Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.

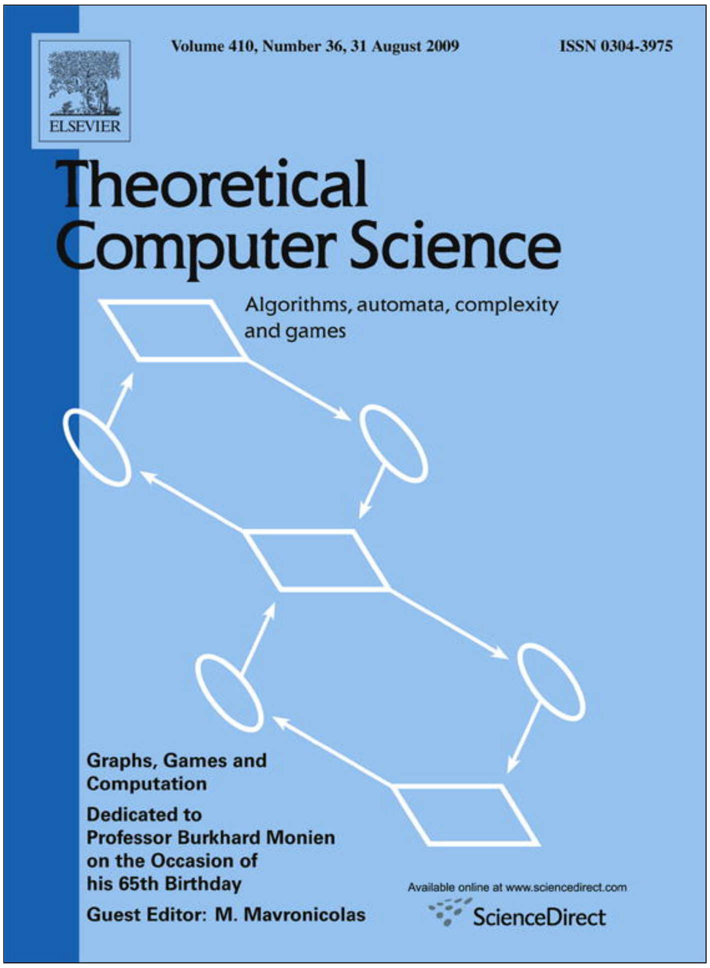

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Author's personal copy

Theoretical Computer Science 410 (2009) 3337–3347

Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

Atomic routing games on maximum congestion☆

Costas Busch [a,*], Malik Magdon-Ismail [b]

[a] *Computer Science Department, Louisiana State University, 286 Coates Hall, Baton Rouge, LA 70803, USA*

[b] *Computer Science Department, Rensselaer Polytechnic Institute, 110 8th Street, Troy, NY 12180, USA*

ARTICLE INFO

Keywords:
Algorithmic game theory
Congestion game
Routing game
Nash equilibrium
Price of anarchy

ABSTRACT

We study atomic routing games on networks in which players choose a path with the objective of minimizing the *maximum congestion* along the edges of their path. The social cost is the global maximum congestion over all edges in the network. We show that the *price of stability* is 1. The *price of anarchy*, *PoA*, is determined by topological properties of the network. In particular, $PoA = O(\ell + \log n)$, where ℓ is the length of the longest path in the player strategy sets, and n is the size of the network. Further, $\kappa - 1 \le PoA \le c(\kappa^2 + \log^2 n)$, where κ is the length of the longest cycle in the network, and c is a constant.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

A fundamental issue in the management of large scale communication networks is to route the packet traffic so as to optimize the network performance. Our measure of network performance is the worst bottleneck (most used link) in the system. We model network traffic as finite, unsplittable packets (atomic flow) [36,41], where each packet's path is controlled independently by a selfish player. The Nash equilibrium (NE) is a natural outcome for a game with selfish players — a stable state in which no player can unilaterally improve her situation. In the recent literature, the *price of anarchy* (*PoA*) [27,38] and the *price of stability* (*PoS*) [1,2] have become prevalent measures of the quality of the equilibria of uncoordinated selfish behavior relative to coordinated optimal behavior. The former quantifies the worst possible outcome with selfish agents, and the latter measures the minimum penalty in performance required to ensure a stable equilibrium outcome.

We study routing games with N players corresponding to N source–destination pairs of nodes on a network G; that is, every player corresponds to a source–destination pair. The strategy set available to each player is any set of edge-simple paths from the player's source to the destination. For example, the strategy set may consist of all edge-simple paths from the source to the destination in G, or alternatively, the strategy set could be any strict subset of all the possible paths. A *pure* strategy profile is a selection of a single path (strategy) by each player from her respective strategy set. We study pure Nash equilibria. In our context, a pure strategy profile corresponds to a *routing* $\mathbf{p}$, a collection of paths, one for each player. We refer to Nash equilibria in this context as *Nash-routings*. A routing $\mathbf{p}$ causes *congestion* in the network: the congestion C_e on an edge e is the number of paths in $\mathbf{p}$ that uses this edge; the congestion C_{p_i} of a path $p_i \in \mathbf{p}$ is the maximum congestion over all edges on the path; the congestion C of the network is the maximum congestion over all edges in the network. The *dilation* D is the maximum path length in $\mathbf{p}$.

Since a packet is to be delivered along each player's path, a natural choice for social cost is the maximum delay incurred by a packet. It is well known that the packets can be scheduled along the paths in $\mathbf{p}$ in a coordinated way so that the maximum delivery delay is close to $O(C + D)$ [7,12,28,30,37,39]. In heavily congested networks, $C \gg D$, and the maximum delay of

☆ A preliminary version of the paper appears in: Proceedings of the 2nd International Conference on Algorithmic Aspects in Information and Management, AAIM, Hong Kong, China, June 2006, pp. 79–91.

* Corresponding author. Tel.: +1 225 578 7510; fax: +1 225 578 1465.
E-mail addresses: busch@csc.lsu.edu (C. Busch), magdon@cs.rpi.edu (M. Magdon-Ismail).

0304-3975/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.tcs.2009.04.015

Author's personal copy

a packet is governed by the congestion C. Thus, the network congestion is an appropriate social cost — this choice for the social cost is often referred to as the *maximum social cost* [10,11,27,42]. Consider player i with path $p_i \in \mathbf{p}$. We choose C_{p_i} as the player cost, since it is shown in [7] that player i's packet can be delivered in time $O(C_{p_i} + |p_i| \log(C_{p_i} + |p_i|))$, where $|p_i|$ is the path length; this holds for all players simultaneously. This choice of player cost is typically referred to as the *maximum player cost*.

In the literature it is common to use the sum of the player cost (instead of the maximum) [14,15,26,27,33,42], where the cost is proportional to the sum of the congestions on the edges along a player's path. However, as discussed previously, in packet scheduling problems a critical parameter for the packet delivery time is the maximum congested edge along its path. Intuitively, when a packet is being delivered along its path and at some specific moment the packet waits on a particular edge to be cleared of other packets, at the same time, the remaining edges along its path may be cleared of their packets, and thus these edges will not have congestion to delay the packet further. For all these reasons, we have chosen to study the metrics of maximum player and social cost for the atomic routing games that we consider here.

1.1. Contributions

We give the first comprehensive analysis of routing games with maximum player and social cost. We study the quality of pure Nash-routings with respect to the price of stability and anarchy.

In our first result, we establish that there exist *optimal* Nash-routings where the social cost (congestion) is equal to the optimal coordinated cost; in other words, $PoS = 1$, where the price of stability expresses the ratio of the optimal social cost in the Nash-routing with the optimal coordinated cost. We also show that any *best response dynamic*, a sequence of best response moves of players, converges to a Nash-routing in a finite amount of time. Thus, we can easily obtain Nash-routings, starting from arbitrary initial routings.

Theorem 1.1. *For every routing game:*

(i) *There is a pure Nash-routing which is optimal ($PoS = 1$).*
(ii) *Every best response dynamic converges to a Nash-routing in finite time.*

We continue by examining the quality of the worst case Nash-routings. The price of anarchy, *PoA*, expresses the ratio of the social cost in the worst-case Nash-routing to the optimal coordinated cost. We bound the price of anarchy in terms of topological properties of the network. The next result bounds the price of anarchy for arbitrary instances of routing games in terms of the maximum path-lengths in the strategy sets:

Theorem 1.2. *For any routing game where the strategy sets of the players have paths with length at most ℓ, $PoA < 2(\ell + \log n)$.*

Theorem 1.2 gives good bounds for the price of anarchy for networks where it is natural to use paths with short length. For example in the Hypercube and Butterfly [29], if we choose bit-fixing paths, then $\ell = O(\log n)$, which implies that $PoA \le c \log n$, for some constant c.

Our next result characterizes the worst case Nash-routing in terms of the longest cycle of the network. For a graph G, the *edge-cycle number* $\kappa_e(G)$ is the length of the longest edge-simple cycle in G; we will drop the dependence on G when the context is clear.

Theorem 1.3. *For any undirected graph G with edge-cycle number κ_e,*

(i) *there exists a routing game for which $PoA \ge \kappa_e - 1$;*
(ii) *for any routing game, $PoA \le c({\kappa_e}^2 + \log^2 n)$, for some constant c.*

Let m denote the number of edges in the network. Since $\kappa_e \le m$, we have that $PoA \le c \cdot m^2$. In graphs with Euler cycles, $\kappa_e = m$. Therefore, Theorem 1.3 implies that $m - 1 \le PoA \le c \cdot m^2$ (we use c to represent a generic constant).

The lower bound of Theorem 1.3 (part (i)) is obtained by constructing a game instance where the players have their sources and destination on the largest cycle. To prove the upper bound of Theorem 1.3 (part (ii)), we use Theorem 1.2. We first examine 2-connected graphs and then general graphs. For 2-connected graphs, every pair of nodes has two edge-disjoint paths connecting them (Menger's theorem [35]), from which we establish that $\ell \le c \cdot {\kappa_e}^2$. The cycle upper bound for 2-connected graphs follows from Theorem 1.2.

If the graph G is not 2-connected, then the relation $\ell \le c \cdot {\kappa_e}^2$ may not hold. To obtain the result for a general graph G, we decompose G into a tree of 2-connected components. We show that if in G the Nash-routing has network congestion C, then there is some 2-connected component G' which has congestion close to C. At the same time the players in G' are in a *partial* Nash-routing, where many of them are locally-optimal. A generalization of Theorem 1.2 to partial Nash-routings, helps to establish the upper bound of Theorem 1.3.

1.2. Related work

General congestion games were introduced and studied in [36,41]. The application of game theory in computer science, specifically the notion of the price of anarchy was introduced in [27]. Since then, many models have been studied which can be categorized by: the topology of the network; the nature of the player and social costs; the nature of the traffic (atomic

Author's personal copy

	Atomic Flow	Splittable Flow
Pure	[3,6,10,31,41], [24,25,46]*, **Our Work**	[42–45], [4,34]*
Mixed	[3,14,13,15,20–22,17–19,26,27,32,33,38]*	[11], [4,16]*

	Max *SC*	Sum *SC*	Other *SC*	**
Max *pc*	**Our Work**, [6]	[34]*	–	[31]
Sum *pc*	[10,11,42], [4,14,13,15,16,20,26,27,33,38]*	[3,10,43–45], [22,17–19,24,25,46]*	[21,32]*	[31,41]

Fig. 1. A summary of the results on routing games. A vast majority of work has been on the *KP* model [27] consisting of m parallel links. The asterisked work ($[\cdot]^*$) indicates that the analysis is predominantly for some specific network model, (the vast majority of the work is on the parallel link network), or for a specific choice of player strategy sets (e.g. singleton sets). The double asterisked column ($**$) indicates that the work is on existence or convergence to equilibrium, as opposed to quality of equilibria.

or splittable); the nature of the strategy sets; the nature of the equilibria studied (pure or mixed). The work in [27] was the first to study the price of anarchy (*PoA* for parallel link networks (the *KP* model), obtaining a result of $\frac{3}{2}$. Since then, there has been an abundance of literature on the *KP* model, obtaining near optimal bounds on the price of anarchy and the price of stability. The tables in Fig. 1 serve this work. For instance, [13] gives tight bounds, $PoA = O(\frac{\log m}{\log\log m})$ (identical links), and $PoA = O(\frac{\log m}{\log\log\log m})$ (arbitrary links). The first general result for splittable flow with sum player and social costs came in [44], in which they showed the tight bound that $PoA \le \frac{4}{3}$. Our results are similar in spirit except that it considers the maximum player and social costs and our bounds depend on the topology of the network and are tight to within a square root factor. A brief taxonomy of some relevant existing results, according to the kind of flow (atomic or splittable) and equilibria (mixed or pure), and according to the social cost *SC* and player cost *pc* (sum or maximum), are shown in the following two tables in Fig. 1.

Typically, the research in the literature has focused on computing upper and lower bounds on the price of anarchy. The vast majority of the work on maximum social cost has been for parallel link networks, with only a few recent results on general topologies [10,11,42]. Essentially, all of the work has focused on the sum player cost, which corresponds to the sum of the edge-congestions on a path (as opposed to the maximum edge-congestion on the path, which we consider here). Further, the social cost is typically also the sum of congestions along the edges.

We have recently become aware of the work in [6] which is close to our work and was published independently from our work. In [6] the authors consider the maximum congestion metric in general networks with splittable and atomic flow where they examine the convergence to equilibria and their efficiency. They prove the existence and non-uniqueness of equilibria in both the splittable and atomic flow models. For the atomic case they show that best response dynamics always converge and the price of stability is 1 (we prove similar results here). They also prove that finding the best Nash equilibrium that minimizes the social cost is an NP-hard problem. Further, they show that the price of anarchy may be unbounded for specific *edge-congestion functions* (these are functions of the number of paths that use the edge). If the edge-congestion function is polynomial with degree p then they bound the price of anarchy with $O(m^p)$, where m is the number of edges in the graph. In our work we use $p = 1$ and our result in Theorem 1.2 can give a similar upper bound $O(m)$, since the maximum path length ℓ cannot exceed m. However, since ℓ may be smaller than m, our result can give tighter upper bounds on the price of anarchy. Further, our result in Theorem 1.3 can give tighter price of anarchy bounds too, since the square of the maximum edge-simple cycle may be smaller than m. In the splittable case it is shown in [6] that if the users always follow paths with low congestion then the equilibrium achieves optimal social cost.

Another result for general networks which has a brief discussion of the maximum player cost is [31] where the authors focus on parallel link networks, but also give some results for general topologies on convergence to equilibria. In [31], the main content is to establish the existence of pure Nash-routings.

Our work has been recently extended in [9]. That paper considers a model with routing classes $Q_1, \ldots, Q_\psi$, where each class j has a service cost S_j. Each path belongs to exactly one routing class. The player i cost function is $C_i + S_i$, where S_i is the cost of the service class of the path of i, and C_i is the congestion experienced by player i by considering only the paths in the same service class. The social cost function is the maximum player cost. It is shown in [9] that such games stabilize with best response dynamics and the price of stability is 1. The price of anarchy is bounded by $O(\min(C^*, S^*) \cdot \psi \log n)$, where $C^* + S^*$ is the optimal social cost. Such games can be used to provide approximations to the social cost function $C + D$ (recall that D is the path dilation).

Outline of paper

In Section 2 we give some basic definitions. We prove Theorem 1.1 in Section 3. We continue with the proof of Theorem 1.2 in Section 4. The lower bound of Theorem 1.3 is proven in Section 5. In the same section we prove the upper bound of Theorem 1.3 for 2-connected graphs. We give the general version of the upper bound in Section 6. We conclude with a discussion in Section 7.

Author's personal copy

2. Definitions

An instance $\mathcal{R}$ of a *routing (congestion) game* is a tuple $(\mathbf{N}, G, \{\mathcal{P}_i\}_{i\in\mathbf{N}})$, where:

- $\mathbf{N} = \{1, 2, \ldots, N\}$ are the players,
- $G = (V, E)$ is an undirected connected graph with $|V| = n$,
- $\mathcal{P}_i$ is a collection of *edge-simple* paths for player i (where in an edge-simple path edges do not repeat but nodes may repeat). Each path in $\mathcal{P}_i$ is a path in G that has the same source $s_i \in V$ and destination $t_i \in V$. Note that $\mathcal{P}_i$ could be any set of edge-simple paths between s_i and t_i. (For example $\mathcal{P}_i$ may consist of all the edge-simple paths from s_i to t_i, or it may consist of any subset of these edge-simple paths.)

Each path in $\mathcal{P}_i$ is a pure strategy available to player i. A pure strategy profile $\mathbf{p} = [p_1, p_2, \cdots, p_N]$ is a collection of pure strategies (paths), one for each player, where $p_i \in \mathcal{P}_i$. We refer to a pure strategy profile as a *routing*. On a finite network, a routing game is necessarily a finite game.

For any routing $\mathbf{p}$ and any edge $e \in E$, the *edge-congestion* $C_e(\mathbf{p})$ is the number of paths in $\mathbf{p}$ that use edge e. For any path p, the *path-congestion* $C_p(\mathbf{p})$ is the maximum edge-congestion over all edges in p, $C_p(\mathbf{p}) = \max_{e\in p} C_e(\mathbf{p})$. The *network congestion* is the maximum edge-congestion over all edges in E, that is, $C(\mathbf{p}) = \max_{e\in E} C_e(\mathbf{p})$. The *social or global cost* is the network congestion $C(\mathbf{p})$. The *player or local cost* $pc_i(\mathbf{p})$ for player i is her path-congestion, $pc_i(\mathbf{p}) = C_{p_i}(\mathbf{p})$. When the context is clear, we will drop the dependence on $\mathbf{p}$ and use C_e, C_p, C, SC, pc_i.

We use the standard notation $\mathbf{p}_{-i}$ to refer to the collection of paths $\{p_1, \cdots, p_{i-1}, p_{i+1}, \cdots, p_N\}$, and $(p_i; \mathbf{p}_{-i})$ as an alternative notation for $\mathbf{p}$ which emphasizes the dependence on p_i. Player i is *locally-optimal* in routing $\mathbf{p}$ if $pc_i(\mathbf{p}) \le pc_i(p'_i; \mathbf{p}_{-i})$ for all paths $p'_i \in \mathcal{P}_i$; that is, the player does not wish to change its current path, since every alternative path does not yield a lower cost. A routing $\mathbf{p}$ is in a Nash Equilibrium ($\mathbf{p}$ is a *Nash-routing*) if every player is locally-optimal. Nash-routings quantify the notion of a stable selfish outcome.

A routing $\mathbf{p}^*$ is an optimal pure strategy profile if it has minimum attainable social cost: for any other pure strategy profile $\mathbf{p}$, $C(\mathbf{p}^*) \le C(\mathbf{p})$. We quantify the quality and diversity of the Nash-routings by the *price of stability* (*PoS*) and the *price of anarchy* (*PoA*) (sometimes referred to as the coordination ratio). Let $\mathcal{T}$ denote the set of distinct Nash-routings, and let C^* denote the social cost of an optimal routing $\mathbf{p}^*$ (namely, $C^* = C(\mathbf{p}^*)$). Then,

$$PoS = \inf_{\mathbf{p}\in\mathcal{T}} \frac{C(\mathbf{p})}{C^*}, \qquad PoA = \sup_{\mathbf{p}\in\mathcal{T}} \frac{C(\mathbf{p})}{C^*}.$$

3. Existence of optimal Nash-routings

The goal in this section is to establish Theorem 1.1. For routing $\mathbf{p}$, the *congestion vector* $\mathbf{C}(\mathbf{p}) = [m_0(\mathbf{p}), m_1(\mathbf{p}), m_2(\mathbf{p}), \ldots]$, where each component $m_k(\mathbf{p})$ is the number of edges with congestion k. Note that $\sum_k m_k(\mathbf{p}) = m$, where m is the number of edges in the network. The social cost (network congestion) $SC(\mathbf{p})$ is the maximum k for which $m_k > 0$. We define a lexicographic total order on routings as follows. Let $\mathbf{p}$ and $\mathbf{p}'$ be two routings, with $\mathbf{C}(\mathbf{p}) = [m_0, m_1, m_2, \ldots]$, and $\mathbf{C}(\mathbf{p}') = [m'_0, m'_1, m'_2 \ldots]$. Two routings are equal, written $\mathbf{p} =_c \mathbf{p}'$, if and only if $m_k = m'_k$ for all $k \ge 0$; $\mathbf{p} <_c \mathbf{p}'$ if and only if there is some k^* such that $m_{k^*} < m'_{k^*}$ and $\forall k > k^*,\ m_k = m'_k$.

Let $(\mathbf{N}, G, \{\mathcal{P}_i\}_{i\in\mathbf{N}})$ be an instance of a routing game. Since there are only finitely many routings (as a player's path may use any edge at most once), there exists at least one minimum routing w.r.t. the total order $<_c$. There may be many distinct routings all of which are minimum (and have equal congestion vectors). Let $\mathbf{p}^*$ be a minimum routing (which exists); then, for all routings $\mathbf{p}$, $\mathbf{p}^* \le_c \mathbf{p}$. Every minimum routing is optimal; indeed, if $SC(\mathbf{p}) < SC(\mathbf{p}^*)$ for some other routing $\mathbf{p}$, then the maximum index k for which $m_k(\mathbf{p}) > 0$ is smaller than the corresponding index k for $\mathbf{p}^*$, contradicting the fact that $\mathbf{p}^* \le_c \mathbf{p}$.

Fact 3.1. *Every minimum routing (at least one exists) is optimal.*

A greedy move is available to player i if she can obtain a lower path-congestion by changing her current path from p_i to p'_i – the greedy move takes the original routing $\mathbf{p} = (p_i; \mathbf{p}_{-i})$ to $\mathbf{p}' = (p'_i; \mathbf{p}_{-i})$ in which p_i is replaced by p'_i.

Lemma 3.2. *If a greedy move by any player takes* $\mathbf{p}$ *to* $\mathbf{p}'$*, then* $\mathbf{p}' <_c \mathbf{p}$*.*

Proof. Suppose that a greedy move by player i takes $\mathbf{p}$ to $\mathbf{p}'$, and so $C_{p'_i}(\mathbf{p}') < C_{p_i}(\mathbf{p})$. Let $k = C_{p_i}(\mathbf{p})$. Since only player i has changed his path, the only edges with higher congestion in $\mathbf{p}'$ than in $\mathbf{p}$ are edges on the path p'_i. Some edges on p_i decreased in congestion by 1 as a result of the greedy move. In particular, all edges of congestion k on p_i have decreased in congestion by 1, since all edges on p'_i have final congestion less than k. Thus, $m_k(\mathbf{p}') \le m_k(\mathbf{p}) - 1$, since at least one edge of congestion k dropped in congestion and no new edges reached congestion k. To conclude that $\mathbf{p}' <_c \mathbf{p}$, we note that no edge with congestion greater than k has been affected by the greedy move, hence $m_j(\mathbf{p}') = m_j(\mathbf{p})$ for all $j > k$. ∎

Thus, a greedy move decreases the number of high congestion edges, by transferring the congestion to lower congestion edges. Since there are only a finite number of routings, every best response dynamic is finite. By Lemma 3.2, no player can have an available greedy move at a minimum routing, as this would contradict the minimality of the routing. Hence,

Lemma 3.3. *Every minimum routing is an optimal Nash-routing.*

Hence, $PoS = 1$. Theorem 1.1 now follows from Lemmas 3.2 and 3.3.

Author's personal copy

4. Path length bound on price of anarchy

Here, we prove Theorem 1.2. In order to do so we will use the *edge-expansion process*, that we introduce here. Before we describe this technique we need to give some necessary definitions.

Let $\mathcal{R} = (\mathbf{N}, G, \{\mathcal{P}_i\}_{i\in\mathbf{N}})$ be an instance of a routing game. Let $\mathcal{P} = \bigcup_{i\in\mathbf{N}} \mathcal{P}_i$. The *path-length* of $\mathcal{R}$ is $\ell = \max_{p\in\mathcal{P}} |p|$. A *path-cut* for player i is a set of edges E_i such that every path in $\mathcal{P}_i$ must use at least one of the edges in E_i. The congestion of a path-cut, denoted $W(E_i)$ is the minimum congestion of any edge in E_i, $W(E_i) = \min_{e\in E_i} C_e$. We have:

Lemma 4.1. *Let $\mathbf{p} = [p_1, p_2, \cdots, p_N]$ be a routing for which player i is locally-optimal. Then, there is a path-cut E_i for player i with congestion $W(E_i) \geq pc_i - 1$.*

Proof. Since player i is locally-optimal, every path in $\mathcal{P}_i$ must have path-congestion at least $pc_i - 1$. Indeed, if not, then there is a path $p_i' \in \mathcal{P}_i$ with path-congestion at most $pc_i - 2$. If player i switches from p_i to p_i', his cost is at most $pc_i - 1$, which contradicts p_i being locally-optimal for i. For every path $p \in \mathcal{P}_i$, let edge $e(p) \in p$ be an edge with maximum congestion on p $(C(e) \geq pc_i - 1)$. Let $E_i = \cup_{p\in\mathcal{P}_i} e(p)$. Since E_i contains at least one edge from every path in $\mathcal{P}_i$, it is a path-cut for player i and every edge in E_i has congestion at least $pc_i - 1$. Thus, $W(E_i) \geq pc_i - 1$. ■

4.1. Edge-expansion process

If only some players are locally-optimal in a routing $\mathbf{p}$, then we say that $\mathbf{p}$ is a *partial* Nash-routing. Note that a Nash-routing is a special case of a partial Nash-routing. The edge-expansion process that we describe below applies to any partial Nash-routing.

Suppose routing $\mathbf{p}$ has network congestion C, and suppose that at least one player is locally-optimal with player cost C. Let $\mathcal{E}_0$ be the set of edges with congestion $C_0 = C$ that are used by at least one locally-optimal player, and let Π_0 be the set of these locally-optimal players that use at least one edge in $\mathcal{E}_0$. By Lemma 4.1, each player in Π_0 has a path-cut with congestion at least $C_0 - 1$. Let $\mathcal{E}_1$ denote the union of $\mathcal{E}_0$ with all these path-cuts of every player in Π_0. Thus, $\mathcal{E}_0 \subseteq \mathcal{E}_1$ and each edge in $\mathcal{E}_1$ has congestion at least $C_1 = C_0 - 1$. Let Π_1 denote the set of locally-optimal players whose paths in $\mathbf{p}$ use at least one edge in $\mathcal{E}_1$. Note that $\Pi_0 \subseteq \Pi_1$. Each player in Π_1 has player cost at least C_1, since every edge in $\mathcal{E}_1$ has congestion at least C_1.

We repeat this process as follows. Suppose that for $i \geq 1$, edge set $\mathcal{E}_i$ has been constructed as the union of $\mathcal{E}_{i-1}$ with path-cuts for the players in Π_{i-1}, thus every edge in $\mathcal{E}_i$ has congestion at least $C_i = C_{i-1} - 1 = C - i$. We now construct Π_i, the set of locally-optimal players whose paths use at least one edge in $\mathcal{E}_i$; every player in Π_i has player cost at least C_i. By Lemma 4.1, each player in Π_i has a path-cut with congestion $C_i - 1$, and we construct $\mathcal{E}_{i+1}$ to be the union of $\mathcal{E}_i$ with all these path-cuts of the players in Π_i.

Using this inductive-like construction, we obtain a sequence of edge sets, $\mathcal{E}_0 \subseteq \mathcal{E}_1 \subseteq \mathcal{E}_2, \ldots$, with $W(\mathcal{E}_j) \geq C_j = C - j$, and corresponding to each edge set, a set of locally-optimal players $\Pi_0 \subseteq \Pi_1 \subseteq \Pi_2 \cdots$. We continue this inductive-like construction up to edge set $\mathcal{E}_s$ which is the first set for which $|\mathcal{E}_s| \leq 2|\mathcal{E}_{s-1}|$. We will refer to this process as the *edge-expansion process*.

4.2. Edge-expansion properties

Since $|\mathcal{E}_i| \leq \frac{1}{2}n^2$ and each expansion at least doubles the size of the edge set,

Lemma 4.2. $|\mathcal{E}_s| \geq 2^{s-1}$ *and* $1 \leq s < 2\log n$.

Proof. If $s = 1$ there is nothing to prove, so assume that $s > 1$. Since $|\mathcal{E}_k| > 2|\mathcal{E}_{k-1}|$ for $k = 1, \ldots, s-1$, $|\mathcal{E}_k| > 2^k|\mathcal{E}_0|$. Since $|\mathcal{E}_0| \geq 1$, we have $|\mathcal{E}_k| > 2^k$. By construction, $\mathcal{E}_{i-1} \subseteq \mathcal{E}_i$, for $1 \leq i \leq s$; thus, $|\mathcal{E}_s| \geq |\mathcal{E}_{s-1}| \geq 2^{s-1}$. Since $|\mathcal{E}_s| \leq |E| < \frac{1}{2}n^2$, $2^{s-1} < \frac{1}{2}n^2$ implying $s < 2\log n$. ■

In routing $\mathbf{p}$, let $F(C') \subseteq \mathbf{N}$ denote the set of non-locally-optimal players with player cost at least C'. We now establish a relationship between the congestion of a partial Nash-routing and the optimal routing.

Lemma 4.3. $C < 2\ell \cdot \left(C^* + |F(C - 2\log n)|\right) + 2\log n$.

Proof. From the edge-expansion process, each edge in $\mathcal{E}_{s-1}$ has congestion at least C_{s-1}. Let M be the number of times edges in $\mathcal{E}_{s-1}$ are used by the paths in $\mathbf{p}$. Then, $M > C_{s-1} \cdot |\mathcal{E}_{s-1}|$. By construction, in $\mathbf{p}$, the congestion in each of the edges of $\mathcal{E}_{s-1}$ is caused only by the players in $A = \Pi_{s-1} \cup B$, where $B \subseteq F(C_{s-1})$ contains the non-locally-optimal players that use edges in $\mathcal{E}_{s-1}$. Since path lengths are at most ℓ, each player in A can use at most ℓ edges in $\mathcal{E}_{s-1}$. Hence, $C_{s-1} \cdot |\mathcal{E}_{s-1}| < M \leq \ell \cdot |A|$. Since, $|A| \leq |\Pi_{s-1}| + |F(C_{s-1})|$, we obtain, $C_{s-1} < \frac{\ell}{|\mathcal{E}_{s-1}|} \cdot (|\Pi_{s-1}| + |F(C_{s-1})|)$. We now bound $|\Pi_{s-1}|$.

$\mathcal{E}_s$ contains a path-cut for every player in Π_{s-1}, and every such players must use at least one edge in $\mathcal{E}_s$ in any routing, including the optimal routing $\mathbf{p}^*$. Thus, edges in $\mathcal{E}_s$ are used at least $|\Pi_{s-1}|$ times, hence some edge is used at least $|\Pi_{s-1}|/|\mathcal{E}_s|$ times, by the pigeonhole principle. Hence, $C^* \geq |\Pi_{s-1}|/|\mathcal{E}_s|$ (note that $|\mathcal{E}_s| > 0$). By the definition of s, $|\mathcal{E}_s| \leq 2|\mathcal{E}_{s-1}|$. Hence, $|\Pi_{s-1}| \leq 2|\mathcal{E}_{s-1}|C^*$, and $C_{s-1} < 2\ell \cdot \left(C^* + \frac{|F(C_{s-1})|}{2|\mathcal{E}_{s-1}|}\right)$. Since $C_{s-1} = C - (s-1)$ and $2|\mathcal{E}_{s-1}| \geq 2^s$ (Lemma 4.2), we obtain

Author's personal copy

$C < 2\ell \cdot \left(C^* + \frac{|F(C-s+1)|}{2^s}\right) + s - 1$. To conclude, $2^s \geq 2$, and note that $C'' < C'$ implies $F(C') \subseteq F(C'')$, hence $|F(C')|$ is non-increasing in C'. Thus, $|F(C-s+1)| \leq |F(C - 2\log n)|$. ■

Since in a Nash-routing, $|F(C')| = 0,\ \forall C' > 0$, by dividing the result of Lemma 4.3 with C^*, we obtain Theorem 1.2.

5. Basic cycle bounds on price of anarchy

Here, we first give the lower bound (part (i)) of Theorem 1.3 for the price of anarchy; we then prove the upper bound (part (ii)) of Theorem 1.3, for the special case of 2-connected graphs. The next result establishes the lower bound of Theorem 1.3.

Lemma 5.1. *For any graph G, there is a routing game with $PoA \geq \kappa_e(G) - 1$.*

Proof. Let $Q = e_1, \ldots, e_{\kappa_e}$ be an edge-simple cycle in G with length κ_e. We construct a routing game with κ_e players, where player i corresponds to edge $e_i = (u_i, v_i)$ in Q, that is, the source of i is $s_i = u_i$ and the destination $t_i = v_i$. The strategy set of i is the collection of all edge-simple paths from s_i to t_i.

There are two special paths in the strategy set of player i, the *forward path* which is composed solely of the edge (u_i, v_i), and the *backward path* which consists of the remaining edges of cycle Q. Since Q is edge simple, if every player uses his forward path $C = 1$. Thus, the optimal social cost is 1. If on the other hand, all the players use their backward paths (backward routing $\bar{\mathbf{p}}$), then player i uses every edge in Q except e_i exactly once. Thus, the congestion on every edge in Q is $N - 1 = \kappa_e - 1$. Hence, if $\bar{\mathbf{p}}$ is a Nash-routing, then $PoA \geq \kappa_e - 1$.

We will show that $\bar{\mathbf{p}}$ is a Nash-routing by contradiction. Suppose that some player k is not locally-optimal — so player k has lower congestion for some other path p. Since every edge on Q has congestion $\kappa_e - 1$ in routing $\bar{\mathbf{p}}$, at least $\kappa_e - 2$ players other than player k use every edge on Q. Thus, if p uses any edge on Q, then $pc_k(p; \bar{\mathbf{p}}_{-k}) = \kappa_e - 1$, which does not improve its cost, so we conclude that p does not use any edge on Q. Therefore, p has length at least 2 (since $p \neq e_k$ and G is not a multi-graph). Thus, replacing $e_k \in Q$ by p results in a new edge-simple cycle Q' that is strictly longer than Q, a contradiction. Thus, $\bar{\mathbf{p}}$ is a Nash-routing. ■

We now continue with the upper bound on the price of anarchy. An *edge-cut* of a connected graph G is a set of edges whose removal from G partition the graph into at least two node-disjoint connected components. A graph G is *k-connected* if it is a connected graph whose minimum edge-cut has size at least k. By Menger's theorem ([35]), G is k-connected if and only if there are at least k edge-disjoint paths between every two nodes. Let L be the longest edge-simple path length in G.

Lemma 5.2. *If G is 2-connected, then $\kappa_e(G) \geq \sqrt{2L} - \frac{3}{2}$.*

Proof. The proof relies on the observation that the longest path p must have at least $\sqrt{L}$ edges in common with the largest cycle q, since otherwise, we would be able to construct a larger cycle by combing pieces of p and q.

Let u and v be the respective starting and ending nodes in the longest path in G. Since the min (u, v)-cut has size at least two, by Menger's theorem ([35]), there is a pair of edge-disjoint paths p_1, p_2 from u to v; let $l_1 \leq l_2$ be the lengths of these paths respectively. Let p be an edge-simple (u, v)-path with length L. Path p can be decomposed into $2z + 1$ path segments, for some z, as follows,

$$p = \lambda_0 \mu_1 \lambda_1 \mu_2 \lambda_2 \cdots \mu_z \lambda_z,$$

where each μ_i has length at least one and consists only of edges on p_1, and each "excursion" λ_i does not contain any edges on p_1. Since each excursion λ_i connects two (not necessarily distinct) nodes on p_1, it follows that there is an edge-simple cycle composed of λ_i together with the segment of p_1 between these two nodes. The length of this cycle is at least $|\lambda_i|$, so we have that $\kappa_e(G) \geq |\lambda_i|$ for all $i \in [0, z]$. Since path p is edge simple, there can be at most $l_1 + 1$ excursions (as each of the κ_i must contain distinct edges), i.e. $z \leq l_1$. We now compute the length of p as follows,

$$\begin{aligned} L &= \sum_{i=0}^{z} |\lambda_i| + \sum_{i=1}^{z} |\mu_i|, \\ &\leq \sum_{i=0}^{z} \kappa_e(G) + \sum_{i=1}^{z} |\mu_i|, \\ &\leq \kappa_e(G) \cdot (l_1 + 1) + l_1. \end{aligned}$$

Solving for $\kappa_e(G)$, we have that $\kappa_e(G) \geq (L+1)/(l_1+1) - 1$. Since p_1 and p_2 form an edge-simple cycle, $\kappa_e(G) \geq l_1 + l_2 \geq 2l_1$. Combining these inequalities, we have

$$\kappa_e(G) \geq \max\left\{2l_1, \frac{L+1}{l_1+1} - 1\right\}.$$

Author's personal copy

To reach a contradiction, suppose that $\kappa_e(G) < \sqrt{2L} - \frac{3}{2}$. Since $\kappa_e(G) \geq 2l_1$, we have $l_1 < \sqrt{L/2} - \frac{3}{4}$. Therefore,

$$\begin{aligned}\kappa_e(G) &\geq \frac{L+1}{l_1+1} - 1,\\ &> \frac{L+1}{\sqrt{L/2}+\frac{1}{4}} - 1,\\ &= \sqrt{2L} - \tfrac{3}{2} + \epsilon,\end{aligned}$$

where $\epsilon = 9/(\sqrt{32L}+2) > 0$. This contradiction concludes the proof. ■

Lemma 5.2 bounds the longest edge-simple path length in G with respect to $\kappa_e(G)$. Theorem 1.2 bounds the price of anarchy in terms of the longest path ℓ in the players' strategy sets. Since $\ell \leq L$, we obtain the following result, which proves the upper bound of Theorem 1.3 for 2-connected graphs:

Lemma 5.3. *For any routing game on a 2-connected graph G, PoA* $\leq c({\kappa_e}^2(G) + \log n)$, *for some constant c.*

6. Cycle upper bound for general graphs

We now prove the upper bound (part (ii)) of Theorem 1.3 for general graphs. We will bound the price of anarchy with respect to the square of the longest cycle. The main idea behind the result is that any Nash-routing in G can be mapped to a partial Nash-routing on some 2-connected subgraph of G. In this partial Nash-routing, many players are locally-optimal, and we can apply Lemma 4.3 in combination with Lemma 5.1 to obtain the result.

6.1. Block-structure of graphs

Consider an arbitrary connected graph $G = (V, E)$. We say that two subgraphs of G are *adjacent* if the intersection of their node sets is non-empty. It is easy to verify that G contains a 2-connected subgraph if and only if it is not a tree. In a 2-connected subgraph every pair of nodes is in some edge-simple cycle (this is a trivial property of 2-connected graphs). A 2-connected subgraph G' is *maximal* if there is no larger 2-connected subgraph that contains G'. A maximal 2-connected subgraph is also known as a *block*. A *bridge* is an edge whose removal disconnects G. It is well known how to decompose any connected graph G into a tree of blocks and bridges known as the *block-structure* of G. In particular, we denote by $A = (V_A, E_A)$ the subgraph of G consisting of all the blocks, while we denote by $B = (V_B, E_B)$ the subgraph of G consisting of all the bridges.

Let $A_1, \ldots, A_\alpha$ be the blocks of A, where $\alpha \geq 1$, and $A_i = (V_{A_i}, E_{A_i})$. Clearly, any two subgraphs A_i and A_j, $i \neq j$, are not adjacent since otherwise their union would be 2-connected, which contradicts their maximality. Subgraph B consists of one or more disjoint maximal connected components (each containing at least two nodes), which we will denote $B_1, \ldots, B_\beta$. We will refer to the A_i as the *type-a blocks* of G and to the B_i as the *type-b blocks* of G (here we abuse notation since the B_i are not really blocks). Clearly, the type-b blocks are trees. Note that only blocks of different types can be adjacent, since both the type-a and type-b blocks have a maximality property. Moreover, any pair of type-a and type-b blocks can have at most one common node, since otherwise we would be able to construct edge-simple cycles involving nodes in type-a and type-b blocks, which would contradict the maximality of the type-a blocks.

We now define a simple bipartite graph $H = (V_H, E_H)$ that represents the block-structure of G. In $V_H = \{a_1, \ldots, a_\alpha, b_1, \ldots, b_\beta\}$ the nodes a_i, b_j correspond to the type-a block A_i and the type-b block B_j, respectively. Edge $(a_i, b_j) \in E_H$ if and only if the blocks A_i and B_j are adjacent. The bipartition for H is $(\mathcal{A}, \mathcal{B})$, where $\mathcal{A} = \{a_1, \ldots, a_\alpha\}$ and $\mathcal{B} = \{b_1, \ldots, b_\beta\}$. The nodes in H inherit the same type as their corresponding blocks in G. Since G is connected, it follows immediately that H is connected too. Furthermore, H is a tree and we will refer to it as the *block-tree* of G.

6.1.1. Block-subpaths

A node in G can belong to at most one type-a block and one type-b block (possibly simultaneously), since no two blocks of the same type are adjacent. If a node is a member of one block, then we define its type to be the same with the type of the block. If a node belongs to two blocks then we define it to be of type-a (we assign it to the type-a block). An edge belongs to exactly one block and inherits the type of that block.

Let $p = v_1, v_2, \ldots, v_k$, $k > 1$, be an edge-simple path in G. We can write p as a concatenation of subpaths $p = q_1 q_2 \cdots q_r$, where $r \geq 1$, and $|q_i| > 0, \forall i \in \{1, \ldots, r\}$, with the following properties: (i) the subpaths are edge disjoint; (ii) all the nodes of a subpath q_i are in the same block and have the same type (which will also be the type and block of q_i); (iii) the types of the subpaths alternate, i.e. the types of q_i and q_{i+1} are different; (iv) there is no type-a subpath with one node (any type-a subpath with one node can be merged with two adjacent type-b subpaths in the same type-b block). We refer to the q_i as the *block-subpaths* of p. Note that there is a unique block-subpath decomposition for path p.

Since graph H is a tree, and the type-b blocks are also trees, it can be shown that any two edge-simple paths from the same source node to the same destination nodes have to go through the same sequence of type-b edges. Thus, we have:

Lemma 6.1. *Any two edge-simple paths from nodes s to t in G use the same sequence of type-b edges.*

Author's personal copy

6.2. Subgames in blocks

Consider a routing game $\mathcal{R} = (\mathbf{N}, G, \{\mathcal{P}_i\}_{i\in\mathbf{N}})$ in G. Let $\mathbf{p}$ be a routing with network congestion C. Let $\mathbf{p}^*$ denote an optimal routing for $\mathcal{R}$ with congestion C^*. An immediate consequence of Lemma 6.1 is that every path in $\mathbf{p}$ uses the same type-b edges as its corresponding path in $\mathbf{p}^*$, hence:

Lemma 6.2. *Any type-b edge e has the same congestion in* $\mathbf{p}$ *and* $\mathbf{p}^*$*, i.e.* $C_e(\mathbf{p}) = C_e(\mathbf{p}^*) \le C^*$.

By Lemma 6.2, all the edges in $\mathbf{p}$ with congestion higher than C^* must occur in type-a subpaths.

Lemma 6.3. *For path p, if* $C_p(\mathbf{p}) > C^*$*, then p must have a type-a subpath q with* $C_q(\mathbf{p}) = C_p(\mathbf{p})$.

Suppose now that $\mathbf{p}$ is an arbitrary Nash-routing which has network congestion C. For a type-a block Λ, let $\mathbf{p}_\Lambda = \{p_1, \ldots, p_\gamma\}$ denote the paths in $\mathbf{p}$ that use edges in Λ, and denote the respective users as $\mathbf{N}_\Lambda$, where $|\mathbf{N}_\Lambda| = \gamma$. Let $Q_\Lambda = \{q_1, \ldots, q_\gamma\}$ denote the type-a block-subpaths of the paths in $\mathbf{p}_\Lambda$ that are in Λ (q_i is a subpath of p_i).

In block Λ, we define a new routing game $\mathcal{R}_\Lambda = (\mathbf{N}_\Lambda, \Lambda, \{\mathcal{P}_i^\Lambda\}_{i\in N_\Lambda})$, where $\mathcal{P}_i^\Lambda$ contains all the type-a subpaths of $\mathcal{P}_i$ that are in Λ and have the same source and destination as q_i. We refer to $\mathcal{R}_\Lambda$ as the *subgame* of $\mathcal{R}$ for block Λ. Q_Λ is a possible routing for $\mathcal{R}_\Lambda$. If q_i is locally-optimal for player i in Λ, we say that its corresponding path p_i in G is *satisfied* in subgame $\mathcal{R}_\Lambda$. In other words, if path p_i is satisfied in $\mathcal{R}_\Lambda$, player i does not wish to change the choice q_i in Λ. Every player with high player cost (higher than C^*) must be satisfied in a type-a block, since otherwise it would violate Lemma 6.3. Thus:

Lemma 6.4. *If player i has path* p_i *and* $pc_i > C^*$*, then player i is satisfied in some subgame* $\mathcal{R}_\Lambda$ *in a type-a block* Λ*, and player i has congestion* pc_i *in* Λ.

Proof. We show that if the claim is false, then path p_i is not locally-optimal for player i in $\mathcal{R}$, contradicting the fact that $\mathbf{p}$ is a Nash-routing. Indeed, we know from Lemma 6.3 that p_i uses type-a blocks. If none of these type-a block-subpaths are not locally-optimal for their respective subgames, then they can all be switched in favor of paths with strictly lower congestion than C. This will give a valid path for player i with strictly lower congestion than C, hence p_i is not locally-optimal for player i. ■

6.3. Main result

Consider routing game $\mathcal{R} = (\mathbf{N}, G, \{\mathcal{P}_i\}_{i\in\mathbf{N}})$ in G and a Nash-routing $\mathbf{p}$ with congestion $C(\mathbf{p}) = C$. Lemma 6.4, implies that each user is satisfied in some type-a block (not necessarily the same). In any type-a block, the resulting routing in the subgame may be a partial Nash-routing, since some users may not be satisfied in it. We first show that there is a block with high congestion where the number of unsatisfied players is bounded. For a type-a block Λ, let $F_\Lambda(C')$ denote the set of non-locally-optimal players in the subgame $\mathcal{R}_\Lambda$ whose congestion in $\mathcal{R}$ is at least C'. We will use C_Λ to denote the congestion in the block Λ. We have:

Lemma 6.5. *Suppose that* $C > C^* + x(1+\log n)$ *for some* $x > 0$*. Then, there is a type-a block* Λ *with congestion* $C_\Lambda \ge C - x\log n$ *and* $|F_\Lambda(C_\Lambda - x)| \le 2C^*$.

Proof. Let $f_\Lambda = |F_\Lambda(C_\Lambda - x)|$, and suppose that $C > C^* + x(1 + \log n)$. Assume, that every type-a block Λ with congestion $C_\Lambda \ge C - x\log n$ has $f_\Lambda > 2C^*$. Since $\mathbf{p}$ is a Nash-routing, every player with congestion $C > C^*$ is locally-optimal in at least one type-a subgame of $\mathcal{R}$ (Lemma 6.4). Thus, there is at least one type-a block Λ_1 with $C_{\Lambda_1} = C$. We will now root H (the block-tree of G) at the type-a node a_1 which corresponds to Λ_1 and define a type-a tree H_a composed only of the type-a nodes in H. The root of H_a is also Λ_1. By assumption, $f_{\Lambda_1} > 2C^*$. Since $C - x > C^* + x\log n$, these f_{Λ_1} players which are not locally-optimal in subgame $\mathcal{R}_{\Lambda_1}$ have congestion at least $C - x$ and are locally-optimal in some other subgame. Therefore, their paths leave Λ_1 and enter some other type-a block. We introduce the following auxiliary claim:

Claim 6.6. *If K paths leave a type-a block* Λ*, they must use at least* $\lceil K/C^* \rceil$ *distinct edges out of* Λ.

Proof. If not, then one of the exit edges (which is a type-b edge) will have congestion greater than C^*, contradicting Lemma 6.2. □

We now build the rooted tree H_a in an inductive-like way as follows. The root node is a_1. Suppose that α is a node in H_a corresponding to type-a block Λ, with the following two properties:

(i) $C_\Lambda - x > C^*$;
(ii) $f_\Lambda = |F_\Lambda(C_\Lambda - x)| > 2C^*$.

Then, we define three potential children for α as follows. Since there are $f_\Lambda > 2C^*$ players with congestion at least $C_\Lambda - x > C^*$ which are not locally-optimal in subgame $\mathcal{R}_\Lambda$, these f_Λ players must be locally-optimal in some other subgame. Therefore all these f_Λ paths leave Λ and proceed to their respective subgames where they are locally-optimal with congestion at least $C_\Lambda - x$. By Claim 6.6, they use at least three distinct type-b edges e_1, e_2, e_3 in leaving Λ (note that these three edges may be in the same type-b block, but this will not affect the argument). Let p_1, p_2, p_3 be three paths with congestion at least

Author's personal copy

$C_\Lambda - x$ that exit Λ on the edges e_1, e_2, e_3 respectively and continue on to their respective blocks $\Lambda_1, \Lambda_2, \Lambda_3$ in which they are locally-optimal. At least two of these blocks correspond to nodes that are not the parent (if it exists) of α in H_a; these two nodes are two children $c_1(\alpha)$ and $c_2(\alpha)$ of α in H_a (if more than two of these children are different from the parent, we arbitrarily pick two). The depth of a child is one greater than the depth of its parent (the depth of the root is 0). The next few claims give some properties of H_a that will be needed to complete the proof of the lemma.

Claim 6.7. *H_a is a tree.*

Proof. H_a is connected, by construction. Suppose that H_a contains a node-simple cycle. By construction, an edge between nodes α_1, α_2 in H_a implies the existence of an edge-simple path which leaves one type-a block and enters the second. Hence, there is a path that leaves a type-a block (say A_i) and re-enters it. This path can be made edge simple by removing all cycles. Thus, block A_i is not maximally 2-connected, since it can be augmented with nodes on the path which are not members of A_i, a contradiction. □

The nodes in H_a can be viewed as constructed level by level. Each node in H_a that satisfies the two conditions above has exactly two children. Note that a_1 satisfies these two conditions, initiating the construction of H_a. The nodes in H_a which do not satisfy the conditions (*i*) and (*ii*) are leaves. Thus all nodes in H_a are either leaves or have two children. We use the following claim:

Claim 6.8. *A node at depth $d \le \log n$ cannot be a leaf.*

Proof. Let α be a node at depth d, corresponding to type-a block Λ. We show that $C_\Lambda \ge C - d \cdot x$ by induction on d. Certainly when $d = 0$, the claim holds since $C_{\Lambda_1} = C$. Consider $d > 0$. The parent of Λ, Par_Λ, has depth $d - 1$, so $C_{Par_\Lambda} \ge C - (d-1)x$, by the induction hypothesis. Since $d - 1 \le \log n$, by assumption $f_{Par_\Lambda} > 2C^*$ and by construction of the children in H_a, Λ is a block in which some player is locally-optimal in the subgame $\mathcal{R}_\Lambda$ and has congestion at least $C_{Par_\Lambda} - x \ge C - (d-1) \cdot x + x$. Therefore, $C_\Lambda \ge C - d \cdot x$.

Since $d \le \log n$, we conclude that $C_\Lambda - x \ge C - x(1 + \log n) > C^*$ by assumption in the statement of the theorem. Thus, condition (*i*) is satisfied for α to have children. Since $C_\Lambda \ge C - x \log n$, by assumption $f_\Lambda > 2C^*$, hence condition (*ii*) is satisfied for α to have children. Since both conditions are satisfied, α cannot be a leaf node. □

We are now ready to conclude the proof of the lemma by obtaining a contradiction. Since H_a must have a leaf node, we conclude that the depth of H_a is at least $1 + \log n$. Since every node at depth at most $\log n$ has 2 children, we conclude that H_a has at least $2^{\lfloor 1+\log n \rfloor} > 2^{\log n} = n$ leaves. Thus, we have our contradiction since H_a cannot possibly contain more nodes than G. ■

By combining Lemma 4.3 and Lemma 6.5 we obtain the following result which establishes the upper bound of Theorem 1.3.

Lemma 6.9. *$PoA \le c \cdot ({\kappa_e}^2(G) + \log^2 n)$, for some constant c.*

Proof. Let $x = 2 \log n$. If $C \le C^* + x(1 + \log n)$, then there is nothing to prove because $C/C^* \le 1 + 2 \log n(1 + \log n)/C^* \le c \log^2 n$, for some generic constant c. So, suppose that $C > C^* + x(1 + \log n)$. By Lemma 6.5, there exists a type-a block Λ such that $C_\Lambda \ge C - 2 \log^2 n$ and $|F_\Lambda(C_\Lambda - 2 \log n)| \le 2C^*$. By applying Lemma 4.3 to the subgame $\mathcal{R}_\Lambda$ we obtain,

$$C_\Lambda < 2\ell \cdot \left(C^*_\Lambda + |F_\Lambda(C_\Lambda - 2 \log n')|\right) + 2 \log n',$$

where ℓ is the length of the longest edge-simple path in the player strategy sets in $\mathcal{R}_\Lambda$, n' is the number of nodes in Λ and C^*_Λ is the optimal congestion for the subgame $\mathcal{R}_\Lambda$. Note that $n' \le n$, and the subgame $\mathcal{R}_\Lambda$ cannot have a higher optimal congestion than the full game $\mathcal{R}$, hence $C^* \ge C^*_\Lambda$. Since $|F_\Lambda|$ is monotonically non-increasing ($F_\Lambda(C') \subseteq F_\Lambda(C'')$ for $C'' < C'$), we have that:

$$C - 2 \log^2 n \;< 2\ell \cdot \left(C^* + |F_\Lambda(C_\Lambda - 2 \log n)|\right) + 2 \log n \le 2\ell \cdot \left(C^* + 2C^*\right) + 2 \log n.$$

From Lemma 5.2, $\ell \le c{\kappa_e}^2(\Lambda) \le c{\kappa_e}^2(G)$, and so $C \le c \cdot ({\kappa_e}^2(G)C^* + \log^2 n)$. After dividing by C^*, we obtain the desired result. ■

7. Discussion

All our results have been stated for paths that are edge simple. Specifically the strategy set for the players is a set of *edge-simple* paths and the social and player costs are the maximum *edge-congestion* in the network and player path respectively. However, exactly analogous results can be obtained for strategy sets containing *node-simple* paths with the social and player costs being the maximum *node-congestion* in the network and player path respectively. In this case, the bounds on the price of anarchy are in terms of the node-cycle number (the length of the longest node-simple cycle).

We believe that the price of anarchy upper bound can be improved to match the lower bound. Specifically, we leave open the following conjecture: *for any routing game, $PoA \le \kappa_e - 1$.* An interesting future direction is to obtain similar results when

Author's personal copy

the latency functions at each link are more general and not necessarily identical, that is, the network graph is weighted or the edge congestion functions are nonlinear.

Another problem is to study the convergence time to Nash equilibria. It is shown in [6] that computing the best Nash equilibrium with optimal social cost is an NP-hard problem. A question remains about the complexity of finding an arbitrary Nash equilibrium. The optimization problem of computing an optimal coordinated solution that minimizes the social cost can be solved in polynomial time by using a simple flow approach (this method was proposed by one of the reviewers of this paper):

> Let $G = (V, E)$ be a graph and $(s_1, t_1), \dots, (s_N, t_N)$ be the source-sink pairs of the players. Now construct a network G_K, where K is a parameter. Each edge in G receives capacity K. Add two additional vertices s, t, and add the edges (s, s_i), (t_i, t) for $(i = 1, \dots, N)$ with respective capacity one. Now find the smallest integer K such that G_K admits a (maximum) s-t-flow with value equal to N. This can be done in polynomial time (with binary search). Moreover, as the capacities are integers, we can augment the flow always integrally, which yields that the flow on any edge is equal to the number of players using it.

The optimization problem of minimizing the maximum congestion has also been studied in the context of oblivious routing [5,8,23,40]. It is an interesting problem to investigate whether such methods of computing optimal coordinated solutions can be extended to compute Nash equilibria.

Acknowledgments

We would like to thank the reviewers of this journal for their thorough and constructive comments.

References

[1] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, Tim Roughgarden, The price of stability for network design with fair cost allocation, SIAM Journal on Computing 38 (4) (2008) 1602–1623.
[2] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, Tom Wexler, Near-optimal network design with selfish agents, in: Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC), San Diego, California, USA, June 2003, pp. 511–520.
[3] B. Awerbuch, Y. Azar, A. Epstein, The price of routing unsplittable flow, in: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), Baltimore, MD, USA, ACM, May 2005, pp. 57–66.
[4] Baruch Awerbuch, Yossi Azar, Yossi Richter, Dekel Tsur, Tradeoffs in worst-case equilibria, in: Klaus Jansen, Roberto Solis-Oba (Eds.), Proceedings of Approximation and Online Algorithms, First International Workshop, Springer, 2004, pp. 41–52.
[5] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, H. Racke, Optimal oblivious routing in polynomial time, in: Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC), San Diego, CA, ACM Press, June 2003, pp. 383–388.
[6] Ron Banner, Ariel Orda, Bottleneck routing games in communication networks, IEEE Journal on Selected Areas in Communications 25 (6) (2007) 1173–1179. Also published in INFOCOM'06.
[7] P. Berenbrink, C. Scheideler, Locally efficient on-line strategies for routing packets along fixed paths, in: 10th ACM-SIAM Symposium on Discrete Algorithms (SODA), 1999, pp. 112–121.
[8] Marcin Bienkowski, Miroslaw Korzeniowski, Harald Räcke, A practical algorithm for constructing oblivious routing schemes, in: Proceedings of the 15th Annual ACM Symposium on Parallelism in Algorithms and Architectures, June 2003, pp. 24–33.
[9] Costas Busch, Rajgopal Kannan, Athanasios V. Vasilakos, Quality of routing congestion games in wireless sensor networks, in: Proc. 4th International Wireless Internet Conference (WICON), Maui, Hawaii, November 2008.
[10] George Christodoulou, Elias Koutsoupias, The price of anarchy of finite congestion games, in: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC), Baltimore, MD, USA, ACM, May 2005, pp. 67–73.
[11] José R. Correa, Andreas S. Schulz, Nicolás E. Stier Moses, Computational complexity, fairness, and the price of anarchy of the maximum latency problem, in: Proc. Integer Programming and Combinatorial Optimization, 10th International IPCO Conference, New York, NY, USA, in: Lecture Notes in Computer Science, vol. 3064, Springer, June 2004, pp. 59–73.
[12] Robert Cypher, Friedhelm Meyer auf der Heide, Christian Scheideler, Berthold Vöcking, Universal algorithms for store-and-forward and wormhole routing, in: Proc. of the 28th ACM Symp. on Theory of Computing, 1996, pp. 356–365.
[13] Czumaj,Vocking, Tight bounds for worst-case equilibria, in: ACM Transactions on Algorithms (TALG), vol. 3, 2007.
[14] Artur Czumaj, Piotr Krysta, Berthold Vöcking, Selfish traffic allocation for server farms, in: Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC), Montreal, Quebec, Canada, May 2002, pp. 287–296.
[15] Dimitris Fotakis, Spyros C. Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, Paul G. Spirakis, The structure and complexity of nash equilibria for a selfish routing game, in: Proc. Automata, Languages and Programming, 29th International Colloquium, (ICALP), Malaga, Spain, in: Lecture Notes in Computer Science, vol. 2380, Springer, July 2002, pp. 123–134.
[16] Dimitris Fotakis, Spyros C. Kontogiannis, Paul G. Spirakis, Selfish unsplittable flows, Theoretical Computer Science 348 (2–3) (2005) 226–239.
[17] Martin Gairing, Thomas Lcking, Marios Mavronicolas, Burkhard Monien, The price of anarchy for polynomial social cost, Theoretical Computer Science 369 (2006) 116–135.
[18] Martin Gairing, Thomas Lcking, Marios Mavronicolas, Burkhard Monien, The price of anarchy for restricted parallel links, Parallel Processing Letters (PPL) 16 (1) (2006) 117–131.
[19] Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, Manuel Rode, Nash equilibria in discrete routing games with convex latency functions, Journal of Computer and System Sciences 74 (2008) 1199–1225.
[20] Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, Computing Nash equilibria for scheduling on restricted parallel links, in: Proceedings of the 36th Annual ACM Symposium on the Theory of Computing (STOC), Chicago, Illinois, USA, June 2004, pp. 613–622.
[21] Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, The price of anarchy for polynomial social cost, Theoretical Computer Science 369 (1–3) (2006).
[22] Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, Manuel Rode, Nash equilibria in discrete routing games with convex latency functions, in: Proc. Automata, Languages and Programming: 31st International Colloquium, ICALP, Turku, Finland, in: Lecture Notes in Computer Science, vol. 3142, Springer, July 2004, pp. 645–657.
[23] Chris Harrelson, Kristen Hildrum, Satish Rao, A polynomial-time tree decomposition to minimize congestion, in: Proceedings of the 15th Annual ACM Symposium on Parallelism in Algorithms and Architectures, June 2003, pp. 34–43.

Author's personal copy

[24] M. Hoefer, A. Souza, Tradeoffs and average-case equilibria in selfish routing, in: Proceedings of the 15th European Symposium on Algorithms (ESA 2007), 2007, pp. 63–74.
[25] M. Hoefer, A. Souza, The influence of link restrictions on (random) selfish routing, in: Proceedings of the First Symposium on Algorithmic Game Theory (SAGT 2008), 2008, pp. 22–32.
[26] Elias Koutsoupias, Marios Mavronicolas, Paul G. Spirakis, Approximate equilibria and ball fusion, Theory of Computing System 36 (6) (2003) 683–693.
[27] Elias Koutsoupias, Christos Papadimitriou, Worst-case equilibria, Trier, Germany, in: LNCS, vol. 1563, Springer-Verlag, March 1999, pp. 404–413.
[28] F.T. Leighton, B.M. Maggs, S.B. Rao, Packet routing and job-scheduling in O ($congestion + dilation$) steps, Combinatorica 14 (1994) 167–186.
[29] F. Thomson Leighton, Introduction to Parallel Algorithms and Architectures: Arrays - Trees - Hypercubes, Morgan Kaufmann, San Mateo, 1992.
[30] Tom Leighton, Bruce Maggs, Andrea W. Richa, Fast algorithms for finding O(congestion + dilation) packet routing schedules, Combinatorica 19 (1999) 375–401.
[31] Lavy Libman, Ariel Orda, Atomic resource sharing in noncooperative networks, Telecommunication Systems 17 (4) (2001) 385–409.
[32] Thomas Lücking, Marios Mavronicolas, Burkhard Monien, Manuel Rode, A new model for selfish routing, Theoretical Computer Science 406 (3) (2008) 187–206.
[33] Mavronicolas,Spirakis, The price of selfish routing, Algorithmica 48 (2007).
[34] Vladimir Mazalov, Burkhard Monien, Florian Schoppmann, Karsten Tiemann, Wardrop equilibria and price of stability for bottleneck games with splittable traffic, in: Proceedings of the Workshop on Internet and Network Economics (WINE 2006), 2007, pp. 331–342.
[35] Karl Menger, Zur allgemeinen kurventheorie, Fundamenta Mathematicae 10 (1927) 96–115.
[36] D. Monderer, L.S. Shapely, Potential games, Games and Economic Behavior 14 (1996) 124–143.
[37] Rafail Ostrovsky, Yuval Rabani, Universal O (congestion + dilation + $\log^{1+\varepsilon} N$) local control packet switching algorithms, in: Proceedings of the 29th Annual ACM Symposium on the Theory of Computing, New York, May 1997, pp. 644–653.
[38] Christos Papadimitriou, Algorithms, games, and the Internet, in: ACM, editor, Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC), Hersonissos, Crete, Greece, July 2001, pp. 749–753.
[39] Yuval Rabani, Éva Tardos, Distributed packet switching in arbitrary networks, in: Proceedings of the Twenty-Eighth Annual ACM Symposium on the Theory of Computing, Philadelphia, Pennsylvania, 22–24 May 1996, pp. 366–375.
[40] Harald Räcke, Minimizing congestion in general networks, in: Proceedings of the 43rd Annual Symposium on the Foundations of Computer Science, Nov. 2002, pp. 43–52.
[41] R.W. Rosenthal, A class of games possessing pure-strategy Nash equilibria, International Journal of Game Theory 2 (1973) 65–67.
[42] Tim Roughgarden, The maximum latency of selfish routing, in: Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), New Orleans, Louisiana, (USA), January 2004, pp. 980–981.
[43] Tim Roughgarden, Selfish routing with atomic players, in: Proc. 16th Symp. on Discrete Algorithms (SODA), ACM/SIAM, 2005, pp. 1184–1185.
[44] Tim Roughgarden, Éva Tardos, How bad is selfish routing, Journal of the ACM 49 (2) (2002) 236–259.
[45] Tim Roughgarden, Éva Tardos, Bounding the inefficiency of equilibria in nonatomic congestion games, Games and Economic Behavior 47 (2) (2004) 389–403.
[46] Subhash Suri, Csaba D. Toth, Yunhong Zhou, Selfish load balancing and atomic congestion games, Algorithmica 47 (1) (2007) 79–96.