



EXISTENCE OF NASH EQUILIBRIA ON INTEGER PROGRAMMING GAMES

Margarida Carvalho
Andrea Lodi
João Pedro Predoso

March 2017

CERC_DATA-SCIENCE_2017_003

POLYTECHNIQUE MONTRÉAL

DÉPARTEMENT DE MATHÉMATIQUES ET GÉNIE INDUSTRIEL

Pavillon André-Aisenstadt

Succursale Centre-Ville

C.P. 6079

Montréal - Québec

H3C 3A7

Canada

Téléphone: 514-340-5121 # 3314

Existence of Nash equilibria on integer programming games

MARGARIDA CARVALHO *

ANDREA LODI †

JOÃO PEDRO PEDROSO ‡

August 3, 2017

Abstract

We aim to investigate a new class of games, where each player’s set of strategies is a union of polyhedra. These are called integer programming games. To motivate our work, we describe some practical examples suitable to be modeled under this paradigm. We analyze the problem of determining whether or not a Nash equilibria exists for an integer programming game, and demonstrate that it is complete for the second level of the polynomial hierarchy.

Keywords— Integer Programming games; Nash equilibria; Computational Complexity.

1 Introduction

Game theory is a generalization of decision theory that takes into account the interaction of multiple decision makers which are concerned about finding their “best” strategies subject to the fact that each controls some, but not all, actions that can take place. See Fudenberg and Tirole [6] for an overview in this field.

Our goal is to study a particular class of games called *integer programming games* (IPG), namely, the existence of “solutions” which in this context are called *Nash equilibria* (NE). We highlight three contributions concerning IPGs: the computational complexity study of the problem of deciding the existence of a pure NE and of a NE, and the determination of sufficient conditions to guarantee the existence of NE.

Our paper is structured as follows. Section 1.1 fixes notation and covers the game theory background. Section 1.2 presents examples of integer programming games highlighting the importance of studying the existence of NE. Section 2 provides literature review. Section 3 classifies the computational complexity of the problems related with the existence of NE to IPGs and states sufficient conditions for NE to exist. Finally, we conclude and discuss further research directions in Section 4.

1.1 Definitions and notation

We consider games with a finite set of decision makers $M = \{1, 2, \dots, m\}$, called *players*. Each player $p \in M$ has the set of feasible strategies X^p . We denote the set of all players’ strategies combinations by $X = \prod_{p \in M} X^p$ and the operator $(\cdot)^{-p}$ to denote (\cdot) for all players except player p . We call each $x^p \in X^p$ and $x \in X$ a player p ’s *pure strategy* and a *pure profile of strategies*, respectively. Let player p ’s payoff function be $\Pi^p(x^p, x^{-p})$. Our investigation focuses on simultaneous *non-cooperative complete information* games, *i.e.*, players play simultaneously, they are self-interested and have full information of each other payoffs and strategies. Each player aims to maximize her payoff, which is influenced by other participants’ decisions.

*margarida.carvalho@polymtl.ca. IVADO Fellow, Canada Excellence Research Chair, École Polytechnique de Montréal, Québec, Canada.

†andrea.lodi@polymtl.ca. Canada Excellence Research Chair, École Polytechnique de Montréal, Québec, Canada.

‡jpp@fc.up.pt. Faculdade de Ciências Universidade do Porto and INESC TEC, Porto, Portugal

In other words, each player p 's goal is to select her *best response* against the opponents' strategies x^{-p} by solving the following mathematical programming problem:

$$\underset{x^p \in X^p}{\text{maximize}} \quad \Pi^p(x^p, x^{-p}). \quad (1)$$

A pure profile of strategies $x \in X$ that solves the optimization problem (1) for all players is called *pure equilibrium*. A game may fail to have pure equilibria and, therefore, a broader solution concept for a game must be introduced. To that end, we recall some basic concepts of measure theory. Let Δ^p denote the space of Borel probability measures over X^p and $\Delta = \prod_{p \in M} \Delta^p$. Each player p 's expected payoff for a profile of strategies $\sigma \in \Delta$ is

$$\Pi^p(\sigma) = \int_{X^p} \Pi^p(x^p, x^{-p}) d\sigma. \quad (2)$$

A *Nash equilibrium* (NE) is a profile of strategies $\sigma \in \Delta$ such that

$$\Pi^p(\sigma) \geq \Pi^p(x^p, \sigma^{-p}), \quad \forall p \in M \quad \forall x^p \in X^p. \quad (3)$$

In a NE each player p 's expected payoff from σ cannot be improved by unilaterally deviating to a different strategy¹.

The *support of a strategy* $\sigma^p \in \Delta^p$, denoted as $\text{supp}(\sigma^p)$, is the set of player p 's strategies played with positive probability, *i.e.*, $\text{supp}(\sigma^p) = \{x^p \in X^p : \sigma^p(x^p) > 0\}$. Given $\sigma \in \Delta$, if each player's support size is 1, then it is a pure profile of strategies, otherwise, we call it (strictly) mixed. For the sake of simplicity, whenever the context makes it clear, we use the term (strategy) profile to refer to a pure one.

A game is called *continuous* if each player p 's strategy set X^p is a nonempty compact metric space and the payoff $\Pi^p(x^p, x^{-p})$ is continuous.

A *separable game* is a continuous game with the payoff functions taking the form

$$\Pi^p(x^p, x^{-p}) = \sum_{j_1=1}^{k_1} \dots \sum_{j_m=1}^{k_m} a_{j_1 \dots j_m}^p f_{j_1}^1(x^1) \dots f_{j_m}^m(x^m), \quad (4)$$

where $a_{j_1 \dots j_m}^p \in \mathbb{R}$ and the f_j^p are real-valued continuous functions.

A game is *finite* if the X^p are finite and the $\Pi^p(x^p, x^{-p})$ are arbitrary. Stein *et al.* [18] state that finite games are special cases of separable games. *Normal-form games* (or strategic-form games) are finite games represented through a multidimensional matrix with an entry for each pure profile of strategies $x \in X$, where that entry is an m dimensional vector of the players' payoffs associated with x .

Based on the definition presented by Köppe *et al.* [9], we define an *integer programming game* (IPG) as a game with $X^p = \{x^p : A^p x^p \leq b^p, \quad x_i^p \in \mathbb{N} \text{ for } i = 1, \dots, B_p\}$, where A^p is a $k_p \times n_p$ matrix (with $n_p \geq B_p$), b^p a column vector of dimension k_p , and the payoff functions $\Pi^p(x^p, x^{-p})$ are continuous and can be evaluated in polynomial time. Note that IPGs contain mathematical programming problems in the special case of a single player.

Observe that any finite game can be modeled as an IPG: associate a binary variable for each player's strategy (which would model the strategy selected), add a constraint summing the decision variables up to one (this ensures that exactly one strategy is selected) and formulate the players' payoffs according to the payoff values for combinations of the binary variables²; moreover, these payoff functions are continuous, since we can endow X with the discrete metric, which makes any function automatically continuous. In an IPG, strategies' sets can be unbounded and thus not compact, which makes the IPG class not contained in that of continuous games. If the strategies' sets in an IPG are nonempty and bounded, then the X^p are finite unions of convex polyhedra which are compact metric spaces and thus, the IPG is a continuous game. Figure 1 displays the connections between these game classes.

¹The equilibrium conditions (3) only reflect a player p deviation to strategy in X^p and not in Δ^p , because a strategy in Δ^p is a convex combination of strategies in X^p , and thus cannot lead to a better payoff than one in X^p .

²In specific, a player's payoff is a summation over all pure profiles of strategies, where each term is the product of the associated binary variables and the associated payoff.




Figure 1: Game classes.

1.2 Examples

Next, we describe two games: the *knapsack game* which is the simplest purely integer programming game that one could devise, the *competitive lot-sizing game* which has practical applicability in production planning.

Knapsack game. A simple and natural IPG would be one with linear payoff functions for the players. Under this setting, each player p aims to solve

$$\max_{x^p \in \{0,1\}^n} \sum_{i=1}^n v_i^p x_i^p + \sum_{k=1, k \neq p}^m \sum_{i=1}^n c_{k,i}^p x_i^p x_i^k \quad (5a)$$

$$\text{s. t.} \quad \sum_{i=1}^n w_i^p x_i^p \leq W^p, \quad (5b)$$

where the parameters v_i^p , $c_{k,i}^p$, w_i^p and W^p are real numbers. This model can describe situations where m entities aim to decide in which of n projects to invest such that each entity budget constraint (5b) is satisfied and the associated payoffs are maximized (5a). The second term of the payoff (5a) reflects the opponents' influence: if player p and player k select project i ($x_i^p = x_i^k = 1$) then, player p earns $c_{k,i}^p > 0$ or loses $c_{k,i}^p < 0$.

In Carvalho [2] mathematical programming tools are used to compute some refined equilibria of this game.

Competitive lot-sizing game. The competitive lot-sizing game is a Cournot competition played through T periods by a set of firms (players) that produce the same good; see [2] for details. Each player has to plan her production as in the lot-sizing problems (see [17]) but, instead of satisfying a known demand in each period of the time horizon, the market price depends on the total quantity of the good that is sold in the market. Each player p has to decide how much will be produced in each time period t (production variable x_t^p) and how much will be placed in the market (variable q_t^p). For producing a positive quantity, player p must pay a fixed and a proportional amount (setup and variable costs, respectively). A producer can build inventory (variable h_t^p) by producing in advance and incurring in an inventory cost. In this way, we obtain the following model for each player (producer) p :

$$\max_{y^p \in \{0,1\}^T, x^p, q^p, h^p} \sum_{t=1}^T P_t(q_t) q_t^p - \sum_{t=1}^T F_t^p y_t^p - \sum_{t=1}^T C_t^p x_t^p - \sum_{t=1}^T H_t^p h_t^p \quad (6a)$$

$$\text{s. t.} \quad x_t^p + h_{t-1}^p = h_t^p + q_t^p \quad \text{for } t = 1, \dots, T \quad (6b)$$

$$0 \leq x_t^p \leq M_t^p y_t^p \quad \text{for } t = 1, \dots, T \quad (6c)$$

where F_t^p is the setup cost, C_t^p is the variable cost, H_t^p is the inventory cost and M_t^p is the production capacity for period t ; $P_t(q_t) = a_t - b_t \sum_{j=1}^m q_t^j$ is the unit market price. The payoff function (6a) is player p 's total profit; constraints (6b) model product conservation between periods; constraints (6c) ensure that the quantities produced are non-negative and whenever there is production ($x_t^p > 0$), the binary variable y_t^p is set to 1 implying the payment of the setup cost F_t^p . In this game, the players influence each other through the unit market price function $P_t(q_t)$.

2 Literature review

Nash [13] defined the most widely accepted concept of solution for non-cooperative games, the Nash equilibrium. From the definition given in the previous section, a NE associates a probability distribution to each player’s set of strategies such that no player has incentive to unilaterally deviate from that NE if the others play according with the equilibrium. In other words, in an equilibrium, simultaneously each player is maximizing her expected payoff given the equilibrium strategies of the other players. In a pure NE only a single strategy of each player has positive probability assigned (*i.e.*, probability 1).

The state-of-the-art game theory tools are confined to “well-behaved” continuous games, where payoff functions and strategy sets meet certain differentiability and concavity conditions, and normal-form games.

The class of continuous games contains a wide range of relevant games. Glicksberg [8] proved that every continuous game has a NE. However, literature focuses in continuous games for which the strategies sets are convex and payoffs are quasi-concave, since by Debreu, Glicksberg and Fan’s famous theorem there is a pure NE which computation can be reduced to a constrained problem by the application of the Karush-Kuhn-Tucker(KKT) conditions to each player’s optimization problem. In this context, an outlier is the work in [18], where the support size of general separable games is studied. Note that the tools mentioned above are not valid, for separable games since they may fail to satisfy concavity conditions.

Finite games have received wide attention in game theory. Nash [13] proved that any finite game has a NE. Daskalis *et al.* [5] proved that computing a NE is PPAD-complete, which is believed to be a class of hard problems, since it is unlikely that *PPAD* (Polynomial Parity Arguments on Directed graphs) is equal to the polynomial time complexity class *P*. Nisan *et al.* [14] describe general algorithms to compute Nash equilibria, which failed to run in polynomial time. We refer the interested reader to the surveys and state-of-art algorithms collected in von Stengel [19]. Currently, some of those algorithms are available on GAMBIT [12], the most up-to-date software for the computation of NE for normal-form games.

On the other hand, enumerating all players’ feasible strategies (as in finite games) for an IPG can be impractical, or the players’ strategies in an IPG might lead to non well-behaved games, for example where the player’s maximization problems are non-concave. This shows that the existent tools and standard approaches for finite games and convex games are not directly applicable to IPGs.

The literature in IPGs is scarce and often focus in the particular structure of specific games. Kostreva [10] and Gabriel *et al.* [7] propose methods to compute NE for IPGs, however it lacks a computational time complexity guarantee and a practical validation through computational results. Köppe *et al.* [9] present a polynomial time algorithm to compute pure NE under restrictive conditions, like number of players fixed and sum of the number of player’s decision variables fixed, to name few. There are important real-world IPGs, in the context of *e.g.*, electricity markets [16], production planning [11], health-care [3]; this highlights the importance of exploring such game models.

3 Existence of Nash equilibria

It can be argued that players’ computational power is bounded and thus, since the space of pure strategies is simpler and contained in the space of mixed strategies – *i.e.*, the space of Borel probability measures – pure equilibria are more plausible outcomes for games with large sets of pure strategies. In this way, it is important to understand the complexity of determining a pure equilibrium to an IPG.

According with Nash famous theorem [13] any finite game has a Nash equilibrium. Since a purely integer bounded IPGs is a finite game, it has a NE. However, Nash theorem does not guarantee that the equilibrium is pure, which is illustrated in the following example.

Example 1 (No pure Nash equilibrium.) Consider the two-player game such that player A solves

$$\underset{x^A}{\text{maximize}} \quad 18x^A x^B - 9x^A \quad \text{subject to} \quad x^A \in \{0, 1\}$$

and player B :

$$\underset{x^B}{\text{maximize}} \quad -18x^A x^B + 9x^B \quad \text{subject to} \quad x^B \in \{0, 1\}.$$

Under the profile $(x^A, x^B) = (0, 0)$ player B has incentive to change to $x^B = 1$; for the profile $(x^A, x^B) = (1, 0)$ player A has incentive to change to $x^A = 0$; for the profile $(x^A, x^B) = (0, 1)$ player A has incentive to change to $x^A = 1$; for the profile $(x^A, x^B) = (1, 1)$ player B has incentive to change to $x^B = 0$. Thus there is no pure NE. The (mixed) NE is $\sigma^A = \sigma^B = (\frac{1}{2}, \frac{1}{2})$ with expected payoff of zero for both players.

In Section 3.1, we classify both the computational complexity of deciding if there is a pure and a mixed NE for an IPG. It will be shown that even with linear payoffs and two players, the problem is Σ_2^P -complete. Then, in Section 3.2, we state sufficient conditions for the game to have finitely supported Nash equilibria.

3.1 Complexity of the existence of NE

The complexity class Σ_2^P contains all decision problems that can be written in the form $\exists x \forall y P(x, y)$, that is, as a logical formula starting with an existential quantifier followed by a universal quantifier followed by a Boolean predicate $P(x, y)$ that can be evaluated in polynomial time; see Chapter 17 in Papadimitriou's book [15].

Theorem 1 *The problem of deciding if an IPG has a pure NE is Σ_2^P -complete.*

Proof. The decision problem is in Σ_2^P , since we are questing if there is a solution in the space of pure strategies such that for any unilateral deviation of a player, her payoff is not improved (and evaluating the payoff value for a profile of strategies can be done in polynomial time).

It remains to prove Σ_2^P -hardness, that is all problems in Σ_2^P can be reduced in polynomial time to the problem of deciding if an IPG has a pure NE. The following problem is Σ_2^P -complete (see Caprara *et al.* [1]) and we will reduce it to a problem of deciding if an IPG has a pure NE:

DeNegre bilevel Knapsack Problem - DN

INSTANCE Non-negative integers n, A, B , and n -dimensional non-negative integer vectors a and b .

QUESTION Is there a binary vector x such that $\sum_{i=1}^n a_i x_i \leq A$ and for all binary vectors y with $\sum_{i=1}^n b_i y_i \leq B$, the following inequality is satisfied

$$\sum_{i=1}^n b_i y_i (1 - x_i) \leq B - 1?$$

Our reduction starts from an instance of DN. We construct the following instance of IPG.

- The game has two players, $M = \{Z, W\}$.
- Player Z controls a binary decision vector z of dimension $2n + 1$; her set of feasible strategies is

$$\begin{aligned} \sum_{i=1}^n a_i z_i &\leq A \\ z_i + z_{i+n} &\leq 1 \quad i = 1, \dots, n \\ z_{2n+1} + z_{i+n} &\leq 1 \quad i = 1, \dots, n. \end{aligned}$$

- Player W controls a binary decision vector w of dimension $n + 1$; her set of feasible strategies is

$$Bw_{n+1} + \sum_{i=1}^n b_i w_i \leq B. \quad (10)$$

- Player Z 's payoff is $(B - 1)w_{n+1}z_{2n+1} + \sum_{i=1}^n b_i w_i z_{i+n}$.
- Player W 's payoff is $(B - 1)w_{n+1} + \sum_{i=1}^n b_i w_i - \sum_{i=1}^n b_i w_i z_i - \sum_{i=1}^n b_i w_i z_{i+n}$.

We claim that in the constructed instance of IPG there is an equilibrium if and only if the DN instance has answer YES.

(Proof of if). Assume that the DN instance has answer YES. Then, there is x satisfying $\sum_{i=1}^n a_i x_i \leq A$ such that $\sum_{i=1}^n b_i y_i (1 - x_i) \leq B - 1$. Choose as strategy for player Z , $\hat{z} = (x, \overbrace{0, \dots, 0}^n, 1)$ and for player W $\hat{w} = (\overbrace{0, \dots, 0}^n, 1)$. We will prove that (\hat{z}, \hat{w}) is an equilibrium. First, note that these strategies are guaranteed to be feasible for both players. Second, note that none of the players has incentive to deviate from (\hat{z}, \hat{w}) :

- Player Z 's payoff is $B - 1$, and $B - 1 \geq \sum_{i=1}^n b_i w_i$ holds for all the remaining feasible strategies w of player W .
- Player W 's has payoff $B - 1$ which is the maximum possible given \hat{z} .

(Proof of only if). Now assume that the IPG instance has answer YES. Then, there is a pure equilibrium (\hat{z}, \hat{w}) .

If $\hat{w}_{n+1} = 1$, then, by (10), $\hat{w} = (\overbrace{0, \dots, 0}^n, 1)$. In this way, since player Z maximizes her payoff in an equilibrium, $\hat{z}_{2n+1} = 1$, forcing $\hat{z}_{i+n} = 0$ for $i = 1, \dots, n$. The equilibrium inequalities (3), applied to player W , imply that, for any of her feasible strategies w with $w_{n+1} = 0$, $B - 1 \geq \sum_{i=1}^n b_i w_i (1 - \hat{z}_i)$ holds, which shows that DN is a YES instance with the leader selecting $x_i = \hat{z}_i$ for $i = 1, \dots, n$.

If $\hat{w}_{n+1} = 0$, under the equilibrium strategies, player Z 's payoff term $(B - 1)\hat{w}_{n+1}z_{2n+1}$ is zero. Thus, since in an equilibrium player Z maximizes her payoff, it holds that $\hat{z}_{i+n} = 1$ for all $i = 1, \dots, n$ with $\hat{w}_i = 1$. However, this implies that player W 's payoff is non-positive given the profile (\hat{z}, \hat{w}) . In this way, player W would strictly improve her payoff by unilaterally deviating to $w = (\overbrace{0, \dots, 0}^n, 1)$. In conclusion, w_{n+1} is never zero in a pure equilibrium of the constructed game instance. \square

Extending the existence property to mixed equilibria would increase the chance of an IPG to have a NE, and thus, a solution. Next theorem shows that the problem remains Σ_2^P -complete.

Theorem 2 *The problem of deciding if an IPG has a NE is Σ_2^P -complete.*

Proof. Analogously to the previous proof, the problem belongs to Σ_2^P .

It remains to show that it is Σ_2^P -hard. We will reduce the following Σ_2^P -complete to it (see [1]):

Dempe Ritsch Problem - DR

INSTANCE Non-negative integers n , A , C and C' , and n -dimensional non-negative integer vectors a and b .

QUESTION Is there a value for x such that $C \leq x \leq C'$ and for all binary vectors satisfying $\sum_{i=1}^n b_i y_i \leq x$, the following inequality holds

$$Ax + \sum_{i=1}^n a_i y_i \geq 1?$$

Our reduction starts from an instance of DR . We construct the following instance of IPG.

- The game has two players, $M = \{Z, W\}$.

- Player Z controls a non-negative variable z and a binary decision vector (z_1, \dots, z_{n+1}) ; her set of feasible strategies is

$$\begin{aligned} \sum_{i=1}^n b_i z_i &\leq z \\ z_i + z_{n+1} &\leq 1, & i = 1, \dots, n \\ z &\leq C'(1 - z_{n+1}) \\ z &\geq C(1 - z_{n+1}). \end{aligned}$$

- Player W controls a non-negative variable w and binary decision vector (w_1, \dots, w_n) .
- Player Z 's payoff is $Az + \sum_{i=1}^n a_i z_i w_i + z_{n+1}$.
- Player W 's payoff is $z_{n+1}w + \sum_{i=1}^n b_i w_i z_i$.

We claim that in the constructed instance of IPG there is an equilibrium if and only if the DR instance has answer YES.

(Proof of if). Assume that the DR instance has answer YES. Then, there is x such that $C \leq x \leq C'$ and $Ax + \sum_{i=1}^n a_i y_i \geq 1$ for a y satisfying $\sum_{i=1}^n b_i y_i \leq x$. As strategy for player Z choose $\hat{z} = C'$ and $(\hat{z}_1, \dots, \hat{z}_n, \hat{z}_{n+1}) = (y_1, \dots, y_n, 0)$; for player W choose $\hat{w} = 0$ and $(\hat{w}_1, \dots, \hat{w}_n) = (y_1, \dots, y_n)$. We prove that (\hat{z}, \hat{w}) is an equilibrium. First, note that these strategies are guaranteed to be feasible for both players. Second, note that none of the players has incentive to deviate from (\hat{z}, \hat{w}) :

- Player Z 's payoff cannot be increased, since it is equal or greater than 1 and for $i = 1, \dots, n$ such that $\hat{z}_i = 0$ the payoff coefficients are zero.
- Analogously, player W 's payoff cannot be increased, since for $i = 1, \dots, n$ such that $\hat{w}_i = 0$ the payoff coefficients are zero and the payoff coefficient of $\hat{z}_{n+1}\hat{w}$ is also zero.

(Proof of only if). Assume that DR is a NO instance. Then, for any x in $[C, C']$ the leader is not able to guarantee a payoff of 1. This means that in the associated IPG, player Z has incentive to choose $z = 0$ and $(z_1, \dots, z_n, z_{n+1}) = (0, \dots, 0, 1)$. However, this player Z 's strategy leads to a player W 's unbounded payoff. In conclusion, there is no equilibrium. \square

In the proof of Theorem 2, it is not used the existence of a mixed equilibrium to the constructed IPG instance. Therefore, it implies Theorem 1. The reason for presenting these two theorems is because in Theorem 1, the reduction is a game where the players have finite sets of strategies, while in Theorem 2, in the reduction, a player has an unbounded set of strategies.

3.2 Conditions for the existence of NE

Glicksberg [8] and Stein et al. [18] provide results on the existence and characterization of equilibria for continuous and separable games (recall the definitions in Section 1.1), which we will apply to IPGs. Their proofs rely on the fact that the payoff functions have the form (4), enabling to describe the “degree” of interdependence among players.

In an IPG, each player p 's strategy set X^p is a nonempty compact metric space if X^p is bounded and nonempty. This together with the fact that in Section 1.1 we assumed that each player's payoff is continuous, allow us to conclude the following:

Lemma 1 *Every IPG such that X^p is nonempty and bounded is a continuous game.*

Given that every continuous game has a NE [8],

Theorem 3 *Every IPG such that X^p is nonempty and bounded has a Nash equilibrium.*

Applying Stein et al. [18] results, we obtain the following:

Theorem 4 For any Nash equilibrium σ of a separable IPG, there is a Nash equilibrium τ such that each player p mixes among at most $k_p + 1$ pure strategies and $\Pi^p(\sigma) = \Pi^p(\tau)$.

Proof. Apply Theorem 2.8 of [18] to a separable IPG. \square

If in an IPG each player's set of strategies X^p is bounded and the payoff takes the form (4), IPG is separable. Assuming that these two conditions are satisfied (so that Theorem 3 and Theorem 4 hold) is not too strong when modeling real-world applications. In other words, the players' strategies are likely to be bounded due to limitations in the players' resources, which guarantees that an IPG has an equilibrium (Theorem 3). For instance, recall the knapsack game and the competitive lot-sizing game from Section 1.2 in which each player's set of strategies is bounded. In the knapsack game, payoffs are linear, thus by Theorem 4, we deduce that the bound on the equilibria supports for each player is $n + 1$.

Interesting IPGs, the competitive lot-sizing game (recall Section 1.2), has quadratic payoff functions that can be written in the form (4).

Corollary 1 Let IPG be such that X^p is nonempty and bounded, and

$$\Pi^p(x^p, x^{-p}) = c^p x^p + \sum_{k \in M} (x^k)^\top Q_k^p x^p, \quad (12)$$

where $c^p \in \mathbb{R}^{n_p}$ and Q_k^p is a $n_k \times n_p$ real matrix. Then, for any Nash equilibrium σ there is a Nash equilibrium τ such that each player p mixes among at most $1 + n_p + \frac{n_p(n_p-1)}{2}$ pure strategies and $\Pi^p(\sigma) = \Pi^p(\tau)$.

Proof. In order to write player p 's payoff in the form (4), there must be a function $f_{j_p}^p(x^p)$ for $1, x_1^p, \dots, x_{n_p}^p, x_1^p x_1^p, x_1^p x_2^p, \dots, x_1^p x_{n_p}^p, x_2^p x_2^p, \dots, x_{n_p}^p x_{n_p}^p$; thus, $k_p = 1 + n_p + \frac{n_p(n_p-1)}{2}$ in Theorem 4. \square

The thesis [2] presents an algorithmic approach that uses the fact that we can restrict our investigations to finitely supported NE.

4 Conclusions and further directions

Literature in non-cooperative game theory lacks the study of games with diverse sets of strategies with practical interest. This paper is a first attempt to address the computational complexity and existence of equilibria to integer programming games.

We classified the game's complexity in terms of existence of pure and mixed equilibria. For both cases, it was proved that the problems are Σ_2^P -complete. However, if the players' set of strategies is bounded, the game is guaranteed to have an equilibrium. Chen *et al.* [4] proved that computing a NE for a finite game is PPAD-complete even with only two players. Thus, recalling Figure 1, computing a NE to a separable IPG is PPAD-hard. Even when there are equilibria, the computation of one is a PPAD-hard problem, which is likely to be a class of hard problems. Furthermore, the PPAD class does not seem to provide a tight classification of the computational complexity of computing an equilibrium in IPGs. In fact, the PPAD class has its root in finite games that are an easier class of games, in comparison with general IPGs. Note that for IPGs, verifying if a profile of strategies is an equilibrium implies solving each player's best response optimization, which can be a NP-complete problem, while for finite games this computation can be done efficiently. In this context, it would be interesting to explore the definition of a "second level PPAD" class, that is, a class of problems for which a solution could be verified in polynomial time if there was access to a NP oracle.

In this paper, we also determined sufficient conditions for the existence of equilibria on IPGs. Moreover, these theoretical results enabled us to conclude that the support of a NE is finite. This is a key result in the correctness of the algorithm that computes an equilibrium for an IPG presented in [2]. Future work in this context should address the question of determining all equilibria, computing an equilibrium satisfying

a specific property (*e.g.*, computing the equilibrium that maximizes the social welfare, computing a non-dominated equilibrium) and equilibria refinements or new solution concepts under a games with multiple equilibria. From a mathematical point of view, the first two questions embody a big challenge, since it seems to be hard to extract problem structure to the general IPG class of games. The last question raises another one, which is the possibility of considering different solution concepts to IPGs.

References

- [1] Alberto Caprara, Margarida Carvalho, Andrea Lodi, and Gerhard J Woeginger. A study on the computational complexity of the bilevel knapsack problem. *SIAM Journal on Optimization*, 24(2):823–838, 2014.
- [2] Margarida Carvalho. *Computation of equilibria on integer programming games*. PhD thesis, Faculdade de Ciências da Universidade do Porto, 2016.
- [3] Margarida Carvalho, Andrea Lodi, João Pedro Pedroso, and Ana Viana. Nash equilibria in the two-player kidney exchange game. *Mathematical Programming*, pages 1–29, 2016.
- [4] Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In *Foundations of Computer Science, 2006. FOCS '06. 47th Annual IEEE Symposium on*, pages 261–272, Oct 2006.
- [5] Constantinos Daskalakis, Paul Goldberg, and Christos Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [6] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.
- [7] Steven A. Gabriel, Saule Ahmad Siddiqui, Antonio J. Conejo, and Carlos Ruiz. Solving discretely-constrained Nash-Cournot games with an application to power markets. *Networks and Spatial Economics*, 13(3):307–326, 2013.
- [8] I. L. Glicksberg. A Further Generalization of the Kakutani Fixed Point Theorem, with Application to Nash Equilibrium Points. *Proceedings of the American Mathematical Society*, 3(1):170–174, 1952.
- [9] Matthias Köppe, Christopher Thomas Ryan, and Maurice Queyranne. Rational generating functions and integer programming games. *Oper. Res.*, 59(6):1445–1460, November 2011.
- [10] Michael M. Kostreva. Combinatorial optimization in Nash games. *Computers & Mathematics with Applications*, 25(10 - 11):27 – 34, 1993.
- [11] Hongyan Li and Joern Meissner. Competition under capacitated dynamic lot-sizing with capacity acquisition. *International Journal of Production Economics*, 131(2):535 – 544, 2011.
- [12] Richard D. Mckelvey, Andrew M. McLennan, and Theodore L. Turocy. Gambit: Software Tools for Game Theory. Version 16.0.0. <http://www.gambit-project.org>, 2016.
- [13] John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, September 1951.
- [14] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [15] Christos M. Papadimitriou. *Computational complexity*. Addison-Wesley, Reading, Massachusetts, 1994.
- [16] M.V. Pereira, S. Granville, M.H.C. Fampa, R. Dix, and L.A. Barroso. Strategic bidding under uncertainty: a binary expansion approach. *Power Systems, IEEE Transactions on*, 20(1):180–188, Feb 2005.
- [17] Yves Pochet and Laurence A. Wolsey. *Production Planning by Mixed Integer Programming (Springer Series in Operations Research and Financial Engineering)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- [18] Noah D. Stein, Asuman Ozdaglar, and Pablo A. Parrilo. Separable and low-rank continuous games. *International Journal of Game Theory*, 37(4):475–504, 2008.
- [19] Bernhard von Stengel, editor. *Economic Theory: Special Issue of on Computation of Nash Equilibria in Finite Games*, volume 42. Springer Berlin/Heidelberg, January 2010.