



The Price of Anarchy for Instantaneous Dynamic Equilibria

Lukas Graf^(✉)  and Tobias Harks

Institute of Mathematics, Augsburg University, 86135 Augsburg, Germany
{lukas.graf,tobias.harks}@math.uni-augsburg.de

Abstract. We consider flows over time within the deterministic queueing model and study the solution concept of instantaneous dynamic equilibrium (IDE) in which flow particles select at every decision point a currently shortest path. The length of such a path is measured by the physical travel time plus the time spent in queues. Although IDE have been studied since the eighties, the efficiency of the solution concept is not well understood. We study the price of anarchy for this model and show an upper bound of order $\mathcal{O}(U \cdot \tau)$ for single-sink instances, where U denotes the total inflow volume and τ the sum of edge travel times. We complement this upper bound with a family of quite complex instances proving a lower bound of order $\Omega(U \cdot \log \tau)$.

Keywords: Dynamic flows · Flows over time · Price of anarchy

1 Introduction

Dynamic flows have gained substantial interest over the last decades in modeling dynamic network systems such as urban traffic or the Internet. A widely used model for describing dynamic flows is based on the fluid queueing model due to Vickrey [23]. There is a directed graph $G = (V, E)$, where edges $e \in E$ are associated with a queue with positive rate capacity $\nu_e \in \mathbb{R}_+$ and a physical transit time $\tau_e \in \mathbb{R}_+$. If the total inflow into an edge $e = vw \in E$ exceeds the rate capacity ν_e , a queue builds up and arriving flow particles need to wait in the queue before they are forwarded along the edge. The total travel time along e is thus composed of the waiting time spent in the queue plus the physical transit time τ_e .

Due to the decentralized nature of the above mentioned applications, the physical flow model needs to be complemented by a *behavioral model* prescribing the actions of flow particles. Most works in the transportation science literature as well as recent works in the mathematics and computer science literature adopt the *full information model*, i.e., all flow particles have complete information on the state of the network for all points in time (including the future evolution of all flow particles) and based on this information travel along a shortest path. This leads to the concept of *dynamic equilibrium* (Nash equilibrium) and has been analyzed in the transportation science literature for decades, see

Friesz et al. [7], Meunier and Wagner [18], Zhu and Marcotte [24] and the more recent works by Koch and Skutella [16] and Cominetti, Correa and Larré [3]. The full information assumption has been justified by assuming that the game is played repeatedly and a dynamic equilibrium is then an attractor of a learning process. In light of the wide-spread use of navigation devices, this concept may not be completely realistic anymore, because drivers are informed in real-time about the current traffic situation and, if beneficial, reroute instantaneously no matter how good or bad that route was in hindsight. This aspect is also discussed in Marcotte et al. [17], Hamdouch et al. [12] and Unnikrishnan and Waller [22].

Instead of the (classical) dynamic equilibrium, we consider in this paper *instantaneous dynamic equilibria (IDE)*, where for every point in time and at every decision node, flow only enters those edges that lie on a currently shortest path towards the respective sink. This concept assumes far less information (only the network-wide queue length which are continuously measured) and leads to a distributed dynamic using only present information that is readily available via real-time information. IDE have been proposed already in the late 80's (cf. Ran and Boyce [19, § VII-IX], Boyce, Ran and LeBlanc [2, 20], Friesz et al. [8]) and it is known that IDE do exist under quite general conditions, see Graf, Harks and Sering [11].

Price of Anarchy. In comparison to dynamic equilibrium, an IDE flow behaves quite differently and several fundamental aspects of IDE are not well understood. There are, for instance, simple single-source single-sink instances in which the unique IDE flow exhibits cycling behavior, that is, some flow particles travel along cycles before they reach the sink. This behavior is impossible for dynamic equilibria as every particle chooses a path once and never gets into a cycle. This raises the question of the (time) price of anarchy of IDE flows.

Question (PoA): Assuming single-sink instances with constant inflow rates for a finite time interval, what is the maximum time needed so that every flow particle reaches the sink?¹

1.1 Our Results and Proof Techniques

We study the termination time of IDE flows for single-sink instances and derive the first quantitative upper bound on the termination of IDE flows. Our bound is parameterized in the numbers U and τ denoting the total flow volume injected into the sources and the sum of physical travel times, respectively. We denote by $\text{PoA}(U, \tau)$ the price of anarchy over the family of instances parameterized by U and τ .

Theorem 1: For multi-source single-sink networks, any IDE flow over time terminates after at most $\mathcal{O}(U\tau)$ time. Moreover, $\text{PoA}(U, \tau) \in \mathcal{O}(U\tau)$.

¹ For multi-sink instances, it is known that IDE flows may cycle forever, thus, the termination time and the PoA is infinity in this case.

We prove this bound by first deriving a general termination bound for acyclic graphs. Using this bound, we then show that there exist so-called sink-like subgraphs that can effectively be treated as acyclic graphs. This way, we can argue that at all times a sufficiently large flow volume enters the current sink-like subgraph and, by the bound for acyclic graphs, reaches the sink within the claimed time. The proof technique and the bound itself are completely different to those for dynamic equilibria in [4].

We then turn to lower bounds on the termination time (PoA) of IDE flows.

Theorem 4: For $(U, \tau) \in \mathbb{N}^* \times \mathbb{N}^*$ with $U \geq 2\tau$, we have $\text{PoA}(U, \tau) \in \Omega(U \log \tau)$.

The lower bound is based on a quite complex instance (see Fig. 2) that works roughly as follows. We combine two gadgets: A “cycling gadget” consisting of a large cycle made of edges with capacity $\approx U$ and a “blocking gadget” consisting of paths with low capacity and length of about τ connecting the nodes on the cycle to the sink node. An IDE flow within this graph can then alternate between two different phases: A “charging phase”, wherein the main amount of flow travels once around the big cycle, loosing a small amount of flow to each of the paths leading towards the sink, and a “blocking phase”, in which the particles traveling along the paths form queues again and again in just the right way as to keep the main amount of flow traveling around on the large cycle without loosing any more flow. In order to derive a lower bound on the price of anarchy we then augment this instance in such a way that the optimal flow can just bypass the two gadgets and reach the sink in constant time while any IDE flow gets diverted into the cycling gadget.

1.2 Further Related Work

The concept of flows over time was studied by Ford and Fulkerson [5]. Shortly after, Vickrey [23] introduced a game-theoretic variant using a deterministic queueing model. Since then, dynamic equilibria have been studied extensively in the transportation science literature, see Friesz et al. [8]. New interest in this model was raised after Koch and Skutella [16] gave a novel characterization of dynamic equilibria in terms of a family of static flows (thin flows). Cominetti, Correa and Omar [3] refined this characterization and Koch and Sering [21] incorporated spillbacks in the fluid queueing model.

Regarding the price of anarchy of dynamic equilibria, Koch and Skutella [16] derived the first results on the price of anarchy for dynamic equilibria, which were recently improved by Correa, Cristi and Oosterwijk [4] devising a tight bound of $\frac{e}{e-1}$, provided that a certain monotonicity conjecture holds. Israel and Sering [14] investigated the price of anarchy for the model with spillbacks. Bhaskar, Fleischer and Anshelevich [1] devised Stackelberg strategies in order to improve the efficiency of dynamic equilibria. Recently, Frascaria and Olver [6] considered a flexible departure choice model from an optimization point of view

and derived insights into devising tolls for improving the performance of dynamic equilibria.

Ismaili [13] considered a discrete version of IDEs and investigated the price of anarchy. He used the utilitarian social cost (not the makespan as we do) and derived lower bounds of order $\Omega(|V| + n)$ for the setting that only simple paths are allowed. Here n denotes the number of discrete players in the game. For general multi-commodity instances allowing also cycles, he proves that the price of anarchy is unbounded. Similarly, Graf, Harks and Sering [11] showed that for the continuous version multi-commodity IDE flows may cycle forever and, thus, the price of anarchy is infinity. For IDE flows in single-sink networks, on the other hand, they showed that termination is always guaranteed. However, due to the non-constructive nature of their proof they could not derive any explicit bound on the termination time or the price of anarchy for those instances.

2 Model

Let $\mathcal{N} = (G, (\nu_e)_{e \in E}, (\tau_e)_{e \in E}, (u_v)_{v \in V \setminus \{t\}}, t)$ be a network consisting of a directed graph $G = (V, E)$, edge capacities $\nu_e \in \mathbb{N}^*$, edge travel times $\tau_e \in \mathbb{N}^*$,² a single sink node $t \in V$ reachable from every other node and for every node $v \in V \setminus \{t\}$ a corresponding integrable (network) inflow rate $u_v : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$. The idea then is that, at all times $\theta \in \mathbb{R}_{\geq 0}$ infinitesimal small agents enter the network at node v at a rate according to $u_v(\theta)$ and start traveling through the graph towards the common sink t . Such a dynamic can be described by a *flow over time*, a tuple $f = (f^+, f^-)$ where $f^+, f^- : E \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ are integrable functions. For any edge $e \in E$ and time $\theta \in \mathbb{R}_{\geq 0}$ the value $f_e^+(\theta)$ describes the (*edge*) *inflow rate* into e at time θ and $f_e^-(\theta)$ is the (*edge*) *outflow rate* from e at time θ .

For any such flow over time f we define the *cumulative (edge) in- and outflow rates* F^+ and F^- as

$$F_e^+(\theta) := \int_0^\theta f_e^+(\zeta) d\zeta \quad \text{and} \quad F_e^-(\theta) := \int_0^\theta f_e^-(\zeta) d\zeta,$$

respectively. The queue length of edge e at time θ is then defined as

$$q_e(\theta) := F_e^+(\theta) - F_e^-(\theta + \tau_e). \tag{1}$$

We call such a flow f a *feasible flow* for a given set of network inflow rates $u_v : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ for each node $v \in V \setminus \{t\}$, if it satisfies the following constraints (2) to (5). The *flow conservation constraints* are modeled for all nodes $v \neq t$ as

$$\sum_{e \in \delta_v^+} f_e^+(\theta) - \sum_{e \in \delta_v^-} f_e^-(\theta) = u_v(\theta) \quad \text{for all } \theta \in \mathbb{R}_{\geq 0}, \tag{2}$$

² Throughout this paper we will restrict ourselves to integer travel times and edge capacities to make the statements and proofs cleaner. However, all results can be easily applied to instances with rational travel times and capacities by simply rescaling the instance appropriately. Note, however, that all bounds will then scale accordingly.

where $\delta_v^+ := \{vu \in E\}$ and $\delta_v^- := \{uv \in E\}$ are the sets of outgoing edges from v and incoming edges into v , respectively. For the sink node t we require

$$\sum_{e \in \delta_t^+} f_e^+(\theta) - \sum_{e \in \delta_t^-} f_e^-(\theta) \leq 0 \tag{3}$$

and for all edges $e \in E$ we always assume

$$f_e^-(\theta) = 0 \text{ f.a. } \theta < \tau_e. \tag{4}$$

Finally we assume that the queues operate at capacity which can be modeled by

$$f_e^-(\theta + \tau_e) = \begin{cases} \nu_e, & \text{if } q_e(\theta) > 0 \\ \min\{f_e^+(\theta), \nu_e\}, & \text{if } q_e(\theta) \leq 0 \end{cases} \text{ for all } e \in E, \theta \in \mathbb{R}_{\geq 0}. \tag{5}$$

Termination Time for Flows over Time. We will now introduce some additional notation in order to formally define the termination time of a feasible flow. Since termination is only relevant for flows with finitely lasting inflow rates, from here on we will always assume that there exists some time θ_0 , such that the supports of all network inflow rates u_v are contained in $[0, \theta_0]$.

Following [21], for any feasible flow f and every edge $e \in E$ we define the *edge load* function F_e^Δ that gives us for any time θ the total amount of flow currently on edge e (either waiting in its queue or traveling along the edge):

$$F_e^\Delta : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}, \theta \mapsto F_e^+(\theta) - F_e^-(\theta).$$

The function $F^\Delta(\theta) := \sum_{e \in E} F_e^\Delta(\theta)$ then gives the *total amount of flow in the network at time θ* . It is a straightforward calculation to show that after θ_0 the function F^Δ is monotonically decreasing.

Lemma 1. *Let f be a feasible flow. Then for all $\theta_2 \geq \theta_1 \geq \theta_0$, we have $F^\Delta(\theta_2) \leq F^\Delta(\theta_1)$. In particular, for $\hat{\theta} \geq \theta_0$ with $F^\Delta(\hat{\theta}) = 0$, we have $F^\Delta(\hat{\theta}) = 0$ for all $\theta \geq \hat{\theta}$.*

This motivates the following definition of termination time.

Definition 1. A feasible flow over time f *terminates* if it satisfies

$$\inf \{ \theta \geq \theta_0 \mid F^\Delta(\theta) = 0 \} < \infty.$$

We then say that $\Theta := \inf \{ \theta \geq \theta_0 \mid F^\Delta(\theta) = 0 \}$ is the *termination time* of f or f *terminates* by time Θ . Lemma 1 then implies $F^\Delta(\theta) = 0$ for all $\theta > \Theta$.

IDE Flows and their PoA. Following [11] we define an IDE flow as a feasible flow with the property that whenever a particle arrives at a node $v \neq t$, it can only enter an edge that is the first edge on a currently shortest v - t path. Here, the *current* or *instantaneous travel time* is defined for any edge e and time θ as

$$c_e(\theta) := \tau_e + \frac{q_e(\theta)}{\nu_e}. \tag{6}$$

We then define time dependent node labels $\ell_v(\theta)$ corresponding to current shortest path distances from v to the sink t . For the sink t we set $\ell_t(\theta) = 0$ for all times $\theta \in \mathbb{R}_{\geq 0}$, while for all other nodes $v \in V \setminus \{t\}$ and $\theta \in \mathbb{R}_{\geq 0}$ we recursively define $\ell_v(\theta) = \min_{e=vw \in E} \{\ell_w(\theta) + c_e(\theta)\}$. We say that an edge $e = vw$ is *active* at time θ , if $\ell_v(\theta) = \ell_w(\theta) + c_e(\theta)$ and we denote the set of active edges by $E_\theta \subseteq E$. We call a v - t path P an *active v - t path at time θ* , if all edges of P are active for i at θ or, equivalently, $\sum_{e \in P} c_e(\theta) = \ell_v(\theta)$. For differentiation we call paths that are minimal with respect to the transit times τ *physical shortest paths*.

Definition 2. A feasible flow over time f is an *instantaneous dynamic equilibrium (IDE)*, if for all $\theta \in \mathbb{R}_{\geq 0}$ and $e \in E$ it satisfies

$$f_e^+(\theta) > 0 \Rightarrow e \in E_\theta. \tag{7}$$

Since in an IDE flow particles act selfishly and without cooperation we should expect that the termination time of an IDE flow is not optimal. To quantify this difference between termination times of IDE flows and optimal flows we will use the price of anarchy, which we define as follows: For any instance $\mathcal{N} = (G, (\nu_e)_{e \in E}, (\tau_e)_{e \in E}, (u_v)_{v \in V \setminus \{t\}}, t)$ we define the worst case termination time of an IDE flow in \mathcal{N} as

$$\Theta_{\text{IDE}}(\mathcal{N}) := \sup \{ \Theta \text{ termination time of } f \mid f \text{ an IDE flow in } \mathcal{N} \}$$

and the optimal termination time in \mathcal{N} as

$$\Theta_{\text{OPT}}(\mathcal{N}) := \inf \{ \Theta \text{ termination time of } f \mid f \text{ a feasible flow in } \mathcal{N} \}.$$

Definition 3. For any pair of whole numbers (U, τ) we define the *Price of Anarchy (PoA)* for instances with total flow volume U and total edge length τ as

$$\text{PoA}(U, \tau) := \sup \left\{ \frac{\Theta_{\text{IDE}}(\mathcal{N})}{\Theta_{\text{OPT}}(\mathcal{N})} \mid \mathcal{N} \text{ s.th. } \sum_{v \in V \setminus \{t\}} \int_0^{\theta_0} u_v(\theta) d\theta = U, \sum_{e \in E} \tau_e = \tau \right\}.$$

Remark 1. At first it might seem strange that the PoA depends only on U and τ while being independent of the capacities ν_e . However, this is only the case here because we always assume that all capacities are at least 1 throughout this paper. In order to transfer our results to networks with arbitrary capacities one has to rescale the network and, in particular, replace U by $\frac{1}{\nu_{\min}} U$, where $\nu_{\min} := \min \{ \nu_e \mid e \in E \}$.

3 Upper Bounds

In this section we will show an upper bound for the termination time of IDE flows in terms of $\tau(G) := \sum_{e \in E} \tau_e$ and $U := \sum_{v \in V \setminus \{t\}} \int_0^{\theta_0} u_v(\theta) d\theta$. From this we can then derive an upper bound for the PoA. However, before we can turn to these general termination results we first have to look at acyclic graphs and

give a termination bound for *all* feasible flows in such networks in terms of U and $\tau(P_{\max})$, where the latter denotes the physical length of a longest v - t path. Even though the bound may seem obvious its proof requires a quite lengthy and careful analysis which can be found in the full version of this paper [10].

Lemma 2. *In an acyclic network, every feasible flow over time terminates before $\theta_0 + \tau(P_{\max}) + U$.*

Similarly to the proof of termination in [11, Theorem 4.6] we will apply our result for feasible flows in acyclic graphs to IDE flows in general graphs by using the fact ([11, Lemma 4.4]) that, whenever the total flow volume in a subgraph is small enough, only the physically shortest paths in this subgraph can be active. Since these edges form an acyclic subgraph, for an IDE flow we can then apply Lemma 2. For the following proof we will look at a particular type of subgraph, which we will call a sink-like subgraph: a subgraph containing all physically shortest paths from its nodes towards the sink, with a sufficiently low flow volume at the beginning of some interval as well as a low inflow into this subgraph over the course of said interval.

Definition 4. An induced subgraph $T \subseteq G$ is a *sink-like subgraph* on an interval $[\theta_1, \theta_2]$ with $\theta_1 \geq \theta_0$ if the following two properties hold:

- For each node $v \in V(T)$ all physically shortest v - t paths are contained in T .
- T satisfies $\text{vol}_T(\theta_1, \theta_2) := \sum_{e \in E(T)} F_e^\Delta(\theta_1) + \sum_{e \in \delta_T^-} \int_{\theta_1}^{\theta_2} f_e^-(\theta) d\theta < \frac{1}{2}$.

Here $\delta_T^- := \{vw \in E \mid v \notin V(T), w \in V(T)\}$ denotes the set of edges entering the subgraph T . Using [11, Lemma 4.4] we can show that inside a sink-like subgraph only physically shortest paths towards the sink can be active (see [10] for the full proof).

Lemma 3. *Let T be a sink-like subgraph on an interval $[\theta_1, \theta_2]$. Then during this interval only physically shortest paths towards t can be active.*

Together with Lemmas 1 and 2 this implies that any IDE flow will terminate once the whole graph is sink-like.

Corollary 1. *Let f be an IDE flow and $\tilde{\theta} \geq \theta_0$ such that the whole graph G is sink-like at time $\tilde{\theta}$. Then, the flow terminates before $\tilde{\theta} + \tau(P_{\max}) + \frac{1}{2}$.*

To get an upper bound on the termination time of an IDE flow it now suffices to find a large enough time horizon such that it contains at least one point in time where the whole graph is sink-like. To determine such a time, we first show that if we have a sink-like subgraph over a sufficiently long period of time, we can extend this subgraph to a larger sink-like subgraph over a slightly smaller subinterval. Note, that the proof of [11, Theorem 4.6] uses a similar strategy, but is non-constructive and, therefore, only establishes the existence of a termination time without revealing anything about the length of this time. Thus, a more thorough analysis is needed here.

Lemma 4. Let $T \subsetneq G$ be an induced subgraph, v the closest node to t not in T and T' the subgraph of G induced by $V(T) \cup \{v\}$. Let θ_1 be some time after θ_0 , $\theta_2 := \theta_1 + \sum_{e \in E \setminus E(T)} (\tau_e + \frac{1}{2\nu_e})$ and $\theta'_2 := \theta_1 + \sum_{e \in E \setminus E(T')} (\tau_e + \frac{1}{2\nu_e})$.

If T is sink-like on $[\theta_1, \theta_2]$, then T' is a sink-like subgraph on $[\theta_1, \theta'_2]$.

Proof. Since it is clear that T' fulfills the first property of being sink-like (by the choice of v), we only need to show that $\text{vol}_{T'}(\theta_1, \theta'_2) \leq \text{vol}_T(\theta_1, \theta_2)$, from which the lemma follows immediately (as T is sink-like on $[\theta_1, \theta_2]$). More precisely we will show that the flow volume on edges between v and T (i.e. edges in $E(T') \setminus E(T) = (\delta_v^+ \cap \delta_T^-) \cup (\delta_v^- \cap \delta_T^+)$) at time θ_1 as well as the inflow into v over the interval $[\theta_1, \theta'_2]$ is already accounted for by the inflow into T on the interval $[\theta_1, \theta_2]$ via edges from v to T . This is formalized in the following three claims:

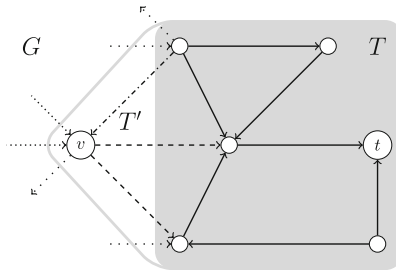


Fig. 1. A sink-like subgraph T and a closest node $v \in V \setminus V(T)$ as in the statement of Lemma 4. By Claim 1 all flow on the dash-dotted edge from T to v will reach v before time θ_v . By Claim 2 between θ_1 and θ_v all flow reaching v (either via the dotted or via the dash-dotted edges) will travel towards T from there (i.e. enter one of the dashed edges). By Claim 3 the dashed edges will never carry a larger flow volume than $\frac{1}{2}$ between θ_1 and θ_v and all flow particles using these edges within this time interval will reach T before θ_2 .

Claim 1. All flow on edges from T to v (dash-dotted edges in Fig. 1) at time θ_1 reaches v before $\theta_v := \theta_1 + \sum_{e \in E \setminus E(T')} (\tau_e + \frac{1}{2\nu_e}) + \sum_{e \in \delta_v^- \cap \delta_T^+} (\tau_e + \frac{1}{2\nu_e})$, i.e.

$$F_e^\Delta(\theta_1) \leq \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta \text{ for all } e \in \delta_T^+ \cap \delta_v^-.$$

Claim 2. All flow reaching v (from T or $G \setminus T'$, i.e. via the dash-dotted or via the dotted edges in Fig. 1) between θ_1 and θ_v will enter an edge towards T (dashed edges in Fig. 1), i.e.

$$\sum_{e \in \delta_v^-} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta = \sum_{e \in \delta_v^+ \cap \delta_T^-} \int_{\theta_1}^{\theta_v} f_e^+(\theta) d\theta.$$

Claim 3. For any edge from v to T (dashed edges in Fig. 1) the total amount of flow currently traveling on this edge at any time $\theta \in [\theta_1, \theta_v]$ is less than $\frac{1}{2}$, i.e.

$$F_e^\Delta(\theta) < \frac{1}{2} \text{ for all } e \in \delta_v^+ \cap \delta_T^- \text{ and } \theta \in [\theta_1, \theta_v].$$

Additionally all this flow will reach T before θ_2 , i.e.

$$F_e^\Delta(\theta_1) + \int_{\theta_1}^{\theta_v} f_e^+(\theta) d\theta \leq \int_{\theta_1}^{\theta_2} f_e^-(\theta) d\theta \text{ for all } e \in \delta_v^+ \cap \delta_T^-.$$

From Claims 1 to 3 we then directly get

$$\begin{aligned} & \sum_{e \in E(T') \setminus E(T)} F_e^\Delta(\theta_1) + \sum_{e \in \delta_{T'}^-} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta \\ &= \sum_{e \in \delta_T^+ \cap \delta_v^-} F_e^\Delta(\theta_1) + \sum_{e \in \delta_v^+ \cap \delta_T^-} F_e^\Delta(\theta_1) + \sum_{e \in \delta_T^- \setminus \delta_v^+} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta + \sum_{e \in \delta_v^- \setminus \delta_T^+} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta \\ &\stackrel{\text{Cl. 1}}{\leq} \sum_{e \in \delta_T^+ \cap \delta_v^-} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta + \sum_{e \in \delta_v^+ \cap \delta_T^-} F_e^\Delta(\theta_1) + \sum_{e \in \delta_T^- \setminus \delta_v^+} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta + \sum_{e \in \delta_v^- \setminus \delta_T^+} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta \\ &\stackrel{\text{Cl. 2}}{=} \sum_{e \in \delta_v^+ \cap \delta_T^-} F_e^\Delta(\theta_1) + \sum_{e \in \delta_T^- \setminus \delta_v^+} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta + \sum_{e \in \delta_v^- \cap \delta_T^+} \int_{\theta_1}^{\theta_v} f_e^+(\theta) d\theta \\ &\stackrel{\text{Cl. 3}}{\leq} \sum_{e \in \delta_T^- \setminus \delta_v^+} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta + \sum_{e \in \delta_v^+ \cap \delta_T^-} \int_{\theta_1}^{\theta_2} f_e^-(\theta) d\theta, \end{aligned}$$

implying $\text{vol}_{T'}(\theta_1, \theta'_2) \leq \text{vol}_{T'}(\theta_1, \theta_v) = \sum_{e \in E(T')} F_e^\Delta(\theta_1) + \sum_{e \in \delta_{T'}^-} \int_{\theta_1}^{\theta_v} f_e^-(\theta) d\theta \leq \sum_{e \in E(T)} F_e^\Delta(\theta_1) + \sum_{e \in \delta_T^-} \int_{\theta_1}^{\theta_2} f_e^-(\theta) d\theta = \text{vol}_T(\theta_1, \theta_2) < \frac{1}{2}$. This shows that T' is indeed sink-like on $[\theta_1, \theta'_2]$. The proofs of the claims are relatively straightforward calculations (see [10]) using [11, Lemma 4.4] and a strengthened version of [11, Lemma 4.2]. Note, that the proofs have to be done in reverse order, as the proof of Claim 2 uses Claim 3 and the proof of Claim 1 uses both Claims 2 and 3. \square

Theorem 1. For multi-source single-sink networks, any IDE flow over time terminates before $\hat{\theta} := \theta_0 + 2U \sum_{e \in E} (\tau_e + \frac{1}{2\nu_e}) + \tau(P_{\max}) + \frac{1}{2}$.

Proof. Starting with the subgraph consisting only of the sink node t (which trivially contains all shortest paths towards t) and iteratively applying Lemma 4 we immediately get

Claim 4. If, after time θ_0 , the sink node t has a total cumulative inflow of less than $\frac{1}{2}$ for some interval of length $\sum_{e \in E} (\tau_e + \frac{1}{2\nu_e})$, then the whole graph is sink-like at the beginning of this interval. \blacksquare

Since all flow reaching t vanishes from the network there can be at most $2U$ (pairwise disjoint) intervals of length $\sum_{e \in E} (\tau_e + \frac{1}{2\nu_e})$ with inflow of at least $\frac{1}{2}$ into t . Thus, there must be some time $\tilde{\theta} \leq 2U \sum_{e \in E} (\tau_e + \frac{1}{2\nu_e})$ which is the beginning of an interval of length $\sum_{e \in E} (\tau_e + \frac{1}{2\nu_e})$ with total inflow of less than $\frac{1}{2}$ into t . So, by Claim 4, the whole graph is sink-like at $\tilde{\theta}$, which, by Claim 1, implies that the flow terminates before $\tilde{\theta} + \tau(P_{\max}) + \frac{1}{2} \leq \hat{\theta}$. \square

Remark 2. This means that for any single-sink network any IDE flow terminates within $\mathcal{O}(U\tau(G))$.

Since Θ_{OPT} is trivially bounded below by $\theta_0 + 1$ this immediately leads to the following upper bound on the PoA for IDE flows:

Theorem 2. *For any pair of integers (U, τ) we have $\text{PoA}(U, \tau) \in \mathcal{O}(U\tau)$.* \square

4 Lower Bounds on the Termination Time of IDE Flows

It is easy to see that a general bound for the termination time cannot be better than $\mathcal{O}(U + \tau(G))$, since any feasible flow in the network consisting of one source node with an inflow rate of U over the interval $[0, 1]$, one sink node and a single edge between the two nodes with capacity 1 and some travel time τ terminates by $U + \tau(G)$. In the following, we will construct a family of instances, parameterized by $K, L \in \mathbb{N}^*$, that provide a lower bound on the termination time in single-sink networks of order $\Omega(U \cdot \log(\tau(G)))$ – which is strictly larger than $\mathcal{O}(U + \tau(G))$.

For any given pair of positive integers $K, L \in \mathbb{N}^*$ the instance is of the form sketched in Fig. 2, with u_{3K+1} as source node and t as its sink. The graph has a “width” (i.e. length of the horizontal paths from u_1 to u_{3K+1}) of $\approx 3^{K+1}$ and a “height” (length of the vertical paths from nodes u_i, v_i and w_i to t) of $\approx K3^{K+1}$. All edges on the horizontal path (including the one edge back to u_1) have a capacity of $2U$ with $U \approx L3^{K+1}$, while all the edges on the vertical paths have capacities of either 1 or 3.

If we let flow enter at node u_{3K+1} at a rate of $2U$ over the interval $[-0.5, 0]$, we will observe the following behavior: At first, all flow enters the direct downwards path towards the sink t until a queue of length 1 has built up on the first edge of this path. After that almost all flow will enter the edge towards u_1 and some flow will go downwards to keep the queue length constant. Assuming $U \gg 1$, most of the flow will travel towards u_1 and arrive there one time step later with a slightly lower inflow rate and over a slightly shorter interval as at u_{3K+1} . At u_1 the same flow split happens again: First all flow enters the edge to u'_1 until a queue of sufficient length to induce a waiting time of 1 has formed, then most of flow travels towards the next node v_1 . Similarly, at all the following nodes on the horizontal path, this pattern repeats, i.e., a small amount of flow (of volume ≈ 3) starts traveling downwards while most of the flow is diverted further to the right. Finally the main block of flow arrives at u_{3K+1} and is diverted back to u_1 (having lost a total volume of $\approx 3^{K+1}$ to the downwards paths). By the time this

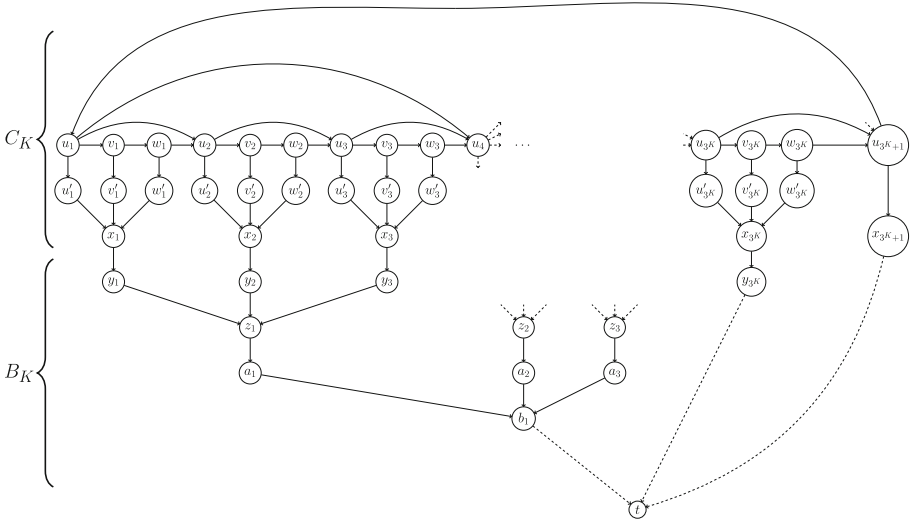


Fig. 2. A network with a total edge travel time of $\approx 3^{2K}$ that, given an inflow volume of $\approx L3^{K+1}$, has a termination time of more than $KL3^{K+1}$.

flow arrives at u_1 again, the flow particles that traveled along the edges $u_1u'_1$, $v_1v'_1$ and $w_1w'_1$ join up again at node x_1 in such a way that they form a queue of length ≈ 3 on the edge y_1z_1 . This queue is long enough to divert the main block of flow away towards u_2 (over the direct edge u_1u_2 of length 3). This pattern, again, repeats at all subsequent nodes u_i until the main flow finally arrives at node $u_{3\kappa+1}$ (having lost no additional flow volume) and is again diverted back to u_1 . This time, the flow particles from the queues on the edges y_1z_1, y_2z_1 and y_3z_1 met at node z_1 and now form a queue of length ≈ 9 on edge a_1b_1 . Thus, our main flow can now be diverted away directly to node u_4 and so on. This way, the main amount of flow can travel along the horizontal path for $\approx K$ times without losing a significant amount of flow until all the flow on the vertical paths finally reaches the sink t . After that the pattern described until now repeats. Thus, flow remains in the network until at least time $\approx 3^{K+1}K \frac{U}{3^{K+1}} = 3^{K+1}KL$.

Theorem 3. *Given any pair $K, L \in \mathbb{N}^*$, there exists an instance $G_{K,L}$ with $\tau(G_{K,L}) \in \mathcal{O}(3^{2K})$ and $U_{K,L} \in \mathcal{O}(L3^K)$ such that there exists an IDE flow that does not terminate before $LK(3^{K+1} + 1)$.*

Proof. The detailed construction as well as the formal proof of its correctness can be found in [10]. □

In order to derive a lower bound on the price of anarchy from Theorem 3, we will slightly modify the instance used there in such a way that the termination time of the worst case IDE flow remains approximately the same, while there exists an optimal termination time in $\mathcal{O}(1)$.

Theorem 4. For any pair of positive integers (U, τ) satisfying $U \geq 2\tau$, we have $\text{PoA}(U, \tau) \in \Omega(U \log \tau)$.

Proof. Wlog assume that there exist positive integers $L \geq 3^K$ such that $U = (L + 3^K)3^K$ and $\tau = 3^{2K}$. Then we take the graph $G_{K,L}$ constructed in the proof of Theorem 3 and modify it by adding two new vertices s and v and four new edges as indicating in Fig. 3.

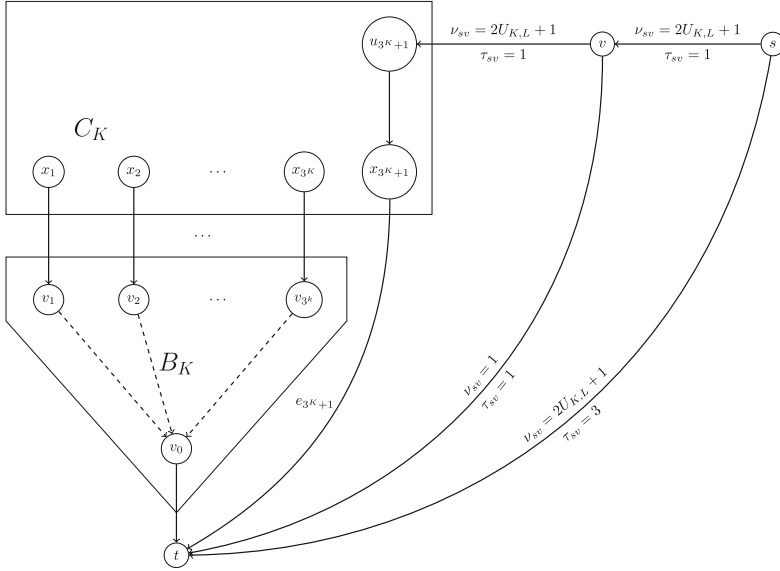


Fig. 3. The instance from the proof of Theorem 3 (on the left), two additional nodes s and v as well as four new edges.

We now use s as our new source node and a constant inflow rate of $2U_{K,L} + 1$ over the interval $[-0.5 - \varepsilon, 0]$, where $\varepsilon := \frac{4 + \tau_{e_{3^{K+1}}}}{2U_{K,L}} \leq \frac{1}{2}$. The optimal termination time of the resulting network \mathcal{N} is then at most 3, since we can just send all flow on the direct edge from s to t , where the last particle arrives 3 time steps later.

In an IDE flow, however, all flow particles travel first to v (as the path $s-v-t$ is shorter than the path $s-t$), where they enter the edge towards t . This continues until a queue of length $\tau_{e_{3^{K+1}}} + 3$ has built up at edge vt and the edge towards $u_{3^{K+1}}$ becomes active, which happens at time $\theta = 0.5$. From there on, flow splits between the two edges, entering edge vt at a rate of 1 and edge $vu_{3^{K+1}}$ at a rate of $2U_{K,L}$ throughout the interval $[0.5, 1]$. Thus, over the interval $[1.5, 2]$ flow arrives at node $u_{3^{K+1}}$ at a rate of $2U_{K,L}$. Continuing with the flow described before the statement of Theorem 3 is then again an IDE flow. This shows that $\Theta_{\text{IDE}}(\mathcal{N}) \geq LK(3^{K+1} + 1)$.

From $\tau(\mathcal{N}) = \tau(G_{K,L}) + 4 \in \mathcal{O}(3^{2K})$ and $U_{\mathcal{N}} = \left(\frac{1}{2} + \varepsilon\right)(2U_{K,L} + 1) \in \mathcal{O}(L3^K + 3^{2K}) = \mathcal{O}(L3^K)$ we get $3^{2K} \in \Omega(\tau(\mathcal{N}))$ and $L3^K \in \Omega(U_{\mathcal{N}})$, which implies $\Theta_{\text{IDE}}(\mathcal{N}) \geq LK(3^K + 1) \in \Omega(U_{\mathcal{N}} \log \tau(\mathcal{N}))$. Thus, in particular, $\text{PoA}(U, \tau) \geq \frac{\Theta_{\text{IDE}}(\mathcal{N})}{\Theta_{\text{OPT}}(\mathcal{N})} \in \Omega(U \log \tau)$. \square

Remark 3. Expanding the network constructed in the proof of Theorem 4 into an acyclic network results in an instance with constant optimal termination time, but IDE termination time of $\tau(P_{\max}) \gg \tau(P_{\min})$, where $\tau(P_{\min})$ is the physical length of a shortest path from the source to the sink node.

Together with the upper bound from Lemma 2 this implies the following bounds for the IDE price of anarchy for acyclic networks:

$$\text{PoA}|_{\text{acyclic}} \in \Omega(\tau(P_{\max})) \cap \mathcal{O}(U + \tau(P_{\max})).$$

5 Conclusions and Open Questions

We studied the efficiency of IDE flows and derived the first upper and lower bounds on the time price of anarchy of IDE flows. These bounds are of order $\mathcal{O}(U\tau)$ and $\Omega(U \log \tau)$, respectively. Comparing these bounds to the constant bound of $\frac{e}{e-1}$ for dynamic equilibria (cf. Correa, Cristi and Oosterwijk [4]) shows in some sense a “price of shortsightedness”. While instantaneous dynamic equilibria may be significantly less efficient than dynamic equilibria, in many situations this might be a price one has to pay as the full information needed for dynamic equilibria might just not be available.

Generally, it would be interesting to test the different equilibria on real instances and see how their efficiency compares there. A large-scale computational study seems more feasible for IDEs compared to dynamic equilibria, as already calculating a single α -extension is much more difficult for the full information model, while it is easy for IDE flows using a simple water-filling procedure. Indeed, for calculating a single extension phase the only positive result for dynamic equilibria is based on a recent work of Kaiser [15] showing that for series-parallel graphs a single phase can be computed in polynomial time. The question of whether a finite number of such extensions is enough to compute a complete equilibrium flow is still open for dynamic equilibria, while we were able to answer this question positively for IDE flows in an upcoming paper [9].

Acknowledgments. We thank Kathrin Gimmi for sparking the initial idea that allowed us to prove the explicit upper bound on the termination time of IDE flows. We are also grateful to the anonymous reviewers for their valuable feedback. Finally, we thank the Deutsche Forschungsgemeinschaft (DFG) for their financial support.

References

1. Bhaskar, U., Fleischer, L., Anshelevich, E.: A Stackelberg strategy for routing flow over time. *Games Econom. Behav.* **92**, 232–247 (2015)

2. Boyce, D.E., Ran, B., LeBlanc, L.J.: Solving an instantaneous dynamic user-optimal route choice model. *Transp. Sci.* **29**(2), 128–142 (1995)
3. Cominetti, R., Correa, J.R., Larré, O.: Dynamic equilibria in fluid queueing networks. *Oper. Res.* **63**(1), 21–34 (2015)
4. Correa, J., Cristi, A., Oosterwijk, T.: On the price of anarchy for flows over time. In: *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, New York, NY, USA*, pp. 559–577 (2019). Association for Computing Machinery
5. Ford, L.R., Fulkerson, D.R.: *Flows in Networks*. Princeton University Press, Princeton (1962)
6. Frascaria, D., Olver, N.: Algorithms for flows over time with scheduling costs. arXiv (2019). <https://arxiv.org/abs/1912.00082>, to appear in IPCO 2020
7. Friesz, T.L., Bernstein, D., Smith, T.E., Tobin, R.L., Wie, B.-W.: A variational inequality formulation of the dynamic network user equilibrium problem. *Oper. Res.* **41**(1), 179–191 (1993)
8. Friesz, T.L., Luque, J., Tobin, R.L., Wie, B.-W.: Dynamic network traffic assignment considered as a continuous time optimal control problem. *Oper. Res.* **37**(6), 893–901 (1989)
9. Graf, L., Harks, T.: A finite time combinatorial algorithm for instantaneous dynamic equilibrium flows (2020). <https://arxiv.org/abs/2007.07808>. (working paper)
10. Graf, L., Harks, T.: The price of anarchy for instantaneous dynamic equilibrium flows. <https://arxiv.org/abs/2007.07794> (2020)
11. Graf, L., Harks, T., Sering, L.: Dynamic flows with adaptive route choice. *Math. Program.*, 309–335 (2020). <https://doi.org/10.1007/s10107-020-01504-2>
12. Hamdouch, Y., Marcotte, P., Nguyen, S.: A strategic model for dynamic traffic assignment. *Netw. Spat. Econ.* **4**(3), 291–315 (2004)
13. Ismaili, A.: Routing games over time with FIFO policy. In: Devanur, N.R., Lu, P. (eds.) *WINE 2017*. LNCS, vol. 10660, pp. 266–280. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-71924-5_19. There is also a version available on arXiv: <http://arxiv.org/abs/1709.09484>
14. Israel, J., Sering, L.: The impact of spillback on the price of anarchy for flows over time. In: Harks, T., Klimm, M. (eds.) *SAGT 2020*. LNCS, vol. 12283, pp. 114–129. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-57980-7_8
15. Kaiser, M.: Computation of dynamic equilibria in series-parallel networks. *Math. Oper. Res.* (2020, forthcoming)
16. Koch, R., Skutella, M.: Nash equilibria and the price of anarchy for flows over time. *Theory Comput. Syst.* **49**(1), 71–97 (2011)
17. Marcotte, P., Nguyen, S., Schoeb, A.: A strategic flow model of traffic assignment in static capacitated networks. *Oper. Res.* **52**(2), 191–212 (2004)
18. Meunier, F., Wagner, N.: Equilibrium results for dynamic congestion games. *Transp. Sci.* **44**(4), 524–536 (2010)
19. Ran, B., Boyce, D.E.: *Dynamic Urban Transportation Network Models: Theory and Implications for Intelligent Vehicle-highway Systems*. LNE. Springer, Berlin (1996)
20. Ran, B., Boyce, D.E., LeBlanc, L.J.: A new class of instantaneous dynamic user-optimal traffic assignment models. *Oper. Res.* **41**(1), 192–202 (1993)
21. Sering, L., Vargas-Koch, L.: Nash flows over time with spillback. In: *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM (2019)
22. Unnikrishnan, A., Waller, S.: User equilibrium with recourse. *Netw. Spat. Econ.* **9**(4), 575–593 (2009)

23. Vickrey, W.S.: Congestion theory and transport investment. *Am. Econ. Rev.* **59**(2), 251–60 (1969)
24. Zhu, D., Marcotte, P.: On the existence of solutions to the dynamic user equilibrium problem. *Transp. Sci.* **34**(4), 402–414 (2000)