Computing the price of anarchy in atomic network congestion games (invited talk) ⋆

Nicolas Markey

IRISA, CNRS & Inria & Univ. Rennes, France

Abstract. Network congestion games are a simple model for reasoning about routing problems in a network. They are a very popular topic in algorithmic game theory, and a huge amount of results about existence and (in)efficiency of equilibrium strategy profiles in those games have been obtained over the last 20 years.
In particular, the price of anarchy has been defined as an important notion for measuring the inefficiency of Nash equilibria. Generic bounds have been obtained for the price of anarchy over various classes of networks, but little attention has been put on the effective computation of this value for a given network. This talk presents recent results on this problem in different settings.

1 Atomic network congestion games

Congestion games have been introduced by Rosenthal in 1973 [28,29] as a model for reasoning about traffic-routing or resource-sharing problems [27,36]. *Network* congestion games [22,30] are played by n players on a weighted graph $G = \langle V, E, l\rangle$, where l labels the edges of G with non-decreasing functions (called *latency functions*): each player has to select a route in this graph from their source state to their target state; for each edge they use, they have to pay $l(k)$, where k is the total number of players using the same edge. These games are called *atomic* as opposed to non-atomic ones, where each player can route arbitrarily-small parts of their load along different paths. In the sequel, we focus on the *symmetric* setting, where all players have the same source- and target vertices.

Example 1. Figure 1 represents an atomic network congestion game: it is made of a graph, with a state s_0 as the source state of the four players (represented by the four tokens in that state), and s_1 as their target state. Edges $s_0 \to l$ and $r \to s_1$ are labelled with the identity function $k \mapsto k$, meaning that the cost of using each of these edges is equal to the number of players using it; in this example, the cost of the other edges is a constant.

Figure 2 displays a round of this game, where one player takes the path $\pi_l\colon s_0 \to l \to s_1$, and the other three take $\pi_r\colon s_0 \to r \to s_1$; the cost for the

⋆ This paper is based on joint works with Nathalie Bertrand, Aline Goeminne, Suman Sadhukhan, Ocan Sankur, and benefited from discussions with Arthur Dumas and Stéphane Le Roux. These works have received fundings from ANR projects TickTac and BisoUS.

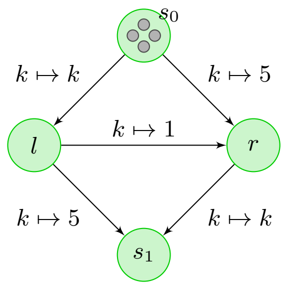

Fig. 1. An example of a congestion game with four players.

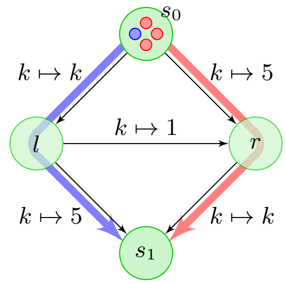

Fig. 2. One player going to s_1 via l, and three players via r.

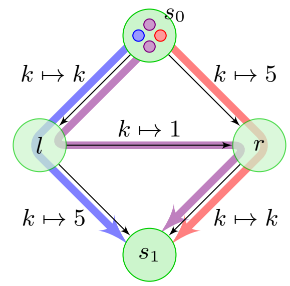

Fig. 3. One player going via l, one via r, and two using the transverse edge.

former player is 6, while it is 8 for the other three. Figure 3 shows another profile of strategies, with one player taking π_l, one player taking π_r, and the remaining two using $\pi_{lr} \colon s_0 \to l \to r \to s_1$. In that case, edges $s_0 \to l$ and $r \to s_1$ are used by three players, hence they have a cost of 3; the first two players thus have to pay 8, while the cost for the other two is 7. $\lhd$

In symmetric atomic network congestion games, all agents are indistinguishable. A strategy profile can thus be represented as a vector $\mathsf{p} = (p_\pi)_{\pi \in \mathsf{Paths}}$, where Paths is the set of all paths from source to target, and p_π is the number of players selecting path π in that profile. For a strategy profile $\mathsf{p} = (p_\pi)_{\pi \in \mathsf{Paths}}$ and a path ρ, we write $\mathsf{p} + \rho$ to denote the strategy profile obtained from p by incrementing p_ρ; $\mathsf{p} - \rho$ is defined similarly, assuming that ρ occurs in p. The *flow* of a strategy profile p is the vector $\mathsf{flow}(\mathsf{p}) = (q_e)_{e \in E}$ such that for each edge e, q_e is the number of players using edge e in the profile p. The cost a following path ρ in a strategy profile p is then defined as $\mathsf{cost}_\mathsf{p}(\rho) = \sum_{e \in \rho} l(q_e)$. The social cost of a strategy profile is then defined as follows: $\mathsf{cost}(\mathsf{p}) = \sum_{\rho \in \mathsf{Paths}} p_\rho \cdot \mathsf{cost}_\mathsf{p}(\rho) = \sum_{e \in E} q_e \cdot l(q_e)$.

In congestion games, two kinds of behaviours of the players are of particular interest:

- the *social optima*, which minimise the *total cost* of the whole set of players: formally, p is a social optimum if $\mathsf{cost}(\mathsf{p})$ is less than or equal to $\mathsf{cost}(\mathsf{p}')$ for any profile p';
- the *Nash equilibria*, which correspond to selfish behaviours, each player aiming at minimising their *individual cost* given the strategies of the other players: formally, p is a Nash equilibrium if for any ρ occurring in p and any ρ', it holds $\mathsf{cost}_\mathsf{p}(\rho) \leq \mathsf{cost}_{\mathsf{p}+\rho'-\rho}(\rho')$.

Social optima represent collaborative behaviours that would be played if the players had to share the total cost, while Nash equilibria correspond to selfish behaviours.

Example 1 (contd). Consider again the network congestion game of Fig. 1. It is easy to check that the social optimum is achieved when two players take π_l and the other two take π_r: the total cost is 28.

Finding Nash equilibria is more difficult. First notice that the social optimum obtained above is not a Nash equilibrium: if one of the players taking π_l decided to take π_{lr}, their cost would be 6 instead of 7. (Notice that this does not improve the social optimum since the two players taking π_r would see their cost increase).

The strategy profile of Fig. 3 is a Nash equilibrium: it can be checked that no players can lower their cost by switching to another path; the total cost (also called *social cost*) of this Nash equilibrium is 30. Notice that another Nash equilibrium exists in this game: when all the players take π_{lr}, each of them pays a cost of 9, and switching to one of the other two paths would give the same cost; the social cost of this Nash equilibrium is 36. $\triangleleft$

In his seminal papers [28,29], Rosenthal shows that congestion games always admit Nash equilibria, by exhibiting a *potential function*, which decreases when any player individually switches to a better strategy. As shown in Example 1, Nash equilibria are in general not unique, and they may have different social costs. Koutsoupias and Papadimitriou defined the *price of anarchy* as the ratio between the social cost of the worst Nash equilibrium and the social optimum [23]: it measures how bad selfish behaviours can be compared to an optimal collaborative solution. Symmetrically, with a more optimistic point of view, the *price of stability* is the ratio between the best Nash equilibrium and the social optimum, and measures the minimal cost of moving from a collaborative solution to a selfish one [2].

An impressive amount of results have been obtained about (network) congestion games during the last 25 years:

- numerous variations on the model have been studied: atomic vs. non-atomic players [32], weighted players [16] with splittable or unsplittable flows [31], sequential or simultaneous choices [26], ... Several restrictions on the graphs have been used to obtain certain results, such as series-parallel graphs [15,19] or graphs with only parallel links [23,25]. Mixed-strategy Nash equilibria have also been been considered, see e.g. [10,25].
- the complexity of computing *some* Nash equilibrium has been studied. Since Nash equilibria always exist and are made of one path per player, the problem is in TFNP; it is actually in FPTIME if all players have the same source and target states, and PLS-complete otherwise [14]. The existence of a Nash equilibrium whose social cost lies below (resp. above) a given threshold (which is the decision problem associated to the computation of a best (resp. worst) Nash equilibrium) is NP-complete [34], even for series-parallel graphs. We call this problem *constrainted Nash-equilibrium problem* in the sequel;
- in some cases, the price of anarchy can be proven to be bounded, independently of the graph: this is the case for instance in atomic network congestion games with linear latency functions, where the price of anarchy can be at

most 5/2 [6,10]; under the series-parallel restriction, the price of anarchy is bounded by 2 [19].

2 Computing the prices of anarchy and stability

In contrast to the topics listed above, the problem of computing the prices of anarchy and stability of a *given* network has received little attention [11,12]. We address this problem by developing techniques to compute best and worst Nash equilibria and social optima for atomic network congestion games in three different settings: first, for series-parallel networks with linear latency functions, we compute a representation of Nash equilibria and social (local) optima for any number of players; second, in a more *dynamic* (sometimes called *sequential*) setting where the players can adapt their route along the way, for piecewise-affine latency functions, we develop an exponential-space algorithm for solving the constrained Nash-equilibrium problem; third, we extend these results to timed network games.

2.1 Series-parallel networks with linear latency functions

Using the notations above, we see strategy profiles as $|\mathsf{Paths}|$-dimensional vectors $(p_\pi)_{\pi\in\mathsf{Paths}}$, and their associated flows as $|E|$-dimensional vectors $(q_e)_{e\in E}$.

The flow of the social optima are the flows $(q_e)_{e\in E}$ such that $\sum_{e\in E} q_e \cdot l(q_e)$ is minimal over all flows. As we assume that l is linear, we get a quadratic expression to minimize, which is intractable. Instead, we focus on social *local* optima, which are only optimal among the profiles obtained by changing a single strategy. We prove that a flow $(q_e)_{e\in E}$ is locally optimal if, and only if, for all paths ρ and ρ' such that $q_e > 0$ for all $e \in \rho$, it holds

$$\text{for all } \rho, \rho', \text{ if } q_e > 0 \text{ for all } e \in \rho, \text{ then } \sum_{e\in\rho\setminus\rho'} l(2q_e - 1) \leq \sum_{e\in\rho'\setminus\rho} l(2q_e + 1).$$

Similarly, flows of Nash equilibria p characterized as vectors $(q_e)_{e\in E}$ satisfying

$$\text{for all } \rho, \rho', \text{ if } \rho \in \mathsf{p}, \text{ then } \sum_{e\in\rho\setminus\rho'} l(q_e) \leq \sum_{e\in\rho'\setminus\rho} l(q_e + 1).$$

It follows:

Theorem 2 ([8,33]). *For any atomic network congestion game G with linear latency functions, the sets $\mathsf{SLO}(G)$ of all social local optima, and $\mathsf{NE}(G)$ of all Nash equilibria, as well as the sets of corresponding flows, are semi-linear [17].*

When restricting to series-parallel graphs (which are graphs obtained from the trivial single-edge graph by series- and parallel compositions) [35]), we can prove that both sets of flows expand in a unique direction δ: in other terms, they can be written as $B \cup \bigcup_{i\in I k\geq 0} b_i + k \cdot \delta$, for finite sets B and I.

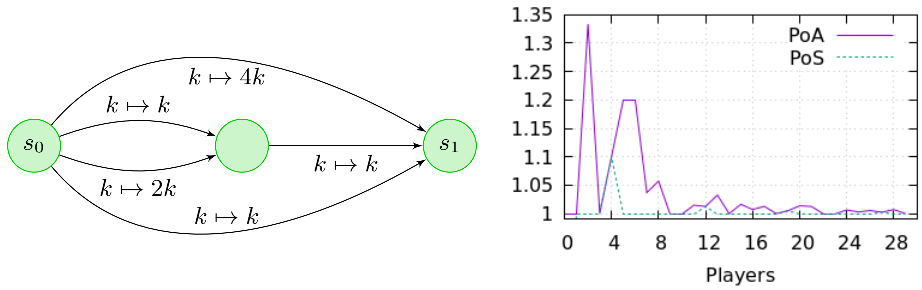

Fig. 4. Evolution of prices of anarchy and stability as a function of the number of players

One consequence of this is that, in our setting, the prices of anarchy and stability converges to 1 when the number of players tends to infinity. It also allows us to compute finite representations of the sets of all social local optima and Nash equilibria, which in turn can be used to compute the prices of anarchy and stability. Figure 4 displays an example of such a computation on a 4-path example for up to 30 players.

2.2 Dynamic network congestion games

Dynamic network congestion games are a refined version of network congestion games with two main changes: first, the number of players using a given edge is now measured *synchronously*, considering that the load of an edge depends on time [20,3]; second, the players do not choose simple paths but adaptive (pure) strategies, that depend on the other players' past moves: this provides a way of reacting to strategy deviations *during the course of the game* [26,5,13].

Example 1 (contd). Consider again the network congestion game of Example 1, and in particular the strategy profile of Figure 3: the edge $r \to s_1$ is used by strategies π_r and π_{lr}, but considering that the players move synchronously, this edge is traversed at the second step in π_r and at the third step in π_{lr}; there is no congestion effect between the single player following path π_r and the two players following π_{lr}; in that case, all three of them will pay a cost of 6.

Strategies may now depend on what the other players have been playing (but the players are still anonymous). Considering again the example of Fig. 3, the players can now decide to go to l, and depending on the number of players in l and in r, choose to go directly to s_1 or take the edge $l \to r$. $\triangleleft$

We call *blind strategies* the special case of strategies that we used in the previous section, which follow a single path independently of the other players' moves. We prove:

Theorem 3 ([7,33]). *Any dynamic network congestion game admits a (blind) Nash equilibrium.*

This is proven in two steps: using a potential function inspired from that of [28,29], we prove that dynamic network congestion games admit blind Nash equilibria (considering only blind deviations); we then prove that blind Nash equilibria are also plain Nash equilibria.

Actually, there exist networks in which there are more Nash equilibria in the dynamic setting than in the classical one. Figure 5 displays an example of a network congestion game in which, for three players, there is a Nash equilibrium with social cost 36, while all blind Nash equilibria have cost at least 37 [7,33]. It follows that the price of anarchy is in general higher (and the price of stability is lower) in the dynamic setting than when restricting to blind strategies.

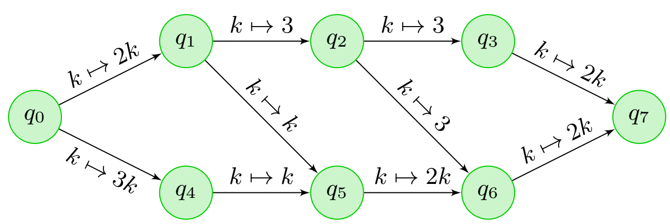

Fig. 5. A network congestion game in which blind Nash equilibria are sub-optimal

Since Nash equilibria always exist, the prices of anarchy and stability are always defined. They can be computed by solving the constrained social-optimum (resp. Nash equilibrium) problem, which asks whether some strategy profile (resp. Nash equilibrium) has social cost satisfying a given linear inequality.

Theorem 4 ([7,33]). *The constrained social-optimum problem can be solved in polynomial space. The constrained Nash-equilibrium problem can be solved in exponential space. The prices of anarchy and stability of a given atomic network congestion game can be computed in exponential space.*

Notice that such complexity results cannot be achieved by just building the explicit game on the graph of configurations, since this graph would have size $|V|^n$ (with n given in binary).

In this dynamic setting, it is usually more relevant to consider subgame-perfect equilibria, which rule out non-credible threats. It is open whether subgame perfect always exist in dynamic network congestion games. Adapting techniques developed in [9], we get:

Theorem 5 ([7,33]). *The constrained subgame-perfect-equilibrium problem can be solved in double-exponential time.*

3 Timed network games

Several works have proposed to add a real-time dimension to network congestion games. In some of those extensions, congestion affects time (instead of cost), with the aim of minimizing the total time to reach the target state [21,20,24]. For the case where congestion affects the cost, the class of timed network games is introduced in [3,4]: following the semantics of timed automata [1], in a timed network game the players can decide to take a transition, or to spend time in the state they are visiting. Transitions are immediate and have no cost, but their availability depends on time; states are decorated with latency functions indicating the cost for spending one time unit in that state, as a function of the load of that state.

Example 6. Figure 6 represents a timed network game: vertices are labelled with their latency functions, and edges are decorated with intervals indicating the time at which they can be traversed.

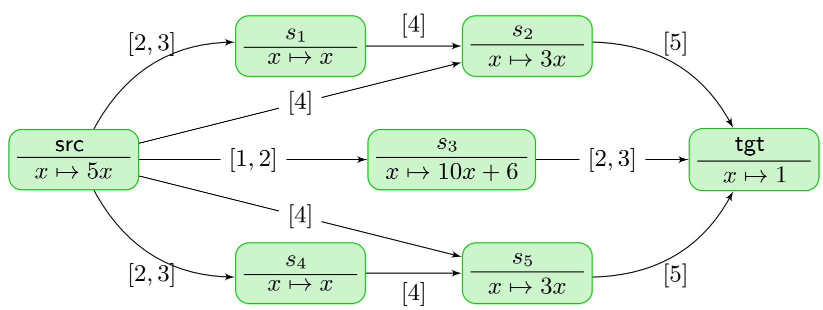

Fig. 6. Example of a timed network game

An example of a (blind) strategy π_1 in that timed network game consists in letting 1.3 time units elapse in src, then go to state s_3, and go to tgt at time 2. Another example π_2 could propose to stay in src until time 4, then go to s_5, and reach tgt at time 5.

If two players follow those strategies π_1 and π_2, the first one will have a total cost of 24.2 (namely, $(5 \times 2) \times 1.3 + (10 \times 1 + 6) \times 0.7$) when entering tgt, while the second one will have a total cost of 29.5 (in details, $(5 \times 2) \times 1.3 + (5 \times 1) \times 2.7 + (3 \times 1) \times 1$). $\lhd$

In this framework, *boundary* blind strategies, which are blind strategies that always take transitions at one of the boundaries of their timing constraints, have been identified as an important subclass of strategies: any timed network game has a boundary social optimum, and a boundary Nash equilibrium (among blind strategies). However, best and worst blind Nash equilibria need not be boundary [3].

Focusing on the computation of best and worst Nash equilibria, the results of the previous section can be extended in discrete time, for non-blind strategies, as follows:

Theorem 7 ([18]). *In* discrete-time *timed network games, the constrained social-optimum problem can be solved in polynomial space, and the constrained Nash-equilibrium problem can be solved in exponential space. The prices of anarchy and stability of a given timed network game can be computed in exponential space.*

References

1. R. Alur and D. L. Dill. A theory of timed automata. *Theor. Computer Science*, 126(2):183–235, 1994.
2. E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In FOCS'04, p. 295–304. IEEE Comp. Soc. Press, 2004.
3. G. Avni, S. Guha, and O. Kupferman. Timed network games. In MFCS'17, LIPIcs 84, p. 37:1–37:16. Leibniz-Zentrum für Informatik, 2017.
4. G. Avni, S. Guha, and O. Kupferman. Timed network games with clocks. In MFCS'18, LIPIcs 117, p. 23:1–23:18. Leibniz-Zentrum für Informatik, 2018.
5. G. Avni, T. A. Henzinger, and O. Kupferman. Dynamic resource allocation games. In SAGT'16, LNCS 9928, p. 153–166. Springer, 2016.
6. B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In STOC'05, p. 57–66. ACM Press, 2005.
7. N. Bertrand, N. Markey, S. Sadhukhan, and O. Sankur. Dynamic network congestion games. In FSTTCS'20, LIPIcs 182, p. 40:1–40:16. Leibniz-Zentrum für Informatik, 2020.
8. N. Bertrand, N. Markey, S. Sadhukhan, and O. Sankur. Semilinear representations for series-parallel atomic congestion games. In FSTTCS'22, LIPIcs 250, p. 32:1–32:20. Leibniz-Zentrum für Informatik, 2022.
9. Th. Brihaye, V. Bruyère, A. Goeminne, J.-F. Raskin, and M. Van den Bogaard. The complexity of subgame perfect equilibria in quantitative reachability games. In CONCUR'19, LIPIcs 140, p. 13:1–13:16. Leibniz-Zentrum für Informatik, 2019.
10. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In STOC'05, p. 67–73. ACM Press, 2005.
11. R. Colini-Baldeschi, R. Cominetti, P. Mertikopoulos, and M. Scarsini. When is selfish routing bad? The price of anarchy in light and heavy traffic. *Operations Research*, 68(2):411–434, 2020.
12. R. Cominetti, V. Dose, and M. Scarsini. The price of anarchy in routing games as a function of the demand. *Mathematical Programming*, 2023. To appear.
13. J. R. Correa, J. de Jong, B. de Keizer, and M. Uetz. The inefficiency of Nash and subgame-perfect equilibria for network routing. *Mathematics of Operations Research*, 44(4):1286–1303, 2019.
14. A. Fabrikant, Ch. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In STOC'04, p. 604–612. ACM Press, 2004.
15. D. Fotakis. Congestion games with linearly-independent paths: convergence time and price of anarchy. In SAGT'08, LNCS 4997, p. 33–45. Springer, 2008.

16. D. Fotakis, S. Kontogiannis, and P. G. Spirakis. Selfish unsplittable flows. *Theor. Computer Science*, 348(2-3):226–239, 2005.
17. S. Ginsburg and E. Spanier. Semigroups, Presburger formulas, and languages. *Pacific Journal of Mathematics*, 16(2):285–296, 1966.
18. A. Goeminne, N. Markey, and O. Sankur. Non-blind strategies in timed network congestion games. In FORMATS'22, LNCS 13465, p. 183–199. Springer, 2022.
19. B. Hao and C. Michini. Inefficiency of pure nash equilibria in series-parallel network congestion games. In WINE'22, LNCS 13778, p. 3–20. Springer, 2022.
20. M. Hoefer, V. S. Mirrokni, H. Röglin, and S.-H. Teng. Competitive routing over time. *Theor. Computer Science*, 412(39):5420–5432, 2011.
21. R. Koch and M. Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory of Computing Systems*, 49(1):71–97, 2011.
22. S. Kontogiannis and P. G. Spirakis. Atomic selfish routing in networks: a survey. In WINE'05, LNCS 3828, p. 989–1002. Springer, 2005.
23. E. Koutsoupias and Ch. H. Papadimitriou. Worst-case equilibria. In STACS'99, LNCS 1563, p. 404–413. Springer, 1999.
24. E. Koutsoupias and K. Papakonstantinopoulou. Contention issues in congestion games. In ICALP'12, LNCS 7392, p. 623–635. Springer, 2012.
25. M. Mavronicolas and P. G. Spirakis. The price of selfish routing. *Algorithmica*, 48(1):91–126, 2007.
26. R. Paes Leme, V. Syrgkanis, and É. Tardos. The curse of simultaneity. In ITCS'12, p. 60–67. ACM Press, 2012.
27. A. C. Pigou. *The economics of welfare.* MacMillan and Co., 1920.
28. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
29. R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.
30. T. Roughgarden. Routing games. In *Algorithmic Game Theory*, chapter 18, p. 461–486. Cambridge University Press, 2007.
31. T. Roughgarden and É. Tardos. How bad is selfish routing? *J. of the ACM*, 49(2):236–259, 2002.
32. T. Roughgarden and É. Tardos. Bounding the inefficiency of equilibria in non-atomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.
33. S. Sadhukhan. *A Verification Viewpoint on Network Congestion Games.* PhD thesis, Université Rennes 1, France, 2021.
34. H. Sperber. How to find Nash equilibria with extreme total latency in network congestion games? *Mathematical Methods of Operations Research*, 71(2):245–265, 2010.
35. J. Valdes, R. E. Tarjan, and E. L. Lawler. The recognition of series-parallel digraphs. In STOC'79, p. 1–12. ACM Press, 1979.
36. J. G. Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1(3):325–362, 1952.