Doing Good with Spam is Hard *

Martin Hoefer, Lars Olbrich, and Alexander Skopalik

Department of Computer Science, RWTH Aachen University, Germany

Abstract. We study economic means to improve network performance in the well-known game theoretic traffic model due to Wardrop. We introduce two sorts of spam flow - auxiliary and adversarial flow - that have no intrinsic value. Auxiliary/adversarial flows are a separate commodity with the sole objective to minimize/maximize the latency at the induced Wardrop equilibrium of the selfish flow. By this means a separate access to the edges is not necessary and the latency of the regulating flow does not distort the arising latency cost. We present networks where auxiliary flow is able to improve the network performance. However, we show that the optimal auxiliary flow is NP-hard to compute and not approximable within a factor of less then $\frac{4}{3}$. The minimal amount of auxiliary flow needed to induce the best possible equilibrium is even hard to approximate by any subexponential factor. These hardness results are complemented by proving NP-hardness for the optimal adversarial flow. All hardness results hold even for single-commodity networks.

1 Introduction

Wardrop's traffic model is a well-studied model for routing with important applications in road traffic and computer networks. In this model, we are given a network equipped with non-decreasing non-negative latency functions mapping flow on the edges to latency. For each of several commodities a fixed demand has to be routed between a source-sink pair. The cost of a flow assignment is the weighted sum of travel times between the source and target nodes. A flow that minimizes the total latency is called *(socially) optimal*. A common interpretation of the Wardrop model is that flow is controlled by an infinite number of selfish users each of which carries an infinitesimal amount of flow. Each user aims at minimizing its path latency. An allocation in which no user can improve its situation by unilaterally deviating from its current path is called *Wardrop equilibrium*. In general a Wardrop equilibrium is not socially optimal, i.e, it does not minimize the total latency. The inefficiency of selfish flows has been extensively studied in previous work [3, 23, 24, 26].

We study a means of reducing the inefficiency of selfish flow applicable in scenarios with no central control. There have been several approaches to this problem in the literature, most prominently taxing, Stackelberg routing, and

* Supported by the DFG GK/1298 "AlgoSyn", by the German Israeli Foundation (GIF) under contract 877/05, and by DFG through UMIC Research Centre at RWTH Aachen University.

network design, but there are some problems with these approaches in large networks without strong centralized control. Taxing requires to collect possibly different taxes at each edge, a process that requires an infrastructure that can be costly or impossible to establish. In addition, a look at classical taxing procedures from a user perspective reveals that, albeit taxes improve the latency of the networks, they do not improve the disutility of users for a large set of networks [7]. In Stackelberg routing the idea is to put a fraction of selfish flow under centralized control and reroute this flow such that the total latency of all flow is optimized. Here the underlying assumption that a central control agency can directly manipulate the selfish demand is quite strong. Finally, network design requires to manipulate the network structure, which is clearly a strong assumption of centralized control in a large network.

In this paper, we consider a means of control motivated by the concept of spam in the Internet. We introduce two sorts of *spam flows*, which we call auxiliary and adversarial flow. The demand value of these flows is given independently in addition to the given selfish flow demand. Spam flow can be seen as a separate altruistic or malicious commodity that tries to influence the routing decisions of selfish players without directly taking control over (parts of) the players or the network. The goal is to route the spam flow in such a way that the induced equilibrium minimizes/maximizes the total latency of the selfish flow. The routed packets solely alter the latency of the used edges. They have no value and are essentially spam. Therefore we assume that the latency of spam flow does not contribute to the social cost.

Our results We first present networks where auxiliary flow eradicates the inefficiency of the Wardrop equilibrium (Section 2). However, it turns out that both the *optimal auxiliary flow* of given value and the *minimal amount of an optimal auxiliary flow* are NP-hard to compute (Subsection 3.1 and 3.2). Further, we prove that for auxiliary flow there is no polynomial time approximation with a factor of less than $\frac{4}{3}$. The minimal amount of the optimal auxiliary flow cannot be approximated by any subexponential factor. These results are complemented by proving NP-hardness for adversarial flow (Subsection 3.3).

Related Work The game theoretic traffic model considered in this paper was introduced by Wardrop [29]. Beckmann et al. [2] observe that such an equilibrium flow is an optimal solution to a related convex program. They give existence and uniqueness results for traffic equilibria (see also [9] and [24]). Dafermos and Sparrow [9] show that the equilibrium state can be computed efficiently under some assumptions on the latency functions.

The inefficiency of Wardrop equilibria is a well-known phenomenon [20], which is exemplified by Braess paradox [3]. Bounding the inefficiency of equilibria, however, has only recently been considered, initiated by Koutsoupias and Papadimitriou [18], and for the Wardrop model by Roughgarden and Tardos [24, 26].

One of the most prominent approaches to eradicate the inefficiency of Wardrop equilibria is taxing. The effectiveness of taxes has been observed by Pigou [20] and generalized by Beckmann et al. [2]. They show that *marginal cost pricing* completely eliminates the inefficiency of selfish routing. Major results for taxes for heterogeneous users can be found in [8], [10], [11] and [14]. Cole et al. [7] consider taxes that minimize the total user disutility (latency plus tax) at equilibrium. They show that for linear latency functions marginal cost pricing does not improve the cost of Wardrop equilibria and prove tight inapproximability results for optimal taxes.

Korilis et al. [16] consider the problem of a Stackelberg leader, who in a first phase can fix the routes for a certain fraction of the demand. In a second phase, selfish users enter the system and route their own flow on top of the leader demand. The objective of the leader is to minimize the resulting cost of the total (both leader and selfish) flow. Roughgarden [22] shows that it is weakly NP-hard to compute the optimal leader strategy even for parallel links with linear latency functions. Kaporis and Spirakis [13] show that for single-commodity networks the minimal fraction of flow needed by the leader to induce optimal cost, can be computed in polynomial time. Sharma and Williamson [27] compute the minimum fraction of users that must be centrally routed to improve the quality of the resulting Wardrop equilibrium. Subsequent papers [28, 15, 4] consider Stackelberg routing in different variants for more general networks.

Roughgarden [25] studies designing networks that exhibit good performance when used selfishly and proves tight inapproximability results.

Other approaches for coping with selfishness are, for example, proposed by Korilis et al. [17] who add capacity to the resources and Cocchi et al. [6] who study the role of various pricing policies in networks with selfish users.

While the fundamental assumptions is that all agents act selfishly, large systems often display forms of altruism or spite. In these cases, some agents' goals is to improve or to harm the global outcome instead of optimizing their personal objective function. Babaioff et al. [1] and Roth [21] study the existence of equilibria for these games, and quantify the impact of malicious players on the equilibrium. Chen and Kempe [5] proved that equilibria exist for any population of selfish, altruistic and spiteful agents.

2 Preliminaries and Initial Results

We first define the classical Wardrop model originally introduced in [29] and then introduce our additional spam flow. We are given a directed graph $G = (V, E)$ with vertex set V, edge set E, a set of commodities $[k] = \{1, \ldots, k\}$ specified by source-sink pairs $(s_i, t_i) \in V \times V$, and flow demands $d_i > 0$. The edges are equipped with non-decreasing, continuous latency functions $\ell_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$.

Let $\mathcal{P}_i$ denote the available paths of commodity i, i. e., all paths connecting s_i and t_i, and let $\mathcal{P} = \bigcup_{i \in [k]} \mathcal{P}_i$. A non-negative path flow vector $(f_P)_{P \in \mathcal{P}}$ is *feasible* if it satisfies the flow demands $\sum_{P \in \mathcal{P}_i} f_P = d_i$ for all $i \in [k]$. Throughout this paper, we will consider only feasible path flow vectors. For single commodity networks we drop the index i and normalize the demand to one. A path flow vector $(f_P)_{P \in \mathcal{P}}$ induces an edge flow vector $f = (f_e)_{e \in E}$ with $f_e = \sum_{i \in [k]} \sum_{P \in \mathcal{P}_i : e \in P} f_P$. The latency of an edge $e \in E$ is given by $\ell_e(f_e)$ and the latency of a path P is given by the sum of the edge latencies $\ell_P(f) = \sum_{e \in P} \ell_e(f_e)$. The latency cost of a flow is defined as $C(f) = \sum_{P \in \mathcal{P}} \ell_P(f) f_P = \sum_{e \in E} \ell_e(f_e) f_e$. A flow f of minimal latency cost is called *(socially) optimal*.

Additionally to the given selfish flow, we introduce two kinds of spam flow - auxiliary and adversarial flow (δ_e). Let $\delta > 0$ denote the spam flow and its demand. The objective of the spam flow is to minimize/maximize the latency cost of the induced equilibrium of the selfish flow. The routed spam has no intrinsic value and hence does not contribute to the latency cost. Given the routes of the spam flow and the selfish flow, the latency cost equals $C(f, \delta) = \sum_{e \in E} \ell_e(f_e + \delta_e) f_e$. If not specified further, we refer by flow to the selfish flow. Finally, we call the tuple $\Gamma = (G, (s,t), \delta)$ an *instance*.

A flow vector is considered stable when no fraction of the flow can improve its sustained cost by moving unilaterally to another path. Such a stable state is generally known as *Nash equilibrium*. In our model a flow is stable if and only if all used paths within a commodity have the same minimal latency, whereas unused paths may have larger latency. We call such a flow *Wardrop equilibrium*.

Definition 1 *Given an instance Γ and fixed routes for the spam δ, a feasible flow vector f is at* Wardrop equilibrium *if for every commodity $i \in [k]$ and paths $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ it holds that $\ell_{P_1}(f + \delta) \le \ell_{P_2}(f + \delta)$.*

Observation 1 *If f is at Wardrop equilibrium then all used paths in commodity i have equal latency $L_i(f, \delta)$ and the latency cost can be expressed as $C(f, \delta) = \sum_{i \in [k]} L_i(f, \delta) \cdot d_i$ ([24, 29]).*

Note that the spam commodity δ is not composed of stabilizing selfish users. Instead, the aim is to allocate this flow in a coordinated way to influence the cost of the Wardrop equilibrium. Our optimization problem is similar to Stackelberg routing [16]. In particular, it can be formulated as a bilevel problem, where in a first phase spam flow is allocated to the routes. In a second phase the selfish flow stabilizes at Wardrop equilibrium depending on the allocation in the first phase. The resulting latency of the selfish flow is to be optimized by the allocation of spam flow in the first place.

Let us note two initial observations about auxiliary flow. Figure 1 yields our first observation.

Observation 2 *There are networks in which auxiliary flow eradicates the inefficiency of selfish routing.*

One can easily modify the network in Figure 1, such that even an arbitrary small amount of auxiliary flow does the job.

Observation 3 *Adding auxiliary flow to selfish flow increases the path latency in series-parallel graphs. Since the cost at equilibrium equals the path latency L, auxiliary flow of arbitrary value does not improve the latency cost at equilibrium.*

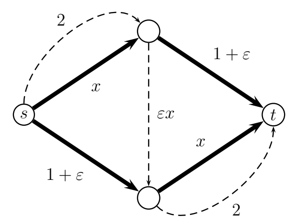

Fig. 1. In absence of spam flow, the selfish flow uses only the zig-zag-path at equilibrium. Routing spam over the dashed edges, the selfish flow splits half-half among the bold edges and reaches the social optimum.

3 Computational Complexity

In this section, we discuss the computational complexity of problems related to auxiliary and adversarial flow.

In the decision problem Optimal-Flow we are given a single-source selfish routing instance, an amount of auxiliary flow, and a cost value C. The problem is to decide if there is a routing of the auxiliary flow such that the latency cost of the equilibrium is at most C.

In the decision problem Threshold-Flow we are given a single-source selfish routing instance and an amount of auxiliary flow δ. The problem is to decide if there is a routing of the auxiliary flow such that the latency cost of the equilibrium is less or equal than the latency cost of the equilibrium induce by any auxiliary flow $\delta' > \delta$.

In the decision problem Worst-Flow we are given a single-source selfish routing instance, an amount of adversarial flow, and a cost value C. The problem is to decide if there is a routing of the adversarial flow such that the latency cost of the equilibrium is at least C.

3.1 Complexity of Optimal-Flow

Observation 2 shows that auxiliary flow can improve the cost of Wardrop equilibria. Here, we show that computing the optimal routing for the auxiliary flow is NP-hard.

Theorem 1. Optimal-Flow *is* NP-*hard.*

Proof. Our proof is based on the proof given in [7] to show that taxing to minimize total disutility is hard. We reduce from the problem 2 Directed Disjoint Path (2DDP) which is known to be NP-hard [12]. An instance $I = (G, (s_1, t_1), (s_2, t_2))$ is a directed graph G and two pairs of nodes (s_1, t_1) and (s_1, t_2). An instance I belongs to 2DDP, that is $I \in$ 2DDP, if and only if there

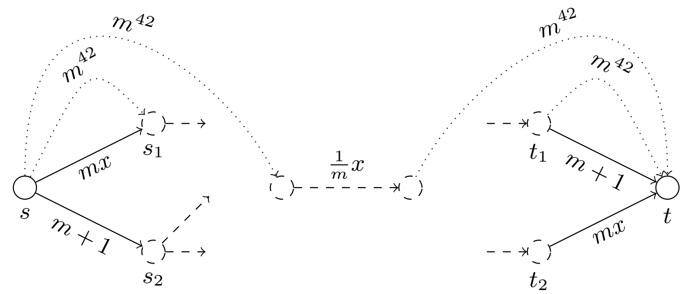

Fig. 2. This figure outlines the construction of G'. The dashed edges are the edges of G and the dotted edges are the edges in P. The edges are labeled with their latency functions

exist two node disjoint paths in G from s_1 to t_1 and from s_2 to t_2, respectively. Without loss of generality, we assume that there exist paths from s_1 to t_1 and from s_2 to t_2, respectively.

Given an instance $I = (G, (s_1, t_1), (s_2, t_2))$ with $G = (V, E)$ and $|E| = m$, we construct a single commodity selfish routing game $\Gamma = (G', (s, t), \delta)$ with auxiliary flow of $\delta = 3m^2$ that has the following properties: If and only if $I \in$ 2DDP, optimal auxiliary flow yields a Wardrop equilibrium with social cost of less than $C = \frac{3}{2}m + 2$.

We construct $G' = (V', E')$ as follows: $V' = V \cup \{s, t\}$ and $E' = E \cup \{(s, s_1), (s, s_2), (t_1, t)(t_2, t)\} \cup P$ with $P = \{(s, u), (v, t) \mid \text{for all } (u, v) \in E\}$. The latency function of each edge $e \in E$ is $\ell_e(x) = \frac{1}{m}x$, for the edges $e \in \{(s, s_1), (t_2, t)\}$ it is $\ell_e(x) = mx$, for the edges $e \in \{(s, s_2), (t_1, t)\}$ it is $\ell_e(x) = m + 1$, and for all edges $e \in P$ it is $\ell_e(x) = m^{42}$. Note that in equilibrium no selfish flow is assigned to an edge $e \in P$ because latency of m^{42} is much larger than the latency of any s-t-path that does not include an edge $e \in P$.

If $I \in$ 2DDP, there exist two disjoint paths from s_1 to t_1 and from s_2 to t_2, respectively, in G'. Let $D \subseteq E$ be the set of edges of these two paths. An auxiliary flow that assigns, for all $(u, v) \in E \setminus D$, flow of at least $3m$ to each of the edges $(s, u), (v, t) \in P$, and (u, v) essentially forces the selfish flow to use the two disjoint paths only. The latency for flow demand d' on such a path is at least $md' + m + 1$ and at most $md' + m \cdot \frac{1}{m} + m + 1$. Thus, in equilibrium the maximal flow demand on each of the two paths is bounded by $\frac{m+1}{2m+1}$. Therefore, the latency of a path in a resulting Wardrop equilibrium is at most $\frac{3}{2}m + 2$ and the latency cost is at most C.

If $I \notin$ 2DDP, we show that there is no auxiliary flow that induces an equilibrium flow with social cost of less than $2m$. We distinguish several cases by the usage of the four edges incident to s and t. It suffices to show that there is an used path with latency of at least $2m$.

1. If a flow uses a path starting with (s, s_2) and ending with (t_1, t), this path has latency of at least $2m + 2$.
2. If a flow uses only paths starting with (s, s_1) and ending with (t_2, t), it has cost of at least $2m$.
3. If a flow uses only paths starting with (s, s_1) and ending with (t_2, t) or (t_1, t), the latency from s_1 to t must be the same on all paths. Therefore every path has latency of at least $2m + 1$.
4. If a flow uses only paths starting with (s, s_1) or (s, s_2) and ending with (t_2, t), the same argument holds.
5. If a flow uses at least one path starting with (s, s_1) and ending with (t_1, t) and at least one path starting with (s, s_2) and ending with (t_2, t), there exists a vertex v^* that is contained in both paths. All path segments from s to v^* and from v^* to t must have the same latency. Thus, every path has latency of at least $2m + 2$.

Thus, the optimal auxiliary flow induces an equilibrium with social cost less or equal C in Γ if and only if $I \in$ 2DDP. □

Note that the decision in the previous instances is whether the cost of the selfish flow can be reduced to a cost of at most $C = \frac{3}{2}m + 2$. If this is impossible, for every flow the cost is at least $2m$. Now suppose there is a polynomial time approximation algorithm, which computes a $\left(\frac{4}{3} - \epsilon\right)$-approximation for optimizing the cost of selfish flow. Then, such an algorithm could be used to decide 2DDP using the previously outlined set of instances. We therefore get the following corollary. Note that a $\frac{4}{3}$-approximation for linear latencies is trivially obtained by routing no auxiliary flow at all [24].

Corollary 4 *For every $\epsilon > 0$ it is* NP*-hard to approximate* Optimal-Flow *on instances with linear latency functions to a factor of $\frac{4}{3} - \epsilon$.*

In addition, note that in the NP-hardness reduction the auxiliary flow is much larger than the demand of selfish flow. However, we can show that the result even holds, if the auxiliary flow is only a (polynomially small) fraction of the selfish demand.

Theorem 2. Optimal-Flow *is* NP*-hard and even* NP*-hard to approximate to a factor of $\frac{4}{3} - \epsilon$ for every $\epsilon > 0$ on instances with linear latency functions and auxiliary flow $\delta \in \mathcal{O}\left(\frac{d}{poly(m)}\right)$.*

Proof. Again, we reduce from 2DDP. Given an instance I and an ϵ, we construct a selfish routing game Γ as described in the proof of Theorem 1. We use $k = 3m^2 \cdot \lceil \epsilon^{-1} \rceil$ copies $\Gamma_1, \ldots, \Gamma_k$ of this game to create a new game Γ' as follows. We add a source vertex s^* and a target vertex t^*. The vertex s^* is connected to each source vertex s'_i of Γ_i (for all $1 \leq i \leq k$) by an edge (s^*, s_i) with the latency

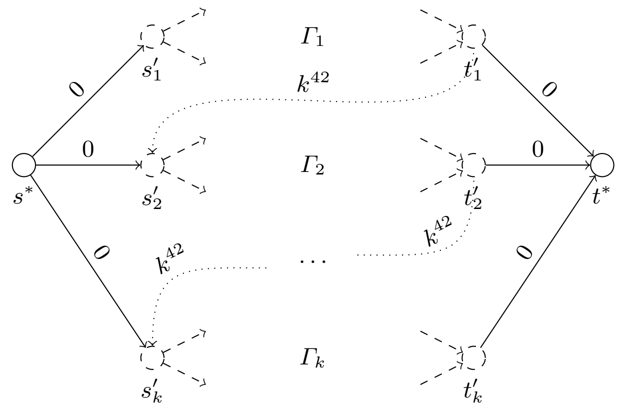

Fig. 3. The network contains $k = 3m^2 \cdot \lceil \epsilon^{-1} \rceil$ copies of the network G' of the proof of Theorem 1. Between s^* and t^* there is a demand of k.

function $\ell_{(s^*,s_i)}(x) = 0$. Likewise, there is an edge with $\ell_{(t'_i,t^*)}(x) = 0$ from each vertex t'_i to t^*. Additionally, for every $i \in \{1, \ldots, k-1\}$, there is an edge from t'_i to s'_{i+1} with $\ell_{(t'_i,s_{i+1})}(x) = k^{42}$. The demand of the selfish flow is $d = k$ and the auxiliary flow is limited to $3m^2$ and $C = d \cdot \frac{3}{2}m + 2$.

If $I \in$ 2DDP, there is an auxiliary flow that yields an equilibrium flow with social costs of at most $d \cdot (\frac{3}{2}m + 2)$: We assign auxiliary flow of at most $3m^2$ between the vertices s'_i and t'_i in each copy Γ_i as described in the proof of Theorem 1. We assign the same amount of flow to the edges $\{(s^*, s'_1), (t'_1, s'_2), \ldots, (t'_{k-1}, s'_k), (t'_k, t^*)\}$ to obtain a flow of at most $3m^2$ from s^* to t^*. In the resulting Wardrop equilibrium, there is a flow of 1 that is assigned to each copy Γ_i and the edges that connect it to s^* and t^*. Each of these flows has cost of at most $\frac{3}{2}m + 2$. Thus the social cost sum up to at most $d \cdot (\frac{3}{2}m + 2)$.

If $I \notin$ 2DDP, then the latency cost of the selfish flow is more than $d \cdot 2m$. Note, that in equilibrium the selfish flow never chooses an edge that connects two of the copies because it has latency of k^{42} and there is always a s^*-t^*-path with lower latency. Therefore, there is at least one copy Γ_i in which flow of at least 1 is routed from s'_i to t'_i. As shown in the proof of Theorem 1, the latency of the s'_i-t'_i-paths at least $2m$. Since the flow is a Wardrop equilibrium, every path between s'_j and t'_j for every $1 \leq j \leq k$ has latency of at least $2m$. Thus, the latency cost sums up to more than $d \cdot 2m$. □

3.2 Complexity of Threshold-Flow

The previous result showed that it is computationally difficult to compute the best possible auxiliary flow. In this section we show that it is even hard to approximate the minimal amount of auxiliary flow that is needed to achieve the best possible Wardrop equilibrium.

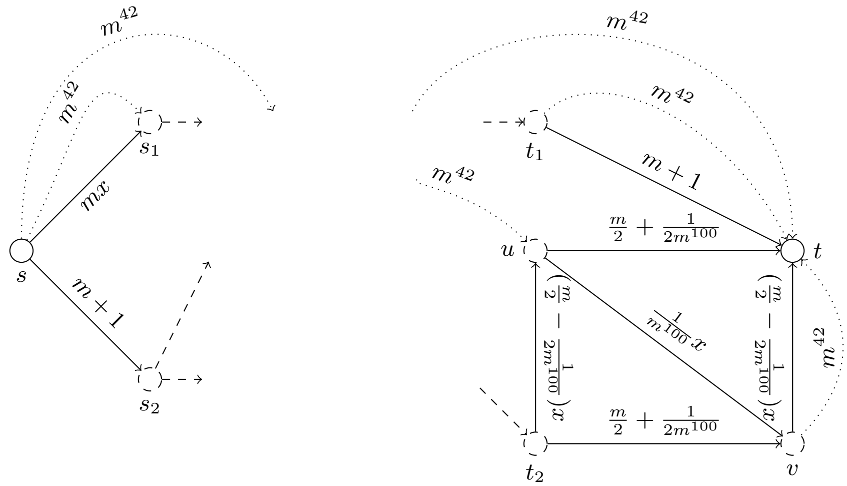

Fig. 4. This figure outlines the modified construction of G' for the proof of Theorem 3.

Note that this result strongly contrasts the corresponding result of Kaporis and Spirakis [13] for Stackelberg routing. In Stackelberg routing the minimal fraction of flow needed by the Stackelberg leader to induce optimal cost can be computed in polynomial time for single commodity networks.

Theorem 3. Threshold-Flow *is* NP*-hard.*

Proof. Again, we reduce from 2 Directed Disjoint Path (2DDP). Given an instance $I = (G, (s_1, t_1), (s_2, t_2))$ with $G = (V, E)$ and $|E| = m$, we construct a single commodity selfish routing game whose optimal auxiliary flow has demand of at most $3m^3$ if and only if $I \in$ 2DDP. Construct $\Gamma = (G', (s, t), \delta)$ as described in the proof of Theorem 1 and modify it as follows. Remove the edge (t_2, t) and replace it with the following gadget. Add the vertices u and v and the edges $(t_2, u), (u, v), (u, t), (t_2, v), (v, t)$. Latency functions are $\ell_e(x) = (\frac{m}{2} - \frac{1}{2m^{100}})x$ for the edges $e \in \{(t_2, u), (v, t)\}$ and $\ell_e(x) = \frac{m}{2} + \frac{1}{2m^{100}}$ for the edges $e \in \{(u, t), (t_2, v)\}$ and $\ell_{(u,v)}(x) = \frac{1}{m^{100}}x$. We add additional edges (s, u) and (v, t) with latency m^{42} and add them to the set P (cf. proof of Theorem 1).

Observe that for routing flow demand $d' \leq \frac{2m^{101}+2}{3m^{101}+1}$ from t_2 to t, it is optimal to leave all selfish flow on the zig-zag path, which generates latency md' and also yields an equilibrium. Observe that the optimum assignment of selfish flow that is achievable by (marginal cost) taxing might split the flow along all three possible paths from t_2 to t. However, the resulting latency of such a flow is larger here, as the auxiliary flow is accounted in the latency of selfish flow. For more flow than d', splitting the flow and assigning $\frac{d'}{2}$ to the edges (t_2, u), (t_2, v), (u, t), and (v, t), yields a better latency. This flow and its improved latency can be achieved using a sufficiently large auxiliary flow along edge (u, v).

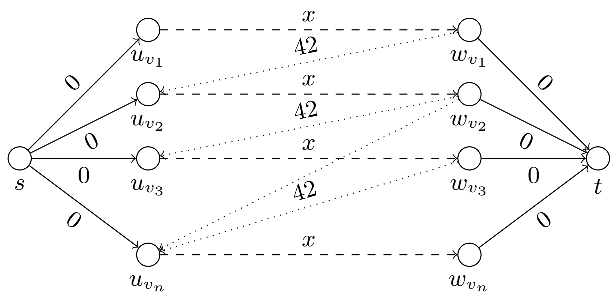

Fig. 5. This figure depicts the corresponding graph G' for an instance G for the problem HAMILTON. The dashed edges correspond to vertices in G and the dotted edges correspond to edges in G'.

If $I \in$ 2DDP, then optimal auxiliary flow of demand of at most $3m^3$ is sufficient to obtain the best possible Wardrop equilibrium. Note that for large m only close to $\frac{1}{2}$ selfish flow is routed through the gadget from t_2 to t. Therefore, it is not necessary to route auxiliary flow over the edge (u, v).

If $I \notin$ 2DDP, then optimal auxiliary flow yields a Wardrop equilibrium in which the whole selfish demand is routed from s via s_1 and t_2 to t. The optimal auxiliary flow must block edge (u, v). Even to motivate one selfish player to use (u, t), it needs to route demand of more than $\frac{m^{101}}{2} - \frac{1}{2}$ over the edge (u, v). $\square$

Note that one can easily replace the term m^{100} in the latency functions of our gadget with any arbitrarily large constant that can be represented by a polynomial number of bits in the input size. In particular, assuming that the numbers in our instance are represented in binary coding, it can be replaced by 2^m. Then, for $m \geq 2$, a $\frac{2^m}{6m^3}$-approximation algorithm of THRESHOLD-FLOW in the above instances can decide 2DDP. Thus, we have the following corollary.

Corollary 5 *For any constant $\epsilon > 0$, it is* NP*-hard to approximate* THRESHOLD-FLOW *by a factor of* $2^{m(1-\epsilon)}$.

3.3 Complexity of WORST-FLOW

We have seen that the optimal auxiliary flow is NP-hard to compute. We now turn to the computational complexity of computing the optimal adversarial flow.

Theorem 4. WORST-FLOW *is* NP*-hard.*

Proof. We reduce from the NP-hard problem HAMILTON. A graph $G \in$ HAMILTON if and only if G contains a Hamiltonian path. Given a directed graph $G = (V, E)$ with $|V| = n$ and $|E| = m$ and two vertices $x, y \in V$, we construct a selfish routing game $\Gamma = (G', (s, t), \delta)$ that has the property that the cost maximizing adversarial flow induces social cost of at least $C = \frac{1}{n} + \delta$ if and only if

$G \in$ Hamilton. We construct $G' = (V', E')$ as follows: For every node v in G there is a a pair of nodes u_v, w_v in G' and, additionally we have a source and a sink node s and t. That is $V = \{s,t\} \cup \{u_v, w_v \mid \forall v \in V\}$.

There are edges from s to all u nodes, from each node u_v to w_v and from all w nodes to t. The selfish flow will use only these edges. Additionally, we have edges (with high latency) that connect a node w_v with a node $u_{v'}$ if there is an edge from v to v' in the graph G for $v \in V - \{x\}$. To summarize $E' = S' \cup U' \cup W'$ with $S' = \{(u_v, w_v) \mid \forall v \in V\}$, $U' = \{(s, u_v), (w_v, t) \mid \forall v \in V\}$, and $W' = \{w_v, u_{v'} \mid \forall (v, v') \in E \text{ and } v' \in V - \{x\}\}$. For all edges $e \in S'$ we set $\ell_e(x) = x$, for all edges $e \in U'$ we set $\ell_e(x) = 0$, and for all edges $e \in W'$ we set $\ell_e(x) = 42$. Note that the selfish flow never uses edges $e \in W'$ and therefore, assigns flow to the n paths s, u_v, w_v, s (for all $v \in V$). Without adversarial flow, the equilibrium flow is equally distributed among these paths and the social costs are $n\frac{1}{n^2} = \frac{1}{n}$.

Assume $G \in$ Hamilton and $x = v_{i_1}, \ldots, v_{i_n} = y$ is a Hamiltonian path in G. Then it is possible to assign adversarial flow of δ to all edges $e \in S'$ by choosing the path $s, u_{v_{i_1}}, w_{v_{i_1}}, u_{v_{i_2}}, w_{v_{i_2}}, \ldots, u_{v_{i_n}}, w_{v_{i_n}},, t$. Note, that the edges between the w and u vertices exist by construction. Because all non constant edges carry the maximal amount of adversarial flow and this flow maximizes the social costs which are $m(\frac{1}{n} + \delta) \cdot \frac{1}{n} = \frac{1}{n} + \delta$.

Consider a graph $G \notin$ Hamilton. Then there is no path in G' from s to t that visits all vertices $e \in U'$. Therefore, there is at least one edge with adversarial flow less than δ. Thus, the latency cost of the equilibrium flow is strictly less than $\frac{1}{n} + \delta$. □

4 Conclusions

We have initiated the study of spam flow in non-atomic routing games. We considered the computational complexity of several problems related to auxiliary and adversarial flow. Both, Optimal-Flow and Worst-Flow turned out to be NP-hard. Moreover, Optimal-Flow and Threshold-Flow are hard to approximate, which strongly contrasts the results for the analogous problem of the “Price of Optimum” in Stackelberg routing [13]. Further research on algorithms and corresponding complexity issues regarding spam that improves or deteriorates latency cost may well be worthwhile.

References

[1] M. Babaioff, R. Kleinberg, C. Papadimitriou. Congestion games with malicious players. In *Proc. 8th EC*, 103–112, 2007.

[2] M. Beckmann, C. B. McGuire, C. B. Winston. Studies in the Economics of Transportation. *Yale University Press*, 1956.

[3] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258-268, 1968.

[4] V. Bonifaci, T. Harks. G.Schäfer. Stackelberg Routing in Arbitrary Networks. In *Proc. 4th WINE*, 239–250, 2008.

[5] P. Chen, D. Kempe. Altruism, selfishness, and spite in traffic routing. In *Proc. 9th EC*, 140–149, 2008.

[6] R. Cocchi, S. Shenker, D. Estrin, L. Zhang. Pricing in computer networks: motivation, formulation, and example. *IEEE/ACM Transactions on Networking*, 1(6):614–627, 1993.

[7] R. Cole, Y. Dodis, T. Roughgarden. How much can taxes help selfish routing? In *Proc. 4th EC*, 98–107, 2003.

[8] R. Cole, Y. Dodis, T. Roughgarden. Pricing edges for heterogeneous selfish users. In *Proc. 35th STOC*, 521–530, 2003.

[9] S. Dafermos, F. T. Sparrow. The Traffic Assignment Problem for a General Network. *J Research National Bureau of Standards*, Series B, 73(2):91–118, 1969.

[10] L. Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Computer Science*, 348:217–225, 2005.

[11] L. Fleischer, K. Jain, M. Mahdian. Taxes for heterogeneous selfish users in a multicommodity network. In *Proc. 45th FOCS*, 277–285, 2004.

[12] S. Fortune, J. E. Hopcroft, J. Wyllie. The Directed Subgraph Homeomorphism Problem. *Theoretical Computer Science*, 10:111–121, 1980.

[13] A.C. Kaporis, P.G. Spirakis. The price of optimum in stackelberg games on arbitrary single commodity networks and latency function. In *Proc. 18th SPAA*, 2006.

[14] G. Karakostas, S.G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *Proc. 45th FOCS*, 268–276, 2004.

[15] G. Karakostas, S.G. Kolliopoulos. Stackelberg strategies for selfish routing in general multicommodity networks. *Algorithmica*, to appear, 2008.

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

[17] Y.A. Korilis, A.A. Lazar, A. Orda. Capacity allocation under noncooperative routing. *IEEE Transactions on Automatic Control*, 42:309–325, 1997.

[18] E. Koutsoupias, C. Papadimitriou. Worst-Case Equilibria. In *Proc. 16th STACS*, 404–413, 1999.

[19] V.S.A. Kumar, A. Marathe. Improved results for Stackelberg scheduling strategies. In *Proc. 29th ICALP*, 2002.

[20] A. C. Pigou. *The Economics of Welfare*. Macmillan, 1920.

[21] A. Roth. The Price of Malice in Linear Congestion Games. In *Proc. 4th WINE*, 118–125, 2008.

[22] T. Roughgarden. Stackelberg scheduling strategies. In *Proc. 33rd STOC*, 104–113, 2001.

[23] T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. In *Proc. 34th STOC*, 428–437, 2002.

[24] T. Roughgarden, É. Tardos. How Bad is Selfish Routing? *J ACM*, 49(2):236–259, 2002.

[25] T. Roughgarden. On the Severity of Braess's Paradox: Designing Networks for Selfish Users is Hard. *JCSS*, 72(5):922–953, 2006.

[26] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[27] Y. Sharma, D. P. Williamson. Stackelberg thresholds in network routing games or the value of altruism. In *Proc. 8th EC*, 93–102, 2007.

[28] C. Swamy. The effectiveness of Stackelberg strategies and tolls for network congestion games. In *Proc. 18th SODA*, 1133–1142, 2007.

[29] J. G. Wardrop. Some Theoretical Aspects of Road Traffic Research. In *Proc. of the Institute of Civil Engineers*, Pt. II:325–378, 1952.