

Incentive Stackelberg Mean-payoff Games [★]

Anshul Gupta¹, Sven Schewe¹, Ashutosh Trivedi²,
Maram Sai Krishna Deepak², and Bharath Kumar Padarthy²

¹ University of Liverpool, UK

² Indian Institute of Technology Bombay, India

Abstract. We introduce and study incentive equilibria for multi-player mean-payoff games. Incentive equilibria generalise well-studied solution concepts such as Nash equilibria and leader equilibria. Recall that a strategy profile is a Nash equilibrium if no player can improve his payoff by changing his strategy unilaterally. In the setting of incentive and leader equilibria, there is a distinguished player—called the leader—who can assign strategies to all other players, referred to as her followers. A strategy profile is a leader strategy profile if no player, except for the leader, can improve his payoff by changing his strategy unilaterally, and a leader equilibrium is a leader strategy profile with a maximal return for the leader. In the proposed case of incentive equilibria, the leader can additionally influence the behaviour of her followers by transferring parts of her payoff to her followers. The ability to incentivise her followers provides the leader with more freedom in selecting strategy profiles, and we show that this can indeed improve the leader’s payoff in such games. The key fundamental result of the paper is the existence of incentive equilibria in mean-payoff games. We further show that the decision problem related to constructing incentive equilibria is NP-complete. On a positive note, we show that, when the number of players is fixed, the complexity of the problem falls in the same class as two-player mean-payoff games. We present an implementation of the proposed algorithms, and discuss experimental results that demonstrate the feasibility of the analysis.

1 Introduction

The classical mean-payoff games [30,7] are two-player zero-sum games that are played on weighted finite graphs, where two players—Max and Min—take turn to move a token along the edges of the graph to jointly construct an infinite play. The objectives of the players Max and Min are to respectively maximise and minimise the limit average reward associated with the play. Mean-payoff games enjoy a special status in verification, since μ -calculus model checking and parity games can be reduced in polynomial-time to solving mean-payoff games. Mean-payoff objectives can also be considered as quantitative extensions [15,13]

[★] This work was supported by the EPSRC through grant EP/M027287/1 (Energy Efficient Control), by DARPA under agreement number FA8750-15-2-0096 and by the US National Science Foundation (NSF) under grant numbers CPS-1446900.

of classical Büchi objectives, where we are interested in the limit-average share of occurrences of accepting states rather than merely in whether or not infinitely many accepting states occur. For a broader discussion on quantitative verification, in general, and the transition from the classical qualitative to the modern quantitative interpretation of deterministic Büchi automata, we refer the reader to Henzinger’s survey on quantitative reactive modelling and verification [15].

We focus on multi-player extension of mean-payoff games where a finite number of players control various vertices and move a token along the edges to collectively produce an infinite run. There is a player-specific reward function that, for every edge of the graph, gives an immediate reward to each player. The payoff to a player associated with a play is the limit average of the rewards in the individual moves. The most natural question related to the multi-player game setting is to find an optimal ‘stable’ *strategy profile* (a set of strategies, one for each player). Broadly speaking, a strategy profile is stable, if no player has an incentive to deviate from it. Nash equilibria [19] and leader equilibria [29,13] (also known as *Stackelberg equilibria*) are the most common notions of stable strategy profiles for multi-player games.

A strategy profile is called a Nash equilibrium if no player can improve his payoff by unilaterally changing his strategy. In a setting where we have a distinguished player (called the leader) who is able to suggest a strategy profile to other players (called followers), a strategy profile is stable if no follower can improve his payoff by unilaterally deviating from the profile. A leader equilibrium is a stable strategy profile that maximises the reward for the leader.

In this paper, we introduce and study a novel notion of stable strategy profiles for multi-player mean-payoff games that we call incentive Stackelberg equilibria (or incentive equilibria for short). In this setting, the leader has more powerful strategies, where she not only puts forward strategies that describe how the players move, but also gives non-negative incentives to the followers for compliance. These incentives are then added to the overall rewards the respective follower would receive in each move of the play, and deduced from the overall reward of the leader. Like for leader equilibria, a strategy profile is stable if no *follower* has an incentive to deviate. An *incentive equilibrium* is a stable strategy profile with maximal reward for the leader.

Using incentive equilibria has various natural justifications. The techniques we discussed here can be applied where distributed development of a system is considered. That is, when several rational components interact among themselves along with a rational controller and they try to optimise their individual objectives and specifications. Our techniques can be applied to maximise utility of a central controller while also complying with individual component specifications. Transferring utilities is also quite natural where the payoffs on the edges directly translate to the gains incurred by individual components. These techniques can also be used to maximise social optima where rational controller follow the objective of maximising joint utility. We now discuss two simple examples that exemplify the role that incentives can play to achieve good stable solutions of multi-player mean-payoff games.




Fig. 1. (a) Incentive equilibrium beats leader equilibrium and Nash equilibrium and (b) Incentive equilibrium gives much better system utilisation.

Example 1. Consider the multi-player mean-payoff game shown in Figure 1(a). Here we have three players: Player 1, Player 2 (the leader), and Player 3. The vertex labelled 1 is controlled by Player 1, while the vertex labelled 2 is controlled by Player 2. All other vertices are controlled by Player 3. We further annotate the rewards of various players on the edges of the graph by giving a triple, where the reward of the players 1, 2, and 3 are shown in that order. We omit the labels when the rewards of all players are 0. An incentive equilibrium would be given by (a strategy profile leading to) the play $\langle 1, 2, 3^\omega \rangle$, where the leader pays an incentive of 1 to Player 1 for each step and 0 to Player 3. By doing this, she secures a payoff of 8 for herself. The reward for the players 1 and 3 in this incentive equilibrium are each 1 and -9 , respectively. A leader equilibrium would result in the play $\langle 1, 2, 5^\omega \rangle$ with rewards of 1 for Player 1 and the leader and -2 for Player 3: when the leader cannot pay any incentive to Player 1, then the move from Vertex 2 to Vertex 3 will not be part of a stable strategy. The only Nash equilibrium in this game would result in the play $\langle 1, 4^\omega \rangle$ with the rewards of 1 for Player 1, 0 for the leader, and -1 for Player 3. This example shows how leader benefits from her additional choices in leader and incentive equilibria. \square

Example 2. Consider the multi-player mean-payoff game shown in the Figure 1(b) with five players—Player 0 (or: leader) and Player 1 to 4 (followers). For $i \in \{1, 2, 3, 4\}$, Player i controls the vertex labelled i in the game and gets a reward of 1 whenever token is at vertex i . (To keep the rewards on the edges, one could encode this by giving this reward whenever vertex i is entered.) Player 0 gets a reward of 1 in all of these vertices. The payoff of all other players is 0 in all other cases. Notice that the only play defined by Nash or leader equilibria in this example is $\langle (1, 5, 6)^\omega \rangle$, which provides a payoff of $\frac{1}{3}$ to Player 0 and Player 1, and a payoff of 0 to all other players. For incentive equilibria, however, the leader can give an incentive of $\frac{1}{12}$ to all followers when they follow the play $\langle (1, 2, 3, 4)^\omega \rangle$. It is easy to see that such a strategy profile is incentive stable. The leader will then receive a payoff of $\frac{2}{3}$, i.e., her payoff from the cycle, 1, minus the incentives given to the other players, $4 \cdot \frac{1}{12}$. All other players receive a payoff of $\frac{1}{3}$, consisting of the payoff from the cycle, $\frac{1}{4}$, plus the incentive they receive from the leader, $\frac{1}{12}$. Notice that this payoff is not only better from the leader's point-of-view, the other players are also better off in this equilibrium. \square

In both examples, we saw that the incentive equilibria are strictly better than Nash and leader equilibria. It is not a coincidence—note that leader reward from any Nash equilibrium cannot be greater than her reward from any leader equilibrium, as in the case of leader strategy profiles, leader can select from a wider range of strategy profiles. Thus, if compared to a Nash equilibrium, a leader equilibrium can only be superior w.r.t. the leader reward. Similarly, a leader equilibrium cannot beat an incentive equilibrium, as here also, leader can select from a wider range of strategy profiles (‘leader stable’ strategy profiles can be viewed as an ‘incentive stable’ strategy profiles with 0 incentives). It again implies that leader reward from any leader equilibrium cannot be greater than her reward from any incentive equilibrium.

Related Work. Ummels and Wojtczak [27,28] considered Nash equilibria for mean-payoff games and showed that the decision problem of finding a Nash equilibria is NP-complete for pure (not allowing randomisation) strategy profiles, while the problem is undecidable for arbitrary randomised strategies. Gupta and Schewe [13] have extended these results to leader equilibria. The undecidability result of [28] for Nash equilibria in arbitrary randomised strategies can be easily extended to leader equilibria. For this reason, we focus on non-randomised strategies throughout this paper. Leader equilibria were introduced by von Stackelberg [29] and were further studied in [11]. The strategy profiles we study here are inspired from [12] and are studied in detail for infinite games in [13]. Incentive equilibria have recently been introduced for bi-matrix games [14], but have, to the best of our knowledge, not been used in infinite games. Two-player mean-payoff games were first studied in [9] and were shown to be positionally determined. They can be solved in pseudo-polynomial time [30,6], smoothed polynomial time [3], PPAD [10] and randomised subexponential [2] time. Their decision problem is also known to be in $UP \cap co-UP$ [30,16].

Contributions. The key contribution of the paper is the concept of incentive equilibria to system analysis in general and to multi-player mean-payoff games in particular. We show that the complexity of finding incentive equilibria is same as that for finding leader equilibria [13] for multi-player mean-payoff games: it is NP-complete in general, but, for a fixed number of players, it is in the same complexity class as solving two-player mean-payoff games (2MPGs). In other words, solving two-player mean-payoff games is the most expensive step involved. We have implemented an efficient version of the optimal strategy improvement algorithm from [24] as a backbone, and equipped it with a logarithmic search to expand it from the qualitative evaluation (finding mean partitions) of mean-payoff games to their quantitative evaluation. We construct incentive equilibria by implementing a constraint system that gives necessary and sufficient conditions for a strategy profile to be (1) stable and (2) provide optimal leader return among them. The evaluation of the constraint system involves evaluating a bounded number of calls to the linear programming solver. The contribution of the paper is therefore two-fold—first to conceptualise incentive equilibria in multi-player mean-payoff games (Section 2-3), and second to present a tool (Section 4) deriving optimal return for the leader by evaluating a number of multi-player games.

2 Incentive equilibrium

We introduce the concept of incentive equilibria for multi-player mean-payoff games. These games are played among multiple players on a multi-weighted finite directed graph arena where a distinguished player, called the leader, is able to put forward a strategy profile. She will follow the strategy she assigned for herself, while all other players, called her followers, will comply with the strategy she suggested, unless they benefit from unilateral deviation. The leader is further allowed to incentivise the behaviour of her followers by sharing her payoff with them, in order to make compliance with the strategy she has put forward sufficiently attractive. This, in turn, may improve the leader's payoff. Before we define incentive equilibria, let us recall a few key definitions.

Definition 1 (Multi-player Mean-Payoff Games). *A multi-player mean-payoff game (MMPG) arena \mathcal{G} is a tuple $(P, V, (V_p)_{p \in P}, v_0, E, (r_p)_{p \in P})$ where*

- P is a finite set of players with a distinguished leader player $l \in P$,
- V is a finite set of vertices with a distinguished initial vertex $v_0 \in V$,
- $(V_p)_{p \in P}$ is a partition of V characterising vertices controlled by players,
- $E \subseteq V \times V$ is a set of edges such that for all $v \in V$ there is $v' \in V$ with $(v, v') \in E$, and
- $(r_p)_{p \in P}$ is a family of reward functions $r_p : E \rightarrow \mathbb{Q}$, that for each player $p \in P$, assigns reward for player p associated with that edge.

A finite play $\pi = \langle v_0, v_1, \dots, v_n \rangle$ of the game \mathcal{G} is a sequence of vertices such that v_0 is the initial vertex, and for every $0 \leq i < n$, we have, $(v_i, v_{i+1}) \in E$. An infinite play is defined in an analogous manner. A multi-player mean-payoff game is played on a game arena \mathcal{G} among several players by moving a token along the edges of the arena. The game begins by placing a token on the initial vertex. Each time the token is on the vertex controlled by a player $p \in P$, the player p chooses an outgoing edge and moves the token along this edge. The game continues in this fashion forever, and the players thus construct an infinite play of the game. The (raw) payoff $r_p(\pi)$ of a player $p \in P$ associated with a play $\pi = \langle v_0, v_1, \dots \rangle$ is the limit average reward of the path, given as $r_p(\pi) \stackrel{\text{def}}{=} \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} r_p((v_i, v_{i+1}))$. We refer to this value as the raw payoff of the player p to distinguish it from the payoff for the player that also includes the incentive given to the player by the leader.

A strategy of a player is a recipe for the player to choose the successor vertex. It is given as a function $\sigma_p : V^* V_p \rightarrow V$ such that $\sigma_p(\pi)$ is defined for a finite play $\langle v_0, \dots, v_n \rangle$ when $v_n \in V_p$ and it is such that $(v_n, \sigma_p(\pi)) \in E$. A family of strategies $\bar{\sigma} = (\sigma_p)_{p \in P}$ is called a strategy profile. Given a strategy profile $\bar{\sigma}$, we write $\bar{\sigma}(p)$ for the strategy of player $p \in P$ in $\bar{\sigma}$. A strategy profile $\bar{\sigma}$ defines a unique play $\pi_{\bar{\sigma}}$, and therefore a raw payoff $r_p(\bar{\sigma}) = r_p(\pi_{\bar{\sigma}})$ for each player p . We write $\Sigma_p^{\mathcal{G}}$ for the set of strategy of player $p \in P$ and $\Pi^{\mathcal{G}}$ for the set of strategy profiles in a game arena \mathcal{G} . When the game arena is clear from the context, we omit it from the superscript.

For a strategy profile $\bar{\sigma}$, a player $p \in P$, and a strategy σ' of p , we write $\bar{\sigma}_{p,\sigma'}$ for the strategy profile $\bar{\sigma}'$ such that $\bar{\sigma}'(p) = \sigma'$ and $\bar{\sigma}'(p') = \bar{\sigma}(p')$ for all $p' \in P \setminus \{p\}$. We can now define Nash and leader (aka Stackelberg) equilibria.

Definition 2 (Nash Equilibria and Leader Equilibria). A strategy profile $\bar{\sigma}$ is a Nash equilibrium if no player would gain from unilateral deviation, i.e., for all $p \in P$ we have $r_p(\bar{\sigma}) \geq r_p(\bar{\sigma}_{p,\sigma'})$ for all $\sigma' \in \Sigma_p$. A strategy profile $\bar{\sigma}$ is a leader strategy profile if no player, except for the leader, would gain from unilateral deviation, i.e., for all $p \in P \setminus \{l\}$ we have $r_p(\bar{\sigma}) \geq r_p(\bar{\sigma}_{p,\sigma'})$ for all $\sigma' \in \Sigma_p$. A leader equilibrium is a maximal leader strategy profile.

We next define an *incentive strategy profile* as a strategy profile which satisfies the stability requirements of the leader equilibria and allows the leader to give incentives to the followers. We refer to an optimal strategy profile in this class that provides maximal reward to the leader as an *incentive equilibrium*.

An incentive to a player p is a function $\iota_p: V^*V_p \rightarrow \mathbb{R}_{\geq 0}$ from the set of histories to incentives. Incentives can be extended to infinite play $\pi = \langle v_0, v_1, \dots \rangle$ in the usual mean-payoff fashion: $\iota_p(\pi) \stackrel{\text{def}}{=} \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \iota_p(v_0 \dots v_{n-1})$. The overall payoff $\rho_p(\pi)$ to a follower in run π is the raw payoff plus all incentives, $\rho_p(\pi) \stackrel{\text{def}}{=} r_p(\pi) + \iota_p(\pi)$, while the overall payoff of the leader $\rho_l(\pi)$ is her raw payoff after deducting all incentives, $\rho_l(\pi) \stackrel{\text{def}}{=} r_l(\pi) - \sum_{p \in P \setminus \{l\}} \iota_p(\pi)$.

We extend the notion of a strategy profile in the presence of incentives as a pair $(\bar{\sigma}, \bar{\iota})$, where $\bar{\sigma}$ is a strategy profile assigned by the leader, in which the leader pays an incentive given by the incentive profile $\bar{\iota} = (\iota_p)_{p \in P \setminus \{l\}}$. We write $\bar{\iota}_p$ for the incentive for player $p \in P \setminus \{l\}$. We write $\bar{\iota}_p(\bar{\sigma})$ for the incentive to player p for the unique run $\pi_{\bar{\sigma}}$ under incentive profile $\bar{\iota}$. In any incentive strategy profile $(\bar{\sigma}, \bar{\iota})$, no player but the leader may benefit from deviation. An optimal strategy profile among this class would form an *incentive equilibrium*.

Definition 3 (Incentive Equilibria). A strategy profile $(\bar{\sigma}, \bar{\iota})$ is an incentive strategy profile, if no follower can improve his overall payoff from a unilateral deviation, i.e., for all players $p \in P \setminus \{l\}$ we have that $r_p(\bar{\sigma}) + \bar{\iota}_p(\bar{\sigma}) \geq r_p(\bar{\sigma}_{p,\sigma'}) + \bar{\iota}_p(\bar{\sigma}_{p,\sigma'})$ for all $\sigma' \in \Sigma_p$. An incentive strategy profile $(\bar{\sigma}, \bar{\iota})$ is an incentive equilibrium if the leader's total payoff for this profile is maximal among all incentive strategy profiles, i.e., for all $(\bar{\sigma}', \bar{\iota}')$ we have that $r_l(\bar{\sigma}) - \sum_{p \in P \setminus \{l\}} \bar{\iota}_p(\bar{\sigma}) \geq r_l(\bar{\sigma}') - \sum_{p \in P \setminus \{l\}} \bar{\iota}'_p(\bar{\sigma}')$.

An incentive strategy profile is a *leader strategy profile* if all incentives are 0, and a *Nash strategy profile* if, in addition, $\bar{\sigma}$ is also a Nash equilibrium. We write SP, ISP, LSP, and Nash SP for the set of strategy profiles, incentive strategy profiles, leader strategy profiles, and Nash strategy profiles respectively. It is clear that $\text{Nash SP} \subseteq \text{LSP} \subseteq \text{ISP} \subseteq \text{SP}$. This observation, together with Example 1(a) yield the following result.

Theorem 1 (IE \geq LE \geq NE). Incentive equilibria do not provide smaller return than leader equilibria, and leader equilibria do not provide smaller return than Nash equilibria. Moreover, there are games for which the leader reward from three equilibria are different.




Fig. 2. Gadget reducing incentive equilibrium to leader equilibrium. Here we assume that ℓ is the leader player and $1, \dots, n$ are n followers. Here we label a state by ℓ if it is controlled by leader, and by p if it is controlled by a player p . Every state p of the original game arena is replaced by $2n + 2$ states in the gadget in the manner shown above. Here reward vector b_i corresponds to reward $-h$ for the leader, reward $+h$ to follower i , and reward 0 to all other followers.

3 Existence, Construction, and Complexity

This section is dedicated to the existence, construction, and complexity of incentive equilibria. Similar results were developed by Gupta and Schewe [13] for leader equilibria for multi-player mean-payoff games. If we are only interested in complexity classes, the results in the following section can be obtained by a simple reduction to leader equilibria [13] using a gadget (Fig. 2) that replaces each state by intuitively circling through the followers and allowing the leader to transfer a large part h (for *high* incentive) of her utility to each of the followers³. Following this approach leads, however, to an increase in the size of the linear programs, simply because we would incur a blow-up of the game by a factor of $2|P|$. To avoid this unnecessary blowup, we adapt the proofs from [13] that result in a more efficient algorithm.

Proof Sketch. We first introduce a canonical class of incentive strategy profiles—the *perfectly-incentivised strategy profiles* (PSPs)—that corresponds to the Stackelberg version of the classic subgame perfection. Keep in mind that not all perfectly-incentivised strategy profiles (PSPs) are valid incentive strategy profiles (ISPs). On the other hand, we show that every ISP has a corresponding PSP (which is also an ISP) with the same leader reward. Thanks to this result, in order to construct incentive equilibrium it suffices to consider PSPs that are also ISPs.

³ When translating a strategy in the presence of incentives (σ, ι) , the translation can be done by simulating the runs. The translation of the strategy profile σ is straight forward; it determines the choices at the respective last states from the gadgets. The translation of the incentives refer to the choices within the gadgets. They can be obtained by letting the leader make the decision to transfer h to follower p if, and only if, the sum of the incentives this follower p has collected in the game on the original MMPG is at least h higher than the sum of the utilities the leader has so far transferred to p in the gadgets. If passing through a gadget is counted as one step, all lim inf values agree on the original and its simulation. The back translation is even more direct: it suffices to wait till the end of each gadget, and then assign incentives accordingly.

Further, we show that, for PSPs that are ISPs, it suffices to find a maximum in a *well behaved* class of strategy profiles: strategy profiles where every edge has a limit share of the run—by showing that the supremum of general strategies cannot be higher than the supremum of these well behaved ones. We then show how to construct well behaved PSPs that are ISPs based on a family of constraint systems that depend on the occurring and recurring vertices on the play. At the same time, we show that no general ISP that defines a play with this set of occurring and recurrent vertices can have a higher value. The set of occurring and recurrent vertices can be guessed and the respective constraint system can be build and solved in polynomial time, which also provides inclusion of the related decision problem in NP.

3.1 Perfectly-incentivised strategy profiles

We define a canonical form of an incentive equilibrium with this play that we call *perfectly-incentivised strategy profiles* (PSP). In a PSP, a deviator (a deviating follower) is punished, and the leader incentivises all other followers to collude against the deviator. While the larger set of strategies and plays that define them (when compared to Nash and leader equilibria) lead to a better value, this incentive scheme leads to a higher stability: the games are subgame perfect relative to the leader. A strategy profile $(\bar{\sigma}, \bar{t})$ is a subgame perfect incentive strategy profile, if every reachable subgame is also an incentive strategy profile. This term adjusts the classic notion of subgame perfect equilibria to our setting. In a subgame perfect Nash equilibrium, it is required that the subgame started on each history also forms a Nash equilibrium. Note that the leader is allowed to benefit from deviation in our setting.

The means to obtain subgame perfection after deviation is to make all players harm the most recent deviator. Thus, we essentially resort to a two-player game. For a multi-player mean-payoff game \mathcal{G} , we define, for each follower p , the two-player mean-payoff game (2MPG) \mathcal{G}_p where p keeps his reward function, while all other players have the same antagonistic reward $-r_p$. Two-player mean-payoff games are memoryless determined, such that every vertex v has a value, which we denote by $r_p(v)$. This value clearly defines a minimal payoff of a follower: when he passes by a vertex in a play, then he cannot expect an outcome below $r_p(v)$, as he would otherwise deviate.

PSP strategy profiles are in the tradition of reward and punish strategy profiles [13]. In any 'reward and punish' strategy profile, the leader facilitates the power of all remaining followers to punish a deviator. If a player p chooses to deviate from the strategy profile at history h , the game would turn into a two-player game, where all the other followers and the leader forsake their own interests, and jointly try to 'punish' p . That is, player p may still try to maximise his reward and his objective remains exactly the same, but the rewards of the rest of the players have changed to negative of the reward of player p . As they form a coalition with the joint objective to harm p , this is an ordinary two-player mean-payoff game that starts at the vertex $\text{last}(h)$.

For a strategy profile $\bar{\sigma}$ and a history h , we call h a deviating history, if it is not a prefix of $\pi_{\bar{\sigma}}$. We denote by $\text{dev}(h, \bar{\sigma})$ the last player p , who has deviated from his or her strategy $\bar{\sigma}_p$ on a deviating history h . A perfectly-incentivised strategy profile is defined as a strategy profile (PSP) $(\bar{\sigma}, \bar{\iota})$ with the following properties. For all prefixes h and h' of $\pi_{\bar{\sigma}}$ and for all followers p , it holds that $\bar{\iota}_p(h) = \bar{\iota}_p(h')$. We also refer to this value by $\bar{\iota}_p$. For deviator histories h' , the incentive $\bar{\iota}_p(h')$ is 0 except for the following cases. On every deviating history h with deviating player $p = \text{dev}(h, \bar{\sigma})$, the player p' who owns the vertex $v = \text{last}(h)$ follows the strategy from the 2MPG \mathcal{G}_p . If, under this strategy, player p' selects the successor v' at a vertex v in the 2MPG \mathcal{G}_p (and thus $\bar{\sigma}_{p'}(h) = v'$), p' is a follower, and $p' \neq p$, then player p' receives an incentive, such that $r_{p'}(v, v') + \bar{\iota}_{p'}(h \cdot v') = r_{\max} + 1$.

Note that, technically, the leader punishes herself in this definition. This is only to keep definitions simple; she is allowed to have an incentive to deviate, and the subgame perfection does not impose a criterion upon her. Note also that a PSP is not necessarily an incentive strategy profile, as it does not guarantee anything about $\pi_{\bar{\sigma}}$. The following theorem states the importance of PSPs in constructing incentive equilibrium.

Theorem 2. *Let $(\bar{\sigma}, \bar{\iota})$ be an ISP that defines a play $\pi_{\bar{\sigma}}$. Then we can define a PSP $(\bar{\sigma}, \bar{\iota})$, that is also an ISP, with the same reward, and defines the same play.*

The proof of this theorem follows from Lemma 1 and Lemma 2.

Lemma 1. *Let $(\bar{\sigma}', \bar{\iota}')$ be a strategy profile that defines a play $\pi_{\bar{\sigma}'}$, which contains precisely the reachable vertices Q . Let $(\bar{\sigma}', \bar{\iota}')$ satisfy that, \forall followers $p \in P \setminus \{l\}$ and all vertices $v \in Q \cap V_p$ owned by p we have that $\bar{\iota}_p(\bar{\sigma}') + r_p(\bar{\sigma}') \geq r_p(v)$. Then we can define a PSP $(\bar{\sigma}, \bar{\iota})$ with the same reward, which defines the same play.*

Proof. We note that a PSP $(\bar{\sigma}, \bar{\iota})$ is fully defined by the play $\pi_{\bar{\sigma}}$ and the $\bar{\iota}$ restricted to the prefixes of $\pi_{\bar{\sigma}}$. We now define the PSP $(\bar{\sigma}, \bar{\iota})$ with the following property: $\pi_{\bar{\sigma}} = \pi_{\bar{\sigma}'}$, that is the play of the PSP equals the play defined by the ISP we started with. For all followers p and all prefixes h of $\pi_{\bar{\sigma}}$, we have $\bar{\iota}_p(h) = \bar{\iota}'_p(h)$. It is obvious that $(\bar{\sigma}', \bar{\iota}')$ and $(\bar{\sigma}, \bar{\iota})$ yield the same reward for all followers and the same reward for the leader. We now assume for contradiction that the resulting PSP is not an incentive strategy profile. If this is the case, then a follower p must benefit from deviation at some history h . Let us start with the case that h is a deviator history. In this case, the reward for p upon not deviating is $r_{\max} + 1$, while it is the outcome of some game upon deviation, which is clearly bounded by r_{\max} .

We now turn to the case that h is not a deviator history, and therefore a prefix of $\pi_{\bar{\sigma}}$. Let p be the owner of $v = \text{last}(h)$. If p is the leader, we have nothing to show. If p is a follower and does not have an incentive to deviate in $(\bar{\sigma}, \bar{\iota})$, we have nothing to show. If p is a follower and has an incentive to deviate in $(\bar{\sigma}, \bar{\iota})$, we note that his payoff after deviation would be bounded from above by $r_p(v)$. Thus, he does not have an incentive to deviate (contradiction). \square

Lemma 2. *Let $(\bar{\sigma}, \bar{\iota})$ be an ISP that defines a play $\pi_{\bar{\sigma}}$, which contains precisely the vertices Q . Then, for all followers $p \in P \setminus \{l\}$ and all vertices $v \in Q \cap V_p$ owned by p , we have that $\bar{\iota}_p(\bar{\sigma}) + r_p(\bar{\sigma}) \geq r_p(v)$.*

Proof. Assume that this is not the case for a follower p and a vertex $v \in Q$ owned by p . Then p would benefit upon deviating when visiting v . \square

3.2 Existence and construction of incentive equilibria

We say that a strategy profile $\bar{\sigma}$ is *well-behaved* if in the resulting play $\pi_{\bar{\sigma}}$, the frequency (ratio) of occurrence of every edge of the game arena occurs has a limit, i.e, each edge here occurs with a limit probability (the limes inferior and superior of the share of its occurrence on $\pi_{\bar{\sigma}}$ are equal). Such notion of well-behaved strategy profiles were also defined in [13] for the case of leader equilibria. We first construct optimal ISPs among well behaved PSPs, and then show that no ISP give a better payoff for leader.

Let $\bar{\sigma}$ is a well-behaved perfectly-incentivised strategy profile and let Q be the set of vertices visited in $\pi_{\bar{\sigma}}$ and $S \subseteq Q$ be the set of vertices that are visited infinitely often (note that S is strongly connected). Let $p_{(s,t)}$ be the limit ratio (frequency) of occurrence of an edge $(s, t) \in E \cap S \times S$ in $\pi_{\bar{\sigma}}$ and let p_v be the limit ratio of each vertex $v \in S$.

Thanks to the proof of Lemma 1, the following constraint system (linear program) characterises the necessary and sufficient conditions for the well-behaved perfectly-incentivised strategy profile $\bar{\sigma}$ to be an ISP.

1. $p_v = 0$ if $v \in V \setminus S$ and $p_v \geq 0$ if $v \in S$.
2. $p_e = 0$ if $e \in E \setminus S \times S$ and $p_e \geq 0$ if $e \in E \cap S \times S$
3. $\sum_{v \in V} p_v = 1$
4. $p_s = \sum_{(s,t) \in E} p_{(s,t)}$ for all $s \in S$ and $p_t = \sum_{(s,t) \in E} p_{(s,t)}$ for all $t \in S$
5. $\bar{l}_p + \sum_{e \in E} p_e r_p(e) \geq \max_{v \in Q} (r_p(v))$ where $r_p(v)$ is the value at vertex v in the 2MPG \mathcal{G}_p characterising minimum payoff expected by player p .

The constraints presented above are quite self-explanatory. Constraints 1 and 2 state that the limit ratio of occurrence of a vertex and edge is positive only when it is visited infinitely often. Constraint 3 expresses that the sum of ratio of occurrence of vertices is equal to 1, while constraint 4 expresses the fact the limit ratio of a vertex should be equal to limit ratios of all incoming edges, and equal to limit ratio of all outgoing edges from that vertex. The last constraint stems from the proof of Lemma 1 combined with the observation that reward $r_p(\bar{\sigma})$ of a player p in $\bar{\sigma}$ is simply $\sum_{e \in E} p_e r_p(e)$, that is, it is the weighted sum of the raw rewards of the individual edges. Before we define the objective function, we state a simple corollary from the proof of Lemma 1.

Corollary 1. *Every well behaved PSP that is an ISP satisfies these constraints, and every well behaved strategy profile $(\bar{\sigma}, \bar{l})$, whose play $\pi_{\bar{\sigma}}$ satisfies these constraints, defines a PSP, which is then an ISP.*

Note that the resulting PSP is an ISP even if $(\bar{\sigma}, \bar{l})$ is not. This is because the satisfaction of the constraints are enough for the final contradiction in the proof of Lemma 1.

Construction of incentive equilibria. The objective of the leader is obviously to maximise $r_l(\bar{\sigma}) - \sum_{p \in P \setminus \{l\}} \bar{r}_p = \sum_{e \in E} p_e r_l(e) - \sum_{p \in P \setminus \{l\}} \bar{r}_p$. Once we have this linear programming problem, it is simple to determine a solution in polynomial time [17,18]. We first observe that it is standard to construct a play defining a PSP from a solution. A key observation is that, if the linear program detailed above for sets Q of reachable vertices and S of vertices visited infinitely often has a solution, then there is a well behaved reward and punish strategy profile that meets this solution.

Theorem 3. *Non-well behaved PSPs that are also ISPs cannot provide better rewards for the leader than those from well behaved PSPs that are also ISPs.*

Proof. Corollary 1 shows that there exists a well defined constraint system obeyed by all well behaved PSPs that are also ISPs with a set Q of reachable vertices and a set S of recurrent vertices. Let us assume for contradiction that there is a reward and punish strategy profile $(\bar{\sigma}, \bar{r})$ that defines a play $\pi_{\bar{\sigma}}$ with the same sets Q and S of reachable and recurrent vertices, respectively, that provides a strictly better reward $r_l(\bar{\sigma}) - \sum_{p \in P \setminus \{l\}} \bar{r}_p$, which exceeds the maximal reward obtained by the leader in well behaved PSPs that are also ISPs by some $\varepsilon > 0$.

We now construct a well behaved PSPs that are also ISPs and that also provides a better return. First, we take a \bar{r}' with $\bar{r}_p = \bar{r}'_p$ for all followers p . This allows us to focus on the raw rewards only. Let k be some position in $\pi_{\bar{\sigma}}$ such that, for all $i \geq k$, only positions in the infinity set S of $\pi_{\bar{\sigma}}$ occur. Let π be the tail $v_k v_{k+1} v_{k+2} \dots$ of $\pi_{\bar{\sigma}}$ that starts in position k . Obviously $r_p(\pi) = r_p(\bar{\sigma})$ holds for all players $p \in P$. We observe that, for all $\delta > 0$, there is an $l \in \mathbb{N}$ such that, for all $m \geq l$, $\frac{1}{m} \sum_{i=0}^{m-1} r_p((v_i, v_{i+1})) > r_p(\pi) - \delta$ holds for all $p \in P$, as otherwise the limes inferior property would be violated. We now fix, for all $a \in \mathbb{N}$, a sequence $\pi_a = v_k v_{k+1} v_{k+2} \dots v_{k+m_a}$, such that $v_{k+m_a+1} = v_k$ and $\frac{1}{m} \sum_{i=0}^{m_a-1} r_p((v_i, v_{i+1})) > r_p(\pi) - \frac{1}{a}$ holds for all $p \in P$. Let $\pi_0 = v_0 v_1 \dots v_{k-1}$. We now select $\pi' = \pi_0 \pi_1^{b_1} \pi_2^{b_2} \pi_3^{b_3} \dots$, where the b_i are natural numbers big enough to guarantee that $\frac{b_i \cdot |\pi_i|}{|\pi_{i+1}| + |\pi_0| + \sum_{j=1}^i b_j \cdot |\pi_j|} \geq 1 - \frac{1}{i}$ holds. Letting b_i grow this fast ensures that the payoff, which is at least $r_p(\pi) - \frac{1}{i}$ for all players $p \in P$, dominates till the end of the first iteration⁴ of $|\pi_{i+1}|$. The resulting play belongs to a well behaved (as the limit exists) strategy profile, and can thus be obtained by a well behaved PSP by Corollary 1. It thus provides a solution to the linear program from above, which contradicts our assumption. \square

Consequently, it suffices to guess the optimal sets Q of vertices that occur and S of vertices that occur infinitely often to obtain a constraint system that describes an incentive equilibrium, which is well behaved and a PSP—and therefore subgame perfect.

⁴ Including the first iteration of π_{i+1} is a technical necessity, as a complete iteration of π_{i+1} provides better guarantees, but without the inclusion of this guarantee, the π_j 's might grow too fast, preventing the existence of a limes.

Corollary 2. *The decision problems ‘is there a (subgame perfect) incentive equilibrium with leader reward $\geq r'$ ’ is in NP, and the answer to these two questions is the same.*

Note that, if we have a fixed number of players, the number of possible constraint systems is polynomial. Like in [13], there are only polynomially many (for n vertices and k followers $O(n^k)$ many) second parts (the constraints on the follower rewards) of the constraint systems. For them, it suffices to consider the most liberal sets Q (which is unique) and S (the SCCs in the game restricted to Q , at most n). For a fixed number of players, finding incentive equilibria is therefore in the same class as solving 2MPGs. We adapt the NP hardness proof for leader equilibrium in mean-payoff games from [13] by reducing the 3SAT satisfiability formula over n atomic propositions with m conjuncts to solve a MMPG with $2n + 1$ players and $5m + 4n + 2$ vertices with payoffs 0 and 1 only.

Theorem 4. *The problem of deciding whether an incentive equilibrium σ with reward $r_l(\sigma) \geq 1 - 1/n$ of the leader exists in games with rewards in $\{0, 1\}$, is NP-complete.*

4 Experimental results

We have implemented a tool [1] in C++ to evaluate the performance of the proposed algorithms for multi-player mean-payoff games (MMPG) for a small number of players. We implemented an algorithm from [24] to find mean values at the vertices. We then infer and solve a number of constraint systems. We describe our main algorithm here.

4.1 Algorithm specific details

We first evaluate MMPGs using reduction to solving underlying 2MPGs. We then infer and solve a number of linear programming problems to find a solution. For few number of players, the number of different solutions to these games is usually small, and, consequently, the number of linear programming problems to solve is small, too. In order to find the individual mean partition, we use an algorithm from [24], that finds 0-mean partitions, and expand it quantitatively to find the value of 2MPGs. We recall that for 2MPG both players have optimal memoryless strategies. Under such strategies, the game will follow a ‘lasso path’ from every starting vertex: a finite (and possibly empty) path, followed by a cycle, which is repeated infinitely many times. The value of a game position is defined by the average of the edge weights on this cycle.

In our context, the edge weights are either 0 or 1. The values of the vertices are therefore fractions $\frac{a}{l}$ with $0 \leq a \leq l \leq n$, where l is the length of the cycle, and a is the number of ‘accepting’ events in the DBA that refers to the objective of the respective player, i.e., the edges with value 1, occurring on this cycle.

An α -mean partition of a 2MPG is the subset of vertices, for which the return is $\geq \alpha$. Conceptually, to find the $\frac{a}{l}$ -mean partition, one would simply subtract $\frac{a}{l}$

from the weight of every edge and look for the 0-mean partition. However, to stay with integers, it is better to use integer values on the edges, e.g., by replacing the 0s by $-a$, and the 1s by $l - a$. For games with n vertices, there are only $O(n^2)$ values for the fraction $\frac{a}{l}$ to consider, as optimal memoryless strategies always lead to lasso paths and only the cycle at the end of the lasso determines the values for a and l , where $0 < a < l \leq n$.

We start by narrowing down the set of values by classifying the mean partition in a logarithmic search. After determining the $\frac{1}{2}$ mean partition, we know which values are < 0.5 and ≥ 0.5 , respectively. The two parts of the game can then be analysed further, determining the $\frac{1}{4}$ and $\frac{3}{4}$ mean partition, respectively. After s such partitionings, all values in a partition of the game are either known to be in an $[k \cdot 2^{-s}, (k+1) \cdot 2^{-s}[$ interval for some $k < 2^s - 1$, or in the interval $[1 - 2^{-s}, 1]$. We stop to bisect when the size p of a partition is at most 2^s . In this case, the respective interval has $f \leq p$ fractions with a denominator $\leq p$. We determine them, store them in a balanced tree, and use it to determine the correct value of all vertices of the partition in $\lceil \log_2 f \rceil$ steps.

The number of different values of nodes in a 2MPG is usually small, and certainly it would be much smaller than the number of vertices in the game. Consequently, the number of constraint systems is also small for a small number of players. We use this algorithm to evaluate a number of randomly created three player mean-payoff games, where the player take turns. We consider three players – player 1, player 2 and a leader and two different evaluations on the same game graph. We first see how each player fares when they try to maximise their return against a coalition of all other players, including the leader. In the first evaluation, leader forms a coalition with player 1 (minimiser) against player 2 (maximiser) on the payoffs defined for player 2. We find the different possible mean values at the nodes in this evaluation, using the algorithm from above. In the second evaluation, leader forms a coalition with player 2 (minimiser) against player 1 (maximiser) on the payoffs defined for player 1. We also note the different possible mean values at the nodes in this evaluation, using again the algorithm from above. The resultant two-player games provide the constraints for the linear programming problems that we presented in Section 3.2. These different values form the different thresholds that we have to consider. We now consider all possible combinations of these different threshold values for the followers and determine the vertices that comply with them.

4.2 Experimental results

Experiments indicate that our implementation of the algorithm can solve examples of size 100 nodes and 10 players within 30 minutes. The algorithm is, of course, much faster for the games with two or three players. Figures 3 and 4 show the experimental results for the following two problem classes.

- Recall the example from Figure 1(b). We generalise this example for token ring graph parameterised by 2 variables, n and d . It has ' n ' nodes on the inner cycle, each of which correspond to ' n ' different players and each of these ' n '




Fig. 3. The figure shows results for a generalisation of example 2 for multiple players with n nodes in the inner cycle and $n-1$ in outer cycles where n is number of players.

nodes is also present on another cycle of length ' d '. The weights are set such that, all players except the leader get ' $1/n$ ' if they chose the inner ring and get ' $1/d$ ' if they chose their respective outer ring. The leader reward is ' 1 ' in the inner ring and ' $1/d$ ' in all the other rings. The data supports the pen-and-paper analysis that incentives are useful iff $n > d > n(n-1)/(2n-1)$ holds. Figure 3 shows the leader reward for this example and the running time of our tool to compute it.

- Figure 4 (left plot) shows the difference between incentive equilibrium and leader equilibrium for randomly generated 3 player MMPGs, while the right plot shows similar results on random graphs, where the number of players range from 3 to 10.

The evaluation results confirm that the leader reward increases significantly in incentive equilibria when compared to leader equilibria.

5 Discussion

The main contribution of this paper is the introduction of incentive equilibria in multi-player mean-payoff games and the implementation of our techniques in a tool. We study how a rational leader might improve over her outcome by paying small incentives to her followers. At first, it may not seem to be a rational move of the leader, but close insight would show how a leader might improve her reward in this way. The incentive equilibria are seen as an extension to

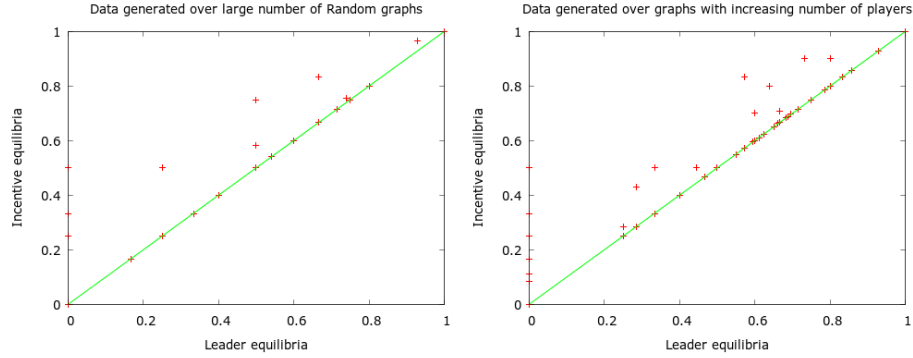


Fig. 4. The left figure shows the results for randomly generated MMPGs with 3 players, while the right one is for randomly generated MMPGs with 3 to 10 players.

leader equilibria, where a rational leader, by giving an incentive to every other player in the game, can derive an optimal strategy profile. We believe that these techniques are helpful for the leader when maximising the return for a single player and would also be instrumental in defining stable rules and optimising various outcomes. The evaluation results from Section 4 show that the results are significantly better for the leader in an incentive equilibrium as compared to her return in a leader equilibrium.

References

1. MMPGSOLVER. <http://www.cse.iitb.ac.in/~trivedi/mmpgsolver/>.
2. H. Björklund and S. Vorobyov. A combinatorial strongly subexponential strategy improvement algorithm for mean-payoff games. *Discrete Applied Math.*, 155(2):210–229, 2007.
3. E. Boros, K. M. Elbassioni, M. Fouz, V. Gurvich, K. Makino, and B. Manthey. Stochastic mean-payoff games: Smoothed analysis and approximation schemes. In *Proceedings of ICALP 2011*, LNCS 6755, pages 147–158, 2011.
4. T. Brihaye, V. Bruyère, and J. De Pril. Equilibria in quantitative reachability games. In *Proceedings of CSR 2010*, LNCS 6072, pages 72–83, 2010.
5. T. Brihaye, J. De Pril, and S. Schewe. Multiplayer Cost Games with Simple Nash Equilibria. In *Proceedings of LFCS 2013*, LNCS 7734, pages 59–73, 2013.
6. L. Brim, J. Chaloupka, L. Doyen, R. Gentilini, and J.-F. Raskin. Faster algorithms for mean-payoff games. *Formal Methods in System Design*, 38(2):97–118, 2011.
7. K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Mean-payoff parity games. In *Proceedings of LICS 2005*, pages 178–187, 2005.
8. K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Games with Secure Equilibria. In *Theor. Comput. Sci.*, pages 67–82, 2006.
9. A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.
10. Kousha Etessami and Mihalis Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM J. Comput.*, 39(6):2531–2597, 2010.

11. J. W. Friedman. *Oligopoly and the theory of games*. North-Holland Pub. Co., 1977. Advanced textbooks in economics.
12. J. W. Friedman. A Non-cooperative Equilibrium for Supergames. In *The Review of Economic Studies*, pages 1–12, 1971.
13. A. Gupta and S. Schewe. Quantitative Verification in Rational Environments. In *Proc. of TIME*, pages 123–131, 2014.
14. A. Gupta and S. Schewe. It Pays to Pay in Bi-Matrix Games – a Rational Explanation for Bribery. In *Proc. of AAMAS*, pages 1361–1369, 2015.
15. T. A. Henzinger. Quantitative reactive modeling and verification. *Computer Science - R&D*, 28(4):331–344, 2013.
16. M. Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68(3):119–124, November 1998.
17. N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proc. of STOC*, pages 302–311, 1984.
18. L. G. Khachian. A polynomial algorithm in linear programming. *Dokl. Akad. Nauk SSSR*, 244:1093–1096, 1979.
19. J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
20. M. J. Osborne and A. Rubinstein. *A course in game theory*. The MIT Press, Cambridge, USA, 1994. electronic edition.
21. M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.
22. A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *Proceedings of POPL*, pages 179–190, 1989.
23. P. J. G. Ramadge and W. M. Wonham. The control of discrete event systems. *Proc. of the IEEE*, 77(2):81–98, 1989.
24. S. Schewe. An optimal strategy improvement algorithm for solving parity and payoff games. In *Proc. of CSL 2008*, LNCS 5213, pages 368–383, 2008.
25. S. Schewe. From parity and payoff games to linear programming. In *Proc. of MFCS 2009*, LNCS 5734, pages 675–686, 2009.
26. M. Ummels. Rational behaviour and strategy construction in infinite multi-player games. In *Proc. of FSTTCS 2006*, pages 212–223, 2006.
27. M. Ummels. The complexity of nash equilibria in infinite multi-player games. In *Proc. of FoSSaCS 2008*, LNCS 4962, pages 20–34, 2008.
28. M. Ummels and D. Wojtczak. The Complexity of Nash Equilibria in Limit-Average Games. In *Proc. of CONCUR 2011*, pages 482–496, 2011.
29. H. von Stackelberg. *Marktform und Gleichgewicht*. J. Springer, 1934.
30. U. Zwick and M. S. Paterson. The complexity of mean-payoff games on graphs. *Theoretical Computer Science*, 158(1–2):343–359, 1996.