On the Computational Complexity of Qualitative Coalitional Games

Michael Wooldridge and Paul E. Dunne *

Dept of Computer Science, University of Liverpool, Liverpool L69 7ZF, United Kingdom

Abstract

We study coalitional games in which agents are each assumed to have a goal to be achieved, and where the characteristic property of a coalition is a set of *choices*, with each choice denoting a set of goals that would be achieved if the choice was made. Such *qualitative coalitional games* (QCGs) are a natural tool for modelling goal-oriented multiagent systems. After introducing and formally defining QCGs, we systematically formulate fourteen natural decision problems associated with them, and determine the computational complexity of these problems. For example, we formulate a notion of coalitional stability inspired by that of the core from conventional coalitional games, and prove that the problem of showing that the core of a QCG is non-empty is D_1^p-complete. (As an aside, we present what we believe is the first "natural" problem that is proven to be complete for D_2^p.) We conclude by discussing the relationship of our work to other research on coalitional reasoning in multiagent systems, and present some avenues for future research.

Key words: Multiagent systems, Coalitional Games, Computational Complexity

1 Introduction

The question *which coalition should I join* is a fundamental problem facing both natural and artificially rational agents in any cooperative encounter [30]. The field of *cooperative game theory* has developed a number of solution concepts which attempt to provide plausible answers to this question [22, pp.255–312]. Perhaps the best known such solution concept is that of the *core* [22, pp.257–274].

Concepts such as the core are usually presented in terms of *coalitional games*, which may be formulated in a number of different ways, depending on the characteristics of the domain they are intended to model. Perhaps the best-known and

* Corresponding author.
Email address: ped@csc.liv.ac.uk (Paul E. Dunne).

most widely studied model of coalitional games is that of the *coalitional game with transferable payoff*, which is simply a structure of the form $\langle Ag, \nu \rangle$, where Ag is a set of agents and ν is a *characteristic function*, $\nu : 2^{Ag} \to \mathbb{R}$, which assigns to every possible coalition a numeric value, the idea being that this value can then be distributed between members of the coalition [22, p.257]. Despite their apparent simplicity, conventional coalitional games have proved to be of great value in understanding the nature of coalitions and coalition formation, and have been successfully used in a number of studies of coalition formation in multiagent systems [32–35,29].

Such games, expressed in terms of real-valued characteristic functions, are inherently *quantitative* in nature. That is, they require us explicitly to assign a numeric value to every coalition. While for many domains this is both feasible and entirely appropriate, for many others, it is not. While some multiagent domains are naturally and easily captured in terms of utilities, others are more appropriately characterised in terms of *goals* that agents wish to be achieved. With such an approach, an agent will regard any given outcome as either "good" (its goal is satisfied), or "bad" (its goal is not satisfied). Of course, utilitarian approaches are obviously richer in terms of the domains and problems they can capture; but this does not imply that their use should always be preferred over a goal-oriented one. Such *qualitative* coalitional game models are arguably less well-known and less widely studied in the game theory and multiagent systems literature, but nevertheless, there is a body of work on this subject. For example, *effectivity functions* provide one such qualitative model of coalitional games [1]. An effectivity function E is typically defined with respect to a set of agents Ag, (as with conventional coalitional games, above) and a set of *outcome states*, S, which intuitively capture the possible outcomes of the game. For every coalition $C \subseteq Ag$, an effectivity function E then determines a set of subsets of S, with the usual interpretation being that if $S' \in E(C)$, then the coalition C is *effective* for S', in the sense that there exists a collective strategy for C such that, if C followed this strategy, the outcome of the game would fall in the set S' (although the members of C could not force the outcome into any *specific* member of S'). Although effectivity functions are beginning to find their way into multiagent systems research (notably by way of the semantics to Pauly's *Coalition Logic* [26]), to date, there has been relatively little work on such models within this community.

In short, then, our aim in the present paper is to formulate and investigate some computational problems associated with "qualitative" coalitional games. Specifically, we study games in which agents are not assigned utility values over outcomes, but rather are assumed to have goals that they desire to be achieved. An agent in such a game will be *satisfied* with some outcome if its goals are achieved in this outcome. Coalitions of agents have different choices available to them, where each choice will achieve some subset of the overall set of possible goals. Despite their apparently simple structure, such *qualitative coalitional games* (QCGs) allow us to formulate and investigate a wide range of natural decision problems in coalitional reasoning. The key contribution of this paper is a comprehensive study of the complexity of these decision problems.

We begin, in the following section, with a summary of the relevant concepts from the theory of computational complexity that are required for the remainder of the paper. We then formally introduce qualitative coalitional games, discuss the issue of how such structures can be *succinctly represented*, and proceed to formulate fourteen natural decision problems associated with QCGs. For each of these, we begin by motivating the problem, then give a precise definition and examples to illustrate it, and finally, classify its computational complexity. As a flavour of the kind of problem we investigate, we define a notion of coalitional stability closely corresponding to that of the core in conventional coalitional games, and we prove that this problem is complete for D^p (the "difference class" introduced in [24]). (As an aside, we show that one of our problems is complete for D_2^p, and to the best of our knowledge, this is the first time a decision problem has been found to be complete for this class that was not specifically constructed for this purpose.) We conclude by discussing related work in the game theory and multiagent systems literature, and by pointing to some avenues for future research.

2 Concepts from Computational Complexity

Although we provide a summary of the main concepts and definitions from complexity theory that we use, we emphasise that a detailed presentation is beyond the scope of this paper. We refer the reader to [13,16,23] for details. We begin with some general comments on notation.

Notational Conventions

We use the symbols $\top$ and $\bot$ as the Boolean constants for truth and falsity, respectively. In general, upper case Greek letters – Φ, Ψ, etc – are used as meta-language variables ranging over formulae of propositional logic. In addition to the standard Boolean operations of conjunction ($\wedge$), disjunction ($\vee$), implication ($\Rightarrow$), and negation ($\neg$), some constructions will make of the binary *exclusive-or* function, which we denote by $\oplus$. For a propositional formula $\Phi(x_1, \ldots, x_n)$ defined over the variables $X_n = \langle x_1, \ldots, x_n \rangle$, given $Z \subseteq X_n$, we denote by $\Phi[Z]$ the result of evaluating Φ under the instantiation $x_i = \top$ if $x_i \in Z$, and $x_i = \bot$ if $x_i \notin Z$. Thus, $\Phi[Z]$ is equivalent to the value of $\Phi(\zeta_1, \ldots, \zeta_n)$ where the tuple $\zeta = \langle \zeta_1, \ldots, \zeta_n \rangle$ describes the characteristic vector from $\{\top, \bot\}^n$ for Z with respect to X_n. For example, given the propositional formula $\Psi(x_1, x_2, x_3) = x_1 \wedge (x_2 \vee x_3)$, the expression $\Psi[x_2]$ evaluates to $\bot \wedge (\top \vee \bot)$, which in turn evaluates to $\bot$.

Finally, note that where no ambiguity arises, we frequently omit explicit indication of conjunction, ($\wedge$), writing $\varphi\psi$ rather than $\varphi \wedge \psi$.

The Polynomial Hierarchy

We start from the complexity classes P (of languages/problems that may be recognised/solved in deterministic polynomial time), and NP (of languages/problems that may be recognised/solved in non-deterministic polynomial time). If $\mathcal{C}$ and $\mathcal{C}'$ are complexity classes, then we denote by $\mathcal{C}^{\mathcal{C}'}$ the class of languages/problems that are in $\mathcal{C}$ assuming the availability of an oracle for languages/problems in $\mathcal{C}'$ [23, pp.415–417]. Thus, for example, $\mathrm{NP}^{\mathrm{NP}}$ denotes the class of languages/problems that may be recognised/solved in non-deterministic polynomial time, assuming the presence of an oracle for languages/problems in NP. A language that is complete for $\mathrm{NP}^{\mathrm{NP}}$ would thus be NP-complete even if we had "unit cost" answers to NP-complete problems (such as SAT – propositional logic satisfiability). We define the *polynomial hierarchy* with reference to these concepts [23, pp.423–429].

Formally, we begin by defining

$$\Delta_0^p = \Sigma_0^p = \Pi_0^p = \mathrm{P}.$$

Thus Δ_0^p, Σ_0^p, and Π_0^p all denote the class of languages/problems that may be recognised/solved in deterministic polynomial time. We then inductively define the remaining tiers of the hierarchy, as follows:

$$\Delta_{u+1}^p = \mathrm{P}^{\Sigma_u^p} \qquad \Sigma_{u+1}^p = \mathrm{NP}^{\Sigma_u^p} \qquad \Pi_{u+1}^p = \text{co-}\Sigma_{u+1}^p$$

Thus Δ_1^p is in fact the same as Δ_0^p, while Σ_1^p is the class NP, Π_1^p is the class co-NP, $\Sigma_2^p = \mathrm{NP}^{\mathrm{NP}}$, and $\Pi_2^p = \text{co-}\mathrm{NP}^{\mathrm{NP}}$. Similarly, a problem that is in the class Δ_2^p is one that could be solved by a deterministic polynomial time algorithm, assuming the availability of an NP oracle.

We employ as "canonical" complete problems for Σ_k^p (respectively Π_k^p) the decision problems SAT [23, p.77] (respectively UNSAT) (for $k = 1$), and QSAT_2^Σ (respectively QSAT_2^Π) (for $k = 2$) [23, p.428]. An instance of the latter is a propositional formula, $\Phi(X_n, Y_n)$ defined over two disjoint sets of n propositional variables $X_n = \langle x_1, \ldots, x_n\rangle$ and $Y_n = \langle y_1, \ldots, y_n\rangle$. For QSAT_2^Σ, the instance $\Phi(X_n, Y_n)$ is accepted if there is an instantiation, α_X of X_n, under which every instantiation, β_Y of Y_n, satisfies $\Phi(X_n, Y_n)$, i.e., $\exists\alpha_X\forall\beta_Y, \Phi(\alpha_X, \beta_Y) = \top$. Similarly, $\Phi(X_n, Y_n)$ is accepted as an instance of QSAT_2^Π if $\forall\alpha_X\exists\beta_Y, \Phi(\alpha_X, \beta_Y) = \top$. That these problems are complete for Σ_2^p (respectively Π_2^p) was proved by Wrathall [38].

We also deal with the "difference" class, D^p, introduced in [24], together with D_k^p — a natural generalisation of this class. See [23, p.412] and [16, p.93] for overviews of the class D^p and its relatives. Formally, a language L is in the class D^p if there are languages $L_1 \in \mathrm{NP}$ and $L_2 \in \text{co-NP}$ such that $L = L_1 \cap L_2$. More generally, $L \in \mathrm{D}_k^p$ if there are languages $L_1 \in \Sigma_k^p$ and $L_2 \in \Pi_k^p$ for which $L = L_1 \cap L_2$. Note that, as pointed out in [23, pp.415–416], the class D_1^p can be understood as the class of languages that may be recognised by oracle machines that make two queries to an NP oracle, where we want the answer to the first query to be "yes", and the second to be "no". It follows that $\mathrm{D}_1^p \subseteq \Delta_2^p$. In fact, we can be more precise, and observe that $\mathrm{D}_1^p \subseteq \Delta_2^{p[2]}$, and similarly $\mathrm{D}_k^p \subseteq \Delta_{k+1}^{p[2]}$ for the difference classes D_k^p, where $k > 1$. Finally, also note that $\mathrm{NP} \cup \text{co-NP} \subseteq \mathrm{D}_1^p$ (and similarly for $k > 1$). The relationship of the classes D_k^p to the polynomial hierarchy (insofar as this relationship is currently understood) is illustrated in Figure 1.

While a number of natural complete problems have been identified for D_1^p, the corresponding classes at higher levels have been largely ignored [1] . This is, perhaps, unsurprising given the comparative lack of (natural) complete problems for Σ_k^p and Π_k^p for $k \geq 2$. The decision problem SAT-UNSAT is, perhaps the paradigm example of a D^p-complete problem, and the following is an obvious generalisation of this.

$\underline{\mathrm{QSAT}_k^{\Sigma}\text{-}\mathrm{QSAT}_k^{\Pi}}$:
Instance: A pair

$$\langle \Phi_1(X_1, X_2, \ldots, X_k), \Phi_2(X_1, X_2, \ldots, X_k) \rangle$$

of propositional formulae, each defined over k disjoint sets of n propositional variables.
Answer: "Yes" if Φ_1 is a positive instance of QSAT_k^{Σ} and Φ_2 is a positive instance of QSAT_k^{Π}, i.e.,

$$\exists \alpha_1 \forall \alpha_2 \cdots \mathrm{Q} \alpha_k \Phi_1(\alpha_1, \ldots, \alpha_k) = \top$$

and

$$\forall \alpha_1 \exists \alpha_2 \cdots \bar{\mathrm{Q}} \alpha_k \Phi_2(\alpha_1, \ldots, \alpha_k) = \top$$

where the quantifier Q (respectively, $\bar{\mathrm{Q}}$) is $\exists$ (respectively, $\forall$) if k is odd and $\forall$ (respectively, $\exists$) otherwise.

[1] That is, in terms of the study of specific languages. The class D_1^p, however, gives rise to a range of classes – the Boolean Hierarchy $\mathrm{BH}(k)$ contained within Δ_2^p – of considerable importance in work in structural complexity, c.f. [17]. The analogue $\mathrm{BH}_3(k)$ a subset of Δ_3^p – which contains D_2^p – plays a significant role in [6].

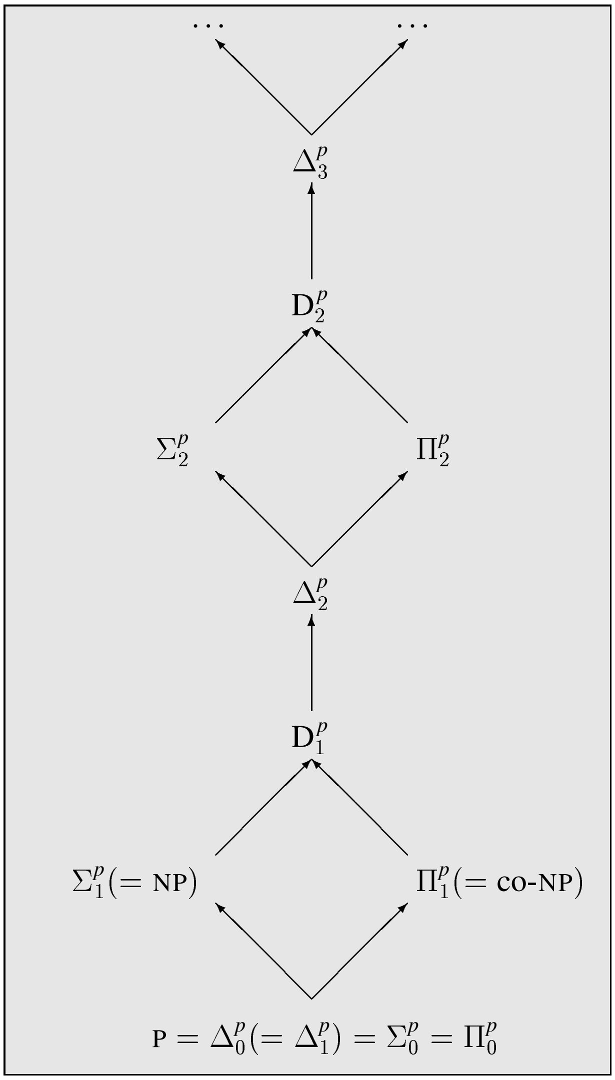

Fig. 1. The relationship of the polynomial hierarchy to the difference classes D_k^p, insofar as this relationship is understood at the time of writing. An arrow from class $\mathcal{C}$ to $\mathcal{C}'$ indicates that $\mathcal{C}'$ includes $\mathcal{C}$, i.e., $\mathcal{C} \subseteq \mathcal{C}'$; transitive inclusion is not illustrated.

For $k = 1$, we denote this problem by its better known name, that is, SAT-UNSAT [23, p.413]. Now, the following result may be trivially established by adapting the argument presented for $k = 1$ in [24].

Fact 1 *For all $k \geq 1$, the problem* QSAT_k^Σ-QSAT_k^Π *is D_k^p-complete.*

3 Qualitative Coalitional Games

We are now ready to introduce the formal framework of QCGs. First, the systems we study contain a (non-empty, finite) set $Ag = \{1, \ldots, n\}$ of *agents*. Each agent $i \in Ag$ is assumed to have associated with it a (finite) set G_i of *goals*, drawn from a set of overall possible goals G. The intended interpretation is that the members of 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 G_i were 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 G_i: it will be *satisfied* if *at least one* member of 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 G_i as possible.

A *coalition*, typically denoted by C, is simply a set of agents, i.e., a subset of Ag. The *grand coalition* is the set of all agents, Ag. 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

$$V : 2^{Ag} \to 2^{2^G}.$$

Thus, in saying that $G' \in V(C)$ for some coalition $C \subseteq Ag$, we are saying that one choice available to the coalition C is to bring about *all* of the goals in G'. Notice that, although they are obviously similar in type to effectivity functions as discussed above [1], we do not intend them to be interpreted as such. We give detailed comments on the relationship of QCG characteristic functions to effectivity functions in section 5.

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:

$$C \subseteq C' \text{ implies } V(C) \subseteq V(C')$$

Thus this constraint would say that monotonically increasing coalitions have montonically increasing sets of choices available to them – adding an agent to a coalition never reduces the choices available to the coalition. This constraint makes sense for some domains, but by no means all. For example, if we interpreted the choices available to a coalition as those that the coalition could enforce no matter what those agents outside the coalition did, then this constraint would certainly be appropriate (as is the case under the conventional reading of effectivity functions [1, p.21]). However, our notion of choices available to a coalition is intended to be more general that this. For example, in modelling certain legal scenarios, the addition of an agent to a coalition might prevent the coalition from making choices that would otherwise be available to it; clearly, in modelling such scenarios, the monotonicity property is not appropriate.

Collecting these components together, a *qualitative coalitional game* (QCG) is an $n+3$-tuple:

$$\Gamma = \langle G, Ag, G_1, \ldots, G_n, V \rangle$$

where

- $G = \{g_1, \ldots, g_m\}$ is a set of *possible goals*;
- $Ag = \{1, \ldots, n\}$ is a set of *agents*;
- $G_i \subseteq G$ is a set of goals for each agent $i \in Ag$, the intended interpretation being that any of the goals in G_i would satisfy i – but i is indifferent between the members of G_i; and
- $V : 2^{Ag} \to 2^{2^G}$ is a *characteristic function*, which for every coalition $C \subseteq Ag$ determines a set $V(C)$ of *choices*, the intended interpretation being that if $G' \in V(C)$, then one of the choices available to coalition C is to bring about *all* the goals in G' simultaneously.

We say a set of goals G' *satisfies* agent i if $G' \cap G_i \neq \emptyset$; we say that G' satisfies $C \subseteq Ag$ if it satisfies every member of C. Also, we say that G' is *feasible* for coalition C if $G' \in V(C)$.

Some examples are called for.

Example 2 *Voting procedures and related social choice mechanisms may be modelled as* QCG*s. Consider the following scenario:*

> *Three agents, a_1, a_2, and a_3, must choose between two outcomes, p and q. Agent a_1's preference is for outcome q, while the other agents prefer p. The mechanism used to choose an outcome is a simple majority vote.*

This scenario can easily be modelled by a QCG *with $Ag = \{a_1, a_2, a_3\}$, $G = \{p, q\}$, $G_1 = \{q\}$, $G_2 = G_3 = \{p\}$, and*

$$V(C) = \begin{cases} \{\{p\}, \{q\}\} & \textit{if } |C| \geq 2 \\ \emptyset & \textit{otherwise}. \end{cases}$$

This QCG *has a number of properties which are worth pointing out. First, it is* monotonic, *in the sense that monotonically increasing coalitions have monotonically increasing sets of choices. Second, we note that singleton coalitions have no choices available to them. The coalition $\{a_2, a_3\}$ has a feasible choice which satisfies both members: $\{p\}$. No coalition that counts a_1 as a member has a choice which satisfies all its members.*

The following example illustrates some slightly different properties of QCGs.

Example 3 *Consider a* QCG Γ_1, *in which* $Ag = \{a_1, a_2, a_3\}$, $G = \{g_1, g_2, g_3\}$, $G_i = \{g_i\}$, *and V is defined as follows:*

$$V(C) = \begin{cases} \{\{g_1, g_2\}, \{g_1, g_3\}, \{g_2, g_3\}\} & \textit{if } |C| = 2 \\ \emptyset & \textit{otherwise.} \end{cases}$$

The only coalitions that can achieve any *goals in* Γ_1 *are those of cardinality two. The characteristic function in this case is thus* not *monotonic, as adding a third member to a coalition* reduces *the choices of a group. The goal set* $\{g_1, g_3\}$ *is feasible for the coalition* $\{a_1, a_3\}$, *and also satisfies all members of this coalition. However, while this goal set is feasible for the coalition* $\{a_1, a_2\}$, *and satisfies* a_1, *it fails to satisfy* a_2. *The goal set* $\{g_1, g_2\}$ *satisfies* a_2 *but is not feasible for this singleton coalition. The set* $\{g_1, g_2, g_3\}$ *satisfies every member of* $\{a_1, a_2, a_3\}$ *but is not feasible for this coalition.*

3.1 Succinct Representations of QCGs

If we want to investigate the complexity of decision problems associated with QCGs, then we need to consider the *representation* of QCGs in input instances, and in particular, the issue of representing characteristic functions. An obvious, naive approach would be to enumerate the characteristic function as a set of ordered pairs, i.e., as the set $\{(I, O) : I \in 2^{Ag}, O = V(I)\}$. But there is a substantial difficulty here: the size of this set will be exponential in the size of $Ag \cup G$. This means that:

- such a representation would be utterly infeasible in practice; and
- such a representation is so large that it renders comparisons to this input size meaningless, since stating that we have an algorithm that runs in (say) time linear in the size of such a representation only actually means that it runs in time exponential in the size of $G \cup Ag$ – which is of no practical value whatsoever.

Such *extensive* representations would thus trivialise questions about the complexity of reasoning about QCGs, and hide their true complexity. What we therefore require is a *succinct* representation of characteristic functions – ideally, one such that the size of representation of the characteristic function is *polynomial* in the size of $G \cup Ag$.

Our starting point is the observation that characteristic functions may be viewed as *propositional logic functions*. That is, with any $V : 2^{Ag} \to 2^{2^G}$ there is an associated propositional function $f_V(Ag, G)$ such that $G' \in V(C)$ if and only if $f_V(\alpha_C, \beta_{G'})$ evaluates to $\top$, where the total instantiations α_C of Ag and $\beta_{G'}$ of G are such that $a_i = \top$ if and only if $a_i \in C$, and $g_j = \top$ if and only if $g_j \in G'$. In consequence we may represent V, for instances of decision problems in which QCGs are a part, by employing a suitable representation for its associated propositional logic function f_V.

There are, of course, a number of standard encoding formalisms that have been used for presenting elements of B_n, (the set of n argument propositional functions), e.g., truth-tables, Disjunctive and Conjunctive Normal Form, Binary Decision Diagrams [4,5], formulae over some finite (complete) basis, Boolean networks (also known as "straight-line programs"), etc. Of course, mechanisms such as truth-tables have space requirements comparable to the naive extensive form representation outlined earlier. Similarly, Disjunctive Normal Form and Conjunctive Normal Form require exponential (in n) length representations for some quite simple functions such as Majority and Parity, cf. [9, Chapter 5, p. 359]. Using Boolean networks as a representation for propositional functions offers one significant advantage via the Schnorr-Fischer-Pippenger simulation [31,11], which demonstrates that the (sequence of) characteristic functions of any language L in P, i.e., the sequence $\langle f_n \rangle_L$ for which $f_n(w) = \top$ if $|w| = n$ and $w \in L$, can be described by a corresponding sequence of polynomial size networks.

We adopt an approach whereby the propositional function f_V defining the characteristic function, V, is presented as a *formula* Ψ of propositional logic over propositional variables Ag and G, that is, we have a propositional variable for each agent and each goal. The *size* of the formula, $\Phi(X_n)$, denoted $|\Phi(X_n)|$, is the total number of occurrences of *literals* in $\Phi(X_n)$, a literal being either a positive (x_i) or negated ($\neg x_i$) instance of a propositional variable. We further assume that $\Phi(X_n)$ is described using only *binary* Boolean operations and unary negation ($\neg$), thus formulae may not employ as basic operations functions defined over arbitrarily many arguments [2]. Given a formula Ψ (over variables Ag, G), and sets $C \subseteq Ag$ and $G' \subseteq G$ we then require that

$$\Psi[C, G'] = \top \text{ if and only if } G' \in V(C).$$

We note that this representation is such that given any Ψ (representing a characteristic function V), C, and G', determining whether $\Psi[C, G'] = \top$ (and hence whether $G' \in V(C)$) can clearly be done in deterministic polynomial time.

Postponing for a moment any discussion of the extent to which such a representation satisfies our requirements, consider the following concrete example of this representation.

Example 4 *Consider the following formula of propositional logic:*

$$\Psi(Ag, G) = E_2^3(Ag) \ \wedge \ E_2^3(G)$$

[2] While for associative operations such as $\wedge$, we employ a short form notation $\wedge_{i=1}^{n} y_i$ this is purely for descriptive convenience rather than indicative of n-argument conjunction being a basic operation.

where

$$E_2^3(z_1, z_2, z_3) = (z_1 \vee z_2)(z_1 \vee z_3)(z_2 \vee z_3)(\neg z_1 \vee \neg z_2 \vee \neg z_3)$$

This formula clearly represents the characteristic function from the QCG Γ_1*, above.*

Hereafter, when we refer to a QCG, we mean a structure $\Gamma = \langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$, with Ψ being a propositional formula representing a characteristic function, and other components as previously defined. For clarity of exposition, we will sometimes refer to characteristic functions in their functional form V in the text; but of course we always use the logical form in proofs of complexity.

The remainder of this section contains a (rather technical) discussion of the extent to which this representation of characteristic functions satisfies our requirements as stated above, and investigates some issues surrounding the representation. This discussion will be primarily relevant to readers with an interest in the technical aspects of our results, and is not required reading for those who simply wish to acquaint themselves with the decision problems we consider and the associated results; such readers may wish to skip directly to section 4.

It is well-known that propositional logic formulae can be treated as a restricted class of Boolean networks, cf. [9, Chapter 1, p. 23], and although it is considered unlikely that the simulation obtained in [31,11] extends to formulae, this does not affect our results. The reductions through which our complexity classifications are obtained employ problems with a propositional formula as part of the instance, e.g., SAT, QSAT_2^{Π}. It will be clear that the size of the formulae constructed when forming a QCG in these reductions is always polynomial in the size of the formula provided as an instance of SAT, QSAT_2^{Π}, etc. In addition, noting the interpretation of formulae as a restricted Boolean network model, the computational complexity of SAT is unchanged if we assume that its instances are Boolean *networks* rather than formulae, and thus our results hold regardless of which of these two standard methods are used to encode propositional functions. We employ formulae as our method as, arguably, this renders some presentational aspects with greater clarity.

It is, of course, the case that irrespective of the encoding form used to describe n variable propositional functions – whether truth-tables, formulae, or networks – there will inevitably be some functions whose encoding requires exponentially many bits in n. This property can be shown by considering the notion of a *representation language* for propositional functions.

Definition 5 *A language $L \subseteq \{0,1\}^*$ is a* reasonable representation language for propositional functions *if*

a. Every $w \in L$ is associated with a propositional function f_n of n arguments, $\eta(w)$

b. *For every propositional function f_n of n arguments there is* at least one $w \in L$ *for which $\eta(w) = f_n$, i.e., such that w is associated with f_n.*
c. *The language $L^{(n)} = \{1^n 0w \ : \ w \in L$ and $\eta(w)$ has n arguments$\}$ is in* P*, i.e., given $1^n 0w \in \{0,1\}^*$ there is a deterministic polynomial-time algorithm that decides if $w \in L$ and depends on exactly n propositional variables.*
d. *There is an algorithm that given $\langle w, n, \alpha \rangle$ with $\alpha \in \langle \top, \bot \rangle^n$ determines if $1^n 0w \in L^{(n)}$ and returns the Boolean value $f_n(\alpha)$ where $\eta(w) = f_n$ with the algorithm taking time polynomial in $n + |w|$.*

So, with truth-tables, a sequence of 2^n bits would describe successive values from $\langle \bot, \ldots, \bot \rangle$ through $\langle \top, \ldots, \top \rangle$; for a network with M 2-input gate operations, a description is given by a sequence of $n + M$ triples $h_k = \langle \theta, l_k, r_k \rangle$ where θ is a Boolean operation if $k > n$ and 0 otherwise, with l_k and r_k being the indices of the tuples supplying the operands for the operation.

Given a reasonable representation language for propositional functions, L, the complexity measures $C_L(f_n)$, (the length of the *shortest* encoding of f_n in the language L), and $C_L(n)$, (the worst-case number of bits needed by the representation L to describe any n argument propositional function), are defined as follows [3] .

$$C_L(f_n) \hat{=} \min_{w \in L^{(n)} \ : \ \eta(w) = f_n} \ |w|$$
$$C_L(n) \hat{=} \max_{f_n \ : \ f_n \text{ is a } n\text{-argument function}} \ C_L(f_n)$$

The following result is now readily established.

Fact 6 *We recall that B_n denotes the set of all distinct n argument propositional functions. If L is* any *reasonable representation language for propositional functions,*

a. $\quad C_L(n) \ \geq \ 2^n - 1.$

b. $\quad \forall \, \epsilon > 0 \ \lim_{n \to \infty} \dfrac{|\{f_n \ : \ C_L(f_n) \ \leq \ (1-\epsilon)2^n\}|}{|B_n|} \ = \ 0$

Proof: Both parts are immediate from the fact that $|B_n| = 2^{2^n}$, however using at most $k(n)$ bits no more than $\sum_{i=0}^{k(n)} 2^i = 2^{k(n)+1} - 1$ distinct $w \in L^{(n)}$ can be

---

[3] For any reasonable representation language, L, $C_L(f_n)$ and $C_L(n)$ are recursive functions, e.g., given f_n described as a 2^n bit truth-table, $C_L(f_n)$ could be determined by enumerating binary words, w, in order of increasing length until the first such is found having $\eta(w)(\alpha) = f_n(\alpha)$ for all $\alpha \in \langle \top, \bot \rangle^n$. Similarly $C_L(n)$ is found by repeating the process just outlined for each f_n.

formed.[4] ∎

While the information-theoretic bound of Fact 6 indicates that any representation formalism for propositional functions must employ exponential length encodings for some cases, this fact does not justify committing to mechanisms such as truth-tables, (which are *always* exponential in the number of propositional variables); nor does it imply that representing a propositional function as a formula must always be exponential. The use of formulae, networks etc., at least provides the capability to describe concise presentations of propositional functions when such are possible. We note that, despite the existential proof of Fact 6, there has yet to be identified any explicitly defined case for which superpolynomial network or formula size results have been demonstrated in models which allow an arbitrary logically complete basis to be used.

As a final issue regarding representation via formulae we consider a further complication that would arise were one to require a *monotonicity* property with respect to the sets of feasible outcomes. More precisely we revisit the following class of QCGs.

Definition 7 *A* QCG $\Gamma = \langle G, Ag, G_1, \ldots, G_n, V\rangle$ *is* coalition monotonic *if whenever $G' \in V(C)$ then $G' \in V(C')$ for all $C' \supseteq C$.*

That is, any feasible choice, G' available to a coalition C is also a feasible choice for coalitions, C', of which C is a subset. We note that this does *not* necessarily mean that G' should *satisfy* C' if it satisfies C: it might well be the case that in extending C to C' some of the new agents do not have any of their goals met by those available in G'.

We first recall some results from the field of Boolean function complexity.

Fact 8 *Let $X_n = \{x_1, \ldots, x_n\}$ be a set of n propositional variables, $\Phi(X_n)$ a propositional formula, and $U_2 = \{\wedge, \vee, \neg\}$ the (logically complete) basis consisting of binary conjunction, binary disjunction, and unary negation. We recall that $|\Phi(X_n)|$ is the total number of occurrences of literals in $\Phi(X_n)$.*

a. There is a formula $\Psi(X_n)$ logically equivalent to $\Phi(X_n)$, employing only the operations in U_2 and with $|\Psi(X_n)| = O(|\Phi(X_n)|^{\log_3 10}) \leq O(|\Phi(X_n)|^{2.096})$.
b. Any formula $\Phi(X_n)$ using only the operations in U_2 may be transformed to a logically equivalent formula $\Psi(X_n)$ in which negation ($\neg$) is applied only to propositional variables, $x_i \in X_n$, and with $|\Psi(X_n)| \leq 2|\Phi(X_n)|$.
c. The translation from $\Phi(X_n)$ to an equivalent formula $\Psi(X_n)$ meeting the conditions specified by (a) and (b) can be carried out by an algorithm whose

[4] This bound of $\Omega(2^n)$ bits does *not* contradict the (optimal) upper bound results on network size of $2^n/n$ proved by Lupanov [19]: Lupanov's bound refers to the number of *operations* (gates) not to the total number of *bits* used in describing the network.

running time is polynomial in $|\Phi(X_n)|$.

Proof: Part (a) is from [27]. Part (b) is an easy induction via De Morgan's Laws – $\neg(x \wedge y) \equiv \neg x \vee \neg y$, $\neg(x \vee y) \equiv \neg x \wedge \neg y$. Finally, part (c) follows by combining the obvious linear time algorithm to effect (b) and the fact that the construction given by [27] is realised by a polynomial time procedure. ∎

We say that a formula that satisfies the conditions on $\Psi(X_n)$ in Fact 8(b) is in *standard form* and, additionally note that the exponent 2.096 in Fact 8 is close to optimal: for bases permitting the operation exclusive-or ($\oplus$), Khrapchenko [18] shows that with $\Phi(X_n) = x_1 \oplus x_2 \cdots \oplus x_n$ the size of any equivalent U_2 formula is at least n^2.

Fact 8 provides sufficient background to obtain the following characterisation of coalition monotonic QCGs.

Theorem 9 *A* QCG $\Gamma = \langle G, Ag, G_1, \ldots, G_n, V\rangle$ *is coalition monotonic if and only if* $f_V(Ag, G)$ *has a representation as a* standard form *propositional formula,* $\Phi(Ag, G)$ *in which there are* no occurrences *of the literal* $\neg a_i$ *for any* $a_i \in Ag$.

Proof:

($\Rightarrow$) Suppose $\Gamma = \langle G, Ag, G_1, \ldots, G_n, V\rangle$ is coalition monotonic and consider the propositional function $f_V(Ag, G)$ that describes V. For $C \subseteq Ag$ and $G' \subseteq G$ define the propositional formulae π_n, δ_n, and γ_m by

$$\begin{aligned}\pi_n(C) &= \textstyle\bigwedge_{a_i \in C}\ a_i\\ \delta_n(C) &= \textstyle\bigwedge_{a_i \in C}\ a_i\ \wedge\ \bigwedge_{a_i \notin C}\ \neg a_i\\ \gamma_m(G') &= \textstyle\bigwedge_{g_j \in G'} g_j\ \wedge\ \bigwedge_{j \notin G'} \neg g_j\end{aligned}$$

Then $f_V(Ag, C)$ satisfies

$$f_V(Ag, G) \equiv \bigvee_{C \subseteq Ag} \delta_n(C) \wedge \left(\bigvee_{G' \in V(C)} \gamma_m(G') \right)$$

Since Γ is coalition monotonic,

$$f_V(Ag, G) \equiv \bigvee_{C \subseteq Ag} \bigvee_{D \subseteq Ag/C} \pi_n(C \cup D) \wedge \left(\bigvee_{G' \in V(C)} \gamma_m(G') \right)$$

which is,

$$\bigvee_{C \subseteq Ag} \pi_n(C) \wedge \left(\bigvee_{D \subseteq Ag/C} \pi_n(D) \right) \wedge \left(\bigvee_{G' \in V(C)} \gamma_m(G') \right)$$

For any $C \subseteq Ag$, the expression $\vee_{D \subseteq Ag/C}\ \pi_n(D)$ is a tautology[5] i.e. equivalent to the constant value $\top$. It follows, therefore, that $f_V(Ag, G)$ when Γ is coalition monotonic, may be described by the propositional formula,

$$\bigvee_{C \subseteq Ag} \pi_n(C) \ \wedge \ \left(\bigvee_{G' \in V(C)} \gamma_m(G') \right)$$

which is in standard form and has no occurrences of the literal $\neg a_i$, as required.

($\Leftarrow$) For the converse implication, let $\Gamma = \langle G, Ag, G_1, \ldots, G_n, V \rangle$ be a QCG for which V is described by a formula $\Phi(Ag, G)$ in standard form with no occurrences of literals $\neg a_i$. Since the binary operations used in defining $\Phi(Ag, G)$ are just $\wedge$ and $\vee$, $\Phi(Ag, G)$ may be translated into an equivalent DNF formula $\Psi(Ag, G)$ without introducing any literals of the form $\neg a_i$. Suppose $q(Ag, G)$ is a product term in this DNF formula. Then $q(Ag, G) = \pi(C) \wedge \mu(G)$ for some $C \subseteq Ag$ and $\mu(G)$ a product of literals over G. Letting G' be the subset of these that occur positively in $\mu(G)$, it is certainly the case that $G' \in V(C)$: $\Phi[C, G'] = \Psi[C, G'] = \top$ and $\Phi(Ag, G)$ is a propositional formula defining V. In addition, for any $C' \supset C$, $q[C', G'] = q[C, G'] = \top$ and thence $\Phi[C', G'] = \top$ so that $G' \in V(C')$, i.e., Γ is coalition monotonic.

∎

Theorem 9 shows that it is possible to characterise coalition monotonic QCGs via *suitably presented* propositional formulae describing $f_V(Ag, G)$. For several of the decision problems we examine in the main body of this paper, one question that arises is the extent to which the complexity classification proved for general cases is affected if the QCGs in question are restricted to those which are coalition monotonic. Suppose that $\Phi(Ag, G)$ is the propositional formula given in an instance of a decision problem concerning such QCGs. It could be insisted upon that $\Phi(Ag, G)$ is given in standard form with no occurrences of $\neg a_i$: Theorem 9 guarantees the existence of such a form. Alternatively we could allow $f_V(Ag, G)$ to be described – by formula, network or other reasonable representation – without imposing any further restriction. The first of these has the significant advantage that it is easy to *validate* that an instance has the correct form. Except when we allow an extensive representation such as truth-tables, it may be extremely hard to validate that an *unrestricted* representation of $f_V(Ag, G)$ does indeed describe a coalition monotonic QCG: a naive quadratic (in terms of the truth-table size) algorithm allows this for extensive representations, however, to validate an arbitrary propositional formula or Boolean network as defining a coalition monotonic QCG is at least as difficult as deciding unsatisfiability. For example, consider the equivalent problem of deciding if a DNF formula, $\Phi(X_n)$ is a propositional tautology with instances limited to those for which $\Phi(\perp, \perp, \ldots, \perp) = \top$:[6] then $\Phi(X_n)$ defines a monotonic propo-

[5] Noting the convention that the *empty* conjunction evaluates to $\top$, hence $\pi_n(\emptyset) \equiv \top$.

[6] An easy proof shows that this restriction still defines a co-NP–complete problem.

sitional function if and only if $\Phi(X_n)$ is a propositional tautology. These issues would seem to argue strongly that if we wish to consider the complexity of decision questions specifically addressing coalition monotonic QCGs then these ought to be presented as indicated in Theorem 9 (whether as formula or network). We conclude our technical discussion of representation issues for QCGs by illustrating that this requirement may also be problematic.

Consider the following coalition monotonic QCG, MATCH, involving $n = t^2$ agents:

$$G = \{g_1\} \;\; ; \;\; Ag = \{a_{i,j} \,:\, 1 \leq i \leq t,\; 1 \leq j \leq t\}$$
$$G_{i,j} = \{g_1\}$$

$$V(C) = \begin{cases} g_1 & \text{if } \exists D \subseteq C \text{ and a permutation } \sigma \text{ of } \{1,\ldots,t\} \text{ such that} \\ & a_{i,\sigma(i)} \in D \text{ for each } 1 \leq i \leq t \\ \emptyset & \text{otherwise} \end{cases}$$

With this example,

Fact 10

a. *If $\Phi(Ag, G)$ is any formula in standard form representing the function f_V for the* QCG MATCH *and $\Phi(Ag, G)$ contains no occurrences of $\neg a_{i,j}$, then $|\Phi(Ag, G)|$ is $\Omega(n^{\log n})$.*
b. *The propositional function $f_V(Ag, C)$ can be represented by a Boolean network in the basis U_2 which uses $O(n^{2.5} \log n)$ operations.*

Proof: For (a), consider the sub-formula of $\Phi(Ag, G)$ resulting by setting $g_1 = \top$ and simplifying. If $\Phi(Ag, G)$ is in standard form with $a_i \in Ag$ appearing only as a positive literal, then $\Phi(Ag, \top)$ contains no incidences of negation and employs only the operations of binary conjunction and disjunction. Furthermore the resulting formula represents the propositional function, equivalent to

$$PM(Ag) = \bigwedge_{\sigma \,:\, \sigma \text{ is a permutation of } \{1,\ldots,t\}} a_{i,\sigma(i)}$$

The results of Razborov [28] give the lower bound stated on the size formulae over the (incomplete) basis $\{\wedge, \vee\}$ representing this function. [7] The upper bound in (b) on the size of unrestricted Boolean networks follows from Hopcroft and Karp [15]

[7] The result of [28] is actually rather stronger, since it applies to *networks* containing only $\{\wedge, \vee\}$ operations and without negated arguments.

via the simulation of [31,11]. ∎

Notwithstanding the issue raised by results such as Fact 10, for those decision questions where coalition monotonic QCGs are considered subsequently, it will be assumed that the propositional formula $\Phi(Ag, G)$ representing $f_V(Ag, G)$ is given in standard form and contain no occurrences of the literal $\neg a_i$: this requirement ensures that the encoding of an instance may be efficiently validated as describing a coalition monotonic QCG. We observe, at this juncture, while in some cases complexity results for the special case follow directly from the reductions used for general QCGs, (e.g., the problem "Successful Coalition"), in other cases, (e.g., the problem "Trivial Game"), this is rather more involved, and we present a separate reduction for coalition monotonic QCGs in order to avoid unnecessarily complicating the general case.

When we wish to distinguish between the general case of a decision problem, Q for QCGs and its restriction to coalition monotonic QCGs we shall employ the notation Q^{mono} to indicate the latter class of instances.

4 The Complexity of Qualitative Coalitional Games

In this section, we obtain results categorising the computational complexity of fourteen naturally defined decision problems relating to QCGs. Informally, the problems we examine are as follows.

- SUCCESSFUL COALITION (SC)
  Given a particular QCG and a coalition in this QCG, is there a feasible choice available to the coalition that will satisfy all its members? (Notice that we refer to *successful* rather than *winning* coalitions: this is because we are dealing with cooperative games, and it is entirely possible that more than one coalition is successful.)
- SELFISH SUCCESSFUL COALITION (SSC)
  Given a particular QCG and a coalition in this QCG, is it the case that this coalition has a feasible choice that will satisfy an agent if and only if it is a member of the coalition?
- MINIMAL COALITION (MC)
  Given a particular QCG and a coalition in this QCG, is it the case that no strict subset of this coalition is successful?
- UNATTAINABLE GOAL SET (UGS)
  Given a particular QCG and a coalition in this QCG, is it the case there is a set of goals that would satisfy every member of this coalition, but that this set of goals does not represent a feasible choice for the coalition?
- CORE MEMBERSHIP (CM)
  Given a particular QCG, a coalition in this QCG, and a set of goals in this QCG,

is it the case that this goal set is in the core of this coalition, that is to say: (i) the goal set satisfies every member of the coalition; (ii) the goal set represents a feasible choice for the coalition; and (iii) the coalition is minimal? (Below, we comment on the relationship of this notion to the conventional definition of the core [22, pp.258–263].)

- CORE NON-EMPTY (CNE)
  Given a particular QCG and a coalition in this QCG, is there *some* set of goals that is in the core of this coalition?
- VETO PLAYER (VP)
  Given a particular QCG and two agents i and j in this QCG, is it the case that i is a veto player for j in this QCG [22, p.261], i.e., that i is a member of every coalition that has a feasible choice satisfying j?
- MUTUAL DEPENDENCE (MD)
  Given a particular QCG and a coalition C in this QCG, is it the case that C is a subset of every coalition with a feasible choice that satisfies one of the members of C?

These decision problems are concerned with properties of given coalitions with respect to a particular QCG. We can, however, also formulate problems regarding properties of *goal sets*. Thus, we have the following two problems.

- GOAL REALISABILITY (GR)
  Given a particular QCG and goal set in this QCG, does there exist some coalition such that this goal set both satisfies every member of the coalition, and represents a feasible choice for it?
- NECESSARY GOAL (NG)
  Given a particular QCG and goal set in this QCG, is it the case that for every coalition and feasible choice that this coalition could make to satisfy every member of the coalition, the given goal set is included in this choice? That is, is it the case that the goal set will necessarily be brought about as a "side effect" of any coalition forming and bringing about their goals?

Finally we note four decision problems concerned with general properties of a given QCG:

- EMPTY GAME (EG)
  Given a QCG, is it the case that *no* coalition in this QCG is successful, i.e., that no coalition in this QCG has a feasible choice available which satisfies every member of the coalition?
- TRIVIAL GAME (TG)
  Given a QCG, is it the case that *every* coalition in this QCG is successful, i.e., that every coalition in this QCG has a feasible choice available which satisfies every member of the coalition?
- GLOBAL UNATTAINABILITY (GU)
  Given a QCG, does there exist a set of goals in this QCG such that for every

Problem	Description	Complexity	Q^{mono}	Reference
SC	SUCCESSFUL COALITION	NP-complete	NP-complete	Thm. 12, Cor. 17
SSC	SELFISH SUCCESSFUL COALITION	NP-complete	NP-complete	Cor. 14, 17
UGS	UNATTAINABLE GOAL SET	NP-complete	NP-complete	Cor. 16, 17
MC	MINIMAL COALITION	co-NP-complete	co-NP-complete	Thm. 22, Cor.23
CM	CORE MEMBERSHIP	co-NP-complete	co-NP-complete	Thm. 25, Cor 26
CNE	CORE NON-EMPTINESS	D^p-complete	D^p-complete	Thm. 27, Cor 28
VP	VETO PLAYER	co-NP-complete	-	Thm. 30
MD	MUTUAL DEPENDENCE	co-NP-complete	-	Thm. 33
GR	GOAL REALISABILITY	NP-complete	P	Cor. 19, 20
NG	NECESSARY GOAL	co-NP-complete	-	Cor. 32
EG	EMPTY GAME	co-NP-complete	co-NP-complete	Thm. 35, Cor.36
TG	TRIVIAL GAME	Π_2^p-complete	Π_2^p-complete	Thm. 38, Cor.39
GU	GLOBAL UNATTAINABILITY	Σ_2^p-complete	NP	Thm. 41, Cor.42
IG	INCOMPLETE GAME	D_2^p-complete	-	Thm. 44

Table 1
Decision Problems in QCGs and their complexity

coalition in the QCG, should the goal set satisfy every member of the coalition, then the goal set is not feasible for the coalition?

- INCOMPLETE GAME (IG)
  Given a QCG, is it the case that this QCG is trivial *but* has some globally unattainable goal set?

The main results of this paper, as relating to these problems, are summarised in Table 1 (recall that we distinguish between the general case of a decision problem Q for QCGs, and its restriction to *coalition monotonic* QCGs, with the notation Q^{mono} indicating the latter class of instances).

Before moving on to the technical substance of our results, we discuss the motivation underlying our problem selection in less formal terms.

4.1 Discussion and Motivation

The decision problems on QCGs introduced above can be viewed as falling into three broad categories: questions concerning properties of given coalitions, (e.g., Successful Coalition, Minimal Coalition); questions regarding properties of particular subsets of goals, (e.g., Goal Realisability); and, finally, questions about the behaviour of particular QCGs, (e.g., Empty Game, Trivial Game). In this section, we motivate our interest in these specific decision problems.

In many ways, Successful Coalition incorporates the most basic question that is of interest with respect to any given QCG. Given the framework of a QCG, the Successful Coalition problem in effect addresses the issue of whether it might be considered "worthwhile" for a given coalition to form, in the sense that there is a goal set available to this coalition that is both feasible and satisfies each member. Of course, given that a particular coalition is successful in this sense, we cannot be certain that this coalition will form; but we *can* be certain that an *unsuccessful* coalition will *not* form – because, by definition, the formation of such a coalition would leave at least one member unsatisfied. In this sense, success is a necessary, but not sufficient condition for coalition formation in QCGs.

On first inspection, the closely related Selfish Successful Coalition problem may seem less well-motivated: assuming that an agent's (principal) aim is to enlist in a coalition with whose support it can realise a goal it wishes to be satisfied, why ought it to be concerned with the status of agents outside the coalition, and in particular, whether such might be satisfied with a particular feasible goal set? Of course, in many scenarios, an agent will be indifferent to the level of satisfaction achieved by non-members: our contention, however, is that such scenarios do not encompass all settings that might usefully be modelled within a QCG environment. Although we do not explicitly consider in this article concepts of preferences between distinct feasible goal sets, in modelling distinct societal attitudes one could posit a number of classes of goal set with respect to the views not only of a coalition's members but also with respect to how these might be seen by those agents excluded from membership. Thus, we have the two cases considered below: goal sets which are simply required to be feasible and satisfying and, the subset of these comprising goal sets which satisfy *exactly* the members of the coalition, leaving all others agents unsatisfied. As a "real" example of the latter case, one might consider goals corresponding to executing particular applications within a multi-user system: as a result of limited shareable resources or computational power, only particular subsets of these applications may run simultaneously, and some applications may require support from others. In this context one might have a collection of agents each with desired goals of carrying out some application. As might be the case with "hostile" agents seeking to damage the system functionality (e.g., viruses, Trojan horses, etc.), an agent may wish to form a coalition with "like minded" agents, so that the effect of each executing one of its satisfied goals is to render the system unusable to other agents (since they cannot execute any of their desired applications).

Such a scenario assumes that agents may seek to form coalitions not only to advance their own interests, but also to damage the interests of others. While such actions may stem from "rational" motivations, (e.g., economic interests resulting in cartels and monopoly dealing), there is also the possibility, well recognised in human social psychology, that such behaviour is motivated by spite or by a sense of *schadenfreude*: the fact that others are discomfited by one's choices increasing the desirability of a particular feasibly satisfying choice. There is a huge range of questions arising from what one might term sociopathic and obstructive agent behaviour, and with a few exceptions, (addressing non-cooperative dialogue processes [10,12]), analytic computational models of these issues have been largely neglected: the decision problem Selfish Successful Coalition considers one very basic aspect [8] of such behaviour. Of course, we do not argue that this behaviour is desirable or that the development of such mechanisms should be an objective of multiagent systems research – far from it. But it *is* possible and clearly does occur in human societies: understanding such properties may be a significant aspect in defending against their potential consequences.

Turning to Minimal Coalition, this, of course, is a problem that has been well-studied in the quantitative setting, where by allying with a strict subset of its feasible coalition an agent may be able to attain a better pay-off than it enjoys with the coalition intact. While this utilitarian consideration does not arise within our qualitative framework, the issue of whether a coalition is minimal remains of some interest. For example, attitudes that may encourage an agent to seek out a selfish successful coalition, as discussed earlier, might prompt agents to seek minimal coalitions. In addition, there are closely related ideas in political science, which suggest that the coalitions which actually form will tend to be the *smallest* coalitions capable of being successful [39, pp.82–85].

Our examination of the issues of success and minimality leads to an analogue of the "core" for qualitative coalitional games. We consider the role of the core in our setting in rather more detail prior to presenting the technical analysis of Section 4.3.

Finally, with respect to decision questions addressing properties of given coalitions, we have the problem Unattainable Goal Set, asking whether there is any satisfying choice for the coalition which fails to be feasible. Although within a typical application of our model one would be primarily interested in the question of whether some satisfying set was *feasible*, (i.e., the decision problem associated with Successful Coalition), rather than than the possibility of infeasible satisfying sets, there are several scenarios in which Unattainable Goal Set is important. Consider QCGs within which some coalitions define *negative* instances of UGS. By virtue of this property, such coalitions have available the greatest degree of latitude in deciding which set of goals to bring about: any goal set that satisfies each member can be feasibly achieved. In extending the basic QCG framework of this article to incorporate notions of preference and choice between different feasible sets, coalitions that can attain any combination of satisfying goals have an advantage in that their

[8] The motivation of agents trying to form a "selfish coalition" is "rational" in the sense that a prerequisite for a feasible goal set is that each member is satisfied. We mention, and no more than this, that a rather more extreme and "irrational" aim might be to seek a coalition and feasible goal set with the express purpose of leaving particular non-coalition members unsatisfied, even if this can only be achieved at the cost of the agent itself being unsatisfied.

preferences can always be realised, regardless of whatever ranking scheme may be used to compare different goal sets. In contrast, for coalitions which define positive instance of UGS, there will be at least one preference ordering under which the coalition can never bring about its most desired goal set.

The main problem we consider relating to properties of given sets of goals is that of Goal Realisability which asks of a given a set of goals whether there is any coalition these satisfy and are feasible for. This is a natural "dual" problem to Successful Coalition: instead of seeking a feasible and satisfying set for a given coalition, we look for a coalition for which a given set of goals is feasible and satisfying. In many instances an external observer may be concerned less with the success or failure of specific coalitions than they are with whether some particular collection of objectives *could* be realised. Thus, given a set of goals, the primary interest lies in identifying a coalition for which these goals represent a feasible choice and whose members are satisfied. Of course, assuming suitable coalitions are discovered, this then raises a question that is outside the scope of the present article, namely that of mechanisms for choosing between different coalitions and negotiation strategies for bringing any of these into being.

We now come to the final group of problems considered: those whose instances are a single QCG and ask whether some property of interest holds of it. The problem Empty Game can be seen as addressing whether a given QCG exhibits a particular type of "redundancy", namely that of possessing no successful coalition. There are a number of ways in which one could interpret QCGs with this property, but perhaps the most important is that cooperation of any kind is not likely to occur, and will be unstable if it does, since there is no incentive for any coalition to form or hold together.

In contrast, the problem Trivial Game deals with the opposite extreme: whether the given QCG is such that *every* coalition would succeed. Again one could interpret "trivial" QCGs as indicating an element of redundancy in the specification: no multi-agent coalition has any reason to emerge, since each individual agent without assistance is capable of bringing about some satisfying and feasible goal set.

We note, in passing, that it is straightforward to show that the question of whether all *singleton* coalitions are successful is of equivalent complexity to deciding if a *single* coalition succeeds, and more generally, the same is true of questions in which polynomially many (in terms of n) coalitions are specified, e.g., all coalitions containing at most k members for some constant k. Just as empty games might be interpreted as indicating over demanding agent requirements, so too one might interpret trivial games as combining an overabundance of feasible outcomes with too readily satisfied agent goals: the first ensuring an extensive choice of feasible outcomes for each coalition; the second increasing the likelihood of any one agent being satisfied by a given set of goals.

The problems Trivial Game and Empty Game address two extremes concerning coalitional properties in QCGS: whether every (respectively, no) coalition succeeds. The question with which Global Unattainability deals is, in effect, the complementary problem to that posed by Trivial Game but with respect to *sets of goals* rather than *sets of agents*: asking whether there is any goal set which is never both feasible and satisfying is the complement problem of asking whether *every* goal set satisfies and is feasible for at least one coalition. Viewed in this way, Global Unattainability is closely related to the problem Goal Realisability discussed above, since it is equivalent to asking of a QCG whether there is a goal set for which the answer regarding its realisability is negative. In terms of the summary results in Table 1 we have two immediate consequences of this interpretation: given that Goal Realisability is NP-complete, its complement problem is co-NP-complete, and thus we easily obtain a Σ_2^p algorithm for deciding whether some goal set is either infeasible or unsatisfactory for every coalition. A second point of interest concerns the coalition monotonic variant of this problem. As is evident from Table 1 most of the decision questions are insensitive to the imposition of a monotonicity constraint – a property which one might advance as a further reason for not insisting upon such a restriction – the significant exception to this being Goal Realisability. Since the coalition monotonic variant is in P, a complexity class that is closed under complementation, this results in the related coalition monotonic Global Unattainability being decidable by an NP method. We observe that QCGs which define *negative* instances of Global Unattainability are likely to be some practical interest: irrespective of which set of goals an external observer may wish to bring about, for such QCGs one is guaranteed to have at least one coalition that could be motivated to accomplish this.

Our final problem – Incomplete Game – may appear on the surface to be of less interest. The motivation for our study of Incomplete Game is in considering the relationship between the capability of a QCG from the contrasting views of which *coalitions* can succeed and which *goal sets* can be effected. One would not, of course, generally expect that QCGs in which every coalition succeeds, (i.e., positive instances of Trivial Game), to be such that every goal set can be realised, (i.e., also *negative* instances of Global Unattainability). It might reasonably be conjectured that the extremes represented by Trivial Game and Global Unattainability would lead to questions about whether a given QCG satisfies some combination of the two being no more difficult to decide than either separately. For the particular conjunction of properties considered within Incomplete Game, the result proved in Theorem 44 offers evidence that such expectations are unlikely to be satisfied. We observe that the particular question asked of QCGs by Incomplete Game could, in very informal terms, be thought of as deciding whether a QCG is such that every coalition can achieve *something* but there are *some things* which *no* coalition can achieve.

Having completed our informal overview of some of the motivating issues underlying the problems considered, we now proceed to the technical results.

Recall that a coalition is *successful* if that coalition has a feasible choice satisfying all members of the coalition.

Example 11 *Returning to the example* QCG Γ_1*, above, the coalition* $\{a_1, a_3\}$ *is successful, as it can choose* $\{g_1, g_3\}$ *as the goal set to bring about, and this choice satisfies all members of the coalition. Similarly,* $\{a_1, a_2\}$ *and* $\{a_2, a_3\}$ *are successful. All other coalitions, however, are unsuccessful.*

Formally, the decision problem is as follows:

SUCCESSFUL COALITION: (SC)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ and coalition $C \subseteq Ag$.
Answer: "Yes" if $\exists G' \subseteq G$ s.t. $\forall i \in C$, $G_i \cap G' \neq \emptyset$ and $\Psi[C, G'] = \top$

Theorem 12 SUCCESSFUL COALITION *is* NP*-complete.*

Proof: For membership, the following NP algorithm decides the problem: guess a subset G' of G and verify both that $\Psi[C, G'] = \top$ and for each $i \in C$ that $G_i \cap G' \neq \emptyset$. Both steps can obviously be done in time polynomial in the size of Ψ, and so the problem is in NP.

For NP-hardness, we reduce from SAT [23, p.171]. An instance of SAT is given by a propositional logic formula $\Phi(x_1, \ldots, x_n)$, the aim being to answer "yes" if there is some valuation to the Boolean variables $x_1, \ldots, x_n$ that satisfies the formula. We create an instance $\langle \Gamma_\Phi, C\rangle$ of SC as follows. Set $Ag = \{a_1, \ldots, a_n\}$ so that there is an agent a_i for each propositional variable x_i of Φ. The goals for a_i are $G_i = \{g_i^\top, g_i^\perp\}$, and $G = \cup_{i=1}^n G_i$. The coalition considered in the instance of SC formed is the set of all agents, i.e., $C = Ag$. Finally Ψ is the formula

$$\left(\bigwedge_{i=1}^n a_i\right) \;\wedge\; \Phi(g_1^\top \vee \neg g_1^\perp, \ldots, g_i^\top \vee \neg g_i^\perp, \ldots, g_n^\top \vee \neg g_n^\perp) \wedge \bigwedge_{i=1}^n (g_i^\top \oplus g_i^\perp)$$

i.e., the expression $g_i^\top \vee \neg g_i^\perp$ is substituted for each occurrence of the variable x_i in Φ. It is obvious that given Φ this instance $(\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle, Ag)$ of SC can be constructed in time polynomial in the length of Φ. We now prove that Ag is a successful coalition for the QCG $\Gamma_\Phi = \langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ if and only if $\Phi(x_1, \ldots, x_n)$ is satisfiable:

($\Rightarrow$) Assume that Ag is successful. Then there is a subset G' of G satisfying $G_i \cap G' \neq \emptyset$ for each $a_i \in Ag$ and with $\Psi[Ag, G'] = \top$. Now since $\Psi[Ag, G'] = \top$ it must be the case that for each a_i, $|G' \cap G_i| = 1$, for otherwise the sub-expression

$\wedge_{i=1}^{n}(g_i^\top \oplus g_i^\perp)[G']$ evaluates to $\perp$. Thus, $\Psi[Ag, G'] = \top$ is equivalent to

$$\Phi((g_1^\top \vee \neg g_1^\perp), \ldots, (g_i^\top \vee \neg g_i^\perp), \ldots, (g_n^\top \vee \neg g_n^\perp))[G'] = \top$$

and we construct a satisfying assignment for $\Phi(x_1, \ldots, x_n)$ simply by fixing x_i to the value of $(g_i^\top \vee \neg g_i^\perp)$ as determined by G'.

($\Leftarrow$) Assume that $\Phi(x_1, \ldots, x_n)$ is satisfiable, letting $\langle \alpha_1, \ldots, \alpha_n \rangle$ be a satisfying instantiation. We construct $G' \subseteq G$ such that $\Psi[Ag, G'] = \top$ and with $G_i \cap G' \neq \emptyset$ for each a_i, thus demonstrating that the coalition Ag is successful. We form G' as follows: if $\alpha_i = \top$ then $g_i^\top$ is included in G'; if $\alpha_i = \perp$ then $g_i^\perp$ is included in G'. In this way $|G' \cap G_i| = 1$ for each i. Now consider the the value of $\Psi[Ag, G']$. Since exactly one goal state from G_i belongs to G' it is certainly the case that the sub-expression $\wedge_{i=1}^{n}(g_i^\top \oplus g_i^\perp)[G']$ evaluates to $\top$. It is easy to see, however, that

$$(g_i^\top \vee \neg g_i^\perp)[G'] = \begin{cases} \top & \text{if } g_i^\top \in G', \text{ i.e., } \alpha_i = \top \\ \perp & \text{if } g_i^\perp \in G', \text{ i.e., } \alpha_i = \perp \end{cases}$$

so that $(g_i^\top \vee \neg g_i^\perp)[G'] = \alpha_i$ and thence $\Psi[Ag, G'] = \top$, completing the proof that Ag is a successful coalition.

We deduce that $\Phi(x_1, \ldots, x_n)$ is satisfiable if and only if Ag is a successful coalition for the QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$ and that the decision problem SC is therefore NP-hard. ∎

Noting that C being successful does not preclude agents outside C having goals satisfied by a subset G' attesting to the success of C, we also introduce the notion of a successful *selfish* coalition, as a coalition C for which there is some $G' \in V(C)$ that satisfies *only* the members of C. This suggests the following decision problem.

SELFISH SUCCESSFUL COALITION: (SSC)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$ and coalition $C \subseteq Ag$.
Answer: "Yes" if $\exists G' \subseteq G$ s.t. $\Psi[C, G'] = \top$ and for which $\forall i \in Ag, G_i \cap G' \neq \emptyset$ if and only if $i \in C$.

Example 13 *The selfish successful coalitions in Γ_1 are in fact exactly the successful ones, that is to say $\{a_1, a_2\}$, $\{a_1, a_3\}$, $\{a_2, a_3\}$. Consider for example the coalition $\{a_1, a_2\}$: the only feasible choice available to this coalition which satisfies all members is $\{g_1, g_2\}$, but this choice does not satisfy a_3.*

Corollary 14 SELFISH SUCCESSFUL COALITION *is* NP-*complete.*

Proof: Membership in NP follows using a similar algorithm to that of Theorem 12 with, however, an additional verification step to ensure that $G_i \cap G' = \emptyset$ whenever $i \notin C$. For NP-hardness exactly the same reduction from SAT applies: the coalition considered in the constructed instance comprising all agents ($C = Ag$) which, if successful, is trivially *selfishly* successful. ∎

The next problem, UNATTAINABLE GOAL SET, concerns whether there exists an *in*feasible set of goals, that would, nevertheless, satisfy each member of a given coalition. Formally, the decision problem is as follows.

UNATTAINABLE GOAL SET: (UGS)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ and coalition $C \subseteq Ag$.
Answer: "Yes" if $\exists G' \subseteq G$ s.t. $\forall i \in C$, $G_i \cap G' \neq \emptyset$ but $\Psi[C, G'] = \bot$.

Example 15 *The successful coalitions in Γ_1, (i.e., $\{a_1, a_2\}$, $\{a_1, a_3\}$, and $\{a_2, a_3\}$), have no associated unattainable goal sets. For example, considering coalition $\{a_1, a_2\}$, this coalition only has one associated goal set that would satisfy all members – $\{g_1, g_2\}$ – and this is clearly a feasible choice for this coalition. However, all other coalitions have unattainable goal sets. For example, the goal set $\{g_1, g_2, g_3\}$ would satisfy the coalition $\{a_1, a_2, a_3\}$, but this goal set is not a feasible choice for this coalition.*

Corollary 16 UNATTAINABLE GOAL SET *is* NP-*complete.*

Proof: Membership in NP follows using a similar algorithm to that of Theorem 12: non-deterministically select $G' \subseteq G$, and verify that G' satisfies each member of C but is not feasible. For NP-hardness a reduction from SAT is used in which given $\Phi(x_1, \ldots, x_n)$ the QCG Γ_Φ defined from this has Ag and G_i as defined in the proof of Theorem 12, $\Psi(Ag, G)$, however, is now given by,

$$\left(\bigwedge_{i=1}^{n} a_i\right) \ \wedge\ \neg\Phi(g_1^\top \vee \neg g_1^\bot, \ldots, g_n^\top \vee \neg g_n^\bot)$$

It is easy to see that a satisfying instantiation $\alpha = \langle \alpha_1, \ldots, \alpha_n\rangle$ leads to $G' \subseteq G$ (in which $g_i^\top \in G'$ if and only if $\alpha_i = \top$) that satisfies Ag but yields $\Psi[Ag, G'] = \bot$. On the other hand, if $G' \subseteq G$ satisfies Ag but has $\Psi[Ag, G'] = \bot$, then the instantiation $\alpha_i = \top$ (if $g_i^\top \in G'$), $\alpha_i = \bot$ (if $g_i^\top \notin G'$) must satisfy Φ so that with $C = Ag$, a goal set is unattainable if and only if it maps to a satisfying instantiation of Φ. ∎

Our reductions in the preceding results indicate that limiting instances to coalition monotonic QCGs will not render the decision problems any easier.

Corollary 17 SC^{mono}*,* SSC^{mono} *and* UGS^{mono} *are all* NP-*complete.*

Proof: Simply use identical reductions to those employed in Theorem 12, and Corollaries 14, 16 with the additional stage that the instance, $\Phi(x_1, \ldots, x_n)$ of SAT must be transformed into standard form.[9] From Theorem 9, the QCG defined via

[9] Of course, if $\Phi(x_1, \ldots, x_n)$ is in CNF this additional stage is redundant.

$\Psi(Ag, G)$ in each reduction is coalition monotonic, since the only occurrences of a_i are in the sub-formula $\wedge_{i=1}^{n} a_i$. ∎

A superficially related problem to UGS involves a set of goals G' being given as part of a problem instance, and asks whether there is *any* coalition for which G' is both feasible and satisfies every member: the problem GOAL REALISABILITY.

<u>GOAL REALISABILITY</u>: (GR)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ and $G' \subseteq G$.
Answer: "Yes" if there exists a coalition $C \subseteq Ag$ for which $\forall i \in C$, $G_i \cap G' \neq \emptyset$ and $\Psi[C, G'] = \top$.

Example 18 *In Γ_1, $\{g_1, g_2\}$ is a realisable goal set, because this set is a feasible choice for the coalition $\{a_1, a_2\}$ and satisfies every member of this coalition. The goal sets $\{g_1, g_3\}$ and $\{g_2, g_3\}$ are also realisable goal sets.*

Corollary 19 GOAL REALISABILITY *is* NP*-complete.*

Proof: Membership in NP follows easily from the following algorithm: guess a coalition C and then verify that $\Psi[C, G'] = \top$ and for each $i \in C$, $G_i \cap G' \neq \emptyset$. For hardness, we do a reduction from SAT that is very similar to that of Theorem 12. From $\Phi(x_1, \ldots, x_n)$, an instance $\langle\Gamma_\Phi, G'\rangle$ of GR is formed with $Ag = \{a_1^\top, a_1^\perp, \ldots, a_n^\top, a_n^\perp\}$, $G_i^\top = G_i^\perp = \{g_i\}$, $G' = G = \cup_{i=1}^{n}\{g_i\}$. We then define Ψ as

$$\Phi(g_1(a_1^\top \vee \neg a_1^\perp), \ldots, g_n(a_n^\top \vee \neg a_n^\perp)) \wedge \bigwedge_{i=1}^{n}(a_i^\top \oplus a_i^\perp)$$

so that, by a similar argument to that of Theorem 12, Φ is satisfiable if and only if there is a coalition C (which will contain exactly n members) to realise the set of goals $G' = G = \{g_1, \ldots, g_n\}$. ∎

In contrast to the cases SC, SSC, and UGS whose complexity is unaffected by restricting to coalition monotonic QCGs, for GR^{mono} this variation becomes significantly easier.

Corollary 20 GR^{mono} *is polynomial-time decidable.*

Proof: Let $\Gamma = \langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ and $G' \subseteq G$ form an instance of GR^{mono}, recalling our convention that $\Psi(Ag, G)$ is presented in accordance with the strictures given by Theorem 9. For each $g_j \in G'$ let $C_j \subseteq Ag$ be the set $\{a_i \ : \ g_j \in G_i\}$, i.e., the set of agents for whom g_j is a goal with which they would be satisfied. Consider the coalition $C = \cup_{g_j \in G'} C_j$. The coalition C is satisfied by G' and is the maximal such subset of Ag. We claim that $\langle\Gamma, G'\rangle$ is accepted as an instance of GR^{mono} if and only if $\Psi[C, G'] = \top$. It is obviously the case that if $\Psi[C, G'] = \top$ then G' is both feasible for and satisfies C, i.e if $\Psi[C, G'] = \top$ then $\langle \Gamma, G' \rangle$ is accepted. On the other hand suppose that $\langle \Gamma, G' \rangle$ is a positive instance of GR^{mono} and let C' be any coalition witnessing this fact, that is C' is satisfied by G' and $\Psi[C', G'] = \top$. It must hold that $C' \subseteq C$ for otherwise there is some $a_k \in C'$ that does not occur in C. Any such agent, however, cannot have its goals (G_k) met by any goal in G': from the construction of C, $G_k \cap G' \neq \emptyset$ would give $a_k \in C$. Now since Γ is coalition monotonic it follows that $\Psi[C', G'] = \top$ would yield $\Psi[C, G'] = \top$ as required. It remains only to observe that the coalition $C = \cup_{g_j \in G'} C_j$ can be formed from the instance $\langle \Gamma, G' \rangle$ efficiently and that the test $\Psi[C, G'] = \top$ simply involves evaluating a propositional formula in standard form. ∎

We say a coalition is *minimal* if no strict subset of this coalition is successful.

Example 21 *Returning again to Γ_1, coalition $\{a_1, a_3\}$ is minimal.*

The notion of minimality is important because it implies a kind of *internal stability* for a coalition (cf. [22, p.281]). That is, in a minimal coalition, there is no incentive for subsets of the coalition to defect away from the coalition, as, by definition, such sub-coalitions cannot be successful.

MINIMAL COALITION: (MC)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$ and coalition $C \subseteq Ag$.
Answer: "Yes" if $\forall C' \subset C, \forall G' \subseteq G$, if $\forall i \in C'\ G' \cap G_i \neq \emptyset$, then $\Psi[C', G'] = \bot$

Theorem 22 MINIMAL COALITION *is co-*NP*-complete.*

Proof: For membership in co-NP, simply verify that for each $C' \subset C$ and each $G' \subseteq G$, that $\Psi[C', G'] = \top$ implies $\exists i \in C' : G_i \cap G' = \emptyset$. Clearly, this can be done via a co-NP computation.

To show that the problem is co-NP-hard we employ a reduction from UNSAT. Given an instance $\Phi(x_1, \ldots, x_n)$ of UNSAT, we form an instance $\langle \Gamma_\Phi, C \rangle$ of MC by setting $Ag = \{a_1, \ldots, a_n, a_{n+1}\}$, $G_i = \{g_i\}$ and Ψ defined by

$$\Phi(a_1, \ldots, a_n) \wedge \bigwedge_{i=1}^{n+1} (a_i \Rightarrow g_i)$$

We prove that the coalition $C = Ag$ is a minimal coalition for Γ_Φ if and only if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable.

($\Rightarrow$) Assume that $\langle \alpha_1, \ldots, \alpha_n \rangle$ satisfies $\Phi(x_1, \ldots, x_n)$. Consider the coalition, $C' \subset$

$C = Ag$ defined through

$$C' = \begin{cases} \{a_i : \alpha_i = \top\} & \text{if } \exists i \text{ with } \alpha_i = \top \\ \{a_{n+1}\} & \text{if } \forall i\ \alpha_i = \bot \end{cases}$$

In addition consider the goal set $G' = \{g_i : 1 \leq i \leq n+1 \text{ and } a_i \in C'\}$. We have $C' \subset C = Ag$, since in the former case $a_{n+1} \notin C'$ and, in the latter, only a_{n+1} is in C'. The construction of Ψ ensures that the sub-expression $\wedge_{i=1}^{n+1}(a_i \Rightarrow g_i)$ within $\Psi[C', G']$ evaluates to $\top$. Similarly the expression $\Phi(a_1, \ldots, a_n)[C]$ takes the value $\Phi(\alpha_1, \ldots, \alpha_n)$, i.e., $\top$. Thus if $\Phi(x_1, \ldots, x_n)$ is satisfiable then Ag is not a minimal coalition, and thence if Ag *is* minimal then Φ must be unsatisfiable.

($\Leftarrow$) Assume that $C' \subset Ag$, $G' \subseteq G$ are such that $\Psi[C', G'] = \top$ with $G_i \cap G' \neq \emptyset$ for each $a_i \in C'$, i.e., that $C = Ag$ is *not* a minimal coalition. It is clearly the case that setting $x_i = \top$ if $a_i \in C'$, and $x_i = \bot$ if $a_i \notin C'$ must define a satisfying instantiation of $\Phi(x_1, \ldots, x_n)$ since this is exactly the instantiation under which the sub-expression $\Phi(a_1, \ldots, a_n)$ of Ψ is satisfied. Thus if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable then $C = Ag$ must be a minimal coalition.

We deduce that $\Phi(x_1, \ldots, x_n)$ is unsatisfiable if and only if Ag is a minimal coalition, whence it follows that MC is co-NP-hard as claimed. ∎

Corollary 23 MC^{mono} *is co-*NP*–complete.*

Proof: Membership in co-NP is via an identical argument to that of Theorem 22. To prove co-NP–hardness we again use a reduction from UNSAT. Given $\Phi(x_1, \ldots, x_n)$ an instance of UNSAT, form $\langle \Gamma_\Phi, C\rangle$ an instance of MC^{mono} as follows. Set $Ag = \{a_1, \ldots, a_n\}$, $G_i = \{g_i\}$, $G = \cup_{i=1}^{n} G_i$. and

$$\Psi(Ag, G) = \Phi(g_1, \ldots, g_n) \vee \left(\bigwedge_{i=1}^{n} a_i g_i\right)$$

Finally to complete the instance of MC^{mono} we set $C = Ag$. It is obviously the case for any $G' \subseteq G$, that G' satisfies the empty coalition. It follows that, $C = Ag$ can be a *minimal* coalition if and only if $\Psi[\emptyset, G'] = \bot$ for all $G' \subseteq G$, i.e. if and only if $\Phi(g_1, \ldots, g_n)$ is unsatisfiable. ∎

As we noted above, saying that C is minimal does *not* imply that it is successful. So, although a minimal coalition has a kind of internal stability, there may in fact be no reason for such a coalition to form, as it may not be successful. Thus, as attempts to characterise stable coalitions, neither success nor minimality seem entirely appropriate, although both seem to have relevant aspects. This motivates us to introduce a notion of stability based on the concept of the core.

4.3 *The Core*

We noted in our introduction that a key question facing any naturally or artificially rational agent is: Which coalition should I join? In our setting, the first pre-requisite that any potential coalition should satisfy is that it is successful. For otherwise, such a coalition would not form: the unsuccessful members of the coalition would look elsewhere for a coalition that satisfied them. But this is only part of the story. A successful coalition might still not be stable; it may be subject to defections. An agent might have an incentive to defect from a coalition C if it could collude with others members to form a strict subset of C that could (feasibly) satisfy its goals.

This motivates the definition of a solution concept that loosely corresponds to the core in regular coalitional games [22, pp.257–274]. Intuitively, the core of a coalition is the set of feasible choices for that coalition from which the members of that coalition have no incentive to deviate. More formally, we say a set of goals G' is in the core of a coalition C if C is minimal, G' is feasible for C, and in addition G' satisfies every member of C. The core of a coalition will thus be non-empty if that coalition is both minimal and successful. Notice that this definition, when applied to the grand coalition, implies that the grand coalition is the *uniquely* successful coalition, and thus is the only coalition that a rational agent would choose to join.

It may be argued that, formulated in this way, our concept of the core is weaker than the definition of the core when presented in terms of conventional coalitional games. In the conventional formulation, a coalition is viewed as being stable if no subset of this coalition could break away to achieve an outcome that they all *strictly* preferred. Of course, the notion of "strict preference" does not really exist in QCGs as formulated here: the only preferences an agent has are with respect to outcomes that satisfy one of its goals over those that do not satisfy any. But, in this way, the notion of stability as captured in the QCG definition of the core is *stronger* than the conventional notion: it ensures that, not only will a subset of the coalition do *no better* by defecting, but they will in fact do *worse* – because, by definition, they will not be successful. One might ask why this is important: why are we concerned about a subset breaking away to achieve an outcome that was no better than that which they could achieve in the larger coalition. The answer is that, if the subset coalition is successful, then there is *no incentive* for the subset to stay with the larger coalition: there is nothing impelling them to stay with the larger coalition, and in this sense, the larger coalition can be said to be unstable. In our formulation, however, we provide such an incentive: a subset will not deviate because they will be guaranteed to do worse – they will be unsuccessful.

The first decision problem we study associated with the core is that of determining whether, given a particular goal set for a particular coalition, that goal set is in the core of that coalition. (As an aside, note that the core is usually framed in terms of the grand coalition, i.e., the set of all agents – in the interests of generality, we define the concept in terms of an arbitrary coalition, but of course our results hold true for the grand coalition.)

<u>CORE MEMBERSHIP</u>: (CM)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$, coalition $C \subseteq Ag$, and goal set $G' \subseteq G$.
Answer: "Yes" if G' is in the core of C, i.e., $\Psi[C, G'] = \top$, $\forall i \in C\ G_i \cap G' \neq \emptyset$ *and* $\forall C' \subset C, \forall G'' \subseteq G$ if $\forall i \in C', G_i \cap G'' \neq \emptyset$ then $\Psi[C', G''] = \bot$.

Example 24 *With respect to* QCG Γ_1, *the goal set* $\{g_1, g_2\}$ *is in the core of the coalition* $\{a_1, a_2\}$, *as (i) the coalition* $\{a_1, a_2\}$ *is minimal; (ii)* $\{g_1, g_2\}$ *is a feasible choice for* $\{a_1, a_2\}$*; and (iii)* $\{g_1, g_2\}$ *satisfies every member of* $\{a_1, a_2\}$.

Theorem 25 CORE MEMBERSHIP *is co-*NP*-complete.*

Proof: For membership in co-NP, having confirmed that $\Psi[C, G'] = \top$ and $G_i \cap G' \neq \emptyset$ for all $i \in C$, it then suffices to verify that for each $C' \subset C$ and $G'' \subseteq G$ it holds that if $(\forall\ i \in C' : G_i \cap G'' \neq \emptyset)$ then $(\Psi[C', G''] = \bot)$, a test which for any fixed C and G'' can be performed in deterministic polynomial-time.

To prove co-NP-hardness we use a reduction from UNSAT. Given an instance $\Phi(x_1, \ldots, x_n)$ of the latter, form an instance $\langle \Gamma_\Phi, C, G' \rangle$ of CM with $Ag = \{a_1, \ldots, a_{n+1}\}$, $G_i = \{g_i\}$, $G = \cup_{i=1}^{n+1} G_i$, and Ψ defined as

$$\left(\Phi(a_1, \ldots, a_n) \wedge \bigwedge_{i=1}^{n+1} (a_i \Rightarrow g_i) \wedge (\neg a_{n+1})\right) \vee \left(\bigwedge_{i=1}^{n+1} (a_i g_i)\right)$$

Finally we set $C = Ag$ and $G' = G$. By arguments similar to those in the proof of Theorem 22 this instance of CM is accepted if and only if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable. ∎

Corollary 26 CM^{mono} *is co-*NP*–complete.*

Proof: Given, $\Phi(X_n)$ an instance of UNSAT form the coalition monotonic QCG, Γ_Φ which has $Ag = \{a_1, \ldots, a_n\}$, $G_i = \{g_i\}$, $G = \cup_{i=1}^{n} G_i$ and $\Psi(Ag, G)$ defined via

$$\Psi(Ag, G) \ = \ \left(\bigwedge_{i=1}^{n} a_i g_i\right) \ \vee \ \Phi(g_1, \ldots, g_n)$$

To complete the instance of CM^{mono} we set $C = Ag$ and $G' = G$. We then have that G' is in the core of Ag if and only if the empty coalition is unsuccessful: should some goal set $G'' \subseteq G$ be such that $\Psi[\emptyset, G'']$ then we have $\Phi[G''] = \top$. Hence $\Phi(X_n)$ is unsatisfiable if and only G' is in the core of Ag in the coalition monotonic

QCG Γ_Φ. ∎

The next problem we consider is that of whether the core of a given coalition is non-empty.

CORE NON-EMPTY: (CNE)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ and coalition $C \subseteq Ag$.
Answer: "Yes" if C is both successful and minimal.

Theorem 27 CORE NON-EMPTY *is D^p-complete.*

Proof: To establish membership in D^p, we must exhibit two languages, L_1 and L_2, such that: (i) $L_1 \in$ NP; (ii) $L_2 \in$ co-NP; and (iii) CORE NON-EMPTY $= L_1 \cap L_2$ [23, pp.412–415]. Define L_1 by

$$L_1 = \{x : \text{SUCCESSFUL COALITION}(x)\}$$

By Theorem 12, $L_1 \in$ NP as desired. Define L_2 by

$$L_2 = \{x : \text{MINIMAL COALITION}(x)\}$$

By Theorem 22, $L_2 \in$ co-NP as desired. It only remains to note that by definition,

$$L_1 \cap L_2 = \{x : \text{CORE NON-EMPTY}(x)\}.$$

Thus CORE NON-EMPTY is in D^p.

To complete the proof, we must show that the problem is D^p-complete. Note that we cannot deduce D^p-completeness by arguing "L_1 is NP-complete and L_2 is co-NP-complete implies $L_1 \cap L_2$ is D^p-complete." An easy counterexample to such a claim is given by taking $L_1 = \{x : \text{SAT}(x)\}$ and $L_2 = \{x : \text{UNSAT}(x)\}$ so that $L_1 \cap L_2 = \emptyset$. We thus we give a reduction from the decision problem SAT-UNSAT [23, p.415], instances of which comprise a pair $\langle \Phi_1(x_1, \ldots, x_n), \Phi_2(x_1, \ldots, x_n)\rangle$ of propositional formulae. An instance is accepted if $\Phi_1(x_1, \ldots, x_n)$ is satisfiable *and* $\Phi_2(x_1, \ldots, x_n)$ is unsatisfiable. D^p-completeness of SAT-UNSAT was proved in [24].

Let $\langle \Phi_1(x_1, \ldots, x_n), \Phi_2(x_1, \ldots, x_n)\rangle$ be an instance of SAT-UNSAT. The instance $\langle \Gamma_{\langle \Phi_1, \Phi_2\rangle}, C\rangle$ of CNE we construct from this has $Ag = \{a_1, \ldots, a_n, a_{n+1}\}$, $G_i = \{g_i^\top, g_i^\bot\}$, and $G = \cup_{i=1}^{n+1} G_i$. We define Ψ as

$$\left(\bigwedge_{i=1}^{n+1} a_i \;\Rightarrow\; \Psi_1\right) \wedge \left(\bigvee_{i=1}^{n+1} \neg a_i \;\Rightarrow \Psi_2\right)$$

in which Ψ_1 is

$$\Phi_1((g_1^\top \vee \neg g_1^\perp), \ldots, (g_n^\top \vee \neg g_n^\perp)) \wedge g_{n+1}^\top \wedge \bigwedge_{i=1}^{n+1} (g_i^\top \oplus g_i^\perp)$$

and Ψ_2 is

$$\Phi_2(a_1, \ldots, a_n) \wedge \bigwedge_{i=1}^{n+1} (a_i \Rightarrow g_i^\top)$$

Finally we set $C = Ag$ to form the instance of CNE.

We now claim that C in the instance $\langle\langle G, Ag, G_1, \ldots, G_{n+1}, \Psi\rangle, C\rangle$ described is a minimal successful coalition if and only if Φ_1 is satisfiable *and* Φ_2 is unsatisfiable. Using an argument similar to that in the proof of Theorem 12 we see that $C = Ag$ is successful if and only if $\Phi_1(x_1, \ldots, x_n)$ is satisfiable. In the same way, following the argument used in proving Theorem 22, C is *minimal* if and only if $\Phi_2(x_1, \ldots, x_n)$ is unsatisfiable. We note that for $C = Ag$ the formula defining $\Psi[C, G']$ is equivalent to

$$\Phi_1((g_1^\top \vee \neg g_1^\perp), \ldots, (g_n^\top \vee \neg g_n^\perp)) \wedge \; g_{n+1}^\top \; \wedge \bigwedge_{i=1}^{n+1} (g_i^\top \oplus g_i^\perp)[C, G']$$

whereas for $C \subset Ag$, $\Psi[C, G']$ reduces to

$$\Phi_2(a_1, \ldots, a_n) \wedge \bigwedge_{i=1}^{n+1} (a_i \Rightarrow g_i^\top)[C, G']$$

which are the formulae constructed in the proofs of Theorems 12 and 22 respectively. ∎

Corollary 28 CNE^{mono} *is* D^p*-complete.*

Proof: Given $\langle\Phi_1, \Phi_2\rangle$ an instance of SAT-UNSAT, form the instance $\langle\Gamma_\Phi, C\rangle$ of CNE^{mono} with $Ag = \{a_1, \ldots, a_n\}$, $G_i = \{g_i, h_i\}$ and

$$\Psi(Ag, G) \;=\; \left(\bigwedge_{i=1}^{n} a_i g_i\right) \wedge \Phi_1(h_1, \ldots, h_n) \;\vee\; \Phi_2(g_1, \ldots, g_n)$$

Finally we set $C = Ag$. Now suppose that Φ_1 is satisfiable, then C is successful with respect to some goal set G' with $\{g_1, \ldots, g_n\} \subseteq G'$; similarly if Φ_2 is unsatisfiable, then Ag is minimal as $C = \emptyset$ does not succeed. On the other hand if Ag is both minimal and successful then we find a satisfying instantiation for Φ_1 via the set $G' \cap \{h_1, \ldots, h_n\}$ with $\Psi[Ag, G'] = \top$. In addition, since Ag is minimal it follows that $\Psi[\emptyset, G'] = \bot$ for all $G' \subseteq G$, from which it is deduced that Φ_2 must be unsatisfiable. ∎

4.4 Dependencies Between Agents: Veto Players

Suppose we are given two agents $i, j \in Ag$, and asked whether i is *essential* for the accomplishment of j's goals; that is, whether i is a member of every coalition that can satisfy one of j's goals (this is a useful generalisation of the notion of veto player in conventional coalitional games [22, p.261]). If this is the case, then j could prevent i from accomplishing its goals: we say that i is a *veto player* for j. Formally, i is a veto player for j if for all $C \subseteq Ag$ and $G' \in V(C)$, if $G' \cap G_j \neq \emptyset$ then $i \in C$. It should be noted that j need not be a member of C.

Example 29 *With respect to* QCG Γ_1, *no member of Ag is a veto player for any other agent. However, consider the* QCG Γ_2, *with* $Ag = \{a_1, a_2, a_3\}$, $G = \{g_1, g_2, g_3\}$, $G_i = \{g_i\}$, *and characteristic function V defined as follows:*

$$V(C) = \begin{cases} \{g_2\} & \textit{if } C = \{a_1\} \\ \emptyset & \textit{if } C = \{a_2\} \\ \{g_3\} & \textit{if } C = \{a_1, a_2\} \\ \{g_1, g_2, g_3\} & \textit{otherwise.} \end{cases}$$

In this QCG, *agent a_3 is a veto player for agents a_1 and a_2 (as well as itself). There are no other veto players in this* QCG.

The notion of a veto player is important because it characterises a *dependency* between agents: if you are a veto player for me, then I am dependent upon you for the successful accomplishment of my goals. There is no way that they can be achieved without your cooperation, because you must be present in any coalition that is capable of bringing about my goals. This notion of a veto player is related to the concept of pivot players and power indices (such as the Shapley-Shubik index) as these concepts are studied in conventional coalitional games [39, pp.85–89]. A pivot player is a player than can transform a non-winning coalition into a winning one. The Shapley-Shubik index attempts to measure the "power" that an agent wields in a cooperative game, by considering the extent to which this agent can transform a non-winning coalition into a winning one.

Of course, just because agent i is a veto player for agent j, this does not mean that i will necessarily exercise this implicit power, or exploit the dependency. However, the dependency may well influence the strategic reasoning of the agents involved. For example, knowing that it is dependent upon i for its satisfaction, agent j might prefer to enter coalitions that also satisfied i (because choosing otherwise might give i reason to exercise its power by not satisfying agent j).

The first decision problem associated with veto players is as follows.

VETO PLAYER: (VP)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$ and agents $veto \in Ag$ and $sat \in Ag$.
Answer: "Yes" if for all $C \subseteq Ag$, $G' \subseteq G$, if $G_{sat} \cap G' \neq \emptyset$ and $\Psi[C, G'] = \top$ then $veto \in C$.

Theorem 30 VETO PLAYER *is co-*NP*-complete.*

Proof: For membership in co-NP we simply have to verify for each $C \subseteq Ag$ and $G' \subseteq G$ that

$$(\Psi[C, G'] \text{ and } (G_{sat} \cap G' \neq \emptyset)) \text{ implies } (veto \in C)$$

To show that the problem is co-NP-hard, we can reduce UNSAT to VP as follows. Given $\Phi(x_1, \ldots, x_n)$ an instance of UNSAT, the instance $\langle \Gamma_\Phi, veto, sat \rangle$ of VP formed has $Ag = \{a_1, \ldots, a_n\} \cup \{v, w\}$; $G_i = \{g_i\}$, and Ψ given by

$$\neg\Phi(a_1, \ldots, a_n) \;\Rightarrow\; (v \wedge w \wedge g_w)$$

Finally we designate *veto* as the agent v and *sat* as the agent w. It is now easy to show that that $\Phi(x_1, \ldots, x_n)$ is unsatisfiable if and only if every coalition $C \subseteq Ag$ and set of goals G' for which $\Psi[C, G']$ holds has the property that should $g_w \in G'$ – i.e., G' satisfies w – then v is in the coalition C. ∎

The notion of a *necessary goal* is a natural counterpart to that of veto players. The idea is that a goal set will be necessary if this goal set is a "side effect" of any coalition achieving their goals. More formally, G' is necessary if, whenever G'' is a feasible choice of some coalition, which also satisfies every member of this coalition, then $G' \subseteq G''$. Thus if G' is necessary, then whichever choice a coalition makes that satisfies all its members, this choice will include G'. The decision problem is as follows.

NECESSARY GOAL: (NG)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$ and $G' \subseteq G$.
Answer: "Yes" if for every coalition $C \subseteq Ag$ and $G'' \subseteq G$: if $\forall i \in C$, $G_i \cap G'' \neq \emptyset$ and $\Psi[C, G''] = \top$ then $G' \subseteq G''$.

Example 31 *The* QCG Γ_1 *has no necessary goals. But consider a* QCG Γ_3 *with* $Ag = \{a_1, a_2\}$*,* $\Gamma = \{g_1, g_2, g_3\}$*,* $\Gamma_1 = \{g_1, g_2\}$*,* $\Gamma_2 = \{g_3\}$*, and characteristic* *function V defined as follows:*

$$V(C) = \begin{cases} \{\{g_1, g_3\}, \{g_3\}\} & \textit{if } C = \{a_1\} \\ \{\{g_2, g_3\}\} & \textit{if } C = \{a_2\} \\ \{\{g_2, g_3\}\} & \textit{if } C = \{a_1, a_2\} \end{cases}$$

In this QCG*, goal g_3 is necessary, but neither g_1 nor g_2 are necessary.*

Corollary 32 NECESSARY GOAL *is co-*NP*-complete.*

Proof: Membership is co-NP is immediate from the algorithm which tests for each $C \subseteq Ag$ and $G'' \subseteq G$ whether

$$\left[\left(\bigwedge_{i \in C} (G_i \cap G'' \neq \emptyset)\right) \text{ and } \Psi[C, G'']\right] \text{ implies } (G' \subseteq G'')$$

That NG is co-NP-hard follows from a reduction from UNSAT whereby given $\Phi(x_1, \ldots, x_n)$ an instance $\langle \Gamma_\Phi, G' \rangle$ of NG is created in which $Ag = \{a_1, \ldots, a_n\}$, $G_i = \{g_i, g_{n+1}\}$, $G = \cup_{i=1}^n G_i$, $G' = \{g_{n+1}\}$, and Ψ defined as

$$(\neg \Phi(a_1, \ldots, a_n)) \Rightarrow \left(g_{n+1} \wedge \bigwedge_{i=1}^{n} \neg g_i\right)$$

With this construction, $\{g_{n+1}\}$ is a necessary goal set for $\langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$ if and only if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable. ∎

If you are a veto player for me, then this might appear to put me in a weak position with respect to you – because I am absolutely reliant upon you for the satisfaction of my goals. Unless, of course, the situation is reciprocal: that is, unless *you* are also dependent upon *me* in return. This gives rise to the notion of a *mutually dependent* coalition, in which everybody is dependent upon everybody else. Formally, coalition C will be mutually dependent if:

$$\forall C' \subseteq Ag, \forall G' \in V(C') \text{ if } G' \text{ satisfies at least one member of } C \text{ then } C \subseteq C'$$

Notice that saying that C are mutually dependent implies that C are a necessary component of any coalition to achieve their goals; but it does not say that they are a *successful* coalition. They may not be able to cooperate so as to satisfy their goals jointly, or they may require the cooperation of other agents to achieve all their goals.

MUTUAL DEPENDENCE: (MD)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ and coalition $C \subseteq Ag$.
Answer: "Yes" if for all $C' \subseteq Ag$, and $G' \subseteq G$, if G' satisfies at least one member of C and $\Psi[C', G'] = \top$ then $C \subseteq C'$.

Theorem 33 MUTUAL DEPENDENCE *is co-*NP*-complete.*

Proof: Membership follows easily by verifying for each $C' \subseteq Ag$ and $G' \subseteq G$ that if $\Psi[C', G'] = \top$ and $\exists i \in C : G_i \cap G' \neq \emptyset$ then $C \subseteq C'$.

Hardness is proved by a reduction from UNSAT. Given $\Phi(x_1, \ldots, x_n)$ we form $\langle \Gamma_\Phi, C\rangle$ by setting $Ag = \{a_1, \ldots, a_n, a_{n+1}\}$, $G_i = \{g_i\}$, and Ψ defined as

$$\neg\Phi(a_1, \ldots, a_n) \;\Rightarrow\; (a_{n+1} g_{n+1}).$$

Finally, fix $C = \{a_{n+1}\}$. By arguments similar to those already seen, it can be shown that $\Phi(x_1, \ldots, x_n)$ is unsatisfiable if and only if for any coalition C' and goal set G' for which $\Psi[C', G'] = \top$, should it be the case that $g_{n+1} \in G'$ – i.e., the member a_{n+1} of the coalition $C = \{a_{n+1}\}$ is satisfied – then $\{a_{n+1}\} = C \subseteq C'$. ∎

4.5 General Properties of Qualitative Coalitional Games

Finally, we consider some general properties of QCGs, that is, properties which are considered independently of any coalition or set of goals. The first such property we will consider is whether a QCG is *empty*. A QCG is said to be empty if it contains no successful coalition. If a QCG is empty, then intuitively, there is no point in *any* coalition forming, as no coalition can succeed in satisfying the goals of all its members. To put it another way, in any coalition, there will always be at least one agent with unsatisfied goals, and hence any coalition we might care to consider will be inherently unstable. The decision problem is as follows.

EMPTY GAME: (EG)
Instance: QCG $\Gamma = \langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$.
Answer: "Yes" if no coalition in Γ is successful.

Example 34 *The* QCG*s Γ_1–Γ_3, introduced in examples above, are all non-empty. But consider* QCG *Γ_4, with $Ag = \{a_1, a_2, a_3\}$, $G = \{g_1, g_2, g_3\}$, $G_i = \{g_i\}$, and V defined as follows.*

$$V(C) = \{g_i : i \in Ag, i \notin C\}$$

In other words, a coalition in Γ_4 can achieve only the goals of agents that are not members. It follows that Γ_4 is empty.

Theorem 35 EMPTY GAME *is co-*NP*-complete.*

Proof: For membership in co-NP it suffices to use the co-NP algorithm which verifies for every choice of $C \subseteq Ag$ and $G' \subseteq G$ that if $\forall i \in C : G_i \cap G' \neq \emptyset$ then $\Psi[C, G'] = \bot$

To show the problem is co-NP-hard, we give a reduction from UNSAT. Given an instance $\Phi(x_1, \ldots, x_n)$ of UNSAT, the instance Γ_Φ of EG has $Ag = \{a_1, \ldots, a_n, a_{n+1}\}$, $G_i = \{g_i\}$, and Ψ defined as:

$$\Phi(a_1, \ldots, a_n) \wedge a_{n+1} \wedge \bigwedge_{i=1}^{n+1} (a_i \Rightarrow g_i)$$

By a similar argument to that of Theorem 22 it follows that the QCG so defined is empty if and only if $\Phi(x_1, \ldots, x_n)$ is unsatisfiable. ∎

Corollary 36 EG^{mono} *is co-*NP*–complete.*

Proof: Membership in co-NP is immediate from the argument of Theorem 35. To prove co-NP–hardness, use a reduction from UNSAT wherein an instance $\Phi(x_1, \ldots, x_n)$ is mapped to a coalition monotonic QCG, Γ_Φ, having $Ag = \{a_1, \ldots, a_n\}$, $G_i = \{g_i, h_i\}$ and $\Psi(Ag, G)$ as

$$\left(\bigwedge_{i=1}^{n} a_i\right) \Phi(g_1, \ldots, g_n) \;\wedge\; \bigwedge_{i=1}^{n} (g_i \vee \neg h_i)(\neg g_i \vee h_i)$$

The only coalition that *could* be successful is $C = Ag$. In order for this to happen, some subset G' of G must give $\Psi[Ag, G'] = \top$, i.e., $\Phi(x_1, \ldots, x_n)$ is satisfiable. On the other hand *if* $\Phi(x_1, \ldots, x_n)$ is satisfied by α then Ag is successful using the subset G_α of G defined by $\{g_i \;:\; \alpha_i = \top\} \cup \{h_i \;:\; \alpha_i = \bot\}$. We note that the subset $\{h_1, \ldots, h_n\}$ of G in the argument above is used in the term $\wedge_{i=1}^{n} (g_i \vee \neg h_i)(\neg g_i \vee h_i)$ to ensure that if a satisfying instantiation of $\Phi(x_1, \ldots, x_n)$ requires x_i to be assigned the value $\bot$ (so that g_i would not be one of the goals that a_i could achieve) then in order for a_i to be satisfied and $\Psi[Ag, G']$ to be $\top$ requires $h_i = \top$ so that $(g_i \vee \neg h_i)(\neg g_i \vee h_i)[Ag, G'] = \top$. ∎

An obvious companion problem is whether or not a QCG is *trivial*: that is, whether *every* coalition is successful. If an agent discovered that the QCG it was playing was trivial, it might regard this as good news, as it certainly means that it can satisfy its goals. However, in trivial QCGs, there is no incentive for anything other singleton coalitions to form: why should an agent join a larger coalition when it can accomplish its goals in isolation? In this sense, the concept of a trivial game is rather like that of an *inessential game* in cooperative game theory, as it implies that every agent can do as well on its own as in any larger coalition [20, p.154] [10] . The decision problem is as follows.

TRIVIAL GAME: (TG)
Instance: QCG $\Gamma = \langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$.
Answer: "Yes" if every coalition in Γ is successful.

Example 37 *Consider* QCG Γ_5*, with* $Ag = \{a_1, a_2, a_3\}$*,* $G = \{g_1, g_2, g_3, g_4\}$*,* $G_i = \{g_i\}$*, and V defined as follows.*

$$V(C) = \{g_i : i \in C\}$$

In this QCG*, any coalition will be successful.*

Theorem 38 TRIVIAL GAME *is* Π_2^p*-complete.*

Proof: Membership in Π_2^p is immediate by observing that $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ is trivial if and only if

$$\forall C \subseteq Ag\ \exists\, G' \subseteq G : (\forall i \in C\ G_i \cap G' \neq \emptyset \text{ and } \Psi[C, G'] = \top)$$

the conditions on C and G' being easily verifiable by a deterministic polynomial-time algorithm.

To prove Π_2^p-hardness, we use a reduction from QSAT$_2^\Pi$. We recall that an instance of QSAT$_2^\Pi$ is a propositional formula $\Phi(x_1, \ldots, x_n, y_1, \ldots, y_m)$ defined over two disjoint sets of variables $X = \langle x_1, \ldots, x_n\rangle$ and $Y = \langle y_1, \ldots, y_m\rangle$ and, without loss of generality, it may be assumed that $m = n$. An instance is accepted if $\forall\alpha_X\exists\beta_Y\ \Phi(\alpha_X, \beta_Y) = \top$, i.e., no matter how the variables X are instantiated (α_X), there is some assignment (β_Y) to the variables Y that will result in Φ being satisfied. Given an instance $\Phi(X, Y)$ of QSAT$_2^\Pi$ with $X = \langle x_1, \ldots, x_n\rangle$ and $Y = \langle y_1, \ldots, y_n\rangle$, we form an instance Γ_Φ of TG by setting $Ag = \{a_1, \ldots, a_n\}$, $G_i = \{g_i^\top, g_i^\perp\}$, $G = \cup_{i=1}^n G_i$, and Ψ to be

$$\Phi(a_1, \ldots, a_n, g_1^\top \vee \neg g_1^\perp, \ldots, g_i^\top \vee \neg g_i^\perp, \ldots, g_n^\top \vee \neg g_n^\perp) \wedge \bigwedge_{i=1}^n (g_i^\top \oplus g_i^\perp)$$

[10] A conventional coalitional game $\langle Ag, \nu : 2^{Ag} \to \mathbb{R}\rangle$ is inessential if for all $C \subseteq Ag$.

$$\nu(C) = \sum_{i \in C} \nu(\{i\})$$

We claim that the instance $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ so formed is a positive instance of TG if and only if $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Π}. To see this, first note that any $C \subseteq Ag$ may be associated with a unique instantiation α_C of X_n via $a_i \in C$ if and only if $\alpha_i = \top$. It therefore suffices to show that given any $C \subseteq Ag$ and successfully realised goal set $G' \subseteq G$ we can construct an instantiation β of Y_n under which $\Phi(\alpha_C, \beta) = \top$; similarly given any instantiation α_C of X_n and associated instantiation β of Y_n under which $\Phi(\alpha_C, \beta) = \top$ we may construct a goal set $G' \subseteq G$ that the associated coalition C realises successfully.

($\Rightarrow$) Assume that $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ is a positive instance of TG and let $C \subseteq Ag$ be any coalition with $G' \subseteq G$ an associated set of goals for which this coalition succeeds. For α_C the instantiation of X_n defined via C we form an instantiation $\beta = \langle\beta_1, \ldots, \beta_n\rangle$ of Y_n via,

$$\beta_i = \begin{cases} \top & \text{if } \ g_i^{\top} \in G' \\ \bot & \text{if } \ g_i^{\bot} \in G' \end{cases}$$

Noting that since $\Psi[C, G'] = \top$ so the term $\wedge_{i=1}^{n}(g_i^{\top} \oplus g_i^{\bot})[G'] = \top$ and therefore this instantiation of Y_n is well defined, it follows that

$$\Phi(\alpha_C, \beta) \; = \; \Psi[C, G'] \; = \; \top$$

Thus, if $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ is a positive instance of TG, then $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Π}.

($\Leftarrow$) Suppose that $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Π} and let $\alpha = \langle\alpha_1, \ldots, \alpha_n\rangle$ be any instantiation of X_n with $\beta = \langle\beta_1, \ldots, \beta_n\rangle$ an instantiation of Y_n for which $\Phi(\alpha, \beta) = \top$. For C_α the coalition corresponding to α, consider the subset G' of G given by: for each $a_i \in C_\alpha$, $g_i^{\top} \in G'$ if $y_i = \top$; $g_i^{\bot} \in G'$ if $y_i = \bot$. With this selection we get

$$\Psi[C_\alpha, G'] \; = \; \Phi(\alpha, \beta) = \top$$

and for each $a_i \in C_\alpha$ exactly one of its goals in $\{g_i^{\top}, g_i^{\bot}\}$ is satisfied. It follows that if $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Π} then $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ is a positive instance of TG.

Hence the problem is Π_2^p-hard, which completes the proof that TG is Π_2^p-complete. ∎

If one considers TG^{mono} one might expect this restriction to result in a decision problem whose complexity is "no worse" than NP: simply test if the empty coalition, $C = \emptyset$, is successful, in which event any G' for which $\Psi[\emptyset, G'] = \top$ will, by virtue of monotonicity, be such that $\Psi[C, G'] = \top$ for all $C \subseteq Ag$. It turns out that this expectation is ill-founded, and that TG^{mono} is no easier than TG: the argument just outlined to move the complexity of TG^{mono} down to NP ignores the fact that successful coalitions must identify goal sets which are both feasible *and* satisfy each member. The latter requirement means that we cannot deduce success of *all* coalitions from that of the empty coalition alone. The fact that $\Psi[\emptyset, G'] = \top$ implies $\Psi[C, G'] = \top$ necessitates the use of a rather more subtle reduction in order to establish our next result. Namely,

Corollary 39 $\textsc{tg}^{mono}$ *is* Π_2^p*-complete.*

Proof: That $\textsc{tg}^{mono} \in \Pi_2^p$ is immediate via the same argument as that demonstrating $\textsc{tg} \in \Pi_2^p$. To establish $\textsc{tg}^{mono}$ is Π_2^p-hard, we again use a reduction from $\textsc{qsat}_2^{\Pi}$.

Let $\Phi(X_n, Y_n)$ be an instance of $\textsc{qsat}_2^{\Pi}$ presented, without loss of generality, in $\textsc{cnf}$. We form an instance $\Gamma_\Phi = \langle G, Ag, G_1, \ldots, G_{n+1}, \Psi \rangle$ of $\textsc{tg}^{mono}$ as follows. We set $Ag = \{a_1, \ldots, a_n, a_{n+1}\}$ and

$$G = \{g_1, \ldots, g_n, g_{n+1}, c_1, \ldots, c_n, c_{n+1}, d_1, \ldots, d_n\}$$

For each i with $1 \leq i \leq n$, G_i contains the single goal $\{g_i\}$. The set of goals for a_{n+1} is,

$$G_{n+1} \ = \ \{c_1, \ldots, c_n, c_{n+1}, d_1, \ldots, d_n, g_{n+1}\}$$

Finally, $\Psi(Ag, G)$ is

$$\left(\bigwedge_{i=1}^{n+1} (a_i \vee \neg c_i) \wedge (c_i \vee \neg g_i) \right) \wedge \Phi(c_1, \ldots, c_n, d_1, \ldots, d_n)$$

We note that this is in the form required by Theorem 9 and hence the instance described does define a coalition monotonic $\textsc{qcg}$. We claim this to be a positive instance of $\textsc{tg}^{mono}$ if and only if $\Phi(X_n, Y_n)$ is a positive instance of $\textsc{qsat}_2^{\Pi}$.

Suppose that $\Phi(X_n, Y_n)$ is accepted as an instance of $\textsc{qsat}_2^{\Pi}$. Consider any instantiation α of X_n and the coalition $C_\alpha \subseteq Ag/\{a_{n+1}\}$ for which $a_i \in C_\alpha$ if $\alpha_i = \top$, $a_i \notin C_\alpha$ if $\alpha_i = \bot$. It is clearly the case that the mapping so defined gives a bijection between instantiations α of X_n and coalitions that do not contain a_{n+1}. Now consider the instantiation β under which $\Phi(\alpha, \beta) = \top$. Define the subset $G_{\alpha,\beta}$ of G by $c_i \in G_{\alpha,\beta}$ and $g_i \in G_{\alpha,\beta}$ if $\alpha_i = \top$, $d_i \in G_{\alpha,\beta}$ if $\beta_i = \top$. Certainly $G_{\alpha,\beta}$ satisfies the coalition C_α since it contains the goal g_i. In addition, however, $\Psi[C_\alpha, G_{\alpha,\beta}] = \top$, since $a_i = c_i = \top$ for each $a_i \in C_\alpha$ and

$$\Phi(c_1, \ldots, c_n, d_1, \ldots, d_n)[G_{\alpha,\beta}] = \Phi(\alpha, \beta) = \top$$

from the choice of β. This indicates that from $\Phi(X_n, Y_n)$ a positive instance of QSAT_2^{Π} every coalition not containing a_{n+1} in the coalition monotonic QCG, Γ_Φ is successful. To complete this part of the proof it remains only to observe that all coalitions $C_\alpha \cup \{a_{n+1}\}$ are also successful by employing the goal set $G_{\alpha,\beta} \cup \{c_{n+1}, g_{n+1}\}$. In total, if $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Π} then the coalition monotonic QCG Γ_Φ is a positive instance of TG^{mono}.

On the other hand, suppose that Γ_Φ is accepted as an instance of TG^{mono} and let $C \subseteq Ag/\{a_{n+1}\}$ be a coalition with $G' \subseteq G$ a subset of goals for which $\Psi[C, G'] = \top$ and $G_i \cap G' \neq \emptyset$ for each $a_i \in C$. First observe that these properties indicate $c_i \in G'$ if $a_i \in C$: otherwise the term $(c_i \vee \neg g_i)$ could only take the value $\top$ with $g_i \notin G'$ with the result that a_i does not have its goal satisfied. Similarly, if $a_i \notin C$ then $c_i \notin G'$ for in that case we would have $(a_i \vee \neg c_i)[C, G'] = \bot$, i.e., C would not succeed with the goal set G'. We deduce, therefore, that $c_i \in G'$ if and only if $a_i \in C$. Thus, each $C \subseteq Ag$ defines a distinct instantiation of $\langle c_1, \ldots, c_n \rangle$ and from the assumption $\Psi[C, G'] = \top$, the subset of $\{d_1, \ldots, d_n\}$ within G' must induce a setting under which $\Phi(c_1, \ldots, c_n, d_1, \ldots, d_n)[C, G'] = \top$. From this we can construct, for each instantiation α of X_n an instantiation β of Y_n under which $\Phi(\alpha, \beta) = \top$. It follows that if Γ_Φ is a positive instance of TG^{mono} then $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Π}, so completing the argument that TG^{mono} is Π_2^p-complete. ∎

There are some points worth commenting upon regarding the proof above. First, we observe that, in one sense, the presence of an agent a_{n+1} is arguably redundant: the same reduction applies with a structure in which a_{n+1} and $\{c_{n+1}, g_{n+1}\}$ do not occur, i.e., with $G = \cup_{i=1}^{n}\{g_i, c_i, d_i\}$. Adopting this choice would, however, create an instance wherein G contained goals – $\{c_1, \ldots, c_n, d_1, \ldots, d_n\}$ – that were of no *direct* interest to any of the agents, since none represent choices that an agent would be satisfied by. While our definitions do not *require* that every goal in G be a satisfactory choice for at least one $a_i \in Ag$, each of our previous reductions has constructed instances in which no goal is "redundant", thus by introducing a_{n+1} we are able to maintain this. A second point concerns one interesting consequence of the structure we exploit in the reduction, but when we do allow notionally redundant goals. Thus, consider the coalition monotonic QCG Γ formed from an instance $\Phi(X_n, Y_n)$ of QSAT_2^{Π}, in which $Ag = \{a_1, \ldots, a_n\}$, $G_i = \{g_i\}$, $G = \cup_{i=1}^{n} \{c_i, d_i, g_i\}$ and $\Psi(Ag, G)$ is

$$\left(\bigwedge_{i=1}^{n} (a_i \vee \neg c_i) \wedge (c_i \vee \neg g_i)\right) \wedge \Phi(c_1, \ldots, c_n, d_1, \ldots, d_n)$$

If we consider a feasible and satisfying goal set G' for some coalition $C \subseteq Ag$ this has two properties: G' always contains goals (c_i) which are of "no interest" to $a_i \in C$; G' satisfies *only* the agents in C and no other, i.e., in terms of our class of decision problems considered above, C is *selfishly successful*. It follows, therefore, that even if we modify our notion of "trivial game" to encompass one in which every coalition is not only successful but also selfishly so, then the resulting decision question remains Π_2^p-complete even in the context of coalition monotonic QCGS.

Next, we consider the problem of whether any set of goals will be unattainable. The idea is that a set of goals will be unattainable if there is no coalition that can bring this set of goals about while at the same time satisfying all members of the coalition. It follows that such a set of goals will never be achieved, as no coalition has any incentive to achieve them – for if a coalition chose to bring about an unattainable goal set, then by definition, some members of the coalition will not have their own goals achieved. The decision problem (which we formulate existentially) is as follows.

GLOBAL UNATTAINABILITY: (GU)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$
Answer: "Yes" if there exists $G' \subseteq G$ such that for every coalition $C \subseteq Ag$ should $(\forall i \in C\ G_i \cap G' \neq \emptyset)$ then $\Psi[C, G'] = \bot$.

Example 40 *With respect to* QCG Γ_1*, the goal set* $\{g_1, g_2, g_3\}$ *is globally unattainable, but no other goal sets in this* QCG *are globally unattainable. In* QCG*s* Γ_2 *and* Γ_4*, there are no unattainable goal sets. In* Γ_3*, the goal set* $\{g_1, g_2, g_3\}$ *is unattainable.*

Theorem 41 GLOBAL UNATTAINABILITY *is* Σ_2^p*-complete.*

Proof: Membership in Σ_2^p follows from the fact that an instance $\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$ of GU is accepted if

$$\exists G' \subseteq G,\ \forall C \subseteq Ag,\ [(\forall\ i \in C, G_i \cap G' \neq \emptyset) \Rightarrow (\Psi[C, G'] = \bot)]$$

and the relationship to be verified can be tested in deterministic polynomial time.

To show that GU is Σ_2^p-hard we present a reduction from the problem QSAT$_2^\Sigma$, instances of which are propositional formulae $\Phi(X, Y)$ (as with QSAT$_2^\Pi$) these being accepted if there is an instantiation (α) of X such that for all instantiations (β) of Y we have $\Phi(\alpha, \beta) = \top$, i.e., $\exists\alpha\forall\beta\Phi(\alpha, \beta)$. This is of course the complement problem to QSAT$_2^\Pi$, and is the canonical Σ_2^p-complete problem.

Given $\Phi(X, Y)$ as an instance of QSAT$_2^\Sigma$, (again assuming without loss of generality that $|X| = |Y| = n$), we form an instance Γ_Φ of GU in which $Ag = \{a_1, \ldots, a_n, a_{n+1}\}$, $G_i = \{g_{n+1}\}$ (when $1 \leq i \leq n$), $G_{n+1} = \{g_1, \ldots, g_n\}$, $G = \cup_{i=1}^{n+1}\{g_i\}$, and with Ψ given by:

$$(g_{n+1}(\neg\Phi(g_1, \ldots, g_n, a_1, \ldots, a_n))) \vee \left(a_{n+1}(\neg g_{n+1})\left(\bigvee_{i=1}^{n} g_i\right)\right) \vee \left(\bigwedge_{i=1}^{n+1}(\neg g_i \wedge \neg a_i)\right)$$

We claim that $\Phi(X_n, Y_n)$ is accepted as an instance of QSAT_2^{Σ} if and only if the instance Γ_Φ of GU formed from it is accepted. We observe that with any instantiation, α of X_n we can associate a unique subset $G^{(\alpha)}$ of $\{g_1, \ldots, g_n\}$.

($\Rightarrow$) Assume that $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Σ} with $\alpha = \langle \alpha_1, \ldots, \alpha_n \rangle$ an instantiation of X_n for which any instantiation $\beta = \langle \beta_1, \ldots, \beta_n \rangle$ of Y_n yields $\Phi(\alpha, \beta) = \top$. Consider the goal set $G' = \{g_{n+1}\} \cup \{g_i : \alpha_i = \top\}$. We claim that for any coalition, $C \subseteq Ag$,

$$(\forall\, i \in C, G_i \cap G' \neq \emptyset) \Rightarrow (\Psi[C, G'] = \bot)$$

If $a_{n+1} \in C$, then since $g_{n+1} \in G'$, the definition of Ψ yields $\Psi[a_{n+1} \cup C', G'] = \Psi[C', G']$. That is, with $g_{n+1} \in G'$, it suffices to consider only $C \subseteq Ag/\{a_{n+1}\}$. Define β_i to be $\top$ if $a_i \in C$. For such coalitions,

$$\Psi[C, G'] \;=\; \neg\Phi(\alpha, \beta_1, \ldots, \beta_n)$$

This evaluates to $\bot$ by our choice of α and the assumption that $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Σ}. Thus, if $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Σ} then the QCG constructed is a positive instance of GU.

($\Leftarrow$) Assume $\langle G, Ag, G_1, \ldots, G_{n+1}, \Psi \rangle$ is a positive instance of GU as witnessed by the goal set $G' \subseteq G$. It must be the case that $g_{n+1} \in G'$, for otherwise we can choose $C = \{a_{n+1}\}$ (if $G' \neq \emptyset$) or $C = \emptyset$ (if $G' = \emptyset$) as coalitions which are satisfied by G' and for which $\Psi[C, G'] = \top$. Both of these contradict the choice of G'. Note that for the case $C = \emptyset$ we recall that the *empty* conjunction by convention has value $\top$ We may therefore assume that $g_{n+1} \in G'$, in which case – defining α_i to be $\top$ if $g_i \in G'$, gives $\Psi[C, G'] = \neg\Phi(\alpha, a_1, a_2, \ldots, a_n)[C]$. Consider any $C \subseteq Ag/\{a_{n+1}\}$. We first note that any such C is either empty (and so trivially satisfied by G') or each $a_i \in C$ is satisfied by g_{n+1}. By the choice of G' we must have $\Psi[C, G'] = \bot$ for all $C \subseteq Ag/\{a_{n+1}\}$. So for each β defined via $\beta_i = \top$ if $a_i \in C$, $\neg\Phi(\alpha, \beta) = \bot$ and thus we have identified an instantiation (α) of X_n for which given any instantiation (β) of Y_n,

$$\neg\Phi(\alpha, \beta) = \bot \text{ equivalently } \exists\alpha\forall\beta\; \Phi(\alpha, \beta) = \top$$

Thus if $\langle G, Ag, G_1, \ldots, G_{n+1}, \Psi \rangle$ is a positive instance of GU then $\Phi(X_n, Y_n)$ is a positive instance of QSAT_2^{Σ}.

It follows that GU is Σ_2^p-complete. ∎

In contrast to the situation regarding TG^{mono} where Corollary 39 shows that restriction to coalition monotonic QCGs does not result in a reduction in complexity, in the case of GU^{mono} we do obtain such a reduction.

Corollary 42 GU^{mono} *is in* NP.

Proof: Recall from Corollary 20 that GR$^{mono} \in$ P. We now obtain an NP algorithm for GU^{mono} from our earlier observation that that GU can be interpreted as asking of an instance Γ whether there exists *any* goal set, G', for which $\langle \Gamma, G' \rangle$ is not accepted as an instance of GR. Thus, given $\Gamma = \langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$ an instance of GU^{mono} we use an NP procedure to guess $G' \subseteq G$ followed by checking in deterministic polynomial time whether $\langle \Gamma, G' \rangle$ is accepted as an instance of GR^{mono} using the method described in the proof of Corollary 20. The instance Γ is accepted as an instance of GU^{mono} if and only if $\langle \Gamma, G' \rangle$ is *not* accepted as an instance of GR^{mono}. ∎

Finally, we consider whether a game is *incomplete*. A game will be incomplete if every coalition is successful, but there are, nevertheless, unattainable goal sets.

Example 43 *The* QCG Γ_5 *is incomplete, as although every coalition is winning, any goal set including* g_4 *is unattainable.*

Formally, the decision problem is as follows.

INCOMPLETE GAME: (IG)
Instance: QCG $\langle G, Ag, G_1, \ldots, G_n, \Psi \rangle$
Answer: "Yes" if every $C \subseteq Ag$ is successful *but* there is some $G' \subseteq G$ such that for every coalition $C \subseteq Ag$ if $(\forall i \in C\ G_i \cap G' \neq \emptyset)$, then $\Psi[C, G'] = \perp$.

Our final result, relating to this decision problem, involves the complexity class D_2^p. This result, in addition to its relation to the main concerns of this paper, may be of some interest from the perspective of computational complexity theory, since it introduces a what we believe to be the first "natural" problem that is known to be D_2^p-complete, i.e., the first D_2^p-complete problem that was not explicitly contrived for the purpose of being D_2^p-complete.

Theorem 44 INCOMPLETE GAME *is D_2^p-complete.*

Proof: Membership in D_2^p is immediate by observing that if

$$L_1 = \{x : \text{GLOBAL UNATTAINABILITY}(x)\}$$
$$L_2 = \{x : \text{TRIVIAL GAME}(x)\}$$

then by definition

$$L_1 \cap L_2 = \{x : \text{INCOMPLETE GAME}(x)\}.$$

We have already shown L_1 to be Σ_2^p-complete (Theorem 41) and L_2 to be Π_2^p-complete (Theorem 38), hence IG $\in D_2^p$.

Recall from Fact 1 that QSAT_2^{Σ}-QSAT_2^{Π} is D_2^p-complete. To establish hardness, we present a reduction from QSAT_2^{Σ}-QSAT_2^{Π} to IG. Given an instance $\langle \Phi_1(X_n, Y_n), \Phi_2(X_n, Y_n)\rangle$ of QSAT_2^{Σ}-QSAT_2^{Π}, we form an instance $\Gamma_{\langle \Phi_1,\Phi_2\rangle} = \langle G, Ag, G_1, \ldots, G_{n+1}, \Psi\rangle$ of IG as follows. Set $Ag = \{a_1, \ldots, a_n, a_{n+1}\}$; $G_i = \{g_i^{\top}, g_i^{\perp}, g_{n+1}\}$ (for $1 \leq i \leq n$), $G = \cup_{i=1}^{n} G_i$ with $G_{n+1} = G/\{g_{n+1}\}$. Finally, fix Ψ as,

$$(a_{n+1} \Rightarrow \Psi_{\text{TG}}) \wedge (g_{n+1} \Rightarrow \Psi_{\text{GU}}) \wedge ((\neg g_{n+1}) \Rightarrow \bigwedge_{i=1}^{n}(a_i \Rightarrow (g_i^{\top} \vee g_i^{\perp})))$$

where Ψ_{TG} is,

$$(\neg g_{n+1})\Phi_2(a_1, \ldots, a_n, g_1^{\top} \vee \neg g_1^{\perp}, \ldots, g_n^{\top} \vee \neg g_n^{\perp})$$

and Ψ_{GU} is,

$$(\neg a_{n+1})(\neg\Phi_1(g_1^{\top} \vee \neg g_1^{\perp}, \ldots, g_n^{\top} \vee \neg g_n^{\perp}, a_1, \ldots, a_n))$$

We claim that $\Gamma_{\langle \Phi_1,\Phi_2\rangle}$ is accepted as an instance of IG if and only if $\langle \Phi_1, \Phi_2\rangle$ is accepted as an instance of QSAT_2^{Σ}-QSAT_2^{Π}.

($\Rightarrow$) Assume the former. Then every coalition $C \subseteq \{a_1, \ldots, a_{n+1}\}$ is successful. In particular, every coalition C for which $a_{n+1} \in C$ is successful. For any such C, $\Psi[C, G']$ reduces to (since we cannot have $g_{n+1} \in G'$)

$$\left(\Phi_2(a_1, \ldots, a_n, g_1^{\top} \vee \neg g_1^{\perp}, \ldots, g_n^{\top} \vee \neg g_n^{\perp}) \wedge \bigwedge_{i=1}^{n}(a_i \Rightarrow (g_i^{\top} \vee g_i^{\perp}))\right)[C, G']$$

For which the only satisfying and feasible goal sets are subsets of $\cup_{i=1}^{n}\{g_i^{\top}, g_i^{\perp}\}$. Now, exactly as in Theorem 38, for every instantiation α of X_n we may construct (from a feasible and satisfying goal set $G' \subseteq \cup_{i=1}^{n}\{g_i^{\top}, g_i^{\perp}\}$) an instantiation β of Y_n for which $\Phi_2(\alpha, \beta) = \top$: thus if $\Gamma_{\langle \Phi_1,\Phi_2\rangle}$ is a positive instance of IG then Φ_2 is a positive instance of QSAT_2^{Π}.

Similarly, under the assumption that $\Gamma_{\langle \Phi_1,\Phi_2\rangle}$ is a positive instance of IG, there must be some $G' \subseteq G$, that fails to be a feasible goal set for any coalition that it satisfies. In particular such G' must fail to be feasible for (satisfied) coalitions $C \subseteq Ag/\{a_{n+1}\}$. We first observe that g_{n+1} must belong to G': for suppose that G' attests to $\Gamma_{\langle \Phi_1,\Phi_2\rangle}$ being a positive instance of GU but $g_{n+1} \notin G'$. If we consider coalitions $C \subseteq Ag/\{a_{n+1}\}$ with such G' we see that $\Psi[C, G']$ is equivalent to the term $\wedge_{i=1}^{n}(a_i \Rightarrow (g_i^{\top} \vee g_i^{\perp}))[G']$. For any goal set $G' \subseteq G/\{g_{n+1}\}$ it is immediate that we can construct a coalition $C \subseteq Ag/\{a_{n+1}\}$ that is satisfied by G' yet results in $\Psi[C, G'] = \top$. Thus given $g_{n+1} \in G'$ and that we may consider only $C \subseteq Ag/\{a_{n+1}\}$, under these constraints $\Psi[C, G']$ reduces to the expression

$$\neg\Phi_1(g_1^{\top} \vee \neg g_1^{\perp}, \ldots, g_n^{\top} \vee \neg g_n^{\perp}, a_1, \ldots, a_n)[C, G']$$

Observe that given $g_{n+1} \in G'$ we can, without loss of generality, assume that $|G' \cap G_i| = 1$ for each $1 \leq i \leq n$: should G' contain both $\{g_i^{\top}, g_i^{\perp}\}$ for some i, then removing $g_i^{\perp}$ does not affect whether a_i is satisfied and does not alter the value of $g_i^{\top} \vee \neg g_i^{\perp}$. Similarly, if neither are present $g_i^{\top}$ can be added to G': a_i if present in C is already satisfied through the presence of g_{n+1} in G'. Letting $\alpha = \langle \alpha_1, \ldots, \alpha_n \rangle$ be the instantiation of X_n induced by G', we see that $\Psi[C, G']$ equals $\neg\Phi_1(\alpha, a_1, \ldots, a_n)[C]$ which by our assumptions regarding G' must evaluate to $\perp$ for every $C \subseteq Ag/\{a_{n+1}\}$. This, however, implies that $\Phi_1(X_n, Y_n)$ is a positive instance of QSAT$_2^{\Sigma}$. We have thus shown that if $\Gamma_{\langle \Phi_1, \Phi_2 \rangle}$ is a positive instance of IG then $\langle \Phi_1, \Phi_2 \rangle$ is a positive instance of QSAT$_2^{\Sigma}$-QSAT$_2^{\Pi}$.

($\Leftarrow$) For the reverse direction, suppose that $\langle \Phi_1, \Phi_2 \rangle$ is a positive instance of QSAT$_2^{\Sigma}$ − QSAT$_2^{\Pi}$. First let $\alpha = \langle \alpha_1, \ldots, \alpha_n \rangle$ be the instantiation of X_n that witnesses to Φ_1 being a positive instance of QSAT$_2^{\Sigma}$. Consider the goal set G' containing g_{n+1} together with $g_i^{\top}$ (if $\alpha_i = \top$) and $g_i^{\perp}$ (if $\alpha_i = \perp$). For any $C \subseteq Ag$ for which which $a_{n+1} \in C$, it is easy to check that $\Psi[C, G'] = \perp$. For any $C \subseteq Ag/\{a_{n+1}\}$, $\Psi[C, G']$ reduces to the term $\neg\Phi_1(\alpha, a_1, \ldots, a_n)[C]$.

Although C is always satisfied by G', we have from our construction and assumptions regarding α that this expression evaluates to $\perp$. Thus if Φ_1 is a positive instance of QSAT$_2^{\Sigma}$ then $\Gamma_{\langle \Phi_1, \Phi_2 \rangle}$ is a positive instance of GU.

For the final stage we have $\Phi_2(X_n, Y_n)$ is a positive instance of QSAT$_2^{\Pi}$, from which we will show that $\Gamma_{\langle \Phi_1, \Phi_2 \rangle}$ is a positive instance of TG. Consider any $C \subseteq Ag$. If $a_{n+1} \notin C$ then we can see that C is successful by considering subsets $G' \subseteq G/\{g_{n+1}\}$. For such choices, $\Psi[C, G']$ is equivalent to

$$\bigwedge_{i=1}^{n} (a_i \Rightarrow (g_i^{\top} \vee g_i^{\perp}))[C, G']$$

for which G' satisfying C and with $\Psi[C, G'] = \top$ can easily be found. We note that $C = \emptyset$ is always satisfied and feasible for every $G' \subseteq G/\{g_{n+1}\}$. It remains to show that every coalition C with $a_{n+1} \in C$ is successful. For such C and G' we see that

$$\Psi[C, G'] = \Psi_{\mathrm{TG}}[C, G'] \wedge \bigwedge_{i=1}^{n} (a_i \Rightarrow (g_i^{\top} \vee g_i^{\perp}))[C, G']$$

writing α_C for the instantiation of X_n defined by $C/\{a_{n+1}\}$ this expression is

$$\Phi_2(\alpha_C, Y_n) \wedge \bigwedge_{i=1}^{n} (a_i \Rightarrow (g_i^{\top} \vee g_i^{\perp}))$$

Since Φ_2 is a positive instance of QSAT$_2^{\Pi}$ there is some $\beta = \langle \beta_1, \ldots, \beta_n \rangle$ for which $\Phi_2(\alpha_C, \beta) = \top$. We can now fix a satisfied and feasible goal set for C by including the goal $g_i^{\top}$ if $\beta_i = \top$ and the goal $g_i^{\perp}$ if $\beta_i = \perp$. The resulting set is non-empty (it contains exactly n goals), satisfies a_{n+1} and every $a_i \in C/\{a_{n+1}\}$. We deduce that if Φ_2 is a positive instance of QSAT$_2^{\Pi}$ then $\Gamma_{\langle \Phi_1, \Phi_2 \rangle}$ is a positive instance of TG.

Classical Coalitional Game	QCG
$\langle Ag, v\rangle$; $v : 2^{Ag} \to \mathbb{R}^+$	$\langle G, Ag, G_1, \ldots, G_n, \Psi\rangle$
Payoff profile $x \in \mathbb{R}^n$	$G' \subseteq G$
$S \subseteq Ag$, $x(S) = \sum_{i \in S} x_i$	$C \subseteq Ag$, $sat[C, G'] = \wedge_{i \in C}(G_i \cap G' \neq \emptyset)$
S-feasible profile x:	C-feasible goal set G':
$x(S) = v(S)$	$(sat[C, G'] \wedge \Psi[C, G']) = \top$.

Table 2
Conventional coalitional games and qualitative coalitional games.

In summary, $\langle \Phi_1, \Phi_2\rangle$ is accepted as an instance of QSAT$_2^{\Sigma}$-QSAT$_2^{\Pi}$ if and only if $\Gamma_{\langle \Phi_1, \Phi_2\rangle}$ is accepted as an instance of IG.

We conclude that IG is D_2^p-complete. ∎

5 Related Work

We begin by comparing our work to the conventional notion of a coalitional game, as formulated in game theory [22, Part IV]. Recall that a conventional coalitional game is a pair $\langle Ag, \nu\rangle$, where Ag is a set of agents (as in our framework), and $\nu : 2^{Ag} \to \mathbb{R}$ is a characteristic function, which assigns to every coalition a numeric value, intuitively representing payoff that may be distributed between the members of that coalition. A brief summary of the relationship of concepts in QCGs to those in conventional cooperative games is given in Table 2. Given these basic structures, cooperative game theory attempts to answer such questions as which coalitions might be formed by rational agents, and how the payoff received by a coalition might be "reasonably" divided between the members of that coalition. With respect to the former question, concepts such as imputations, the core, stable sets, the kernel, and the nucleolus have been formulated. These concepts represent progressively richer attempts to formalise when a coalition is stable (in the sense that there would be no incentive for a rational agent to do other than remain a member of the coalition). With respect to the latter, concepts such as the Shapley value [22, p.289] have been developed, which attempt to answer the question of how much an agent should receive based on an analysis of how much that agent contributes to a coalition.

As we noted above, the fundamental distinction between our work and that in cooperative game theory is that we assume a *qualitative* representation of what an agent wants. That is, rather than assuming a quantitative characterisation of a game (where we associate numeric utility values with outcomes, and a probability distribution over outcomes), we model an agent as having a goal that it desires to be achieved, where this goal is denoted by a set of equally preferred possibilities. Haddawy and Hanks give a comparison of utility-based and goal-based (qualitative) approaches (in the context of decision-theoretic planning) [14]. We emphasise again that we are not implying the superiority of the QCG approach over alternatives – far from it. Rather, we argue that each approach has its own merits. For some applications, a qualitative, goal based framework will be more natural and easier to conceptualise, while for others, a quantitative, utility/probability-based approach will be more appropriate. We do not advocate one over the other. While goal-based approaches have some limitations, we note that there are many well-documented problems with attempting to formulate domains in precise numeric terms, and in such domains, a qualitative approach may be more appropriate (see, e.g., [25, pp.107–110] for a discussion).

A number of authors have taken ideas from cooperative game theory and attempted to apply them in multiagent systems. Sandholm *et al* identify three key issues that have been addressed [30, pp.210–211]:

- *Coalition structure generation*:
  The partitioning of a group of agents into coalitions, where the overall partition is a *coalition structure*.
- *Solving the optimization problem of each coalition*:
  Solving the "joint problem" of a coalition, i.e., finding the best way to maximise the utility of the coalition itself.
- *Dividing the value of the solution for each coalition*:
  Deciding "who gets what" in the payoff.

Sandholm and colleagues developed algorithms to find optimal coalition structures (i.e., partitions of agents) with worst case guarantees (that is, within some given ratio bound k of optimal) [30]. They showed that finding the optimal coalition structure (where optimal is defined as maximising the sum of the values of each coalition in the structure) is NP-complete [30, pp.224–225]. They were able to show that for a ratio bound $k = a$ (where a is the number of agents) their algorithm required searching 2^{a-1} nodes, and that in a precise sense, this is the best that could be expected of such an algorithm. They also went on to establish how more extensive search might be used to lower the bound k. In earlier work, Shehory and Kraus developed algorithms for coalition structure formation in which agents were modelled as having different capabilities, and were assumed to benevolently desire some overall task to be accomplished, where this task had some complex (plan-like) structure [32–34]. They noted the NP-complete nature of the problems they tackled.

With respect to the specific issue of computational complexity, a somewhat smaller body of work exists on the complexity of solution concepts from cooperative game theory. Bilbao and colleagues survey the complexity of a number of problems in cooperative game theory [2]. They focussed on settings in which the characteristic function was given by various combinatorial structures (such as minimum cost spanning trees) – although such structures cannot represent all characteristic functions. Given such representations, they state that establishing membership of the core ranges from polynomial-time computable to co-NP complete, while determining whether the core is empty is NP-complete in the worst case, computing the Shapley value of such a game is in general #P-complete, (and hence as hard as counting satisfying assignments of propositional logic formulae), while computing the nucleolus of a game is in general NP-hard.

Conitzer and Sandholm also investigated the complexity of determining non-emptiness of the core [7]. They considered a representation of *superadditive* characteristic functions via *utility possibility sets*. For a given coalition, a utility possibility set captures the possible utilities that members of a coalition can *guarantee* themselves by cooperating with one another (these are somewhat similar to the sets of choices in QCGs). With the assumption of superadditivity, a succinct representation of utility possibility sets can be achieved – intuitively, one needs only a basis of utility possibilities; utility possibilities of coalitions can be derived from these by applying superadditivity. With such a representation, Conitzer and Sandholm proved that establishing non-emptiness of the core is NP-complete (both with and without transferable utility). They point out that the assumption of superadditivity, actually *strengthens* their results overall, as this implies that the problem in general (i.e., not assuming superadditivity) must be at least NP-hard.

Some complexity results have also been derived for related problems. Tennenholtz and Moses investigated the complexity of the *cooperative goal achievement* (CGA) problem [36,21]. This problem may be understood as follows. Given a set of agents, each with their own abilities an goals, is it possible for the agents to cooperate in such a way that all their goals are achieved. Representing agents in such systems as non-deterministic finite state machines (where the non-determinism is used to capture the fact that agents may have a number of possible actions available to them at any moment), Tennenholtz and Moses show that the complexity of the CGA problem is PSPACE-complete (assuming that agents have complete information about the initial state of the system). Somewhat related is the COOPSAT problem studied by d'Inverno and colleagues [8]. In this problem, we are given a goal to be achieved, a set of agents, (each with different abilities), and preferences about who they will cooperate with. We are asked whether there is some configuration of the agents that is capable of bringing about the goal; they show that this problem is NP-complete (given a simple representation of capabilities and preferences).

Finally, we note that our representation of characteristic functions is close to the notion of an *effectivity function* in the social choice literature [1,26]. An effectivity function is generally formulated as a function

$$E : 2^{Ag} \to 2^{2^S}$$

where S is a set of *possible states*. The intuition is that if $S' \in E(C)$ for coalition

$C \subseteq Ag$, then this means that one of the choices available to the coalition C is to enforce *one of the states in* S', but not to be able to choose between these states. Thus we would say that a coalition C was able to force a property P if these was some choice $S' \in E(C)$ such that P was true in every state in S' [26]. With a little additional technical machinery, we can make precise the relationship between effectivity functions and our characteristic functions. Assume that we have a set of states S, an effectivity function E, as above, and a set G of goals, as in QCGs. We write $s \models g$ to indicate that goal $g \in G$ is satisfied in state s (the use of the satisfaction relation symbol "$\models$" is deliberately suggestive). If $S' \subseteq S$, then we write G/S' to denote the set of goals that are satisfied in all members of S'.

$$G/S' = \{g : \forall s \in S', s \models g\}$$

Then the characteristic function $V : 2^{Ag} \to 2^{2^G}$ corresponding to E is defined as follows.

$$V(C) = \{G/S' : S' \in E(C)\}$$

So moving from effectivity functions to characteristic functions is straightforward, assuming we have the relation "$\models$" available to tell us which goal is satisfied in which state. Moving from characteristic functions back to effectivity functions, however, is in general not directly possible, for obvious reasons. Nevertheless, it is clear that characteristic functions and effectivity functions are intimately related: they represent different ways of modelling essentially the same types of scenario. However, for the questions we wish to investigate, the technical machinery of QCGs is somewhat simpler than that of effectivity functions – we do not require states and the satisfaction relation. Abdou and Keiding give a detailed analysis of the use of effectivity functions, with particular emphasis on social choice theory [1].

6 Conclusions

In this article our first aim has been to define a framework for modelling scenarios wherein sets of agents (coalitions) wish to achieve particular goals and may require the assistance of other agents in the system in order to facilitate this; as such, our work is central to the multiagent research paradigm [3,37] An important aspect of our model is that the utility of a specific set of goals to a given coalition is interpreted solely in terms of these goals being both feasible for the coalition and such that each member will be satisfied: a purely qualitative measure. This contrasts with the widely studied models in classical coalitional game theory in which coalitions are assessed with respect to some real-valued payoff profile. By representing characteristic functions via formulae of propositional logic, we are able to make meaningful complexity assessments regarding a number of natural decision questions that arise. We have shown that such questions range in complexity from NP-complete to completeness within complexity classes which (under the standard complexity-theoretic assumptions) lie strictly between the second and third levels of the polynomial-time hierarchy.

While we have not focussed on the issue of computationally tractable cases – a topic that is a concern of current work in progress – we note that the propositional representation suggests that many of the decision problems falling within NP $\cup$ co-NP may well prove to be tractable for those cases where the representing formula falls within a class for which polynomial-time satisfiability methods exist, e.g., Horn clauses,[11] 2-CNF, etc. If this is indeed the case, an open question of some interest is to assess to what extent such classes define "realistic" contexts for QCGs.

Acknowledgments

The work reported in this article was carried out with the support of EPSRC Grant GR/R60836/01.

References

[1] J. Abdou and H. Keiding. *Effectivity Functions in Social Choice Theory*. Kluwer Academic Publishers: Dordrecht, The Netherlands, 1991.

[2] J. Bilbao, J. Fernández, and J. López. Complexity in cooperative game theory. Manuscript.

[3] A. H. Bond and L. Gasser, editors. *Readings in Distributed Artificial Intelligence*. Morgan Kaufmann Publishers: San Mateo, CA, 1988.

[4] R. E. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Trans. on Computers*, C-35(6):677–691, 1986.

[5] R. E. Bryant. Symbolic manipulation with ordered binary-decision diagrams. *ACM Computing Surveys*, 24(3):293–318, 1992.

[6] R. Chang and J. Kadin. The boolean hierarchy and the polynomial hierarchy: A closer connection. *SIAM Journal of Computing*, 25(2):340–354, 1996.

[7] V. Conitzer and T. Sandholm. Complexity of determining nonemptiness of the core. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 613–618, Acapulco, Mexico, 2003.

[11] In the case where $\Psi(Ag, G)$ is given as a Horn clause formula, the decision problem Empty Game becomes trivial: every instance is rejected by virtue of the fact that $\Psi[\emptyset, \emptyset] = \top$, i.e. the empty coalition is satisfied by any subset G' of G, and the choice $G' = \emptyset$ will be feasible for $C = \emptyset$ when $\Psi(Ag, G)$ is a conjunction of Horn clauses.

[8] M. d'Inverno, M. Luck, and M. Wooldridge. Cooperation structures. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, pages 600–605, Nagoya, Japan, 1997.

[9] P. E. Dunne. *The Complexity of Boolean Networks*. The Academic Press: London, England, 1988.

[10] P.E. Dunne. Prevarication in dispute protocols. *Proc. Ninth International Conference on A.I. and Law*, (ICAIL' 03), Edinburgh, June 2003, ACM Press, pages 12–21

[11] M. Fischer and N. J. Pippenger. Relations among complexity measures. *Journal of the ACM*, 26:361–381, 1979.

[12] D. Gabbay and J. Woods. More on Non-Cooperation in Dialogue Logic *Logic Jnl. of the IGPL*, 9:305–323, 2001

[13] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of* NP*-Completeness*. W. H. Freeman: New York, 1979.

[14] P. Haddawy and S. Hanks. Utility models for goal-directed decision-theoretic planners. *Computational Intelligence*, 14(3):392–429, 1998.

[15] J. E. Hopcroft and R. M. Karp An $n^{5/2}$ algorithm for maximum matching in bipartite graphs. *SIAM Journal of Computing*, 2:225–231, 1975

[16] D. S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science Volume A: Algorithms and Complexity*, pages 67–161. Elsevier Science Publishers B.V.: Amsterdam, The Netherlands, 1990.

[17] J. Kadin. The polynomial time hierarchy collapses if the boolean hierarchy collapses. *SIAM Journal of Computing*, 17(6):1263–1282, 1988.

[18] V. M. Khrapchenko. Methods of determining lower bounds for the complexity of Π–schemes. *Math. Notes of the Acad. Sci., USSR*, 10:474–479, 1971.

[19] O. B. Lupanov. On a method of circuit synthesis. *Izvestia VUZ (Radiofizika)*, 1:120–140, 1958.

[20] P. Morris. *Introduction to Game Theory*. Springer-Verlag: Berlin, Germany, 1994.

[21] Y. Moses and M. Tennenholtz. Multi-entity models. In K. Furukawa, D. Michie, and S. Muggleton, editors, *Machine Intelligence 14*, pages 65–90, Tokyo, Japan, 1995.

[22] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press: Cambridge, MA, 1994.

[23] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley: Reading, MA, 1994.

[24] C. H. Papadimitriou and M. Yannakakis. The complexity of facets (and some facets of complexity). In *Proceedings of the Fourteenth ACM Symposium on the Theory of Computing (STOC-82)*, pages 255–260, San Francisco, CA, 1982.

[25] S. Parsons. *Qualitative Methods for Reasoning Under Uncertainty*. The MIT Press: Cambridge, MA, 2001.

[26] M. Pauly. A modal logic for coalitional power in games. *Journal of Logic and Computation*, 12(1):149–166, 2002.

[27] V. R. Pratt. The effect of basis on the size of Boolean expressions. In *Proceedings of the Sixteenth Symposium on Foundations of Computer Science (FOCS)*, pages 119–121, 1975.

[28] A. A. Razborov. A lower bound on the monotone complexity of the logical permanent. *Math. Notes of the Acad. Sci., USSR*, 37:485–493, 1985.

[29] T. Sandholm. Distributed rational decision making. In G. Weiß, editor, *Multiagent Systems*, pages 201–258. The MIT Press: Cambridge, MA, 1999.

[30] 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.

[31] C. P. Schnorr. The network complexity and turing machine complexity of finite functions. *Acta Informatica*, 7:95–107, 1976.

[32] O. Shehory and S. Kraus. Coalition formation among autonomous agents: Strategies and complexity. In C. Castelfranchi and J.-P. Müller, editors, *From Reaction to Cognition — Fifth European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW-93 (LNAI Volume 957)*, pages 56–72. Springer-Verlag: Berlin, Germany, 1995.

[33] O. Shehory and S. Kraus. Task allocation via coalition formation among autonomous agents. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*, pages 655–661, Montréal, Québec, Canada, August 1995.

[34] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, 1998.

[35] O. Shehory and S. Kraus. Feasible formation of stable coalitions among autonomous agents in non-super-additive environments. *Computational Intelligence*, 15(33):218–251, 1999.

[36] M. Tennenholtz and Y. Moses. On cooperation in a multi-entity model: Preliminary report. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, pages 918–923, Detroit, MI, 1989.

[37] M. Wooldridge. *An Introduction to Multiagent Systems*. John Wiley & Sons, 2002.

[38] C. Wrathall. Complete sets and the polynomial-time hierarchy. *Theoretical Computer Science*, 3:23–33, 1976.

[39] F. C. Zagare. *Game Theory: Concepts and Applications*. Sage Publications: Beverly Hills, 1984.