Towards Lightning Network Channel Randomization

Kanishkar K* Guillem García† Martí Llinés†
Julián Salas† Guillermo Navarro-Arribas†

April 10, 2026

Preprint notice: This is a preprint of the paper: Kanishkar, K., García, G., Llinés, M., Salas, J., Navarro-Arribas, G. (2026). Towards Lightning Network Channel Randomization. In: Modeling Decisions for Artificial Intelligence. MDAI 2025. Lecture Notes in Computer Science, vol 15957. Springer, Cham. https://doi.org/10.1007/978-3-032-00891-6_14

Abstract

Payment channel networks like the Lightning Network (LN) boost cryptocurrency scalability but might have some privacy risks, as establishing a direct channel reveals user or entity relationships. This paper introduces a novel method using edge local differential privacy (LDP) applied during public channel creation to protect channel existence privacy. Our approach adds statistical uncertainty to the network graph, providing plausible deniability about direct links between specific users while ensuring channels remain usable for routing payments. We formally introduce edge LDP for this context, define relevant utility metrics focused on multi-hop payment feasibility, cost, and node centrality, and evaluate the privacy versus utility trade-offs using a snapshot of the Bitcoin LN.

1 Introduction

This work focuses on the privacy associated with the existence of a payment channel in payment networks built on top of cryptocurrency systems. The most notable example is the Lightning Network (LN) over Bitcoin. The idea behind the LN and such payment networks is to solve the scalability problems of blockchain-based systems such as Bitcoin. Once two parties establish a payment channel, they can exchange payments without having to relay them to the

*Department of Information and Communications Engineering, Universitat Autònoma de Barcelona; SASTRA Deemed University.

†Department of Information and Communications Engineering, Universitat Autònoma de Barcelona

blockchain. Only the opening and closing (whether cooperatively or not) of the channel are committed to the blockchain, allowing it to scale up the number and speed of payments. An important advantage of payment networks, such as the LN, is that they allow multi-hop payments. That is, node A can route a payment to node C through a third node B. This is usually done when there is not a direct channel between A and B (but there is one between A and C, and another between C and B). For this to work, all nodes advertise their channels publicly through the network.

There are different privacy considerations in payment networks. One of them is concerning the existence of the payment channel itself. Some users might wish to perform payments to other users but keep this fact as private as possible. Establishing a direct channel between two users clearly reveals the payment intention. LN allows setting up so-called private channels. Those are channels that are not publicly advertised, through the network. They are private but cannot be used to route payments. That is, the nodes cannot obtain fees from the use of the channel for routing other payments. Our proposal does allow channels to be publicly advertised and used for routing payments, but their existence can be considered private.

We propose a novel paradigm to establish channels in a payment network such as the LN. The main idea is to make channels private by enforcing local differential privacy regarding the existence of a given channel. So observing the graph composed of the network nodes and their channels, there is enough uncertainty about the intended channel existence. From the user point of view, it means that if user A wants to create a channel with user B to perform payments in the future, they might create such a channel or create a channel with another user C, and have the payment routed through other nodes (such a C). This implies that the channel creation should follow some procedure to ensure local differential privacy while keeping some properties in the overall network and from the user's perspective. Such properties are payment feasibility and cost, overall channel density, or node centrality.

In this paper, we introduce the use of edge local differential privacy to provide privacy to the existence of a channel in a payment network. To consider this approach, we introduce utility metrics based on the payment feasibility and performance, on the payment itself, and on the expected node revenues given its centrality. All these metrics have a special focus on multi-hop payments.

The paper is organized as follows. Section 2 introduces the Lightning Network and the edge local differential privacy method, showing the motivation of our proposal. In Section 3 and Section 4 we show the results of applying the edge differential privacy to a specific snapshot of the Bitcoin Lightning Network. Finally, Section 5 concludes the paper.

2 Preliminaries

We outline here the characteristics of the payment channels and edge local differential privacy. Our work can be applied to generic payment networks, but

we contextualize it in the LN network, given its actual relevance.

2.1 Lightning Network

The Lightning Network was initially proposed for Bitcoin [11] as a solution to the limitations that Bitcoin imposes on transaction throughput, latency, and cost. It is considered as a layer-2 protocol since it operates on top of the first layer, which in this case is Bitcoin.

Broadly speaking, its fundamental component is the payment channel, a construct established between two parties via an initial on-chain transaction. Once a channel is open, the involved parties can conduct a virtually unlimited number of off-chain transactions between themselves nearly instantaneously and with minimal fees. These transactions merely update the balance distribution within the channel without broadcasting each one to the main blockchain. The underlying blockchain is primarily utilized for the initial channel funding, the final settlement when a channel is closed, and dispute resolution, thereby significantly reducing the load on the main chain.

Most notably, the LN extends this concept into a payment network. Payments can be securely and trustlessly routed across multiple intermediary nodes in the network, even if the sender and receiver do not share a direct payment channel. Each node in the route receives a small fee to promote its participation. The payment path is source-routed, so the source node determines the path based on different attributes associated with the path.

Channels are publicly advertised using the LN gossip protocol [4]. The main information published for a payment channel directly used to route a payment is, apart form its existence, its *capacity* and *fees*. Capacity is the overall capacity of the channel, which is distributed among the balance of each node. Such balance is private and not advertised to avoid payment tracking. The fees are the amount that a node takes to route a payment through its channel. Each node advertises a base fee, and a proportional fee (proportional to the payment amount). There are other attributes associated with the path-finding algorithm such as the time lock, channel age, or success probability. In general, most implementations use a variation of the Dijkstra algorithm using a combination of the above mentioned information to determine a cost for each channel (see [6] for more detail on pahtfinding strategies). It is however common in the literature to simplify such payment path finding algorithm as a common Dijkstra algorithm [10].

For our purposes, a payment network is modeled as a graph $G = (V, E)$, where a network node is represented as a vertex $v \in V$, and a channel between two nodes as an edge $e \in G$. We note that sometimes it is convenient to model the payment network as a directed graph since some attributes of a channel will be different depending on its direction in a multi-hop payment. We do not get into such details to keep the exposition simple and focus on the private nature of the channel. Moreover, we will assume a uniform capacity, balance distribution, and fees for each channel.

The Lightning Network has the concept of *private* or *unannounced* channels. Those are channels that are not publicly announced using the gossip protocol.

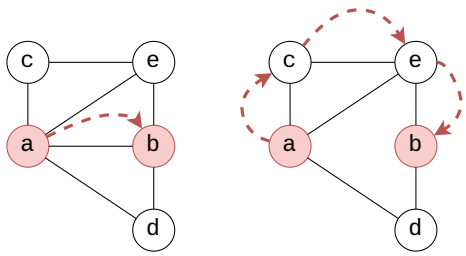

Figure 1: Example of multi-hop payment with and without a direct channel.

They can be effectively considered private, but this also means that they cannot be used by a source node to route a payment so easily. If all channels in the network were private, performing multi-hop payments as it is done now could become unfeasible. Our proposal aims to provide another way to ensure some degree of privacy in channels without having to rely on hiding them.

The existence of a payment path between nodes a and b in the left part of Figure 1 can reveal, in some sense, a commercial agreement between them (subscription to a service, etc.). The idea is to exploit the possibility of doing multi-hop payments to provide privacy to this agreement. As shown in the right part of the figure, a and b could conduct the same payment without having to create a channel between them. The cost of the payment might increase in terms of time and fees, so in some sense, privacy might come with a price. We believe that in the general case, an edge local differential privacy approach could be applied with a relatively acceptable cost. Nodes will create channels based on randomization of their preferences, so the existence of a channel has some uncertainty associated with the actual intention of the nodes to conduct payments between them. Ideally, the channel creation cost should be the same for each node, and payment costs should not drastically increase.

Other works attempt to anonymize the balance of a channel [3] even in the presence of balance discovery attacks [5], but the existence of the channels is not considered private in such works.

2.2 Edge Local Differential Privacy

To provide formal guarantees of privacy for the existence of channels (more precisely to the intention of creating one of them) we base our solution in Edge Local Differential Privacy. In this section we give its formal definition, which is based on providing Local Differential Privacy (LDP) to the edges of a network.

LDP is a method in which users randomize their values without the need of a trusted data curator (or central authority) to protect their privacy with formal guarantees of ε-Differential Privacy. LDP is defined as follows in [2]:

Definition 1 (Local differential privacy)**.** *A randomized algorithm $\mathcal{A}$ satisfies ε-local differential privacy if for all inputs i, j and all outputs $k \in Range(\mathcal{A})$:*

$$Pr[\mathcal{A}(i) = k] \leq e^{\varepsilon} Pr[\mathcal{A}(j) = k], \tag{1}$$

we say that $\mathcal{A}$ is ε-locally differentially private (ε-LDP).

To obtain Edge Local Differential Privacy from LDP, we just need to consider that the inputs and outputs i, j and k are the value of a given edge in the adjacency matrix of the graph (1 if it is present, 0 if not), and the algorithm $\mathcal{A}$ randomizes such value with given probabilities p_{01} and p_{10} that represent the probability of adding an edge, and the probability of removing it.

It was shown that for specific values of p_{01} and p_{10} the algorithm $\mathcal{A}$ obtained may preserve the sparseness of the original graph and provide ε-Edge Local Differential Privacy [9].

The explicit probabilities for p_{01} and p_{10} that we obtain by some simple calculations from [9] are:

$$(p_{01}, p_{10}) = \left(\frac{1}{e^{\varepsilon} - 1 + \frac{1}{d_G}}, 1 - \frac{e^{\varepsilon}}{e^{\varepsilon} - 1 + \frac{1}{d_G}} \right). \tag{2}$$

Thus, we will use such probabilities to randomize the channels in the Lightning Network.

3 Lightning Network Randomization

The goal of our proposal is to obtain a randomized payment network, modeled as a graph G', where the edges (channels) fulfill edge local differential privacy. That is, given a payment network G, where each user establishes the channel based on their preferences (user A will set a channel with user B in order to conduct direct payments in the future), we can obtain a randomized version G', where the existence of a channel does not directly reveal the the initial intention of conducting payments between the involved nodes.

To build such a graph there are different approaches, which we outline here. Note that the main goal of our proposal is to study the viability of the randomization, and describing specific methods to achieve the randomized graph in practice could be considered in future installments.

- Centralized randomization. The randomization process is computed by a trusted entity based on the desired preferences of each node. The graph can be constructed from scratch by the centralized entity even dynamically [8]. The obvious drawback is having to rely on such a trusted entity.

- Distributed randomization. In this case, we envision a scenario, where nodes agree to perform such randomization. Nodes keep their preferences private and only attempt to establish randomized channels. This randomization will require collaboration between such nodes. Moreover, this could be applied to a subset of nodes, which agree to conduct a randomized protocol between them. In such case, only channels between those nodes will provide differential privacy guarantees.

The goal of the randomization is to have a network G' with some specific characteristics with respect to the original intended network G:

- Payment feasibility: if a payment between two nodes in G is possible, it should also be in G'.
- Payment average cost: the cost of conducting a given payment in G should be similar to the same payment in G'. This will be mainly determined by the fees of the payment.
- Node centrality: a node with high centrality in G should also have high centrality in G'. This has a direct impact on the revenue that such a node can obtain by routing payments, which should be the same regardless of the randomization.

In the following section, we introduce some metrics to analyze these characteristics.

4 Evaluation

To evaluate the viability of our proposal we have conducted different experiments. We have taken a snapshot of the Bitcoin Lightning Network on the 27th of March, 2025, from an LND [6] node. The whole graph has 11254 nodes and 21474 channels. We have applied the edge local differential privacy method explained in Section 2.2 with different initial ε values. The original network is denoted as G_0, and the randomized versions as G_ε. We used a very low $\varepsilon = 0.1$ just to show an extreme case, and then used values $\varepsilon = 1, \ldots, 10$ for more general cases. Table 1 shows this notation and the number of edges and connected components for each network.

Graph	Epsilon	Edges	Components
G_0	–	21474	528
$G_{0.1}$	0.1	21458	230
G_1	1	21305	274
G_2	2	21333	249
G_3	3	21254	259
G_4	4	21338	256
G_5	5	21556	268
G_6	6	21485	341
G_7	7	21600	388
G_8	8	21477	609
G_9	9	21572	714
G_{10}	10	21500	672

Table 1: Details for the original and randomized networks.

As expected, the density of the randomized networks remains similar to the original one. The number of components decreases as the ε decreases. This is also expected, since when the networks tend to a random graph the connection of the graph increases [1].

4.1 Failed payments

We evaluate how a payment will be conducted as compared to the original network. The first question is whether the payment will be possible in the randomized networks. To that end, we replicate a random payment from the original network in the randomized networks (we repeat the experiment 1000 times for each case). A payment might fail because nodes that are connected through some given path in G_0 might not be connected in the other networks. Table 2 shows the average percentage of failed payments in each randomized network. This percentage decreases with the parameter ε because, as previously mentioned, introducing more randomness makes the network more connected. In any case, the percentage is commonly under approximately the 10%. Note that we are only considering payments that do not fail in the original network G_0.

$G_{0.1}$	G_1	G_2	G_3	G_4	G_5	G_6	G_7	G_8	G_9	G_{10}
4.02	4.48	4.97	4.73	4.96	5.27	6.15	6.57	10.82	11.19	9.0

Table 2: Average percentage of failed transaction for the randomized networks.

4.2 Payment cost

We also consider the cost of making a payment, in case such payment is feasible. We select two random nodes from the original network and attempt a multi-hop payment, then we replicate the same payment on the private networks (we repeat the experiment 1000 times, with different random nodes). We compare the number of hops taken by the payment in each network. A similar number of hops will mean that the payment can be executed similarly in the private version as compared to the original network, both in terms of time and cost. We have to emphasize that we simplify this measure by considering the cost to be proportional to the payment path length (see Section 2.1). That is, we are assuming that the payment can take place without constraints given by the actual capacity and balance of the channels and that the cost in fees will be proportional to the number of hops. We believe that these assumptions, even if a bit coarse, can help simplify the evaluation while giving a first impression of the feasibility of the proposal.

Figure 2 shows the average length of the same payment (that is, same source and destination nodes) in the original network and the private versions. Each line corresponds to a given length (number of hops) in the original network. We

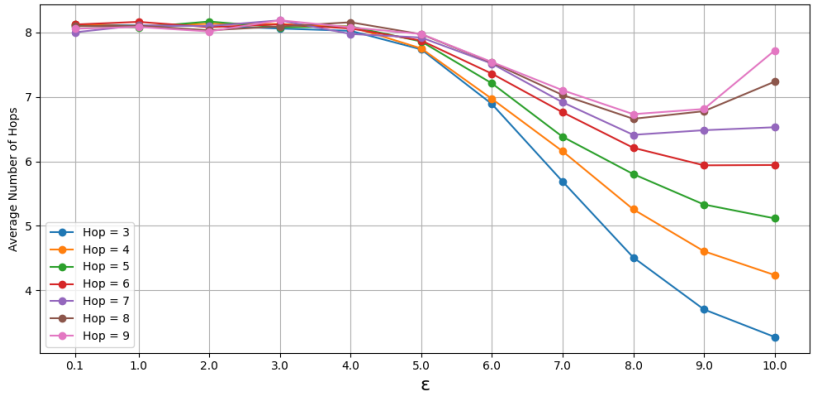

Figure 2: Average path length in randomized networks compared to the original network.

can see that increasing randomization tends to a path length of approximately 8, which seems to be the average path length in the random graph.

4.3 Payment path similarity

We check to what extent, the paths are the same for payments in the original network compared to the randomized versions. That is, in some cases, the multi-hop payment will take the same channels. This means that all the nodes in the payment path are connected in such a way that it replicates the original network topology. This is interesting data giving an idea of the effects of the randomization, and the practical meaning of the differential privacy guarantees.

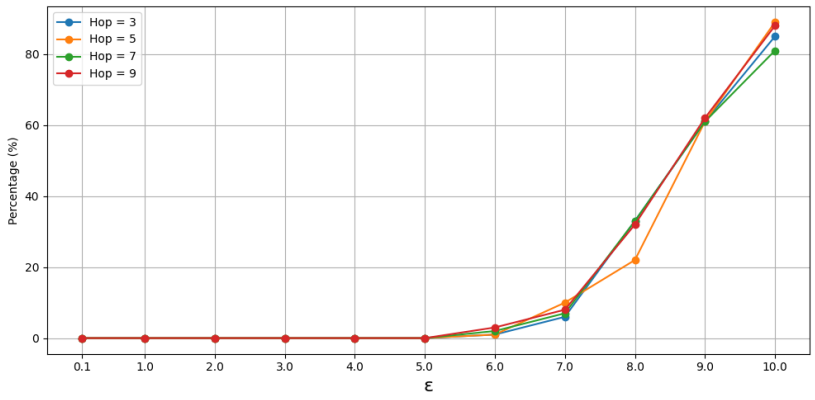

Figure 3: Average percentage of equal path in the original and randomized networks.

Figure 3 shows the average percentage of times that the path in the randomized network is equal to the path in the original network. We see that for $\varepsilon = 5$ and below there are no equal paths.

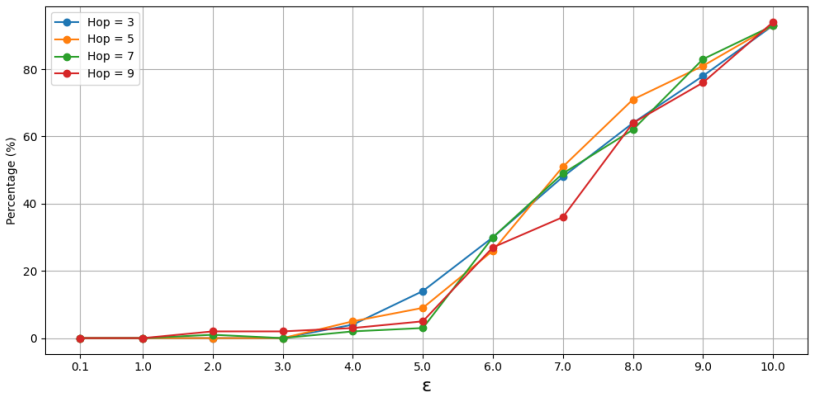

Figure 4: Percentage of paths sharing one node.

We also compute path correlation by finding if at least one node in the randomized network is in the path from the original network. This can be seen as the probability that the path in both the original and the randomized graph have a common node. This is shown in Figure 4.

4.4 Node centrality

We also consider how the centrality varies between the original and the randomized networks. A node with a lot of channels in the network is a node that will potentially gain more revenue by routing payments for other nodes. In this sense, centrality measures give an idea of this potential revenue. A good randomized version of the network will preserve such centrality. Even if the node does not have the same channels (edges) its centrality should be similar to the original.

For this, we are using degree (DC), betweeness (BC), and closeness (CC) centrality measures. We compare the top 1% of the nodes having the highest centrality in the original network to the top 1% nodes in the randomized networks. The comparison is done using the Jaccard index. This is shown in Table 3.

Although the results do not seem as good as expected, note that we are comparing the 1% nodes from each network, which does not give an idea of the variation of the centrality in each case.

	G_{10}	G_9	G_8	G_7	G_6	G_5	G_4	G_3	G_2	G_1	$G_{0.1}$
DC	0.948	0.948	0.837	0.687	0.468	0.153	0.081	0.037	0.004	0.004	0.018
BC	0.794	0.752	0.548	0.387	0.249	0.097	0.066	0.032	0.004	0.004	0.013
CC	0.674	0.468	0.322	0.171	0.102	0.071	0.066	0.023	0.004	0.023	0.000

Table 3: Jaccard Index between the original and randomized networks for the 1% of nodes for their degree (DC), betweeness (BC), and closeness (CC) centrality measures.

5 Conclusions

In this paper, we have considered the definition of a payment network in terms of edge local differential privacy. This allows to introduce the concept of a private channel in a novel way to payment networks, as well as to provide formal guarantees for such privacy. To that end, we apply edge local differential privacy to a payment network while preserving the overall density of the network. That is the number of payment channels. We have evaluated such randomized payment networks by defining metrics in terms of multi-hop payments, showing promising results.

Acknowledgemnts

This research has been partially supported by the project DANGER C062/23 within the Plan de Recuperacion, Transformacion y Resiliencia funded with Next Generation EU funds. Spanish Ministry under Grant PID2021-125962OB-C33 SECURING/NET. Catalan AGAUR under Grant SGR2021-00643.

References

[1] Bollobás, B. (2001). Random Graphs (2nd ed.). Cambridge: Cambridge University Press.

[2] Cormode, G., Jha, S., Kulkarni, T., Li, N., Srivastava, D. and Wang, T. (2018). Privacy at Scale: Local Differential Privacy in Practice. In Proceedings of the 2018 International Conference on Management of Data (Houston, TX, USA) (SIGMOD '18). Association for Computing Machinery, New York, NY, USA, 1655–1658.

[3] Dam, Gijs van, Kadir, Rabiah Abdul (2022) Hiding payments in lightning network with approximate differentially private payment channels. Computers & Security. Volume 115, Elsevier.

[4] Dryja, T., Osuntokun, O., et al. (2025). Basis of Lightning Technology (BOLT) #7: P2P Node and Channel Discovery. Lightning Network Specifications. Retrieved from `https://github.com/lightning/bolts`

[5] Herrera-Joancomartí, J., Navarro-Arribas, G., Ranchal-Pedrosa, A., Pérez-Solà, C., Garcia-Alfaro, J., (2019). On the Difficulty of Hiding the Balance of Lightning Network Channels. In Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security (Asia CCS '19). Association for Computing Machinery, pp. 602–612.

[6] Kumble, S.P., Roos, S.: Comparative Analysis of Lightning's Routing Clients. In: 2021 IEEE International Conference on Decentralized Applications and Infrastructures (DAPPS). pp. 79–84 (2021). https://doi.org/10.1109/DAPPS52256.2021.00014.

[7] Lightning Labs: Lightning Network Daemon (LND). GitHub repository. https://github.com/lightningnetwork/lnd. Accessed 8 Apr 2025.

[8] Paul, S., Salas, J., Torra, V. (2023) Edge Local Differential Privacy for Dynamic Graphs. In Security and Privacy in Social Networks and Big Data. pp. 224–238 Springer. https://doi.org/10.1007/978-981-99-5177-2_13.

[9] Salas, J., González-Zelaya, V., Torra, V., Megías, D. (2023) Differentially Private Graph Publishing Through Noise-Graph Addition. In Modeling Decisions for Artificial Intelligence. pp. 253–264. Springer. https://doi.org/10.1007/978-3-031-33498-6_18.

[10] Sharma, P.K., Gosain, D., Diaz, C., (2023) On the Anonymity of Peer-To-Peer Network Anonymity Schemes Used by Cryptocurrencies. In Proceedings 2023 Network and Distributed System Security Symposium. Internet Society. https://doi.org/10.14722/ndss.2023.23241.

[11] Poon, J., Dryja, T. (2016). The Bitcoin Lightning Network: Scalable Off-Chain Instant Payments. Draft Paper.