The price of anarchy in series-parallel graphs

Pieter Senster Tim van Heugten Ot ten Thije

Delft University of Technology

Abstract

Congestion games model self-interested agents competing for resources in communication networks. The price of anarchy quantifies the deterioration in performance in such games compared to the optimal solution. Recent research has shown that, when the social cost is defined as the maximum cost of all players, specific graph topologies impose a bound on the price of anarchy. We extend this research by providing bounds on the price of anarchy for congestion games on series-parallel networks. First we show that parallel composition does not increase the price of anarchy. This result is then used to show that the price of anarchy is bounded above by both the diameter of the graph and the number of players in the game, and that these bounds are tight. Finally we identify an important aspect of proofs for bounds on the price of anarchy: when a bound is achieved by restricting multiple parameters of the game, one should also prove that this bound cannot be realized using only a subset of these restrictions.

1 Introduction

The optimality of solutions given by most classical optimization algorithms relies critically on the fact that all parties involved are interested in achieving the same objective: to achieve a common optimal solution. However, in many cases individual participants (or *agents*) are interested in optimising their individual objectives. Agents displaying such behaviour are called *selfish*: they only care about their own performance, regardless of what the effect on the entire system may be. The presence of selfish agents into a system can severely degrade its global performance.

For an example of the effects of selfish behaviour, consider internet routing. When downloading a file over the internet from a single source, the packets that constitute the file should be delivered as quickly as possible. The global objective is to minimise the time it takes for the latest packet to arrive, while the individual agents want their packets to be delivered as fast as possible. In order to avoid congestion, some packets will be redirected along routes that take far more time than the shortest possible link. By redirecting some packets, the routing algorithm averts congestion, thus ensuring that the maximum travelling time for a packet remains low.

However, the packets ordered to take the longer route can improve their performance by ignoring the command to reroute and continue along the shortest link anyway. If they act selfishly like this, their *own* delivery times speed up, but the additional load may cause the shortest link to become congested. This in turn negatively affects the delivery time of *all other* packets, thereby decreasing the performance of the system as a whole.

Depending on the problem, there may be different ways to reduce the effect of selfishness on the overall optimality of the solution. In order to compare the relative merits of these different approaches, Papadimitriou et al. introduce the notion of the *price of anarchy* [10], defined as the ratio between the global performance in the worst-case Nash equilibrium and the global performance in the optimal solution. A solution is said to be a Nash equilibrium when no agent can unilaterally improve its result.

This measure is closely related to the approximation ratio for approximation algorithms and the competitive ratio for online algorithms [5]: The approximation ratio quantifies the price we pay in terms of optimality when *heuristics* are used to approximate hard problems in polynomial time. Likewise, the competitive ratio quantifies the price of dealing with *incomplete information* in online situations. Finally, the price of anarchy (or *coordination ratio*) allows us to quantify the effect of the *lack of coordination* between selfish agents.

The price of anarchy varies depending on the exact characteristics of the problem. These characteristics include the topology of the graph, the exact definition of "performance" and various properties of the agents themselves.

In this paper we concentrate on the price of anarchy of network routing modelled as a congestion game. An interesting result in this field is Braess's paradox [3]: the observation that adding an edge to a network can result in a higher price of anarchy, leading to worse performance. Since the price of anarchy can be unbounded in general graphs [12], researchers seek alternative ways to bound the price of anarchy.

Quite some research concentrates on the total latency of the system, defined as the sum of the latencies of all players. For arbitrary latency functions, Roughgarden [12] shows that the price of anarchy is unbounded, while the bound is $\frac{4}{3}$ when only linear latency functions are used.

Roughgarden concentrates on non-atomic network flow. Christodoulou and Koutsoupias [4] on the other hand study networks with linear latency functions for atomic network flow. They show that the price of anarchy is $\frac{5n-2}{2n+1}$ for symmetric games. When the social cost is defined as the maximum cost of a player (as in this paper), they show that the price of anarchy for symmetric games is $\frac{5}{2}$.

Another way to bound the price of anarchy is to restrict the network topology [6, 7]. Epstein et al. follow this approach to determine what network topologies are *efficient*, i.e. for which every Nash equilibrium is socially optimal [6]. They show that series-parallel graphs are efficient only in *bottleneck routing* games, that is, when the social cost is determined by the maximum cost of any player on any edge. This definition of the social cost is particularly relevant in communication networks [1]. They also show that so-called extension-parallel graphs, a restricted subset of series-parallel graphs, are efficient for both bottleneck routing and network congestion games. When the social cost is defined as the total cost of all players, Fotakis [7] shows that the price of anarchy in an extension-parallel graph is bounded by the price of anarchy for non-atomic congestion games.

In a different approach, the price of anarchy is not analysed as a property of the game itself, but rather as a value to be optimized by guiding the agents [2, 8, 13, 14]. In this case a central authority, aware of the optimal solution, tries to guide the agents to follow that solution. This can be accomplished by assigning a fraction of the (non-atomic) agents to a specific route, as is done with Stackelberg routing by Roughgarden in [11], who uses the analysis of Korilis et al. [8]. Another approach is to modify the cost functions in the network. Christodoulou et al. introduce a *coordination mechanism* as a means of bounding the price of anarchy [5]. In an example with selfish task allocation, represented as an asymmetric parallel graph congestion game, they use their mechanism to improve the price of anarchy from an initial value of $\frac{3}{2}$ to $\frac{4}{3} - \frac{1}{3m}$. Unfortunately, their mechanism is not applicable to general (symmetric) singe-commodity congestion games.

Our results We study atomic routing games on series-parallel graphs, with the social cost defined as the maximum of the costs of all players. For games with these properties, we show that the price of anarchy is bounded in both the number of players and the diameter of the graph. We note that Christodoulou et al. [5] present a coordination mechanism that achieves the same bound, while in our research we discover that this bound is a property of the graph topology. Furthermore, we show that parallel composition does not increase the price of anarchy, which is a generalisation of the work by Epstein et al. [6].

Organisation The remainder of this paper is organised as follows. In section 2 we introduce the model and notation used. We then provide the upper bounds on the price of anarchy in series-parallel graphs in section 3 and provide an example which shows that these bounds are tight. Finally, we present our conclusions and suggestions for future work in section 4.

2 The model

We denote a symmetric congestion game with n players as the tuple $\Lambda = (N, \{\Sigma_i\}_{i\in N}, \{c_i(\cdot)\}_{i\in N})$. Each player $i \in \{1, \ldots, n\} = N$ is associated with a set of strategies Σ_i. The cost of the player i is defined by the cost function $c_i(\cdot)$, which maps a strategy $S_i \in \Sigma_i$ to a value. The n-tuple $S = (S_1, \ldots, S_n)$ denotes the joint action taken by the players. We refer to the actions of the players other than i as S_{-i}.

A game Λ is associated with a directed graph $G = (V, E)$, in which edges E represent the resources the players compete for. We consider atomic single-commodity games, so all players have one unsplittable unit of flow they want to route from a source $s \in V$ to a destination $t \in V$. The allowed strategies Σ_i of a player

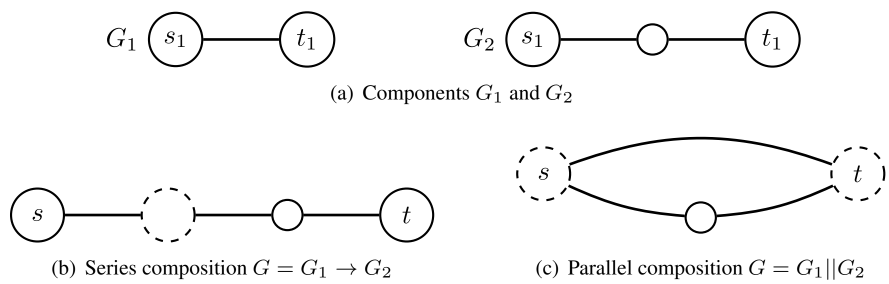

Figure 1: Composition of series-parallel graphs. Nodes with dashed borders were merged in a composition step.

consist of the set of simple paths in G between s and t. Because for the cases handled in this paper we use a symmetric strategy system, we have $\Sigma_i = \Sigma$ for all players.

Each edge has a non-decreasing cost function $\ell_e(k)$, representing the cost (or latency) of the edge when it is used by k players. The cost of a facility is the same for all players using it, and depends only on the number of players k. We denote the number of players using an edge e in a given joint strategy S by $n_e(S)$.

We define the cost $c(S_i, S_{-i})$ of a player to be the sum of the costs of the edges in the path S_i that it uses, given the strategies S_{-i} of the players in the game, so $c(S_i, S_{-i}) = \sum_{e \in S_i} \ell_e(n_e(S_i \cup S_{-i}))$. Given a joint action S_{-i} of all players other than i, the best response S_i of this player is defined as $\text{argmin}_{S_i \in \Sigma} c(S_i, S_{-i})$. The player with the maximum cost defines the social cost $sc_\Lambda(S) = \max_i c(S_i, S_{-i})$. For a given congestion game Λ, we want to minimize this social cost.

In this paper we compare pure Nash equilibria in congestion games to the social optimum OPT. This optimum does not need to be a Nash equilibrium itself, therefore $OPT(\Lambda) = \min_{S \in \Sigma} sc_\Lambda(S)$. The ratio between the worst possible Nash equilibrium and the social optimum is called the price of anarchy [9].

Series-parallel graphs Series-parallel (SP) graphs form a class of graphs with a simple network topology. Each SP graph has two special nodes, called *terminals*: a source s and a sink t. Using the two terminals as reference, an SP-graph G can be defined recursively, as follows:

1. A graph G consisting of only the source s, the sink t and a single edge connecting them is a series-parallel graph.
2. Given two series-parallel graphs G_1 and G_2, their *series composition* $G = G_1 \rightarrow G_2$ is an SP graph as well. The series composition G is obtained by merging the sink t_1 of G_1 with the source s_2 of G_2.
3. Given two series-parallel graphs G_1 and G_2, their *parallel composition* $G = G_1 || G_2$ is an SP-graph. To obtain the parallel composition G, merge the two sources s_1 and s_2 into a single source node s and merge the two sinks in the same way.

Figure 1 illustrates the series and parallel composition steps.

3 Bounding the price of anarchy in series-parallel graphs

In [5], Christodoulou et al. present a coordination mechanism that places an upper bound of n on the price of anarchy in congestion games. In this mechanism, a designer with knowledge of the optimal solution modifies the cost functions in the network, such that the participating agents will only consider the edges used by the optimal solution. However, in general graphs the resulting strategy selection still need not be optimal, since the number of agents choosing each of those edges may still differ from that in the optimal solution. So, while this mechanism is not applicable to general graphs, Christodoulou et al. were able to show that their mechanism could be applied to ensure a bound of n in series-parallel graphs.

In this section we show that this bound is an *inherent property* of series parallel graphs, that is, the bound holds even if no coordination mechanism is used at all. This illustrates the fact that adjusting the cost functions and restricting the topology of a problem can *independently* affect the price of anarchy. Therefore, when multiple aspects of the problem are changed simultaneously, one should always make sure that each of these changes is indeed required to achieve the desired bound, and that the same bound can not be achieved by a subset of the suggested changes.

In order to prove our bound, we first prove that parallel composition does not increase the price of anarchy. Specifically, assume that for any symmetric network routing game Λ on the networks G_1 and G_2, for any joint action of the players, the cost of the best response strategy of any player is at most f times the optimal social cost. In that case we show that for any joint action of the players in the game Λ, the cost of the best response strategy of any player is at most f times the optimal social cost. This statement is made precise in lemma 3.1.

Lemma 3.1. *For $j \in \{1, 2\}$, let $P_{i,j}$ be the best response of any player i to the joint action $(S_j)_{-i}$ of the players in the symmetric network routing game Λ_j on the SP graph G_j. Furthermore, let P_i be the best response of any player i to the joint action S_{-i} of the players in the symmetric network routing game Λ on the graph $G = G_1||G_2$. Then, if $c(P_{i,j}, (S_j)_{-i}) \leq f \cdot OPT(\Lambda_j)$ (where $j \in \{1, 2\}$) it must hold that $c(P_i, S_{-i}) \leq f \cdot OPT(\Lambda)$.*

As mentioned before, intuitively this lemma states that parallel composition does not increase the price of anarchy. The proof given below is a generalisation of the proof given by Epstein et al. [6], who prove this lemma for $f = 1$.

Proof of lemma 3.1. Let S^* be the optimal joint action of the players in G, such that $OPT(\Lambda) = sc_\Lambda(S^*)$. Let T_j^* be the set of players using paths in G_j according to S^*, and let $x_j^* = |T_j^*|$, for $j \in \{1, 2\}$. Finally, let S be any (other) joint action, let T_j be the set of players using paths in G_j according to S, and let $x_j = |T_j|$. Now, there are two cases:

Case 1: $x_1 = x_1^*$. In this case the same number of players selects the upper parallel part (G_1) in solution S as in the optimum solution S^*, and (therefore) the same number of players selects the lower parallel part (G_2) as well.

Let Λ_1 be a symmetric network routing game on the symmetric network G_1, with players T_1 and the original edge cost functions. Define Λ_2 and T_2 analogously for the network G_2. Now let S' and S'' be the joint action induced by S of the players in T_1 and T_2 respectively.

Using the assumption on the network G_1 stated in the lemma, for every player i in the game Λ_1 with a best response P'_i to S'_{-i}, it holds that $c(P'_i, S'_{-i}) \leq f \cdot OPT(\Lambda_1)$. Following the same reasoning for Λ_2 shows that $c(P''_i, S''_{-i}) \leq f \cdot OPT(\Lambda_2)$ for any player i in game Λ_2 with best response P''_i to S''_{-i}.

Because we assumed that the number of players using any component j is the same in both the optimal solution and in any other solution, it follows that $OPT(\Lambda_j) \leq OPT(\Lambda)$ for $j \in \{1, 2\}$. After all, in both games the same number of players have to use each component, so the optimal solution in the composite game Λ cannot have costs on a component j lower than the costs in the optimal solution for the component game Λ_j.

Now let P_i be player i's best response to S_{-i} in Λ. Since P_i is a best response, it holds for any player $i \in T_1$ that $c(P_i, S_{-i}) \leq c(P'_i, S_{-i})$ for every P'_i. Also, because S' is the response induced by S in G_1, it holds that $c(P'_i, S_{-i})$ in Λ is equal to $c(P'_i, S'_{-i})$ in Λ_1.

In combination with the inequalities obtained above, this allows the following derivation:

$$\begin{aligned} c(P_i, S_{-i}) &\leq c(P'_i, S_{-i}) && \text{(1a)} \\ &= c(P'_i, S'_{-i}) && \text{(1b)} \\ &\leq f \cdot OPT(\Lambda_1) && \text{(1c)} \\ &\leq f \cdot OPT(\Lambda) && \text{(1d)} \end{aligned}$$

In this derivation equation 1a holds because P_i is the best response to S_i, equation 1b holds because S' are the actions of the players in T_1 induced by S, equation 1c uses the assumption from the formulation of the lemma and equation 1d is true since $OPT(\Lambda_1) \leq OPT(\Lambda)$.

The same line of reasoning can be used for Λ_2, by substituting Λ_1, S', P'_i and T_1 with Λ_2, S'', P''_i and T_2 respectively. Since the derivation holds for players in both T_1 and T_2, the lemma holds in this case.

Case 2: There exists a network G_j for which $x_j^* > x_j$. In this case one of the components is used by more players in the solution S than in the optimal solution S^*. Conversely, the other component is used by less players in S than in S^*.

W.l.o.g. suppose $x_1^* > x_1$. Consider an arbitrary player i. Let $x'_1 = |T_1 \cup \{i\}|$, which yields $x_1^* \geq x_1 + 1 \geq x'_1$. Let Λ_1 be a symmetric network routing game on the symmetric network G_1 with players T_1 and the original edge cost functions. Now let S' be the joint action of the players in $T_1 \cup \{i\}$ induced by S.

Since $x_1^* \geq x_1'$, we have that $OPT(\Lambda_1) \leq OPT(\Lambda)$. After all, there are at least as many player using component G_1 in game Λ as there are players using that component in Λ_1. Therefore, if $OPT(\Lambda_1)$ were greater than $OPT(\Lambda)$, the players in Λ_1 could simply adapt the cheapest strategies from Λ and come out ahead: Since there are no more players in Λ_1 than in Λ, the cost to pay would be less or equal. This would contradict the optimality of $OPT(\Lambda_1)$, so indeed it must hold that $OPT(\Lambda_1) \leq OPT(\Lambda)$.

By the assumption on the network G_1, for every player i in the game Λ_1 with a best response P_i to S'_{-i}, it holds that $c(P'_i, S'_{-i}) \leq f \cdot OPT(\Lambda_1)$. Using the inequality above it follows that $c(P'_i, S'_{-i}) \leq f \cdot OPT(\Lambda)$.

Now let P_i be player i's best response to S_{-i} in Λ. Since P_i is a best response, it holds for any player in $i \in T_1$ that $c(P_i, S_{-i}) \leq c(P'_i, S_{-i})$. Also, because S' is the response induced by S in G_1, it holds that $c(P'_i, S_{-i})$ in Λ is equal to $c(P'_i, S'_{-i})$ in Λ_1.

In combination with the inequalities obtained above, this allows the same derivation as in case 1:

$$\begin{aligned} c(P_i, S_{-i}) &\leq c(P'_i, S_{-i}) && \text{(2a)} \\ &= c(P'_i, S'_{-i}) && \text{(2b)} \\ &\leq f \cdot OPT(\Lambda_1) && \text{(2c)} \\ &\leq f \cdot OPT(\Lambda) && \text{(2d)} \end{aligned}$$

Again, equation 2a holds because P_i is the best response to S_i, equation 2b holds because S' are the actions of the players in $T_1 \cup \{i\}$ induced by S, equation 2c uses the assumption from the formulation of the lemma and equation 2d is true since $OPT(\Lambda_1) \leq OPT(\Lambda)$. Since player i was selected arbitrarily, the lemma holds in this case.

Since cases 1 and 2 constitute an exhaustive listing of possibilities, we conclude that the lemma holds. □

Now that we have shown that parallel composition does not increase the price of anarchy, we can concentrate on bounds that are dominated by series composition. A property that comes to mind immediately is the diameter of the graph. Since the diameter of a graph grows linearly with the diameter of the serial components, and considering that intuitively the price of anarchy is dependent on the independency of serial components, one would expect the price of anarchy to grow with the diameter of the graph. We show here that the price of anarchy does not grow faster than the diameter of the graph.

Theorem 3.2. *Let Λ be a symmetric network congestion game on an SP network G. Consider the joint action for the worst-case Nash equilibrium $S \in \Sigma$. Let $d(G)$ be the diameter of G. Then $sc_\Lambda(S) \leq d(G) \cdot OPT(\Lambda)$.*

Proof. We prove this theorem by induction on the network size $|E|$. Let Λ be a symmetric network game on an SP network $G(V, E)$. For $|E| = 1$ the claim holds trivially. In what follows we show that this property is preserved under series and parallel compositions.

Series composition. Suppose the network $G = G_1 \to G_2$ is a series composition of the SP networks G_1 and G_2. Let Λ_j be the symmetric network congestion game on the symmetric network G_j with the original players and original edge cost functions, for $j \in \{1, 2\}$. Let S_j be the joined action of the players for the worst-case Nash equilibrium in Λ_j.

Assuming as induction hypothesis that the theorem holds for the components separately, consider the following derivation:

$$\begin{aligned} sc_\Lambda(S) &\leq sc_\Lambda(S_1) + sc_\Lambda(S_2) && \text{(3a)} \\ &\leq d(G_1) \cdot OPT(\Lambda_1) + d(G_2) \cdot OPT(\Lambda_2) && \text{(3b)} \\ &\leq d(G_1) \cdot OPT(\Lambda) + d(G_2) \cdot OPT(\Lambda) && \text{(3c)} \\ &= (d(G_1) + d(G_2))OPT(\Lambda) && \text{(3d)} \\ &= d(G) \cdot OPT(\Lambda) && \text{(3e)} \end{aligned}$$

Inequality 3a follows from the assumption that we are considering the worst-case Nash equilibria; the social cost of S can not be worse than the concatenation of the strategies dominating $sc_\Lambda(S_1)$ and $sc_\Lambda(S_2)$ Otherwise there would have to be a path in either part with a higher cost, which would contradict the definition of worst-case Nash equilibria. Inequality 3b follows from the induction hypotheses, and 3c from the fact that the optimum in either part is not greater than the optimum of the complete graph. Otherwise there would exist a path with lower cost in one of the parts, contradicting the definition of optimum.

Parallel composition. Follows from lemma 3.1, by substituting $f = d(G)$. □

In order to prove that the price of anarchy is bounded by the number of players, we need a measure to bound the social cost for series composition of the graph. The most expensive combination of edges forming a simple s-t-path, using the edge-costs as incurred by S^*, is such a measure. We will first show that this measure is bounded above by $n \cdot OPT(\Lambda)$.

Lemma 3.3. *Let Λ be a symmetric network congestion game on an SP network G with source s and sink t. Let S^* be the optimal joint action, i.e., $OPT(\Lambda) = sc_\Lambda(S^*)$. Then $\max_{P\in\Sigma}\sum_{e\in P}\ell_e(n_e(S^*)) \leq n \cdot OPT(\Lambda)$.*

Proof.

$$\max_{P\in\Sigma}\sum_{e\in P}\ell_e(n_e(S^*)) \leq \sum_{S_i\in S^*}\sum_{e\in S_i}\ell_e(n_e(S^*)) \tag{4a}$$

$$\leq n\cdot \max_{S_i\in S^*}\sum_{e\in S_i}\ell_e(n_e(S^*)) \tag{4b}$$

$$= n\cdot OPT(\Lambda) \tag{4c}$$

Equation 4a states that the most expensive combination of edges, using the costs incurred by applying joint action S^*, is not greater than the sum of all costs incurred by that joint action. This is smaller than the number of players, n, times the most expensive path in the joint action, as described in equation 4b. And, since the most expensive path is the definition of social cost, equation 4c follows from the definitions of social cost and optimum. □

Next we will show that the bound on the social cost of the cost of the most expensive simple s-t-path is maintained under series composition.

Lemma 3.4. *If for the j-th component, with $j \in \{1, 2\}$, it holds that $sc_\Lambda(S_j) \leq \max_{P\in\Sigma_j}\sum_{e\in P}\ell_e(n_e(S_j^*))$ then also for the series composition $G = G_1 \to G_2$ it holds that: $sc_\Lambda(S) \leq \max_{P\in\Sigma}\sum_{e\in P}\ell_e(n_e(S^*))$*

Proof. Suppose the network $G = G_1 \to G_2$ is a series composition of the SP networks G_1 and G_2. Let Λ_j be the symmetric network congestion game on the symmetric network G_j with the original players and original edge cost functions, for $j \in \{1, 2\}$. Let S^* be the optimal joint action, i.e., $OPT(\Lambda) = cost_\Lambda(S^*)$, then:

$$sc_\Lambda(S) \leq sc_\Lambda(S_1) + sc_\Lambda(S_2) \tag{5a}$$

$$\leq \max_{P\in\Sigma_1}\sum_{e\in P}\ell_e(n_e(S_1^*)) + \max_{P\in\Sigma_2}\sum_{e\in P}\ell_e(n_e(S_2^*)) \tag{5b}$$

$$\leq \max_{P\in\Sigma}\sum_{e\in P}\ell_e(n_e(S^*)) \tag{5c}$$

The inequality 5a follows from the assumption that we consider worst-case Nash equilibria. If $sc_\Lambda(S)$ would be worse than the sum of the partial social costs, there must be a player in the composition that suffers a cost on a component that is higher than when the component was considered alone. This would contradict the definition of worst-case Nash equilibria. Inequality 5b follows from the assumption of the lemma. The last inequality follows from concatenating the most expensive combination of edges in both parts. If this would not hold, a more expensive path would have to exist in one of the parts, contradicting the definition of the maximum. □

Now we can continue to define the second bound on the price of anarchy.

Theorem 3.5. *Let Λ be a symmetric network congestion game on an SP network G. Consider the joint action for the worst-case Nash equilibrium S with $S_i \in \Sigma$. Then $sc_\Lambda(S) \leq n \cdot OPT(\Lambda)$, where n is the number of players in G.*

Proof. We prove this lemma by induction on the network size $|E|$. Let Λ be a symmetric network game on an SP network $G(V, E)$. For $|E| = 1$ the claim holds trivially. In what follows we show that this property is preserved under series and parallel compositions.

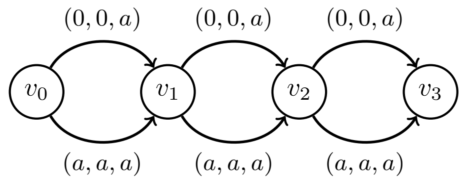

Figure 2: Congestion game for three players with price of anarchy n. In the optimum solution player i takes the i^{th} lower edge and upper edges elsewhere, for a social cost of a. In the worst-case Nash equilibrium two players take the upper edges, and the third player determines the social cost of $3a$. (Example reproduced from [5].)

Series composition. We observe that for the base case, with $|E| = 1$, the condition of lemma 3.4 is fulfilled. Therefore the lemma applies, and can be used to yield:

$$sc_\Lambda(S) \leq \max_{P \in \Sigma} \sum_{e \in P} \ell_e(n_e(S^*)) \tag{6a}$$

By use of lemma 3.3 we obtain the following result:

$$sc_\Lambda(S) \leq \max_{P \in \Sigma} \sum_{e \in P} \ell_e(n_e(S^*)) \tag{7a}$$

$$\leq n \cdot OPT(\Lambda) \tag{7b}$$

So it follows that the theorem holds for series composition.

Parallel composition. Follows from lemma 3.1, by substituting $f = n$. □

We summarize the results of both theorems presented in this paper in the following theorem.

Theorem 3.6. *The price of anarchy for general series parallel graphs is bounded above in both the diameter of the graph and the number of players in the game.*

Proof. Follows trivially from theorem 3.2 and theorem 3.5. □

We use the network in figure 2 to prove that these bounds are tight. The optimal set of strategies results in a social cost of a. On the other hand the worst-case Nash equilibrium for this example gives a social cost of $3 \cdot a$. With $n = 3$ as well as $d = 3$ this example is an instance that realizes the bounds from theorem 3.6. The bounds presented in this paper are therefore tight.

4 Conclusions and future research

In this paper we have shown that the price of anarchy in series-parallel graphs is bounded above in both the diameter of the graph and the number of players. We have also shown that the price of anarchy in series-parallel graphs does not increase under parallel composition. These results hold for atomic games in which the social cost is the maximum cost over all players. Earlier research showed that when we study the total or average cost of non-atomic players, the price of anarchy is unbounded regardless of the topology of the graph [12].

Our results put an upper bound on the price of anarchy in series-parallel graphs. One might ask whether this bound can be improved if we adjust the cost functions of the network edges. Christodoulou et al. [5] showed that no symmetric coordination mechanism is able to reduce the price of anarchy of a general series-parallel graph further than n or d. But although we have shown that their coordination mechanism for series-parallel graphs does not reduce the upper bound on the price of anarchy in general, it could still be the case that a coordination mechanism exists that lowers the price of anarchy for specific instances. Furthermore, it remains an open problem whether a symmetric coordination mechanism exists that achieves that bound in general graphs. More research is warranted on these subjects.

When designing a network, one could also consider enforcing restrictions on the topology of the graph. This gives the designer the guarantee of a bounded price of anarchy, without needing mechanisms to achieve a bound. When considering extension-parallel graphs instead of series-parallel graphs, Epstein et al. [6]

showed that the price of anarchy for such graphs is one. In many applications however, it is infeasible to change the topology of the network. Therefore, if a series-parallel graph topology is given, other methods (for example Stackelberg routing [2, 8, 13, 14]) are needed when a price of anarchy lower than n or d is desired.

In the search for better bounds on the price of anarchy, sometimes several parameters of the underlying game need to be restricted in order to prove a bound for a specific case. We have shown in this paper that in such cases one should be very careful what the exact effects of *subsets* of those restrictions are. In order to fully prove the necessity of the full set of restrictions, it should therefore also be shown that no subset of restrictions is sufficient to prove the bound. In the particular case described in this paper we have shown that the effects of restricting the graph topology have a considerable effect on the price of anarchy of the underlying graph.

Apart from graph topology, we have seen that the price of anarchy also depends heavily on other properties of the game such as the class of cost functions, the topology of the underlying network, the definition of the social cost and whether traffic is splittable. Many combinations of these properties have been documented in the literature, but an overview of the bounds is missing, even though it could be of great help to anyone studying the price of anarchy in a certain network.

References

[1] R. Banner and A. Orda. Bottleneck routing games in communication networks. *IEEE Journal on Selected Areas in Communications*, 25(6):1173–1179, 2007.

[2] Vincenzo Bonifaci, Tobias Harks, and Guido Schäfer. Stackelberg routing in arbitrary networks. In *WINE '08: Proceedings of the 4th International Workshop on Internet and Network Economics*, pages 239–250, Berlin, Heidelberg, 2008. Springer-Verlag.

[3] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Mathematical Methods of Operations Research*, 12(1):258–268, 1968.

[4] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, page 73. ACM, 2005.

[5] G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. *Theoretical Computer Science*, 410(36):3327–3336, 2009.

[6] A. Epstein, M. Feldman, and Y. Mansour. Efficient graph topologies in network routing games. *Games and Economic Behavior*, 66(1):115–125, 2009.

[7] D. Fotakis. Congestion games with linearly independent paths: Convergence time and price of anarchy. *Theory of Computing Systems*, pages 1–24.

[8] Y.A. Korilis, A.A. Lazar, and A. Orda. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Transactions on Networking (TON)*, 5(1):161–173, 1997.

[9] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.

[10] Christos Papadimitriou. Algorithms, games, and the internet. In *STOC '01: Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 749–753, New York, NY, USA, 2001. ACM.

[11] T. Roughgarden. Stackelberg scheduling strategies. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 104–113. ACM New York, NY, USA, 2001.

[12] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[13] C. Swamy. The effectiveness of Stackelberg strategies and tolls for network congestion games. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, page 1142. Society for Industrial and Applied Mathematics, 2007.

[14] H. Von Stackelberg. *Marktform und gleichgewicht*. J. Springer, 1934.