On Scalability Power of Payment Channel Networks

Sajjad Alizadeh, Majid Khabbazian

***Abstract*—Payment channel networks have great potential to scale cryptocurrency payment systems. However, their scalability power is limited as payments occasionally fail in these networks due to various factors. In this work, we study these factors and analyze their imposing limitations. To this end, we propose a model where a payment channel network is viewed as a compression method. In this model, the compression rate is defined as the ratio of the total number of payments entering the network to the total number of transactions that are placed on the blockchain to handle failed payments or (re)open channels.**

We analyze the compression rate and its upper limit, referred to as compression capacity, for various payment models, channel-reopening strategies, and network topologies. For networks with a tree topology, we show that the compression rate is inversely proportional to the average path length traversed by payments. For general networks, we show that if payment rates are even slightly asymmetric and channels are not reopened regularly, a constant fraction of payments will always fail regardless of the number of channels, the topology of the network, the routing algorithm used and the amount of allocated funds in the network. We also examine the impact of routing and channel rebalancing on the network's compression rate. We show that rebalancing and strategic routing can enhance the compression rate in payment channel networks where channels may be reopened, differing from the established literature on credit networks, which suggests these factors do not have an effect.

***Index Terms*—Blockchain, Scalability, Payment Channel Networks, Lightning Network.**

I. Introduction

Most prominent blockchains, including Bitcoin and its variants, can handle merely a handful of transactions per second [12]. In contrast, custodial payment networks like Visa can handle tens of thousands of transactions per second [18].

To alleviate this limitation, payment channel networks[1] were introduced. The pioneering network in this domain, the Lightning Network (LN), was initially proposed in 2016 [18] and subsequently launched in 2018. As of February 2025, LN has a total capacity exceeding $470 million USD and features over 45,000 channels [1].

A channel in LN is essentially an escrow mechanism between two parties (two nodes in the network), enabling them to conduct payments off-chain by adjusting their balance within the channel without immediate blockchain confirmation. Using its collection of channels, LN enables off-chain payments between any two parties in the network. Occasionally, however, a payment fails due to the lack of a feasible route. When this occurs, the payment must inevitably be settled by the slow blockchain itself.

The frequency of payment failures ultimately determines the scalability power of the network: the lower the failure rate, the higher the scalability. To measure scalability, we use a model, where the payment network is viewed as a compression method: the method receives a sequence of n transactions over a period of time and places k transactions on the blockchain; Essentially, the method compresses n transactions into k transactions. In this view, the ratio[2] $\frac{n}{k}$ represents the compression rate of the method and determines the scalability power of the network when applied over a long period.

For example, consider a simple network with a single channel between Alice and Bob. Suppose both Alice and Bob each have a balance of \$10 in the channel, making the total channel capacity—the sum of both balances—\$20. Assume that each payment between Alice and Bob is of value \$1. If the first payment is from Alice to Bob, then the balance of Alice and Bob will change to \$9 and \$11, respectively.

Suppose that in each transaction, there is a 50% probability of Alice being the payer. Then, on average, half of the payments flow from Alice to Bob, while the remaining half traverse in the opposite direction. Despite this split of payments, at some point, a payment is unable to proceed through the channel as one party—the intending payer—has a zero balance in the channel. This is expected to occur after $11^2 = 121$ payments, as shown in the literature [5], [8], [13].

When this happens, the payment faces one of two outcomes: it either fails and must be processed directly by the blockchain, or the channel undergoes a reopening, allowing the payment to proceed through the newly established channel. Regardless of the path taken, a transaction is invariably placed on the blockchain. Should Alice and Bob keep their channel, from time to time a transaction needs to be placed on the blockchain. In the above example, for instance, every 121 payments on average require one blockchain transaction[3]. This effectively yields a compression rate of 121.

In our simple single-channel network, the compression rate depends on the channel capacity; the higher the channel capacity, the higher the compression rate. The compression rate also depends on the payment pattern. For instance, in the single-channel network, if Alice is the payer with probability 51% (instead of 50%), the compression rate reduces drastically. We will show later that this phenomenon is not specific to single-channel networks.

To analyze the network's compression rate, we consider different models for payments and channel reopening. For

Sajjad Alizadeh and Majid Khabbazian are with the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, Canada (e-mail: salizad1@ualberta.ca; mkhabbazian@ualberta.ca).

[1]We will use the term *payment networks*, or simply *networks*, to refer to payment channel networks.

[2]This ratio, or a similar one, has also been adopted as a metric of interest in other works in the literature [5], [12], [13].

[3]This is assuming channel-reopening needs one transaction. In practice, a channel-reopening may need more transactions, say, one to close the channel and one to open it with new balances.

payments, we assume two models commonly used in the literature. In the first, known as the *unbiased model*, each node sends payments at the same rate that it receives them. In the second, referred to as the *biased model*, these rates can differ.

For channel reopening, we consider two strategies: *active* and *dormant*. In the active model, a channel is reopened with equal balances once it becomes saturated (i.e., one of its balances has reached zero). This model represents a scenario where users actively maintain their channels. Conversely, in the dormant model, a channel is never reopened, reflecting a scenario in which users choose not to reopen their channels due to the associated costs.

The compression capacity of networks with certain tree topologies, such as stars and lines, has been well studied under the unbiased payment model and the dormant channel-reopening strategy [8]. In this work, we extend the analysis to active channel-reopening for the entire class of tree topologies. In particular, we show that the compression rate is inversely proportional to the average length of the paths that payments traverse in the network.

In addition, we analyze the compression rate of networks with general topologies under the dormant channel-reopening strategy. We prove that under the biased payment model, the network capacity remains constant regardless of the network topology, the total funds locked in the network, or the employed routing and rebalancing algorithms. This represents a negative result for payment networks.

To circumvent the above negative result, two potential approaches can be considered: 1) Abandon the biased payment model. However, this may not be practical, as even a slight imbalance in payment rates in real-world—a likely scenario—can still lead to the same negative outcome. 2) Adopt an active channel-reopening strategy.

In reality, users may hesitate to reopen channels actively, primarily due to associated blockchain transaction costs. Instead, they might prefer a passive strategy, hoping to receive future payments in the opposite direction to mitigate channel saturation. Our analysis explicitly quantifies the scalability consequences of such a passive approach, demonstrating that even minor payment-rate imbalances will result in a constant compression rate when using the dormant reopening strategy.

Finally, we demonstrate that a well-known result in the credit network literature does not apply to payment channel networks where users actively reopen their channels. The well-established result in the credit network literature (Theorem 1 in [8]) states that channel rebalancing[4] has no impact on the success of a payment if the routing algorithm is *diligent*. A routing algorithm is considered diligent if it can identify a feasible path whenever such paths exist.

Simply put, the above result states that if a payment is successful under one diligent routing algorithm, it will be successful under any other diligent routing algorithm. Any additional efforts, such as rebalancing, are deemed unnecessary: while rebalancing may benefit a particular channel, it imposes an equivalent cost on other channels, rendering the overall process ineffective in improving the compression rate.

In contrast to the above existing result, we show that when users adopt an active channel-reopening strategy instead of a dormant one, different diligent routing algorithms can achieve different compression rates. Moreover, rebalancing can enhance the compression rate of a diligent routing algorithm under an active reopening strategy. This finding justifies the search for improved routing and rebalancing strategies when users adopt an active rebalancing strategy.

In summary, we make the following contributions.

1) We model and define the scalability of payment networks in terms of compression rate and compression capacity, where compression capacity represents the highest achievable compression rate. Using this model, we analyze the scalability of various payment networks under different probabilistic payment models and channel-reopening strategies. This framework can serve as a foundation for future studies on the performance of payment channel networks.
2) The performance of single-channel networks has been extensively studied in the literature (e.g., [5], [13]). Our contribution for this simple network is limited to analyzing the compression rate under the combination of dormant channel-reopening and biased payment models. For completeness, we also provide a unified proof that encompasses this and the existing results in the literature. Furthermore, we prove that the active channel-reopening strategy (defined in Section III) is, in the sense described in Section IV-A, optimal for any payment model.
3) The compression capacity of star and line networks are known under the dormant channel-reopening model [8]. In this work, we find the compression capacity of every tree-topology network under the active channel-reopening model. Our results show that the compression capacity of tree-topology networks is inversely proportional to the average length of the paths that payments traverse in the network.
4) For general networks, we prove that the compression capacity remains constant under the dormant channel-reopening strategy and the biased payment model. This result is independent of the network's size, topology, the amount of funds in the network, and the routing/rebalancing strategies employed. This underscores the importance of regularly reopening channels when payments are (even slightly) biased.
5) We show that, contrary to the existing literature on credit networks, strategic routing and rebalancing can influence the compression rate in payment channel networks where channels are actively reopened.
6) We confirm our analytical results using simulations.

[4]Channel rebalancing is a method aimed at rectifying unbalanced channels by identifying a circular path and strategically reallocating funds to restore balance [17].

II. Related Work

Single channel. We start our analysis by considering a single-channel network. Analysis of this network is straightforward using existing results on random walks. Several existing work [13], [5], [8] have already analyzed the payment success rate for a single channel under various combinations of payment models and channel-reopening strategies. We organize and state these results in the language of compression rate in Proposition 1 as a warm-up before presenting our novel results on multi-channel networks. In addition to organizing the existing results, Proposition 1 analyzes the compression rate under the biased payment model combined with the dormant channel-reopening strategy, a combination that has not been covered in the literature.

Optimal strategies for channel reopening under both biased and unbiased payment models have been previously studied in the literature (e.g., [5]). However, what remains unexplored is the optimal channel-reopening strategy when the payment model is not known a priori. In this work, we demonstrate that the active channel-reopening strategy is optimal, in the sense described in Section IV-A, for any payment model.

Tree Topologies. Dandekar et al. [8] analyzed two specific tree topologies, namely star and line topologies, under the dormant channel-reopening strategy.[5] In addition, under the active channel-reopening strategy, the compression rate for the star topology is known [5]. In this work, we extend this result to the entire class of tree topologies under the active channel-reopening strategy.

Tree topologies have been explored in other works, but primarily for objectives such as cost reduction, which are orthogonal to the focus of this work. For instance, Zeta et al. [3] study the topologies that emerge when players act selfishly in the creation of payment networks. In this context, they examine various topologies, including trees, and identify the parameter space in which these topologies constitute a Nash equilibrium.

General Topologies under the Biased Payment Model. The network is expected to achieve a lower compression rate under the biased payment model compared to the unbiased payment model. In [20], the authors demonstrate (Proposition 1) that network throughput is limited by the unbiased component of payments and can only reach this limit if channel capacities are sufficiently high. Our findings align with this conclusion but present an even more negative perspective: we show that the biased component of payments forces the compression rate to remain constant, regardless of channel capacities in the network.

Impact of Routing and Rebalancing. There are numerous routing and rebalancing algorithms proposed in the literature (e.g. [4], [16]), many essentially aimed at improving the network's compression rate. For instance, Revive [16] is a rebalancing strategy designed to address skewed channels. While simulations indicate performance improvements, these gains primarily stem from the shortcomings of routing algorithms in identifying feasible paths (due to hidden balances). This is because a well-established result by Dandekar et al. [8] demonstrates that rebalancing has no real impact if the routing algorithm is diligent (i.e. is able to find feasible paths).

In this work, we show that the fundamental results of [8] do not extend to non-dormant channel-reopening models, particularly the active channel-reopening strategy. More specifically, we prove that different diligent routing algorithms can achieve different compression rates, and that rebalancing can enhance the compression rate of diligent routing algorithms.

Credit networks. A payment channel network is essentially an embodiment of a credit network, a concept originally introduced independently by De Figueiredo and Barr [9], by Ghosh et al. [10], and by Karlan et al. [15]. Dandekar et al. and Goel et al. [8], [11] modeled and analyzed the steady-state transaction success rates in these networks. Their results can be translated into compression rates. However, the results only apply to the unbiased payment model, and hold only for the dormant channel-reopening strategy. In this work, we extend their results to the entire class of networks with tree topologies under the active channel-reopening strategy, and to the general class of networks under the biased payment model.

Minimum-cost topologies. Given an arbitrary network, a payment flow, and a cost function with certain defined properties, Guasoni et al. [14] show how to reduce the cost of the network by locally restructuring channels, or by adding nodes that serve as a hub for several nodes. The cost of a channel defined in [14] depends on the channel's net flow, with the cost equal to zero if the net flow is zero. In the context of compression rate, however, even channels with a net flow of zero can significantly impact the network's compression rate as these channels, too, become saturated.

In another work, Sali and Zohar [19] provide efficient algorithms for constructing minimal cost-spanning-tree networks. They also show that spanning trees, in general, are efficient. In the context of compression rate, which differs from the measure considered in [19], non-tree topologies can achieve higher rates than tree topologies (e.g., a ring topology compared to a line topology). In addition, we show that, within tree topologies, the compression rate is inversely proportional to the network's average path length.

III. System Model

We model the payment network as a weighted, directed graph $G = (V, E)$, where V denotes the set of nodes and E denotes the set of payment channels. Each physical channel between nodes u and v is represented by two directed edges (u, v) and (v, u). The weight of edge (u, v) corresponds to the balance that node u holds on the channel, representing the maximum transaction amount that u can send to node v. Similarly, the weight of edge (v, u) represents node v's balance on the same channel. We refer to the sum of these two balances as the capacity of the channel. Throughout the paper, unless explicitly stated otherwise, we collectively refer to both directed edges connecting two nodes as a single payment channel.

[5]They also examined certain non-tree topologies under the dormant channel-reopening strategy.

An execution of the payment network is modeled as a tuple (S, P, O, R), where S is the sequence of payments[6], P is the sequence of paths used by each payment, O is the sequence of channels (re)opened, and R is the sequence of rebalancing operations.

Every payment in S is categorized as either failed or successful. For each successful payment $s \in S$, all the channels on the corresponding path(s) $p \in P$ have enough balances to forward the payment toward the destination. After each successful payment s, the balances of all the nodes on the path(s) p are updated accordingly. We call an execution *successful* if every payment in S is successful; that is for every payment in S, the corresponding path(s) in P has enough balance to forward the payment. In our analysis, we disregard the fees paid to nodes for forwarding payments as, in practice, these fees are negligible compared to the payments themselves[7].

Payment model. We adopt a probabilistic payment model represented by a rate matrix Λ, where

$$\forall i, j: \quad \sum_{i,j} \Lambda[i,j] = 1 \quad and \quad \Lambda[i,i] = 0.$$

In this model, at each time step, a pair of nodes (u_i, u_j) are selected with probability $\Lambda[i,j]$, and then node u_i is set to pay a unit payment to node u_j.

In the *unbiased* version of this model, we assume that

$$\sum_i |\sum_j (\Lambda[i,j] - \Lambda[j,i])| = 0,$$

that is each node sends payments at the same rate that it receives them. Conversely, in the *biased* version of the model, we assume the opposite, that is

$$\sum_i |\sum_j (\Lambda[i,j] - \Lambda[j,i])| > 0.$$

Channel reopening. We consider two models that represent different strategies for maintaining channels: the *active model* and the *dormant model*.

In the *active model*, a channel is reopened whenever processing a payment would result in one of its balances becoming zero or negative. Such a payment triggers an immediate reopening of the channel, simultaneously processing the triggering payment on the blockchain and reopening the channel with equal balances allocated to both parties. This model represents a scenario in which users actively maintain their channels. Note that under this model, payment failures cannot occur.

In contrast, in the *dormant model*, a channel is never reopened, even when it becomes saturated. This model represents a scenario in which users do not reopen their channels due to the associated costs and may instead opt for less costly alternatives, such as channel rebalancing.

[6]A *multi-part payment (MPP)* is a single logical payment intentionally split into smaller parts, each potentially routed independently. An MPP is considered one payment in S.

[7]In the Lightning Network, the median fee rate amounts to less than \$0.00004 for every dollar sent [1].

IV. Compression Capacity

In this section, we analyze the compression capacity of various networks using both unbiased and biased probabilistic payment models. We then explore how routing and rebalancing algorithms influence the network's compression capacity. To start, we provide formal definitions, including a definition of compression capacity, and proceed with the analysis of the most basic payment network: a network with a single channel.

Definition 1 (Diligent Routing)**.** *Let $G = (V, E)$ be a directed graph whose edges carry non-negative weights (balances). For nodes $u, v \in V$, a (u,v)-*cut *is a partition $(S, \bar{S})$ of V with $u \in S$ and $v \in \bar{S}$. The weight of this cut is the sum of the weights of all edges that point from S to $\bar{S}$. The* minimum (u,v)-cut *is the cut with the smallest such weight.*

A routing algorithm is diligent *if, for every pair $u, v \in V$ and payment value $\nu > 0$, it always finds a path (or set of paths) that can forward k from u to v whenever the weight of the minimum (u,v)-cut is at least ν.*

Definition 2 (Compression Rate)**.** *Let $\mathcal{E} = (S, P, O, R)$ be a payment network execution, and $S^\dagger \subseteq S$ be the sub-sequence of S consisting of all the failed payments. The compression rate of $\mathcal{E}$, denoted $\mathcal{R}(\mathcal{E})$, is defined as the ratio of the total number of payments in S over the total number of transactions published on the blockchain. Formally,*

$$\mathcal{R}(\mathcal{E}) = \frac{|S|}{\kappa \cdot |O| + \alpha \cdot |S^\dagger|},$$

where κ is the average number of transactions that have to be published on the blockchain to (re)open a single channel, α is the percentage of failed payments on the payment network that end up being placed on the blockchain, and $|S^\dagger|$ is the number of failed payments in S. For instance, if all the failed payments on the payment network are subsequently placed on the blockchain itself, we get $\alpha = 1$. Throughout this paper, we assume that $\alpha = 1$ (or $\alpha \in \theta(1)$, in our asymptotic analysis), and $\kappa = 1$, for simplicity. In practical scenarios, $\alpha \neq 1$ may occur when users opt for an alternative payment method after failure in processing the payment. Similarly, instances where $\kappa \neq 1$ can arise in real-world applications, such as when utilizing channel factories [7].

Remark 1. *Under the* active *channel-reopening model, payments never fail. Therefore, we have $S^\dagger = \varnothing$, and the compression rate simplifies to $\mathcal{R}(\mathcal{E}) = |S|/(\kappa\,|O|)$.*

Remark 2 (Relation to throughput and success-rate metrics)**.** *In Definition 2, we define the* compression rate *R as the number of payments processed divided by the number of on-chain transactions triggered by failures or channel (re)openings. Equivalently, $1/R$ is the* on-chain transactions per payment*; multiplying by the on-chain fee yields the expected per-payment on-chain cost.*

This differs fundamentally from throughput *(payments/s), a time-based rate that can be high even when the network generates many on-chain transactions, and from* success-rate/liquidity *metrics, which measure the probability that a payment is routable but do not account for the on-chain work* *required to sustain that success. For instance, a policy that aggressively* reopens *channels to keep success probability high may indeed deliver a high success rate, yet exhibit a low compression rate because frequent reopenings increase on-chain transactions per payment.*

Definition 3 (Compression Capacity)**.** *Let S be a payment sequence of possibly infinite length. The compression capacity of a network for the payment sequence S is defined as the maximum compression rate achievable by any set of diligent routing, channel-reopening, and rebalancing strategies. Formally, we define $\mathcal{C}(S)$, the compression capacity of the network with respect to S, as*

$$\mathcal{C}(S) = \lim_{i\to\infty} \max_{P,O,R} \mathcal{R}(S{\downarrow}i, P, O, R),$$

where $S{\downarrow}i$ denotes the prefix of S with length i. Similarly, the compression capacity of a network for the payment sequence S and a channel-reopening strategy O is defined as

$$\mathcal{C}(S, O) = \lim_{i\to\infty} \max_{P,R} \mathcal{R}(S{\downarrow}i, P, O, R).$$

When we analyze the network's compression capacity under a probabilistic payment model, we set each element of the payment sequence S randomly according to the given distribution, and we extend the length of S infinitely. Under the probabilistic model, we use the simplified notation $\mathcal{C}(S) = c$ to indicate $Pr(\mathcal{C}(S) = c) = 1$, where the probability is over the choice of S.

A. Warm-up: A single-channel network

As a warm-up exercise, let us consider a network with only one channel between two nodes, say, u and v. By definition, $\Lambda[u, v]$ is then the probability that u pays v.

Let $\mathscr{B}_u$ and $\mathscr{B}_v$ be two integers representing, respectively, the balances of u and v on the channel. Assume that each payment made between u and v is of amount one unit. If only u pays v on this channel (i.e. $\Lambda[u, v] = 1$), then u's balance becomes zero after $\mathscr{B}_u$ payments. Under the active channel-reopening strategy, every time u's balance becomes zero, the channel is reopened *with equal balances $(\mathscr{B}_u+\mathscr{B}_v)/2$ on each side*. Given this strategy, therefore, the network's compression capacity will be $\mathcal{O}(\mathscr{C})$, where $\mathscr{C} = \mathscr{B}_u + \mathscr{B}_v$ denotes the channel's capacity. Under the dormant channel-reopening strategy, on the other hand, the network's compression capacity will be zero as all payments (beyond the first $\mathscr{B}_u$ payments) fail[8].

Let us now consider a more generalized case where $\Lambda[u, v] \neq 1$. The analysis of compression capacity, in this case, follows closely that of the gambler's ruin problem [21], a classical example of random walks.

In the gambler's ruin problem, a gambler starts with n dollars, wins a dollar on each bet with probability p, and loses a dollar on each bet with probability $1-p$. The gambler stops playing if she makes m dollars or goes bankrupt (loses all her money). Proposition 1 explains the connection between the gambler's ruin problem and the network's compression capacity. Except for the case under which the compression capacity becomes $\theta(1)$, the results presented in Proposition 1 have already been shown in the literature. For completeness, here we present a cohesive proof for all these results (using the gambler's ruin problem). We will then proceed to prove that in single-channel networks, the active channel-reopening strategy is asymptotically optimal not only under the biased and unbiased probabilistic models but also for *arbitrary payment sequences*, including those with dependencies between payments.

We first present a lemma, which we will specifically use to prove the fourth case in Proposition 1, as this particular case has not been addressed previously in the literature.

Lemma 1. *Consider a payment network with a single channel of capacity $\mathscr{C}$ between two nodes, say, u and v. Suppose the network works under the dormant channel-reopening model, that is the channel is never reopened. Let S be a sequence of payments, and n_u and n_v denote the number of payments by u and v in S, respectively. Assume that the amount of each payment is one unit.*
Then, the number of payments that will fail in the network is at least $|n_u - n_v| - \mathscr{C}$.

Proof. Let $n'_u \le n_u$ and $n'_v \le n_v$ denote the number of successful payments by u and v, respectively. Since each payment is of amount one unit, we get

$$|n'_u - n'_v| \le \mathscr{C} \tag{1}$$

as otherwise the balance of either u or v becomes negative, which is not possible. Therefore, the number of failed payments is

$$\begin{aligned}
&(n_u + n_v) - (n'_u + n'_v)\\
&= (n_u + n_v) - \left\{\max(n'_u, n'_v) - \min(n'_u, n'_v) + 2 \cdot \min(n'_u, n'_v)\right\}\\
&\ge (n_u + n_v) - \left\{\left|\max(n'_u, n'_v) - \min(n'_u, n'_v)\right| + 2 \cdot \min(n'_u, n'_v)\right\}\\
&= (n_u + n_v) - \left\{|n'_u - n'_v| + 2 \cdot \min(n'_u, n'_v)\right\}\\
&\overset{(1)}{\ge} (n_u + n_v) - \left\{\mathscr{C} + 2 \cdot \min(n'_u, n'_v)\right\}\\
&\ge (n_u + n_v) - \{\mathscr{C} + 2 \cdot \min(n_u, n_v)\}\\
&= \{(n_u + n_v) - 2 \cdot \min(n_u, n_v)\} - \mathscr{C}\\
&= |n_u - n_v| - \mathscr{C}.
\end{aligned}$$

□

Proposition 1 ([5], [8], [13])**.** *Consider a payment network with a single channel of capacity $\mathscr{C}$ between two nodes, say, u and v.*
Then, the network's compression capacities under various payment and channel-reopening models are

- $\mathcal{C}(\texttt{unbiased}, \texttt{active}) = \theta(\mathscr{C}^2)$
- $\mathcal{C}(\texttt{unbiased}, \texttt{dormant}) = \theta(\mathscr{C})$
- $\mathcal{C}(\texttt{biased}, \texttt{active}) = \theta(\mathscr{C})$,
- $\mathcal{C}(\texttt{biased}, \texttt{dormant}) = \theta(1)$

Proof. Let $\mathscr{B}_u$ and $\mathscr{B}_v$, respectively, denote the balances of u and v, where $\mathscr{B}_u + \mathscr{B}_v = \mathscr{C}$. Let X be a random variable equal to the number of payments made on the channel before it becomes saturated. If we set $n = \mathscr{B}_u$ and $m = \mathscr{B}_v$ in the context of the gambler's ruin problem, then the variable X will

[8]The compression rate approaches zero as the number of payments tends to infinity.

be equal to the number of bets made by the gambler before she stops. Using the analysis of the gambler's ruin problem, we then get

$$\mathbb{E}[X] = \mathscr{B}_u \cdot \mathscr{B}_v$$

in the unbiased payment model, and

$$\mathbb{E}[X] = \frac{\mathscr{B}_u}{1-2p_u} - \frac{\mathscr{B}_u + \mathscr{B}_v}{1-2p_u} \left(\frac{\left(\frac{1-p_u}{p_u}\right)^{\mathscr{B}_u} - 1}{\left(\frac{1-p_u}{p_u}\right)^{\mathscr{B}_u+\mathscr{B}_v} - 1} \right) \tag{2}$$

in the biased payment model, where $p_u = \Lambda[u,v]$, $p_v = \Lambda[v,u]$ and it is assumed w.l.g. that $p_u < p_v$.

Under the active channel-reopening model, a channel is reopened once it becomes saturated. Therefore, we have

$$\mathbb{E}[X] = \lim_{N\to\infty} \frac{N}{N_O},$$

almost surely [9], where N and N_O, respectively, denote the total number of payments and the number of channel-reopening. Thus, by definition, the compression rate under this model is

$$\mathcal{R} = \lim_{N\to\infty} \frac{N}{N_O} = \mathbb{E}[X].$$

In the unbiased payment model, the compression rate $\mathcal{R} = \mathbb{E}[X] = \mathscr{B}_u \cdot \mathscr{B}_v$ is maximized when $\mathscr{B}_u = \mathscr{B}_v$. Therefore, in this model, the compression capacity is

$$\mathcal{C}(\texttt{unbiased}, \texttt{active}) \leq \frac{\mathscr{C}^2}{4}. \tag{3}$$

Since this upper bound is attained (by the equal-balance allocation), the compression capacity grows quadratically with the total channel capacity, i.e. $\mathcal{C}(\texttt{unbiased}, \texttt{active}) = \Theta(\mathscr{C}^2)$.

In the biased model, from (2), the second term

$$\frac{\mathscr{B}_u + \mathscr{B}_v}{1-2p_u} \cdot \frac{\left(\frac{1-p_u}{p_u}\right)^{\mathscr{B}_u} - 1}{\left(\frac{1-p_u}{p_u}\right)^{\mathscr{B}_u+\mathscr{B}_v} - 1}$$

is positive and $O(1)$ as $\mathscr{B}_u+\mathscr{B}_v$ grows large, since $\frac{1-p_u}{p_u} > 1$. Thus, we directly obtain $\mathbb{E}[X] = \Theta(\mathscr{B}_u)$. Maximizing this subject to $\mathscr{B}_u + \mathscr{B}_v = \mathscr{C}$ yields $\mathscr{B}_u = \Theta(\mathscr{C})$, and consequently,

$$\mathcal{C}(\texttt{biased}, \texttt{active}) = \Theta(\mathscr{C}).$$

Under the dormant channel-reopening model, a channel is never reopened, even when it becomes saturated. Note that a saturated channel can still carry out a payment if the payment is in the right direction. Additionally, once the channel transitions out of its saturated state, it can handle payments in both directions.

Let Y be a random variable equal to the number of successful payments between two subsequent payment failures. We note that payment failures can happen only if the channel is saturated, i.e., one of the channel balances is equal to zero while the other balance is equal to $\mathscr{C}$. From a saturated state, the expected number of payments the channel can handle before it becomes saturated again can be obtained by setting $n = 1$ and $m = \mathscr{C} + 1$ in the gambler's ruin problem. This yields

$$\mathbb{E}[Y] = \mathscr{C} + 1,$$

in the unbiased payment model. Using the law of large numbers, we then get that the compression capacity in the unbiased payment model is

$$\mathcal{C} = \lim_{N\to\infty} \frac{N}{N_F} = \mathbb{E}[Y] = \mathscr{C} + 1 \in \theta(\mathscr{C}),$$

where N_F denotes the number of failed payments.

Finally, let us consider the biased payment model together with the dormant channel-reopening model. Let N denote the total number of payments, n_u denote the number of payments made by u, and n_v denote the number of payments made by v. Let N_F denote the number of failed payments. By Lemma 1, we have $N_F \geq |n_u - n_v| - \mathscr{C}$. Therefore, the compression capacity is

$$\begin{aligned} \mathcal{C} &= \lim_{N\to\infty} \frac{N}{N_F} \\ &\leq \lim_{N\to\infty} \frac{N}{|n_u - n_v| - \mathscr{C}} \\ &= \lim_{N\to\infty} \frac{N}{N \cdot |p_u - p_v| - \mathscr{C}} \\ &= \frac{1}{|p_u - p_v|} \in \theta(1) \end{aligned}$$

□

Optimal channel-reopening. In the biased payment model, a channel remains unsaturated longer if it is reopened with more balance on one side [13], [5]. Nevertheless, by Proposition 1, the active channel-reopening strategy achieves the highest asymptotic compression rate in the biased model as well as the unbiased payment model. This asymptotic optimality does not extend to all arbitrary payment models, even if all payment values are restricted to one. However, if one is willing to approximately double the channel's capacity (specifically, increasing the capacity from $\mathscr{C}$ to $2\mathscr{C}+2$), the active channel-reopening strategy guarantees the best compression rate compared to any other strategy on the original channel. This holds for arbitrary payment sequences with arbitrary payment values. The following theorem formally states this result.

Theorem 1. *Let $\mathcal{R}_{\mathscr{C}}(\mathcal{E})$ denote the compression rate of a single-channel network with the channel capacity $\mathscr{C}$, for execution $\mathcal{E}$. Then, we get*

$$\begin{aligned} &\forall S, P, O, R: \\ &\mathcal{R}_{\mathscr{C}}(S,P,O,R) \leq \mathcal{R}_{2\cdot\mathscr{C}+2}(S,P,O=\texttt{active},R) \end{aligned} \tag{4}$$

Proof. The theorem claims that the compression rate is not reduced when we enlarge the channel capacity from $\mathscr{C}$ to $2\mathscr{C}+2$ and replace the baseline channel-reopening strategy with the *active* channel-reopening strategy.

Consider two single-channel networks, H_1 and H_2, whose capacities are $\mathscr{C}$ and $2\mathscr{C}+2$, respectively. Network H_2 employs the active channel-reopening strategy. The key insight

[9] By the strong law of large numbers: the intervals between reopenings are i.i.d. with finite mean (as in gambler's ruin), hence their average converges almost surely to $\mathbb{E}[X]$.

behind (4) is that every reopening event in H_2 corresponds to *at least one* unique event in H_1—either a reopening or a payment failure. Hence the total number of reopenings *plus* failures in H_1 is no smaller than the number of reopenings in H_2 (no payment fails in H_2 because of the active strategy). By Definition 2, it follows that the compression rate in H_1 is at most the compression rate in H_2.

Let s_{t_m}, $m \geq 1$, denote the mth payment whose amount is *at least* the payer's remaining balance on the channel in H_2. This payment triggers a blockchain transaction that reopens the channel with balances chosen so that, immediately after processing s_{t_m}, both parties hold equal balances. Consequently, the subsequent payment subsequence $(s_{t_m+1}, \ldots, s_{t_{m+1}})$ accumulates a balance drift of at least $\mathscr{C}+1$, as payment $s_{t_{m+1}}$ will again trigger a channel reopening in H_2.

Since H_1 has capacity $\mathscr{C}$, a drift of $\mathscr{C}+1$ within that interval forces either a payment failure in H_1 *or* a channel reopening in H_1. Consequently, each reopening in H_2 corresponds to a unique reopening or failure in H_1: the $(m+1)$-st reopening of H_2 maps to a unique event in H_1 inside the subsequence $(s_{t_m+1}, \ldots, s_{t_{m+1}})$. □

B. Tree topologies

In this section, we analyze a payment network with multiple channels, focusing specifically on networks with a tree topology. In tree topology, there is a unique path between any pair of nodes, allowing us to examine the network's compression rate independently of routing influences.

Moreover, tree topologies possess other properties that make their analysis interesting. For instance, tree topologies such as stars and lines appear in research on capital-efficient payment networks [2] and in studies on how the network topology evolves when users attempt to minimize their costs by choosing whom to connect with [3].

The compression rate for networks with tree topologies—particularly line and star topologies—has been established in [8] for the dormant channel-reopening model. Additionally, the compression rate for networks with star topologies under the active channel-reopening model is known (Corollary 6.2 in [5]). In Theorem 2, we extend this analysis by deriving the compression rate for arbitrary tree-topology networks under the active channel-reopening model.

Theorem 2 characterizes the compression capacity of networks with tree topology under the active channel-reopening strategy. As a consequence, Corollary 1 establishes that when all channel capacities are identical, the compression capacity scales as $\mathcal{O}\left(\frac{\mathscr{C}^2}{\zeta}\right)$, where ζ denotes the average length of paths traversed by payments. An immediate implication, stated explicitly in Corollary 2, is that the compression capacity of a line network with n nodes is $\mathcal{O}\left(\frac{\mathscr{C}^2}{n}\right)$, while the compression capacity of a complete binary-tree topology is $\mathcal{O}\left(\frac{\mathscr{C}^2}{\log n}\right)$, assuming a uniform payment distribution Λ. Along similar lines, the compression capacity for a star-topology network with n nodes is $\mathcal{O}(\mathscr{C}^2)$, a result previously demonstrated in [5].

Before stating Theorem 2, we introduce two key definitions and clarify their relationship in Lemma 2. For clarity, we use the notation $\langle i|(u,v)|j\rangle$ to indicate that removing the channel $(u,v) \in E$ disconnects nodes i and j in the network. Additionally, we denote the shortest-path distance between two nodes $u, v \in V$ as $\operatorname{dist}(u,v)$.

Definition 4 (Channel Centrality)**.** *The* centrality *of a channel $(u,v) \in E$, denoted by $T_{u,v}$, is defined as the probability that a randomly chosen payment path traverses the channel. Formally:*

$$T_{u,v} = \sum_{\substack{i,j \in V \\ \langle i|(u,v)|j\rangle}} \Lambda[i,j].$$

Definition 5 (Average Payment Path Length)**.** *We define the* average length of paths traversed by payments*, denoted by ζ, as:*

$$\zeta = \sum_{i,j \in V} \operatorname{dist}(i,j) \cdot \Lambda[i,j].$$

Lemma 2. *In a tree topology, we have*

$$\sum_{(u,v) \in E} T_{u,v} = \zeta.$$

Proof.

$$\begin{aligned}
\zeta &= \sum_{i,j \in V} \operatorname{dist}(i,j) \cdot \Lambda[i,j] \\
&= \sum_{i,j \in V} \left(\sum_{\substack{(u,v) \in E \\ \langle i|(u,v)|j\rangle}} 1 \right) \cdot \Lambda[i,j] \\
&= \sum_{i,j \in V} \sum_{\substack{(u,v) \in E \\ \langle i|(u,v)|j\rangle}} \Lambda[i,j] \\
&= \sum_{(u,v) \in E} \sum_{\substack{i,j \in V \\ \langle i|(u,v)|j\rangle}} \Lambda[i,j] \\
&= \sum_{(u,v) \in E} T_{u,v}
\end{aligned}$$

□

Theorem 2. *Consider a network with $n > 2$ nodes. Suppose that the network topology is a tree. Let $\mathscr{C}_{u,v}$ denote the capacity of the channel $(u,v) \in E$.*
Then, under the unbiased probabilistic payment and the active channel-reopening models, the compression capacity of the network is $\mathcal{O}\left(\frac{1}{\sum_{(u,v)\in E} \frac{T_{u,v}}{\mathscr{C}_{u,v}^2}}\right)$.

Proof. Let S be a sequence of payments generated in the unbiased probabilistic payment model. For any sequence S, let $S(u,v)$ denote the sub-sequence of S that consists of all the payments that go through the channel (u,v). Since the topology of the payment network $G = (V,E)$ is a tree, removing a channel (u,v) cuts the network into two subsets: the subset $\mathcal{S}_u$ which contains u, and the subset $\mathcal{S}_v$ which contains v. Any payments between these two subsets must go through the channel (u,v). Conversely, any payment that goes through (u, v) must be between a node in $\mathcal{S}_u$ and a node in $\mathcal{S}_v$. Since the payment model is unbiased, we get

$$\sum_{i \in \mathcal{S}_u, j \in \mathcal{S}_v} \Lambda[i,j] = \sum_{i \in \mathcal{S}_u, j \in \mathcal{S}_v} \Lambda[j,i].$$

Thus, the sequence $S(u, v)$ follows the unbiased probabilistic payment method for the channel (u, v).

Fix an edge (u, v) with capacity $\mathscr{C}_{u,v}$ and consider the subsequence $S(u, v)$ of payments that traverse (u, v). Under the active reopening policy, each reopening resets the channel to fixed balances $\mathscr{B}_u^{(0)} + \mathscr{B}_v^{(0)} = \mathscr{C}_{u,v}$ (equal split in our model). Between reopenings, every traversal of (u, v) changes $\mathscr{B}_u$ by ± 1 with probability $1/2$, thus the balance process is a simple symmetric random walk on $\{0, 1, \ldots, \mathscr{C}_{u,v}\}$ with absorbing states $\{0, \mathscr{C}_{u,v}\}$. The classical gambler's-ruin formula gives the expected number of traversals before absorption as $E[X_{u,v}] = \mathscr{B}_u^{(0)} \mathscr{B}_v^{(0)} \leq \mathscr{C}_{u,v}^2/4$ (equality at the equal split); see the single-channel analysis leading to Eq. (3). Under active reopening, renewal cycles are i.i.d., hence by the strong law of large numbers,

$$\lim_{l \to \infty} \frac{|S\downarrow l(u, v)|}{N_O(S\downarrow l, (u, v))} = E[X_{u,v}] \leq \frac{\mathscr{C}_{u,v}^2}{4}, \tag{5}$$

almost surely, where $N_O(S\downarrow l, (u, v))$ denotes the number of times the channel (u, v) is reopened in the sequence $S\downarrow l$, and $|S\downarrow l(u, v)|$ denotes the number of payments in the sequence $S\downarrow l$ that use the channel (u, v). By the definition, $T_{u,v}$ is the probability that a payment path includes the channel (u, v). Thus, by (5) and the law of larger numbers,

$$\lim_{l \to \infty} \frac{l \cdot T_{u,v}}{N_O(S\downarrow l, (u, v))} \leq \frac{\mathscr{C}_{u,v}^2}{4},$$

with probability one. Thus

$$\frac{4 \cdot T_{u,v}}{\mathscr{C}_{u,v}^2} \leq \lim_{l \to \infty} \frac{N_O(S\downarrow l, (u, v))}{l},$$

and hence

$$\begin{aligned}
\sum_{(u,v) \in E} \frac{4 \cdot T_{u,v}}{\mathscr{C}_{u,v}^2} &\leq \sum_{(u,v) \in E} \lim_{l \to \infty} \frac{N_O(S\downarrow l, (u, v))}{l} \\
&= \lim_{l \to \infty} \sum_{(u,v) \in E} \frac{N_O(S\downarrow l, (u, v))}{l} \\
&= \lim_{l \to \infty} \frac{\sum_{(u,v) \in E} N_O(S\downarrow l, (u, v))}{l}
\end{aligned}$$

with probability one. Therefore

$$\mathcal{C}(S, O = \texttt{active}) \leq \frac{1}{\sum_{(u,v) \in E} \frac{4 \cdot T_{u,v}}{\mathscr{C}_{u,v}^2}}.$$

□

Corollary 1. *Under the conditions of Theorem 2, suppose further that the network is homogeneous, meaning all channels have identical capacity $\mathscr{C}$. Then, the compression capacity of the network is*

$$\mathcal{O}\left(\frac{\mathscr{C}^2}{\zeta}\right),$$

where ζ is the average length of paths traversed by payments.

Proof. By Theorem 2, we have:

$$\begin{aligned}
\mathcal{C}(S, O = \texttt{active}) &\leq \frac{1}{\sum_{(u,v) \in E} \frac{4 \cdot T_{u,v}}{\mathscr{C}_{u,v}^2}} \\
&= \frac{1}{\sum_{(u,v) \in E} \frac{4 \cdot T_{u,v}}{\mathscr{C}^2}} \\
&= \frac{\mathscr{C}^2}{4 \cdot \sum_{(u,v) \in E} T_{u,v}} \\
&= \frac{\mathscr{C}^2}{4 \cdot \zeta},
\end{aligned}$$

where the last equality is by Lemma 2. □

Corollary 2. *Under the conditions of Corollary 1, suppose further that the rate matrix Λ is uniform. Then, the compression capacities for line networks, complete binary-tree networks, and star networks are*

$$\mathcal{O}\left(\frac{\mathscr{C}^2}{n}\right), \quad \mathcal{O}\left(\frac{\mathscr{C}^2}{\log n}\right), \quad \textit{and} \quad \mathcal{O}(\mathscr{C}^2),$$

respectively, where n denotes the number of nodes in the network.

Proof. If Λ is uniform, then ζ is equal to the average path length in the network. The average path lengths in line networks, complete binary-tree networks, and star networks are $\theta(n)$, $\theta(\log n)$, and $\theta(1)$, respectively. The proof then follows directly from Corollary 1. □

C. General topologies

Under the dormant channel-reopening strategy and the unbiased payment model, the compression capacity increases with the number of channels in the network and the capacity of those channels [8]. In contrast, Theorem 3 demonstrates that the compression capacity remains constant when replacing the unbiased payment model with the biased payment model. Notably, this theorem does not assume any specific number or capacity of channels within the network, nor does it depend on the network topology or the routing and rebalancing strategies employed.

Theorem 3 presents a negative outlook for networks where payments are biased and nodes refrain from reopening their channels. In such networks, regardless of the amount of funds locked or the routing and rebalancing algorithm used, the compression rate cannot exceed the constant bound established in Theorem 3. When payments are (even slightly) biased, which is likely to be the case in practice, the only way to overcome this limitation is for nodes to reopen their channels regularly, rather than relying on, for example, better routing or rebalancing strategies or securing large funds in their channels. While active reopening increases the compression rate—and thus reduces aggregate on-chain load and total network fees—it also imposes on-chain costs on individual operators. Some channel owners may therefore rationally refrain from reopening; our results quantify the resulting scalability trade-offs.

On the positive side, should users engage in actively reopening channels, the above narrative can shift. For example, consider a star network with nodes $u_1, u_2, \ldots, u_n$, each connected to a central node v. Suppose that the capacity of each channel between u_i and v is $\mathscr{C}$. By Theorem 3, under the biased payment model, the compression capacity of this network is $\mathcal{O}(1)$ if the nodes never reopen their channels. If the nodes adopt the active channel-reopening model, however, the compression capacity becomes $\Omega(\mathscr{C})$ even under the biased payment model. It is because every time a channel is reopened, it can successfully handle at least $\frac{\mathscr{C}}{2}$ payments.

Lemma 3. *Consider two nodes u and v connected by a channel with capacity $\mathscr{C}$. Node u sends payments to v at rate λ_u, and node v sends payments to u at rate λ_v. Let $\nu^{\max}$ denote the maximum payment value, and let $\overline{\nu}_u$ and $\overline{\nu}_v$ denote the average payment values sent by u and v, respectively. Then, with probability one, the long-run average rate of payment failures, defined as the limit of the number of failures divided by the elapsed time T as $T \to \infty$, is at least*

$$\frac{|\lambda_u \cdot \overline{\nu}_u - \lambda_v \cdot \overline{\nu}_v|}{\nu^{\max}}.$$

Proof. Let ν_u^{tot} and ν_v^{tot} denote the total amount sent by nodes u and v, respectively, over a time interval of length T. Similarly, let ${\nu'_u}^{\text{tot}} \leq \nu_u^{\text{tot}}$ and ${\nu'_v}^{\text{tot}} \leq \nu_v^{\text{tot}}$ denote the total value of successful payments sent. Then we must have

$$|{\nu'_u}^{\text{tot}} - {\nu'_v}^{\text{tot}}| \leq \mathscr{C}, \tag{6}$$

as otherwise the channel balance would become negative.

Each failed payment is at most $\nu^{\max}$. Hence, the long-run average rate of payment failures (number of failures per unit time as $T \to \infty$) is bounded below by:

$$\begin{aligned}
&\lim_{T\to\infty} \frac{(\nu_u^{\text{tot}} + \nu_v^{\text{tot}}) - ({\nu'_u}^{\text{tot}} + {\nu'_v}^{\text{tot}})}{T \cdot \nu^{\max}} \\
&= \lim_{T\to\infty} \frac{(\nu_u^{\text{tot}} + \nu_v^{\text{tot}}) - |{\nu'_u}^{\text{tot}} - {\nu'_v}^{\text{tot}}| - 2\min({\nu'_u}^{\text{tot}}, {\nu'_v}^{\text{tot}})}{T \cdot \nu^{\max}} \\
&\overset{(6)}{\geq} \lim_{T\to\infty} \frac{(\nu_u^{\text{tot}} + \nu_v^{\text{tot}}) - \mathscr{C} - 2\min({\nu'_u}^{\text{tot}}, {\nu'_v}^{\text{tot}})}{T \cdot \nu^{\max}} \\
&\geq \lim_{T\to\infty} \frac{(\nu_u^{\text{tot}} + \nu_v^{\text{tot}}) - \mathscr{C} - 2\min(\nu_u^{\text{tot}}, \nu_v^{\text{tot}})}{T \cdot \nu^{\max}} \\
&= \lim_{T\to\infty} \frac{|\nu_u^{\text{tot}} - \nu_v^{\text{tot}}| - \mathscr{C}}{T \cdot \nu^{\max}} \\
&= \lim_{T\to\infty} \frac{|\nu_u^{\text{tot}}/T - \nu_v^{\text{tot}}/T| - \mathscr{C}/T}{\nu^{\max}} \\
&= \frac{|\lambda_u \overline{\nu}_u - \lambda_v \overline{\nu}_v|}{\nu^{\max}},
\end{aligned}$$

where we used the facts that $\frac{\mathscr{C}}{T} \to 0$ and, with probability one, $\frac{\nu_u^{\text{tot}}}{T} \to \lambda_u \overline{\nu}_u$ and $\frac{\nu_v^{\text{tot}}}{T} \to \lambda_v \overline{\nu}_v$ as $T \to \infty$. □

Theorem 3. *Under the biased payment model and the dormant channel-reopening strategy, the compression capacity of any payment network is bounded by a constant. Specifically, we have*

$$\mathcal{C}(S) \leq \frac{2\,\nu^{\max}}{\sum_i \left|\sum_j \left(\Lambda[i,j] \cdot \overline{\nu}_{i,j} - \Lambda[j,i] \cdot \overline{\nu}_{j,i}\right)\right|},$$

where $\overline{\nu}_{i,j}$ denotes the average payment value from node i to node j, and $\nu^{\max}$ represents the maximum payment value in the network.

Proof. Consider any node u_i in the network. Let us call a payment active (with respect to u_i) if u_i is either the payer or the payee in the payment. From node u_i's perspective, the network can be simplified to a single-channel network with just two nodes: node u_i itself and a virtual node v representing the rest of the network, connected by a virtual channel of capacity $\mathscr{C}_{u_i}$, where $\mathscr{C}_{u_i}$ is the sum of the capacities of the channels owned by u_i.

By Lemma 3, the long-run average rate of active payments failures is at least

$$\frac{|\lambda_u \overline{\nu}_u - \lambda_v \overline{\nu}_v|}{\nu^{\max}} = \frac{\left|\sum_j (\Lambda[i,j] \cdot \overline{\nu}_{i,j} - \Lambda[j,i] \cdot \overline{\nu}_{j,i})\right|}{\nu^{\max}}.$$

Therefore, the long-run total rate of payment failures across the network is at least

$$\frac{1}{2} \sum_i \frac{\left|\sum_j (\Lambda[i,j] \cdot \overline{\nu}_{i,j} - \Lambda[j,i] \cdot \overline{\nu}_{j,i})\right|}{\nu^{\max}}.$$

where the factor $\frac{1}{2}$ corrects for the double counting of a single failure which involves exactly two nodes. The total rate of payments is $\sum_{i,j} \Lambda[i,j] = 1$. Thus by definition, the compression rate is

$$\begin{aligned}
&\frac{1}{\frac{1}{2}\sum_i \left|\sum_j (\Lambda[i,j] \cdot \overline{\nu}_{i,j} - \Lambda[j,i] \cdot \overline{\nu}_{j,i})\right| / \nu^{\max}.} \\
&= \frac{2\nu^{\max}}{\sum_i \left|\sum_j (\Lambda[i,j] \cdot \overline{\nu}_{i,j} - \Lambda[j,i] \cdot \overline{\nu}_{j,i})\right|}.
\end{aligned}$$

□

The following corollaries follow immediately from Theorem 3.

Corollary 3. *Assume the biased payment model and the dormant channel–reopening strategy. If each payment amount is drawn i.i.d. uniformly from the interval $[\nu^{\min}, \nu^{\max}]$ with $0 < \nu^{\min} \leq \nu^{\max}$, then the compression capacity satisfies*

$$\mathcal{C}(S) \;\leq\; \frac{4}{\left(1 + \nu^{\min}/\nu^{\max}\right) \sum_i \left|\sum_j \left(\Lambda[i,j] - \Lambda[j,i]\right)\right|}.$$

Proof. For every ordered pair (i,j), a uniformly drawn payment on $[\nu^{\min}, \nu^{\max}]$ has mean $\overline{\nu}_{i,j} = (\nu^{\min} + \nu^{\max})/2$. Substituting $\overline{\nu}_{i,j}$ and $\nu^{\max}$ into Theorem 3 gives

$$\begin{aligned}
\mathcal{C}(S) &\leq \frac{2\nu^{\max}}{\frac{\nu^{\min} + \nu^{\max}}{2} \sum_i \left|\sum_j (\Lambda[i,j] - \Lambda[j,i])\right|} \\
&= \frac{4}{(1 + \nu^{\min}/\nu^{\max}) \sum_i \left|\sum_j (\Lambda[i,j] - \Lambda[j,i])\right|}.
\end{aligned}$$

□

Corollary 4. *Under the biased payment model and dormant channel-reopening strategy, if all payments have identical value, the compression capacity of the network is bounded by:*

$$\mathcal{C}(S) \leq \frac{2}{\sum_i \left|\sum_j (\Lambda[i,j] - \Lambda[j,i])\right|}.$$

Proof. This is the direct special case of Theorem 3 obtained by setting $\overline{\nu}_{i,j} = \nu^{\max}$ for all ordered pairs (i, j). □

Example 1. *Consider a network with $2m$ nodes u_i, $i \in [1, 2m]$. Suppose the network uses the dormant channel-reopening strategy, and assume*

$$\Lambda[i,j] = \begin{cases} 0.5\alpha & \text{if } i,j \in [1,m]; i \neq j \\ 0.5\alpha & \text{if } i,j \in [m+1,2m]; i \neq j \\ (0.5-\epsilon)\alpha & \text{if } i \in [1,m], j \in [m+1,2m] \\ (0.5+\epsilon)\alpha & \text{if } i \in [m+1,2m], j \in [1,m] \end{cases}$$

where $\Lambda[i,j]$ is the probability that node u_i pays node u_j, $\Lambda[i,i] = 0$, and $\alpha = \frac{1}{m(2m-1)}$ is a scaling factor to assure that $\sum_{i,j} \Lambda[i,i] = 1$. Note that this is a biased payment model when $\epsilon > 0$. Therefore, when $\epsilon > 0$, by Theorem 3, we get

$$\begin{aligned} \mathcal{C}(S) &\leq \frac{2}{\sum_i |\sum_j \Lambda[i,j] - \Lambda[j,i]|} \\ &= \frac{2}{2m^2 \cdot 2\epsilon \cdot \alpha} = \frac{2}{2m^2 \cdot 2\epsilon \cdot \left(\frac{1}{m(2m-1)}\right)} < \frac{1}{\epsilon}. \end{aligned}$$

For instance, if we set $\epsilon = 1\%$, we find that $\mathcal{C}(S) < 100$. Despite achieving a nearly unbiased payment model by setting $\epsilon = 1\%$, the result indicates that the compression rate achieved by any routing and rebalancing algorithms will be less than 100. This holds true for any network, regardless of size, topology, and total funds, including the Lightning Network, which has over ten thousand nodes and total funds exceeding three hundred million dollars. Assuming a maximum throughput of 7 transactions per second (tps) for the Bitcoin blockchain, and assuming that the entire blockchain is dedicated to supporting the Lightning Network, the above limit of $\mathcal{C}(S) < 100$ would translate to a maximum throughput of less than $100 \times 7 = 700$ tps.

Theorem 3 shows that under *biased* payments with *dormant* reopening (i.e., channels are never reopened), the compression rate is bounded by a constant. Avoiding this requires some reopening, but not necessarily a fully *active* policy. We discuss two practical directions:

(i) Hybrid selective/asymmetric reopening + off-chain rebalancing. Rather than reopening everywhere, reopen *only* at persistent source–sink bottlenecks and allocate capacity *asymmetrically* to enable a balancing cycle; then use off-chain cycle rebalancing to propagate the effect. *Example:* in a triangle A–B–C with AB and AC fully on A's side and BC fully on C's side, reopening just AB with funds placed on B's side enables a single off-chain cycle $B \to A \to C \to B$ that balances all three channels, achieving the effect of three reopenings with one.

(ii) On-chain–efficient reopening mechanisms. When supported, techniques such as *splicing* (adjusting capacity/balances within an existing channel) and *channel factories* [6] amortize on-chain work across multiple adjustments. These mechanisms increase the number of payments per on-chain transaction, i.e., R per unit of on-chain activity, improving the cost–benefit of reopening relative to naïve per-channel updates.

D. Impact of routing and rebalancing

It was demonstrated in [8] that if a sequence of transactions is successful with one diligent routing algorithm, it will also be successful with any other diligent routing algorithm, and rebalancing is not required. This finding diminishes the importance of both rebalancing and routing choices, as successful transactions can rely solely on diligence. Our key contribution in this section is to show that while this result holds in the dormant channel-reopening model (commonly used in credit network settings, e.g., those in [8]), it does not generalize to non-dormant models.

The *score vector* of the network is defined as the vector whose i-th entry represents the net balance (total outgoing minus total incoming balance) of node i. The intuition behind the above claim is that, while it is correct that the score vector after a successful payment remains independent of path selection in the dormant channel-reopening strategy (as observed in [8]), it can change depending on the selected paths in non-dormant models. Specifically, in a non-dormant reopening model, different paths may trigger different channel reopenings, thereby altering the network's score vector. For instance, one path may trigger no reopenings, while another path may trigger one, causing the score vector to differ accordingly.

In the following, we substantiate the claim by analytically deriving the compression rates achieved by different diligent routing policies on a concrete example. Because the claim is existential, a single counterexample suffices: it is enough to exhibit one topology and a pair of diligent policies whose compression rates differ once channels are reopened.

To this end, consider the small network shown in Figure 1. As depicted, the channels incident to u_1 and to u_7 each have capacity $4 \cdot \mathscr{C}$, whereas every other channel has capacity $2 \cdot \mathscr{C}$. This instance is minimal yet allows closed-form analysis and is sufficient to witness the separation.

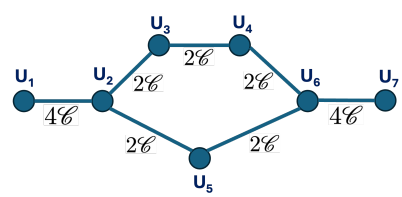

Fig. 1: A network with seven nodes and seven channels. The capacity of each of the channels (u_1, u_2) and (u_6, u_7) is $4\mathscr{C}$; the capacity of each of the remaining channels is $2\mathscr{C}$.

Assume that all the payments are of size one unit, the payments are only between nodes u_1 and u_7—other nodes only forward payments between u_1 and u_7—and that the payments follow the unbiased model.

Consider the following two routing algorithms:

1) **Shortest Path First (SPF) algorithm:** As the name suggests, this algorithm prioritizes the shortest available path to route a payment whenever possible. In the given network, it first attempts to send u_1's payment to u_7 via the path $[u_1, u_2, u_5, u_6, u_7]$. If this attempt fails, it then tries the alternative path $[u_1, u_2, u_3, u_4, u_6, u_7]$.

2) **Longest Path First (LPF) algorithm:** Unlike the Shortest Path First algorithm, this approach prioritizes the longer path $[u_1, u_2, u_3, u_4, u_6, u_7]$ to route the payment first. If this attempt fails, it then tries the shorter path $[u_1, u_2, u_5, u_6, u_7]$.

Note that there are only two distinct paths between u_1 and u_2. As long as at least one of these paths is feasible, the above routing algorithms will successfully transfer the payment; hence, both routing algorithms are diligent. Therefore, by the established result of [8], they achieve the same compression rate in the dormant channel-reopening model.

However, Theorem 4 shows that this result does not extend to the active channel-reopening model. In particular, Theorem 4 shows that the compression rate of the SPF algorithm is 40% higher ($\frac{2}{5} \cdot \mathscr{C}^2$ v.s. $\frac{2}{7} \cdot \mathscr{C}^2$) than that of the LPF algorithm. Moreover, it shows that the compression rate of the SPF algorithm can be improved (by a factor of about two) if the network employs a simple rebalancing algorithm. This contrasts with the result of [8], which proves that rebalancing has no impact on the compression rate in the dormant channel-reopening model.

Theorem 4. *In the network shown in Figure 1, suppose the payments S are unbiased and are only between nodes u_1 and u_7.*
Then we get

- $\lim_{i\to\infty} \mathcal{R}(S{\downarrow}i, \mathit{LPF}, \mathtt{active}, \emptyset) = \frac{2}{7} \cdot \mathscr{C}^2$
- $\lim_{i\to\infty} \mathcal{R}(S{\downarrow}i, \mathit{SPF}, \mathtt{active}, \emptyset) = \frac{2}{5} \cdot \mathscr{C}^2$
- $\lim_{i\to\infty} \mathcal{R}(S{\downarrow}i, \mathit{SPF}, \mathtt{active}, \mathtt{auto\text{-}fill}) = \frac{2\mathscr{C}^2(3\mathscr{C}-2)}{7\mathscr{C}-2}$

where `auto-fill` *is a simple rebalancing strategy that adds one unit of funds (u_2, u_5) and (u_5, u_6) by drawing from (u_2, u_3), (u_3, u_4), and (u_4, u_6) through a circular payment along the cycle $(u_2, u_3, u_4, u_6, u_5)$. It is triggered when (i) (u_2, u_5) or (u_5, u_6) is nearly saturated—i.e., one balance drops to one; and (ii) each of (u_2, u_3), (u_3, u_4), and (u_4, u_6) has at least two units of balance.*

Proof. In the first setting, where we employ the LPF routing algorithm without any rebalancing, each payment takes the long path of

$$P_{long} = [u_1, u_2, u_3, u_4, u_6, u_7].$$

In the active channel-reopening strategy, the channels on P_{long} are reopened with equal balances once they become saturated. When reopened, it takes on average $\left(\frac{4\mathscr{C}}{2}\right)^2 = 4\mathscr{C}^2$ payments for channels (u_1, u_2) and (u_6, u_7) to become saturated again, while it is $\left(\frac{2\mathscr{C}}{2}\right)^2 = \mathscr{C}^2$ payments for the remaining channels on the path P_{long} become saturated. Therefore, the total number of channel-reopenings approaches

$$\lim_{i\to\infty} \mathcal{R}(S{\downarrow}i, \mathtt{LPF}, \mathtt{active}, \emptyset) = \lim_{N\to\infty} \left(\frac{N}{\frac{N}{4\mathscr{C}^2} + \frac{N}{\mathscr{C}^2} + \frac{N}{\mathscr{C}^2} + \frac{N}{\mathscr{C}^2} + \frac{N}{4\mathscr{C}^2}} \right) = \frac{2}{7} \cdot \mathscr{C}^2,$$

as the number of payments N tends to infinity.

Similarly, if we use the SPF routing algorithm with no rebalancing, every payment will go through the path $P_{short} = [u_1, u_2, u_5, u_6, u_7]$ we get

$$\lim_{i\to\infty} \mathcal{R}(S{\downarrow}i, \mathtt{SPF}, \mathtt{active}, \emptyset) = \lim_{N\to\infty} \left(\frac{N}{\frac{N}{4\mathscr{C}^2} + \frac{N}{\mathscr{C}^2} + \frac{N}{\mathscr{C}^2} + \frac{N}{4\mathscr{C}^2}} \right) = \frac{2}{5} \cdot \mathscr{C}^2.$$

Finally, suppose we use the SPF routing algorithm together with the auto-fill rebalancing strategy. In this setting, the channels (u_2, u_3), (u_3, u_4), and (u_4, u_5) never reopen because: (1) no payments between u_1 and u_7 pass through them, and (2) rebalancing occurs only if each of these channels has at least two units of balance, leaving them with at least one unit after rebalancing, thus preventing saturation.

On the other hand, channels (u_2, u_5) and (u_5, u_6) reopen whenever they saturate. At the time of reopening, each long-path channel ((u_2, u_3), (u_3, u_4), and (u_4, u_5)) must have exactly one unit of balance; otherwise, rebalancing would have occurred, avoiding the reopening of short-path channels.

Viewing the cycle $(u_2, u_3, u_4, u_6, u_5)$ as a single logical channel between u_2 and u_6, reopening occurs when one side of this logical channel has balance 1 and the other has balance $4\mathscr{C} - 1$. Immediately after reopening, the balances reset so that the side previously at balance 1 now has $\mathscr{C} + 1$ units, and the opposite side has $3\mathscr{C} - 1$ units. By the gambler's ruin problem, the expected number of payments between reopenings is thus

$$(\mathscr{C} + 1 - 1)(3\mathscr{C} - 1 - 1) = \mathscr{C}(3\mathscr{C} - 2).$$

As before, channels (u_1, u_2) and (u_6, u_7) reopen on average every $4\mathscr{C}^2$ payments. Consequently, we have

$$\lim_{i\to\infty} \mathcal{R}(S{\downarrow}i, \mathtt{SPF}, \mathtt{active}, \mathtt{auto\text{-}fill}) = \lim_{N\to\infty} \left(\frac{N}{\frac{N}{4\mathscr{C}^2} + \frac{2N}{\mathscr{C}(3\mathscr{C}-2)} + \frac{N}{4\mathscr{C}^2}} \right) = \frac{2\mathscr{C}^2(3\mathscr{C} - 2)}{7\mathscr{C} - 2} \approx \frac{6}{7} \cdot \mathscr{C}^2.$$

□

Both Theorems 3 and 4 demonstrate that actively reopening saturated channels significantly improves the compression rate. Specifically, Theorem 3 shows that if payments are even slightly biased and channels are not reopened regularly (e.g., due to associated personal costs), then the network's compression rate remains bounded by a constant, independent of the total locked funds, network topology, and the choice of routing and rebalancing algorithms.

Since practical payment patterns generally exhibit bias, reopening channels is crucial for achieving higher compression rates and reducing the total number of blockchain transactions. However, channel owners may still choose not to reopen channels due to individual cost considerations. Our results quantify this trade-off by showing how the compression rate—and, consequently, the overall blockchain load—changes when channels are not reopened, thereby providing a clear measure of the scalability impact of such a decision.

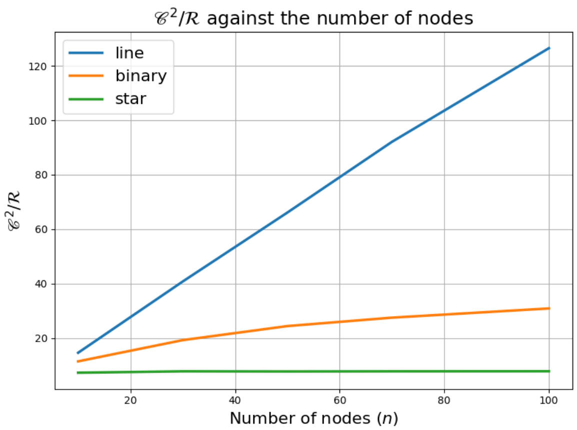

Fig. 2: Validation of Corollary 2: The ratio $\mathscr{C}^2/\mathcal{R}$ as a function of network size (n) for various tree topologies.

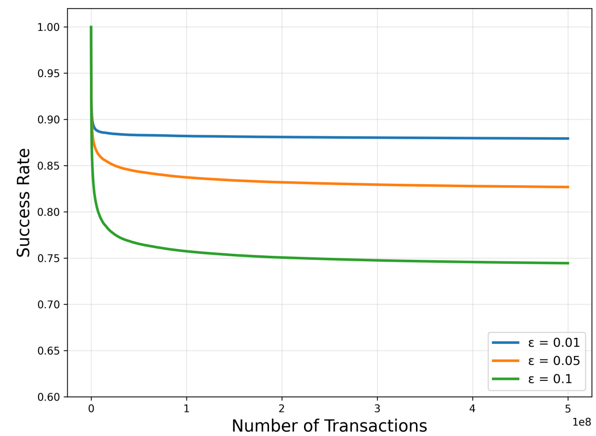

Fig. 3: Payment success rate over time in LN, with each channel's capacity set to \$10.

Theorem 4 shows that when saturated channels are reopened, routing and rebalancing can improve the compression rate, motivating the development of more effective algorithms. This contrasts with the dormant channel-reopening model, where a diligent routing algorithm (i.e., one that finds a feasible path whenever one exists) is sufficient, and additional efforts such as rebalancing have no effect on the compression rate.

V. Simulation Results

In this section, we present simulation results that validate our analytical findings.[10]

Tree Topology. In payment networks with tree-topology, we proved that the compression capacity is $\mathcal{O}(\frac{\mathscr{C}^2}{\zeta})$ when we use the unbiased payment model together with the active channel-reopening model. By Corollary 2, this bounds the compression capacities of line networks, complete binary-tree networks, and star networks to $\mathcal{O}\left(\frac{\mathscr{C}^2}{n}\right)$, $\mathcal{O}\left(\frac{\mathscr{C}^2}{\log n}\right)$, and $\mathcal{O}(\mathscr{C}^2)$, respectively, where n denotes the number of nodes in the network.

To validate these results, we computed the compression rate $\mathcal{R}$ through simulations and plotted in Figure 2 the ratio $\frac{\mathscr{C}^2}{\mathcal{R}}$ as a function of the number of nodes n for the tree topologies described above. In these simulations, we set $\mathscr{C} = 100$, used unit-value payments, and performed 10M payments for each data point. The figure confirms that this ratio increases roughly linearly for line networks, logarithmically for complete binary-tree networks, and remains constant for star networks, closely matching our analytical predictions.

General Topology. In Section IV-C, we proved a negative implication of the dormant reopening strategy: if payments are even slightly asymmetric, the network's compression capacity remains constant irrespective of network topology or total funds locked. The following simulations validate this analytical result.

For this set of simulations, we utilized a snapshot of the Lightning Network taken on May 20, 2024. This snapshot included 16,172 public nodes and 55,705 channels. To determine the capacity of each channel in USD, we used the conversion rate from satoshis to USD on the day of the snapshot, which was 1 USD = 1,435 satoshis. We then rounded the converted values to the nearest integer.

In our simulations, we initially divided each channel's capacity equally between its two nodes, employed a dormant channel-reopening strategy, and used a shortest-path routing algorithm (selecting a shortest path with sufficient channel balances). To avoid infeasible payments between disconnected nodes, we restricted the simulations to the largest connected component of the network.

As in Example 1, we randomly partitioned the nodes into two groups. In every iteration of the simulation, we conducted a single payment. For this, we selected two distinct nodes in the network uniformly at random. If the two nodes belonged to the same partition, we chose the source between the two nodes uniformly at random; otherwise, we chose the node in the first partition as the source with probability $0.5+\epsilon$. Next, we sent a payment of \$1 from the source to the destination. We assumed a small payment amount and that the source was aware of the balances of nodes in the network to aid in payment success. We adopted this approach to demonstrate that even with all this help, the payment success rate will become constant over time due to the combination of payment asymmetry and the dormant channel-reopening strategy.

In the first set of simulations, we reduced the capacity of all channels in the network to only \$10, and we simulated 500M payments. Figure 3 shows the payment success rate for three different values of ϵ, where each point on the curves is the payment success rate till that point. As shown in the figure, the payment success rate over time approaches a constant. For example, when $\epsilon = 0.1$, the steady-state success rate approaches approximately 0.75. Consequently, the compression rate is calculated as $\frac{1}{1-0.75} = 4$, aligning with our analytical result from Example 1, which predicts a compression rate less than $\frac{1}{\epsilon} = 10$ for this case. Assuming the Bitcoin blockchain handles 7 transactions per second (tps), and the entire blockchain is dedicated to LN, the LN throughput would be approximately $4 \cdot 7 = 28$ tps. This throughput remains

[10] All code and data required to reproduce the experiments are available at https://github.com/SecureX-UofA/compression.

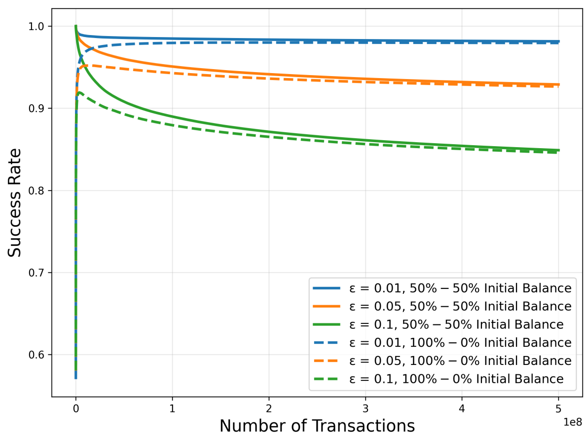

Fig. 4: Payment success rate over time in LN, using the original capacities of each channel.

significantly lower than what custodial payment networks can achieve.

The limited throughput of LN shown above is not due to the small channel capacities ($10) chosen for the first simulation, as demonstrated in our analysis. To confirm this, we ran further simulations using actual LN channel capacities in two scenarios: in the first, all initial channel balances were set to 50%–50%, while in the second, all initial balances were randomly set to 100%–0% (saturated).

Figure 4 presents the results for three different values of ϵ. Predictably, in the saturated mode the success rate started at lower values, but eventually both scenarios converged to similar levels.. As expected, over a sufficiently long period (500M payments), the success rate declines significantly. For example, with $\epsilon = 0.1$, the steady-state success rate stabilizes around 0.84, resulting in a compression rate of approximately $\frac{1}{1-0.84} = 6.25$—again lower than our theoretical upper bound of 10. This compression rate is particularly unimpressive considering that the total liquidity locked in LN exceeds $200M.

Notably, if we adopt an active channel-reopening strategy in the above setting ($\epsilon = 0.1$), our simulations indicate a significant improvement in the compression rate—to approximately 76. This improvement directly translates into fewer blockchain transactions; specifically, the network would require only one blockchain transaction per 76 payments (compared to one per 6.25 without reopening), substantially reducing the overall blockchain load and total transaction fees.

To further strengthen our simulation study, we conducted two additional sets of experiments. In the first, we investigated the effect of varying the network size by retaining only a fraction of the nodes from the current Lightning Network. Specifically, we constructed subgraphs containing either 10% or 50% of the nodes and ran simulations under each configuration. For this purpose, We selected randomly from nodes and maintained connectivity to avoid impossible payments.

In the second set of simulations, we varied the payment amounts by drawing each payment uniformly at random from the intervals $[1, 20]$ and $[1, 100]$, simulating more realistic usage patterns. As shown in Figure 5, in all cases, the success rate (and consequently the compression rate) eventually stabilizes. Specifically, when $\epsilon = 0.1$, the success rate for random payments between 1 and 20 settles around 0.77, translating to a compression rate of approximately $\frac{1}{1-0.77} = 4.35$. This is lower than both the compression rate observed previously and the theoretical upper bound of 10, further supporting our analytical results. Similarly, when using only 10% of the Lightning Network nodes and setting $\epsilon = 0.1$, the success rate converges to 0.8, which corresponds to a compression rate of $\frac{1}{1-0.8} = 5$, again below the theoretical bound.

Impact of Routing/Rebalancing. To confirm the compression rates derived in Theorem 4, we used the network of Figure 1 and the three settings presented in Section IV-D. In all three settings we employed the active channel-reopening model. Recall that in this model a channel is reopened once it becomes saturated. The first setting uses the SPF routing with no rebalancing, the second uses the LPF routing with no rebalancing and the third employs the auto-fill rebalancing with the SPF routing algorithm.

For this simulation, we varied $\mathscr{C}$ from 20 to 100 in increments of 20. For each setting, we generated 10M payments using the unbiased payment model. Figure 6 compares our simulation results with the analytical results proven in Theorem 4. The figure shows that the two results align closely.

As analytically shown, the SPF routing algorithm outperforms the LPF routing algorithm. This demonstrates the impact of strategic routing on the compression rate. We can also see the impact of rebalancing by comparing the compression rates in the second and third settings: both settings use the same SPF routing algorithm, one uses no rebalancing while the other uses the auto-fill rebalancing strategy.

VI. Conclusion and Future Directions

In this work, we studied the scalability of payment networks by analyzing their compression rates—a parameter that quantifies the number of payments a network can process per transaction recorded on the blockchain. We evaluated the compression capacity of various networks under different payment models and channel-reopening strategies.

For networks with a tree topology, we proved for the first time that the compression rate is inversely proportional to the average length of payment paths. In general networks, we demonstrated that the compression rate remains constant with respect to the total funds locked, network topology, and routing/rebalancing algorithms if payments are even slightly biased and users choose not to reopen their channels (e.g., due to associated costs).

Moreover, contrary to a well-established result in the literature, we showed analytically that routing and rebalancing algorithms can indeed affect the compression rate when users actively reopen their channels. Our results, supported by simulations, advocate for an active channel-reopening strategy in payment channel networks to improve scalability.

Several important directions remain open for future research. First, extending our results to more diverse and realistic network topologies and payment-behavior models would enhance practical relevance. Second, comprehensively

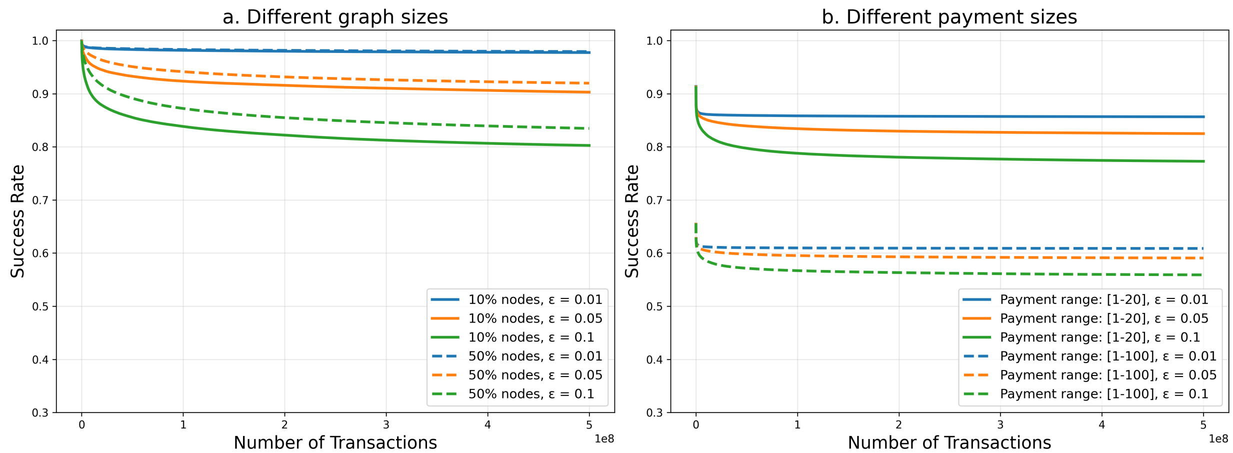

Fig. 5: Payment success rate as a function of the number of transactions under two experimental setups. (a) *Varying network size*: retaining 10% or 50% of current Lightning Network nodes, with fixed channel capacities and constant 1-dollar payments. (b) *Varying payment size*: payments drawn uniformly at random from $[1, 20]$ and $[1, 100]$ units. Each line corresponds to a different ε value.

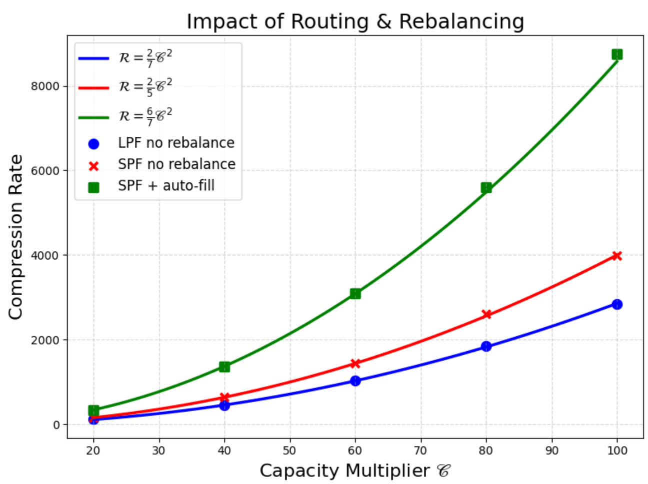

Fig. 6: Simulation results verifying the compression rates established in Theorem 4.

comparing advanced routing and rebalancing algorithms under active reopening scenarios would provide deeper insights into their practical impact on the compression rate. Third, a detailed individual-level cost-benefit analysis to understand channel owners' incentives for reopening, given the personal costs involved, is an important avenue for further investigation. Finally, tightening the theoretical upper bounds derived under conditions of partial or no visibility in routing decisions presents an intriguing challenge for future study.

References

[1] 1ML: Lightning Network Statistics. `https://1ml.com/statistics`, accessed: 2025-02-21

[2] Avarikioti, G., Wang, Y., Wattenhofer, R.: Algorithmic channel design. In: 29th International Symposium on Algorithms and Computation (ISAAC 2018). Schloss-Dagstuhl-Leibniz Zentrum für Informatik (2018)

[3] Avarikioti, Z., Heimbach, L., Wang, Y., Wattenhofer, R.: Ride the lightning: The game theory of payment channels. In: Financial Cryptography and Data Security: 24th International Conference, FC 2020, Kota Kinabalu, Malaysia, February 10–14, 2020 Revised Selected Papers 24. pp. 264–283. Springer (2020)

[4] Avarikioti, Z., Pietrzak, K., Salem, I., Schmid, S., Tiwari, S., Yeo, M.: Hide & seek: Privacy-preserving rebalancing on payment channel networks. In: International Conference on Financial Cryptography and Data Security. pp. 358–373. Springer (2022)

[5] Brânzei, S., Segal-Halevi, E., Zohar, A.: How to charge lightning: The economics of bitcoin transaction channels. In: 58th Annual Allerton Conference on Communication, Control, and Computing, Allerton 2022, Monticello, IL, USA, September 27-30, 2022. pp. 1–8. IEEE (2022)

[6] Burchert, C., Decker, C., Wattenhofer, R.: Scalable funding of bitcoin micropayment channel networks. In: Stabilization, Safety, and Security of Distributed Systems (SSS). Lecture Notes in Computer Science, vol. 10616, pp. 361–377. Springer (2017)

[7] Burchert, C., Decker, C., Wattenhofer, R.: Scalable funding of bitcoin micropayment channel networks. Royal Society open science **5**(8), 180089 (2018)

[8] Dandekar, P., Goel, A., Govindan, R., Post, I.: Liquidity in credit networks: a little trust goes a long way. In: Shoham, Y., Chen, Y., Roughgarden, T. (eds.) Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011. pp. 147–156. ACM (2011)

[9] DeFigueiredo, D.B., Barr, E.T.: Trustdavis: A non-exploitable online reputation system. In: 7th IEEE International Conference on E-Commerce Technology (CEC 2005), 19-22 July 2005, München, Germany. pp. 274–283. IEEE Computer Society (2005)

[10] Ghosh, A., Mahdian, M., Reeves, D.M., Pennock, D.M., Fugger, R.: Mechanism design on trust networks. In: Deng, X., Graham, F.C. (eds.) Internet and Network Economics, Third International Workshop, WINE 2007, San Diego, CA, USA, December 12-14, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4858, pp. 257–268. Springer (2007)

[11] Goel, A., Khanna, S., Raghvendra, S., Zhang, H.: Connectivity in random forests and credit networks. In: Indyk, P. (ed.) Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015. pp. 2037–2048. SIAM (2015)

[12] Goel, A., Ramseyer, G.: Continuous credit networks and layer 2 blockchains: Monotonicity and sampling. In: Proceedings of the 21st ACM Conference on Economics and Computation. pp. 613–635 (2020)

[13] Guasoni, P., Huberman, G., Shikhelman, C.: Lightning network economics: Channels. Michael J. Brennan Irish Finance Working Paper Series Research Paper (21-7) (2021)

[14] Guasoni, P., Huberman, G., Shikhelman, C.: Lightning network economics: Topology. Management Science (2024)

[15] Karlan, D., Mobius, M., Rosenblat, T., Szeidl, A.: Trust and social collateral. The Quarterly Journal of Economics **124**(3), 1307–1361 (2009)

[16] Khalil, R., Gervais, A.: Revive: Rebalancing off-blockchain payment networks. In: Proceedings of the 2017 acm sigsac conference on computer and communications security. pp. 439–453 (2017)
[17] Pickhardt, R., Nowostawski, M.: Imbalance measure and proactive channel rebalancing algorithm for the lightning network. In: 2020 IEEE International Conference on Blockchain and Cryptocurrency (ICBC). pp. 1–5. IEEE (2020)
[18] Poon, J., Dryja, T.: The bitcoin lightning network: Scalable off-chain instant payments (2016)
[19] Sali, Y., Zohar, A.: Optimizing off-chain payment networks in cryptocurrencies. CoRR **abs/2007.09410f** (2020)
[20] Sivaraman, V., Venkatakrishnan, S.B., Ruan, K., Negi, P., Yang, L., Mittal, R., Fanti, G., Alizadeh, M.: High throughput cryptocurrency routing in payment channel networks. In: 17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20). pp. 777–796 (2020)
[21] Song, S., Song, J.: A note on the history of the gambler's ruin problem. Communications for Statistical Applications and Methods **20**(2), 157–168 (2013)