic ability of (permanent) coalition Γ in the *context* of a $\Delta \subseteq \Sigma \setminus \Gamma$, implementing (their part of) a SPTC denoted by $\mathbf{s}$ by putting

$$\langle\langle \Gamma \rangle\rangle^{\mathbf{s}}_{\Delta} \varphi \hat{=} \langle\langle \Gamma \rangle\rangle (\Box \bigcirc \widetilde{\mathbf{s}}_\Delta \Rightarrow \varphi). \tag{5}$$

As expected, we write the dual as $[[\Gamma]]^{\mathbf{s}} \varphi \hat{=} \neg \langle\langle \Gamma \rangle\rangle^{\mathbf{s}} \neg \varphi$. This can be regarded as a temporary coalition upgrade of the semantics of *strategy contexts* from [WvdHW07, BLLM09, WHY11].

Related Work on the Specification of Strategies by Propositional Variables The technique we use for the naming of decisions and strategies is a mixture of folklore an authored work. For instance, the transformation of CGMs involved was key to establishing the equivalence between the [AHK97] alternating transition systems based semantics of ATLs and the [AHK02] CGM-based one in [GJ04]. In automata theory, a similar transformation is known as the passage from automata with labels on the transitions to automata with (those same) labels appearing on the respective transitions' target states; the latter are called *state-based* automata. A different technique of using propositions and quantification for encoding strategies and expressing their existence was proposed in [Pin07]. In [Gue13], we used a similar naming technique to show how validity in (a subset of epistemic) ATL reduces to validity in (epistemic) CTL. In [GD12] we applied the technique to epistemic ATL with strategy contexts, to essentially identify the contribution of individual coalition members and the contribution of the players from the strategy context towards the considered objectives. Model checking complete information ATL_{sc} was shown to be reducible to validity in QCTL^* by introducing atomic propositions to name states and decisions and encoding the CGM's outcome function as a formula in terms of these propositions in [LLM12].

3 Preference in Concurrent Game Models

We assume the preference relations $<_i$ of individual players i to be strict partial orders on infinite plays, with unrelated plays being of the same value to the respective players. In much of the literature preference is a pre-order; a comprehensive discussion on modeling preference can be found in [Han04].

To facilitate algorithmic methods, we require the relations

$$\mathbf{v} \sim_i \mathbf{w} \hat{=} (\forall \mathbf{u} \in R_M^{\mathrm{inf}}(w_I))((\mathbf{u} <_i \mathbf{v} \leftrightarrow \mathbf{u} <_i \mathbf{w}) \wedge (\mathbf{v} <_i \mathbf{u} \leftrightarrow \mathbf{v} <_i \mathbf{w}))$$

of indiscernibility wrt $<_i$, $i \in \Sigma$, to partition $R_M^{\mathrm{inf}}(w_I)$ into finitely many LTL-definable classes.

We denote the sets of the formulas which define these indiscernibility classes by $\Theta_{I,i}$, $i \in \Sigma$, where I stands for *I*nitial like in w_I. We require distinct $\theta', \theta'' \in \Theta_{I,i}$ to define disjoint sets of plays: $M, w_I \models \forall\neg(\theta' \wedge \theta'')$. To incorporate the preference relations $<_i$, we consider *extended* CGMs M of the form

$$\langle W, w_I, \langle Act_i : i \in \Sigma \rangle, o, \langle <_i : i \in \Sigma \rangle, V \rangle \tag{6}$$

Given $<_i$, we assume that the sets $\Theta_{I,i}$, become available automatically. Then tuples of the form

$$\langle W, w_I, \langle Act_i : i \in \Sigma \rangle, o, \langle \langle \Theta_{I,i}, <_i \rangle : i \in \Sigma \rangle, V \rangle \tag{7}$$

with the partial orders $<_i$ now defined on the classes of plays specified by $\theta \in \Theta_{I,i}$ can serve as extended CGMs too. In the sequel we tacitly assume an arbitrary fixed extended CGM with the preferences of the players in it represented either as in (6) or as in (7).

The restriction to LTL-definable properties is meant to provide a match with the expressive power of the TL language. Requiring the classes to be just regular ω-languages is no less reasonable. This more general setting has been investigated in [BBMU12]. That work also presents a compelling variety of ways to combine primitive objectives into compound ones. Temporary coalitions are natural to occur given the ordering of the objectives in the example below.

Example 2 1 is a sworn enemy of 2 but feels for its own life even more strongly. Then, if p_i stands for i *perishes*, 1's preferences can be expressed as follows:

$$(\Diamond p_1 \wedge \Box\neg p_2) <_1 (\Diamond p_1 \wedge \Diamond p_2) <_1 (\Box\neg p_1 \wedge \Box\neg p_2) <_1 (\Box\neg p_1 \wedge \Diamond p_2). \tag{8}$$

With 3 sworn enemies we have:

$$\begin{pmatrix} \Diamond p_1 \wedge \\ \Box\neg p_2 \wedge \\ \Box\neg p_3 \end{pmatrix} <_1 \begin{array}{c} \begin{pmatrix} \Diamond p_1 \wedge \\ \Diamond p_2 \wedge \\ \Box\neg p_3 \end{pmatrix} <_1 \begin{pmatrix} \Box\neg p_1 \wedge \\ \Box\neg p_2 \wedge \\ \Box\neg p_3 \end{pmatrix} <_1 \begin{pmatrix} \Box\neg p_1 \wedge \\ \Diamond p_2 \wedge \\ \Box\neg p_3 \end{pmatrix} \\ \bar{\vee} \\ \begin{pmatrix} \Diamond p_1 \wedge \\ \Box\neg p_2 \wedge \\ \Diamond p_3 \end{pmatrix} <_1 \begin{pmatrix} \Diamond p_1 \wedge \\ \Diamond p_2 \wedge \\ \Diamond p_3 \end{pmatrix} <_1 \begin{pmatrix} \Box\neg p_1 \wedge \\ \Box\neg p_2 \wedge \\ \Diamond p_3 \end{pmatrix} \end{array} <_1 \begin{pmatrix} \Box\neg p_1 \wedge \\ \Diamond p_2 \wedge \\ \Diamond p_3 \end{pmatrix} \tag{9}$$

Now, e.g., 1 and 2 may conspire to eliminate 3, but their relations are bound to sour afterwards.

For non-singleton coalitions Γ, we write $\Theta_{I,\Gamma} \hat{=} \left\{ \bigwedge\limits_{i \in \Gamma} \theta_i : \theta_i \in \Theta_{I,i}, i \in \Gamma \right\}$. We refer to the formulas from $\Theta_{I,\Gamma}$ as Γ*'s objectives*.

In the temporal language, we introduce the binary operators $<_i$ and $\not\not<_i$, $i \in \Sigma$, as follows:

$$\begin{array}{lll} M, \mathbf{v} \models \varphi_1 <_i \varphi_2 & \text{iff} & \mathbf{w}_1 <_i \mathbf{w}_2 \text{ for all } \mathbf{w}_k \in R_M^{\mathrm{inf}}(\mathbf{v}) \text{ s. t. } M, \mathbf{w}_k, |\mathbf{v}| \models \varphi_k,\ k = 1, 2. \\ M, \mathbf{v} \models \varphi_1 \not\not<_i \varphi_2 & \text{iff} & \mathbf{w}_1 \not<_i \mathbf{w}_2 \text{ for all } \mathbf{w}_k \in R_M^{\mathrm{inf}}(\mathbf{v}) \text{ s. t. } M, \mathbf{w}_k, |\mathbf{v}| \models \varphi_k,\ k = 1, 2. \end{array}$$

For the sake of simplicity, we allow only PLTL formulas to be operands of $<_i$ and $\not\leq_i$, with no embedded branching time constructs.

In words, $M, \mathbf{v} \models \varphi_1 <_i \varphi_2$ means that i prefers φ_2 plays to φ_1 ones. Unlike $\neg(\varphi_1 <_i \varphi_2)$, $\varphi_1 \not\leq_i \varphi_2$ holds only if no φ_1 play is preferable to a φ_2 one. We abbreviate $\bigwedge_{i\in\Gamma} \varphi <_i \psi$ and $\bigwedge_{i\in\Gamma} \varphi \not\leq_i \psi$ by $\varphi <_\Gamma \psi$ and $\varphi \not\leq_\Gamma \psi$, respectively.

According to [vBvOR05], the non-temporal archetype of $<_i$ can be traced back to [vW63]. It is just one of 8 preference operators featured in [vBvOR05], the one about the case in which *all* the φ_1-plays and *all* the φ_2-plays are related by the *strict* ordering $<_i$ on individual plays. The remaining 7 operators are defined by changing *all* to *some*, and changing $<_i$ to the non-strict $\lesssim_i$. Similar operators have been investigated in the setting of discrete contact spaces and relational syllogistic in [BTV07, IV12]. In [Lor10], the state space consists of the global strategy profiles. The graded preference relations $\gtrsim_i^k$, $k \leq n$, relate the reference profile to the profiles which are of *quality k or higher* to player i, $k \in \{0, \ldots, n\}$, and are used to define unary modalities in the standard way. The binary operator $\lesssim_i$ is defined by the clause

$$\varphi \lesssim_i \psi \mathrel{\hat{=}} \bigwedge_{k\leq n} \langle \gtrsim_i^k \rangle \varphi \Rightarrow \langle \gtrsim_i^k \rangle \psi.$$

4 Encoding Solution Concepts in ATL*

To conveniently refer to objectives as they stand at the beginning of time $\mathsf{I} \mathrel{\hat{=}} \neg \ominus \top$, we write

$$[\theta] \mathrel{\hat{=}} \diamondminus(\mathsf{I} \Rightarrow \theta)\ .$$

In this section we use the notation for temporary coalitions to express some solution concepts in QCTL* with the preference operator $<_i$.

The Core According to [CEW12], the *core* of a game is a set of decisions which is preferable to all coalitions. Let θ_* define the plays which can be the outcome of carrying out a core SPTC. Then, assuming that θ_* is achieved by SPTC r, we must have, for any alternative SPTC s,

$$(M^T)^{\mathbf{X}(r),\mathbf{X}(s)}_{\mathbf{r},\ \mathbf{s}}, w_I^T \models \forall(\Box\bigcirc\widetilde{\mathbf{r}}_\Sigma \Rightarrow \theta_*) \wedge \bigwedge_{\theta_\circ\in\Theta_{I,\Sigma}} \forall(\Box\bigcirc\widetilde{\mathbf{s}}_\Sigma \Rightarrow \theta_\circ) \Rightarrow \bigwedge_{i\in\Sigma}[\theta_\circ \wedge \neg\theta_*] <_i [\theta_*]$$

Domination with Temporary Coalitions With no assumptions on coalitioning, a strategy $s \in S_{\{i\}}$ is *dominant*, if it brings achievements which, in the view of player i, are better than those guaranteed by any other strategy, regardless of what the other players do. One simple way to generalize this to temporary coalitions is to require the outcome to be better for the coalition members, regardless of what the non-members of the coalitions entered by i in the various parts of plays do. To express that a certain SPTC r is dominant for i, we need to state that any strategy profile s is either the same as r to i, (assuming the resilience of the coalitions in which i participates, or performs worse than r for i:

$$(M^T)^{\mathbf{X}(r),\mathbf{X}(s)}_{\mathbf{r},\ \mathbf{s}}, w_I^T \models \left(\begin{array}{l} \forall\Box \bigwedge_{i\in\Gamma\subseteq\Sigma} (\mathbf{s}_\Gamma \Leftrightarrow \mathbf{r}_\Gamma) \vee \\ \bigwedge_{\theta'\in\Theta_{I,i}} \left(\begin{array}{l} \forall\left(\Box\bigcirc \bigvee_{i\in\Gamma\subseteq\Sigma} \mathbf{s}_\Gamma \Rightarrow [\theta']\right) \Rightarrow \\ \bigvee_{\theta''\in\Theta_{I,i}} [\theta'] <_i [\theta''] \wedge \forall\left(\Box\bigcirc \bigvee_{i\in\Gamma\subseteq\Sigma} \mathbf{r}_\Gamma \Rightarrow [\theta'']\right) \end{array} \right) \end{array} \right)$$

Note that, since the variables from $\mathbf{s}$ denote a SPTC, the variables $\mathbf{s}_\Gamma$ from $\Box\bigcirc \bigvee\limits_{i\in\Gamma\subseteq\Sigma} \mathbf{s}_\Gamma$ can each be true for at most one Γ at any state. The same holds about $\mathbf{r}_\Gamma$. Importantly, the above formula does not state that the players who are joining i for the various steps of plays according to r are interested in doing so.

For an r to be dominant for all $i\in\Sigma$, an arbitrary profile s should either coincide with r, or perform worse than r for all $i\in\Sigma$, i.e., in formulas, $(M^T)^{\mathbf{X}(r),\mathbf{X}(s)}_{\mathbf{r},\ \mathbf{s}}, w_I^T \models D(\mathbf{s},\mathbf{r})$ where

$$D(\mathbf{s},\mathbf{r}) \hat{=} \left(\begin{array}{l} \forall\Box \bigwedge\limits_{\Gamma\subseteq\Sigma} \mathbf{s}_\Gamma \Leftrightarrow \mathbf{r}_\Gamma \vee \\ \bigwedge\limits_{\theta'\in\prod\limits_{i\in\Sigma}\Theta_{I,i}} \left(\begin{array}{l} \bigwedge\limits_{i\in\Sigma} \forall\left(\Box\bigcirc \bigvee\limits_{i\in\Gamma\subseteq\Sigma} \mathbf{s}_\Gamma \Rightarrow [\theta'_i]\right) \Rightarrow \\ \bigvee\limits_{\theta''\in\prod\limits_{i\in\Sigma}\Theta_{I,i}} \left(\bigwedge\limits_{i\in\Sigma} [\theta'_i] <_i [\theta''_i] \wedge \bigwedge\limits_{i\in\Sigma} \forall\left(\Box\bigcirc \bigvee\limits_{i\in\Gamma\subseteq\Sigma} \mathbf{r}_\Gamma \Rightarrow [\theta''_i]\right)\right) \end{array} \right) \end{array} \right)$$

Let $\{\Gamma_0,\ldots,\Gamma_{2^{|\Sigma|-1}}\} \hat{=} \mathscr{P}(\Sigma)$. Then the existence of a dominant profile can be expressed in QCTL* with $<_i$ by the condition

$$M^T, w_I^T \models \exists \mathbf{r}_{\Gamma_0} \ldots \exists \mathbf{r}_{\Gamma_{2^{|\Sigma|-1}}} (\tilde{\delta}(\mathbf{r}) \wedge \forall \mathbf{s}_{\Gamma_0} \ldots \forall \mathbf{s}_{\Gamma_{2^{|\Sigma|-1}}} (\tilde{\delta}(\mathbf{s}) \Rightarrow D(\mathbf{s},\mathbf{r}))).$$

5 Axioms for $<_i$ and $\not\ll_i$

The axioms in this section are valid without the assumption on $<_i$ to be partitioning $R_M^{\inf}(w_I)$ into finitely many LTL-definable indiscernibility classes. $P1$ expresses extensionality and is the form of **K** that applies to binary modalities. The axioms $P2$ state that $<_\Gamma$ and $\not\ll_\Gamma$ are closed under disjunctions on both sides. $P3$ and $P4$ state that $<_i$ is irreflexive and transitive, respectively. The axioms $P5$ state that $<_\Gamma$ and $\not\ll_\Gamma$ trivially hold, if one of the operands is $\bot$. P6 states that players stick to their preferences. Below σ stands for $<$ or $\not\ll$.

$$\varphi_1\sigma_\Gamma\psi_1 \wedge \forall\bigcirc(\varphi_2 \Rightarrow \varphi_1) \wedge \forall\bigcirc(\psi_2 \Rightarrow \psi_1) \Rightarrow \varphi_2\sigma_\Gamma\psi_2 \tag{P1}$$

$$\varphi_1\sigma_\Gamma\psi \wedge \varphi_2\sigma_\Gamma\psi \Leftrightarrow (\varphi_1\vee\varphi_2)\sigma_\Gamma\psi, \qquad \varphi\sigma_\Gamma\psi_1 \wedge \varphi\sigma_\Gamma\psi_2 \Leftrightarrow \varphi\sigma_\Gamma(\psi_1\vee\psi_2) \tag{P2}$$

$$\varphi <_\Gamma \psi \Rightarrow \forall\bigcirc\neg(\varphi\wedge\psi) \qquad \varphi <_\Gamma \psi \Rightarrow \psi \not\ll_\Gamma \varphi \qquad \varphi \not\ll_\Gamma \psi \Rightarrow \neg(\varphi <_\Gamma \psi) \tag{P3}$$

$$\varphi <_\Gamma \psi \wedge \psi <_\Gamma \chi \wedge \exists\bigcirc\psi \Rightarrow \varphi <_\Gamma \chi \tag{P4}$$

$$\bot\sigma_\Gamma\varphi, \qquad \varphi\sigma_\Gamma\bot \tag{P5}$$

$$[\varphi]\sigma_\Gamma[\psi] \Leftrightarrow \forall\Box([\varphi]\sigma_\Gamma[\psi]) \tag{P6}$$

According to [vBvOR05], $P2$ and a variant of $P4$ without the conjunctive member $\exists\bigcirc\psi$ can be traced back to [vW63]. Interestingly, that variant of $P4$ is unsound in our semantics: From $\varphi_1 <_\Gamma \bot$ and $\bot <_\Gamma \varphi_2$, $P4$ with the $\exists\bigcirc\bot$ (which is false) deleted entails $\varphi_1 <_\Gamma \varphi_2$.

The meaning of $[\theta]$ after finite plays which render θ either ultimately failed ($M^T, \mathbf{w} \models \forall\neg[\theta]$) or achieved ($M^T, \mathbf{w} \models \forall[\theta]$) reflects removing from sight objectives that are no longer relevant. The example below illustrates this.

Example 3 Consider the 3-player setting from Example 2 with the ordering of 1's objectives given in (9). Suppose that p_3 occurs. Then the objectives from (9) which include $\Box\neg p_3$ become forfeited. The others simplify to objectives which no longer mention p_3 and are ordered as in (8). Formally, P1 entails that, e.g., $[\Diamond p_1 \wedge \Box\neg p_2 \wedge \Diamond p_3] <_1 [\Diamond p_1 \wedge \Diamond p_2 \wedge \Diamond p_3]$ is equivalent to

$$((\Diamond\!\!\!\!-\, p_3 \vee \boxminus\neg p_3 \wedge \Diamond p_3) \wedge [\Diamond p_1 \wedge \Box\neg p_2]) <_i ((\Diamond\!\!\!\!-\, p_3 \vee \boxminus\neg p_3 \wedge \Diamond p_3) \wedge [\Diamond p_1 \wedge \Diamond p_2]).$$

Now a somewhat longer deduction entails the validity of

$$p_3 \Rightarrow ([\Diamond p_1 \wedge \Box\neg p_2 \wedge \Diamond p_3] <_1 [\Diamond p_1 \wedge \Diamond p_2 \wedge \Diamond p_3] \Leftrightarrow [\Diamond p_1 \wedge \Box\neg p_2] <_1 [\Diamond p_1 \wedge \Diamond p_2]).$$

Similarly, once p_2 occurs too, (8) simplifies to just $[\Diamond p_1] <_1 [\Box\neg p_1]$ because then the continuations of the play do not satisfy $[\Diamond p_1 \wedge \Box\neg p_2]$ and $[\Box\neg p_1 \wedge \Box\neg p_2]$, and $\Diamond p_2$ from $[\Diamond p_1 \wedge \Diamond p_2]$ and $[\Box\neg p_1 \wedge \Diamond p_2]$ simplifies to $\top$. If $\Box(\diamondminus p_2 \wedge \diamondminus p_3 \Rightarrow p_1 \vee \Box\neg p_1)$, *nothing can kill* 1 *but* 2 *or* 3, definitely a solid reason for 1 to hate 2 and 3, is valid in the model, then 1 can relax after p_2 and p_3 occur.

6 Reasoning with Finitely Many Given Objectives

As stated in Section 3, we assume $<_i$ to be partitioning plays into finitely many classes of pairwise indiscernible plays, and we assume these classes to be definable by some given finite set of formulas $\Theta_{I,i}$ for every $i \in \Sigma$. Consider the axioms

$$\exists\bigcirc(\varphi \wedge [\theta]) \Rightarrow (\varphi\sigma_\Gamma\psi \Leftrightarrow (\varphi \vee [\theta])\sigma_\Gamma\psi), \qquad \exists\bigcirc(\psi \wedge [\theta]) \Rightarrow (\varphi\sigma_\Gamma\psi \Leftrightarrow \varphi\sigma_\Gamma(\psi \vee [\theta])) \tag{O1}$$

$$\bigvee_{\theta\in\Theta_{I,\Gamma}} [\theta], \qquad \bigwedge_{\theta_1,\theta_2\in\Theta_{I,\Gamma}} \forall\Box([\theta_1] \Leftrightarrow [\theta_2]) \vee \forall\Box\neg([\theta_1] \wedge [\theta_2]) \tag{O2}$$

$$[\theta_1] <_\Gamma [\theta_2],\ \text{resp. } [\theta_1] \not\lessgtr_\Gamma [\theta_2],\ \text{for } \theta_1, \theta_2 \text{ such that } \theta_1 <_\Gamma \theta_2,\ \text{resp. } \theta_1 \not<_\Gamma \theta_2,\ \text{is given.} \tag{O3}$$

for $\theta, \theta_1, \theta_2 \in \Theta_{I,\Gamma}$, $\Gamma \subseteq \Sigma$, $\sigma \in \{<, \not\lessgtr\}$.

The axioms from Section 5 and this section are complete for $<_i$ and $\not\lessgtr_i$ in CTL* with $<_i$ and $\not\lessgtr_i$, $i \in \Sigma$, relative to validity in (just) CTL*. The full deductive power of the axioms about $<_i$ and $\not\lessgtr_i$ from Section 5 is not really necessary for the case of finitely many given objectives. The completeness proof is based on the possibility to eliminate the occurrences of $<_i$, $\not\lessgtr_i$ and this way show that formulas with $<_i$ have $<_i$- and $\not\lessgtr_i$-free equivalents. Hence the small model property is inherited from CTL* without $<_i$ and $\not\lessgtr_i$.

Lemma 4 *Let φ_1 and φ_2 be* PLTL *formulas, and let $\Gamma \subseteq \Sigma$. Then the formulas below, $\sigma \in \{<, \not\lessgtr\}$, are derivable in* CTL* *by the axioms from Section 5 and this section:*

$$\varphi_1\sigma_\Gamma\varphi_2 \Leftrightarrow \bigvee_{\Theta_1,\Theta_2\subseteq\Theta_{I,\Gamma}} \bigwedge_{k=1,2} \left(\bigwedge_{\theta_k\in\Theta_k} \exists\bigcirc([\theta_k] \wedge \varphi_k) \wedge \bigwedge_{\theta_k\in\Theta_{I,\Gamma}\setminus\Theta_k} \forall\neg\bigcirc([\theta_k] \wedge \varphi_k) \right) \wedge \bigwedge_{\theta_1\in\Theta_1,\theta_2\in\Theta_2} [\theta_1]\sigma_\Gamma[\theta_2] \tag{E_σ}$$

The theorem below follows from the possibility to eliminate arbitrary uses of $<_i$ and $\not\lessgtr_i$ by means of E_σ and resolve uses with the operands being designated objectives using O3.

Theorem 5 (relative completeness) *Every formula in* CTL* *with $<_i$ and $\not\lessgtr_i$ has a $<_i$- and $\not\lessgtr_i$-free equivalent in* CTL*. *The equivalence is derivable by the axioms and rules from Section 5 and this section.*

7 Reasoning with No Model-supplied System of Objectives

This section is about the relative completeness of the axioms about $<_i$ from Section 5, with no assumptions on the properties of indiscernibility wrt $<_i$ in the considered models. The result shows that the satisfiability of any given a formula φ in CTL* in extended CGMs with the indiscernibility induced by

$<_i$ not necessarily having finite index, is equivalent to the satisfiability of φ in an extended CGM with a finite system of objectives which can be determined from φ. These objectives are boolean combinations of the operands of the occurrences of $<_i$ in φ.

We exclude the axioms from Section 6. Instead we add two axioms about the interaction of $<_i$ and $\not\lessdot_i$ with the *separated normal form* [Gab87] of PLTL formulas and the *guarded normal form* in (future) LTL. Let $\sigma \in \{<, \not\lessdot\}$. Let $\{g_0, \dots, g_{2^{|AP|}-1}\}$ be the set of all the conjunctions of the form $\bigwedge\limits_{p \in AP} \varepsilon_p p$ where ε_p is either $\neg$ or nothing for each $p \in AP$. Then

$$\left(\bigwedge_{k'} \ominus\pi'_{k'} \Rightarrow \varphi'_{k'}\right)\sigma_\Gamma\left(\bigwedge_{k''} \ominus\pi''_{k''} \Rightarrow \varphi''_{k''}\right) \Leftrightarrow \bigwedge_{k',k''}(\pi'_{k'} \wedge \pi''_{k''} \Rightarrow \varphi'_{k'}\sigma_\Gamma\varphi''_{k''}) \tag{P7}$$

$$\left(\bigvee_{k<2^{|AP|}} g_k \wedge \bigcirc\varphi^t_k\right)\sigma_\Gamma\left(\bigvee_{k<2^{|AP|}} g_k \wedge \bigcirc\psi^t_k\right) \Leftrightarrow \forall\bigcirc\left(\bigvee_{k<2^{|AP|}} g_k \wedge \varphi^t_k\sigma_\Gamma\psi^t_k\right) \tag{P8}$$

Formulas φ which are consistent with P1,...,P8, are satisfiable in a model where the preference relations partition the set of the infinite plays into finitely many classes of pairwise indiscernible plays. This means that our relative completeness result applies to the class of models of the form (7) with finite systems of objectives too. Note that propositional quantification and $\not\lessdot_i$ are not included in the completeness result below. Proving the exact form of the completeness result involves a finite set of instances of the axioms P1,...,P8, which depends on the formula φ whose satisfiability is considered. We denote their conjunction by Ax_φ.

Theorem 6 (relative completeness) *Let $Ax_\varphi \wedge \varphi$* not *be the negation of a formula that is valid in* CTL*, *assuming that the $<_i$-subformulas in $Ax_\varphi \wedge \varphi$ are treated as atomic propositions. Then φ is satisfiable in an extended CGM of the form (6).*

Theorem 7 (finite systems of objectives) *Assume that φ is as in Theorem 6. Then φ is satisfiable in an extended CGM of the form (7).*

Concluding Remarks

We have proposed a way to use temporal logic for strategic reasoning about temporary coalitions, which are outside the immediate scope of the basic constructs of established logics for strategic reasoning such as ATL and SL as both ATL's $\langle\langle.\rangle\rangle$ and SL's first order language for strategies as the domain of individuals are meant to model long term individual player strategies. The proposed notation builds on the use of propositional variables to denote decisions and, more generally, strategies, and quantifying over strategies, in QCTL*. We extended the notation to capture evolving coalition structure too. Furthermore, we have extended CTL* with a preference operator on objectives and proposed a complete set of axioms for that operator. We have illustrated the use of the notation by specifying temporary coalition variants of the solution concepts of *the core* and *dominant strategies*.

Acknowledgement

This work was partially supported by Contract DN02/15/19.12.2016 "Space, Time and Modality: Relational, Algebraic and Topological Models" with Bulgarian NSF.

References

[AHK97] Rajeev Alur, Tom Henzinger, and Orna Kupferman. Alternating-time Temporal Logic. In *Proceedings of FCS'97*, pages 100–109, 1997.

[AHK02] Rajeev Alur, Tom Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):1–42, 2002.

[BBMU12] Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Concurrent Games with Ordered Objectives. In *FOSSACS 2012*, volume 7213 of *LNCS*, pages 301–315. Springer, 2012.

[BBMU15] Patricia Bouyer, Romain Brenguier, Nicolas Markey, and Michael Ummels. Pure Nash Equilibria in Concurrent Deterministic Games. *Logical Methods in Computer Science*, 11(2), 2015.

[BLLM09] Thomas Brihaye, Arnaud Da Costa Lopes, François Laroussinie, and Nicolas Markey. ATL with Strategy Contexts and Bounded Memory. In *LFCS*, volume 5407 of *LNCS*, pages 92–106, 2009.

[BMS20] Laura Bozzelli, Aniello Murano, and Loredana Sorrentino. Alternating-time temporal logics with linear past. *Theor. Comput. Sci.*, 813:199–217, 2020.

[BTV07] Philippe Balbiani, Tinko Tinchev, and Dimiter Vakarelov. Dynamic Logics of the Region-based Theory of Discrete Spaces. *J. Appl. Non Class. Logics*, 17(1):39–61, 2007.

[CEW12] Georgios Chalkiadakis, Edith Elkind, and Michael Wooldridge. *Computational Aspects of Cooperative Game Theory*. Morgan & Claypool Publishers, 2012.

[FG92] Marcelo Finger and Dov M. Gabbay. Adding a Temporal Dimension to a Logic System. *Journal of Logic, Language and Information*, 1(3):203–233, 1992.

[FKL10] Dana Fisman, Orna Kupferman, and Yoad Lustig. Rational Synthesis. In *TACAS 2010*, volume 6015 of *LNCS*, pages 190–204. Springer, 2010.

[Fre01] Tim French. Decidability of Quantifed Propositional Branching Time Logics. In *Australian Joint Conference on Artificial Intelligence*, volume 2256 of *LNCS*, pages 165–176. Springer, 2001.

[Fre06] Tim French. *Bisimulation Quantifiers for Modal Logics*. Ph.D. Thesis, The University of Western Australia, 2006. Accessed in 2011 from `http://www.csse.uwa.edu.au/~tim/papers/thesis2.pdf`.

[Gab87] Dov M. Gabbay. The Declarative Past and Imperative Future: Executable Temporal Logic for Interactive Systems. In *Temporal Logic in Specification, Altrincham, UK, April 8-10, 1987, Proceedings*, volume 398 of *LNCS*, pages 409–448. Springer, 1987.

[GD12] Dimitar P. Guelev and Catalin Dima. Epistemic ATL with Perfect Recall, Past and Strategy Contexts. In *CLIMA XIII, 2012*, volume 7486 of *LNCS*, pages 77–93. Springer, 2012.

[GJ04] Valentin Goranko and Wojtek Jamroga. Comparing Semantics for Logics of Multi-agent Systems. *Synthese*, 139(2):241–280, 2004.

[GS98] Dov M. Gabbay and Valentin B. Shehtman. Products of Modal Logics, Part 1. *Log. J. IGPL*, 6(1):73–146, 1998.

[Gue13] Dimitar P. Guelev. Reducing Validity in Epistemic ATL to Validity in Epistemic CTL. In *Proceedings 1st International Workshop on Strategic Reasoning, SR*, volume 112 of *EPTCS*, pages 81–89, 2013.

[Han04] Sven Ove Hansson. Preference logic. In Dov M. Gabbay and Franz Guenthner, editors, *Handbook of Philosophical Logic, 2nd Edition*, pages 319–394. Springer, 2004.

[IV12] Nikolay Ivanov and Dimiter Vakarelov. A System of Relational Syllogistic Incorporating Full Boolean Reasoning. *Journal of Logic, Language and Information*, 21(4):433–459, 2012.

[KPV16] Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. Synthesis with rational environments. *Ann. Math. Artif. Intell.*, 78(1):3–20, 2016.

[LLM12] Arnaud Da Costa Lopes, François Laroussinie, and Nicolas Markey. Quantified CTL: Expressiveness and Model Checking - (Extended Abstract). In *CONCUR 2012. Proceedings*, volume 7454 of *LNCS*, pages 177–192. Springer, 2012.

[Lor10] Emiliano Lorini. A Logical Account of Social Rationality in Strategic Games. Technical Report IRIT/RT2010-3FR, UPS-IRIT, 2010. LOFT 9, Accessed from `https://www.irit.fr/publis/LILAC/Reports/LoriniGameTheoryRationality.pdf` in 2020.

[MMV10] Fabio Mogavero, Aniello Murano, and Moshe Y. Vardi. Reasoning about strategies. In *FSTTCS 2010*, volume 8 of *LIPIcs*, pages 133–144, 2010.

[Pin07] Sophie Pinchinat. A Generic Constructive Solution for Concurrent Games with Expressive Constraints on Strategies. In *ATVA 2007, Proceedings*, volume 4762 of *LNCS*, pages 253–267. Springer, 2007.

[vBvOR05] J. van Benthem, S. van Otterloo, and O. Roy. Preference Logic, Conditionals and Solution Concepts in Games. Technical report, Inst. for Logic, Language and Computation, 2005.

[vW63] G. H. von Wright. *The Logic of Preference*. Edinburgh University Press, Edinburgh, 1963.

[WHY11] Farn Wang, Chung-Hao Huang, and Fang Yu. A temporal logic for the interaction of strategies. In *CONCUR*, volume 6901 of *LNCS*, pages 466–481. Springer, 2011.

[WvdHW07] Dirk Walther, Wiebe van der Hoek, and Michael Wooldridge. Alternating-time Temporal Logic with Explicit Strategies. In Dov Samet, editor, *TARK*, pages 269–278. ACM Press, 2007.