

Towards a Concurrent and Distributed Route Selection for Payment Channel Networks

Elias Rohrer¹, Jann-Frederik Laß², and Florian Tschorsch¹

¹ Technical University of Berlin,
{elias.rohrer, florian.tschorsch}@tu-berlin.de

² Humboldt University of Berlin,
lassjann@informatik.hu-berlin.de

Abstract. Payment channel networks use off-chain transactions to provide virtually arbitrary transaction rates. In this paper, we provide a new perspective on payment channels and consider them as a flow network. We propose an extended push-relabel algorithm to find payment flows in a payment channel network. Our algorithm enables a distributed and concurrent execution without violating capacity constraints. To this end, we introduce the concept of capacity locking. We prove that flows are valid and present first results.

1 Introduction

It seems that blockchain-based systems such as Bitcoin [6] will, due to their requirements regarding storage, processing power, and bandwidth, not be able to natively scale to high transaction rates [2]. Off-chain approaches [3, 7], however, offer a way to create long-lived payment channels between two nodes. The payments transferred via a payment channel are processed locally and therefore eliminate the need to commit each individual transaction to the blockchain.

In order to enable payments between any two nodes—whether they are directly connected or not—payment channels form a network in which payments can be routed over more than one hop. Finding a route that can process a certain transaction volume is challenging, though. Related approaches [8] cannot guarantee to utilize the available capacities as they focus on finding a single path. We argue that single-path routing restricts the transferable amount and misses many payment opportunities due to bottleneck capacities in the network. Eventually, failed payments will fall back to on-chain transactions, instead of using the available (and already locked) resources efficiently.

In this paper, we propose to aggregate multiple paths to a *payment flow*, which can in sum provide larger transaction volumes. We believe that algorithms from the domain of flow networks in general and the push-relabel algorithm [4] in particular are appropriate candidates for route selection in payment networks.

Our main contribution is an algorithm for distributed route selection, which is based on the push-relabel algorithm. It can find feasible flows in a payment channel network and is safe for concurrent execution. To this end, we introduce the concept of *capacity locking*. We show that our algorithm guarantees that routes are feasible flows and at the same time does not violate any capacity

constraints. Our first results confirm that the approach is able to handle a high number of flows and transaction volumes. The results emphasize that our approach succeeds in scenarios where single-path routing schemes are bound to fail. In summary, we offer a new perspective on payment channel networks.

The remainder is structured as follows. Sec. 2 discusses related work. Subsequently, Sec. 3 introduces payment flows and describes the basic algorithmic design. Sec. 4 develops a distributed and concurrent route selection algorithm. In Sec. 5, we present and discuss first results, before Sec. 6 concludes the paper.

2 Background and Related Work

Payment channels are a new and unexplored concept. The specifications [5] of the Lightning Network [7], for example, are subject to constant change. For the sake of clarity, we abstract from any specific payment channel design [3, 7].

Routing in a payment channel network poses many challenges, e. g., regarding the routing paradigm (per-hop routing vs. source routing) and the topology (hub-and-spoke vs. peer-to-peer). In this paper, we focus on route selection, i. e., finding a route in a payment channel network that meets certain constraints. Flare [8], a routing system for the Lightning Network, creates a list of candidate routes from the set of channels with sufficient capacity. So far, however, Flare and the Lightning Network opt for single-path routes. In our work, we consider a payment as a flow and elaborate the possibility to aggregate multiple paths.

We identify flow network algorithms as a promising direction to find multi-path routes. While multi-commodity flows address a similar problem, most of the existing approaches require global knowledge and/or a centralized routing coordinator. The approach in [1] allows a distributed and concurrent execution but solves the feasible-flow problem only approximately. Our distributed algorithm, in contrast, guarantees that the selected route is a feasible flow. Moreover, it can be executed concurrently without violating capacity constraints.

3 Payment Flows

Payment flows describe a flow of units between pairs of nodes in a payment channel network. Figure 1 shows an example of a payment channel network in which node s wants to send a payment to node t . We consider the payment channel network as a peer-to-peer network in which nodes communicate directly with each other and build an overlay network congruent with the payment channel network. That is, we aim for a decentralized route selection.

In order to process the payment, a path between s and t must exist. Every path is a concatenation of payment channels. Since payment channels have a capacity, as indicated by the edge labeling in Figure 1, a path’s transaction volume is limited by the smallest payment channel capacity of this path. While we cannot eliminate this limit, we can use multiple paths, which in sum provide a higher transaction volume.

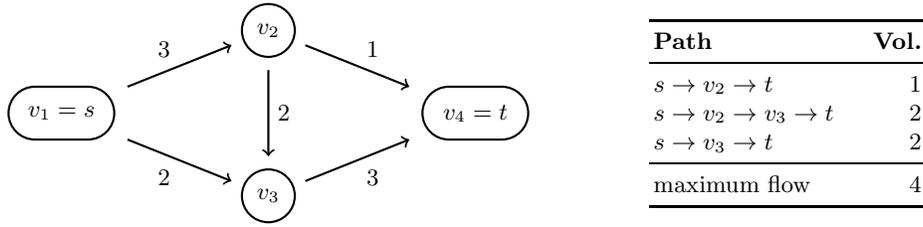


Fig. 1. Payment channel network example.

Determining the maximum transferable amount poses a challenge. For example, simply finding all paths from source to sink and summing up their respective capacities does not suffice; paths may have common edges and thus need to share the respective capacities. For the example in Figure 1, this naive approach would violate payment channel capacities.

The problem of finding the largest payment flow between two nodes s and t in a capacitated flow network is known as the *maximum-flow problem*. Several algorithmic solutions to the maximum-flow problem exist. In the following, we elaborate on the efficient and well-studied *push-relabel* [4] algorithm and adopt it for the route selection of payment flows in payment channel networks.

We consider a network of payment channels as a directed graph $G = (V, E)$ and a non-negative function $c : V \times V \rightarrow \mathbb{R}_{\geq 0}$. We call c the capacity function, which determines a channel's capacity $c(u, v)$ with $u, v \in V$ and $(u, v) \in E$. Moreover, nodes s and t are the source and sink of the flow. The resulting network $F = (G, c, s, t)$ is called a *flow network*.

Definition 1 (pseudo-flow, pre-flow, feasible flow). A *pseudo-flow* on the capacitated graph (G, c) is a mapping $f : V \times V \rightarrow \mathbb{R}$ with the properties:

$$\begin{aligned} f(u, v) &\leq c(u, v), \forall (u, v) \in E && \text{(capacity constraint)} \\ f(u, v) &= -f(v, u), \forall (u, v) \in E && \text{(skew symmetry)} \end{aligned}$$

Note that pseudo-flows do not require incoming and outgoing flows of a node to be equal. Therefore, nodes can hold **excess flow**, denoted by

$$x_f(u) = \sum_{v \in V} f(v, u) - \sum_{v \in V} f(u, v).$$

A *pre-flow* and a *feasible flow* are special kinds of pseudo-flows with one of the following constraints. A **pre-flow** requires

$$x_f(v) \geq 0, \forall v \in V \setminus \{s, t\} \quad \text{(non-negativity constraint)}$$

and a **feasible flow** requires

$$x_f(v) = 0, \forall v \in V \setminus \{s, t\} \quad \text{(conservation constraint)}.$$

Definition 2 (residual capacity and residual graph). The *residual capacity* c_f with regard to the pseudo-flow f of an edge $(u, v) \in E$ is defined as the difference between the edge's capacity and its flow:

$$c_f(u, v) = c(u, v) - f(u, v).$$

Then, the *residual graph* $G_f(V, E_f)$ indicates when changes can be made to flow f in the network $G(V, E)$, where

$$E_f = \{(u, v) \in V \times V : c_f(u, v) > 0\}.$$

Note, that edges (u, v) do not have to be in the original set of edges E .

Definition 3 (height function). A mapping $h : V \rightarrow \mathbb{N}$ is a *height function* for the push-relabel algorithm, if

$$h(s) = |V|, \quad h(t) = 0, \quad h(u) \leq h(v) + 1, \quad \forall (u, v) \in E_f.$$

At the beginning, the generic push-relabel algorithm initializes node heights and flow excess, as well as the edge pre-flow values with 0. Please note that source node s , in contrast to all other nodes, is set to a height $|V|$. Moreover, s 's outgoing edges are saturated according to the height function's third condition. After these initialization steps, the algorithm repeatedly selects a node u as active node and applies one of the two basic operations **push** and **relabel**. Both operations have mutually exclusive conditions, which ensure that either **push** or **relabel** is applicable at a time.

The **push** procedure (cf. Procedure 1) tries to push an excess δ from node u towards a neighbor v with a smaller height. The maximum possible δ is determined as the minimum between the excess flow and the residual capacity of edge (u, v) . Accordingly, edge capacities and excess values are updated to reflect flow changes in the residual graph. The procedure requires that u has excess flow and that an unsaturated edge (u, v) to a neighbor v one level below u exists.

Eventually, node u will saturate all outgoing edges that lead to neighbors on a lower level. In this case, the **relabel** procedure (cf. Procedure 2) "raises" node u to a higher level. The procedure calculates the minimal height of its neighbor nodes and sets u 's height to the level above this minimum. Therefore, the excess of node u is guaranteed to be "pushable" in the next step.

The generic push-relabel algorithms continues until the conditions fail for all nodes. That means, the highest possible transaction volume has been pushed to the sink t and all network excess has been pushed back to the source, i. e., $x_f(v) = 0, \forall v \in V$. At this point, the push-relabel algorithm has transformed the pre-flow into a maximum flow and hence solved the maximum-flow problem.

In a payment channel network, however, it is often not necessary to know the maximum transaction volume. Rather, we want to find a payment flow that can process a certain amount only. This is a slightly different problem, which is known as the *feasible-flow problem*. Fortunately, the push-relabel can easily be modified to solve the feasible-flow problem: in order to find a payment flow from source

Procedure 1 `push(u, v)`

Conditions: $x_f(u) > 0, c(u, v) > 0, h(u) = h(v) + 1$

$$\delta := \min(x_f(u), c_f(u, v))$$

$$f(u, v) := f(u, v) + \delta; f(v, u) := f(v, u) - \delta$$

$$x_f(u) := x_f(u) - \delta; x_f(v) := x_f(v) + \delta$$

Procedure 2 `relabel(u)`

Conditions: $x_f(u) > 0, \forall (u, v) \in E : h(u) \leq h(v)$

$$h(u) := 1 + \min(h(v) : (u, v) \in E)$$

Fig. 2. Push-relabel algorithm [4], which solves the maximum-flow and the feasible-flow problem in flow networks.

s to sink t with a transaction volume d , we can simply insert a new (virtual) node to the payment network. We call it the *pre-source* s' , with a single edge (s', s) and capacity $c(s', s) = d$. The virtual edge caps the transferable amount at exactly d . We can now still apply the push-relabel algorithm, as described before, to find a feasible-flow in this network.

So far, we assumed only one instance of the push-relabel algorithm. If multiple flows ought to be found subsequently in the same network, the initial flow of one instance is the result of the last instance. A generalization for subsequent flows, however, is easily possible. The following section is dedicated to show how the push-relabel algorithm can be adapted to enable route selection for concurrent and distributed payment flows.

4 Concurrent and Distributed Payment Flows

In payment channel networks, it is desirable to allow a concurrent execution of the route selection algorithm. To this end, simply running multiple instances of the push-relabel algorithm in parallel is not enough: one instance for flow f_1 , for example, could consume the reverse edges' residual capacity that belong to another instance for flow f_2 . We call this issue *capacity stealing*.

The problem domain of finding flows f_1, \dots, f_k for k commodities with source-sink pairs $(s_1, t_1), \dots, (s_k, t_k)$ that meet the total capacity constraint

$$F(u, v) = \sum_{i=1}^k f_i(u, v) \leq c(u, v), \forall (u, v) \in E,$$

are known as *multi-commodity flow problems*.

As our main contribution, we propose a modified push-relabel algorithm that allows to find feasible flows in a concurrent multi-commodity scenario. To this end, we introduce the concept of *capacity locking*: flow volumes are accounted for every commodity independently, while still respecting each payment channel's total capacity constraint. The capacities on the reverse edges created by a flow f_1 are therefore *locked* for another flow f_2 , which prevents capacity stealing.

Definition 4 (locked capacities and new residual capacity). Let the *locked capacity* and *total locked capacity* of flow f_i on edge (u, v) be

$$l_i(u, v) = \max(0, f_i(u, v)) \quad \text{and} \quad L(u, v) = \sum_{i=1}^k l_i(u, v).$$

Accordingly, the *residual capacity* is redefined as

$$c_i(u, v) = c(u, v) - L(u, v) + l_i(v, u),$$

which yields an individual residual graph $G_i(V, E_i)$ for each commodity i .

Definition 4 ensures that there is always enough residual capacity on the reverse edges available to push the existing excess back to the source. Except for this augmented definition of the residual capacity, the **locked-push** procedure (cf. Procedure 3) is similar to the original **push** procedure. Note, however, that the modified push-relabel algorithm does not necessarily yield optimal flows in the multi-commodity scenario. It guarantees *validity*, though, which makes it superior compared to other approaches from this domain [1].

In the following, we prove validity for our proposed algorithm. As the skew-symmetry and flow-conservation constraints follow directly from the definition of the algorithm, it suffices to show that it yields flows that respect the total capacity constraint.

Lemma 1. *The total capacity constraint $F(u, v) \leq c(u, v)$, $\forall (u, v) \in E$ is never violated.*

Proof. For a **locked-push** of commodity i on edge (u, v) , the change in flow volume δ is always chosen to be at maximum the remaining residual capacity of the flow on this edge. Accordingly, lock $l_i(u, v)$ cannot be greater than δ . Therefore, the locked capacity never exceeds the edge capacity for each individual edge. It follows that the total capacity constraint is never violated:

$$F(u, v) = \sum_{i=1}^k f_i(u, v) \leq L(u, v) = \sum_{i=1}^k l_i(u, v) \leq \sum_{i=1}^k c_i(u, v) \leq c(u, v). \quad \square$$

In order to execute the modified algorithm in a distributed scenario, the asynchronous distributed algorithm, introduced in [4], is adapted to our needs: each node maintains a local view on flow states, channel capacities, and its neighbors' height. Furthermore, each node maintains routing information and its own height. Then, every node u with positive excess tries to push its excess along an unsaturated outgoing edge to a neighbor v of smaller height. A **locked-push** can only be committed, if v acknowledges u that it is has indeed a smaller height. Alternatively, v can reject the **locked-push** and respond with its actual height. This way, u learns its neighbors' height and can trigger **relabel**, if necessary. After relabeling, u sends height updates to its neighbors. The source and sink node can determine the termination of the algorithm and communicate the result to finalize route selection.

Procedure 3 `locked-push(i, u, v)`

Conditions: $x_i(u) > 0, c_i(u, v) > 0, h_i(u) > h_i(v)$
 $l_i(u, v) := \max(0, f_i(u, v)); l_i(v, u) := \max(0, f_i(v, u))$
 $c_i(u, v) := c(u, v) - L(u, v) + l_i(v, u)$
 $\delta := \min(x_i(u), c_i(u, v))$
 $f_i(u, v) := f_i(u, v) + \delta; f_i(v, u) := f_i(v, u) - \delta$
 $L(u, v) := L(u, v) + \delta; L(v, u) := L(v, u) - \delta$
 $x_i(u) := x_i(u) - \delta; x_i(v) := x_i(v) + \delta$

Fig. 3. Capacity locking enables concurrent push-relabel execution without violating capacity constraints, i. e., capacity stealing.

5 Evaluation

In order to evaluate our approach, we constructed a Watts-Strogatz graph with $\beta = 0.5$, $n = 200$, and a node degree of 10. Channel capacities were generated by uniform random sampling from $[0, 10]$. In the following, we compare the sequential (seq.) and the concurrent (conc.) algorithm.

First, we are interested in the number of flows that each algorithm can handle. To this end, we sampled the transaction volume from $[0, 20]$ and calculated the mean success rate over 10 runs, i. e., the share of successfully found flows. The results, shown on the left of Figure 4, indicate that both algorithms are able to find a large number of flows (relative to the network size). At some point, when network capacities are exhausted, the success rate eventually drops. Single-path approaches, in contrast, achieve in the best case a 0.5 success rate (cf. horizontal line in the plot): while the maximum channel capacity is 10, on average every second transaction volume is in $(10, 20]$ and therefore not feasible with a single path. Effectively, this reduces the utilization of the available capacities by 50%.

Second, we are interested in the transaction volume that we can achieve by aggregating multiple paths. To this end, we set the number of flows to 128, increased the transaction volume, and calculated the mean success rate. The results, shown on the right of Figure 4, suggest that again both variations are able to route relatively large volumes. In more than 50% of the cases, the concurrent algorithm still manages to process all 128 flows for up to a volume of 15 each. This is especially noteworthy, as a single-path approach would not be able to route a single payment with a volume exceeding 10 in our scenario (cf. vertical line in the plot). These first results illustrate that our approach is superior compared to single-path route selection schemes.

6 Conclusion

In this paper, we argued that currently deployed single-path routing schemes for payment channel networks suffer from a number of drawbacks. Most prominently, they utilize the available capacities in the network inefficiently. Eventually, single-path routes will lead to on-chain transactions as a fallback strategy and therefore subvert the idea of payment channels.

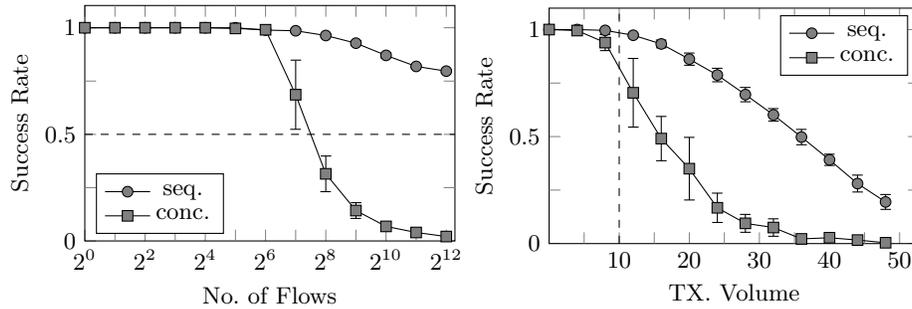


Fig. 4. Flow Network Simulation: mean success rate over 10 runs, dependant on the number of flows and transaction volume. Error bars show the 95% confidence interval.

We addressed this issue by presenting a novel perspective on route selection that considers payment channel networks as flow networks. Flow network algorithms utilize the available capacity by aggregating multiple paths, which allow to route transactions of larger volume. We proposed an extended push-relabel algorithm that finds flows based on local knowledge. Thus, it is suitable for the concurrent and distributed scenario encountered in payment channel networks. We proved the validity of the flows and showed that our algorithm is indeed able to satisfy demands, where single-path based approaches fail.

References

1. Awerbuch, B., Leighton, T.: Improved approximation algorithms for the multi-commodity flow problem and local competitive routing in dynamic networks. In: STOC '94: Proceedings of the 26th Annual ACM Symposium on Theory of Computing. pp. 487–496 (May 1994)
2. Croman, K., Decker, C., Eyal, I., Gencer, A.E., Juels, A., Kosba, A.E., Miller, A., Saxena, P., Shi, E., Sirer, E.G., Song, D., Wattenhofer, R.: On scaling decentralized blockchains - (A position paper). In: BITCOIN '16: Proceedings of the 3rd Workshop on Bitcoin Research. pp. 106–125 (Feb 2016)
3. Decker, C., Wattenhofer, R.: A fast and scalable payment network with bitcoin duplex micropayment channels. In: SSS '15: Proceedings of the 17th International Symposium on Stabilization, Safety, and Security of Distributed Systems. pp. 3–18 (Aug 2015)
4. Goldberg, A.V., Tarjan, R.E.: A new approach to the maximum-flow problem. J. ACM 35(4), 921–940 (1988)
5. Lightning Network: In-progress specifications, <https://github.com/lightningnetwork/lightning-rfc>, accessed on 14.6.2017.
6. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008)
7. Poon, J., Dryja, T.: The bitcoin lightning network: Scalable off-chain instant payments (Jan 2016)
8. Prihodko, P., Zhigulin, S., Sahno, M., Ostrovskiy, A., Osuntokun, O.: Flare: An approach to routing in lightning network (2016)