STACKELBERG ROUTING IN ATOMIC NETWORK GAMES

JOSÉ R. CORREA* AND NICOLÁS E. STIER-MOSES**

* *School of Business, Universidad Adolfo Ibáñez, Santiago, Chile;* `correa@uai.cl`

** *Graduate School of Business, Columbia University, New York, USA;* `stier@gsb.columbia.edu`

Abstract. We consider network games with atomic players, which indicates that some players control a positive amount of flow. Instead of studying Nash equilibria as previous work has done, we consider that players with considerable market power will make decisions before the others because they can predict the decisions of players without market power. This description fits the framework of Stackelberg games, where those with market power are leaders and the rest are price-taking followers. As Stackelberg equilibria are difficult to characterize, we prove bounds on the inefficiency of the solutions that arise when the leader uses a heuristic that approximate its optimal strategy.

Keywords. Atomic Network Games, Atomic Congestion Games, Stackelberg Games, Price of Anarchy, Selfish Routing.

1. Introduction

This article studies the efficiency loss due to selfish behavior in network games with atomic players that can split their flow among different paths. These games can be used to model situations in which players, who control an arbitrary portion of the whole demand, have to ship flow from their origins to their destinations. Competition arises because shipment costs depend on how much flow from all the participants is sent on each link, which creates externalities. Typical applications of these games can be found in the domains of telecommunication, transportation and distribution networks.

Roughgarden and Tardos (2002), Correa, Schulz, and Stier-Moses (2005), and Cominetti, Correa, and Stier-Moses (2006) concentrated on Nash equilibria for atomic games, and proved bounds for their inefficiency with respect to a coordinated solution that achieves a minimal social cost. The problem is that, as exemplified by some instances presented by Cominetti et al., a Nash equilibrium may fail to capture what participants of the game would do if presented with that decision. This is especially the case in situations in which some players have more market power than others. Although a Nash equilibrium demands that a player controlling most of the demand plays a best response to what the rest do, this player can do better by simply predicting what the rest will do for each possible strategy and selecting the best. Indeed, the others players are price-takers, so they cannot influence the market by themselves unless they collude. This observation led us to consider a Stackelberg setting in which the players with market power are *leaders* while those without market power are *followers* (von Stackelberg 1934). In a situation of this type, Stackelberg equilibria may lead to a better representation of players' behavior.

As providing a characterization of Stackelberg equilibria is much more difficult than describing Nash equilibria, we consider a simpler market structure in which there is a single atomic player with market power, and a continuum of nonatomic players represented by the price-taking followers. The atomic player can predict how nonatomic players will react to its routing decision because nonatomic players select shortest paths as they would do under a Wardrop equilibrium (Wardrop 1952).

In Section 2 we introduce the problem and define the linear programming strategy for the leader. This strategy generalizes the LLF strategy defined by Roughgarden (2004) for parallel links networks. Section 3 considers general network topologies and prove bounds on the price of anarchy for some natural strategies that have a similar flavor than the known bounds for the price of anarchy of nonatomic network games. These bounds depend heavily on the type of congestion functions allowed in the network. Furthermore, we prove a counterintuitive phenomena appearing in general networks. Namely, if the leader plays according to the LP strategy, it may end up having larger cost than the optimal flow, even though she routes a smaller amount of flow. This construction allows us to show that the total cost experienced by all players may not approach to the system optimal cost when the fraction controlled by the leader goes to 1. Thus the bound of $1/d$ proved by Roughgarden for parallel links networks (where d is the fraction controlled by the leader) does not hold in general. We conclude Section 3 with a short proof of the latter result.

In Section 4 we impose more structure to the competition and study series-parallel networks. Our main analytical tool in this section is a structural property that establishes that a circulation in a series-parallel network always has a path from the source to the sink with non-negative flow on every arc. This result does not hold in general networks. In this section, we prove that several natural properties about Wardrop and optimal flows (which would lead to approaches to bound the price of anarchy) do not hold already for series-parallel networks. We conclude the section by providing our bounds on the price of anarchy for series-parallel networks.

Related Literature. Koutsoupias and Papadimitriou (1999) proposed to use the worst-case inefficiency of an arbitrary equilibrium in an arbitrary instance of a game as a way to quantify the impact of not being able to impose a coordinated solution. This concept became known as the price of anarchy (Papadimitriou 2001) and has been used to measure the inefficiency introduced by competition in various settings. Starting with the seminal work of Roughgarden and Tardos (2002), many articles used this idea to quantify the inefficiency of equilibria in various versions of network and congestion games under different assumptions.

Using Stackelberg games and strategies to improve the inefficiency of equilibria precedes the idea of measuring that inefficiency with the price of anarchy. Korilis, Lazar, and Orda (1997) considered a system manager that routes a fraction of the flow with the goal of achieving an optimal solution for the system. Later, Kaporis, Politopoulou, and Spirakis (2005), and Kaporis and Spirakis (2006) quantified the minimum demand that the leader is required to control to guarantee optimality. Sharma and Williamson (2006) considered a related question: they quantified the minimum demand that the leader needs to control to guarantee a lower social cost than at a Nash equilibrium.

More in line with our results, Roughgarden (2004) considered parallel-link networks and a Stackelberg leader that controls a fraction of the demand. The objective of this leader is aligned with the system's because the leader tries to minimize the social cost. As computing the optimal strategy is difficult, Roughgarden studies a number of heuristics with good worst-case performances. We are especially interested in the heuristics SCALE and LLF introduced in that article. Anil Kumar and Marathe (2002) extended this research and found an approximation scheme for the best Stackelberg strategy. More recently, Karakostas and Kolliopoulos (2006), and Swamy (2007) proposed extensions to more general network structures. Specifically, Karakostas and Kolliopoulos give an almost tight bound for the price of anarchy when cost functions are affine and the underlying strategy of the leader is SCALE. On the other hand, Swamy considers series-parallel networks and proves a result similar to our Theorem 1 for a generalization of strategy LLF to general networks (this generalization is different from ours). It is important to point out that, although we proved Theorem 1 independently, it was after we learned of Swamy's result.

2. The Model

We consider a network $G = (V, A)$ with a single OD pair (s, t) and a total demand equal to one (unless we explicitly specify another demand value). As we said in the introduction, there is a single atomic player (the leader) that controls a demand of $d \in [0, 1]$ units of flow, while the rest of the demand $1 - d$ is controlled by nonatomic players (the followers). We associate a nondecreasing cost function $c_a(\cdot) : \mathbb{R}_+ \to \mathbb{R}_+$ to every arc $a \in A$ which maps the total flow f_a on an arc a to its per-unit cost $c_a(f_a)$. Players are homogeneous so they all experience the same cost.

The leader's strategy consists on deciding how to route its demand d. If it selects flow x, the followers will take that flow as fixed and will select a shortest path with respect to the corresponding costs. Namely, they will route their collective demand of $1-d$ according to flow y, defined as a Wardrop equilibrium with respect to costs $\hat{c}_a(y_a) := c_a(x_a+y_a)$ where x_a is considered constant from the nonatomic players' perspective. Here, a Wardrop equilibrium is the solution under which all followers take shortest paths between s and t under the prevailing costs (Beckmann et al. 1956). Given an objective for the leader (see next section), an optimal solution for the leader x^* together with the corresponding Wardrop equilibrium y^* is called a Stackelberg equilibrium because neither the leader, nor the followers, have the incentive to deviate from it. Indeed, x^* is best for the leader and the followers cannot unilaterally influence the leader because they are price-takers.

Consider the solution x selected by the leader following a predefined strategy (the optimal one or a heuristic), and its corresponding Wardrop equilibrium y. To quantify the quality of a flow $f := x + y$, we consider its social cost, which is given by $C(f) := \sum_{a \in A} f_a c_a(f_a)$. Additionally, we write $C^x(x+y) := \sum_{a \in A} x_a c_a(x_a+y_a)$ to refer to the share of the cost paid by the leader, and $C^y(x + y) := \sum_{a \in A} y_a c_a(x_a + y_a)$ to refer to the followers' share. A flow f^{OPT} minimizing $C(\cdot)$ among all feasible flows that route a unit demand is called a *social optimum*.

A global measure of the inefficiency of a solution is given by the ratio between its cost and that of the social optimum. Indeed, the price of anarchy is defined as the worst-case inefficiency of the solution arising from the leader's strategy:

$$\max_{\mathcal{I} \in \text{ instances}} \quad \frac{C(x+y)}{C(f^{\text{OPT}})} \,, \tag{1}$$

where x, y, and f^{OPT} correspond to instance $\mathcal{I}$. On occasions, we will also compare solutions to the Wardrop equilibrium when the total demand is one unit, which will be denoted by f^{WE}. For this solution, we assume that all the demand is controlled by nonatomic users.

Strategies for the Leader. Roughgarden (2004) considered a model in which the atomic player's objective is aligned with the system's: the leader selects the flow x that minimizes $C(x + y)$. He proved that finding the corresponding Stackelberg equilibrium is NP-hard, and hence he considered heuristics for the atomic player for which he could prove guarantees.

In a selfish setting, the atomic player cares more about its cost $C^x(\cdot)$, as opposed to the social cost. Because of that, we shall also consider the worst-case inefficiency from the perspective of the atomic player. Indeed, atomic players would prefer strategies that incur in low per-unit costs. We can quantify the performance of a given strategy for the leader by computing:

$$\max_{\text{instances}} \frac{C^x(x+y)/d}{C(f^{\text{OPT}})/1}\,. \tag{2}$$

Although the complexity of finding the Stackelberg equilibrium in the latter case is open, it is likely to be NP-hard as well. As the leader is unlikely to be able to solve its problem to optimality, we analyze the following two simple heuristics. The leader's bounded rationality justifies the use of approximations to the Stackelberg equilibrium of the game.

SCALE: the leader routes proportionally to the social optimum considering all the demand: $x = df^{\text{OPT}}$ (Roughgarden 2004).

LP: the leader routes its flow in the most expensive arcs first because they are the ones that followers are less likely to choose for themselves. This can be achieved by solving the following LP:

$$\max\Big\{\sum_{a\in A} x_a c_a(f_a^{\text{OPT}}) : x_a \leqslant f_a^{\text{OPT}} \text{ and } x \text{ is a feasible flow for demand } d\Big\}, \tag{3}$$

which generalizes the Largest Latency First (LLF) heuristic proposed by Roughgarden (2004). The latter strategy works in parallel-link networks, and consists in sorting the arcs according to $c_a(f_a^{\text{OPT}})$ from high to low and assigning the necessary demand d to the arcs in that order without surpassing the system-optimal flow f_a^{OPT}.

3. General Networks

The nonatomic players in a Stackelberg game with a single OD pair are always better off than the single atomic player. The reason is that as the nonatomic players decide afterwards, they will always route their demand along shortest paths. Hence, the common length of all the paths used by nonatomic players equals $C^y(x+y)/(1-d) \leqslant \min_P\{c_P(x+y) : x_P > 0\} \leqslant C^x(x+y)/d$. Using this inequality and $C^y(x+y)+C^x(x+y) = C(x+y)$, we also get that $C^y(x+y)/(1-d) \leqslant C(x+y) \leqslant C^x(x+y)/d$, which means that the overall per-unit cost is higher than that of nonatomic players but lower than that of the atomic player.

The following simple example illustrates that if one considers arbitrary strategies, they can perform very badly and hurt both the system and the nonatomic players. In particular, the per-unit cost arising from the Stackelberg competition can be much higher than that of a Wardrop equilibrium.

Example 1. *For arbitrary networks,* $\min_Q\{c_Q(x+y) : y_Q > 0\}$ *can be arbitrarily larger than* $\min_Q\{c_Q(f^{\text{WE}}) : f_Q^{\text{WE}} > 0\}$ *for any* d.

Proof. Consider the network depicted in Figure 1, where the atomic player controls a demand of d, there are k horizontal arcs, and the costs not displayed are 0. The Wardrop equilibrium in this network is the flow that splits the demand equally along the k 'direct' paths through a horizontal arc. Their costs are equal to $c_P(f^{\text{WE}}) = \frac{1}{kd}$. If x routes its d units of flow along the zigzag path that takes all horizontal arcs, all paths have cost $c_P(x+y) \geqslant 1$. □

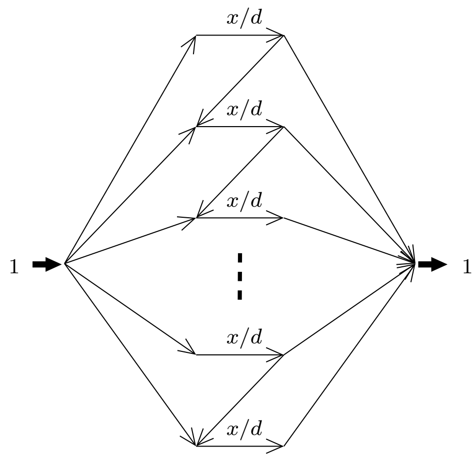

FIGURE 1. Instance for Example 1

For this reason we will restrict our attention to strategies that always route less than a social optimum, that is, $x_a \leqslant f_a^{\text{OPT}}$ for all $a \in A$. Notice that the two strategies defined in the previous section satisfy that property. Using this property, we can prove some bounds that are independent of d but dependent on the set of allowed cost functions. Let us see first that the worst-case inefficiency of the resulting solution cannot be worse than that of a Wardrop equilibrium. Notice that this result holds even though the game is atomic, and Nash equilibria of atomic games can be more inefficient than the worst-case inefficiency of Wardrop equilibria. (We refer the reader to Cominetti et al. 2006 for a discussion of the inefficiency of Nash equilibria in atomic games.)

Proposition 1. *Consider an arbitrary network with cost functions that belong to a set $\mathcal{C}$. If the strategy for the atomic player satisfies that $0 \leqslant x_a \leqslant f_a^{\text{OPT}}$, then $C(x+y) \leqslant (1-\beta(\mathcal{C}))^{-1} C(f^{\text{OPT}})$.*

Proof. Reasoning as in Correa et al. (2005), the fact that y is routed along shortest paths can be expressed as a variational inequality. This implies that for any flow z feasible for the nonatomic player,

$$C(x+y) = \sum_{a \in A} (x_a + y_a) c_a(x_a + y_a) \leqslant \sum_{a \in A} (x_a + z_a) c_a(x_a + y_a).$$

Hence, defining

$$\beta(c) := \max_{x,y,z \in \mathbb{R}_+^3} \frac{(x+z)(c(x+y) - c(x+z))}{(x+y)c(x+y)}$$

and $\beta(\mathcal{C}) := \sup_{c \in \mathcal{C}} \beta(c)$, we get that $\sum_{a \in A} (x_a + z_a) c_a(x_a + y_a) \leqslant \beta(\mathcal{C}) C(x+y) + C(x+z)$. If we set $z = f^{\text{OPT}} - x \geqslant 0$, the claim follows because $C(x+z) = C(f^{\text{OPT}})$. □

Notice that the constant β in the previous proof coincides with that of Correa et al. (2005) because x, y and z are arbitrary reals. Therefore, as long as the leader does not route more flow along each arc than in a social optimum (but otherwise regardless of how it plays), the resulting situation is not worse than for a Wardrop equilibrium in the worst case. Another trivial way to guarantee the same bound is when the leader does not route more flow along each arc than in a Wardrop equilibrium. Indeed, in that case the nonatomic players 'complete' the Wardrop equilibrium, and the bound follows from previous results (Roughgarden 2003; Correa, Schulz, and Stier-Moses 2004). The following result provides a bound on the cost experienced by the leader.

Table 1. Comparison of Props. 1 and 2 as a function of the maximum degree of the allowed polynomials.

Result	Pr. of An.	0	1	2	3	4	...	k
Prop. 1	$(1-\beta(\cdot))^{-1}$	1	4/3	1.626	1.896	2.151		$\frac{(k+1)^{1+1/k}}{(k+1)^{1+1/k}-k}$
Prop. 2	α_k	1	1	1.185	1.688	2.621		$\frac{2}{k+1}\left(\frac{2k}{k+1}\right)^k$

Proposition 2. *Consider an arbitrary network with cost functions that are polynomials of degree at most k. If the strategy for the atomic player satisfies that $0 \leqslant x_a \leqslant f_a^{\text{OPT}}$, then $C^x(x+y) \leqslant \alpha_k C(f^{\text{OPT}})$, where*

$$\alpha_k := \frac{2}{k+1}\left(\frac{2k}{k+1}\right)^k .$$

Proof. As in the proof of Proposition 1, $\sum_a y_a c_a(x_a+y_a) \leqslant \sum_a (f_a^{\text{OPT}} - x_a)c_a(x_a+y_a)$ because $f^{\text{OPT}} - x$ is feasible for the nonatomic players. Rewriting the inequality, we get that $C^x(x+y) \leqslant \sum_a (f_a^{\text{OPT}} - y_a)c_a(x_a + y_a) \leqslant \sum_a (f_a^{\text{OPT}} - y_a)c_a(f_a^{\text{OPT}} + y_a)$. Let us compare the last expression with the cost of the system optimum for monomials of degree k. Letting $\alpha_k := \max_{y,f\in\mathbb{R}_+^2} \frac{(f-y)(f+y)^k}{f^{k+1}}$, we have that $C^x(x+y) \leqslant \alpha_k C(f^{\text{OPT}})$ when cost functions are taken from the set of polynomials of degree up to k. To compute α_k, we solve $\max_y\{(f-y)(f+y)^k : 0 \leqslant y \leqslant f\}$. As the derivative of the objective function is $k(f+y)^{k-1}(f-y)-(f+y)^k$, y has to be equal to $\frac{k-1}{k+1}f$. Plugging y back into the formula, we get the expression in the claim. □

Computing α_k for specific values of k, we get the values shown in Table 1. Although the α_k's increase exponentially on the maximum degree of the polynomials used as cost functions, for small degrees—which is what appears in practice—the inefficiency is small. For example, for affine cost functions, the leader's cost is always bounded from above by the social cost of an optimal solution, regardless of how much demand it controls.

We now present an example to illustrate that Stackelberg equilibria do not behave intuitively. It turns out that if the leader goes from high market power to an absolute monopoly, we observe a discontinuity in the cost experienced by the leader. Notice that the first part of the claim below cannot occur when cost functions are affine or for parallel-link networks (as we shall prove in Proposition 3).

Example 2. *If the leader plays according to strategy LP, its cost $C^x(x+y)$ can be larger that the optimal social cost $C(f^{\text{OPT}})$ (even though the leader may have less amount of flow to route). In addition, $C^x(x+y)$ as function on d can be discontinuous at 1.*

Proof. Let k be an arbitrary (large) integer and n be the maximum integer such that $\left(\frac{k}{k+1}\right)^n(n+1) > 1$. Consider the network depicted in Figure 2, where the atomic player controls a demand of $d = 1 - \frac{1}{2k}$ and $c(x) = \left(\frac{2k}{k+1}\right)^n x^n$. The social optimum routes half the demand on the top path and the other half on the bottom one. Its cost $C(f^{\text{OPT}})$ is $1 + \left(\frac{k}{k+1}\right)^n \leqslant 1 + \frac{1}{n+2}\frac{k+1}{k} \to 1$. As the leader uses strategy LP, it routes $\frac{1}{2}$ on the upper path and $\frac{1}{2} - \frac{1}{2k}$ on the lower one. The nonatomic players route $\frac{1}{2k}$ on the zigzag path, implying that $C^x(x+y) = \frac{1}{2}2 + \left(\frac{1}{2} - \frac{1}{2k}\right)\left(1 + \left(\frac{k}{k+1}\right)^n\right)$. The last expression tends to 3/2 when $k \to \infty$. □

Obviously, the objective is continuous when the flow the leader controls tends to 1 for any fixed instance. The previous example is based on different instances for different values of d. This is in line with the definition of price of anarchy which takes the supremum over *all* instances, as we defined in (1) and (2). The

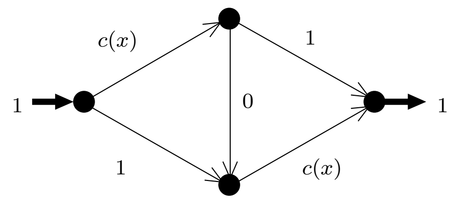

FIGURE 2. Instance for Example 2

previous example shows that the price of anarchy, parametrized with respect to the demand of the leader, is a discontinuous function.

The following proposition was first shown by Roughgarden (2004) for strategy LLF. Here, we present a simpler proof that results from using the variational inequality that describes the behavior of nonatomic players.

Proposition 3 (Roughgarden 2004)**.** *Consider a parallel-link network and an atomic player that uses strategy LP. In that case, the price of anarchy is bounded by* $1/d$*.*

Proof. First we prove that if $x_a > 0$, then $x_a + y_a \leqslant f_a^{\text{OPT}}$. (Intuitively, this statement would seem to be true in general, but Proposition 3 presents a counterexample with a series-parallel network.) With the purpose of getting a contradiction, let us assume the opposite. Therefore, there is an arc $a \in A$ such that $0 < x_a \leqslant f_a^{\text{OPT}}$ and $x_a + y_a > f_a^{\text{OPT}}$. This implies that $y_a > 0$, and that for every arc b such that $x_b + y_b \leqslant f_b^{\text{OPT}}$,

$$c_a(f_a^{\text{OPT}}) < c_a(x_a + y_a) \leqslant c_b(x_b + y_b) \leqslant c_b(f_b^{\text{OPT}}).$$

As b would have to be used by x before a, this implies that $x_b = f_b^{\text{OPT}}$ and $y_b = 0$. Summing over all arcs, we get that $\sum_a (x_a + y_a) > \sum_a f_a^{\text{OPT}}$, which is a contradiction as $x + y$ and f^{OPT} are flows routing the same total demand. Hence, we have that $C^x(x + y) \leqslant C(f^{\text{OPT}})$ because we just proved that arcs used by x do not have more flow than in f^{OPT}.

We continue by using the variational inequality that describes the behavior of the nonatomic players. Considering flow $\frac{1-d}{d}x$, which is feasible for the nonatomic players, we get that $C^y(x+y) \leqslant \sum_a \left(\frac{1-d}{d}x_a\right) c_a(x_a+y_a) = \frac{1-d}{d}C^x(x + y)$. Finally, adding the costs for the atomic and nonatomic players, we get that $C(x + y) \leqslant \frac{1}{d}C(f^{\text{OPT}})$ as we wanted to show. □

4. Series-Parallel Networks

In this section, we restrict our attention to series-parallel networks. This class can be used to model that there are groups of resources that are complements and substitutes of each other. Because of this, series-parallel networks are an order of magnitude more expressive than the class of parallel-link networks, which is what most of the previous research on Stackelberg games for selfish routing considered so far.

A graph is called *series-parallel* when it can be generated recursively from an initial arc s-t using the following two operations. (a) Given an arc a-b, add an intermediate node z and replace the old arc by arcs a-z and z-b. (b) Given an arc a-b, add a parallel arc to it. Although this type of graphs can be arbitrarily complex, it is guaranteed that there won't be an instance like that of Figure 2 (a Braess' paradox network) as a minor. We will use that series-parallel graphs always have paths satisfying the following definition, which is proved below. Note that this is false for general networks; actually, it is not hard to have a counterexample using the graph depicted in Figure 2. Indeed, taking a flow that splits the demand in two halves and another flow that routes everything along the zigzag path, we see that there are instances such that no path is small.

Definition 1. Consider a directed graph $G = (V, A)$ with a single source $s \in V$ and a single sink $t \in V$. Let v and v' be two feasible flows satisfying the same demand. A path P from s to t is called *small* when $v_a \leqslant v'_a$ for all arcs $a \in P$.

Proposition 4. *A small path always exists in a series-parallel graph.*

Proof. Consider the graph $G' = (V, A')$ that includes arcs $a \in A$ such that $v_a \leqslant v'_a$, and reverse arcs $\bar{a}$ of arcs $a \in A$ such that $v_a > v'_a$. Let us first prove by contradiction that there is a path P in G' from s to t. If there is no path, the connected component from s is a cut $S \subseteq V$. All outgoing arcs from S satisfy $v_a > v'_a$ (otherwise a would belong to A'), and all incoming arcs to S satisfy $v_a \leqslant v'_a$ (otherwise a would belong to A'). Summing over the cut, we get that $v_{\Delta^+} - v_{\Delta^-} > v'_{\Delta^+} - v'_{\Delta^-}$, a contradiction because v and v' are feasible flows carrying the same demand.

Therefore there is a path P from s to t in G'. If P takes a backward arc in A', as the network is series-parallel, it has to create a cycle. Removing the cycle and repeating, we get that P only uses forward arcs. This implies that P is small. □

For the case when the leader plays arbitrary strategies, the following results study how $c_P(x + y)$ compares to the cost $C(f^{\text{WE}})$ of an equilibrium, for paths P that are used by nonatomic players.

Proposition 5. *Consider a series-parallel network. Under an arbitrary strategy for the atomic player, the average cost for nonatomic users is smaller than in a Wardrop equilibrium:* $C^y(x + y)/(1 - d) \leqslant C(f^{\text{WE}})$.

Proof. We restrict attention to arcs that are used by some flow ($f_a^{\text{WE}} > 0$ or $x_a > 0$ or $y_a > 0$) because otherwise they are irrelevant. As y routes flow along minimum cost paths,

$$\frac{C^y(x + y)}{1 - d} = \min_{Q: y_Q > 0} c_Q(x + y) \leqslant c_P(x + y).$$

Proposition 4 implies that there is a path P such that $x_a + y_a \leqslant f_a^{\text{WE}}$ in all its arcs. Taking an appropriate decomposition, f^{WE} carries flow along P. Thus, the average cost for nonatomic users is bounded by $c_P(f^{\text{WE}}) = C(f^{\text{WE}})$. □

The next example proves that the property that was used in Proposition 3 needs not be true in general. Notice that the leader uses the strategy LP in the example.

Example 3. *In a series-parallel network,* $0 < x_a \leqslant f_a^{\text{OPT}}$ *does not imply that* $x_a + y_a \leqslant f_a^{\text{OPT}}$.

Proof. Consider the series-parallel network depicted in Figure 3, where the atomic player controls $d = 1/3$ of the flow and the nonatomic one controls $2/3$. Although the atomic user routes less than the system optimum in all arcs, $x + y$ is larger than f^{OPT} in the lower path, where the followers route $7/12$ units of flow. □

In addition, the following example shows a similar property for nonatomic flows. Although a Wardrop equilibrium may not load one arc too much, the leader may force followers to excessively use some arcs. As before, in this example the atomic player uses strategy LP.

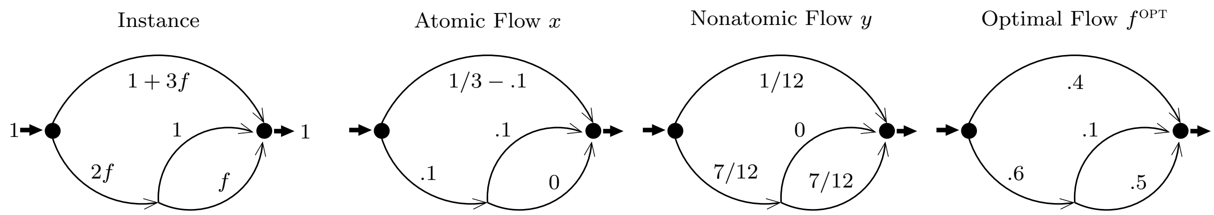

Figure 3. Instance for Example 3

Example 4. *In a series-parallel network, $y_a > 0$ does not imply that $x_a + y_a \leqslant f_a^{\text{WE}}$.*

Proof. Consider the series-parallel network depicted in Figure 4, where the total demand is 1.1 and atomic player controls $d = .45$ while the nonatomic players control the rest. In the Stackelberg game, the lower left arc is used more than in a Wardrop equilibrium. □

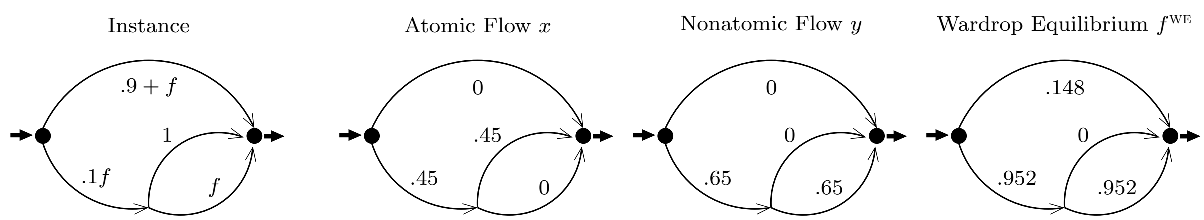

Figure 4. Instance for Example 4

Example 5. *In a series-parallel network, the set of arcs used by a Wardrop equilibrium and by a system optimum may not be included one inside the other.*

Proof. One inclusion follows easily from, e.g., Pigou's network: the system optimum uses both arcs while the equilibrium uses only one (see Roughgarden and Tardos 2002). Let us show that the opposite may happen. Consider the series-parallel network depicted in the graph on the left of Figure 5, where the total demand is 1.01. As can be seen in the picture, the Wardrop equilibrium uses all arcs while the system optimum does not use the lower right arc. □

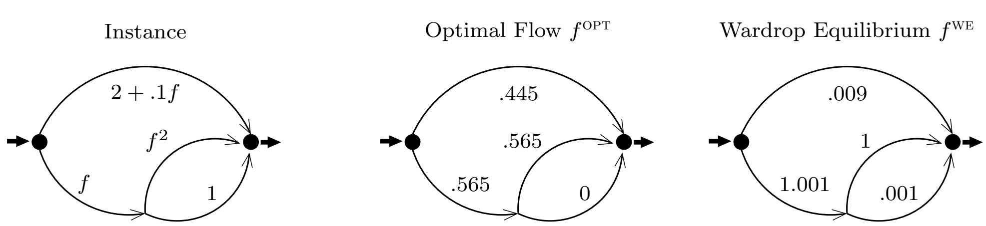

Figure 5. Instance for Example 5

Let us now derive bounds for the cost of paths used by the atomic and nontoamic players. The following results will allow us to bound the inefficiency of a Stackelberg equilibrium.

Lemma 1. *Assume that the atomic player uses strategy LP. If the nonatomic players carry flow along path P ($y_P > 0$), then $c_P(x+y) \leqslant C(f^{\text{OPT}})/d$.*

Proof. Let P^y be a path such that $y_{P^y} > 0$. First we will prove that there is a path P^* such that $x_a + y_a \leqslant f_a^{\text{OPT}}$ and $x_a < f_a^{\text{OPT}}$ for all $a \in P^*$. We proceed as in Proposition 4 by letting S be the connected component from s that includes forward arcs such that $x_a + y_a \leqslant f_a^{\text{OPT}}$ and $x_a < f_a^{\text{OPT}}$, and backward arcs such that $x_a + y_a \geqslant f_a^{\text{OPT}}$. To get a contradiction, let us assume that $t \notin S$. Hence, S is a cut and its forward arcs satisfy that $x_a + y_a > f_a^{\text{OPT}}$ or $x_a \geqslant f_a^{\text{OPT}}$, and its backward arcs satisfy $x_a + y_a < f_a^{\text{OPT}}$. As the total flow across the cut for $x+y$ and f^{OPT} has to be the same, all forward arcs much satisfy that $x_a = f_a^{\text{OPT}}$ and $y_a = 0$ and there cannot be backward arcs. This is a contradiction and therefore, using that the graph is series-parallel, there must be a path P^* with the desired property. Because y is at equilibrium and carries flow along P^y, $\sum_{a \in P^y} c_a(x_a + y_a) \leqslant \sum_{a \in P^*} c_a(x_a + y_a)$. By the monotonicity of $c(\cdot)$, the latter is bounded by $\sum_{a \in P^*} c_a(f_a^{\text{OPT}})$. As x is a maximizer of the LP displayed in (3), the fact that P^* is not saturated implies that it is shortest with respect to $c_a(f_a^{\text{OPT}})$ among those used by x (i.e., $\sum_{a \in P^*} c_a(f^{\text{OPT}}) \leqslant \sum_{a \in Q} c_a(f^{\text{OPT}})$ for any path Q such that $x_Q > 0$).

Note that the cost $C(f^{\text{OPT}}) = \sum_{P \in \mathcal{P}: f_P^{\text{OPT}} > 0} f_P^{\text{OPT}} c_P(f^{\text{OPT}})$ can be viewed as the average cost of paths used by the flow f^{OPT}. Markov's inequality[1] implies that there cannot be more than a fraction d of the paths with cost above $C(f^{\text{OPT}})/d$. Hence, the cheapest path used by flow x has cost at most $C(f^{\text{OPT}})/d$, implying that $c_{P^*}(f^{\text{OPT}}) \leqslant C(f^{\text{OPT}})/d$ and completing the proof. □

Lemma 2. *Assume that the atomic player uses strategy LP. If the atomic player carries flow along path P ($x_P > 0$), then $c_P(x+y) \leqslant 2c_P(f^{\text{OPT}})$.*

Proof. Let P^x be a path such that $x_a > 0$ for all $a \in P^x$ (i.e., there is a flow decomposition for which $x_{P^x} > 0$). Let $P^s = \{a \in P^x : x_a + y_a \leqslant f_a^{\text{OPT}}\}$ and $P^l = \{a \in P^x : x_a + y_a > f_a^{\text{OPT}}\}$, so that $P^x = P^s \cup P^l$ and $c_{P^x}(x+y) = \sum_{a \in P^s} c_a(x_a + y_a) + \sum_{a \in P^l} c_a(x_a + y_a)$. By monotonicity of $c(\cdot)$, $\sum_{a \in P^s} c_a(x_a + y_a) \leqslant \sum_{a \in P^x} c_a(f_a^{\text{OPT}})$. We will complete the proof by showing that $\sum_{a \in P^l} c_a(x_a + y_a)$ is bounded by the same expression. Note that by the choice of x, $y_a > 0$ for all $a \in P^l$. As the network is series-parallel, there is a path P^y such that $P^y \supset P^l$ and $y_a > 0$ for all $a \in P^y$ (otherwise there are two edges e and e' in P^l such that no path from e to e' has y-flow, which would imply that the network cannot be series-parallel). Using the argument in the proof of Lemma 1, we get that

$$\sum_{a \in P^l} c_a(x_a + y_a) \leqslant \sum_{a \in P^y} c_a(x_a + y_a) \leqslant \sum_{a \in P^x} c_a(f_a^{\text{OPT}}).$$

On the other hand, by Prop. 4 there is a path P'' such that $x_a + y_a < f_a^{\text{OPT}}$ for all $a \in P''$. As y is at equilibrium and routes flow along P',

$$\sum_{a \in B_P} c_a(x_a + y_a) \leqslant \sum_{a \in P'} c_a(x_a + y_a) \leqslant \sum_{a \in P''} c_a(x_a + y_a) \leqslant \sum_{a \in P''} c_a(f_a^{\text{OPT}}).$$

To conclude, we prove that $\sum_{a \in P''} c_a(f_a^{\text{OPT}}) \leqslant \sum_{a \in P} c_a(f_a^{\text{OPT}})$. This holds since otherwise we can define:

$$\tilde{x} = x + \varepsilon(e_{P''} - e_P),$$

[1]Markov's inequality says that if X is a RV with positive support, then $P(X > zE(X)) \leqslant 1/z$.

where e_A is the indicator vector of path A. Therefore $\tilde{x}$ is feasible strategy for the Stackelberg player since $x_a > 0$ for all $a \in P$ and $x_a < f_a^{\text{OPT}}$ for all $a \in P''$. This is a contradiction because the objective value of $\tilde{x}$ is larger than that of x. □

The previous two lemmas can be used to bound the price of anarchy. The bound we present below generalizes that of Proposition 3 to series-parallel networks, losing only an additive factor of 1.

Theorem 1. *Assume that we have a Stackelberg game on a series-parallel network and the atomic player uses strategy LP. Then, the price of anarchy is bounded by* $1 + \frac{1}{d}$.

Proof. Lemma 1 implies that $C^y(x+y) = \sum_{P\in\mathcal{P}} y_P c_P(x+y) \leqslant (C(f^{\text{OPT}})/d) \sum_{P\in\mathcal{P}} y_P = (1/d-1)C(f^{\text{OPT}})$, and Lemma 2 implies that $C^x(x+y) = \sum_{P\in\mathcal{P}} x_P c_P(x+y) \leqslant 2\sum_{P\in\mathcal{P}} f_P^{\text{OPT}} c_P(f^{\text{OPT}}) \leqslant 2C(f^{\text{OPT}})$. Putting the two together, we get that $C(x+y) \leqslant (1+1/d)C(f^{\text{OPT}})$. □

The following result proves bounds on the cost of the solution resulting from the leader playing strategy LP. Building on earlier results on the inefficiency of Wardrop equilibria, they can be used to prove that the price of anarchy is not bounded.

Theorem 2. *Assume that we have a Stackelberg game on a series-parallel network and the atomic player uses strategy LP. Then, we have that* (i) $C(x+y) \leqslant 2C(f^{\text{OPT}}) + (1-d)C(f^{\text{WE}})$, *and* (ii) $C(x+y) \leqslant C(f^{\text{OPT}}) + C(f^{\text{WE}})$.

Proof. To see (i) it is enough to use Proposition 5 instead of Lemma 1, in the proof of the last theorem. Let us prove (ii). To this end we first go back to the proof of Lemma 2 where we implicitly showed that

$$\sum_{P\in\mathcal{P}} x_P c_P(x+y) \leqslant \sum_{P\in\mathcal{P}} f_P^{\text{OPT}} c_P(f^{\text{OPT}}) + dL^y,$$

where L^y is the cost of a path P^y such that $y_{P^y} > 0$. As y is at equilibrium, all paths used by y have equal cost, implying that $L^y = C^y(x+y)/(1-d)$. Hence,

$$\begin{aligned}\sum_{P\in\mathcal{P}} (x_P + y_P) c_P(x+y) &\leqslant \sum_{P\in\mathcal{P}} f_P^{\text{OPT}} c_P(f^{\text{OPT}}) + dL^y + \sum_{P\in\mathcal{P}} y_P c_P(x+y) \\ &= C(f^{\text{OPT}}) + L^y = C(f^{\text{OPT}}) + C^y(x+y)/(1-d) \leqslant C(f^{\text{OPT}}) + C(f^{\text{WE}}),\end{aligned}$$

where the last inequality follows from Proposition 5. □

5. Conclusions

In the context of network competition with atomic players, we have presented positive and negative results related to certain heuristics for an atomic player that competes against many small firms that do not control prices. We have proved that as long as it routes less flow in all arcs than in an optimal solution, the worst-case inefficiency is not going to be too bad. Moreover, we have extended earlier results from parallel-link networks to series-parallel ones. Indeed, we have proved that the maximum inefficiency of a heuristic that directs the atomic player to select the resources that are less likely to be used by the nonatomic players is not high either.

Although the complexity of the problem that the leader has to solve is still unknown, in practice the problem is a mathematical program with equilibrium constraints (MPEC). If the leader wants to compute the optimal strategy for a particular instance, there are relatively standard optimization techniques to solve this problem. For a background on MPECs and solution methods, we refer the interested reader to the book by Luo, Pang, and Ralph (1996).

Acknowledgments

The first author gratefully acknowledges the support of CONICYT via project Anillo en Redes ACT08. The second author gratefully acknowledges the support of the Center for International Business Education and Research (CIBER), and of the Center for Excellence in E-Business (CEBIZ), both at Columbia University.

References

Anil Kumar, V. S. and M. V. Marathe (2002). Improved results for Stackelberg scheduling strategies. In P. Widmayer, F. Triguero, R. Morales, M. Hennessy, S. Eidenbenz, and R. Conejo (Eds.), *Automata, Languages and Programming: Proceedings of the 29th International Colloquium (ICALP)*, Málaga, Spain, Volume 2380 of *Lecture Notes in Computer Science*, pp. 776–787. Springer, Heidelberg.

Beckmann, M. J., C. B. McGuire, and C. B. Winsten (1956). *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT.

Cominetti, R., J. R. Correa, and N. E. Stier-Moses (2006). Network games with atomic players. In M. M. Yung (Ed.), *Automata, Languages and Programming: Proceedings of the 33rd International Colloquium (ICALP)*, Venice, Italy, Volume 4051 of *Lecture Notes in Computer Science*, pp. 525–536. Springer, Heidelberg. Part 1.

Correa, J. R., A. S. Schulz, and N. E. Stier-Moses (2004). Selfish routing in capacitated networks. *Mathematics of Operations Research 29*(4), 961–976.

Correa, J. R., A. S. Schulz, and N. E. Stier-Moses (2005). On the inefficiency of equilibria in congestion games. In M. Jünger and V. Kaibel (Eds.), *Proceedings of the 11th International Integer Programming and Combinatorial Optimization Conference (IPCO)*, Berlin, Germany, Volume 3509 of *Lecture Notes in Computer Science*, pp. 167–181. Springer, Heidelberg.

Kaporis, A., E. Politopoulou, and P. Spirakis (2005). The price of optimum in Stackelberg games. *Electronic Colloquium on Computational Complexity (ECCC)*. Report 56.

Kaporis, A. and P. Spirakis (2006). The price of optimum in Stackelberg games on arbitrary single commodity networks and latency functions. In *18th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, Cambridge, MA, pp. 19–28. ACM Press, New York, NY.

Karakostas, G. and S. G. Kolliopoulos (2006). Stackelberg strategies for selfish routing in general multicommodity networks. Technical Report AdvOL2006/08, McMaster University.

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

Koutsoupias, E. and C. H. Papadimitriou (1999). Worst-case equilibria. In C. Meinel and S. Tison (Eds.), *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, Trier, Germany, Volume 1563 of *Lecture Notes in Computer Science*, pp. 404–413. Springer, Heidelberg.

Luo, Z.-Q., J.-S. Pang, and D. Ralph (1996). *Mathematical Programs with Equilibrium Constraints*. Cambridge University Press, Cambridge, U.K.

Papadimitriou, C. H. (2001). Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, Hersonissos, Greece, pp. 749–753. ACM Press, New York, NY.

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

Roughgarden, T. (2004). Stackelberg scheduling strategies. *SIAM Journal on Computing 33*(2), 332–350.

Roughgarden, T. and É. Tardos (2002). How bad is selfish routing? *Journal of the ACM 49*, 236–259.

Sharma, Y. and D. P. Williamson (2006). Stackelberg thresholds in network routing games or The value of altruism. Technical Report TR2006-2042, Cornell University.

Swamy, C. (2007). The effectiveness of Stackelberg strategies and tolls for network congestion games. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, New Orleans, LA. SIAM, Philadelphia, PA. To Appear.

von Stackelberg, H. (1934). *Marktform und Gleichgewicht*. Springer, Wein. English translation: The Theory of the Market Economy, published in 1952 by Oxford University Press.

Wardrop, J. G. (1952). Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers, Part II, Vol. 1*, 325–378.