Selfish Routing of Splittable Flow with Respect to Maximum Congestion ⋆

Rainer Feldmann
Faculty of Computer Science, Electrical Engineering and Mathematics,
University of Paderborn,
email: obelix@upb.de

Keywords: selfish routing, traffic networks, coordination ratio

Abstract. We study the problem of selfishly routing splittable traffic with respect to maximum congestion through a shared network. Our model naturally combines features of the two best studied models in the context of selfish routing: The KP-model [11] and the Wardrop-model [20].
We are given a network with source nodes s_i, sink nodes t_i, $1 \leq i \leq k$, m edges, and a latency function for each edge. Traffics of rate r_i are destined from s_i to t_i. Traffics are splittable and each piece of traffic tries to route in such a way that it minimizes its private latency. In the absence of a central regulation, Nash Equilibria represent stable states of such a system. In a Nash Equilibrium, no piece of traffic can decrease its private latency by unilaterally changing its route. The increased social cost due to the lack of central regulation is defined in terms of the coordination ratio, i.e. the worst possible ratio of the social cost of a traffic flow at Nash Equilibrium and the social cost of a global optimal traffic flow.
In this paper, we show that in the above model pure Nash Equilibria always exist. Then, we analyze the coordination ratio of single-commodity networks with linear latency functions. Our main result is a tight upper bound of $\frac{4}{3}m$, where m is the number of edges of the network, for the coordination ratio of single-commodity networks with linear latency functions. On our way to our main result we analyze the coordination ratio of single-hop networks and show a tight upper bound of $m+\Theta(\sqrt{m})$. A more sophisticated analysis yields an upper bound of $\frac{4}{3}m$ for the coordination ratio of multi-hop networks, which is then used to derive the main result for arbitrary single-commodity linear networks.

1 Introduction

Motivation and Framework: Routing traffic through a shared network is a fundamental problem that has been studied since the early 1950s [20]. While at that time road traffic systems were the focus of interest, nowadays, routing models are re-investigated to model the behavior of communication networks like the internet [11, 16, 5, 3]. Due to the size of the networks a central regulation of flow is usually impossible. In such a scenario, users are assumed to act selfishly in that they try to optimize their individual welfare, e.g. minimize their personal delay. The individual welfare experienced by a user depends on the total traffic flow of all users in the system. A central authority would try to optimize the social welfare. A fundamental question that has recently been studied by several researchers in different routing models [11, 17, 14, 4, 6, 9, 10, 8, 12, 13] is the question of how much the social welfare suffers from the lack of a central authority.
Routing problems with selfish users are typically modeled as a game with non-cooperative agents, which choose routes through the network aiming

⋆ This work has been partially supported by the European Union within the 6th Framework Programme under contract 001907 (DELIS), and by the IST Program of the European Union under contract number IST-2001-33116 (FLAGS)

to maximize their individual welfare. One of the most important concepts in non-cooperative game theory is the concept of Nash Equilibria [15]. A Nash Equilibrium is a stable state in such a system: No user can decrease his individual welfare by unilaterally deviating from his routing strategy (a path through the network). A natural question to ask is the following: What is the largest possible ratio of the social cost of a Nash Equilibrium and the social cost of a global optimal solution? This ratio is termed coordination ratio [11], or price of anarchy [16].

In this paper we study a new model, called WMAX-model, which combines features from the Wardrop-model [20] and the KP-model [11]. In the KP-model routing networks are restricted to single-hop networks with m parallel links, traffics are unsplittable, link latencies are given as the ratio of load and link speed, private costs are defined to be the maximum (expected) latency experienced by a user, and the social cost is defined to be the maximum (expected) link latency in the network. In the Wardrop-model networks are arbitrary, traffics are splittable, link latencies are given as nondecreasing link latency functions, private costs are defined to be the sum of the link latencies on a path used, and the social cost is defined to be the sum of the link congestions.

In the WMAX-model, we allow arbitrary networks, splittable traffics, and arbitrary continuous, nonnegative, nondecreasing edge latency functions as in the Wardrop-model. We define the private cost of a piece of traffic as the maximum edge latency experienced by the traffic, and the social cost as the maximum edge congestion in the network. Another modification of the Wardrop-model has already been studied in [19, 2], where the authors consider minimizing the maximum path latency of a network. In application, it enables to identify bottlenecks in the traffic network and models selfish users that, e.g. in a road traffic system, try to avoid heavily congested roads, but instead prefer a detour along uncongested roads.

Related work: The Wardrop-model has already been studied in the 1950's [20, 1] as a model for road traffic systems. Koutsoupias and Papadimitriou [11] considered the KP-model, and defined the coordination ratio, or "price of anarchy" as termed by Papadimitriou [16].

Inspired by the work in [11], Roughgarden and Tardos re-investigated the Wardrop-model and showed that for networks with linear latency functions the coordination ratio is bounded by the constant $\frac{4}{3}$. Roughgarden [18] showed that the network design problem is computationally hard.

Since its definition the KP-model has been studied by several researchers: Mavronicolas and Spirakis [14] introduced fully mixed Nash Equilibria, i.e. mixed Nash Equilibria with strictly positive probabilities only, and proved that for identical links the worst case coordination ratio when restricted to fully mixed Nash Equilibria is $O(\frac{\log m}{\log \log m})$. Czumaj and Vöcking [4] proved a tight bound of $\Theta(\min\{\frac{\log m}{\log \log \log m}, \frac{\log m}{\log\left(\frac{\log m}{\log c_1/c_m}\right)}\})$ for the coordination ratio of the KP-model with m related links and link speeds in $[c_m, c_1]$. The combinatorial structure and computational complexity of Nash Equilibria in the KP-model were studied in [7, 9, 12]. Lücking et al. [13] defined a modification of the KP-model with quadratic social cost and showed, besides constant bounds for the coordination ratio, that for identical users and identical links the fully mixed Nash Equilibrium is the worst one. Gairing et al. [8] consider a modification, where users have restricted access to links. They present a

polynomial time algorithm which, given an arbitrary routing, computes a Nash Equilibrium with social cost not worse than that of the input.
Roughgarden [19] and Correa et al. [2] study a modification of the Wardrop-model. They define private cost as the maximum path latency in the network Roughgarden [19] proved a bound of $n-1$ for the coordination ratio of single-commodity networks with n nodes in this model. Correa et al. [2] give complexity results for this model and show that the coordination ratio is unbounded for multi-commodity networks.
Results: First, in section 2, we present the terminology used throughout this paper. We then start our analysis of the coordination rato in section 3, showing that the coordination ratio of single-hop networks with linear latency functions is bounded from above by $m+\Theta(m)$. In section 4, we show an upper bound of $\frac{4}{3}m$ for the coordination ratio of multi-hop networks with linear latency functions, which is then used in section 5 to prove an upper bound of $\frac{4}{3}m$ for the coordination ratio of arbitrary networks with linear latency functions.

2 The Model

Throughout the paper we will use the notations $\mathbb{N}$ for the natural numbers, $\mathbb{R}$ for the real numbers, and $[k] = \{1, \ldots, k\}$ for $k \in \mathbb{N}$. An instance of the WMAX-model is a tuple (G, r, l), where $G = (V, E)$ is a directed network with node set V and edge set E. $r = (r_i, s_i, t_i)_{i\in[k]}$, is a vector of k requests, i.e. traffics of rate $r_i \in \mathbb{R}_{>0}$ have to be routed from a source $s_i \in V$ to a destination $t_i \in V$. $l = (l_e(x))_{e\in E}$ is a vector of continuous, nonnegative, and nondecreasing edge latency functions $l_e(\cdot)$. We denote by $\mathcal{P}_i$ the set of paths from s_i to t_i in G. $\mathcal{P} = \cup_{i\in[k]}\mathcal{P}_i$. A flow f is a function $f : \mathcal{P} \to \mathbb{R}_{\geq 0}$, where, for a fixed flow f, we define $f_e := \sum_{P:e\in P} f_P$. A flow f is feasible, if $\sum_{P\in\mathcal{P}_i} f_P = r_i$ for all $i \in [k]$. The latency $l_P(f)$ of a path $P \in \mathcal{P}$ with respect to a flow f is defined as the maximum latency on P: $l_P(f) = \max_{e\in P}\{l_e(f_e)\}$. $L(f) := \max_{e\in E}\{l_e(f_e)\}$. Traffic routing along a path P in a flow f experiences a private latency of $l_P(f)$. The cost $c_e(f_e)$ of an edge $e \in E$ routing traffic f_e is defined as $c_e(f_e) = f_e \cdot l_e(f_e)$. The social cost $C(f)$ of a flow f is given as $C(f) = \max_{e\in E}\{c_e(f_e)\}$. $c_P(f) := \max_{e\in P}\{c_e(f_e)\}$.

Definition 1. *A flow f in G is at Nash equilibrium, if for all $i \in [k]$, $P_1, P_2 \in \mathcal{P}_i$ and all $\delta \in]0, f_{P_1}]$, we have $l_{P_1}(f) \leq l_{P_2}(\tilde{f})$, where*

$$\tilde{f}_P = \begin{cases} f_P - \delta & \text{if } P = P_1, \\ f_P + \delta & \text{if } P = P_2, \text{ and} \\ f_P & \text{if } P \notin \{P_1, P_2\}. \end{cases}$$

With continuous and nondecreasing edge latency functions, letting δ tend to zero we obtain

Lemma 1. *A flow f is at Nash equilibrium iff $\forall\ i \in [k]$, $P_1, P_2 \in \mathcal{P}_i$ with $f_{P_1} > 0$ we have $l_{P_1}(f) \leq l_{P_2}(f)$.*

Thus, in a Nash equilibrium all paths $P \in \mathcal{P}_i$ which are used by request i have equal latency, say $L_i(f)$.
It can be shown easily that in our model a Nash Equilibrium always exists. In contrast to the Wardrop-model Nash Equilibria of an instance may have different social costs.

Definition 2. *Let (G, r, l) be an instance for the* WMAX*-model. Let f^* be an optimal flow for (G, r, l). The coordination ratio of (G, r, l) is defined as*

$\rho(G,r,l) = \sup_{f\ NE}\left\{\frac{C(f)}{C(f^*)}\right\}$. *The coordination ratio of a set of instances $\mathcal{I}$ is defined as* $\rho(\mathcal{I}) = \sup_{(G,r,l)\in\mathcal{I}}\{\rho(G,r,l)\}$.

For the rest of the paper it will be guaranteed that a flow f at Nash Equilibrium *with maximum cost* $C(f)$ exists, so we will use $\max_f$ instead of $\sup_f$ in the definition of the coordination ratio. We will denote a flow f at Nash Equilibrium with maximum cost $C(f)$ a worst flow at Nash Equilibrium, or worst Nash Equilibrium, for short.

3 Coordination Ratio of Single-Hop Networks

In this section we will analyze the coordination ratio of single-hop instances (G,r,l) with m parallel links, i.e. instances with $G = (V,E)$, $V = \{s,t\}$, $E = \{e_1,\ldots,e_m\}$, $e_j = (s,t)$, a traffic rate $r \in \mathbb{R}_{>0}$ of traffic to be send from s to t, and linear latency functions $l_j(x) = a_jx + b_j$, $j \in [m]$. Such instances will be termed as single-hop linear networks. We start this section with two simple lemmas:

Lemma 2. *Let (G,r,l) be a single-hop instance, $l_j(x) = a_jx + b_j$, $j \in [m]$. Let $f = (f_1,\ldots,f_m)$ be a Nash Equilibrium and let $f^* = (f_1^*,\ldots,f_m^*)$ be a global optimum for (G,r,l). Let $f_i = \min_{e\in E}\{f_e\}$, and $f_j = \max_{e\in E}\{f_e\}$. Then, $f_i \leq f_i^*$ and $f_j^* \leq f_j$.*

Sketch of proof: Since the global optimum equalizes the costs $c_i(f_i^*)$ and $c_j(f_j^*)$, but the Nash Equilibrium does not necessarily do so, the maximum flow on a link of the global optimum must be less or equal to the maximum flow of the Nash Equilibrium on any link. Similarly, the minimum flow on a link in the global optimum cannot be smaller than the minimum one in the Nash Equilibrium. □

Lemma 3. *Let (G,r,l) with $l_j(x) = a_jx + b_j$, $j \in [m]$. Then, there exists $(G,1,\tilde{l})$, with $\rho(G,r,l) \leq \rho(G,1,\tilde{l})$.* *(Proof: see full version)*

Consequently, in what follows, we will always assume that $r = 1$. In order not to change the notation we will continue to denote instances of the Wmax-model as (G,r,l). We will now show that the coordination ratio of single-hop linear networks is bounded from above by $R(m) := \frac{2(m-1)^2}{2m-1-\sqrt{4m-3}}$, and that this bound is tight. The proof proceeds in four steps:
First, in lemma 5 we show that the coordination ratio of instances with constant latencies only, i.e. $l_j(x) = b_j$, $1 \leq j \leq m$, is bounded by m.
Then, in lemma 6 we prove that for all $m \geq 2$ there exist instances (G,r,l) achieving a coordination ratio of $R(m)$. In particular, we present instances which have $m-1$ identical constant latency functions and one proportional latency function and achieve the above bound.
In a third step, in lemma 7 we show that for an instance (G,r,l) with arbitrary linear latency functions there is an instance $(G,r,\tilde{l})$ where $\tilde{l}$ has proportional and constant latency functions only and has a coordination ratio not less than the coordination ratio of (G,r,l).
Finally, based on the result of step 3 we prove in theorem 1 that for an arbitrary instance (G,r,l) there always exists an instance $(G,r,\tilde{l})$ where $\tilde{l}$ contains a single proportional latency function $l_e(x) = ax$ and $m-1$ identical constant latency functions $l_e(x) = a$, for some $a \in \mathbb{R}_{>0}$, such that $(G,r,\tilde{l})$ has a coordination ratio not less than the coordination ratio of (G,r,l).

Definition 3. *Define the function $R:\ \mathbb{R}_{>1} \rightarrow \mathbb{R}$ as $R(m) = \frac{2(m-1)^2}{2m-1-\sqrt{4m-3}}$.*

Before we proceed with the main line of the proof we present some simple properties of the R-function. The proofs are given in the full version.

Lemma 4. *Let $R(m)$ as defined in definition 3. Then, for all $m \in \mathbb{R}_{>1}$,*
a) $R(m) = m + \frac{\sqrt{4m-3}-1}{2}$, *b)* $R(m) \leq m + \sqrt{m-1}$,
c) $R(m) \geq m + \sqrt{m-1} - 1$, *d)* $R(m) \leq \frac{4}{3}m$.

Lemma 5. *Let (G, r, l) be a single-hop instance, $l_j(x) = b_j$, $1 \leq j \leq m$. Then, $\rho(G, r, l) \leq m$.*

Proof: Let $b = \min_{1 \leq j \leq m}\{b_j\}$. Wlog we assume that $b_1 = b$. Then, $f = (1, 0, \ldots, 0)$ is a worst Nash Equilibrium for (G, r, l) with cost $C(f) = b$. On the other hand, replacing all constant latencies $l(x) = b_j > b$ with $\tilde{l}(x) = b$ cannot increase the cost of a global optimum, and f still is a Nash Equilibrium. Therefore, for the cost $C(f^*)$ of a global optimum f^* we have $C(f^*) \geq \frac{b}{m}$. Together we obtain $\rho(G, r, l) \leq \frac{b}{b/m} = m$. □

The next lemma shows, that there exist instances (G, r, l) which are worse in terms of coordination ratio compared to the instances with constant latency functions.

Definition 4. *The set of instances $\mathcal{W}$ is defined as*
$$\mathcal{W} = \{(G, r, l) \mid G = (V, E),\ V = \{s, t\},\ E = \{e_1, \ldots, e_m\},\ \text{for some } m \in \mathbb{N},\ e_j = (s, t),\ j \in [m],\ l = (ax, a, \ldots, a),\ a \in \mathbb{R}_{>0}\ \}.$$

Lemma 6. *Let $(G, r, l) \in \mathcal{W}$ be a single-hop instance with m links. Then, $\rho(G, r, l) = R(m)$.* *(Proof: see full version)*

Next, we will prove a lemma that allows us to restrict our further considerations to special instances (G, r, l):

Lemma 7. *Let (G, r, l) be a single-hop instance with m links, $l_j(x) = a_j x + b_j$, $j \in [m]$. Let f be a worst Nash Equilibrium for (G, r, l). Then, there exists a single-hop instance $(G, r, \tilde{l})$ such that*
i) *$\tilde{l}$ contains proportional or constant latency functions only,*
ii) *all constant latency functions are equal to $L(f)$,*
iii) *any worst Nash Equilibrium $\tilde{f}$ for $(G, r, \tilde{l})$ has $\max_{i \in [m]}\{\tilde{f}_i\} = \max_{i \in [m]}\{f_i\}$, and*
iv) *$\rho(G, r, l) \leq \rho(G, r, \tilde{l})$.*

Proof: Let f^* be a global optimum for (G, r, l). For a first step, define $\hat{l}_j(x) := L(f)$, if $f_j \leq f_j^*$, and $\hat{l}_j(x) := \frac{L(f)}{f_j}x$, otherwise.

Then, $\hat{l}(f_j) = L(f)$ for all $j \in [m]$. Thus, f is a Nash Equilibrium for $(G, r, \hat{l})$ with $\hat{L}(f) = L(f)$ and $\hat{C}(f) = C(f)$. It follows that for a worst Nash Equilibrium $\hat{f}$ we must have $\hat{C}(\hat{f}) \geq C(f)$. Moreover, any Nash Equilibrium $\hat{f}$ for $(G, r, \hat{l})$ will have $\hat{L}(f) = L(f)$. Now, consider the global optimum f^* for (G, r, l):

If $\hat{l}_j(x) = L(f)$, then $\hat{l}_j(x) \leq l_j(x)$ for all $x \in [f_j, \infty]$. Since $f_j^* \geq f_j$ we have $\hat{l}_j(f_j^*) \leq l_j(f_j^*)$, and thus $\hat{c}_j(f_j^*) = f_j^* \hat{l}_j(f_j^*) \leq f_j^* l_j(f_j^*) = c_j(f_j^*) = C(f^*)$. If $\hat{l}_j(x) = \frac{L(f)}{f_j}x$, then $\hat{l}_j(x) \leq l_j(x)$ for all $x \in [0, f_j]$. Since $f_j^* < f_j$ we have $\hat{l}_j(f_j^*) < l_j(f_j^*)$, and thus $\hat{c}_j(f_j^*) = f_j^* \hat{l}_j(f_j^*) < f_j^* l_j(f_j^*) = c_j(f_j^*) = C(f^*)$.

Therefore, $\hat{c}_j(f_j^*) \leq C(f^*)$ for all $j \in [m]$. It follows that a global optimum $\hat{f}^*$ for $(G, r, \hat{l})$ must have $\hat{C}(\hat{f}^*) \leq C(f^*)$. Thus, $\rho(G, r, l) = \frac{C(f)}{C(f^*)} \leq \frac{\hat{C}(\hat{f})}{\hat{C}(\hat{f}^*)} = \rho(G, r, \hat{l})$.

Let $\alpha = \max_{i\in[m]}\{f_i\}$. A worst Nash Equilibrium $\hat{f}$ for $(G, r, \hat{l})$ has $\hat{f}_j = f_j$ for all $j \in [m]$ such that $\hat{l}_j(x)$ is proportional, and load $\hat{\alpha} := \sum_{j:\hat{l}(x)=L(f)} f_j$ on a single constant link, say m. If $\hat{\alpha} > \alpha$, let $k = \lfloor \frac{\hat{\alpha}}{\alpha} \rfloor$. Then, by lemma 2, we know that $\hat{f}^*_m \le \hat{f}_m = \hat{\alpha}$, and we can replace k constant latency functions $\hat{l}(x) = L(f)$ by $\tilde{l}(x) = \frac{L(f)}{\alpha}x$. The new instance $(G, r, \tilde{l})$ has $\tilde{C}(\tilde{f}) = \alpha L(f) = C(f)$, $\tilde{C}(\tilde{f}^*) \le \hat{C}(\hat{f}^*)$ by the same argument as above, and obeys the constraints i) to iv). □

The following technical lemma will be used in the proofs of thrs. 1 and 2:

Lemma 8. *Let $k \in \mathbb{N}$. The problem $\mathcal{P}$: $\max\{\sum_{i\in[k]} \frac{1}{\sqrt{a_i}}\}$ s.t. $\sum_{i\in[k]} \frac{1}{a_i} = 1$, $a_i \in]1, \infty]$, $i \in [k]$ has the unique solution $(a_1, \ldots, a_k) = (k, \ldots, k)$ with objective value $\sqrt{k}$. (Proof: see full version)*

Now, we show that we may replace an instance (G, r, l) having at least two non-constant latency functions by an instance $(G, r, \tilde{l})$ having a single non-constant latency function without decreasing the coordination ratio.

Theorem 1. *Let (G, r, l) be a single-hop instance with m links and $l_j(x) = a_j x + b_j$, $j \in [m]$. Then, there exists an instance $(G, r, \tilde{l}) \in \mathcal{W}$ with 1 proportional link, $m - 1$ identical constant links, and $\rho(G, r, \tilde{l}) \ge \rho(G, r, l)$.*

Proof: Let $f = (f_1, \ldots, f_m)$ be a worst Nash Equilibrium for (G, r, l) with latency $L(f)$, and let $f^* = (f^*_1, \ldots, f^*_m)$ be a global optimum for (G, r, l).
By lemma 7 we may assume wlog that $l = (a_1x, \ldots, a_kx, L(f), \ldots, L(f))$, $k \in \{2, \ldots, m\}$, and $f_p > f^*_p$ for all $p \in [k]$ (otherwise, $l_p(x)$ may be replaced by $l_p(x) = L(f)$ as in the proof of lemma 7). Reordering the links gives $0 < a_1 \le a_2 \le \ldots \le a_k$. Since $a_i f_i = L(f)$ for all $i \in [k]$ and $\sum_{i\in[k]} f_i \le 1$, we get $\sum_{i\in[k]} \frac{1}{a_i} \le \frac{1}{L(f)}$. A worst Nash Equilibrium f will use at most one of the links with constant latency function $L(f)$, say $f_{k+1} > 0$. Define $\tilde{l}_1(x) = \frac{L(f)}{f_1+f_{k+1}}x$. Then, $\tilde{f} = (f_1 + f_{k+1}, f_2, \ldots, f_k, 0, \ldots, 0)$ is the unique Nash Equilibrium for $(G, r, \tilde{l})$ with $\tilde{C}(\tilde{f}) = (f_1+f_{k+1})L(f) \ge f_1L(f) = C(f)$. Since $a_1 = \frac{L(f)}{f_1} \ge \frac{L(f)}{f_1+f_{k+1}} =: \tilde{a}_1$, it follows that $\tilde{l}_1(x) \le l_1(x)$ for all $x \in [0, 1]$ which implies that for a global optimum $\tilde{f}^*$ of $(G, r, \tilde{l})$ we get $\tilde{C}(\tilde{f}^*) \le C(f^*)$. Therefore, we may now assume wlog that $l = (a_1x, \ldots, a_kx, L(f), \ldots, L(f))$ with $0 < a_1 \le a_2 \le \ldots \le a_k$, $f_i \in]0, 1]$ for all $i \in [k]$, $\sum_{i\in[k]} f_i = 1$. By scaling we may additionally assume that $L(f) = 1$. We want to show that the coordination ratio of (G, r, l) is upper bounded by the coordination ratio of $(G, r, \tilde{l})$, where $\tilde{l} = (L(f)x, L(f), \ldots, L(f))$.
$L(f) = 1$ implies that $\tilde{l} = (x, 1, \ldots, 1)$. $\tilde{f} = (1, 0, \ldots, 0)$ is the unique Nash Equilibrium for $(G, r, \tilde{l})$ with $\tilde{C}(\tilde{f}) = 1 \overset{L(f)=1}{=} \frac{L(f)f_1}{f_1} = \frac{C(f)}{f_1} \overset{\frac{1}{f_1}=a_1}{=} a_1C(f)$.
Now, let $l = (a_1x, \ldots, a_kx, 1, \ldots, 1)$ for some $k \in [m]$. Let $a_1 = \min_{i\in k}\{a_i\}$. In order to prove that $\rho(G, r, \tilde{l}) \ge \rho(G, r, l)$ we have to show that $\tilde{C}(\tilde{f}^*) \le a_1C(f^*)$, or, equivalently, $C(f^*) \ge \frac{\tilde{C}(\tilde{f}^*)}{a_1}$, for a global optimum $\tilde{f}^*$ of $(G, r, \tilde{l})$. Since the total flow to be routed equals 1, we have $1 = \sum_{i\in[m]} f^*_i = \left(\sum_{i\in[k]} \sqrt{\frac{1}{a_i}}\right) \sqrt{C(f^*)} + (m-k)C(f^*)$, and $1 = \sum_{i\in[m]} \tilde{f}^*_i = \sqrt{\tilde{C}(\tilde{f}^*)} + (m-1)\tilde{C}(\tilde{f}^*)$. By lemma 8, $\sum_{i\in[k]} \sqrt{\frac{1}{a_i}}$ attains its maximum value $\sqrt{k}$ iff all the $a_i's$ are equal to k, in which case $\sqrt{k} = \sqrt{a_1}$. Thus, from the first equation we obtain $\sqrt{a_1}\sqrt{C(f^*)} + (m-k)C(f^*) \ge 1$. Now, suppose to the

contrary that $C(f^*) < \frac{\tilde{C}(\tilde{f}^*)}{a_1}$. Then, $\sqrt{a_1}\sqrt{\frac{\tilde{C}(\tilde{f}^*)}{a_1}} + (m-k)\frac{\tilde{C}(\tilde{f}^*)}{a_1} > 1$, from which, since $a_1 \geq 1$, we deduce that $\sqrt{\tilde{C}(\tilde{f}^*)} + (m-k)\tilde{C}(\tilde{f}^*) > 1$.

But $\sqrt{\tilde{C}(\tilde{f}^*)} + (m-1)\tilde{C}(\tilde{f}^*) = 1$, which gives us a contradiction for all $k \in [m]$. Therefore, $C(f^*) \geq \frac{\tilde{C}(\tilde{f}^*)}{a_1}$. It follows that we can replace the proportional latency functions $l_1(x), \ldots, l_k(x)$ by $\tilde{l}_1(x) = L(f)x$ and $\tilde{l}_i(x) = L(f)$, $2 \leq i \leq k$, without decreasing the coordination ratio. The resulting instance $(G, r, \tilde{l})$ is in $\mathcal{W}$, and $\rho(G, r, \tilde{l}) \geq \rho(G, r, l)$. □

As a corollary we obtain a tight upper bound for the coordination ratio of single-hop networks with m links:

Corollary 1. *Let (G, r, l) with $G = (V, E)$, $V = \{s, t\}$, $E = \{e_1, \ldots, e_m\}$, $e_j = (s, t)$, $l_j(x) = a_j x + b_j$, $1 \leq j \leq m$. Then, $\rho(G, r, l) \leq R(m) = m + \Theta(\sqrt{m})$ and this bound is tight.*

4 Coordination Ratio of Multi-Hop Networks

Consider a 2-hop instance (G, r, l), $G = (V, E)$, $V = \{s, v, t\}$, $E = E_1 \cup E_2$ with $E_1 = \{e_{1,1}, \ldots, e_{1,m_1}\}$, $e_{1,j} = (s, v)$ for all $j \in [m_1]$ and $E_2 = \{e_{2,1}, \ldots, e_{2,m_2}\}$, $e_{2,j} = (v, t)$ for all $j \in [m_2]$. Let $l_{i,j}(x)$ be the latency functions for edges $e_{i,j}$, $i \in [2]$, $j \in [m_i]$.

Assume that in a Nash Equilibrium f all edges of the first hop with strictly positive flow have equal latency, say $L(f)$. Then, f may route the total flow of 1 in an arbitrary manner across the second hop, as long as the latencies $l_e(f_e)$ of the edges $e \in E_2$ do not exceed $L(f)$. If $\alpha = \max_{e \in E_1}\{f_e\}$ then $C_1(f) := \max_{e \in E_1}\{c_e(f_e)\} = \alpha L(f)$.

If, for some $e \in E_2$, $l_e(x) = L(f)$, f may route the total flow along the single edge e of the second hop, thus incurring costs $C(f) = L(f) = \frac{1}{\alpha}C_1(f)$ for some $\alpha \in]0, 1]$. But then, $\rho(G, r, l)$ in general cannot be bounded by the coordination ratio of a single-hop network consisting of the first hop only. Since $\alpha = \max_{e \in E_1}\{f_e\} \geq \frac{1}{m_1}$, a trivial upper bound for the coordination ratio is $\rho(G, r, l) \leq m_1 R(m_1) = m_1^2 + \Theta(m_1^{3/2})$.

In what follows we will show a tight upper bound of $\frac{4}{3}m$ for the coordination ratio of h-hop networks, where m is the number of edges of a cut (hop) at Nash Equilibrium.

We proceed by identifying worst case single-hop instances with bounded maximum flow α at Nash Equilibrium. These worst-case instances are then used to derive a bound of $\frac{4}{3}m$ on the coordination ratio of multi-hop instances. The proofs of the following lemmas can be found in the full version. We start with two generalizations of the R-function from definition 3 in section 3:

Definition 5. *The function $R_0 : \mathbb{R}_{>1} \times]0, 1] \to \mathbb{R}$ is defined as $R_0(m, \alpha) = \frac{2\alpha(m-k)^2}{2(m-k) - k\sqrt{4\alpha(m-k) + \alpha^2 k^2 + \alpha k^2}}$, where, for $\alpha \in]0, 1]$, $k = k(\alpha) \in \mathbb{N}$ is such that $\alpha \in]\frac{1}{k+1}, \frac{1}{k}]$.*

$R_0(m, \alpha)$ is well defined for all $\alpha \in]\frac{1}{m}, 1]$, since $k = k(\alpha) \in [m-1]$.

Lemma 9. *Let $m \in \mathbb{N}$, $\alpha \in]\frac{1}{m}, 1]$. Let $k = k(\alpha) \in \mathbb{N}$ be such that $\alpha \in]\frac{1}{k+1}, \frac{1}{k}]$.*

a) $R_0(m, \alpha) = \alpha m - \alpha k + \frac{1}{2}\alpha^2 k^2 + \alpha k\sqrt{\alpha m - \alpha k + \frac{1}{4}\alpha^2 k^2}$.

b) $R_0(m, \frac{1}{k}) = R(\frac{m}{k})$, where $R(m)$ is the function defined in definition 3.

Definition 6. *Let* $m \in \mathbb{N}$, $p \in [m-1]$, $\alpha \in]\frac{1}{m}, 1]$. *Let* $h = \sqrt{p\frac{1-\alpha}{\alpha}}$.
The functions $R_p : \mathbb{R}_{>1} \times]0,1] \to \mathbb{R}$ *are defined as*
$R_p(m,\alpha) = \frac{2\alpha(m-p-1)^2}{2(m-p-1)-(1+h)\sqrt{4\alpha(m-p-1)+\alpha^2(1+h)^2+\alpha(1+h)^2}},$ *if* $p \le m-2$, *and*
$R_{m-1}(m,\alpha) = \alpha\left(\alpha + \sqrt{(m-1)\alpha(1-\alpha)}\right)$.

Lemma 10. *Let* $m \in \mathbb{N}$, $p \in [m-2]$.

a) $R_p(m,\alpha) = \alpha(m-p-1) + \frac{1}{2}\left(\alpha + \sqrt{p\alpha(1-\alpha)}\right)^2$

$$+ (\alpha + \sqrt{p\alpha(1-\alpha)}) \cdot \sqrt{\alpha(m-p-1) + \tfrac{1}{4}(\alpha + \sqrt{p\alpha(1-\alpha)})^2}.$$

b) $R_p(m, \frac{m}{p+1}) = R(\frac{m}{p+1})$, *where* $R(m)$ *is the function defined in definition 3 in section 3.*

The next lemma shows that there are single-hop networks achieving a coordination ratio of exactly $R_0(m,\alpha)$.

Definition 7. *The set of instances* $\mathcal{W}_0$ *is defined as*

$$\mathcal{W}_0 = \{(G,r,l) \mid G=(V,E),\ V=\{s,t\},\ E=\{e_1,\dots,e_m\},\ \text{for } m \in \mathbb{N},\ e_j=(s,t),\ 1 \le j \le m,\ l=(\underbrace{ax,\dots,ax}_{k\ links},\underbrace{a\alpha,\dots,a\alpha}_{m-k\ links}),\ a \in \mathbb{R}_{>0},\ \alpha \in [\tfrac{1}{m},1],\ k=k(\alpha) \in \mathbb{N} \text{ such that } \alpha \in]\tfrac{1}{k+1},\tfrac{1}{k}]\ \}.$$

Lemma 11. *Let* $(G,r,l) \in \mathcal{W}_0$ *be an instance with* m *links and maximum load* α *in a worst Nash Equilibrium. Then,* $\rho(G,r,l) = R_0(m,\alpha)$.

Similarly, we can find single-hop networks achieving a coordination ratio equal to $R_p(m,\alpha)$:

Definition 8. *The set of instances* $\mathcal{W}_p$ *is defined as*

$$\mathcal{W}_p = \{(G,r,l) \mid G=(V,E),\ V=\{s,t\},\ E=\{e_1,\dots,e_m\},\ \text{for } m \in \mathbb{N},\ e_j=(s,t),\ 1 \le j \le m,\ l=(ax,\underbrace{bx,\dots,bx}_{p\ links},\ \underbrace{a\alpha,\dots,a\alpha}_{m-p-1\ links}),\ a \in \mathbb{R}_{>0},\ \alpha \in [\tfrac{1}{p+1},1],\ b=\tfrac{pa\alpha}{1-\alpha}\}.$$

Lemma 12. *Let* $(G,r,l) \in \mathcal{W}_p$ *be an instance with* m *links and maximum load* α *in a worst Nash Equilibrium. Then,* $\rho(G,r,l) = R_p(m,\alpha)$.

The lemma above states that there exist single-hop instances (G,r,l) achieving a coordination ratio of $\rho(G,r,l) = R_p(m,\alpha)$ for some $p \in \{0,\dots,m-1\}$. These instances have $\max_{e\in E}\{f_e\} = \alpha$ for a worst Nash Equilibrium f. In the following we will state that the ratios $\frac{R_p(m,\alpha)}{\alpha}$ are bounded from above by $\frac{4}{3}m$ for all $p \in \{0,\dots,m-1\}$.

Lemma 13. *Let* $m \in R_{>1}$, $\alpha \in]\frac{1}{m},1]$. *Then,* $R_0(m,\alpha) \le R(\alpha m)$. *Moreover,* $R_0(m,\alpha) = R(\alpha m)$ *iff* $\alpha = \frac{1}{k}$, *where* $k = k(\alpha) \in \mathbb{N}$ *such that* $\alpha \in]\frac{1}{k+1},\frac{1}{k}]$.

Lemma 14. *Let* $m \in \mathbb{N}_{\ge 2}$, $\alpha \in]\frac{1}{m},1]$. *Then,* $\frac{R(\alpha m)}{\alpha} \le \frac{4}{3}m$.

As an immediate corollary of the previous two lemmas we obtain

Corollary 2. *Let* $m \in \mathbb{N}_{\ge 2}$, $\alpha \in]\frac{1}{m},1]$. *Then,* $\frac{R_0(m,\alpha)}{\alpha} \le \frac{4}{3}m$. □

What makes the proof of corollary 2 easy is the fact that $R_0(m,\alpha) \le R(\alpha m)$. Unfortunately, this does not hold for $R_p(m,\alpha)$ and general $p \in [m-2]$. Nevertheless, in lemmas 15 and 16 we state that $\frac{R_p(m,\alpha)}{\alpha}$, can be bounded from above by $\frac{4}{3}m$ as well for all $p \in [m-1]$.

Lemma 15. *Let $m \in \mathbb{N}_{\ge 2}$, $\alpha \in [\frac{1}{m}, 1]$. Then, $\frac{R_{m-1}(m,\alpha)}{\alpha} \le m$.*

Lemma 16. *Let $m \in \mathbb{N}_{\ge 2}$, $p \in [m-2]$, $\alpha \in]\frac{1}{m}, 1]$. Then, $\frac{R_p(m,\alpha)}{\alpha} \le \frac{4}{3}m$.*

We are now ready to prove that the coordination ratio of single-hop instances with maximum load α in a worst Nash Equilibrium is bounded from above by the coordination ratio of an instance from $\mathcal{W}_p$ for some $p \in \{0,\dots,m\}$. This is done in the following

Theorem 2. *Let (G,r,l) be a single-hop instance with m links and $l_j(x) = a_j x + b_j$, $j \in [m]$. Let $f = (f_1,\dots,f_m)$ be a worst Nash Equilibrium for (G,r,l), $\alpha := \max_{e\in E}\{f_e\}$. Then, $\rho(G,r,l) \le \max_{p\in\{0,\dots,m-1\}}\{R_p(m,\alpha)\}$.*

<u>Proof:</u> By lemma 7 we may assume that $l = (a_1x,\dots,a_kx, L(f),\dots,L(f))$ for some $k \in [m]$. Let $0 < a_1 \le a_2 \le \dots \le a_k$. $f_i = \frac{L(f)}{a_i}$, $1 \le i \le k$, and $f_j \ne 0$ for at most one link $j \in \{k+1,\dots,m\}$, say $j = k+1$.
Let $\alpha = \max_{i\in[m]}\{f_i\}$. If $\alpha = f_{k+1}$, then, by lemma 2, $f^*_{k+1} \le f_{k+1}$ for any global optimum f^* for (G,r,l), and we replace $l_{k+1}(x) = L(f)$ by $\hat{l}_{k+1}(x) = \frac{L(f)}{\alpha}x$. f is still a Nash Equilibrium for $(G,r,\hat{l})$ and any global optimum $\hat{f}^*$ has $C(\hat{f}^*) \le C(f^*)$. Thus, $\rho(G,r,l) \le \rho(G,r,\hat{l})$.
Therefore, we may assume that the maximum load α in (G,r,l) is on a proportional link, say on link 1.
If $k = 1$, then $l = (a_1x, L(f),\dots,L(f))$, and $\rho(G,r,l) \le R_0(m,\alpha)$. If $k = m$, then $\rho(G,r,l) = R_{m-1}(m,\alpha) \le m$.
Now, let $k \in \{2,\dots,m-1\}$, and let $p := k-1 \in [m-2]$. Let $\beta := \frac{1}{p}\sum_{i=2}^k f_i$. Since $k \ge 2$, $\beta \ne 0$. Since $\alpha = \max_{i\in[m]}\{f_i\}$, $\beta \le \alpha$. $f_{k+1} < \alpha$ and by a simple scaling argument we can assume wlog that $L(f) = 1$.
Let $b = \frac{1}{\beta}$ and define $\tilde{l} = (a_1x, \underbrace{bx,\dots,bx}_{p\ links}, \underbrace{1,\dots,1}_{m-p-1\ links})$.
Then, $\tilde{f} = (\alpha,\beta,\dots,\beta,f_{k+1},0,\dots,0)$ is a Nash Equilibrium for $(G,r,\tilde{l})$ with $\tilde{C}(\tilde{f}) = \alpha L(f)$.
Let $\tilde{f}^*$ be a global optimum for $(G,r,\tilde{l})$ and suppose to the contrary that $\tilde{C}(\tilde{f}^*) > C(f^*)$.
Since $1 = \sum_{i\in[m]} f^*_i = \left(\frac{1}{\sqrt{a_1}} + \sum_{i=2}^k \frac{1}{\sqrt{a_i}}\right)\sqrt{C(f^*)} + (m-p-1)C(f^*)$ and $1 = \sum_{i\in[m]} \tilde{f}^*_i = \left(\frac{1}{\sqrt{a_1}} + (k-1)\frac{1}{\sqrt{b}}\right)\sqrt{\tilde{C}(\tilde{f}^*)} + (m-p-1)\tilde{C}(\tilde{f}^*)$, $C(\tilde{f}^*) > C(f^*)$ implies that $\frac{1}{\sqrt{a_1}} + \sum_{i=2}^k \frac{1}{\sqrt{a_i}} > \frac{1}{\sqrt{a_1}} + (k-1)\frac{1}{\sqrt{b}}$, which is a contradiction to lemma 8. It follows that $C(\tilde{f}^*) \le C(f^*)$ and thus $\rho(G,r,l) \le \frac{\tilde{C}(\tilde{f})}{\tilde{C}(\tilde{f}^*)} = \rho(G,r,\tilde{l}) = R_p(m,\alpha) \le \max_{p\in\{0,\dots,m-1\}}\{R_p(m,\alpha)\}$.
□

Theorem 2 implies that for an arbitrary single-hop network (G,r,l) with maximum load α there exists a number of links $p \in \{0,\dots,m-1\}$ such that $\rho(G,r,l) \le R_p(m,\alpha)$. We have already seen, that we can bound the fraction $\frac{R_p(m,\alpha)}{\alpha}$ by $\frac{4}{3}m$ for all p. This fact is now used in the following theorem, which gives a tight upper bound on the coordination ratio of multi-hop networks.

Theorem 3. *Let (G,r,l) be an h-hop instance, $V = \{v_0, v_1, \ldots, v_h\}$, $E = E_1 \dot\cup \ \ldots \dot\cup \ E_h$, $E_i = \{e_{i,j} \mid 1 \leq j \leq m_j\}$, $1 \leq i \leq h$, $e_{i,j} = (v_{i-1}, v_i)$, $1 \leq i \leq h$, $1 \leq j \leq m_i$, $s = v_0$, $t = v_h$, $l_{e_{i,j}}(x) = l_{i,j}(x) = a_{i,j}x + b_{i,j}$, $1 \leq i \leq h$, $1 \leq j \leq m_i$. Then $\rho(G,r,l) \leq \frac{4}{3}|E|$.*
Furthermore, there exist instances (G,r,l) such that $\rho(G,r,l) = \frac{4}{3}(|E| - 2)$.

Proof: Let $m := |E|$. For an admissible flow g denote $C_i(g) = \max_{e \in E_i}\{c_e(g_e)\}$, $1 \leq i \leq h$. Let $f = (f_{i,j})$ be a worst flow at Nash Equilibrium for (G,r,l). Then, there exists a hop $q \in [h]$ of G and a number $L(f) \in \mathbb{R}_{\geq 0}$ such that the following conditions hold:
a) $l_{i,j}(f_{i,j}) \leq L(f) \ \ \forall i \in \{1,\ldots,h\},\ j \in \{1,\ldots,m_i\},\ f_{i,j} > 0.$
b) $l_{q,j}(f_{q,j}) = L(f) \ \ \forall j \in \{1,\ldots,m_x\},\ f_{q,j} > 0.$
c) $l_{q,j}(f_{q,j}) \geq L(f) \ \ \forall j \in \{1,\ldots,m_x\},\ f_{q,j} = 0.$

It follows that $C(f) = \max_{e\in E}\{c_e(f_e)\} = \max_{e \in E}\{f_e l_e(f_e)\} \overset{f_e \leq 1}{\leq} L(f)$.
Let $f^* = (f^*_{i,j})$ be a global optimum for (G,r,l). Wlog let f^* be such that f^* is a global optimal flow at every single-hop $i \in [h]$. Let $\alpha \in [\frac{1}{m_p}, 1]$ be the maximum load of a link of hop q at Nash Equilibrium f.
If $\alpha = \frac{1}{m_q}$ then all latency functions at hop q must be equal, from which it follows that $C_q(f^*) = C_q(f) = \frac{1}{m_q}L(f)$. Then, $\rho(G,r,l) \leq \frac{L(f)}{C(f^*)} \leq \frac{L(f)}{C_q(f^*)} = m_q \frac{C_q(f)}{C_q(f^*)} = m_q \leq m$.
Now, let $\alpha > \frac{1}{m_q}$. By thr. 2, we have $\frac{C_q(f)}{C_q(f^*)} \leq \max_{p\in\{0,\ldots,m-1\}}\{R_p(m_q,\alpha)\}$ from which we get the lower bound

$$C_q(f^*) \geq \frac{C_q(f)}{\max_{p\in\{0,\ldots,m-1\}}\{R_p(m_q,\alpha)\}} = \frac{\alpha L(f)}{\max_{p\in\{0,\ldots,m-1\}}\{R_p(m_q,\alpha)\}}.$$

It follows that $\rho(G,r,l) \leq \frac{L(f)}{C_q(f^*)} \leq L(f)\frac{\max_p\{R_p(m_q,\alpha)\}}{\alpha L(f)} = \frac{\max_p\{R_p(m_q,\alpha)\}}{\alpha}$.
By corollary 2 and lemmas 15, 16, it follows that $\rho(G,r,l) \leq \frac{4}{3}m_q \leq \frac{4}{3}m$.
To see that the lower bound of $\frac{4}{3}(m-2)$ holds, let m be any natural number which is divisible by 3, and consider the following 2-hop instance:
$G = (V,E)$, $V = \{s,v,t\}$, $E = E_1 \dot\cup E_2$, $E_1 = \{e_1,\ldots,e_m\}$, $e_j = (s,v)$, $1 \leq j \leq m$, $E_2 = \{e_{m+1}, e_{m+2}\}$, $e_j = (v,t)$, $m+1 \leq j \leq m+2$. Let $\alpha = \frac{3}{m}$, $k = \frac{m}{3}$ and define $l = (\underbrace{x,\ldots,x}_{k\ links}, \underbrace{\alpha,\ldots,\alpha}_{m-k\ links}, \underbrace{\alpha, 0}_{hop\ 2})$
Then, $f = (\alpha,\ldots,\alpha,0,\ldots,0,1,0)$ is a worst Nash Equilibrium for (G,r,l) with $C_1(f) = \alpha^2$, and $C_2(f) = \alpha \geq \alpha^2 = C_1(f)$.
Let f^* be a global optimal solution, and wlog let f^* route optimally at every hop of G. Then, $C_2(f^*) = 0$. Since for the first hop we have $\frac{C_1(f)}{C_1(f^*)} \overset{L.11}{=} R_0(m,\alpha)$, a global optimal solution f^* has $C_1(f^*) = \frac{C_1(f)}{R_0(m,\alpha)} = \frac{\alpha^2}{R_0(m,\alpha)}$.
From this we get $\rho(G,r,l) = \frac{C(f)}{C(f^*)} = \frac{\alpha}{\frac{\alpha^2}{R_0(m,\alpha)}} = \frac{R_0(m,\alpha)}{\alpha} = \frac{R(\alpha m)}{\alpha} = \frac{4}{3}m = \frac{4}{3}(m-2)$, where the fourth equality follows from Lemma 13 ($\alpha = \frac{1}{k}$) and the fifth equality follows from $\alpha = \frac{3}{m}$ and $R(\alpha m) = R(3) = 4$. □
From the proof of theorem 3 we can see that instances almost matching the upper bound can already be found in the set of 2-hop networks.

5 Coordination Ratio of Linear Networks

In this section, we prove the main result of our paper, namely an upper bound of $\frac{4}{3}m$ for the coordination ratio of arbitrary single-commodity networks with linear latency functions.

Theorem 4. *Let (G, r, l) be an arbitrary instance with a single commodity network $G = (V, E)$, and linear latency functions $l_e(x) = a_e x + b_e$ for all $e \in E$. Then, $\rho(G, r, l) \leq \frac{4}{3}|E|$.*

Sketch of proof: Let $m := |E|$. Let us denote by a cut (S, T) of G a partition of the nodes V of G into two disjoint subsets S, T, such that $s \in S$ and $t \in T$. Let $E_{S,T}$ denote the cut edges of (S, T) directed from the set S into the set T.
Let $f = (f_e)_{e \in E}$ be a Nash Equilibrium. Then, there exists a cut (S, T) of G, $E_{S,T} = \{e_1, \ldots, e_p\}$ such that $l_e(f_e) \geq L(f)$ for all $e \in E_{S,T}$ and equality holds iff $f_e > 0$. Similarly to the result for single-hop networks in lemma 3 we may assume wlog that $r = 1$.
In a first case, let us assume that there exists an $e \in E_{S,T}$ with $c_e(f_e) = C(f)$. In this case consider a single-hop network $(\tilde{G}, r, \tilde{l})$ where $\tilde{G} = (\tilde{V}, \tilde{E})$ is defined as $\tilde{V} = \{s, t\}$, $\tilde{E} = \{e_1, \ldots, e_p\}$, $e_i = (s, t)$ for all $i \in [p]$, $\tilde{l}_i(x) := l_i(x)$, $i \in [p]$, where $l_i(x)$ is the latency function of the edge $e_i \in E_{S,T}$ in (G, r, l). Then, $\tilde{f}$ defined as $\tilde{f}_e := f_e$ for all $e \in \tilde{E} = E_{S,T}$ confirms a flow at Nash Equilibrium for $(\tilde{G}, r, \tilde{l})$ with $\tilde{C}(\tilde{f}) = \max_{e \in E_{S,T}}\{c_e(f_e)\} = C(f)$. A global optimal flow $\tilde{f}^*$ for $(\tilde{G}, r, \tilde{l})$ must have $\tilde{C}(\tilde{f}^*) = \max_{e \in \tilde{E}}\{c_e(\tilde{f}^*_e)\} \leq \max_{e \in E_{S,T}}\{c_e(f^*_e)\} = C(f^*)$. Therefore, by corollary 1 and Lemma 4 d) $\rho(G, r, l) = \frac{C(f)}{C(f^*)} \leq \frac{\tilde{C}(\tilde{f})}{\tilde{C}(\tilde{f}^*)} \leq R_0(m, 1) \leq \frac{4}{3}m$.
In the second case, assume that for all $e \in E_{S,T}$ we have $c_e(f_e) < C(f)$. Let (S_0, T_0) be a cut such that E_{S_0,T_0} contains an edge e_0 with $c_{e_0}(f_{e_0}) = C(f)$. Let $q = |E_{S_0,T_0}|$. Now consider a 2-hop network $(\tilde{G}, r, \tilde{l})$ where $\tilde{G} = (\tilde{V}, \tilde{E})$ is defined as $\tilde{V} = \{s, v, t\}$, $\tilde{E} = \tilde{E}_1 \cup \tilde{E}_2$, $\tilde{E}_1 = \{e_{1,1}, \ldots, e_{1,p}\}$, $\tilde{E}_2 = \{e_{2,1}, \ldots, e_{2,q}\}$, $e_{1,i} = (s, v)$ for all $i \in [p]$, $e_{2,i} = (v, t)$ for all $i \in [q]$. $\tilde{l}_{1,i}(x) := l_i(x)$, $1 \leq i \leq p$, where $l_i(x)$ is the latency function of the edge $e_i \in E_{S,T}$, and $\tilde{l}_{2,j}(x) := l_j(x)$, $1 \leq j \leq q$, where $l_j(x)$ is the latency function of the edge $e_j \in E_{S_0,T_0}$.
Define $\tilde{f}$ by setting $\tilde{f}_e := f_e$ for all $e \in \tilde{E}$. Then, $\tilde{l}_e(\tilde{f}_e) = l_e(f_e) \geq L(f)$ for all $e \in \tilde{E}_1$, and equality holds iff $f_e > 0$. For the seond hop, $\tilde{l}_e(\tilde{f}_e) \leq \max_{e \in E_{S_0,T_0}}\{l_e(f_e) \mid f_e > 0\} = L(f)$ for all $e \in \tilde{E}_2$. Therefore, all paths from s to t are latency-blocked with respect to $L(f)$ and f is a Nash Equilibrium. For the cost $\tilde{C}(\tilde{f})$ of $\tilde{f}$ we obtain $\tilde{C}(\tilde{f}) = \tilde{c}_{e_0}(f_{e_0}) = c_{e_0}(f_{e_0}) = C(f)$. On the other hand, a global optimal flow $\tilde{f}^*$ for $(\tilde{G}, r, \tilde{l})$ must have $\tilde{C}(\tilde{f}^*) = \max_{e \in \tilde{E}}\{\tilde{c}_e(\tilde{f}^*_e)\} \leq \max_{e \in E_{S,T} \cup E_{S_0,T_0}}\{c_e(f^*_e)\} \leq \max_{e \in E}\{c_e(f^*_e)\} = C(f^*)$. Therefore, by theorem 3, $\rho(G, r, l) = \frac{C(f)}{C(f^*)} \leq \frac{\tilde{C}(\tilde{f})}{\tilde{C}(\tilde{f}^*)} \leq \frac{4}{3}m$. □

6 Conclusion

We have analyzed the coordination ratio of the WMAX-model, a model in which selfish traffic routes through a shared network aiming to minimize the maximum latency on a link used, while central optimization is done aiming to minimize the maximum congestion in the whole network. First, we have shown that Nash Equilibria always exist in the WMAX-model. Then, we have proved that the coordination ratio for single-hop networks with linear edge latency functions is bounded from above by $m + \Theta(m)$, and that this bound is tight. For multi-hop networks, as well as for arbitrary networks, we have shown a tight upper bound of $\frac{4}{3}m$.
We leave it as an open problem to generalize our results to networks with nonlinear latency functions and to multi-commodity instances.

Theorem 3. *Let (G,r,l) be an h-hop instance, $V = \{v_0, v_1, \ldots, v_h\}$, $E = E_1 \dot\cup \ \ldots \dot\cup \ E_h$, $E_i = \{e_{i,j} \mid 1 \le j \le m_j\}$, $1 \le i \le h$, $e_{i,j} = (v_{i-1}, v_i)$, $1 \le i \le h$, $1 \le j \le m_i$, $s = v_0$, $t = v_h$, $l_{e_{i,j}}(x) = l_{i,j}(x) = a_{i,j}x + b_{i,j}$, $1 \le i \le h$, $1 \le j \le m_i$. Then $\rho(G,r,l) \le \frac{4}{3}|E|$.*
Furthermore, there exist instances (G,r,l) such that $\rho(G,r,l) = \frac{4}{3}(|E|-2)$.

<u>Proof:</u> Let $m := |E|$. For an admissible flow g denote $C_i(g) = \max_{e \in E_i}\{c_e(g_e)\}$, $1 \le i \le h$. Let $f = (f_{i,j})$ be a worst flow at Nash Equilibrium for (G,r,l). Then, there exists a hop $q \in [h]$ of G and a number $L(f) \in \mathbb{R}_{\ge 0}$ such that the following conditions hold:
 a) $l_{i,j}(f_{i,j}) \le L(f) \ \ \forall i \in \{1,\ldots,h\},\ j \in \{1,\ldots,m_i\},\ f_{i,j} > 0.$
 b) $l_{q,j}(f_{q,j}) = L(f) \ \ \forall j \in \{1,\ldots,m_x\},\ f_{q,j} > 0.$
 c) $l_{q,j}(f_{q,j}) \ge L(f) \ \ \forall j \in \{1,\ldots,m_x\},\ f_{q,j} = 0.$

It follows that $C(f) = \max_{e\in E}\{c_e(f_e)\} = \max_{e\in E}\{f_e l_e(f_e)\} \overset{f_e \le 1}{\le} L(f)$.
Let $f^* = (f^*_{i,j})$ be a global optimum for (G,r,l). Wlog let f^* be such that f^* is a global optimal flow at every single-hop $i \in [h]$. Let $\alpha \in [\frac{1}{m_p}, 1]$ be the maximum load of a link of hop q at Nash Equilibrium f.
If $\alpha = \frac{1}{m_q}$ then all latency functions at hop q must be equal, from which it follows that $C_q(f^*) = C_q(f) = \frac{1}{m_q}L(f)$. Then, $\rho(G,r,l) \le \frac{L(f)}{C(f^*)} \le \frac{L(f)}{C_q(f^*)} = m_q \frac{C_q(f)}{C_q(f^*)} = m_q \le m$.
Now, let $\alpha > \frac{1}{m_q}$. By thr. 2, we have $\frac{C_q(f)}{C_q(f^*)} \le \max_{p\in\{0,\ldots,m-1\}}\{R_p(m_q,\alpha)\}$ from which we get the lower bound

$$C_q(f^*) \ge \frac{C_q(f)}{\max_{p\in\{0,\ldots,m-1\}}\{R_p(m_q,\alpha)\}} = \frac{\alpha L(f)}{\max_{p\in\{0,\ldots,m-1\}}\{R_p(m_q,\alpha)\}}.$$

It follows that $\rho(G,r,l) \le \frac{L(f)}{C_q(f^*)} \le L(f)\frac{\max_p\{R_p(m_q,\alpha)\}}{\alpha L(f)} = \frac{\max_p\{R_p(m_q,\alpha)\}}{\alpha}$.
By corollary 2 and lemmas 15, 16, it follows that $\rho(G,r,l) \le \frac{4}{3}m_q \le \frac{4}{3}m$.
To see that the lower bound of $\frac{4}{3}(m-2)$ holds, let m be any natural number which is divisible by 3, and consider the following 2-hop instance:
$G = (V,E)$, $V = \{s,v,t\}$, $E = E_1 \dot\cup E_2$, $E_1 = \{e_1,\ldots,e_m\}$, $e_j = (s,v)$, $1 \le j \le m$, $E_2 = \{e_{m+1}, e_{m+2}\}$, $e_j = (v,t)$, $m+1 \le j \le m+2$. Let $\alpha = \frac{3}{m}$, $k = \frac{m}{3}$ and define $l = (\underbrace{x,\ldots,x}_{k\ links}, \underbrace{\alpha,\ldots,\alpha}_{m-k\ links}, \underbrace{\alpha, 0}_{hop\ 2})$
Then, $f = (\alpha,\ldots,\alpha,0,\ldots,0,1,0)$ is a worst Nash Equilibrium for (G,r,l) with $C_1(f) = \alpha^2$, and $C_2(f) = \alpha \ge \alpha^2 = C_1(f)$.
Let f^* be a global optimal solution, and wlog let f^* route optimally at every hop of G. Then, $C_2(f^*) = 0$. Since for the first hop we have $\frac{C_1(f)}{C_1(f^*)} \overset{L.11}{=} R_0(m,\alpha)$, a global optimal solution f^* has $C_1(f^*) = \frac{C_1(f)}{R_0(m,\alpha)} = \frac{\alpha^2}{R_0(m,\alpha)}$. From this we get $\rho(G,r,l) = \frac{C(f)}{C(f^*)} = \frac{\alpha}{\frac{\alpha^2}{R_0(m,\alpha)}} = \frac{R_0(m,\alpha)}{\alpha} = \frac{R(\alpha m)}{\alpha} = \frac{4}{3}m = \frac{4}{3}(m-2)$, where the fourth equality follows from Lemma 13 ($\alpha = \frac{1}{k}$) and the fifth equality follows from $\alpha = \frac{3}{m}$ and $R(\alpha m) = R(3) = 4$. □
From the proof of theorem 3 we can see that instances almost matching the upper bound can already be found in the set of 2-hop networks.

5 Coordination Ratio of Linear Networks

In this section, we prove the main result of our paper, namely an upper bound of $\frac{4}{3}m$ for the coordination ratio of arbitrary single-commodity networks with linear latency functions.