

1 Condorcet Attack Against Fair Transaction 2 Ordering

3 **Mohammad Amin Vafadar** ✉ 

4 University of Alberta, Canada

5 **Majid Khabbazian** ✉ 

6 University of Alberta, Canada

7 — Abstract —

8 We introduce the *Condorcet attack*, a new threat to fair transaction ordering. Specifically, the attack
9 undermines batch-order-fairness, the strongest notion of transaction fair ordering proposed to date.
10 The batch-order-fairness guarantees that a transaction \mathbf{tx} is ordered before \mathbf{tx}' if a majority of nodes
11 in the system receive \mathbf{tx} before \mathbf{tx}' ; the only exception (due to an impossibility result) is when \mathbf{tx}
12 and \mathbf{tx}' fall into a so-called “Condorcet cycle”. When this happens, \mathbf{tx} and \mathbf{tx}' along with other
13 transactions within the cycle are placed in a batch, and any unfairness inside a batch is ignored.

14 In the Condorcet attack, an adversary attempts to undermine the system’s fairness by imposing
15 Condorcet cycles to the system. In this work, we show that the adversary can indeed impose
16 a Condorcet cycle by submitting as few as two otherwise legitimate transactions to the system.
17 Remarkably, the adversary (e.g., a malicious client) can achieve this even when all the nodes in
18 the system behave honestly. A notable feature of the attack is that it is capable of “trapping”
19 transactions that do not naturally fall inside a cycle, i.e. those that are transmitted at significantly
20 different times (with respect to the network latency). To mitigate the attack, we propose three
21 methods based on three different complementary approaches. We show the effectiveness of the
22 proposed mitigation methods through simulations, and explain their limitations.

23 **2012 ACM Subject Classification** Security and privacy → Distributed systems security

24 **Keywords and phrases** Transaction ordering, fairness, Condorcet cycle

25 **Digital Object Identifier** 10.4230/LIPIcs.AFT.2023.15

26 **1** Introduction

27 The first blockchain application, Bitcoin, emerged in the midst of the financial crisis of 2008,
28 caused in part by the excessive trust placed in centralized institutions. Blockchain technology
29 changed this. In blockchain, there is no central authority or intermediary controlling the entire
30 system. Instead, transactions are validated and included through a consensus mechanism
31 among the participating parties. Decentralization also promotes transparency and reduces
32 the possibility of fraud or corruption since all transactions are publicly recorded and visible
33 to all participants on the network.

34 Despite the decentralized nature of blockchain systems, the ordering of transactions is
35 carried out in a centralized manner; the miner/validator who creates a block determines the
36 ordering of transactions within the block. This gives too much power to a single entity as
37 the success and profitability of a transaction can be determined by the order in which the
38 transaction appears in a block [6, 1, 8, 9, 16]. For instance, when a Non-Fungible Token
39 (NFT) is dropped in a given block, transactions positioned earlier in the block have a higher
40 chance of acquiring the NFT compared to those placed later.

41 To address this issue, several existing works [12, 20, 11, 4, 10, 13] proposed decentralized
42 methods for handling transaction ordering, where instead of a single node, a committee
43 of nodes collectively decide on the ordering of received transactions. At the core of these
44 methods, each node in the system reports a list of transactions in the order the node has



© Mohammad Amin Vafadar and Majid Khabbazian;
licensed under Creative Commons License CC-BY 4.0
5th Conference on Advances in Financial Technologies (AFT 2023).



Editors: Joseph Bonneau and S. Matthew Weinberg; Article No. 15; pp. 15:1–15:21

Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

15:2 Condorcet Attack

45 received them. The system then generates and agrees on a “fair” ordering by taking the
46 reported orderings into account.

47 Finding a fair ordering is not trivial. For instance, suppose that for any two transactions
48 \mathbf{tx}_1 and \mathbf{tx}_2 , we require \mathbf{tx}_1 to be placed before \mathbf{tx}_2 if a large majority of nodes in the system
49 claim to have received \mathbf{tx}_1 before \mathbf{tx}_2 . Despite being a primitive requirement, no method
50 can provide a guarantee due to an impossibility result rooted in social choice theory [2]. As
51 an example, consider a system consisting of three nodes, where each node has received three
52 transactions: \mathbf{tx}_1 , \mathbf{tx}_2 , and \mathbf{tx}_3 . Suppose the nodes report the ordering as $[\mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3]$,
53 $[\mathbf{tx}_2, \mathbf{tx}_3, \mathbf{tx}_1]$, and $[\mathbf{tx}_3, \mathbf{tx}_1, \mathbf{tx}_2]$. In this case, \mathbf{tx}_1 is reported to be before \mathbf{tx}_2 by two
54 nodes (i.e. the majority), \mathbf{tx}_2 is reported to be before \mathbf{tx}_3 by two nodes, and \mathbf{tx}_3 is reported
55 to be before \mathbf{tx}_1 by two nodes. This essentially creates a cycle, referred to as *Condorcet*
56 *cycle* [5], which prevents any final ordering from respecting the views of the majority on how
57 transactions should be ordered.

58 The existing fair ordering methods adopt a relaxed approach to ordering transactions
59 inside a Condorcet cycle. For instance, Cachin et al. in Quick-Fairness [4] do not mention
60 any ordering mechanism for such transactions, and Kelkar et al. in Aequitas [12] suggest a
61 simple alphabetical ordering. This relaxed approach is, perhaps, due to two reasons: 1) it is
62 not possible to guarantee fair ordering of transactions inside a cycle; 2) Condorcet cycles
63 occur infrequently in practice, and when they do occur, they usually involve transactions that
64 are received around the same time by the nodes in the system. Nevertheless, in this work,
65 we show that Condorcet cycles deserve more attention as they can be created “artificially”
66 by adversaries through what we refer to as the *Condorcet attack*. An interesting feature of
67 the Condorcet attack proposed in this work is that it can be conducted by a client outside
68 the system. In particular, the attack can be effectively executed even when all the nodes in
69 the system are honest!

70 As will be explained later, in the Condorcet attack, an adversarial client sends a small
71 number of transactions to different nodes in the system. The adversary chooses the timing and
72 order of these transactions to create a Condorcet cycle that traps many honest transactions
73 in it (a Condorcet cycle with only the adversary’s transactions in it is all but harmless to the
74 system.). This cycle has to be broken by the leader in a leader-based method in order to
75 establish a total ordering. Even if the leader is honest, the act of breaking the cycle could
76 change the order of honest transactions, which would have otherwise been appropriately
77 ordered¹.

78 Defending against the Condorcet attack is not straightforward. It is partly because it is
79 challenging to differentiate between honest transactions and otherwise valid transactions that
80 are submitted with the intention of creating a cycle. It becomes notably more challenging
81 to safeguard the system when, in addition to the adversarial client outside the system, the
82 leader and possibly a fraction of the nodes in the system are adversarial. Nevertheless,
83 in this work, we propose three mitigation techniques based on three different approaches.
84 The proposed techniques complement each other and can work together harmoniously to
85 maximize resistance against the attack.

86 In summary, we make the following contributions. We introduce a framework for a new
87 type of attack (Condorcet attack) against fair transaction ordering. We show that the attack
88 can be highly successful in trapping honest transactions in a cycle. To mitigate the attack,
89 we propose three techniques based on three different complimentary approaches, and show

¹ Kelkar et al. [11] consider it a success for an adversary if the adversary places two transactions into the same cycle when they should not have been.

90 the effectiveness of the technique through simulations.

91 **2 Related Work**

92 The classical approach to mandating fair transaction ordering is through *secure causal*
 93 *ordering*, a method introduced by Birman and Reiter in 1994 [17], and later improved by
 94 Cachin et al. in 2001 [3]. This method uses encryption to conceal the content of transactions
 95 during the ordering process, and allows decryption of transactions only after the order of
 96 transactions is finalized. This prevents an adversary from observing the content of transactions
 97 during the ordering process, thereby effectively eliminating attacks such as the sandwich
 98 attack [16] that rely on inspecting transaction contents. However, the method is unable to
 99 prevent “blind front-running attacks” where, for instance, the adversary’s sole objective is to
 100 order her transaction first (to, for example, get priority in purchasing a token). In addition,
 101 the method cannot prevent attacks based on transactions’ metadata, as metadata (such as
 102 the source of a transaction) is not encrypted.

103 The second approach to mandating fair transaction ordering involves a first-come, first-
 104 served strategy. This approach is complementary to the first approach and has been the
 105 focus of several recent studies. The existing methods that follow this strategy can be
 106 broadly classified into two categories: timestamp-based methods and batch-based methods.
 107 Timestamp-based methods are computationally inexpensive but require synchronized clocks.
 108 Batch-based methods, on the other hand, offer stronger fairness than timestamp-based
 109 methods, but can tolerate fewer adversarial nodes.

110 **Timestamp-based Methods.** An example of a timestamp-based protocol is Pompe [20]
 111 due to Zhang et al. Pompe introduces a notion of fairness called the *ordering linearizability*.
 112 This notion stipulates that if the highest timestamp of a transaction \mathbf{tx} is less than the
 113 lowest timestamp of a transaction \mathbf{tx}' among honest nodes, then \mathbf{tx} must be ordered before
 114 \mathbf{tx}' in the final order of transactions. Although it can enforce ordering linearizability, Pompe
 115 suffers from censorship issues, as noted in [11].

116 Kursawe’s Wendy protocol [13] is another timestamp-based protocol that defines a notion
 117 of fairness called *timed-relative-fairness*. This notion requires that if all honest nodes received
 118 a transaction \mathbf{tx} before time τ , and transaction \mathbf{tx}' after τ , then \mathbf{tx} must be ordered before
 119 \mathbf{tx}' .

120 **Batch-based Methods.** Aequitas [12] by Kelkar et al. is a batch-based method
 121 proposed for fair transaction ordering. Aequitas enforces a fairness notion known as the
 122 γ -*batch-order-fairness*. The notion requires that if two transactions \mathbf{tx} and \mathbf{tx}' are received by
 123 all nodes in a system with n nodes, and γn nodes received \mathbf{tx} before \mathbf{tx}' , then all honest nodes
 124 must output \mathbf{tx} no later than \mathbf{tx}' . Aequitas suffers from high communication complexity of
 125 $\mathcal{O}(n^3)$, and can guarantee only a weak notion of liveness, one of the two pillars of consensus
 126 security.

127 The second batch-based method is Quick-Fairness [4] proposed by Cachin et al. This
 128 method enforces a fairness notion called the κ -*differential order-fairness*. This notion
 129 mandates that if the number of nodes that have received transaction \mathbf{tx} before \mathbf{tx}' exceeds
 130 $\kappa + 2f$ for some $\kappa \geq 0$, then \mathbf{tx} should be ordered no later than \mathbf{tx}' , where f is the maximum
 131 number of adversarial nodes in the system. Kelkar et al. [11] show that this notion of
 132 fairness is indeed a re-parameterized version of the γ -batch-order-fairness notion. They
 133 also demonstrate that the Quick-Fairness protocol satisfies fairness only when all nodes are
 134 honest.

135 Kelkar et al. addressed the shortcomings of Aequitas in their protocol called Themis [11].

15:4 Condorcet Attack

136 Themis satisfies the γ -batch-order-fairness notion, and solves the liveness problem of Aequitas.
137 Moreover, SNARK-Themis variant offers a communication complexity of $\mathcal{O}(n)$ and standard
138 Themis offers a communication complexity of $\mathcal{O}(n^2)$ instead of $\mathcal{O}(n^3)$ offered by Aequitas. In
139 addition, it satisfies a more generalized notion of fairness than the one used in Quick-Fairness
140 and a stronger notion of fairness than those used in the existing time-based methods. For
141 these reasons, in our work, we focus on Themis and Aequitas as the state-of-the-art fair
142 transaction ordering methods.

143 3 Model

144 **System.** We consider a permissioned system with a committee of n nodes. The nodes
145 receive transactions directly from clients, and submit the list of their received transactions
146 together with the order in which the transactions were received to a special node called the
147 leader. The leader collects the lists of transactions from the nodes, and proposes a final
148 ordering using a pre-decided fair-ordering protocol. The leader in the system is not fixed,
149 and can change through a pre-determined protocol.

150 **Fair Ordering.** We adopt the batch-order-fairness from [12, 11], the strongest notion of fair
151 ordering proposed to date. For a parameter $\frac{1}{2} < \gamma \leq 1$, the batch-order-fairness specifies
152 that if a fraction γ of nodes receive a transaction \mathbf{tx} before receiving another transaction
153 \mathbf{tx}' , then \mathbf{tx} must be placed in the order before \mathbf{tx}' , with exceptions allowed only if \mathbf{tx}
154 and \mathbf{tx}' are within the same Condorcet cycle (Condorcet cycles are defined in Section 4).
155 Transactions within a cycle are placed in a batch, and are ordered by a method that we refer
156 to as *batch-ordering scheme*. The existing fair ordering protocols either do not specify a
157 batch-ordering scheme or propose a simple one (e.g., an alphabetical-based scheme [12]).

158 **Network.** The network utilizes public key infrastructure and secure digital signatures for
159 communications. As in [12], we consider two networks: the (standard) internal network
160 (for communication amongst nodes in the system) and the external network (for clients to
161 transmit their transactions to the system).

162 We assume that the network operates under partial synchrony [7], meaning that there
163 is a network delay Δ (not known to the nodes) that limits the amount of time it takes for
164 messages to be delivered between nodes.

165 **Adversary.** We consider an adversary who has control over $f \geq 0$ out of n nodes, and also
166 possesses at least one client capable of submitting transactions to the system. The adversary
167 can deviate arbitrarily from the protocol. The adversary does not have control over the
168 external network, but may have full control over the internal network, hence can delay and
169 reorder messages up to the bound Δ .

170 4 Preliminaries

171 **Graph Terminology.** We use $G = (V, E)$ to denote a graph with the set of vertices V
172 and the set of edges E . In this work, each vertex represents a transaction, therefore, we use
173 the terms vertices and transactions interchangeably. Unless otherwise specified, we use an
174 unweighted and directed graph. In the case of a weighted graph, the weight or cost associated
175 with the edge $(u, v) \in E$ is represented by $w(u, v)$.

176 A *tournament graph* is a directed graph where every pair of distinct vertices is connected
177 by a directed edge in either of two possible directions. A *Strongly Connected Component*
178 (*SCC*) in a graph is a maximal subgraph in which there is a path from every vertex to every

179 other vertex. A *condensation graph* is obtained from the original graph by combining its
 180 SCCs into a single vertex. A *Directed Acyclic Graph (DAG)* is a directed graph that contains
 181 no cycles, meaning it is possible to move from one vertex to another along the directed edges,
 182 but it is not possible to return to the original vertex by following a sequence of directed
 183 edges. A *topological sort* is an ordering of the vertices in a DAG such that for every directed
 184 edge (u, v) , vertex u appears before vertex v . In other words, if there is a directed edge from
 185 vertex u to vertex v , then u must appear before v in the topological sort. A *Hamiltonian*
 186 *Path* is a path in a graph that passes through every vertex exactly once. A *Hamiltonian*
 187 *Cycle* is a cycle in a graph that passes through every vertex exactly once.

188 **Themis.** Themis is the latest ordering method which achieves batch-order-fairness in the
 189 presence of an adversary who controls up to $f < \frac{(2\gamma-1)n}{4}$ nodes out of n nodes. Themis
 190 categorized received transactions into three different categories.

- 191 ■ **Solid Transactions:** A transaction is solid if it has been received by at least $n - 2f$
 192 nodes. A solid transaction is one that has been received by enough honest nodes that the
 193 leader can unambiguously include it in the current proposal while respecting the fairness
 194 guarantees.
- 195 ■ **Blank Transactions:** A transaction is blank if it has not been received by at least
 196 $n(1 - \gamma) + f + 1$ nodes. A blank transaction has not been received by enough nodes
 197 yet, hence excluding it from the current proposal will not violate fairness with respect to
 198 transactions that are included.
- 199 ■ **Shaded Transactions:** A shaded transaction is a transaction that is neither solid nor
 200 blank. A shaded transactions is received by enough nodes to be included to preserve
 201 fairness, but not enough nodes to finalize its position in the current proposal.

202 Themis is a leader-based method and works in three phases, as described below.

- 203 ■ **Phase 1 (Fair Propose):** The Fair Propose phase is the first phase of the algorithm, where
 204 each node proposes a set of transactions and their local orderings to the leader. The
 205 leader then uses the local orderings of $n - f$ nodes to build a dependency graph. In
 206 the dependency graph, an edge from a vertex v_1 to v_2 indicates that the transaction v_1
 207 should be placed before the transaction v_2 . From the dependency graph, the leader then
 208 computes the condensation graph and its topological sorting to output a fair ordering.
- 209 ■ **Phase 2 (Fair Update):** The Fair Update phase is the second phase of the algorithm,
 210 where the leader node updates the ordering for previous proposals. This is necessary
 211 since this is part of the deferred ordering technique, and new transactions may depend
 212 on previously proposed transactions, and these dependencies need to be accounted for in
 213 the ordering. The Fair Update algorithm takes the local transaction orderings of $n - f$
 214 nodes for previously proposed shaded transactions as input and outputs the updated
 215 dependencies.
- 216 ■ **Phase 3 (Fair Finalize):** The Fair Finalize phase is the third and final phase of the
 217 algorithm, where a sequence of proposals is finalized into a final ordering. The Fair
 218 Finalize algorithm updates the graphs for each proposal and computes the condensation
 219 graphs and their topological sorting. It then retrieves the final transaction ordering for
 220 each proposal based on the Hamiltonian cycles of the vertices in the sorted condensation
 221 graphs.

222 **Condorcet Cycles.** As mentioned above, Themis constructs a dependency graph, a
 223 directed graph where each vertex represents a transaction, and an edge from a vertex v_1 to
 224 v_2 indicates that the transaction corresponding to v_1 should be placed before the transaction

225 corresponding to v_2 . We refer to any cycle in this dependency graph as a *Condorcet cycle*.
 226 We note that cycles can occur in a dependency graph because of the Condorcet paradox [11].

227 **5** Condorcet Attack

228 In this section, we present the framework of the Condorcet attack. The attack aims at
 229 trapping honest transactions (i.e., transactions submitted by honest clients) inside a Condorcet
 230 cycle. If there is no effective batch-ordering scheme in place (e.g., if the batch-ordering
 231 scheme is alphabetical-based as suggested in [12]), this can change the ordering of the honest
 232 transactions even when all the nodes in the system are honest.

233 An adversary can take different strategies to impose a Condorcet cycle. For instance,
 234 suppose that the adversary controls f nodes, including the leader, in the system. The
 235 adversary then controls f local orderings, and can manipulate these orderings in a way to
 236 create a cycle. In the simulation section, we show that this strategy can not only create a
 237 cycle but also chain the cycles to involve more honest transactions. Nevertheless, the length
 238 of these cycles is typically small and the chain usually breaks rather quickly. As a result,
 239 this strategy is not effective in trapping distant transactions² (e.g., two transactions whose
 240 times of submission are separated by a multiple of the average network latency).

241 Another strategy, which is the one we take in this work, is to create a Condorcet cycle by
 242 injecting (valid) transactions into the system following a pre-described pattern. This can be
 243 done by an adversarial client outside the system, and can be effective even when all the nodes
 244 in the system are honest. The attack will be more effective in creating cycles and bypassing
 245 potential countermeasures if the adversary controls a fraction of nodes in the system (see
 246 Example 5).

247 The immediate damage of imposing a Condorcet cycle, as mentioned earlier, is that it
 248 can change the true ordering of honest transactions. In addition to this, the attack may be
 249 used to conduct other malicious activities; for instance, the adversary can create a cycle and
 250 then with the help of an adversarial leader can try to place its own transaction in desired
 251 positions in the final ordering.


252 ► **Example 1.** Let $\mathcal{P} = \{P_1, P_2, P_3\}$ be a partition of nodes, where P_1, P_2 and P_3 are three
 253 parts with almost equal size. In this simple example, the adversary \mathcal{C} uses/injects two
 254 transactions \mathbf{A}, \mathbf{B} (i.e., $\mathcal{S} = \{\mathbf{A}, \mathbf{B}\}$). In the initialization phase, \mathcal{C} sends the transaction \mathbf{A} and
 255 then \mathbf{B} to all the nodes in part P_1 , and sends the transaction \mathbf{B} to all the nodes in part P_2 (it
 256 sends no transactions to the nodes in part P_3). Then, after the pause period, \mathcal{C} sends \mathbf{A} to
 257 all the nodes in part P_2 , and transaction \mathbf{A} and \mathbf{B} , in that order, to all the nodes in part P_3 .
 258 Suppose that during the pause phase, the nodes receive three honest transactions $\mathbf{tx}_1, \mathbf{tx}_2$,
 259 and \mathbf{tx}_3 all the in that order. The local ordering of transactions at each node will be then:

$$\begin{aligned}
 260 \quad P_1 &: [\mathbf{A}, \mathbf{B}, \mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3] \\
 P_2 &: [\mathbf{B}, \mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3, \mathbf{A}] \\
 P_3 &: [\mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3, \mathbf{A}, \mathbf{B}]
 \end{aligned}$$

261 Note that without the adversarial client \mathcal{C} disturbing the system (i.e., without transactions \mathbf{A}
 262 and \mathbf{B}), the system would have had an easy job of ordering the honest transactions as all
 263 the nodes in the system have received the honest transactions in the same order, i.e. $[\mathbf{tx}_1,$
 264 $\mathbf{tx}_2, \mathbf{tx}_3]$. Because of the adversary's transactions \mathbf{A}, \mathbf{B} , and \mathbf{C} , however, we have a cycle now

² The analysis of why this occurs is left for future work.

265 as illustrated in Figure 1. In this figure, an edge from a transaction \mathbf{tx} to a transaction \mathbf{tx}'
 266 indicates that the majority of the nodes have received \mathbf{tx} before \mathbf{tx}' .



■ **Figure 1** A Condorcet cycle created using two transactions A and B

267 **Attack Framework.** In this section, we provide a general construction that encompasses
 268 the different variants of the Condorcet attack. Let \mathcal{C} be a client controlled by the adversary,
 269 and \mathcal{S} be a set of arbitrary but valid transactions created by \mathcal{C} . Let \mathcal{P} be a partition of the
 270 nodes in the system. In its general form, the Condorcet attack is executed in three phases:

- 271 ■ **Phase 1 (Initialization):** In this phase, the client \mathcal{C} sends a number of transactions from
 272 the set \mathcal{S} to each node in the system. The set of transactions sent to a node can be
 273 different from that sent to another system. More specifically, the client \mathcal{C} assigns a subset
 274 \mathcal{S}_i of \mathcal{S} (possibly an empty subset) to each part \mathcal{P}_i in the partition \mathcal{P} . It then determines
 275 an ordering for each subset \mathcal{S}_i , and sends the transactions in \mathcal{S}_i to all the nodes in part
 276 \mathcal{P}_i with the determined order.
- 277 ■ **Phase 2 (Pause):** In the second phase, the attacker waits for a specific amount of time,
 278 referred to as the pause time, for the honest transactions to be received by the nodes.
 279 The adversary can trap more transactions within a cycle as the pause time increases.
 280 However, the pause time should be limited to a single consensus round in the system as
 281 the attack should not extend across multiple consensus rounds.
- 282 ■ **Phase 3 (Finalization):** The third and final phase is the finalization phase, where the
 283 attacker completes the Condorcet cycle by sending a new set of transactions to each
 284 part in the partition. More specifically, the client \mathcal{C} assigns a subset \mathcal{S}'_i of \mathcal{S} (typically a
 285 different subset than \mathcal{S}_i , used in the initialization phase) to each part \mathcal{P}_i in the partition
 286 \mathcal{P} . It then determines an ordering for each subset \mathcal{S}'_i , and sends the transactions in \mathcal{S}'_i to
 287 all the nodes in part \mathcal{P}_i with the determined order.

288 ► **Remark 2.** In practice, nodes in the system may receive some honest transactions during
 289 the initialization and/or finalization phases. These transactions may or may not get trapped
 290 in the Condorcet cycle. Based on our simulation results, however, the vast majority of honest
 291 solid transactions during the pause time fall into the Condorcet cycle.

292 ► **Remark 3.** A potential issue that can impact the success of the Condorcet attack is that
 293 the external network may deliver the transactions injected by the adversary out of order.
 294 For instance, in Example 1, the transactions A and B may be received out of order by the
 295 nodes in part \mathcal{P}_1 , in which case a cycle does not occur. If the network is prone to packet
 296 reordering, then to improve its success, the adversary can execute multiple Condorcet attacks
 297 concurrently through what we refer to as *cloning*.

298 **Cloning.** Packet reordering can happen in a network because of various factors such as
 299 network congestion, routing algorithms, and the physical distance between the source and

300 the destination. To conduct a successful Condorcet attack, it is important that nodes receive
 301 the injected packets in the order they were transmitted; a deviation from the intended order
 302 may result in the failure of the attack.

303 To increase the success probability of the attack in the presence of network reordering, the
 304 adversary can send cloned transactions to the nodes: Instead of sending a single transaction
 305 A , the adversary sends multiple clones of the transaction. For instance, in Example 1, the
 306 adversary can send A_1 and A_2 instead of A , and sends B_1 and B_2 instead of B . Essentially, the
 307 adversary interleaves the execution of two Condorcet attacks (for better results, the adversary
 308 can interleave several instances of the attack). Then, if the network does not change the order
 309 of the transactions, the nodes in parts P_1 , P_2 , and P_3 will receive transactions as follows:

$$\begin{aligned}
 P_1 &: [A_1, A_2, B_1, B_2, tx_1, tx_2, tx_3] \\
 P_2 &: [B_1, B_2, tx_1, tx_2, tx_3, A_1, A_2] \\
 P_3 &: [tx_1, tx_2, tx_3, A_1, A_2, B_1, B_2]
 \end{aligned}$$

311 In Section 7.3, we show that cloning can significantly increase the success rate of the Condorcet
 312 attack in the presence of network reordering.

313 **Impact on Current Solutions.** The current fair transaction ordering protocols either do
 314 not offer a batch-ordering scheme (e.g. [4]) or offer a primitive one (e.g. [12]). For instance,
 315 the proposed batch-ordering scheme in Aequitas [12] is alphabetical ordering. Therefore, if
 316 an adversary creates a Condorcet cycle, as in Example 1, the honest transactions will be
 317 ordered alphabetically rather than by the time of their arrival.

318 Themis [11], proposes a more thoughtful batch-ordering scheme. In this scheme, a
 319 Hamiltonian cycle is built and then used to order transactions in the cycle. The latest
 320 version of Themis at the time of writing this work suggests to break the weakest link in
 321 the Hamiltonian cycle in order to convert it into a Hamiltonian path. We use this version
 322 of Themis in our work. In the best-case scenario, the order of honest transactions in the
 323 Hamiltonian cycle is preserved. Even in this case, the final ordering of these transactions
 324 can change because the Hamiltonian cycle has to be converted into a path by breaking the
 325 cycle at one point. It is at this point where honest transactions can be divided into two
 326 groups. The ordering of the honest transactions within each group remains correct, but the
 327 ordering of any two transactions from different groups will be incorrect. Therefore, similar
 328 to [4] and [12], Themis is vulnerable to the Condorcet attack even if all the nodes (including
 329 the leader) in the system are honest.

330 To combat the Condorcet attack, a natural approach is to use a strong batch-ordering
 331 scheme. For instance, in Example 1, we can observe that all the nodes report tx_1 before tx_2 ,
 332 and all the nodes report tx_2 before tx_3 , whereas only two third of the nodes report A before
 333 B . In this example, the weakest link is between adversarial transactions, and breaking it (as
 334 suggested by Themis) does not change the true ordering of the honest transactions. This
 335 solution works for the scenario described in Example 1. However, this solution may not work
 336 in other settings, for example when the adversary controls a faction of nodes in the system
 337 (see Example 5).

338 **6 Mitigation**

339 Despite its simplicity, it is not straightforward to completely defeat the Condorcet attack. In
 340 the following, we present three mitigation techniques based on three different approaches to
 341 hinder an adversary from successfully executing the attack. We elaborate on the strength of
 342 each technique and confirm it through simulations later in Section 7. We also explain the

343 limitation of each technique, i.e. under what settings/assumptions the technique may not be
344 effective.

345 An interesting feature of the proposed mitigation methods is that they do not conflict
346 with each other, thus in practice, they can be applied together for the maximum defense
347 against the attack. Another interesting feature of the proposed mitigation methods is that
348 they can be easily applied to Themis, which is currently the strongest fair-ordering solution
349 in the literature. We elaborate on this when we cover each proposed mitigation.

350 6.1 Ranked Pairs Batch-ordering

351 The approach we take in our first proposed mitigation is to use a strong batch-ordering
352 scheme to order transactions within a batch. Formally, a batch-ordering scheme is a method
353 that takes as input a strongly connected (possibly weighted) directed graph $G = (V, E)$,
354 and returns an ordering of the vertices V . The strongly connected graph represents the
355 transactions that are in a batch/cycle.

356 The candidate for our batch-ordering scheme is *ranked pairs*, an electoral system developed
357 by Nicolaus Tideman in 1987 [19]. Ranked pairs satisfies many natural and well-studied
358 axiomatic properties in social choice theory³ and is resistant to certain manipulations
359 including adding, deleting and changing a fraction of orderings reported by nodes [15]. In
360 ranked pairs, the ordering is essentially achieved by choosing a maximal subset E' of E in
361 the inputted graph $G = (V, E)$ with high weights such that $G' = (V, E')$ is a DAG. The DAG
362 is then used to establish an ordering of the vertices V .

363 More specifically, our ranked pairs batch-ordering scheme takes as input a weighted
364 directed graph $G = (V, E)$. Let $E_1 = E$. In step i , $i \geq 1$, the algorithm selects an edge
365 $(u, v) \in E_i$ with the highest weight⁴. It then sets the order $u \prec v$, unless this violates the
366 transitivity of the orders decided in previous steps. Finally, it sets $E_{i+1} \leftarrow E_i \setminus \{(u, v)\}$, and
367 terminates if $E_{i+1} = \emptyset$.

368 We note that the idea in the above batch-ordering scheme is to establish an ordering
369 using the strongest edges in G . This will be an effective defense against the Condorcet attack
370 if the ordering of the honest transactions has “strong support” in the system. In a special
371 case where all the nodes are honest, and all support/report the same ordering of honest
372 transactions, the Condorcet attack can be fully prevented as stated in the following theorem.

373 **► Proposition 4.** *Suppose that the Condorcet attack succeeds in creating a Condorcet cycle.*
374 *Let $\mathbf{tx}_1, \mathbf{tx}_2, \dots, \mathbf{tx}_m$ be the set of honest transactions in the Condorcet cycle. Suppose that*
375 *all the nodes in the system are honest and report \mathbf{tx}_i before \mathbf{tx}_j for every $1 \leq i < j \leq m$.*
376 *Then the proposed ranked pairs batch-ordering scheme returns the true ordering of the honest*
377 *transaction, that is it orders \mathbf{tx}_i before \mathbf{tx}_j for every $1 \leq i < j \leq m$.*

378 **Proof.** Let $G = (V, E)$ be the graph with V representing the transactions in the Condorcet
379 cycle, and the weight of each edge $(u, v) \in E$, represented as $w(u, v)$, be equal to the number
380 of nodes that reported u before v . Let u_1, u_2, \dots, u_m be the vertices in V that represent the
381 honest transactions. Let $E_f \subseteq E$ be the set of all edges with the full support of the nodes,

³ besides Schulze, ranked pairs is the only existing electoral system that satisfies anonymity, Condorcet criterion, resolvability, Pareto optimality, reversal symmetry, monotonicity, and independence of clones [18].

⁴ When there are multiple edges with the highest weight, one can be chosen according to a fixed tie-breaking method.

15:10 Condorcet Attack

382 that is

$$383 \quad E_f = \{e \in E \mid w(e) = n\},$$

384 where n is the number of nodes in the system. Since all the nodes in the system have the
 385 same view on the ordering of the honest transactions, we get that $(u_i, u_j) \in E_f$ for every
 386 $1 \leq i < j \leq m$. We note that the sub-graph $G' = (V, E_f)$ of G is cycle free, as otherwise
 387 there will be a cycle in the ordering of individual nodes. The ranked pairs batch-ordering
 388 algorithm first chooses all the edges in E_f before proceeding with other edges in E . When
 389 the algorithm covers all the edges in E_f the true ordering of the honest transactions will be
 390 set, and cannot be changed by the remaining steps of the algorithm. ◀

391 **Limitation.** Proposition 4 considers an ideal scenario where 1) all the nodes are honest,
 392 and 2) they all report the honest transaction in the same order. If one of the above two
 393 conditions does not hold, however, the Condorcet attack may be able to create a cycle (see
 394 the following example).

395 ▶ **Example 5.** Consider a system with $n = 5$ nodes. Let $\mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3$ be three honest
 396 transactions. An adversarial client \mathcal{C} can create a Condorcet cycle of the form

$$\begin{aligned} N_1 &: [A_1, A_2, A_3, A_4, \mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3] \\ N_2 &: [A_2, A_3, A_4, \mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3, A_1] \\ 397 \quad N_3 &: [A_3, A_4, \mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3, A_1, A_2] \\ N_4 &: [A_4, \mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3, A_1, A_2, A_3] \\ N_5 &: [\mathbf{tx}_3, \mathbf{tx}_2, \mathbf{tx}_1, A_1, A_2, A_3, A_4] \end{aligned}$$

398 where A_1, A_2, A_3, A_4 are the transactions submitted by \mathcal{C} . Note that all the nodes, except
 399 Node 5, report the order $[\mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3]$, while node 5 reports $[\mathbf{tx}_3, \mathbf{tx}_2, \mathbf{tx}_1]$ (Node 5 is either
 400 controlled by the adversary or is an honest node who has simply received the transactions
 401 in this order). If we run the proposed ranked pairs batch-ordering scheme on this cycle,
 402 the returned order of honest transactions may be incorrect. It is because the edge between
 403 any pair of transactions has a weight of 4 in the dependency graph. As a result, an edge
 404 between two honest transactions such as \mathbf{tx}_1 and \mathbf{tx}_2 may be eliminated in the ranked pairs
 405 method, which would result in \mathbf{tx}_2 and \mathbf{tx}_3 to be ordered before \mathbf{tx}_1 . As for Themis, if we
 406 use the proposed method by Yannis Manoussakis [14] (as suggested by Themis), we get the
 407 Hamiltonian cycle $(A_1, A_2, A_3, A_4, \mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3)$. All the edges in this cycle have the identical
 408 weight of four, hence there is no distinct weakest edge. Therefore, Themis may remove any
 409 of the edges in the cycle. If the removed edge is between two honest transactions, the final
 410 ordering of honest transactions would be incorrect. We remark that both the ranked pairs
 411 batch-ordering scheme and Themis would order honest transactions correctly if Node 5 order
 412 honest transactions as $[\mathbf{tx}_1, \mathbf{tx}_2, \mathbf{tx}_3]$. This example, therefore, shows that the adversary
 413 has more power in modifying the order of honest transactions if (in addition to injecting
 414 transactions) it controls a number of nodes in the system (e.g. Node 5 in this example).

415 ▶ **Remark 6.** To use the proposed ranked pairs batch-ordering scheme in Themis, we can
 416 simply replace the Hamiltonian-based batch-ordering scheme of Themis with the ranked pairs
 417 batch-ordering scheme in the FairFinalize algorithm. We remark that the weight information
 418 of the dependency graph is available within the FairFinalize algorithm, thus this replacement
 419 is possible.

420 6.2 Post-decryption Resolution

421 In secure causal ordering, as mentioned earlier, transactions are ordered while they are
422 encrypted, and get decrypted only once a total ordering is committed [17, 3]. This prevents
423 an adversary from observing the contents of transactions while they are being ordered, hence
424 eliminating those front-running attacks (e.g. the sandwich attack [16]) that must examine
425 the content of transactions.

426 To mitigate the Condorcet attack, we propose to maintain the above strategy, except we
427 leave the ordering of transactions inside a Condorcet cycle to after they are decrypted. Note
428 that after the decryption of these transactions, an adversary cannot impose a change to the
429 ordering as 1) there is already a consensus on the set of transactions that must be included,
430 thus the adversary cannot add or remove any transaction to the set; 2) the ordering of the
431 transactions is performed locally at each node using a pre-determined algorithm. In other
432 words, it is too late for the adversary to manipulate the ordering of transactions, although
433 the contents of transactions are disclosed.

434 Once the transactions within a cycle are decrypted, their contents are disclosed, enabling
435 them to be partitioned into independent groups (i.e., transactions inside different groups are
436 independent of each other). Each group can then be ordered independent of the others. By
437 implementing this measure, the adversary is unable to manipulate the ordering of honest
438 transactions if the adversary's transactions are independent of honest transactions. This
439 is because the adversary's transactions will not fall within any group that includes honest
440 transactions. Note that we still need to order the groups themselves (i.e. which group comes
441 first, which comes second, and so on). As transactions across various groups have no effect
442 on one another, the groups can be safely ordered using a pre-determined algorithm such as
443 ranked pairs as described in Section 6.1.

444 ► **Remark 7.** In the Themis protocol, we can apply the above post-decryption resolution
445 method within the FairFinalize algorithm: If transactions A and B are independent, the edge
446 between them in the dependency graph can be safely removed.

447 **Limitation.** The post-decryption resolution prevents the adversary from manipulating
448 the order of honest transactions if the adversary's transactions are independent of the honest
449 transactions. In certain scenarios, however, the adversary may be able to create dependencies.
450 For instance, consider a situation where a popular NFT is dropping in a block currently
451 being formed. Given the high demand, many transactions are transmitted with the intention
452 of acquiring this NFT. Recognizing this, the adversary can execute the Condorcet attack
453 by using transactions that fall within the same dependency group as those attempting to
454 acquire the NFT.

455 Another limitation of the post-decryption resolution is the computational burden it places
456 on the system to identify dependencies between transactions.

457 6.3 Broadcast

458 In the Condorcet attack, the adversary follows a well-structured three-phase strategy: in
459 the first phase, the adversary sends a set of transactions, then pauses in the second phase,
460 and then finishes the attack by sending another round of transactions in the third phase.
461 The idea behind our third mitigation technique is to disturb/break the above pattern by
462 broadcasting transactions inside the system as soon as they arrive at an honest node. Because
463 of the broadcast, the adversary's transactions that were submitted in the first phase will
464 propagate in the system, which can nullify the adversary's target in the third phase since the

15:12 Condorcet Attack

465 transactions that the adversary transmits in the third phase have already been received by
466 the nodes (thus their order has already been decided by the nodes).

467 In Section 7.5, we observe that this strategy proves highly effective in mitigating the
468 suggested Condorcet attack. However, it is important to note that this strategy does incur
469 increased communication overhead as a drawback. For instance, in Themis, nodes transmit
470 transactions only to the leader as opposed to broadcasting in the network by themselves.
471 Therefore, when applied to Themis, the above strategy will increase Themis’s communication
472 overhead (although it does not increase Themis’s quadratic communication complexity).

473 **Limitation.** The main limitation of the above broadcast-based mitigation technique is
474 that it will be ineffective if the adversary has strong control over the internal network. For
475 instance, in Themis and Aequitas, it is assumed that the adversary controls all message
476 delivery in the internal network, and can delay messages up to a bound Δ . If Δ is large
477 enough (e.g., if it is larger than the duration of the Condorcet attack) then the adversary
478 can circumvent the proposed mitigation by delaying all the broadcast transactions so they
479 are delivered only after the attack is complete.

480 7 Simulation

481 To assess the impact of the Condorcet attack, as well as the effectiveness of the proposed
482 mitigation methods, we conduct a series of experiments through simulations. In this section,
483 we present the results of these experiments.

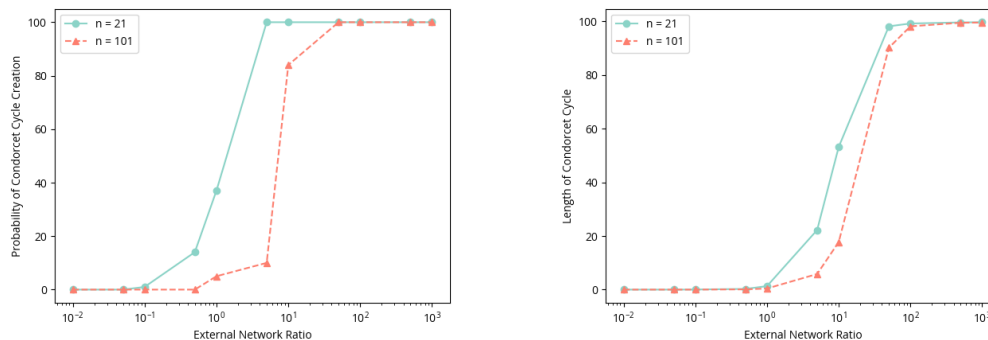
484 **Environments.** Our simulation encompasses four environments. The first environment
485 captures the honest setting, where all the nodes and clients are honest, thereby eliminating
486 the possibility of a Condorcet attack. Even in this environment, Condorcet cycles can occur.
487 Therefore, we are interested to know if our proposed ranked pairs batch-order scheme can
488 more effectively order transactions within a cycle than the Hamiltonian-cycle-based scheme
489 used in Themis.

490 In the second environment, all the nodes in the system are honest, but there is an external
491 adversary, who conducts the Condorcet attack from outside the system. In this environment,
492 we are interested to evaluate the success rate and impact of the Condorcet attack (i.e., how
493 many honest transactions the adversary can trap within a cycle).

494 In the third environment, we introduce packet reordering to the external network. We
495 evaluate the impact of this on the success rate of the Condorcet attack. We also observe how
496 the cloning method can help the adversary to improve its success rate.

497 The last environment that we consider is similar to the second environment, except this
498 time we guard the system using the proposed mitigation methods. In this environment, we
499 measure the impact of the Condorcet attack in order to examine the strength of the proposed
500 mitigation methods.

501 **Clients.** We use a sending process to submit all the clients’ transactions to the system.
502 The sending process transmits transactions in sequence at discrete times t_i , $i \geq 0$. At each
503 time instance, the process sends (n copies of) the transaction of a given client to all the n
504 nodes in the system. Each copy of the transaction will arrive at its destination node with a
505 random delay drawn independently from a distribution named `NetworkDist`. We refer to
506 this distribution as the network latency. We use another distribution, `GenerationDist`, to
507 determine the delay between two consecutive time instances (i.e. $t_{i+1} - t_i$ follows the
508 `GenerationDist` distribution). Similar to [11], we set both `GenerationDist` and
509 `NetworkDist` to exponential distributions with means of one and r , respectively. We refer to



(a) The chance of a Condorcet cycle

(b) Number of transactions in cycles

■ **Figure 2** Condorcet cycles in the honest environment

510 r as *external network ratio*. One can think of r as the expected number of clients who
 511 transmit transactions within a time frame equal to the average network latency.

512 **Themis Variant.** In our simulations, we use the practical Themis variant with the
 513 communication complexity of $\mathcal{O}(n^2)$, instead of the the SNARK-Themis variant. In our
 514 simulations, all transactions are eventually received by each node in a single round. Therefore,
 515 the choice of γ does not have any impact on the simulation results (hence, we simply set
 516 $\gamma = 1$). We used the latest version of Themis, which breaks the Hamiltonian cycle by
 517 removing the weakest link. The weakest link is the link that has the least weight or support
 518 in the Hamiltonian cycle. To construct a Hamiltonian cycle, we used the proposed method
 519 by Yannis Manoussakis [14] as suggested by Themis.

520 7.1 Honest Environment

521 **Honest Environment Setting.** In this environment, all the nodes and clients are honest,
 522 and consequently, there is no Condorcet attack. Nevertheless, as shown in Figure 2, Condorcet
 523 cycles can occur particularly when the external network ratio is greater than one.

524 To obtain the results plotted in Figure 2, we varied the external network ratio from 0.01
 525 to 1000. For each given network ratio, and each network size of $n = 21$ and $n = 101$, we
 526 conducted 100 simulation runs. In each run, the sending process transmitted 100 transactions
 527 (at 100-time instances drawn from the `GenerationDist` distribution). Once every node
 528 received all the transmitted transactions, we proceeded to generate the dependency graph
 529 using the Themis algorithm. By examining the graph (i.e. extracting strongly connected
 530 components) we then identified all the Condorcet cycles.

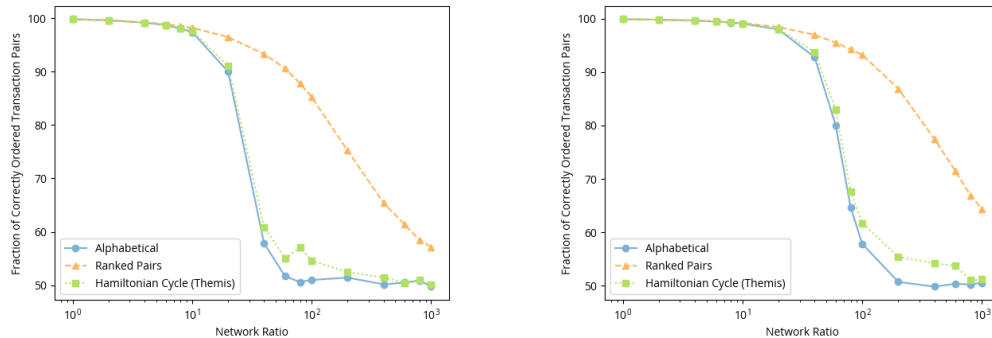
531 **Cycle Length.** An interesting observation from Figure 2 is that when the external
 532 network ratio is less than about one, Condorcet cycles rarely occur. As the external network
 533 ratio becomes larger than one, however, Condorcet cycles start to appear. For high values of
 534 the external network ratio, as depicted in Figure 2, Condorcet cycles not only occur frequently
 535 but also include many of the transmitted transactions. Overall, this observation suggests a
 536 critical threshold at which the system’s behavior, with respect to creating Condorcet cycles,
 537 significantly changes.

538 **Condorcet Cycles Categories.** We refer to Condorcet cycles that are not created by
 539 an adversary as *natural Condorcet cycles*. Conversely, we call a Condorcet cycle adversarial
 540 if it is created by an adversary. In Section 6.1, we proposed a ranked pairs batch-ordering

541 scheme to handle the ordering of transactions within an adversarial Condorcet cycle. Later
 542 in this section, we demonstrate that the proposed scheme indeed alleviates the severity of
 543 the Condorcet attack.

544 **Ranked Pairs Performance.** Here, we show (Figure 3) that the proposed ranked pairs
 545 batch-ordering scheme is also a good candidate for ordering transactions within a natural
 546 Condorcet cycle. Consequently, even in an honest environment, we can improve fairness in
 547 ordering transactions by replacing the existing batch-ordering schemes (i.e., the alphabetical
 548 scheme, and the Hamiltonian-based scheme of Themis) with the proposed ranked pairs
 549 batch-ordering scheme.

550 **Batch Ordering-Schemes Performance Comparison.** In Figure 3, the external
 551 network ratio (the x -axis) ranges from 1 to 1000; this is the range in which Condorcet
 552 cycles naturally occur. The y -axis shows the fraction of transaction pairs that are ordered
 553 correctly according to their transmission time. Each data point in Figure 3 is the average
 554 of values obtained over 100 simulation runs. The data presented in this figure demonstrate
 555 the superiority of the proposed ranked pairs batch-ordering scheme for two network sizes of
 556 $n = 21$ and $n = 101$.



(a) The number of nodes is $n = 21$

(b) The number of nodes is $n = 101$

■ **Figure 3** Fraction of correctly ordered transactions in the honest environment

557 7.2 Adversarial Environment

558 **Adversarial Environment Setting.** In the existing adversarial environments in the
 559 literature, there is often at least one (typically up to $f = \theta(n)$) adversarial node in the system.
 560 In our adversarial environment, in contrast, all the nodes in the system can be honest. There
 561 is, however, an adversarial client in our environment who orchestrates the Condorcet attack
 562 from outside the system.

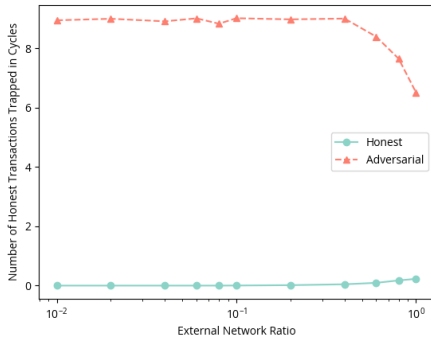
563 In this section, we evaluate the performance of the Condorcet attack in this environment.
 564 In particular, we measure the success rate of the attack in the number of honest transactions
 565 it can trap within a cycle. The measurement is carried out for external network ratios r less
 566 than one, as natural Condorcet cycles are rare in this regime, particularly when $r \ll 1$. This
 567 allows us to assess the strength of the attack in creating cycles in a setting where Condorcet
 568 cycles do not naturally happen.

569 In our simulation, we simply use two adversarial transactions to create the Condorcet cycle
 570 as described in Example 1. We set the pause time of the Condorcet attack to $\tau \in \{10, 50\}$
 571 times the mean of the `GenerationDist` distribution. This means that, on average, τ honest

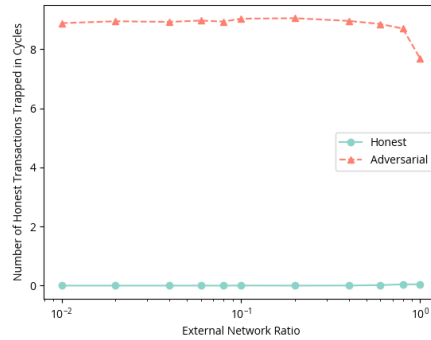
572 transactions are transmitted to the system during the pause time.

573 In parallel to the transmissions of honest transactions, the two adversarial transactions
 574 are transmitted to create a Condorcet cycle. Once all transactions are received by the nodes,
 575 we calculate two separate dependency graphs: one considering the adversarial transactions,
 576 and one ignoring them. By comparing these two dependency graphs, we then assess the
 577 impact of the attack on the final ordering.

578 **Condorcet Attack Performance.** Figures 4 and 5 show the average number of the
 579 honest transactions that the attack can trap within cycles over two different settings: $\tau = 10$
 580 and $\tau = 50$. As shown, for a wide range of external network ratios, the attack can trap
 581 nearly all the honest transactions that are transmitted during the pause time (about 9
 582 honest transactions in the setting $\tau = 10$, and nearly 49 honest transactions in the setting
 583 $\tau = 50$). This demonstrates the strength of the attack, considering that, on average τ honest
 584 transactions are submitted to the system during the pause time (and the attack traps nearly
 585 all of them).




(a) $\tau = 10, n = 21$




(b) $\tau = 10, n = 101$

■ **Figure 4** Number of honest transactions trapped in Condorcet cycles for $\tau = 10$.



(a) $\tau = 50, n = 21$



(b) $\tau = 50, n = 101$

■ **Figure 5** Number of honest transactions trapped in Condorcet cycles for $\tau = 50$.

586 **7.3 Network Reordering**

587 In the Condorcet attack, the adversary sends a sequence of transactions in a particular order
 588 to create a cycle. The external network may, however, change the order of transactions
 589 transmitted, which can, in turn, reduce the attack’s success rate. To evaluate this, we
 590 performed simulations over a network which changes the order of two consecutively
 591 transmitted transactions with probability $0 \leq p \leq 0.5$. For each value of p , we performed
 592 1000 runs of simulations. The success rate of the attack was set to the fraction of runs in
 593 which the attack successfully trapped the honest transactions in a Condorcet cycle.

594 Using the above setting, we conducted two instances of the Condorcet attack. The first
 595 instance uses two adversarial transactions A and B as in Example 1, and takes the following
 596 pattern:

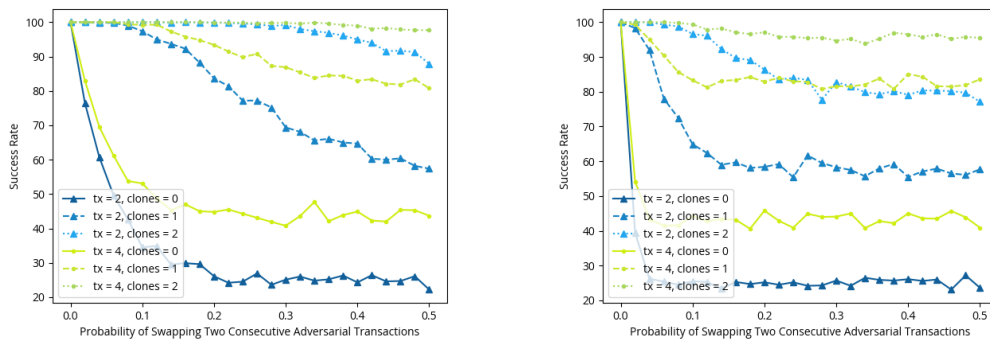
- P_1 : A, B, Pause
- 597 P_2 : B, Pause, A
- P_3 : Pause, A, B

598 As illustrated in Figure 6, this instance is sensitive to network reordering (the success
 599 rate of the attack drops quickly with p). As shown in the figure, the attack’s success rate
 600 increases when we use the second instance of cloning described below.

601 In our second instance (denote as $tx = 4$ in Figure 6), the adversary partitions nodes
 602 into four parts P_1, P_2, P_3 and P_4 , and uses four transactions (A, B, C and D) instead of two,
 603 in the following pattern:

- P_1 : A, B, Pause, C, D
- 604 P_2 : B, C, Pause, D, A
- P_3 : C, D, Pause, A, B
- P_4 : D, A, Pause, B, C

605 This instance of the Condorcet attack is more robust against network reordering as
 606 demonstrated in Figure 6. As in the first instance, the success rate of the instance can be
 607 boosted using the cloning method. In particular, note that the second instance together with
 608 a single clone is almost fully resistant to network transaction reordering.



(a) The number of nodes is $n = 21$

(b) The number of nodes is $n = 101$

■ **Figure 6** Impact of network reordering on the success of the Condorcet attack.

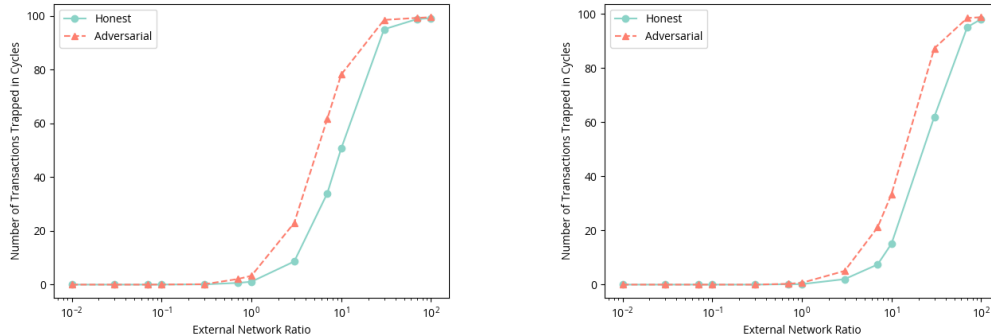
609 7.4 A Non-Injective Condorcet Attack

610 Injecting transactions into the system is a key component of the proposed Condorcet attack.
 611 Without this component, an adversary has limited power in creating cycles even when the
 612 adversary controls the leader and a faction of all the nodes in the system.

613 To illustrate the above point, we conducted simulations over two networks with sizes:
 614 $n = 21$ and $n = 101$. In our simulation, the adversary controls the maximum fraction of
 615 nodes, including the leader, allowed by Themis (a quarter of nodes minus one). All these
 616 nodes report the order of their received transactions in reverse, in a strategy to create
 617 Condorcet cycles⁵. The external network ratio is varied from 0.01 to 100 to capture a wide
 618 range of network conditions. The total number of transmitted transactions is set to 100.

619 To evaluate the impact of the above strategy in creating cycles, we created two dependency
 620 graphs. The first graph represents the scenario where the adversarial nodes reverse their
 621 orderings, whereas the second graph represents the scenario where the adversarial nodes
 622 report the true ordering. Figure 7 shows the results of our simulation.

623 **Non-Injective Condorcet Attack Performance.** As shown in Figure 7, the
 624 adversary's attempts to create cycles are largely unsuccessful in the region where the
 625 external network ratio is less than one. We note that in this region, the average temporal
 626 gap between two different transaction transmissions is more than the average network
 627 latency. In particular, when $r \ll 1$ (i.e., when transactions are transmitted far apart in time
 628 with respect to the network latency), honest nodes in the system have a clear view of the
 629 true ordering of transactions. In this region, the adversary is all but powerless in creating
 630 cycles⁶, as evident in Figure 7. In contrast, in the same region, an external adversary can
 631 create a cycle using the proposed Condorcet attack, even when all the nodes in the system
 632 are honest.



(a) The number of nodes is $n = 21$

(b) The number of nodes is $n = 101$

■ **Figure 7** The non-injective attack has limited power in creating cycles.

⁵ We note that this may not be an optimum strategy to create Condorcet cycles. Nevertheless, we believe that an optimum strategy (which may be computationally intractable) may not be significantly more successful than the adopted strategy. We leave the validation of this claim for future work.

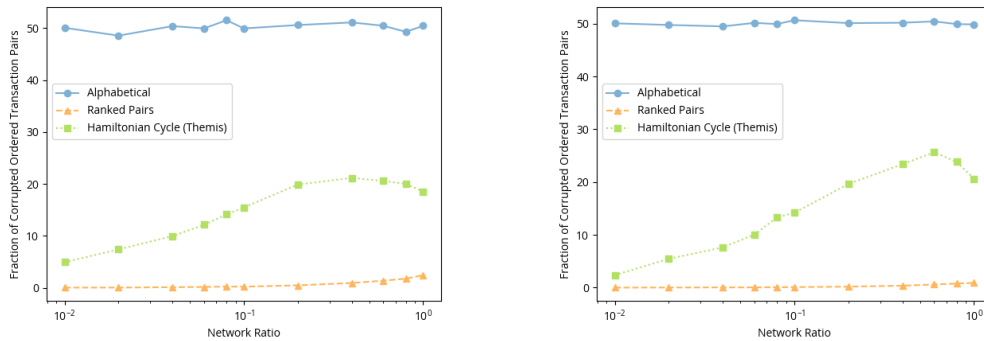
⁶ When $r > 1$ (i.e., in the region where Condorcet cycles naturally emerge) the adversary achieves some degree of success in creating larger cycles than naturally occur.

633 **7.5 Mitigation**

634 In this section, we evaluate the performance of our mitigation methods in preventing or
 635 minimizing the impact of the Condorcet attack.

636 **Ranked-pairs-based Mitigation Method.** To evaluate the effectiveness of this
 637 mitigation, we conducted a simulation over two network sizes of $n = 21$ and $n = 101$. We
 638 set the pause time of the attack to 10 times the mean of `GenerationDist`, and set the total
 639 number of honest transactions to 20. We varied the external network ratio r from 0.001 to
 640 1. Recall that in this range of external network ratio (i.e., $r < 1$), Condorcet cycles do not
 641 emerge naturally; rather they are created by the Condorcet attack. To evaluate the true
 642 impact of our ranked-pairs mitigation method, therefore, we focused on this region.

643 **Ranked Pairs Mitigation Performance.** Figure 8 compares the performance of our
 644 proposed ranked-pairs-based mitigation method to the Hamiltonian-based method used in
 645 Themis, and the simple alphabetical method. The results show that the proposed ranked-pairs
 646 method achieves a low error rate, indicating that it can effectively order honest transactions
 647 correctly even when they fall in a Condorcet cycle. In contrast, the Themis algorithm’s error
 648 rate increases as the network ratio increases, and reaches as high as about 25%. The error
 649 rate in the case of alphabetical ordering is 50%. Note that a random ordering method can,
 650 on average, correctly orders 50% of all the pairs of transactions. In this sense, the worst-case
 651 transaction ordering error is 50%, which is the case for the alphabetical method (this method
 652 is essentially a random ordering method).

(a) The number of nodes is $n = 21$ (b) The number of nodes is $n = 101$

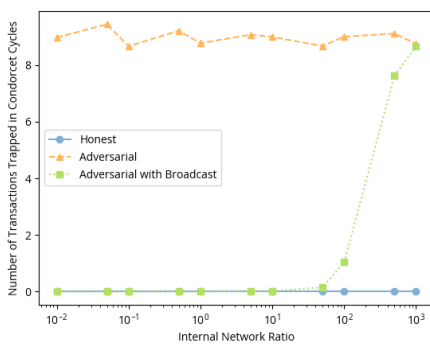
■ **Figure 8** The performance of the proposed ranked pairs-based mitigation method

653 **The Broadcast-based Mitigation Method.** To evaluate the effectiveness of the
 654 broadcast-based mitigation method, we conducted simulations using two network sizes: $n = 21$
 655 and $n = 101$. We introduced a new exponential distribution called `InternalNetworkDist`,
 656 which captures the random delays experienced by messages within the internal network.
 657 Specifically, we sample from `InternalNetworkDist` to determine the delay between sending
 658 a transaction from one node to another node. This is in contrast to `NetworkDist`, which is
 659 used to determine the random delays between a client and a node in the external network.

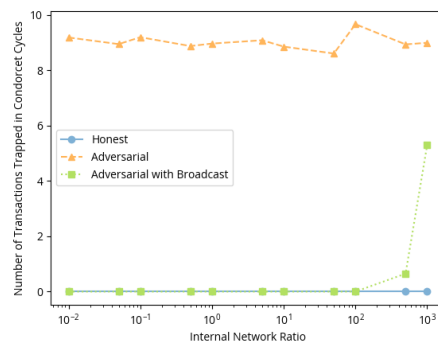
660 In our simulation, we set the mean of `InternalNetworkDist` to r' . We refer to r' as the
 661 *internal network ratio*. In our simulations, we set τ to 10 times the mean of `GenerationDist`
 662 (i.e. $\tau = 10 \cdot r$), and set the total number of honest transactions to 20. We fixed the external
 663 network ratio to $r = 0.1$, to ensure that no natural Condorcet cycles were created, and varied
 664 the internal network ratio r' from 0.01 to 1000.

665 **Broadcast Environment Categories.** We analyzed the number of honest transactions
 666 trapped in a Condorcet cycle under three different settings. In the first setting, referred
 667 to as the “honest setting”, nodes did not broadcast and the adversary did not conduct a
 668 Condorcet attack. In the second setting, nodes still did not broadcast, but the adversary
 669 attempted a Condorcet attack. Finally, in the last setting, the adversary launched an attack
 670 while the nodes employed the broadcasting method to mitigate it.

671 **Broadcast Mitigation Performance.** Figure 9 shows the result of our simulations
 672 in the above three settings. The results demonstrate that the proposed broadcast-based
 673 mitigation is highly effective in preventing the adversary from creating a Condorcet cycle
 674 and trapping honest transactions. This can be attributed to two key factors: Firstly, the
 675 mitigation strategy disrupts the completion of the pause phase, thereby preventing honest
 676 transactions from being trapped in a Condorcet cycle. When the internal network ratio
 677 r' is smaller than the pause time, almost no transactions are trapped. Interestingly, even
 678 when r' exceeds the pause time, the adversary cannot achieve the same level of performance.
 679 It is because the broadcast of transactions with the internal network can still somewhat
 680 disturb the ordering of adversarial transactions. This reduces the success rate of the attack
 681 as the specific ordering of adversarial transactions is crucial for creating a Condorcet cycle.
 682 If, on the other hand, the adversary has enough control over the internal network to delay
 683 transactions as much as the pause time, it can circumvent the proposed broadcast-based
 684 mitigation as the adversary can enforce the ordering of its transactions within the internal
 685 network by delaying all the messages.



(a) The number of nodes is $n = 21$



(b) The number of nodes is $n = 101$

■ **Figure 9** The performance of the proposed broadcast mitigation method

686 **8 Conclusion**

687 Condorcet cycles can occur naturally. While these natural cycles may not significantly disrupt
 688 fairness in the system since transactions falling within these cycles are typically received
 689 around the same time, the artificial creation of Condorcet cycles can lead to significant
 690 unfairness in the system. In this paper, we showed that even with all nodes in the system
 691 behaving honestly, it is relatively simple to generate such artificial cycles. Furthermore, we
 692 demonstrated that these created cycles possess significant power, as they can trap transactions
 693 submitted at widely different times that would not naturally fall within a cycle.

694 To address this attack, we proposed three mitigation methods using different approaches.
 695 These methods complement one another and can be employed collectively to fortify the
 696 defensive measures against the attack. Through simulations, we showcased that despite their
 697 described limitations, the proposed mitigation methods can substantially reduce the adverse
 698 impact of the Condorcet attack.

699 **References**

-
- 700 **1** Carsten Baum, James Hsin-yu Chiang, Bernardo David, Tore Kasper Frederiksen, and Lorenzo
 701 Gentile. Sok: Mitigation of front-running in decentralized finance. *IACR Cryptol. ePrint*
 702 *Arch.*, page 1628, 2021. URL: <https://eprint.iacr.org/2021/1628>.
- 703 **2** Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia, editors.
 704 *Handbook of Computational Social Choice*. Cambridge University Press, 2016. doi:10.1017/
 705 CBO9781107446984.
- 706 **3** Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient
 707 asynchronous broadcast protocols. In Joe Kilian, editor, *Advances in Cryptology - CRYPTO*
 708 *2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA,*
 709 *August 19-23, 2001, Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages
 710 524–541. Springer, 2001. doi:10.1007/3-540-44647-8_31.
- 711 **4** Christian Cachin, Jovana Micic, Nathalie Steinhauer, and Luca Zanolini. Quick order fairness.
 712 In *Financial Cryptography and Data Security - 26th International Conference, FC 2022,*
 713 *Grenada, May 2-6, 2022, Revised Selected Papers*, Lecture Notes in Computer Science, pages
 714 316–333. Springer, 2022. doi:10.1007/978-3-031-18283-9_15.
- 715 **5** M. d. Condorcet. Essay on the application of analysis to the probability of majority decisions.
 716 *Paris: Imprimerie Royale*, 1785.
- 717 **6** Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz
 718 Breidenbach, and Ari Juels. Flash boys 2.0: Frontrunning, transaction reordering, and
 719 consensus instability in decentralized exchanges. *CoRR*, abs/1904.05234, 2019. URL: <http://arxiv.org/abs/1904.05234>, arXiv:1904.05234.
- 721 **7** Cynthia Dwork, Nancy A. Lynch, and Larry J. Stockmeyer. Consensus in the presence of
 722 partial synchrony. *J. ACM*, 35(2):288–323, 1988. doi:10.1145/42282.42283.
- 723 **8** Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. Sok: Transparent dishonesty:
 724 Front-running attacks on blockchain. In Andrea Bracciali, Jeremy Clark, Federico Pintore,
 725 Peter B. Rønne, and Massimiliano Sala, editors, *Financial Cryptography and Data Security*
 726 *- FC 2019 International Workshops, VOTING and WTSC, St. Kitts, St. Kitts and Nevis,*
 727 *February 18-22, 2019, Revised Selected Papers*, volume 11599 of *Lecture Notes in Computer*
 728 *Science*, pages 170–189. Springer, 2019. doi:10.1007/978-3-030-43725-1_13.
- 729 **9** Lioba Heimbach and Roger Wattenhofer. Sok: Preventing transaction reordering manipulations
 730 in decentralized finance. *CoRR*, abs/2203.11520, 2022. arXiv:2203.11520, doi:10.48550/
 731 arXiv.2203.11520.
- 732 **10** Mahimna Kelkar, Soubhik Deb, and Sreeram Kannan. Order-fair consensus in the
 733 permissionless setting. In Jason Paul Cruz and Naoto Yanai, editors, *APKC '22: Proceedings of*

- 734 *the 9th ACM on ASIA Public-Key Cryptography Workshop, APKC@AsiaCCS 2022, Nagasaki,*
735 *Japan, 30 May 2022*, pages 3–14. ACM, 2022. doi:10.1145/3494105.3526239.
- 736 11 Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. Themis: Fast,
737 strong order-fairness in byzantine consensus. *IACR Cryptol. ePrint Arch.*, page 1465, 2021.
738 URL: <https://eprint.iacr.org/2021/1465>.
- 739 12 Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine
740 consensus. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology -*
741 *CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa*
742 *Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes*
743 *in Computer Science*, pages 451–480. Springer, 2020. doi:10.1007/978-3-030-56877-1_16.
- 744 13 Klaus Kursawe. Wendy, the good little fairness widget: Achieving order fairness for blockchains.
745 In *AFT '20: 2nd ACM Conference on Advances in Financial Technologies, New York, NY,*
746 *USA, October 21-23, 2020*, pages 25–36. ACM, 2020. doi:10.1145/3419614.3423263.
- 747 14 Yannis Manoussakis. A linear-time algorithm for finding hamiltonian cycles in tournaments.
748 *Discret. Appl. Math.*, 36(2):199–201, 1992. doi:10.1016/0166-218X(92)90233-Z.
- 749 15 David C. Parkes and Lirong Xia. A complexity-of-strategic-behavior comparison between
750 schulze’s rule and ranked pairs. In Jörg Hoffmann and Bart Selman, editors, *Proceedings*
751 *of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto,*
752 *Ontario, Canada*. AAAI Press, 2012. URL: [http://www.aaai.org/ocs/index.php/AAAI/](http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/5075)
753 [AAAI12/paper/view/5075](http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/5075).
- 754 16 Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How
755 dark is the forest? In *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco,*
756 *CA, USA, May 22-26, 2022*, pages 198–214. IEEE, 2022. doi:10.1109/SP46214.2022.9833734.
- 757 17 Michael K. Reiter and Kenneth P. Birman. How to securely replicate services. *ACM*
758 *Transactions on Programming Languages and Systems*, 16(3):986–1009, 1994. doi:10.1145/
- 759 177492.177745.
- 760 18 Markus Schulze. A new monotonic, clone-independent, reversal symmetric, and condorcet-
761 consistent single-winner election method. *Soc. Choice Welf.*, 36(2):267–303, 2011. doi:
- 762 10.1007/s00355-010-0475-4.
- 763 19 T. N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and*
764 *Welfare*, 4(3):185–206, September 1987. doi:10.1007/bf00433944.
- 765 20 Yunhao Zhang, Srinath T. V. Setty, Qi Chen, Lidong Zhou, and Lorenzo Alvisi. Byzantine
766 ordered consensus without byzantine oligarchy. In *14th USENIX Symposium on Operating*
767 *Systems Design and Implementation, OSDI 2020, Virtual Event, November 4-6, 2020*, pages
768 633–649. USENIX Association, 2020. URL: [https://www.usenix.org/conference/osdi20/](https://www.usenix.org/conference/osdi20/presentation/zhang-yunhao)
769 [presentation/zhang-yunhao](https://www.usenix.org/conference/osdi20/presentation/zhang-yunhao).