Towards Tractable Boolean Games

Paul E. Dunne
Department of Computer Science
Uinversity of Liverpool
Liverpool L69 3BX, UK
ped@csc.liv.ac.uk

Michael Wooldridge
Department of Computer Science
Uinversity of Liverpool
Liverpool L69 3BX, UK
mjw@csc.liv.ac.uk

ABSTRACT

Boolean games are a compact and expressive class of games, based on propositional logic. However, Boolean games are computationally complex: checking for the existence of pure Nash equilibria in Boolean games is Σ_2^p-complete, and it is co-NP-complete to check whether a given outcome for a Boolean game is a pure Nash equilibrium. In this paper, we consider two possible avenues to tractability in Boolean games. First, we consider the development of alternative solution concepts for Boolean games. We introduce the notion of k-bounded Nash equilibrium, meaning that no agent can benefit from deviation by altering fewer than k variables. After motivating and discussing this notion of equilibrium, we give a logical characterisation of a class of Boolean games for which k-bounded equilibria correspond to Nash equilibria. That is, we precisely characterise a class of Boolean games for which all Nash equilibria are in fact k-bounded Nash equilibria. Second, we consider classes of Boolean games for which computational problems related to Nash equilibria are easier than in the general setting. We first identify some restrictions on games that make checking for beneficial deviations by individual players computationally tractable, and then show that certain types of *socially desirable* equilibria can be hard to compute even when the standard decision problems for Boolean games are easy. We conclude with a discussion of related work and possible future work.

Categories and Subject Descriptors

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

General Terms

theory

Keywords

Boolean games, complexity, game theory, logic, Nash equilibria

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), June, 4–8, 2012, Valencia, Spain.

1. INTRODUCTION

Game-theoretic solution concepts such as Nash equilibria were originally formulated independently of considerations of whether or how they might be practically computed. Since solution concepts typically attempt to capture a notion of optimal choice in strategic settings, it is therefore not at all surprising that solution concepts are hard to compute in practice for many natural and important classes of games (see, e.g., [3, 12, 4, 15, 8]). The problem of classifying exactly the complexity of solution concepts in various settings has been an area of significant research activity over the past two decades. Most notably, the problem of classifying the complexity of computing mixed strategy Nash equilibria in 2-person strategic form games turned out to be one of the major challenges in complexity theory in the first decade of the 21st century [4].

Given that solution concepts are very often computationally complex, at least two possible routes to tractability suggest themselves:

- First, we can try to identify useful classes of games or representations of games for which solution concepts can be easily computed. In cooperative game theory, for example, the marginal contribution net representation allows for the efficient computation of the Shapley value solution concept [10].
- Second, we can develop alternative solution concepts, which lend themselves to efficient computation. For example, the notion of ϵ-Nash equilibrium has been developed, which relaxes the strict notion of Nash equilibrium by requiring that no player can gain more than ϵ in a deviation from a given strategy profile [13, p.83].

Our aim in the present paper is to consider these possibilities in the context of *Boolean games* [9, 2, 6, 7]. Boolean games are a simple, compact, and expressive class of games based on propositional logic. In a Boolean game, each player i has under its unique control a set of Boolean variables Φ_i, drawn from an overall set of Boolean variables Φ. Player i is at liberty to assign values to these variables as it chooses. The strategies or choices available to i correspond to all possible Boolean assignments that can be made to these variables. The outcome of a Boolean game is a valuation for the variables Φ, which will be composed from the individual assignments made by the players in the game to their variables. In addition, each player i has a goal that it desires to be achieved: the goal is represented as a Boolean formula γ_i, and this goal formula may contain variables under the control of other players. A player is satisfied with an overall outcome if that outcome satisfies its goal γ_i, and is unsatisfied otherwise. The fact that the achievement of one agent's goal may depend on the choices of other agents is what gives Boolean games their strategic character. Now, Boolean games are computationally complex: checking whether a Boolean game has a pure strategy Nash equilibrium is Σ_2^p-complete, while checking whether a particular outcome is a pure Nash equilibrium is co-NP-complete. If Boolean games are to find applications, then the issue of intractability must surely be addressed.

Our contribution in the present paper is twofold. Following the discussion above, we first formally define and investigate the notion of *k-bounded Nash equilibrium* for Boolean games. An outcome is a k-bounded Nash equilibrium if no agent has any incentive to deviate by flipping the value of at most k variables. We consider the computational aspects of k-bounded equilibria, and then prove a logical characterisation of those classes of Boolean games for which k-bounded Nash equilibria and standard pure strategy Nash equilibria coincide. We then move on to consider classes of Boolean games for which the computation of pure strategy Nash equilibria is easier than the general case. Again, we give a logical characterisation of such cases. We conclude with discussion and some issues for future research.

2. GAMES AND SOLUTION CONCEPTS

Let $\mathcal{G}$ be a class of games, and let Ω be the set of *outcomes* for these games; for the purposes of this discussion, it does not matter exactly what the games and outcomes are. A *solution concept* for $\mathcal{G}$ can be understood as a function:

$$\sigma : \mathcal{G} \to \mathbf{2}^{\Omega}.$$

That is, a solution concept identifies with every game a subset of outcomes; intuitively, those that are "rational" according to the solution concept in question.

The obvious computational problems associated with such a solution concept σ for a class of games $\mathcal{G}$ are as follows:

- NON-EMPTINESS:
  Given some $G \in \mathcal{G}$, is it the case that $\sigma(G) \neq \emptyset$?
- MEMBERSHIP:
  Given $G \in \mathcal{G}$ and $\omega \in \Omega$, is it the case that $\omega \in \sigma(G)$?
- COMPUTATION:
  Given $G \in \mathcal{G}$, exhibit some ω such that $\omega \in \sigma(G)$.

In the present paper, we will be concerned largely with the first two of these problems. Now, for many important classes of games, these problems are computationally hard [3, 12, 4, 15, 8]. As mentioned in the introduction, there are at least two ways to approach this problem:

1. *Develop alternative solution concepts, which lend themselves to being computed efficiently.* For example, the notion of ϵ-Nash equilibrium has been developed, which relaxes the strict notion of Nash equilibrium by requiring that no player can gain more than ϵ in a deviation from a strategy profile [13, p.83].
2. *Try to identify useful classes of games or representations of games for which solution concepts can be easily computed.* In cooperative game theory, for example, the marginal contribution net representation allows for the efficient computation of the Shapley value solution concept [10].

With respect to the first proposal, let us make the discussion a little more formal. Suppose we have a solution concept σ for a class of games $\mathcal{G}$ such that the associated computational problems (NON-EMPTINESS, MEMBERSHIP, COMPUTATION) are intractable (NP-hard or worse). Then we might try to develop an alternative solution concept $\hat{\sigma}$, which "approximates" σ, but which is tractable. Note that here we mean "approximate" in the informal everyday sense, rather than the formal sense of approximation algorithms and FPTAS [1] (although of course looking for FPTAS would be a very natural approach). Now, how might σ and its "approximation", $\hat{\sigma}$, be related? We can consider two natural properties, as follows.

- We say $\hat{\sigma}$ is a *sound approximation* of a solution concept σ for a class of games $\mathcal{C} \subseteq \mathcal{G}$ if

  $$\forall G \in \mathcal{C} : \quad \hat{\sigma}(G) \subseteq \sigma(G)$$

  i.e., the solutions proposed by $\hat{\sigma}$ are a subset of those proposed by σ.
- We say $\hat{\sigma}$ is a *complete approximation* of a solution concept σ for a class of games $\mathcal{C} \subseteq \mathcal{G}$ if

  $$\forall G \in \mathcal{C} : \quad \hat{\sigma}(G) \supseteq \sigma(G)$$

  i.e., the solutions proposed by $\hat{\sigma}$ are a superset of those proposed by σ.

In the limit, where $\hat{\sigma}$ is a sound and complete approximation of σ w.r.t. the class $\mathcal{C} = \mathcal{G}$ of all games, then σ and $\hat{\sigma}$ would be identical, and it would then be no easier to compute $\hat{\sigma}$ than σ. A typical situation, with respect to the class $\mathcal{G}$ of all games, is that we will have approximate solution concepts $\hat{\sigma}$ that are complete (all solutions according to σ are solutions according to $\hat{\sigma}$) but not sound (not all solutions according to $\hat{\sigma}$ are solutions according to σ). This is exactly the situation with ϵ-Nash equilibrium, for example: all "exact" Nash equilibria are ϵ-Nash equilibria, but in general, (i.e., where $\epsilon > 0$), not all ϵ-Nash equilibria will be "exact" Nash equilibria. The fact that an approximate solution concept is complete but not sound captures our intuitions about relaxing the requirements for optimality inherent in exact solution concepts: an approximate solution concept will often admit more solutions than its exact counterpart, thus (we hope) making approximate solutions easier to find.

If we have an approximate solution concept $\hat{\sigma}$, then one interesting and important question is the following: can we identify a class of games $\mathcal{C} \subset \mathcal{G}$ such that $\hat{\sigma}$ is a sound and complete approximation to σ with respect to $\mathcal{C}$, *even though $\hat{\sigma}$ is not sound and complete with respect to the class of all games $\mathcal{G}$*? If we can do this, and the class $\mathcal{C}$ corresponds to games that are of practical value, then this means that the approximate solution concept is in fact all we need: we do not need to look for exact solutions σ, since these will in any case be given by $\hat{\sigma}$.

In the present paper, we will focus on *bounded* approximations to Nash equilibria for Boolean games, which will be complete but not sound with respect to the class of all games, and we will identify classes of games for which the bounded solution concept is both sound and complete.

3. BOOLEAN GAMES

We now present the formal framework of propositional logic and Boolean games that we use throughout the remainder of this paper. Our presentation is fairly standard [9, 2, 6, 7].

Propositional Logic: Let $\mathbb{B} = \{\top, \bot\}$ be the set of Boolean truth values, with "$\top$" being truth and "$\bot$" being falsity. We will abuse notation a little by using $\top$ and $\bot$ to denote both the syntactic constants for truth and falsity respectively, as well as their semantic counterparts. Let $\Phi = \{p, q, \ldots\}$ be a (finite, fixed, non-empty) vocabulary of Boolean variables, and let $\mathcal{L}$ denote the set of (well-formed) formulae of propositional logic over Φ, constructed using the conventional Boolean operators ("$\wedge$", "$\vee$", "$\rightarrow$", "$\leftrightarrow$", and "$\neg$"), as well as the truth constants "$\top$" and "$\bot$". Where $\varphi \in \mathcal{L}$, we let $vars(\varphi)$ denote the (possibly empty) set of Boolean variables occurring in φ (e.g., $vars(p \wedge q) = \{p, q\}$). A *valuation* is a total function $v : \Phi \rightarrow \mathbb{B}$, assigning truth or falsity to every Boolean variable. We write $v \models \varphi$ to mean that the propositional formula φ is true under, or satisfied by, valuation v, where the satisfaction relation "$\models$" is defined in the standard way. Let $\mathcal{V}$ denote the set of all valuations over Φ. We write $\models \varphi$ to mean that φ is a tautology. We denote the fact that $\models \varphi \leftrightarrow \psi$ by $\varphi \equiv \psi$. We use some additional definitions, as follows:

- A *literal*, ℓ, is either (i) a Boolean variable or the negation of a Boolean variable, or (ii) a Boolean constant (i.e., a member of $\mathbb{B}$) or the negation of a Boolean constant.
- A *clause*, C, is a disjunction of literals, i.e., a formula of the form $C = \ell_1 \vee \cdots \vee \ell_m$.
- A *Horn clause* is a clause in which at most one literal is not negated.
- A formula is in *Conjunctive Normal Form* (CNF) if it is a conjunction of clauses, i.e., is of the form $\varphi = C_1 \wedge \cdots \wedge C_l$, where each C_i, $(1 \leq i \leq l)$ is a clause.
- A CNF formula $\varphi = C_1 \wedge \cdots \wedge C_l$ is in u-CNF, if each clause C_i $(1 \leq i \leq l)$ contains at most u literals.
- A CNF formula $\varphi = C_1 \wedge \cdots \wedge C_l$ is in u-clause CNF if it contains no more than u clauses (i.e., $l \leq u$).

  Notice that there is an important difference between u-CNF and u-clause CNF: the former constrains the number of literals permitted in a clause, but does not constrain the number of clauses permitted in a formula; while the latter constrains the number of clauses permitted in a formula, but does not constrain the number of literals that appear in clauses.
- A CNF formula $\varphi = C_1 \wedge \cdots \wedge C_l$ is said to be in Horn clause form if for all $1 \leq i \leq l$, the clause C_i is a Horn clause.

The *satisfiability problem* for formulae φ is the problem of determining whether there exists a valuation v such that $v \models \varphi$. For arbitrary CNF formulae, this problem is of course NP-complete; for 2-CNF formulae, and for Horn clause formulae, the satisfiability problem is decidable in polynomial time (see, e.g., [11]).

Agents and Variables: The games we consider are populated by a set $N = \{1, \ldots, n\}$ of *agents* – the players of the game. Each agent is assumed to have a *goal*, characterised by an $\mathcal{L}$-formula: we write γ_i to denote the goal of agent $i \in N$. Agents $i \in N$ each *control* a (possibly empty) subset Φ_i of the overall set of Boolean variables. By "control", we mean that i has the unique ability within the game to set the value (either $\top$ or $\bot$) of each variable $p \in \Phi_i$. We will require that $\Phi_i \cap \Phi_j = \emptyset$ for $i \neq j$, and that $\Phi_1 \cup \cdots \cup \Phi_n = \Phi$ (i.e., $\Phi_1, \ldots, \Phi_n$ partition Φ).

When playing a Boolean game, the primary aim of an agent i will be to choose an assignment of values for the variables Φ_i under its control so as to satisfy its goal γ_i. The difficulty is that γ_i may contain variables controlled by other agents $j \neq i$, who will also be trying to choose values for their variables Φ_j so as to get their goals satisfied; and their goals in turn may be dependent on the variables Φ_i. A *choice* for agent $i \in N$ is a function $v_i : \Phi_i \rightarrow \mathbb{B}$, i.e., an allocation of truth or falsity to all the variables under i's control. Let $\mathcal{V}_i$ denote the set of choices for agent i.

Outcomes: An *outcome* is a collection of choices, one for each agent. Formally, an outcome is a tuple $(v_1, \ldots, v_n) \in \mathcal{V}_1 \times \cdots \times \mathcal{V}_n$. Notice that an outcome defines a value for all variables, and we will often think of outcomes as valuations, for example writing $(v_1, \ldots, v_n) \models \varphi$ to mean that the valuation defined by the outcome $(v_1, \ldots, v_n)$ satisfies formula $\varphi \in \mathcal{L}$.

Boolean Games: A Boolean game, G, is a $(2n+2)$-tuple:

$$G = \langle N, \Phi, \Phi_1, \ldots, \Phi_n, \gamma_1, \ldots, \gamma_n \rangle$$

where $N = \{1, \ldots, n\}$ is a set of agents, $\Phi = \{p, q, \ldots\}$ is a finite set of Boolean variables, $\Phi_i \subseteq \Phi$ is the set of Boolean variables under the unique control of $i \in N$, and $\gamma_i \in \mathcal{L}$ is the goal of agent $i \in N$.

Utility: We now introduce a model of utility for our games. The basic idea is that an agent will strictly prefer all outcomes in which it gets its goal achieved over all outcomes where it does not. We capture this in utility functions $u_i(\cdots)$ defined over outcomes $(v_1, \ldots, v_n)$:

$$u_i(v_1, \ldots, v_n) = \begin{cases} 1 & \text{if } (v_1, \ldots, v_n) \models \gamma_i \\ 0 & \text{otherwise.} \end{cases}$$

Nash Equilibrium: Let $(v_1, \ldots, v_i, \ldots, v_n)$ be an outcome. We say that a player i has a *beneficial deviation* if there exists a choice $v_i' \in \mathcal{V}_i$ for i such that $u_i(v_1, \ldots, v_i', \ldots, v_n) > u_i(v_1, \ldots, v_i, \ldots, v_n)$. In this case, v_i' serves as a witness to the beneficial deviation. It will be useful to refer to the following fact later.

OBSERVATION 1. *Suppose $(v_1, \ldots, v_i, \ldots, v_n) \in \mathcal{V}_1 \times \cdots \times \mathcal{V}_i \times \cdots \times \mathcal{V}_n$ and $v_i' \in \mathcal{V}_i$. Then v_i' is a beneficial deviation for i from $(v_1, \ldots, v_i, \ldots, v_n)$ iff:*

1. $(v_1, \ldots, v_i, \ldots, v_n) \not\models \gamma_i$ *and*
2. $(v_1, \ldots, v_i', \ldots, v_n) \models \gamma_i$.

We then say an outcome $(v_1, \ldots, v_n)$ is a Nash equilibrium if no player has a beneficial deviation. We denote the Nash equilibrium outcomes of a game G by $NE(G)$. It may of course be that $NE(G) = \emptyset$.

Referring back to our discussion in the preceding section, there are two obvious decision problems relating to Nash equilibria in Boolean games: NON-EMPTINESS (given a game G, is it the case that $NE(G) \neq \emptyset$?) and MEMBERSHIP (given a game G and an outcome $(v_1, \ldots, v_n)$ for G, is it the case that $(v_1, \ldots, v_n) \in NE(G)$?). Unfortunately, it is known that both of these problems are computationally complex [2]:

PROPOSITION 1. *The* NON-EMPTINESS *problem for Boolean games is Σ_2^p-complete; the* MEMBERSHIP *problem for Boolean games is co-NP-complete.*

4. K-BOUNDED EQUILIBRIA

We will now define a new class of equilibria for Boolean games, which we will call *k-bounded equilibria*. To motivate this new class, take an agent i who is trying to decide whether he has any beneficial deviation from an outcome $(v_1, \ldots, v_i, \ldots, v_n)$. From Proposition 1, this problem is in general NP-complete. Intuitively, player i is trying to decide whether he can get his goal achieved by flipping the value of some subset of his variables. The complexity in this problem arises because i must consider $2^{|\Phi_i|}$ sets of variables in this evaluation. It follows that our agent's task will be simpler if we can eliminate some of these sets of variables from player i's consideration.

Where $\{v_i, v_i'\} \subseteq \mathcal{V}_i$ are choices for player $i \in N$, let us define the *distance* between them as being the number of variables that have different values in the two valuations; we denote this value by $\delta(v_i, v_i')$:

$$\delta(v_i, v_i') = |\{x \in \Phi_i \mid v_i(x) \neq v_i'(x)\}|.$$

We sometimes refer to $\delta(v_i, v_i')$ as the *size* of the deviation v_i'. Now, given an outcome $(v_1, \ldots, v_i, \ldots, v_n)$ and a value $k \in \mathbb{N}, k > 1$ we will say a player i has a *k-bounded beneficial deviation* if there is some $v_i' \in \mathcal{V}_i$ such that:

1. $\delta(v_i, v_i') \leq k$; and
2. $u_i(v_1, \ldots, v_i', \ldots, v_n) > u_i(v_1, \ldots, v_i, \ldots, v_n)$.

We will say an outcome $(v_1, \ldots, v_i', \ldots, v_n)$ is a *k-bounded Nash equilibrium* if no player has a k-bounded beneficial deviation. Let the set of k-bounded Nash equilibria of game G be denoted by $NE_k(G)$.

EXAMPLE 1. *Consider the Boolean game,*

$$G^{(t)} = \langle \{a_1, a_2\}, \Phi, \Phi_1, \Phi_2, \gamma_1, \gamma_2 \rangle$$

in which $\Phi_1 = \{x_1, \ldots, x_t\}$, $\Phi_2 = \{y_1, \ldots, y_t\}$ $(t \geq 1)$ and

$$\gamma_1 = \left(\bigvee_{i=1}^{t} y_i \right) \wedge \left(\bigwedge_{j=1}^{t} x_j \right)$$

$$\gamma_2 = \left(\bigvee_{i=1}^{t} x_i \right) \wedge \left(\bigwedge_{j=1}^{t} y_j \right)$$

The outcome v in which $x_1 = y_1 = \top$, and $x_i = \bot$, $y_i = \bot$ for all $i \neq 1$, is a k-bounded equilibrium for all $k < t-1$. It is not, however, a Nash equilibrium: neither goal is satisfied but by changing the $t-1$ variables under its control to $\top$ both agents can realise their goals.

How does the notion of k-bounded equilibrium relate to the general concept of Nash equilibrium? We have:

PROPOSITION 2. *The solution concept of k-bounded Nash equilibrium is a complete but unsound approximation for pure strategy Nash equilibria in Boolean games. Formally:*

1. *There exist Boolean games G and bounds $k \in \mathbb{N}, k \geq 1$ such that $NE_k(G) \not\subseteq NE(G)$.*
2. *For all Boolean games G and bounds $k \in \mathbb{N}, k \geq 1$, we have $NE(G) \subseteq NE_k(G)$.*

PROOF. Example 1 illustrates point (1) (in fact, it is easily seen that the construction of Example 1 shows that for all $k \geq 2$ there are games in which $NE_{k-1}(G) \subset NE_k(G)$, i.e. k-bounded equilibria are more general than $(k-1)$-bounded). For point (2), observe that if an outcome is stable in the general sense, then no player has any beneficial deviation; and in particular, no player has any beneficial deviation of size $\leq k$. $\square$

Now, for the notion of k-bounded equilibrium to be of anything other than purely theoretical interest, it must reduce the complexity of the reasoning task faced by agents. Of course, it is easy to see that, with respect to worst case asymptotic analysis, k-bounded Nash equilibria present no advantages over Nash equilibria, since if we set the bound k to

$$k = \max\{|\Phi_i| \mid i \in N\}$$

then k-bounded Nash equilibrium collapses to the standard notion of Nash equilibrium for Boolean games. However, we will now show that nevertheless, by constraining possible deviations v_i' such that $\delta(v_i, v_i') \leq k$, we can dramatically reduce the search space of possible deviations. Formally, where G is a game containing a player i, let $v_i \in \mathcal{V}_i$ be a choice for i, and let $k \in \mathbb{N}, k \geq 1$ be a bound, then we have:

$$|\{v_i' \in \mathcal{V}_i \mid \delta(v_i, v_i') \leq k\}| = \sum_{j=1}^{k} \binom{|\Phi_i|}{j}.$$

From standard combinatorics, it follows that if we set the deviation bound k so that $k < |\Phi_i|/2$ (for example), then the search space will for a beneficial deviation will be less than half the search space for general deviations. We thus have an exponential reduction in the size of the search space when looking for k-bounded beneficial deviations, compared to the case for pure Nash equilibria in general.

In fact, from recent results of Szeider [14] it turns out that this upper bound can often be signifcantly improved.

PROPOSITION 3. *Let G be a Boolean game in which every goal formula, γ_i, is expressed as a CNF formula containing at most t clauses, where t is a arbitrary but fixed natural number, and v be an outcome. Deciding if v is a k-bounded equilibrium is fixed-parameter tractable (with respect to the parameter k), i.e., there is a decision algorithm whose running time is bounded above by $f(k)$ $poly(|G|)$, where $|G|$ is the number of bits needs to encode the game G, and $f(\cdots)$ is a function whose value depends* only *on k.*

PROOF. Szeider [14] shows that instances of the so-called "k-FLIP SAT" problem, whereby given a CNF formula, F, and assignment, α to its variables, it is required to decide if there is some assignment β satisfying F and having $\delta(\beta, \alpha) \leq k$, may be decided in $(f(k) + t^3)poly(|F|)$ steps. The "k-FLIP SAT" problem, however, is exactly that of deciding if γ_i, after simplification to $\gamma_i^{\Phi_i}$, i.e., the CNF formula arising by applying the variable settings to $\Phi \setminus \Phi_i$, is satisfiable by changing the values of (exactly) k variables. We can thus decide if i has a k-bounded beneficial deviation just by considering each $1 \leq l \leq k$, deciding if $\gamma_i^{\Phi_i}$ is a positive instance of l-FLIP SAT with respect to the assignment v_i of values currently set in Φ_i. In total, v will be accepted as a k-bounded equilibrium if no agent suceeds in identifying a beneficial deviation via this process. □

Now, it might seem that k-bounded equilibria are much weaker than general Nash equilibria, in the sense that a player may well have a beneficial deviation from an outcome in the general case, but not if we restrict ourselves to k-bounded Nash equilibria. But in fact, in some classes of games, this is not an issue: the following proposition shows that, in certain useful and important classes of games, *general Nash equilibria and k-bounded Nash equilibria coincide.*

PROPOSITION 4. *The solution concept of k-bounded Nash equilibrium is a sound and complete approximation of pure strategy Nash equilibria for the class of Boolean games in which every agent has a goal formula that is logically equivalent to a propositional formula in k-clause CNF. In other words, for all Boolean games*

$$G = \langle N, \Phi, \Phi_1, \dots, \Phi_n, \gamma_1, \dots, \gamma_n \rangle,$$

if there exists some $k \in \mathbb{N}, k \geq 1$ such that for every agent $i \in N$ there exists a k-clause CNF formula γ_i' such that $\gamma_i \equiv \gamma_i'$, then we have $NE(G) = NE_k(G)$.

PROOF. Consider an arbitrary outcome $(v_1, \dots, v_i, \dots, v_n)$ and a player $i \in N$. We claim that player i has a k-bounded beneficial deviation iff the player has a general beneficial deviation. The left-to-right implication is obvious. For the right-to-left implication, suppose the player has a beneficial deviation, call it v_i'. Then from Observation 1, we know that:

- $(v_1, \dots, v_i, \dots, v_n) \not\models \gamma_i$
- $(v_1, \dots, v_i', \dots, v_n) \models \gamma_i$

From the conditions of the Proposition, we can infer:

- $(v_1, \dots, v_i, \dots, v_n) \not\models \gamma_i'$
- $(v_1, \dots, v_i', \dots, v_n) \models \gamma_i'$.

The existence of a k-bounded beneficial deviation may then be seen as follows. Since γ_i' is in k-clause CNF, it is the conjunction of no more than k clauses: $\gamma_i' = C_1 \wedge \cdots \wedge C_k$, where each C_i is a disjunction of literals. The assignment v_i' need only satisfy at most one literal from each clause C_i, and so to satisfy γ_i', (and hence γ_i), we only need to flip at most k variables compared to v_i. □

Now, the implication of this result is that if agents have goals that are logically equivalent to k-clause CNF formulae, then *we don't need to consider arbitrary deviations: k-bounded deviations are all we need.* It follows that in such games, the search space for possible deviations can be dramatically reduced, compared to general Nash equilibria.

As an aside, observe that *every* propositional logic formula can be converted into CNF. So, if we start with a Boolean game in which goals are arbitrary propositional formulae, we can translate each formula into an equivalent CNF form, and take k to be the largest number of clauses of any CNF formula in the resulting game. If we end up with $k < \max\{|\Phi_i| \mid i \in N\}$, then this tells us that we can rule out a potentially large fraction of the search space when looking for beneficial deviations, as we only have to consider k-bounded Nash equilibria. However, in the worst case, translation to CNF can result in an exponential blow-up in the number of clauses in the formula: $k = O(2^{|\Phi|})$. In such cases, we would clearly obtain no benefit from k-bounded equilibria.

5. TRACTABLE BOOLEAN GAMES

An alternative to developing solution concepts that are easy (or at least, easier) to compute than exact solutions is to consider classes of games for which exact solution concepts are easy to compute. In our case, consider the problem of determining whether a player $i \in N$ has a beneficial deviation in a Boolean game. This involves the player considering $2^{|\Phi_i|}$ choices, to see whether it can get its goal achieved through making one of these. So, to what extent can we identify classes of games for which checking for beneficial deviations is computationally easy? Well, as a starting point, the following is very easy to see:

OBSERVATION 2. *Let $\mathcal{C} \subseteq \mathcal{L}$ be a class of propositional logic formulae with a polynomial time satisfiability problem. The* MEMBERSHIP *problem for $G = \langle N, \Phi, \Phi_1, \dots, \Phi_n \gamma_1, \dots, \gamma_n \rangle$ with $\{\gamma_1, \dots, \gamma_n\} \subseteq \mathcal{C}$ is decidable in polynomial time, and the* NON-EMPTINESS *problem is in NP.*

So, for example, if the goal formulae of all players are expressed in, e.g., Horn clause form, or 2-CNF, then checking whether an outcome is a Nash equilibrium or not is polynomial time decidable, and checking whether there exists a stable outcome is in NP. However, as we will now see, we can in fact significantly strengthen this result; to do this, however, we need some further notation and terminology.

First, where $S = \{1, \dots, k\}$ is a subset of players and $\{v_1, \dots, v_k\}$ is a collection of choices, one for each player $i \in S$, define a function w by:

$$w_{\{v_1,\dots,v_k\}}(x) = \begin{cases} v_1(x) & \text{if } x \in \text{dom } v_1 \\ v_2(x) & \text{if } x \in \text{dom } v_2 \\ \cdots & \cdots \\ v_k(x) & \text{if } x \in \text{dom } v_k \end{cases}$$

Where φ is a propositional logic formula, and $\{v_1, \dots, v_k\}$ is a collection of choices, one for each player in $S = \{1, \dots, k\}$, then we will denote by $\varphi[v_1, \dots, v_k]$ the formula obtained from φ by systematically replacing each variable x such that

$$x \in \text{dom } v_1 \cup \cdots \cup \text{dom } v_k$$

by the Boolean value $w_{\{v_1,\dots,v_k\}}(x)$. Finally, the *reduction* of $\varphi[v_1, \dots, v_k]$ will be denoted by $\varphi^*[v_1, \dots, v_k]$, and is defined to be the propositional logic formula obtained from $\varphi[v_1, \dots, v_k]$ by carrying out the following two steps:

1. deleting any clause containing $\top$ or $\neg\bot$;
2. deleting $\bot$ or $\neg\top$ from any clause in which they occur.

If the resulting formula contains any empty clause (i.e., all literals have been deleted from the original clause), then we replace the whole formula by $\bot$.

The soundness of these simplification steps is clear from basic propositional reasoning.

Now, given a player i's goal γ_i, we will denote the *range* of γ_i by $rng(\gamma_i)$, and define this to be the following set of formulae:

$$rng(\gamma_i) = \{\gamma_i^*[v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n] \mid (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n) \in \prod_{i \in N \setminus \{i\}} \mathcal{V}_i\}.$$

Thus, intuitively, $rng(\gamma_i)$ is the set of formulae that could be obtained from γ_i by simplifying it under all possible combinations of choices made by other players. Notice that if $\psi \in rng(\gamma_i)$, then the only variables occurring in ψ will be controlled by i, i.e., for each $\psi \in rng(\gamma_i)$, we have $vars(\psi) \subseteq \Phi_i$.

Now, let $\mathcal{C} \subseteq \mathcal{L}$ be a class of propositional logic formulae, and let $G = \langle N, \Phi, \Phi_1, \ldots, \Phi_n, \gamma_1, \ldots, \gamma_n \rangle$ be a Boolean game. Then we say *the range of G is in $\mathcal{C}$* if:

$$rng(\gamma_i) \subseteq \mathcal{C} \quad \text{for all } i \in N.$$

Given this, we can strengthen Observation 2 as follows.

PROPOSITION 5. *Let $\mathcal{C} \subseteq \mathcal{L}$ be a class of propositional logic formulae with a polynomial time satisfiability problem, and let $G = \langle N, \Phi, \Phi_1, \ldots, \Phi_n, \gamma_1, \ldots, \gamma_n \rangle$ be a game such that the range of G is in $\mathcal{C}$. Then the* MEMBERSHIP *problem for G is decidable in polynomial time and the* NON-EMPTINESS *problem is in NP.*

Notice Proposition 5 is *not* the same as Observation 2, above (although it is related). It is a much stronger result: it *does not* require that the goal formulae γ_i are in a tractable form; only that *the range of the goal formulae are tractable*, (i.e., the formulae obtained by simplifying the goal formulae under the possible choices for all other players). This is a very different, and much more powerful result than that of Observation 2. In particular, for every player $i \in N$, the only constraints it imposes on the goal formulae of player i relate to the variables actually controlled by player i; essentially no constraints are placed on the variables of players $j \neq i$. Thus we obtain:

PROPOSITION 6. *For the following classes of games, the* MEMBERSHIP *problem is decidable in polynomial time, while the* NON-EMPTINESS *problem is NP-complete:*

1. *Games in which for all players $i \in N$, if $\gamma_i = C_1 \wedge \cdots \wedge C_l$, then we have $|vars(C_j) \cap \Phi_i| \leq 2$ for all $1 \leq j \leq l$.*
2. *Games in which for all players $i \in N$, if $\gamma_i = C_1 \wedge \cdots \wedge C_l$, then for all $1 \leq j \leq l$, the clause C_j contains at most one unnegated element of Φ_i.*

PROOF. We will do the proof for point (1); the second point is similar. So consider MEMBERSHIP. Take an arbitrary outcome $(v_1, \ldots, v_{i-1}, v_i, v_{i+1}, \ldots, v_n)$.

Since for all players $i \in N$, if $\gamma_i = C_1 \wedge \cdots \wedge C_l$, then we have $|vars(C_j) \cap \Phi_i| \leq 2$ for all $1 \leq j \leq l$, then it follows that $\gamma_i^*[v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n]$ is in 2-CNF, and contains only variables in Φ_i. From Proposition 1 we can see that i has a beneficial deviation from

$$(v_1, \ldots, v_{i-1}, v_i, v_{i+1}, \ldots, v_n)$$

if there exists a choice $v_i' \in \mathcal{V}_i$ such that

$$v_i' \models \gamma_i^*[v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n].$$

Since $\gamma_i^*[v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n]$ is in 2-CNF, this check can be done in polynomial time. We now prove NP-completeness of the NON-EMPTINESS problem. Membership is obvious by "guess and check". For hardness, we reduce SAT. Let φ be a SAT instance. We construct a game G_φ, satisfying the conditions of the proposition, such that $NE(G_\varphi) \neq \emptyset$ iff φ is satisfiable. Assume φ has l clauses, $\varphi = C_1 \wedge \cdots \wedge C_l$, and $vars(\varphi) = \{x_1, \ldots, x_k\}$. We introduce one additional variable, z. We define $k+1$ agents, with:

- player $1 \leq i \leq k$ controlling variable x_i and having goal y; and
- player $k+1$ controlling variable y and having goal

$$\gamma_{k+1} = (\varphi \wedge z) \vee \neg(y \leftrightarrow z).$$

- player $k+2$ controls variable z and has goal

$$\gamma_{k+2} = (\neg\varphi) \wedge (y \leftrightarrow z)$$

(Notice that γ_{k+1} and γ_{k+2} can trivially be translated into the form required by the Proposition.) We claim that φ is satisfiable iff $N(G_\varphi) \neq \emptyset$. If φ is satisfiable, then take any satisfying assignment for φ and set $y = z = \top$. We claim that this outcome is stable: for observe that in this case, players $1 \leq i \leq k+1$ have their goal achieved, and so cannot benefit by deviating; and player $k+2$ does not get his goal achieved, but has no beneficial deviation. Now we claim that if φ is unsatisfiable then $NE(G_\varphi) = \emptyset$. For consider any outcome. First observe that since φ is unsatisfiable, any outcome will falsify φ. Thus, player $k+2$ will have his goal achieved if the outcome assigns y and z the same value; if the outcome does not give them the same value then player $k+2$ has a beneficial deviation. However, if player $k+2$ deviates to give y and z the same value, then player $k+1$ would have a beneficial deviation. Hence no outcome can be a Nash equilibrium. □

5.1 Equilibria that Maximise Social Welfare

The results above indicate that, in the event that $rng(\gamma_i)$ falls within a class of formulae that have a polynomial time satisfiability problem, then determining whether a given outcome is an equilibrium for the instance is tractable. It may, however, often be the case that we do not merely wish to accept *any* equilibrium, but would, if possible, prefer to identify one which satisfies some notion of "optimality". In the case of Boolean games a natural way of distinguishing between equilibria is to associate each with its *social welfare*, which we define as the total number of agents whose goal is satisfied. As a very simple example of a Boolean game in which there are equilibria in which this measure varies significantly, consider $n = 2m$ agents, $N = \langle a_1, \ldots, a_m, b_1, \ldots, b_m \rangle$ each of which has control over exactly one propositional variable from

$$\Phi = \langle x_1, \ldots, x_m, y_1, \ldots, y_m \rangle$$

so that i determines the value of x_i and b_i that of y_i. Suppose the goal formula for i is $x_i \wedge y_i$, and that for b_i is $x_i \vee \neg y_i$. It is easily checked that the outcome in which all variables are assigned $\perp$ is a Nash equilibrium for this game that satisfies *only* the goals of the b_i agents. On the other hand, the assignment in which all variables are assigned $\top$ is also a Nash equilibrium, but one which that satisfies the goals of *all* agents. Intuitively, the latter seems preferable to the former as a solution.[1] This motivates the following decision question, which we present, at first, in its most general form:

> GOAL MAXIMIZATION (GM):
> Given a Boolean game, G involving n agents and $t \in \mathbb{N}$ such that $1 \leq t \leq n$, is there an outcome v for which $v \in \mathcal{N}(G)$ *and* $|\{\gamma_i \ : \ v \models \gamma_i\}| \geq t$?

It is not hard to show that, even if G admits a polynomial time process for deciding membership, this is not sufficient to ensure GM is tractable.

PROPOSITION 7. *GM is NP-complete even if instances are restricted to those in which every goal formula is in 2-CNF, which we will denote* GM^2.

PROOF. That GM is in NP follows by simply guessing an outcome, v, and checking that v satisfies at least t goals. For NP-hardness, we recall that the so-called MAX-2-SAT problem – deciding if a given 2-CNF, F, has an assignment that satisfies at least K of its clauses – is NP–complete. We show that MAX-2-SAT is polynomially reducible to GM^2.

Given $\langle F, K\rangle$ an instance of MAX-2-SAT in which F uses n variables, $\{x_1, \ldots, x_n\}$ and has r clauses, form an instance, $\langle G_{\langle F,K\rangle}, t\rangle$ of GM^2 as follows: $G_{\langle F,K\rangle}$ has $r+1$ agents, with $\gamma_i = C_i$ the i'th clause of F for $1 \leq i \leq r$, and $\gamma_{r+1} \equiv \top$. Set $\Phi = \{x_1, \ldots, x_n\}$, $\Phi_i = \emptyset$ when $1 \leq i \leq r$ and $\Phi_{r+1} = \Phi$. To complete the instance t is fixed to $K+1$. We claim that $\langle G_{\langle F,K\rangle}, t\rangle$ is a positive instance of GM^2 iff $\langle F, K\rangle$ is a positive instance of MAX-2-SAT. Trivially, any assigment α to $\langle x_1, \ldots, x_n\rangle$ that satisfies at least K clauses of F, immediately yields an outcome that achieves the goals of the corresponding K agents and since γ_{r+1} is always satisfied this outcome satisfies $t = K+1$ goals. Furthermore the outcome is in $\mathcal{N}(G_{\langle F,K\rangle})$: i $(1 \leq i \leq r)$ cannot deviate (no variables are under its control) and a_{r+1} has no reason to deviate (its goal is already satisfied). Conversely, should there be an outcome $v \in \mathcal{N}(G_{\langle F,K\rangle})$ satisfying at least $t = K+1$ goals then, since this outcome must satisfy γ_{r+1} it follows that at least K goals from $\{\gamma_1, \ldots, \gamma_r\}$ are satisfied so that the corresponding assignment witnesses $\langle F, K\rangle$ as a positive instance of MAX-2-SAT. □

5.2 Utility as Reachability

The forms taken by utility functions as described are somewhat restrictive in that these fail to model the possibility that (assuming, say, negotiation with other agents can take place) despite not having its goal satisfied with a particular outcome, an agent may, in fact be "close to" realising its intended goal. To make this idea more precise, suppose we define the *t-reachable utility* of an outcome to i (denoted w_i^t to distinguish from our standard notion of utility u_i) by

$$w_i^t(v) = \begin{cases} 1 - \min\{\ r \ : \ r \leq t \ \text{ and } \exists v' \text{ with } \delta(v,v') = r \text{ and } v' \models \gamma_i\}/|\Phi| \\ 0 \quad \text{if no suitable } v' \text{ exists} \end{cases}$$

Notice that for any outcome, v, $w_i^t(v) \geq w_i^0(v) = u_i(v)$ and captures the behaviour that there is a *possibility* of γ_i being achieved (even though this may not be completely within i's control).

It turns out, although this new form appears superficially computationally more demanding, if we limit attention to k-bounded deviations then we can still identify tractable versions of the membership problem.

[1]Endriss *et al.* discuss taxation-based mechanisms that incentivise players to socially desirable Nash equilibria [7].

PROPOSITION 8. *Let $\mathcal{C}$ be any class of propositional formulae for which k-FLIP SAT is fixed-parameter tractable wrt to parameter k. Let $t \in \mathbb{N}$ be fixed, G be a Boolean game with agent utilities captured through w_i^t. If all goal formulae are in the class $\mathcal{C}$ then for any outcome, $v \in NE_t(G)$ is polynomial time decidable.*

PROOF. Suppose that v is an outcome. By definition $v \in NE_t(G)$ iff no i has a t-bounded beneficial deviation, i.e. letting $v//\alpha$ denote the outcome obtained from v by the values currently assigned to Φ_i in v being replaced by α, it follows that $v \in NE_t(G)$ if and only if for each i:

$$\forall\ \alpha \in \langle \top, \bot\rangle^{|\Phi_i|}\ \delta(v, v//\alpha) \leq t \ \Rightarrow\ w_i^t(v//\alpha) \leq w_i^t(v)$$

In order to test if i has a beneficial deviation (under the new notion of utility) it is first necessary to compute $w_i^t(v)$, i.e. to determine if by changing the values of $0, 1, \ldots, t$ variables in the outcome v it is the case that γ_i can be achieved. This, however, is simply the r-FLIP SAT problem and the goal formulae are restricted to those within some fixed-parameter tractable (with parameter r) class. Thus we can compute $w_i^t(v)$ efficiently and it remains only to test if there is a t-bounded deviation which improves upon this, i.e whether by altering at most t variables within the control of i (leaving the values assigned to other variables unchanged), it is possible to to construct an outcome v' for which $w_i^t(v') > w_i^t(v)$. this, however, can (at worst) be carried out just by enumerating through the $O(|\Phi_i|^t)$ possible deviations. □

We note, however, that very simple negotiation protocols, even when only two agents are involved may lead to problematic situations. For example consider the following. We have a Boolean game, $G = \langle\{a_1, a_2\}, \Phi, \Phi_1, \Phi_2, \gamma\rangle$ in which the following protocol, which call the *agreed-1-flip* protocol is used: starting from initial assignments $\underline{a}$ (for Φ_1) and $\underline{b}$ (for Φ_2) a_1 proposes a variable of Φ_2 to flip at the same time as a_2 proposes a variable of Φ_1 for a_1 to flip. If γ (the common goal for both agents) is not satisfied the process continues. Although this mechnaism is very basic, it turns out – even for γ being simply a conjunction of literals, that starting from an ill-chosen initial assignment can lead to exponentially many negotiation rounds taking place. That is,

PROPOSITION 9. *There are 2-player Boolean games, $G = \langle\{a_1, a_2\}, \Phi, \Phi_1, \Phi_2, \gamma\rangle$ for which, all of the following hold:*

1. *There is* exactly *one outcome, $v = (v_1, v_2)$ belonging to $NE(G)$ and such that $u_1(v) = u_2(v) = 1$.*
2. *There are initial valuations v'_1 for Φ_1 and v'_2 for Φ_2 with which the agreed-1-flip protocol will require $\Omega(2^{|\Phi|/2})$ rounds in order to reach this unique equilibrium state.*

PROOF. Immediate from Dunne [5]: initial and final valuations correspond to points on the $|\Phi|$-dimensional hypercube, with the effect of a single negotiation round being to move from the current point to one at (Hamming) distance two from it. The argument in [5] constructs examples where the minimum numer of hyperedges to be traversed in moving from initial to final valuation according to this protocol is bounded below by $(77/128)2^{|\Phi|/2}$. □

6. CONCLUSIONS

Boolean games lie at the intersection of logic, game theory, and computer science, and since they were introduced in 2001, they have attracted steadily increasing attention within the multi-agent systems community in particular. However, a standard criticism of Boolean games is that they are computationally complex, which on reflection is not surprising, given that they are, ultimately, games played on propositional logic formulae. If Boolean game are to find wider application, it will surely therefore be necessary to consider possible routes to tractability in Boolean games. In this paper, we have explored two such routes. First, we considered the idea of relaxing the conditions of pure strategy Nash equilibrium so that we only require that no agent can benefit by flipping no more than k variables. We saw that this very natural idea was, for a certain class of Boolean games (where player's goals are represented as k-clause CNF formulae) in fact sufficient to capture all Nash equilibria; we also saw that k-bounded equilibria lead to a smaller search space than pure Nash equilibria in general. We also saw fixed parameter tractability results for k-bounded equilibria. Next, we considered possible classes of Boolean games for which the corresponding game-theoretic questions were computationally tractable. We identified a condition on goal formulae that leads to tractability: if the *range* of all goal formulae lies within a tractable class of Boolean formula, the corresponding decision problems are much simpler.

In terms of related work, a great deal of work in the algorithmic game theory community has addressed the issues of: (i) the complexity of computing equilibria; (ii) approximate solution concepts such as ϵ-Nash equilibria; and (iii) consideration of classes of games for which computation of solution concepts is tractable; see, for example, the references in [3, 12, 4, 15, 8]. However, this body of work differs from the work presented here in that Boolean games have a very distinctive logical form. Moreover, most work on Boolean games (including the present paper) has considered only *pure* Nash equilibria, while mixed equilibria have received most attention in the algorithmic game theory community. This is of course not surprising, given that the existence of mixed Nash equilibria is guaranteed in finite games, while pure Nash equilibria are not guaranteed to exist in many classes of games.

For future work, several issues suggest themselves. First, it would be interesting to consider mixed Nash equilibria in the context of Boolean games; indeed, it is perhaps surprising that this issue has not been considered previously. In particular, it will be interesting to see how the now-famous PPAD results of [4] manifest themselves in Boolean games. Second, since there is clearly a very close relationship between pure Nash equilibria in Boolean games and the SAT problem for propositional logic, it would be interesting to explore this further, and investigate the extent to which SAT solvers (and QBF/QSAT solvers) can be used to find or verify equilibrium outcomes. Finally, as our results with k-bounded equilibria have demonstrated, it is sometimes possible to have approximate solution concepts that correspond to exact solution concepts on Boolean games in which goal formulae are in certain logical normal forms. It would be interesting to consider this issue further, to see whether other approximate solution concepts "correspond" to their exact counterparts on certain classes of games.

7. REFERENCES

[1] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation*. Springer-Verlag: Berlin, Germany, 1999.

[2] E. Bonzon, M.-C. Lagasquie, J. Lang, and B. Zanuttini. Boolean games revisited. In *Proceedings of the Seventeenth European Conference on Artificial Intelligence (ECAI-2006)*, Riva del Garda, Italy, 2006.

[3] V. Conitzer and T. Sandholm. Complexity results about Nash equilibria. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, pages 765–771, Acapulco, Mexico, 2003.

[4] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC-2006)*, Seattle, WA, 2006.

[5] P. E. Dunne. Extremal behaviour in multiagent contract negotiation. *Journal of AI Research*, 23:41–78, 2004.

[6] P. E. Dunne, S. Kraus, W. van der Hoek, and M. Wooldridge. Cooperative boolean games. In *Proceedings of the Seventh International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2008)*, Estoril, Portugal, 2008.

[7] U. Endriss, S. Kraus, J. Lang, and M. Wooldridge. Designing incentives for boolean games. In *Proceedings of the Tenth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2011)*, Taipei, Taiwan, 2011.

[8] G. Gottlob, G. Greco, and F. Scarcello. Pure nash equilibria: Hard and easy games. *Journal of AI Research*, 24:357–406, 2005.

[9] P. Harrenstein, W. van der Hoek, J.-J.Ch. Meyer, and C. Witteveen. Boolean games. In J. van Benthem, editor, *Proceeding of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK VIII)*, pages 287–298, Siena, Italy, 2001.

[10] S. Ieong and Y. Shoham. Marginal contribution nets: A compact representation scheme for coalitional games. In *Proceedings of the Sixth ACM Conference on Electronic Commerce (EC'05)*, Vancouver, Canada, 2005.

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

[12] R. Savani and B. von Stengel. Exponentially many steps for finding a Nash equilibrium in a bimatrix game. In *Proceedings of the 45th IEEE Symposium on Foundations of Computer Science (FOCS-2004)*, Rome, Italy, 2004.

[13] Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press: Cambridge, England, 2008.

[14] S. Szeider. The parameterized complexity of k-flip local search for SAT and MAX SAT. *Discrete Optimization*, 8:139–145, 2011.

[15] B. von Stengel. Equilibrium computation for two-player games in strategic and extensive forms. In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 53–78. Cambridge University Press: Cambridge, England, 2007.