The curse of sequentiality in routing games*

José Correa[1], Jasper de Jong[2], Bart de Keijzer[3], and Marc Uetz[2]

[1] Universidad de Chile, Santiago, Chili
correa@uchile.cl
[2] Universiteit Twente, Enschede, The Netherlands
{j.dejong-3,m.uetz}@utwente.nl
[3] Sapienza University of Rome, Rome, Italy
dekeijzer@dis.uniroma1.it

Abstract. In the "The curse of simultaneity", Paes Leme at al. show that there are interesting classes of games for which sequential decision making and corresponding subgame perfect equilibria avoid worst case Nash equilibria, resulting in substantial improvements for the price of anarchy. This is called the sequential price of anarchy. A handful of papers have lately analysed it for various problems, yet one of the most interesting open problems was to pin down its value for linear atomic routing (also: *network congestion*) games, where the price of anarchy equals 5/2. The main contribution of this paper is the surprising result that the sequential price of anarchy is unbounded even for linear symmetric routing games, thereby showing that sequentiality can be arbitrarily worse than simultaneity for this important class of games. Complementing this unboundedness result we solve an open problem in the area by establishing that the (regular) price of anarchy for linear symmetric routing games equals 5/2. Additionally, we prove that in these games, even with two players, computing the outcome of a subgame perfect equilibrium is NP-hard.

1 Introduction

The concept of the price of anarchy, introduced by Koutsoupias and Papadimitriou [9], has spurred a lot of research over the past 15 years that has contributed significantly to establish the area of algorithmic game theory. Not only Nash equilibria, but also alternative equilibrium concepts have been addressed. One recent and interesting example of the latter is the *sequential price of anarchy* (*SPoA*), recently introduced by Paes Leme et al. [13], that aims at understanding the quality of subgame perfect equilibrium outcomes of a game.

Similar to the price of anarchy (*PoA*) [9], the sequential price of anarchy measures the cost of decentralization. However, while the price of anarchy compares the quality of a worst case Nash equilibrium to the quality of an optimal

* This work is partially supported by the EU FET project MULTIPLEX no. 317532, the ERC StG Project PAAl 259515, the Google Research Award for Economics and Market Algorithms, the EU-IRSES project EUSACOU, and the Millennium Nucleus Information and Coordination in Networks ICM/FIC RC130003.

solution, the sequential price of anarchy considers the possible outcomes of a game where players choose their strategies sequentially in some arbitrary order. It then compares the quality of the outcome of the worst possible subgame perfect equilibrium [15] to the quality of an optimal solution. Note that for games with perfect information, subgame perfect equilibria coincide with sequential equilibria as introduced by Kreps and Wilson [10]. In that sense, subgame perfect equilibria are indeed the "right" equilibrium concept for the most natural sequential routing games. It turns out that there are interesting examples of games where this notion leads to improved worst case guarantees, and in this sense avoid the "curse of simultaneity" [13] inherent in some simultaneous move games. Indeed, for a handful of games, the *SPoA* has indeed been proven to be lower than the *PoA* [13,7,8], while for others, this is not the case [1,5].

In this paper we consider one of the most basic types of congestion game, namely the atomic network routing game with linear latencies. Here, the *PoA* has long been known to be equal to 5/2 [3,6], while de Jong and Uetz [8] recently showed that the *SPoA* is less than the *PoA* for a small number of players leading them to conjecture that the *SPoA* is at most 5/2. Our main result is to disprove this conjecture, and to thereby establish a sharp contrast between the *PoA* and the *SPoA* in network routing games. Indeed, we prove that even in the symmetric case, i.e. when all players share the same origin and destination, the *SPoA* is not bounded by any constant and can be as large as $\Omega(\sqrt{n})$, with n being the number of players.

The crucial part of our proof is to come up with a "contingency plan of actions" for every player and every possible move of all previous players that indeed leads to a subgame perfect outcome. This is generally very difficult, since the strategies of the players are of exponential size. We are however able to design a plan leading to an unbounded *SPoA* that can be described in a succinct manner: The core idea, that we believe may be of independent interest, is to design a master plan of actions that all players are supposed to follow, together with a punishing action that players only apply when some previous player deviates from the master plan. The main technical difficulty is to design a construction such that the punishing actions do not lead to a higher cost for the player applying it, so that subgame perfection is achieved.

To complement the previous result, we resolve an open problem posed by Bhawalkar et al. [4] about the *PoA* for symmetric atomic network routing games with linear latencies. Indeed, we prove that this equals 5/2, as it is the case for the nonsymmetric network case [3] and the symmetric case for general congestion games [6] (not necessarily networked).

Finally, we prove a number of additional results for the symmetric two player case. We start by observing that even for just two players subgame perfect equilibria are more complex than Nash equilibria. In particular, the corresponding outcome is generally not a Nash equilibrium of the simultaneous game, as opposed to the *crowding games* studied by Milchtaich [12]. Furthermore, we show that computing the outcome of a subgame perfect equilibrium is in general NP-hard. Although we know from [13] that computing subgame perfect equilibria is

PSPACE-complete in general congestion games, that reduction requires a non-constant number of players. Our result shows that the problem remains at least NP-hard even when the number of players is two. To conclude, we pin down the exact sequential price of anarchy for the symmetric two player case, showing that it equals 7/5. This constitutes an improvement over the 3/2 upper bound in the more general non-symmetric case [8], but is higher than the straightfoward 4/3 bound for the *PoA*.

2 Model and Notation

Throughout we consider a special case of *atomic congestion games*, namely, symmetric atomic network routing games with linear latency functions. The input of an instance $I \in \mathcal{I}$ consists of a directed graph $G = (V, E)$, with designated source and target nodes $s, t \in V$, and for each arc $e \in E$ a linear latency function with coefficient d_e. There are n players that all want to travel from s to t, so that the possible actions of all players consist of all directed (s,t)-paths in G. Note that all players have the same set of actions at their disposal, hence the term *symmetric*. We will denote by m the number of arcs $|E|$. We refer to the possible paths a player can choose the *actions* and to a vector of paths, one for each player, $A = (A_1, \ldots, A_n)$ as an *outcome* or *action profile*.

The cost of a player i for choosing a specific (s,t)-path A_i depends on the number of players on each arc on that path. Specifically, for an outcome $A = (A_1, \ldots, A_n)$, let $n_e(A) := \sum_{i=1}^{n} |A_i \cap \{e\}|$ denote the number of players using arc e, then the cost of that arc for each player using it equals $n_e(A)d_e$, and therefore the cost for player i, choosing path A_i, is defined as[4]

$$c_i(A) = \sum_{e \in A_i} d_e \cdot n_e(A)\,.$$

This induces the *social cost* $C(A) = \sum_{i=1}^{n} c_i(A)$, i.e., the sum of the costs of the players.[5]

A pure Nash equilibrium is an outcome A in which no player can decrease her costs by unilaterally deviating, i.e. switching to an action that is different from A_i. The price of anarchy *PoA* [9] measures the quality of any Nash equilibrium relative to the quality of a globally optimal allocation, *OPT*. Here *OPT* is an outcome minimizing $C(\cdot)$. More specifically, for an instance I,

$$\mathit{PoA}(I) = \max_{\mathit{NE} \in \mathit{NE}(I)} \frac{C(\mathit{NE})}{C(\mathit{OPT})}, \tag{1}$$

where $\mathit{NE}(I)$ denotes the set of all Nash equilibria for instance I. The price of anarchy of a class of instances $\mathcal{I}$ is defined by $\mathit{PoA}(\mathcal{I}) = \sup_{I \in \mathcal{I}} \mathit{PoA}(I)$.

[4] Our upper bound on the *SPoA* for two players also holds with affine functions.

[5] Note that we consider a *utilitarian* social cost function. This is one of the standard models, yet different than the *egalitarian* makespan objective as studied, e.g., in [9].

In this paper our goal is to evaluate the quality of *subgame perfect equilibria* of an induced extensive form game that we call the *sequential* version of the game [11,15]. In the sequential game, players choose an action from the set of (s,t)-paths, but instead of doing so simultaneously, they choose their actions in an arbitrary predefined order $1, 2, \ldots, n$, so that the i-th player must choose action A_i, observing the actions of players preceding i, but of course not observing the actions of the players succeeding her.[6] A strategy S_i then specifies for player i the full contingency plan of actions she would choose for each potential choice of actions $A_{<i} = (A_1, \ldots, A_{i-1})$ chosen by her predecessors. We use $S_i(A_{<i})$ to denote the action that i plays under strategy S_i when $A_{<i}$ is the vector of actions chosen by players $1, \ldots, i-1$. We refer to a choice of strategies $S = (S_1, \ldots, S_n)$ by each of the players as a *strategy profile*. Note the explicit distinction between action (profile) and strategy (profile). The *outcome resulting from S* is then the set of actions chosen by the players when they play according to the strategy profile S.

Subgame perfect equilibria, defined by Selten [15], are defined as strategy profiles S that induce pure Nash equilibria in any subgame of the extensive form game. In other words, a strategy profile S is a subgame perfect equilibrium if for all i and for any choice of actions $A_{<i}$ of players $1, \ldots, i-1$, player i cannot decrease her cost by switching to an action different from $S_i(A_{<i})$, in the *subgame* where the actions of $1, \ldots, i-1$ are fixed to $A_{<i}$, and $i, \ldots, n$ play strategies $(S_i, \ldots, S_n)$.

Subgame perfect equilibria reflect farsighted strategic behaviour of players that observe the state of the game and reason strategically about choices of subsequent players, always choosing the action that will minimize their individual cost. Analogous to (1), the sequential price of anarchy of an instance I is defined by

$$SPoA(I) = \max_{SPE \in SPE(I)} \frac{C(SPE)}{C(OPT)}, \tag{2}$$

where $SPE(I)$ denotes the set of all outcomes of subgame perfect equilibria of instance I. The sequential price of anarchy of a class of instances $\mathcal{I}$ is defined as in [13] by $SPoA(\mathcal{I}) = \sup_{I \in \mathcal{I}} SPoA(I)$. Throughout the paper, when the class of instances is clear from the context, we write *PoA* and *SPoA*.

Extensive form games can be represented in a *game tree* (see Figure 1 for an example), with the nodes on one level representing the possible states of the game that a single player can encounter, and the arcs emanating from any node representing the possible actions of that player in the given state. The nodes of the game tree are called *information sets* or *states*. We will refer to a state by a pair $(A_{<i}, i)$ where $A_{<i}$ is the choice of actions of the players $1, \ldots, i-1$ in that state, and i is the next player who has to choose her action. Since we deal with a game with perfect information, subgame perfect equilibria correspond to *sequential equilibria* (see [10]), and can be computed by backward induction. In

[6] However, since players are fully rational and fully informed, at equilibrium they anticipate the others' behavior and therefore make optimal choices anticipating the followers actions.

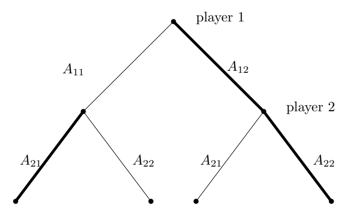

Fig. 1. Game tree for a symmetric sequential game with two players. The nodes are the states. Note that A_{11} and A_{21} are actions of players one and two respectively, but denote the same action (recall that we have a symmetric game). The same holds for for A_{12} and A_{22}. Fat lines denote a subgame perfect strategy $S = (S_1, S_2)$ where $S_1 = A_{12}$, $S_2(A_{11}) = A_{21}$ and $S_2(A_{12}) = A_{22}$. The outcome resulting from S would be (A_{12}, A_{22}), i.e., the rightmost path of the game tree.

particular, it is known that subgame perfect equilibria always exist; see e.g. [14]. Note however that, if S is a subgame perfect equilibrium, the resulting outcome A need not be a Nash equilibrium of the corresponding strategic form game, as will also be witnessed in the next section.

3 Warm-up: the two-player case

As a way to illustrate the difficulties behind arguing about subgame perfect equilibria in general we focus for the moment on the two player case and point out two phenomena that showcase the fundamental difference between the concept of subgame perfect equilibrium and that of Nash equilibrium.

First we derive a simple instance in which the resulting actions of a subgame perfect equilibrium do not correspond to a Nash equilibrium. This contrasts with the case of parallel links [8] and even with the so-called crowding games [12]. Based on this particular instance we additionally prove that the sequential price of anarchy for the two player case equals 7/5. This exceeds the price of anarchy (which equals 4/3), but it is smaller than the sequential price of anarchy for the asymmetric case (which equals 3/2 [8]).[7] Secondly, we show that even in the two-player case, computing the outcome of a subgame perfect equilibrium is NP-hard.

3.1 The sequential price of anarchy

Consider the two-player instance depicted in Figure 2, with five vertices and eight arcs. The vertices $1, 2, \dots, 5$ are numbered from left to right and from top to bottom so that $s = 1$ and $t = 5$. The linear latency functions are given by the numbers next to the respective arcs.

It can be easily verified that the following is a subgame perfect equilibrium:

[7] In [8] a lower bound example is given for general congestion games which can be easily transformed to network routing games.

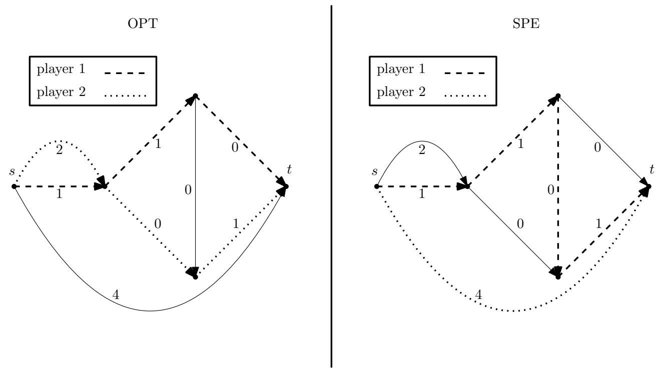

Fig. 2. Lower bound example for 2 players. Numbers are arc latencies.

- Player 1 chooses path $(1,2,3,4,5)$.
- Player 2 chooses:
  - $(1,5)$ if player 1 chooses $(1,2,3,4,5)$,
  - $(1,2,4,5)$ if player 1 chooses $(1,2,3,5)$,
  - $(1,2,3,5)$ if player 1 chooses $(1,2,4,5)$,
  - Any (best response) path for all remaining choices of player 1.

In this equilibrium outcome player 1 chooses the dashed path on the right, that is vertices $(1,2,3,4,5)$, while player 2 chooses the dotted path on the right, which is simply the straight arc going from 1 to 5. Interestingly, one may think that player 1 has an incentive to deviate to the path $(1,2,3,5)$ since the cost of going straight from 3 to 5 is 0. However, if player one does this, player two would pick path $(1,2,4,5)$ and therefore player 1's cost would still be 3. This implies that indeed the outcome of the subgame perfect equilibrium is not a Nash equilibrium. Note furthermore that player 1's cost is 3 and player 2's is 4, for a total social cost of 7, while in the socially optimal situation, depicted to the left of the figure, the social cost is 5. So in particular this instance shows that the *SPoA* is at least $7/5$.

In the above, the subgame perfect equilibrium is not unique. However, the latencies can be slightly perturbed so uniqueness is achieved, while the cost of the equilibrium remains arbitrarily close to 7 and that of the optimum remains arbitrarily close to 5. To this end consider the same instance but changing the latency of the $(1,2)$ arc of latency 2 to $2+\epsilon$, that of the $(1,5)$ arc from 4 to $4+\epsilon$, and those of arcs $(2,4)$ and $(3,5)$ from 0 to ϵ.

With the latter observation not only the sequential price of anarchy but also the *sequential price of stability*[8] equals $7/5$ in the two-player case. This is because it is possible to prove a matching upper bound, even for the more general class of symmetric affine congestion games. The proof of this upper

[8] Just like the price of stability as defined in [2], the sequential price of stability is the ratio of the outcome of the best subgame perfect equilibrium over the optimum.

bound is a bit tedious, and can be found in the full version of this paper. It uses a proof technique based on linear programming, but is nonetheless fundamentally different from the technique used in [8] (where linear programming is also used to derive upper bounds on the *SPoA*).

3.2 Hardness of computing subgame perfect equilibria

Notice that the encoding of subgame perfect strategies can, in general, require super-polynomial space in terms of the input size of a network routing game. This is even the case for two players, for example if the first player has a super-polynomial number of possible actions, i.e., (s, t)-paths. Then, for each of these potential actions of player one, a subgame perfect equilibrium needs to prescribe the respective actions taken by player two. We head for a meaningful statement, however, with respect to the input size of a network routing game, and not the output. Therefore we consider the computational problem to only output the *outcome* resulting from a subgame perfect equilibrium. This exactly corresponds to a single path in the game tree, which for two players has depth two. This outcome has polynomial size, as it is just one path per player. The problem to compute such an equilibrium path in the game tree, however, turns out to be hard.

Theorem 1. *Computing an action profile resulting from a subgame perfect equilibrium of symmetric linear network routing games or symmetric affine congestion games is (strongly)* NP*-hard for any number of players $n \geq 2$.*

The proof is by a reduction from the hamiltonian path problem, and is deferred to the full version of the paper. Moreover, we can also show NP-completeness of the decision problem that asks if in a two-player game, the cost of the first player is below some threshold k in a subgame perfect equilibrium.

4 The n-player case

Our main result is as follows.

Theorem 2. *The sequential price of anarchy of symmetric linear network routing games is unbounded.*

We prove the theorem by constructing a sequence of lower bound instances where the sequential price of anarchy gets arbitrarily large. Intuitively, the construction of these instances works as follows. To obtain a *SPoA* of x, an instance consists of x segments. In OPT, the majority of players chooses only a single free resource per segment, while in the worst case subgame perfect equilibrium, the majority of players form groups of $\sqrt{x}$ players who choose the same sets of $\sqrt{x}$ resources per segment. Any player who deviates from this strategy is punished by some of her successors. The tricky part of the construction is to make sure that all punishing strategies are credible. This is achieved in the following way: There

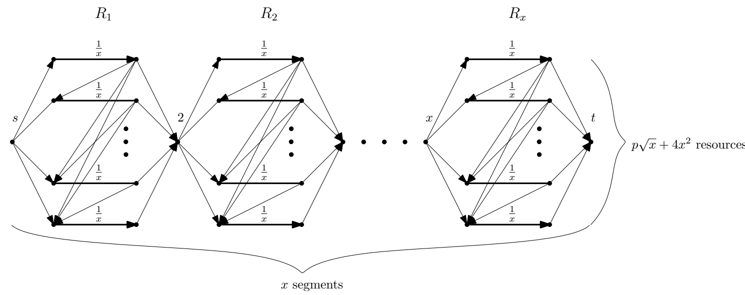

Fig. 3. A lower bound instance of a network routing game. Players travel from s to t.

are slightly more players than disjoint strategies. As an effect, the last player has to necessarily share every arc in her chosen action with one other player. That will result in the situation that this player can credibly "threaten" any other player j by choosing the arcs that player j chooses, if player j does not stick to a certain action. More generally, we extend this idea so that a whole group of players can force a common predecessor into a certain action. This is achieved in such a way that the "concerted" threatening is not too expensive for every single threatener, but very expensive for the common predecessor.

Definition of instance Γ_x. Formally, in order to obtain a sequential price of anarchy of x, where $x \geq 4$ is a square number, we construct the following instance Γ_x: Let p be a sufficiently large integer. There are $n = p\sqrt{x} + 5x^2$ players. The network consists of x *segments* $R_i, i \in [x]$. Segment R_i consists of $2(1 + p\sqrt{x} + 4x^2)$ nodes $\{i, (2i, 1), (2i, 2), \ldots, (2i, p\sqrt{x} + 4x^2), (2i + 1, 1), (2i + 1, 2), \ldots, (2i + 1, p\sqrt{x} + 4x^2), i + 1\}$. Note that node $i + 1$ is in both segments R_i and R_{i+1}. There is an arc with latency 0 from node i to node $(2i, j)$ for all $j \in \{1, \ldots, (p\sqrt{x} + 4x^2)\}$. There is an arc with latency $1/x$, from $(2i, j)$ to $(2i + 1, j)$ for all $j \in \{1, \ldots, (p\sqrt{x} + 4x^2)\}$. There is an arc with latency 0 from $(2i + 1, j)$ to $i + 1$ for all $j \in \{1, \ldots, (p\sqrt{x} + 4x^2)\}$. There is an arc with latency 0 from $(2i + 1, j)$ to $(2i, k)$ for all $j \in \{1, \ldots, (p\sqrt{x} + 4x^2)\}$ and for all $k \in \{j, \ldots, (p\sqrt{x} + 4x^2)\}$. Note that between any nodes $i, i + 1$, there exist $2^{p\sqrt{x}+4x^2}$ different paths: one for every subset of arcs with latency $1/x$ of segment R_i. For brevity, when we refer from now on to *arcs*, we mean the arcs of which the latency function is not identically zero, i.e., arcs with latency $1/x$.
Node 1 is the source s, and node $x + 1$ is the sink t. Now any feasible action of a player consists of at least one arc from each segment $R_i, i \in [x]$. This example is shown in Figure 3.

In the remainder of the section, we say that in a state $(A_{<i}, i)$, an arc e is *free* if no player in $[i - 1]$ has chosen e in her action, i.e., there does not exist an $i' \in [i - 1]$ such that $e \in A_{i'}$.

Optimal social cost of Γ_x. In the optimal outcome A^*, each player chooses exactly one arc from each segment, and players share arcs as little as possible.

Straightforward counting based on the above definitions yields that the optimal social cost is $C(A^*) = p\sqrt{x} + 3x^2 + (2x^2)2 = p\sqrt{x} + 7x^2$.

Definition of strategy profile S for Γ_x. In order to describe our worst-case subgame perfect equilibrium strategy, we first define the following actions, relative to the state in which a player must choose her action:

- *Greedy*: In each segment, choose the single arc chosen by the fewest number of players. In case of ties, the tie-breaking rule as described below is used.
- *Punish*(j) (for $j \in [n]$): Denote by R a segment where all arcs chosen by player j are chosen by less than x players from $[j]$. Denote by e an arc from R that is chosen by the largest number of players among the arcs chosen by j (breaking ties in a consistent way). The action *Punish*(j) is then defined as choosing e in R, and any free arc in each other segment.
- *Fill*: Choose $\sqrt{x}$ free arcs in each segment.
- *Copy*: Choose exactly the same arcs as the previous player.

Note that the above actions are defined relative to a given state in the game. The actions *Greedy* and *Copy* are well-defined for each state, while the actions *Punish*(j) and *Fill* only exist for a subset of the states.

Using these actions, we now define our subgame perfect equilibrium $S = (S_1, \ldots, S_n)$ for Γ_x. For each state $(A_{<i}, i)$, strategy S_i prescribes to play an action $S_i(A_{<i})$, which is determined as follows.

if every player $j \in [i-1]$ plays according to S_j **then**
  if i has at least $5x^2$ successors **then**
    if i is the first player, or if the previous $\sqrt{x}-1$ players chose *Copy* **then**
      Fill
    else
      Copy
  else
    Greedy
else
  if exactly 1 player $j \in [i-1]$ does not play according to S_j **then**
    if j has chosen less than x^2 arcs in each segment **then**
      if S_j prescribed j to choose *Fill* or *Copy* **then**
        if there exists a segment such that all arcs e chosen by j contain less than x players in total **then**
          Punish(j)
        else
          Greedy
  else
    Greedy

Tie-breaking rule:
When the strategy S_i prescribes that a player i chooses an arc chosen by the smallest number of players, and a set E' of multiple arcs have this property, the following tie-breaking rule is used: All predecessors of i are ordered. The set of all players that deviate from S comes first in this ordering. After that comes the

set of all other players. Within these two sets, the players are ordered by index from high to low. Now the arcs are ordered as follows: Arc e is ordered before e' iff the set of players on e is lexographically less than the set of players on e' according to the ordering on the players just defined. Finally, ties are broken by choosing the first arc in this order, among the arcs in E'.

Example 1. As an example to clarify the tie-breaking rule, consider the following situation: Say player 5 has to choose 2 arcs among arc set $\{a, b, c, d\}$, which are chosen by the smallest number of players. Players 1 and 3 have deviated from S. Player 1 has chosen (among arcs $\{a, b, c, d\}$) arcs b and c, player 2 has chosen arcs c and d, player 3 has chosen arcs a and d, and player 4 has chosen arcs a and b. Thus, the players are ordered $3, 1, 4, 2$ and the arcs are ordered d, a, c, b, so player 5 chooses arcs a and d. ◁

It is straightforward to see that S is a well-defined sequential strategy profile, i.e., whenever any of the actions *Greedy*, *Copy*, *Fill*, or *Punish*(j) is prescribed by S, it is possible for a player to choose this action. An extensive discussion of this fact is deferred to the full version of the paper.

Social cost of S. If each player i chooses the action prescribed by S_i, then the social cost is at least $(p\sqrt{x})(\sqrt{x}\sqrt{x}) + 3x^2 + (2x^2)2 = p\sqrt{x}x + 7x^2$. We see that $\lim_{p\to\infty} C(t)/C(s^*) = \lim_{p\to\infty}(px\sqrt{x} + 7x^2)/(p\sqrt{x} + 7x^2) = x$.

Checking that S is a subgame perfect equilibrium. For a state $(A_{<i}, i)$, an action A_i is said to be *subgame perfect with respect to a sequential strategy profile S* iff choosing A_i minimizes i's cost when players 1 to $i-1$ play $A_{<i}$, and players $i+1$ to n play according to S.

We now show that S is a subgame perfect equilibrium. This is done by showing that for any state $(A_{<i}, i)$, action $S_i(A_{<i})$ is subgame perfect with respect to S.

Lemma 1. *For each state $(A_{<i}, i)$ of Γ_x, action $S_i(A_{<i})$ is subgame perfect with respect to S.*

Proof. For each of the possible actions *Greedy*, *Fill*, *Punish*(j) (where $j \in [i-1]$), and *Copy*, that S_i may prescribe to player i in state $(A_{<i}, i)$, we prove that deviating from this prescription will not decrease the cost of player i, on the assumption that all succeeding players $i+1, \ldots, n$ play according to S.

- Suppose player i is prescribed by S_i to play *Fill* or *Copy*. Then no player in $[i-1]$ has deviated from S. Therefore, (assuming that all succeeding players play according to S as well) the cost of player i when she does not deviate is x. If player i does deviate, then the subsequent players will play *Punish*(i), which makes sure that in each segment one of the arcs chosen by i gets chosen by at least x players. Her utility will therefore be at least x. Thus, deviating is not beneficial for player i.
- Suppose player i is prescribed by S_i to play *Greedy*. Then (assuming that players $i+1, \ldots, n$ all play according to S) observe that by definition of S, players $i+1, \ldots, n$ play *Greedy*, even if player i deviates from playing *Greedy*.

We denote by A^* the outcome that results if i does not deviate from S_i. We show that if i does deviate, then in each segment, i's costs at least as high as in A^*. Let $j \in [x]$ and consider segment R_j. Let e_i and e_n denote the arcs from R_i chosen by respectively player i and player n in A^*. Denote by R^* the set of arcs in R_j chosen by players $i, \ldots, n$ in A^*.

We denote by c the latency of e_n in A^*. Any arc $e \in R^*$ has latency either c/x or $(c-1)/x$. (If it were higher, then the last player who chose e would have chosen e_n, because she plays greedily.) Specifically the latency of e_i is at most c/x. Also, any arc $e \in R_i$ that is not in R^* is chosen by at least $c-1$ players of $[i-1]$. (If this were false, then in A^* player n would have chosen e instead of e_n.)

Now consider outcome A' which occurs when player i deviates from S_i. If player i chooses any arc e'_i that is not in R_*, then this arc has latency at least c/x. We now show that if e'_i is in R^*, then it has latency at least c/x as well. In that case, if any player $i' \in \{i+1, \ldots, n\}$ chooses an arc not in R^* then all arcs in R^* would yield cost at least c/x. (Because, if there would be an arc $e' \in R^*$ with cost $(c-1)/x$, then the tie breaking rule dictates that i' would have chosen e'_i instead of e'.) However, if all players $i, \ldots, n$ choose an arc in R^*, then player n has cost at least c/x. Combining this with the tie-breaking rule, we conclude that e'_i has a latency of at least c/x as well. Therefore, in all cases the costs of player n' do not decrease by deviating.

– Suppose player i is prescribed by S_i to play *Punish*(j) for some $j \in [i-1]$. Let us compute first the cost of i if she would follow this prescription (assuming that players $i+1, \ldots, n$ all play according to S). Then observe that by definition of S, there is a number of other players succeeding i that play *Punish*(j) as well. Let k be this number of players. So: $\{j+1, \ldots, i+k\}$ is the set of players that play *Punish*(j). Let $\ell = |\{j+1, \ldots, i+k\}|$. Players $\{i+k+1, \ldots, n\}$ play *Greedy*, again by definition of S. Players in $[j-1]$ together occupy at most $j-2+\sqrt{x}$ arcs in each segment. Player j occupies at most x^2 arcs in each segment. Players $j+1, \ldots, i+k$ all choose *Punish*(j), so they each occupy 1 arc per segment. The total number of arcs occupied per segment by players in $[i+k]$ is therefore $j-2+x^2+\ell+\sqrt{x}$. Therefore, there are at least $F := (p-1)\sqrt{x}+3x^2-j-\ell+2$ free arcs per segment after the first $i+k$ players have chosen their action. The set $i+k+1, \ldots, n$ is of size $G := p\sqrt{x}+5x^2-j-\ell$. We see that $G/F \leq 2$ so the *Greedy* players will choose only those free arcs. (I.e., by the tie-breaking rule the *Greedy* players will not choose arcs of player i). Therefore, player i's utility is exactly $2-1/x$ if she plays *Punish*(j). (This holds because in $x-1$ segments, i chooses 1 free arc that will not be chosen by any of her successors as we have shown. In the remaining segment, i chooses an arc that player j has chosen, which will be chosen by precisely x players.)

Suppose next that i deviates from playing *Punish*(j). In that case, all succeeding players will play *Greedy*. We prove that in each segment, i's costs are at least $2/x$, so that her total cost is at least 2. All players in $[j-1]$ together occupy at least $j-1$ arcs per segment. This implies that in state

$(A_{<i}, i)$ in each segment there are at least $j-1$ occupied arcs and at most $p\sqrt{x} + 4x^2 - j + 1$ free arcs. The number of players succeeding i is $p\sqrt{x} + 5x^2 - i \geq p\sqrt{x} + 4x^2 - j + x$, where the inequality holds because $i \leq j + x^2 - x$ (because by the definition of S, there are at most $x^2 - x$ players choosing *Punish*(j)). Therefore, there exist players among the *Greedy* players who choose in each segment an arc that is occupied by at least one player. The tie-breaking rule for the *Greedy* action then makes sure that the first such a *Greedy* player chooses in each segment an arc on which i is the sole player, in case such an arc exists. Therefore, when i deviates, her cost in each segment is at least $2/x$.

□

It follows from the lemma above that S is a subgame perfect equilibrium of Γ_x. The proof of Theorem 2 therefore follows directly from all of the above.

Although the *SPoA* is not bounded by any constant, it is not hard to see that it is trivially upper bounded by the number of players n. In fact our construction shows a lower bound of $SPoA \geq \Omega(\sqrt{n})$. To see this, we choose $p = x\sqrt{x}$. Then $n = x^2 + 5x^2 = 6x^2$ which yields $x = \sqrt{n/6}$. Now, $SPoA \geq (x^3 + 7x^2)/(x^2 + 7x^2) \geq x^3/(8x^2) = x/8 = \sqrt{n}/(8\sqrt{6})$. There may exist a choice of p which yields an even worse lower bound.

4.1 The price of anarchy

In this section we focus on the regular (i.e., non-sequential) price of anarchy of symmetric network routing games with linear latencies, and show that it equals 5/2. This resolves an open question regarding the price of anarchy of congestion games [4]. Surprisingly, the lower bound that we provide is conceptually simpler than the one previously provided for the more general class of (non-network) affine congestion games [6].

Theorem 3. *The price of anarchy of symmetric linear and affine network routing games is* 5/2.

Proof sketch. It is well known that the price of anarchy of the more general class of affine (non-symmetric) congestion games is 5/2 [3,6]. Thus, it suffices to prove that the price of anarchy of symmetric linear network routing games is at least 5/2.

To this end we construct the following family of instances. For 3 players, the instance (along with the optimal and equilibrium strategies) is depicted in Figure 4. In general, let n be the number of players and consider an instance in which there are n *principal* disjoint paths from the source s to the sink t. These paths are all composed of $2n-1$ arcs (and thus $2n$ nodes, s being the first and t being the last), so we denote by $e_{i,j}$ the j-th arc of the i-th path, for $i = 1, \ldots, n$ and $j = 1, \ldots, 2n-1$, and by $v_{i,j}$ the j-th node of the i-th path, for $i = 1, \ldots, n$ and $j = 1, \ldots, 2n$. There are $n \cdot (n-1)$ additional *connecting* arcs that connect these paths: there is an arc from $v_{i,2k+1}$ to $v_{i-1,2k}$ for $k = 1, \ldots, n-1$, where

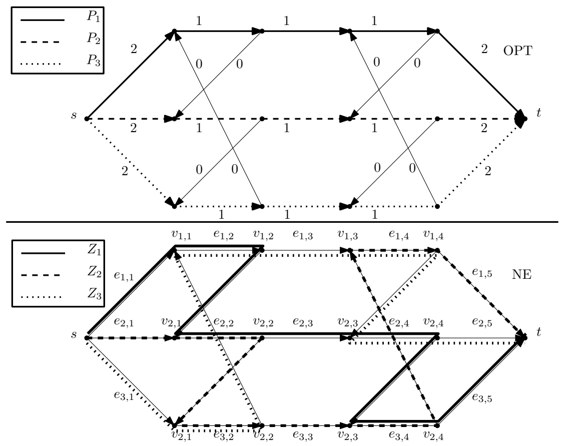

Fig. 4. A lower bound instance for the *PoA*. Players travel from s to t.

$i-1$ is taken $\mod (n)$. This defines the network. The latencies on the arcs are set as follows. Arcs $e_{i,1}$ (that start from s) have latency 2, arcs $e_{i,2n-1}$ (that end in t) have latency 2, while arcs $e_{i,j}$ with $1 < j < 2n-1$ have latency 1. All connecting arcs have latency zero.

It is easy to check that the optimal solution in this instance is to route one player in each of the principal paths, as demonstrated in the top part of Fig. 4. On the other hand, a Nash equilibrium arises if each of the players make use of the connecting arcs and use only a segment of at most three consecutive arcs on each principal path, as demonstrated in the bottom part of Fig. 4. Straightforward calculations then show that this Nash equilibrium has a social cost that is $5/2$ times worse than the optimal social cost, when we take n (i.e., both the number of principal paths and the number of players) to infinity. $\square$

5 Discussion and open problems

The central result of this paper states that the sequential price of anarchy is unbounded for symmetric affine network routing games. One property that stands out in our constructions is that they admit multiple subgame perfect equilibria. In fact, there even exists a subgame perfect equilibrium that induces an optimal strategy profile, and the existence of a poorly performing subgame perfect equilibrium relies crucially on tie breaking: Whenever a player is indifferent between two strategies, we essentially let the player choose the strategy that results in the worst social welfare. However, if we consider generic games, i.e., admitting a unique subgame perfect equilibrium, we do not know whether the sequential price of anarchy can be made arbitrarily high. A closely related problem is to derive the *sequential price of stability* of symmetric linear network routing games.

As for our bound on the (regular) price of anarchy: We emphasize that the existing upper bound of 5/2 for general affine congestion games holds even for coarse correlated equilibria, which contains the sets of pure, mixed, and correlated equilibria. Therefore, our last result on the price of anarchy implies that also for symmetric affine network routing games, the price of anarchy for mixed, correlated, and coarse correlated equilibria is 5/2. An open problem is to characterize the pure price of anarchy for symmetric network affine congestion games on *undirected* graphs.

Acknowledgments. We thank Mathieu Faure for stimulating discussions and particularly for pointing out a precursor of the instance depicted in Figure 2. We thank Marco Scarsini and Victor Verdugo for discussions on the price of anarchy of the symmetric atomic network game. We also thank Éva Tardos for allowing us to (partially) recycle their catchy paper title.

References

1. A. Angelucci, V. Bilò, M. Flammini, and L. Moscardelli. On the sequential price of anarchy of isolation games. In *Proceedings of the 19th COCOON*, 17–28, 2013.
2. E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th FOCS*, 295–304. IEEE, 2004.
3. B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th STOC*, 57–66, 2005.
4. K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: the price of anarchy, universal worst-case examples, and tightness. *ACM Transactions on Economics and Computation*, 2(4), Article 14, 2014.
5. V. Bilo, M. Flammini, G. Monaco, and L. Moscardelli. Some anomalies of farsighted strategic behavior. In *Proceedings of the 10th WAOA*, 229–241, 2013.
6. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th STOC*, 67–73, 2005.
7. J. de Jong, M. Uetz, and A. Wombacher. Decentralized throughput scheduling. In *Proceedings of the 8th CIAC*, 134–145, 2013.
8. J. de Jong and M. Uetz. The sequential price of anarchy for atomic congestion games. In *Proceedings of the 10th WINE*, 429–434, 2014.
9. E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th STACS*, 404–413, 1999.
10. D. M. Kreps and R. B. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.
11. H. W. Kuhn. Extensive games and the problem of information. *Annals of Mathematical Studies*, 28:193–216, 1953.
12. I. Milchtaich. Crowding Games are Sequentially Solvable. *International Journal of Game Theory* 27:501–509, 1998.
13. R. Paes Leme, V. Syrgkanis, and É. Tardos. The curse of simultaneity. In *Proceedings of the 3rd ITCS*, 60–67, 2012.
14. M. J. Osborne. *An Introduction to Game Theory.* Oxford University Press, 2003.
15. R. Selten. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit: Teil 1: Bestimmung des dynamischen Preisgleichgewichts. *Zeitschrift für die gesamte Staatswissenschaft*, 121(2):301–324, 1965.