Survivable Payment Channel Networks

Yekaterina Podiatchev, Ariel Orda and Ori Rottenstreich

***Abstract*—Payment channel networks (PCNs) are a leading method to scale the transaction throughput in cryptocurrencies. Two participants can use a bidirectional payment channel for making multiple mutual payments without committing them to the blockchain. Opening a payment channel is a slow operation that involves an on-chain transaction locking a certain amount of funds. These aspects limit the number of channels that can be opened or maintained. Users may route payments through a multi-hop path and thus avoid opening and maintaining a channel for each new destination. Unlike regular networks, in PCNs capacity depends on the usage patterns and, moreover, channels may become unidirectional. Since payments often fail due to channel depletion, a protection scheme to overcome failures is of interest. We define the stopping time of a payment channel as the time at which the channel becomes depleted. We analyze the mean stopping time of a channel as well as that of a network with a set of channels and examine the stopping time of channels in particular topologies. We then propose a scheme for optimizing the capacity distribution among the channels in order to increase the minimal stopping time in the network. We conduct experiments and demonstrate the accuracy of our model and the efficiency of the proposed optimization scheme.**

***Index Terms*—Blockchain, Payment Channels, Network Algorithms.**

I. Introduction

A. Background

Blockchain cryptocurrency networks such as Bitcoin and Ethereum replace a trusted third party with a network of mutually mistrusting peers by aiming to achieve a global consensus between all participants [1]. Payments are recorded on a public ledger providing decentralization, transparency and immutability [1]. Yet this comes at the cost of poor scalability, high transaction fees and large computational overhead. Such consensus-based protocols do not scale well because of limited block size and constant block addition rate. While Visa and Mastercard can handle thousands of payments per second, Bitcoin and Ethereum allow rates of just tens of transactions per second [2].

Offchain networks are a leading solution for the scalability problem [2]. Lightning [3] is an available second layer network for Bitcoin. Offchain networks typically share the same concept, namely: the ability to take payment operations outside of the blockchain in a secure way by backing them up by the blockchain.

Payment Channel. A payment channel provides a way for two users to transfer funds to each other back and forth repeatedly without relying on frequent blockchain interactions. In fact, the users only need to interact with the blockchain when opening, closing or settling disagreements on the channel. The security of the payment channel stems from the collateral funds (namely, the capacity of the channel C) that are locked in a joint blockchain account. In case of dishonest behavior of one party, the counterpart can claim the entire locked collateral as compensation for the breach [3]. Thus, by making it economically irrational for either party to breach the agreed-upon terms, we obtain a powerful deterrent against attempting to cheat the system. The process of making a payment on the channel can be visualized as the transfer of coins from one side of the channel to the other. For example, imagine a payment channel between Alice and Bob, where Alice has 3 coins and Bob has 2 coins, as illustrated in Fig. 1a. If Alice transfers a coin to Bob, a coin from Alice's side of the channel is moved to Bob's side, resulting in the balance illustrated in Fig. 1b. If Bob decides to transfer a coin back to Alice, the coin is simply moved back to Alice's side. Note that, while transactions between Alice and Bob occur by transferring coins back and forth within the payment channel, the total capacity of the channel remains constant [3]. This process can continue indefinitely, as long as there are enough coins at the originating side of the channel.

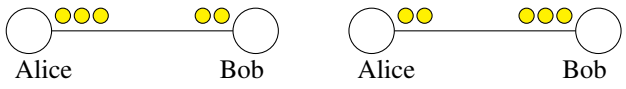

(a) Before the transaction Alice has 3 coins and Bob has 2 coins.

(b) After the transaction Alice has 2 coins and Bob has 3 coins.

Fig. 1: Illustration of the influence of a transaction of one coin from Alice to Bob.

Payment Channel Network. While payment channels are great for frequent transactions between two users, in many cases, transactions to a certain destination are a one-time occurrence, hence opening a channel for every new payment destination defeats the purpose of avoiding on-chain transactions, because the very act of opening a new payment channel requires interaction with the blockchain. Luckily, transactions can be performed on a multi-hop path securely [3]. That way, existing channels can be utilized to facilitate transactions between users that do not share a direct payment channel.

Channel Depletion. A payment channel becomes unidirectional if all the funds accumulate at one of its sides. This may happen even if the expected transaction value in the two directions are equal. In this case, an expensive on-chain transaction may be needed to replenish the funds [4]. A depleted channel may lead to transaction failures if an unaware node tries to use it as one of the intermediate links on a multi-hop path. Thus, balancing out the channels is of great importance for providing appealing quality of service to end users.

Yekaterina Podiatchev, Ariel Orda and Ori Rottenstreich are with the Technion - Israel Institute of Technology, Haifa, Israel. (Emails: yekaterinapo@campus.technion.ac.il, ariel@ee.technion.ac.il, or@technion.ac.il).

B. Contributions

As mentioned in Subsection I-A, payment channels can facilitate transactions only when there are enough funds on the originating side of the channel. When a channel reaches a balance of zero on one of its sides it is said to be depleted. Depleted channels can cause transaction failures and expensive on-chain fees for replenishing the balance on the depleted side. The goal of this work is to increase the success ratio of transactions in Payment channel networks (PCNs) by increasing the survivability level of the network. We refer to survivability in the sense of the first channel failure in the network. We aim to reduce the likelihood of network disconnection by prolonging the time until disconnection occurs.

We distinguish between payment channels according to the distribution of the size of the transactions performed on them. More specifically, we distinguish between balanced channels and channels with drift, and we focus mainly on balanced channels. We define the stopping time of a channel to be the time it takes until the channel becomes depleted. We establish a lower bound on its expected value and attempt to optimize it by redistributing the funds in the network.

The paper is organized as follows. We begin by introducing some useful concepts, definitions and identities that will be used throughout the paper (Subsection II-A). Then, we derive a lower bound on the expected stopping time of a single balanced payment channel with a known transaction variance (Subsection II-B). Next, we calculate and optimize the stopping time for a channel with drift for a fixed-size payment distribution, draw general conclusions about the optimization results, and propose a mapping from a general payment distribution to the distribution we examined (Subsection II-C). Then, we present an approach for calculating the stopping time of a set of channels, assuming uniform capacities and known transition probabilities (Section III). Next, we proceed to optimize the distribution of funds over a set of balanced channels so as to maximize the minimal expected stopping time of a channel in the set using the results of Subsection II-A (Section IV). We evaluate our findings on a few synthetic topologies and on a 2022 snapshot of the Lightning network. We also examine a (more realistic) distributed approach for optimization (Section V). Then, we derive lower bounds for the expected stopping time of a channel in various topologies such as chains and cycles (Section VI). Finally, we present related work on the subject of payment channel networks (Section VII).

II. The Stopping time of a Payment Channel

A. Motivation and Definitions

In this section, we examine under which conditions and how often a payment channel becomes depleted. A payment channel has a constant amount of funds that are distributed between its two sides. When a user wants to perform a transaction over the channel, the desired amount of funds is simply moved from one side to the other. If a transaction exceeds the balance on the originating side then we assume it is performed partially and the remaining amount is canceled. If a payment channel has a balance of zero on one of its sides then it becomes depleted in one direction and no transaction can be performed in that direction. Since a depleted channel may result in transaction failures or necessitate a costly on-chain transaction to replenish it, it is of interest to extend the time until depletion occurred.

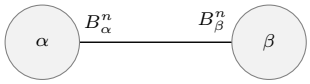

Fig. 2: A Payment Channel connecting node α and β. The balance on the side of each node is written next to it.

Definitions. Consider a payment channel in discrete-time between some nodes α and β, as illustrated in Fig. 2. We divide time into equal duration time steps and accumulate all the transaction requests over that time period. The accuracy of this representation may be increased to the desired extent by choosing sufficiently short time steps. In addition, transaction success rate is defined as the fraction of successful transactions out of the total attempted transactions, that is, it is not affected by the length of the time intervals between each attempt. Denote the balances on nodes α and β at the end of time step n as B^n_α and B^n_β, respectively. Denote the random variable that represents the total required balance change on β's side at time n as X^n; note that, due to the nature of payment channels, the corresponding balance change on α's side is $-X^n$. As was stated before, transactions that go over the balance limit are partially performed:

$$B^n_\beta = min(max(X^n + B^{n-1}_\beta, 0), B^0_\beta + B^0_\alpha)$$

Definition 1 (Stopping Time of a Payment Channel)**.** *It is the first time step n, for which $B^n_\beta = 0$ or $B^n_\beta = B^0_\beta + B^0_\alpha$*

When a payment channel in a PCN becomes depleted, it will become impossible to perform certain transfers, and often there will be a need for a costly on-chain transaction to rebalance the channel. Thus, exploring what affects the stopping time of a payment channel is of interest.

Definition 2 (Stochastic (Random) Process)**.** *It is a collection of random variables indexed by some set $\mathcal{T}$, often interpreted as time.*

Definition 3 (Stochastic (Random) Discrete Process)**.** *It is a random process for which the set $\mathcal{T}$ is a discrete set.*

Definition 4 (Stopping Time)**.** *Let $W = \{W^n\}_{n\geq 1}$ be a stochastic process. A random time, T, is said to be a stopping time with respect to W, if for each $n \geq 1$, the event $\{T = n\}$ is completely determined by (at most) the total information $\{W^1, ..., W^n\}$ known up to time n.*

In this work, $\mathcal{T}$ denotes the index set over which the stochastic process is defined. In contrast, T denotes a stopping time, which is a random variable. In other words, $\mathcal{T}$ is the entire set of possible times, and T is a particular random time within that set.

To calculate the expected stopping time of a given channel, we define a discrete-time random walk S^n for which the stopping time coincides with the stopping time of the channel. $S^n \triangleq B^0_\beta + \sum_{i=0}^{n} X^n$ is a random walk with increments X^n

and initial value $S^0 = B^0_\beta$ which satisfies this condition for a stopping time defined as $T = inf\{n \geq 0 \mid S^n \leq 0 \text{ or } S^n \geq B^0_\alpha + B^0_\beta\}$.

To evaluate the expected stopping time, we use Wald's identities. Notice that the distribution of the stopping time of the random walk that we defined is identical to the distribution of a stopping time of a shifted random walk, $\tilde{S}^n = S^n - S^0$, with correspondingly shifted stopping criteria $-a \triangleq -B^0_\beta$ and $b \triangleq B^0_\alpha$. Wald's identities are defined for a random walk with an initial value 0, thus we will use $\tilde{S}^n$ from now on.

We shall employ Wald's identities, stated as follows. Let $\{X^i\}_{i\in\mathbb{N}}$ be a sequence of independent identically distributed random variables with a common expected value μ and variance σ^2. Let n be a stopping time with respect to $S_k = \sum_{i=1}^k X^i$ for which the expectation is finite. The following properties have been shown in [5] as Theorem 7.1 and a special case of Theorem 7.2. Those properties might seem intuitive, but their proof is nontrivial, since the number of random variables in the sum is itself a random variable, and moreover, it depends on $\{X^i\}_{i\in\mathbb{N}}$. A detailed explanation that includes a counter-example can be found in [6].

Property 1 (Wald's First Identity)**.** *If μ is finite, then $E(S_n) = E(n) \cdot \mu$.*

Property 2 (Wald's Second Identity)**.** *If σ^2 is finite and $\mu = 0$, then $Var(S_n) = \sigma^2 \cdot E(n)$.*

B. Stopping Time of a Balanced Payment Channel

We begin by exploring a balanced channel. A network of balanced channels may be obtained, for example by correctly routing transactions on a circulation network, a concept which was explored in [4], and also occurs naturally in particular topologies by employing symmetry considerations, some of which we explored in Section VI. We shall address the case of imbalanced channels later, in Subsection II-C.

Lemma 1. *The expected stopping time of a channel for which $\{X^n\}_{n\in\mathbb{N}}$ are identically randomly distributed with expectation $\mu = 0$ variance σ^2 satisfies:*

$$E[T] \geq \frac{min(B^0_\alpha, B^0_\beta)^2}{\sigma^2} \tag{1}$$

Proof. Applying Chebyshev's inequality to $\tilde{S}^T$ with a constant K implies:

$$P[|\tilde{S}^T - E[\tilde{S}^T]| \geq K] \leq \frac{Var(\tilde{S}^T)}{K^2}$$

By Property 1:

$$E[\tilde{S}^T] = E[X^i]E[T] = 0 \cdot E[T]$$

Thus:

$$P[|\tilde{S}^T| \geq K] \leq \frac{E[(\tilde{S}^T)^2]}{K^2}.$$

By choosing $K = min(B^0_\alpha, B^0_\beta)$ we obtain $P[|(\tilde{S}^T)^2| \geq K] = 1$ by the definition of $\tilde{S}^T$.

$$1 \leq \frac{E[(\tilde{S}^T)^2]}{min(B^0_\alpha, B^0_\beta)^2}$$
$$E[(\tilde{S}^T)^2] \geq min(B^0_\alpha, B^0_\beta)^2$$

Combining this result with Property 2 implies:

$$min(B^0_\alpha, B^0_\beta)^2 \leq E[(\tilde{S}^T)^2] = \sigma^2 \cdot E[T]$$
$$E[T] \geq \frac{min(B^0_\alpha, B^0_\beta)^2}{\sigma^2} \qquad \square$$

The expected stopping time of any channel in any graph can be bounded from below using Eq. 1 if its expectation is 0 *and its variance is known.* We see that by Eq. 1, the lower bound on the expected stopping time becomes larger when $min(B^0_C, B^0_A)$ increases or σ^2 decreases.

The exact expected stopping time can be calculated for certain distributions of X_i. We provide an example in the following lemma.

Lemma 2. *Consider a payment channel where the payment distribution satisfies $P(X^T = 1) = p = 0.5$ and $P(X^T = -1) = q = 0.5$. If the initial balances on the nodes are a and b, then the expected stopping time $E[T]$ is given by:*

$$E[T] = a \cdot b.$$

Proof. We choose to examine the side of the channel with an initial balance of a. The channel becomes depleted when the balance shifts by $-a$ or by b from the initial balance. The stopping time coincides with the stopping time of S_n where $S_n = S_0 + \sum_{i=1}^n X_i$, where $S_0 = 0$ and the stopping criteria are $S_n = -a$ or $S_n = b$.

We apply Wald's identities to find $E[T]$. By Property 1,

$$-a \cdot P(S_T = a) + b \cdot P(S_T = b) = E[S_T] = 0 \cdot E[T] = 0.$$

By substituting $P(S_T = b) = 1 - P(S_T = a)$ we obtain an expression for $P(S_T = a)$:

$$P(S_T = a) = \frac{b}{a+b},$$

and in a similar way for $P(S_T = b)$:

$$P(S_T = b) = \frac{a}{a+b}.$$

By using Property 2:

$$a^2 \cdot P(S_T = a) + b^2 \cdot P(S_T = b) = Var(S_n) = E[T].$$

Combining those results we obtain:

$$E[T] = a^2 \frac{b}{a+b} + b^2 \cdot \frac{a}{a+b} = a \cdot b.$$

$\square$

C. Stopping Time for Fixed Size Payment Distribution with Drift

We examine two special cases of payment distribution over a single channel, namely the distribution of the transfers in each direction of the channel. In the first case, there is a transaction of one unit in some direction at each time step. In the second case, there may be no transactions at all.

Lemma 3. *The stopping time of a channel for which the payment distribution satisfies $P(X^T = 1) = p \neq 0.5$ and $P(X^T = -1) = q = 1 - p$ and the initial balances on the nodes are a and b, is:*

$$E[T] = \frac{1}{p-q} \cdot \left(-a \cdot \frac{1-(q/p)^b}{(q/p)^{-a} - (q/p)^b} + b \cdot \frac{(q/p)^{-a} - 1}{(q/p)^{-a} - (q/p)^b} \right)$$

The proof can be obtained by recursively expressing the expected stopping time starting from the initial channel balance using the expected stopping time starting from the possible balances at the next step. Note that the lemma does not hold for $p = 1 - q = 0.5$ because the expression $\frac{1}{p-q}$ is not defined for these values. The case where $p = q = 0.5$ falls under the category of a balanced payment channel and is discussed in length in Subsection II-B. The exact expression for $E[T]$ for the specific case discussed here is provided in Lemma 2.

Lemma 4. *Consider a discrete random process $\tilde{W}$ with independent identically distributed values and stopping time $\tilde{T}$ that is completely determined by the sum of values of $\tilde{W}$ up to time $\tilde{T}$ [1]. Let W be a discrete random process defined by:*

$$W^i = \begin{cases} \tilde{W}^i & \textit{with probability } p \\ 0 & \textit{with probability } 1-p \end{cases}$$

with independent sampling for each index i. Let the stopping time T of W be defined as the first time when $\tilde{W}$ would stop for the series of values of W up to that point.
Then, the expected stopping time of W is:

$$E[T] = E[\tilde{T}]/p.$$

Proof. Define the following process:

$$I_i = \begin{cases} 1 & \text{in case } W^i \text{ is sampled from } \tilde{W}^i \\ 0 & \text{otherwise} \end{cases}$$

T, the stopping time of W, is also a valid stopping time with respect to I. In other words, it is valid to define the stopping time of I as the time when the underlying process W stops. By applying Property 1 to $S_n = \sum_{i=1}^{n} I_i$ and T, we obtain that:

$$E[S_T] = E[T] \cdot E[I_1].$$

The sub-series that consists of the values from $\tilde{W}$ that are chosen in W has the same expected stopping time as $\tilde{W}$ since $\{\tilde{W}^i\}$ are i.i.d.. S_T counts the number of steps sampled from this subseries until the stopping time of W. Since by definition this event occurs when $\tilde{W}$ would stop for the subseries of values from $\tilde{W}$ in W, this sum is equal to the stopping time of the aforementioned subseries. Thus, we obtain that:

$$E[S_T] = E[\tilde{T}].$$

In addition, note that T is the stopping time of W and $E[I_1] = p$. Given those assignments of $E[S_T]$ and $E[I_1]$, we deduce that:

$$E[T] = E[\tilde{T}]/p.$$

□

[1] An example of such a stopping time may be the first time the sum of values up to that time reaches a certain threshold.

Lemma 5. *Consider a channel with initial balances C_1, C_2 and a payment distribution satisfying:*

$$X^n = \begin{cases} 1 & \textit{with probability } p_1 \cdot (1-p_2) \\ 0 & \textit{with probability } p_1 \cdot p_2 + (1-p_1) \cdot (1-p_2) \\ -1 & \textit{with probability } p_2 \cdot (1-p_1) \end{cases} \tag{2}$$

Then the stopping time of the channel is:

$$E[T] = \frac{1}{(\tilde{p}-\tilde{q}) \cdot (p_1 \cdot (1-p_2) + p_2 \cdot (1-p_1))} \cdot \left(-C_1 \cdot \frac{1-(\tilde{q}/\tilde{p})^{C_2}}{(\tilde{q}/\tilde{p})^{-C_1} - (\tilde{q}/\tilde{p})^{C_2}} + C_2 \cdot \frac{(\tilde{q}/\tilde{p})^{-C_1} - 1}{(\tilde{q}/\tilde{p})^{-C_1} - (\tilde{q}/\tilde{p})^{C_2}} \right) \tag{3}$$

where:

$$\tilde{p} = p_1 \cdot (1-p_2)/(p_1 \cdot (1-p_2) + p_2 \cdot (1-p_1))$$
$$\tilde{q} = p_2 \cdot (1-p_1)/(p_1 \cdot (1-p_2) + p_2 \cdot (1-p_1))$$

and the optimal capacity distribution is:

$$C_1 = \frac{\ln(\frac{(\tilde{q}/\tilde{p})^C - 1}{C \cdot \ln(\tilde{q}/\tilde{p})})}{\ln(\tilde{q}/\tilde{p})}, C_2 = C - C_1 \tag{4}$$

Recall that the channel capacity refers to the total amount of funds available in the channel, which is distributed between its sides.

Proof. The distribution of X^n is equivalent to sampling from a Bernoulli distribution with $p = \tilde{p}$ with probability $p_1 \cdot (1-p_2) + p_2 \cdot (1-p_1)$ and 0 otherwise. Thus, according to Lemma 4, the stopping time can be obtained by dividing the corresponding stopping time of the Bernoulli distribution, computed using Lemma 3, by $p_1 \cdot (1-p_2) + p_2 \cdot (1-p_1)$. The optimal capacity distribution is obtained by comparing the derivative of the stopping time to 0. □

Capacity Distribution Optimization for a Fixed Size Demand Distribution. In order to come up with a heuristic for distributing the total capacity in a channel with drift, we consider the specific distribution of X^n defined in Lemma 5.

Definition 5 (Capacity Distribution Skew)**.** *It is the distance of a given capacity distribution from the balanced capacity distribution with the same total capacity.*

Definition 6 (Payment Drift Factor)**.** *It indicates the level of drift in the channel payments, i.e., the probability of $\Delta B > 0$ given $\Delta B \neq 0$.*

We use the above definitions to simplify the discussion about the effect of the tendency of transactions to transfer

funds in one direction over the other on the initial distribution of funds within the channel. We focus on the payment drift factor rather than on the probability of a transaction in a certain direction in order to account for payment distributions that include a nonzero probability of no transaction at all. Referring to capacity distribution skew instead of simply the balance is more expressive because it directly shows how far the capacity distribution is from the balanced state in which the capacity is split equally between the two sides of the channel.

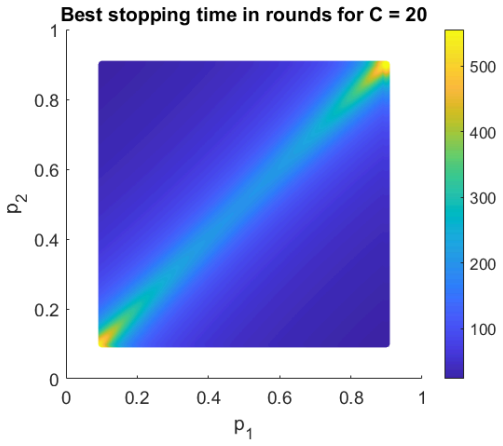

Fig. 3: Optimal stopping time as a function of transfer probabilities p_1 and p_2 as calculated using Eq. 3

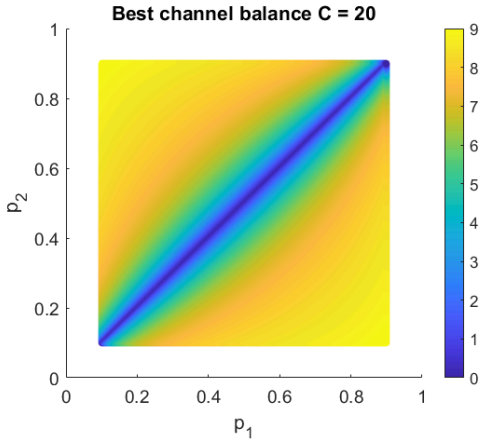

Fig. 4: Optimal capacity distribution skew for stopping time as a function of transfer probabilities p_1 and p_2 as calculated using Eq. 4

For example, in the examined private case, the Payment Drift Factor is given by $\frac{p_1 \cdot (1-p_2)}{p_1 \cdot (1-p_2) + p_2 \cdot (1-p_1)}$. Here, p_1 is the probability of a transaction in the direction that increases ΔB and p_2 refers to the other direction.

We present the optimal capacity distribution calculated using Eq. 4 in Fig. 4, and the corresponding stopping times calculated using Eq. 3 in Fig. 3.

Fig. 3 is a heat map in which the x-axis, namely p_1, is the probability of a transaction in one direction, and the y-axis, namely p_2, is the respective probability in the other direction. The color depicts the optimal stopping time for a given total capacity C. The symmetry around $y = x$ is due to the arbitrary choice of labeling p_1 and p_2, which does not affect the stopping time. The symmetry around $y = 1 - x$ is less obvious, yet it becomes clearer by describing the transferred amount in each direction as a constant transfer of one unit and a negative transfer of one unit with the complementary probability. The constant transfer of one unit in each direction is canceled out and the results are complementary probabilities in the opposite directions. One can verify from Fig. 5 that the total effect on the channel is the same.

Fig. 4 represents the absolute value of the shift of the capacity of each side of the channel from the balanced capacity distribution for which the best stopping time is achieved.

Fig. 6 presents the optimal normalized capacity distribution skew as a function of the payment drift factor for different values of total capacity C. As can be seen, the steepness of the curve increases with C. We thus conclude that *as the total capacity of the channel increases, so does the sensitivity of the optimal capacity distribution to changes in the payment drift factor.* In other words, larger capacity channels benefit more from the proposed optimization of their capacity distribution skew.

Capacity Distribution Optimization for General Distributions. Since it is overly complex to obtain an analytical solution for every payment distribution and it is not feasible to obtain the entire payment distribution from the data, we propose to use the results presented in Subsection II-C by mapping the actual payment distribution on the distribution referred to in Lemma 5.

First, we empirically calculate the probability of transfer in any direction, namely p. Then, we filter out the steps in which no transfer happened and present a mapping from the first and second moments of the remaining transfers, namely $m_1 \triangleq E[X]$ and $m_2 \triangleq E[X^2]$, to a normalization factor, namely x, and a transfer probability in a certain direction, namely q.

To that end, we express the moments using q and x:

$$m_1 = x \cdot (2 \cdot q - 1) \tag{5}$$

$$m_2 = x^2 \tag{6}$$

From Eq. 6, we obtain $x = \sqrt{m_2}$. By employing this in Eq. 5 we obtain $q = 0.5 \cdot (\frac{m_1}{\sqrt{m_2}} + 1)$. Note that $q \leq 1$ because $m_1^2 \leq m_2$ since $m_2 - m_1^2 = \sigma^2 \geq 0$.

III. Expected Minimal Stopping Time of a Set of Channels

Till now we considered the stopping time of a single channel. In this section, we consider the expected time of the first channel failure in a set of channels. In a tree topology, this event coincides with the topology becoming disconnected, i.e., the event of having at least one pair of nodes with no connecting path. In other topologies, a single failure might not affect the connectivity; however, users are not aware of the exact balance of channels that they do not maintain, thus even when a transaction has a path on which it will succeed, the user may unknowingly choose a path that includes a depleted channel, which would lead to failure and in turn affect the success rate. Therefore, it is of interest to postpone the first failure even when it does not render the topology disconnected.

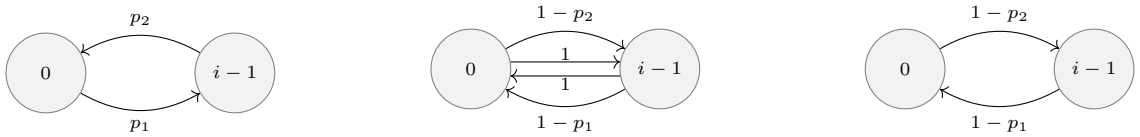

(a) original channel

(b) each direction divided into a deterministic transfer and a transfer with the complimentary probability in the opposite direction

(c) equivalent transfer after deterministic transfers cancel each other

Fig. 5: Illustration of the equivalence between a channel and its complementary channel.

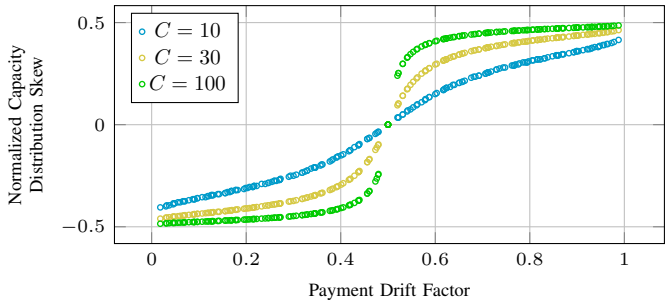

Fig. 6: Optimal capacity distribution skew as a function of payment drift factor for different values of total channel capacity C

Denote the stopping times of the channels $\{T_i\}_1^M$ where i represents the index of the channel. The expected minimal stopping time can be expressed as follows:

$$T = min_i(\{T_i\}_1^M)$$

We cannot easily use the results on the expected stopping time of a single channel because the order of the expectation and the minimum operators cannot be swapped. Therefore we seek to analyze the stopping times of all channels jointly. [7] presents a way to calculate the joint (minimal) stopping time of several different irreducible Markov chains. We can make use of these results by associating a Markov chain with each of the channels.

Lemma 6. *The joint (minimal) stopping time of N independent channels with discrete capacity distribution options, initial balances $j_1, ..., j_N$ and capacity C, is given by $m_{j_1...j_N,C}$ where M is obtained by:*

$$vec(M) = (I_{C^{N+1}} - (I_C \otimes P^{(1)} \otimes \cdots \otimes P^{(N)}) \cdot (I_{C^{N+1}} - E))^{-1} \cdot \mathbb{1}_{C^{N+1}} \quad (7)$$

where $P^{(i)}$ is the transition matrix of channel i such that C is the failure state and $E = diag[\delta]$ where $\delta = vec([\delta_{j_1...j_N,i}])$. The parameters used in Eq. 7 are summarized in Table I.

Proof. A channel can be associated with a Markov chain representing its balance distribution, or more specifically the balance at one of its ends. The states of the chain will consist of all the possible balances of the channel when it is not depleted and a special failure state that is associated with zero and full capacity. The transaction probabilities will be the probabilities to go from a certain channel balance to another balance or failure state. Failure state leads to all other states with a certain probability to keep the Markov chain irreducible. The exact probabilities of leaving the failure state do not affect the joint stopping time since we are not interested in what happens after the failure state is reached for the first time. The initial state of the Markov chain is the initial balance of the channel. The state transition graph for the case of a single channel is presented in Fig. 7. □

TABLE I: Summary of main notations (Section III)

Term	Meaning
$vec(M)$	The vectorization of matrix M by concatenating the columns of M
$m_{h_1...h_N,j}$	The mean first time node j is reached when each channel i starts from node h_i
C	The capacity of each channel
$[\delta_{j_1...j_N,i}]$	A matrix of zeros and ones of the same size as M in which an element is equal to 1 in case its first index includes the second index
E	A diagonal matrix where the diagonal is the vectorization of $\delta_{j_1...j_N,i}]$
$P^{(i)}$	The transition matrix associated with channel i
$\mathbb{1}_n$	vector of ones of size n
I_n	Identity matrix of size n
$\otimes$	Kronecker product

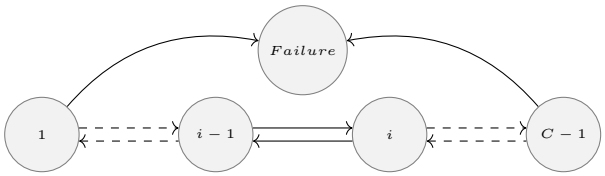

Fig. 7: transition diagram for a single channel with total capacity of C in case the balance can only increase or decrease by 1. Transitions from the Failure state are not included

M is a two-dimensional matrix where the height index $h_1...h_N$ represents the initial state of each of the N channels, and the width index j represents the target state. Accordingly, the last column of M holds the stopping times for each combination of initial states. An example of M for two channels with capacity 2 is:

$$\begin{bmatrix} m_{11,1} & m_{11,2} \\ m_{12,1} & m_{12,2} \\ m_{21,1} & m_{21,2} \\ m_{22,1} & m_{22,2} \end{bmatrix} = \begin{bmatrix} 0 & 3 \\ 0 & 0 \\ 0 & 0 \\ 1 & 0 \end{bmatrix}$$

In the example above we can see that the expected time of reaching the failure state, 2, when both channels start at state 1, is stored in $m_{11,2}$ and is equal to 3.

We verified the validity of Lemma 6 by empirically evaluating the stopping time of two independent channels with

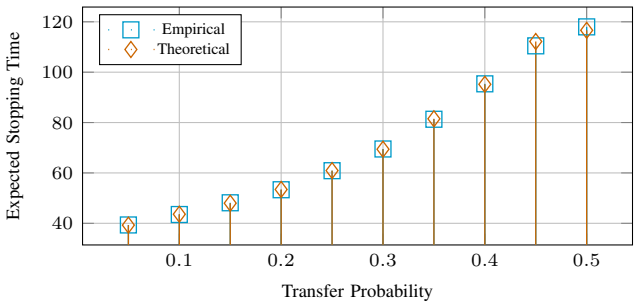

Fig. 8: The expected time until the first channel failure in a set of two channels with capacity 20 as a function of the transfer probability in one direction.

capacity 20 where in each time step a transaction of size 1 is transferred with a given probability p in one direction and with probability $1-p$ in the other direction. We compared the results for various transfer probabilities to the theoretical value obtained by Eq. 7 in Fig. 8.

Although M gives us all the information we need, in practice, it is not feasible to calculate M for even moderately large PCNs because Eq. 7 includes matrices of size $C^{N+1} \times C^{N+1}$ a figure that becomes fairly large pretty quickly. For example for 10 channels of capacity 10 each, $C^{N+1} = 10^{11}$. In addition, it is difficult to calculate the transition probability matrix accurately if there is not sufficient data. In addition to the computational limitations related to M, we assumed that the channels are not correlated, which is not a reasonable assumption since transactions are often performed using a multi-hop path that influences a few channels at the same time. Furthermore, this approach requires the assumption that all the channels have the same capacity.

Another interesting observation is that *this method can be used to accurately calculate the stopping time of a single channel with or without drift in case the channel has a discrete set of possible capacity distributions*.

IV. Minimal Expected Channel Stopping Time Optimization

Motivation. When a payment channel becomes disconnected, certain transfers cannot be performed, and in many cases there is a need for a costly on-chain transaction to rebalance the channel. As mentioned, users do not know the exact balance of channels that they do not maintain; thus, even when a transaction has a path on which it will succeed, the user may unknowingly choose a path that includes a depleted channel, which would lead to failure. Therefore, it is of interest to postpone the first failure even when it does not render the topology disconnected. Since we saw that the expected stopping time of a channel may depend on the distribution of transfers over it as well as its initial balance, it is possible to increase the time until the first failure by re-distributing the funds over the channels. It is not clear how to use the results of Section III for this task since the model we used there does not sustain different channel capacities and has some major drawbacks regarding accuracy and space complexity. Therefore, we propose to consider a simpler metric for maximization towards increasing the expected minimal stopping time.

TABLE II: Summary of main notations (Section IV)

Symbol	Meaning
K	The sum of funds of all the channels in the topology
C_i	The total capacity of channel i
b_i	maximum value of $\min(B_A^2, B_B^2)$ given $B_A + B_B = C_i$

Consider the set of edges $\{E_i\}_{i=1}^M$ composing a given network, and assume that each edge E_i has a total capacity of C_i and a variance of balance change of σ_i^2. Our aim is to re-distribute the funds in order to maximize the smallest lower bound on the expected stopping time given that the total amount of funds in all the channels is K.

Definition 7 (Minimal Expected Channel Stopping Time)**.** *Given a set of channels $\{E_i\}_{i=1}^M$ with corresponding stopping times $\{T_i\}_{i=1}^M$, the minimal expected channel stopping time is the smallest expected stopping time of all the expected times of all the channels in the set, namely, $\min_{i\in(1,M)} E[T_i]$.*

Lemma 7. *The minimal expected channel stopping time of a set of balanced channels $\{E_i\}_{i=1}^M$ with corresponding transaction variances, $\{\sigma_i^2\}_{i=1}^M$, and balances $\{B_{\alpha_i}^0, B_{\beta_i}^0\}_{i=1}^M$, is bounded from below by*

$$\min_{i\in(1,M)} \frac{\min(B_{\alpha_i}^0, B_{\beta_i}^0)^2}{\sigma_i^2}.$$

Proof. The lemma is established by assigning the lower bounds on the stopping time of each channel (Eq. 1) into the definition of the minimal expected stopping time. □

First, note that, given the capacity of a channel E_i, in order to maximize the lower bound on its stopping time we need to maximize $\min(B_{\alpha_i}, B_{\beta_i})^2$. We thus denote $b_i = \max \min(B_{\alpha_i}, B_{\beta_i})^2 = \max_{B_{\alpha_i}} \min(B_{\alpha_i}, (C_i - B_{\alpha_i}))^2 = 0.25 \cdot C_i^2$ and formulate the following optimization problem:

$$\begin{aligned} &\max_{b_i} \min_i \quad \tfrac{b_i}{\sigma_i^2} \\ &\text{s.t.} \qquad \textstyle\sum_i C_i = \sum_i 2 \cdot \sqrt{b_i} = K \end{aligned}$$

Lemma 8. *The minimal lower bound on the expected stopping time of a set of balanced channels with a constant total capacity is maximized when all lower bounds are equal.*

Proof. Assume by way of contradiction that the distribution of funds over the channels is optimal but the lower bounds on the average stopping times are not equal. Since not all lower bounds are equal, there must exist a channel E_i for which the lower bound is larger than the minimal lower bound across the channels. Due to the continuity of the lower bound with respect to b_i, there is some $\Delta b > 0$ that can be deducted from E_i's total balance and still maintain the above property. The minimal lower bound can be improved by re-distributing Δb among all other channels for which the lower bound is minimal and E_i equally because the lower bound on the stopping time of a channel E_j is a strictly increasing function of b_j. □

With Lemma 8 at hand, instead of an optimization problem we get a set of equations, as follows:

$$\begin{aligned} &\tfrac{b_i}{\sigma_i^2} = \tfrac{b_j}{\sigma_j^2} && \forall i,j \\ &\textstyle\sum_i C_i = \sum_i 2 \cdot \sqrt{b_i} = K \end{aligned} \tag{8}$$

By comparing every channel E_i to E_1 we have $\frac{b_i}{\sigma_i^2} = \frac{b_1}{\sigma_1^2}$. We can then express b_i in terms of b_1:

$$b_i = \frac{\sigma_i^2}{\sigma_1^2} \cdot b_1 \tag{9}$$

Now, we are left with only one unknown variable, namely b_1. We can write the following total capacity condition equation:

$$\sum_i 2 \cdot \sqrt{b_i} = \frac{2 \cdot \sqrt{b_1}}{\sigma_1} \cdot \sum_j \sigma_j = K$$

Isolating b_1 and assigning to Eq. 9, we get:

$$b_1 = \frac{K^2 \cdot \sigma_1^2}{4 \cdot \left(\sum_j \sigma_j\right)^2}$$

$$b_i = \frac{K^2 \cdot \sigma_1^2}{4 \cdot \left(\sum_j \sigma_j\right)^2} \cdot \frac{\sigma_i^2}{\sigma_1^2} = \frac{K^2 \cdot \sigma_i^2}{4 \cdot \left(\sum_j \sigma_j\right)^2}$$

Finally, we obtain the following lower bound on $E(T_i)$:

$$E[T_i] \geq \frac{b_i}{\sigma_i^2} = \frac{1}{\sigma_i^2} \cdot \frac{K^2 \cdot \sigma_i^2}{4 \cdot \left(\sum_j \sigma_j\right)^2} = \frac{K^2}{4 \cdot \left(\sum_j \sigma_j\right)^2} \tag{10}$$

as well as the following expression for C_i:

$$C_i = 2 \cdot \sqrt{b_i} = 2 \cdot \sqrt{\frac{K^2 \cdot \sigma_i^2}{4 \cdot \left(\sum_j \sigma_j\right)^2}} = K \cdot \frac{\sigma_i}{\sum_j \sigma_j} \tag{11}$$

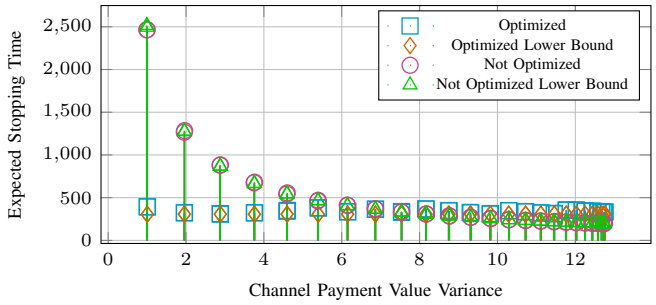

Fig. 9: Chain of 50 nodes: Average expected stopping time of a channel as a function of the edge transaction value variance. The minimal expected stopping time is the smallest point in the graph, which happens to be on the right side.

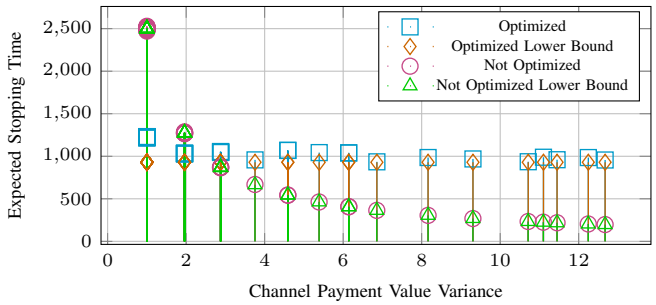

Fig. 10: Random tree of 50 nodes: Average expected stopping time of a channel as a function of the edge transaction value variance. The minimal expected stopping time is the smallest point in the graph, which happens to be on the right side.

As can be seen, the capacity of a channel depends on its variance, while the optimized lower bounds of the stopping time do not. The redistribution of funds cancels out the differences in the lower bounds of the stopping time, as expected. To evaluate how accurate the above lower bound is in practice and to evaluate how effective the optimization is, we simulated a series of random transactions of 5 tokens over a chain, as well as a random tree of 50 nodes, where each channel had a capacity of 100 split equally between its sides. We ran each simulation for 10^6 time steps. For each channel in the topology, we calculated the average time until depletion by dividing the number of simulation time steps by the number of channel depletion events encountered during the simulation. When a channel depletion occurred during the simulation, we immediately replenished the funds to the original state of the channel. In the tree and chain topologies there is only one simple path between every two nodes, hence the routing in this case is trivial. The results are presented in Fig. 9 and Fig. 10 respectively. Both figures depict the expected stopping time of each channel in the respective topology based on its variance before and after optimization, as well as the corresponding lower bounds. As can be seen, the optimization goal, namely the minimal expected stopping time, which is the lowest value in the graph, improved by several folds (more so in the tree topology). Moreover, the stopping time of channels, with a variance higher than 8 in the case of the chain and 3 in the case of the random tree, was improved by the optimization. The optimized stopping times are uniformly spread with respect to the variance and the lower bound is close to the actual stopping time. The improvement is more noticeable in the case of the random tree and the resulting stopping time is twice as high. Since the number of channels in a tree is always $n-1$, the sum of all the capacities is the same in both cases, hence we conclude that a chain is not an efficient topology.

Fig. 11 is a histogram of channels based on the transaction value variances of channels in both the chain and random tree topologies. The chain topology has a higher number of edges with low variance, while the random tree topology has a higher number of edges with high variance. The higher occurrence of channels with high variance in the tree topology allows for a more effective redistribution of funds during the optimization process. In contrast, the chain topology, which predominantly consists of channels with low variance, has a limited potential for capacity redistribution and subsequent improvement in stopping time. Hence we conclude that *topologies that facilitate higher variance of the transaction values are preferred*.

In addition, we evaluated the average time until a certain number of failures occurs in a random tree topology in Fig. 12. As can be seen, the average time improved by at least 50% for any number of failures, and the duration between consecutive failures increased as well.

Lightning evaluation. In addition, we evaluated the channel stopping time with and without optimization on the real topology of the Lightning network taken in December 2022. Lightning is the payment channel network of Bitcoin. The topology had 11268 nodes and 61966 edges. The transfer probabilities we used were derived from a traffic simulator from 2019 described in [8]. The simulator provides a series of transactions for a given topology. The size of each transaction was chosen to be 60000. For each channel in the topology,

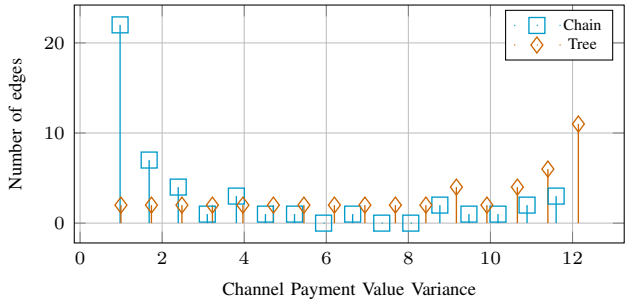

Fig. 11: The number of edges that possess a given transaction value variance

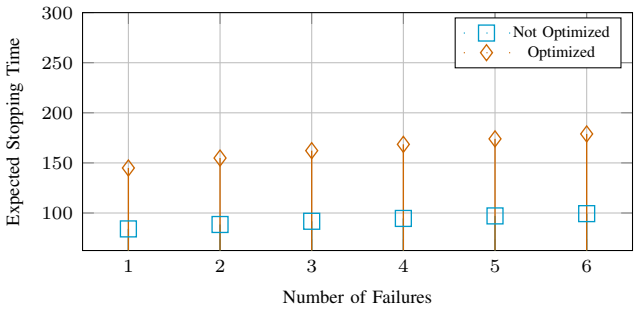

Fig. 12: Average of 20 Random trees of 50 nodes: Average expected time until x failures as a function of x.

we calculated the average time until depletion by dividing the number of simulation time steps by the number of channel depletion events encountered during the simulation. When a channel depletion occurred during the simulation, we immediately replenished the funds to the original state of the channel. We routed each transaction on the shortest path in terms of hop count. Since channels in Lightning have different capacities, and the stopping time is influenced both by the capacity and the channel variance, we chose to present the expected stopping time as a function of the lower bound instead of the variance. We present the expected stopping time and its lower bound before optimization in Fig. 13a. Moreover, we present the effect of the optimization on the expected stopping time and the lower bound in Fig. 13b. We averaged the acquired data over 30 equal-sized bins. As expected for points above the lower bound samples the average stopping time is larger than the lower bound. Since we ran the experiment for 10^7 steps, we focused on the channels where the lower bound is up to 10^5, to obtain a sufficiently accurate estimate of the expected stopping time and to avoid sampling bias due to the 10^7 limit. We can see that the empirical evaluation follows the lower bound closely and similarly to what we saw in the synthetic topologies: here too, after the optimization, the stopping times and the lower bound are evenly spread. The minimal expected stopping time (lowest point in the graph) improved by about $4 \cdot 10^5$ time steps.

V. Minimal Expected Channel Stopping Time Optimization: Distributed Version

Although distributing the total amount of funds between the channels seems like a reasonable approach, it might not be practical or easy to implement it, because usually each user decides how much to invest in opening payment channels, hence we cannot guarantee that users will comply with the optimized proposal. Furthermore, although the total amount of funds in the network should stay the same in order to guarantee that each user will invest the sought amount, just fulfilling this requirement is not enough. Accordingly, we would like to explore what each user can do in order to increase the stopping time of the channels that it operates.

We start by reformulating the optimization problem in order to account for a predetermined investment by each user in the network:

$$\begin{aligned} &\max_{B_i,B_j} \min_{(i,j)\in E} && \frac{min(B_i,B_j)^2}{\sigma^2_{i,j}} \\ &\text{s.t.} && \sum_{j\in N(i)} B_j = K_i \forall i \end{aligned}$$

Even though the reformulated problem has stricter constraints on the funds distribution among channels than the global version, it elevates an implicit constraint on the balance distribution within each channel. Indeed, till now we assumed that we could impose on each channel to split the capacity equally between its ends, however in the distributed version each end determines its investment independently, hence we can no longer make that assumption. Therefore, while in the global version there was a single (closed-form) solution, now there might be multiple optimal solutions.

Since the lower bound we found in Eq. 1 depends on both ends of the edge, it may seem hard to find an optimal solution for each node separately. However, we can decouple the problem by presenting the following reformulation of the lower bound. Note that B_α and B_β are non-negative.

$$E[T] \geq \frac{min(B^0_\alpha, B^0_\beta)^2}{\sigma^2} = min(\frac{B^2_\alpha}{\sigma^2}, \frac{B^2_\beta}{\sigma^2})$$

The minimal lower bound on the expected stopping time in the network can be expressed as:

$$\min_{(i,j)\in E} min(\frac{B^2_{i,j}}{\sigma^2_{i,j}}, \frac{B^2_{j,i}}{\sigma^2_{i,j}}) = \min_{n\in Nodes} min_{j\in N(n)}(\frac{B^2_{n,j}}{\sigma^2_{n,j}})$$

Where $N(i)$ is the neighborhood of node i. Note that $\{\frac{B^2_{n,j}}{\sigma^2_{n,j}}\}_{j\in N(n)}$ don't overlap, thus the minimum can be maximized by maximizing the minimum of each subset.

Lemma 9. *Under the condition that the total investment of each node is kept constant, the minimal expected channel stopping time of a set of balanced channels is maximized when each node n maximizes the minimum value of $\frac{B^2_{n,j}}{\sigma^2_{n,j}}$ for all neighboring nodes $j \in N(n)$.*

The capacity distribution proposed in Lemma 9 is optimal with respect to the minimal expected stopping time. However, it is not the only solution and there is room for further optimization with respect to other optimization goals. For example, since every node sets its balance on each edge independently of the balance at the other end of the channel, some edges may have different balances at their ends. Since the lower bound on the stopping time depends on the minimum of the balances, the extra balance on the larger end can be used to increase the stopping time of another channel without detrimenting the lower bound on the first channel.

We examined the practical impact of the optimization proposed in Lemma 9 on the minimal expected stopping time,

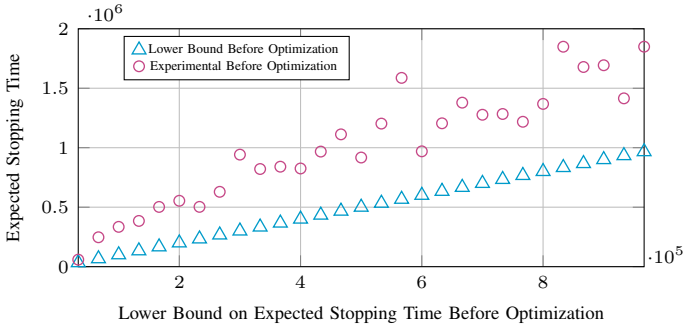

(a) Lightning Network: Average expected stopping time of channels with similar non-optimized lower bounds.

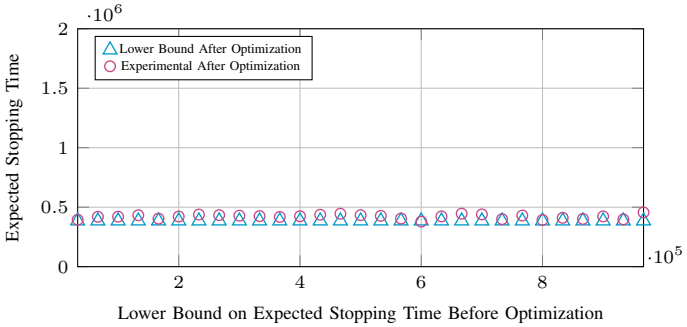

(b) Lightning Network: Optimized average expected stopping time of channels with similar non-optimized lower bounds.

Fig. 13: Lightning Network: Expected Stopping time before and after capacity optimization as the function of the lower bound on the expected stopping time before optimization. The minimal expected stopping time is the smallest point in each of the graphs.

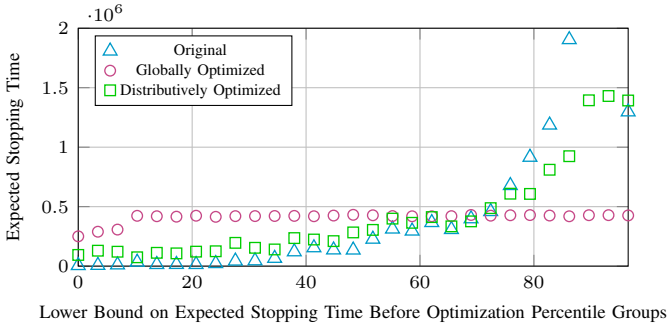

Fig. 14: Lightning Network: Expected stopping time before optimization, after global optimization and after distributed optimization, as the function of the lower bound on the expected stopping time before optimization percentile group.

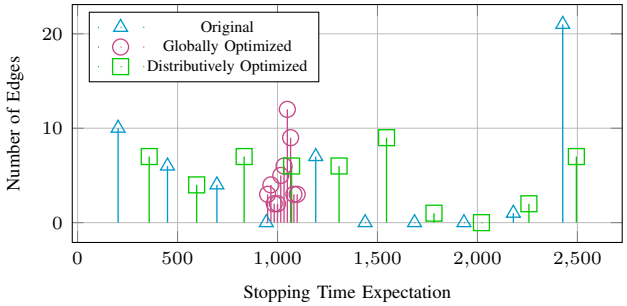

Fig. 15: The number of edges that possess a given expected stopping time in a random tree of 50 nodes before optimization, after global optimization and after distributed optimization.

as well as the overall effect on the network, in the topology of the Lightning network (as of December 2022). Here again, for each channel in the topology, we calculated the average time until depletion by dividing the number of simulation time steps by the number of channel depletion events encountered during the simulation. When a channel depletion occurred during the simulation, we immediately replenished the funds to the original state of the channel. We routed each transaction on the shortest path in terms of hop counts. Fig. 14 displays the average expected stopping time as a function of the lower bound before optimization in the Lightning network percentile group. As can be seen, the expected stopping time consistently improved in the lower percentiles but not as prominently as in the global case and, accordingly, the detriment to the higher percentiles was less substantial. We thus conclude that, *even though our optimization focused on the minimal expected stopping time, it has a positive effect on other channels with relatively low stopping time*. Fig. 15 presents a histogram of the expected stopping times without optimization, under global optimization and under distributed optimization, in a random tree of 50 nodes. Here again, we see that the effect of the distributed optimization on the minimal expected stopping time is smaller than in the case of global optimization. In addition, we see that the distributed optimization does not equalize the expected stopping times across the network but rather performs small local improvements.

VI. Stopping Time of a Channel in Particular Topologies

Let us examine the stopping time of a channel in various topologies. The topology along with the routing scheme determine the transactions on each channel. That is, different topologies may lead to different stopping times given the same payments. In this section we examine a few basic topologies, including two that were previously examined for PCNs (in [9] and [10]), namely the star and the ring.

Denote the balance of node i on the channel connecting i and j at time n as $B^n_{i\to j}$ and the corresponding requested increment from time $n-1$ to time n as $\Delta B^n_{i\to j}$. Denote the random variable representing the payment from node i to node j at time n as $Y^n_{i\to j}$. Across all following examples, we assume that the distribution of transfers from a user to every other user does not change over time and we denote its expectation and variance by $\mu_{Y_{i\to j}}$ and $\sigma^2_{Y_{i\to j}}$ correspondingly. If, in addition, the distribution of the payments is the same across all users, we use μ_Y and σ^2_Y to refer to the common variance and expectation. This assumption is useful in the sense that it allows us to compare various topologies without further knowledge about the transaction distribution. Note that the lower bound that we developed in Eq. 1 does not require this assumption and it depends on the transaction value variance, which can be estimated empirically. Notations are summarized in Table III.

TABLE III: Summary of main notations (Section VI)

Symbol	Meaning
$Y^n_{i\to j}$	Random variable which represents the transferred amount from node i to node j at time n
$\sigma^2_{Y_{i\to j}}$	The variance of $Y^n_{i\to j}$
$\mu^2_{Y_{i\to j}}$	The expectation of $Y^n_{i\to j}$
$B^n_{i\to j}$	The balance of node i on the channel connecting i and j at time n
$\Delta B^n_{i\to j}$	The requested increment to $B^n_{i\to j}$ from time $n-1$ to time n

A. Topology I: Tree

Consider a payment channel $i \leftrightarrow j$ in a tree topology, an illustration of the channel is presented in Fig. 16. In a tree structure, there is only one simple path between a pair of nodes. The edge $i \leftrightarrow j$ partitions the tree into two sub-trees,

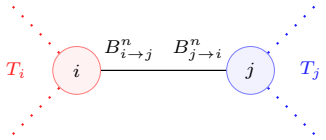

Fig. 16: An edge between node i and j in a tree. The edge divides the tree into two sub-trees: A subtree with node i marked as T_i and a subtree with node j marked as T_j. Balances of the channel between i and j appear near each of the two nodes.

T_i and T_j correspondingly. All possible transactions can be divided into 4 groups, namely: transactions within T_i and T_j, and transactions from nodes in T_i to nodes in T_j and vice versa. Transactions within a sub-tree do not pass through the examined edge, transactions from T_i to T_j pass on $i \leftrightarrow j$ in the direction from i to j and vice versa. Let us examine the distribution of the change in the balance of channel $i \leftrightarrow j$ at time n on the side of node i:

$$\Delta B^n_{i\to j} \triangleq B^n_{i\to j} - B^{n-1}_{i\to j} = \sum_{k\in T_i}\sum_{l\in T_j} Y^n_{k\to l} - \sum_{k\in T_j}\sum_{l\in T_i} Y^n_{k\to l}$$

Due to linearity of expectation:

$$\begin{aligned} E[\Delta B^n_{i\to j}] &= \sum_{k\in T_i}\sum_{l\in T_j} E[Y^n_{k\to l}] - \sum_{k\in T_j}\sum_{l\in T_i} E[Y^n_{k\to l}] \\ &= \sum_{k\in T_i}\sum_{l\in T_j} \mu_Y - \sum_{k\in T_j}\sum_{l\in T_i} \mu_Y \end{aligned}$$

Since $Y^n_{k\to l}$ are independent, the variance of $\Delta B^n_{i\to j}$ is the sum of the variances of $Y^n_{k\to l}$.

$$Var[\Delta B^n_{i\to j}] = \sum_{k\in T_i}\sum_{l\in T_j} Var[Y^n_{k\to l}] + \sum_{k\in T_j}\sum_{l\in T_i} Var[Y^n_{k\to l}]$$

In case all $Y_{k\to l}$ are identically randomly distributed, we denote the number of nodes on the first sub-tree T_i as M_i and the number of nodes on the second sub-tree T_j as M_j. Since the number of potential transfers from T_j to T_i is equal to the number of potential transfers from T_i to T_j and they all have the same distribution, the expectation of $\Delta B^n_{i\to j}$ is 0 even though μ_Y is not necessarily 0.

$$E[\Delta B^n_{i\to j}] = \sum_{k\in T_i}\sum_{l\in T_j} \mu_Y - \sum_{k\in T_j}\sum_{l\in T_i} \mu_Y = 0$$

The variance of $\Delta B^n_{i\to j}$ is:

$$Var[\Delta B^n_{i\to j}] = 2 \cdot M_i \cdot M_j \cdot \sigma^2_Y \tag{12}$$

Since the expectation of $\Delta B^n_{i-1\to i}$ is 0, we can use Eq. 1 to bound the expectation of the stopping time of a channel in a tree from below:

$$E[T] \geq \frac{\min^2(B^0_{i\to j}, B^0_{j\to i})}{2 \cdot M_i \cdot M_j \cdot \sigma^2_Y} \tag{13}$$

B. Topology II: Chain

Consider a chain of M nodes connected by $M-1$ payment channels, as illustrated in Fig. 17.

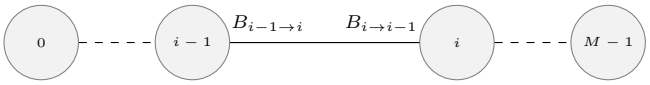

Fig. 17: A Chain of M nodes numbered from 0 to $M-1$. The balances on the channel between $i-1$ and i are written next to the corresponding nodes

A chain is a tree, therefore we can use Eq. 13. The channel $i \leftrightarrow i-1$ divides the chain into two sub-chains of size $M_i = M - i$ and $M_{i-1} = i$, assigning those into Eq. 13 we get:

$$E[T] \geq \frac{\min^2(B^0_{i\to i-1}, B^0_{i-1\to i})}{2 \cdot \sigma^2_Y \cdot i \cdot (M-i)} \tag{14}$$

In a chain, the stopping time of a channel depends on the location inside the chain. Channels closer to the middle of the chain are more frequently used as intermediate channels and therefore are more likely to fail. This opens up room for optimization, as we saw in Section IV.

C. Topology III: Star

Consider a star of M nodes connected by $M-1$ payment channels, as illustrated in Fig. 18.

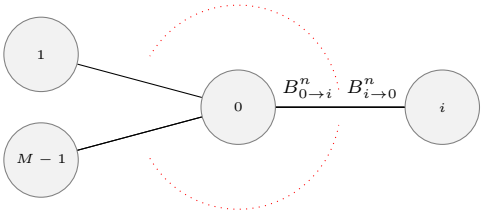

Fig. 18: A Star composed of an inner node 0 and $M-1$ leaf nodes numbered from 1 to $M-1$ (such that $i \in [1, M-1]$). The balances on the channel between i and 0 are written next to the corresponding nodes.

A star is a tree, thus we use Eq. 13. The channel $i \leftrightarrow 0$ divides the star into a sub-tree of size $M_i = 1$ which contains the leaf node i, and a sub-tree that contains the remaining $M_0 = M-1$ nodes. Together with Eq. 13 we get:

$$E[T] \geq \frac{\min^2(B^0_i, B^0_0)}{2 \cdot (M-1) \cdot \sigma^2_Y} \tag{15}$$

D. Topology IV: Clique

Consider a clique of M nodes connected by M^2 payment channels. In a clique, the shortest path between any two nodes consists of a single edge. Assuming all transfers are performed on the shortest path, the variance and expectation of the balance change, is identical to the corresponding parameters in a chain of two nodes. Thus, by assigning to Eq. 14, we get:

$$E[T] \geq \frac{\min^2(B^0_{i\to j}, B^0_{j\to i})}{2 \cdot \sigma_Y^2} \tag{16}$$

E. Topology V: Cycle

Consider a cycle of M nodes labeled from 0 to $M-1$ as illustrated in Fig. 19.

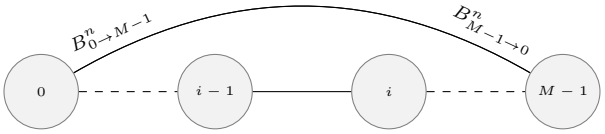

Fig. 19: A cycle of M nodes

Since there are two paths between each pair of nodes, the transfers are not well-defined by the payments. We thus assume that the source node chooses a path randomly with a probability of 0.5.

Since the cycle of nodes is symmetric, we examine a single channel and extend the conclusions to the rest of the channels. For simplicity, we look at $M-1 \leftrightarrow 0$. First, we define $Z_{k\to l}$ to be a random variable that is equal to the value of the transferred amount from node k to node l in case the transfer is performed in the direction that includes the channel $0 \leftrightarrow M-1$, and it is equal to 0 otherwise. The change to the balance of $0 \leftrightarrow M-1$ can be expressed using $Z_{k\to l}$ as:

$$\Delta B_{M-1\to 0} = \sum_{k<l} Z^n_{k\to l} - \sum_{k>l} Z^n_{k\to l}$$

The above holds because between each pair of nodes k and l, $k < l$, there is only one path that includes the channel $0 \leftrightarrow M-1$. If the transfer is done from k to l then it is passed from 0 to $M-1$ and therefore increases $\Delta B_{M-1\to 0}$; else, it is passed in the opposite direction and it decreases that value. If the chosen transfer direction does not include the channel $0 \leftrightarrow M-1$, it does not affect it. An illustration of the effect of the transfer direction on the examined channel can be seen in Fig. 20.

Since we assume that the direction of a transfer does not depend on its amount, $E[Z^n_{k\to l}] = 0.5 \cdot E[Y^n_{k\to l}]$. From linearity of expectation:

$$E[\Delta B_{M-1\to 0}] = \sum_{k<l} 0.5 \cdot E[Y^n_{k\to l}] - \sum_{k>l} 0.5 \cdot E[Y^n_{k\to l}] = 0.$$

The variance of $Z^n_{k\to l}$ can be calculated using the law of total variance. Note that $Z^n_{k\to l}$ given the transfer direction equals

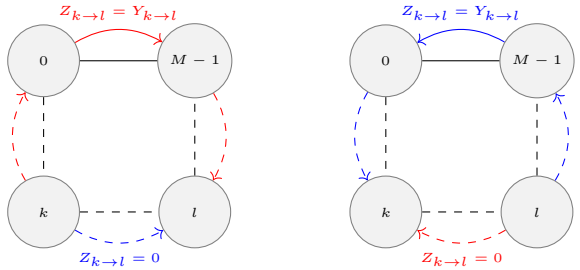

(a) transfer from k to l (b) transfer from l to k

Fig. 20: Illustration of the influence of the transfer direction on $\Delta B_{M-1\to 0}$. There are two possible directions of transfer between each pair of nodes. Clockwise path is marked in red, counter clockwise path in blue. As can be seen, in each subfigure only one direction influences the $0 \leftrightarrow M-1$ channel.

$Y_{k\to l}$ if the transfer direction includes $0 \leftrightarrow M-1$ and equals 0 otherwise.

$$\begin{aligned} Var(Z^n_{k\to l}) &= E[Var(Z^n_{k\to l}|\text{ transfer direction})] \\ &\quad + (E[Z^n_{k\to l}|\text{ transfer direction}]) \\ &= 0.5 \cdot Var(0) + 0.5 \cdot Var(Y^n_{k\to l}) \\ &\quad + 0.5 \cdot E[0]^2 + 0.5 \cdot E[Y^n_{k\to l}]^2 \\ &= 0.5 \cdot Var(Y^n_{k\to l}) \end{aligned}$$

Since $Z^n_{k\to l}$ are independent of each other, the variance of $\Delta B_{M-1\to 0}$ is the sum of the relevant $Z^n_{k\to l}$:

$$\begin{aligned} Var[\Delta B_{M-1\to 0}] &= \sum_{k<l} Var[Z^n_{k\to l}] + \sum_{k>l} Var[Z^n_{k\to l}] \\ &= \sum_{k<l} 0.5 \cdot Var[Y^n_{k\to l}] + \sum_{k>l} 0.5 \cdot Var[Y^n_{k\to l}] \\ &= 0.5 \cdot 0.5 \cdot M \cdot (M-1) \cdot \sigma^2 \\ &= 0.25 \cdot M \cdot (M-1) \cdot \sigma^2 \end{aligned}$$

Since the expectation of the increment is 0 and its variance is known, we can calculate the expected failure time of the channel using Eq. 1:

$$E[T] \geq \frac{4 \cdot \min^2(B^0_{i\to i-1}, B^0_{i-1\to i})}{\sigma^2 \cdot M \cdot (M-1)} \tag{17}$$

As expected, the lower bound on $E[T]$ does not depend on the location inside the cycle.

VII. Related Work

Offchain Payment Channels. Layer two protocols, built on top of (layer one) blockchains, have been implemented for major blockchain networks, such as Lightning [3] for Bitcoin. In most offchain networks, transactions between users that are not connected by a payment channel are performed by intermediate transactions through other nodes on a multi-hop path; however, other approaches exist.

Virtual Payment Channels. Virtual payment channels are payment channels that are built on top of regular payment channels and use them as arbitrators, similar to how regular payment channels utilize the blockchain. Virtual payment channels alleviate some shortcomings of regular payment

channels, such as reliance on intermediary nodes being constantly online, and they improve the speed and privacy of the transaction confirmation. However, they do not solve the issue of channel depletion. For our purposes, we treat virtual channels as regular channels.

Payment channel virtualization was introduced by Dziembowski et al. [11] for blockchain architectures that provide a Turing-complete scripting language, such as Ethereum, making it unsuitable for blockchain architectures with limited scripting capabilities, such as Bitcoin. Jourenko et al. [12] and Aumayr et al. [13] further optimized the construction of virtual payment channels by providing an implementation that does not have this limitation, enabling virtual payment channels on a broader range of blockchain systems. Xie et al. [14] used a bidirectional locking mechanism to lock the collateral on payment channels to increase the capacity of the virtual channel. Jia et al. [15] presented a cross-chain virtual payment channel scheme implemented a blockchain system that supports Turing-complete scripting language.

Routing Schemes. Several studies focused on routing schemes. Sivaraman et al. [4] optimized success ratio and volume by packetizing transactions and spreading packets across time and routes, using conjecture control. Mazumdar and Ruj [16] introduced a secure and privacy-preserving atomic multipath payment protocol, allowing to securely split transactions into smaller parts to increase the success rate. Mazumdar et al. [17] proposed an efficient privacy preserving routing algorithm while maintaining a high success ratio. Wang et al. [18] optimized the success volume by dividing transactions into "mice" and "elephants" and routing elephant transactions using a modified max-flow algorithm. Bagaria et al. [19] and Rahimpour et al. [20] proposed techniques to safely construct redundant payment paths to improve latency and throughput, while Awathare et al. [21] proposed to do so by re-routing transactions from intermediate nodes around a unidirectional channel to avoid failure. Hong et al. [22] optimized multi-path payment routing using limited capacity information. Xie et al. [23] proposed a multi-path routing algorithm based on the Ant Colony Optimization (ACO).

Topology and Demand. Khamis et al. [9] optimized the topology given the demand matrix by assigning nodes from the demand matrix to nodes in a given topology or by constructing a new topology with a limited number of edges. Khamis et al. [24] optimized the demand matrix itself by finding an equivalent matrix with the same user impact using fewer transactions, thus reducing the number of performed transactions. Lange et al. [25] and Davis et al. [26] explored the question of which connection point is preferred for joining the network. Sivaraman et al. [27] studied the role of network topology and channel imbalance on credit network throughput. Wu et al. [28] proposed to partition the nodes of the network into interconnected clusters surrounding supernodes such that each supernode pools and manages all funds within its cluster thus increasing scalability and liquidity. Rohrer et al. [29] proposed an algorithm that enables a distributed and concurrent execution of transactions without violating capacity constraints. To this end, they introduced the concept of capacity locking and used it in an extended push-relabel algorithm.

Rebalancing Schemes. Some studies focused on actively rebalancing the channels by performing transactions with the sole purpose of increasing the future success rate. Pickhardt et al. [30] defined an imbalance measure and proposed a greedy rebalancing algorithm. Khalil et al. [31] proposed a leader-based rebalancing scheme. Sahoo et al. [32] introduced a secure rebalancing model that maintains the safe state of every honest participant. Even so, a rebalancing action does not come for free as it is usually done via additional transactions that may carry additional fees. Recently, Kotzer et al. [33] showed the occurrence of the Braess paradox in payment channel networks with regards to routing fees. Sometimes, establishing new payment channels increases the expected fees despite the higher routing flexibility they allow.

VIII. Conclusions and Discussion

Overview. We examined the stopping time of payment channels. An overview of our contributions is summarized in Table IV. First, we looked at one channel in isolation and established a lower bound that depends on the initial channel balance and the variance of the payments over it for the case of a balanced channel, and calculated it for various topologies. Then we explored a private case for a channel with drift, calculated and maximized its stopping time, and analyzed the results in order to draw conclusions for the general case.

Next, we explored an approach for calculating the expected first failure time of a set of channels. For future work, we believe it should be possible to extend the results of Section III by expanding the model to incorporate channels with different capacities. This extension is feasible because [7] claims to provide a result for multiple random walks on different subgraphs.

Next, we attempted to optimize the lower bounds of a set of channels by increasing the smallest lower bond on the expected stopping time. We observed that the stopping time of a channel depends on its capacity as well as on its transaction variance. Since the latter highly depends on the location of the channel and the topology, by redistributing the total capacity the minimal expected stopping time of the channels in the network can be improved.

Through experiments, we indicated that this approach has a positive effect on channels with relatively low stopping times as well, and not only on the minimal expected stopping time in the network. First, we examined a global approach: although this is not very feasible to implement due to the distributed nature of the network, it indicates the extent to which capacity redistribution can improve the minimal expected stopping time. Next, we introduced a (more realistic) distributed approach, in which each node redistributes the funds available to it among the channel ends it is connected to. In both cases, we focused on balanced channels (transaction-wise). Although this assumption does not always hold, it arises in several topologies, some of which we explored in Section VI, due to basic symmetry assumptions; moreover, it holds in circulation networks (explored in [4]) given proper routing. Both approaches rely on the knowledge of the transaction size variance on the edges. This information can be obtained by

TABLE IV: Contributions overview

		Target				
		Expected stopping time of a single balanced channel	Expected stopping time of a single channel with drift	Expected stopping time of a single channel in a topology	Minimal expected stopping time of a set of channels	Expected minimal stopping time of a set of channels
Action	Calculation	Lower bound [Eq. 1]	Fixed size payment distribution [Lemma 5]	Lower bounds Tree [Eq. 13] Chain [Eq. 14] Star [Eq. 15] Clique [Eq. 16] Cycle [Eq. 17]	Lower bound [Lemma 7]	Assuming channels are independent: exact but costly [Lemma 6]
	Capacity distribution optimization	-	Optimal [Eq. 4]	-	Optimized lower bound [Eq. 10] Distributed strategy [Lemma 9]	-

monitoring the transactions on each edge and keeping a running average of the square values of the passing transactions.

Discussion and Future Work. Throughout this work, we have focused on survivability in the sense of the first channel failure in the network. However, other definitions exist. For example, it can be of interest to explore the first time the network becomes disconnected as opposed to the first channel depletion. Even though a single depleted channel may still lead to transaction failures, the existence of multiple depleted channels may more severely affect the transaction failure rate.

In addition, our proposed optimization may be undermined by malicious attacks by enforcing a very unbalanced transaction distribution on certain channels. Overcoming such vulnerabilities is an interesting topic for future work.

Furthermore, our approach can be integrated into a broader research landscape by complementing existing routing and rebalancing strategies. Indeed, it can be applied alongside any routing scheme and thus be easily combined with the routing strategy of choice. Additionally, our method can enhance rebalancing schemes by reducing the frequency of rebalancing operations and the entailed overhead. Exploring such integrations and their impacts is another interesting direction for future research.

Other directions for future work may include to further explore imbalanced channels and extend the context to a set of channels. In addition, the distributed approach that we considered is not the only approach that can achieve the goal of maximizing the minimal expected stopping time in the network, and we believe that the exploration of further approaches is a worthy topic for future work.

IX. Acknowledgment

We would like to thank Isaac Keslassy and Gal Mendelson for their helpful insights and suggestions.

References

[1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Decentralized Business Review*, 2008.

[2] G. Wood *et al.*, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, pp. 1–32, 2014.

[3] J. Poon and T. Dryja, "The Bitcoin Lightning network: Scalable off-chain instant payments," bitcoinlightning.com/wp-content/uploads/2018/03/lightning-network-paper.pdf, 2016.

[4] V. Sivaraman, S. B. Venkatakrishnan, M. Alizadeh, G. C. Fanti, and P. Viswanath, "Routing cryptocurrency with the Spider network," in *ACM HotNets*, 2018.

[5] J. Wolfowitz, "The efficiency of sequential estimates and Wald's equation for sequential processes," *The Annals of Mathematical Statistics*, vol. 18, no. 2, pp. 215–230, 1947.

[6] S. Chewi, https://inst.eecs.berkeley.edu/~ee126/fa17/wald.pdf, 2017.

[7] R. Patel, A. Carron, and F. Bullo, "The hitting time of multiple random walks," *SIAM Journal on Matrix Analysis and Applications*, vol. 37, no. 3, pp. 933–954, 2016.

[8] F. Béres, I. A. Seres, and A. A. Benczúr, "A cryptoeconomic traffic analysis of Bitcoin's Lightning network," *CoRR*, vol. abs/1911.09432, 2019.

[9] J. Khamis and O. Rottenstreich, "Demand-aware channel topologies for off-chain payments," in *International Conference on COMmunication Systems NETworkS (COMSNETS)*, 2021.

[10] G. Avarikioti, G. Janssen, Y. Wang, and R. Wattenhofer, "Payment network design with fees," in *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, 2018.

[11] S. Dziembowski, L. Eckey, S. Faust, and D. Malinowski, "Perun: Virtual payment hubs over cryptocurrencies," in *IEEE Symposium on Security and Privacy (SP)*, 2019.

[12] M. Jourenko, M. Larangeira, and K. Tanaka, "Lightweight virtual payment channels," in *International Conference on Cryptology and Network Security*, 2020.

[13] L. Aumayr, M. Maffei, O. Ersoy, A. Erwig, S. Faust, S. Riahi, K. Hostáková, and P. Moreno-Sanchez, "Bitcoin-compatible virtual channels," in *IEEE Symposium on Security and Privacy (SP)*, 2021.

[14] S. Xie, L. Xiao, D. Han, K. Xie, X. Li, and W. Liang, "HCVC: A High-Capacity Off-Chain Virtual Channel Scheme Based on Bidirectional Locking Mechanism," *IEEE Transactions on Network Science and Engineering*, vol. 11, no. 5, pp. 3995–4006, 2024.

[15] X. Jia, Z. Yu, J. Shao, R. Lu, G. Wei, and Z. Liu, "Cross-Chain Virtual Payment Channels," *IEEE Transactions on Information Forensics and Security*, pp. 3401–3413, 2023.

[16] S. Mazumdar and S. Ruj, "CryptoMaze: Privacy-preserving splitting of off-chain payments," *IEEE Trans. Dependable Secur. Comput.*, vol. 20, no. 2, pp. 1060–1073, 2023.

[17] S. Mazumdar, S. Ruj, R. G. Singh, and A. Pal, "Hushrelay: A privacy-preserving, efficient, and scalable routing algorithm for off-chain payments," in *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2020.

[18] P. Wang, H. Xu, X. Jin, and T. Wang, "Flash: Efficient dynamic routing for offchain networks," in *ACM CoNEXT*, 2019.

[19] V. Bagaria, J. Neu, and D. Tse, "Boomerang: Redundancy improves latency and throughput in payment-channel networks," in *International Conference on Financial Cryptography and Data Security*, 2020.

[20] S. Rahimpour and M. Khabbazian, "Spear: Fast multi-path payment with redundancy," in *ACM Conference on Advances in Financial Technologies*, 2021.
[21] N. Awathare, Suraj, Akash, V. J. Ribeiro, and U. Bellur, "Rebal: Channel balancing for payment channel networks," in *International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*, 2021.
[22] H.-J. Hong, S.-Y. Chang, and X. Zhou, "Auto-tune: An efficient autonomous multi-path payment routing algorithm for Payment Channel Networks," *Computer Networks*, vol. 225, 2023.
[23] M. Xie and X. He, "PACO: Efficient Routing in Payment Channel Networks," in *International Conference on Computer Science and Blockchain (CCSB)*, 2023, pp. 26–31.
[24] J. Khamis, S. Schmid, and O. Rottenstreich, "Demand matrix optimization for offchain payments in blockchain," in *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2021.
[25] K. Lange, E. Rohrer, and F. Tschorsch, "On the impact of attachment strategies for payment channel networks," in *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2021.
[26] V. Davis and B. Harrison, "Learning a scalable algorithm for improving betweenness in the lightning network," in *International Conference on Blockchain Computing and Applications (BCCA)*, 2022.
[27] V. Sivaraman, W. Tang, S. Bojja Venkatakrishnan, G. Fanti, and M. Alizadeh, "The effect of network topology on credit network throughput," *SIGMETRICS Perform. Eval. Rev.*, vol. 49, no. 3, p. 59–60, 2022.
[28] J. Wu and S. Jiang, "On increasing scalability and liquidation of lightning networks for blockchains," *IEEE Transactions on Network Science and Engineering*, vol. 9, no. 4, pp. 2589–2600, 2022.
[29] E. Rohrer, J.-F. Laß, and F. Tschorsch, "Towards a concurrent and distributed route selection for payment channel networks," in *Data Privacy Management, Cryptocurrencies and Blockchain Technology*, 2017.
[30] R. Pickhardt and M. Nowostawski, "Imbalance measure and proactive channel rebalancing algorithm for the Lightning Network," in *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2020.
[31] R. Khalil and A. Gervais, "Revive: Rebalancing off-blockchain payment networks," in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2017.
[32] S. S. Sahoo, M. M. Hosmane, and V. K. Chaurasiya, "A secure payment channel rebalancing model for layer-2 blockchain," *Internet of Things*, vol. 22, 2023.
[33] A. Kotzer and O. Rottenstreich, "Braess paradox in layer-2 blockchain payment networks," in *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2023.

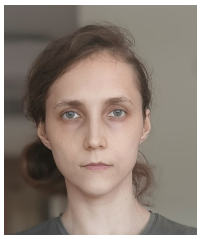

Yekaterina Podiatchev received her B.Sc. degree in Electrical Engineering from the Technion, Haifa, Israel, in 2019, and completed her M.Sc. degree at the Technion Viterbi Department of Electrical and Computer Engineering in 2024.

Ori Rottenstreich is an associate professor at the Taub department of Computer Science and the Viterbi department of Electrical and Computer Engineering of the Technion, Haifa, Israel. Previously, he was a Postdoctoral Research Fellow at Princeton university. Ori received his B.Sc. degree in Computer Engineering and Ph.D. degree in Electrical Engineering from the Technion.

Ariel Orda (Fellow, IEEE) received the B.Sc. (summa cum laude), M.Sc., and D.Sc. degrees in electrical engineering from Technion, Haifa, Israel, in 1983, 1985, and 1991, respectively. He is the Herman and Gertrude Gross Professor of communication at the Viterbi Department of ECE, Technion, and served as its dean during 2014-2018. His research interests include the application of game theory to networks and distributed systems, network routing, survivability, QoS provisioning, and wireless networks. He served as the program co-chair for IEEE INFOCOM 2002, WiOpt 2010 and Netgcoop 2020, and as the general chair for Netgcoop 2012 and ACM SIGCOMM 2022. He was an editor of the IEEE/ACM Transactions on Networking and Computer Networks. He received several awards for research, teaching, and service.