On the (Not So) Surprising Impact of Multi-Path Payments on Performance and Privacy in the Lightning Network

Charmaine Ndolo[1] and Florian Tschorsch[1,2]

[1] Humboldt-Universität zu Berlin, Berlin, Germany
c.n.ndolo@hu-berlin.de
[2] Technische Universität Dresden, Dresden, Germany
florian.tschorsch@tu-dresden.de

Abstract. The Lightning network (LN) addresses Bitcoin's scalability issues by providing fast and private payment processing. In order to mitigate failures caused by insufficient channel capacities, LN introduced multi-path payments. To the best of our knowledge, the effect of multi-path payments remains unclear. In this paper, we therefore study the impact of multi-path payments on performance and privacy. We identify metrics quantifying the aforementioned properties and utilise them to evaluate the impact of multi-path payments. To this end, we develop a simulator implementing pathfinding in LN using single and multi-path payments as well as various pathfinding algorithms. We find that, while the success rate of multi-path payments is up to 20% higher, the impact of multi-path payments on performance otherwise remains within limits. On the other hand, the impact on privacy appears to be greater, e.g., multi-path payments are more likely to encounter an on-path adversary and the relationship anonymity is more likely to be compromised by colluding intermediate hops. However, multi-path payments are less likely to be deanonymised based on the path lengths.

1 Introduction

Layer 2 solutions such as the Lightning network (LN) [21] offer a solution to Bitcoin's scalability problem by means of a payment channel network (PCN). A PCN is a network of off-chain payment channels, each between two parties, in which funds can move in either direction as long as both parties agree. LN facilitates fast payment processing by limiting the need for global consensus to a subset of states [21]. In addition to speed, privacy is a focal objective in LN leading to various privacy-enhancing measures. For instance, while channel capacities are announced to the public network, the individual endpoints' balances are kept private. LN also supports multi-hop payments allowing the routing of payments over multiple intermediate nodes.

The final publication will be available in the proceedings of the 7th International Workshop on Cryptocurrencies and Blockchain Technology (CBT 2023).

Finding such paths is an essential part of LN and is delegated to the sender of a payment. i.e., multi-hop payments are source-routed. The pathfinding algorithm in LN is typically accomplished using Dijkstra's shortest path algorithm [5] and a fee-based cost function.[3] Given the uncertainty over balances, a path chosen by a sender, may fail due to insufficient balances along the way. To ameliorate this challenge, a lot of effort has been devoted to the question of efficient routing, e.g., [2, 22, 26] and notably [20] whose authors proposed the selection of paths based on the probability of a payment succeeding.

It has been shown that LN generally performs well with lower payment volumes, but suffers from degradation with increasing payment volumes due to a lack of channels with sufficiently high capacity [3,31]. As a response, the network introduced *multi-path payments (MPP)* and *atomic multi-path payments (AMP)* as an alternative payment scheme [6,8,18]. Such payments allow splitting a payment amount into multiple payment parts of lesser value and thereby maximise on the entire available flow between sender and receiver. A crucial difference between MPPs and AMPs is that the former use the same payment hash for all parts and AMPs are atomic.

To the best of our knowledge, the relationship between performance and privacy in conjunction with multi-path payments remains open. Privacy in LN has been shown to be susceptible to various attacks such as balance-revealing [4, 9] and deanonymisation attacks [11,15,23]. It is probable that multi-path payments may heighten privacy concerns due to payment data traversing the network on multiple occasions, e.g., with respect to correlation attacks.

In this work, we study the impact of multi-path payments on performance and privacy in LN empirically using network simulations. We include fee and probability-based pathfinding in our analysis as the pathfinding algorithm plays a role in the outcome of key routing parameters. Among others, we find that high-volume payments are more likely to succeed as multi-path payments which are also less likely to be deanonymised based on the path lengths. Where applicable, we contextualise our results with earlier research. The main contributions of this work can be summarised as follows:

1. we identify various metrics to quantify performance and privacy in the LN;
2. we compare the single and multi-path payments w.r.t the identified metrics in combination with fee and probability-based pathfinding algorithms; and
3. we implement an LN simulator providing us with empirical, simulation-based results on the impact of multi-path on performance and privacy in LN.

The remainder of this paper is structured as follows. Sec. 2 provides an overview of LN as well as our methodology including the identified metrics for performance and privacy. We present and discuss our results in Sec. 3 and Sec. 4. We summarise related literature in Sec. 5 and conclude in Sec. 6.

[3] `https://github.com/lightning/bolts/blob/master/07-routing-gossip.md#requirements-9`

2 Network Model

Once a (bidirectional) channel in LN has been established between two parties, an arbitrary number of payments can be made between them. By opening a channel, a fixed number of funds is committed – known as the *capacity* – which can be disposed of freely within the channel. Transactions between the two channel endpoints alter the parties' *balances*, i.e., each node's share of the channel's capacity. Motivated by privacy concerns, node balances are kept private.

For reasons mainly related to practicability, LN is not a complete network in which every pair of nodes has a channel. Instead, channels form a PCN that enables routing payments between parties via multiple intermediate hops. Multi-hop routing requires that there must be a set of channels linking the sender and the recipient and essentially boils down to a shortest path problem. The PCN is therefore commonly modelled and reasoned about as a (directed) graph [3, 28].

Definition 1 (The Lightning network graph). *The LN graph is a directed multigraph $G = (V, E)$ where V is the set of Lightning nodes and E the multiset of payment channels in the network.*

Note that while channels are bidirectional, attributes such as each node's fee structure make it necessary to distinguish the direction of an edge in G when reasoning about pathfinding. Hence, it is necessary to define G as a directed graph. G is a multigraph as, in practice, a pair of nodes may have more than one channel between them.

Payments in Lightning are source-routed, i.e., the sender is responsible for finding a path to the recipient, and is typically accomplished using some form of Dijkstra's shortest path algorithm [5]. The LN specification, Basis of Lightning Technology (BOLT) [1], defines the edge weights using a fee-minimising cost function based on channel capacities, fees and locking duration. At the time of writing, routing nodes in LN charge two types of fees – a *base fee* that is due regardless of the amount in question as well as a *proportional fee* that is scaled by the amount to be forwarded. Given the cost function, a shortest path search algorithm is expected to return the cheapest path between two nodes. Due to the uncertainty over balances, this cost function often leads to failed payments as a result of insufficient liquidity [20].

Based on the observation that channels with higher capacity provide a higher chance of success, Pickhardt et al. propose to select paths based on the *success probability* [20], which is the product of the involved channels' individual success probabilities. The lower the ratio of payment amount and channel capacity, the higher the success probability. In this case, a shortest path search algorithm returns the channel with the highest probability of success.

Further design details of payment channels in general and PCNs in particular, e.g., on the atomicity of multi-hop payments, can be found in [1, 13, 21, 29].

2.1 Performance Metrics

In the following, we describe metrics quantifying the performance of LN and PCNs in general. While some of the described metrics can already be found in existing literature, we identify additional metrics that encapsulate the *utility* and *usability* of the network from a user's perspective.

We begin with the *success rate* as a rudimentary measure of performance which we define as the ratio of successful payments and the total number of payments [16]. Furthermore, we study the amount of *transaction fees* due for successful multi-hop payments. In the case of multi-part payments, we consider each part individually. The *path length* is a well-known measure of network topologies and quantifies the number of edges (channels) a payment traverses before arriving at its recipient. Similar to the transaction fees, we consider the partial payment paths independently. The path length is only relevant to successful payments as failed payments do not have a complete path between sender and receiver. As a final performance metric, we use the *number of payment attempts* as an indicator of routing efficiency. We define the number of payment attempts as the number of Hashed Timelock Contracts (HTLCs) triggered by a single payment before it eventually succeeds or fails. We define the number of attempts for a multi-path payment as the sum of all parts' attempts.

2.2 Privacy Metrics

Although LN (and layer 2 solutions in general) strives for improved privacy, compared to on-chain solutions, recent works have identified shortcomings in the privacy provided by LN [13, 23, 29, 30]. We compile measures quantifying privacy in LN in what follows.

Unless stated otherwise, we assume an *honest but curious (HBC)* adversary. Such an adversary is a legitimate participant in LN who will not deviate from the protocol but will try to infer as much as possible from observed messages [17]. Given that an HBC adversary has limited capabilities, we consider these properties to be a lower bound on privacy in the network.

Observation rate: We quantify how often an *on-path* adversary observes payments using what we call the *observation rate*. The observation rate is the proportion of the number of payments that encountered an adversary in any of their attempted paths and the total number of payments. This metric has previously been studied in [29]. In the case of multi-path payments, we define the observation rate as the proportion of payments that include such a node in at least one of their parts' attempted paths.

A high observation rate is a result of either a high number of watchers or, more plausibly based on the properties of LN channel graph [12, 23, 27], routing hubs that forward a great number of payments.

Sender and Receiver Inference: After analysing the length of payment paths in LN, Kappos et al. set up a formula defining the probability that a node's predecessor and successor in a payment's path are the respective payment endpoints [10]. The formula is based on the path length probability distribution and estimates the probabilities Pr_s^{succ} and Pr_s^{fail} of correctly identifying the sender of successful and failed payments [10]. The probabilities Pr_r^{succ} and Pr_r^{fail} for the payment's destination are calculated analogously.

Shorter, unsuccessful paths lead to the highest probabilities whereas longer, successful paths exhibit the lowest probabilities. We propose to extend this measure to multi-path payments by handling payment parts as individual payments.

Relationship Anonymity: Tikhomirov et al. examine the probability of a successful path being vulnerable to a confirmation attack [29], i.e., a path $s, i_1, ..., i_n, d$ in which the hops i_1 and i_n are under adversarial control.

We extend this measure to payments by characterising a payment as vulnerable if at least one of its paths is vulnerable in that both the first and last hops are controlled by an adversary. The measure is identical to the one in [29] for single payments as they have exactly one path if successful. Assuming an on-path adversary is able to correlate payments, e.g., using common identifiers such as the condition to fulfil an HTLC, successfully deanonymising one payment path is sufficient to determine the sender-recipient pairs for the remaining parts.

Path Diversity: We introduce *path diversity* as a further measure of privacy in the LN, i.e., we want to identify how (dis)similar the paths of a multi-path payment are with respect to the routing nodes and edges. It does not make much sense to examine single payments in this context as there is only one path for each such payment. Path diversity is desirable from a privacy standpoint in order to reduce the number of payments being observed by the same node so as to, e.g., hamper correlation attacks. A lack of path diversity is also suggestive of an (over)reliance on some nodes and edges which is unhealthy for a network in respect to resilience. However, path diversity also means more nodes are involved in delivering a given payment likely leading to an increase of the observation rate.

We propose to quantify the path diversity for a set of payment paths using the *effective path diversity (EPD)* measure defined by Rohrer et al. [25]. The EPD is the degree to which a set of paths between the same source s and destination d share common intermediate nodes and edges. It is an aggregation of path diversities for a set of paths between a given node pair (s, d) and defined as

$$EPD = 1 - e^{-\lambda k_{sd}} \in [0, 1)\,,\ k_{sd} = \sum_{i=1}^{k} D_{min}(P_i), \tag{1}$$

where k is the number of paths and $D_{min}(P_i)$ is the minimum diversity of path i when measured against all previously selected paths. The constant λ scales the impact of k_{sd} based on the utility of added diversity. Lower marginal utility is indicated by a high value of λ (> 1) whereas a low value of λ represents a higher marginal utility. We argue that lower values of λ are more representative of the significance of diverse paths in LN.

2.3 Network Simulations

In order to analyse the effects of the different payment types combined with different pathfinding approaches, we developed a tool to simulate pathfinding and payment delivery in LN using algorithms similar to those used in practice.

The simulator is publicly available in our accompanying Git repository.[4] The simulator reads LN snapshots to reconstruct the PCN according to the network model in Def. 1. It supports probability-weighted [20] and fee-weighted pathfinding[5] as well as single and multi-path payments. The simulator implements a *trial-and-error* loop and attempts to deliver a failed payment by looking for an alternative path until the set of possible paths is exhausted. In case of ambiguity in the Lightning specification, e.g., on the maximum number of parts a payment may be split into, we followed `LND`'s implementation as it is the most commonly used client [32]. To this end, the simulator only attempts to split payments greater than 10k sat and into at most 16 parts.

We utilised a channel graph dated 15^{th} May 2023, which contains $18,820$ nodes and $134,838$ edges and was collected from our own well-synchronised LN node. We discarded nodes and edges without essential data for the simulation such as fee structure and reduced the graph to its largest strongly connected component leading to a graph with $14,453$ nodes and $134,782$ edges.

Given the private nature of node balances, the simulator splits the channel capacity into two balance values following a uniform distribution at the beginning of a simulation (see [20] for discussions on the distribution of capacities) and updates the node balances after every successful payment delivery. We simulate various payment volumes following the categorisation of payments in [7] as actual volumes are unknown, i.e., *micro* payments, *medium* payments and *macro* payments. We chose not to make assumptions about patterns between transacting parties and simulated $5,000$ transactions between random sender-receiver pairs for each selected amount ranging between 100 sat and 10 million (m) sat. We repeated each simulation scenario multiple times with different seeds for the random number generator, i.e., for each set of $5,000$ sender-receiver pairs, the channel graph was initialised before simulating payment delivery of each amount with all four combinations of pathfinding algorithm and payment type.

While the results in this work are based on an implementation of MPP in that the same payment identifier is used for all parts, the results can be generalised to AMP, e.g., by assuming an attacker is able to identify related parts. Furthermore, the simulator ensures that all multi-path payments are atomic, i.e., either all parts succeed or no funds are moved at all.

3 Impact on Performance

We present and discuss our results pertaining to the performance of LN based on the metrics presented in Sec. 2.1. All in all, our simulations confirmed either what previous works have already established or what we can expect given what we know about the network. We omit some charts due to space constraints but provide interactive versions of all charts in our accompanying repository.[6]

[4] `https://github.com/cndolo/lightning-simulator`

[5] `https://github.com/lightning/bolts/blob/master/07-routing-gossip.md#htlc-fees`

[6] `https://cndolo.github.io/lightning-simulator`

3.1 Success Rate

We observe that the choice of pathfinding method is not significant for the success of the simulated payments. Instead, the payment amount – limited by channel capacities – is the decisive factor. The type of payment plays a secondary role for the success rate in that larger payments are more successful when routed as multi-path payments. Most failures in LN are due to an insufficient maximum flow between sender and receiver, i.e., there is no path between sender and receiver where all of the path's edges have enough capacity to forward the requested amount [3]. Clearly, multi-path payments are at an advantage over single payments when it comes to utilising the maximum flow because a payment can be delivered via multiple paths.

As the payment amount increases, probability-weighted payments begin to show a very slight advantage of up to 2% over their respective fee-weighted counterparts. Multi-path payments start to separate themselves at payment volumes $\geq$10k sat when payments may actually be split and succeed approximately 20% more often than their single counterparts. However, less than 2% of the payments worth 5m sat and greater succeeded regardless of payment type.

3.2 Transaction Fees

In general, we noticed that the amount of absolute fees increases with the payment amount. Furthermore, paths selected based on the success probability are more expensive than fee-weighted paths. This is not surprising and indeed expected given that fee-weighted pathfinding selects paths by minimising the total amount of fees whereas probability-weighted pathfinding disregards the fees.

While multi-path payments mostly incur slightly higher fees than comparable single payments, the additional cost of splitting a payment seems to be negligible. The difference is partly due to the *base* fees charged by some nodes in LN. In the absence of base fees, we expect close to no difference between single and multi-path payments provided that all parts take paths with similar fee policies. Tochner et al. found that fee policies in LN mostly follow the same structure [30].

The necessity of the base fee in LN has been questioned [19], however, it is yet to be eliminated completely. At the time of writing, 50% of the channels in LN have adopted this proposal and do not charge a base fee.[7] The impact of base fees becomes clearer when looking at the charged fees relative to the payment volume. Lower payment volumes such as 100 sat were the most expensive regardless of the pathfinding method or payment type.

3.3 Path Length

The distribution of the successful paths' lengths is depicted in Fig. 1. At lower payment volumes, some fee-weighted paths are significantly longer than probability-weighted paths with some even at the maximum permitted hop count of 20.

[7] According to `https://lnrouter.app/graph/zero-base-fee`.

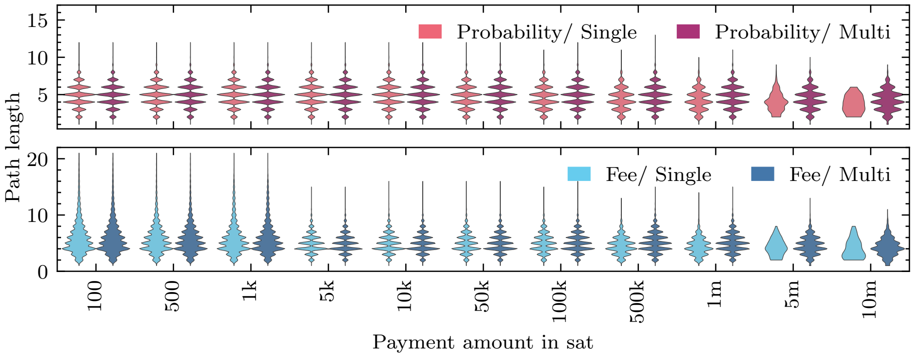

Fig. 1. The distribution of successful payments' path lengths.

For payment volumes of up to 100k sat, all combinations yield a constant median path length of 5. The median path length drops to 4 at payment volumes ≥500k sat for single and ≥5m sat for multi-path payments. In anticipation of discussions on privacy in Sec. 4.2, shorter paths have a negative impact on privacy in LN. The results presented here align with findings presented in [10].

3.4 Payment Attempts

The total number of payment attempts recorded during simulation of the different combinations for various amounts is shown in Fig. 2. The total number of attempts remains almost constant for all simulated amounts. As the payment volume increases, we notice that the number of attempts made by single payments gradually decreases for payments greater than 50k sat. In cases where no capable routes are found, payments fail without recording any attempts thus leading to the decline in the number of attempts for single payments. This claim is supported by the results in Sec. 3.1 where we recorded an almost zero success rate for the highest payment amounts. When looking at the percentage of successful attempts, we find that, while the total number of attempts remains almost constant, most of the HTLCs initiated by multi-path payments are not fulfilled in contrast to single payments.

Furthermore, probability-weighted pathfinding requires marginally fewer attempts, which becomes evident as the payment amount increases. This is because of the fundamental premise that probability-weighted pathfinding prefers channels with endpoints that are more likely to have sufficient routing balance leading to fewer iterations of the trial-and-error loop.

3.5 Insights

- The payment volume plays the most significant role in the success of a payment. We observe a previously confirmed inverse relationship between the

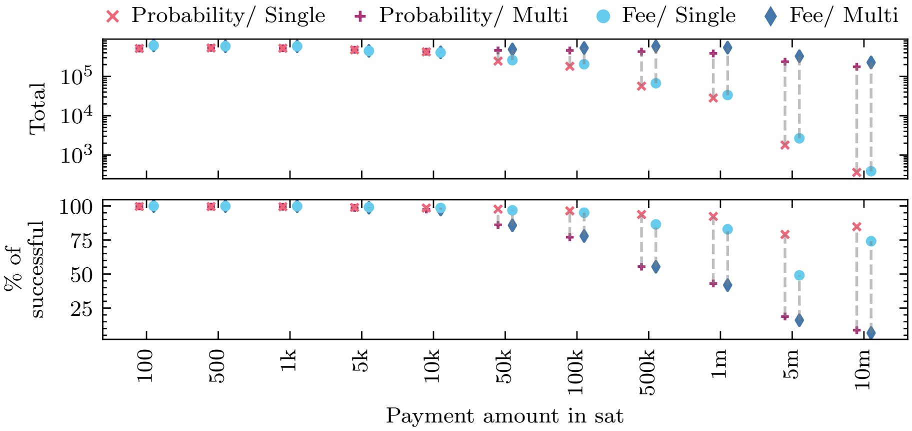

Fig. 2. The total number of HTLC attempts and the percentage of successful attempts.

payment volume and success rate [3, 31]. The type of payment plays a secondary role with multi-path payments able to deliver up to 20% more high-volume payments than single payments.

- The impact of multi-path payments on fees proved to be marginal. In light of the gradual elimination of base fees, we expect the additional fees accrued by multi-path payments to diminish. Furthermore, our simulations show that the pathfinding algorithm plays a vital role in the accumulated fees as routes computed based on the success probability are significantly more expensive than fee-weighted payments. While the main result is not unexpected, we have been able to quantify that probability-weighted paths charge between 3% and 10% more than fee-weighted paths in relative fees.
- Our results on path lengths indicate that it is not quite determined by the payment type but more by the pathfinding approach. The payment amount in question plays a minor role although the difference between the median path lengths for different amounts is not significant.
- We find that the number of additional payments triggered by multi-path payments is reasonable, however, a quick glance at the relative values shows that more and more of these attempts are futile as the payment volume increases. The heightened success rate comes at the price of more network activity. Furthermore, we establish that probability-weighted pathfinding is more efficient than fee-weighted pathfinding with regard to the number of payment attempts.

4 Impact on Privacy

As the outcome of some of the discussed privacy measures is heavily dependent on an adversary's standing in the network, we executed our simulations with two different adversary selection strategies on the same set of payments. Similar

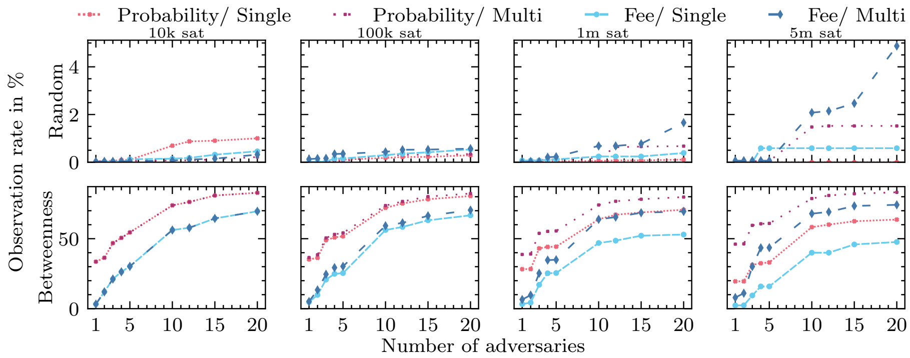

Fig. 3. The observation rate for successful payments for selected payment volumes, various adversary selection strategies and a varying number of adversarial nodes.

to [11] and [14], we characterised up to 20 nodes as adversaries based on betweenness centrality and random selection. Note that the betweenness centrality was computed without weights.

4.1 Observation rate

Fig. 3 shows the observation rate for successful payments using two different adversary selection strategies. Unsurprisingly, the random selection of adversaries results in a very low observation rate whereas central nodes observe a high number of payments. These results are to be expected given the underlying scale-free topology [3, 23] of the channel graph and hint at the presence of routing hubs.

The observation rate is higher for multi-path payments and is highest with probability-weighted pathfinding. With only 15 adversarial nodes, over 70% of the payments were observed by a central adversary. These findings are indicative of a relation between centrality and capacity because probability-weighted pathfinding deliberately looks for high capacity channels (in proportion to the payment's value).

An explanation for the higher observation rate when using multi-path payments is the triggered payment attempts (cf. Fig. 2). Accordingly, multi-path payments have a higher observation rate than single payments and are otherwise identical with regard to the different pathfinding methods.

From a privacy standpoint, the presence of hubs is indeed problematic as these central watchers observe a fair share of the payments allowing them to gather an abundant amount of information. For instance, such a node could profile users in the case of regular payments of a certain amount taking the same (sub-)path. The importance of path diversity in the network becomes evident.

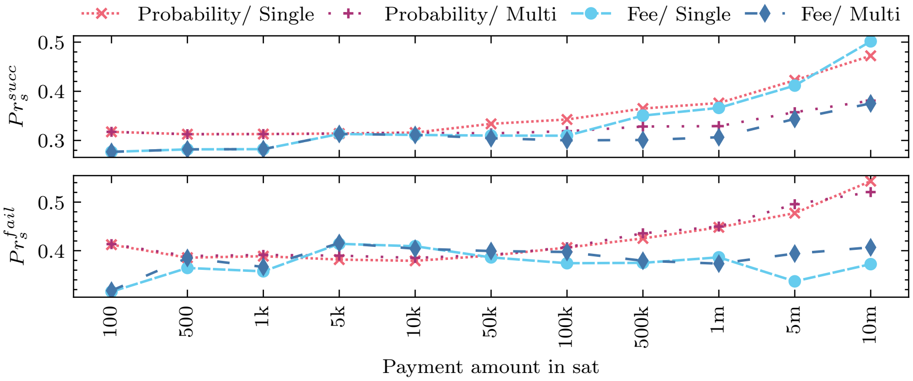

Fig. 4. The probability Pr_s that a node's predecessor is the payment's sender.

4.2 Sender and Receiver Inference

Having studied the odds of encountering an adversary on a payment's path, we examine what information can be gained by such an observation. The probabilities of correctly deanonymising the sender are depicted in Fig. 4. Recall that payment parts are handled individually and that the probabilities for receiver deanonymisation Pr_r are equal to the corresponding Pr_s.

Successful, single payments are more likely to be deanonymised with the probabilities increasing for higher payment amounts. It may seem counter-intuitive that the sender of multiple payment parts is harder to deanonymise but it is resultant of the individual parts' path lengths and how often they occur. There are fewer successful single payments at higher volumes and given that the density of single payments' path lengths around the median increases (see Fig. 1), the probability of a path being of that length rises. Besides, this measure does not try to correlate observed payments.

We also observe that the probabilities Pr_s^{succ} generally increase as the payment volume increases. Given that we know from Sec. 3.3 that the path lengths not only tend to get shorter as the payment volume increases but also same-length paths become more common, Pr_s^{succ} is expected to increase. As every additional edge increases the risk of payment failure, the availability of short, liquid paths is a desirable property for a PCN like Lightning. However, precisely this property has a negative impact on the anonymity.

The odds shift slightly when looking at the probabilities Pr_s^{fail} for failed paths in that fee-weighted paths are easier to deanonymise than probability-weighted paths for payments >1k sat and <100k sat. Outside of this range, probability-weighted paths continue to be more likely to be deanonymised. We also find that Pr_s^{fail} for the two payment types do not differ greatly when using the same pathfinding approach and establish that the pathfinding method is decisive for the sender/receiver inference of failed payments.

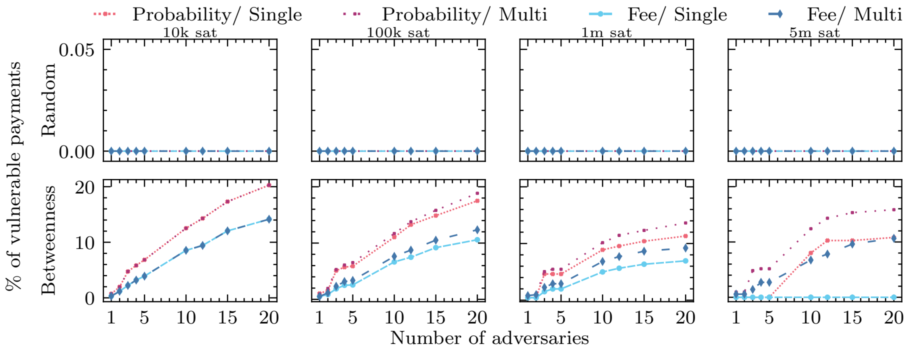

Fig. 5. The percentage of payments in which the first and last hop of at least one part are adversarial.

4.3 Relationship Anonymity

The percentage of payments vulnerable to deanonymisation based on colluding intermediate hops is shown in Fig. 5. Similar to [29], we find that a random selection of adversaries does not compromise the relationship anonymity. On the other hand, the impact of central adversaries is already evident at just a handful of adversaries. 20 adversaries, corresponding to $\ll 1\%$ of the node population, are sufficient to compromise up to 20% of the payments.

As the payment amount increases, the overall share of vulnerable payments decreases, however, the anonymity of multi-path payments is more likely to be compromised. As shown in Fig. 6, the number of payment parts increases thereby increasing the attack surface (in comparison to single payments). With respect to the different pathfinding algorithms, probability-weighted paths are more vulnerable to such a confirmation attack. We hypothesise that this is because its search optimises for shorter paths with relatively high capacity which places well-connected nodes with high-capacity channels at an advantage. This claim is supported by the slight decrease in vulnerable payments as the payment amount increases driving the shortest path search algorithm to deviate to less-preferred paths.

We conclude that the pathfinding approach plays a significant role for relationship anonymity as probability-weighted paths are clearly more susceptible to attacks. Multi-path payments are also more vulnerable to deanonymisation attacks than single payments in case of compromised intermediaries.

4.4 Path Diversity

We applied the EPD measure to every set of paths taken by multi-path payments using different values of λ reflecting the utility of diverse paths and depict the results in Fig. 7. The number of paths k is the number of parts a payment was split into. In general, and regardless of the pathfinding approach, the utilised paths exhibit diversity ranging between 0 and 0.6 for the smallest and greatest tested λ respectively. For payments below 500k sat and 1m sat, the EPD values for fee and probability pathfinding are all 0, which can be attributed to the fewer number of parts needed to complete payments. These scores imply that paths contain very few disjoint hops, e.g., detours around a bottleneck channel, and are otherwise identical.

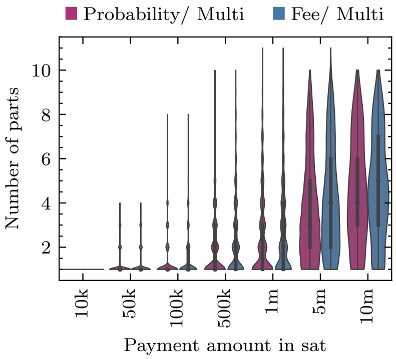

Fig. 6. The number of parts successful multi-path payments were delivered in.

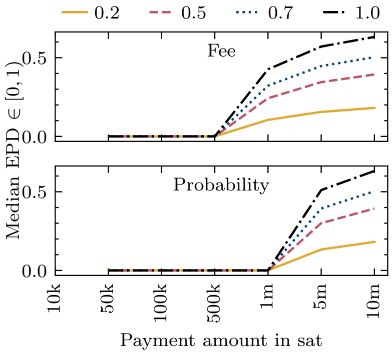

Fig. 7. The median EPD for multi-path payments using various λ.

The EPD values increase as the payment amount increases and show a clear correlation between the number of parts and the diversity. As visible in Fig. 6, the number of parts increases as the payment amount increases. Although higher values of λ result in higher EPD values, the overall progression of the EPDs remains the same. We argue that lower values of λ best signify the utility of diverse paths with respect to the privacy of the Lightning channel graph.

4.5 Insights

- High-centrality nodes observe a fair share of the payments especially when using probability-weighted pathfinding. Furthermore, the observation rate for multi-path payments is higher than that of single payments. The higher observation rate is due to the number of triggered HTLCs and payment parts. A tug-of-war between performance and privacy appears to be evident.
- Our results point to the influence of the payment type on the sender and receiver inference of successful payments based on path length probabilities. We established that successful single payments are more likely to be compromised by a predecessor/successor attack while the pathfinding method plays a bigger role for failed payments.
- In contrast, we find that relationship anonymity is more likely to be compromised by colluding intermediate nodes when using multi-path payments.

- We find that a higher number of payment parts has a positive impact on the path diversity which can lead to better privacy, e.g., by bypassing correlation attacks. Simultaneously, the presence of more diverse paths means a payment has a higher chance of traversing different observation points leading to a higher share of vulnerable paths.

5 Related Work

Given the abundance of literature on PCNs and the LN in particular, we limit ourselves to relevant literature that deals with the interaction between utility and privacy in Lightning. To the best of our knowledge, no prior work covers these properties with regard to multi-path payments.

In an early work on LN, Martinazzi studies the structural properties of the channel graph shortly after its mainnet launch and finds, among others, that the network is resilient to random failures [14]. More recent works also present findings on the structural properties of LN [3, 23, 27, 30]. For example, Béres et al. find that it is possible to deduce transacting parties based on the short path lengths PCN [3]. Kappos et al. study the privacy offered by LN [10], revealing privacy attacks and formalising the findings presented in [3].

Tang et al. study the interplay between privacy and utility in PCNs and show fundamental limits of the established trade-off [28]. They argue that PCNs must choose either low privacy or low utility and cannot offer profound privacy and utility simultaneously. Additionally, Malavolta et al. prove the trade-off between privacy and concurrency in PCNs [13] and show that PCNs can only achieve non-blocking progress at the expense of privacy.

The authors of [18] investigate the utility of multi-path payments for the successful delivery of payments in LN and show that splitting payments leads to a higher success rate. The authors of [20] investigate payment splitting and discuss when and how to split a payment. They conclude that payment splitting is only beneficial for large payments.

Similar to our work, multiple previous works follow an empirical approach and are based on network simulators [3, 10, 11, 24]. However, they either make simplifying assumptions about the routing algorithm, payment distribution or updates to the channel graph during simulation. In addition, the simulator developed in this work contributes to the state of the art by implementing support for multiple payment schemes as well as different pathfinding algorithms.

6 Conclusion

We identified performance and privacy metrics for PCNs and studied the impact of multi-path payments on performance and privacy in LN empirically. In part, we confirm earlier results such as the relationship between the success rate and payment volume. Having recorded notable differences in the fees, number of payment attempts, and path lengths, we find that the choice of pathfinding algorithm has a greater impact on performance than on privacy. Our results indicate that the impact of multi-path payments on performance generally remains within limits, although multi-path payments have a higher success rate. Our results point to a greater impact of multi-path payments on privacy, e.g., the higher chance of encountering an on-path adversary. Remarkably, while such payments are more susceptible to confirmation attacks, they are also less likely to be deanonymised by a simple predecessor/successor attack than single payments. In summary, multi-path payments are especially useful for the delivery of high-volume payments which, however, comes with concerns on privacy. Both payment types showed weaknesses with regard to privacy due to the structure of the channel graph making it difficult to mark one superior to the other.

References

1. Bolt: Basis of lightning technology (lightning network specifications), https://github.com/lightning/bolts
2. Bagaria, V.K., Neu, J., Tse, D.: Boomerang: Redundancy improves latency and throughput in payment-channel networks. In: Financial Cryptography and Data Security - 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10-14, 2020 Revised Selected Papers. Lecture Notes in Computer Science, vol. 12059, pp. 304–324. Springer (2020)
3. Béres, F., Seres, I.A., Benczúr, A.A.: A cryptoeconomic traffic analysis of bitcoins lightning network (2019), http://arxiv.org/abs/1911.09432
4. Biryukov, A., Naumenko, G., Tikhomirov, S.: Analysis and probing of parallel channels in the lightning network. In: Financial Cryptography and Data Security - 26th International Conference, FC 2022, Grenada, May 2-6, 2022, Revised Selected Papers. Lecture Notes in Computer Science, vol. 13411, pp. 337–357. Springer (2022)
5. Dijkstra, E.W.: A note on two problems in connexion with graphs. Numerische Mathematik **1**, 269–271 (1959)
6. Eckey, L., Faust, S., Hostáková, K., Roos, S.: Splitting payments locally while routing interdimensionally. IACR Cryptol. ePrint Arch. p. 555 (2020), https://eprint.iacr.org/2020/555
7. Ersoy, O., Roos, S., Erkin, Z.: How to profit from payments channels. In: Financial Cryptography and Data Security - 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10-14, 2020 Revised Selected Papers. Lecture Notes in Computer Science, vol. 12059, pp. 284–303. Springer (2020)
8. Fromknecht, C., Osuntokun, O.: Bolt 21: Atomic multi-path payments, https://github.com/cfromknecht/lightning-rfc/blob/bolt-amp/21-atomic-multi-path-payments.md
9. Herrera-Joancomartí, J., Navarro-Arribas, G., Ranchal-Pedrosa, A., Pérez-Solà, C., García-Alfaro, J.: On the difficulty of hiding the balance of lightning network channels. In: Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security, AsiaCCS 2019, Auckland, New Zealand, July 09-12, 2019. pp. 602–612. ACM (2019)
10. Kappos, G., Yousaf, H., Piotrowska, A.M., Kanjalkar, S., Delgado-Segura, S., Miller, A., Meiklejohn, S.: An empirical analysis of privacy in the lightning network. In: Financial Cryptography and Data Security - 25th International Conference, FC 2021, Virtual Event, March 1-5, 2021, Revised Selected Papers, Part I. Lecture Notes in Computer Science, vol. 12674, pp. 167–186. Springer (2021)

11. Kumble, S.P., Epema, D.H.J., Roos, S.: How lightning's routing diminishes its anonymity. In: ARES 2021: The 16th International Conference on Availability, Reliability and Security, Vienna, Austria, August 17-20, 2021. pp. 13:1–13:10. ACM (2021)
12. Lin, J., Primicerio, K., Squartini, T., Decker, C., Tessone, C.J.: Lightning network: a second path towards centralisation of the bitcoin economy (2020), https://arxiv.org/abs/2002.02819
13. Malavolta, G., Moreno-Sanchez, P., Kate, A., Maffei, M., Ravi, S.: Concurrency and privacy with payment-channel networks. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017. pp. 455–471. ACM (2017)
14. Martinazzi, S.: The evolution of lightning network's topology during its first year and the influence over its core values (2019), https://arxiv.org/abs/1902.07307
15. Nisslmueller, U., Foerster, K., Schmid, S., Decker, C.: Toward active and passive confidentiality attacks on cryptocurrency off-chain networks. In: Proceedings of the 6th International Conference on Information Systems Security and Privacy, ICISSP 2020, Valletta, Malta, February 25-27, 2020. pp. 7–14. SCITEPRESS (2020)
16. Papadis, N., Tassiulas, L.: Blockchain-based payment channel networks: Challenges and recent advances. IEEE Access **8**, 227596–227609 (2020)
17. Paverd, A.J., Martin, A.C.: Modelling and automatically analysing privacy properties for honest-but-curious adversaries (2014)
18. Piatkivskyi, D., Nowostawski, M.: Split payments in payment networks. In: Data Privacy Management, Cryptocurrencies and Blockchain Technology - ESORICS 2018 International Workshops, DPM 2018 and CBT 2018, Barcelona, Spain, September 6-7, 2018, Proceedings. Lecture Notes in Computer Science, vol. 11025, pp. 67–75. Springer (2018)
19. Pickhardt, R., Richter, S.: Optimally reliable & cheap payment flows on the lightning network (2021), https://arxiv.org/abs/2107.05322
20. Pickhardt, R., Tikhomirov, S., Biryukov, A., Nowostawski, M.: Security and privacy of lightning network payments with uncertain channel balances (2021), https://arxiv.org/abs/2103.08576
21. Poon, J., Dryja, T.: The bitcoin lightning network: Scalable off-chain instant payments (Jan 2016), https://lightning.network/lightning-network-paper.pdf
22. Rohrer, E., Laß, J., Tschorsch, F.: Towards a concurrent and distributed route selection for payment channel networks. In: Data Privacy Management, Cryptocurrencies and Blockchain Technology - ESORICS 2017 International Workshops, DPM 2017 and CBT 2017, Oslo, Norway, September 14-15, 2017, Proceedings. Lecture Notes in Computer Science, vol. 10436, pp. 411–419. Springer (2017)
23. Rohrer, E., Malliaris, J., Tschorsch, F.: Discharged payment channels: Quantifying the lightning network's resilience to topology-based attacks. In: 2019 IEEE European Symposium on Security and Privacy Workshops, EuroS&P Workshops 2019, Stockholm, Sweden, June 17-19, 2019. pp. 347–356. IEEE (2019)
24. Rohrer, E., Tschorsch, F.: Counting down thunder: Timing attacks on privacy in payment channel networks. In: AFT '20: 2nd ACM Conference on Advances in Financial Technologies, New York, NY, USA, October 21-23, 2020. pp. 214–227. ACM (2020)
25. Rohrer, J.P., Jabbar, A., Sterbenz, J.P.G.: Path diversification: A multipath resilience mechanism. In: 7th International Workshop on Design of Reliable Communication Networks, DRCN 2009, Washington, DC, USA, October 25-28, 2009. pp. 343–351. IEEE (2009)

26. Roos, S., Moreno-Sanchez, P., Kate, A., Goldberg, I.: Settling payments fast and private: Efficient decentralized routing for path-based transactions. In: 25th Annual Network and Distributed System Security Symposium, NDSS 2018, San Diego, California, USA, February 18-21, 2018. The Internet Society (2018)
27. Seres, I.A., Gulyás, L., Nagy, D.A., Burcsi, P.: Topological analysis of bitcoin's lightning network. In: Mathematical Research for Blockchain Economy, 1st International Conference, MARBLE 2019, Santorini, Greece, May 6-9, 2019. pp. 1–12. Springer (2019)
28. Tang, W., Wang, W., Fanti, G., Oh, S.: Privacy-utility tradeoffs in routing cryptocurrency over payment channel networks. Proc. ACM Meas. Anal. Comput. Syst. **4**(2), 29:1–29:39 (2020)
29. Tikhomirov, S., Moreno-Sanchez, P., Maffei, M.: A quantitative analysis of security, anonymity and scalability for the lightning network. In: IEEE European Symposium on Security and Privacy Workshops, EuroS&P Workshops 2020, Genoa, Italy, September 7-11, 2020. pp. 387–396. IEEE (2020)
30. Tochner, S., Schmid, S., Zohar, A.: Hijacking routes in payment channel networks: A predictability tradeoff (2019), `https://arxiv.org/abs/1909.06890`
31. Waugh, F., Holz, R.: An empirical study of availability and reliability properties of the bitcoin lightning network (2020), `https://arxiv.org/abs/2006.14358`
32. Zabka, P., Foerster, K., Schmid, S., Decker, C.: Empirical evaluation of nodes and channels of the lightning network. Pervasive Mob. Comput. **83**, 101584 (2022)