

Temporal Qualitative Coalitional Games

Thomas Ågotnes
Department of Informatics,
University of Bergen
PB. 7800, N-5020 Bergen,
Norway
agotnes@ii.uib.no

Wiebe van der Hoek
Department of Computer
Science, University of
Liverpool
Liverpool L69 7ZF, UK
wiebe@csc.liv.ac.uk

Michael Wooldridge
Department of Computer
Science, University of
Liverpool
Liverpool L69 7ZF, UK
mjw@csc.liv.ac.uk

ABSTRACT

Qualitative Coalitional Games (QCGs) are a version of coalitional games in which an agent's desires are represented as goals which are either satisfied or unsatisfied, and each choice available to a coalition is a set of goals, which would be jointly satisfied if the coalition made that choice. A coalition in a QCG will typically form in order to bring about a set of goals that will satisfy all members of the coalition. In this paper, we introduce and study *Temporal QCGs* (TQCGs), i.e., games in which a sequence of QCGs is played. In order to represent and reason about such games, we introduce a linear time temporal logic of QCGs, known as $\mathcal{L}(TQCG)$. We give a complete axiomatization of $\mathcal{L}(TQCG)$, use it to investigate the properties of TQCGs in a small example, identify its expressive power, establish its complexity, characterise classes of TQCGs with formulas from our logical language, and formulate several (temporal) solution concepts for TQCGs.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems;
I.2.4 [Knowledge representation formalisms and methods]

General Terms

Theory

Keywords

Coalitional games, repeated games, logic

1. INTRODUCTION

The study of repeated games now forms a major component of the game theory literature [9, pp.133–161]. Perhaps the best-known example of such a repeated game is the iterated prisoner's dilemma, which has for example been studied both analytically [3, pp.353–358] and by means of competitions [2].

Given the role of game theory as a theoretical underpinning to the multi-agent systems field [11], it seems that repeated games are

of particular importance to the field. By-and-large, we are not interested in building multi-agent systems that will operate in a “one-shot” fashion: we typically want them to operate over time, often without a pre-defined termination time. Moreover, given the important role that *coalitional games* play in multi-agent systems [12, 13], it seems that repeated coalitional games are also likely to be of significance. However, comparatively little research has considered repeated coalitional games, or coalitional games played over time [7].

Our aim in this paper is to study an iterated version of *Qualitative Coalitional Games* (QCGs) [15], a variation of coalitional games in which agent's desires are represented as goals which are either satisfied or unsatisfied, and each choice available to a coalition is a set of goals, which would be jointly satisfied if the coalition made that choice. A coalition in a QCG will typically form in order to bring about a set of goals that will satisfy all members of the coalition. In this paper, we introduce and study *Temporal QCGs* (TQCGs): games in which QCGs are played repeatedly. In order to represent and reason about such games, we introduce $\mathcal{L}(TQCG)$, a linear time temporal logic of QCGs. We give a complete axiomatization of $\mathcal{L}(TQCG)$, demonstrate its expressive power with respect to a type of simulation between TQCGs, establish the computational complexity of satisfiability for TQCGs, investigate the properties of TQCGs by characterising them as formulae in $\mathcal{L}(TQCG)$ and finally characterise some solution concepts of TQCGs in $\mathcal{L}(TQCG)$.

We begin, in the following section, with a short introduction to QCGs. As part of this introduction, in section 2.1 we define a logic for expressing properties of individual QCGs: this logic will serve as the “state language” or “assertion language” for the temporal QCG logic. In section 3, we introduce TQCGs. We begin with a short informal motivation, then give the temporal language for expressing properties of TQCGs, present an example to illustrate the idea of TQCGs and the role of the temporal language, investigate the expressive power of the language by means of a simulation relation between TQCG structures, give a complete axiomatization of the temporal language, and then investigate the axiomatic characterisation of various classes of TQCGs.

2. QUALITATIVE COALITIONAL GAMES

We give a brief introduction to Qualitative Coalitional Games (QCGs): details may be found in [15]. A QCG contains a (non-empty, finite) set $\mathcal{A} = \{1, \dots, m\}$ of *agents*. Each agent $i \in \mathcal{A}$ is assumed to have associated with it a (finite) set \mathcal{G}_i of *goals*, drawn from a set of overall possible goals \mathcal{G} . The intended interpretation is that the members of \mathcal{G}_i represent all the individual rational outcomes for i – intuitively, the outcomes that give it “better than zero utility”. That is, agent i would be happy if *any* member of \mathcal{G}_i were

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.
Copyright 2006 ACM 1-59593-303-4/06/0005 ...\$5.00.

achieved – then it has “gained something”. But, in QCGs, we are not concerned with preferences over individual goals. Thus, at this level of modelling, i is *indifferent* among the members of \mathcal{G}_i : it will be *satisfied* if *at least one* member of \mathcal{G}_i is achieved, and *unsatisfied* otherwise. Note that cases where more than one of an agent’s goals are satisfied are not an issue – an agent’s aim will simply be to ensure that at least one of its goals is achieved, and there is no sense of an agent i attempting to satisfy as many members of \mathcal{G}_i as possible.

A *coalition*, typically denoted by C , is simply a set of agents, i.e., a subset of \mathcal{A} . The *grand coalition* is the set of all agents, \mathcal{A} . We assume that each possible coalition has available to it a set of possible *choices*, where each choice intuitively characterises the outcome of one way that the coalition could cooperate. We model the choices available to coalitions via a *characteristic function* with the signature $\mathcal{V} : 2^{\mathcal{A}} \rightarrow 2^{2^{\mathcal{G}}}$. Thus, in saying that $G \in \mathcal{V}(C)$ for some coalition $C \subseteq \mathcal{A}$, we are saying that one choice available to the coalition C is to bring about *exactly* the goals in G . At this point, the reader might expect to see some constraints placed on characteristic functions. For example, at first sight the following *monotonicity* constraint might seem natural: $C \subseteq C'$ implies $\mathcal{V}(C) \subseteq \mathcal{V}(C')$. Although such a constraint is entirely appropriate for many scenarios, there are cases where such a constraint is not appropriate¹.

Bringing these components together, a qualitative coalitional game (QCG) is a tuple:

$$\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle \text{ where}$$

- \mathcal{A} is a finite, non-empty set of *agents*;
- \mathcal{G} is a finite, non-empty set of possible *goals*;
- $\mathcal{G}_i \subseteq \mathcal{G}$ is the set of goals for agent $i \in \mathcal{A}$; and
- $\mathcal{V} : 2^{\mathcal{A}} \rightarrow 2^{2^{\mathcal{G}}}$ is the characteristic function of the game.

EXAMPLE 1. Let Γ_1 be the following QCG for a collection of agents and a collection of goals $\{g_1, \dots\}$. Agent 1 is satisfied with g_1 and g_4 , and agent 2 is satisfied with g_2 and g_3 . The characteristic function, where C_1, C_2, C_3, C_4 are different coalitions:

$$\begin{aligned} \mathcal{V}(C_1) &= \{ \{g_1, g_2\} \} & \mathcal{V}(C_2) &= \{ \{g_2, g_3\}, \{g_1\} \} \\ \mathcal{V}(C_3) &= \{ \{g_5, g_6\} \} & \mathcal{V}(C_4) &= \{ \{g_2, g_3\}, \{g_1\}, \{g_4\} \} \end{aligned}$$

We will make use of Γ_1 in later examples.

2.1 A Logic for QCGs

A logic for expressing properties of individual QCGs has not been formalised before. We now introduce such a logic. This logic will later be used as the assertion language, or state language, for the temporal logic we develop in section 3. The language is defined in two parts: \mathcal{L}_c is the *satisfaction language*, and is used to express properties of choices made by agents. The basic constructs in this language are of the form sat_i , meaning “agent i is satisfied”. The overall language $\mathcal{L}(QCG)$ is used for expressing properties of QCGs themselves. The main construct in this language is of the form $\langle C \rangle \varphi$, where φ is a formula of the satisfaction language, and means that C have a choice such that this choice makes φ true. For example, $\langle 3 \rangle (sat_1 \wedge sat_4)$ will mean that 3 has a choice that simultaneously satisfies agents 1 and 4.

Formally, the grammar φ_c defines the *satisfaction language* \mathcal{L}_c , while φ_q defines the *QCG language* $\mathcal{L}(QCG)$.

$$\begin{aligned} \varphi_c &::= sat_i \mid \neg \varphi_c \mid \varphi_c \vee \varphi_c \\ \varphi_q &::= \langle C \rangle \varphi_c \mid \neg \varphi_q \mid \varphi_q \vee \varphi_q \end{aligned}$$

¹For example, consider a legal scenario in which certain coalitions are forbidden by monopoly or anti-trust laws.

where $i \in \mathcal{A}$ and $C \subseteq \mathcal{A}$. (We note some similarities between our logical language $\mathcal{L}(QCG)$ and Pauly’s language for axiomatizing judgement aggregation procedures [10], although the motivation and use of the languages are quite different.)

We use the usual derived propositional connectives ($\wedge, \rightarrow, \leftrightarrow$) for both languages \mathcal{L}_c and $\mathcal{L}(QCG)$, and in addition write $[C]\varphi$ to abbreviate $\neg \langle C \rangle \neg \varphi$. The formula $[C]\varphi$ will be defined to be true exactly when φ is a *necessary* consequence of the coalition C making a choice; φ will be true no matter *which* choice the coalition makes. When $C = \{a\}$ is a singleton, we sometimes write $\langle a \rangle$ and $[a]$ for $\langle C \rangle$ and $[C]$.

When $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ is a QCG, $G \subseteq \mathcal{G}$ and $\varphi \in \mathcal{L}_c$, $\Gamma, G \models_Q \varphi$ is defined as follows:

$$\begin{aligned} \Gamma, G \models_Q sat_i &\text{ iff } \mathcal{G}_i \cap G \neq \emptyset \\ \Gamma, G \models_Q \neg \psi &\text{ iff not } \Gamma, G \models_Q \psi \\ \Gamma, G \models_Q \psi_1 \vee \psi_2 &\text{ iff } \Gamma, G \models_Q \psi_1 \text{ or } \Gamma, G \models_Q \psi_2 \end{aligned}$$

When $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ is a QCG and φ is a $\mathcal{L}(QCG)$ formula, $\Gamma \models_Q \varphi$ is defined as follows:

$$\begin{aligned} \Gamma \models_Q \langle C \rangle \psi &\text{ iff there is a } G \in \mathcal{V}(C) \text{ such that } \Gamma, G \models_Q \psi \\ \Gamma \models_Q \neg \psi &\text{ iff not } \Gamma \models_Q \psi \\ \Gamma \models_Q \psi_1 \vee \psi_2 &\text{ iff } \Gamma \models_Q \psi_1 \text{ or } \Gamma \models_Q \psi_2 \end{aligned}$$

EXAMPLE 2. Let Γ_1 be as in Example 1. Then:

$$\begin{aligned} \Gamma_1 \models_Q \langle C_1 \rangle (sat_1 \wedge sat_2) \\ \Gamma_1 \models_Q (\langle C_2 \rangle sat_1 \wedge \langle C_2 \rangle sat_2) \wedge \neg (\langle C_2 \rangle (sat_1 \wedge sat_2)) \\ \Gamma_1 \models_Q \neg (\langle C_3 \rangle sat_1 \vee \langle C_3 \rangle sat_2) \end{aligned}$$

Summarising, the satisfaction of agents is evaluated against a set of goals, while Boolean combinations of expressions referring to choices of coalitions are evaluated on a QCG Game Γ . The latter combinations will be the atomic assertions in our temporal framework of Section 3.

2.2 Expressive Power of L(QCG)

We look at the properties of QCGs which are definable in our language. It is clear from our language definition that what $\mathcal{L}(QCG)$ can express is which coalition can satisfy which set of agents currently. Note that we are not interested in *how* the coalitions make certain sets of agents satisfied, nor *why* an agent is satisfied (i.e., which goal satisfied him). We will now demonstrate that the properties of QCGs we can express in the language $\mathcal{L}(QCG)$ are exactly the properties closed under a notion of *QCG-simulation*. In other words, the language can not differentiate two games Γ and Γ' iff they *QCG-simulate* each other.

Obviously, equivalence of models transcends mere isomorphism. In particular, the semantics of performing a *choice* seem to depend only on which agents are satisfied by the choice. For example, one could imagine a mapping between “equivalent” goals of two models, maybe collapsing two goals of one model into one goal of the other. However, such a relation between models does not capture all instances of equivalent models. What is needed is a relation between *sets* of goals. This motivates the following definition of a QCG-simulation as a *relation* between two models. It is only necessary to relate goals which can actually be chosen by some coalition. Furthermore, it only makes sense to relate models which are defined over the same set of agents.

A relation

$$Z \subseteq \bigcup_{C \subseteq \mathcal{A}} (\mathcal{V}(C) \times \mathcal{V}'(C))$$

is a *QCG-simulation* between two QCGs $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ and $\Gamma' = \langle \mathcal{A}, \mathcal{G}', \mathcal{G}'_1, \dots, \mathcal{G}'_n, \mathcal{V}' \rangle$ iff the following conditions hold for all coalitions C .

1. If GZG' then $G \cap \mathcal{G}_i = \emptyset$ iff $G' \cap \mathcal{G}'_i = \emptyset$, for all i (the *satisfaction condition*)
2. For every $G \in \mathcal{V}(C)$ there is a $G' \in \mathcal{V}'(C)$ such that GZG' (Z is *total*)
3. For every $G' \in \mathcal{V}'$ there is a $G \in \mathcal{V}(C)$ such that GZG' (Z is *surjective*)

If there exist a QCG-simulation between two games Γ and Γ' , we write $\Gamma \rightleftharpoons \Gamma'$. If $\Gamma \rightleftharpoons \Gamma'$, we can simulate any choice in one model with a choice in the other, and vice versa. This notion of simulation is somewhat similar to the notion of “alternating simulation” between alternating transition systems in [1].

EXAMPLE 3. Let Γ_2 be the QCG with the same agents as in Γ_1 (Example 1), goals f_1, f_2, \dots such that agent 1 is satisfied in f_1 and f_3 and agent 2 is satisfied in f_2, f_3 and f_4 , and the following characteristic function:

$$\begin{aligned} \mathcal{V}(C_1) &= \{ \{f_3\} \} & \mathcal{V}(C_2) &= \{ \{f_2\}, \{f_1\} \} \\ \mathcal{V}(C_3) &= \{ \{f_5\} \} & \mathcal{V}(C_4) &= \{ \{f_1\}, \{f_2\}, \{f_4\} \} \end{aligned}$$

Then $\Gamma_1 \rightleftharpoons \Gamma_2$. The relation Z consisting of the following pairs is a QCG-simulation between Γ_1 and Γ_2 .

$$\begin{array}{lll} \langle \{g_1, g_2\}, \{f_3\} \rangle & \langle \{g_2, g_3\}, \{f_2\} \rangle & \langle \{g_1\}, \{f_1\} \rangle \\ \langle \{g_5, g_6\}, \{f_5\} \rangle & \langle \{g_2, g_3\}, \{f_2\} \rangle & \langle \{g_2, g_3\}, \{f_4\} \rangle \\ \langle \{g_1\}, \{f_1\} \rangle & \langle \{g_4\}, \{f_1\} \rangle & \end{array}$$

Note that Z is not a function, nor the inverse of a function.

We write $\Gamma \equiv \Gamma'$ iff $\forall \varphi \in \mathcal{L}(QCG) [\Gamma \models_Q \varphi \Leftrightarrow \Gamma' \models_Q \varphi]$.

THEOREM 1. *Satisfaction is invariant under QCG-simulation:*

$$\Gamma \rightleftharpoons \Gamma' \quad \Rightarrow \quad \Gamma \equiv \Gamma'$$

PROOF. Let $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ and $\Gamma' = \langle \mathcal{A}, \mathcal{G}', \mathcal{G}'_1, \dots, \mathcal{G}'_n, \mathcal{V}' \rangle$ with $\Gamma \rightleftharpoons \Gamma'$. First, we show that

$$GZG' \Rightarrow (\Gamma, G \models_Q \psi \Leftrightarrow \Gamma', G' \models_Q \psi) \quad (1)$$

for any ψ by induction over ψ . For the base case, let $\psi = sat_i$. $\Gamma, G \models_Q \psi$ iff $\mathcal{G}_i \cap G \neq \emptyset$ iff, by the satisfaction condition, $\mathcal{G}'_i \cap G' \neq \emptyset$ iff $\Gamma', G' \models_Q \psi$. The inductive step (negation and disjunction) is straightforward. We now show that

$$\Gamma \models_Q \varphi \Leftrightarrow \Gamma' \models_Q \varphi$$

for any φ by induction on φ . For the base case, let $\varphi = \langle C \rangle \psi$. For the direction to the right, if $\Gamma \models_Q \varphi$ then there is a $G \in \mathcal{V}(C)$ such that $\Gamma, G \models_Q \psi$. By totality of Z , there is a $G' \in \mathcal{V}'(C)$ such that GZG' . By (1), $\Gamma', G' \models_Q \psi$, and thus $\Gamma' \models_Q \varphi$. The direction to the left is symmetric: if $\Gamma' \models_Q \varphi$ there is a $G' \in \mathcal{V}'(C)$ such that $\Gamma', G' \models_Q \psi$; by surjectivity of Z there is a $G \in \mathcal{V}(C)$ such that GZG' ; and by (1) $\Gamma, G \models_Q \psi$ and thus $\Gamma \models_Q \varphi$. The inductive step (negation and disjunction) is straightforward. \square

The obvious question now is whether every pair of equivalent models are connected by a QCG-simulation. The answer is “yes”.

THEOREM 2. *Let Γ, Γ' be defined over the same set of agents:*

$$\Gamma \rightleftharpoons \Gamma' \quad \Leftarrow \quad \Gamma \equiv \Gamma'$$

PROOF. Let $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$ and $\Gamma' = \langle \mathcal{A}, \mathcal{G}', \mathcal{G}'_1, \dots, \mathcal{G}'_n, \mathcal{V}' \rangle$ with $\Gamma \equiv \Gamma'$. With any coalition C and any choice $G \in \mathcal{V}(C)$, associate the set $S_G^C = \{i : G \cap \mathcal{G}_i \neq \emptyset\}$ of agents satisfied if C chooses G . Similarly for Γ' : $T_{G'}^C = \{i : G' \cap \mathcal{G}'_i \neq \emptyset\}$ for any $G' \in \mathcal{V}'(C)$.

We define a QCG-simulation $Z : \Gamma \rightleftharpoons \Gamma'$ as follows: for every coalition C and pair of choices $G \in \mathcal{V}(C)$, $H \in \mathcal{V}'(C)$,

$$GZH \Leftrightarrow S_G^C = T_H^C$$

We must show that Z is total, i.e., that if $G \in \mathcal{V}(C)$, then there is a $H \in \mathcal{V}'(C)$ such that $S_G^C = T_H^C$. Suppose not: assume that $i \in S_G^C$ and $i \notin T_H^C$ for all $H \in \mathcal{V}'(C)$ (the argument is similar when $i \notin S_G^C$ and $i \in T_H^C$ for some $H \in \mathcal{V}'(C)$). Then $\Gamma \models_Q \langle C \rangle sat_i$ and $\Gamma' \models_Q \neg \langle C \rangle sat_i$, which contradicts the fact that $\Gamma \equiv \Gamma'$.

Similarly, we must show that Z is surjective, i.e., that if $H \in \mathcal{V}'(C)$, then there is a $G \in \mathcal{V}(C)$ such that $S_G^C = T_H^C$. Suppose not: assume that $i \in T_H^C$ and $i \notin S_G^C$ for all $G \in \mathcal{V}(C)$ (the argument is similar when $i \notin T_H^C$ and $i \in S_G^C$ for some $G \in \mathcal{V}(C)$). Then $\Gamma' \models_Q \langle C \rangle sat_i$ and $\Gamma \models_Q \neg \langle C \rangle sat_i$, which contradicts the fact that $\Gamma \equiv \Gamma'$.

Finally, we show that the satisfaction condition holds. If GZH , then $G \cap \mathcal{G}_i \neq \emptyset$ iff $i \in S_G^C$ iff, by the definition of Z , $i \in T_H^C$ iff $H \cap \mathcal{G}'_i \neq \emptyset$. \square

2.3 Axiomatisation for QCGs

We define a Hilbert style axiomatisation of qualitative coalitional games, and prove its soundness and completeness. We name our axiomatisation for QCGs $\mathbf{K}(QCG)$. This name emphasises the close resemblance to the modal system \mathbf{K} , which also indicates that our logic, in a sense, is a weakest basic system for QCGs, to which more sophisticated constraints can easily be added — such extensions are the topic of Section 4. The system $\mathbf{K}(QCG)$ over the language $\mathcal{L}(QCG)$ is defined as follows, where φ, ψ are arbitrary $\mathcal{L}(QCG)$ formulae, α, β are arbitrary \mathcal{L}_c formulae and C an arbitrary coalition:

$Prop^-$	If φ is an $\mathcal{L}(QCG)$ -instance of a propositional tautology, then φ is provable
K^-	$[C](\alpha \rightarrow \beta) \rightarrow ([C]\alpha \rightarrow [C]\beta)$ is provable
MP^-	If $\varphi, \varphi \rightarrow \psi$ are provable, then ψ is provable
Nec^-	If α is an (\mathcal{L}_c) instance of a propositional tautology, then $[C]\alpha$ is provable

It is easy to see that the deduction theorem holds for $\mathbf{K}(QCG)$.

We will need the following properties of $\mathbf{K}(QCG)$. The proofs are straightforward for readers familiar with modal logic.

LEMMA 1. *Let $\alpha, \beta \in \mathcal{L}_c$:*

1. $\vdash_{\mathbf{K}(QCG)} \langle C \rangle (\alpha \wedge \beta) \rightarrow \langle C \rangle \alpha$
2. $\vdash_{\mathbf{K}(QCG)} \langle C \rangle (\alpha \vee \beta) \rightarrow (\langle C \rangle \alpha \vee \langle C \rangle \beta)$
3. $\vdash_{\mathbf{K}(QCG)} (\langle C \rangle \alpha \wedge [C](\alpha \rightarrow \beta)) \rightarrow \langle C \rangle \beta$

THEOREM 3 (SOUNDNESS & COMPLETENESS). *For any $\Phi \subseteq \mathcal{L}(QCG)$, $\varphi \in \mathcal{L}(QCG)$: $\Phi \models_Q \varphi \Leftrightarrow \Phi \vdash_{\mathbf{K}(QCG)} \varphi$*

PROOF. For soundness (the direction to the left), it is easy to see that the axioms are valid, and that the rules preserve logical consequence.

For completeness, let $\Psi \subseteq \mathcal{L}(QCG)$ be $\mathbf{K}(QCG)$ consistent. We show that Ψ is satisfied by some QCG. Let \mathcal{A} be the set of agents and let $n = |\mathcal{A}|$. Let Δ be a $\mathcal{L}(QCG)$ maximal and $\mathbf{K}(QCG)$ consistent set containing Ψ (the proof of existence of such a set is the standard proof of Lindenbaum’s lemma). We now construct $\Gamma = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_n, \mathcal{V} \rangle$, intended to satisfy Ψ , as follows:

- $\mathcal{G} = \{sat_1, \dots, sat_n\}$
- $\mathcal{G}_i = \{sat_i\}$, for each i

- $X \in \mathcal{V}(C) \Leftrightarrow \langle C \rangle \xi_X \in \Delta$, for any $X \subseteq \mathcal{G}$, where

$$\xi_X \equiv \bigwedge_{sat_i \in X} sat_i \wedge \bigwedge_{i \in \mathcal{A}, sat_i \notin X} \neg sat_i$$

We show that

$$Gamma \models_Q \gamma \Leftrightarrow \gamma \in \Delta$$

for any γ by structural induction over γ . For the base case, $\gamma = \langle C \rangle \alpha$ for some $\alpha \in \mathcal{L}_c$. Again, we use induction on the structure of α . For the (nested) base case, let $\alpha = sat_i$. For the direction to the right, if $\Gamma \models_Q \gamma$ then there is an $X \in \mathcal{V}(C)$ such that $\Gamma, X \models_Q \alpha$, i.e., there is an $X \subseteq \mathcal{G}$ such that $\langle C \rangle \xi_X \in \Delta$ and $X \cap \{sat_i\} \neq \emptyset$. Thus, $sat_i \in X$, and by Lemma 1.1, $\gamma = \langle C \rangle sat_i \in \Delta$. For the direction to the left, let $\langle C \rangle sat_i \in \Delta$. Let

$$\chi_i = \bigvee_{S \subseteq \mathcal{A}} \xi_{(S \cup \{sat_i\})}$$

$sat_i \rightarrow \chi_i$ is a \mathcal{L}_c instance of a propositional tautology, so $[C](sat_i \rightarrow \chi_i) \in \Delta$ by *Nec*. By Lemma 1.3, $\langle C \rangle \chi_i \in \Delta$. By Lemma 1.2,

$$\bigvee_{S \subseteq \mathcal{A}} \langle C \rangle \xi_{(S \cup \{sat_i\})} \in \Delta$$

and thus $\langle C \rangle \xi_{S \cup \{sat_i\}} \in \Delta$ for some $S \subseteq \mathcal{A}$. It follows that $S \cup \{sat_i\} \in V(C)$, and since $\Gamma, (S \cup \{sat_i\}) \models_Q sat_i$ we get that $\Gamma \models_Q \langle C \rangle sat_i$ which concludes the proof of the direction to the left in the innermost induction proof. Both the inner and the outer induction steps (negation and disjunction) are straightforward. \square

Note that the completeness proofs demonstrates that we do not need to deal with multiple satisfaction of an agent's goal: in fact, one (symbol for a) goal for each agent is enough to reason about abilities for satisfaction!

3. TEMPORAL QCGS

In principle there are many ways to temporalise QCGs. As a first investigation, we assume a linear time model, in which, at each time point, a (possibly different) QCG Γ is played. A *temporal qualitative coalitional game* (TQCG) is then a triple

$$M = \langle S, \sigma, Q \rangle \text{ where:}$$

- S is a set of *states*;
- $\sigma : \mathbb{N} \rightarrow S$ associates a state $\sigma(u)$ with every natural number time point $u \in \mathbb{N}$; and
- $Q : S \rightarrow \mathbf{Q}$, where \mathbf{Q} is the class of all QCGs, is a function associating a qualitative coalitional game $Q(s) = \langle \mathcal{A}^s, \mathcal{G}^s, \mathcal{G}_1^s, \dots, \mathcal{G}_n^s, \mathcal{V}^s \rangle$ with every state s .

We will make just one requirement of TQCGs: that the set of agents and overall goals remains the same in all states. Formally, $\forall s, t \in S: \mathcal{A}^s = \mathcal{A}^t$ and $\mathcal{G}^s = \mathcal{G}^t$. This does not mean that an agent's goals must remain fixed, however: we allow for the possibility that an agent has different goals in different states. We also admit the possibility of a coalition having different choices in different states. Since the sets of agents and overall goals are fixed across all states, we will simply denote these by \mathcal{A} and \mathcal{G} respectively, omitting the state index.

3.1 A Logic for TQCGs

To express properties of TQCGs, we extend the QCG language $\mathcal{L}(QCG)$ with the standard temporal operators of linear-time temporal logic: \bigcirc – “next”, \diamond – “eventually”, \square – “always in the

future”, and \mathcal{U} – “until” [8]. Formally, the language $\mathcal{L}(TQCG)$ is defined by the grammar φ_t .

$$\varphi_t ::= \langle C \rangle \varphi_c \mid \neg \varphi_t \mid \varphi_t \vee \varphi_t \mid \varphi_t \mathcal{U} \varphi_t \mid \bigcirc \varphi_t$$

We again assume the usual derived propositional connectives, in addition to $\diamond \varphi$ for $\top \mathcal{U} \varphi$ and $\square \varphi$ for $\neg \diamond \neg \varphi$. Moreover, we define $\square^* \varphi$ as $(\varphi \wedge \square \varphi)$ (φ is true now and always in the future), and $\diamond^* \varphi = \neg \square^* \neg \varphi$ (φ is true now or sometime in the future).

When $M = (S, \sigma, Q)$ is a TQCG, $u \in \mathbb{N}$, and φ is a $\mathcal{L}(TQCG)$ formula, the satisfaction relation $M, u \models_T \varphi$ is defined as follows (the cases for negation and disjunction are defined as usual):

$$M, u \models_T \varphi \text{ iff } Q(\sigma(u)) \models_Q \varphi, \text{ when } \varphi \in \mathcal{L}(QCG)$$

$$M, u \models_T \bigcirc \psi \text{ iff } M, u + 1 \models_T \psi$$

$$M, u \models_T \psi_1 \mathcal{U} \psi_2 \text{ iff there is some } i \text{ such that } M, u + i \models_T \psi_2 \text{ and for all } 0 < j < i \text{ } M, u + j \models_T \psi_1$$

For instance, the following formula of $\mathcal{L}(TQCG)$ means that eventually, agent 3 can always choose to satisfy agents 1 and 4 simultaneously: $\diamond \square \langle 3 \rangle (sat_1 \wedge sat_4)$.

We will henceforth use $\mathcal{L}(TQCG)$ to refer to both the language, and the logic we have defined over this language.

3.2 An Example

We illustrate the logic by a small example. We focus here on temporal properties of goal satisfaction, rather than on contrasting the power of different coalitions (i.e., on which coalitions are likely to form). The latter is discussed in detail in Section 4.

We model the following situation by a temporal qualitative coalitional game. Two agents 1 and 2 both need to use the same resource, say a web service, from time to time. Sometimes an agent needs *read* access, and sometimes it needs *write* access. The integrity of the web service is violated if at the same time either i) both read and write access are granted (inconsistent reads), ii) two write accesses are granted (inconsistent writes) or iii) no read access and no write access are granted (inefficiency).

Let $M = (S, \sigma, Q)$ be a TQCG where S is some infinite set of states, and σ and Q are such that the following holds for $Q(\sigma(k)) = \langle \mathcal{A}, \mathcal{G}, \mathcal{G}_1^{\sigma(k)}, \mathcal{G}_2^{\sigma(k)}, \mathcal{G}_{sys}^{\sigma(k)}, \mathcal{V}^{\sigma(k)} \rangle$ for any $k \geq 0$:

- $\mathcal{A} = \{1, 2, sys\}$. We model the agents as players 1 and 2, and the web service as player *sys* (“the system”).
- $\mathcal{G} = \{r, w_1, w_2, ok\}$. That each of these goals are achieved means that right now:

$$\begin{aligned} r &: \text{every client is granted read access} \\ w_i &: \text{agent } i \text{ is granted write access} \\ ok &: \text{the integrity of the system is not violated} \end{aligned}$$

$$\bullet \mathcal{G}_1^{\sigma(k)} = \begin{cases} \{w_1\} & \text{if } k \bmod 5 = 0 \\ \{r, w_1\} & \text{otherwise} \end{cases}$$

Agent 1 needs to have write access at every fifth point in time. At any other point in time, it is happy as long as it is not left idle, i.e., if it has either read or write access.

$$\bullet \mathcal{G}_2^{\sigma(k)} = \begin{cases} \{w_2\} & \text{if } k \bmod 3 = 0 \\ \{r, w_2\} & \text{otherwise} \end{cases}$$

Agent 2's goals are similar to agent 1's, except that it needs write access at every third instead of fifth time point.

- $\mathcal{G}_{sys}^{\sigma(k)} = \{ok\}$. The system is satisfied if the integrity is not violated. Note that $\mathcal{G}_{sys}^{\sigma(k)}$ does not depend on k ; the system's goal does not vary over time.

- $\mathcal{V}^{\sigma(k)}(sys) = \left\{ \begin{array}{l} \emptyset, \{w_1, ok\}, \{w_2, ok\}, \{r, ok\}, \\ \{w_1, w_2\}, \{w_1, r\}, \{w_2, r\}, \{w_1, w_2, r\} \end{array} \right\}$.

The web service can satisfy certain sets of goals. These sets does not necessarily include the goal that the integrity is not violated. We have implicitly defined what is the desired behaviour of the system: each choice involving *ok* implements a choice in which the integrity invariant is not violated. Note that the choices available to the system do not vary over time. For reasons of space, in this example we are not bothered about $\mathcal{V}^s(C)$ when C is a coalition different from $\{sys\}$.

The following properties hold in M , 1.

1. $\Box \langle sys \rangle sat_{sys}$. The system can maintain integrity.
2. $\Box (\langle sys \rangle sat_1 \wedge \langle sys \rangle sat_2)$. Agent 1 can always be satisfied by the system, and the same for agent 2.
3. $\Box \langle sys \rangle (sat_1 \wedge sat_2)$. Agents 1 and 2 can always be simultaneously satisfied by the system.
4. $\Diamond \neg \langle sys \rangle (sat_1 \wedge sat_2 \wedge sat_{sys})$. The system cannot always satisfy agents 1 and 2 simultaneously without violating the integrity of the system.
5. $\Box \langle sys \rangle \neg sat_1$. The system can keep agent 1 unsatisfied forever.
6. $\Box \Diamond \langle sys \rangle (\neg sat_1 \wedge \neg sat_2 \wedge sat_{sys})$. It is infinitely often the case that the system can make agents 1 and 2 unsatisfied at the same time without violating integrity (this happens at multiples of fifteen).
7. $\langle sys \rangle (\neg sat_1 \wedge \neg sat_2 \wedge sat_{sys}) \mathcal{U} \neg \langle sys \rangle (sat_1 \wedge sat_2 \wedge sat_{sys})$. At some point in the future (i.e., $u = 15$), the system is unable to jointly satisfy agents 1 and 2 without violating integrity. Up until that time, *sys* is always able to make agents 1 and 2 jointly *unsatisfied* (note that we evaluate the formula in M , 1).
8. $\psi \wedge \Box (\psi \rightarrow \overbrace{\Box (\neg \psi \wedge \dots \Box (\neg \psi \wedge \Box \psi) \dots)}{\text{"}\Box(\neg\psi\wedge\text{" 14 times}}))$ where $\psi = \langle sys \rangle (\neg sat_1 \wedge \neg sat_2 \wedge sat_{sys})$. The system can make agents 1 and 2 jointly unsatisfied without violating integrity at time points which are multiples of fifteen, and at no other time points.

As a final point, observe that from a logical point of view, the situations at time points 3 and 5 are indistinguishable:

$$Q(\sigma(3)) \equiv Q(\sigma(5))$$

This once again demonstrates that our logic abstracts away from *how* a coalition satisfies individuals: obviously, to satisfy agent 1 for instance, *sys* has to make different choices in $\sigma(3)$ from those in $\sigma(5)$.

3.3 Expressive Power of TQCGs

The notion of simulation for QCGs (Section 2.2) can be naturally lifted to the temporal case. When $M = (S, \sigma, Q)$ and $M' = (S', \sigma', Q')$ are TQCGs and $k \geq 0$, we define

$$\begin{aligned} M, k \equiv_T M', k &\Leftrightarrow Q(\sigma(k)) \equiv Q'(\sigma'(k)) \\ M \equiv_T M' &\Leftrightarrow \forall_{n \geq 0} M, n \equiv_T M', n \end{aligned}$$

The notion of elementary equivalence for TQCGs over the language $\mathcal{L}(TQCG)$ can be defined as follows. $M, k \equiv M', k$ iff, for every $\varphi \in \mathcal{L}(TQCG)$, $M, k \models_T \varphi$ iff $M', k \models_T \varphi$. $M \equiv M'$ iff $M, k \equiv M', k$ for every $k \geq 0$.

THEOREM 4. For all TQCGs M, M' : $M \equiv_T M' \Leftrightarrow M \equiv M'$

Note that in the temporal case, the fact that $M, k \equiv_T M', k$ is not sufficient for $M, k \equiv M', k$ to hold.

3.4 Satisfiability

The *satisfiability* problem for $\mathcal{L}(TQCG)$ is as follows: given a formula $\varphi \in \mathcal{L}(TQCG)$, does there exist a TQCG M and $u \in \mathbb{N}$ such that $M, u \models \varphi$?

THEOREM 5. The sat. probl. for $\mathcal{L}(TQCG)$ is PSPACE-complete.

PROOF. Membership of PSPACE follows from the fact that satisfiability for LTL+ K_n (the fusion of LTL and multi-modal K) is PSPACE-complete [5]. Any $\mathcal{L}(TQCG)$ formula is also a formula of LTL+ K_n , interpreting sat_i as Boolean variable. (The reverse is not the case, of course.) But the relationship is more than merely syntactic: for all $\varphi \in \mathcal{L}(TQCG)$:

φ is $\mathcal{L}(TQCG)$ -satisfiable iff φ is LTL+ K_n satisfiable

(Notice that we are here quantifying over $\mathcal{L}(TQCG)$, formulae, not LTL+ K_n formulae.) Given an LTL+ K_n interpretation that satisfies $\varphi \in \mathcal{L}(TQCG)$, it is straightforward to extract from this a TQCG that satisfies φ .

For PSPACE-hardness, we reduce LTL satisfiability [14]. First, let φ^\dagger denote the result of systematically replacing each Boolean variable p that occurs in LTL formula φ with a symbol sat_p . Next, we define a transformation τ , from LTL formulae to $\mathcal{L}(TQCG)$, as follows:

$$\tau(\varphi) = \begin{cases} [1](\varphi^\dagger) & \text{where } \varphi \text{ is propositional} \\ \# \tau(\psi) & \text{where } \varphi = \# \psi \text{ and } \# \in \{\neg, \Box\} \\ \tau(\psi) \# \tau(\chi) & \text{where } \varphi = \psi \# \chi \text{ and } \# \in \{\vee, \mathcal{U}\} \end{cases}$$

Finally, given an LTL formula φ , the $\mathcal{L}(TQCG)$ instance φ^τ we create is:

$$\varphi^\tau = (\langle 1 \rangle \top) \wedge (\Box \langle 1 \rangle \top) \wedge \tau(\varphi)$$

We claim that φ is LTL satisfiable iff φ^τ is $\mathcal{L}(TQCG)$ satisfiable; the proof is an easy induction. The key point is that the choice sets of agent 1 in any TQCG satisfying φ^τ define an appropriate valuation for propositional variables in a corresponding LTL interpretation satisfying φ , and *vice versa* (remember that $[1]\varphi$ iff φ holds for all of 1's choices). The first two conjuncts in the definition of φ^τ ensure that such a choice set always exists. \square

3.5 Axiomatisation for TQCGs

The system $\mathbf{K}(TQCG)$ over the language $\mathcal{L}(TQCG)$ is defined as follows, where φ, ψ are arbitrary $\mathcal{L}(TQCG)$ formulae, A, B are arbitrary $\mathcal{L}(QCG)$ formulae, α, β are arbitrary \mathcal{L}_c formulae and C an arbitrary coalition. For simplicity, we write \vdash_T instead of $\vdash_{\mathbf{K}(TQCG)}$ for derivability in $\mathbf{K}(TQCG)$.

$Prop^-$	If A is an ($\mathcal{L}(QCG)$) instance of a propositional tautology, then $\vdash_T A$
K^-	$\vdash_T [C](\alpha \rightarrow \beta) \rightarrow ([C]\alpha \rightarrow [C]\beta)$
MP^-	If $\vdash_{\mathbf{K}(QCG)} A$ and $\vdash_{\mathbf{K}(QCG)} A \rightarrow B$, then $\vdash_T B$
Nec^-	If α is an (\mathcal{L}_c) instance of a propositional tautology, then $\vdash_T [C]\alpha$
A1	$\vdash_T \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$
A2	$\vdash_T \Box\neg\varphi \leftrightarrow \neg\Box\varphi$
A3	$\vdash_T \Box(\varphi \rightarrow \psi) \rightarrow (\Box\varphi \rightarrow \Box\psi)$
A4	$\vdash_T \Box\varphi \rightarrow (\Box\varphi \wedge \Box\varphi)$
A5	$\vdash_T \Box(\varphi \rightarrow \Box\varphi) \rightarrow (\Box\varphi \rightarrow \Box\varphi)$
U1	$\vdash_T \varphi \mathcal{U} \psi \rightarrow \Diamond\psi$
U2	$\vdash_T \varphi \mathcal{U} \psi \leftrightarrow \Box\psi \vee (\Box\varphi \wedge \Box(\varphi \mathcal{U} \psi))$
$Prop$	If φ is an ($\mathcal{L}(TQCG)$) instance of a propositional tautology, then $\vdash_T \varphi$
MP	If $\vdash_T \varphi$ and $\vdash_T \varphi \rightarrow \psi$, then $\vdash_T \psi$
Nec	If $\vdash_T \varphi$ then $\vdash_T \Box\varphi$

Axioms $Prop^-$ and K^- and rules MP^- and Nec^- say that every $\mathbf{K}(QCG)$ -theorem is also a $\mathbf{K}(TQCG)$ -theorem. The subsystem consisting of axioms $A1-U2$ and rules $Prop-Nec$ is a version (with $\mathcal{L}(QCG)$ formulae in place of atomic propositions) of an axiomatisation of linear time logic proved to be sound and complete in [6].

THEOREM 6 (SOUNDNESS & COMPLETENESS). *For any $\varphi \in \mathcal{L}(TQCG)$: $\vdash_T \varphi \Leftrightarrow \models_T \varphi$*

PROOF. The logic $\mathbf{K}(TQCG)$ is what Finger and Gabbay [4] calls a *temporalisation* of $\mathbf{K}(QCG)$: the language of $\mathbf{K}(TQCG)$ has atomic $\mathbf{K}(QCG)$ formulae in place of atomic propositions; the semantic structures of $\mathbf{K}(TQCG)$ identifies a semantic structure for $\mathbf{K}(QCG)$ at each time point used to interpret $\mathbf{K}(QCG)$ formulae; and the axioms/rules of $\mathbf{K}(TQCG)$ are the axioms/rules of the temporal logic for temporal formulae in addition to axioms/rules of $\mathbf{K}(QCG)$ for $\mathbf{K}(QCG)$ formulae.

Finger and Gabbay show that the temporalisation of a sound and complete system is sound and complete. It should be noted that our definition of $\mathbf{K}(TQCG)$ differs from the definition of a temporalisation in [4] by the following. First, we do not have past-time operators in our language. The expressive power is nevertheless the same [6]. Second, we use a slightly different temporal axiomatisation. Neither of these differences change the soundness and completeness proof in [4] in any significant degree. The theorem thus follows immediately from Theorem 3. \square

4. CHARACTERIZING TQCGS

In this section, we investigate the axiomatic characterisation of various classes of TQCG. As usual, in saying that a formula scheme φ characterises a property P of models, we mean that φ is valid in a model M iff M has property P ; if only the right-to-left part of this biconditional holds, then we say property P implies φ . Also note that for an $\mathcal{L}(TQCG)$ formula φ , to say that φ is valid in a class of models, is the same as saying that $\Box^* \varphi$ is valid in that class.

Basic Correspondences

Let $h^s(C)$ denote the set of all agents that could possibly be satisfied (not necessarily jointly) by coalition C in state s :

$$h^s(C) = \{i : i \in \mathcal{A} \ \& \ \exists G \in \mathcal{V}^s(C), \mathcal{G}_i^s \cap G \neq \emptyset\}$$

The “ h ” here is for “happiness”: we think of $h^s(C)$ as all the agents that C could possibly make happy in s . Thus the semantic property $i \in h^s(C)$ is a counterpart to the syntactic expression $\langle C \rangle sat_i$.

The first property on models that we consider is the *persistence of happiness (PH)*: if coalition C can make i happy in a state s , they can make i happy in the state immediately following s .

$$\forall u \in \mathbb{N}, (i \in h^{\sigma(u)}(C)) \rightarrow (i \in h^{\sigma(u+1)}(C)) \quad (PH)$$

We have the following characterisation.

LEMMA 2. $\langle C \rangle sat_i \rightarrow \bigcirc \langle C \rangle sat_i$ characterises PH.

In the same way, we can characterise the persistence of *unhappiness*: property *PU* says that if C cannot make i happy in a state s , then they cannot make i happy in the state t that immediately follows s .

$$\forall u \in \mathbb{N}, (i \notin h^{\sigma(u)}(C)) \rightarrow (i \notin h^{\sigma(u+1)}(C)) \quad (PU)$$

LEMMA 3. $\neg \langle C \rangle sat_i \rightarrow \bigcirc \neg \langle C \rangle sat_i$ characterises PU.

Now consider the following two constraints. The first, *EH*, says that eventually, C will be able to make i happy.

$$\exists u \in \mathbb{N}, (i \in h^{\sigma(u)}(C)) \quad (EH)$$

Notice that in the terminology of reactive systems, this is a *fairness* or *response* property [8, p.288]: it implies that something (i being made happy) can happen infinitely often. (Of course, the fact that C can make i happy infinitely often does not mean they will do so.)

LEMMA 4. $\Diamond^* \langle C \rangle sat_i$ characterises EH.

The obvious counterpart to *EH* is of course the property *EU*, which states that, eventually, C will be unable to satisfy i .

$$\exists u \in \mathbb{N}, (i \notin h^{\sigma(u)}(C)) \quad (EU)$$

LEMMA 5. $\Diamond^* \neg \langle C \rangle sat_i$ characterises EU.

Combining these properties, we get the following.

LEMMA 6. *PH and EH together imply $\Diamond^* \Box^* \langle C \rangle sat_i$, while properties PU and EU together imply $\Diamond^* \Box^* \neg \langle C \rangle sat_i$.*

Finally, we consider *safety* properties. The constraint *AH* says that C can always make i happy, while the constraint *AU* says that C can never make i happy.

$$\forall s \in S, (i \in h^s(C)) \quad (AH) \quad \forall s \in S, (i \notin h^s(C)) \quad (AU)$$

The characterizations are as follows. (Note that there are some obvious implications between these and other properties that we do not list explicitly – e.g., *AH* implies both *EH* and *PH*.)

LEMMA 7. $\langle C \rangle sat_i$ characterises AH, and $\neg \langle C \rangle sat_i$ char. AU.

Basic Properties of Choice Sets

Three obvious constraints that we might consider relate to whether or not a particular coalition C has any “real” choices. The first, *ECS*, says that C never has any choices.

$$\forall s \in S, \mathcal{V}^s(C) = \emptyset \quad (ECS)$$

The second says that C always has a meaningful choice.

$$\forall s \in S, \exists G \in \mathcal{V}^s(C), G \neq \emptyset \quad (NECS)$$

The third says that C can choose everything .

$$\forall s \in S, \mathcal{G} \in \mathcal{V}^s(C) \quad (CCS)$$

LEMMA 8. *Any model that satisfies ECS also satisfies AU, and so ECS implies $\neg \langle C \rangle sat_i$, while any model that satisfies CCS also satisfies AH, and so CCS implies $\langle C \rangle sat_i$.*

Note that *NECS* alone does not have any characterization: however, when combined with other properties, below, we will see that it has a role.

Static Goal Sets and Choices

Another two simple properties are that the goal sets for each agent and the choice sets for each coalition are guaranteed to remain unchanged. We get the following two constraints, stating that agent i 's goal sets are static (constraint *SGS*) and that coalition C 's choices remain static (*SC*).

$$\forall s, s' \in S, (\mathcal{G}_i^s = \mathcal{G}_i^{s'}) \quad (SGS)$$

$$\forall s, s' \in S, (\mathcal{V}^s(C) = \mathcal{V}^{s'}(C)) \quad (SC)$$

Taken separately, there does not seem too much we can say about static goal sets and static choice sets. However, taken together, we get the following.

LEMMA 9. Any model satisfying both *SGS* and *SC* also satisfies *PH* and *PU*, and as a consequence, *SGS* and *SC* together imply $\langle C \rangle sat_i \leftrightarrow \bigcirc \langle C \rangle sat_i$.

Note that we do not immediately derive a characterisation here. It is perfectly well possible that $\langle C \rangle sat_i \leftrightarrow \bigcirc \langle C \rangle sat_i$ is true in a model M not just because all agents' goals and all coalitions' choices stay fixed, but because there is an intricate interplay going on between for instance an agent changing some of his goals, while at the same time, the coalition C 'synchronously' changing its options. Note that in our example of Section 3.2 for instance, both (*SGS*) and (*SC*) are true for $C = \{sys\}$ and $i = sys$, so that, indeed, $\langle \{sys\} \rangle sat_{sys} \leftrightarrow \bigcirc \langle \{sys\} \rangle sat_{sys}$. On the other hand, taking $C = \{sys\}$ and $i = 1$, we don't have (*SGS*) and (*SC*), although we still have $\langle \{sys\} \rangle sat_1 \leftrightarrow \bigcirc \langle \{sys\} \rangle sat_1$.

Dynamic Goal Sets

There are several properties we can investigate with respect to goal sets. First, suppose that agent i 's goal set is guaranteed to *monotonically decrease* over time. Roughly, this condition means that every agent is guaranteed to get *no easier* to satisfy over time. Formally, this condition on a model M is defined by the following property.

$$\forall u \in \mathbb{N} (\mathcal{G}_i^{\sigma(u+1)} \subseteq \mathcal{G}_i^{\sigma(u)}) \quad (MDGS)$$

LEMMA 10. Any model satisfying *SC* and *MDGS* will satisfy *PU*, and hence *SC* and *MDGS* together imply $\neg \langle C \rangle sat_i \rightarrow \bigcirc \neg \langle C \rangle sat_i$.

Suppose we this condition is *strict*, so that an agent i is guaranteed to get strictly harder to satisfy over time. This condition is defined by the following further constraint, in addition to *MDGS*.

$$\forall u \in \mathbb{N} \quad (\mathcal{G}_i^{\sigma(u)} = \emptyset) \vee (\exists v \in \mathbb{N} : (v > u) \wedge (\mathcal{G}_i^{\sigma(v)} \subset \mathcal{G}_i^{\sigma(u)})) \quad (SMDGS)$$

We get the following.

LEMMA 11. Any model satisfying *SC*, *MDGS*, and *SMDGS* will also satisfy *PU* and *EU*, and so *SC*, *MDGS*, and *SMDGS* together imply $\diamond^* \square^* \neg \langle C \rangle sat_i$.

Now suppose agent i has monotonically *increasing* goal sets: that is, agent i gets *no harder* to satisfy over time.

$$\forall u \in \mathbb{N}, (\mathcal{G}_i^{\sigma(u)} \subseteq \mathcal{G}_i^{\sigma(u+1)}) \quad (MIGS)$$

LEMMA 12. Any model satisfying both *SC* and *MIGS* will satisfy constraint *PH*, and hence *SC* and *MIGS* together imply $\langle C \rangle sat_i \rightarrow \bigcirc \langle C \rangle sat_i$.

The associated strictness constraint is as follows.

$$\forall u \in \mathbb{N} \quad (\mathcal{G}_i^{\sigma(u)} = \mathcal{G}) \vee (\exists v \in \mathbb{N} : (v > u) \wedge (\mathcal{G}_i^{\sigma(u)} \subset \mathcal{G}_i^{\sigma(v)})) \quad (SMIGS)$$

We might expect that *SC*, *MIGS*, and *SMIGS* together imply the validity of the formula scheme $\diamond^* \square^* \langle C \rangle sat_i$, but this is not the case. A counter example is given by a model that satisfies the empty choice set property (*ECS*) for coalition C , as described above. If we add the constraint that the choices for C are non-empty (*NECS*), however, then we get the following.

LEMMA 13. Any model that satisfies *NECS*, *SC*, *MIGS*, and *SMIGS* also satisfies *PH* and *EH*, and hence the following formula scheme will be valid in any model satisfying *NECS*, *SC*, *MIGS*, and *SMIGS*: $\diamond^* \square^* \langle C \rangle sat_i$.

Dynamic Choices

We can also consider the ways in which the choices available to coalitions may change over time. Analogously to *MIGS* and *MDGS*, we can define properties *MICS* and *MDCS*, which say that the sets of choices available to coalition C monotonically increase and decrease respectively.

$$\forall u \in \mathbb{N}, (\mathcal{V}^{\sigma(u)}(C) \subseteq \mathcal{V}^{\sigma(u+1)}(C)) \quad (MICS)$$

$$\forall u \in \mathbb{N}, (\mathcal{V}^{\sigma(u+1)}(C) \subseteq \mathcal{V}^{\sigma(u)}(C)) \quad (MDCS)$$

Notice that taken together, these two conditions imply static choice sets (*SC*). Alone, the properties do not have any characterisation, but axioms emerge when we make assumptions about goal sets.

LEMMA 14. (1) Any model satisfying *MICS* and *SGS* will satisfy constraint *PH*, and hence *MICS* and *SGS* together imply $\langle C \rangle sat_i \rightarrow \bigcirc \langle C \rangle sat_i$.

(2) Any model satisfying *MDCS* and *SGS* will satisfy constraint *PU*, and hence *MDCS* and *SGS* together imply $\neg \langle C \rangle sat_i \rightarrow \bigcirc \neg \langle C \rangle sat_i$.

The associated strictness condition for increasing choice sets is:

$$\forall u \in \mathbb{N}, \forall G_1 \in \mathcal{V}^{\sigma(u)}(C) (G_1 = \mathcal{G}) \vee (\exists v \in \mathbb{N}, \exists G_2 \in \mathcal{V}^{\sigma(v)}(C), (v > u) \wedge (G_1 \subset G_2)) \quad (SMICS)$$

LEMMA 15. Any model satisfying *MICS*, *SGS*, and *SMICS* or *MDCS*, *MIGS*, and *SMICS* will also satisfy constraints *PH* and *EH*, and hence *MICS*, *SGS*, and *SMICS* together imply $\diamond^* \square^* \langle C \rangle sat_i$.

We omit here the analysis for the case of monotonically *decreasing* choice sets.

Solution Concepts

In [15], a range of different solution concepts were defined for QCGs. It should be clear that many of the solution concepts of [15] can be characterised via formulae of $\mathcal{L}(QCG)$. For example, a basic solution concept is that of a *successful* coalition – one that has a choice available such that this choice satisfies all its members [15, p.47]. We can characterise this via a predicate $succ(C)$, as follows.

$$succ(C) \equiv \langle C \rangle \left(\bigwedge_{i \in C} sat_i \right)$$

Similarly, the notion of a minimal coalition (one such that no subset of the coalition is successful [15, p.51]) may be captured as follows.

$$min(C) \equiv \bigwedge_{C' \subset C} \neg succ(C')$$

Thus the core of a coalition being non-empty [15, p.54] may be captured as follows:

$$cne(C) \equiv (succ(C) \wedge min(C))$$

The idea of agent i being a veto player for agent j [15, p.57] is defined by:

$$veto(i, j) \equiv \bigwedge_{C \subseteq \mathcal{A}} (\langle C \rangle sat_j \rightarrow \neg \langle C \setminus \{i\} \rangle sat_j)$$

And finally, the idea of a coalition being mutually dependent [15, p.58] is captured as follows:

$$md(C) \equiv \bigwedge_{i \neq j \in C} veto(i, j)$$

How might these concepts be extended into the temporal dimension of TQCGs and $\mathcal{L}(TQCG)$? It should first be clear that each concept has four different temporal versions, corresponding to prefixing the formula characterising it with one of the following four, increasingly powerful temporal operators:

$$\diamond \quad \square \diamond \quad \diamond \square \quad \square$$

Thus, for example, $\square \diamond succ(C)$ means that coalition C are successful *infinitely often* – no matter which time point we pick, there will be a subsequent time point at which C are successful. (Using the terminology of reactive systems [8], we might then say that C are hence *fairly successful*.) Similarly, a *temporally strong* form of coalitional stability is captured by the formula $\square cne(A)$: if this formula is satisfied in a TQCG, then, it can be argued, the only coalition that will ever form is the grand coalition.

It is potentially more interesting, however, to study a richer interplay between temporal and QCG dimensions. For example, from agent *is* point of view, perhaps the only really interesting issue is whether at every time point there is some stable coalition, containing this agent.

$$tstable(i) \equiv \square \bigvee_{C \subseteq A: i \in C} cne(C)$$

From the point of view of a coalition C , which seeks to form, the notion of a *stable government* seems relevant: a stable government is a coalition that can always satisfy its “electorate”.

$$sg(C) \equiv \square \langle C \rangle \left(\bigwedge_{i \in A} sat_i \right)$$

This can of course be strengthened, requiring C to in addition be an internally stable coalition.

$$sg'(C) \equiv \square (cne(C) \wedge \langle C \rangle \left(\bigwedge_{i \in A} sat_i \right))$$

With respect to mutual dependence, one possibility, captured by the formula $\square md(C)$, is that a coalition is *always* mutually dependent. However, we can capture a weaker type of mutual dependence as follows:

$$wmd(C) \equiv \bigwedge_{i \neq j \in C} \diamond veto(i, j)$$

We draw two conclusions. The first is that the language $\mathcal{L}(TQCG)$ is well suited to capturing such solution concepts: it makes it possible to express elegantly concepts that would be difficult to understand were they expressed at the semantic level. The second is that extending QCGs into the temporal dimension adds an entirely new level of richness to their structure, which as these examples suggest, demands further study.

5. CONCLUSION

Qualitative Game Structures were introduced in [15] as a model to deal with one of the driving questions in cooperation: ‘Which coalition should I join?’. The design of such games is motivated by a principle of economy: rather than associating a utility to every choice, the emphasis is on *satisfaction* of agents, which is triggered or not by *choices* of coalitions.

The logical analysis of such games, as underlines the idea that we have a basic and simple notion of coalitional games with QCGs: a natural language for it gives rise to an axiomatisation almost identical to the simplest modal logic **K**. We established several technical results for this language, and were then able to lift them to the era of temporal QCGs.

The possible directions for further research are multiple. First of all, the properties of TQCGs that we characterised in Section 4 are only the most straightforward. Even in static games, there are interesting conditions to be investigated (see the *monotonicity* property mentioned in Section 2, for example). Second, our way of temporalising QCGs also only reflects a simple case. It would be interesting to add temporal structure to the games themselves, and reason about what agents can achieve over time, by applying suitable *strategies*, rather than making ‘one-shot choices’. In addition, *finite horizon* versions of TQCGs might also be worth investigating: for example, if an agent is only concerned about being satisfied *once*, then it might be prepared to join a coalition that does not satisfy it throughout a game, as long as, in the final state of the game, the coalition *does* satisfy it. Such strategising is not possible or appropriate in infinite horizon games.

Acknowledgements

The research reported in this paper was carried out when the first author was visiting the Department of Computer Science, University of Liverpool, supported by grant 166525/V30 from the Norwegian Research Council.

6. REFERENCES

- [1] R. Alur, T.A. Henzinger, O. Kupferman, and M.Y. Vardi. Alternating refinement relations. In *Proc. of CONCUR'98*, vol. 1466 of LNCS, pp. 163–178. Springer-Verlag, 1998.
- [2] R. Axelrod. *The Evolution of Cooperation*. Basic Books: New York, 1984.
- [3] K. Binmore. *Fun and Games: A Text on Game Theory*. D. C. Heath and Company: Lexington, MA, 1992.
- [4] M. Finger and D. M. Gabbay. Adding a temporal dimension to a logic system. *Journal of Logic, Language, and Information*, 1:203–233, 1992.
- [5] D. M. Gabbay, A. Kurucz, F. Wolter, and M. Zakharyashev. *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier, 2003.
- [6] D. Gabbay, A. Pnueli, S. Shelah, and J. Stavi. On the temporal analysis of fairness. In *POPL '80*, pp. 163–173, 1980. ACM Press.
- [7] M. Klusch and A. Gerber. Dynamic coalition formation among rational agents. *IEEE Int. Syst.*, 17(3):42–47, 2002.
- [8] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer-Verl.: Berlin, Germany, 1992.
- [9] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press: Cambridge, MA, 1994.
- [10] M. Pauly. *Axiomatising Judgement Aggregation Procedures in a Minimal Logical Language*. Unpublished paper.
- [11] T. Sandholm. Distributed rational decision making. In G. Weiß, editor, *Multiagent Systems*, pages 201–258. The MIT Press: Cambridge, MA, 1999.
- [12] T. Sandholm, K. Larson, M. Andersson, O. Shehory, and F. Tohmé. Coalition structure generation with worst case guarantees. *Artificial Intelligence*, 111(1–2):209–238, 1999.
- [13] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1–2):165–200, 1998.
- [14] A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. *Journal of the ACM*, 32(3):733–749, 1985.
- [15] M. Wooldridge and P. E. Dunne. On the computational complexity of qualitative coalitional games. *Artificial Intelligence*, 158(1):27–73, 2004.