

Stable Flows over Time

Ágnes Cseh* Jannik Matuschke* Martin Skutella*

Preprint 023-2012

July 10, 2012

Abstract

In this paper, the notion of stability is extended to network flows over time. As a useful device in our proofs, we present an elegant preflow-push variant of the Gale-Shapley algorithm that operates directly on the given network and computes stable flows in pseudo-polynomial time both in the static flow and the flow over time case. We show periodical properties of stable flows over time on networks with infinite time horizon. Finally, we discuss the influence of storage at vertices, with different results depending on the priority of the corresponding holdover edges.

1 Introduction

In the *stable marriage problem* every vertex of an undirected bipartite graph G represents either a woman or a man. Their sympathy for each person of the opposite gender is expressed by their *preference lists*: the more beloved person has the higher rank. A marriage scheme is a matching M on G . We say that such a scheme is *stable*, if there is no pair of participants willing to leave their partners in order to marry each other. More formally: an edge uv *blocks* M , if u and v both are unpaired or prefer each other to their partners in M . A matching M is *stable*, if there is no blocking edge in G . Gale and Shapley [7] were the first to state that a stable matching always exists. Their well-known deferred-acceptance algorithm finds a stable marriage in strongly polynomial time.

One of the most advanced extensions of the stable marriage problem is the *stable allocation problem*. It was introduced by Baïou and Balinski [1] in 2002. Here we talk about jobs and machines instead of men and women. Edges have capacity and participants have quota on their matched edges. This quota stands for the time a job needs to get done, and the time a machine is able to work in total. Edges are used for assigning jobs to machines such that none of the machines spends more time on a job than the edge capacity allows them.

The goal is to find a feasible set of contracts such that no machine-job pair exists where both could improve their states by breaking the scheme. Edge e is *blocking*, if it is unsaturated and neither end vertex of e could fill up its quota with at least as good edges as e . An allocation x is *stable*, if none of the edges of G is blocking. Baïou and Balinski [1] give two algorithms to solve the problem: while their augmenting path algorithm runs in strongly polynomial time, the refined Gale-Shapley algorithm is more efficient in simple cases, e. g., on instances where all jobs get one of their best choices it only needs sublinear time. Dean et al. [2, 3] succeeded to speed up the first method relying on sophisticated data structures such as dynamic trees. They also extended the second one to the case of irrational data.

*TU Berlin, Institut für Mathematik, Straße des 17. Juni 136, 10623 Berlin, Germany,
e-mail: {cseh,matuschke,skutella}@math.tu-berlin.de

Stable allocations have been further generalized to *stable flows* by Ostrovsky [9] and Fleiner [4]. Fleiner also gave a constructive proof for the existence of a stable flow in every network by solving a stable allocation problem in a modified network. In an instance of stable flow, special vertices are designated as terminals, and each non-terminal vertex has a preference list of its incoming and outgoing edges that specifies from which edges it prefers to receive flow and which edges it prefers to send flow along. Stable flows are well-suited for modeling real-world market situations, as they capture the different trading preferences of vendors and customers.

We extend this model to the setting of *flows over time*. Flows over time were introduced by Ford and Fulkerson [5, 6]. In addition to the stable flow instance we have transit times on the edges and a time horizon specifying the end of the process. We prove existence of stable flows over time and, moreover, show that, for the case of an infinite (or sufficiently large) time horizon, there is a stable flow over time which converges to a static stable flow. By introducing time to the stable flow setting, one can achieve a considerably more realistic and more interesting description of real market situations. With flows over time and transit times on the edges we can model transportation problems or illustrate distances amongst the vendors.

Structure of the paper In Section 2, we introduce the stable flow problem and present an elegant version of the Gale-Shapley algorithm that operates directly on the given network. We also give a simple proof of its correctness. In Section 3, we present the stable flow over time problem and show existence of stable flows even in the case of networks with infinite time horizon. We conclude the section by analyzing how different variations of storage at vertices can influence the stability of flows in a network.

2 Stable flows

Although most bipartite matching problems can be easily interpreted as network flow problems, stability was defined for flows only in 2008 by Ostrovsky [9]. He proved the existence and some basic properties of stable flows. Two years later, Fleiner [4] came up with a generalized setting and further results. In this paper, we use this general setting and build upon Fleiner’s remarkable achievements.

2.1 Basic notions

We consider a network (D, c) , where D is a directed graph whose vertices $V(D)$ are partitioned into a set of terminals $S \subseteq V(D)$ and non-terminal vertices $V(D) \setminus S$. Moreover, there is a capacity function $c : E(D) \rightarrow \mathbb{R}_{>0}$ on the edges. The digraph D might contain multiple edges as well as loops. This concession forces us to modify slightly the structure of preference lists: each vertex $v \in V(D)$ sets up a strictly ordered list of the neighboring edges instead of the neighboring vertices. For convenience, we consider the orderings on incoming and outgoing edges as two separate lists. The set of these lists are denoted by O . Vertex v prefers uv to wv , if uv has a lower number on v ’s preference list than wv . In this case we say that uv *dominates* wv at v and denote it by $uv <_v wv$. The same notation is used for outgoing edges.

Definition 2.1 (flow). *A flow f in network (D, c) is a function $f : E(D) \rightarrow \mathbb{R}_{\geq 0}$ such that the following properties hold:*

1. $f(uv) \leq c(uv)$ for all edges $uv \in E(D)$;

$$2. \sum_{uv \in E(D)} f(uv) = \sum_{vw \in E(D)} f(vw) \text{ for all vertices } v \in V(D) \setminus S.$$

We would like to emphasize that we do not distinguish sources and sinks in S , their role is the same: they are the vertices in D that do not have to obey the Kirchhoff law.

Definition 2.2 (stable flow). *A blocking walk of flow f is a directed walk $W = (v_1, e_1, \dots, e_{k-1}, v_k)$ such that all of the following properties hold:*

1. *each edge e_i , $i = 1, \dots, k - 1$, is unsaturated;*
2. *$v_1 \in S$ or there is an edge $e' = v_1u$ such that $f(e') > 0$ and $e_1 <_{v_1} e'$;*
3. *$v_k \in S$ or there is an edge $e'' = wv_k$ such that $f(e'') > 0$ and $e_{k-1} <_{v_k} e''$.*

A network flow is stable, if there is no blocking walk in the graph.

The network can be seen as a market situation where the vertices are the traders and the edges connecting them are possible deals. All participants rank their partners for arbitrary reasons: e. g., quality, price or location. These rankings are the preference lists of our instance. Notice that edges do not necessarily correspond to deals involving the same product and there may be cycles on the graph, even of length two. An unsaturated walk is a possible deal between vendors v_1 and v_k . If they are suppliers or consumers (terminals) or can improve their situation using the unsaturated walk, then they will agree to send some flow along it and break the existing scheme.

It was shown by Fleiner [4] that each instance \mathcal{I} of the stable flow problem can be converted into an equivalent instance \mathcal{I}' of the stable allocation problem such that every stable flow corresponds to a stable allocation and vice versa. Since there always exists a stable allocation [1], the existence of stable flows directly follows from the equivalence of stability on \mathcal{I} and \mathcal{I}' . The construction shows that stable flows can be found in polynomial time.

Theorem 2.3 (Fleiner, 2010 [4]). *There is a stable flow on every instance (D, c, O) .*

Theorem 2.4 (Fleiner, 2010 [4]). *For a fixed instance, the value of every stable flow is the same. Moreover, each edge incident to a terminal vertex has the same value in every stable flow.*

The stable flow problem can be seen as a generalization of the stable allocation problem. We introduce two terminal vertices to G and connect s with all vertices representing jobs, and t with all vertices representing a machine. The new edges get the quota of their non-terminal end vertex as capacity. Now we orient all edges from s to the jobs, from the jobs to the machines and from the machines to t in order to get a directed network. On this network all stable flows induce a stable allocation on the original graph and vice versa.

2.2 Algorithms to find stable flows

The stable allocation problem can be solved in polynomial time [1, 3]. As mentioned in the introduction, the augmenting path algorithm is not always the most efficient way to find stable allocations: the Gale-Shapley algorithm terminates faster in some cases. It can be run on \mathcal{I}' in order to give a stable allocation, which yields a stable flow on \mathcal{I} . Note that this holds for irrational data as well. The fact that this method can be directly applied to instance \mathcal{I} is briefly mentioned by Fleiner [4]. In the following we will show how to interpret the direct application of the Gale-Shapley algorithm on the

network as a preflow-push-type algorithm and prove its correctness. We will provide two variants, a basic preflow-push variant that is easy to understand and one that resembles the alternating proposal/refusal scheme of the original Gale-Shapley algorithm.

Proposal and refusal pointers For each vertex $v \in V(D)$, the algorithm maintains two pointers $p[v]$ and $r[v]$. The first pointer, $p[v]$, iterates through v 's list of outgoing edges from the highest to the lowest priority edge. It points to that edge which v is currently willing to offer more flow along. Likewise, $r[v]$ iterates through v 's list of incoming edges from the lowest to the highest priority edge. It points to that edge which v is going to refuse next, if necessary. For technical reasons we introduce one more element to each preference list: after passing through all neighbors $p[v]$ reaches a state encoded by $p[v] = 0$. This means that v cannot submit any more offers. Likewise $r[v] = 0$ initially, as v has no intention to refuse flow in the beginning.

Initialization The algorithm starts by saturating all edges leaving the terminal set, i. e., $f(sv) = c(sv)$ for all $sv \in E(D)$ with $s \in S$. We define the *excess* of a vertex v w.r.t. f by $\text{ex}(v, f) := \sum_{uw \in \delta^-(v)} f(uw) - \sum_{vw \in \delta^+(v)} f(vw)$, where $\delta^-(v)$ denotes the set of incoming edges, while $\delta^+(v)$ stands for the outgoing ones. Note that f initially is not a feasible flow, as $\text{ex}(v, f) > 0$ for some non-terminal vertices $v \in V(D) \setminus S$ —we will call such vertices *active*.

Preflow-push variant The algorithm iteratively selects an active vertex $v \in V(D)$ and pushes as much flow as possible along $p[v]$, advancing the proposal pointer whenever the edge is saturated or an already refused edge is encountered. If $p[v]$ reaches the 0-state before all excessive flow has been pushed out of the vertex, it continues by decreasing the flow on the incoming edge $r[v]$, advancing the refusal pointer whenever the flow on the edge reaches 0. After a push operation, the excess of the vertex is 0 and another active vertex is selected. The algorithm terminates once there is no active vertex left. For a pseudo-code listing of this preflow-push approach see Algorithm 1.

Simultaneous variant An alternative variant that resembles the Gale-Shapley algorithm more closely can be obtained by performing alternating rounds of proposal and refusal steps, respectively, on all active vertices simultaneously (cf. Algorithm 2). This variant of the algorithm will prove useful when analyzing stable flows in a time-expanded network later in this paper.

Algorithm 1 Preflow-Push Algorithm

```

Initialize  $p, r$ . Saturate all edges leaving  $S$ .
while  $\exists v \in V(D) \setminus S : \text{ex}(v, f) > 0$  do
  while  $\text{ex}(v, f) > 0$  do
    if  $p[v] \neq 0$  then
      PROPOSE( $p[v]$ )
    else
      REFUSE( $r[v]$ )
    end if
  end while
end while

```

Algorithm 2 Simultaneous Push Algorithm

Initialize p, r . Saturate all edges leaving S .

```
while  $\exists v \in V(D) \setminus S : \text{ex}(v, f) > 0$  do
  for all  $v \in V(D) : p[v] \neq 0$  do
    PROPOSE( $p[v]$ )
  end for
  for all  $v \in V(D) : p[v] = 0$  do
    REFUSE( $r[v]$ )
  end for
end while
```

procedure PROPOSE($e = (v, w)$)

if ($r[w] >_w e$ or $w \in S$) and $f(e) < c(e)$ **then**

$f(e) := \min(f(e) + \text{ex}(v, f), c(e))$

else

ADVANCE($p[v]$)

end if

end procedure

procedure REFUSE($e = (v, w)$)

if $r[w] \neq 0$ and $f(e) > 0$ **then**

$f(e) := \max(f(e) - \text{ex}(w, f), 0)$

else

ADVANCE($r[w]$)

end if

end procedure

A special execution of Algorithm 1 gives the McVitie-Wilson algorithm [8] for the stable marriage problem. We can determine the choice of active vertices while running Algorithm 1 on the flow instance defined by the stable matching instance. At initialization the source sends one unit of flow to each vertex symbolizing a man. The active vertex chosen arbitrarily in the first step is one of them. He proposes along his best edge, the asked lady accepts the offer and sends the flow further to the sink. Now the second man will be chosen arbitrarily, he proposes along his best edge. If any lady gets more than one offers, her vertex enters the active set and she needs to be chosen next. Afterwards, the refused man must play the role of the selected vertex, and so on. This way we run the deferred-acceptance algorithm by taking the men one by one to the instance, always setting up a current stable matching.

Theorem 2.5. *If c is integral, both algorithms return an integral stable flow in at most $O(\sum_{e \in E} c(e))$ iterations.*

The proof is split into three parts:

Claim 1. *Throughout the course of the algorithms, f is integral.*

Proof. We prove this by induction. The claim is true after initialization as the capacities of all edges are integral. Thus, before a call of REFUSE or PROPOSE, the excess of the corresponding vertex is integral as well. This implies that the flow value of the corresponding edge is changed by an integral amount. \square

Claim 2. *Both algorithms terminate after $O(\sum_{e \in E} c(e))$ steps.*

Proof. In each call of PROPOSE, the flow value of the corresponding edge is increased by an integral amount (by Claim 1), or the pointer p is advanced. Likewise, in each call of REFUSE, the flow value of the corresponding edge is decreased by an integral amount, or the pointer r is advanced. Once REFUSE is called for some edge uv , the flow value cannot be increased by a PROPOSE call anymore. Thus, there can be only at most $O(\sum_{e \in E} c(e))$ calls of PROPOSE and REFUSE. \square

Claim 3. *The algorithms return a stable flow.*

Proof. After termination, f is a feasible flow, as the excess of every non-terminal vertex is 0. Now suppose there is a blocking walk in the network. There are two reasons for leaving unsaturated edges in the network: either the edge was refused or there was no proposal along it using all its capacity, it stayed at least partly unexamined. We will study which case can come up at which position in the blocking walk.

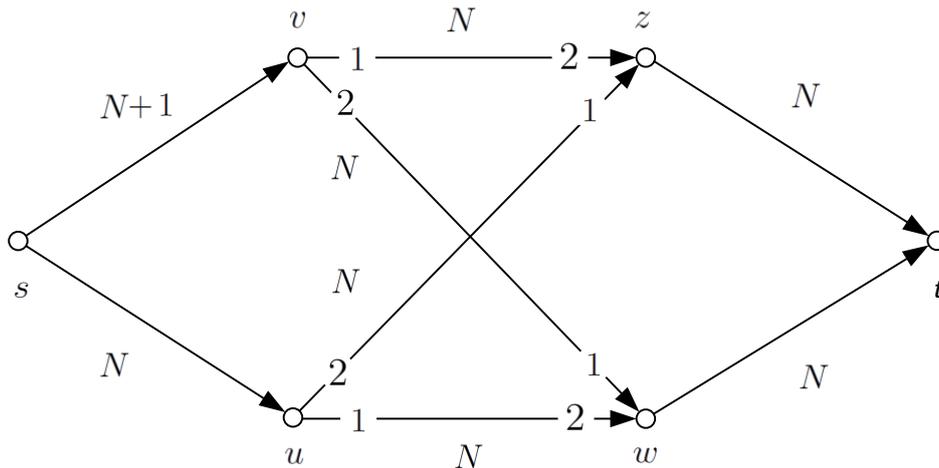
If an unsaturated walk starts at a terminal vertex, then the first edge of it has been refused, since $s \in S$ must try to fill all its adjacent edges with maximum capacity. If the walk ends at a terminal vertex, there was no full proposal along that edge in the algorithm, since terminal vertices do not refuse any flow. If the blocking walk starts at a non-terminal vertex, then there must be a dominating edge starting at the same vertex and having nonzero value. This proves that the unsaturated edge has been refused, because vertices submit offers along edges in their order in the preference list of the start vertex. A similar argument can be applied to the end vertex of the blocking walk: there must be a dominating edge ending at the same vertex and having nonzero value. The unsaturated edge can not be a refused one, since we always refuse the worst edges.

The argument above shows that the blocking walk must start with a refused edge and end with a not fully proposed one. This means that along the walk there has to be at least one refused edge uv and an at least partly unexamined one vw . This implies that vertex v refused flow although it has not filled up its outgoing quota, contradiction. \square

Corollary 2.6. *If c is rational, both algorithms return a rational stable flow.*

Proof. Any instance with rational capacity vector can be transformed to an instance with integral capacity vector by multiplying all capacities with their smallest common denominator. \square

A simple network setting can illustrate how large capacities may cause a long running time. Note that this example is the flow extension of the allocation problem described by Baïou and Balinski [1].



In the example above $S = \{s, t\}$ and N is an arbitrary large number. After saturating sv and su the simultaneous push algorithm saturates uw and vz and proposes along vw with one unit. This offer will be accepted by w , forcing it to refuse one flow unit along

uw . This way u has to submit an offer to z , that needs to accept it and reject a flow unit from v . An alternating cycle of new offers and refusals will be made along v, w, u, z as long as there is any flow to refuse along uw and vz . This means in total N augmentations along the cycle.

3 Flows over time

3.1 Basic notions

We are given a network (D, c, τ) consisting of a directed graph D , some terminal vertices $S \subseteq V(D)$ and a capacity function $c : E(D) \rightarrow \mathbb{R}_{>0}$ on the edges. The last element is the *transit time* function: $\tau : E(D) \rightarrow \mathbb{Z}_{\geq 0}$. Besides these an instance contains a *time horizon* $T \in \mathbb{Z}_{>0}$ as well. Loops and multiple edges are allowed in D .

Definition 3.1 (flow over time). *Functions $f_e : \{0, 1, \dots, T - 1\} \rightarrow \mathbb{R}_{\geq 0}$ for each edge $e \in E(D)$ form a flow over time or dynamic flow with time horizon T , if they fulfill all of the following requirements:*

1. $f_e(\theta) = 0$ for $\theta \geq T - \tau(e)$
This ensures that flow can be sent from u along an edge uv , only if there is enough time left for it to reach v .
2. $f_e(\theta) \leq c(e)$ for all $e \in E(D)$ and $\theta \in \{0, 1, \dots, T - 1\}$
Capacity constraints hold all the time.
3. $\sum_{e \in \delta^-(v)} \sum_{\xi \leq \theta - \tau(e)} f_e(\xi) = \sum_{e \in \delta^+(v)} \sum_{\xi \leq \theta} f_e(\xi)$ for all $v \in V(D) \setminus S$ and $\theta \in \{0, 1, \dots, T - 1\}$
Flow conservation is fulfilled at every point in time.

Any dynamic flow problem can be converted into an ordinary flow problem with the help of the *time-expanded network*. We create T copies of $V(D)$, with v_i denoting the $(i + 1)$ th copy of vertex i . For every edge $uv \in E(D)$ and every $i \in \{0, \dots, T - 1 - \tau(uv)\}$, we connect u_i with $v_{i+\tau(uv)}$. The vertices of D^T with a fixed index i form a *timeslot*. Though this construction reduces the flow over time concept to the static setting without transit times, the price of this simplification is a considerably larger time-expanded network D^T , having a size linear in T and thus exponential in the input size. This does not always cause a problem: in the maximum dynamic flow problem there is always a temporally repeated maximum flow that can be found in polynomial time [5, 6].

We extend the flow over time instance to stable flows over time, by introducing preference lists. They have the same behavior as in the stable flow problem and they do not change in time. The notion of stability can be extended to this instance the following way:

Definition 3.2 (stable flow over time). *A directed walk $W = (v_1, e_1, \dots, e_{k-1}, v_k)$ is a blocking walk of flow over time f , if there is a certain point in time $0 \leq \theta \leq T - 1$ such that all the following properties hold:*

1. each edge e_i is unsaturated at time $\theta + \sum_{j=1}^{i-1} \tau(e_j)$
2. $v_1 \in S$ or there is an edge $e' = v_1u$ such that $f_{e'}(\theta) > 0$ and $e_1 <_{v_1} e'$

3. $v_k \in S$ or there is an edge $e'' = wv_k$ such that $f_{e''} \left(\theta + \sum_{j=1}^{k-1} \tau(e_j) - \tau(e'') \right) > 0$ and $e_{k-1} <_{v_k} e''$

A blocking walk W can be interpreted in a similar way as in the static case: if the two participants symbolized by the end vertices agree that sending some flow along W would improve their situation, then the scheme will be broken by them.

Note that a flow over time is stable, if and only if the corresponding flow in the time-expanded network is stable in the classical sense.

Theorem 3.3. *For every (D, c, τ, O) and time horizon $T \in \mathbb{Z}_{>0}$ there is a stable flow over time.*

Proof. A stable static flow exists in the time-expanded network. \square

Corollary 3.4. *In every stable flow over time with $T \in \mathbb{Z}_{>0}$ the terminal vertices send and receive the same amount of flow in a fixed timeslot on all edges incident to them.*

Proof. This follows from Theorem 2.4. \square

Note that these two statements hold even if the preference lists may change in time.

The following lemma gives important structural insights on the stable flow computed by the simultaneous push algorithm (Algorithm 2), which will prove useful in the following sections.

Lemma 3.5. *Let $v \in V(D) \setminus S$ and $i < j$. At any step of the Algorithm 2, the following statements hold:*

1. $p[v_i] \leq_v p[v_j]$
2. If $p[v_i] =_v p[v_j] \neq 0$, then $f(p[v_i]) \leq f(p[v_j])$.
3. $r[v_i] \geq_v r[v_j]$
4. If $r[v_i] =_v r[v_j] \neq 0$, then $f(r[v_i]) \geq f(r[v_j])$.

Proof. We prove the lemma by induction on the algorithm. Clearly, all statements are true after initialization. We show that they also stay true after the end of each proposal loop. The statement for refusal loops can be shown analogously.

First observe that if $p[v_i] <_v p[v_j]$ before the loop, then (1) or (2) cannot be invalidated. Thus, let $p[v_i] =_v v_i w_{i+\tau(e)}$ and $p[v_j] =_v v_j w_{j+\tau(e)}$ for some $e = vw \in E(D)$.

(1) cannot be invalidated by a proposal loop: $p[v_i]$ will only be advanced if either $f(p[v_i]) = c(p[v_i])$ or $r[w_k] \leq_w p[v_i]$ at the beginning of the loop. In the former case, $f(p[v_j]) \geq f(p[v_i]) = c(p[v_j])$, in the latter case $r[w_l] \leq_w r[w_k] \leq_w p[v_j]$ by (3), and thus in either case $p[v_j]$ will be advanced as well.

In order to see that (2) cannot be invalidated either, observe that the inflow of v_j is at least the inflow of v_i : If $e' = uv \in E(D)$, then edge $u_{i-\tau(e')}v_i$ exists in D^T only if edge $u_{j-\tau(e')}v_j$ exists as well, and by (1) and (2), $f(u_{i-\tau(e')}v_i) \leq f(u_{j-\tau(e')}v_j)$ in this case. On the other hand, the the outflow of v_j on edges other than $p[v_j]$ is at most the same as on edges of v_i : Let $e' = (v, u)$ with $e' \neq p[v_j]$. If $e' >_v p[v_j]$, then $f(v_j u_{j+\tau(e')}) = 0$. If $e' <_v p[v_j] =_v p[v_i]$, then either $f(v_i u_{i+\tau(e')}) = c(e) \geq f(v_j u_{j+\tau(e')})$ or $r[u_{i+\tau(e)}] \leq_u e'$, which implies $f(v_j u_{j+\tau(e')}) \leq f(v_i u_{i+\tau(e')})$ by (3) and (4). Thus, $f(p[v_j])$ is set to a value at least as large as $f(p[v_i])$.

(3) and (4) stay true during a proposal loop as no r -pointer and no flow value of any edge $r[u]$ for any vertex $u \in V(D^T)$ is modified. \square

3.2 Infinite time

In this section we will prove the existence of a stable flow even if the time horizon is infinite. This flow can be constructed by applying the simultaneous push algorithm on D^∞ (under the assumption that it can apply the PROPOSE and REFUSE steps on all vertices simultaneously). Even more, after a certain point in time, the stable flow is identical to a temporal repetition of the stable flow computed by the same algorithm in the static network D .

We define D_i^∞ to be the subgraph induced by the vertices v_j with $j > i$. We will run Algorithm 2 in parallel on D and D^∞ , assuming we can execute the PROPOSE and REFUSE steps simultaneously on all vertices of D^∞ . Let f be the flow values in D and f' be the flow values in D^∞ that occur throughout the run of the algorithm, and let p, r and p', r' be the corresponding pointers, respectively.

Our main theorem in this section states that there is a point in time, from which on all computations in the infinite time expanded network correspond one-to-one to those in the static network.

Theorem 3.6. *There is a time $0 \leq i < \infty$ such that throughout the course of the algorithm, for every $j \geq i$, $f'(u_j v_k) = f(uv)$ for all edges $uv \in E(D)$ and $p'[v_j] =_v p[v]$ and $r'[v_j] =_v r[v]$ for all $v \in V(D)$.*

Proof. Let $\tau_{\max} := \max_{e \in E(D)} \tau(e)$. We will first prove by induction that after K iterations of the algorithm the statement of the theorem is true for $i_K := K \cdot \tau_{\max}$. Clearly, this is true after initialization, as all edges leaving the terminal set are saturated and all pointers are at their initial state.

Now assume that after iteration K , the state of all flow variables and pointers in $D_{i_K}^\infty$ is identical to that in D . Now let $v \in V(D)$ be any vertex that is subject to a PROPOSE operation in D . Note that for $j > i_{K+1}$, vertex v_j is only incident to edges in $D_{i_K}^\infty$, as $i_{K+1} = i_K + \tau_{\max}$. Therefore, $\text{ex}(v_j, f') = \text{ex}(v, f)$. As $p'[v_j] =_v p[v]$ and $r'[w_{j+\tau(e)}] =_w r[w]$, the PROPOSE call has the same effect on $p'[v_j]$ and $f'(p[v_j])$ as it has on $p[v]$ and $f(p[v])$. An analogous statement holds for calls of REFUSE.

We now have shown that up to iteration K , the statement of the theorem holds for i_K . Now let K_0 be the number of iterations performed by Algorithm 2 before f is a stable flow in D . After iteration K_0 , no flow values or pointers in D are changed. Let $i := i_{K_0+1}$. We will show that after iteration K_0 , also no flow values or pointers in D_i^∞ are changed, which concludes the proof of the theorem.

After iteration K_0 , all vertices in D_i^∞ are inactive. Thus, their state can only change if the flow value on an edge $v_k w_l$ for $k < i$ and $l \geq i$ is increased. This can only happen, if $f'(v_k w_l) < c(e)$, as well as $r'[w_l] >_w v_k w_l$, and $p[v_k] =_v v w$. Note that after iteration K_0 , $f'(v_{k+\tau_{\max}} w_{l+\tau_{\max}}) = f'(v_k w_l)$ and $r'[w_{l+\tau_{\max}}] =_w r'[w_l] >_w v w$ by choice of $i = i_{K_0+1}$. Now, by Lemma 3.5, $p[v_{k+\tau_{\max}}] \leq_v p[v_k]$, and as $v_{k+\tau_{\max}} w_{l+\tau_{\max}}$ is neither refused nor saturated, $p[v_{k+\tau_{\max}}] =_v v w =_v p[v_k]$. Again by Lemma 3.5, $f'(v_k w_l) \leq f'(v_{k+\tau_{\max}} w_{l+\tau_{\max}})$ throughout the remaining course of the algorithm, i. e., $f'(v_k w_l)$ cannot be increased before $v_{k+\tau_{\max}}$ becomes active. Thus, $f(v_k w_l)$ cannot be increased. \square

Corollary 3.7. *Algorithm 2 constructs a stable flow f' in D^T such that $f'(v_i w_{i+\tau(vw)}) = f(vw)$ for all $vw \in E(D)$ and all $j \geq (K_0 + 1) \cdot \tau_{\max}$, where K_0 is the number of iterations performed by the same algorithm to compute the stable flow f in D .*

3.3 Storing at vertices

The third point in the definition of a flow over time requires flow conservation in every timeslot. A different model allows storing at vertices: a vendor may delay the shipment of goods at his convenience, as long as the flow arrives the terminal vertices within the time horizon. More formally, we generalize the definition of the *excess* of vertex v at time θ as the amount of stored goods at vertex v at time θ :

$$\text{ex}_f(v, \theta) := \sum_{e \in \delta^-(v)} \sum_{\xi \leq \theta - \tau(e)} f_e(\xi) - \sum_{e \in \delta^+(v)} \sum_{\xi \leq \theta} f_e(\xi)$$

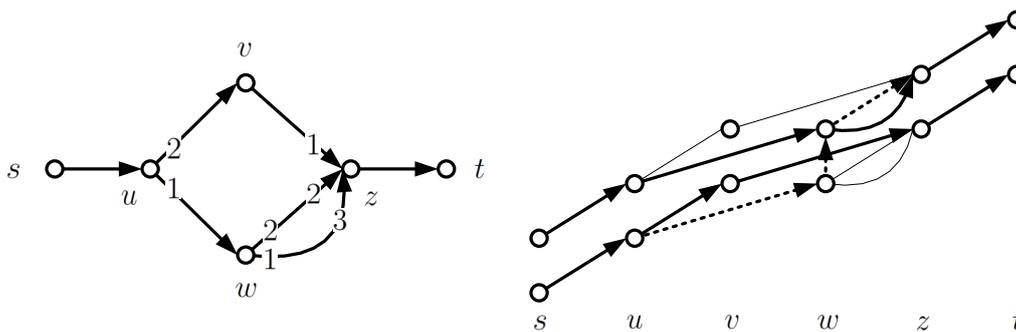
While *strict flow conservation* requires $\text{ex}_f(v, \theta) = 0$ for all non-terminal vertices and times in $\{0, \dots, T-1\}$, *weak flow conservation* allows $\text{ex}_f(v, \theta) \geq 0$ except for time $T-1$, where $\text{ex}_f(v, T-1) = 0$ must hold for all $v \in V(D) \setminus S$.

The time-expanded network can be adapted for weak flow conservation by introducing so-called *holdover edges* $v_i v_{i+1}$ of infinite capacity for each $v \in V(D)$ and $i \in \{0, \dots, T-2\}$. Setting the ranks of these holdover edges in the preference lists allows the investigation of interesting variations. One intuitive view comes from the fact that storing goods is expensive and people are impatient, thus letting the edges $v_i v_{i+1}$ be the last on v_i 's preference list and the first on v_{i+1} 's preference list. In the following we will investigate the four main cases with weak flow conservation, given by varying first and last preferences on the end vertices of holdover edges. In each case we study how the sets of stable flows obeying the different flow conservation rules are related to each other.

We adapt the definition of blocking walk for the case of intermediate storing. Apart from a starting time θ we need the exact time θ_i when the walk in D leaves a vertex v_i . A walk in D together with a sequence $\{\theta, \theta_1, \theta_2, \dots, \theta_k\}$ uniquely determines a walk on D^T .

Example 3.8. *Introducing waiting at vertices stable flows may lose their stability, independent from the rank of holdover edges.*

Proof. Consider the following example:



All edges in the static network have unit capacity and unit transit time, s and t are the non-terminal vertices. The unobvious preferences are shown on the edges. The time-expanded network with $T = 5$ stands on the right side, showing only the vertices and edges that may be used in a feasible flow. On the static network above there are three stable flows, each of them uses a different incoming edge of z . If waiting is not allowed, then the dynamic flow shown by the thick edges on the second network is stable, since all unsaturated paths are dominated by it at one end. Adding holdover edges to the time-expanded network breaks the scheme: the walk $\{u_1, w_2, w_3, z_4\}$ (denoted by dashed lines) blocks the flow over time. Note that this is independent of the rank of the holdover edges, since we only need dominance at the ends of the blocking walk. \square

Theorem 3.9. *If T is sufficiently large and holdover edges always stand on the first place on preference lists, then there is no stable flow with value 0 on all holdover edges.*

Proof. Suppose there is a stable flow on the time-expanded network that does not use any of the holdover edges. Consider all the copies of an arbitrary non-terminal vertex v : If there are two of them that have a positive value on all incoming (and outgoing) edges in total, then holdover edges between them form a walk that blocks the flow. This way a stable flow over time passes through an arbitrary non-terminal vertex maximal once. Apart from the terminal-terminal edges that do not affect stability—a flow with this property has positive value on at most $2|V(D) \setminus S|$ edges in the time-expanded network.

These maximum $2|V(D) \setminus S|$ edges have to ensure that there is no unsaturated walk between terminal vertices. If the length (w.r.t. transit times) of the shortest walk is denoted by $\tau(W_{min})$, then $T - \tau(W_{min}) + 1$ is a lower bound for the number of disjoint walks between terminal vertices. In order to saturate at least one edge along these walks, the inequality $2|V(D) \setminus S| \geq T - \tau(W_{min}) + 1$ must hold, otherwise at least one copy of the shortest walk is unsaturated, hence it blocks the flow. \square

Theorem 3.10. *If $T \in \mathbb{Z}_{>0}$ and the holdover edges stand on the last place on preference lists, then the stable flow produced by the simultaneous push algorithm has value 0 on all holdover edges.*

Proof. By contradiction assume, there is a holdover edge with positive flow value at some vertex v . W.l.o.g., assume i to be minimal with $f(v_i v_{i+1}) > 0$. Since the time horizon is finite, there also is a latest point in time j such that $f(v_j v_{j+1}) > 0$.

Since $f(v_i v_{i+1}) > 0$, the pointer $p[v_i]$ must already have passed all outgoing non-holdover edges of v_i , offering as much flow as the corresponding vertices are willing to accept. Therefore, by Lemma 3.5, the flow on each of these edges must be at least the flow on the corresponding edges of v_{j+1} (if they exist). Thus, the total outflow of v_i exceeds the total outflow of v_{j+1} by at least $f(v_i v_{i+1})$. On the other hand, $r[v_{j+1}]$ points either to 0 or the incoming holdover edge, and thus v_{j+1} does not refuse any flow on the regular incoming edges. Again by Lemma 3.5, the flow on those edges is at least the flow on the corresponding edges of v_i . Thus, the total inflow of v_{j+1} exceeds the total outflow of v_i by at least $f(v_j v_{j+1})$. Putting this together yields $\text{inflow}(v_{j+1}) - \text{outflow}(v_{j+1}) \geq \text{inflow}(v_i) + f(v_j v_{j+1}) - \text{outflow}(v_i) + f(v_i v_{i+1}) > 0$, contradicting flow conservation. \square

In the other two cases, when holdover edges are the best for one of the end vertices and the worst for the other one, storing can be used or avoided depending on the network. There are even networks where all stable flows have positive value on some holdover edges.

Conclusion and open questions

We introduced stable flows over time, extending the concept of stability to network flows over time. As initial results, we proved the existence of stable flows both for finite and infinite time horizon. In both cases, a stable flow can be computed in pseudo-polynomial time by applying a preflow-push algorithm operating directly on the flow network. We also showed that the possibility of storage at non-terminal vertices has an effect on the set of stable flows, depending on the preference given to holdover edges in the time expanded network.

Although this paper provides first structural insights, many questions brought up by the definition of stable flows over time remain open, most prominently the complexity of

the stable flow problem: Is there a polynomial time algorithm for finding a stable flow? As of today, it is not even clear whether a stable flow over time can be encoded in polynomial space. First research in this direction indicates that at least the concept of (generalized) temporally repeated flows cannot be applied directly.

Also, further extension of stability can be studied on the new setting, e. g., special edges or ties on preference lists, as can be extensions of the flow over time model, e. g., connections to earliest arrival flows. Finally, we conjecture a stronger form of Theorem 3.10: If holdover edges have lowest priority at both start and end vertex, no stable flow uses storage at non-terminal vertices.

Acknowledgments This work was supported by the Berlin Mathematical School and by the DFG Research Center MATHEON in Berlin.

References

- [1] Mourad Baïou and Michel Balinski. The stable allocation (or ordinal transportation) problem. *Math. Oper. Res.*, 27(3):485–503, 2002.
- [2] Brian C. Dean, Michel X. Goemans, and Nicole Immorlica. Finite termination of "augmenting path" algorithms in the presence of irrational problem data. pages 268–279, 2006.
- [3] Brian C. Dean and Siddharth Munshi. Faster algorithms for stable allocation problems. *Algorithmica*, 58(1):59–81, 2010.
- [4] Tamás Fleiner. On stable matchings and flows. In *Proceedings of the 36th International Workshop on Graph-Theoretic Concepts in Computer Science, WG'10*, pages 51–62, Berlin, Heidelberg, 2010. Springer-Verlag.
- [5] Lester R. Ford and Delbert R. Fulkerson. Constructing Maximal Dynamic Flows from Static Flows. *Operations Research*, 6(3):419–433, 1958.
- [6] Lester R. Ford and Delbert R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
- [7] David Gale and Lloyd Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 1:9–14, 1962.
- [8] David G. McVitie and Leslie B. Wilson. The stable marriage problem. *Commun. ACM*, 14(7):486–490, July 1971.
- [9] Michael Ostrovsky. Stability in supply chain networks. *American Economic Review*, 98(3):897–923, June 2008.