The Computational Complexity of Trembling Hand Perfection and Other Equilibrium Refinements

Kristoffer Arnsfelt Hansen[1], Peter Bro Miltersen[1,⋆], and Troels Bjerre Sørensen[2,⋆⋆]

[1] Department of Computer Science
Aarhus University, Åbogade 34
8200 Århus N, Denmark
{arnsfelt,bromille}@cs.au.dk

[2] Department of Computer Science
University of Warwick, Gibbet Hill Rd
Coventry CV4 7AL, United Kingdom
trold@dcs.warwick.ac.uk

Abstract. The king of refinements of Nash equilibrium is trembling hand perfection. We show that it is **NP**-hard and Sqrt-Sum-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form with integer payoffs is trembling hand perfect. Analogous results are shown for a number of other solution concepts, including proper equilibrium, (the strategy part of) sequential equilibrium, quasi-perfect equilibrium and CURB.

The proofs all use a reduction from the problem of comparing the minmax value of a three-player game in strategic form to a given rational number. This problem was previously shown to be **NP**-hard by Borgs *et al.*, while a Sqrt-Sum hardness result is given in this paper. The latter proof yields bounds on the algebraic degree of the minmax value of a three-player game that may be of independent interest.

1 Introduction

Celebrated recent results [10,6,13] concern the computational hardness of *finding* a Nash equilibrium of a given finite game in strategic form, i.e., a game given by a finite payoff matrix for each of the players. In contrast, the problem of *deciding* whether a *given* strategy profile of a game in strategic form is a Nash equilibrium is trivial to solve efficiently. This latter fact can be regarded as an important feature of Nash equilibrium as a scientific concept: It is feasible to verify or falsify that a particular pure strategy profile we observe "in nature" is

⋆ Research supported by Center for Algorithmic Game Theory, funded by the Carlsberg Foundation.

⋆⋆ Research supported by EPSRC award EP/G064679/1 and by the Centre for Discrete Mathematics and its Applications (DIMAP).

S. Kontogiannis, E. Koutsoupias, P.G. Spirakis (Eds.): SAGT 2010, LNCS 6386, pp. 198–209, 2010.

in equilibrium. The main message of the present paper is that this feature is not shared by standard *refinements* of Nash equilibrium.

Arguably [9], the most important refinement of Nash equilibrium for games in strategic form is Selten's [25] notion of *trembling hand perfection.* The set of trembling hand perfect equilibria is in general a subset of the Nash equilibria of a game and many "unreasonable" Nash equilibria are not trembling hand perfect, thus justifying the notion. However, we prove that this added degree of rationality of the solution concept comes at a cost. We prove: *It is* **NP**-*hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form is trembling hand perfect.* In particular, unless **P**=**NP**, there is no polynomial time algorithm for deciding if a given equilibrium of a given three-player game in strategic form is trembling hand perfect. In contrast to the above hardness result, one may efficiently determine if a given equilibrium of a *two*-player game is trembling hand perfect. Indeed, for the two-player case, an equilibrium is trembling hand perfect if and only if it is undominated [9] and this can be checked by linear programming in polynomial time.

The hardness result is extended to a number of other refinements, including properness [24], sequential equilibrium[1] [21] and quasi-perfect equilibrium [8] of extensive form games, and the discrete solution concept CURB (Closed Under Rational Behavior) [1], where the proof yields **coNP**-hardness. In all cases, the hardness result is shown for games with three players. As is the case with trembling hand perfection, CURB sets of two-player game can be verified and found in polynomial time, using linear programming techniques [3].In contrast, we do not know if the two-player case is easy for properness and quasi-perfection, and leave this as an open problem.

After establishing the **NP**-hardness result, we next ask if the problem of deciding whether an equilibrium is trembling hand perfect (or satisfies any of the other refinements notions we consider) is even *in* **NP**. An **NP**-membership result would be somewhat beneficial for the status of an equilibrium concept as a useful scientific concept, as it would mean that we can at least, with some ingenuity, verify that a situation is in equilibrium, even if we can not in general falsify this efficiently. For deciding trembling hand perfection, it seems that an obvious nondeterministic algorithm would be to guess and verify a *lexicographic belief structure* and appeal to the characterizations of Blume *et al* [4] and Govindan and Klumpp [15] of trembling hand perfection in terms of these. However, it is not clear if the real numbers involved in such a belief structure can be represented as polynomial length strings over a finite alphabet in a way that yields to efficient verification. To argue that it is fact *not* possible to do so using current knowledge, we apply the notion of Sqrt-Sum hardness introduced by Etessami and Yannakakis [12]. In particular, we show that deciding trembling hand perfection (and all the other refinements considered) is Sqrt-Sum hard and therefore not in **NP** unless Sqrt-Sum is in **NP**. Hence, devising a compact representation of belief structures witnessing trembling hand perfection would solve a long standing open problem of numerical analysis.

[1] To be precise, the "strategy part" of a sequential equilibrium.

The hardness proofs all use a reduction from the problem of comparing the minmax value of a game in a strategic form to a given rational number. This problem was previously shown to be **NP**-hard by Borgs *et al.*, [5], while the SQRT-SUM hardness result is given in this paper. The latter proof yields bounds on the algebraic degree of the minmax value of a three-player game that may be of independent interest.

1.1 Related Work

As mentioned, there is a lot of work on hardness of finding equilibria when the game is given as input, or checking whether equilibria with certain properties exists when the game is given as input. In contrast, we are not aware of much previous work on the complexity of determining whether a *given* equilibrium satisfies a refined stability notion. An exception is Etessami and Lochbihler [11] who show that it is **NP**-hard to determine if a given strategy in a symmetric game in strategic form is an evolutionarily stable strategy.

2 NP-hardness of Trembling Hand Perfect and Proper Equilibrium

We recall the definitions of Selten [25]. For a motivation and discussion of the solution concept, we refer to the excellent monograph of van Damme [9].

Definition 1 (ϵ-perfect equilibrium). *A strategy profile σ is an ϵ-perfect equilibrium iff it assigns strictly positive probability to all pure strategies, and only pure strategies that are best replies get probability more than ϵ.*

Definition 2 (Trembling hand perfect equilibrium). *A strategy profile σ is a trembling hand perfect equilibrium iff is the limit point of a sequence of ϵ-perfect equilibria with $\epsilon \to 0+$.*

Theorem 1. *It is* **NP***-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form is trembling hand perfect.*

Proof. Our proof is a reduction from the problem of approximately computing minmax values of 3-player games with 0-1 payoffs. The minmax value of a 3-player game is the smallest number v so that player 2 and player 3 can guarantee, using *uncorrelated* mixed strategies, that player 1 does not get an expected payoff larger than v. The problem of approximately computing this value was recently shown to be **NP**-hard by Borgs *et al* [5]. In particular, it follows from Borgs *et al.* that the following *promise problem* MINMAX is **NP**-hard:
MINMAX:

1. YES-instances: Pairs (G, r) for which the minmax value for Player 1 in the 3-player game G is strictly smaller than the rational number r.
2. NO-instances: Pairs (G, r) for which the minmax value for Player 1 in G is strictly greater than r.

In fact, by multiplying the payoffs of the game with the denominator of r, we can without loss of generality assume that r is an integer. We now reduce MINMAX to deciding trembling hand perfection.

Let G be a three-player game in strategic form and let r be an integer. We define G' be the game where the strategy space of each player is as in G, except that it is extended by a single pure strategy, $\perp$. The payoffs of G' are defined as follow. The payoff to Players 2 and 3 are 0 for all strategy combinations. The payoff to Player 1 is r for all strategy combinations where at least one player plays $\perp$. For those strategy combinations where no player plays $\perp$, the payoff to player 1 is the same as it would have been in the game G. Obviously, $\mu = (\perp, \perp, \perp)$ is a Nash equilibrium of G'.

We claim that if the minmax value for Player 1 in G is strictly smaller than r, then μ is a trembling hand perfect equilibrium of G'. Indeed, let (τ_2, τ_3) be a minmax strategy profile of Players 2 and 3 in G. Let τ be any profile of G' where Players 2 and 3 play (τ_2, τ_3). Also, let u be the strategy profile of G' where each player mixes all pure strategies uniformly. Now define

$$\sigma_k = (1 - \frac{1}{k} - \frac{1}{k^2})\mu + \frac{1}{k}\tau + \frac{1}{k^2}u$$

We have that σ_k is a fully mixed strategy profile of G' converging to μ as $k \to \infty$. Also, for sufficiently large k, the strategies of μ are best replies to σ_k. This follows from the fact that Players 2 and 3 are indifferent about the outcome and the fact that Player 1 gets payoff r by playing $\perp$ while he gets a payoff strictly smaller than r for large values of k by playing any other strategy. We conclude, using Theorem 2.2.5 in van Damme [9], that μ is trembling hand perfect, as desired.

On the other hand, we claim that if the minmax value for Player 1 in G is strictly greater than r, then μ is a *not* a trembling hand perfect equilibrium of G'. Indeed, let $(\sigma_{k,1}, \sigma_{k,2}, \sigma_{k,3})_k$ be any sequence of fully mixed strategy profiles converging to $(\perp, \perp, \perp)$. Since $\sigma_{k,2}$ and $\sigma_{k,3}$ do not put all their probability mass on $\perp$, Player 1 has a reply to $(\sigma_{k,2}, \sigma_{k,3})$ with an expected payoff strictly greater than r. Therefore, $\perp$ is not a best reply of Player 1 to $(\sigma_{k,2}, \sigma_{k,3})$ and we conclude that $(\perp, \perp, \perp)$ is not trembling hand perfect.

That is, we have reduced the promise problem MINMAX to deciding trembling hand perfection and are done. □

We now refine the proof so that it applies to *proper equilibrium.* Proper equilibrium was introduced by Myerson [24] as a further refinement of trembling hand perfect equilibrium. For a motivation and discussion of the solution concept, we refer to the excellent monograph of van Damme [9] or the survey of Hillas and Kohlberg [18].

Definition 3 (ϵ-proper equilibrium). *A strategy profile σ is an ϵ-proper equilibrium iff it assigns strictly positive probability to all pure strategies, and the following condition holds: Given two pure strategies, p_i and p_j, of the same player. If p_i is a worse reply against σ than p_j, then σ must assign a probability to p_i that is at most ϵ times the probability it assign to p_j.*

Definition 4 (Proper equilibrium). *A strategy profile σ is a proper equilibrium iff is the limit point of a sequence of ϵ-proper equilibria with $\epsilon \to 0+$.*

Theorem 2. *It is* **NP***-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form is proper.*

Proof. We only need to make minor changes to the proof of **NP**-hardness of trembling hand perfection to get the same result for proper equilibria. Construct the game in the same way, with a new strategy $\perp$ for each player. Define the strategy $\tau_{1,k}$ for Player 1 to be a permutation of $(1 - \sum_i k^{-i}, k^{-1}, k^{-2}, \ldots, k^{-n+1})$, such that worse replies against (τ_2, τ_3) get more negative powers of k. In case two pure strategies are equal against (τ_2, τ_3), compare against the uniform mix u of Players 2 and 3, again with the worse reply getting the more negative powers of k. This can be achieved by sorting the strategies of Player 1 lexicographically on payoff against (τ_2, τ_3) and the uniform strategy u, and then assigning powers in decreasing order to the lower indices. Define $\tau_k = (\tau_{1,k}, \tau_2, \tau_3)$, $\nu_k = (\tau_{1,k}, u, u)$, and $\mu = (\perp, \perp, \perp)$. Now define

$$\sigma_k = (1 - \frac{1}{k} - \frac{1}{k^2})\mu + \frac{1}{k}\tau_k + \frac{1}{k^2}\nu_k$$

σ_k is fully mixed of all finite k. Furthermore, if the minmax value for Player 1 in G is less than r, then for any sufficiently large k, better replies of Player 1 gets k' times higher probability than worse replies, thus satisfying the condition for being a $\frac{1}{k'}$-proper equilibrium, with $k' = k/(1 - \sum_i k^{-i})$. Since σ_k tends towards μ as $k \to \infty$, we therefore have that μ is a proper equilibrium. If the minmax value for Player 1 in G is greater than r, μ is not even trembling hand perfect, and therefore not proper either. μ is therefore proper if and only if the minmax value for Player 1 in G is less than r. □

3 NP-hardness of Refinements of Nash Equilibria for Extensive form Games

An extensive form game is given by a finite tree with payoffs for each player at the leaves, information sets partitioning nodes of the tree and with some of the nodes having predefined moves of chance. An information set is a collection of nodes of the same player, where the player cannot distinguish between them. This can be used to model information hidden from the player, both as actively hidden information in a game over time, and as a way of modelling simultaneous moves. A player is said to have *perfect recall* if for each of the player's information sets, all nodes in the set share the same sequence of actions and information sets of the player on the path from the root to the nodes. A game is said to be of perfect recall, if all players have perfect recall. This is a standard assumption to make, and one that the game produced by our reduction will satisfy.

Actions of a player are denoted by labels on edges of the tree. A *behavior strategy* assigns probabilities to actions such that it forms probability distributions over the actions for each of the information sets. A Nash equilibrium in behavior strategies is a profile of behavior strategies so that no player wants to deviate, given that other players play according to the profile. As is the case of games in strategic form, it is straightforward to verify in polynomial time that a given profile is a Nash equilibrium. For details, see e.g., Koller, Megiddo and von Stengel [20] or any textbook on game theory.

The most important refinement of Nash equilibrium for game in sequential form is the notion of *sequential equilibrium* due to Kreps and Wilson [21] is based on the notion of *beliefs*. Formally, a belief of a player is a probability distribution on each of his information sets. Intuitively, the belief should indicate the subjective probability of the player of being in each of the nodes in the information set, given that he has arrived at this information set. An *assessment* (ρ, μ) is a strategy profile ρ, and a *belief profile* μ: a belief for each of the players. A sequential equilibrium is an assessment which is (1) *consistent* and (2) *a sequential best reply against itself*, the former notion capturing that the beliefs are sensible given the strategies, and the latter notion capturing that the strategies are sensible given the beliefs. We define these two notions formally next.

We first define consistency for *fully mixed* strategy profiles, i.e., ones where every action in every information set has a strictly positive probability of being taken. For such a strategy profile, the *induced belief profile* is the unique one consistent with the strategy profile: The strategies being played out against each other induces a probability distribution on possible plays; the induced belief assigns to information set u the conditional probability distribution on u derived from this probability distribution. This is well-defined as at most one node in u may be reached during each particular play (due to the perfect recall property) and u has a non-zero probability of being reached (as the strategies are fully mixed). The contribution of Kreps and Wilson is a generalization of this consistency notion to strategy profiles where some of the information sets may be reached with probability 0: For this general case, we say that an assessment is consistent if it is the limit point of a sequence of consistent assessments with fully mixed strategy profiles.

We next define what it means to be a sequential best reply against itself. For each player, a strategy profile will assign an expected value to each node of the tree, which is the expectation over the leaves given than play starts at that node and follows the given probabilities of play. Given a belief as well, we can assign an expected value to each action, being the expected value of the node reached by taking the action for each node, weighted by the probability given by the belief. An assessment is said to be a sequential best reply against itself, if all players only assign positive probability to actions with maximal expected payoff, given the strategy profile and belief.

Note that a sequential equilibrium is an assessment, i.e., a behavior strategy profile and a belief profile. Our NP-hardness result applies when the input is the "strategy-part" only.

Theorem 3. *Given a pure strategy profile of an extensive form three-player game, it is* **NP***-hard to decide if it is part of an assessment that is a sequential equilibrium.*

Proof. The reduction is again similar to that of trembling hand perfection. Given a game G in strategic form, construct an extensive form game G' where the players choose an action of G in turn, but without revealing the choice to the other players. Player 2 chooses first, then Player 3, and finally Player 1. Each Player now has a single information set, and the game is strategically equivalent to G. Now give each player a new action $\perp$. If this action is chosen, the game ends immediately without having the remaining players choose and action. If $\perp$ is chosen by either player, the payoff to Player 1 is r, otherwise it is simply the payoff from G.

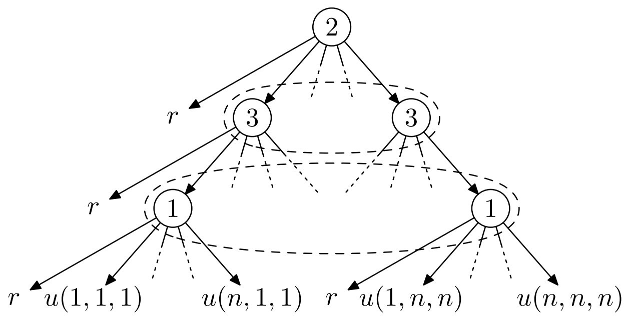

Fig. 1. The extensive form game G'

We now argue that $\mu = (\perp, \perp, \perp)$ is part of a sequential equilibrium iff the minmax value of Player 1 in G is less than r.

Define τ to be some strategy profile where Players 2 and 3 play minmax against Player 1, and let u be the strategy profile with all players playing the uniform distribution. As in the previous proofs, let

$$\sigma_k = (1 - \frac{1}{k} - \frac{1}{k^2})\mu + \frac{1}{k}\tau + \frac{1}{k^2}u$$

σ_k is fully mixed of all finite k. Furthermore, if the minmax value for Player 1 in G is less than r, then for any sufficiently large k, $\perp$ will be the unique best reply of Player 1 against σ_k. This also means that the expected value for Player 1 of choosing $\perp$ given the induced belief of σ_k will be strictly higher than for all other actions, and this will also hold for the limit.

On the other hand, if the minmax value is greater than r, no strategy of Players 2 and 3 will make $\perp$ be the best reply of Player 1. Therefore, no belief (consistent with a strategy of Players 2 and 3) will give a maximal expected payoff to Player 1 playing $\perp$.

μ is therefore part of a sequential equilibrium if and only if the minmax value for Player 1 in G is less than r. □

Theorem 3 begs the following question: Can one check in polynomial time if an entire assessment (a strategy profile and a belief profile) given as input is a sequential equilibrium? Kohlberg and Reny [19] present a finite-step algorithm performing this task, but as they state it, their algorithm is exponential. It is not clear to us if this problem is in **P** or if it is **NP**-hard and we consider this an interesting open problem. It is interesting to note that this is in some contrast with the situation for strategic form games: Perfect and proper equilibrium of strategic form games can also be "backed up" by belief structures [4,15] and if a rational-valued belief structure is given as part of the input, it is straightforward to verify the equilibrium condition (however, as we argue in Section 5, a given perfect strategy profile may require belief structures with no polynomial-size representation - in particular, using algebraic numbers of very high degree may be necessary).

A refinement of (the strategy part of) sequential equilibrium is *quasi-perfect* equilibrium [8]. Despite the fact that quasi-perfect equilibrium is a lesser known refinement that sequential equilibrium, it has been argued strongly by Mertens [22] (see also [18]) that quasi-perfect equilibrium is the "right" equilibrium notion of extensive form games. We omit the technically involved definition of quasi-perfection, but note that it is straightforward to check that the reduction in the proof of Theorem 3 maps "yes"-instance to equilibria that are not only sequential but also quasi-perfect. Since quasi-perfect equilibrium refines sequential equilibrium, we also have that the reduction maps "no"-instances to equilibria that are not quasi-perfect. Therefore we have the following corollary.

Corollary 1. *It is* **NP***-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in extensive form is quasi-perfect.*

4 coNP-hardness of CURB Sets

A set valued solution concept is Strategy Sets Closed Under Rational Behavior (CURB) [1].

Definition 5 (CURB set). *In an m-player game, a family of sets of pure strategies, $S_1, S_2, \ldots, S_m$ with S_i being a subset of the strategy set of player i, is* closed under rational behavior *(CURB) iff for all pure strategies x of Player i so that x is a best reply to some product distribution on $S_1 \times S_2 \times \cdots \times S_{i-1} \times S_{i+1} \times \cdots \times S_m$, we have that $x \in S_i$.*

CURB sets are guaranteed always to exist, as the set of all pure strategies is trivially CURB, as there are no pure strategies outside the set. The CURB condition is usually paired with a minimality condition, so as not to get unnecessarily large solutions. This minimality condition would be the obvious place to look for **coNP**-hardness, but we show here that simply checking the CURB condition is **coNP**-hard. This also implies that it is not obvious that minimality should even be contained in **coNP**.

Theorem 4. *It is* **coNP***-hard to check whether a set of n pure strategies of each player is CURB in an $(n+1) \times n \times n$ strategic form game with integer payoffs.*

Proof. We again reduce from MINMAX, so let G be a three-player game in strategic form and let r be an integer. We define G' be the game where the strategy space of each player is as in G, except that Player 1 gets an additional strategy, $\perp$. The payoffs of G' are defined as follow. The payoff to Players 2 and 3 are 0 for all strategy combinations. The payoff to Player 1 is r, if he plays $\perp$, and otherwise the payoff to player 1 is the same as it would have been in the game G.

Now, the minmax value of G is less than r iff the set of all pure strategies except $\perp$ is CURB in G'. Indeed, if the minmax value of G is less than r, then Player 1's best reply to the optimal treat of Players 2 and 3 is $\perp$ in G'. The set of all pure strategies except $\perp$ is therefore *not* CURB. If the minmax value of G is greater than r, then $\perp$ is never a best reply in G, and the set of all other strategies is CURB. $\square$

5 Sqrt-Sum-hardness

Sqrt-Sum is the following decision problem [16,14]: Given positive integers $a_1, a_2, \ldots, a_n, k$, decide whether $\sum_{i=1}^{n} \sqrt{a_i} < k$.

Though it is not unlikely that this problem is in **P**, we do not even know if it is in **NP** at the moment. A decision problem is called Sqrt-Sum-hard if Sqrt-Sum reduces to it by a polynomial time many-one reduction. Etessami and Yannakakis [12] pioneered the use of Sqrt-Sum-hardness to argue that certain problems are hard "given current state of the art". It is important to notice that unlike **NP**-hardness, Sqrt-Sum-hardness should *not* be used as an indication that a problem is *actually* hard, only as an indication that we do not know if it is easy. In this section we show Sqrt-Sum-hardness of the minmax value of a 3-player game and thus by the previously described reductions give evidence that it is not possible to decide the refined solution concepts in **NP** "given current state of the art".

Lemma 1. *For every pair of probability distributions x and y on $\{1, \ldots, n\}$ there exists another probability distribution z such that $x_i y_i \leq z_i^2$ for all i.*

Proof. If $x_i y_i = 0$ for all i we may pick z arbitrarily. Otherwise, define $w_i = \sqrt{x_i y_i}$ for all i. By the Cauchy-Schwarz inequality we have

$$\sum_{i=1}^{n} w_i = \sum_{i=1}^{n} \sqrt{x_i}\sqrt{y_i} \leq \sqrt{\sum_{i=1}^{n} x_i}\sqrt{\sum_{i=1}^{n} y_i} = 1 \ .$$

We may thus obtain the required z by letting $z_j = w_j / (\sum_{i=1}^{n} w_i)$. $\square$

Given positive numbers $a_1, \ldots, a_n$ define the payoff to player 1 in an $n \times n \times n$ game $G(a_1, \ldots, a_n)$ by letting $u_1(i, j, k) = -1/a_i$ if $i = j = k$ and $u_1(i, j, k) = 0$ otherwise.

Proposition 1. *The minmax value for player 1 in the game $G(a_1, \ldots, a_n)$ is $-1/(\sum_{i=1}^{n} \sqrt{a_i})^2$.*

Proof. If player 2 and player 3 play strategies p and q, player 1 may obtain payoff $\max_i -p_i q_i/a_i$. Let v be the minmax value for player 1. For optimal strategies for player 2 and 3 we may assume by Lemma 1 that $p = q$, and furthermore we must then have that $v = -p_i^2/a_i$ for all i, and thus $p_i = \sqrt{-v}\sqrt{a_i}$ for all i. Summing over i gives

$$1 = \sum_{i=1}^{n} p_i = \sum_{i=1}^{n} \sqrt{-v}\sqrt{a_i} = \sqrt{-v}\sum_{i=1}^{n}\sqrt{a_i} \ .$$

Squaring and rearranging gives

$$v = -\frac{1}{(\sum_{i=1}^{n}\sqrt{a_i})^2} \ ,$$

as stated. □

Theorem 5. *Deciding whether the minmax value for player 1 in a $n \times n \times n$ game is less than a given rational k is* SQRT-SUM *hard.*

Proof. Deciding whether $\sum_{i=1}^{n}\sqrt{a_i} < k$ reduces to decide for the minmax value v for player 1 in the game $G(a_1, \ldots, a_n)$ whether $v < -\frac{1}{k^2}$ by Proposition 1. □

Corollary 2. *It is SQRT-SUM hard to determine whether a given pure equilibrium in a 3-player game in strategic form with integer payoffs is trembling-hand perfect or proper and whether a given pure equilibrium in a 3-player game in extensive form with integer payoffs is quasi-perfect or the strategy part of a sequential equilibrium. It is also* SQRT-SUM *hard to test whether a given set of pure strategies is not CURB. In particular, neither of these problems are in* **NP** *unless* SQRT-SUM *is in* **NP**.

Finally, we show that our reduction can also be used to give lower bounds on the algebraic degree of the minmax value of a 3-player game. Such a result is interesting for computational reasons: They indicate that if we want to compute the *exact* minmax value of a 3-player game and want to represent the exact irrational but algebraic answer in, say, a standard representation such as Thom encoding [7], exponential space is needed even to represent the output.

For providing the lower bound of the algebraic degree of the minmax value we use basic results from the theory of field extensions.

Proposition 2. *The algebraic degree of the minmax value for player 1 in a $n \times n \times n$ game can be 2^{n-1}.*

Proof. Let $a_1, \ldots, a_n$ be arbitrary relatively prime positive integers, and let v be the minmax value of the game $G(a_1, \ldots, a_n)$. We shall calculate the degree $[\mathbf{Q}(v) : \mathbf{Q}]$ of the field extension $\mathbf{Q}(v)$ of $\mathbf{Q}$. It is well known that for relatively prime positive integers $a_1, \ldots, a_n$ we have $[\mathbf{Q}(\sqrt{a_1}, \ldots, \sqrt{a_n}) : \mathbf{Q}] = 2^n$ (*e.g.* [23, Example 11.5]). Furthermore, we have $\mathbf{Q}(\sqrt{a_1} + \cdots + \sqrt{a_n}) = \mathbf{Q}(\sqrt{a_1}, \ldots, \sqrt{a_n})$. By Proposition 1 we have that $-1/\sqrt{v} = \sum_{i=1}^{n}\sqrt{a_i}$, and thus $[\mathbf{Q}(\sqrt{v}) : \mathbf{Q}] = 2^n$. Finally using $[\mathbf{Q}(\sqrt{v}) : \mathbf{Q}] = [\mathbf{Q}(\sqrt{v}) : \mathbf{Q}(v)][\mathbf{Q}(v) : \mathbf{Q}] \leq 2[\mathbf{Q}(v) : \mathbf{Q}]$ the result follows. □

One can give an almost matching upper bound using the general tool of quantifier elimination for the first order theory of the reals.

Proposition 3. *The algebraic degree of the minmax value for player 1 in a $n \times n \times n$ game is $2^{O(n)}$.*

Proof. We may describe the minmax value by a first order formula $P(v)$ with free variable v, as $P(v) := A(v) \wedge B(v)$, where

$$A(v) := (\exists p, q \in \mathbf{R}^n) \bigwedge_{i=1}^{n} \left(\sum_{j=1}^{n} \sum_{k=1}^{n} u_1(i,j,k) p_j q_k \leq v \right) \wedge C(p,q) \ ,$$

$$B(v) := (\forall p, q \in \mathbf{R}^n) \bigvee_{i=1}^{n} \left(\sum_{j=1}^{n} \sum_{k=1}^{n} u_1(i,j,k) p_j q_k \geq v \right) \wedge C(p,q) \ ,$$

and $C(p,q) := (\bigwedge_{i=1}^{n} p_i \geq 0) \wedge (\sum_{i=1}^{n} p_i = 1) \wedge (\bigwedge_{i=1}^{n} q_i \geq 0) \wedge (\sum_{i=1}^{n} q_i = 1)$.

We note that the degree of each polynomial in the formula is at most 2. Thus applying the quantifier elimination procedure of Basu, Pollack and Roy [2] to each of the formulas $A(v)$ and $B(v)$ yields equivalent quantifier free formulas $A'(v)$ and $B'(v)$ wherein each polynomial is of degree $2^{O(n)}$. It follows that $A'(v) \wedge B'(v)$ is quantifier free formula equivalent to $P(v)$ wherein each polynomial are univariate polynomials in v of degree $2^{O(n)}$. Now, since the actual minmax value v is an isolated solution to this formula, it must satisfy one of the polynomial equations involving a nonconstant polynomial with equality. We can thus conclude it must be a root of a polynomial of degree $2^{O(n)}$. □

Remark 1. The above bound is especially relevant for the special case of $k \times n \times n$ games, where k is considered a constant [17]. For this case one may find the minmax value by considering all $k \times k \times k$ subgames and the minmax value of those. This also means that for fixed k one can in polynomial time compute the Thom encoding of the minmax value of a given $k \times n \times n$ game, employing general algorithms for the first-order theory of the reals [2].

References

1. Basu, K., Weibull, J.W.: Strategy subsets closed under rational behavior. Economics Letters 2(36), 141–146 (1991)
2. Basu, S., Pollack, R., Roy, M.F.: On the combinatorial and algebraic complexity of quantifier elimination. Journal of the ACM 43(6), 1002–1045 (1996)
3. Benisch, M., Davis, G.B., Sandholm, T.: Algorithms for closed under rational behavior (CURB) sets. Journal of Artificial Intelligence Research (Forthcoming 2010)
4. Blume, L., Brandenburger, A., Dekel, E.: Lexicographic probabilities and equilibrium refinements. Econometrica 59, 81–98 (1991)
5. Borgs, C., Chayes, J.T., Immorlica, N., Kalai, A.T., Mirrokni, V.S., Papadimitriou, C.H.: The myth of the folk theorem. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing, pp. 365–372. ACM, New York (2008)

6. Chen, X., Deng, X.: Settling the complexity of two-player Nash equilibrium. In: Proceedings of 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), pp. 261–272 (2006)
7. Coste, M., Roy, M.: Thom's lemma, the coding of real algebraic numbers and the computation of the topology of semi-algebraic sets. Journal of Symbolic Computation 5(1-2), 121–129 (1988)
8. van Damme, E.: A relation between perfect equilibria in extensive form games and proper equilibria in normal form games. International Journal of Game Theory 13, 1–13 (1984)
9. van Damme, E.: Stability and Perfection of Nash Equlibria, 2nd edn. Springer, Heidelberg (1991)
10. Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a Nash equilibrium. In: Procedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC 2006), pp. 71–78 (2006)
11. Etessami, K., Lochbihler, A.: The computational complexity of evolutionarily stable strategies. International Journal of Game Theory 31(1), 93–113 (2008)
12. Etessami, K., Yannakakis, M.: Recursive markov decision processes and recursive stochastic games. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 891–903. Springer, Heidelberg (2005)
13. Etessami, K., Yannakakis, M.: On the complexity of Nash equilibria and other fixed points (extended abstract). In: Proc. 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), pp. 113–123 (2007)
14. Garey, M.R., Graham, R.L., Johnson, D.S.: Some NP-complete geometric problems. In: Proceedings of the 8th Annual ACM Symposium on Theory of Computing, STOC 1976, Hershey, PA, May 3-5, pp. 10–22. ACM Press, New York (1976)
15. Govindan, S., Klumpp, T.: Perfect equilibrium and lexicographic beliefs. International Journal of Game Theory 31(2), 229–243 (2003)
16. Graham, R.L.: 10 - Problems and Solutions. In: P73: Euclidian Minimum Spanning Trees. Bulletin of the EATCS, vol. 24, pp. 205–206. EATCS (October 1984)
17. Hansen, K.A., Hansen, T.D., Miltersen, P.B., Sørensen, T.B.: Approximability and parameterized complexity of minmax values. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 684–695. Springer, Heidelberg (2008)
18. Hillas, J., Kohlberg, E.: Foundations of strategic equilibria. In: Aumann, R.J., Hart, S. (eds.) Handbook of Game Theory, ch. 42, vol. 3, pp. 1597–1663. Elsevier Science, Amsterdam (2002)
19. Kohlberg, E., Reny, P.J.: Independence on relative probability spaces and consistent assessments in game trees. Journal of Economic Theory 75(2), 280–313 (1997)
20. Koller, D., Megiddo, N., von Stengel, B.: Efficient computation of equilibria for extensive form games. Games and Economic Behavior 14, 247–259 (1996)
21. Kreps, D.M., Wilson, R.: Sequential equilibria. Econometrica 50(4), 863–894 (1982)
22. Mertens, J.F.: Two examples of strategic equilibrium. Games and Economic Behavior 8, 378–388 (1995)
23. Morandi, P.: Field and Galois Theory. Graduate Texts in Mathematics, vol. 167. Springer, Heidelberg (1996)
24. Myerson, R.B.: Refinements of the Nash equilibrium concept. International Journal of Game Theory 15, 133–154 (1978)
25. Selten, R.: A reexamination of the perfectness concept for equilibrium points in extensive games. International Journal of Game Theory 4, 25–55 (1975)