Sink Equilibria and Convergence

Michel Goemans*, Vahab Mirrokni† and Adrian Vetta‡

Abstract. We introduce the concept of a *sink equilibrium*. A sink equilibrium is a strongly connected component with no out-going arcs in the strategy profile graph associated with a game. The strategy profile graph has a vertex set induced by the set of pure strategy profiles; its arc set corresponds to transitions between strategy profiles that occur with non-zero probability. (Here our focus will just be on the special case in which the strategy profile graph is actually a best response graph; that is, its arc set corresponds exactly to best response moves that result from myopic or greedy behaviour.) We argue that there is a natural convergence process to sink equilibria in games where agents use pure strategies. This leads to an alternative measure of the social cost of a lack of coordination, the *price of sinking*, which measures the worst case ratio between the value of a sink equilibrium and the value of the socially optimal solution. We define the *value of a sink equilibrium* to be the expected social value of the steady state distribution induced by a random walk on that sink.

We illustrate the value of this measure in three ways. Firstly, we show that it may more accurately reflects the inefficiency of uncoordinated solutions in competitive games when the use of pure strategies is the norm. In particular, we give an example (a valid-utility game) in which the game converges to solutions which are a factor n worse than socially optimal. The price of sinking is indeed n, but the price of anarchy is close to 1. Secondly, sink equilibria always exist. Thus, even in games in which pure strategy Nash equilibria (PSNE) do not exist, we can still calculate the price of sinking. Thirdly, we show that bounding the price of sinking can have important implications for the speed of convergence to socially good solutions in games where the agents make best response moves in a random order.

We present two examples to illustrate our ideas.

(i) *Unsplittable Selfish Routing (and Weighted Congestion Games)*: we prove that the price of sinking for the weighted unsplittable flow version of the selfish routing problem (for bounded-degree polynomial latency functions) is at most $O(2^{2d}d^{2d+3})$. In comparison, we give instances of these games without any PSNE. Moreover, our proof technique implies fast convergence to socially good (approximate) solutions. This is in contrast to the negative result of Fabrikant, Papadimitriou, and Talwar [2] showing the existence of exponentially long best-response paths.

(ii) *Valid-Utility Games*: we show that for valid-utility games the price of sinking is at most $n+1$; thus the worst case price of sinking in a valid-utility game is between n and $n+1$. We use our proof to show fast convergence to constant factor approximate solutions in basic-utility games.

*Massachusetts Institute of Technology. Email: `goemans@math.mit.edu`

†Massachusetts Institute of Technology. Email: `mirrokni@theory.csail.mit.edu`

‡McGill University. Email: `vetta@math.mcgill.ca`

In addition, we present a hardness result which shows that, in general, there might be states that are exponentially far from any sink equilibrium in valid-utility games. We prove this by showing that the problem of finding a sink equilibrium (or a PSNE) in valid-utility games is PLS-complete.

1. Introduction

A standard approach in analysing the performance of systems controlled by non-cooperative agents is by the examination of Nash equilibria. Of particular interest is the *price of anarchy*[1] in a game [8]. This gives one measure of the cost to society of the inherent *lack of coordination* in a game. There are, however, several drawbacks in the use of Nash equilibria. For example, one issue relates to use of non-randomized (pure) and randomized (mixed) strategies. Often pure strategy Nash equilibria may not exist, yet the use of a randomized (mixed) strategy is unrealistic in many games. This necessitates the need for an alternative solution concept in evaluating such games. Another issue arises from the observation that Nash equilibria represent "stable" points in a system. Therefore (even if pure Nash equilibria exist), they are a more acceptable solution concept if it is likely that the system does converge to such stable points. In particular, the use of Nash equilibria seems more valid in games in which Nash equilibria arise when agents iteratively engage in selfish behaviour. However, in many games it is not the case that repeated selfish behaviour always leads to Nash equilibria. In these games it also seems that another measure of the cost of the lack of coordination would be useful. Observe that these issues are particularly important in games in which the use of pure strategies and repeated moves are the norm, for example, auctions. We remark that for many practical games these properties are the rule rather than the exception (and this observation motivates much of the work in this paper). For these games, then, it is not sufficient to just study the value of the social function at Nash equilibria.

In this paper we introduce a new solution concept in a game, namely *sink equilibria*. We model the behaviour of agents using a graph, called the state graph (or strategy profile graph) whose vertex set is the set of strategy states (or strategy profiles). We assume that evolution of the game over time can be described by walks on this graph. Here, we also assume that the only arcs of the state graph are arcs that correspond to moves of the players that may occur with non-zero probability. Thus, solutions or stable outcomes will be given by the long-run behaviour of such random walks. In particular, eventually these walks must lead to a set of states that have the following two properties:

- These states form a strongly connected component in the state graph.
- The strongly connected component has no outgoing arcs in the state graph.

These strongly connected components are sink equilibria. They are stable in that once we reach such a component we will never leave it. They include PSNE as a special case, but unlike PSNE they are guaranteed to exist in all such games. As with Nash equilibria, we can use sink equilibria to measure the cost to society of the lack of coordination. In particular, here we will consider an analogue of the price of anarchy termed the *price of sinking*. This is the worst case ratio of the social value of a sink equilibrium compared to the optimal social solution. The social value of a sink equilibrium is measured by the expected value of the stationary distribution of a random walk on the states in the sink.

[1]The price of anarchy is the worst case ratio between the social value of an optimal solution and a Nash equilibrium.

We formally define the price of sinking in Section 2. For any game the arc set and their associated probabilities in the state profile graph may vary dramatically. As mentioned, we will focus on perhaps the simplest case: the best response graph associated with myopic players. Here, the arc set consists only of those arcs that correspond to a best response move of some player. We will also assume that, at a given state, each player is equally likely to be selected to move. Thus our random walk will be a uniform random walk on the best response graph. We call sink equilibria in such graphs *myopic sink equilibria*, and refer to the *price of sinking myopically*. We will omit the "myopic" term when the context is clear. We remark that the assumption of myopic behaviour is very restrictive and unrealistic in many situations. Consequently, further investigation into the general case is important. This would allow for an examination into alternate behaviours such as non-myopic behaviour, long-term planning, and simultaneous moves. We content ourselves, here, with considering the basic case of myopic behaviour with non-simultaneous moves for several reasons though. Firstly, it allows us to introduce sink equilibria in a clear manner, without having to deal with the complexities (both practical and game-theoretic) of alternative behaviours. For example, given a game how do you justify non-uniform moves, realistically incorporate forward planning, or assign probabilities to simultaneous moves etc. Moreover, even finding simple, realistic examples of games with non-myopic behaviours is not a straight-forward task. In addition, mathematically there appears to be no intrinsic additional difficulty in tackling the general case, and so the ideas and techniques presented here should also be useful in examining games with non-myopic behaviours.

We illustrate the usefulness of our measure in Section 3 where we present an n-agent valid-utility game which always converges to states with social value a factor n worse than optimal. Indeed, the price of sinking for this game is n. However the price of anarchy is almost 1. Thus, the price of anarchy gives us a misleading confidence in the social quality of an outcome that will result from selfish behaviour.

As well as being perhaps a more appropriate solution concept than PSNE in many games, the existence of sink equilibria has several nice implications. Since sink equilibria always exist, the price of sinking can always be calculated[2] even in games without PSNE. Unlike PSNE, sink equilibria also possess natural convergence properties. In particular, the techniques used to bound the price of sinking may often also give bounds on the speed of convergence of random walks to sink equilibria and/or approximate solutions. We study two examples in Section 4:

(1) *Unsplittable Selfish Routing (and Weighted Congestion Games).* We present instances of the weighted unsplittable flow version of the selfish routing problem that possess no PSNE. However, we show that, for polynomial latency functions of degree at most d, the price of sinking is $O(2^{2d}d^{2d+3})$. In addition, our proof technique implies fast convergence to good (approximate) solutions. This may be compared to the negative result by Fabrikant, Papadimitriou, and Talwar [2] showing the existence of exponentially long best-response paths to PSNE. For example, consider the case of linear latency functions. Here, it is known that PSNE exist [4], but it may be the case that the number of best response moves needed for convergence to a PSNE is exponential. Our results show that after a small number of random best response moves the social value of the flow is within a constant factor of the optimal solution.

[2]Of course, actually doing so may not be easy!

(2) *Valid-Utility Games.* Our second example concerns the class of valid-utility games; specific example in this class include marking sharing games [5], caching games [3], traffic routing games, facility location games, and multiple item auctions [14]. Here we show that the price of sinking is at most $n+1$; thus the worst case price of sinking in a valid-utility game is between n and $n+1$. Again, our methods signify fast convergence to approximate solutions. In particular, for basic-utility games, the expected social value of any state after $n \log n$ random best response moves is at least half of optimum.

We also present a hardness result concerning sink equilibria. In section 5 we show that in general it is a PLS-complete problem to find a sink equilibria (or PSNE) in valid-utility games. This implies the existence of exponentially long best response paths to any sink equilibrium in some valid-utility games.

We conclude this introduction with a very brief discussion on related work. In order to deal with the stability and convergence problems of Nash equilibria, equilibrium concepts other than Nash equilibria have been studied in the economics literature. Among these concepts are stable equilibria [7], stochastic adjustment models [6], iterative elimination of dominated strategies, the set of undominated strategies etc. Convergence and strategic stability of equilibria in evolutionary game theory is a also central subject of study for many economists. However, in their studies the most important factor is typically the stability of equilibria, and not measurements of the social value of equilibria. In [9], we began our investigation into games in which pure strategy moves are the norm.

2. Sink Equilibria

A *strategic game* $\mathcal{G}$ is defined as a tuple $\mathcal{G}(U, \{F_i | i \in U\}, \{\alpha_i() | i \in U\})$ where (i) U is the set of n players or agents, (ii) F_i is a family of feasible *(pure) strategies* or *actions* for player i and (iii) $\alpha_i : \Pi_{i \in U} F_i \rightarrow \mathbb{R}^+ \cup \{0\}$ is the (private) *payoff* or *utility* function for agent i, given the set of strategies of all players. Player i's strategy is denoted by $s_i \in F_i$, and we let $\mathcal{F} := \Pi_{i \in U} F_i$ be the set of all possible strategy profiles. In the games we consider, there will be a *social utility function*, usually denoted by $\gamma : \Pi_{i \in U} F_i \rightarrow \mathbb{R}$, defined on all strategy profiles in a strategic game. The social value of the optimal solution is denoted by OPT. Our main focus is on the social quality of outcomes produced by selfish agents.

A *strategy profile* or a (strategy) *state*, denoted by $S = (s_1, s_2, \ldots, s_n)$, is the collection of strategies chosen by the players. We let $S \oplus s'_i := (s_1, \ldots, s_{i-1}, s'_i, s_{i+1}, \ldots, s_k)$, that is, the strategy profile obtained from S if agent i changes its strategy from s_i to s'_i. In order to model the selfish behavior of players, we use the underlying *strategy profile graph* or *state graph.* Each vertex in the state graph represents a state $S = (s_1, s_2, \ldots, s_n)$. As noted, in this paper the arcs in the state graph will correspond to best-response moves by the players. Hence we have, for each player i an arc from S to $S \oplus \hat{s}_i$, where $\hat{s}_i$ is the best response of agent i at state S. (This model can be justified in extensive games with complete information, and is used in the economics literature extensively in the context of studying convergence in games.) In many games with iterative moves, the evolution of game-play may then be naturally modeled by a path in the state graph. Such a path may or may not converge to a pure strategy Nash equilibrium (PSNE); a PSNE of a strategic game is a strategy profile in which each player plays mutual best responses (that is, a vertex in the state graph for which the best response move of each agent corresponds to a self-loop). Clearly it may be the case that there are no PSNE. So we may ask what happens in such games. Specifically, does some concept of stability or equilibrium exist? The answer is yes, and we now describe such an "equilibrium".

Consider the strongly connected components of the state graph. If we contract the strongly connected components to singletons then we obtain an acyclic graph. The sink nodes in this graph (nodes with out-degree equal to zero) correspond to strongly connected components with no out-going arcs in the state graph. We call such a strongly connected component a *(myopic) sink equilibrium*. The reason for this terminology is clear: if a best-response walk ever reaches a node in a sink equilibrium then it will never leave that set of nodes. In addition, a long enough random walk in the state graph will converge to a sink equilibrium with probability arbitrarily close to 1.

We denote by $\mathcal{Q}$ the set of sink equilibria in a game. We remark that the union of states in sink equilibria correspond to the set of recurrent states in a Markov chain that only has non-zero transitional probabilities on arcs in the state graph. In a random sequence of best responses of agents, we independently choose an agent uniformly at random at each step and let this agent play its best response (if the agent has more than one best-response move, we may assume that the agent arbitrarily chooses a move from the collection of best-response moves). When this walk reaches a state in some sink we then follow a *random walk* over the states in that sink. For a sink $Q \in \mathcal{Q}$, let $\pi_Q : Q \rightarrow \mathbb{R}^+ \cup \{0\}$ be the steady state distribution of the random walk over states in Q. Let $\gamma(S)$ measure the social value of a state S. The (expected) social value of a sink equilibrium $Q \in \mathcal{Q}$, denoted by $\Gamma(Q)$, is the expected social value of states given by the steady distribution of the random walk over the states of Q, i.e., $\Gamma(Q) = \sum_{S \in Q} \pi_Q(S)\gamma(S)$ We then define, the *price of sinking (myopically)* for a maximization social function as

$$\text{Price of Sinking} = \frac{\text{OPT}}{\min\limits_{Q \in \mathcal{Q}} \Gamma(Q)} = \frac{\text{OPT}}{\min\limits_{Q \in \mathcal{Q}} \sum_{S \in Q} \pi_Q(S)\gamma(S)}$$

In other words, the price of sinking is the worst ratio between the expected social value of a sink equilibrium and the social value of the optimum. Similarly, the price of sinking for a minimization problem is $\max\limits_{Q \in \mathcal{Q}} \Gamma(Q)/\text{OPT}$. Moreover, we have an analogous definitions for the price of sinking for general strategy profile graphs with alternate arc sets. Given that sink equilibria are stable solutions in such games, this may be a more realistic measure of the cost of the lack of coordination than the price of anarchy.

3. Price of Sinking vs. Price of Anarchy

In this section, we present an n-agent (valid-utility) game in which the price of sinking and the price of anarchy give very different pictures as to the consequences of non-cooperative behavior. In particular, the price of anarchy will be close to 1, suggesting that no form of mechanism design is required to enforce socially good solutions. However, every possible outcome of the game will result in a solution whose value is a factor n smaller than that of the optimal social solution. The collection of strategies (groundset) available to of agent i is $\{y_i, x_i^1, x_i^2, \ldots, x_i^n\}$, where $i = 0, 1, \ldots, n-1$. For motivation, we can think of strategy y_i as a socially responsible strategy for agent i. In contrast, all the strategies $\{x_i^1, x_i^2, \ldots, x_i^n\}$ can be viewed as socially irresponsible strategies. Moreover, we will see that in any situation one of these n irresponsible strategies provides a better payoff for agent i than acting responsibly. Consequently, there is an incentive for every agent to act anti-socially with extreme consequences for the social outcome. In contrast, the price of anarchy is oblivious to this incentive for anti-social behavior. The reason being that the payoffs to each agent are intrinsically linked to the behavior of the other agents. Any specific irresponsible strategy may be beneficial in certain circumstances but typically (given the other agents responses) that specific strategy has smaller payoff than the responsible strategy. Consequently, under randomized strategies, playing an irresponsible strategy is likely to lead to low private returns. Thus mixed strategy Nash equilibria will require that most agents behave responsibly, blissfully ignoring the fact that in *every* possible situation *each* agent has an incentive to behave irresponsible.

The family of feasible strategies F_i for each agent i is the set of singletons of his ground set and the empty set, i.e., $F_i = \{s \subseteq V_i : |s| \leq 1\}$. Let $X_i = \{x_i^1, x_i^2, \ldots, x_i^n\}$ and $X = \cup_i X_i$. Let $S = (s_1, s_2, \ldots, s_n)$ be a collection of subsets $s_i \subseteq V_i$ for all $i = 0, 1, \ldots, n-1$. For a collection $S = (s_1, \ldots, s_n)$, we let $S^{\mathsf{U}} = \cup_{i \in U} s_i$. We construct a non-decreasing, submodular social utility function γ on $\Pi_{i \in U} V_i$ in the following manner.

$$\gamma(S) = \begin{cases} |S^{\mathsf{U}} \backslash X| & \text{if } S^{\mathsf{U}} \cap X = \emptyset \\ |S^{\mathsf{U}} \backslash X| + 2 & \text{otherwise} \end{cases}$$

We now need to specify the private utilities of each agent at any state. In order to define the payoff functions, we define a function $i^*(S)$ for each strategy profile S. We set $i^*(S) = null$ for any strategy profile S in which no player plays an irresponsible strategy. If in a strategy profile S, some players play irresponsibly, $i^*(S)$ is the index of one of the players who plays irresponsibly. In addition, we would like $i^*(S)$ to satisfy the following property: given the strategies of the other agents, any agent i can always choose some irresponsible strategy which forces $i^*(S) = i$. Clearly, this will give agents an incentive to act irresponsibly when using pure strategies. In order to complete the description of the function i^*, let $\chi_{ij}(S)$ be the indicator variable for the event that agent i plays the irresponsible strategy x_i^j. That is

$$\chi_{ij}(S) = \begin{cases} 1 & \text{if } x_i^j \in S^{\mathsf{U}} \\ 0 & \text{otherwise.} \end{cases}$$

Next let

$$i^*(S) = \begin{cases} null & \text{if } S^{\mathsf{U}} \cap X = \emptyset \text{ (No-one plays irresponsibly)} \\ i_l & \text{if } \cup_i (S^{\mathsf{U}} \cap X_i) \neq \emptyset \textit{ and} \\ & l = [\sum_{i \in U} (\sum_{j=1}^n j \cdot \chi_{ij}(S)) \mod k] \end{cases}$$

Observe that if $i^*(S) = null$ then i can play the irresponsible strategy $s_i' = \{x_i^i\}$, thus forcing $i^*(S \oplus s_i') = i$. Moreover, there always exists a strategy $s_i' = \{x_i^p\}$ such that if i plays $s_i' = \{x_i^p\}$ then $i^*(S \oplus s_i') = i$. We are now ready to give a payoff function α_i for each agent i.

$$\alpha_i(S) = \begin{cases} 0 & \text{if } y_i \notin s_i \text{ and } i \neq i^*(S) \\ 1 & \text{if } y_i \in s_i \text{ and } i \neq i^*(S) \\ 2 & \text{if } y_i \notin s_i \text{ and } i = i^*(S) \\ 3 & \text{if } y_i \in s_i \text{ and } i = i^*(S). \end{cases}$$

So agent i gets utility 1 for playing the responsible strategy and another 2 units of utility if $i = i^*(S)$. We will see in Section 4.2 that this is a valid-utility game with a non-decreasing social utility function. Thus we may apply the following result from [14].

Theorem 3.1. *A valid-utility game with a non-decreasing social utility function has a price of anarchy at most* 2. □

If fact, it is easy to see that the price of anarchy in this game actually tends to 1 as the number of agents increases. In particular, a socially optimal solution has $n-1$ of the agents playing their responsible strategies and exactly one of the agents plays an irresponsible strategy. Such an outcome has value $n+1$. Moreover, note that by playing responsibly an agent can guarantee that they receive 1 unit of utility. Thus, it must be the case that in a Nash equilibrium[3] every agent has an expected payoff of at least 1. Since $\gamma(S) \geq \sum_{i \in U} \alpha_i(S)$ for any state S, we have that the expected social value of a Nash equilibrium is at least n. Thus the price of anarchy is at most $1 + \frac{1}{n}$.

Now we consider the price of sinking in this game. Given any strategy profile S, the best response of each agent is to play the specific irresponsible strategy that gives it a payoff of 2. To see this, note that agent i always has a move that sets $i^*(S') = i$. Thus a responsible strategy y_i is never a best-response strategy. In fact, the best response of each player is to play an irresponsible strategy to get the payoff of 2, thus forcing to the payoffs of the other players using irresponsible strategies to 0. It follows that there is a unique sink equilibrium consisting of every strategy profile in which each agent plays an irresponsible strategy. Thus, every state in the sink has social value exactly two. Hence the price of sinking is exactly $\frac{n+1}{2}$. We remark that even if we start at an optimal solution and then allow each agent to make just one single best-response move in turn then we end up with a solution of value 2! Moreover, we can then never leave this sink if players play their myopic best responses.

Notice also that we could alter the payoffs in the game slightly so that the payoff resulting from the first irresponsible move is $1+\delta$ rather than 2. Clearly the price of sinking is then $\frac{n+\delta}{1+\delta}$ whilst the price of anarchy is $1 + \frac{\delta}{n}$. Thus we have

Lemma 3.2. *There are valid-utility games, with non-decreasing social utility functions, having a price of sinking of almost n and a price of anarchy of almost* 1. □

Consequently the price of anarchy underestimates the social cost of the lack of coordination by a factor n. The reason for this is that the good strategy always gives a good return. Any bad strategy can give a high return but only in a small number of situations, thus any bad strategy performs badly against randomized strategies and players tend to play their good strategies in a mixed Nash equilibria. This type of issue often arises in games, and explains why the price of anarchy may often significantly under-estimate the social cost of the lack of coordination in such games.

Finally, note that this game has no PSNE so focusing here upon sink equilibria is essential. Surprisingly, Lemma 3.2 is also almost tight; we will show in Section 4 that the price of sinking in a valid-utility game is at most $n+1$.

---

[3]One Nash equilibrium is the following. Each agent i plays strategy y_i with probability p and each bad strategy with probability $\frac{1-p}{n}$. It is easy to check that letting $p = \sqrt[n-1]{\frac{1}{2}(1 - \frac{1}{n-1})}$ gives a Nash equilibrium.

4. Price of Sinking and Convergence

Recall that PSNE are special cases of sink equilibria. We have already seen that games in which agents repeatedly react to the other agent's strategies via the use of pure strategy best responses will converge to sink equilibria and not necessarily to PSNE. Moreover, many classes of games have instances for which no PSNE exists. In these games, we can still measure the cost to society of the lack of coordination using the price of sinking. Moreover, in bounding the price of sinking for sink equilibria we may obtain bounds on the expected social value of states after a random sequence of best responses.

4.1. **Unsplittable Selfish Routing and Weighted Congestion Games.** Consider the "unsplittable flow" version of the selfish routing game. We have a directed network $G = (V, E)$ with a flow dependent latency function $\lambda_e : \mathbb{R} \to \mathbb{R}^+ \cup \{0\}$ on each arc $e \in E$. There is a set U of n agents; agent i wishes to route flow at a rate r_i from a source s_i to a sink t_i. Each agent aims to incur as small a latency as possible. In the unsplittable flow version, an agent may not split its flow. Hence each agent picks a unique $s_i - t_i$ path and routes all its flow along the path. The latency of an agent is equal to its traffic size multiplied by the sum of the latencies of arcs along the path that it chooses. The latency of an arc e is a non-decreasing and non-negative function of the total load on arc e. In this paper, we consider bounded-degree polynomial latency functions. In particular, for an arc e, we let $\lambda_e(x) = \sum_{0 \le j \le d} a_{e,j} x^j$ be a non-negative and non-decreasing delay function for arc e. For a strategy profile $\mathcal{P} = (P_1, P_2, \dots, P_n)$ where P_i is a $s_i - t_i$ path, let the load of arc e be $f_e = \sum_{i: e \in P_i} r_i$. Then, the latency of agent i is $l_i(f) = r_i \sum_{e \in P_i} \lambda_e(f_e)$ and the total latency of flow f is $l(f) = \sum_{i \in U} l_i(f) = \sum_{e \in E(G)} \lambda_e(f_e) f_e$.

Recently Awerbuch, Azar, and Epstein [1] proved that the price of anarchy in such games is exactly 2.618 for linear latency functions and is at most $O(2^d d^{d+1})$ for polynomial latency functions of degree at most d. They extended their results to mixed Nash equilibria, since the existence of pure Nash equilibria for these games with polynomial latency functions was not known. For linear latency function Fotakis, Kontogiannis, and Spirakis [4] proved that the game is a potential game. Here, we exhibit an instance of this game with quadratic latency functions that does not possess any PSNE. This, in turn, provides additional motivation for analyzing the price of sinking in these games. Our example is shown in Figure 1. It depicts a network with 4 vertices and 6 arcs. Arcs are labeled from 1 to 6. The latency functions of arcs are $\lambda_1(x) = x + 33$, $\lambda_2(x) = 13x$, $\lambda_3(x) = 3x^2$, $\lambda_4(x) = 6x^2$, $\lambda_5(x) = x^2 + 44$, and $\lambda_6(x) = 47x$. There are two agents with traffic $r_1 = 1$ and $r_2 = 2$. The source of both agents is vertex 1 ($s_1 = s_2 = 1$) and the destination of both agents is vertex 4 ($t_1 = t_2 = 4$). There are four source-destination paths: $P_1 = (6)$, $P_2 = (3, 5)$, $P_3 = (3, 4, 2)$, and $P_4 = (1, 2)$ where the numbers within the parentheses are the labels of arcs on the path. It is not hard to check that the weighted unsplittable selfish routing game on this network has no PSNE. There is one sink equilibrium, namely the set of strategy profiles $\{(P_1, P_2), (P_3, P_2), (P_3, P_4), (P_1, P_4)\}$.

The key to obtaining bounds on the price of sinking is that any agent making a best-response move cannot cause too much cumulative harm to the other agents. Consequently, if an agent can make a move that significantly increases its private welfare, then the overall social welfare must rise. This will be an important factor in allowing us to prove that we have a low price of sinking in these routing games.

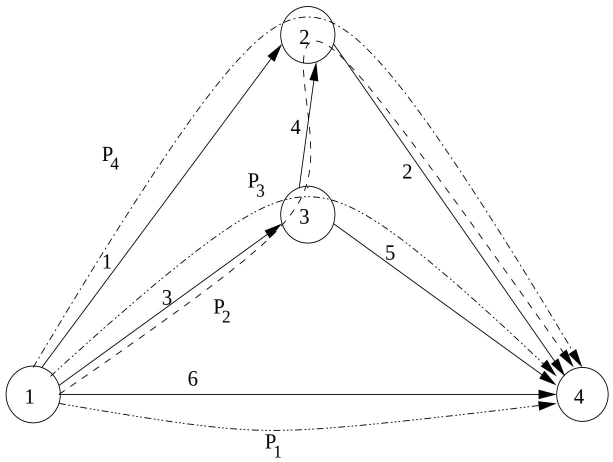

FIGURE 1. A routing game without PSNE.

Theorem 4.1. *The price of sinking for a weighted unsplittable selfish routing game (or a weighted congestion game) is at most* $O(2^{2d}d^{2d+3})$.

Proof. We need the following three lemmas for the proof.

Lemma 4.2. *Let f be the flow corresponding to the current strategy profile $\mathcal{P} = (P_1, \ldots, P_n)$. Suppose agent i changes its flow path from P_i to P_i', to give a new flow f_i'. Then $l(f_i') \leq l(f) + (d+1)l_i(f_i') - l_i(f)$. In particular, if agent i decreases its latency by changing to P_i', then $l(f_i') \leq l(f) + dl_i(f) \leq (d+1)l(f)$.*

Proof. The latency incurred by agent i is then

$$l_i(f_i') = r_i \sum_{e \in P_i'} \sum_{0 \leq j \leq d} a_{e,j}(f_{i,e}')^j = r_i \left(\sum_{e \in P_i' \cap P_i} \sum_{0 \leq j \leq d} a_{e,j} f_e^j + \sum_{e \in P_i' - P_i} \sum_{0 \leq j \leq d} a_{e,j}(f_e + r_i)^j \right).$$

Note that for $e \in P_i' - P_i$, we have $f_{i,e}' = f_e + r_i$. Moreover, we know that

$$l(f_i') \leq l(f) + (l_i(f_i') - l_i(f)) + \sum_{e \in P_i' - P_i} \left(\sum_{0 \leq j \leq d} (a_{e,j} {f_{i,e}'}^j) - (a_{e,j} f_e^j) \right) (f_{i,e}' - r_i),$$

the last term corresponding to the increase in latency for agents other than i due to the rerouting of agent i. We can get an upper bound on the increase in latencies faced by the other agents by noting that

$$\begin{aligned}
&\sum_{e\in P_i'-P_i}\left(\sum_{0\le j\le d}(a_{e,j}{f_{i,e}'}^{j})-(a_{e,j}f_e^j)\right)(f_{i,e}'-r_i)\\
&= \sum_{e\in P_i'-P_i}\sum_{0\le j\le d}(a_{e,j}({f_{i,e}'}^{j}-f_e^j)f_e)\\
&= \sum_{e\in P_i'-P_i}\left(\sum_{0\le j\le d}a_{e,j}(f_{i,e}'-f_e)\left(\sum_{1\le t\le j}{f_{i,e}'}^{j-t}f_e^{t-1}\right)f_e\right)\\
&< \sum_{e\in P_i'-P_i}\left(\sum_{0\le j\le d}a_{e,j}r_i\left(\sum_{1\le t\le j}(f_e+r_i)^{j-1}\right)(f_e+r_i)\right)\\
&\le r_i\sum_{e\in P_i'-P_i}\left(\sum_{0\le j\le d}ja_{e,j}(f_e+r_i)^j\right)\\
&\le dl_i(f_i').
\end{aligned}$$

Thus, the total latency after agent i changes its strategy is at most $l(f)+(d+1)l_i(f_i')-l_i(f)$. Since, $l_i(f_i')\le l_i(f)$, this shows that $l(f_i')\le l(f)+dl_i(f)\le(d+1)l(f)$. □

Lemma 4.3. *Let f be the flow corresponding to the current strategy profile. Consider the following random process: choose an agent i at random and let it play its best response. If f' is the new flow after this change, then* $\mathbf{E}[l(f')|f]\le(1+\frac{d}{n})l(f)$.

Proof. Let f_i' be the flow after agent i plays its best response to f. Then, using Lemma 4.2, we have:

$$\begin{aligned}
\mathbf{E}[l(f')|f] &= \frac{1}{n}\sum_{i\in U}l(f_i')\\
&\le \frac{1}{n}\sum_{i\in U}(l(f)+dl_i(f))\\
&= \frac{1}{n}(nl(f)+dl(f))\\
&= (1+\frac{d}{n})l(f).
\end{aligned}$$

□

The third lemma we need is below. Its proof is inspired by the work of Azar et al. [1].

Lemma 4.4. *Let f be the flow corresponding to the current strategy profile. Consider the following random process: choose an agent i at random and let it play its best response. If f' is the new flow after this change, then either* $\mathbf{E}[l(f')|f]\le(1-\frac{1}{2n})l(f)$, *or* $l(f)\le O(2^{2d}(d+1)^{2d+2})\text{OPT}$.

Proof. Assume that the best response of agent i is to switch from path P_i to P_i' resulting in the flow f_i'. Thus, $\mathbf{E}[l(f')|f]=\frac{1}{n}\sum_{i\in U}l(f_i')$. We consider the following two cases:

Case 1: $\sum_{i\in U} 2(d+1)l_i(f_i') \le \sum_{i\in U} l_i(f)$. In this case, by Lemma 4.2,

$$\begin{aligned}
\mathbf{E}[l(f')|f] &= \frac{1}{n}\sum_{i\in U} l(f_i') \\
&\le \frac{1}{n}\sum_{i\in U}\big(l(f) + (d+1)l_i(f_i') - l_i(f)\big) \\
&\le \frac{1}{n}\left(\sum_{i\in U} l(f) + \sum_{i\in U}\frac{1}{2}l_i(f) - \sum_{i\in U} l_i(f)\right) \\
&= \frac{1}{n}(nl(f) - \frac{1}{2}l(f)) \\
&= (1 - \frac{1}{2n})l(f).
\end{aligned}$$

Thus, we obtain $\mathbf{E}[l(f')|f] \le (1-\frac{1}{2n})l(f)$.

Case 2: $\sum_{i\in U} 2(d+1)l_i(f_i') > \sum_{i\in U} l_i(f)$. Let $\mathcal{P}^* = (P_1^*, \dots, P_n^*)$ be the optimal solution and let f^* be the flow corresponding to $\mathcal{P}^*$. Set $J^*(e) = \{i : e \in P_i^*\}$. Let f_i^* be the flow resulting from the switch of agent i from P_i to P_i^*. Since P_i' is i's best response, we have $l_i(f_i^*) \ge l_i(f_i')$. Thus, in this case, $\sum_{i\in U} 2(d+1)l_i(f_i^*) \ge \sum_{i\in U} l_i(f) = l(f)$. Consequently,

$$\begin{aligned}
l(f) &\le \sum_{i\in U} 2(d+1)l_i(f_i^*) \\
&\le (2d+2)\sum_{i\in U} r_i \sum_{e\in P_i^*} \lambda_e(f_e + r_i) \\
&= (2d+2)\sum_{i\in U} r_i \sum_{e\in P_i^*}\sum_{j=0}^{d} a_{e,j}(f_e+r_i)^j \\
&= (2d+2)\sum_{e}\sum_{j=0}^{d}\sum_{i\in J^*(e)} a_{e,j}(f_e+r_i)^j r_i.
\end{aligned}$$

The rest of the proof of this case is based on the proof of Lemmas A1, A2, and A3 in [1]. First, we use the following inequality from [1]: $(x+y)^d \leq cx^d + (y(\frac{d}{\ln c}+1))^d$ for any $c>1$. Thus, we get:

$$\begin{aligned}
l(f) &\leq (2d+2)\sum_e \sum_{j=0}^{d} \sum_{i\in J^*(e)} a_{e,j}(f_e+r_i)^j r_i \\
&\leq (2d+2)\sum_e \sum_{j=0}^{d} a_{e,j} \sum_{i\in J^*(e)} \left(cf_e^j r_i + \left(\frac{j}{\ln c}+1\right)^j r_i^{j+1}\right) \\
&\leq (2d+2)\sum_e \sum_{j=0}^{d} a_{e,j} \left(cf_e^j f_e^* + \left(\frac{d}{\ln c}+1\right)^d f_e^{*j+1}\right) \\
&= c(2d+2)\sum_e \sum_{j=0}^{d} a_{e,j} f_e^j f_e^* + (2d+2)\left(\frac{d}{\ln c}+1\right)^d \sum_e \sum_{j=0}^{d} a_{e,j} f_e^{*j+1} \\
&= c(2d+2)\sum_e \sum_{j=0}^{d} a_{e,j} f_e^j f_e^* + (2d+2)\left(\frac{d}{\ln c}+1\right)^d \sum_e \lambda_e(f_e^*) f_e^*
\end{aligned}$$

where the second inequality comes from the fact that $\sum_{i\in J^*(e)} r_i^d \leq f_e^{*d}$ and the function $f(x) = (\frac{x}{\ln c}+1)^x$ is an increasing function for $x \geq 0$. Holder's inequality states:

$$\sum_j a_j^\alpha b_j^{1-\alpha} \leq \left(\sum_j a_j\right)^\alpha \left(\sum_j b_j\right)^{1-\alpha}.$$

Applying this, with $a_j = a_{e,j} f_e^{j+1}$, $b_j = a_{e,j} f_e^{*j+1}$, $\alpha = \frac{j}{j+1}$, yields

$$
\begin{aligned}
l(f) \le c(2d+2) \sum_e \sum_{j=0}^{d} a_{e,j} f_e^j f_e^* + (2d+2)\left(\frac{d}{\ln c}+1\right)^d \sum_e \lambda_e(f_e^*) f_e^* \\
\le \quad & 2c(d+1) \sum_{j=0}^{d} \left(\sum_e a_{e,j} f_e^{j+1}\right)^{j/(j+1)} \left(\sum_e a_{e,j} f_e^{*j+1}\right)^{1/(j+1)} \\
& +2(d+1)\left(\frac{d}{\ln c}+1\right)^d \sum_e \lambda_e(f_e^*) f_e^* \\
\le \quad & 2c(d+1) \sum_{j=0}^{d} \left(\sum_e \lambda_e(f_e) f_e\right)^{j/(j+1)} \left(\sum_e \lambda_e(f_e^*) f_e^*\right)^{1/(j+1)} \\
& +2(d+1)\left(\frac{d}{\ln c}+1\right)^d \sum_e \lambda_e(f_e^*) f_e^* \\
\le \quad & 2c(d+1) \sum_{j=0}^{d} \left(\sum_e \lambda_e(f_e) f_e\right)^{d/(d+1)} \left(\sum_e \lambda_e(f_e^*) f_e^*\right)^{1/(d+1)} \\
& +2(d+1)\left(\frac{d}{\ln c}+1\right)^d \sum_e \lambda_e(f_e^*) f_e^* \\
\le \quad & 2c(d+1)^2 \left(\sum_e \lambda_e(f_e) f_e\right)^{d/(d+1)} \left(\sum_e \lambda_e(f_e^*) f_e^*\right)^{1/(d+1)} \\
& +2(d+1)\left(\frac{d}{\ln c}+1\right)^d \sum_e \lambda_e(f_e^*) f_e^*
\end{aligned}
$$

where the fourth inequality is from the inequality $x^{\alpha} y^{1-\alpha} \ge x^{\alpha'} y^{1-\alpha'}$ for $x \ge y > 0$ and $1 \ge \alpha \ge \alpha' \ge 0$ with $x = \sum_e \lambda_e(f_e) f_e$ and $y = \sum_e \lambda_e(f_e^*) f_e^*$. By letting

$$x = \frac{l(f)^{\frac{1}{d+1}}}{\mathrm{OPT}^{\frac{1}{d+1}}},$$

we get

$$x^{d+1} \le 2c(d+1)^2 x^d + 2(d+1)\left(\frac{d}{\ln c}+1\right)^d.$$

After dividing both sides by x^d, we get:

$$x \le 2c(d+1)^2 + 2(d+1)\left(\frac{\frac{d}{\ln c}+1}{x}\right)^d.$$

We claim that if we set $c = 2 - \epsilon$ for $\epsilon = \frac{1}{d+1}\left(\frac{1}{2(d+1)}\right)^d$, then we have $x \le 4(d+1)^2$. Assume for contradiction that $x > 4(d+1)^2$. Then,

$$4(d+1)^2 < x \le 4(d+1)^2 - 2\epsilon(d+1)^2 + 2(d+1)\left(\frac{\frac{d}{\ln c}+1}{x}\right)^d.$$

Thus,

$$\begin{aligned}(d+1)\epsilon &< \left(\frac{\frac{d}{\ln c}+1}{x}\right)^d \\ &\le \left(\frac{2d+1}{x}\right)^d \qquad\qquad [\text{since } \ln c > 0.5] \\ &< \left(\frac{2d+1}{4(d+1)^2}\right)^d \\ &< \left(\frac{1}{2(d+1)}\right)^d \\ &= (d+1)\epsilon\end{aligned}$$

which is a contradiction. Therefore, by setting $c = 2-\epsilon$, we get $x \le 4(d+1)^2$. Hence, $l(f) = x^{d+1}\text{OPT} \le O(2^{2d}(d+1)^{2d+2})\text{OPT}$. □

From Lemma 4.4, we can bound the price of sinking as follows. Consider a sink Q. Let f_0 be a flow in Q. Consider a random walk starting from f_0 in which we let a random agent play his best response at each step. Let $f_0, f_1, f_2, \ldots, f_N$ be a sequence of observed flows in Q. Recall that the value for sink Q is equal to $\Gamma(Q) = \sum_{S\in Q}\pi_Q(S)l(f_S)$ where f_S is the flow corresponding to the state S and π_Q is the steady distribution for the random walk on Q. Since Q is strongly connected, this is equal to $\Gamma(Q) = \lim_{N\to\infty}\frac{\sum_{0\le j\le N}\mathbf{E}[l(f_j)]}{N}$. In order to upper bound this value, it is sufficient to upper bound $\mathbf{E}[l(f_j)]$ for each $0 \le j \le N$. Lemma 4.4 shows that there exists a state in any sink Q with total latency less than $O(2^{2d}(d+1)^{2d+2})\text{OPT}$. Note that, as Q is strongly connected the value of the sink is independent of the choice of f_0. Therefore, we can set f_0 such that $l(f_0) \le c' 2^{2d}(d+1)^{2d+2}\text{OPT}$. Let c_i be the coin toss of step i in the random walk. More precisely, we want to upper bound $a_j = E_{c_1,c_2,\ldots,c_j}[l(f_j)]$. By Lemma 4.4 and Lemma 4.3, we have

- Either $E_{c_{j+1}}[l(f_{j+1})|f_j] \le (1-\frac{1}{2n})l(f_j)$ or $l(f_j) \le c'2^{2d}(d+1)^{2d+2}\text{OPT}$.
- $E_{c_{j+1}}[l(f_{j+1})|f_j] \le (1+\frac{d}{n})l(f_j)$

Let E_1 be the event that $l(f_j) \le c'2^{2d}(d+1)^{2d+2}$ and E_2 be the event that $l(f_j) > c'2^{2d}(d+1)^{2d+2})\text{OPT}$. Let p be the probability that event E_2 happens. Furthermore, let $Y = \mathbf{E}[l(f_j)|E_1] \le c'2^{2d}(d+1)^{2d+2}$ and $X = \mathbf{E}[l(f_j)|E_2]$. Thus, $a_j = \mathbf{E}[l(f_j)] = pX + (1-p)Y$. Now,

$$\begin{aligned}a_{j+1} &= \mathbf{E}[l(f_{j+1})] \\ &\le p\left(1-\frac{1}{2n}\right)X + (1-p)\left(1+\frac{d}{n}\right)Y \\ &\le \left(1-\frac{1}{2n}\right)(pX+(1-p)Y) + \frac{2d+1}{2n}Y \\ &\le \left(1-\frac{1}{2n}\right)a_j + \frac{2d+1}{2n}Y \\ &\le \left(1-\frac{1}{2n}\right)a_j + \frac{2d+1}{2n}c'2^{2d}(d+1)^{2d+2}\text{OPT}.\end{aligned}$$

Combining the above recurrence relation and $a_0 \le l(f_0) \le 2c'2^{2d}(d+1)^{2d+3}\text{OPT}$, we can prove $a_{j+1} \le 2c'2^{2d}(d+1)^{2d+3}\text{OPT}$ by induction. Thus, $E_{c_1,c_2,\dots,c_j}[l(f_j)] \le O(2^{2d}(d+1)^{2d+3}\text{OPT})$. Hence, the price of sinking is at most $O(2^{2d}(d+1)^{2d+3})$ by the linearity of expectation. As $(d+1)^{2d+3} = O(d^{2d+3})$, we have the desired bound. □

We can also use the lemmas used in the proof of Theorem 4.1 to bound the rate of convergence to states with good social value in unsplittable (weighted) selfish routing games. We can prove that starting from a flow of latency C, after $O(n \log \frac{C}{\text{OPT}})$ random best responses, the expected social value is less than 70 OPT for linear latency functions, and is less than $O(2^{2d}d^{2d+3})\text{OPT}$ for polynomial latency functions of degree at most d. This is in contrast with the negative convergence result of Fabrikant, Papadimitriou, and Talwar [2], in which they exhibit exponentially long best-response paths to PSNE (or sink equilibria) in these games. Our bounds show that, even though convergence to PSNE (or sink equilibria) may be exponential, a random sequence of best responses of agents converges to a state with good social value after polynomial number of best responses. Here, we prove a tighter bound for convergence in the weighted unsplittable selfish routing game with linear latency functions. We assume that the latency function of arc e is a linear function. In particular, we let the latency function for arc $e \in E(G)$ be $\lambda_e(x) = a_e x + b_e$ with $a_e, b_e \ge 0$.

Theorem 4.5. *In the weighted unsplittable selfish routing game with linear latency functions, starting from any state with total latency C the expected latency of the flow after $O(n \log \frac{C}{\text{OPT}})$ random best responses is at most* 70 OPT *for any $\epsilon > 0$.*

Proof. Let f be the current flow, and suppose agent i changes its flow path from P_i to P'_i, to give a new flow f'_i. From Lemma 4.2, $l(f'_i) \le l(f) + 2l_i(f'_i) - l_i(f)$. We will use the following refinement to Lemma 4.4.

Lemma 4.6. *Let f be the flow corresponding to the current strategy profile. Consider the following random process: choose an agent i at random and let it play its best response. If f' is the new flow after this change, then either $\mathbf{E}[l(f')|f] \le (1 - \frac{1}{2n})l(f)$, or $l(f) \le 23.32$* OPT.

Proof. Assume that the best response of agent i is to switch from path P_i to P'_i resulting in the flow f'_i. Thus, $\mathbf{E}[l(f')|f] = \frac{1}{n}\sum_{i\in U} l(f'_i)$. We consider the following two cases:
Case 1: $\sum_{i\in U} 4l_i(f'_i) \le \sum_{i\in U} l_i(f)$. In this case, similar to Case 1 of the proof of Lemma 4.4, it follows that $\mathbf{E}[l(f')|f] \le (1 - \frac{1}{2n})l(f)$.
Case 2: $\sum_{i\in U} 4l_i(f'_i) > \sum_{i\in U} l_i(f)$. Let $\mathcal{P}^* = (P^*_1, \dots, P^*_n)$ be the optimal solution and let f^* be the flow corresponding to $\mathcal{P}^*$. Set $J^*(e) = \{i : e \in P^*_i\}$. Let f^*_i be the flow resulting from the switch of agent i from P_i to P^*_i. Since P'_i is i's best response, we have $l_i(f^*_i) \ge l_i(f'_i)$. In this case, we can apply the method

of Azar et al. [1] as follows:

$$
\begin{aligned}
l(f) &= \sum_i (r_i \sum_{e\in P_i} (a_e f_e + b_e)) \\
&\le \sum_{i\in U} 4 l_i(f_i') \\
&\le \sum_{i\in U} 4 l_i(f_i^*) && \text{[since player } i \text{ play his best response to } f_i'] \\
&\le \sum_{i\in U} 4 r_i \sum_{e\in P_i^*} (a_e(f_e + r_i) + b_e) \\
&= 4 \sum_{i\in U} \sum_{e\in P_i^*} \left[(a_e f_e + b_e) r_i + a_e r_i^2\right] \\
&\le 4 \sum_{e} \sum_{i: e\in P_i^*} \left[(a_e f_e + b_e) r_i + a_e r_i^2\right].
\end{aligned}
$$

It follows that

$$
\begin{aligned}
l(f) &\le 4 \sum_e f_e^* (a_e f_e + b_e) + 4 \sum_e a_e f_e^{*2} \\
&= 4 \sum_e f_e^* a_e f_e + 4 \sum_e (a_e f_e^* + b_e) f_e^* \\
&= 4 \sum_e f_e^* a_e f_e + 4 \text{ OPT} \\
&\le 4 \sqrt{\left(\sum_e (\sqrt{a_e} f_e)^2\right)\left(\sum_e (\sqrt{a_e} f_e^*)^2\right)} + 4 \text{ OPT} \text{[Cauchy-Schwartz inequality]} \\
&= 4 \sqrt{\left(\sum_e a_e f_e^2\right)\left(\sum_e a_e f_e^{*2}\right)} + 4 \text{ OPT} \\
&\le 4 \sqrt{\sum_e (a_e f_e + b_e) f_e \sum_e (a_e f_e^* + b_e) f_e^*} + 4 \text{ OPT} \\
&= 4 \sqrt{l(f)\text{OPT}} + 4 \text{ OPT}.
\end{aligned}
$$

By setting $x = \frac{l(f)}{\text{OPT}}$, we have $x \le 4(\sqrt{x}+1)$. This gives $x \le 23.32$. Hence, in this case, $l(f) \le 23.32$ OPT. □

Proof of Theorem 4.5. Let $a_0 = C$ be the social value of the initial flow. Assume that at each step we choose an agent at random and let it play its best response. Let a_j be the expected latency of the flow after j's step. From Lemma 4.6, we have for any $j \ge 0$, $a_j \le 23.32$ OPT or $a_{j+1} \le a_j(1 - \frac{1}{2n})$. Moreover, from Lemma 4.3, $a_{j+1} \le a_j(1 + \frac{1}{n})$ for any $j \ge 0$. Now, let p be the probability that $a_j > 23.32$ OPT. Let X be the expected value of a_j given that $a_j > 23.32$OPT and Y be the expected value of a_j given that

$a_j \leq 23.32\text{OPT}$. Thus,

$$\begin{aligned} a_{j+1} &\leq p\left(1-\frac{1}{2n}\right)X + (1-p)(1+\frac{1}{n})Y \\ &\leq \left(1-\frac{1}{2n}\right)(pX+(1-p)Y) + \frac{3}{2n}Y \\ &\leq \left(1-\frac{1}{2n}\right)a_j + \frac{69.96}{2n}\text{OPT}. \end{aligned}$$

It follows that

$$a_j \leq a_{j-i}\left(1-\frac{1}{2n}\right)^i + \frac{69.96}{2n}\text{OPT}\frac{\left(1-\left(1-\frac{1}{2n}\right)^i\right)}{\frac{1}{2n}}$$

for $i \leq j$. As a result, $a_j \leq a_0\left(1-\frac{1}{2n}\right)^j + 69.96\ \left(1-\left(1-\frac{1}{2n}\right)^j\right)\text{OPT} \leq C\left(1-\frac{1}{2n}\right)^j + 69.96\ \text{OPT}$. Thus, for $j \geq n\log\frac{1}{\epsilon}\log\frac{C}{\text{OPT}}$, we get $a_j \leq (69.96+\epsilon)\ \text{OPT}$. Therefore, after $O(n\log\frac{C}{\text{OPT}})$ steps the expected value of a_j is at most 70 OPT. □

Finally, we note that all our results on the price of sinking and convergence for weighted unsplittable selfish routing games extend to weighted congestion games. Weighted *congestion games* are the generalization of weighted unsplittable selfish routing game in which the family of feasible strategies of players are an arbitrary family of subsets of arcs (and not necessarily paths from a source to a destination). Our proofs do not rely on the fact that the feasible strategy is a path. Therefore, all our results hold for general weighted congestion games.

4.2. Valid-Utility Games.

Here we define the class of valid-utility games; see [14] for more details. A function f of the form $2^V \to \mathbb{R}^+ \cup \{0\}$ is called a set function on the ground set V. A set function $f : 2^V \to \mathbb{R}^+ \cup \{0\}$ is submodular if for any two sets $A, B \subseteq V$, $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$. A set function f, is non-decreasing if $f(X) \leq f(Y)$ for any $X \subseteq Y \subseteq V$. In valid-utility games, for each player i, there exists a ground set V_i. We denote by V the union of ground sets of all players, i.e., $V = \cup_{i\in U} V_i$. The feasible strategy set F_i of player i is a subset of the power set, 2^{V_i}, of V_i. Thus, the strategy s_i of player i is a subset of V_i ($s_i \subseteq V_i$). The empty set, denoted $\emptyset_i$ for player i, corresponds to player i taking no action.

Given a collection of strategies $S = (s_1, \ldots, s_n)$, where s_i is a subset of the ground set V_i ($s_i \subseteq V_i$), the set $\mathcal{H}_S = \{(i,j) : i \in U, j \in s_i\}$ is called *the pair set* for the collection S. Note that S may or may not be a feasible strategy profile. Given a function $f : \Pi_{i\in U} 2^V \to \mathbb{R}^+ \cup \{0\}$, the *corresponding set function* f^s of f is a set function of the form $2^{\mathcal{H}} \to \mathbb{R}^+ \cup \{0\}$ where $\mathcal{H} = \{(i,j) : i \in U, j \in V\}$ and $f^s(\mathcal{H}_S) = f(S)$. In other words, for a set $A \subseteq \mathcal{H}$, $f^s(A) = f((a_1, a_2, \ldots, a_n))$ if $a_i = \{j : (i,j) \in A\}$. Here, we also assume that the social function γ is of the form $\Pi_{i\in U} 2^V \to \mathbb{R}^+ \cup \{0\}$ rather than just of the form $\Pi_{i\in U} F_i \to \mathbb{R}^+ \cup \{0\}$.

Let $\mathcal{G}(U, \{F_i | i \in U\}, \{\alpha_i() | i \in U\})$ be a non-cooperative strategic game where $F_i \subseteq 2^{V_i}$ is a family of feasible strategies for player i. Let $V = \cup_{i\in U} V_i$ and let the social function be $\gamma : \Pi_{i\in U} 2^V \to \mathbb{R}^+ \cup \{0\}$. Then $\mathcal{G}$ is a *valid-utility game* if it satisfies the following properties:

(1) Submodular and Non-decreasing Social Function: The corresponding set function, γ^s, of γ over the set $\mathcal{H} = \{(i,j) : i \in U, j \in V\}$, is submodular and non-decreasing.

(2) Vickrey Condition: The payoff of a player is at least the difference in the social function when the player participates versus when it does not participate, i.e., $\alpha_i(S) \geq \gamma'_{s_i}(S \oplus \emptyset_i)$. In *basic-utility* games we always have $\alpha_i(S) = \gamma'_{s_i}(S \oplus \emptyset_i)$.
(3) Cake Condition: For any strategy profile, the sum of the payoffs of players should be less than or equal to the social function for that strategy profile, i.e., for any strategy profile S, $\sum_{i \in U} \alpha_i(S) \leq \gamma(S)$.

This framework encompasses a wide range of games including the facility location games, traffic routing games, auctions [14], market sharing games [5], and distributed caching games [3]. In [14] it was shown that the price of anarchy (for mixed Nash equilibria) in valid-utility games is at most 2. While proving theorems about valid-utility and basic-utility games, we use the following notation: given $S = (s_1, \ldots, s_n)$ and $S' = (s'_1, \ldots, s'_n)$, we define $S \cup S' := (s_1 \cup s'_1, \ldots, s_n \cup s'_n)$. Also we define $S \cup s'_i := (s_1, s_2, \ldots, s_{i-1}, s_i \cup s'_i, s_{i+1}, \ldots, s_n)$.

Here we prove bounds on the worst-case price of sinking in valid-utility games. First, we show that our bad example in Section 3 is a valid-utility game. Thus the price of sinking in valid-utility games can be as bad as n. Then, we will prove that this lower bound for valid-utility games is almost tight. In particular, we will show that the price of sinking in a valid-utility game is at most $n + 1$.

In order to prove that the bad example in Section 3 is a valid-utility game, we need to verify three conditions:

1) **Non-decreasing and Submodular Social Function::** First, it is clear that the corresponding set function of the social function γ^s is non-decreasing. To show its submodularity, we use an equivalent definition of submodular functions: A set function f is submodular if for any two subsets A and B such that $A \subset B$ and for any element $i \notin B$, $f(A \cup \{i\}) - f(A) \geq f(B \cup \{i\}) - f(B)$. Thus, in order to prove that γ^s is submodular, it is enough to prove that for two (possibly infeasible) strategy profiles $S = (s_1, \ldots, s_n)$ and $S' = (s'_1, \ldots, s'_n)$ such that $s_i \subseteq s'_i$ for all $i \in U$, by adding a new element j to the strategy of any player i the increase in γ^s for S is not less than the increase for S'. First, we consider the case that $j = x_i^t$. If $S'^{\cup} \cap X = \emptyset$ then $S^{\cup} \cap X = \emptyset$, and thus $\gamma'_{x_i^t}(S' \oplus \emptyset_i) = 2$ and $\gamma'_{x_i^t}(S \oplus \emptyset_i) = 2$. If $S'^{\cup} \cap X \neq \emptyset$ then $\gamma'_{x_i^t}(S' \oplus \emptyset_i) = 0 \leq \gamma'_{x_i^t}(S \oplus \emptyset_i)$. Hence if $j = x_i^t$, the desired condition for submodularity holds. Also, if $j = y_i$ it is implied that $\gamma'_{y_i}(S' \oplus \emptyset_i) = 1$ if and only if $S'^{\cup} \cap \{y_i\} = \emptyset$, otherwise $\gamma'_{y_i}(S' \oplus \emptyset_i) = 0$. It follows that $\gamma'_{y_i}(S' \oplus \emptyset_i) \leq \gamma'_{y_i}(S \oplus \emptyset)$. Therefore, γ^s is submodular.
2) **Vickrey Condition::** If player i plays y_i then she gets 1 and the social value changes by 1. If player i plays an element of X_i and increases the social value by 2, then she is the only player who plays an irresponsible strategy. Thus, $i = i^*(S)$ and so she receives those two utility units. Otherwise the playing of an element of X_i has no effect on the social value. Thus, the Vickrey condition is trivially satisfied.
3) **Cake Condition::** It is straightforward to check that $\sum_{i \in U} \alpha_i(S) = \gamma(S)$ and the cake condition holds.

Now, we prove that this bound is almost tight.

Lemma 4.7. *Given a strategy profile $T = (t_1, \ldots, t_n)$ in a valid-utility game, let the best response of agent i be s_i. Set $T^i = (t_1, \ldots, t_{i-1}, s_i, t_{i+1}, \ldots, t_n)$. Then $\sum_{i \in U} \alpha_i(T^i) \geq \text{OPT} - \gamma(T)$.*

Proof. Let $\Omega = (\sigma_1, \ldots, \sigma_n)$ be the optimum state. Let

$$\Omega^i = (\sigma_1, \sigma_2, \ldots, \sigma_i, \emptyset_{i+1}, \emptyset_{i+2}, \ldots, \emptyset_n).$$

Given that s_i is a best-response strategy, we have $\alpha_i(T^i) \geq \gamma'_{\sigma_i}(T \oplus \emptyset_i)$. Combining this with the submodularity of γ, we obtain

$$\begin{aligned}
\sum_{i \in U} \alpha_i(T^i) &\geq \sum_{i \in U} \gamma'_{\sigma_i}(T \oplus \emptyset_i) \\
&= \sum_{i \in U} (\gamma(T \oplus \sigma_i) - \gamma(T \oplus \emptyset_i)) \\
&\geq \sum_{i \in U} (\gamma(T \cup \sigma_i) - \gamma(T)) \\
&\geq \sum_{i \in U} (\gamma(T \cup \Omega^i) - \gamma(T \cup \Omega^{i-1})) \\
&= \gamma(T \cup \Omega) - \gamma(T).
\end{aligned}$$

Since γ is non-decreasing, it follows that $\sum_{i \in U} \alpha_i(T^i) \geq \text{OPT} - \gamma(T)$. □

Theorem 4.8. *The price of sinking in a valid-utility game is at most $n + 1$.*

Proof. Consider a sink equilibrium Q. Let $T = (t_1, \ldots, t_n)$ be a state in Q. Let the best response of agent i be s_i at state T, and set $T^i = (t_1, \ldots, t_{i-1}, s_i, t_{i+1}, \ldots, t_n)$. Let Y be the expected social value of the state after a random best-response move from T. By the cake property and Lemma 4.7, we have

$$\begin{aligned}
Y &= \frac{1}{n} \sum_{i \in U} \gamma(T^i) \\
&\geq \frac{1}{n} \sum_{i \in U} \alpha_i(T^i) \\
&\geq \frac{1}{n} (\text{OPT} - \gamma(T)).
\end{aligned}$$

Observe that the price of sinking is equal to the expected social value on a sufficiently long random walk. Now take a long random walk $T_0, T_1, \ldots, T_k$. Let e_i be the expected value of $\gamma(T_i)$ where the expectation is over the random coin tosses of the random walk. We know that as i tends to ∞, $\Gamma(Q) = e_i$. We need to prove that $e_i \geq \frac{1}{n+1}\text{OPT}$ as i tends to ∞. Let $p_{i,y}$ be the probability that $\gamma(T_i) = y$. Thus, $e_i = \sum_y p_{i,y} y$ and $e_{i+1} = \sum_y p_{i,y} \mathbf{E}[\gamma(T_{i+1}) | \gamma(T_i) = y]$. The above inequality shows that $\mathbf{E}[\gamma(T_{i+1}) | \gamma(T_i) = y] \geq \frac{1}{n}(\text{OPT} - y)$. Therefore,

$$\begin{aligned}
e_{i+1} &\geq \frac{1}{n} \sum_y p_{i,y} (\text{OPT} - y) \\
&= \frac{1}{n} (\text{OPT} - \sum_y p_{i,y} y) \\
&= \frac{1}{n} (\text{OPT} - e_i).
\end{aligned}$$

Hence, $e_{i+1} \geq \frac{1}{n}\text{OPT} - \frac{e_i}{n}$. Since as i goes to ∞, $\Gamma(Q) = e_i = e_{i+1}$, we get $\Gamma(Q) \geq \frac{1}{n}\text{OPT} - \frac{\Gamma(Q)}{n}$. Therefore, $\Gamma(Q) \geq \frac{1}{n+1}\text{OPT}$ as desired. $\square$

Thus the worst case price of sinking in a valid-utility game is between n and $n+1$.

4.3. Basic Utility Games.

For basic utility games (examples include service provider and facility location games [14]) the situation is much better. These games are potential games, thus, the only sink equilibria are PSNE. Hence, the price of sinking in a basic-utility game is equal to the price of anarchy for PSNE which is at most 2. Using similar techniques to those of Theorem 4.8, we can prove that in basic-utility games, the expected social value of a state after $O(n \log \frac{1}{\epsilon})$ random best responses is at least $\frac{1}{2} - \epsilon$ of the optimal social value, for any $\epsilon > 0$.

Theorem 4.9. *In basic-utility games, for any constant $\epsilon > 0$, there exists a constant c such that the expected social value of a state after $cn \log \frac{1}{\epsilon}$ random best responses is at least $\frac{1}{2} - \epsilon$ of the optimum. Moreover, for any constant $\epsilon' > 0$, there exist constants $\epsilon, c' > 0$ such that after $c'n \log n \log \frac{1}{\epsilon}$ random best responses, the social value is at least $\frac{1}{2} - \epsilon'$ of the optimum with high probability.*

Proof. Let $\Omega = (\sigma_1, \ldots, \sigma_n)$ denote an optimal state, and $T = (t_1, t_2, \ldots, t_n)$ be a strategy profile of agents. Let T^i be the strategy profile resulting from T after agent i plays its best response in T and let $\Omega^i = (\sigma_1, \ldots, \sigma_i, \emptyset_{i+1}, \ldots, \emptyset_n)$. Let $Y = \frac{1}{n}\sum_{i \in U} \gamma(T^i)$ be the expected social value of the state after a random agent plays its best response. Our goal is to lower bound Y.

To do so, using submodularity, basicness and the cake condition we get:

$$
\begin{aligned}
nY - n\gamma(T) &= \sum_{i \in U} (\gamma(T^i) - \gamma(T)) \\
&= \sum_{i \in U} (\gamma(T^i) - \gamma(T \oplus \emptyset_i)) - \sum_{i \in U} (\gamma(T) - \gamma(T \oplus \emptyset_i)) \\
&= \sum_{i \in U} \alpha_i(T^i) - \sum_{i \in U} \alpha_i(T) \quad \text{[by basicness]} \\
&\geq \sum_{i \in U} \alpha_i(T^i) - \gamma(T) \quad \text{[by cake condition]} \\
&\geq \sum_{i \in U} \alpha_i(T^i \oplus \sigma_i) - \gamma(T) \quad \text{[since } i \text{ plays his best response in } T^i] \\
&= \sum_{i \in U} \gamma'_{\sigma_i}(T \oplus \emptyset_i) - \gamma(T) \quad \text{[by basicness]} \\
&\geq \sum_{i \in U} (\gamma(T \oplus \sigma_i) - \gamma(T \oplus \emptyset_i)) - \gamma(T) \\
&\geq \sum_{i \in U} (\gamma(T \cup \Omega^i) - \gamma(T \cup \Omega^{i-1})) - \gamma(T) \quad \text{[by submodularity]} \\
&= \gamma(T \cup \Omega) - \gamma(T) - \gamma(T) \quad \text{[since it is a telescopic summation]} \\
&\geq \text{OPT} - 2\gamma(T) \quad \text{[since } \gamma \text{ is non-decreasing]}.
\end{aligned}
$$

The above inequalities show that $Y \geq \frac{n-2}{n}\gamma(T) + \frac{1}{n}\text{OPT}$. Let Y_0 be the actual social value of the initial state. At each step, a random agent is picked and plays its best response. Thus, if Y_i is the social value of the state after step i, then $\mathbf{E}[Y_i|Y_{i-1} = y] \geq (\frac{n-2}{n})y + \frac{1}{n}\text{OPT}$. Let $p_{yy'}$ be the probability that $Y_{i-1} = y'$ given that $Y_{i-2} = y$. Thus, $\mathbf{E}[Y_{i-1}|Y_{i-2} = y] = \sum_{y'} p_{yy'}y'$. Therefore,

$$\begin{aligned}
\mathbf{E}[Y_i|Y_{i-2} = y] &= \sum_{y'} p_{yy'}\mathbf{E}[Y_i|Y_{i-2} = y, Y_{i-1} = y'] \\
&\geq \sum_{y'} p_{yy'}((\frac{n-2}{n})y' + \frac{1}{n}\text{OPT}) \\
&= (\frac{n-2}{n})\mathbf{E}[Y_{i-1}|Y_{i-2} = y] + \frac{1}{n}\text{OPT} \\
&\geq (\frac{n-2}{n})((\frac{n-2}{n})y + \frac{1}{n}\text{OPT}) + \frac{1}{n}\text{OPT} \\
&= (\frac{n-2}{n})^2 y + \frac{1}{n}\text{OPT}(1 + (\frac{n-2}{n})).
\end{aligned}$$

Thus, $\mathbf{E}[Y_i|Y_{i-2} = y] \geq (\frac{n-2}{n})^2 y + \frac{1}{n}\text{OPT}(1 + (\frac{n-2}{n}))$. Similarly, we can prove that $\mathbf{E}[Y_i|Y_0 = y_0] \geq (\frac{n-2}{n})^i y_0 + \frac{1}{n}\text{OPT}(1 + (\frac{n-2}{n}) + \ldots + (\frac{n-2}{n})^{i-1})$. Since $y_0 \geq 0$, $\mathbf{E}[Y_i] \geq \frac{\text{OPT}}{2}(1 - (1 - \frac{2}{n})^i)$.

This proves that for a sufficiently large constant c and by setting $i = cn \log \frac{1}{\epsilon}$, the expected social value after $cn \log \frac{1}{\epsilon}$ best responses is at least $\frac{1}{2} - \epsilon$ of the optimum. Moreover, since in basic-utility games the social value is non-decreasing as agents play their best responses, we claim that for a sufficiently large $c' = cc'' > 0$ and a sufficiently small $\epsilon > 0$, after $c'n \log n \log \frac{1}{\epsilon}$ random best responses, with high probability the social value is at least $\frac{1}{2} - \epsilon'$ of the optimum. The reason is that we can partition the best response walk of length $c'n \log n \log \frac{1}{\epsilon}$ into $c'' \log n$ best-response walks of length $cn \log \frac{1}{\epsilon}$, and after each of these subwalks, the expected social value is at least $\frac{1}{2} - \epsilon$ of the optimum. Thus, by Markov inequality, with a constant probability after each of the subwalks of length $cn \log \frac{1}{\epsilon}$, the expected social value is at least $(\frac{1}{2} - \epsilon')$ of the optimum. Hence, after $c'n \log n \log \frac{1}{\epsilon}$ best responses, the social value is at least $\frac{1}{2} - \epsilon'$ of the optimum with high probability. □

5. A Hardness Result

In this section, we prove that finding a sink equilibrium (or a PSNE if it exists) in some instances of valid-utility games is PLS-complete. We prove this using a tight PLS reduction from the Max-Cut problem. This, in turn, has some implications on the convergence to sink equilibria of these games.

Theorem 5.1. *Finding a sink equilibrium is PLS-complete for some instances of valid-utility games.*

Proof. We give a reduction from the Max-Cut local search problem. Consider an instance $G = (V, E)$, with edge weights, of the Max-Cut problem. Suppose our local search operation is the moving of a vertex from one side of the bipartition to the other; this operation is called *swapping*.

We create a game corresponding to this instance as follows. There is an agent v for each vertex v in the graph. The groundset of agent v is $V_v = \{e_L, e_R : e \in \delta(v)\}$. Thus the size of the groundset of agent v is twice the degree of v. A strategy for agent v is then just a subset of this groundset. Now, given the strategy profile $(T_v : v \in V)$, what is the payoff to each agent?

Take an edge $e = (u, v)$. Agent v receives a payoff of value w_e if it plays e_L but agent u does not, and vice versa. If both agents play e_L then they each receive a payoff of value $\frac{1}{2}w_e$. If neither agents plays e_L then they both receive nothing. The same payoff scheme arises with the element e_R. Hence the payoff to agent v is the sum of the payoffs he receives from each of the elements in its groundset. The social function is just the sum of payoffs of agents. We claim that this game is a valid-utility game. Submodularity follows from the construction. By definition $\sum_i \alpha_i(S) = \gamma(S)$, so the cake condition holds. Finally, if agent i plays a subset $S_i \subset V_i$ then the increase in the social value is exactly the sum of the values of those elements in S_i that have not been played by another agent. The payoff to agent i is the value of these elements *plus* half the value of the other elements in S_i. So the Vickrey condition holds.

In order for the game to model the Max-Cut problem we just need to restrict the set of feasible strategies for each agent. In fact, we will only allow two feasible strategies (in addition to the null strategy) per agent. Agent v has the feasible strategies $L_v = \{e_L : e \in \delta(v)\}$ and $R_v = \{e_R : e \in \delta(v)\}$. The motivation for this is clear; the former strategy corresponds to placing v on the left side of the partition, the latter strategy corresponds to placing v on the right side of the partition. Since it is never in an agents interest to play the null strategy, a best response move for agent v is either staying in its current side of the partition or *swapping* to the other side of the partition.

Let $W = \sum_{e \in E} w_e$, and suppose that the strategy profile $T = (T_v : v \in V)$ induces a cut (A, B) in the graph. Then the social function is $\gamma(T) = W + \sum_{e \in \delta(A)} w_e$, and private utility of agent v is $a_v(T) = \frac{1}{2}\sum_{e \in \delta(v), e \notin \delta(A)} w_e + \sum_{e \in \delta(v), e \in \delta(A)} w_e$. Evidently, the social objective is to maximize the cut. Moreover, this reduction equates local improvements in the Max-Cut local search problem with best responses of the valid-utility game. If the local improvement corresponds to node v swapping sides, then the increase in the value of a cut equals the increase in the social function. This, in turn, is exactly twice the increase in the private payoff of node v arising from the swap. The theorem follows. $\square$

Using the above reduction from the Max-Cut problem to a valid-utility game, we can prove the following:

Corollary 5.2. *In some instances of valid-utility games, there exist states that are exponentially far from any sink equilibrium.* $\square$

References

[1] B. Awerbuch and Y. Azar and A. Epstein, "The price of routing unsplittable flow", *STOC*, 2005.

[2] A. Fabrikant and C. Papadimitriou and K. Talwar, "On the complexity of pure equilibria", *STOC*, 2004.

[3] L. Fliescher, M. Goemans, V. Mirrokni and M. Sviridenko, "Almost tight approximation algorithms for maximising general assignment problems", submitted, 2005.

[4] D. Fotakis and S. Kontogiannis and P. Spirakis, "Selfish unsplittable flow", *ICALP*, 2004.

[5] M. Goemans, L. Li, V. Mirrokni, and M. Thottan, "Market sharing games applied to content distribution in ad-hoc networks", *MOBIHOC*, 2004.

[6] M. Kandori, G. Mailath and R. Rob, "Learning, mutuation, and long-run equilibria in games", *Econometrica*, **61**, pp29-56, 1993.

[7] E. Kohlberg and J. Mertens, "On the strategic stability of equilibria", *Econometrica*, **54(5)**, pp1003-1037, 1986.

[8] E. Koutsoupias and C. Papadimitriou, "Worst-case equilibria", *STACS*, 1999.

[9] V. Mirrokni and A. Vetta, "Convergence issues in competitive games", *RANDOM-APPROX*, 2004.

[10] C. Papadimitriou, "Algorithms, games and the internet", *STOC*, 2001.

[11] C. Papadimitriou and A. Schaffer and M. Yannakakis, "On the strategic stability of equilibriaOn the complexity of local search", *STOC*, 1990.

[12] T. Roughgarden and E. Tardos, "How bad is selfish routing?", *J. ACM*, **49(2)**, pp236-259, 2002.

[13] A. Schaffer and M. Yannakakis, "Simple local search problems that are hard to solve", *SIAM Journal on Computing*, **20(1)**, pp56-87, 1991.

[14] A. Vetta, "Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions", *FOCS*, 2002.