FULL LENGTH PAPER

Series B



# A finite time combinatorial algorithm for instantaneous dynamic equilibrium flows

Lukas Graf<sup>1</sup> · Tobias Harks<sup>1</sup>

Received: 22 April 2021 / Accepted: 5 January 2022 / Published online: 10 February 2022 © The Author(s) 2022

# Abstract

Instantaneous dynamic equilibrium (IDE) is a standard game-theoretic concept in dynamic traffic assignment in which individual flow particles myopically select en route currently shortest paths towards their destination. We analyze IDE within the Vickrey bottleneck model, where current travel times along a path consist of the physical travel times plus the sum of waiting times in all the queues along a path. Although IDE have been studied for decades, several fundamental questions regarding equilibrium computation and complexity are not well understood. In particular, all existence results and computational methods are based on fixed-point theorems and numerical discretization schemes and no exact finite time algorithm for equilibrium computation is known to date. As our main result we show that a natural extension algorithm needs only finitely many phases to converge leading to the first finite time combinatorial algorithm computing an IDE. We complement this result by several hardness results showing that computing IDE with natural properties is NP-hard.

**Mathematics Subject Classification** 05C21 (flows in graphs)  $\cdot$  90C35 (Programming involving graphs or networks)  $\cdot$  68Q25 (Analysis of algorithms and problem complexity)

 Lukas Graf lukas.graf@math.uni-augsburg.de
 Tobias Harks tobias.harks@math.uni-augsburg.de

Lukas Graf and Tobias Harks: The research of the authors was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - HA 8041/1-1 and HA 8041/4-1. Parts of the results of this paper will appear in less detailed form in the Proceedings of the 22nd Conference on Integer Programming and Combinatorial Optimization, 2021.

<sup>&</sup>lt;sup>1</sup> Institute of Mathematics, Augsburg University, Augsburg 86135, Germany

# **1** Introduction

Flows over time or dynamic flows are an important mathematical concept in network flow problems with many real world applications such as dynamic traffic assignment, production systems and communication networks (e.g., the Internet). In such applications, flow particles that are sent over an edge require a certain amount of time to travel through each edge and when routing decisions are being made, the dynamic flow propagation leads to later effects in other parts of the network. A key characteristic of such applications, especially in traffic assignment, is that the network edges have a limited flow capacity which, when exceeded, leads to congestion. This phenomenon can be captured by the *fluid queueing model* due to Vickrey [25]. The model is based on a directed graph G = (V, E), where every edge *e* has an associated physical transit time  $\tau_e \in \mathbb{R}_+$  and a maximal rate capacity  $\nu_e \in \mathbb{R}_+$ . If flow enters an edge with higher rate than its capacity the excess particles start to form a queue at the edge's tail, where they wait until they can be forwarded onto the edge (cf. Fig. 1). Thus, the total travel time experienced by a single particle traversing an edge *e* is the sum of the time spent waiting in the queue of *e* and the physical transit time  $\tau_e$ .

This physical flow model then needs to be enhanced with a *behavioral model* prescribing the actions of flow particles. There are two main standard behavioral models in the traffic assignment literature known as dynamic equilibrium (DE) (cf. Ran and Boyce [20, § V-VI]) and instantaneous dynamic equilibrium (IDE) ([20, § VII-[X]). Under DE, 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. The full information assumption is usually justified by assuming that the game is played repeatedly and a DE is then an attractor of a learning process. The behavioral model of IDE is based on the idea that drivers are informed in real-time about the current traffic situation and, if beneficial, reroute instantaneously no matter how good or bad that route will be in hindsight. Thus, at 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 lengths which are continuously measured) and leads to a distributed dynamic using only present information that is readily available via real-time information. IDE has been proposed already in the late 80's (cf. Boyce, Ran and LeBlanc [1,21] and Friesz et al. [8]).

A line of fairly recent works starting with Koch and Skutella [17] and Cominetti, Correa and Larré [3] derived very elegant combinatorial characterizations of DE for the fluid queueing model of Vickrey. They derived a complementarity description of DE flows via so-called *thin flows with resetting* which leads to an  $\alpha$ -extension property

Fig. 1 An edge e = vw. As the inflow rate at node v exceeds the edge's capacity, a queue forms at its tail



stating that for any equilibrium up to time  $\theta$ , there exists  $\alpha > 0$  so that the equilibrium can be extended to time  $\theta + \alpha$ . An extension that is maximal with respect to  $\alpha$  is called a *phase* in the construction of an equilibrium and the existence of equilibria on the whole  $\mathbb{R}_+$  then follows by a limit argument over the phases. In the same spirit, Graf, Harks and Sering [10] established a similar characterization for IDE flows and also derived an  $\alpha$ -extension property.

For both models (DE or IDE), it is an open question whether for constant inflow rates and a finite time horizon, a *finite number of phases* suffices to construct an equilibrium, see [3,10,17]. This problem remains even unresolved for single-source single-sink series-parallel graphs as explicitly mentioned by Kaiser [16]. Proving finiteness of the number of phases would imply an exact finite time algorithm. Such an algorithm is not known to date neither for DE nor for IDE.<sup>1</sup> More generally, the computational complexity of equilibrium computation is widely open.

### 1.1 Our contribution and proof techniques

In this paper, we study IDE flows and derive algorithmic and computational complexity results. As our main result we settle the key question regarding finiteness of the  $\alpha$ -extension algorithm.

Theorem 3.7: For single-sink networks with piecewise constant inflow rates for a finite time horizon, there is an  $\alpha$ -extension algorithm computing an IDE after finitely many extension phases. This implies the first finite time combinatorial exact algorithm computing IDE within the Vickrey model.

The proof of our result is based on the following ideas. We first consider the case of *acyclic* networks and use a topological order of vertices in order to schedule the extension phases in the algorithm. The key argument for the finiteness of the number of extension phases is that for a single node v and any interval with linearly changing distance labels of nodes closer to the sink and constant inflow rate into v this flow can be redistributed to the outgoing edges in a finite number of phases of constant outflow rates from v. We show this using the properties (derivatives) of suitable edge label functions for the outgoing edges (see the graph in Fig. 3). The overall finiteness of the algorithm follows by induction over the nodes and time. We then generalize to arbitrary single-sink networks by considering *dynamically* changing topological orders depending on the current set of active edges. Finally, a closer inspection of the proofs also enables us to give an explicit upper bound on the number of extension steps in the order of

$$\mathcal{O}\left(P\left(2(\Delta+1)^{4^{\Delta}+1}\right)^{2L\cdot|E|\cdot|V|\cdot T/\tau_{\min}^{2}}\right)$$

<sup>&</sup>lt;sup>1</sup> Algorithms for DE or IDE computation used in the transportation science literature are *numerical*, that is, only approximate equilibrium flows are computed given a certain numerical precision, see the related work for a more detailed comparison.

where *P* is the number of constant phases of the network inflow rates,  $\Delta$  the maximum out degree in the network, *T* the termination time, *L* an upper bound on the absolute value of the derivatives of the distance labels depending on the network inflow rates and the edge capacities of the given network and  $\tau_{\min}$  the shortest physical transit time of all edges of the given network.

We then turn to the computational complexity of IDE flows. Our first result here is a lower bound on the output complexity of any algorithm. We construct an instance in which the unique IDE flow oscillates with a changing periodicity (see Fig. 4).

Theorem 4.1: There are instances for which the output complexity of an IDE flow is not polynomial in the encoding size of the instance, even if we are allowed to use periodicity to reduce the encoding size of the flow.

We also show that several natural decision problems about the existence of IDE flows with certain properties are NP-hard.

Theorem 4.9: The following decision problems are all NP-hard:

- Given a specific edge: Is there an IDE using/not using this edge?
- Given some time horizon T: Is there an IDE that terminates before T?
- Given some  $k \in \mathbb{N}$ : Is there an IDE with at most k phases?

The proof is a reduction from 3SAT, wherein for any given 3SAT-formula we construct a network (see Fig. 10) with the following properties: If the 3SAT-formula is satisfiable there exists a quite simple IDE flow, where all flow particles travel on direct paths towards the sink. If, on the other hand, the 3SAT-formula is unsatisfiable all IDE flows in the corresponding network lead to congestions diverting a certain amount of flow into a separate part of the network. Placing different gadgets in this part of networks then allows for the reduction to various decision problems involving IDE flows.

## 1.2 Related work

The concept of flows over time was studied by Ford and Fulkerson [6]. Shortly after, Vickrey [25] 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 [17] gave a novel characterization of dynamic equilibria in terms of a family of static flows (thin flows) which was further refined by Cominetti, Correa and Larré in [3]. Using this new approach Sering and Skutella [23] considered dynamic equilibria in networks with multiple sources or multiple sinks, Correa, Cristi, and Oosterwijk [5] derived a bound on the price of anarchy for dynamic equilibria and Sering and Vargas-Koch [24] incorporated spillbacks in the fluid queuing model. In a very recent work, Kaiser [16] showed that the thin flows needed for the extension step in computing dynamic equilibria can be determined in polynomial time for

series-parallel networks. Several of these papers [3,5,16,23] also explicitly mention the problem of possible non-finiteness of the extension steps.

In the traffic assignment literature, the concept of IDE was studied by several papers such as Ran and Boyce [20, § VII-IX], Boyce, Ran and LeBlanc [1,21], Friesz et al. [8]. These works develop an optimal control-theoretic formulation and characterize instantaneous user equilibria by Pontryagin's optimality conditions. For solving the control problem, Boyce, Ran and LeBlanc [1] proposed to discretize the model resulting in finite dimensional NLP whose optimal solutions correspond to approximative IDE. While this approach only gives an approximative equilibrium, there are further difficulties. The control-theoretic formulation is actually not compatible with the deterministic queueing model of Vickrey. In Boyce, Ran and LeBlanc [1], a differential equation per edge governing the cumulative edge flow (state variable) is used. The right-hand side of the differential equation depends on the exit flow function which is assumed to be differentiable and strictly positive for any positive inflow. Both assumptions (positivity and differentiability) are not satisfied for the Vickrey model. For example, flow entering an empty edge needs a strictly positive time after which it leaves the edge again, thus, violating the strict positiveness of the exit flow function. More importantly, differentiability of the exit flow function is not guaranteed for the Vickrey queueing model. Non-differentiability (or equivalently discontinuity w.r.t. the state variable) is a well-known obstacle in the convergence analysis of a discretization of the Vickrey model, see for instance Han et al. [11]. It is a priori not clear how to obtain convergence of a discretization scheme for an arbitrary flow over time (disregading equilibrium properties) within the Vickrey model. And while a recent computational study by Ziemke et al. [26] shows some positive results with regards to convergence for DE, Otsubo and Rapoport [19] report "significant discrepancies" between the continuous and a discretized solution for the Vickrey model. To overcome the discontinuity issue, Han et al. [11] reformulated the model using a PDE formulation. They obtained a discretized model whose limit points correspond to dynamic equilibria of the continuous model. The algorithm itself, however, is numerical in the sense that a precision is specified and within that precision an approximate equilibrium is computed. The overall discretization approach mentioned above stands in line with a class of numerical algorithms based on fixed point iterations computing approximate equilibrium flows within a certain numerical precision, see Friesz and Han [7] for a recent survey.

The long term behavior of dynamic equilibria with infinitely lasting constant inflow rate at a single source was studied by Cominetti, Correa and Olver [4]. They introduced the concept of a steady state and showed that dynamic equilibria always reach a stable state provided that the inflow rate is at most the capacity of a minimal *s*-*t* cut.

Ismaili [14,15] considered a discrete version of DE and IDE, respectively. He investigated the computational complexity of computing best responses for DE showing that the best-response optimization problem is not approximable, and that deciding the existence of a Nash equilibrium is complete for the second level of the polynomial hierarchy. In [15] a sequential version of a discrete routing game is studied and PSPACE hardness results for computing an optimal routing strategy are derived. For further results regarding a discrete packet routing model, we refer to Cao et al. [2], Scarsini et al. [22], Harks et al. [12] and Hoefer et al. [13].

## 2 Model and the extension-algorithm

Throughout this paper we always consider networks  $\mathcal{N} = (G, (v_e)_{e \in E}, (\tau_e)_{e \in E}, (u_v)_{v \in V \setminus \{t\}}, t)$  given by a directed graph G = (V, E), edge capacities  $v_e \in \mathbb{Q}_{>0}$ , edge travel times  $\tau_e \in \mathbb{Q}_{>0}$ , and a single sink node  $t \in V$  which is reachable from anywhere in the graph. Every other node  $v \in V \setminus \{t\}$  has a corresponding (network) inflow rate  $u_v : \mathbb{R}_{\geq 0} \to \mathbb{Q}_{\geq 0}$  indicating for every time  $\theta \in \mathbb{R}_{\geq 0}$  the rate  $u_v(\theta)$  at which the infinitesimal small agents enter the network at node v and start traveling through the graph until they leave the network at the common sink node t. We will assume that these network inflow rates are right-constant step functions with bounded support and finitely many, rational jump points and denote by  $P \in \mathbb{N}^*$  the total number of jump points for all network inflow rates.

A flow over time in  $\mathcal{N}$  is a tuple  $f = (f^+, f^-)$  where  $f^+, f^- : E \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ 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  $\theta$  and  $f_e^-(\theta)$  is the (edge) outflow rate from e at time  $\theta$ . For any such flow over time f we define the cumulative (edge) inand outflow rates  $F^+$  and  $F^-$  by

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

respectively. The queue length of edge e at time  $\theta$  is then defined as

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

Such a flow f is called a *feasible flow* for the given set of inflow rates  $u_v : \mathbb{R}_{\geq 0} \to \mathbb{Q}_{\geq 0}$ , if it satisfies the following constraints (2) to (5). The *flow conservation constraints* are modeled for all nodes  $v \neq t$  as

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

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

$$\sum_{e \in \delta_l^+} f_e^+(\theta) - \sum_{e \in \delta_l^-} f_e^-(\theta) \le 0$$
(3)

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

$$f_e^-(\theta) = 0 \text{ for all } \theta < \tau_e.$$
(4)

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

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

D Springer

Following the definition in [10] we call a feasible flow an IDE flow if whenever a particle arrives at a node  $v \neq t$ , it can only ever enter an edge that is the first edge on a currently shortest *v*-*t* path. In order to formally describe this property we first define the *current* or *instantaneous travel time* of an edge *e* at  $\theta$  by

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

We then define time dependent node labels  $\ell_v(\theta)$  corresponding to current shortest path distances from v to the sink t. For  $v \in V$  and  $\theta \in \mathbb{R}_{>0}$ , define

$$\ell_{v}(\theta) := \begin{cases} 0, & \text{for } v = t \\ \min_{e=vw \in E} \{\ell_{w}(\theta) + c_{e}(\theta)\}, & \text{else.} \end{cases}$$
(7)

We say that an edge e = vw is *active* at time  $\theta$ , if  $\ell_v(\theta) = \ell_w(\theta) + c_e(\theta)$ , denote the set of active edges by  $E_{\theta} \subseteq E$  and call the subgraph  $G[E_{\theta}]$  induced by these edges the *active subgraph*.

**Definition 2.1** 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

$$f_e^+(\theta) > 0 \Rightarrow e \in E_\theta.$$
(8)

During the computation of an IDE we also need the concept of *partial* flows/IDE that are only defined up to a certain point in time. First, a *partial flow over time* is a tupel  $(f^+, f^-)$  such that for every edge e we have two integrable functions  $f_e^+ : [0, a_e) \to \mathbb{R}_{\geq 0}$  and  $f_e^- : [0, a_e + \tau_e) \to \mathbb{R}_{\geq 0}$  for some non-negative number  $a_e$ , satisfying constraints (5) and (4) for all  $\theta < a_e$ . Such a flow is a *feasible (partial)* flow up to  $\hat{\theta}$  at node v if

- The edge outflow rates for all edges leading towards v are defined at least up to time θ̂, i.e. a<sub>e</sub> + τ<sub>e</sub> ≥ θ̂ for all e ∈ δ<sup>-</sup><sub>v</sub>,
- The edge inflow rates for all edges leaving v are defined up to time θ̂, i.e. a<sub>e</sub> = θ̂ for all e ∈ δ<sup>+</sup><sub>v</sub> and
- Constraint (2) or constraint (3), respectively, holds at v for all  $\theta < \hat{\theta}$ .

A partial flow is a *feasible (partial) flow up to time*  $\hat{\theta}$ , if it is a feasible partial flow up to time  $\hat{\theta}$  at every node. We call such a flow a *partial IDE up to time*  $\hat{\theta}$ , if additionally (8) holds for all edges and all times before  $\hat{\theta}$ . If the given network is acyclic, we can even speak of a *partial IDE up time*  $\hat{\theta}$  at some node v, denoting a feasible flow up to time  $\hat{\theta}$  at v, at least up to time  $\hat{\theta}$  for all nodes lying on some path from v to the sink t and satisfying (8) for all  $e \in \delta_v^+$  and  $\theta < \hat{\theta}$ .

Note that, while the edge inflow rates of a feasible partial flow up to  $\hat{\theta}$  are defined only on  $[0, \hat{\theta})$ , this already determines the queue length functions and, therefore, the instantaneous edge travel times on  $[0, \hat{\theta}]$ . In particular, for such a flow we can speak about active edges at time  $\hat{\theta}$  even though the flow itself is not yet defined at that time. In [10, Section 3] the existence of IDE flows in single-sink networks is proven by the following almost constructive argument: A partial IDE up to some time  $\hat{\theta}$  can always be extended for some additional proper<sup>2</sup> time interval on a node by node basis (starting with the nodes closest to the sink *t*). The existence of IDE for the whole  $\mathbb{R}_{\geq 0}$ then follows by a limit argument. This leads to a natural algorithm for computing IDE flows in single-sink networks, which we make explicit here as Algorithm 1, wherein  $b_v^-$  denotes the gross node inflow rate at node *v* defined by setting

$$b_v^-(\theta) := \sum_{e \in \delta_v^-} f_e^-(\theta) + u_v(\theta)$$

for all  $v \in V \setminus \{t\}$  and  $\theta \in [\hat{\theta}, \hat{\theta} + \tau_{\min})$ , where  $\tau_{\min} := \min\{\tau_e | e \in E\} > 0$ .

Algorithm	1:	<b>IDE-Construction</b> Algorithm from	[10]	1
-----------	----	--	------	---

<b>Input</b> : A single-sink network $\mathcal{N}$ with piecewise constant network inflow rates						
<b>Output</b> : An IDE flow $f$ in $\mathcal{N}$						
1 Let f be the zero flow and $\theta \leftarrow 0$						
2 while not all flow particles have reached the sink t do						
/* $f$ is a partial IDE up to time $ heta$ */						
Let $t = v_1 < v_2 < \cdots < v_n$ be a topological order w.r.t. $G[E_{\theta}]$						
4 <b>for</b> $i = 2,, n$ <b>do</b>						
5 Compute $b_{v_i}^-(\theta)$ and determine a constant distribution of this inflow to edges in $\delta_{v_i}^+$ such that						
the used edges remain active for some proper interval						
6 end for						
7 Determine the largest $\alpha \ge 0$ such that all $b_{\nu}^{-}$ are constant on $(\theta, \theta + \alpha)$ and the set of active						
edges does not change						
8 Extend f up to time $\theta + \alpha$ with constant edge inflow rates and set $\theta \leftarrow \theta + \alpha$						
9 end while						

For the extension at a single node v in line 5 we can use a solution to the following convex optimization problem, which can be determined in polynomial time using a simple water filling procedure (see appendix for more details):

$$\min \sum_{e=vw\in\delta_v^+\cap E_\theta} \int_0^{x_e} \frac{g_e(z)}{v_e} + \partial_+ \ell_w(\theta) dz \qquad (\text{OPT-}b_v^-(\theta))$$
  
s.t. 
$$\sum_{e\in\delta_v^+\cap E_\theta} x_e = b_v^-(\theta), \quad x_e \ge 0 \text{ for all } e \in \delta_v^+ \cap E_\theta,$$

where  $g_e$  denotes the right side derivative of the queue length function  $q_e$  depending on the inflow rate into e, i.e.  $g_e(z) := z - v_e$ , if  $q_e(\theta) > 0$  and  $g_e(z) := \max\{z - v_e, 0\}$ , otherwise. The right side derivatives  $\partial_+ \ell_w(\theta)$  exist because we only need them for nodes w closer to the sink with respect to the current topological order. And for those we already determined (constant) edge inflow rates for all outgoing edges for

<sup>&</sup>lt;sup>2</sup> We call an interval [a, b) proper if a < b.

some additional proper time interval beginning with  $\theta$ . The integrand of the objective function is, thus, the right derivative of the shortest instantaneous travel time towards the sink when entering edge vw at time  $\theta$  and assuming a constant inflow rate of z into this edge starting at time  $\theta$ . Using this observation one can show (cf. [10, Lemma 3.1]) that any solution to  $(OPT-b_v^-(\theta))$  corresponds to a flow distribution to active edges so that for every edge  $e = vw \in \delta_v^+ \cap E_\theta$  the following condition is satisfied

$$\begin{aligned}
f_e^+(\theta) &> 0 \implies \partial_+ \ell_v(\theta) = \partial_+ c_e(\theta) + \partial_+ \ell_w(\theta) \\
f_e^+(\theta) &= 0 \implies \partial_+ \ell_v(\theta) \le \partial_+ c_e(\theta) + \partial_+ \ell_w(\theta).
\end{aligned}$$
(9)

Because the network inflow rates as well as all already constructed edge inflow rates are piecewise constant and the node label functions as well as the queue length functions are continuous, (9) ensures that the used edges will remain active for some proper time interval.

It is, however, not obvious whether a finite number of such extension phases suffices to construct an IDE flow for all of  $\mathbb{R}_{\geq 0}$ . Since IDE flows always have a finite termination time in single-sink networks ([10, Theorem 4.6]) it is at least enough to extend the flow for some finite time horizon (in [9] we even provide a way to explicitly compute such a time horizon). This leaves the possibility of continuously decreasing lengths of the extension phases as possible reason for Algorithm 1 not to terminate within finite time, e.g. some sequence of extension phases of lengths  $\alpha_1, \alpha_2, \ldots$  such that  $\sum_{i=1}^{\infty} \alpha_i$  converges to some point strictly before the IDE's termination time (see Remark 4.7 for an example where we can in fact achieve arbitrarily small extension phases). Thus, the question of whether IDE flows can actually be computed was left as an open question in [10]. A first partial answer was found in [18], where finite termination was shown for graphs obtained by series composition of parallel edges. In the following section we give a full answer by showing that the  $\alpha$ -extension algorithm terminates for all single-sink networks.

## 3 Finite IDE-construction algorithm

In this chapter we will show that IDE flows can be constructed in finite time using Algorithm 1 or slight variations thereof. We will first show this only for acyclic networks since there we can use a single constant order of the nodes for the whole construction. Building on that, we will then prove the general case by showing that we can always compute IDE flows while changing the node order only finitely many times.

#### 3.1 Acyclic networks

For each extension step, Algorithm 1 takes a partial IDE and determines a networkwide constant extension of all edge in- and outflow rates. This flow distribution then continues until an event (change of gross node inflow rate or change of the set of active edges) *anywhere* in the network requires a new flow split. In [10] such a maximal extension is called a *phase* of the constructed IDE. After each phase, one then has to determine a new topological order with respect to the active subgraph at the beginning of the next phase.

For an acyclic network, we can instead use a single static topological order of the nodes with respect to the whole graph, which is then in particular a topological order with respect to any possible active subgraph. This allows us to rearrange the order of the extension steps: Considering the nodes according to the fixed topological order, at each node, we then already know the gross node inflow rate for the whole interval  $[\theta, \theta + \tau_{min})$  as well as the flow distribution for all nodes closer to the sink over the same time interval. Thus, we have enough information to determine a (possibly infinite) sequence of extensions covering the whole interval  $[\theta, \theta + \tau_{min})$ , where each extension is defined through constant flow distributions at this node. Within this sequence, each extension lasts until an event at the current node happens, which forces us to compute a new flow distribution. We call such a maximal extension using one constant flow distribution at a single node a *local phase*. The restructured version of the extension algorithm is formalized in Algorithm 2.

Algorithm 2: IDE-Construction Algorithm for acyclic networks Input: An acyclic single-sink network  $\mathcal{N}$  with piecewise constant network inflow rates **Output:** An IDE flow f in  $\mathcal{N}$ 1 Choose  $T \in \mathbb{Q}$  large enough such that all IDE flows in  $\mathcal{N}$  terminate before T 2 Let f be the zero flow,  $\theta \leftarrow 0$  and  $t = v_1 < \cdots < v_n$  a topological order 3 for  $k = 0, \ldots, \lfloor T/\tau_{\min} \rfloor$  do /\* f is a partial IDE up to time  $\theta = k \tau_{\min}$  \*/ 4 for i = 2, ..., n do Compute the piecewise constant gross node inflow function  $b_{v_i}^-$  for the interval  $[\theta, \theta + \tau_{\min})$ 5 Distribute this inflow for the whole interval to active edges in  $\delta_{v_i}^+$  using maximal local phases 6 of constant flow distribution end for 7  $\theta \leftarrow \theta + \tau_{\min}$ 9 end for

**Observation 3.1** For acyclic networks both variants of the general algorithm (Algorithms 1 and 2) construct the same IDE provided that they use the same tie-breaking rules. Thus, showing that one of them terminates in finite time, also proves the same for the other variant.

Using the water filling procedure (Algorithm 4) we can compute an IDE compliant flow distribution with constant edge-inflow rates at a node  $v_i$  for any interval wherein the inflow into node  $v_i$  is constant, the labels on all the nodes w with  $v_i w \in \delta_{v_i}^+$ change linearly and the set of active edges leaving  $v_i$  remains constant (see Fig. 2 for an example of such a situation). Thus, it suffices to show that in line 6 we can always cover the extension interval  $[\theta, \theta + \tau_{\min})$  with a finite number of local phases. We will show this by induction over  $k \in \mathbb{N}_0$  and  $i \in [n]$  using the following key lemma:

**Lemma 3.2** Let  $\mathcal{N}$  be a single-sink network on an acyclic graph with some fixed topological order on the nodes, v some node in  $\mathcal{N}$  and  $\theta_1 < \theta_2 \leq \theta_1 + \tau_{\min}$  two points in time. If f is a partial flow over time in  $\mathcal{N}$  such that



**Fig. 2** The situation in Lemma 3.2: We have an acyclic graph with some topological order on the nodes (here from left to right) and a partial IDE up to some time  $\theta_2$  for all nodes closer to the sink *t* than *v* and up to some earlier time  $\theta_1$  for *v* and all nodes further away than *v* from *t*. Additionally, over the interval  $[\theta_1, \theta_2)$  the edges leading into *v* have a constant outflow rate and the nodes  $w_i$  all have affine label functions  $\ell_{w_i}$ . The edges  $vw_i$  start with some current queue lengths  $q_{vw_i}(\theta_1) \ge 0$ 

- f is a partial IDE up to time  $\theta_2$  for all nodes closer to the sink t than v with respect to the fixed topological order,
- *f* is a partial IDE up to time  $\theta_1$  for all other nodes,
- $b_v^-$  is constant during  $[\theta_1, \theta_2)$  and
- The labels at the nodes reachable via direct edges from v are affine functions on  $[\theta_1, \theta_2)$ ,

then we can extend f to a partial IDE up to time  $\theta_2$  at v using a finite number of local phases.

**Proof** We want to show that a finite number of maximal constant extensions of the flow at node v using the water filling algorithm is enough to extend the given flow for the whole interval  $[\theta_1, \theta_2)$  at node v. So, let f be the flow after an, a priori, infinite number of extension steps getting us to a partial IDE up to some  $\hat{\theta} \in (\theta_1, \theta_2]$  at node v.

Let  $\delta_v^+ = \{vw_1, \dots, vw_p\}$  be the set of outgoing edges from v. Then, by the lemma's assumption, the label functions  $\ell_{w_i} : [\theta_1, \theta_2) \to \mathbb{R}_{\geq 0}$  are affine functions and, since we extended f at node v up to  $\hat{\theta}$ , the queue length functions  $q_{vw_i}$  are well defined on the interval  $[\theta_1, \hat{\theta})$ . Thus, for all  $i \in [p]$  we can define functions

$$h_i : [\theta_1, \hat{\theta}) \to \mathbb{R}_{\geq 0}, \theta \mapsto \tau_{vw_i} + \frac{q_{vw_i}(\theta)}{v_{vw_i}} + \ell_{w_i}(\theta)$$

🖉 Springer



**Fig. 3** A possible flow distribution from the node v in five local phases for the situation depicted in Fig. 2. The first six pictures show the flow split for these five local phases. The graph at the bottom shows the corresponding functions  $h_i$ . The bold gray line marks the graph of the function  $\ell_v$ . The second, third and fifth local phase all start because an edge becomes newly active (edges  $vw_3$ ,  $vw_1$  and  $vw_3$  again, respectively). The fourth local phase starts because the queue on the active edge  $vw_1$  runs empty. By observation (ii) these are the only two possible events which can trigger the beginning of a new local phase. Edge  $vw_2$  is inactive for the whole time interval and – as stated in Claim 1 – has a convex graph. Also, note the slope changes of the functions  $h_i$  and  $\ell_v$  in accordance with Claim 2

such that  $h_i(\theta)$  is the shortest current travel time to the sink t for a particle entering edge  $vw_i$  at time  $\theta$ . Then, for any edge  $vw_i \in \delta_v^+$  and any time  $\theta \in [\theta_1, \hat{\theta})$  we have

$$vw_i \in E_\theta \iff h_i(\theta) = \min\{h_i(\theta) | j \in [p]\} = \ell_v(\theta).$$
 (10)

We start by stating two important observations and then proceed by showing two key-properties of the functions  $h_i$  and  $\ell_v$ , which are also visualized in Fig. 3:

- (i) The functions h<sub>i</sub> are continuous and piece-wise linear. In particular they are differentiable almost everywhere and their left and right side derivatives ∂<sub>−</sub>h<sub>i</sub> and ∂<sub>+</sub>h<sub>i</sub>, respectively, exist everywhere. The same holds for the function ℓ<sub>v</sub>.
- (ii) A new local phase begins at a time  $\theta \in [\theta_1, \hat{\theta})$  if and only if at least one of the following two events occurs at time  $\theta$ : An edge  $vw_i$  becomes newly active or the queue of an active edge  $vw_i$  runs empty.

**Claim 1** If an edge  $vw_i$  is inactive during some interval  $(a, b) \subseteq [\theta_1, \hat{\theta}]$  the graph of  $h_i$  is convex on this interval.

**Claim 2** For any time  $\theta$  define  $I(\theta) := \{i \in [p] | h_i(\theta) = \ell_v(\theta)\}$ . Then, we have

$$\min\{\partial_{-}h_{i}(\theta)|i\in I(\theta)\} \le \partial_{+}\ell_{v}(\theta).$$
(11)

If, additionally, no edge becomes newly active at time  $\theta$ , we also have

$$\partial_{-}\ell_{\nu}(\theta) \le \partial_{+}\ell_{\nu}(\theta). \tag{12}$$

**Proof of Claim 1** By the lemma's assumption  $\ell_{w_i}$  is linear on the whole interval. For an inactive edge  $vw_i$  its queue length function consists of at most two linear sections: One where the queue depletes at a constant rate of  $-v_e$  and one where it remains constant 0. Thus,  $h_i$  is convex as sum of two convex functions for any interval, where  $vw_i$  is inactive.

**Proof of Claim 2** To show (11), let I' be the set of indices of edges active immediately after  $\theta$ , i.e.

$$I' := \{ i \in I(\theta) | \partial_{+}h_{i}(\theta) = \partial_{+}\ell_{v}(\theta) \}.$$

Since the total outflow from node v is constant during  $[\theta_1, \hat{\theta})$  and flow may only enter edges  $vw_i$  with  $i \in I'$  after  $\theta$ , there exists some  $j \in I'$ , where the inflow rate into  $vw_j$ after  $\theta$  is the same or larger than before. But then we have  $\partial_+h_j(\theta) \ge \partial_-h_j(\theta)$  and, thus,

$$\min\{\partial_{-}h_{i}(\theta)|i \in I(\theta)\} \leq \min\{\partial_{-}h_{i}(\theta)|i \in I'\} \leq \partial_{-}h_{i}(\theta) \leq \partial_{+}h_{i}(\theta) = \partial_{+}\ell_{v}(\theta).$$

If, additionally, no edge becomes newly active at time  $\theta$ , all edges  $vw_i$  with  $i \in I'$  have been active directly before  $\theta$  as well implying

$$\partial_{-}\ell_{v}(\theta) = \min\{\partial_{-}h_{i}(\theta)|i \in I(\theta)\} \stackrel{(11)}{\leq} \partial_{+}\ell_{v}(\theta).$$

We also need the following observation which is an immediate consequence of the way the water filling algorithm determines the flow distribution (see Observation 6.1) combined with the lemma's assumption that all label functions  $\ell_{w_i}$  have constant derivative during the interval  $[\theta_1, \theta_2)$ .

**Claim 3** There are uniquely defined numbers  $\ell_{I,J}$  for all subsets  $J \subseteq I \subseteq [p]$  such that  $\ell'_v(\theta) = \ell_{I,J}$  within all local phases, where  $\{vw_i | i \in I\}$  is the set of active edges in  $\delta_v^+$  and  $\{vw_i | i \in J\}$  is the subset of such active edges that also have a non-zero queue during this local phase.

Using these properties we can now first show a claim which implies that the smallest  $\ell_{I,J}$  can only be the derivative of  $\ell_v$  for a finite number of intervals. Inductively the same then holds for all of the finitely many  $\ell_{I,J}$ . The proof of the lemma finally concludes by observing that an interval with constant derivative of  $\ell_v$  can contain only finitely many local phases.

**Claim 4** Let  $(a_1, b_1), (a_2, b_2) \subseteq [\theta_1, \hat{\theta})$  be two disjoint maximal non-empty intervals with constant  $\ell'_v(\theta) =:$ c. If  $b_1 < a_2$  and  $\ell'_v(\theta) \ge c$  for all  $\theta \in (b_1, a_2)$  where the derivative exists, then there exists an edge  $vw_i$  such that

- 1. The first local phase of  $(a_2, b_2)$  begins because  $vw_i$  becomes newly active and
- 2. This edge is not active for any time in the interval  $[a_1, a_2)$ .

In particular, the first local phase of  $(a_1, b_2)$  is not triggered by  $vw_i$  becoming active.

**Proof of Claim 4** Since we have  $\partial_+ \ell_v(a_2) = c$ , Claim 2 implies that there exists some edge  $vw_i$  with  $h_i(a_2) = \ell_v(a_2)$  and  $\partial_- h_i(a_2) \le c$ . As  $(a_2, b_2)$  was chosen to be maximal and  $\ell'_v(\theta) \ge c$  holds almost everywhere between  $b_1$  and  $a_2$ , we have  $\partial_- \ell_v(a_2) > c$ . Thus,  $vw_i$  was inactive before  $a_2$ .

Now let  $\tilde{\theta} < a_2$  be the last time before  $a_2$ , where  $vw_i$  was active. By Claim 1 we know then that  $h'_i(\theta) \leq c$  holds almost everywhere on  $[\tilde{\theta}, a_2]$ . At the same time we have  $\ell'_v(\theta) \geq c$  almost everywhere on  $[a_1, a_2]$  and  $\ell'_v(\theta) > c$  for at least some proper subinterval of  $[b_1, a_2]$ , since the intervals  $(a_1, b_1)$  and  $(a_2, b_2)$  were chosen to be maximal. Combining these two facts with  $\ell_v(a_2) = h_i(a_2)$  implies  $\ell_v(\theta) < h_i(\theta)$  for all  $\theta \in [\tilde{\theta}, a_2) \cap [a_1, a_2)$ . As both functions are continuous we must have  $\tilde{\theta} < a_1$ . Thus,  $vw_i$  is inactive for all of  $[a_1, a_2)$ .

This claim directly implies that the lowest derivative of  $\ell_v$  during  $[\theta_1, \hat{\theta}]$  only appears in a finite number of intervals, as each of these intervals has to start with a different edge becoming newly active. But then, iteratively applying this claim for the intervals between these intervals shows that any derivative of  $\ell_v$  can only appear in a finite number of intervals. Since, by Claim 3,  $\ell'_v$  can only attain a finite number of values, this implies that  $[\theta_1, \hat{\theta}]$  consists of a finite number of intervals with constant derivative of  $\ell_v$ .

**Claim 5** Let  $(a, b) \subseteq [\theta_1, \hat{\theta})$  be an interval during which  $\ell'_v$  is constant. Then (a, b) contains at most 2 p local phases, where p denotes the out-degree of v.

**Proof of Claim 5** By Claim 1 an edge that changes from active to inactive during the interval (a, b) will remain inactive for the rest of this interval. Thus, at most p local phases can start because an edge becomes newly active. By Claim 4 if a local phase begins because the queue on an active edge  $vw_i$  runs empty at time  $\theta$ , we have  $\partial_+h_i(\theta) > \partial_-h_i(\theta) = \partial_-\ell_v(\theta) = \partial_+\ell_v(\theta)$  meaning that this edge will become inactive. Thus, at most p local phases start because the queue of an active edge runs empty. Since by observation (ii) these are the only ways to start a new local phase, we conclude that there can be no more than 2p local phases during (a, b).

Combining Claims 4 and 5 we see that  $[\theta_1, \hat{\theta})$  only contains a finite number of local phases and, thus, we achieve  $\hat{\theta} = \theta_2$  with finitely many extensions.

With this lemma the proof of the following theorem is straightforward.

**Theorem 3.3** For any acyclic single-sink network with piecewise constant networkinflow rates an IDE can be constructed in finite time using Algorithm 2.

**Proof** First, note that by [9, Theorem 1] for any given single-sink network  $\mathcal{N}$  there exists an (easily computable) time T such that all IDE in  $\mathcal{N}$  terminate before T. This makes the first line of Algorithm 2 possible. Thus, it remains to show that in line 6 a finite number of local phases always suffices. We show this by induction over  $\theta$  and  $i \in [n]$ , i.e. we can assume that the currently constructed flow f is a partial IDE up to time  $\theta$  for all nodes  $v_j$ ,  $j \ge i$  and up to time  $\theta + \tau_{\min}$  for all nodes  $v_i$ , j < i with only a finite number of (local) phases. In particular, this means that we can partition the interval  $[\theta, \theta + \tau_{\min})$  into a finite number of proper subintervals such that within each such subinterval there is a constant gross node inflow rate into node  $v_i$  and the labels at all the vertices w with  $v_i w \in \delta_{v_i}^+$  change linearly. Then, by Lemma 3.2, we can distribute the flow at node  $v_i$  to the outgoing edges using a finite number of local phases for each of these subintervals. Note that, aside from the queue lengths on the edges leaving  $v_i$ , the so distributed flow has no influence on the flow distribution in later subintervals and, in particular, does not influence the partition into subintervals or the flow distribution at nodes closer to t than  $v_i$ . Thus, we can distribute the outflow from  $v_i$  for the whole interval  $[\theta, \theta + \tau_{\min})$  using only a finite number of local phases. 

Closer inspection of the proofs above also allows us to derive a rough but explicit bound on the number of phases the constructed IDE flow can have.

**Proposition 3.4** For any acyclic single-sink network with piecewise constant networkinflow rates the number of phases of any IDE flow constructed by Algorithm 2 is bounded by

$$\mathcal{O}\left(P\left(2(\Delta+1)^{4^{\Delta}+1}\right)^{|V|T/\tau_{\min}}\right),$$

where  $\Delta := \max\{ |\delta_v^+| | v \in V \}$  the maximum out-degree in the given network and P is the number of intervals with constant network inflow rates.

**Proof** First, we look at an interval  $[\theta_1, \theta_2)$  and a single node v as in Lemma 3.2. Here we can use Claim 4 to bound the number of intervals of constant derivative of  $\ell_v$  by

$$\left(\left|\delta_{v}^{+}\right|+1\right)^{\left|\left\{(I,J)|J\subseteq I\subseteq [\left|\delta_{v}^{+}\right|]\right\}\right|} \leq \left(\left|\delta_{v}^{+}\right|+1\right)^{4\left|\delta_{v}^{+}\right|}$$

each of them containing at most  $2 |\delta_v^+|$  local phases by Claim 5. Together this shows that any such interval will be subdivided into at most  $2(\Delta + 1)^{4^{\Delta}+1}$  local phases. Thus, whenever we execute line 6 of Algorithm 2 every currently existing (local) phase may be subdivided further into at most  $2(|\Delta|+1)^{4^{\Delta}+1}$  local phases. Consequently, for every pass of the outer for-loop the number of local phases can be multiplied by at most  $\prod_{v \in V} (2(|\Delta|+1)^{4^{\Delta}+1})$  in total during the extension over the interval  $[\theta, \theta + \tau_{\min})$ . Combining this with the at most *P* phases triggered by changing network inflow rates results in the bound of

$$\mathcal{O}\left(P\left(2(\Delta+1)^{4^{\Delta}+1}\right)^{|V|T/\tau_{\min}}\right).$$

#### 3.2 General single-sink networks

We now want to extend this result to general single-sink networks, i.e. we want to show that Algorithm 1 terminates within finite time not only for acyclic graphs, but for all graphs. We first note that the requirement for input-graphs of Algorithm 2 to be acyclic is somewhat too strong. It is actually enough to have some (static) order on the nodes such that it is always a topological order with respect to the active subgraph. That is, for a general single-sink network we can still apply Algorithm 2 to determine an IDE-extension with finitely many phases for any interval during which we have such a static node ordering. Thus, Algorithm 1 will also use finitely many extension phases for each interval with such a static ordering. This observation gives rise to Algorithm 3, another slight variant of Algorithm 1.

We will prove that this algorithm does indeed construct an IDE for arbitrary singlesink networks within finite time by first showing that this algorithm is a special case of the original algorithm. Thus, it is correct and uses only a finite number of phases for any interval in which the topological order does not change. We then conclude the proof by showing that it is indeed enough to change the topological order a finite number of times for any given time horizon.

Lemma 3.5 Algorithm 3 is a special case of Algorithm 1. In particular it is correct.

**Proof** As in Algorithm 2 the existence of an upper bound *T* on the termination time of all IDE flows for a given single-sink network is guaranteed by [9, Theorem 1]. Next, note that  $\tilde{E}$  is clearly always acyclic (except in lines 11 and 13) which guarantees that we can always find a topological order with respect to  $\tilde{E}$ . We now only need to show that such an ordering is also a topological order with respect to the active edges, i.e. that for any time  $\theta$  we have  $E_{\theta} \subseteq \tilde{E}$ . For this we will use the following observation

# Algorithm 3: IDE-Construction Algorithm for general single-sink networks

I	<b>nput</b> : A single-sink network $\mathcal{N}$ with piecewise constant network inflow rates								
C	<b>Output</b> : An IDE flow $f$ in $\mathcal{N}$								
1 C	Choose T large enough such that all IDE flows in $\mathcal{N}$ terminate before T								
2 L	2 Let f be the zero flow, $\theta \leftarrow 0$ and $\tilde{E} \leftarrow E_0$								
3 D	3 Determine a topological order $t = v_1 < v_2 < \cdots < v_n$ w.r.t. the edges in $\tilde{E}$								
4 W	4 while $\theta < T$ do								
	/* $f$ is a partial IDE up to time $ heta$ */								
5	for $i = 2, \ldots, n$ do								
6	Compute $b_{v_i}^-(\theta)$ and determine a constant distribution of this inflow to edges in $\delta_{v_i}^+$ such that								
	the used edges remain active for some proper interval								
7	end for								
8	Determine the largest $\alpha \ge 0$ such that all $b_v^-$ are constant on $(\theta, \theta + \alpha)$ and the set of active								
	edges does not change								
9	Extend f up to time $\theta + \alpha$ with constant edge inflow rates and set $\theta \leftarrow \theta + \alpha$								
10	if $E_{\theta} \setminus \tilde{E} \neq \emptyset$ then								
11	Define $\tilde{E} \leftarrow \tilde{E} \cup E_{\theta}$ .								
12	while there exists a cycle C in $\tilde{E}$ do								
13	Remove an edge $e = xy$ with the largest value $\ell_y(\theta) - \ell_x(\theta)$ of all edges in C								
14	end while								
15	Determine a topological order $t = v_1 < v_2 < \cdots < v_n$ w.r.t. the edges in $\tilde{E}$								
16	end if								
rend while									

**Claim 6** Any edge xy removed from  $\tilde{E}$  in line 13 of Algorithm 3 satisfies  $\ell_x(\theta) < \ell_y(\theta)$ .

**Proof** Let  $C \subseteq \tilde{E}$  be a cycle containing the removed edge xy. Since  $\tilde{E}$  was acyclic before we added the newly active edges in line 11, this cycle also has to contain some currently active edge vw. This gives us

$$\sum_{e=uz \in C \setminus \{vw\}} (\ell_z(\theta) - \ell_u(\theta)) = \sum_{e=uz \in C} (\ell_z(\theta) - \ell_u(\theta)) - (\ell_w(\theta) - \ell_v(\theta))$$
$$= 0 - \ell_w(\theta) + \left(\ell_w(\theta) + \tau_{vw} + \frac{q_{vw}(\theta)}{v_{vw}}\right)$$
$$= \tau_{vw} + \frac{q_{vw}(\theta)}{v_{vw}} \ge \tau_{\min}.$$

Thus, *C* contains at least one edge uz with  $\ell_z(\theta) - \ell_u(\theta) > 0$  and, by the way it was chosen, this then holds in particular for edge xy.

This claim immediately implies that in line 13 we only remove inactive edges and that, afterwards, we still have  $E_{\theta} \subseteq \tilde{E}$ .

**Lemma 3.6** For any single-sink network there exists some constant C > 0 such that for any time interval of length C the set  $\tilde{E}$  changes at most |E| times during this interval in Algorithm 3.

**Proof** The proof of this lemma mainly rest on the following claim stating that for any fixed network we can bound the slope of the node labels of any feasible flow in this network by some constant.

**Claim 7** For any given network there exists some constant L > 0 such that for all feasible flows, all nodes v and all times  $\theta$  we have  $|\ell'_v(\theta)| \leq L$ .

**Proof** First note that for any node v we can bound the maximal inflow rate into this node by some constant  $L_v$  as follows:

$$\sum_{e \in \delta_v^-} f_e^-(\theta) + u_v(\theta) \stackrel{(5)}{\leq} \sum_{e \in \delta_v^-} v_e + \max\{u_v(\theta) | \theta \in \mathbb{R}_{\geq 0}\} = :L_v.$$

Using flow conservation (2) this, in turn, allows us to bound the inflow rates into all edges  $e \in \delta_v^+$  and, thus, the rate at which the queue length and the current travel time on these edges can change:

$$-1 \le c'_e(\theta) \stackrel{(1),(6)}{\le} \frac{f_e^+(\theta)}{\nu_e} \le \frac{L_v}{\nu_e} =: L_e.$$

Since this rate of change is also lower bounded by -1 setting  $L := \sum_{e \in E} \max\{1, L_e\}$  proves the claim, as for all nodes v and times  $\theta$  we then have

$$\ell'_{v}(\theta) \Big| \leq \sum_{e \in E} |c'_{e}(\theta)| \leq \sum_{e \in E} L_{e} = L.$$

Now, from Claim 6 we know that, whenever we remove an edge xy from  $\tilde{E}$  at time  $\theta$  we must have  $\ell_x(\theta) < \ell_y(\theta)$ . But at the time where we last added this edge to  $\tilde{E}$ , say at time  $\theta' < \theta$ , it must have been active (since we only ever add active edges to  $\tilde{E}$ ) and, thus, we had  $\ell_x(\theta') = \ell_y(\theta') + c_{xy}(\theta') \ge \ell_y(\theta') + \tau_{\min}$ . Therefore, the difference between the labels at x and y has changed by at least by  $\tau_{\min}$  between  $\theta'$  and  $\theta$ . Claim 7 then directly implies  $\theta - \theta' \ge \frac{\tau_{\min}}{2L}$ . So, for any time interval of length at most  $\frac{\tau_{\min}}{2L}$  each edge can be added at most once to  $\tilde{E}$ . Since  $\tilde{E}$  only ever changes when we add at least one new edge to it, setting  $C := \frac{\tau_{\min}}{2L}$  proves the lemma.

**Theorem 3.7** For any single-sink network with piecewise constant network-inflow rates an IDE can be constructed in finite time using Algorithm 3.

**Proof** By Lemma 3.5 Algorithm 3 is a special case of Algorithm 1. Thus, for any interval with static  $\tilde{E}$  it produces the same flow as Algorithm 2. In particular, by Theorem 3.3, for any such interval the constructed flow consists of finitely many phases. Finally, Lemma 3.6 shows that the whole relevant interval [0, T] can be partitioned into a finite number of intervals with static set  $\tilde{E}$ . Consequently, Algorithm 3 constructs an IDE with finitely many phases and, thus, terminates within finite time.

As in the acyclic case we can again also extract an explicit upper bound on the number of phases.

**Proposition 3.8** For any single-sink network with piecewise constant network inflow rates the number of phases of any IDE flow constructed by Algorithm 3 is bounded by

$$\mathcal{O}\left(P\left(2(\Delta+1)^{4^{\Delta}+1}\right)^{2L\cdot|E|\cdot|V|\cdot T/\tau_{\min}^{2}}\right),$$

where, again,  $\Delta := \max\{|\delta_v^+| | v \in V\}$  is the maximum out-degree in the given network, P is the number of intervals with constant network inflow rates and L the bound on the slopes of the label functions from Claim 7.

**Proof** For any time interval with fixed node order Algorithm 3 is equivalent to Algorithm 2 and, thus, the bound from Proposition 3.4 applies. Also note, that in Algorithm 2 we could change the node order after every time step of length  $\tau_{min}$  without any impact on correctness or the bound on the number of phases (as long as we always choose an order which is a topological order with respect to the active edges). As, by Lemma 3.6, the node order in Algorithm 3 changes at most  $2L \cdot |E|/\tau_{min}$  times during any unit time interval, replacing T by  $2L \cdot |E| \cdot T/\tau_{min}$  in the bound for Algorithm 2 yields a valid bound for the number of phases of Algorithm 3.

**Remark 3.9** If presented with rational input data (i.e. rational capacities, node inflow rate, current queue lengths, current distance labels and slopes of distance labels of neighbouring nodes) the water filling procedure Algorithm 4 again produces a rational solution to (OPT- $b_v^-(\theta)$ ) (i.e. rational edge inflow rates) which then, in turn, results in a rational maximal extension length  $\alpha$ . Thus, Algorithm 3 can be implemented as an exact combinatorial algorithm.

Since DE and IDE coincide for parallel link networks and for DE paths can always be replaced by single edges, the above theorem also implies the following result for DE. Note, however, that, while to the best of our knowledge this result has never explicitly been stated elsewhere, it seems very likely that it could also be shown in a more direct way for this very simple graph class.

**Corollary 3.10** On parallel paths networks Dynamic Equilibria can be constructed in finite time using the natural extension algorithm.

# 4 Computational complexity of IDE-flows

While Theorem 3.7 shows that IDE flows can be constructed in finite time, the bound provided in Proposition 3.8 is clearly superpolynomial. We now want to show that in some sense this is to be expected. Namely, we first look at the output complexity of any such algorithm, i.e. how complex the structure of IDE flows can be. Then we show that many natural decision problems involving IDE are actually NP-hard.



**Fig. 4** A network (top left picture) where constant inflow rate of 2 over [0, U] leads to an IDE with  $\Omega(U)$  different phases. The following pictures show the first states of the network, which are described in general in Table 1

# 4.1 Output complexity and steady state

In this section we call an open interval  $(a, b) \subseteq \mathbb{R}_{\geq 0}$  a *phase* of a feasible flow f, if it is a maximal interval with constant in- and outflow rates for all edges. Then it seems reasonable to expect of any algorithm computing feasible flows that its output has to contain in some way a list of the flow's phases and corresponding in- and outflow rates. In particular, the number of phases of a flow is a lower bound for the runtime of any algorithm determining that flow. This observation allows us to give an exponential lower bound for the output complexity and therefore also for the worst case runtime of any algorithm determining IDE flows. This remains true even if we only look at acyclic graphs and allow for our algorithm to recognize simple periodic behaviour and abbreviate the output accordingly.

**Theorem 4.1** The worst case output complexity of calculating IDE flows is not polynomial in the encoding size of the instance, even if we are allowed to use periodicity to reduce the encoding size of the determined flow. This is true even for series parallel graphs.

**Proof** For any given  $U \in \mathbb{N}^*$  consider the network pictured in Fig. 4 with a constant inflow rate of 2 at *s* over the interval [0, *U*]. This network can clearly be encoded in  $\mathcal{O}(\log U)$  space. The unique (up to changes on a set of measure zero) IDE is displayed up to time  $\theta = 6.5$  in Fig. 4 and described for all times in Table 1. As this pattern is clearly non-periodic and continues up to time  $\theta = U$ , it exhibits  $\Omega(U)$  distinct phases. This proves the theorem.

**Remark 4.2** In [4, Section 5.2] Cominetti et al. sketch a family of instances of size  $O(d^2)$  where a dynamic equilibrium flow exhibits an exponential number of phases (of order  $\Omega(2^d)$ ) before it reaches a stable state.

$\theta =$	$f_{vt}^+(\theta)$	$f_{wx}^+(\theta)$	$q_{vt}(\theta)$		$q_{wx}(\theta)$		$f_{sv}^+(\theta)$	$f_{sw}^+(\theta)$
$4k + 2^{-k} - 1$	0	2	$2 - 2^{-k}$	$\searrow$	$1 - 2^{-k}$	7	2	0
$4k + 2^{-k}$	2	0	$1 - 2^{-k}$	7	$2 - 2^{-k}$	$\searrow$	2	0
$4k + 2^{-k} + 1$	2	0	$2 - 2^{-k}$	7	$1 - 2^{-k}$	$\searrow$	0	2
4k + 2	2	0	$2 - 2^{-k}$	7	0	$\rightarrow$	0	2
$4k + 2^{-k} + 2$	0	2	$3 - 2^{-k}$	$\mathbf{n}$	0	7	0	2

Table 1 Phases of the (unique) IDE in the instance of Fig. 4

For all  $k \in \mathbb{N}_0$  the table includes the (constant) inflow rates into edges on the intervals  $(4k + 2^{-k} - 1, 4k + 2^{-k})$ ,  $(4k + 2^{-k}, 4k + 2^{-k} + 1)$ ,  $(4k + 2^{-k} + 1, 4k + 2)$ ,  $(4k + 2, 4k + 2^{-k} + 2)$  and  $(4k + 2^{-k} + 2, 4k + 2^{-(k+1)} + 3)$  as well as the queue lengths on the edges vt and wx at the beginning of these intervals and the rate of change for the queue lengths over the following interval ( $\nearrow$  stands for an increase at rate 1,  $\searrow$  for a decrease at rate -1 and  $\rightarrow$  for no change)

The network constructed in the above proof can also be used to gain some insights into the long term behavior of IDE flows, i.e. how such flows behave if the inflow rates continue forever. In order to analyze this long term behavior of dynamic equilibrium flows Cominetti et al. define in [4, Section 3] the concept of a steady state:

**Definition 4.3** A feasible flow f with forever lasting constant inflow rate reaches a *steady state* if there exists a time  $\tilde{\theta}$  such that after this time all queue lengths stay the same forever i.e.

$$q_e(\theta) = q_e(\tilde{\theta})$$
 f.a.  $e \in E, \theta \ge \tilde{\theta}$ .

For dynamic equilibrium flows Cominetti et al. then show that the obvious necessary condition that the inflow rate is at most the minimal total capacity of any s-t cut is also a sufficient condition for any dynamic equilibrium in such a network to eventually reach a steady state ([4, Theorem 3]). We will show that this is not true for IDE flows - even if we consider a weaker variant of steady states:

**Definition 4.4** A feasible flow f reaches a *periodic state* if there exists a time  $\tilde{\theta}$  and a periodicity  $p \in \mathbb{R}_{\geq 0}$  such that after time  $\tilde{\theta}$  all queue lengths change in a periodic manner, i.e.

$$q_e(\theta + kp) = q_e(\theta)$$
 f.a.  $e \in E, \theta \ge \tilde{\theta}, k \in \mathbb{N}^*$ .

Note that, in particular, every flow reaching a stable state also reaches a periodic state (with arbitrary periodicity).

**Theorem 4.5** There exists a series parallel network with a forever lasting constant inflow rate u at a single node s, satisfying  $u \leq \sum_{e \in \delta_X^+} v_e$  for all s-t cuts X, where no *IDE* ever reaches a periodic state.

**Proof** Consider the network constructed in the proof of Theorem 4.1, i.e. the one pictured in Fig. 4, but with a constant inflow rate of 2 at *s* for all of  $\mathbb{R}_{\geq 0}$ . A minimal



Fig. 5 The dynamic equilibrium flow for the network constructed in the proof of Theorem 4.5

cut is  $X = \{s, v, w\}$  with  $\sum_{e \in \delta_X^+} v_e = 2$ . The unique IDE flow is still the one described in Table 1 and, thus, never reaches a periodic state.

**Remark 4.6** In contrast the (again unique) dynamic equilibrium for the network from Fig. 4 is displayed in Fig. 5 and does indeed reach a steady state at time  $\theta = 4$ .

**Remark 4.7** The network considered in the proof of Theorem 4.5 also shows that we can in fact achieve arbitrarily short extension phases even within quite simple networks. Namely, the gross node inflow rate at node x is of the following form

$$b_x^-(\theta) = \begin{cases} 0, & \text{if } \theta \in [4k+3, 4k+3+2^{-k}] \text{ for some } k \in \mathbb{N}_0 \\ 1, & \text{else.} \end{cases}$$

Thus, the flow distribution at node x requires phases of lengths  $2^{-k}$  for any  $k \in \mathbb{N}_0$ . Note however, that these ever smaller getting phases are far enough apart so as to still allow us to reach any finite time horizon within a finite number of extension phases (as it is guaranteed by Theorem 3.3).

#### 4.2 NP-hardness

We will now show that the decision problem whether in a given network there exists an IDE with certain properties is often NP-hard – even if we restrict ourselves to only single-source single-sink networks on acyclic graphs. Note, however, that due to the non-uniqueness of IDE flows this does not automatically imply that computing *any* IDE must be hard.

We first observe that the restriction to a single source can be made without loss of generality.

**Lemma 4.8** For any multi-source single-sink network  $\mathcal{N}$  with piecewise constant inflow rates with finitely many jump points there exists a (larger) single-source single-sink network  $\mathcal{N}'$  with constant inflow rate such that

- (a) The encoding size of  $\mathcal{N}'$  is linearly bounded in that of  $\mathcal{N}$ ,
- (b) If  $\mathcal{N}$  is acyclic, so is  $\mathcal{N}'$ ,
- (c)  $\mathcal{N}$  is a subnetwork of  $\mathcal{N}'$  (except for the sources),
- (d) The restriction map composed with some constant translation is a one-to-one correspondence between the IDE-flows in N' and those in N:

$$\{ IDE in \mathcal{N}' \} \rightarrow \{ IDE in \mathcal{N} \}, f \mapsto f |_{\mathcal{N}}(-c).$$

**Proof** This can be accomplished by using the construction from the proof of [10, Theorem 6.3], which clearly satisfies all four properties.  $\Box$ 

**Theorem 4.9** The following decision problems are NP-hard:

- (i) Given a network and a specific edge: Is there an IDE not using this edge?
- (ii) Given a network and a specific edge: Is there an IDE using this edge?
- (iii) Given a network and a time horizon T: Is there an IDE that terminates before T?
- (iv) Given a network and some  $k \in \mathbb{N}$ : Is there an IDE consisting of at most k phases?

All these decision problems remain NP-hard even if we restrict them to single-source instances with constant inflow rate on acyclic graphs. Problem (iv) becomes NP-complete if we restrict k by some polynomial in the encoding size of the whole instance.

**Proof** We will show this theorem by reducing the NP-complete problem 3 SAT to the above problems. The main idea of the reduction is as follows: For any given instance of 3 SAT we construct a network which contains a source node for each clause with three outgoing edges corresponding to the three literals of the clause. Any satisfying interpretation of the 3 SAT-formula translates to a distribution of the network inflow to the literal edges, which leads to an IDE flow that passes through the whole network in a straightforward manner. If, on the other hand, the formula is unsatisfiable every IDE flow will cause a specific type of congestion which will divert a certain amount of flow into a different part of the graph. This part of the graph may contain an otherwise unused edge or a gadget which produces many phases (e.g. the graph constructed for the proof of Theorem 4.1) or a long travel time (e.g. an edge with very small capacity).

We start by providing two types of gadgets: One for the clauses and one for the variables of a 3SAT-formula. The clause gadget *C* (see Fig. 6) consists of a source node *c* with a constant network inflow rate of 12 over some interval of length 1 and three edges with capacity 12 and travel time 1 connecting *c* to the nodes  $\ell_1$ ,  $\ell_2$  and  $\ell_3$ , respectively. This gadget will later be embedded into a larger network in such a way that the shortest paths from the nodes  $\ell_1$ ,  $\ell_2$  and  $\ell_3$  to the sink *t* all have the same length. Thus, the flow entering the gadget at the source node *c* can be distributed in any way over the three outgoing edges. We will have a copy of this gadget for any clause of the given 3SAT-formula with the three nodes  $\ell_1$ ,  $\ell_2$  and  $\ell_3$  corresponding to the three literals of the respective clause. Setting a literal to true will than correspond to sending a flow volume of at least 4 towards the respective node.

The variable gadget V (see Fig. 7) has two nodes x and  $\neg x$  over which flow can enter the gadget. From both of these nodes there is a path consisting of two edges of length 1 leading towards a common node z, from where another edge of length and capacity 1 leads to node z'. From there the gadget will be connected to the sink node t somewhere outside the gadget. The path from  $\neg x$  to z has infinite capacity<sup>3</sup>, while the path from x to z consists of one edge with capacity 1 followed by one edge of

<sup>&</sup>lt;sup>3</sup> Throughout this construction whenever we say that an edge has "infinite capacity" by that we mean some arbitrary capacity high enough such that no queues will ever form on this edge. Since the network we construct will be acyclic such capacities can be constructed inductively similarly to the constant  $L_e$  in the proof of Claim 7

**Fig. 6** The clause gadget *C* consists of a source node and three edges leaving it, each with capacity 12 and travel time 1. If embedded in a larger network in such a way that the shortest paths from  $\ell_1, \ell_2$  and  $\ell_3$  to *t* all have the same length (and no queues during the interval [0, 1]), the inflow at node s can be distributed in any way among the three edges. In particular, it is possible to send all flow over only one of the three edges. In any distribution there has to be at least one edge which carries a flow volume of at least 4

Fig. 7 The variable gadget V. The edges xy and zz' have capacity 1, all other edges have infinite capacity. The travel times on all (solid) edges are 1 while the dashed lines represent paths with a length such that the travel time from  $s_2$  to t is the same as from y over z and z' to t. If flow enters this gadget at any rate over a time interval of length one at either x or  $\neg x$  all flow will travel over the edge zz'to the sink t. If, on the other hand, at both x and  $\neg x$  a flow of volume at least 4 enters the gadget over an interval of length 1 a flow volume of more than 1 will be diverted towards  $s_2$ 



infinite capacity with a node y between the two edges. The first edge can be bypassed by a path of length 3 and infinite capacity. From the middle node y there is also a path leaving the gadget towards t via some node  $s_2$  outside the gadget. This path has a total length of one more than the path via z and z' to t.

We will have a copy of this gadget for every variable of the given 3SAT-formula. Similarly to the clause gadget we will interpret the variable x to be set to true if a flow of volume at least 4 traverses node x and the variable to be set to false if a flow volume of at least 4 passes through node  $\neg x$ . If both happens at once, i.e. both x and  $\neg x$  each are traversed by a flow of volume at least 4 over the span of a time interval of length 1, we interpret this as an inconsistent setting of the variables. In this case a flow of volume more than 1 will leave the gadget via the edge  $y_{s_2}$  during the unit length time interval three time steps later. To verify this, assume that the flow enters at nodes xand  $\neg x$  during [0, 1]. Then the flow entering through  $\neg x$  will start to form a queue on edge zz' two time steps later. This queue will have reached a length of at least 2 at time 3 and, thus, still has a length of at least 1 at time 4. The flow entering through x at first only uses edge xy until a queue of length 2 has build up there. After that, flow will only enter this edge at a rate of 1 to keep the queue length constant, while the rest of the flow travels through the longer path towards y. This flow (of volume at least 1) as well as some non-zero amount of flow from the queue on edge xy will arrive at node y during the interval [3, 4]. Because of the queue on edge zz' all of this flow (of volume more than 1) will be diverted towards  $s_2$ . If, on the other hand, flow travels through only one of these two nodes over the course of an interval of length 1 than all this flow will be forced to travel to t via z. The third option, i.e. flow entering the gadget through both nodes but with a volume of less than 4 at at least one of them, will not be relevant for the further proof.

We can now transform a 3 SAT-formula into a network as follows: Take one copy of the clause gadget C for every clause of the formula (each with an inflow rate of 12 during the interval [0, 1] at its respective node c), one copy of the variable gadget V for every variable and connect them in the obvious way with edges of infinite capacity and unit travel time (e.g. if the first literal of some clause is  $\neg x_1$  connect the node  $\ell_1$ of this clause's copy of C with the node  $\neg x$  of the variable  $x_1$ 's copy of V and so on). Then add a sink node t and connect the nodes z' of all variable gadgets to t via edges of travel time 1 and infinite capacity. Finally, connect the node  $s_2$  (which is the same for all variable gadgets) to t by first an edge  $s_2v$  of travel time 1 and then another edge vt of travel time 2 and infinite capacity. The resulting network (see Fig. 8) has an IDE flow not using edge  $s_2 v$  if and only if the 3 SAT-formula is satisfiable: Namely, if the formula is satisfiable, take one satisfying interpretation and define a flow as follows: In every clause gadget choose one literal satisfied by the chosen interpretation and send all flow from this gadget over this literal's corresponding edge. This ensures that in the variable gadgets all flow will enter through only one of the two possible entry nodes x and  $\neg x$  and, as noted before, will then leave the gadget exclusively over node z'. If, on the other hand, the 3SAT-formula is unsatisfiable every IDE flow will sent a flow volume of more than 1 over edge  $s_2 v$  during the interval [4, 5] since in this case any flow has to have at least one variable gadget where flow volumes of at least four enter at node x as well as node  $\neg x$  (otherwise such a flow would correspond to a satisfying



**Fig. 8** Schematic representation of the whole network corresponding to a 3SAT-formula with clauses  $C_1, \ldots, C_k$  in variables  $x_1, \ldots, x_n$ . The triangles are clause gadgets (cf. Fig. 6), the rectangles are variable gadgets (cf. Fig. 7)

interpretation of the given 3SAT-formula). This shows that the first problem stated in Theorem 4.9 is NP-hard.

In order to show that the other problems are NP-hard as well, we will introduce a third type of gadget: The indicator gadget I (see Fig. 9). We can construct such a gadget for any given single-source single-sink network  $\mathcal{N}$  with constant inflow rate over the interval  $[0, \theta_0]$  at its source node. It consists of a new source node  $s_1$  with the same inflow rate as  $\mathcal{N}$ 's source node shifted by 5 time steps. The node  $s_1$  is connected to the sink node t (outside the gadget) by two paths: One through the network  $\mathcal{N}$ (entering it at its original source node  $s_{\mathcal{N}}$  and leaving it from its sink node  $t_{\mathcal{N}}$ ) and one through two additional nodes  $s_2$  and v and an edge of capacity and travel time 1 between them. All other edges outside  $\mathcal{N}$  have infinite capacity. The two outgoing edge from  $s_1$  both have a length of  $\theta_0$ . The path through the gadget has length one more than the path via  $s_2$  and v. The node  $s_2$  has a constant network inflow rate of 1 starting at time 4 and ending at time  $5 + \theta_0$ . When embedding this gadget into a larger network (with sink t) the gadget is connected to the larger network by one or more incoming edges into  $s_2$ . Fig. 9 The indicator gadget I for a single-source single-sink network  $\mathcal{N}$  with network inflow rate  $u_{\mathcal{N}} \mathbb{1}_{[0,\theta_0]}$ . All bold edges have infinite capacity, the edge  $s_2 v$  has capacity 1. The edges  $s_1 s_N$  and  $s_1 s_2$  both have travel time  $\theta_0$ , edge  $s_2 v$  has a travel time of 1 and the edges  $t_{N}t$  and vt can have any travel time such that the shortest  $s_1$ -t path through  $\mathcal{N}$  has a length of exactly one more than the  $s_1$ -t path using edge  $s_2 v$ . If within the interval [4, 5] a flow of volume more than 1 arrives at  $s_2$ over the dashed edge, all flow entering the network at  $s_1$  will travel trough  $\mathcal{N}$  (it will arrive at that sub-networks source node  $s_{\mathcal{N}}$  at a rate of  $u_{\mathcal{N}}$  during the interval  $[5 + \theta_0, 5 + 2\theta_0]$ ). If, on the other hand, a flow volume of at most 1 reaches s2 via the dashed edge up to time  $5 + \theta_0$ all flow originating at  $s_2$  will bypass  $\mathcal{N}$  using edge  $s_2 v$  and  $\mathcal{N}$ will forever remain empty



If no flow ever enters the gadget via this edge, all flow generated at  $s_1$  will travel through the path containing  $s_2v$ . If, on the other hand, a flow of volume more than 1 comes through this edge before the inflow at node  $s_1$  starts, all the flow generated there will travel through the subnetwork  $\mathcal{N}$ . Adding this gadget to the network constructed from the 3SAT-formula as described above results in a network with the following properties (see Fig. 10 for an example):

- If the 3SAT-formula is satisfiable there exists an IDE flow where the subnetwork  $\mathcal{N}$  inside gadget *I* is never used but edge  $s_1s_2$  is used.
- If the 3 SAT-formula is unsatisfiable every IDE flow will be such that its restriction to the subnetwork  $\mathcal{N}$  inside *I* is a (time shifted) IDE flow in the original stand alone network  $\mathcal{N}$  and the edge  $s_1s_2$  is never used.

Accordingly, if for example we use the network from Fig. 4 as sub-network we have a reduction from 3 SAT to the fourth problem from Theorem 4.9. Any network  $\mathcal{N}$  gives us a reduction to the second problem (with edge  $s_1s_2$  as the special edge). And just an edge with a very small capacity allows a reduction to the third problem. Alternatively, one could also use a network wherein flow gets caught in cycles for a long time before it reaches the sink as, for example, the network constructed to prove the lower bound on the termination time of IDE in [9].



**Fig. 10** The whole network for the 3SAT-formula  $(x_1 \lor x_2 \lor \neg x_3) \land (x_1 \lor \neg x_2 \lor x_4) \land (\neg x_1 \lor x_3 \lor x_4)$ . The bold edges have infinite capacity, while all other edges have capacity 1. The solid edges have a travel time of 1, the dashdotted edges may have variable travel time (depending on the subnetwork  $\mathcal{N}$ )

**Remark 4.10** Combining a construction similar to the one above with the single-source multi-sink network constructed in the proof of [10, Theorem 6.3] to show that multi-commodity IDE flows may cycle forever, shows that the problem to decide whether a given multi-sink network has an IDE terminating in finite time is NP-hard as well.

**Remark 4.11** The above construction also shows the following aspect of IDE flows: While a network may trivially contain edges that are never used in any IDE, edges that are only used in some IDE flows and edges that are used in every IDE, there can also be edges that are either not used at all or used for some flow volume of at least c, but never with any flow volume strictly between 0 and c.

# 5 Conclusions and open questions

We showed that Instantaneous Dynamic Equilibria can be computed in finite time for single-sink networks by applying the natural  $\alpha$ -extension algorithm. The obtained explicit bounds on the required number of extension steps are quite large and we do not think that they are tight.

We then turned to the computational complexity of IDE flows. We gave an example of a small instance which only allows for IDE flows with rather complex structure, thus, implying that the worst case output complexity of any algorithm computing IDE flows has to be exponential in the encoding size of the input instances. Furthermore, we showed that several natural decision problems involving IDE flows are NP-hard by describing a reduction from 3 SAT.

One common observation that can be drawn from many proofs involving IDE flows (in this paper as well as in [10] and [9]) is that they often allow for some kind of

*local analysis* of their structure – something which seems out of reach for Dynamic Equilibrium flows. This local argumentation allowed us to analyse the behavior of IDE flows in the rather complex instance from Sect. 4.2 by looking at the local behavior inside the much simpler gadgets from which the larger instance is constructed. At the same time, this was also a key aspect of the positive result in Sect. 3 where it allowed us to use inductive reasoning over the single nodes of the given network. We think that this local approach to the analysis of IDE flows might also help to answer further open questions about IDE flows in the future. One such topic might be a further investigation of the computational complexity of IDE flows. While both our upper bound on the number of extension steps as well as our lower bound for the worst case computational complexity are superpolynomial bounds, the latter is at least still polynomial in the termination time of the constructed flow, which is not the case for the former. Thus, there might still be room for improvement on either bound.

**Acknowledgements** We are grateful to the anonymous reviewers for their valuable feedback on this paper. Additionally, we thank the Deutsche Forschungsgemeinschaft (DFG) for their financial support. Finally, we want to thank the organizers and participants of the 2020 Dagstuhl seminar on "Mathematical Foundations of Dynamic Nash Flows", where we had many helpful and inspiring discussions on the topic of this paper.

Funding Open Access funding enabled and organized by Projekt DEAL.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

## Appendix: The waterfilling algorithm

In order to find a possible extension of a given partial IDE at a single node we have to determine a solution to  $(OPT-b_v^-(\theta))$ , i.e. find a distribution of the flow coming into this node to outgoing active edges in such a way that all used edges remain active for some proper time interval. As shown in [10] this can be done by a simple water filling procedure ([10, Algorithm 1 (electronic supplementary material)]), which we will restate here for the convenience of the reader. The basic idea of this procedure is to first determine for every outgoing active edge vw and all possible future constant edge inflow rates z the right side derivative of the resulting shortest instantaneous travel time towards the sink for particles starting with this edge, i.e.

$$\frac{g_e(z)}{\nu_e} + \partial_+ \ell_w(\theta),$$

where  $g_e(z):=z - v_e$ , if  $q_e(\theta) > 0$  and  $g_e(z):= \max\{z - v_e, 0\}$ , otherwise. Seen as functions in z these are continuous monotonic increasing functions starting with a constant part (if the edge has no queue to begin with) followed by an affine linear part.

Thus, they can always be written in the form

$$k(z) = \begin{cases} \beta, & z \le \gamma \\ \beta + \frac{1}{\alpha}(z - \gamma), & z \ge \gamma \end{cases}$$

for appropriately chosen constants  $\alpha$ ,  $\beta$ ,  $\gamma$ . The goal is now to distribute the current gross node inflow rate to the outgoing edges such that for all edges getting a non-zero part of this flow rate their respective functions *k* evaluated at these rates coincide, while the *k* functions of all other edges are at least as high when evaluated at a flow rate of 0. This can be accomplished by ordering the edges with increasing value of *k* for inflow rate 0 and then simultaneously filling the available node inflow into the edges with currently lowest value of *k* until all flow is distributed. This is exactly what is accomplished by Algorithm 4.

Algorithm 4: Water filling procedure for flow distributionInput : A number  $b_v^-(\theta) \ge 0$  and functions  $k_i : \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$  with  $\alpha_i > 0$  for $i = 1, \ldots, p := |\delta_v^+ \cap E_\theta|$  and  $\beta_1 \le \beta_2 \le \cdots \le \beta_p$ .Output: Values  $z_i \ge 0$  such that  $\sum_{i=1}^p z_i = b_v^-(\theta)$  and for some  $r' \le p$ satisfying  $k_0(z_0) = \cdots = k_{r'}(z_{r'}) \le \beta_{r'+1}$ ,  $z_i > 0$  for  $i \le r'$  and  $z_i = 0$ for i > r'.1 Find the maximal  $r \in \{0, 1, \ldots, p\}$  with  $\sum_{i=1}^r \max\{z|k_i(z) \le \beta_r\} \le b_v^-(\theta)$ 2 if r < p and  $\sum_{i=1}^r \max\{z|k_i(z) \le \beta_{r+1}\}$ ,  $i \le r$ 3 Set  $z_i \leftarrow \begin{cases} \max\{z|k_i(z) \le \beta_r\}, \quad i \le r \\ 0 & i > r + 1 \end{cases}$ 4 else5 Set  $z_i \leftarrow \begin{cases} \max\{z|k_i(z) \le \beta_r\}, \quad i \le r \\ 0 & i > r \end{cases}$  and  $b' \leftarrow b_v^-(\theta) - \sum_{i=1}^p z_i$ 6 Set  $z_i \leftarrow z_i + \frac{\alpha_i}{\sum_{j=1}^{r-1} \alpha_j} b'$  for all  $i \le r$ .7 end if8 return  $z_1, \ldots, z_p$ 

The correctness of this approach has been proven in [10, electronic supplementary material, Lemma 1].

**Observation 6.1** The flow distribution obtained by using Algorithm 4 at a given node only depends on the set of active edges, which subset of those currently has a non-zero queue, the gross node inflow rate and the label functions  $\ell_{w_i}$ .

**Observation 6.2** If all input data for Algorithm 4 (i.e.  $b_v^-(\theta)$  as well as all  $\alpha_i$ ,  $\beta_i$  and  $\gamma_i$ ) is rational, so is the output (i.e. the  $z_i$ ).

## References

- Boyce, D.E., Ran, B., LeBlanc, L.J.: Solving an instantaneous dynamic user-optimal route choice model. Transp. Sci. 29(2), 128–142 (1995)
- Cao, Z., Chen, B., Chen, X., Wang, C.: A network game of dynamic traffic. In: Daskalakis, C., Babaioff, M., Moulin, H. (eds.) Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26–30, 2017, pp. 695–696. ACM (2017)
- Cominetti, R., Correa, J., Larré, O.: Dynamic equilibria in fluid queueing networks. Oper. Res. 63(1), 21–34 (2015). https://doi.org/10.1287/opre.2015.1348
- Cominetti, R., Correa, J., Olver, N.: Long-term behavior of dynamic equilibria in fluid queuing networks. Oper. Res. (2021)
- 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 '19, pp. 559–577, New York, NY, USA. Association for Computing Machinery. ISBN 9781450367929. https://doi.org/10.1145/ 3328526.3329593 (2019)
- 6. Ford, L.R., Fulkerson., D.R.: Flows in Networks. Princeton University Press (1962)
- Friesz, T.L., Han, K.: The mathematical foundations of dynamic user equilibrium. Transp. Res. Part B: Methodol. 126, 309–328 (2019)
- 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)
- Graf, L., Harks, T.: The price of anarchy for instantaneous dynamic equilibrium flows. In: Chen, X., Gravin, N., Hoefer, M., Mehta, R. (eds.) Web and Internet Economics, pp. 237–251. Springer, Cham (2020)
- Graf, L., Harks, T., Sering, L.: Dynamic flows with adaptive route choice. Math. Program. 183(1), 309–335 (2020)
- Han, K., Friesz, T.L., Yao, T.: A partial differential equation formulation of Vickrey's bottleneck model, part ii: numerical analysis and computation. Transp. Res. Part B: Methodol. 49, 75–93 (2013)
- Harks, T., Peis, B., Schmand, D., Tauer, B., Vargas-Koch, L.: Competitive packet routing with priority lists. ACM Trans. Econ. Comput. 6(1), 4:1-4:26 (2018)
- Hoefer, M., Mirrokni, V.S., Röglin, H., Teng, S.-H.: Competitive routing over time. Theor. Comput. Sci. 412(39), 5420–5432 (2011)
- Ismaili, A.: Routing games over time with fifo policy. In: Devanur, N.R., Lu, P. (eds). Web and Internet Economics, pp. 266–280. Springer (2017). ISBN 978-3-319-71924-5. URL https://link.springer.com/ chapter/10.1007/978-3-319-71924-5\_19. There is also a version available on arXiv: 1709.09484
- 15. Ismaili, A.: The complexity of sequential routing games. CoRR, arXiv:1808.01080 (2018)
- Kaiser, M.: Computation of dynamic equilibria in series-parallel networks. Math. Oper. Res. forthcoming (2020)
- Koch, R., Skutella, M.: Nash equilibria and the price of anarchy for flows over time. Theory Comput. Syst. 49(1), 71–97 (2011)
- 18. Kraus, G.: Calculation of IDE-flows in SP-networks. Master's Thesis (unpublished) (2020)
- Otsubo, H., Rapoport, A.: Vickrey's model of traffic congestion discretized. Transp. Res. Part B: Methodol. 42(10), 873–889 (2008)
- Ran, B., Boyce, D.E.: Dynamic Urban Transportation Network Models: Theory and Implications for Intelligent Vehicle-Highway Systems, Lect. Notes Econ. Math. Syst. Springer, Berlin (1996)
- 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)
- Scarsini, M., Schröder, M., Tomala, T.: Dynamic atomic congestion games with seasonal flows. Oper. Res. 66(2), 327–339 (2018)
- Sering, L., Skutella, M.: Multi-source multi-sink nash flows over time. In: Borndörfer, R., Storandt, S. (eds). 18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2018), volume 65 of OASIcs - OpenAccess Ser. Inform., pp. 12:1–12:20, Dagstuhl, Germany (2018). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-096-5. http://drops.dagstuhl.de/opus/volltexte/2018/9717
- Sering, L., Vargas-Koch, L.: Nash flows over time with spillback. In: Proc. 30th Annual ACM-SIAM Sympos. on Discrete Algorithms, ACM (2019)
- 25. Vickrey, W.S.: Congestion theory and transport investment. Am. Econ. Rev. 59(2), 251-60 (1969)

 Ziemke, T., Sering, L., Vargas-Koch, L., Zimmer, M., Nagel, K., Skutella, M.: Flows over time as continuous limits of packet-based network simulations. submitted to The 23rd Euro Working Group on Transportation (2020)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.