Consensus Games

Julian Zappala, Natasha Alechina, and Brian Logan

School of Computer Science
University of Nottingham
Nottingham, NG8 1BB, UK
{jxz,nza,bsl}@cs.nott.ac.uk

Abstract. Consensus Games (CGs) are a novel approach to modelling coalition formation in multi-agent systems inspired by threshold models and quorum sensing found in sociological and biological systems. In a consensus game, each agent's degree of commitment to the coalitions in which it may participate is expressed as a quorum function, and an agent is willing to participate in a coalition if and only if a quorum consensus can be achieved by all the agents participating in the coalition. The computational complexity of several decision problems associated with CGs is analysed and tractable algorithms for problems such as determining whether a coalition is a consensus coalition are given.

1 Introduction

If the potential of multi-agent systems is to be realised it is essential that agents are able to cooperate and coordinate their behaviours [1]. A crucial issue for agents when choosing between alternative cooperative options is therefore determining which coalition to join [2].

Coalition formation has traditionally been modelled using game theoretic techniques; such models often necessitate strong economic assumptions. These include the existence of a common, transferable utility, that valuations for each coalition can be known, and that mechanisms for fair distribution of coalitional gains are in place. The multi-agent community in particular have investigated coalition formation in circumstances where these economic assumptions cannot easily be applied. Examples of such models include effectivity functions [3], where agents are interested in enforcing a particular state, and qualitative coalitional games [2], where agents are interested in achieving at least one of their goals. An assumption common to all of this work is that, for a coalition to form, all member-agents must somehow 'agree' to its formation. In other words, for a coalition to form it is necessary that there is a *consensus* amongst the coalition's members regarding the coalition's formation.

Consensus has been extensively studied in the biological literature. For example, many natural systems, including bacteria [4], ants [5], bees [6], and fish [7] exhibit a behaviour known as quorum sensing. Through a process termed the *quorum response*, the probability of an individual selecting a particular behaviour is increasing in the proportion of individuals already having made that choice. The macroscopic behaviour of this self-organising system resembles one in which individuals converge upon consensus. At a higher level of complexity, threshold models have been used to explain collective behaviours in sociological systems. For example, in [8], a model is described in which individuals are faced with the binary decision of whether or not to participate in a riot. In this model each individual possesses an idiosyncratic threshold representing the minimum proportion of others which must participate in order that the given individual will also participate. A population of n individuals is considered with uniformly distributed thresholds $\{\frac{0}{n}, \frac{1}{n}, \ldots, \frac{n-1}{n}\}$; the scenario begins with a single instigator performing some riotous display; this behaviour then incites the next, and so forth until mass rioting ensues. Similar models have been found to describe a variety of social phenomena including segregation in urban housing [9, 10], and the adoption of consumer trends [11].

In this paper we propose *consensus games* (CGs), a novel model of consensual coalition formation for multi-agent systems which draws inspiration from a consensus mechanism found in biological and sociological systems. We extend the model proposed in [8] beyond binary choice decisions to the more general problem of coalition formation. For each coalition of which they may be a member, each agent holds a threshold representing the proportion of agents from that coalition which must support the formation of the coalition in order that the agent will also support the coalition. There is consensus about the formation of a particular coalition only where all member-agents support the formation of that coalition.

The remainder of the paper is structured as follows. In the next section we introduce the notion of a consensus game. In section 3 we define the notion of a *strong consensus coalition* and in section 4 we consider collective rationality. In section 5 we establish the complexity of some decision problems for consensus games. In section 6 we define another notion of consensus, *weak consensus*, and establish complexity of the corresponding decision problems. We review related work in section 7 and conclude in section 8.

2 Consensus Games

In this section we give a formal definition of a consensus game.

Definition 1. *A consensus game (CG) is a tuple $\Gamma = \langle G, q \rangle$ where:*

G *is a finite set of agents,* $\{1, \ldots, n\}, n \geq 2$.
q *is a quorum function. It is a partial function* $q : G \times 2^G \rightarrow [0, 1]$ *which takes an agent* $i \in G$ *and a coalition of agents* H *such that* $i \in H$ *and returns a number in the interval* $[0, 1]$.

The value of the quorum function for a coalition indicates the agent's 'degree of support' for the formation of that coalition.

For an agent $i \in H \subseteq G$ the quorum function $q(i, H)$ gives the minimum proportion of agents in H which must support the formation of the coalition H in order that i will support the formation of this coalition. Where $q(i, H) = 0$ agent i unconditionally supports the formation of the coalition H, where $0 < q(i, H) \leq \frac{|H|-1}{|H|}$ the agent conditionally supports the formation of the coalition H; where $\frac{|H|-1}{|H|} < q(i, H) \leq 1$ the agent objects to the formation of the coalition H. We will use an abbreviation

$q\#(i, H)$ to denote the number of other agents from H which have to support H in order for i to support H. Formally, $q\#(i, H)$ is the minimal natural number k such that $q(i, H) \leq k/|H|$. We will also denote the number of agents $i \in H$ with $q\#(i, H) = k$ by $n_k(H)$.

3 Strong Consensus

In this section we define the main solution concept of consensus games, a strong consensus coalition. A *strong consensus coalition* H is a coalition where for each agent $i \in H$ the quorum threshold $q(i, H)$ is satisfied in the sense that H contains at least $q\#$ other agents with strictly lower $q\#$ values.

Definition 2. *A coalition H is a strong consensus coalition if the following conditions hold:*

- $n_0(H) \neq 0$
- *if* $n_k(H) \neq 0$*, then* $\Sigma_{j<k} n_j(H) \geq k$

Note that the definition implies that if H is a strong consensus coalition, then $n_{|H|}(H) = 0$.

Consider the following example.

Example 1. Alice (A) and Bob (B) are considering whether to get married. Bob no longer wishes to be a bachelor and is keen to be married. Alice is not opposed to the idea of marrying Bob provided that she knows that Bob also wants to marry her, otherwise Alice will happily continue her single lifestyle. Alice's and Bob's positions can be formalised as the consensus game $\Gamma_1 = \langle G, q \rangle$ where:

$$G = \{A, B\}$$

$$q(i, H) = \begin{cases} 0 & \text{if } i = B \text{ and } H = \{A, B\} \\ 0.5 & \text{if } i = A \text{ and } H = \{A, B\} \\ 1 & \text{if } i = B \text{ and } H = \{B\} \\ 0 & \text{if } i = A \text{ and } H = \{A\} \end{cases}$$

In Example 1, Bob unconditionally supports the formation of the grand coalition (of all agents); Alice conditionally supports formation of this coalition provided that one other agent (Bob) also supports forming this coalition. Alice also unconditionally supports formation of the singleton coalition {A}, whereas Bob objects to the formation of the singleton coalition {B}. The grand coalition in this example is a strong consensus coalition.

Next we show that there is an alternative definition of a strong consensus coalition as a fixed point of a function which intuitively corresponds to agents indicating their support for a coalition.

Consider the following function $f_H : 2^G \rightarrow 2^G$ defined relative to $H \subseteq G$:

$$i \in f_H(Q) \text{ iff } i \in H \text{ and } |Q \cap H \setminus \{i\}| \geq q(i, H) \times |H|$$

This function takes as its input a set $Q \subseteq G$ and returns the set of agents in H whose quorum thresholds are satisfied by the membership of $Q \cap H$. If $Q = \emptyset$, f_H will contain only the agents i with $q(i, H) = 0$, if Q is the set of agents which have unconditional support for H, then $f_H(Q)$ will contain the agents i with $q\#(i, H) \leq |Q|$, and so on.

We are going to show that a coalition H is a strong consensus coalition if and only if it is the least fixed point of f_H. First we need the following auxiliary result:

Proposition 1. *The function f_H is guaranteed to possess at least one fixed point.*

Proof. Existence of at least one fixed point is guaranteed for monotonic functions; Knaster-Tarski theorem [12].

For monotonicity it must be shown that $\forall Q, Q' \subseteq G$ where $Q \subseteq Q'$ it is always the case that $f_H(Q) \subseteq f_H(Q')$.

Let $Q \subseteq Q'$. We need to show that for every i, if $i \in f_H(Q)$, then $i \in f_H(Q')$. Assume $i \in f_H(Q)$. By the definition of f_H, $i \in H$ and $|Q \cap H \setminus \{i\}| \geq q(i, H) \times |H|$. Since $Q \subseteq Q'$, $|Q' \cap H \setminus \{i\}| \geq |Q \cap H \setminus \{i\}|$ hence $|Q' \cap H \setminus \{i\}| \geq q(i, H) \times |H|$ and $i \in f_H(Q')$. □

The least fixed point of a function can be established by recursive calls to the function starting with the empty set of agents as an argument; each invocation of f_H will be referred to as a *round*. If H can achieve least fixed point consensus, then it will be achieved in at most $|H|$ rounds.

Now we can prove

Theorem 1. *H is a strong consensus coalition if and only if it is the least fixed point of f_H.*

Proof. Assume H is a strong consensus coalition. Then $n_0(H) \neq 0$ and $|f_H(\emptyset)|$ = $|\{i \in H : q\#(i, H) = 0\}|$ = n_0, so $f_H(\emptyset) \neq \emptyset$. Similarly, at each round $k > 1$, $f_H^k(\emptyset)$ = $f_H^{k-1}(\emptyset) \cup \{i \in H : q\#(i, H) \leq |f_H^{k-1}(\emptyset)|\}$ (where f_H^k is k applications of f_H). By the definition of a strong consensus coalition, the set $\{i \in H : q\#(i, H) \leq |f_H^{k-1}(\emptyset)|\}$ is always non-empty until $f_H^{k-1}(\emptyset) = |H|$, so H is the least fixed point of f_H.

Assume H is the least fixed point of f_H. Then the first condition in the definition of strong consensus coalition is satisfied because $f_H(\emptyset) \neq \emptyset$ hence there are agents i with $q(i, H) = 0$. To show that for every k, if $n_k(H) \neq 0$, then there are at least k agents in H with strictly lower $q\#$ values, consider an agent i with $q\#(i, H) = k$. Since H is the least fixed point of f_H, at some round m, $i \in f_H^m(\emptyset)$. This means that $|f_H^m(\emptyset)| \geq k$, and since $|f_H^m(\emptyset)|$ is the number of agents with lower $q\#$ values, $\Sigma_{j<k} n_j(H) \geq k$. □

4 The q-Minimal Core

The core is a popular solution concept in game theory, aggregating equilibria that are both individually and collectively rational. In traditional, quantitative game theoretic models, rational behaviour is typified by agents' maximising some notion of utility. Agents are said to act with individual rationality where each agent will achieve their maximum reward given the behaviour of the other agents; this corresponds to the Nash equilibrium of the game. Collectively rational outcomes are those where no subset of agents can achieve a higher reward through unilateral defection.

By contrast, CGs are qualitative; they do not describe a pay-off or reward structure. Instead, individually rational behaviour is associated with the quorum thresholds of the agents for the coalitions of which they may be a member. It is individually rational for an agent to support the formation of some coalition only if the number of agents already known to be supporting that coalition is a least as great as the proportion specified by the agent's quorum threshold for that coalition.

Notions of equilibria constituting collectively rational outcomes for CGs are harder to define. One natural way of distinguishing between coalitions is in terms of the effort required to reach consensus. For some coalitions, strong consensus may be established in a single round, while for others it may require as many as $|H|$ rounds. The number of rounds required to reach consensus for a given coalition can be taken as a quantitative measure of the ease with which the agents reach consensus and the stability of the resulting coalition.

Let $rounds(H)$ be the number of rounds necessary for strong consensus to be established for some coalition H. A *q-minimal consensus coalition* is a strong consensus coalition of agents $H \subseteq G$ for which no subset of agents, $H' \subset H$, is also a strong consensus coalition converging in strictly fewer rounds than $rounds(H)$. The q-Minimal Core aggregates q-minimal consensus coalitions.

5 Complexity of CGs

To characterise the computational complexity of CGs we consider three natural decision problems associated with coalition formation. These decision problems address three fundamental questions for CGs regarding the presence or otherwise of equilibria within a game. The first considers the problem of verification, the second, that of existence whist the third considers the matter of non-existence:

Consensus Coalition (CC): Can a given coalition reach consensus?
Consensus Coalition Exists (CE): Does there exist some coalition which can reach consensus?
No Consensus Coalition (NC): Is there no coalition which can reach consensus?

These decision problems are considered, first for strong consensus coalitions and subsequently for q-minimal consensus coalitions. To begin, the representational scheme and abstract model of computation for these analyses are established.

5.1 Representation

When considering the representation of CGs the structure of most interest is the quorum function, q. The quorum threshold has to be specified for each coalition $H \in \mathcal{P}(G) \setminus \emptyset$ and for each agent $i \in H$. To do this, each coalition $H \subseteq G$ is represented as a set of pairs $(i, q(i, H))$. The overall representation is the set $R = \{rep(H) \mid H \in \mathcal{P}(G) \setminus \emptyset\}$, where $rep(H) = \{(i, q(i, H)) \mid i \in H, q(i, H) \in [0, 1]\}$. Note that the size of R is exponential in the number of agents n.

It is assumed that R is implemented as a random access data structure, hence, the following results are given for the non-deterministic random access machine (NRAM) model of computation [13].

5.2 Complexity of strong consensus decision problems

STRONG CONSENSUS COALITION (SCC).

Given a CG $\Gamma = \langle G, q \rangle$ and a coalition $H \subseteq G$, can H reach strong consensus?

A deterministic algorithm must verify that H is the least fixed point of f_H. Algorithm 1 runs in time which is polynomial (linear) in n and therefore lies within $P(n)$. Note that an $O(n \times log(n))$ algorithm, which runs in constant space, can also be obtained by sorting H on $q(i, H)$.

Algorithm 1 Can H reach strong consensus.

```
function SCC(R, H)
    array support[|H| + 1] ← {0, ..., 0}
    for all (i, q) ∈ H do
        k ← ⌈q × |H|⌉
        support[k] ← support[k] + 1
    s ← support[0]
    for k from 1 to |H| do
        if k ≤ s then
            s ← s + support[k]
        else
            return false
    return true
```

STRONG CONSENSUS COALITION EXISTS (SCE).

Given a CG $\Gamma = \langle G, q \rangle$ is there some $H \subseteq G$ which can reach strong consensus?

A deterministic algorithm would iterate over R, checking for each $H \subseteq G$ whether H is a strong consensus coalition. Hence the problem is in $O(n2^n)$.

A non-deterministic algorithm first guesses an index of a coalition $H \in R$ and then checks that that H can reach strong consensus. This can be done in time linear in n using Algorithm 1. This gives a non-deterministic linear time algorithm for a random access machine. Hence, the problem is in $NP(n)$ for NRAM.

NO STRONG CONSENSUS COALITION (SNC).

Given a CG $\Gamma = \langle G, q \rangle$ is there no $H \subseteq G$ which can reach strong consensus?

A deterministic algorithm must verify that $\neg\exists H \subseteq G$ such that H is the least fixed point of f_H. Hence the problem is in $O(n2^n)$.

The problem of verifying that there exists some coalition which can reach strong consensus is in $NP(n)$. Therefore the complement of that problem, verifying that there exists no coalition which can reach strong consensus is in *co*-$NP(n)$ for NRAM.

q-MINIMAL STRONG CONSENSUS COALITION (QM-SCC).

A deterministic algorithm must verify that H is the least fixed point of f_H and that $\neg\exists H' \subset H$ such that H' is the least fixed point of $f_{H'}$ and H' reach consensus in strictly fewer rounds than H. Algorithm 2 computes the number of rounds required to encounter the least fixed point of f_H. The algorithm has time complexity $O(n)$ and so is in $P(n)$.

Algorithm 2 Number of rounds for H.

```
function rounds(rep(H))
    array support[|H| + 1] ← {0, ..., 0}
    for all (i, q) ∈ H do
        k ← ⌈q × |H|⌉
        support[k] ← support[k] + 1
    r ← 0
    i1 ← 1
    i2 ← s ← support[0]
    while i1 ≤ i2 do
        r ← r + 1
        for k from i1 to i2 do
            s ← s + support[k]
        i1 ← i2 + 1
        i2 ← s
    return r
```

Algorithm 3 then verifies that a given coalition, H is a q-minimal strong consensus coalition, by iterating over all subsets of H. Hence the problem is in $O(n2^n)$.

Algorithm 3 Is H is a q-minimal strong consensus coalition.

```
function QM-SCC(H, R)
    if ¬SCC(R, H) then
        return false
    for all H' ⊂ H ∈ R do
        if SCC(R, H) ∧ rounds(H') < rounds(H) then
            return false
    return true
```

A non-deterministic algorithm to solve the complement of this problem (decide whether a coalition is *not* a q-minimal consensus coalition) first checks whether H is a strong consensus coalition (and returns true if it is not); if H is a strong consensus coalition, it will guess an index of an coalition $H' \subset H$ and check that H' is a strong consensus coalition which converges in fewer rounds than for H (and returns false if it does). So the problem of deciding whether a coalition is *not* a q-minimal strong consensus coalition is in $NP(n)$ on NRAM. Hence deciding whether a coalition is a q-minimal strong consensus coalition is in *co*-$NP(n)$ for NRAM.

q-MINIMAL STRONG CONSENSUS COALITION EXISTS (QM-SCE).

A deterministic algorithm must verify that $\exists H \subseteq G$ such that H is the least fixed point of f_H and that $\forall H' \subset H, \neg\exists H'$ such that H' is the least fixed point of $f_{H'}$ and H' reach consensus in strictly fewer rounds than H. Hence the problem is in $O(n2^n)$.

If there exists some coalition which is a strong consensus coalition then either that coalition itself, or some subset of that coalition will be a q-minimal strong consensus coalition. Therefore, this decision problem can be solved in a manner similar to SCE and so is in $NP(n)$ for NRAM.

NO q-MINIMAL STRONG CONSENSUS COALITION (QM-SNC).

A deterministic algorithm must verify that $\neg\exists H \subseteq G$ such that H is the least fixed point of f_H and that $\forall H' \subset H, \neg\exists H'$ such that H' is the least fixed point of $f_{H'}$ and H' reach consensus in strictly fewer rounds than H. Hence the problem is in $O(n2^n)$.

If there exists some coalition which is a strong consensus coalition then either that coalition itself, or some subset of that coalition will be a q-minimal strong consensus coalition. Therefore, this decision problem can be solved in a manner similar to SNC and so is in *co*-$NP(n)$ for NRAM.

6 Weak consensus

The results from the previous section suggest that strong consensus is reasonably easy (takes time linear in the number of agents) to reach. However, it may seem that there is a notion of consensus that is even easier to reach, and which has the same intuitive appeal as strong consensus. Namely, it is a notion of consensus corresponding to no agent *objecting* to the formation of a coalition.

Definition 3. *H is a weak consensus coalition if no agent $i \in H$ has $q(i, H) > \frac{|H|-1}{|H|}$.*

Interestingly, this notion of consensus also has a fixed point characterisation:

Theorem 2. *H is a weak consensus coalition iff H is the greatest fixed point of f_H.*

Proof. The only if direction: if there exists $i \in H$ such that $q(i, H) > \frac{|H|-1}{|H|}$, then by the definition of f_H, $i \notin f_H(H)$ so $f_H(H) \neq H$.

For the if direction, assume that for all $i \in H$, $q(i, H) > \frac{|H|-1}{|H|}$. First we show that $f_H(H) = H$. By the definition of f_H, any $i \in f_H(H)$ iff $i \in H$ and $q(i, H) \leq \frac{|H|-1}{|H|}$. Since the latter holds for all $i \in H$, we have $i \in f_H(H)$ iff $i \in H$ hence H is a fixed point.

It is also the greatest fixed point, since the definition implies that $f_H(H') \subseteq H$ for any H', so $f_H(H') \subset H'$ for any $H' \supset H$. □

An example is called for:

Example 2. Consider again Example 1. However, let us now assume that both Alice and Bob conditionally support getting married. Alice's and Bob's positions can be formalised as the consensus game $\Gamma_2 = \langle G, q \rangle$ where:

$$G = \{A, B\}$$
$$q(i, G') = \begin{cases} 0.5 & \text{if } i = A \text{ and } G' = \{A, B\} \\ 0.5 & \text{if } i = B \text{ and } G' = \{A, B\} \\ 1 & \text{otherwise} \end{cases}$$

In Example 2, both Alice and Bob will support formation of the grand coalition, and so get married, if the other also supports this. As neither Alice nor Bob object to the formation of the grand coalition, it is a weak consensus coalition.

Finally, it is easy to show that while every strong consensus coalition is also a weak consensus coalition (since no strong consensus coalition can contain an agent that objects), the converse is not the case: there exists weak consensus coalitions which are not strong consensus coalitions. Since $f_H(J) = H$ for any $J \supseteq H$, it follows that if H is the least fixed point of f_H, then it is also the greatest fixed point of f_H. One such coalition is the grand coalition in Example 2.

6.1 The H-Minimal Core

The agents in a weak consensus coalition reach consensus in a single round. The analogue of a q-minimal consensus coalition is therefore not informative for weak consensus.

Instead we adopt a similar approach to the qualitative model of the core introduced in [2] for qualitative coalitional games (QCG). A coalition is in the qualitative core of a QCG if and only if that coalition is stable and no subset of that coalition is also stable. Consider two stable coalitions, $H \subseteq G$ and $H' \subset H$. Qualitatively speaking, agents in the coalition H' can do no better by forming H' than they would by participating in the larger coalition H; however, by the same rationale agents in H' will do no worse by forming this smaller coalition. In [2] it is argued that the existence of the stable coalition H' undermines the stability of H; there is nothing impelling H to remain together.

Building on this qualitative definition of the core where collective rationality is associated with minimality, we define an *H-minimal consensus coalition* as a weak consensus coalition of agents $H \subseteq G$ for which no subset, $H' \subset H$, of agents is also a weak consensus coalition. Following [2], the H-minimal core of a CG is defined as containing only H-minimal consensus coalitions. The H-minimal core aggregates consensus coalitions which are collectively rational in the sense that they are immune to defection by some agents $H' \subset H$.

6.2 Complexity of weak consensus decision problems

In the remainder of this section, we look at decision problems for weak consensus and H-minimal consensus coalitions.

WEAK CONSENSUS COALITION (WCC).

This first decision problem considers the complexity of determining if a given coalition is a weak consensus coalition. Given a CG $\Gamma = \langle G, q \rangle$ and a coalition $H \subseteq G$, will H reach weak consensus?

A deterministic algorithm must verify that $\neg \exists i \in H$ such that $q(i, H) > \frac{|H|-1}{|H|}$ (from Theorem 2). It simply iterates through $rep(H)$ checking that no agent i has $q(i, H) > \frac{|H|-1}{|H|}$. Hence the problem is in $O(n)$.

WEAK CONSENSUS COALITION EXISTS (WCE).

Given a CG $\Gamma = \langle G, q \rangle$ will any $H \subseteq G$ reach weak consensus?

A deterministic algorithm must verify that $\exists H \in R$ such that $\neg \exists i \in H$ such that $q(i, H) > \frac{|H|-1}{|H|}$. Hence the problem is in $O(n2^n)$.

A non-deterministic algorithm first guesses an index of a coalition $H \in R$ and then checks that that H can reach weak consensus. This can be done in time linear in n. This gives a non-deterministic linear time algorithm for a random access machine. Hence, the problem is in $NP(n)$ for NRAM.

NO WEAK CONSENSUS COALITION (WNC).

Given a CG $\Gamma = \langle G, q \rangle$ can no $H \subseteq G$ reach weak consensus?

A deterministic algorithm must verify that $\neg \exists H \in R$ such that $\neg \exists i \in H$ such that $q(i, H) > \frac{|H|-1}{|H|}$. Hence the problem is in $O(n2^n)$.

The problem of verifying that there exists some coalition which can reach weak consensus is in $NP(n)$ (from WCE). Therefore the complement of that problem, verifying that there exists no coalition which can reach weak consensus consensus is in *co*-$NP(n)$ for NRAM.

***H*-MINIMAL WEAK CONSENSUS COALITION (GM-WCC).**

Given a CG $\Gamma = \langle G, q \rangle$ and a coalition $H \subseteq G$, is H a H-minimal weak consensus coalition?

A deterministic algorithm must verify that H a weak consensus coalition and that $\neg \exists H' \subset H$ such that H' is also a weak consensus coalition. Hence the problem is in $O(n2^n)$.

A non-deterministic algorithm to solve the complement of this problem (decide whether a coalition is *not* a H-minimal weak consensus coalition) first checks whether H is a weak consensus coalition (and returns true if it is not); if H is a weak consensus coalition, it will guess an index of an coalition $H' \in R \setminus H$ and check that H' is a weak consensus coalition. So the problem of deciding whether a coalition is *not* an H-minimal weak consensus coalition is in NP(n) on NRAM. Hence deciding whether a coalition is an H-minimal weak consensus coalition is in *co*-$NP(n)$ for NRAM.

***H*-MINIMAL WEAK CONSENSUS COALITION EXISTS (GM-WCE).**

Given a CG $\Gamma = \langle G, q \rangle$ is some $H \subseteq G$ a H-minimal weak consensus coalition?

A deterministic algorithm must verify that $\exists H \subseteq G$ such that H is a weak consensus coalition and that $\neg\exists H' \subset H$ such that H' is also a weak consensus coalition. Hence the problem is in $O(n2^n)$.

If there exists some coalition which is a weak consensus coalition then either that coalition itself, or some subset of that coalition will be an H-minimal weak consensus coalition. Therefore, this decision problem can be solved in a manner similar to WCE and so is in $NP(n)$ for NRAM.

NO H-MINIMAL WEAK CONSENSUS COALITION (GM-WNC).

Given a CG $\Gamma = \langle G, q \rangle$ is there no $H \subseteq G$ a H-minimal weak consensus coalition?

A deterministic algorithm must verify that $\neg\exists H \subseteq G$ such that H an H-minimal weak consensus coalition. Hence the problem is in $O(n2^n)$.

If there exists some coalition which is a weak consensus coalition then either that coalition itself, or some subset of that coalition will be a H-minimal weak consensus coalition. Therefore, this decision problem can be solved in a manner similar to WNC and so is in co-$NP(n)$ for NRAM.

7 Related Work

The central proposition of CGs is that agents' choices are conditioned by the number of other agents also making some choice; this resembles the premise of congestion games [14], a type of potential game [15]. Although in these games players do not seek to make agreements or form coalitions, they do aim to minimise costs which in turn are dependent on the choices of the other players. The strategies of the players each describe a set of primary factors, or resources, utilised in that strategy. The cost to each player for the use of each resource is dependent upon the number of other players also choosing a strategy requiring that resource. It has been shown [14] that every congestion game possesses at least one Nash equilibrium in pure strategies.

CGs also resemble aspects of anonymous games [16] in which the individual utility of participation in some coalition is independent of the identities of the agents concerned. In such situations other factors, including the size of the coalition become determinants of an agent's choice. In general, however, CGs are non-anonymous therefore, for example, an agent could object ($q(i, H) = 1$) to participation in any coalition in which some other, specific, agent participates.

The notion that an agent's behaviour may influence that of others is also found in imitation games [17] where two players take the roles of leader and follower; through the payoff structure the follower is motivated to act in consensus with the leader. In the terminology of CGs the leader has a quorum threshold of zero, and the follower has a quorum threshold of at most one half. McLennan and Tourky [17] find that the complexity of computing Nash equilibria in such games is no less than for the general two-player case, i.e., is $PPAD$-complete [18]. The problem of computing (existence of) equilibria in CGs is thus no more complex than it is for normal form games. The problem of verifying equilibria in CGs can be solved in time which is polynomial in the number of agents, whereas, the problem of verifying equilibria for normal form games requires time which is polynomial in the product of the number of players and of alternative strategies [19].

The notion of weak consensus resembles aspects of hedonic games [20]. For weak consensus the quorum thresholds of agent i can be understood as partitioning the set of coalitions containing i into those which i is willing to join and those which it is not. These latter coalitions are identified as those to which the agent objects. Agents in hedonic games have a complete, reflexive and transitive preference relation over those coalitions to which they may belong. Hence, for a CG under weak consensus a corresponding, simple, hedonic game can be constructed where agents are indifferent amongst those coalitions $H \subseteq G$ where $q(i, H) \leq \frac{|H|-1}{|H|}$ but strictly prefer these to coalitions where this inequality does not hold.

CGs have some similarities to Qualitative Coalitional Games [2]. It is therefore interesting to compare CGs and QCGs, especially with respect to the size of representation and the complexity of similar decision problems.

A QCG Γ may be represented as an $(n+3)$ tuple $\Gamma = \langle \mathsf{G}, \mathsf{Ag}, \mathsf{G}_1 \ldots \mathsf{G}_n, \mathsf{V} \rangle$ where $\mathsf{G}_i \subseteq \mathsf{G}$ represents each agent's $\mathsf{i} \in \mathsf{Ag}$ set of goals and $\mathsf{V} : 2^{\mathsf{Ag}} \rightarrow 2^{2^{\mathsf{G}}}$ is the characteristic function of the game mapping each possible coalition of agents to the sets of goals that coalition can achieve. In QCGs:

- A set of goals $\mathsf{G}' \subseteq \mathsf{G}$ is *feasible* for a coalition $\mathsf{C} \subseteq \mathsf{Ag}$ if $\mathsf{G}' \in \mathsf{V}(\mathsf{C})$.
- A set of goals $\mathsf{G}' \subseteq \mathsf{G}$ *satisfies* an agent $\mathsf{i} \in \mathsf{C} \subseteq \mathsf{Ag}$ if $\mathsf{G}' \cap \mathsf{G}_i \neq \emptyset$.
- A coalition $\mathsf{C} \subseteq \mathsf{Ag}$ is *successful* if there exists some set of goals $\mathsf{G}' \subseteq \mathsf{G}$ such that G' is feasible for C and G' satisfies at least all agents $\mathsf{i} \in \mathsf{C}$.

There is no immediate direct correspondence between goals in QCGs and quorum thresholds in CGs, apart from an intuition that the agent may have a lower quorum threshold for a coalition if it is more likely to achieve the agent's goals. However, the parallels between successful coalitions in QCGs and consensus coalitions in CGs are clearer: both capture the notion of it being rational for an agent to join a coalition. The worst case size of the game representation for QCGs is the characteristic function where each coalition can enforce any subset of goals. There are 2^n coalitions and 2^m subsets of goals, so the worst case size of V is $O(2^{n+m})$. For CGs, the representation is exponential in n. Complexity results for QCGs in [2] are given as a function of the size of representation, where the characteristic function is represented by a propositional formula Ψ (which as noted may be exponential in the number of agents and goals, but generally will be more concise than a naive representation of V). The successful coalition problem is NP in the size of the representation. This corresponds to our SCC and WCC problems which are linear in the number of agents (hence also in the size of representation).

8 Conclusions and future work

Consensus games are a novel approach to the problem of coalition formation in multi-agent systems. We have presented two solution concepts for CGs, corresponding to the least and greatest fixed points of f_H. Weak consensus resembles the threshold model presented by Chwe in [21] whilst strong consensus resembles the model presented by

Granovetter in [8]. The two solution concepts of strong and weak consensus represent different approaches to reaching consensus. Weak consensus requires fewer rounds (namely, one) to reach consensus compared to strong consensus, and allows consensus to be reached for a strictly larger range of quorum values: all strong consensus coalitions are weak consensus coalitions but not vice versa.

CGs bring together these two approaches to threshold modelling and, in particular, they focus on the special case of consensus, an important aspect of threshold models which has not been previously considered. CGs apply threshold models to the problem of coalition formation and so develop notions of collective rationality and coalitional stability for threshold models; problems which have not been addressed by previous work.

Our work is the first to consider computational aspects of threshold models. We have analysed the complexity of several decision problems for CGs under both individually and collectively rational solution concepts for both strong and weak notions of consensus; these results are summarised in Table 1. Whilst these membership results do not require technically demanding proofs they do represent important and fundamental computational questions concerning the verification, existence and non-existence of stable coalitions in CGs. In particular, we have shown that verification of both strong and weak consensus coalitions can be achieved in time that is polynomial (linear) in the number of agents. These results may help to explain why it is that threshold behaviours are so prevalent in both biological and sociological systems.

Table 1. Summary of complexity results (upper bounds). Note that we assume random access to indices in R, so the complexity classes are for NRAM.

	Strong Consensus	q-minimal	Weak Consensus	H-minimal
Verification (CC)	$P(n)$	*co*-$NP(n)$	$P(n)$	*co*-$NP(n)$
Existence (CE)	$NP(n)$	$NP(n)$	$NP(n)$	$NP(n)$
Non-existence (NC)	*co*-$NP(n)$	*co*-$NP(n)$	*co*-$NP(n)$	*co*-$NP(n)$

CGs as presented here treat the problem of coalition formation in an abstract sense. It is often the case that coalition formation in multi-agent systems is directed toward the achievement of goals. It would be interesting to study decision problems for CGs extended to include representations of collective action and heterogeneous goals. The work of Chwe [21] considers the mechanics and epistemics of threshold models where communication is not necessarily global; this also seems a fertile area for future work.

References

1. Kraus, S.: Negotiation and cooperation in multi-agent environments. Artificial Intelligence **94** (1997) 79–97

2. Wooldridge, M., Dunne, P.E.: On the computational complexity of qualitative coalitional games. Artificial Intelligence **158** (2004) 27–73
3. Abdou, J., Keiding, H.: Effectivity functions in social choice. Kluwer Academic (1991)
4. Jacob, E., Becker, I., Shapira, Y., Levine, H.: Bacterial linguistic communication and social intelligence. TRENDS in Microbiology **12** (2004) 366–372
5. Pratt, S.C., Sumpter, D.J.T., Mallon, E.B., Franks, N.R.: An agent-based model of collective nest choice by the ant temnothorax albipennis. Animal Behaviour **70** (2005) 1023–1036
6. Seeley, T., Visscher, P.: Group decision making in nest-site selection by honey bees. Apidologie **35** (2004) 101–116
7. Ward, A.J.W., Sumpter, D.J.T., Couzin, I.D., Hart, P.J.B., Krause, J.: From the cover: Quorum decision-making facilitates information transfer in fish shoals. Proceedings of the National Academy of Sciences **105** (2008) 6948
8. Granovetter, M.: Threshold models of collective behavior. American Journal of Sociology **83** (1978) 1420–1443
9. Schelling, T.C.: Models of segregation. The American Economic Review **59** (1969) 488–493
10. Schelling, T.C.: Micromotives and Macrobehaviour. W.W.Norton (1978)
11. Granovetter, M., Soong, R.: Threshold models of interpersonal effects in consumer demand. Journal of Economic Behavior & Organization **7** (1986) 83–99
12. Tarski, A.: A lattice-theoretical fixpoint theorem and its applications. Pacific Journal of Mathematics **5** (1955) 285–309
13. van Emde Boas, P. In: Machine models and simulations. MIT Press, Cambridge, MA, USA (1990) 1–66
14. Rosenthal, R.: A class of games possessing pure-strategy nash equilibria. International Journal of Game Theory **2** (1973) 65–67
15. Monderer, D., Shapley, L.: Potential games. Games and economic behavior **14** (1996) 124–143
16. Daskalakis, C., Papadimitriou, C.: Computing equilibria in anonymous games. In: Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07), IEEE (2007) 83–93
17. McLennan, A., Tourky, R.: Simple complexity from imitation games. Games and Economic Behavior **68** (2010) 683 – 688
18. Chen, X., Deng, X.: Settling the complexity of two-player Nash equilibrium. In: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), IEEE (2006) 261–272
19. Tadjouddine, E.: Complexity of verifying game equilibria. In Burkhard, H.D., Lindemann, G., Verbrugge, R., Varga, L., eds.: Multi-Agent Systems and Applications V. Volume 4696 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2007) 103–112
20. Bogomolnaia, A., Jackson, M.O.: The stability of hedonic coalition structures. Games and Economic Behavior **38** (2002) 201 – 230
21. Chwe, M.: Structure and strategy in collective action. American Journal of Sociology **105** (1999) 128–156