



HAL
open science

Semilinear Representations for Series-Parallel Atomic Congestion Games

Nathalie Bertrand, Nicolas Markey, Suman Sadhukhan, Ocan Sankur

► **To cite this version:**

Nathalie Bertrand, Nicolas Markey, Suman Sadhukhan, Ocan Sankur. Semilinear Representations for Series-Parallel Atomic Congestion Games. FSTTCS 2022 - 42nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, Dec 2022, Chennai, France. pp.1-20, 10.4230/LIPIcs.FSTTCS.2022.32 . hal-03937259

HAL Id: hal-03937259

<https://hal.science/hal-03937259v1>

Submitted on 13 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Semilinear Representations for Series-Parallel Atomic Congestion Games

Nathalie Bertrand ✉ 

Univ Rennes, Inria, CNRS, IRISA, France

Nicolas Markey ✉ 

Univ Rennes, Inria, CNRS, IRISA, France

Suman Sadhukhan ✉

Univ Rennes, Inria, CNRS, IRISA, France

University of Haifa, Israel

Ocan Sankur ✉ 

Univ Rennes, Inria, CNRS, IRISA, France

Abstract

We consider atomic congestion games on series-parallel networks, and study the structure of the sets of Nash equilibria and social local optima on a given network when the number of players varies. We establish that these sets are definable in Presburger arithmetic and that they admit semilinear representations whose all period vectors have a common direction. As an application, we prove that the prices of anarchy and stability converge to 1 as the number of players goes to infinity, and show how to exploit these semilinear representations to compute these ratios precisely for a given network and number of players.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory; Theory of computation → Solution concepts in game theory

Keywords and phrases congestion games, Nash equilibria, Presburger arithmetic, semilinear sets, price of anarchy

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2022.32

Funding This research was partially supported by ISF grant no. 1679/21.

1 Introduction

Network congestion games are used to model situations in which agents share resources such as routes or bandwidth [31], and have applications in communication networks (e.g. [1]). We consider the *atomic* variant of these games, where each player controls one unit of flow and must assign it to a path in the network. All players using an edge then incur a cost (a.k.a. latency) that is a nondecreasing function of the number of players using the same edge. Since all players try to minimize their own cost, this yields a noncooperative multiplayer game. It is well known that Nash equilibria exist in these games [31] but that they can be inefficient, that is, a global measure such as the total cost, or the makespan may not be minimized by Nash equilibria [30].

To quantify this inefficiency, [27] introduced the notion of *price of anarchy (PoA)*, which is the ratio of the cost of the worst Nash equilibrium and the social optimum. Here, social optimum refers to the sum of the individual costs. A tight bound of $\frac{5}{2}$ on this ratio was given in [3, 9]. Various works have studied bounds on the PoA for restricted classes of graphs or types of cost functions; see [28]. While the price of anarchy is interesting to understand behaviors that emerge in a system from a worst-case perspective, the best-case is also interesting if, for instance, the network designer is able to select a Nash equilibrium.



© Nathalie Bertrand, Nicolas Markey, Suman Sadhukhan, and Ocan Sankur; licensed under Creative Commons License CC-BY 4.0

42nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2022).

Editors: Anuj Dawar and Venkatesan Guruswami; Article No. 32; pp. 32:1–32:20



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The *price of stability* (*PoS*) is thus the ratio between the cost of the best Nash equilibrium and the social cost, and was studied in [2]; a bound of $1 + \sqrt{3}/3$ was given in the atomic case; see [7, 8, 2].

These bounds have been studied for restrictions of this problem such as particular classes of graphs. One such class is *series-parallel networks* which are built using single edges, or parallel and serial composition of smaller series-parallel networks (see e.g. [34]). On these networks with affine cost functions, [24] reports that PoA is between $\frac{27}{19}$ and 2, where a previous known lower bound due to [19] was $\frac{15}{11}$. [19] also proves an upper bound on PoA for *extension-parallel* networks which are a strict subclass of series-parallel networks.

While tight bounds are known for atomic network congestion games and even for particular subclasses, this does not help one to evaluate the price of anarchy of a given specific game with a given number of players. In fact, the upper bounds mentioned above are obtained by building particular networks, and these are shown to be tight by exhibiting families of instances in which both the networks and the number of players vary. One of our objectives is to provide tools to analyze a given network congestion game, by computing both ratios precisely for varying numbers of players. We are interested both in the case of a given number of players, and in the case of the limit behavior. Note that we are considering a hard problem since computing *extreme* Nash equilibria, that is, best and worst ones, is NP-hard in networks with only three and two players respectively [33].

In this paper, we consider series-parallel networks with linear cost functions and establish interesting properties of their Nash equilibria and social optima. We start with the observation that Nash equilibria and *locally* social optima can be expressed in Presburger arithmetic; it follows that these sets admit *semilinear* representations [22]. Our main result is that that these semilinear sets have a particular structure. In fact, the flows (i.e. edge loads) induced by Nash equilibria and local social optima admit semilinear representations with a common direction for all period vectors, which is moreover efficiently computable. We call this direction the *characteristic vector* of the network. Intuitively, this vector determines how the flow evolves in Nash equilibria and social local optima as the number of players increases.

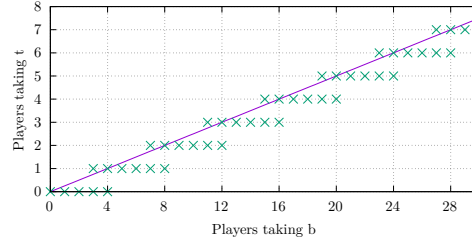

We believe that the form of these representations and the characteristic vector have an interest on their own. We give one application of these representations here, namely, that the PoA and the PoS both tend to 1 as the number of players goes to infinity in series-parallel networks with linear cost functions with positive coefficients. This result was proven recently [37]; we provide here new techniques and thus an alternative proof. Observe that a similar result holds in nonatomic congestion games [11, 10] (see Section 6 for a discussion).

We also illustrate how these semilinear representations allow one to study the evolution of PoA and PoS for a given network. The computation of these representations is an expensive step, but once this is done, thanks to the particular form of the representations, for any n , one can easily query the exact value for PoA and PoS in the network instantiated with n players. One can thus analyze both ratios precisely and specifically for a given network as a function of n , while the limits will always be 1.

Illustrating Example

Let us illustrate our results on a simple example. Consider the network with two parallel links in Figure 1a. There are n players who would like to go from **src** to **tgt**, by taking either the bottom edge (b) with the cost function $x \mapsto x$, or the top edge (t) with cost function $x \mapsto 4x$. The cost function determines the cost each player pays when taking a given edge, and it is a function of the total number of players using the same edge. For instance, for $n = 4$, assume

that 3 players use b , and 1 uses t . Then, each player using b pays a cost of 3, while the only player using t pays 4. Here, a strategy profile can be seen as a pair $(k, n - k)$ determining how many players take t and how many take b .




(a) A network with two parallel links where each player can choose either the top edge (t) with cost $x \mapsto 4x$, or the bottom edge (b) with cost $x \mapsto x$. (b) Nash equilibria in the network on the left for varying total numbers of players. There is a point at coordinates (x, y) if a strategy profile assigning x players to b , and y players to t is a Nash equilibrium.

■ **Figure 1** Analysis of a simple network congestion game.

Intuitively, in Nash equilibria, the number of players taking b should be *roughly* four times the number of them taking t ; so $\frac{4}{5}$ of them should take b , and $\frac{1}{5}$ of them should take t . This would make sure that both edges have identical cost, and make profitable deviations impossible. Although this intuition holds in the nonatomic case, players cannot always be split with this proportion in the atomic case, as in the case of $n = 4$ above; and there are indeed equilibria that do not match this proportion exactly. Figure 1b shows the Nash equilibria in this game, while the line with direction $(4, 1)$ shows the ideal distribution (as in the nonatomic case). Not all Nash equilibria are on this line, but one can notice that they do form a tube around this line that go in the same direction. Formally, our results determine that the Nash equilibria form the semilinear set $B_{NE} + \vec{\delta} \cdot \mathbb{N}$, where $\vec{\delta} = (4, 1)$, and $B_{NE} = \{(0, 0), (1, 0), (2, 0), (3, 0), (4, 0), (3, 1)\}$. In other terms, it is the union of the integer points of six lines with the same direction vector $\vec{\delta}$.

Similarly, we describe the set of *locally* social optimal profiles and show that it admits a semilinear representation. Locally optimal profiles are those in which the social cost cannot be decreased by changing the strategy of a single player; the formal definition is given in Section 3.1. It turns out that these have a structure very similar to that of Nash equilibria. In our example, local optima are given by $B_{SO} + \vec{\delta} \cdot \mathbb{N}$, thus with the same vector $\vec{\delta}$ as above, and with $B_{SO} = \{(0, 0), (1, 0), (2, 0), (2, 1)\}$.

The particular structures of these semilinear sets we obtain allow us to compute the prices of anarchy and stability for any given number of players. In fact, given n , one can easily find, in a set $B + \vec{\delta} \cdot \mathbb{N}$, all strategy profiles with n players. So one can compute the worst and the



■ **Figure 2** Prices of anarchy and stability as a function of n .

best Nash equilibria, as well as the social optimum for any given n . The plot in Fig. 2 shows how PoA and PoS evolve as n increases, and was calculated from the previous representations. This plot illustrates our objective of analyzing the inefficiency of these games precisely for varying n . Even in this simple example, PoA and PoS can significantly vary depending on n , and they are far from the known tight bounds for the whole class of networks. One can notice in Fig. 2 that both PoA and PoS seem to converge to 1 as n increases. This is indeed the case and is a consequence of our results (Theorem 14). We believe that this approach can allow one to better understand the specific network under analysis. Section 5 contains more examples.

Paper Overview. We provide formal definitions in Section 2. Section 3 characterizes the form of semilinear representations of local social optima; and Section 4 proves that of Nash equilibria. Section 5 shows how to use these representations to compute PoA and PoS. We provide more discussion on related work in Section 6, and present conclusions in Section 7.

2 Preliminaries

2.1 Network Congestion Games and Series-Parallel Arenas

A *network* is a weighted graph $\mathcal{A} = \langle V, E, \text{orig}, \text{dest}, \text{wgt}, \text{src}, \text{tgt} \rangle$, where V is a finite set of *vertices*, E is a finite set of *edges*, $\text{orig}: E \rightarrow V$ and $\text{dest}: E \rightarrow V$ indicate the origin and destination of each edge, $\text{wgt}: E \rightarrow \mathbb{N}_{>0}$ is a *weight function* assigning *positive* weights to edges, and src and tgt are, respectively, a *source* and a *target* states. Let $\text{In}(v)$ and $\text{Out}(v)$ denote, respectively, the set of incoming and outgoing edges of v . We restrict to acyclic networks in which all vertices are reachable from src and tgt is reachable from all vertices.

A *path* π of \mathcal{A} is a sequence $e_1 e_2 \dots e_n$ of edges with $\text{dest}(e_i) = \text{orig}(e_{i+1})$ for all $1 \leq i \leq n-1$. For an edge e and a path $\pi = e_1 e_2 \dots e_n$, we write $e \in \pi$ if $e = e_i$ for some $1 \leq i \leq |\pi|$. We will sometimes see paths as sets of edges and apply set operations such as intersection and set difference. Let $\text{Paths}_{\mathcal{A}}(s, t)$ denote the set of simple paths from s to t in \mathcal{A} , and let $\text{Paths}_{\mathcal{A}} = \text{Paths}_{\mathcal{A}}(\text{src}, \text{tgt})$.

In this work, we consider *series-parallel networks* [34]. These are built inductively from single edges using *serial* and *parallel* composition. Two networks $\mathcal{A}_1 = \langle V_1, E_1, \text{orig}_1, \text{dest}_1, \text{wgt}_1, \text{src}_1, \text{tgt}_1 \rangle$ and $\mathcal{A}_2 = \langle V_2, E_2, \text{orig}_2, \text{dest}_2, \text{wgt}_2, \text{src}_2, \text{tgt}_2 \rangle$ are *composed in series* to a new network denoted by $\mathcal{A}_1; \mathcal{A}_2 = \langle V, E, \text{orig}, \text{dest}, \text{wgt}, \text{src}, \text{tgt} \rangle$ obtained by taking the disjoint union of \mathcal{A}_1 and \mathcal{A}_2 , and merging the vertices tgt_1 and src_2 , and setting $\text{src} = \text{src}_1$, $\text{tgt} = \text{tgt}_2$. Two networks $\mathcal{A}_1 = \langle V_1, E_1, \text{orig}_1, \text{dest}_1, \text{wgt}_1, \text{src}_1, \text{tgt}_1 \rangle$ and $\mathcal{A}_2 = \langle V_2, E_2, \text{orig}_2, \text{dest}_2, \text{wgt}_2, \text{src}_2, \text{tgt}_2 \rangle$ are *composed in parallel* to a new network denoted by $\mathcal{A}_1 \parallel \mathcal{A}_2 = \langle V, E, \text{orig}, \text{dest}, \text{wgt}, \text{src}, \text{tgt} \rangle$ obtained by taking the disjoint union of \mathcal{A}_1 and \mathcal{A}_2 , and merging src_1 and src_2 , then merging tgt_1 and tgt_2 , and setting $\text{src} = \text{src}_1$, $\text{tgt} = \text{tgt}_1$. A network is said to be *series-parallel* if it is either a single edge, or it is a serial or parallel composition of series-parallel graphs.

A *network congestion game* (NCG) is a pair $\mathcal{G} = \langle \mathcal{A}, n \rangle$ where $\mathcal{A} = \langle V, E, \text{orig}, \text{dest}, \text{wgt}, \text{src}, \text{tgt} \rangle$ is a network, and $n \in \mathbb{N}$ is the number of players in the game. We consider the symmetric case where all players start at src and want to reach tgt . A *strategy* of a player is a path in $\text{Paths}_{\mathcal{A}}$. The setting can be seen as a one-shot game in which each player selects a strategy simultaneously. In our study, we do not need to identify players, we thus represent strategy profiles by counting how many players choose each strategy. That is, a *strategy profile* is a tuple $\vec{p} = (p_{\pi})_{\pi \in \text{Paths}_{\mathcal{A}}}$ where p_{π} is the number of players taking path π . In this case, the number of players is given by $\|\vec{p}\| = \sum_{\pi \in \text{Paths}_{\mathcal{A}}} p_{\pi}$; thus \vec{p} is a strategy profile in

the game $\langle \mathcal{A}, \|\vec{p}\| \rangle$. Let $\mathfrak{S}(\mathcal{A})$ denote the set of all strategy profiles, and $\mathfrak{S}_n(\mathcal{A})$ the set of strategy profiles with n players, that is, $\{\vec{p} \in \mathfrak{S}(\mathcal{A}) \mid n = \|\vec{p}\|\}$. For $\pi \in \text{Paths}_{\mathcal{A}}$, let $\vec{p} + \pi$ (resp. $\vec{p} - \pi$) denote the strategy profile obtained by incrementing (resp. decrementing) p_π by one.

Another useful notion we use is the *flow* of a strategy profile, which consists in the projection of a strategy profile to edges. Formally, given a strategy profile \vec{p} , $\text{flow}(\vec{p}) = (q_e)_{e \in E}$ where $q_e = \sum_{\pi \in \text{Paths}_{\mathcal{A}}: e \in \pi} p_\pi$, that is, q_e is the number of players that use the edge e in the profile \vec{p} . This vector satisfies the following flow equations:

$$\forall v \in V \setminus \{\text{src}, \text{tgt}\}, \quad \sum_{e \in \text{In}(v)} q_e = \sum_{e \in \text{Out}(v)} q_e. \quad (1)$$

We refer to a vector $\vec{q} = (q_e)_{e \in E}$ with nonnegative coefficients satisfying (1) as a *flow*; and denote by $\mathcal{F}(\mathcal{A})$ the set of all flows. Observe that $\mathcal{F}(\mathcal{A})$ is the image of $\mathfrak{S}(\mathcal{A})$ by flow . For a flow \vec{q} , let $\|\vec{q}\| = \sum_{e \in E: \text{orig}(e) = \text{src}} q_e$, which is the number of players. Let $\mathcal{F}_n(\mathcal{A})$ define the set of flows with n players as follows: $\mathcal{F}_n(\mathcal{A}) = \{\vec{q} \in \mathcal{F} \mid n = \|\vec{q}\|\}$. Observe that this corresponds to a flow of size n , and that several strategy profiles can project to the same flow.

For a strategy profile \vec{p} and $\pi \in \text{Paths}_{\mathcal{A}}$, each player using the path π incurs a cost equal to $\text{cost}_\pi(\vec{p}) = \sum_{e \in \pi} \text{wgt}(e) \cdot \text{flow}_e(\vec{p})$, where $\text{flow}_e(\vec{p})$ is the number of players using edge e in the strategy profile \vec{p} . The *social cost* of a strategy profile \vec{p} is the sum of the costs for all players, i.e., $\text{soccost}(\vec{p}) = \sum_{\pi \in \text{Paths}_{\mathcal{A}}} p_\pi \cdot \text{cost}_\pi(\vec{p})$. The *social optimum* of the game $\mathcal{G} = \langle \mathcal{A}, n \rangle$ is $\text{opt}(\mathcal{G}) = \min_{\vec{p} \in \mathfrak{S}_n(\mathcal{A})} \text{soccost}(\vec{p})$. A strategy profile $\vec{p} \in \mathfrak{S}_n$ in a game $\mathcal{G} = \langle \mathcal{A}, n \rangle$ is *socially optimal* if $\text{soccost}(\vec{p}) = \text{opt}(\mathcal{G})$.

Observe that the cost of a path in a strategy profile, and the social cost of a strategy profile, are determined by the flow of that profile. Thus, we define the social cost of a flow \vec{q} as $\text{soccost}(\vec{q}) = \sum_{e \in E} q_e^2 \cdot \text{wgt}(e)$ (in fact, q_e players use strategies that include e , and each of them pays $q_e \text{wgt}(e)$ for crossing this edge). A flow $\vec{q} \in \mathcal{F}_n$ is socially optimal if $\text{soccost}(\vec{q}) = \text{opt}(\mathcal{G})$.

A strategy profile \vec{p} is a Nash equilibrium if no player can reduce their cost by unilaterally changing strategy, i.e., if

$$\forall \pi \in \text{Paths}_{\mathcal{A}}, p_\pi > 0 \rightarrow \forall \pi' \in \text{Paths} \setminus \{\pi\}, \text{cost}_\pi(\vec{p}) \leq \text{cost}_{\pi'}(\vec{p}'), \quad (2)$$

where \vec{p}' is defined by $p'_\pi = p_\pi - 1$, $p'_{\pi'} = p_{\pi'} + 1$, and $p'_\tau = p_\tau$ for all other paths τ . In fact, $\text{cost}_\pi(\vec{p})$ is the cost of a player playing π in the profile \vec{p} , while $\text{cost}_{\pi'}(\vec{p}')$ is their cost in the new profile \vec{p}' obtained by switching from π to π' .

Let $\text{NE}(\mathcal{A})$ denote the set of strategy profiles satisfying (2), that is, the set of Nash equilibria, and let $\text{NE}_n(\mathcal{A})$ denote the set of Nash equilibria for n players. The *price of anarchy* is the ratio of the social cost of the worst Nash equilibrium, and the social optimum: $\text{PoA}(\langle \mathcal{A}, n \rangle) = \max_{\vec{p} \in \text{NE}_n} \text{soccost}(\vec{p}) / \text{opt}(\langle \mathcal{A}, n \rangle)$. The *price of stability* is the ratio of the social cost of the best Nash equilibrium, and the social optimum: $\text{PoS}(\langle \mathcal{A}, n \rangle) = \min_{\vec{p} \in \text{NE}_n} \text{soccost}(\vec{p}) / \text{opt}(\langle \mathcal{A}, n \rangle)$.

2.2 Presburger Arithmetic and Semilinear Sets

We recall the definition and some basic properties of semilinear sets; see e.g. [23] for more details. A set $S \subseteq \mathbb{N}^m$ is called *linear* if there is a *base vector* $\vec{b} \in \mathbb{N}^m$ and a finite set of *period vectors* $P = \{\vec{\delta}_1, \vec{\delta}_2, \dots, \vec{\delta}_p\}$ such that $S = \vec{b} + \vec{\delta}_1 \cdot \mathbb{N} + \vec{\delta}_2 \cdot \mathbb{N} + \dots + \vec{\delta}_p \cdot \mathbb{N}$, that is $S = \{\vec{b} + \lambda_1 \vec{\delta}_1 + \dots + \lambda_p \vec{\delta}_p \mid \lambda_1, \dots, \lambda_p \in \mathbb{N}\}$. Such a linear set S will be denoted as $L(\vec{b}, P)$.

A set $S \subseteq \mathbb{N}^m$ is said to be *semi-linear* if it is a finite union of linear sets. Therefore, a semi-linear set S can be written in the form $S = \cup_{i \in I} L(\vec{b}_i, P_i)$ where I is a finite set, P_i 's are finite sets of period vectors, and the \vec{b}_i are the base vectors of the same dimension.

Note that in a linear set $L(\vec{b}, P)$, P can be empty, which corresponds to a singleton set. Thus, finite sets are semilinear; and the union of any semilinear set with a finite set is semilinear. Furthermore, each semilinear set admits a *non-ambiguous* representation in the sense that $S = \cup_{i \in I} L(\vec{b}_i, P_i)$ such that each P_i is linearly independent and $L(\vec{b}_i, P_i) \cap L(\vec{b}_j, P_j) = \emptyset$ for all $i \neq j \in I$ [16, 25].

Presburger arithmetic is the first-order theory of integers without multiplication. It is well-known that any set expressible in Presburger arithmetic is *semilinear* [22]. So, in order to show that a set is semilinear, one can either exhibit its semilinear representation, or show that it is expressible in Presburger arithmetic.

3 Local Social Optima

In this section, our goal is to obtain a representation of social optima in a given network congestion game as a function of the number n of players. Characterizing the social optimum directly by a formula brings two difficulties. First, expressing that a flow \vec{q} is optimal would require to quantify over all flows \vec{q}' and writing that $\text{soccost}(\vec{q}) \leq \text{soccost}(\vec{q}')$, so such a formula contains universal quantifiers. Second, the formula is quadratic since $\text{soccost}(\vec{q}) = \sum_{e \in E} q_e^2 \cdot \text{wgt}(e)$, so this cannot be represented by a semilinear set.

Here, we introduce the notion of *local optimality* which allows us to circumvent both difficulties, providing semilinear representations which, moreover, allow us to compute the global optimum.

3.1 Locally-Optimal Profiles

Let us fix a series-parallel network \mathcal{A} . Intuitively, a strategy profile is *locally-optimal* if the social cost cannot be reduced by exchanging one path for another. Formally, $\vec{p} \in \mathfrak{S}_n(\mathcal{A})$ is locally-optimal if for all $\pi, \pi' \in \text{Paths}_{\mathcal{A}}$ with $p_\pi > 0$, $\text{soccost}(\vec{p}) \leq \text{soccost}(\vec{p} - \pi + \pi')$. By extension, a flow $\vec{q} \in \mathcal{F}_n(\mathcal{A})$ is locally-optimal if it is the image of a locally-optimal strategy profile. Observe that the (global) social optimum is locally-optimal.

In the following lemma, we see paths π, π' as sets of edges.

► **Lemma 1.** *In a network congestion game $\langle \mathcal{A}, n \rangle$, a flow \vec{q} is locally-optimal if, and only if, for all $\pi, \pi' \in \text{Paths}_{\mathcal{A}}$ such that $\forall e \in \pi, q_e > 0$,*

$$\sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot (2q_e - 1) \leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (2q_e + 1). \quad (3)$$


We define $\text{LocOpt}(\mathcal{A})$ to be the set of locally-optimal flows, and $\text{LocOpt}_n(\mathcal{A})$ those with n players. It follows from Lemma 1 that $\text{LocOpt}(\mathcal{A})$ and $\text{LocOpt}_n(\mathcal{A})$ are expressible in Presburger arithmetic, and are thus semilinear. We will now characterize the form of the semilinear set describing $\text{LocOpt}(\mathcal{A})$ by proving that it admits a single and computable period vector $\vec{\delta}$, that is, $\text{LocOpt}(\mathcal{A})$ can be written as $B \cup \bigcup_{i \in I} L(\vec{b}_i, \vec{\delta})$, where B is a finite set of flows, I is a finite set of indices, and $(\vec{b}_i)_{i \in I}$ are the base vectors.

3.2 Large Numbers of Players

To simplify the proof of the characterization of the period vector $\vec{\delta}$, we would like to consider instances in which (3) holds for *all* paths $\pi, \pi' \in \text{Paths}_{\mathcal{A}}$, that is, we would like to get rid of the assumption on π in Lemma 1. It turns out that the assumption that $q_e > 0$ for all edges e

of π holds whenever the number of players is large enough. Moreover, we do not lose generality by focusing on these instances; in fact, as we will see, a semilinear representation with the same period vector \vec{d} for *all* locally-optimal profiles can be derived once this representation is established for instances with large numbers of players.

The next lemma shows that all edges are used in locally-optimal profiles whenever the number of players is sufficiently large. This property is specific to series-parallel graphs and may not hold in a more general network, as the following example shows.



■ **Figure 3** A network with an edge $v \rightarrow v'$ that is never used in any locally-optimal profile.

► **Example 2.** Consider the network of Fig. 3. We claim that the edge from v to v' cannot be used in any locally-optimal profile. Write a , b and c for the number of players taking paths $\pi_a : \text{src} \rightarrow v \rightarrow \text{tgt}$, $\pi_b : \text{src} \rightarrow v \rightarrow v' \rightarrow \text{tgt}$ and $\pi_c : \text{src} \rightarrow v' \rightarrow \text{tgt}$, respectively. Observe (by contradiction) that if there are at least two players, then paths π_a and π_c will be taken by at least one player. Writing Eq. (3) for π_a and π_c yields $-1 \leq a - c \leq 1$. Assuming $b > 0$, we can also apply Eq. (3) to π_b and π_a , and get $-2a + 8b + 6c - 5 \leq 0$. It follows $8b + 4c - 7 \leq 0$. This cannot be preserved when the number of players grows. Hence π_b cannot be used in any locally-optimal profile. ◀

► **Lemma 3.** *In all series-parallel networks \mathcal{A} , there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, all flows $\vec{q} \in \text{LocOpt}_n(\mathcal{A})$ are such that $q_e > 0$ for all $e \in E$.*

Proof. We inductively define a value $n_0(\mathcal{A})$ for each series-parallel network \mathcal{A} , such that $n_0(\mathcal{A}) \geq 4$ (for technical reasons) and satisfying the following property:

$$\forall k \geq 1, \forall n \geq k \cdot n_0(\mathcal{A}), \forall \vec{q} \in \text{LocOpt}_n(\mathcal{A}), \forall e \in E, q_e \geq k.$$

The lemma follows by taking $k = 1$.

If \mathcal{A} is a single edge, we define $n_0(\mathcal{A}) = 4$, and the property trivially holds.

If $\mathcal{A} = \mathcal{A}_1; \mathcal{A}_2$, then we define $n_0(\mathcal{A}) = \max(n_0(\mathcal{A}_1), n_0(\mathcal{A}_2))$. Observe that if \vec{q} is locally-optimal in \mathcal{A} , then each $\vec{q}|_{\mathcal{A}_i}$ is locally-optimal in \mathcal{A}_i , and they have the same number of players. So the property holds by induction.

The non-trivial case is when $\mathcal{A} = \mathcal{A}_1 \parallel \mathcal{A}_2$. Let $m = \max(n_0(\mathcal{A}_1), n_0(\mathcal{A}_2))$, and $n_0(\mathcal{A}) = 2|E|Wm^2$, where E is set of edges of \mathcal{A} and W is the largest weight of \mathcal{A} (notice that $W \geq 1$). Then $m \geq 4$ and $n_0(\mathcal{A}) \geq 4$. Take $n \geq kn_0(\mathcal{A})$, and $\vec{q} \in \text{LocOpt}_n$.

We show that for all $i \in \{1, 2\}$, we have $\|\vec{q}|_{\mathcal{A}_i}\| \geq kn_0(\mathcal{A}_i)$. Towards a contradiction, assume for example that $\|\vec{q}|_{\mathcal{A}_1}\| < kn_0(\mathcal{A}_1)$. Then $\|\vec{q}|_{\mathcal{A}_2}\| > n - kn_0(\mathcal{A}_1) \geq 2k|E|W(m^2 - m)$ (because $kn_0(\mathcal{A}_1) \leq km \leq 2k|E|Wm$).

For all pair $(\pi_1, \pi_2) \in \text{Paths}_{\mathcal{A}_1} \times \text{Paths}_{\mathcal{A}_2}$, by (3) we have that $\sum_{e \in \pi_2} \text{wgt}(e)(2q_e - 1) \leq \sum_{e \in \pi_1} \text{wgt}(e)(2q_e + 1)$. Take an arbitrary edge e_0 of π_2 ; then e_0 does not appear in π_1 , and we get

$$2q_{e_0} \leq 2W \sum_{e \in \pi_1} q_e + \sum_{e \in \pi_2} \text{wgt}(e) + \sum_{e \in \pi_1} \text{wgt}(e) < 2W|E| \cdot kn_0(\mathcal{A}_1) + 2W|E| \leq 2W|E|(km + 1).$$

Now, take $k' = 2k|E|W(m-1)$. Then $k' \geq 1$, and $\|\vec{q}|_{\mathcal{A}_2}\| > k'm \geq kn_0(\mathcal{A}_2)$. By induction hypothesis applied to \mathcal{A}_2 , we have $q_{e_0} \geq k'$, which implies

$$2k|E|W(m-1) \leq W|E|(km+1)$$

hence $m \leq 2 + \frac{1}{k} \leq 3$, which is a contradiction since $m \geq 4$. \blacktriangleleft

Let $\text{LocOpt}_{\geq n_0}(\mathcal{A}) = \bigcup_{n \geq n_0} \text{LocOpt}_n(\mathcal{A})$. By the previous lemma, all $\vec{q} \in \text{LocOpt}_{\geq n_0}(\mathcal{A})$ satisfy (3) for all pairs of paths π, π' . Observe that $\text{LocOpt}(\mathcal{A})$ and $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ have the same period vectors. In fact, $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ differs from $\text{LocOpt}(\mathcal{A})$ only by a finite set; so, given a semilinear representation $\bigcup_{i \in I} L(\vec{b}_i, P_i)$ for the former, one can obtain a representation for the latter as $B \cup \bigcup_{i \in I} L(\vec{b}_i, P_i)$ where B is the finite difference between the two. Therefore, establishing that $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ has a single period vector suffices to prove the same result for $\text{LocOpt}(\mathcal{A})$. Note also that I is non-empty here since the set $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ is infinite.

The following lemma shows that the period vectors of $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ assign the same cost to all paths. Intuitively, this is because if the cost along two paths were different in the period vector, then by adding a large number of copies of the period vector to its base vector, one could amplify this difference and obtain a vector that is not locally-optimal.

► **Lemma 4.** *In a series-parallel network \mathcal{A} , for all period vectors $\vec{d} \in \mathbb{N}^E$ of a semilinear representation of $\text{LocOpt}_{\geq n_0}(\mathcal{A})$, there exists $\kappa \geq 0$ such that for all $\pi \in \text{Paths}_{\mathcal{A}}$, we have $\sum_{e \in \pi} \text{wgt}(e) \cdot d_e = \kappa$.*

3.3 A Unique Period Vector: The Characteristic Vector

We now establish that $\text{LocOpt}(\mathcal{A})$ admits a unique period vector. For any $\kappa \in \mathbb{R}$, we study the following system $\mathcal{E}(\kappa)$ of equations with unknowns $\{q_e\}_{e \in E}$:

$$\forall \pi \in \text{Paths}_{\mathcal{A}}, \sum_{e \in \pi} \text{wgt}(e) \cdot q_e = \kappa, \quad (4)$$

$$\forall v \in V \setminus \{\text{src}, \text{tgt}\}, \sum_{e \in \text{In}(v)} q_e - \sum_{e \in \text{Out}(v)} q_e = 0. \quad (5)$$

Note that all period vectors of $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ satisfy the above. In fact, (4) comes from Lemma 4, and (5) is the set of flow equations (1).

The following lemma states that $\mathcal{E}(\kappa)$ has a unique solution whenever we fix κ .

► **Lemma 5.** *For a series-parallel network \mathcal{A} , for each $\kappa \in \mathbb{R}$, the system $\mathcal{E}(\kappa)$ admits a unique solution.*

The system $\mathcal{E}(\kappa)$ is actually the characterization of the flows of Nash equilibria in the non-atomic congestion games, and the unicity of the solution of this equation system is known (see [12]). We represent the system $\mathcal{E}(\kappa)$ in the matrix form as $M_{\mathcal{A}} \cdot X = \kappa \vec{b}$, where $X = (q_e)_{e \in E}$ is the vector of unknowns, and \vec{b} a $\{0, 1\}$ -column vector. Lemma 5 means that $M_{\mathcal{A}}$ admits a left-inverse $M_{\mathcal{A}}^{-1}$. It follows that all period vectors can be written as $\kappa M_{\mathcal{A}}^{-1} \vec{b}$ for some κ . Since $M_{\mathcal{A}}$ and \vec{b} have integer coefficients, $M_{\mathcal{A}}^{-1} \vec{b}$ is a vector with rational coefficients.

► **Definition 6 (Characteristic Vector).** *Let κ_0 denote the least rational number such that $\kappa_0 M_{\mathcal{A}}^{-1} \vec{b}$ is integer. We define $\vec{\delta} = \kappa_0 M_{\mathcal{A}}^{-1} \vec{b}$, called the characteristic vector of \mathcal{A} .*

Since period vectors of $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ have natural number coefficients, any κ corresponding to a period vector is also a natural. We have that all period vectors are integer multiples of $\vec{\delta}$. In fact, if $\vec{d} = \kappa M_{\mathcal{A}}^{-1} \vec{b}$ is a period vector, then $\vec{d} = \frac{\kappa}{\kappa_0} \vec{\delta}$, and $\frac{\kappa}{\kappa_0}$ is an integer since otherwise its denominator would divide $\vec{\delta}$, which would contradict the minimality of κ_0 .


► **Corollary 7.** *Consider a series-parallel network \mathcal{A} , and its characteristic vector $\vec{\delta}$. There exist finite sets of vectors B and $\{\vec{b}_i \mid i \in I\}$ such that $\text{LocOpt}(\mathcal{A}) = B \cup \bigcup_{i \in I} L(\vec{b}_i, \vec{\delta})$.*

Proof. Since each linear set can be assumed to have linearly-independent period vectors (Section 2.2) all linear sets included in $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ can be written in the form of $\vec{b} + m\vec{\delta}\mathbb{N}$,

We show that $m = 1$. By Lemma 3, \vec{b} has only positive coefficients, so it satisfies (3) for all pairs of paths. We show that $\vec{b} + k\vec{\delta}$ also satisfies (3) for all $k \geq 0$, which proves that $L(\vec{b}, \vec{\delta})$ is included in $\text{LocOpt}_{\geq n_0}(\mathcal{A})$. For all paths π, π' , consider (3) by adding an identical term to both sides:

$$\begin{aligned} \sum_{e \in \pi \setminus \pi'} \text{wgt}(e)(2q_e - 1) + 2k \sum_{e \in \pi} \text{wgt}(e)\delta_e &\leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (2q_e + 1) + 2k \sum_{e \in \pi'} \text{wgt}(e)\delta_e \\ \sum_{e \in \pi \setminus \pi'} \text{wgt}(e)(2q_e - 1) + 2k \sum_{e \in \pi \setminus \pi'} \text{wgt}(e)\delta_e &\leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (2q_e + 1) + 2k \sum_{e \in \pi' \setminus \pi} \text{wgt}(e)\delta_e \\ \sum_{e \in \pi \setminus \pi'} \text{wgt}(e)(2(q_e + k\delta_e) - 1) &\leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (2(q_e + k\delta_e) + 1). \end{aligned}$$

Therefore, $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ can be written in the form $\bigcup_{i \in I} L(\vec{b}_i, \vec{\delta})$; and since $\text{LocOpt}(\mathcal{A})$ differs from $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ by a finite set, it can be represented by $B \cup \bigcup_{i \in I} L(\vec{b}_i, \vec{\delta})$ where $B = \text{LocOpt}(\mathcal{A}) \setminus \text{LocOpt}_{\geq n_0}(\mathcal{A})$. ◀



■ **Figure 4** Network whose set of local optima admit several period vectors.

► **Remark 8.** Note that we chose here to study the semilinear representations of locally-optimal flows, rather than locally-optimal strategy profiles. The latter are also expressible in Presburger arithmetic, thus also admit a semilinear representation. However, the set of locally-optimal strategy profiles admit, in general, several linearly independent period vectors, so their representation is more complex, and more difficult to use, for instance, to compute the global optimum for given number n of players.

To see this, consider the example of Fig. 4. Consider the strategy profile \vec{p} with $p_{ac} = p_{bd} = 1$ and $p_{ad} = p_{bc} = 0$; and \vec{p}' such that $p'_{ac} = p'_{bd} = 0$ and $p'_{ad} = p'_{bc} = 1$. For all $k \geq 0$, both $k\vec{p}$ and $k\vec{p}'$ are socially optimal, but they are linearly independent. In larger networks, there can be a larger number of period vectors due to similar phenomena. Notice however that $\text{flow}(\vec{p}) = \text{flow}(\vec{p}')$, that is, the projections of these period vectors to their flows are identical, and are, in fact, equal to the characteristic vector.

4 Nash Equilibria

In this section, we show how to compute a semilinear representation of the flows of Nash equilibria, which will allow us to compute the costs of the best and the worst equilibria.

32:10 Semilinear Representations for Series-Parallel Atomic Congestion Games

The following lemma is a characterization of Nash equilibria that follows from (2).

► **Lemma 9.** *Given a network \mathcal{A} , a strategy profile \vec{p} is a Nash equilibrium if, and only if,*

$$\forall \pi, \pi' \in \text{Paths}, p_\pi > 0 \implies \sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot q_e \leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (q_e + 1), \quad (6)$$

where $\vec{q} = \text{flow}(\vec{p})$.

It follows from the previous lemma that Nash equilibria, but also their flows, are definable in Presburger arithmetic.

► **Lemma 10.** *The sets $\text{NE}(\mathcal{A})$ and $\text{flow}(\text{NE}(\mathcal{A}))$ are semilinear.*

The study of the semilinear representation is similar to what was done for locally-optimal profiles. We prove that the period vectors of $\text{flow}(\text{NE}(\mathcal{A}))$ are colinear to $\vec{\delta}$, thus establishing the form of their semilinear representation. In particular, we use Lemma 5 and show that the period vectors are multiples of the characteristic vector. Here, we consider $\text{flow}(\text{NE}(\mathcal{A}))$ rather than $\text{NE}(\mathcal{A})$ due to Remark 8 which also holds for Nash equilibria.

The following lemma shows that a semilinear representation of $\text{flow}(\text{NE}(\mathcal{A}))$ can be deduced from that of $\text{NE}(\mathcal{A})$ and follows by the definition of flow . In particular, all period vectors of the former are projections of the period vectors of the latter by flow .

► **Lemma 11.** *For any semilinear representation $B \cup \bigcup_{i \in I} L(\vec{b}_i, P_i)$ of $\text{NE}(\mathcal{A})$, we have $\text{flow}(\text{NE}(\mathcal{A})) = \text{flow}(B) \cup \bigcup_{i \in I} L(\text{flow}(\vec{b}_i), \text{flow}(P_i))$, where $\text{flow}(X) = \{\text{flow}(\vec{p}) \mid \vec{p} \in X\}$.*

As in the case of locally-optimal strategy profiles, we establish that the period vectors of $\text{flow}(\text{NE}(\mathcal{A}))$ for series-parallel networks have constant cost along all paths.

► **Lemma 12.** *For a series-parallel network \mathcal{A} , and any period vector \vec{q} of $\text{flow}(\text{NE}(\mathcal{A}))$, there exists κ , such that for all $\pi \in \text{Paths}_{\mathcal{A}}$, $\sum_{e \in \pi} \text{wgt}(e) \cdot q_e = \kappa$.*

Proof. We use structural induction on the series-parallel network. The base case is when the network is a single edge, which trivially satisfies the property.

Consider a network $\mathcal{A} = \langle V, E, \text{orig}, \text{dest}, \text{wgt}, s, t \rangle$ with $\mathcal{A} = \mathcal{A}_1; \mathcal{A}_2$ with $\mathcal{A}_1 = \langle V_1, E_1, \text{orig}_1, \text{dest}_1, \text{wgt}_1, s_1, t_1 \rangle$ and $\mathcal{A}_2 = \langle V_2, E_2, \text{orig}_2, \text{dest}_2, \text{wgt}_2, s_2, t_2 \rangle$.

Let \vec{q} be a period vector of $\text{flow}(\text{NE}(\mathcal{A}))$. By Lemma 11, there exists a period vector \vec{p} of $\text{NE}(\mathcal{A})$, with $\vec{q} = \text{flow}(\vec{p})$. Let us show that each $\vec{q}|_{\mathcal{A}_i}$ is a period vector of $\text{flow}(\text{NE}(\mathcal{A}_i))$. It suffices to show that $\text{prj}_{\mathcal{A}_i}(\vec{p})$ is a period vector of $\text{NE}(\mathcal{A}_i)$: since $\vec{q}|_{\mathcal{A}_i} = \text{flow}(\text{prj}_{\mathcal{A}_i}(\vec{p}))$, we can then conclude by Lemma 11.

For some base vector \vec{b} , we have that $\vec{b} + k\vec{p} \in \text{NE}(\mathcal{A})$ for all $k \geq 0$. By Lemma 16, $\text{prj}_{\mathcal{A}_i}(\vec{b} + k\vec{p}) = \text{prj}_{\mathcal{A}_i}(\vec{b}) + k \cdot \text{prj}_{\mathcal{A}_i}(\vec{p}) \in \text{NE}(\mathcal{A}_i)$, thus $\text{prj}_{\mathcal{A}_i}(\vec{p})$ is indeed a period vector.

Let us call $\vec{q}^i = \vec{q}|_{\mathcal{A}_i}$. By the induction hypothesis, there exist constants κ_1, κ_2 such that

$$\forall \pi \in \text{Paths}_{\mathcal{A}_i}, \sum_{e \in \pi} \text{wgt}(e) \cdot q_e^i = \kappa_i.$$

It follows that

$$\forall \pi \in \text{Paths}_{\mathcal{A}}, \sum_{e \in \pi} \text{wgt}(e) \cdot q_e = \sum_{e \in \pi_1} \text{wgt}(e) \cdot q_e^1 + \sum_{e \in \pi_2} \text{wgt}(e) \cdot q_e^2 = \kappa_1 + \kappa_2,$$

where $\pi_1 \pi_2$ denotes the decomposition of π such that $\pi_i \in \text{Paths}_{\mathcal{A}_i}$.

Consider now the case $\mathcal{A} = \mathcal{A}_1 \parallel \mathcal{A}_2$. Let \vec{q} be a period vector of $\text{flow}(\text{NE}(\mathcal{A}))$. By Lemma 11, there exists a period vector \vec{p} of $\text{NE}(\mathcal{A})$, with $\vec{q} = \text{flow}(\vec{p})$. Let us show that each $\vec{q}|_{\mathcal{A}_i}$ is a period vector of $\text{flow}(\text{NE}(\mathcal{A}_i))$. The argument is identical to the first case, using Lemma 17 in place of Lemma 16. It suffices to show that $\vec{p}|_{\mathcal{A}_i}$ is a period vector of $\text{NE}(\mathcal{A}_i)$ since we can then conclude by Lemma 11. For some base vector \vec{b} , we have that $\vec{b} + k\vec{p} \in \text{NE}(\mathcal{A})$ for all $k \geq 0$. By Lemma 17, $(\vec{b} + k\vec{p})|_{\mathcal{A}_i} \in \text{NE}(\mathcal{A}_i)$, and since $(\vec{b} + k\vec{p}) = \vec{b}|_{\mathcal{A}_i} + k\vec{p}|_{\mathcal{A}_i}$, $\vec{p}|_{\mathcal{A}_i}$ is indeed a period vector.

By the induction hypothesis, there exist κ_1, κ_2 such that

$$\forall \pi \in \text{Paths}_{\mathcal{A}_i}, \sum_{e \in \pi} \text{wgt}(e) \cdot q_e = \kappa_i.$$

We need to prove that $\kappa_1 = \kappa_2$. Assume otherwise; for instance, $\kappa_1 < \kappa_2$. Then, for paths $\pi_1, \pi_2 \in \text{Paths}_{\mathcal{A}}$, we have $\text{cost}_{\pi_1}(\vec{q}) < \text{cost}_{\pi_2}(\vec{q})$. Consider a period vector \vec{p} of $\text{NE}(\mathcal{A})$ with $\text{flow}(\vec{p}) = \vec{q}$. Observe that \vec{p} and its multiples are also a Nash equilibria (by Lemma 9). Then, there must exist $\pi_2 \in \text{Paths}_{\mathcal{A}_2}$ such that $p_{\pi_2} > 0$; if not, paths in $\text{Paths}_{\mathcal{A}_2}$ would become profitable deviations for $k\vec{p}$ for large k . Take such a $\pi_2 \in \text{Paths}_{\mathcal{A}_2}$. By (6), for all $\pi_1 \in \text{Paths}_{\mathcal{A}_1}$, and all $k \geq 0$, we have

$$\sum_{e \in \pi_2 \setminus \pi_1} \text{wgt}(e)kq_e \leq \sum_{e \in \pi_1 \setminus \pi_2} \text{wgt}(e)(kq_e + 1).$$

Since π_1 and π_2 are disjoint, we get

$$\sum_{e \in \pi_2} \text{wgt}(e)kq_e \leq \sum_{e \in \pi_1} \text{wgt}(e)(kq_e + 1),$$

hence $k(\kappa_2 - \kappa_1) \leq \sum_{e \in \pi_1} \text{wgt}(e)$, which is a contradiction for large k . This concludes the proof. \blacktriangleleft

It follows from Lemma 5 that all period vectors of $\text{flow}(\text{NE}(\mathcal{A}))$ are multiples of $\vec{\delta}$, the characteristic vector of \mathcal{A} .

► **Corollary 13.** *For a series-parallel network \mathcal{A} , the set $\text{flow}(\text{NE}(\mathcal{A}))$ admits a semilinear representation in the form $B \cup \bigcup_{i \in I} L(\vec{b}^i, m_i \vec{\delta})$, for a finite set B , and finitely-many base vectors \vec{b}^i and natural numbers m_i .*

We do not know whether the m_i in the above corollary can be different from 1. We did not encounter such a case in our experiments, and we conjecture that $\vec{\delta}$ is the only period vector of $\text{flow}(\text{NE}(\mathcal{A}))$.

An immediate consequence of Corollaries 7 and 13 is that the prices of anarchy and stability converge to 1 for a fixed series-parallel network, when the number of players goes to infinity. This is intuitively due to the fact that both $\text{LocOpt}(\mathcal{A})$ and $\text{flow}(\text{NE}(\mathcal{A}))$ have the same direction $\vec{\delta}$. This result already appeared recently in [37]; our setting provides an alternative proof.

► **Theorem 14.** *For series-parallel networks \mathcal{A} , $\lim_{n \rightarrow \infty} \text{PoA}(\langle \mathcal{A}, n \rangle) = \lim_{n \rightarrow \infty} \text{PoS}(\langle \mathcal{A}, n \rangle) = 1$.*

Proof. Consider $\text{worst}_n = \max_{\vec{p} \in \text{NE}_n(\mathcal{A})} \text{soccost}(\vec{p})$, $\text{best}_n = \min_{\vec{p} \in \text{NE}_n(\mathcal{A})} \text{soccost}(\vec{p})$, and $\text{opt}_n = \min_{\vec{p} \in \mathcal{F}_n(\mathcal{A})} \text{soccost}(\vec{p})$. We show that $\lim_{n \rightarrow \infty} \text{worst}_n / \text{opt}_n = \lim_{n \rightarrow \infty} \text{best}_n / \text{opt}_n = 1$.

As we already argued, the social cost only depends on the flow, so worst_n and best_n can be computed by maximizing or minimizing the social cost among $\text{flow}(\text{NE}(\mathcal{A}))$ restricted to n players. The optimum can be computed by minimizing over $\text{LocOpt}_n(\mathcal{A})$ since the global optimum is also locally optimal.

32:12 Semilinear Representations for Series-Parallel Atomic Congestion Games

Let $\text{flow}(\text{NE}(\mathcal{A})) = B \cup \bigcup_{i \in I} L(\vec{b}^i, m_i \vec{\delta})$, and $\text{LocOpt}(\mathcal{A}) = B' \cup \bigcup_{i \in I'} L(\vec{b}'^i, \vec{\delta})$. Consider $n > \max_{\vec{b} \in B \cup B'} \|\vec{b}\|$, so that all Nash equilibria with n players belong to some $L(\vec{b}^i, m_i \vec{\delta})$, and similarly, all local optima with n players are in some $L(\vec{b}'^i, \vec{\delta})$. Note that if a strategy profile \vec{p} belong to $L(\vec{b}^i, m_i \vec{\delta})$, there exists $k \in \mathbb{N}$ with $\vec{p} = \vec{b}^i + k m_i \vec{\delta}$, which implies that $\|\vec{p}\| - \|\vec{b}^i\| \equiv 0 \pmod{m_i \|\vec{\delta}\|}$, and $k = \frac{\|\vec{p}\| - \|\vec{b}^i\|}{m_i \|\vec{\delta}\|}$. We have that

$$\text{worst}_n = \max \left\{ \text{soccost} \left(\vec{b}^i + \vec{\delta} \cdot \frac{n - \|\vec{b}^i\|}{\|\vec{\delta}\|} \right) \mid i \in I, n - \|\vec{b}^i\| \equiv 0 \pmod{m_i \|\vec{\delta}\|} \right\}.$$

Similarly,

$$\text{opt}_n = \min \left\{ \text{soccost} \left(\vec{b}'^i + \vec{\delta} \cdot \frac{n - \|\vec{b}'^i\|}{\|\vec{\delta}\|} \right) \mid i \in I', n - \|\vec{b}'^i\| \equiv 0 \pmod{\|\vec{\delta}\|} \right\}.$$

Consider $(i_0, i'_0) \in I \times I'$ such that worst_n is maximized for $i_0 \in I$, and opt_n is minimized for $i'_0 \in I'$ for infinitely many n . Let $(\alpha_k)_{k \geq 0}$ denote the increasing sequence of indices n such that this is the case. We are going to show that the limit of the sequence $(\text{worst}_{\alpha_k} / \text{opt}_{\alpha_k})_{k \in \mathbb{N}}$ is 1 (independently of i_0 and i'_0), which yields the result.

We have

$$\begin{aligned} \text{worst}_{\alpha_k} &= \text{soccost} \left(\vec{b}_{i_0} + \frac{\alpha_k - \|\vec{b}^i\|}{\|\vec{\delta}\|} \vec{\delta} \right) \\ &= \sum_{e \in E} \text{wgt}(e) \left(\frac{\delta_e}{\|\vec{\delta}\|} \right)^2 \alpha_k^2 + 2 \sum_{e \in E} \text{wgt}(e) \frac{\delta_e (b_e^i - \delta_e \|\vec{b}^i\| / \|\vec{\delta}\|)}{\|\vec{\delta}\|} \alpha_k \\ &\quad + \sum_{e \in E} \text{wgt}(e) \left(b_e^i - \delta_e \cdot \frac{\|\vec{b}^i\|}{\|\vec{\delta}\|} \right)^2. \end{aligned}$$

We have $\text{worst}_{\alpha_k} = A \alpha_k^2 + o(\alpha_k)$ where $A = \sum_{e \in E} \text{wgt}(e) \left(\frac{\delta_e}{\|\vec{\delta}\|} \right)^2$.

A similar expression can be obtained for opt_{α_k} since it has the same form as worst_{α_k} . In particular, the first term is again A . We can obtain that $\text{opt}_{\alpha_k} \geq A \alpha_k^2 - \frac{2E \|\vec{b}'^i\|}{\|\vec{\delta}\|}$. Hence,

$$1 \leq \frac{\text{worst}_{\alpha_k}}{\text{opt}_{\alpha_k}} \leq \frac{A \alpha_k^2 + o(\alpha_k)}{A \alpha_k^2 + o(1)} = 1 + o\left(\frac{1}{\alpha_k}\right).$$

It follows that $\lim_{n \rightarrow \infty} \text{PoA}(\langle \mathcal{A}, n \rangle) = 1$.

This also implies that $\lim_{n \rightarrow \infty} \text{PoS}(\langle \mathcal{A}, n \rangle) = 1$. ◀

5 Computation of PoA and PoS

Semilinear Representations

Let us show how the semilinear representation for $\text{LocOpt}(\mathcal{A})$ can be computed. First, the characteristic vector $\vec{\delta}$ can be computed by solving the homogeneous equation system $\mathcal{E}(0)$ using a symbolic solver (so as to obtain a rational solution), and multiplying the unique solution by the gcd of its coefficients. Next, one can incrementally construct the semilinear representation using an integer-arithmetic solver to find the linear sets and the finite set B . This can be done with quantifier-free formulas only. Although integer linear programming is already NP-hard [6, 35], available solvers are efficient for small instances.

At a given iteration, assume that the current subset of $\text{LocOpt}(\mathcal{A})$ is $B \cup \bigcup_{i \in I} L(\vec{b}^i, \vec{\delta})$. We write a linear quantifier-free formula $\phi(\vec{q})$ with free variables a flow \vec{q} which requires that \vec{q} is locally optimal (by Lemma 1), and that \vec{q} is not included in the current set. We already saw that the former is a Presburger formula. The latter constraints can be written as $\bigwedge_{\vec{b} \in B} \vec{q} \neq \vec{b} \wedge \bigwedge_{i \in I} \vec{q} \notin L(\vec{b}^i, \vec{\delta})$. Here, $\vec{q} \notin L(\vec{b}^i, \vec{\delta}) \equiv \neg(\exists k. \vec{q} = \vec{b}^i + k\vec{\delta})$ but the existentially quantified k can be determined from the number of players of $\vec{q}, \vec{b}^i, \vec{\delta}$, so this can be simplified as follows:

$$(\exists k. \vec{q} = \vec{b}^i + k\vec{\delta}) \equiv (\|\vec{q}\| - \|\vec{b}^i\|) \equiv 0 \pmod{\|\vec{\delta}\|} \wedge \bigwedge_{e \in E} q_e = b_e^i + (\|\vec{q}\| - \|\vec{b}^i\|)\delta_e / \|\vec{\delta}\|,$$

where $\|\vec{\delta}\|$ and $\|\vec{b}^i\|$ are fixed numbers.

If $\phi(\vec{q})$ is not satisfiable, then the current representation is complete. Otherwise, a model satisfying the above formula gives a new vector \vec{q} that is locally optimal. To determine whether \vec{q} should belong to B , or whether $L(\vec{q}, \vec{\delta})$ is to be added to our set, we simply check if $\vec{q} + \vec{\delta}$ satisfies the condition of Lemma 1: if this is the case, then we keep the linear set, otherwise we add \vec{q} to B . In fact, for all $k \geq 1$, $\vec{q} + k\vec{\delta}$ has the same set of paths π satisfying $\forall e \in \pi, (q_e + k\delta_e) > 0$, so this check is sufficient. Since the set admits a finite semilinear representation, this procedure terminates.

Let us explain the computation of $\text{flow}(\text{NE}(\mathcal{A}))$. Let $\text{NE}(\vec{p})$ denote the linear constraints of (6). Assume that we currently have a subset of $\text{flow}(\text{NE}_{\mathcal{A}})$ in the form $B \cup \bigcup_{i \in I} L(\vec{b}^i, m_i \vec{\delta})$. The following formula $\psi(\vec{p}, \vec{q})$ with free variables \vec{p}, \vec{q} is satisfiable iff some Nash equilibrium \vec{p} (with $\text{flow}(\vec{p}) = \vec{q}$) is not in the set:

$$\psi(\vec{p}, \vec{q}) = \text{NE}(\vec{p}) \wedge \text{flow}(\vec{p}) = \vec{q} \wedge \bigwedge_{i \in I} \vec{q} \notin L(\vec{b}^i, m_i \vec{\delta}) \wedge \bigwedge_{\vec{b} \in B} \vec{q} \neq \vec{b}.$$

Assume that this is satisfiable, and let \vec{p}, \vec{q} be a model. We need to check whether \vec{q} is to be added to B , or whether $L(\vec{q}, m_i \vec{\delta})$, for some m_i , is to be included. We use the following properties of the period vectors of $\text{NE}(\mathcal{A})$. Let us first define $S_{\vec{\delta}} = \{\pi \in \text{Paths}_{\mathcal{A}} \mid \sum_{e \in \pi} \text{wgt}(e)\delta_e \leq \min_{\pi' \in \text{Paths}_{\mathcal{A}}} \sum_{e \in \pi'} \text{wgt}(e)\delta_e\}$. Intuitively, $S_{\vec{\delta}}$ is the set of paths in $\text{Paths}_{\mathcal{A}}$ with minimum cost in the profile $\vec{\delta}$.

► **Lemma 15.** *Let $\vec{p}, \vec{p}' \in \text{NE}(\mathcal{A})$ such that $\text{flow}(\vec{p}') = m\vec{\delta}$. We have $L(\vec{p}, \vec{p}') \subseteq \text{NE}(\mathcal{A})$ iff for all $\pi \in \text{Paths}_{\mathcal{A}}$, $\pi \notin S_{\vec{p}} \Rightarrow p'_\pi = 0$.*

Proof. Assume $L(\vec{p}, \vec{p}') \subseteq \text{NE}(\mathcal{A})$. Then, for all π such that $p_\pi > 0$ or $p'_\pi > 0$, we have

$$\begin{aligned} \forall \pi' \in \text{Paths}(\mathcal{A}), \sum_{e \in \pi' \setminus \pi} (q_e + q'_e) \text{wgt}(e) &\leq \sum_{e \in \pi' \setminus \pi} (q_e + q'_e + 1) \text{wgt}(e). \\ &\Leftrightarrow \\ \forall \pi' \in \text{Paths}(\mathcal{A}), \sum_{e \in \pi' \setminus \pi} q_e \text{wgt}(e) &\leq \sum_{e \in \pi' \setminus \pi} (q_e + 1) \text{wgt}(e), \end{aligned}$$

where the equivalence holds by Lemma 4. This already holds for $\pi \in \text{Paths}(\mathcal{A})$ such that $p_\pi > 0$. Thus any path π such that $p'_\pi > 0$ must satisfy $\pi \in S_{\vec{p}}$, which is equivalent to the above. So any period vector \vec{p}' for the base \vec{p} satisfies $\bigwedge_{\pi \notin S_{\vec{p}}} p'_\pi = 0$. Conversely, for any such vector \vec{p}' , $\vec{p} + k\vec{p}'$ is a Nash equilibrium for all $k \geq 0$. ◀

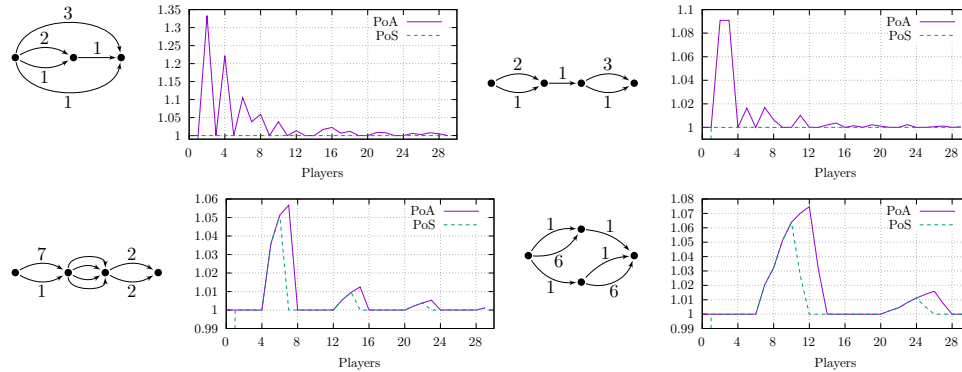
Given the pair \vec{p}, \vec{q} , we write another query to guess m, \vec{p}' such that $\text{NE}(\vec{p}') \wedge \text{flow}(\vec{p}') = m\vec{\delta} \wedge \bigwedge_{\pi \notin S_{\vec{p}}} p'_\pi = 0$. If this is satisfiable, then we query again the solver to find the smallest such m , and keep the set $L(\vec{q}, m\vec{\delta})$. Otherwise \vec{q} is added to B .

Price of Anarchy and Stability

Given a semilinear representation $B \cup \bigcup_{i \in I} L(\vec{b}^i, \vec{\delta})$ of $\text{LocOpt}(\mathcal{A})$, observe that there is only a finite number of vectors with a given n number of players. So in order to compute the *global* social optimum with n players, we iterate over all vectors with n players in this representation, and keep the minimal social cost. We first iterate over $\{\vec{b} \in B \mid \|\vec{b}\| = n\}$ and consider the vector with the least social cost among those (this set can be empty). Second, for each linear set $L(\vec{b}^i, \vec{\delta})$ such that $\frac{n - \|\vec{b}^i\|}{\|\vec{\delta}\|}$ is an integer, we compute the vector $\vec{b}^i + \frac{\|\vec{b}^i\| - n}{\|\vec{\delta}\|} \vec{\delta}$, compute its social cost, and keep it if it is less than the previous value.

We compute the social costs of the best and the worst Nash equilibria similarly on the semilinear representation of $\text{flow}(\text{NE}_{\mathcal{A}})$.

Figure 5 shows the plots of the PoA and PoS computed for four examples. We used the Python `sympy` package for solving linear equations, and the Z3 SMT solver for integer arithmetic queries. The characteristic vector was always easy to compute since it is computable in polynomial time in the size of the linear equation system. However, the number of base vectors can be large and depends on the weights used in the network. Our prototype is currently limited in scalability but it allows us to explore small yet non-trivial networks.



■ **Figure 5** Plots for PoA and PoS on four series-parallel networks. Unlabeled edges have weight 1.

6 More on Related Works

Inefficiency in congestion networks were mentioned in [30], and equilibria were first mathematically studied in [36]. Network congestion games are mainly studied in two settings which have different mathematical properties: the *nonatomic* case, where one considers a large number of players, each of which contributes an infinitesimal amount to congestion; and the *atomic* case, as we do, where there is a discrete number of players involved.

Existence and properties of Nash equilibria in the nonatomic case were established in [4]. The price of anarchy of the nonatomic case was studied in [32] which gives a tight bound of $\frac{4}{3}$ for networks with affine cost functions. [28, Chapter 18] presents a survey of these results. It is shown in [17] that Nash equilibria in atomic network congestion games can be found in polynomial time, by reduction to maximum flow in the symmetric case, that is, when all players share the same source and target vertices. The problem in the non-symmetric case is however complete for the class PLS, Polynomially Local Search, and is believed to be intractable [26, 29]. In extension-parallel networks, best-response procedures are shown to converge in linear time in [19]; but this does not extend to general series-parallel graphs.

The complexity of finding *extremal* Nash equilibria, that is, the best and the worst ones is however higher. Finding such equilibria is NP-hard for the makespan objective with varying sizes [20]. For the makespan objective and unit sizes, finding a Nash equilibrium minimizing the makespan in series-parallel networks with linear cost functions is strongly NP-hard when the number of players is part of the input, while a polynomial-time greedy algorithm allows one to find a worst Nash equilibrium [21]. For the total cost objective (as in this paper), NP-hardness holds for both best and worst equilibria for three and two players respectively [33]. Note that in our work, we are interested in computing such equilibria (or their costs) for arbitrarily large numbers of players. In the more general case of network congestion games with *dynamic* strategies which allow players to choose each move according to the current state of the game, doubly exponential-time algorithms were given for computing such equilibria in [5] when the number of players is encoded in binary.

It is possible to efficiently compute the social optimum in atomic congestion games by transforming the cost function, and reducing the problem to the computation of Nash equilibria [15]. In our case, this transformation would yield affine cost functions. This direction could be exploited to compute the costs of socially optimal profiles using semilinear representations for Nash equilibria, if these could be extended to affine costs. The behaviors of PoA for large numbers of players have been studied before. [18] considers congestion games with large numbers of players, and shows that the PoA of atomic congestion games converges to the PoA of the nonatomic game; the result holds for games with affine cost functions and positive coefficients (as in our case). A consequence is that, although the PoA for the atomic case is often larger than that in the nonatomic case, this difference vanishes in the limit, and thus the upper bound of $\frac{4}{3}$ holds for the atomic case in the limit. In [13], the limit of atomic congestion games is considered in a setting where either players participate in the game with given probabilities that tend to 0, or they have weights that tend to 0; in both cases, the limit of the PoA for mixed equilibria is equal to that in a corresponding nonatomic game. Asymptotic PoA bounds (of ≈ 1.35188) are provided in [14] for symmetric atomic congestion games with affine cost functions, by restricting to specific strategies called k -uniform. In the nonatomic case, the works [11, 10] establish that the limit of the PoA is 1. Paper [12] considers nonatomic congestion games as a function of the demand and studies continuous derivability properties of the PoA function.

7 Conclusion

Our results provide theoretical tools that allow us to have a better understanding of the structure of Nash equilibria and social optima in atomic congestion games over series-parallel networks. An immediate question is how the semilinear representations would change if we allow affine cost functions rather than linear ones. However, extending further our approach to nonlinear cost functions would not be immediate since these sets would no longer be definable in Presburger arithmetic.

Although the characteristic vector is easy to compute, the exact computation of the whole semilinear representations is costly and currently does not scale to large networks. One might investigate how this computation can be rendered more efficient in practice. Another possible direction is to explore more efficient approximation algorithms using just the characteristic vector, and perhaps a subset of the base vectors.

References

- 1 Eitan Altman, Anurag Kumar, and Yezekael Hayel. A potential game approach for uplink resource allocation in a multichannel wireless access network. In *Proceedings of the Fourth International ICST Conference on Performance Evaluation Methodologies and Tools*, pages 1–9, 2009.
- 2 E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 295–304, 2004. doi:10.1109/FOCS.2004.68.
- 3 Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, STOC'05*, pages 57–66, New York, NY, USA, 2005. Association for Computing Machinery. doi:10.1145/1060590.1060599.
- 4 Martin J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. RAND Corporation, Santa Monica, CA, 1955.
- 5 Nathalie Bertrand, Nicolas Markey, Suman Sadhukhan, and Ocan Sankur. Dynamic network congestion games. In *40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2020)*, Goa, India, December 2020. URL: <https://hal.archives-ouvertes.fr/hal-02980833>.
- 6 I. Borosh and L. B. Treybig. Bounds on positive integral solutions of linear diophantine equations. *Proceedings of the American Mathematical Society*, 55(2):299–304, 1976. URL: <http://www.jstor.org/stable/2041711>.
- 7 Ioannis Caragiannis, Michele Flammini, Christos Kaklamanis, Panagiotis Kanellopoulos, and Luca Moscardelli. Tight bounds for selfish and greedy load balancing. *Algorithmica*, 61(3):606–637, November 2011. doi:10.1007/s00453-010-9427-8.
- 8 George Christodoulou and Elias Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games,. In *Proceedings of the 13th Annual European Conference on Algorithms, ESA'05*, pages 59–70, Berlin, Heidelberg, 2005. Springer-Verlag. doi:10.1007/11561071_8.
- 9 George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing, STOC'05*, pages 67–73, New York, NY, USA, 2005. Association for Computing Machinery. doi:10.1145/1060590.1060600.
- 10 Riccardo Colini-Baldeschi, Roberto Cominetti, Panayotis Mertikopoulos, and Marco Scarsini. When is selfish routing bad? the price of anarchy in light and heavy traffic. *Oper. Res.*, 68(2):411–434, 2020. doi:10.1287/opre.2019.1894.
- 11 Riccardo Colini-Baldeschi, Roberto Cominetti, and Marco Scarsini. Price of anarchy for highly congested routing games in parallel networks. *Theory Comput. Syst.*, 63(1):90–113, 2019. doi:10.1007/s00224-017-9834-1.
- 12 Roberto Cominetti, Valerio Dose, and Marco Scarsini. The price of anarchy in routing games as a function of the demand. *Mathematical Programming*, 2021. To appear. doi:10.1007/s10107-021-01701-7.
- 13 Roberto Cominetti, Marco Scarsini, Marc Schröder, and NE Stier-Moses. Convergence of large atomic congestion games. *arXiv preprint*, 2020. arXiv:2001.02797.
- 14 Jasper de Jong, Walter Kern, Berend Steenhuisen, and Marc Uetz. The asymptotic price of anarchy for k-uniform congestion games. In *International Workshop on Approximation and Online Algorithms*, pages 317–328. Springer, 2017.
- 15 Alberto Del Pia, Michael Ferris, and Carla Michini. Totally unimodular congestion games. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 577–588. SIAM, 2017.
- 16 Samuel Eilenberg and Marcel Paul Schützenberger. Rational sets in commutative monoids. *Journal of Algebra*, 13:173–191, 1969.

- 17 Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure nash equilibria. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing, STOC '04*, pages 604–612, New York, NY, USA, 2004. Association for Computing Machinery. doi:10.1145/1007352.1007445.
- 18 Michal Feldman, Nicole Immorlica, Brendan Lucier, Tim Roughgarden, and Vasilis Syrgkanis. The price of anarchy in large games. In *48th ACM Symposium on Theory of Computing (STOC) 2016*, June 2016. URL: <https://www.microsoft.com/en-us/research/publication/price-anarchy-large-games/>.
- 19 Dimitris Fotakis. Congestion games with linearly independent paths: Convergence time and price of anarchy. In *Proceedings of the 1st International Symposium on Algorithmic Game Theory, SAGT '08*, pages 33–45, Berlin, Heidelberg, 2008. Springer-Verlag. doi:10.1007/978-3-540-79309-0_5.
- 20 Dimitris Fotakis, Spyros Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul Spirakis. The structure and complexity of nash equilibria for a selfish routing game. *Theoretical Computer Science*, 410(36):3305–3326, 2009. Graphs, Games and Computation: Dedicated to Professor Burkhard Monien on the Occasion of his 65th Birthday. doi:10.1016/j.tcs.2008.01.004.
- 21 Elisabeth Gassner, Johannes Hatzl, Sven O. Krumke, Heike Sperber, and Gerhard J. Woeginger. How hard is it to find extreme nash equilibria in network congestion games? *Theoretical Computer Science*, 410(47):4989–4999, 2009. doi:10.1016/j.tcs.2009.07.046.
- 22 Seymour Ginsburg and Edwin H. Spanier. Bounded algol-like languages. *Transactions of the American Mathematical Society*, 113(2):333–368, 1964. URL: <http://www.jstor.org/stable/1994067>.
- 23 Christoph Haase. A survival guide to presburger arithmetic. *ACM SIGLOG News*, 5(3):67–82, July 2018. doi:10.1145/3242953.3242964.
- 24 Bainian Halo and Carla Michini. The price of anarchy in series-parallel network congestion games. Technical report, University of Wisconsin-Madison, 2021.
- 25 Ryuichi Ito. Every semilinear set is a finite union of disjoint linear sets. *J. Comput. Syst. Sci.*, 3(2):221–231, May 1969. doi:10.1016/S0022-0000(69)80014-0.
- 26 David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988. doi:10.1016/0022-0000(88)90046-3.
- 27 Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009. doi:10.1016/j.cosrev.2009.04.003.
- 28 Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- 29 Christos H. Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43(3):425–440, 1991. doi:10.1016/0022-0000(91)90023-X.
- 30 A. C. Pigou. *Economics of welfare*. McMillan, 1920.
- 31 Robert W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3:53–59, 1973.
- 32 Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.
- 33 Heike Sperber. How to find nash equilibria with extreme total latency in network congestion games? In *2009 International Conference on Game Theory for Networks*, pages 158–163, 2009. doi:10.1109/GAMENETS.2009.5137397.
- 34 Jacobo Valdes, Robert E. Tarjan, and Eugene L. Lawler. The recognition of series parallel digraphs. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing, STOC '79*, pages 1–12, New York, NY, USA, 1979. Association for Computing Machinery. doi:10.1145/800135.804393.
- 35 Joachim von zur Gathen and Malte Sieveking. A bound on solutions of linear integer equalities and inequalities. *Proceedings of the American Mathematical Society*, 72(1):155–158, 1978.

- 36 John Glen Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the institution of civil engineers*, 1(3):325–362, 1952.
- 37 Zijun Wu, Rolf H Moehring, Chunying Ren, and Dachuan Xu. A convergence analysis of the price of anarchy in atomic congestion games. *Mathematical Programming*, 2022. To appear. doi:10.1007/s10107-022-01853-0.

A Proofs of Section 3

A.1 Proofs of Section 3.1 and 3.2

► **Lemma 1.** *In a network congestion game $\langle \mathcal{A}, n \rangle$, a flow \vec{q} is locally-optimal if, and only if, for all $\pi, \pi' \in \text{Paths}_{\mathcal{A}}$ such that $\forall e \in \pi, q_e > 0$,*

$$\sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot (2q_e - 1) \leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (2q_e + 1). \quad (3)$$

Proof. Observe that flow \vec{q} is locally optimal iff for all paths π, π' such that $\forall e \in \pi, q_e > 0$, the vector \vec{q}' defined by

$$q'_e = \begin{cases} q_e - 1 & \text{if } e \in \pi \setminus \pi', \\ q_e + 1 & \text{if } e \in \pi' \setminus \pi, \\ q_e & \text{otherwise,} \end{cases}$$

satisfies $\text{soccost}(\vec{q}) \leq \text{soccost}(\vec{q}')$. Given such paths π, π' , let us thus write this inequality as

$$\begin{aligned} \sum_{e \in E} \text{wgt}(e) \cdot q_e^2 &\leq \sum_{e \in E} \text{wgt}(e) \cdot q_e'^2 \\ &= \sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot (q_e - 1)^2 + \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (q_e + 1)^2 + \sum_{e \in \pi \cap \pi'} \text{wgt}(e) \cdot q_e^2. \end{aligned}$$

This is equivalent to

$$\sum_{e \in \pi \setminus \pi' \cup \pi' \setminus \pi} \text{wgt}(e) \cdot q_e^2 \leq \sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot (q_e^2 - 2q_e + 1) + \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (q_e^2 + 2q_e + 1),$$

hence to

$$\sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot (2q_e - 1) \leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (2q_e + 1). \quad \blacktriangleleft$$

► **Lemma 4.** *In a series-parallel network \mathcal{A} , for all period vectors $\vec{d} \in \mathbb{N}^E$ of a semilinear representation of $\text{LocOpt}_{\geq n_0}(\mathcal{A})$, there exists $\kappa \geq 0$ such that for all $\pi \in \text{Paths}_{\mathcal{A}}$, we have $\sum_{e \in \pi} \text{wgt}(e) \cdot d_e = \kappa$.*

Proof. Consider a linear set $L(\vec{b}, \vec{d})$ in $\text{LocOpt}_{\geq n_0}(\mathcal{A})$ and two paths π_1 and π_2 .

Applying Lemma 3 to $\vec{b} \in \text{LocOpt}_{\geq n_0}(\mathcal{A})$, we have that $b_e > 0$ for all $e \in E$. For any $k \geq 0$, the flow $\vec{q} = \vec{b} + k\vec{d}$ is locally optimal and has $q_e > 0$ for all $e \in \pi_2$. By Lemma 1,

$$\sum_{e \in \pi_2 \setminus \pi_1} \text{wgt}(e) \cdot (2(b_e + kd_e) - 1) \leq \sum_{e \in \pi_1 \setminus \pi_2} \text{wgt}(e) \cdot (2(b_e + kd_e) + 1)$$

which rewrites as

$$\sum_{e \in \pi_2 \setminus \pi_1} \text{wgt}(e) \cdot (2b_e - 1) - \sum_{e \in \pi_1 \setminus \pi_2} \text{wgt}(e) \cdot 2(b_e + 1) + 2k \left(\sum_{e \in \pi_2} \text{wgt}(e) d_e - \sum_{e \in \pi_1} \text{wgt}(e) d_e \right) \leq 0.$$

Since this holds for any $k \geq 0$, we must have $\sum_{e \in \pi_2} \text{wgt}(e)d_e - \sum_{e \in \pi_1} \text{wgt}(e)d_e \leq 0$. The converse inequality can be obtained using similar arguments, hence $\sum_{e \in \pi_2} \text{wgt}(e)d_e = \sum_{e \in \pi_1} \text{wgt}(e)d_e$. ◀

B Proofs of Section 4

► **Lemma 9.** *Given a network \mathcal{A} , a strategy profile \vec{p} is a Nash equilibrium if, and only if,*

$$\forall \pi, \pi' \in \text{Paths}, p_\pi > 0 \implies \sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot q_e \leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (q_e + 1), \quad (6)$$

where $\vec{q} = \text{flow}(\vec{p})$.

Proof. The lemma is a reformulation of (2): for $\pi, \pi' \in \text{Paths}$ with $p_\pi > 0$, $\text{cost}_\pi(\vec{p}) \leq \text{cost}_{\pi'}(\vec{p})$ can be written as

$$\sum_{e \in \pi} \text{wgt}(e) \cdot q_e \leq \sum_{e \in \pi'} \text{wgt}(e) \cdot q'_e,$$

where \vec{q} and \vec{q}' are the respective flows of \vec{p} and \vec{p}' . This is equivalent to

$$\sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot q_e \leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot q'_e = \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (q_e + 1). \quad \blacktriangleleft$$

► **Lemma 10.** *The sets $\text{NE}(\mathcal{A})$ and $\text{flow}(\text{NE}(\mathcal{A}))$ are semilinear.*

Proof. We show that both sets can be expressed in Presburger arithmetic. This follows from Lemma 9 since the following formula with free variables $\{q_e \mid e \in E\} \cup \{p_\pi \mid \pi \in \text{Paths}(\mathcal{A})\}$ expresses (6) in Presburger arithmetic:

$$\phi = \bigwedge_{\pi, \pi' \in \text{Paths}} \left(p_\pi > 0 \implies \sum_{e \in \pi \setminus \pi'} \text{wgt}(e) \cdot q_e \leq \sum_{e \in \pi' \setminus \pi} \text{wgt}(e) \cdot (q_e + 1) \right) \wedge \bigwedge_{e \in E} \left(q_e = \sum_{\pi \in \text{Paths}: e \in \pi} p_\pi \right).$$

Here we use the fact that Paths is finite so that the above is a well-defined formula. Now, existentially quantifying $\{q_e\}_{e \in E}$ in ϕ yields a formula describing $\text{NE}(\mathcal{A})$. Existentially quantifying $\{p_\pi\}_{\pi \in \text{Paths}(\mathcal{A})}$ in ϕ yields $\text{flow}(\text{NE}(\mathcal{A}))$. ◀

Our next results prove that the “projections” of the Nash equilibria of series-parallel networks onto their constituent subnets still are Nash equilibria. We first formally define those projections.

For a network $\mathcal{A} = \mathcal{A}_1; \mathcal{A}_2$ and $\vec{p} \in \mathfrak{S}(\mathcal{A})$, for $i \in \{1, 2\}$, let us define $\text{prj}_{\mathcal{A}_i}(\vec{p}) \in \mathfrak{S}(\mathcal{A}_i)$ where for each $\pi_i \in \text{Paths}_{\mathcal{A}_i}$, $\text{prj}_{\mathcal{A}_i}(\vec{p})_{\pi_i} = \sum_{\pi_{3-i} \in \text{Paths}_{\mathcal{A}_{3-i}}} p_{\pi_1 \pi_2}$. Thus, $\text{prj}_{\mathcal{A}_i}(\vec{p})_{\pi_i}$ is the number of players that cross the path π_i in the profile \vec{p} .

For a network $\mathcal{A} = \mathcal{A}_1 \parallel \mathcal{A}_2$, and a vector $\vec{p} \in \mathfrak{S}(\mathcal{A})$, for $i \in \{1, 2\}$, let us denote $\vec{p}|_{\mathcal{A}_i} \in \mathfrak{S}(\mathcal{A}_i)$ obtained by restricting \vec{p} to $\text{Paths}_{\mathcal{A}_i}$. Similarly, for a vector $\vec{q} \in \mathcal{F}(\mathcal{A})$, let $\vec{q}|_{\mathcal{A}_i}$ the restriction of \vec{q} to the edges of \mathcal{A}_i .

► **Lemma 16.** *Consider a network $\mathcal{A} = \mathcal{A}_1; \mathcal{A}_2$. Then, for all $\vec{p} \in \mathfrak{S}(\mathcal{A})$, we have*

$$\vec{p} \in \text{NE}(\mathcal{A}) \Leftrightarrow \forall i \in \{1, 2\}, \text{prj}_{\mathcal{A}_i}(\vec{p}) \in \text{NE}(\mathcal{A}_i).$$

32:20 Semilinear Representations for Series-Parallel Atomic Congestion Games

Proof. Consider $\vec{p} \in \mathfrak{S}(\mathcal{A})$. Observe that all $\pi \in \text{Paths}_{\mathcal{A}}$ can be written as $\pi = \pi_1\pi_2$ where $\pi_i \in \text{Paths}_{\mathcal{A}_i}$, and that $\text{cost}_{\pi}(\vec{p}) = \text{cost}_{\pi_1}(\text{prj}_{\mathcal{A}_1}(\vec{p})) + \text{cost}_{\pi_2}(\text{prj}_{\mathcal{A}_2}(\vec{p}))$.

Assume that $\vec{p} \in \text{NE}(\mathcal{A})$. Consider $\pi_1 \in \text{Paths}_{\mathcal{A}_1}$ such that $\text{prj}_{\mathcal{A}_1}(\vec{p})_{\pi_1} > 0$. Then, there must exist a path $\pi_2 \in \text{Paths}_{\mathcal{A}_2}$ such that $p_{\pi_1\pi_2} > 0$. By (2), we have that for all $\pi' \in \text{Paths}_{\mathcal{A}}$, we have $\text{cost}_{\pi_1\pi_2}(\vec{p}) \leq \text{cost}_{\pi'}(\vec{p}')$ where $\vec{p}' = \vec{p} - \pi_1\pi_2 + \pi'$. In particular, for all $\pi'_1 \in \text{Paths}_{\mathcal{A}_1}$, we have $\text{cost}_{\pi_1\pi_2}(\vec{p}) \leq \text{cost}_{\pi'_1\pi_2}(\vec{p}')$, i.e.,

$$\text{cost}_{\pi_1}(\text{prj}_{\mathcal{A}_1}(\vec{p})) + \text{cost}_{\pi_2}(\text{prj}_{\mathcal{A}_2}(\vec{p})) \leq \text{cost}_{\pi'_1}(\text{prj}_{\mathcal{A}_1}(\vec{p}')) + \text{cost}_{\pi_2}(\text{prj}_{\mathcal{A}_2}(\vec{p}')).$$

The second terms of both sides are equal since $\text{prj}_{\mathcal{A}_2}(\vec{p}) = \text{prj}_{\mathcal{A}_2}(\vec{p}')$. It follows that

$$\text{cost}_{\pi_1}(\text{prj}_{\mathcal{A}_1}(\vec{p})) \leq \text{cost}_{\pi'_1}(\text{prj}_{\mathcal{A}_1}(\vec{p}')) = \text{cost}_{\pi'_1}(\text{prj}_{\mathcal{A}_1}(\vec{p}) - \pi_1 + \pi'_1),$$

which means that $\text{prj}_{\mathcal{A}_1}(\vec{p}) \in \text{NE}(\mathcal{A}_1)$. The argument is symmetric for \mathcal{A}_2 .

Conversely, assume that for all $i \in \{1, 2\}$, $\text{prj}_{\mathcal{A}_i}(\vec{p}) \in \text{NE}(\mathcal{A}_i)$. Consider any path $\pi \in \text{Paths}_{\mathcal{A}}$ such that $p_{\pi} > 0$, and write it as $\pi = \pi_1\pi_2$. Then, for both $i \in \{1, 2\}$, $\text{prj}_{\mathcal{A}_i}(\vec{p})_{\pi_i} > 0$. Take any other path $\pi' = \pi'_1\pi'_2 \in \text{Paths}_{\mathcal{A}}$. Because each $\text{prj}_{\mathcal{A}_i}(\vec{p})$ is a Nash equilibrium, we have

$$\text{cost}_{\pi_i}(\text{prj}_{\mathcal{A}_i}(\vec{p})) \leq \text{cost}_{\pi'_i}(\text{prj}_{\mathcal{A}_i}(\vec{p}) - \pi_i + \pi'_i).$$

Summing up these inequations, we get

$$\begin{aligned} \text{cost}_{\pi_1\pi_2}(\vec{p}) &\leq \text{cost}_{\pi'_1}(\text{prj}_{\mathcal{A}_1}(\vec{p}) - \pi_1 + \pi'_1) + \text{cost}_{\pi'_2}(\text{prj}_{\mathcal{A}_2}(\vec{p}) - \pi_2 + \pi'_2) \\ &= \text{cost}_{\pi'_1}(\text{prj}_{\mathcal{A}_1}(\vec{p} - \pi_1\pi_2 + \pi'_1\pi'_2)) + \text{cost}_{\pi'_2}(\text{prj}_{\mathcal{A}_2}(\vec{p} - \pi_1\pi_2 + \pi'_1\pi'_2)) \\ &= \text{cost}_{\pi'}(\vec{p} - \pi + \pi'). \end{aligned}$$

Hence π' is not a profitable deviation, whichever $\pi' \in \text{Paths}_{\mathcal{A}}$. ◀

► **Lemma 17.** Consider a network $\mathcal{A} = \mathcal{A}_1 \parallel \mathcal{A}_2$. Then, for all $\vec{p} \in \text{NE}(\mathcal{A})$, we have that for all $i \in \{1, 2\}$, we have $\vec{p}|_{\mathcal{A}_i} \in \text{NE}(\mathcal{A}_i)$.

Proof. Consider $\vec{p} \in \text{NE}(\mathcal{A})$, and $i \in \{1, 2\}$. Then, for all $\pi, \pi' \in \text{Paths}_{\mathcal{A}_i}$ such that $p_{\pi} > 0$, $\text{cost}_{\pi}(\vec{p}) \leq \text{cost}_{\pi'}(\vec{p} - \pi + \pi')$. Since the only paths sharing edges with π and π' are in $\text{Paths}_{\mathcal{A}_i}$, we have that $\text{cost}_{\pi}(\vec{p}|_{\mathcal{A}_i}) \leq \text{cost}_{\pi'}(\vec{p}|_{\mathcal{A}_i} - \pi + \pi')$. Hence $\vec{p}|_{\mathcal{A}_i} \in \text{NE}(\mathcal{A}_i)$. ◀

► **Remark 18.** Notice that contrary to Lemma 16, Lemma 17 is not an equivalence: a 2-player strategy profile involving the “shortest” path of \mathcal{A}_1 with the “shortest” path of \mathcal{A}_2 need not yield a Nash equilibrium.