

1 Boosting Liquidity in Payment Channel Networks 2 with Online Admission Control

3 **Anonymous author**

4 Anonymous affiliation

5 — Abstract —

6 Payment channel networks (PCNs) are a promising technology to improve the scalability of crypto-
7 currencies. PCNs, however, face the challenge that the frequent usage of certain routes may deplete
8 channels in one direction, and hence prevent further transactions. In order to reap the full potential
9 of PCNs, recharging and rebalancing mechanisms are required to provision channels, as well as an
10 admission control logic to decide which transactions to reject in case capacity is insufficient. This
11 paper presents a formal model of this optimisation problem. In particular, we consider an online
12 algorithms perspective, where transactions arrive over time in an unpredictable manner. Our main
13 contributions are competitive online algorithms which come with provable guarantees over time. We
14 empirically evaluate our algorithms on randomly generated transactions to compare the average
15 performance of our algorithms to our theoretical bounds. We also show how this model and approach
16 differs from related problems in classic communication networks.

17 **2012 ACM Subject Classification** Applied computing → Electronic commerce; Theory of computa-
18 tion → Mathematical optimization; Theory of computation → Online algorithms

19 **Keywords and phrases** Payment channel networks, Blockchain, Optimisation, Rebalancing, Online
20 Algorithms, Layer 2

21 **Digital Object Identifier** 10.4230/LIPIcs.CVIT.2016.23

22 **1** Introduction


23 Blockchain consensus protocols are notoriously inefficient: for instance, Bitcoin can only
24 support 7 transactions per second on average which makes it unrealistic to use in everyday
25 situations. Payment channel networks like Bitcoin's Lightning Network [15] and Ethereum's
26 Raiden [1] have been proposed as scalability solutions to blockchains. Instead of sending
27 transactions to the blockchain and waiting for the entire blockchain (which can comprise of
28 millions of users) to achieve consensus, any two users that wish to transact with each other
29 can simply open a payment channel between themselves. Opening a payment channel requires
30 an initial funding transaction on the blockchain where both users lock some funds only to
31 use in the channel. Once a payment channel is opened, the channel acts as a local, two-party
32 ledger: payments between the users of channel simply involve decreasing the balance of the
33 payer by the payment amount, and increasing the balance of the payee correspondingly.
34 As these local transactions only involve exchanging signatures between the two users and
35 do not involve the blockchain at all, they can be almost instantaneous. As long as there
36 is sufficient balance, payments can occur indefinitely between two users, until the users
37 decide to close the channel. This would involve going back to the blockchain and takes, in
38 the worst case, a small constant number of transactions. Thus, with only a small constant
39 number of on-chain transactions, any two users can potentially make arbitrarily many costless
40 transactions between themselves.

41 Apart from joining a payment channel network to efficiently transact with other users, an
42 additional financial incentive to joining the network is to profit from forwarding transactions.
43 Any two users that are not directly connected can transact with each other in a multi-hop
44 fashion as long as they are connected by a path of payment channels. To incentivise the
45 intermediary nodes on the path to forward the payment, the network typically allows these



© Anonymous author(s);
licensed under Creative Commons License CC-BY 4.0
42nd Conference on Very Important Topics (CVIT 2016).

Editors: John Q. Open and Joan R. Access; Article No. 23; pp. 23:1–23:22



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

46 nodes to charge a transaction fee. Thus, it is common for users to join the network specifically
 47 to play the role of an intermediary node that routes transactions, creating channels and
 48 fees optimally and selecting the most profitable transactions to maximise their profit from
 49 transaction fees [4, 8].

50 However, greedily accepting and routing incoming transactions could rapidly deplete
 51 a user’s balance in their channels. In particular, if certain routes are primarily used in
 52 one direction, their channels can get depleted, making it impossible to forward further
 53 transactions. Accounting for this problem can be non-trivial since demand patterns are hard
 54 to predict and often confidential.

55 To resolve this issue, PCNs typically support two mechanisms:

- 56 ■ *On-chain recharging*: A user can close and reopen a depleted channel with more funds
 57 on-chain.
- 58 ■ *Off-chain rebalancing*: An alternative solution is to extend the lifetime of a depleted
 59 channel without involving the blockchain, by finding a cycle of payment channels in the
 60 network to shift funds from one channel to another.

61 Both cases, however, entail a cost. Intermediaries need to consider the tradeoff between
 62 admitting transactions and potential recharging and rebalancing costs. This decision making
 63 process is especially important to big routers which are the primary maintainers of payment
 64 channel networks like the Lightning Network.

65 In this work, we focus on the problem of admission control, recharging and rebalancing in
 66 a single payment channel from the perspective of an intermediary node that seeks to route
 67 as many transactions as possible with minimal costs. Specifically, we address the following
 68 research question:

69
 70 *Can we design efficient online algorithms for deciding when to accept/reject transactions,
 71 and when to recharge or rebalance in a single payment channel?*

72
 73 We seek to address this problem with as few restrictions on user actions in order to
 74 ensure that our work remains realistic. Thus, we assume a fixed PCN topology with some
 75 recharging and rebalancing costs, and a global fee function that is linear in the transaction
 76 size. We also assume users incur a rejection cost in the form of opportunity cost when they
 77 reject to route a transaction.

78 We are interested in robust solutions which do not depend on any knowledge or assump-
 79 tions on the demand. Accordingly, we assume that transactions can arrive in an arbitrary
 80 order at a channel, and aim to design online algorithms which provide worst-case guarantees.
 81 We are in the realm of competitive analysis, and assume that an adversary with knowledge
 82 of our algorithms chooses the most pessimal online transaction sequence. Our objective is to
 83 optimise the *competitive ratio* [6]: we compare the performance of our online algorithms (to
 84 which the transaction sequence is revealed over time) with the optimal offline algorithm that
 85 has access to the entire transaction sequence in advance.

86 1.1 Our contributions

87 We initiate the study of a fundamental resource allocation problem in payment channel
 88 networks, from an online algorithms perspective. Our main result is a competitive online
 89 algorithm to admit transaction streams arriving at both sides of a payment channel, and
 90 also to recharge and rebalance the channel, in order to maximise the throughput over the
 91 channel while accounting for costs. In particular, our algorithm achieves a competitive ratio

of $7 + 2\lceil \log C \rceil$ where $C + 1$ is the length of the rebalancing cycle used to replenish the funds on the channel off-chain. We also provide lower bounds on the amount of funds needed in a channel in order to ensure our algorithm is c -competitive for $c < \frac{\log C}{\log \log C}$.

In order to prove our main theorem, we decompose the problem into two simpler sub problems that may also be of independent interest:

1. *Sub problem 1:* The first and most restrictive sub problem considers a transaction stream coming only from one direction across a payment channel, and users do not have the option to reject incoming transactions. We present a 2-competitive algorithm for this problem, which is optimal in the sense that no deterministic online algorithm can achieve a lower competitive ratio.
2. *Sub problem 2:* As a relaxation, our second sub problem allows users to reject transactions although all transactions are still restricted to come from one direction along a payment channel. We show that our algorithm achieves a competitive ratio of $2 + \frac{\sqrt{5}-1}{2}$ for this sub problem. We stress that our lower bound of 2 we achieve in sub problem 1 also holds in this sub problem, hence our competitive ratio of $2 + \frac{\sqrt{5}-1}{2}$ is close to optimal.

All intermediate and main results are summarised in Table 1. The algorithms and analysis designed to address these sub problems are eventually used as building blocks for our main algorithm and main theorem.

We complement our theoretical worst-case analysis by performing an empirical evaluation of the performance of our algorithm on randomly generated transaction sequences. We observe that our algorithms perform much better on average compared to our theoretical worst-case bound.

Sub problem	Competitive ratio
Unidirectional stream without rejection	2
Unidirectional stream with rejection	$2 + \frac{\sqrt{5}-1}{2}$
Bidirectional stream	$7 + 2\lceil \log C \rceil$

■ **Table 1** Summary of the theoretical results in our paper. The first column presents each sub problem we analyse in our paper and the second column shows the competitive ratio achieved by our algorithms for each sub problem

1.2 Related work

Maintaining balanced payment channels. As channel balances are typically private, classic transaction routing protocols on payment channel networks like Flare [16], SilentWhispers [13] and SpeedyMurmurs [17] focus mainly on throughput and ignore the issue of balance depletion. Recently, several works shift the focus on maintaining balanced payment channels for as long as possible while ensuring liveness of the network. Revive [10] initiated the study rebalancing strategies, Spider[19] uses multi-path routing to ensure high transaction throughput while maintaining balanced payment channels, the Merchant[20] utilises fee strategies to incentivise the balanced use of payment channels, and [11] uses estimated payment demands along channels to plan the amount of funds to inject into a channel during channel creation, to just give a few examples. Our work focuses on minimising costs incurred in the process of handling transactions across a channel and thus we also indirectly seek to maintain balanced payment channels. Moreover, in contrast to previous works which typically assume some

127 form of offline knowledge of the transaction flow in the network, we provide an algorithm
128 which comes with provable worst-case guarantees.

129 *Off-chain rebalancing.* Off-chain rebalancing has been studied as a cheaper alternative
130 to refunding a channel by closing and reopening it on the blockchain. In the Lightning
131 Network, there are already several off-chain rebalancing plugins for c-lightning¹ and lnd².
132 An automated approach to performing off-chain rebalancing using the imbalance measure as
133 a heuristic has been proposed in [14]. Our work similarly studies when to rebalance payment
134 channels, however we make the decision in tandem with other decisions like accepting
135 or rejecting transactions. Recently, [5] and [10] propose a global approach to off-chain
136 rebalancing where demand for rebalancing cycles is aggregated across the entire network and
137 translated to an LP which is subsequently solved to obtain an optimal rebalancing solution.
138 These approaches are orthogonal and complementary to ours as our focus concerns decision
139 making in a single payment channel and not the entire network.

140 *Online algorithms for payment channel networks.* Online algorithms for payment channel
141 networks have also been studied in [3] and [9]. Avarikioti et al. [3] establish impossibility
142 results against certain classes of adversaries, however they only consider a limited problem
143 setting where their algorithms can only accept or reject transactions (with constant rejection
144 cost). Fazli et al. [9] considers the problem of optimally scheduling on-chain recharging given
145 a sequence of transactions. In contrast to previous work, our work considers a more general
146 problem setting where our algorithms can not only accept or reject transactions, but also
147 recharge and rebalance channels off chain. We also extend the cost of rejection to take into
148 account the size of the transaction.

149 *Relationship to classic communication networks.* Admission control problems such as
150 online call admission [2, 12] are fundamental and have also received much attention in
151 the context of communication networks. However, in classic communication networks the
152 available capacity of a link in one direction is independent of the flows travelling in the other
153 direction, and moreover, link capacities are only consumed by the currently allocated flows.
154 In contrast, the capacities of links in payment-channel networks are permanently reduced by
155 transactions flowing in one direction, but can be topped up by flows travelling in the other
156 direction. The resulting rebalancing opportunity renders the underlying algorithmic problem
157 significantly different.

158 **2** Model

159 *Payment channels.* We model the payment channel network as an undirected graph $G =$
160 (V, E) . A payment channel between users ℓ (left) and r (right) in the network is an edge
161 $(\ell, r) \in E$. We denote the balance of user ℓ (resp. r) in the channel (ℓ, r) by $b(\ell)$ (resp. $b(r)$).
162 The *capacity* of the channel is the total amount of funds locked in the channel. That is, for a
163 channel (ℓ, r) , the capacity of (ℓ, r) is $b(\ell) + b(r)$. A left-to-right transaction of amount x
164 decreases ℓ 's balance by x and increases r 's balance by x and vice versa for a right-to-left
165 transaction of x .

166 *Recharging and rebalancing payment channels.* When a user in a channel does not have
167 sufficient funds to accept a transaction, the user can either reject the transaction, recharge the
168 channel, or rebalance the channel. Recharging the channel happens on-chain and corresponds

¹ <https://github.com/lightningd/plugins/tree/master/rebalance>

² <https://github.com/bitromortac/lndmanage>


169 to closing the payment channel on the blockchain and opening a new channel with more
 170 funds. In contrast, rebalancing the channel happens entirely off-chain (refer to Figure 1 for
 171 an example). Here, users find a cycle of payment channels to shift funds from one of their
 172 other channels to refund the depleted channel.

173 *Transactions.* We consider a transaction sequence $X_t = (x_1, \dots, x_t)$, $x_i \in \mathbb{R}^+$, that arrives
 174 at a payment channel online. Each transaction x_i has both a value and a direction along a
 175 payment channel. The value of a transaction is simply the amount that is being transferred.
 176 The direction of a transaction across a payment channel (ℓ, r) determines who is the sender
 177 and who is the receiver. When we have a sequence of transactions that go in both directions
 178 along a payment channel, we use \vec{x} to denote a transaction that goes from left-to-right and
 179 \overleftarrow{x} to denote a transaction that goes from right-to-left. We say a user, wlog ℓ , *accepts* a
 180 transaction of size x coming from the left to right direction along the channel (ℓ, r) if ℓ agrees
 181 to forward x to r . Similarly, we say a user ℓ *rejects* a transaction x coming from the left to
 182 right direction along the channel (ℓ, r) when ℓ does not forward the transaction to r . When it
 183 is clear which channel and direction we are referring to, we simply say ℓ accepts or rejects x .

184 *Costs.* We consider three types of costs in our problem setting:

- 185 **1. Rejecting transactions:** For a user ℓ , the revenue in terms of transaction fees from
 186 forwarding a payment of size x is $Rx + f_2$, where $R, f_2 \in \mathbb{R}^+$. Consequently, the cost of
 187 rejecting a transaction of size x is simply the opportunity cost of gaining revenue from
 188 accepting the transaction, i.e. $Rx + f_2$.
- 189 **2. On-chain recharging:** For any user ℓ , the cost of recharging a channel on-chain is
 190 $F + f_1$, where F is some function of the amount of funds ℓ puts into the new channel (this
 191 captures the opportunity cost of locking in the funds in the channel) and $f_1 \in \mathbb{R}^+$ is an
 192 auxiliary cost independent of F which captures the on-chain recharging transaction fee.
- 193 **3. Off-chain rebalancing:** For any user ℓ , the cost of off-chain rebalancing for an amount
 194 x is $C \cdot (Rx + f_2)$, where C is the length of the cycle along which funds are sent -1 . In
 195 the example of off-chain rebalancing in Figure 1, the length of the rebalancing cycle is 3
 196 and thus $C = 2$.

197 Let us denote by OFF the optimal offline algorithm and ON an online deterministic
 198 algorithm. We denote by $\text{COST}_{\text{ON}}(X_t)$ (resp. $\text{COST}_{\text{OFF}}(X_t)$) the total cost of ON (resp.
 199 OFF) given the transaction sequence X_t .



■ **Figure 1** Example of off-chain rebalancing with users ℓ , h , and r . The graph on the right depicts the channel balances after off-chain rebalancing.

Competitive ratio. We say an online algorithm ON is c -competitive if for every transaction sequence X_t generated by the adversary,

$$\text{COST}_{\text{ON}}(X_t) \leq c \cdot \text{COST}_{\text{OFF}}(X_t)$$

200 *Main problem.* Our main problem is to design a competitive deterministic online algorithm
 201 that determines when to accept/reject transactions and when to recharge or rebalance the
 202 channel given a bidirectional stream of transactions across a payment channel. More precisely,
 203 we consider a stream of transactions that can arrive from both right to left or left to right in

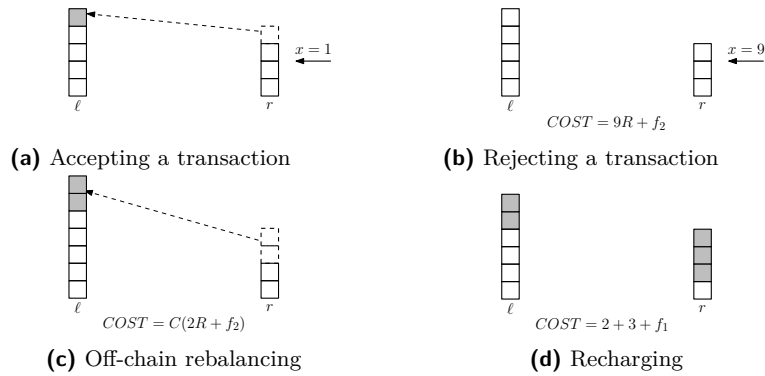


Figure 2 Example of actions users ℓ and r can take in the general bidirectional stream setting. Each square represents 1 coin.

204 a given payment channel (ℓ, r) . ℓ (resp. r) can choose to accept or reject transactions coming
 205 in the left-to-right (resp. right-to-left) direction in the stream. Either user would incur a cost
 206 of $Rx + f_2$ for rejecting a transaction of size x . Both users can also recharge the channel
 207 on-chain at any point, incurring a cost of $F_\ell + F_r + f_1$ where F_ℓ and F_r are functions of the
 208 funds put into the channel by ℓ and r respectively. Since transactions are streaming in both
 209 directions in this model, both users would incur costs in this setting. Thus, we seek to design
 210 an algorithm that minimises the cost of the *entire* channel. Refer to Figure 2 for examples of
 211 the actions that a user can take in our main bidirectional transaction stream setting.

212 To this end, we give a formal definition of two sub problems of decreasing restrictiveness
 213 on user actions. We present these sub problems as the algorithms and analysis used to solve
 214 these sub problems are used in developing the algorithm and analysis for our main problem.

215 *Unidirectional stream without rejection.* In this model, transactions stream only in one
 216 direction along a given payment channel. Here, we assume users cannot reject incoming
 217 transactions. Formally, given a channel (ℓ, r) and a transaction stream from wlog left to right,
 218 user ℓ only accepts a transaction x if $b(\ell) > x$. Otherwise, ℓ has to recharge the channel
 219 on-chain with more funds, incurring a cost of $F + f_1$ where F is some function of the amount
 220 of funds ℓ adds to the channel. As we only consider transactions streaming in one direction,
 221 only one user would incur costs in this setting (the user that has to decide whether to accept
 222 or reject transactions).

223 *Unidirectional stream with rejection.* In this model, we still restrict the transaction
 224 stream from wlog left to right in a payment channel (ℓ, r) . However, in addition to accepting
 225 transactions and recharging, ℓ can now also reject transactions, incurring a rejection cost of
 226 $Rx + f_2$ for a transaction of size x .

227 3 Algorithmic Building Blocks

228 Before we describe and analyse the performance of our algorithms in the various problem
 229 settings, we first introduce two algorithmic building blocks that we use extensively in our
 230 work. The first building block is an algorithm FUNDS. It takes a sequence of transactions as
 231 an input and returns the amount of funds that an optimal algorithm uses on this sequence.
 232 The purpose of the algorithm is to track the funds OFF has in their channel assuming that
 233 the sequence of transactions ends at this point. For the first two sub problems we show how
 234 to compute FUNDS. For the main problem, we propose a dynamic programming approach in

■ **Algorithm 1** (γ, δ) -recharging

```

Initialise:  $F_{tracker}, X \leftarrow 0, \emptyset$ 
1 for transaction  $x$  in order of arrival do
2   concatenate  $x$  to  $X$ 
3    $F'_{tracker} \leftarrow \text{FUNDS}(X)$ 
4   if  $F'_{tracker} > F_{tracker}$  then
5      $F_{tracker} \leftarrow F'_{tracker} + \delta$ 
6     recharge to  $\gamma F_{tracker}$ 

```

235 Appendix E. The second building block is a general recharging online algorithm that calls
 236 FUNDS as a subroutine and uses the output to decide when and how much to recharge the
 237 channel. The intuition behind the recharging online algorithm is to recharge whenever the
 238 amount of funds in OFF's channel "catches up" to the amount of funds ON has in their
 239 channel.

240 *Building block 1: tracking funds of OFF.* For a given transaction sequence $X_t =$
 241 (x_1, \dots, x_t) , let us denote $A(X_i)$ to be the amount of funds OFF would use in the channel
 242 if OFF gets the sequence $X_i = (x_1, \dots, x_i)$ (i.e. the length i prefix of X_t) as input. By
 243 appending subsequent transactions x_{i+1}, \dots, x_t from X_t to X_i , we can view $A(X_i)$ as a partial
 244 solution to the online optimisation problem that gets updated with any new transaction.
 245 In the unidirectional transaction stream (with or without rejection) setting, $A(X_i)$ refers
 246 to the funds a user locks into a payment channel. In the bidirectional transaction stream
 247 setting, $A(X_i)$ refers to the total balance of both users in the channel. We assume that given
 248 an input sequence X_t , $\text{FUNDS}(X_t)$ performs the necessary computations and returns $A(X_t)$.
 249 For our main problem, computing $\text{FUNDS}(X_t)$ is generally NP-hard, but we can approximate
 250 it to a constant factor, see [18] for more details.

251 *Building block 2: using tracking for recharging.* In Algorithm 1, we describe an online
 252 (γ, δ) -recharging algorithm ON that uses FUNDS as a subroutine to decide when and how
 253 much to recharge the channel. ON is run by one user (wlog ℓ) in a payment channel (ℓ, r) .
 254 ON calls FUNDS after each transaction to check if the new transaction sequence results in
 255 a significant increase in the amount of funds OFF has in their channel. Whenever ON
 256 notices that OFF's funds have increased above a threshold (Algorithm 1), ON recharges the
 257 channel with an amount of $\gamma(A(X_i) + \delta)$ where $A(X_i)$ is the amount of funds OFF has in
 258 their channel.

259 Let us denote $A_t := \max_{i \leq t} A(X_i)$. Now we state and prove (in Appendix C.1 and Ap-
 260 pendix C.2) two important properties of the (γ, δ) -recharging algorithm.

261 ► **Lemma 1.** *Algorithm 1 with parameters (γ, δ) ensures that ON always has at least γ times*
 262 *the amount of funds OFF has and ensures that ON incurs a cost of at most $\gamma(A_t + \delta) + f_1 \cdot \lceil \frac{A_t}{\delta} \rceil$.*

263 Next, we show a simple lower bound in terms of A_t for the cost of OFF given a sequence
 264 of transactions X_t .

265 ► **Lemma 2.** *If $A_t > 0$, then $\text{COST}_{\text{OFF}}(X_t)$ is at least $A_t + f_1$.*

266 4 Unidirectional transaction stream without rejection

267 In this section we consider the first sub problem where, given a payment channel (ℓ, r) ,
 268 transactions stream along the channel in only one direction (wlog left to right). Moreover, ℓ

269 has to accept an incoming transaction of size x and forward it to r if ℓ 's balance $b(\ell) \geq x$.
 270 Otherwise, ℓ needs to recharge the channel on-chain (and accept the transaction after).

271 The optimal offline algorithm OFF follows a simple strategy: since it knows the entire
 272 stream of transactions in advance, it makes a single recharging action at the beginning of the
 273 transaction sequence X_t of size $\sum_{i=1}^t x_i$. The cost incurred by OFF is thus $f_1 + \sum_{i=1}^t x_i$.

274 Now, we present a 2-competitive online algorithm ON for this sub problem (see Al-
 275 gorithm 5 in Appendix A). ON uses (γ, δ) -recharging with parameters $\gamma = 1$ and $\delta = f_1$. The
 276 algorithm accepts all transactions and the recharging ensures that ON always has enough
 277 funds.

278 **► Theorem 3.** *The algorithm described above is 2-competitive in the unidirectional transaction*
 279 *stream without rejection.*

280 In addition, we note that ON is optimal in this setting. The next theorem (with proof
 281 in Appendix C.4) proves that no deterministic algorithm can achieve a strictly smaller
 282 competitive ratio compared to ON. In particular, our proof shows that ON cannot lock too
 283 much funds into the channel, otherwise ON's cost is too high, but if ON locks too little
 284 funds, it needs to recharge often.

285 **► Theorem 4.** *There is no deterministic algorithm that is c -competitive for $c < 2$ in the*
 286 *unidirectional transaction stream without rejection sub problem.*

287 **5 Unidirectional transaction stream with rejection**

288 In this section we consider the second sub problem where transactions are still streaming
 289 along a given payment channel (ℓ, r) in one direction (wlog left to right). This time though,
 290 a user can choose to reject incoming transactions. We describe an algorithm (detailed in
 291 Appendix A as Algorithm 6) with competitive ratio $2 + \frac{\sqrt{5}-1}{2}$. We note that the competitive
 292 ratio for this setting is larger than the competitive ratio we achieve in the previous setting
 293 as OFF has a wider range of decisions.

294 Let us call a transaction of size x *big* if $x > Rx + f_2$ and *small* otherwise. We first observe
 295 that OFF in this setting always rejects big transactions.

296 **► Lemma 5.** *OFF rejects all big transactions in the unidirectional transaction stream with*
 297 *rejection.*

298 Thus, the strategy of OFF in this setting is to simply reject all big transactions. Moreover,
 299 if there are sufficiently many small transactions in the sequence to offset the cost of re-
 300 charging, OFF makes a single recharging action at the beginning of the sequence of size
 301 $\sum_{x \in X_t, x \text{ is small}} x$ for a cost of $f_1 + \sum_{x \in X_t, x \text{ is small}} x$.

302 The online algorithm performs $(1, \frac{\sqrt{5}-1}{2} f_1)$ -recharging and it accepts a transaction x if it
 303 has enough funds and x is small. The following theorem (with proof in Appendix C.6) states
 304 that ON is $(2 + \frac{\sqrt{5}-1}{2})$ -competitive in this problem setting.

305 **► Theorem 6.** *The algorithm described above is $(2 + \frac{\sqrt{5}-1}{2})$ -competitive in the unidirectional*
 306 *transaction stream with rejection sub problem.*

307 Before analysing the optimality of ON, we first observe, as a simple corollary of Theorem 4,
 308 that the lower bound of 2 also holds for this sub problem.

309 **► Corollary 7.** *There is no deterministic algorithm that is c -competitive for $c < 2$ in the*
 310 *unidirectional transaction stream with rejection sub problem.*

311 We conjecture that no other deterministic algorithm can perform better than ON in this
312 setting. Moreover, we sketch an approach to prove the conjecture in Appendix B.

313 ► **Conjecture 8.** *There is no deterministic algorithm that is c -competitive for $c < 2 + \frac{\sqrt{5}-1}{2}$
314 in the unidirectional transaction stream with rejection setting.*

315 6 Bidirectional transaction stream

316 In this section, we consider the most general problem setting, where for a given payment
317 channel (ℓ, r) , transactions stream along the channel (ℓ, r) in both directions. A user ℓ (resp.
318 r) can accept or reject incoming transactions that stream from left to right (resp. right to
319 left). Either user would incur a cost of $Rx + f_2$ for rejecting a transaction of size x . ℓ does
320 not need to take any action when encountering transactions that stream from right to left as
321 they simply increase the balance of ℓ in the channel (ℓ, r) . Both users can also decide at any
322 point to recharge their channel on-chain, or rebalance their channel off-chain.

323 Our main online algorithm ON for the bidirectional transaction stream setting is detailed
324 in Algorithm 4. For simplicity, we assume that $R = 0$ in the rejection cost. This means that
325 the cost of rejecting a single transaction of size x is simply f_2 , and rebalancing an amount of
326 x off-chain now only incurs a cost of Cf_2 . Our algorithm is run by both users on a payment
327 channel and is composed of three smaller algorithms: the first is a recharging algorithm to
328 determine when and how much to recharge the channel on-chain. The second algorithm
329 (Algorithm 2) decides whether to accept or reject new transactions and when to perform
330 off-chain rebalancing. The last algorithm (Algorithm 3) describes how to store the funds
331 received from the other user of the channel.

332 $(4 + 2\lceil \log C \rceil, f_1)$ -recharging. ON runs an on-chain recharging algorithm similar to
333 Algorithm 1 (see Algorithm 4 and Algorithm 4 in Algorithm 4) but with parameters $\gamma =$
334 $4 + 2\lceil \log C \rceil$ and $\delta = f_1$. Since we are in the bidirectional transaction stream setting, FUNDS
335 returns the amount of funds OFF has inside the entire channel (i.e. $b(\ell) + b(r)$) given a
336 transaction sequence.

337 Let us look at the period between the on-chain recharging instances of ON. From
338 Algorithm 4 in Algorithm 4, we know that ON ensures that it has more than $4 + 2\lceil \log C \rceil$
339 times more funds than OFF locked in the channel. These funds are distributed in the
340 following way: ON initialises $\lceil \log C \rceil + 2$ "buckets" on each end of the channel. We denote
341 set of left-side buckets as B^ℓ and it consists of $B_s^\ell, B_1^\ell, \dots, B_{\lceil \log C \rceil}^\ell, B_o^\ell$. Likewise, the set of
342 right-side buckets is B^r and it consists of $B_s^r, B_1^r, \dots, B_{\lceil \log C \rceil}^r, B_o^r$.

343 After recharging, users decide how to distribute funds in the channel, so the buckets
344 B_s^ℓ and B_s^r are filled with $2F_{tracker}$ funds. Buckets B_o^ℓ and B_o^r are empty (0 funds). Other
345 buckets contain $F_{tracker}$ funds.

346 Looking ahead, the funds in the i -th bucket on both sides are used to accept transactions
347 x with a size in the interval $[\frac{F_{tracker}}{2^i}, \frac{F_{tracker}}{2^{i-1}})$. The funds in B_s are used to accept transactions
348 with a size less than $\frac{F_{tracker}}{C}$. Finally, B_o stores excess funds coming from payments from the
349 other side when all other buckets are full.

350 *Transaction handling.* When a transaction arrives at the channel, based on the direction
351 of the transaction, either ℓ or r executes Algorithm 2 to decide whether to accept the
352 transaction. Wlog let us assume ℓ encounters transaction \vec{x} . If $\frac{F_{tracker}}{2^i} < x \leq \frac{F_{tracker}}{2^{i-1}}$ for
353 some $i \in [\lceil \log C \rceil]$ and B_i^ℓ has sufficient funds, the funds from B_i^ℓ are used to accept the
354 transaction. If B_i^ℓ lacks sufficient funds for accepting x , ℓ rejects x .

355 Now, we consider the case where $x \leq \frac{F_{tracker}}{C}$. If B_s^ℓ has sufficient funds, ℓ uses the funds
356 from B_s^ℓ to accept x . If B_s^ℓ has insufficient funds to accept x , ℓ performs off-chain rebalancing

357 with an amount such that after deducting x from B_s^ℓ , there would still be $2F_{tracker}$ funds left
 358 in B_s^ℓ . ℓ subsequently accepts x . The required funds for off-chain rebalancing are transferred
 359 from B_o^r and B_s^r (see Algorithm 2 and Algorithm 2 in Algorithm 2). Whenever $B_o^\ell > 0$ and
 360 some bucket in B^ℓ gets under its original capacity, funds are reallocated from B_o^ℓ to fill the
 361 bucket. Figure 3 in Appendix D depicts an example of how funds are used from different
 362 buckets to accept transactions.

■ **Algorithm 2** Decision on transaction

```

1 DECIDE( $F_{tracker}, x, B^{sdr}, B^{rcv}$ )
2    $Status \leftarrow \text{Accept}$ 
3   if  $\frac{F_{tracker}}{2^i} < x \leq \frac{F_{tracker}}{2^{i-1}}$  and  $x \leq B_i^{sdr}$  then
4      $\text{Accept } x$ 
5      $X \leftarrow \min(F_{tracker}, B_i^{sdr} - x + B_o^{sdr})$ 
6      $B_o^{sdr} \leftarrow \max(0, B_i^{sdr} - x + B_o^{sdr} - F_{tracker})$ 
7      $B_i^{sdr} \leftarrow X$ 
8   else if  $x_i \leq \frac{F_{tracker}}{C}$  and  $x \leq B_s^{sdr}$  then
9      $\text{Accept } x$ 
10     $X \leftarrow \min(2F_{tracker}, B_s^{sdr} - x + B_o^{sdr})$ 
11     $B_o^{sdr} \leftarrow \max(0, B_s^{sdr} - x + B_o^{sdr} - 2F_{tracker})$ 
12     $B_s^{sdr} \leftarrow X$ 
13  else if  $x_i \leq \frac{F_{tracker}}{C}$  and  $x > B_s^{sdr}$  then
14    Do off-chain rebalancing to fill  $B_s$  and pay  $f_2C$ .
15     $B_o^{rcv} \leftarrow B_o^{rcv} - (2F_{tracker} - B_s^{sdr})$ .
16     $B_s^{rcv} \leftarrow B_s^{rcv} - x$ .
17     $\text{Accept } x$ 
18     $B_s^{sdr} \leftarrow 2F_{tracker}$ .
19  else
20     $\text{Reject } x$ 
21     $Status \leftarrow \text{Reject}$ 
22  return ( $B^{sdr}, B^{rcv}, Status$ )

```

363 *Handling funds coming from the other side.* When a transaction x is accepted by wlog ℓ ,
 364 ON calls Algorithm 3 to distribute the transferred funds among r 's buckets in the following
 365 way: r first uses x to fill B_s^r up to its capacity of $2F_{tracker}$ (see Algorithm 3 in Algorithm 3). If
 366 there are still funds left, r refills the B_i^r buckets in descending order from $i = \lceil \log C \rceil$ to $i = 1$.
 367 Intuitively, the reason why buckets are refilled in descending order is due to our simplified
 368 cost model for this problem where we assume the cost of rejection for any transaction is
 369 f_2 . Thus, rejecting three small transactions size x costs thrice as much as rejecting a larger
 370 transaction of size $3x$. Finally, if there are still some funds left, they are added to B_o^r .

371 Our main theorem (with proof in Appendix C.7) shows that our main algorithm is
 372 $7 + 2\lceil \log C \rceil$ competitive.

373 ► **Theorem 9.** *Algorithm 4 is $7 + 2\lceil \log C \rceil$ competitive.*

374 Finally, we also analyse in the next lemma (with proof in Appendix C.8) how much
 375 funds ON needs to lock in the channel to have a chance to be c -competitive. We make the
 376 construction for A , the amount of funds that OFF locked in the channel. Observe that

■ **Algorithm 3** Handling funds coming from the other side

```

1 HANDLEFUNDS( $F_{tracker}, x, B$ )
2    $X \leftarrow \min(2F_{tracker}, B_s + x)$ 
3    $x \leftarrow \max(x + B_s - 2F_{tracker}, 0)$ 
4    $B_s \leftarrow X$ 
5   for  $i \in [\lceil \log C \rceil]$  in decreasing order do
6     if  $x > 0$  then
7        $X \leftarrow \min(F_{tracker}, B_i + x)$ 
8        $x \leftarrow \max(x + B_i - F_{tracker}, 0)$ 
9        $B_i \leftarrow X$ 
10   $B_o \leftarrow B_o + x$  return ( $B$ )

```

■ **Algorithm 4** Main algorithm

```

Initialise: left side buckets  $B^\ell$ 
Initialise: right side buckets  $B^r$ 
Initialise: tracker  $F_{tracker}, X \leftarrow 0, \emptyset$ 
1 for transaction  $x$  in order of arrival do
2   concatenate  $x$  to  $X$ 
3    $F'_{tracker} \leftarrow \text{FUNDS}(X)$ 
4   if  $F'_{tracker} > F_{tracker}$  then
5      $F_{tracker} \leftarrow F'_{tracker} + f_1$ 
6     recharge to  $2(2 + \lceil \log C \rceil)F_{tracker}$ 
7    $sdr, rcv \leftarrow \ell, r$ 
8   if  $x$  is from right to left then
9      $sdr, rcv \leftarrow r, \ell$ 
10   $B^{sdr}, B^{rcv}, Status \leftarrow \text{DECIDE}(F_{tracker}, x, B^{sdr}, B^{rcv})$ 
11  if  $Status == \text{Accept}$  then
12  |  $B^{rcv} \leftarrow \text{HANDLEFUNDS}(F_{tracker}, x, B^{rcv})$ 

```

377 OFF would rather reject transactions that have average size $> \frac{A}{C}$ than perform off-chain
378 rebalancing to accept them.

379 ► **Lemma 10.** *For any A , if ON's cost for rejection is at most c times OFF's cost for*
380 *rejection (for $c < \frac{\log C}{\log \log C}$), any deterministic ON needs to lock at least $\sigma = A \cdot \left(\frac{\frac{1}{c+1} \log C}{\log c+1} + 1 \right)$*
381 *funds in the channel.*

382 ► **Theorem 11.** *There is no deterministic c -competitive algorithm for $c \in o(\sqrt{\log C})$.*

383 **Proof.** From Lemma 10 for any A , ON needs $A \cdot \left(\frac{\frac{1}{c+1} \log C}{\log c+1} + 1 \right)$ funds to have its rejection
384 cost c -competitive. But ON also needs to lock some funds in the channel. The total cost is
385 then $c + A \left(\frac{\frac{1}{c+1} \log C}{\log c+1} + 1 \right)$, which is bigger than $\mathcal{O}(\sqrt{\log C})$. ◀

386 7 Empirical Evaluation

387 *Methodology.* We consider the performance of Algorithm 4 on randomly generated transaction
388 sequences. We compare it with the optimal offline algorithm OFF. Since computing the

389 optimal solution is NP-hard [18], we use dynamic programming to compute the cost (see the
390 algorithm in Appendix E).

391 *Average performance of ON.* We sample 50 random transaction sequences of length 50
392 each. In each sequence, transaction sizes are sampled independently from a folded Gaussian
393 with mean 0 and standard deviation 3, and then we assign its direction (left-to-right or
394 right-to-left) uniformly at random. Finally, we quantise the size of the transaction to the
395 closest integer. We run both OFF and ON on the generated sequences and compute five
396 important metrics.

397 We present our results in Table 2. As we can see from the cost of ON vs OFF in Table 2,
398 the competitive ratio is generally lower than the $7 + 2\lceil \log C \rceil$ bound as suggested by our
399 conservative worst-case analysis in Theorem 9. In addition, we notice that when we use some
400 heuristics to make further minor modifications to ON, we achieve even better performance.
401 We compare the average-case performance of ON and OFF with these modified algorithms
402 in Appendix F, and also on sequences sampled from different distributions. Finally, we also
403 perform an empirical case study of the Lightning Network which we present in Appendix G.

Param		OFF					ON				
C	f_2	Cost	$A(X)$	Accept rate	Off-chain rebalancing	Rechargings	Cost	$A(X)$	Accept rate	Off-chain rebalancing	Rechargings
2	0.5	15.02	6.4	0.78	0.8	1	63.3	44.26	0.50	0.9	2.18
8	0.5	15.21	6.38	0.77	0	1	87.79	69.06	0.50	0	2.04
2	2	23.6	14.26	0.95	5.36	1	127.02	100.2	0.91	0.38	5.86
8	2	24.5	13.9	0.92	0	1	184.32	156.6	0.9	0	5.84

■ **Table 2** Comparison between OFF and ON for $f_1 = 3$ and $R = 0$. $A(X)$ is the total amount of funds in the channel. “Accept rate” shows the fraction of transactions that were accepted. “Off-chain rebalancing” shows how much funds on was moved along the channel using off-chain rebalancing. “Rechargings” shows the number of rechargings performed. Note that OFF recharges only once.

404 8 Conclusion

405 This paper presents competitive strategies to maintain minimise cost while maximising
406 liquidity and transaction throughput in a payment channel. Our algorithms come with formal
407 worst-case guarantees, and also perform well in realistic scenarios in simulations.

408 We believe that our work opens several interesting avenues for future research. On the
409 theoretical front, it would be interesting to close the gap in the achievable competitive ratio,
410 and to explore the implications of our approach on other classic online admission control
411 problems. Furthermore, while in our work we have focused on deterministic algorithms, it
412 would be interesting to study the power of randomised approaches in this context, or to
413 consider different adversarial models.

414 — References —

- 415 1 Raiden network. <https://raiden.network/>, 2017.
- 416 2 James Aspnes, Yossi Azar, Amos Fiat, Serge Plotkin, and Orli Waarts. On-line routing of
417 virtual circuits with applications to load balancing and machine scheduling. *Journal of the
418 ACM (JACM)*, 44(3):486–504, 1997.
- 419 3 Georgia Avarikioti, Kenan Besic, Yuyi Wang, and Roger Wattenhofer. Online payment
420 network design. *CoRR*, abs/1908.00432, 2019. URL: <http://arxiv.org/abs/1908.00432>,
421 [arXiv:1908.00432](https://arxiv.org/abs/1908.00432).

- 422 4 Zeta Avarikioti, Lioba Heimbach, Yuyi Wang, and Roger Wattenhofer. Ride the lightning: The
423 game theory of payment channels. In Joseph Bonneau and Nadia Heninger, editors, *Financial*
424 *Cryptography and Data Security - 24th International Conference, FC 2020, Kota Kinabalu,*
425 *Malaysia, February 10-14, 2020 Revised Selected Papers*, volume 12059 of *Lecture Notes in*
426 *Computer Science*, pages 264–283. Springer, 2020. doi:10.1007/978-3-030-51280-4_15.
- 427 5 Zeta Avarikioti, Krzysztof Pietrzak, Iosif Salem, Stefan Schmid, Samarth Tiwari, and Michelle
428 Yeo. HIDE & SEEK: privacy-preserving rebalancing on payment channel networks. *CoRR*,
429 abs/2110.08848, 2021. URL: <https://arxiv.org/abs/2110.08848>, arXiv:2110.08848.
- 430 6 Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. cambridge
431 university press, 2005.
- 432 7 Christian Decker. Lightning network research; topology, datasets. [https://github.com/](https://github.com/lnresearch/topology)
433 [lnresearch/topology](https://github.com/lnresearch/topology). Accessed: 2022-04-01. doi:10.5281/zenodo.4088530.
- 434 8 Oguzhan Ersoy, Stefanie Roos, and Zekeriya Erkin. How to profit from payments channels,
435 2019. arXiv:1911.08803.
- 436 9 MohammadAmin Fazli, Seyed Moeen Nehzati, and MohammadAmin Salarkia. Building stable
437 off-chain payment networks. *CoRR*, abs/2107.03367, 2021. URL: [https://arxiv.org/abs/](https://arxiv.org/abs/2107.03367)
438 [2107.03367](https://arxiv.org/abs/2107.03367), arXiv:2107.03367.
- 439 10 Rami Khalil and Arthur Gervais. Revive: Rebalancing off-blockchain payment networks. *IACR*
440 *Cryptol. ePrint Arch.*, page 823, 2017. URL: <http://eprint.iacr.org/2017/823>.
- 441 11 Peng Li, Toshiaki Miyazaki, and Wanlei Zhou. Secure balance planning of off-blockchain
442 payment channel networks. In *39th IEEE Conference on Computer Communications,*
443 *INFOCOM 2020, Toronto, ON, Canada, July 6-9, 2020*, pages 1728–1737. IEEE, 2020.
444 doi:10.1109/INFOCOM41043.2020.9155375.
- 445 12 Tamás Lukovszki and Stefan Schmid. Online admission control and embedding of service
446 chains. In *International Colloquium on Structural Information and Communication Complexity*,
447 pages 104–118. Springer, 2015.
- 448 13 Giulio Malavolta, Pedro Moreno-Sanchez, Aniket Kate, and Matteo Maffei. Silentwhispers:
449 Enforcing security and privacy in decentralized credit networks. In *24th Annual Network and*
450 *Distributed System Security Symposium, NDSS 2017, San Diego, California, USA, February*
451 *26 - March 1, 2017*. The Internet Society, 2017.
- 452 14 Rene Pickhardt and Mariusz Nowostawski. Imbalance measure and proactive channel rebalan-
453 cing algorithm for the lightning network. In *IEEE International Conference on Blockchain*
454 *and Cryptocurrency, ICBC 2020, Toronto, ON, Canada, May 2-6, 2020*, pages 1–5. IEEE,
455 2020. doi:10.1109/ICBC48266.2020.9169456.
- 456 15 Joseph Poon and Thaddeus Dryja. The bitcoin lightning network: Scalable off-chain instant
457 payments. <https://lightning.network/lightning-network-paper.pdf>, 2015.
- 458 16 Pavel Prihodko, S. N. Zhigulin, Mykola Sahnno, A B Ostrovskiy, and Olaoluwa Osuntokun.
459 Flare : An approach to routing in lightning network white paper. 2016.
- 460 17 Stefanie Roos, Pedro Moreno-Sanchez, Aniket Kate, and Ian Goldberg. Settling payments
461 fast and private: Efficient decentralized routing for path-based transactions. In *25th Annual*
462 *Network and Distributed System Security Symposium, NDSS 2018, San Diego, California,*
463 *USA, February 18-21, 2018*. The Internet Society, 2018.
- 464 18 Stefan Schmid, Jakub Svoboda, and Michelle Yeo. Weighted packet selection for rechargeable
465 links: Complexity and approximation, 2022. URL: <https://arxiv.org/abs/2204.13459>,
466 doi:10.48550/ARXIV.2204.13459.
- 467 19 Vibhaalakshmi Sivaraman, Shaileshh Bojja Venkatakrishnan, Kathleen Ruan, Parimarjan
468 Negi, Lei Yang, Radhika Mittal, Giulia C. Fanti, and Mohammad Alizadeh. High throughput
469 cryptocurrency routing in payment channel networks. In Ranjita Bhagwan and George Porter,
470 editors, *17th USENIX Symposium on Networked Systems Design and Implementation, NSDI*
471 *2020, Santa Clara, CA, USA, February 25-27, 2020*, pages 777–796. USENIX Association,
472 2020.

■ **Algorithm 5** Unidirectional transaction stream without rejection

```

Initialise: tracker  $F_{tracker}, X \leftarrow 0, \emptyset$ 
Initialise: balance  $b = 0$ 
1 for transaction  $x$  in order of arrival do
2   concatenate  $x$  to  $X$ 
3    $F'_{tracker} \leftarrow \text{FUNDS}(X)$ 
4   if  $F'_{tracker} > F_{tracker}$  then
5      $F_{tracker} \leftarrow F'_{tracker} + f_1$ 
6     recharge to  $F_{tracker}$ 
7   Accept  $x$ 

```

■ **Algorithm 6** Unidirectional transaction stream with rejection

```

Initialise: tracker  $F_{tracker}, X \leftarrow 0, \emptyset$ 
Initialise: balance  $b = 0$ 
1 for transaction  $x$  in order of arrival do
2   concatenate  $x$  to  $X$ 
3    $F'_{tracker} \leftarrow \text{FUNDS}(X)$ 
4   if  $F'_{tracker} > F_{tracker}$  then
5      $F_{tracker} \leftarrow F'_{tracker} + \frac{\sqrt{5}-1}{2} f_1$ 
6     recharge to  $F_{tracker}$ 
7   if  $b \geq x$  and  $x$  is small then
8     Accept  $x$ 
9   else
10    Reject  $x$ 

```

473 20 Yuup van Engelshoven and Stefanie Roos. The merchant: Avoiding payment channel depletion
474 through incentives. *2021 IEEE International Conference on Decentralized Applications and*
475 *Infrastructures (DAPPS)*, pages 59–68, 2021.

476 **A Algorithms**

477 **B Optimality of deterministic algorithms in a unidirectional stream**
478 **with rejection**

479 Sketch proof of Conjecture 8.

480 **Sketch.** The best ON algorithm needs to recharge the channel when OFF does. If it
481 recharges the channel later, it incurs cost that OFF is not incurring, so the competitive ratio
482 worsens. If it recharges sooner, there exists a sequence that either forces OFF to waste funds
483 or incur a big cost.

484 After OFF recharges, it can reconsider and accept previously rejected transactions, but
485 ON needs to reject them. Now, the situation is similar as in the case without rejection.
486 ON needs to recharge, but it already paid for some rejections whereas OFF pays only for
487 recharging and accepting very small transactions.

488 ON disadvantaged in this way cannot achieve a better competitive ratio than $2 + \frac{\sqrt{5}-1}{2}$ ◀

C Omitted proofs

C.1 Proof of Lemma 1

Proof. The first part of the claim follows from the fact that the moment $A(X_i) > F_{tracker}$ for some i , $F_{tracker}$ gets updated to $A(X_i) + \delta > A(X_i)$ and ON recharges the channel to $\gamma F_{tracker} > \gamma A(X_i)$.

For the second part of the claim, we note that the cost incurred by ON is simply the total amount of funds added to the channel with an additional cost of f_1 each time ON recharges the channel on-chain. The amount of funds locked in the channel for ON is always at most $A_t + \delta$ and the times when ON recharges the channel occurs whenever OFF increases its funds by an amount of at least δ . Thus, the number of rechargings for ON that can occur is at most $\lceil \frac{A_t}{\delta} \rceil$ with a cost of f_1 for each recharging instance. The total cost incurred by ON is therefore $\gamma(A_t + \delta) + f_1 \cdot \lceil \frac{A_t}{\delta} \rceil$. ◀

C.2 Proof of Lemma 2

Proof. We first note that the sequence of costs for OFF is monotonically increasing, i.e. $\text{COST}_{\text{OFF}}(X_t) \leq \text{COST}_{\text{OFF}}(X_{t+1})$. This comes from the fact that any action of OFF at step i of the sequence can only increase its cost (i.e. either rejecting x_{i+1} or recharging the channel and then accepting x_{i+1}), or it does not change the cost at all (i.e. by accepting x_{i+1} without recharging).

Since $A_t > 0$, we know that OFF recharged on-chain at some point to an amount A_t for a recharging cost of $A_t + f_1$. Since the sequence of costs for OFF is monotonically increasing, $\text{COST}_{\text{OFF}}(X_t) \geq A_t + f_1$. ◀

C.3 Proof of Theorem 3

Proof. From Lemma 1, setting $\gamma = 1$ and $\delta = f_1$ gives ON a cost of at most $A_t + f_1 + f_1 \cdot \lceil \frac{A_t}{f_1} \rceil$. Since $f_1 \cdot \lceil \frac{A_t}{f_1} \rceil \leq f_1 \cdot (\frac{A_t}{f_1} + 1) = A_t + f_1$, the cost of ON is at most $2(A_t + f_1)$. From Lemma 2, we know that the cost of OFF is at least $A_t + f_1$. Thus, ON is 2-competitive. ◀

C.4 Proof of Theorem 4

Proof. We prove the theorem by contradiction. For the sake of contradiction, suppose that there exists a c -competitive algorithm ON for $c = 2 - \varepsilon$ for some $\varepsilon > 0$. Consider the following sequence of transactions: $\frac{\varepsilon}{3}, \frac{\varepsilon}{3}, \frac{\varepsilon}{3}, \dots$. We note that when the sequence of transactions is of length k , the cost of OFF is $f_1 + k \cdot \frac{\varepsilon}{3}$ as the optimal solution is to recharge the channel at the start of the sequence to the total sum of the transactions in the sequence.

For ON to remain $(2 - \varepsilon)$ -competitive after processing the first transaction, ON locked at most $f_1 - \varepsilon f_1 + \frac{\varepsilon}{3}$ in the channel ($\text{COST}_{\text{ON}}(X_1) = 2f_1 - \varepsilon f_1 + 2\frac{\varepsilon}{3}$).

We generalize the above idea and show that either ON has always smaller amount than $f_1 - \frac{\varepsilon}{3}$ in the channel or at some point it has at least $f_1 - \frac{\varepsilon}{3}$. In both cases, we derive a contradiction to the $(2 - \varepsilon)$ competitive ratio of ON.

First, suppose that ON always recharges to at most $f_1 - \varepsilon'$ for some $\varepsilon' > 0$. Then after t transactions, the number of rechargings is at least $\lceil \frac{t \frac{\varepsilon}{3}}{f_1 - \varepsilon'} \rceil$. So $\text{COST}_{\text{ON}}(X_t) \geq t \cdot \frac{\varepsilon}{3} + f_1 \lceil \frac{t \frac{\varepsilon}{3}}{f_1 - \varepsilon'} \rceil$. Setting $t = \frac{3k(f_1 - \varepsilon')}{\varepsilon}$ for some k gives $\text{COST}_{\text{OFF}}(X_t) = f_1 + k(f_1 - \varepsilon')$ and $\text{COST}_{\text{ON}}(X_t) = k(f_1 - \varepsilon') + k f_1$, but since $\lim_{k \rightarrow \infty} \frac{2k f_1 - k \varepsilon'}{(k+1)f_1 - k \varepsilon'} = \frac{2f_1 - \varepsilon'}{f_1 - \varepsilon'}$ for any ε' , then the competitive ratio is at least 2.

530 Now, suppose that t is the first time that after processing a transaction, ON has at least
 531 $f_1 - \frac{\varepsilon}{3}$ locked in the channel. At time t , $\text{COST}_{\text{OFF}}(X_t) = f_1 + t\frac{\varepsilon}{3}$. Cost of ON is $f_1 + t\frac{\varepsilon}{3}$
 532 for funds locked in the channel plus any additional recharging cost. But since it is the first
 533 time ON recharged by more than f_1 , the cost for recharging is $f_1 \lceil \frac{t\frac{\varepsilon}{3}}{f_1 - \varepsilon'} \rceil \geq f_1 + \frac{t\varepsilon}{3}$ for some
 534 other positive ε' . So again, $\text{COST}_{\text{ON}}(X_t) \geq t\frac{\varepsilon}{3} + f_1 + t\frac{\varepsilon}{3} + f_1 = 2(f_1 + t\frac{\varepsilon}{3})$ which is twice of
 535 $\text{COST}_{\text{OFF}}(X_t)$.

536 In both cases the cost of ON is at least twice that of OFF which contradicts the
 537 assumption that ON is $(2 - \varepsilon)$ -competitive. ◀

538 C.5 Proof of Lemma 5

539 **Proof.** Accepting a transaction x incurs a cost of x for increasing funds. Rejecting a
 540 transaction x incurs a cost of $Rx + f_2$. So any big transaction should be rejected. ◀

541 C.6 Proof of Theorem 6

542 **Proof.** From Lemma 5, OFF rejects big transactions. Thus, ON should also reject these
 543 transactions.

544 While $A_t = 0$, both OFF and ON reject all transactions in the sequence and both incur
 545 the same cost. The moment $A_t > 0$, we know that OFF recharged with an amount at least
 546 A_t to accept all small transactions in the sequence. Thus $\text{COST}_{\text{OFF}}(X_t) \geq A_t + f_1$.

547 At this time ON would have rejected the small transactions in the sequence for at most
 548 a cost of $A_t + f_1$ together with some additional recharging cost. From Lemma 1, we know
 549 that the recharging cost for ON is at most

$$550 \quad A_t + \frac{\sqrt{5}-1}{2}f_1 + \left\lceil \frac{A_t}{\frac{\sqrt{5}-1}{2}f_1} \right\rceil f_1 \leq A_t + \frac{\sqrt{5}-1}{2}f_1 + \frac{A_t}{\frac{\sqrt{5}-1}{2}} + f_1.$$

551 Summing up both costs, we get

$$552 \quad \text{COST}_{\text{ON}}(X_t) \leq A_t + f_1 + A_t + \frac{\sqrt{5}-1}{2}f_1 + \frac{A_t}{\frac{\sqrt{5}-1}{2}} + f_1$$

$$553 \quad = \left(2 + \frac{\sqrt{5}-1}{2}\right) \text{COST}_{\text{OFF}}(X_t)$$

554

555 ◀

556 C.7 Proof of our main Theorem

557 Here, we detail the full proof of Theorem 9

558 **Proof.** We know that for any i , $\text{COST}_{\text{OFF}}(X_i) \geq \text{COST}_{\text{OFF}}(X_{i-1})$. From Lemma 1 and
 559 Lemma 2, we know that cost of ON for recharging (in Algorithm 4) is at most $(5 +$
 560 $2\lceil \log C \rceil)\text{COST}_{\text{OFF}}(X_t)$. Let t_1 and t_2 (with $t_2 > t_1$) be the two consecutive times ON
 561 recharges, then we show that the cost of ON for rebalancing and rejection is smaller than
 562 $2(\text{COST}_{\text{OFF}}(X_{t_2}) - \text{COST}_{\text{OFF}}(X_{t_1}))$. Then $\text{COST}_{\text{ON}}(X_t) \leq (7 + 2\lceil \log C \rceil)\text{COST}_{\text{OFF}}(X_t)$.

563 For every strategy of OFF and any two consecutive recharging times t_1 and t_2 , we show
 564 that the rebalancing and rejection cost of ON between times t_1 and t_2 is at most twice that
 565 of OFF as defined by the strategy. Having the strategy of OFF, we split the time between
 566 rechargings even further, into epochs; we will show that the competitive ratio of 2 holds for
 567 every epoch.

568 The left epoch starts with the first transaction that makes some bucket in B^ℓ non-full
 569 (smaller than the original amount); the left epoch ends either before ON recharges, or $B_o^\ell > 0$.
 570 In a left epoch, every transaction from the right side is accepted; non-fullness of some buckets
 571 on one side means $B_o > 0$ on the other side. The right epoch is defined similarly, but since
 572 the epochs are disjoint, we can prove the statement for a left epoch only.

573 For transactions below $\frac{F_{tracker}}{C}$, we argue that the cost of ON is at most the cost of OFF.
 574 ON accepts everything, so ON pays only for rebalancing. OFF either rebalances too, in
 575 which case the cost is the same as ON; or it rejected some transactions. Since ON starts with
 576 $2F_{tracker}$ funds in B_s and refills the bucket with the highest priority, this means OFF rejected
 577 some transactions summing to at least $F_{tracker}$. There are at least C of them, so OFF's cost
 578 is also above Cf_2 . If there is a counterexample containing a small transaction that OFF
 579 rejects, then we can modify it to a counterexample where the transaction is increased to
 580 $\frac{F_{tracker}}{C}$. So we can show the ratio in the case that no small transactions are coming.


581 Now that we have the strategy for OFF (decisions before rebalancing), we define some
 582 variables that track the competitive ratio. We will look at incoming transactions, and prove
 583 that the competitive ratio is always below 2. We say that a transaction x belongs to a bucket
 584 B_i if $\frac{F_{tracker}}{2^i} < x \leq \frac{F_{tracker}}{2^{i-1}}$. A transaction is red if it is rejected by ON and accepted by OFF
 585 and it is blue if it is accepted by ON and rejected by OFF. Let ρ_i (β_i) be the number of red
 586 (blue) transactions in the bucket B_i . We can disregard transactions for which ON and OFF
 587 make the same decision. If the transaction is rejected by both, it improves the ratio. If the
 588 transaction is accepted by both, it can be simulated by decreasing $F_{tracker}$.

589 We prove by induction that $\sum_{k \leq j} \rho_k \leq 2 \sum_{k \leq j} \beta_k$ for all j . We show that if the equality
 590 holds and OFF has enough funds to accept incoming transactions, ON accepts too.

591 Let us examine $j = \lceil \log C \rceil$. We know that the ratio between any two transaction sizes
 592 in $B_{\lceil \log C \rceil}$ is less than 2, so any two red transactions are bigger than one blue. Moreover,
 593 all funds that arrived from the right side were put into $B_{\lceil \log C \rceil}$ (if it is not full). So if
 594 $\rho_{\lceil \log C \rceil} = 2\beta_{\lceil \log C \rceil}$, ON has at least the amount of funds in $B_{\lceil \log C \rceil}$ OFF has. To continue
 595 the induction for buckets with smaller indices, we reassign some red or blue transactions to
 596 different buckets. If $\rho_{\lceil \log C \rceil} < 2\beta_{\lceil \log C \rceil}$, we move at most one red and some blue transactions
 597 (in decreasing order of size) to $B_{\lceil \log C \rceil - 1}$, stopping just before $\rho_{\lceil \log C \rceil} \geq 2\beta_{\lceil \log C \rceil}$.

598 For general j , we know that, due to the reassignment, in every bucket smaller than j ,
 599 ON rejected exactly twice the number of transactions OFF did. Moreover, OFF needs to
 600 use at least the same amount of funds to accept red transactions compared to funds needed
 601 by ON to accept blue ones. Now, in the bucket B_j holds $\rho_j = 2\beta_j$. Again, we pair every
 602 two red transactions to one blue, such that the sum of red is bigger than blue. Before the
 603 reassignment, the ratio between any two transactions is at most 2. The reassignment (if
 604 occurred) moved at most one red and at least one blue that is smaller than any original
 605 transaction in the bucket, so we can pair the moved red to moved blue. In transactions in
 606 buckets in j and bigger, OFF used more funds than ON. Any funds that arrived from the
 607 right side were put into some bucket in j or below, so if OFF has enough funds to accept,
 608 ON has too.

609 The same argument holds for a right-epoch, and we note that epochs are disjoint and
 610 cover the entire transaction sequence between times t_1 and t_2 . Since we chose the consecutive
 611 recharging times t_1 and t_2 arbitrarily, the rebalancing and rejection cost of ON between any
 612 two consecutive rechargings is at most twice that of OFF within the same period. Therefore,
 613 Algorithm 4 is $7 + 2\lceil \log C \rceil$ competitive. ◀



■ **Figure 3** An example of how funds are transferred across a payment channel and how buckets are refilled. Both ℓ and r start with full buckets. The first transaction is in the left-to-right direction and is transferred using funds from B_2^ℓ to B_o^r . The second transaction is in the right-to-left direction and is small, thus funds from B_s^r are used. B_s^r is immediately refilled using funds from B_o^r . The third transaction is in the left-to-right direction and uses funds from B_1^ℓ .

614 C.8 Proof of Lemma 10

615 **Proof.** We describe an epoch: OFF starts with A funds left, then some transactions are sent
 616 from left to right and finally one transaction of size A is sent right to left. If the funds in the
 617 channel of ON is smaller than σ , then the cost of ON is more than c times that of OFF.

618 One epoch consists of at most $\frac{\log C}{\log c+1} + 1$ phases. In phase i (starting from $i = 0$), there
 619 are $(c+1)^i$ transactions of size $\frac{A}{(c+1)^i}$. OFF always accepts all transactions in the latest
 620 phase. If at the end of any phase, the cost of ON is more than c times of OFF, then a
 621 transaction of size A is sent back and another epoch starts. Observe that for $c < \frac{\log C}{\log \log C}$,
 622 after rebalancing the epoch ends too (because the cost is c times bigger). We can assume
 623 this cannot happen, ON does not perform off-chain rebalancing.

624 We compute how much funds ON needs to stay within the competitive ratio until
 625 the last phase (where transactions of sizes $\frac{A}{C}$ are sent). After phase i , OFF accepted
 626 $(c+1)^i$ transactions and rejected $\frac{(c+1)^i - 1}{c}$, so OFF can reject up to $(c+1)^i - 1$ transactions
 627 among $\frac{(c+1)^{i+1} - 1}{c}$. So ON has to accept at least $\frac{(c+1)^{i+1} - 1 - c(c+1)^i + c}{c} = (c+1) \frac{(c+1)^{i-1} - 1}{c}$
 628 transactions.

629 The size of transactions is decreasing, so optimally, ON accepts transactions when they
 630 are needed. So in phase $i+1$ it needs to accept $(c+1) \frac{(c+1)^i - 1}{c} - (c+1) \frac{(c+1)^{i-1} - 1}{c} = (c+1)^i$
 631 transactions. Of course, it needs to accept the transaction in the phase 0.

632 The cost of transactions accepted by ON in phase $i > 0$ is $(c+1)^{i-1} \frac{A}{(c+1)^i} = \frac{A}{c+1}$. To
 633 maintain the competitive ratio ON needs to accept transactions worth $A + A \frac{\log C}{\log c+1}$ in total.
 634 Independently of the cost of OFF and ON before, there can be one epoch after another
 635 where the ratio is worse than c , so at the end, the ratio would be above c . ◀

D Helpful diagrams

E Computing the cost of OFF using dynamic programming

In this section, we describe a dynamic programming algorithm ON that solves the main problem. We assume that the size of transactions is integer (moreover the sum of transactions should be small).

Let $\text{COST}_i(F_\ell, F_r)$ be the minimum cost for rejecting and off-chain rebalancing in processing sequence X_i that ends with $b(\ell) = F_\ell$ and $b(r) = F_r$. (For values of F_ℓ and F_r smaller than 0, we define it to be ∞). $\text{COST}_i(F_\ell, F_r)$ can be derived from COST_{i-1} given the decision on the i 'th transaction.

Let us assume wlog that i 'th transaction is from ℓ to r . OFF has three choices when encountering x_i . The first option is to reject x_i , the the cost is $A_1 = \text{COST}_{i-1}(F_\ell, F_r) + Rx_i + f_2$. The second option is to accept x_i which gives cost $A_2 = \text{COST}_{i-1}(F_\ell + x_i, F_r - x_i)$. The last option is to off-chain rebalance before x_i and then accept x_i . Note that any off-chain rebalancing before rejecting or accepting (while having enough funds) can be postponed. This gives cost $A_3 = \min_{a \leq F_\ell + x_i} \text{COST}_{i-1}(a, F_r + F_\ell - a) + C \cdot R(F_\ell + x_i - a) + Cf_2$. OFF chooses the best option, that means $\text{COST}_i(F_\ell, F_r) = \min \{A_1, A_2, A_3\}$

We handle right to left transaction in the same way.

Given the previous, ON computes COST_t for all valid pairs (F_ℓ, F_r) and the final cost is

$$\text{COST}_{\text{OFF}}(X_i) = \min_{F_\ell, F_r} \text{COST}_t(F_\ell, F_r) + F_\ell + F_r + f_1$$

To bound the time complexity of ON, we observe some bounds for $S = F_\ell^* + F_r^*$, where F_ℓ^* and F_r^* are the values of F_ℓ and F_r achieving the minimal cost. Observe that $S \leq \sum_{i \leq t} x_i$, it is not worth to have more money than the sum of the trasactions. We can strenghten the inequality and instead of $\sum_{i \leq t} x_i$, we can compute the minimal amount needed to accept every transaction. The other option is to reject everything, so we know that $f_1 + S \leq \sum_{i \leq t} Rx_i + f_2$.

Now we can prove the theorem about the described algorithm ON.

► **Theorem 12.** *ON computes the optimal cost in time $\mathcal{O}(tS^3)$, where S is the bound on the maximal funds in the channel and t the number of transactions.*

Proof. In the dynamic programming, we take into account all possible decisions ON can make, by this, correctness follows.

The algorithm ON tries all possible amounts between 1 and S and starting distributions. There are S^2 of them. While computing one value, it needs to look at at most S precomputed values. And it needs to do it at most t times. ◀

Using dynamic programming for calculating OFF has two advantages. First, we can easily recover the decisions of OFF. Secondly, dynamic programming provides us with optimum solution for all subsequences of X_t . This is useful for implementing Algorithm 4.

F Heuristics to improve the average-case performance of our online algorithm

We notice in our experiments that ON seems to overcharge the channel. This is most noticeable when we observe the effect of C on the performance of ON. From Table 2, increasing C in a range of medium (not too small) values does not change OFF's cost

noticeably. In contrast, both the average cost and total amount of locked funds of ON grows with C . This is due to the fact that ON uses $(4 + 2\lceil \log C \rceil, f_1)$ -recharging to ensure that it always has significantly more funds than OFF, even though a big fraction of these funds remain unspent. ON is also limited by the fact that it does not borrow funds from other buckets when a bucket is depleted. For instance, ON always charges B_s to $2F_{tracker}$ and only uses these funds to accept transactions that are smaller than $\frac{F_{tracker}}{C}$. Thus, as C increases, the number of transactions that fall into the B_s bucket decreases and the funds in B_s remain unspent.

These observations motivate us to design a less pessimistic version of ON that we expect will perform better than ON. We introduce ON-I which is a slightly altered version of ON: ON-I follows the $(\lceil \log C \rceil, f_1)$ -recharging algorithm and does not divide the funds into separate buckets. Instead, ON-I accepts all the transactions smaller than $F_{tracker}$ as long as it has funds. Otherwise, if the transaction is small ($< \frac{F_{tracker}}{C}$) it off-chain rebalances to fill the bucket and accepts the transaction. Similar to ON, ON-I rejects a transaction if it is larger than $\frac{F_{tracker}}{C}$.

Our empirical results in Table 3 confirms that the average cost of ON-I is significantly smaller than ON. The acceptance rate of ON-I is slightly smaller than ON, which is expected as ON-I does not have separate funds for each range of transactions (as defined by the buckets), and as a result might miss some transactions. We observe that ON-I performs more off-chain rebalancings compared to ON on average because ON-I does not reserve separate funds for small transactions. However, one issue with both ON and ON-I is that they recharge the channel too often (as soon as $F_{tracker} < A(X_t)$). As can be seen from Table 2 ($f_2 = 2$), both algorithms perform more than 5 rechargings on average for transaction sequences of length 50. This increases the cost of both algorithms significantly as each recharging instance incurs a cost of at least f_1 .

We thus design another version of ON-I to address the aforementioned problem. ON-II works exactly as ON-I except that it does not recharge the channel as frequently as ON-I does. ON-II only recharges the channel if $\alpha \cdot F_{tracker} < A(X)$, where $\alpha > 1$ is some constant that controls the how often the algorithm recharges the channel and can be fine-tuned empirically based on f_1 and f_2 . If we set $\alpha = 1$, ON-II becomes equivalent to ON-I and has higher acceptance rate. This is favorable when f_2 is large and f_1 is small. Conversely, by increasing α , ON-II recharges the channel less frequently but the acceptance rate falls. This is favorable when f_1 is large and f_2 is small. In our experiments, we observe that for the case $f_2 = 2, C = 2$, when all the other parameters are as Table 2, $\alpha = 2$ yields the lowest average cost. Thus, in our evaluation of ON-II, we use $\alpha = 2$ and from Table 2 we note that this choice of α halves the number of rechargings compared to ON-I, which consequently leads to lower average cost. Additionally, we note that the total amount of funds in the channel of ON-II is close to OFF.

F.0.0.1 Varying the distribution of the generated sequences.

We also evaluate how the performance of our algorithms is affected by the variance of the transaction size. We sample 50 sequences each of length 50 with each transaction in the sequence independently sampled from the folded normal distribution with mean 0 and standard deviation σ , for a range of σ values across $[3, 20]$. We then observe the cost of ON and OFF. As can be seen from in Figure 4a, the cost of both algorithms rises as σ increases. This is due to the fact that increasing the variance of the sampled transactions reduces the probability of getting a similarly sized transaction coming from the other side, thus increasing the speed at which the balance on one side gets depleted. We note, however,

724 that Figure 4a shows that even for large values of σ , ON's average cost remains a lot smaller
725 than the worst case upper bound of $7 + 2\lceil \log C \rceil$.

726 Another factor we look at is how the asymmetry of the transaction flow along a channel
727 can affect the performance of our algorithms. To do so, we generate 50 sequences each of
728 length 50, and sample the size of each transaction from a folded normal distribution with
729 mean 0 and standard deviation 3. We then sample the direction of the transactions according
730 to a Bernoulli distribution with parameter p , where p represents the probability of sampling
731 a left-to-right transaction. We see from Figure 4b that the cost of all algorithms decrease
732 as p increases from 0 to 0.5. As p increases from 0.5 to 1, the cost function increases again.
733 This conforms to our intuition that extremely asymmetric sequences are harder to handle as
734 the lack of sufficiently many transactions from one side just increases the speed at which the
735 balance on the other side gets depleted.

Param		OFF					ON				
C	f_2	Cost	$A(X)$	Accept rate	Off-chain re-balancing	Rechargings	Cost	$A(X)$	Accept rate	Off-chain re-balancing	Rechargings
2	0.5	15.02	6.4	0.78	0.8	1	63.3	44.26	0.50	0.9	2.18
8	0.5	15.21	6.38	0.77	0	1	87.79	69.06	0.50	0	2.04
2	2	23.6	14.26	0.95	5.36	1	127.02	100.2	0.91	0.38	5.86
8	2	24.5	13.9	0.92	0	1	184.32	156.6	0.9	0	5.84

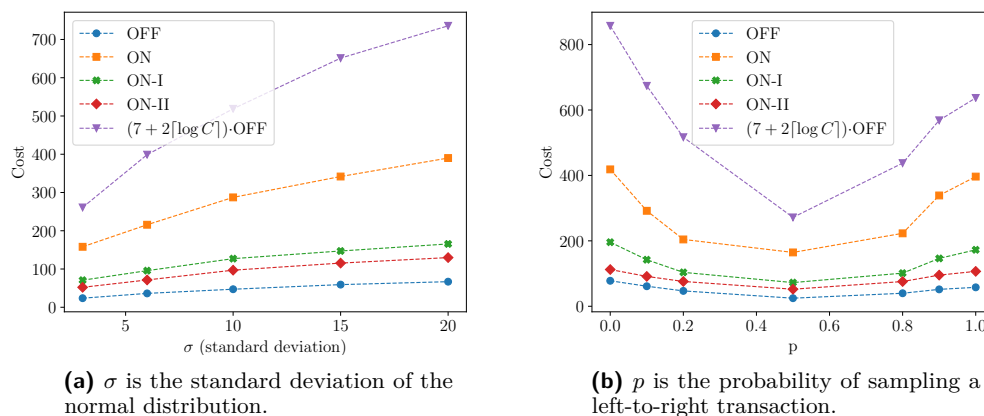
Param		ON-I					ON-II				
C	f_2	Cost	$A(X)$	Accept rate	Off-chain re-balancing	Rechargings	Cost	$A(X)$	Accept rate	Off-chain re-balancing	Rechargings
2	0.5	30.74	6.86	0.44	11.18	2.18	26.52	5.56	0.41	10.22	1.2
8	0.5	39.98	19.86	0.48	0	2.04	32.94	15.9	0.46	0	1.18
2	2	66.16	14.42	0.84	25.48	5.86	60.38	11.3	0.78	27.16	2.2
8	2	81.35	42.72	0.89	1.5	5.84	58.38	33.3	0.83	1.56	2.16

■ **Table 3** Comparison between the performance of OFF, ON, ON-I and ON-II on randomly generated transaction streams. The result is averaged over 50 sequences each of length 50. The size of each transaction is independently sampled from the folded normal distribution with mean 0 and standard deviation 3, then quantised to the closest integer. We set $f_1 = 3$ and $R = 0$. $A(X)$ is the total amount of funds in the channel from recharging the channel. "Accept rate" shows the average fraction of transactions that were accepted. "Off-chain rebalancing" shows how much funds on average was moved along the channel using off-chain rebalancing. "Rechargings" shows the average number of rechargings performed. Note that since OFF knows the entire sequence in advance, it only recharges the channel once at the beginning of each sequence.

736 G Case study: Lightning Network

737 We also ran a case study of the Lightning network. We first run our experiments with
738 realistic parameters taken from Lightning Network data. In the Lightning Network, f_1 is the
739 on-chain transaction fee (roughly 1000 satoshi) which is a lot larger than f_2 , the base fee one
740 receives when forwarding a payment (around 1 satoshi).

741 From analysis of the Lightning Network (We use snapshot from September 2021) [7], we
742 know that the average cycle length is 4.15 (after excluding roughly 10% of vertices that are
743 not part of any cycle). That means the value C in Theorem 9 is just slightly above 4. Details
744 are in Table 4.



■ **Figure 4** Average cost of our algorithms over 50 randomly generated transaction streams each of length 50. In both figures the size of each transaction is sampled from the folded normal distribution with mean 0. The standard deviation of the normal distribution is fixed to 3 in the right figure, however in the left figure the standard deviation is varying on the x-Axis. We use the parameters $f_1 = 3, f_2 = 2, R = 0, C = 4, \alpha = 2$.

Cycle length	≤ 4	5	6	7	N.A.
Frequency	49,424(77.44%)	7,758(12.16%)	469(0.73%)	12(0.02%)	6,157(9.65%)

■ **Table 4** Frequency of the length of shortest cycle between all users in the Lightning Network. The last column shows the frequency of channels that are not part of any cycle (N.A. not applicable) The average cycle length is 4.15