

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

The Inefficiency of Nash and Subgame Perfect Equilibria for Network Routing

José Correa

University of Chile, Santiago, Chile, correa@uchile.cl

Jasper de Jong

University of Twente, Enschede, The Netherlands, j.dejong-3@utwente.nl

Bart de Keijzer

University of Essex, Colchester, United Kingdom, b.dekeijzer@essex.ac.uk

Marc Uetz

University of Twente, Enschede, The Netherlands, m.uetz@utwente.nl

This paper provides new bounds on the quality of equilibria in finite congestion games with affine cost functions, specifically for atomic network routing games. It is well known that the price of anarchy equals exactly $5/2$ in general. For symmetric network routing games, it is at most $(5n - 2)/(2n + 1)$, where n is the number of players. The paper answers to two open questions for congestion games. First, we show that the price of anarchy bound $(5n - 2)/(2n + 1)$ is tight for symmetric network routing games, thereby answering a decade-old open question. Secondly, we ask if sequential play and subgame perfection allows to evade worst-case Nash equilibria, and thereby reduces the price of anarchy. This is motivated by positive results for congestion games with a small number of players, as well as recent results for other resource allocation problems. Our main result is the perhaps surprising proof that subgame perfect equilibria of sequential symmetric network routing games with linear cost functions can have an unbounded price of anarchy. We complete the picture by analyzing the case with two players: We show that the sequential price of anarchy equals $7/5$, and that computing the outcome of a subgame perfect equilibrium is NP-hard.

Key words: Nash and subgame perfect equilibrium, price of anarchy, congestion games, network routing

MSC2000 subject classification: Primary: 91A05, 91A06, 91A10, 91B32, 91B52; Secondary: 05C57, 90B10

OR/MS subject classification: Primary: Games/group decisions > noncooperative, Transportation >

Models > Network, Decision Analysis > Sequential; Secondary: Mathematics > combinatorics

History:

1. Introduction The concept of the price of anarchy (PoA), introduced by Koutsoupias and Papadimitriou [30, 31], has spurred a significant amount of research over the past two decades and has contributed to establishing the area algorithmic game theory. In its original version, one relates the outcome of a worst-case Nash equilibrium of a noncooperative game with selfish players to that of an optimal, centrally computable solution, using as metric some global objective function defined on the set of possible outcomes of the game. It is well known that the lack of central coordination, translated into solutions that arise as Nash equilibria, can lead to substantial efficiency losses. This remains true even if the global objective function is the sum of the objective functions of all players.

In their highly influential work, Roughgarden and Tardos [39] prove that the price of anarchy for “non-atomic” network routing games with affine cost functions is exactly $4/3$. The equilibrium concept in non-atomic network routing is Wardrop’s equilibrium, meaning that only minimum cost paths can have nonzero flow [41]. Wardrop equilibria possess a simple characterization through variational inequalities that has been extensively used to derive price of anarchy results, e.g. [14]. Furthermore, they can be thought of as the limiting case of a Nash equilibrium where the total flow is routed by infinitely many players, each responsible for an infinitesimally small amount of flow. This variational inequality approach was shortly after used to pin down, to $5/2$, the price of anarchy of the “atomic” counterpart of the network routing problem [7, 6, 12].

This paper focuses on finite congestion games with affine cost functions, and specifically its special case, atomic network routing with linear cost functions. A congestion game consists of set of n players, and a set of m resources, and each player is equipped with a feasibility system of subsets of the resources that describes the feasible choices that are available for that player. The cost of a resource is an affine, non-negative function that depends on the number of players that choose that resource. A solution consists of a choice of a feasible subset for each of the players, and the total cost of a player is the sum of the costs of all chosen resources. The total cost of the system is the sum of the costs of all players. In the special case of atomic network routing, the resources are arcs in a digraph, and each of the n players is interested in establishing one origin-destination path. The game is symmetric if all players have the same feasibility system, and translated into network routing, that means that all players have the same origin and destination. For general congestion games with affine cost functions, the price of anarchy (PoA) is known to be equal to $5/2$ for $n \geq 3$ players [7, 6, 12], and it is at most $(5n - 2)/(2n + 1)$ if the game is symmetric [12].

This paper first closes a gap in the literature by showing that $(5n - 2)/(2n + 1)$ is indeed the exact price of anarchy for symmetric network routing games. The major part of this paper, however, is dedicated to the question what effect sequential decision making has on the quality of the resulting equilibria. Let us next elaborate on the relevance of this question.

One motivation to address the effect of sequential decision making lies in a set of recent publications, led by Paes Leme et al. [37]. They made the proposal to consider sequential versions of finite, n -player games, and to analyze subgame perfect equilibrium outcomes of such sequential games, instead of Nash equilibria of the simultaneous move strategic form games. Roughly speaking, subgame perfect equilibria represent the rational outcomes of a sequential game where all players act farsighted and strategically [40]. The driver behind this proposal is the insight that some of the Nash equilibria of the strategic form game cannot be realized as subgame perfect equilibria of the sequential version of the game. That leads to the notion of the *sequential price of anarchy*, introduced by Paes Leme et al. [37], which could be smaller than the price of anarchy. Similar to the price of anarchy (PoA) [30, 31], the sequential price of anarchy ($SPoA$, for short) 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 solution, the sequential price of anarchy considers the possible outcomes of a game where players choose their strategies sequentially, in some arbitrary (albeit fixed) order. It then relates the quality of the outcome of the worst possible subgame perfect equilibrium taken over all possible orders of the players to the quality of an optimal solution. Note that for games with perfect information which we consider here, subgame perfect equilibria can be computed via backwards induction. Although this is conceptually simple, it raises an important challenge when compared to the understanding of Nash or Wardrop equilibria, the main difficulty being the absence of a simple characterization such as the variational inequality mentioned earlier.

It turns out that there are interesting examples of games where sequential play and subgame perfection can lead to substantially improved guarantees on the price of anarchy, and in this sense allow to avoid what has been called the “curse of simultaneity” [37]. Indeed, for a handful of games, the $SPoA$ has been proven to be lower than the PoA [37, 16, 15]. For some others, however, this turns out to be opposite [2, 1, 11, 10].

The motivation to analyze the sequential price of anarchy also for congestion games, and specifically for atomic network routing games, is two-fold.

First, de Jong and Uetz [15] considered general congestion games with affine cost functions and showed that the sequential price of anarchy ($SPoA$) is equal to 1.5 for two players, and equal to $2 + \frac{63}{488}$ for three players. That means in particular that worst-case Nash equilibria are not realizable by farsighted players in sequential play. In [15] it is further conjectured that the $SPoA$ for affine congestion games is below $5/2$, for any number of players. Yet the only known upper bound on the $SPoA$ for an arbitrary number of players is -trivially- the number of players, n .

Second, in practical situations it often seems that sequential decision making is a more realistic model in comparison to simultaneous single shot games. This is certainly true for applications such as traffic and load balancing, where autonomous decisions are typically taken on top of an

existing and known load of the resources, but future decisions of other players are to be expected. The model that we consider is an abstraction of such concrete applications, and specifically asks what the effect is if all n players act sequentially and fully farsighted. We believe it is important to understand the effect of sequential play and subgame perfection in this basic setting.

Our results. As a first result, this paper answers a question that was posed, among others, by Christodoulou and Koutsoupias [12] and Bhawalkar et al. [8, 9], and that had remained an open problem for more than a decade. Indeed, it was open if the price of anarchy bound $5/2$ can also be attained (asymptotically) with *symmetric* network routing games, and if the upper bound $(5n - 2)/(2n + 1)$ of Christodoulou and Koutsoupias [12] is tight. We close this gap, by deriving a lower bound equal to $(5n - 2)/(2n + 1)$ for symmetric, atomic network routing games with linear cost functions. We thereby exactly match the upper bound that was known for the case of symmetric (not necessarily network) congestion games [12]. This asymptotically matches the upper bound $5/2$ that is known for $n \geq 3$ players for the asymmetric network routing case [6, 12].

As to the quality of equilibria for sequentially played congestion games, we show the following. First, we pin down the exact sequential price of anarchy for the symmetric network routing case with $n = 2$ players, showing that it equals $7/5$. This constitutes an improvement over the $3/2$ bound that is known for the more general, non-symmetric case [15]. For a large number of players however, and in sharp contrast to what was conjectured in [15], we show that congestion games, and even symmetric network routing games, have an unbounded sequential price of anarchy. Indeed, we prove that even in the symmetric network routing case, i.e., when all players share the same origin and destination, the sequential price of anarchy can be as large as $\Omega(\sqrt{n})$. In view of earlier results that compare Nash to subgame perfect equilibria [37, 16, 15], this establishes an unexpectedly sharp separation between Nash and subgame perfect equilibria for network routing games. Subsequent recent work by Groenland and Schäfer [21] suggests that the presence of ties seems to be pivotal for the existence of such worst case instances.

The crucial part of the proof of our main result is to come up with a “contingency plan of actions” for every player and every possible profile of actions of all previous players that indeed leads to a subgame perfect outcome. This is generally 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 can be achieved.

To some extent this construction may seem somewhat unrealistic from a practical viewpoint, yet the main point is to show that such a construction is possible *at all* for atomic congestion games.

We finally also address the computational complexity of subgame perfection, showing that the computation of the outcome of a subgame perfect equilibrium is an NP-hard problem. Although we know from Paes Leme et al. [37] that computing subgame perfect equilibria can be PSPACE-hard in general congestion games, that reduction requires a non-constant number of players. Our result shows that the problem remains at least NP-hard for symmetric network routing games, and even when the number of players is only two. This contrasts with the fact that a pure Nash equilibrium can be efficiently computed in these games through a min cost flow computation [17]. Interestingly, our analysis also shows that outcomes of subgame perfect equilibria are generally not Nash equilibria of the simultaneous single shot game, as opposed to e.g. *crowding games* as studied by Milchtaich [34]. Conditions under which this *is* actually the case for network routing games have recently been given by Groenland and Schäfer [21].

Outline. We first discuss in Section 2 the related literature with respect to the sequential price of anarchy. We then formally introduce the model along with some necessary preliminary definitions and insights in Section 3. Section 4 addresses the (non-sequential) atomic network routing game, and shows that the regular price of anarchy PoA equals $(5n - 2)/(2n + 1)$ when the game is symmetric. In Section 5, we start discussing the sequential network routing game for the case when there are only $n = 2$ players. We prove that the sequential price of anarchy is precisely $7/5$, and show that computing an outcome from a subgame perfect equilibrium is NP-hard. In Section 6, we prove the main result of this paper, namely that the sequential price of anarchy $SPoA$ of symmetric network routing games with linear cost functions is generally unbounded. In Section 7, we conclude with some remarks and open problems.

2. Related work on the sequential price of anarchy. Network routing games and congestion games form a cornerstone class of problems in the field of algorithmic game theory. Many variants and extensions on these games have been proposed, and the body of literature on these games is so large that it is impossible to survey it or summarize its most important articles in this section. We refer to the recent survey by Fotakis [19], which discusses various results on the price of anarchy and the price of stability, including more general variants such as weighted congestion games, nonlinear cost functions, etc. In the remainder of this section, we focus on recent work on sequential games and subgame perfect equilibria, as the present paper falls most naturally into this line of research.

The idea to consider subgame perfect equilibria of sequential games as an alternative to rule out some potentially “unrealistic” Nash equilibria was initiated by Paes Leme et al. [37]. They

introduced the notion of the sequential price of anarchy $SPoA$ and proved it to be relevant for three classes of games: In consensus games, players choose one of two alternatives, and gain utility by picking the same alternative as other players. The authors show that consensus games have a $SPoA$ of 1 as opposed to an unbounded PoA . In machine cost sharing games with fair cost allocation, each player chooses a set of machines, and players distribute the cost of each machine equally. The authors show that the $SPoA$ is $\Theta(\log n)$, an improvement over the PoA which is equal to n for these games (where n is the number of players). In unrelated machine scheduling games, each player controls a job and needs to choose a machine to run the job on. Processing times of the jobs are machine dependent, and jobs are interested in the total load of the selected machine. Unrelated machine scheduling games have an unbounded PoA but the $SPoA$ lies in between $\Omega(\sqrt{n})$ and $m \cdot 2^n$, where m is the number of machines. The paper [37] also contains the result that subgame perfect equilibria are PSPACE-hard to compute in general.

Later, Bilò et al. [11] showed that the improvements obtained by Paes Leme et al. [37] should be relativized, as they only hold for the class of generic games¹. Specifically, in the non-generic case, consensus games as well as the related class of cut games have a $SPoA$ of 3, while machine cost sharing games with fair cost allocation have a $SPoA$ between $\Omega(n + 1 - H_n)$ and n , and unrelated machine scheduling games have a $SPoA$ between $2^{\Omega(\sqrt{n})}$ and 2^n . Note that the latter upper bound is also an improvement for the generic case.

Closest to our work is the study of congestion games with few players by de Jong and Uetz [15]. That paper triggered the question on the actual value of the price of anarchy for atomic congestion games that is ultimately answered here. The authors of [15] show that for two players the $SPoA$ equals $3/2$ and for three players it equals $2 + 63/488 \approx 2.13$. Hence, in those cases the $SPoA$ is less than the PoA which equals 2 and $5/2$, respectively [12]. It was conjectured in [15] that the $SPoA$ for affine congestion games is below $5/2$, but Kolev [29] showed that when the number of players is four or more, the $SPoA$ already exceeds 2.5.

We conclude this introductory section with a brief account of models for which the effect of sequential play on the quality or existence of equilibria has been analyzed.

- A set of authors that first addressed the effect of sequential decisions in network routing games are Olver [35], Harks et al. [24], Harks and Vegh [25], as well as Farzad et al. [18]. Without going into too much detail, their models differ in that the cost of a given player depends on preceding players but not on future players, just like in traffic situations where players enter streets in sequence, and a player only influences players that are behind.

¹In a generic game the players are never faced with a tie, for any strategy profile of the other players. Thus all decisions are uniquely determined and in particular uniqueness of subgame perfect equilibria is guaranteed.

- Hassin and Yovel [26] consider machine scheduling and focus on the special case of related machines. The authors prove that for this model, the $SPoA$ is at most 2 while the PoA is $\Theta(\ln n / \ln \ln n)$. Moreover, the $SPoA$ decreases to $4/3$ when each machine always executes the shortest jobs first. Their analysis relies on well known approximation results for list scheduling by Graham [20].
- The effect of sequential play was also analyzed for a *throughput scheduling* model by de Jong et al. [16]. There, each player owns a set of machines, and players select jobs with individual values that must be scheduled in given time windows. The authors show that for identical machine scheduling, the $SPoA$ is $e/(e-1)$, an improvement over the PoA , which is 2.
- Avni et al. [5] study another variation of congestion games. The difference is that in their model, players make multiple, incremental decisions, and each time choose a single resource that extends the current solution. These decisions are in scheduled phases, and within each phase, the players make decisions in $k \leq n$ subsequent turns, where at least one player makes a choice in each turn. In some sense, this model mixes sequential and simultaneous decision making, and the main results of [5] are on the existence and non-existence of (pure strategy) subgame perfect equilibria.
- Rahn and Schäfer [38] study tree graph coordination games, where players are nodes in a graph. Each player chooses a color, and obtains utility equal to the number of neighbors sharing her color. The authors prove that the $SPoA$ is 2 while the PoA is unbounded.
- Yet another context is that of item bidding games. Here players bid on a set of public projects, aiming to have their personal preferred project selected. Lucier et al. [33] prove, among other things, that even though the PoA can be as large as $n-1$, any subgame perfect equilibrium in the sequential version is optimal, yielding a $SPoA$ of 1.
- Isolation games that were studied by Angelucci et al. [1, 2]. These are games in which players have to choose a point in a metric space, and are interested in maximizing the distance to any other player. The authors show that, depending on the definition of the players' utility functions (sum or minimum distance) and social cost (utilitarian or egalitarian), the $SPoA$ is higher than the PoA in some cases and lower in other cases.
- Variants of sequential matching games were studied by Haeringer and Wooders [23] and the very recent paper of Kawase et al. [27, 28]. Haeringer and Wooders [23] characterize Nash and subgame perfect equilibria in a game where firms sequentially propose workers to be assigned to them, after which a second stage starts in which the workers make their acceptance and rejection decisions. Kawase et al. [27] studies the problem to compute a subgame perfect equilibrium in a similar type of game where the difference is that a worker immediately makes his decision after each offer. They show various tractability results and PSPACE-hardness results.

- Subsequent to this paper, recent work by Groenland and Schäfer [21, 22] suggest to bridge the gap between far-sighted behaviour and myopic greedy best response behaviour. They assume that for some parameter k , each player chooses a best response with respect to the action profile that arises when the k subsequent players play subgame perfectly. The authors establish conditions under which the action profiles resulting from such behaviour correspond to Nash equilibria, thereby allowing to import known results for the price of anarchy PoA . In particular, for a special class of graphs their work shows that the absence of ties guarantees that subgame perfect equilibria are also Nash equilibria, and hence have a bounded sequential price of anarchy.

3. Model and notation. Throughout the paper, we consider a special class of finite, atomic *congestion games*, namely symmetric atomic network routing games with linear cost 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 cost (or latency) function $c_e(y) = a_e y$ with linear cost coefficient a_e . That is, the cost of an arc equals a_e times the number of players using it. There are n players that all want to travel from the same origin s to destination t , so that the possible actions of all players consist of all directed (s, t) -paths in G . Hence, all players have the same set of actions and the game is symmetric. We will denote by m the number of arcs $|E|$. We refer to the path that a player i chooses the *action* A_i of player i , and $A = (A_1, \dots, A_n)$ is a vector with one action for each player, also called 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 $A = (A_1, \dots, A_n)$, 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 $a_e n_e(A)$, and therefore the cost for player i , choosing path A_i , is defined as

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

For the sake of strengthening some of our results, sometimes we assume that the cost functions are *affine* instead of linear². Affine cost functions are of the form

$$c_i(A) = \sum_{e \in A_i} (a_e \cdot n_e(A) + b_e),$$

where $b_e \geq 0$ can be thought of as a player-independent activation cost associated to an arc $e \in E$.

² Note that our main results are lower bounds on the price of anarchy and they already apply under linear cost functions.

We define the *social cost* of an action profile A as $C(A) = \sum_{i=1}^n c_i(A)$, i.e., the sum of the costs of the players. In other words, we consider a *utilitarian* social cost function. This is a standard model, yet note that it differs from the *egalitarian* makespan objective as studied, e.g., in [30, 31].

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 [31, 30] 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 ,

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

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

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 [32, 40]. 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 $1, 2, \dots, n$, so that the i -th player must choose action A_i , observing the actions of players preceding i , but of course not knowing the actions of the players succeeding her. However, since players are fully rational and fully informed, at equilibrium they anticipate the successors' behavior and therefore make optimal strategic choices by fully anticipating all followers' actions.

A strategy S_i then specifies for player i the full contingency plan of actions she would choose for each potential choice of actions $A_{<i} := (A_1, \dots, A_{i-1})$ 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 $1, \dots, i-1$. We refer to a choice of strategies $S = (S_1, \dots, S_n)$ 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 [40], are defined as strategy profiles S that induce 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 $1, \dots, i-1$, player i cannot decrease her cost by switching to an action different from $S_i(A_{<i})$, in the *subgame* where the actions of $1, \dots, i-1$ are fixed to $A_{<i}$, and players $i+1, \dots, n$ play strategies (S_{i+1}, \dots, 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




FIGURE 1. Game tree for a symmetric sequential two player game. 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 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 is (A_{12}, A_{22}) , the rightmost path of the game tree.

action that will minimize their individual cost. Analogous to (1), the sequential price of anarchy of an instance I is defined by

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

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 [37] 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 level i representing the possible states of the game that a single player i 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 $1, \dots, i-1$ 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 can be computed by backwards induction. In particular, it is known that subgame perfect equilibria always exist; see e.g. [36]. 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 discussed in the next section.

4. The price of anarchy for symmetric network routing In this section we first focus on the regular price of anarchy PoA as defined in (1) for symmetric network routing games with linear cost functions. We show that it equals $(5n-2)/(2n+1)$, n being the number of players. This resolves an open question regarding the price of anarchy of network congestion games of Christodoulou and Koutsoupias [12], see also [8]. Indeed, it is well known that the price of anarchy of the more general class of linear symmetric (not necessarily network) congestion games is at most $(5n-2)/(2n+1)$ [12]. It was not clear, however, whether this bound is can also be attained with




FIGURE 2. A lower bound instance for the PoA for $n = 3$. Players travel from s to t .

symmetric linear network routing games. Here we settle that issue and give a matching lower bound of $(5n - 2)/(2n + 1)$. Surprisingly, the lower bound is conceptually simpler than the one previously provided for the more general class of (non-network) affine congestion games [12].

THEOREM 1. *The price of anarchy of symmetric linear network routing games with n players equals $(5n - 2)/(2n + 1)$.*

Proof. We construct the following family of instances. For 3 players, the instance (along with the optimal and equilibrium strategies) is depicted in Figure 2. In general, let n be the number of players and consider an instance in which there are n disjoint paths from the source s to the sink t , which we refer to as *principal* paths. 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 $i = 1, \dots, n$ and $j = 1, \dots, 2n - 1$, and by $v_{i,j}$ the j -th node of the i -th path, for $i = 1, \dots, n$ and $j = 1, \dots, 2n$. 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 $k = 1, \dots, n - 1$, where $i - 1$ is taken mod (n) . This defines the network. The cost coefficients on the arcs are set as follows. Arcs $e_{i,1}$ (that start from s) have cost coefficient $a_e = 2$, arcs $e_{i,2n-1}$ (that end in t) have cost coefficient $a_e = 2$, and arcs $e_{i,j}$ with $1 < j < 2n - 1$ have cost coefficient $a_e = 1$. All connecting arcs have cost coefficient $a_e = 0$.

It is easy to check that the optimal solution in this instance is to route one player in each of the principal paths. Since in this solution no two players intersect in any arc, its cost can be computed

as $n(2 + 2n - 3 + 2) = n(2n + 1)$. On the other hand a Nash equilibrium is obtained when each player k follows the following path: she starts with arcs $e_{k,1}$, $e_{k,2}$, then uses all arcs of the form $e_{k+j,2j}$, $e_{k+j,2j+1}$, $e_{k+j,2j+2}$ for $j = 1, \dots, n - 2$, and finishes with arcs $e_{k+n-1,2n-2}$, $e_{k+n-1,2n-1}$, and uses all the required connecting arcs. Here, the additions on the principal paths index are taken mod (n) . In this solution every arc $e_{i,j}$ with j even is used exactly twice while every other arc is used only once. Thus the social cost is $n(n - 1) \cdot 4 + n(n - 2) \cdot 1 + 2n \cdot 2 = n(5n - 2)$. It immediately follows that the *PoA* of the instance is at least $(5n - 2)/(2n + 1)$.

The remainder of the proof consists in checking that the latter path choices indeed result in a Nash equilibrium. Consider a player, which by symmetry we may assume is player 1, and let us evaluate possible deviations from the current path which we call the zig-zag path Z . Note that in any path an arc $e_{i,2j-1}$ is always followed by $e_{i,2j}$ for $2 \leq j \leq n - 1$, and that player 1 evaluates the joint cost of these two arcs as either 5 (if not in the zig-zag path) or 3 (if in the zig-zag path). Now assume player 1 follows a path P not intersecting Z in arcs of the form $e_{i,2j-1}$ for $j = 1, \dots, n$, then the cost of this path is at least $(n - 2) \cdot 5 + 8 = 5n - 2$ (the extra 8 comes from the arc starting in s and the arc ending in t), and thus the deviation is not profitable. Therefore we may assume that P and Z do intersect in arcs of the form $e_{i,2j-1}$. Since these arcs are in Z they actually are of the form $e_{i,2i-1}$. Assume there are two such arcs, and consider two of these intersection arcs $e_{i,2i-1}$ and $e_{k,2k-1}$ such that there is no other intersection arc on the path segment of P that runs from $e_{i,2i-1}$ to $e_{k,2k-1}$. The cost of the restricted Z path between nodes $v_{i,2i}$ and $v_{k,2k-1}$ is $3 + 5(k - i - 1)$, whereas path P has to cost at least $5(k - i - 1)$ just to get to a node of the form $v_{l,2k-1}$ plus what it needs to pay to get to the principal path k , which is at least 3. The total cost is thus at least $3 + 5(k - i - 1)$ implying that the deviation P restricted to the segment between any two edges intersecting with Z is not profitable. Finally we consider the subpath between s and the first such intersection, say $e_{i,2i-1}$ (and symmetrically between the last and t). In this case the cost of the restricted Z path between nodes s and $v_{i,2i}$ is $6 + (i - 2)5$, whereas the cost of P is at least $4 + 3 + (i - 2)5$. Here, the 4 comes from the first arc and the 3 from the second arc. Again the deviation cannot be profitable. We thus conclude that Z is indeed a best response and thus we have a Nash equilibrium. \square

5. Sequential games and equilibria for $n = 2$ players As a way to illustrate the difficulties in arguing about subgame perfect equilibria in general, we first focus on the special case with only two players. Doing that, we point out two phenomena that clarify 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 is in contrast to the case of parallel links [15] and

also to the crowding games as studied by Milchtaich [34]. This effect is the major complication in analyzing subgame perfect equilibria, and it seems it can be circumvented if the underlying game is *generic*, meaning that no player is ever confronted with a tie; see [34, 11], and [21] for network routing games in particular.

Based on this particular instance we additionally prove that the sequential price of anarchy for the two player case equals $7/5$. This is smaller than the regular price of anarchy which is equal to $8/5$ [12], and also smaller than the sequential price of anarchy for the asymmetric network routing case, which equals $3/2$ [15].³ Secondly, we show that even in the two-player case, computing the outcome of a subgame perfect equilibrium is NP-hard.

5.1. The sequential price of anarchy Consider the two-player instance depicted in Figure 3, with five vertices and eight arcs. The vertices 1, 2, 3, 4, 5 are numbered from left to right and from top to bottom so that $s = 1$ and $t = 5$. The linear cost coefficients a_e are given by the numbers next to the respective arcs e .




FIGURE 3. Lower bound example for 2 players. Each arcs is labeled with the coefficients of the linear cost function associated to it. There are two parallel arcs emanating from s : The purpose of the arc with the higher coefficient is to cause player 2 to be indifferent between the two arcs in case player 1 uses the arc with the lower coefficient. This makes sure that there exists an SPE that results in the the action profile depicted on the right hand side.

There are two parallel arcs between vertices 1 and 2. In the discussion below, whenever we specify that a player takes a path that includes an arc from 1 to 2, that player uses the parallel arc with

³ In [15] a lower bound example is given for general congestion games which can be easily transformed to network routing games.

the lower of the two coefficients. It can be easily verified that the following is a subgame perfect equilibrium:

- 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, path $(1, 2, 3, 4, 5)$, while player 2 chooses the dotted path on the right, which is simply the arc going directly 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 instance, the subgame perfect equilibrium is not unique. However, the latencies can be slightly perturbed so that 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 change the cost coefficient of arc $(1, 2)$ from 2 to $2 + \epsilon$, that of arc $(1, 5)$ 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*⁴ 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 (not necessarily network) congestion games. It uses a proof technique based on linear programming, but is nonetheless different from the proof used in [15] where linear programming is also used to derive upper bounds on the *SPoA*. The main result of this section is the following.

THEOREM 2. *The sequential price of anarchy of symmetric linear network routing games and symmetric affine congestion games is $7/5$ when there are two players only.*

Proof. It follows from the example above that the *SPoA* for two-player linear network routing games is at least $7/5$, so it only remains to show that $7/5$ is also an upper bound. The proof goes by

⁴ Just like the price of stability as defined in [4, 3], the sequential price of stability is the ratio of the outcome of the best subgame perfect equilibrium over the optimum.

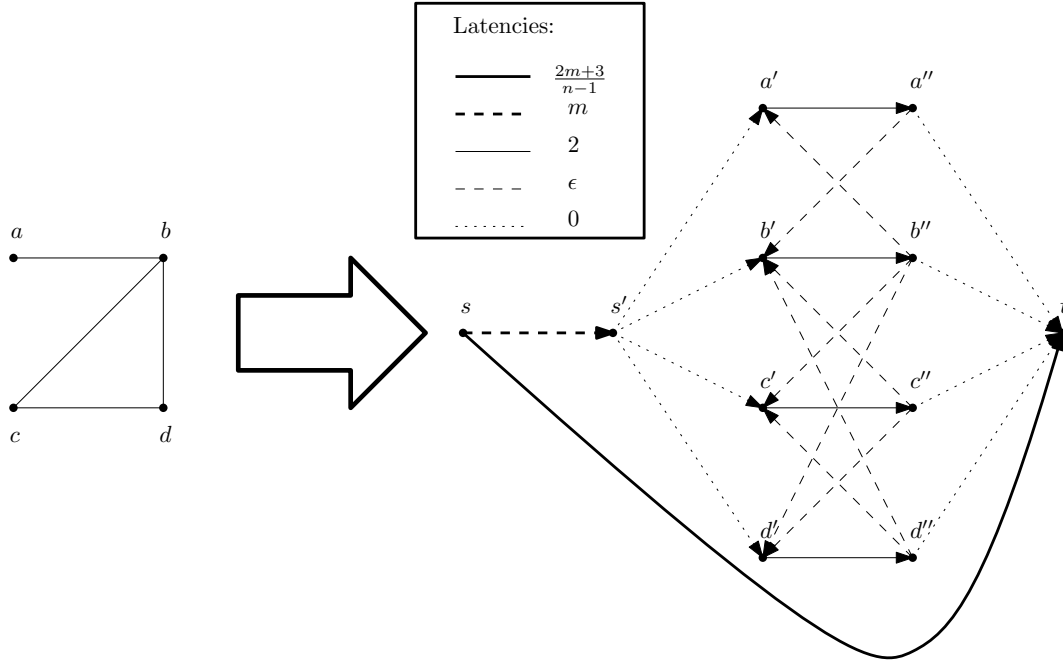
establishing a set of valid inequalities that must be fulfilled for any outcome of a subgame perfect equilibrium (A_1, A_2) and any outcome (A_1^*, A_2^*) that minimizes the total cost. After applying several combinatorial arguments, we arrive at the matching upper bound of $7/5$ via linear programming. The details of these arguments that complete the proof of Theorem 2, together with a combinatorial argument to show that the underlying linear program has $7/5$ as an optimal solution, can be found in Lemma 2 in the appendix. \square

5.2. 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 consists of just one path in the network G per player. The problem to compute such an equilibrium path in the game tree, however, turns out to be hard.

THEOREM 3. *Computing an action profile resulting from a subgame perfect equilibrium is (strongly) NP-hard for any number of players $n \geq 2$.*

Proof. We prove the theorem by a reduction from the Hamiltonian path problem, where the task is to determine whether a given undirected graph has a Hamiltonian path, i.e., a path that visits each vertex of the graph exactly once. Consider any instance of Hamiltonian path on undirected graph $G = (V, E)$, where we assume without loss of generality that G is connected and has at least one edge. We construct the following game: There are n players. There are two copies v', v'' for each node $v \in V$. There is also a source node s , a sink node t , and a node s' . We define $m = 2|V| + 1$, and $\epsilon = 1/(|E| - 1)$. There is an arc with cost coefficient m from s to s' , and an arc with cost coefficient $(2m + 3)/(n - 1)$ from s to t . For each $v \in V$, there is an arc with cost coefficient 0 from s' to v' , an arc with cost coefficient 2 from v' to v'' , and an arc with cost coefficient 0 from v'' to t . Moreover, for each arc $(u, v) \in E$, there is an arc with cost coefficient ϵ from u'' to v' . This reduction is shown in Figure 4.

We claim that in an outcome resulting from a subgame perfect equilibrium, player 1 chooses all arcs (v', v'') that correspond to all $v \in V$. If the graph is Hamiltonian she will choose these arcs exactly in the order of a Hamiltonian path, and otherwise will have to traverse at least one arc (v', v'') twice. Moreover, all subsequent players choose the arc (s, t) .

FIGURE 4. Reducing a Hamiltonian path instance to an n -player network routing game.

Let us argue that indeed, this is the outcome of a subgame perfect equilibrium. First note that if for at least one node v , player 1 does *not* choose arc (v', v'') , then there is some successor that will choose the path (s, s', v', v'', t) . Let us call this player j . Then j has a cost of at most $2m + 2$, because all other players will play (s, t) . The latter is due to the high cost of $3m$ that the third player would have if she would choose arc (s, s') (while choosing (s, t) would guarantee her a lower cost of at most $2m + 3$). So, in this case, player 1 has a cost of at least $2m + 2$.

On the other hand, if player 1 chooses all arcs (v', v'') , then for any succeeding player, choosing any path using an arc (v', v'') would yield her a cost of at least $2m + 4$. In that case, choosing (s, t) is always a better option for any succeeding player, because doing so guarantees that her cost is at most $2m + 3$.

Now, suppose there exists a Hamiltonian path in G . Let us look at the cost of player 1 if she chooses the arcs (v', v'') that correspond to all $v \in V$, in the order of the Hamiltonian path. Then player 1's cost is at most $m + 2|V| + \epsilon|V - 1| < 2m$, because, as we showed, succeeding players will choose (s, t) .

If no Hamiltonian path exists, then for player 1, any path that contains all arcs (v', v'') uses at least $|V|$ arcs with cost coefficient ϵ , since she will have to use at least one arc (v', v'') twice. This yields player 1 a cost of at least $m + 2|V| + \epsilon|V|$.

From the above we conclude the following: if there is a Hamiltonian path in G , then there is an action A_1 for player 1 that gives her a cost of $m + 2|V| + \epsilon|V - 1|$ when the succeeding players all play a subgame perfect strategy profile. Moreover, playing any action other than A_1 will give player

1 a strictly higher cost when the succeeding players all play a subgame perfect strategy profile. Therefore, if there exists a Hamiltonian path, player 1 plays A_1 in a subgame perfect equilibrium, and have a cost of $m + 2|V| + \epsilon|V - 1|$. If there is no Hamiltonian path in G , then in any subgame perfect equilibrium, player 1 has a cost of at least $m + 2|V| + \epsilon|V|$.

Hence, if we were able to compute the outcome of a subgame perfect equilibrium in polynomial time, we could verify in polynomial time if the cost of player 1 equals $m + 2|V| + \epsilon|V - 1|$ or not. This would allow to decide the Hamiltonian path problem in polynomial time. \square

Recall that membership in NP is not clear in general, because it is generally not clear how to verify subgame perfectness in polynomial time. However if we define a decision problem SPE-DEC that asks if the first player can guarantee herself a cost at most equal to some threshold k in a subgame perfect equilibrium, we can also show the following.

THEOREM 4. *Decision problem SPE-DEC is NP-complete for the case of two players.*

Proof. By Theorem 3, *SPE-DEC* is NP-hard, so we only need to prove that *SPE-DEC* is contained in NP. We use the subgame perfect action profile (A_1, A_2) as a certificate for a yes-instance. We can verify in polynomial time by a shortest path algorithm, that A_2 , the action chosen by player two is subgame perfect. For player 1 we do not have a way to verify that A_1 is a subgame perfect action, but we do not need to: we simply verify if the cost of player 1 is indeed at most k for the action profile (A_1, A_2) . If yes, player 1 can guarantee herself cost at most k by choosing A_1 , so in any subgame perfect equilibrium, player 1 will have cost at most that much. Hence we can verify the validity of the certificate for any yes-instance in polynomial time. \square

6. The sequential price of anarchy for symmetric network routing games. We now prove our main result, namely that the *SPoA* of symmetric network routing games is unbounded.

THEOREM 5. *The sequential price of anarchy of symmetric network routing games with linear cost functions is not bounded by any constant.*

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. There are slightly more players than disjoint (s, t) -paths. As an effect of our construction, it will turn out that 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




FIGURE 5. A lower bound instance of a network routing game. Players travel from s to t .

is not too expensive for every single threatener, but very expensive for the common predecessor. Altogether, the goal of the construction is to incentivize a large number of players to choose a set of arcs, much larger than in the optimal outcome, so as to drive the sequential price of anarchy to infinity. The challenging aspect is to make sure that this is indeed a subgame perfect equilibrium.

6.1. 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 $\{i, (2i, 1), (2i, 2), \dots, (2i, p\sqrt{x} + 4x^2), (2i+1, 1), (2i+1, 2), \dots, (2i+1, p\sqrt{x} + 4x^2), i+1\}$. Note that node $i+1$ is in both segments R_i and R_{i+1} . The arc set is defined as follows.

- There is an arc e with cost coefficient $a_e = 0$ from node i to node $(2i, j)$ for all $j \in \{1, \dots, (p\sqrt{x} + 4x^2)\}$.
- There is an arc e with cost coefficient $a_e = 1/x$, from $(2i, j)$ to $(2i+1, j)$ for all $j \in \{1, \dots, (p\sqrt{x} + 4x^2)\}$.
- There is an arc e with cost coefficient $a_e = 0$ from $(2i+1, j)$ to $i+1$ for all $j \in \{1, \dots, (p\sqrt{x} + 4x^2)\}$.
- There is an arc e with cost coefficient $a_e = 0$ from $(2i+1, j)$ to $(2i, k)$ for all $j \in \{1, \dots, (p\sqrt{x} + 4x^2)\}$ and for all $k \in \{j+1, \dots, (p\sqrt{x} + 4x^2)\}$.

Note that between any two nodes $i, i+1$, there exist $2^{p\sqrt{x}+4x^2}$ different paths: one for every subset of arcs with cost coefficient $1/x$ of segment R_i . For brevity, when we refer from now on to *arcs*, we mean the arcs of which the cost coefficient is not identically zero, i.e., arcs with cost coefficient $1/x$.

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 5.

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 $e \in A_{i'}$.

6.2. 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

$$C(A^*) = p\sqrt{x} + 3x^2 + (2x^2)2 = p\sqrt{x} + 7x^2.$$

6.3. Definiton of strategy profile S for Γ_x In order to describe our worst-case subgame perfect equilibrium strategy, we first define the following actions. Note that these 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.

- *Greedy*: In each segment, choose the single arc chosen by the least 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 in total 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.

Using these actions, we now define our subgame perfect equilibrium $S = (S_1, \dots, S_n)$ 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 in Algorithm 1.

6.4. Tie-breaking rule. In Algorithm 1, when the strategy S_i prescribes that a player i chooses an arc chosen by the largest number of players, any arbitrary fixed tie-breaking rule is used. Secondly, when the strategy S_i prescribes that a player i chooses an arch 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 lexicographically 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' .

Algorithm 1 Definition of subgame perfect equilibrium S (for player i).

```

1: if Every player  $j \in [i - 1]$  plays according to  $S_j$  then
2:   if  $i$  has at least  $5x^2$  successors then
3:     if  $i$  is the first player, or if the previous  $\sqrt{x} - 1$  players chose Copy then
4:       Fill
5:     else
6:       Copy
7:     else
8:       Greedy
9:   else
10:  if exactly 1 player  $j \in [i - 1]$  does not play according to  $S_j$  then
11:    if  $j$  has chosen less than  $x^2$  arcs in each segment then
12:      if  $S_j$  prescribed  $j$  to choose Fill or Copy then
13:        if there exists a segment such that all arcs  $e$  chosen by  $j$  contain less than  $x$  players
14:          in total then
15:            Punish( $j$ )
16:          else
17:            Greedy
18:        else
19:          Greedy

```

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 $\{e_1, e_2, e_3, e_4\}$, which are chosen by the smallest number of players. Players 1 and 3 have deviated from S . Player 1 has chosen, among arcs $\{e_1, e_2, e_3, e_4\}$ the two arcs e_2 and e_3 , player 2 has chosen arcs e_3 and e_4 , player 3 has chosen arcs e_1 and e_4 , and player 4 has chosen arcs e_1 and e_2 . Thus, the players are ordered 3, 1, 4, 2 and the arcs are ordered e_4, e_1, e_3, e_2 , so player 5 chooses arcs e_4 and e_1 . \triangleleft

Observe first that S is well-defined:

- In any state, S_i prescribes i to play either *Greedy*, *Copy*, *Fill*, or *Punish*(j) for some $j \in [i - 1]$.
- In any state, the actions *Greedy* and *Copy* always exist.
- Whenever the action *Fill* is prescribed, then by line 1, no player in $[i - 1]$ has deviated from S . Combined with line 2 this means that in each segment there are at most $p\sqrt{x}$ arcs that are chosen by at least one player in $[i - 1]$. Therefore, there is guaranteed to be enough free arcs in each segment, so the action *Fill* exists in any state where S_i prescribes i to choose *Fill*.

- Whenever the action $Punish(j)$ is prescribed, for some $j \in [i - 1]$, then from line 12 and line 2 it follows that $j \in [p\sqrt{x} - x^2]$. Also, from line 10 it follows that in each segment, the total number of arcs chosen by players in $[j - 1]$ (who all play $Fill$ and $Copy$) is at most $p\sqrt{x} - x^2 + \sqrt{x} < p\sqrt{x}$, and player j chooses at most x^2 arcs in each segment by line 11. Moreover, it follows from line 10 that players in $\{j + 1, \dots, i - 1\}$ have not deviated from S and have therefore chosen only one arc in each segment (as they all have played $Punish(j)$). From line 13 it follows that the number of players between j and i is at most x^2 , so the total number of arcs chosen in each segment, by players in $\{j + 1, \dots, i - 1\}$, is at most x^2 . Lastly, line 11 certifies that player j occupies at most x^2 arcs in each segment. So, in a single segment, the total number of arcs used by players in $[i - 1]$ is at most $p\sqrt{x} + 3x^2$. Thus, when S_i prescribes $Punish(j)$, each segment has a free arc. Moreover, by line 13, there is a segment in which all arcs chosen by player j are chosen by less than x players from $[j]$. Therefore, $Punish(j)$ exists in any state where S_i prescribes i to choose $Punish(j)$.

6.5. Social cost of S . If each player i chooses the action prescribed by S_i , then the social cost is at least

$$(p\sqrt{x})(\sqrt{x}\sqrt{x}) + 3x^2 + (2x^2)2 = p\sqrt{x}x + 7x^2.$$

If we denote by A the outcome induced by S , we see that

$$\lim_{p \rightarrow \infty} C(A)/C(A^*) = \lim_{p \rightarrow \infty} (px\sqrt{x} + 7x^2)/(p\sqrt{x} + 7x^2) = x.$$

6.6. 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 .

The following lemma and corollary shows that S is a subgame perfect equilibrium.

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 , under the assumption that all succeeding players $i + 1, \dots, n$ 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 cost 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 $i + 1, \dots, n$ all play according to S as well, observe that by definition of S , players $i + 1, \dots, n$ 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 are 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_j chosen by respectively player i and player n in A . Denote by R the set of arcs in R_j chosen by players i, \dots, n in A .

We denote by c the number of players choosing arc e_n in A . Any arc $e \in R$ has cost 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 cost of e_i is at most c/x . Also, any arc $e \in R_j$ 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 cost at least c/x . We now show that if e'_i is in R , then it has cost at least c/x as well. In that case, if any player $i' \in \{i + 1, \dots, n\}$ chooses an arc not in R then all arcs in R would yield cost at least c/x : 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, \dots, 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 cost of at least c/x as well. Therefore, in all cases the costs of player i does 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 $i + 1, \dots, n$ 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: $\{j + 1, \dots, i + k\}$ is the set of players that play $Punish(j)$. Let $\ell = |\{j + 1, \dots, i + k\}|$. Players $\{i + k + 1, \dots, n\}$ 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 $j + 1, \dots, i + k$ all choose $Punish(j)$, so they each occupy one 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 $i + k + 1, \dots, n$ is of size $G := p\sqrt{x} + 5x^2 - j - \ell$. We see that

$$\frac{G}{F} = \frac{p\sqrt{x} + 5x^2 - j - \ell}{(p - 1)\sqrt{x} + 3x^2 - j - \ell + 2} \leq \frac{p\sqrt{x} + 6x^2 - j - \ell}{(1/2)\sqrt{x} + 3x^2 - (1/2)j - (1/2)\ell} = 2,$$

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 cost is exactly $2 - 1/x$ if she plays $Punish(j)$: In $x - 1$ segments, i chooses one 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 now 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 $p\sqrt{x} + 4x^2 - j$ 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 since $i \leq j + x^2 - x$, which is true 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 $2/x$. \square

Now we can finish the proof of Theorem 5.

COROLLARY 1. *The strategy profile S is a subgame perfect equilibrium of Γ_x .*

Proof. Recalling the definition of a subgame perfect equilibrium, we have to prove that for all $i \in [n]$ and for each state $(A_{<i}, i)$ it holds that the action $S_i(A_{<i})$ that i plays under S minimizes her cost in the subgame corresponding to state $(A_{<i}, i)$, when the remaining players play according to S . This follows directly from Lemma 1 and from the way we defined subgame perfection of an action with respect to strategy profile S . \square

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 . This holds as for nonnegative, affine cost functions, n is the maximum multiplicative gap between the congestion of any given resource, for any two arbitrary action profiles. Our construction above shows actually a lower bound $SPoA \geq \Omega(\sqrt{n})$. To see this, we choose $p = x\sqrt{x}$. Then $n = x^2 + 5x^2 = 6x^2$ which yields $x = \sqrt{n/6}$. Now, $SPoA \geq (x^3 + 7x^2)/(x^2 + 7x^2) \geq x^3/(8x^2) = x/8 = \sqrt{n}/(8\sqrt{6})$. If instances can be constructed where the sequential price of anarchy is asymptotically worse than \sqrt{n} , remains open.

7. Discussion. The central result of this paper states that the sequential price of anarchy is generally unbounded for symmetric affine network routing games. One property that stands out in our lower bound constructions is that it admits 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 bad tie breaking. Recent results for a special case, namely so-called “EP graphs”, shows that indeed, the sequential price of anarchy is no longer unbounded if we consider *generic* games, i.e., those games admitting a unique subgame perfect equilibrium; see [21]. We do not know, however, if the sequential price of anarchy is bounded for *all* symmetric congestion games that are generic. That question relates to determining the *sequential price of stability* of symmetric network routing games.

Finally, note that our main result is asymptotic for large n . In view of concrete applications, it may be interesting to pin down the exact value for small numbers of n ; here we did that for $n = 2$.

Acknowledgments Jasper de Jong was supported by the Dutch 4TU Applied Mathematics Institute. José Correa was supported by research grant FONDECYT 1160079. All authors acknowledge travel funds from the Dutch Mathematics cluster DIAMANT. This work was done while the third author was with the Sapienza University of Rome and the University of Liverpool. We thank Mathieu Faure for stimulating discussions and particularly for pointing out a precursor of the instance depicted in Figure 3. We thank Marco Scarsini and Victor Verdugo for discussions on the price of anarchy of the symmetric atomic network game. We thank Éva Tardos for allowing us to recycle their paper title for the extended abstract of this paper that appeared in the proceedings of the 11th Conference on Web and Internet Economics [13]. Finally, we express our gratitude to the reviewers for their critical but constructive comments.

References

- [1] Angelucci A, Bilò V, Flammini M, Moscardelli L (2013) On the sequential price of anarchy of isolation games. Du DZ, Zhang G, eds., *Computing and Combinatorics (COCOON): 19th International Conference, Proceedings*, volume 7936 of *Lecture Notes in Computer Science*, 17–28 (Springer).
- [2] Angelucci A, Bilò V, Flammini M, Moscardelli L (2015) On the sequential price of anarchy of isolation games. *Journal of Combinatorial Optimization* 29(1):165–181.
- [3] Anshelevich E, Dasgupta A, Kleinberg J, Tardos É, Wexler T, Roughgarden T (2004) The price of stability for network design with fair cost allocation. *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 295–304 (IEEE).
- [4] Anshelevich E, Dasgupta A, Kleinberg J, Tardos É, Wexler T, Roughgarden T (2008) The price of stability for network design with fair cost allocation. *SIAM Journal on Computing* 38(4):1602–1623.
- [5] Avni G, Henzinger T, Kupferman O (2016) Dynamic resource allocation games. Gairing M, Savani R, eds., *Algorithmic Game Theory (SAGT): 9th International Symposium, Proceedings*, volume 9928 of *Lecture Notes in Computer Science*, 153–166 (Springer).

- [6] Awerbuch B, Azar Y, Epstein A (2005) The price of routing unsplittable flow. *Proceedings of the 37th Symposium on Theory of Computing (STOC)*, 57–66 (ACM).
- [7] Awerbuch B, Azar Y, Epstein A (2013) The price of routing unsplittable flow. *SIAM Journal on Computing* 42(1):160–177.
- [8] Bhawalkar K, Gairing M, Roughgarden T (2010) Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. de Berg M, Meyer U, eds., *Algorithms – ESA 2010: 18th Annual European Symposium, Proceedings, Part II*, volume 6347 of *Lecture Notes in Computer Science*, 17–28 (Springer).
- [9] Bhawalkar K, Gairing M, Roughgarden T (2014) Weighted congestion games: The price of anarchy, universal worst-case examples, and tightness. *Transactions on Economics and Computation* 2(4):14:1–14:23.
- [10] Bilò V, Flammini M, Monaco G, Moscardelli L (2013) Some anomalies of farsighted strategic behavior. Kaklamakis C, Pruhs K, eds., *Approximation and Online Algorithms (WAOA): 10th International Workshop, Revised Selected Papers*, volume 8447 of *Lecture Notes in Computer Science*, 229–241 (Springer).
- [11] Bilò V, Flammini M, Monaco G, Moscardelli L (2015) Some anomalies of farsighted strategic behavior. *Theory of Computing Systems* 56(1):156–180.
- [12] Christodoulou G, Koutsoupias E (2005) The price of anarchy of finite congestion games. *Proceedings of the 37th Symposium on Theory of Computing (STOC)*, 67–73 (ACM).
- [13] Correa J, de Jong J, de Keijzer B, Uetz M (2015) The curse of sequentiality in routing games. Markakis E, Schäfer G, eds., *Web and Internet Economics (WINE): 11th International Conference, Proceedings*, volume 9470 of *Lecture Notes in Computer Science*, 258–271 (Springer).
- [14] Correa J, Schulz A, Stier Moses N (2004) Selfish routing in capacitated networks. *Mathematics of Operations Research* 29:961–976.
- [15] de Jong J, Uetz M (2014) The sequential price of anarchy for atomic congestion games. Liu T, Qi Q, Ye Y, eds., *Web and Internet Economics (WINE): 10th International Conference, Proceedings*, volume 8877 of *Lecture Notes in Computer Science*, 429–434 (Springer).
- [16] de Jong J, Uetz M, Wombacher A (2013) Decentralized throughput scheduling. Spirakis P, Serna M, eds., *Algorithms and Complexity (CIAC): 8th International Conference, Proceedings*, volume 7878 of *Lecture Notes in Computer Science*, 134–145 (Springer).
- [17] Fabrikant A, Papadimitriou C, Talwar K (2004) The complexity of pure Nash equilibria. *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, 604–612 (ACM).
- [18] Farzad B, Olver N, Vetta A (2008) A priority-based model of routing. *Chicago Journal of Theoretical Computer Science* 2008(1).
- [19] Fotakis D (2015) A selective tour through congestion games. *Algorithms, Probability, Networks, and Games - Scientific Papers and Essays Dedicated to Paul G. Spirakis on the Occasion of His 60th Birthday*, volume 9295 of *Lecture Notes in Computer Science*, 223–241 (Springer).

- [20] Graham R (1969) Bounds on multiprocessing timing anomalies. *SIAM Journal on Applied Mathematics* 17(2):416–429.
- [21] Groenland C, Schäfer G (2018) The curse of ties in congestion games with limited lookahead. *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (to appear)* (IFAAMAS).
- [22] Groenland C, Schäfer G (2018) The curse of ties in congestion games with limited lookahead. *ArXiv/CoRR* abs/1804.07107.
- [23] Haeringer G, Wooders M (2011) Decentralized job matching. *International Journal of Game Theory* 40(1):1–28.
- [24] Harks T, Heinz S, Pfetsch M (2009) Competitive online multicommodity routing. *Theory of Computing Systems* 45:533–545.
- [25] Harks T, Vegh L (2007) Selfish routing with online demands. Jansen J, Pralat P, eds., *Proc. 4th Workshop on Combinatorial and Algorithmic Aspects of Networking (CAAN 2007)*, volume 4852 of *Lecture Notes in Computer Science*, 27–45 (Springer).
- [26] Hassin R, Yovel U (2015) Sequential scheduling on identical machines. *Operations Research Letters* 43(5):530–533.
- [27] Kawase Y, Yamaguchi Y, Yokoi Y (2018) Computing a subgame perfect equilibrium of a sequential matching game. *Proceedings of the 3rd Conference on Economics and Computation (EC) (to appear)* (ACM).
- [28] Kawase Y, Yamaguchi Y, Yokoi Y (2018) Computing a subgame perfect equilibrium of a sequential matching game. *ArXiv/CoRR* abs/1804.10353.
- [29] Kolev K (2016) *Sequential price of anarchy for atomic congestion games with limited number of players*. Master’s thesis, Universiteit Twente.
- [30] Koutsoupias E, Papadimitriou C (1999) Worst-case equilibria. Meinel C, Tison S, eds., *STACS 99: 16th Annual Symposium on Theoretical Aspects of Computer Science, Proceedings*, volume 1563 of *Lecture Notes in Computer Science*, 404–413 (Springer).
- [31] Koutsoupias E, Papadimitriou C (2009) Worst-case equilibria. *Computer Science Review* 3(2):65–69.
- [32] Kuhn HW (1953) Extensive games and the problem of information. Kuhn HW, Tucker AW, eds., *Contributions to the Theory of Games, Vol II*, 193–216 (Princeton University Press).
- [33] Lucier B, Singer Y, Syrgkanis V, Tardos É (2013) Equilibrium in combinatorial public projects. Chen Y, Immorlica N, eds., *Web and Internet Economics (WINE): 9th International Conference, Proceedings*, volume 8289 of *Lecture Notes in Computer Science*, 347–360 (Springer).
- [34] Milchtaich I (1998) Crowding games are sequentially solvable. *International Journal of Game Theory* 27(4):501–509.

- [35] Olver N (2006) *The Price of Anarchy and a Priority-Based Model of Routing*. Master’s thesis, McGill University, Montréal, Canada.
- [36] Osborne M (2003) *An Introduction to Game Theory* (Oxford University Press).
- [37] Paes Leme R, Syrgkanis V, Tardos É (2012) The curse of simultaneity. *Proceedings of the 3rd Conference on Innovations in Theoretical Computer Science (ITCS)*, 60–67 (ACM).
- [38] Rahn M, Schäfer G (2015) Efficient equilibria in polymatrix coordination games. Italiano G, Pighizzini G, Sannella D, eds., *Mathematical Foundations of Computer Science (MFCS) 2015: 40th International Symposium, Proceedings, Part II*, volume 9235 of *Lecture Notes in Computer Science*, 529–541 (Springer).
- [39] Roughgarden T, Tardos É (2002) How bad is selfish routing? *Journal of the ACM* 49(2):236–259.
- [40] Selten R (1965) Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit: Teil 1: Bestimmung des dynamischen Preisgleichgewichts. *Zeitschrift für die gesamte Staatswissenschaft* 121(2):301–324.
- [41] Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers* 1(3):325–362.

Appendix. Proof of upper bound for Theorem 2.

We here give the missing proof of the $7/5$ upper bound for sequentially played affine congestion games with $n = 2$ players.

LEMMA 2. *The sequential price of anarchy of symmetric affine congestion games is at most $7/5$ when there are $n = 2$ players only.*

Because we consider non-network congestion games for this upper bound proof, we use the term *resources* instead of *arcs*. Consider any symmetric affine congestion game with 2 players, m resources, affine cost functions (c_1, \dots, c_m) and the set $R \subseteq 2^{[m]}$ such that the action set of both players is R . We first show that we may assume without loss of generality that all the cost functions are of the form $c_e(y) = y$. This is because we can transform any symmetric affine congestion game as follows: First, by scaling all the cost functions simultaneously by an appropriate constant, we may assume that each cost function c_e is of the form $c_e(y) = a_e y + b_e$, with a_e and b_e being natural numbers. (It is straightforward to do so when the original coefficients of the affine cost functions are rational. In case the original coefficients are irrational, additionally one must first turn such an instance into an equivalent one where the coefficients are approximated by rational numbers.) For a resource e where $a_e \neq 1$ or $b_e \neq 0$ we may replace each action $A' \in R$ that contains e by two new strategies: $A^1 = (A' \setminus e) \cup \{e'_j : j \in [a_e]\} \cup \{e''_j : j \in [b_e]\}$ and $A^2 = (A' \setminus e) \cup \{e'_j : j \in [a_e]\} \cup \{e'''_j : j \in$

$\{b_e\}$, where for all j , the resources e'_j, e''_j and e'''_j are newly introduced into the game, and all these newly introduced resources have cost functions of the form $c_e(y) = y$. There is a natural mapping from the old outcomes to the new outcomes: for player 1, action A' is mapped to A^1 , and for player 2, action A' is mapped to A^2 . Note that, since the transformed game is still symmetric, action A^2 can be chosen by player 1 and action A^1 can be chosen by player 2. However, the two players will not both choose action A^1 if $b_e > 0$ since player 2 could then improve her cost by choosing the action A^2 , and a symmetric conclusion can be made if both players would choose action A^2 . Thus, by this transformation, the sequential price of anarchy can only increase. This shows that we may assume that the cost functions have the form $c_e(y) = y$.

Now, let $A^* = (A_1^*, A_2^*)$ be an outcome that minimizes the social cost, and let $A = (A_1, A_2)$ be the outcome resulting from a subgame perfect equilibrium $S = (S_1, S_2)$. We may assume w.l.o.g. that $c_1(A^*) \leq c_2(A^*)$, because the action sets are symmetric. We first define some parameters that we will use throughout this proof:

- Let x be the minimum cardinality of a set of resources in the action set R .
- Let $\alpha \in \mathbb{R}$ be such that $|A_1^* \cap A_2^*| = \alpha x$.
- Let β be such that $|A_1^* \cap A_1| = \beta |A_1|$. Note that this implies that $|A_2^* \cap A_1| \leq (1 - \beta)|A_1| + \alpha x$.
- Let γ be such that $C(A^*) = (2 + \gamma)x$. Observe that γ is positive, as $2x$ is a lower bound on $C(A^*)$.
- Let δ be such that $c_1(A^*) = (1 + \delta)x$. It holds that $\delta \leq \gamma/2$ because $c_1(A^*) \leq c_2(A^*)$. The definition of δ implies that $c_2(A^*) = (1 + \gamma - \delta)x$.

We derive several upper bounds on $C(A)$ that are expressed in terms of α, β, γ , and δ .

Observe first that each player experiences a cost of at most $2x$ under A , because there exists an action of cardinality x of which each resource is chosen by at most two players.

PROPOSITION 1.

$$c_1(A) \leq 2x, \quad c_2(A) \leq 2x.$$

This gives us a straightforward upper bound of $4x$ on $C(A)$. Note that in case $\gamma \geq 6/7$, we obtain that

$$\frac{C(A)}{C(A^*)} \leq \frac{4}{20/7} = \frac{7}{5},$$

so it remains to prove the claim for the case that $\gamma \in [0, 6/7]$, which we assume from now on.

We prove a second upper bound on $C(A)$ next. By subgame perfection of S it holds that $c_1(A) \leq c_1(A_1^*, S_2(A_1^*))$. Also, $c_2(A_1^*, S_2(A_1^*)) \leq c_2(A_1^*, A_2^*) = (1 + \gamma - \delta)x$. The number x is defined as the smallest cardinality of an action in R , so $S_2(A_1^*)$ intersects with A_1^* in at most $(\gamma - \delta)x$ resources. We combine the latter with the fact that the cardinality of A_1^* is $(1 + \delta - \alpha)x$, and we conclude that $c_1(A_1^*, S_2(A_1^*)) \leq (1 + \delta - \alpha)x + (\gamma - \delta)x$. Therefore:

PROPOSITION 2.

$$c_1(A) \leq (1 + \gamma - \alpha)x.$$

Combining this with Proposition 1 gives us that

$$C(A) \leq (1 + \gamma - \alpha)x + 2x = (3 + \gamma - \alpha)x. \quad (3)$$

We prove two additional upper bounds on $C(A)$ next. By subgame perfection of S it holds that $c_2(A) \leq c_2(A_1, A_1^*)$ and it holds that $c_2(A) \leq c_2(A_1, A_2^*)$. The cost $c_2(A_1, A_1^*)$ can be upper bounded as follows.

$$\begin{aligned} c_2(A_1, A_1^*) &\leq |A_1^*| + |A_1^* \cap A_1| \leq (1 + \delta - \alpha)x + \beta|A_1| \\ &\leq (1 + \delta - \alpha)x + \beta c_1(A) \\ &\leq (1 + \delta - \alpha)x + \beta(1 + \gamma - \alpha)x, \end{aligned}$$

where we use Proposition 2 for the last inequality. The cost $c_2(A_1, A_2^*)$ can be upper bounded as follows.

$$\begin{aligned} c_2(A_1, A_2^*) &\leq |A_2^*| + |A_2^* \cap A_1| \leq (1 + \gamma - \delta - \alpha)x + (1 - \beta)|A_1| + \alpha x \\ &\leq (1 + \gamma - \delta)x + (1 - \beta)|A_1| \\ &\leq (1 + \gamma - \delta)x + (1 - \beta)c_1(A) \\ &\leq (1 + \gamma - \delta)x + (1 - \beta)(1 + \gamma - \alpha)x. \end{aligned}$$

where we use Proposition 2 for the last inequality.

Combining the above two with Proposition 2 gives us

$$\begin{aligned} C(A) &= c_1(A) + c_2(A) \\ &\leq (1 + \gamma - \alpha)x + (1 + \delta - \alpha)x + \beta(1 + \gamma - \alpha)x \\ &= (1 + \delta - \alpha)x + (1 + \beta)(1 + \gamma - \alpha)x \end{aligned} \quad (4)$$

and

$$\begin{aligned} C(A) &= c_1(A) + c_2(A) \\ &\leq (1 + \gamma - \alpha)x + (1 + \gamma - \delta)x + (1 - \beta)(1 + \gamma - \alpha)x \\ &= (1 + \gamma - \delta)x + (2 - \beta)(1 + \gamma - \alpha)x. \end{aligned} \quad (5)$$

Combining (3), (4), and (5), we conclude that the sequential price of anarchy of our game is at most

$$\frac{\min \left\{ \begin{array}{l} (3 + \gamma - \alpha)x, \\ (1 + \delta - \alpha)x + (1 + \beta)(1 + \gamma - \alpha)x, \\ (1 + \gamma - \delta)x + (2 - \beta)(1 + \gamma - \alpha)x \end{array} \right\}}{(2 + \gamma)x} = \frac{\min \left\{ \begin{array}{l} (3 + \gamma - \alpha), \\ (1 + \delta - \alpha) + (1 + \beta)(1 + \gamma - \alpha), \\ (1 + \gamma - \delta) + (2 - \beta)(1 + \gamma - \alpha) \end{array} \right\}}{(2 + \gamma)}$$

We can obtain a concrete upper bound on the sequential price of anarchy of the complete class two-player symmetric affine congestion games when we maximize the latter expression subject to the constraints $\gamma \in [0, 6/7], \beta \in [0, 1], \delta \in [0, 1/2\gamma]$. The variable α can be eliminated, as it is clear that the maximum is attained when $\alpha = 0$. This results in the optimization problem

$$\max \left\{ \frac{\min\{(3 + \gamma), (1 + \delta) + (1 + \beta)(1 + \gamma), (3 - \beta)(1 + \gamma) - \delta\}}{(2 + \gamma)}, : 0 \leq \gamma \leq \frac{6}{7}, 0 \leq \beta \leq 1, 0 \leq \delta \leq \frac{\gamma}{2} \right\}.$$

Numerically solving this program gives that the solution is $7/5$, attained when we take $\gamma = \beta = 1/2, \delta = 1/4$, although this does not comprise a formal proof. However, it is possible to prove formally that the solution does not exceed $7/5$, by showing that the optimal solution to the following optimization problem does not exceed zero:

$$\max \left\{ \min \left\{ \begin{array}{l} (3 + \gamma), \\ (1 + \delta) + (1 + \beta)(1 + \gamma), \\ (1 + \gamma - \delta) + (2 - \beta)(1 + \gamma) \end{array} \right\} - \frac{7}{5}(2 + \gamma) : 0 \leq \gamma \leq \frac{6}{7}, 0 \leq \beta \leq 1, 0 \leq \delta \leq \frac{\gamma}{2} \right\}. \quad (6)$$

We introduce an additional variable z that we use in order to eliminate the min-expression in the objective function.

$$\max \left\{ \begin{array}{l} z - \frac{7}{5}(2 + \gamma) \\ : z \leq 3 + \gamma, z \leq (1 + \delta) + (1 + \beta)(1 + \gamma), \\ z \leq (1 + \gamma - \delta) + (2 - \beta)(1 + \gamma), \\ 0 \leq \gamma \leq \frac{6}{7}, 0 \leq \beta \leq 1, 0 \leq \delta \leq \frac{\gamma}{2} \end{array} \right\}.$$

There are still two constraints in this program that are non-linear, because they contain the terms $\gamma\beta$ and $-\gamma\beta$ respectively. We introduce a new variable ξ that we restrict to lie in $[0, 6/7]$ and

substitute the latter two terms by ξ and $-\xi$ respectively. This results in the following linear program, with an “enlarged” feasible region:

$$\begin{aligned} \max \quad & \left\{ z - \frac{7}{5}(2 + \gamma) \right. \\ & : z \leq 3 + \gamma, z \leq 2 + \delta + \beta + \gamma + \xi, z \leq 3 + 3\gamma - \delta - \beta - \xi, \\ & \left. 0 \leq \gamma \leq \frac{6}{7}, 0 \leq \beta \leq 1, 0 \leq \delta \leq \frac{\gamma}{2}, 0 \leq \xi \leq 6/7 \right\}. \end{aligned}$$

Because the set of feasible points of this linear program is larger, the solution to this linear program is an upper bound to the solution of (6). The exact solution to this linear program can be obtained by known algorithms, and turns out to be 0, as we needed to show. \square