

# Cryptographically Blinding the Mempool: A Systematic Review of Zero-Knowledge Based Architectures and Commit-Reveal Scheme in Decentralized Exchanges (DEXs)

\*Gboraloo A. W.<sup>1</sup>; Eke B.<sup>2</sup>; Onuodu F. E.<sup>3</sup>

<sup>1</sup> Department of Computer Science, Ken Saro-Wiwa Polytechnic, Bori, Rivers State, Nigeria.

<sup>2,3</sup> Department of Computer Science, University of Port Harcourt, Rivers State, Nigeria.

\*Corresponding Author

DOI: <https://dx.doi.org/10.51584/IJRIAS.2026.11060106>

Received: 01 June 2026; Accepted: 06 June 2026; Published: 26 June 2026

## ABSTRACT

Decentralized exchanges (DEXs) have emerged as a foundational component of blockchain-based financial systems, enabling trustless asset trading without centralized intermediaries. However, the transparency of public mempools introduces significant vulnerabilities, including front-running, sandwich attacks, transaction reordering, and broader information asymmetry. In response, Cryptographic mechanisms such as Zero Knowledge (ZK) based architectures and commit reveal schemes are increasingly proposed as a solution to these vulnerabilities. This research systematically reviews the structural transparency paradox and cryptographic architectures in Decentralized Exchange based Automated Market Makers (DEX-AMM), evaluate their effectiveness in mitigating Maximal Extractable Values (MEVs), analyze computational complexity trade-offs including proof generation/verification costs, gas overhead, latency, and throughput, and identify why commit-reveal may offer superior practical viability despite zk-proofs' stronger theoretical privacy guarantees. A comprehensive search was conducted across arXiv, IEEE Xplore, ACM Digital Library, Scopus, Web of Science, Google Scholar including grey literatures for studies published between 2021 to 2026. Findings indicate that ZK-based approaches provide strong cryptographic privacy guarantees but often incur computational overhead and integration complexity, zk-rollups provide strong validity guarantees through cryptographic proofs, but their practical security depends heavily on the sequencer layer used by ( zkSync, StarkEx, and Loopring) which is responsible for transaction ordering, which can censor, delay, reorder transactions or cause failure of execution, while Commit-reveal schemes may be superior for real-world DEXs due to their constant time hash-based complexity ( $O(1)$ ), lower gas costs, sub-second latency, and simpler implementation, despite requiring two-transaction UX friction, which can be mitigated through wallet automation. The computational efficiency advantage of commit-reveal becomes critical as DEX transaction complexity increases, where zk-circuit depth grows exponentially. Future research should prioritize optimizing zk-circuit efficiency, developing zk-commit-reveal hybrids system that balance cryptographic strength with computational practicality, and advancing hash-based commit-reveal schemes with UX improvements. DEX developers should prioritize commit-reveal for latency-sensitive applications and zk-proofs only when strongest cryptographic privacy is mandatory.

**Keywords:** Automated Market Maker (AMM), Zero-Knowledge Proofs, Commit-Reveal Scheme, Maximal Extractable Value (MEV), PulseChain Blockchain, Sandwich Attacks, Mempool Privacy, Transaction Leakage, Transaction Lineage, PRISMA 2021.

## INTRODUCTION

Blockchain technology has fundamentally transformed digital infrastructures by enabling decentralization, transparency, and tamper resistant systems without reliance on centralized intermediaries (Nakamoto, 2008). One of the most prominent applications of blockchain is decentralized finance (DeFi), within which

Decentralized Exchanges (DEXs) facilitate peer-to-peer trading through automated smart contracts. In particular, Automated Market Makers (AMMs) have become the dominant design paradigm for DEXs, replacing traditional order books with liquidity pools governed by deterministic pricing algorithms (Buterin, 2017; Angeris & Chitra, 2020), governed by constant product market rules  $x \cdot y \leq z$  (Chainlink, 2023; Maciej, 2021). AMMs enable instant token swaps directly from users' non-custodial wallets, providing continuous liquidity and user anonymity, thereby protecting retail participants from regulatory and custodial risks associated with CEXs (CFI Team, 2020; Kaur, 2023). Despite their advantages in accessibility and decentralization, a major architectural challenge exists within the public blockchain networks: the transparency paradox (Xu et al., 2022, 2023) where DEXs operates.

Public blockchains like Ethereum and Solana achieve Byzantine fault tolerance by using public verification of transactions, thus all pending smart contract calls are broadcasted to an open access unconfirmed transaction waiting area called the Mempool (Coinfinity, 2023), prior to confirmation, exposing sensitive trading information such as trade size, timing, and intent.

This universal exposure creates a toxic ecosystem, where arbitrage searchers, algorithmic bots, and blockchain validators can easily audit the public mempool to inspect the precise trade intent, slippage tolerance, and token amounts of pending orders (Xu et al., 2023). Such total visibility enables adversarial transaction reordering, front-running, sandwich attacks and censorship, leading to predatory Maximal Extractable Value (MEV) that systematically extract financial value directly from retail trades (Daian et al., 2020; Rui, 2019), which collectively contributes to information asymmetry in trading environments (Daian et al., 2020; Qin et al., 2022). These issues undermine the fairness, degrade user trust, and introduce inefficiencies in market execution of decentralized exchange automated Market Maker Systems.

To address these challenges, the modern blockchain industry has been heavily centered around the development of various mitigation mechanisms. Among these are Zero-Knowledge (ZK) cryptography (such as zk-SNARKs and zk-STARKs) and Layer-2 Rollup architectures (like zkSync, Starknet and ZKSwap) that facilitate compression and off-chain transaction processing, verified on-chain through the use of cryptographic proofs (Corey, 2024; Victor, 2024), enabling the verification of transaction validity without revealing underlying data, thereby preserving confidentiality while maintaining correctness (Ben-Sasson et al., 2014; Bowe et al., 2018). In parallel, commit reveal schemes aim to conceal transaction details during an initial commitment phase and disclose them only at execution phase, reducing the risk of pre-trade exploitation (Bonneau et al., 2015).

While both approaches offer promising solutions, they exhibit distinct trade-offs. ZK rollup architectures depend on a central off-chain sequencer or prover to batch transactions, which requires access to the unencrypted private transaction data in order to generate the mathematical proof. This reliance implies that transaction details are not kept private, but are subject to sequencer-level censorship and front-running, and does not prevent sandwich attacks (Daian et al., 2020), while ZK-based systems often involve significant computational overhead, complex circuit design, and integration challenges within existing blockchain infrastructures. Conversely, commit reveal architecture schemes are relatively simple to implement but remain susceptible to timing analysis, delayed execution risks, and partial information leakage. Recent research has therefore explored hybrid architectures that combine cryptographic privacy guarantees with protocol-level execution control to mitigate these limitations (Kelkar et al., 2020; Zhou et al., 2023).

Given the rapid evolution of this research area and the diversity of proposed solutions, there is a need for a structured and comprehensive synthesis of existing work. This study conducts a systematic review of zero-knowledge based architectures and commit reveal schemes architecture in DEXs using the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework (Page et al., 2021). This systematic review contributes to both academic research and practical DEX development in four key ways:

1. **Taxonomical contribution:** We provide the first comprehensive categorization of zk-based architectures and commit-reveal privacy mechanisms specifically for Decentralized Exchange (DEX) applications, enabling researchers to systematically compare architectural variants and their MEV mitigation properties.

2. **Computational complexity analysis:** By explicitly analyzing the  $O(n^3)$ – $O(n^4)$  vs.  $O(1)$  complexity trade-off, we clarify when commit-reveal's computational efficiency may outweigh zk-proofs' stronger cryptographic privacy, offering a nuanced perspective beyond the assumption that "zk is always better."
3. **Practical guidance for DEX developers:** Our synthesis of gas costs, latency, throughput, and UX trade-offs provides actionable criteria for mechanism selection based on application priorities (e.g., latency-sensitive high-frequency trading vs. privacy-critical large trades).
4. **Research direction identification:** By identifying deployment barriers and implementation gaps, we highlight open problems requiring future research, including zk-circuit optimization, UX improvements for commit-reveal, and hybrid zk-commit-reveal architectures.

## RESEARCH METHODOLOGY

This research follows the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) guidelines (Page et al., 2021) which have been adapted to a structured systematic review methodology for scientific reproducibility and comprehensive coverage, and for structural clarity. The methodology section synthesizes relevant studies on zero-knowledge (ZK) based architectures and commit–reveal schemes in decentralized exchanges (DEXs) which is structured around four key stages: identification, screening, eligibility assessment, and inclusion.

1. **Research Design:** A qualitative systematic review design was employed to analyze existing research contributions on privacy-preserving mechanisms in blockchain-based trading systems. The review focuses on studies published between 2021 and 2026 to capture recent advancements in ZK proofs (e.g., zk-SNARKs, zk-STARKs) and commit reveal protocols within DEXs and Automated Market Maker (AMM) environments.
2. **Data Sources and Search Strategy:** A comprehensive search was conducted across major academic and preprint databases, including: IEEE Xplore, ACM Digital Library, Scopus, web of Science, Google Scholar and arXiv, and Blog/Grey Literature including: Speedrun Ethereum (developer tutorials), Shutter Network Blog (protocol documentation), Chainlink Articles (technical guides), Blockworks (industry analysis), CloudLogic (developer guides), StarkWare Blog (protocol updates) and Medium technical posts (community implementations). Search queries were developed using a combination of logic operators and free-text keywords targeting three conceptual domains: (1) privacy preserving technologies, (2) DEX/AMM applications, and (3) MEV/front-running problems. The search was refined to query sensitivity and specificity, adjusted the terms to balance recall and precision (McGuinness, 2020). The rationale for blog inclusion is that blockchain research evolves rapidly, with practical implementation details often published first in developer blogs before peer reviews.
3. **The String:** The search query follows Boolean logic: ("zero-knowledge" OR "zk-rollup" OR "zk-proof" OR "zk-SNARK" OR "zk-STARK" OR "commit-reveal" OR "commit-reveal scheme" OR "hash commitment" OR "threshold encryption" OR "encrypted mempool") AND ("DEX" OR "decentralized exchange" OR "AMM" OR "automated market maker" OR "DeFi" OR "decentralized finance") AND ("MEV" OR "maximal extractable value" OR "miner extractable value" OR "front-running" OR "transaction privacy" OR "mempool privacy" OR "sandwich attack" OR "sniping"), while the **Blog/Grey Literature Queries:** The blog and grey literature queries includes: Q1: "commit-reveal DEX MEV protection technical implementation blog" Q2: "zk-rollup DEX privacy blog post 2024 2025 technical" Q3: "threshold encryption DEX privacy blog Shutter network CloudLogic" Q4: "blockchain developer blog MEV protection strategies commit-reveal" Q5: "Speedrun Ethereum commit-reveal Solidity DEX front-running"
4. **The Time frame:** The time frame is between 2021 to 2026 (last 5 years), reflecting the emergence of MEV as a research topic following Daian et al.'s (2020) foundational work on flash boys and MEV

**Language:** English only, due to resource constraints and the predominance of English-language blockchain research

**Document types:** Conference proceedings, journal articles, preprints, technical reports

The 2021 start date was selected because MEV research expanded significantly after Daian et al. (2020) introduced the term "Miner Extractable Value," and privacy-preserving DEX mechanisms gained traction with the rise of zk-rollups in 2021 to 2022 (Qin et al., 2022).

### Eligibility Criteria

- a. **Inclusion Criteria (Academic AND Grey Literature) :** The records retrieved were screened at various levels. Studies were included if they met all of the following criteria: (1) Original research article peer-reviewed conference/journal paper, preprint, or technical report ), Grey literature (technical blogs/company documentation): Only if they meet all blog criteria below: (2) contains complete code implementation (public repository or tutorial with full code) **OR quantitative performance metrics** (gas costs, latency, throughput measurements), blog2: **DEX-specific** (not general blockchain privacy, smart contract privacy, or identity privacy), blog3: written by technical author (developer, protocol engineer, researcher) vs. marketing/promotional content, blog4: References academic papers (with citations) OR production deployments (mainnet systems), blog5: Reproducible (public GitHub repository, live testnet, or detailed step-by-step tutorial), (3) focus on privacy-preserving mechanisms zero-knowledge proofs OR commit-reveal schemes specifically applied to DEXs/AMMs, (3) Addresses MEV mitigation, front-running prevention, or transaction privacy as a primary or secondary objective, (4) Includes implementation details, empirical evaluation, or performance metrics (gas costs, latency, throughput, MEV reduction) and (5) published between 2021 to 2026 and in English with full text available.
- b. **Exclusion Criteria:** Retrieved records were also excluded if they contain **any** of the following criteria: (1) review articles, survey papers, book chapters, tutorials, or opinion pieces, (2) duplicate publications or overlapping versions of the same study, (3) Purely conceptual proposals without implementation, evaluation, or performance analysis, (4) Non-DEX blockchain privacy applications (e.g., zk-payments, zk-identity, privacy coins without DEX component), (5) Studies published before 2021 or in non-English languages and (6) Abstracts without full text available.

The exclusion of purely conceptual proposals was necessary because the research question also focus on computational complexity trade-offs and practical deployment status, which require empirical data (Chen et al., 2024). Non-DEX privacy applications (EC4) were excluded to maintain scope specificity, as privacy mechanisms for payments or identity have different performance requirements than DEX transaction privacy.

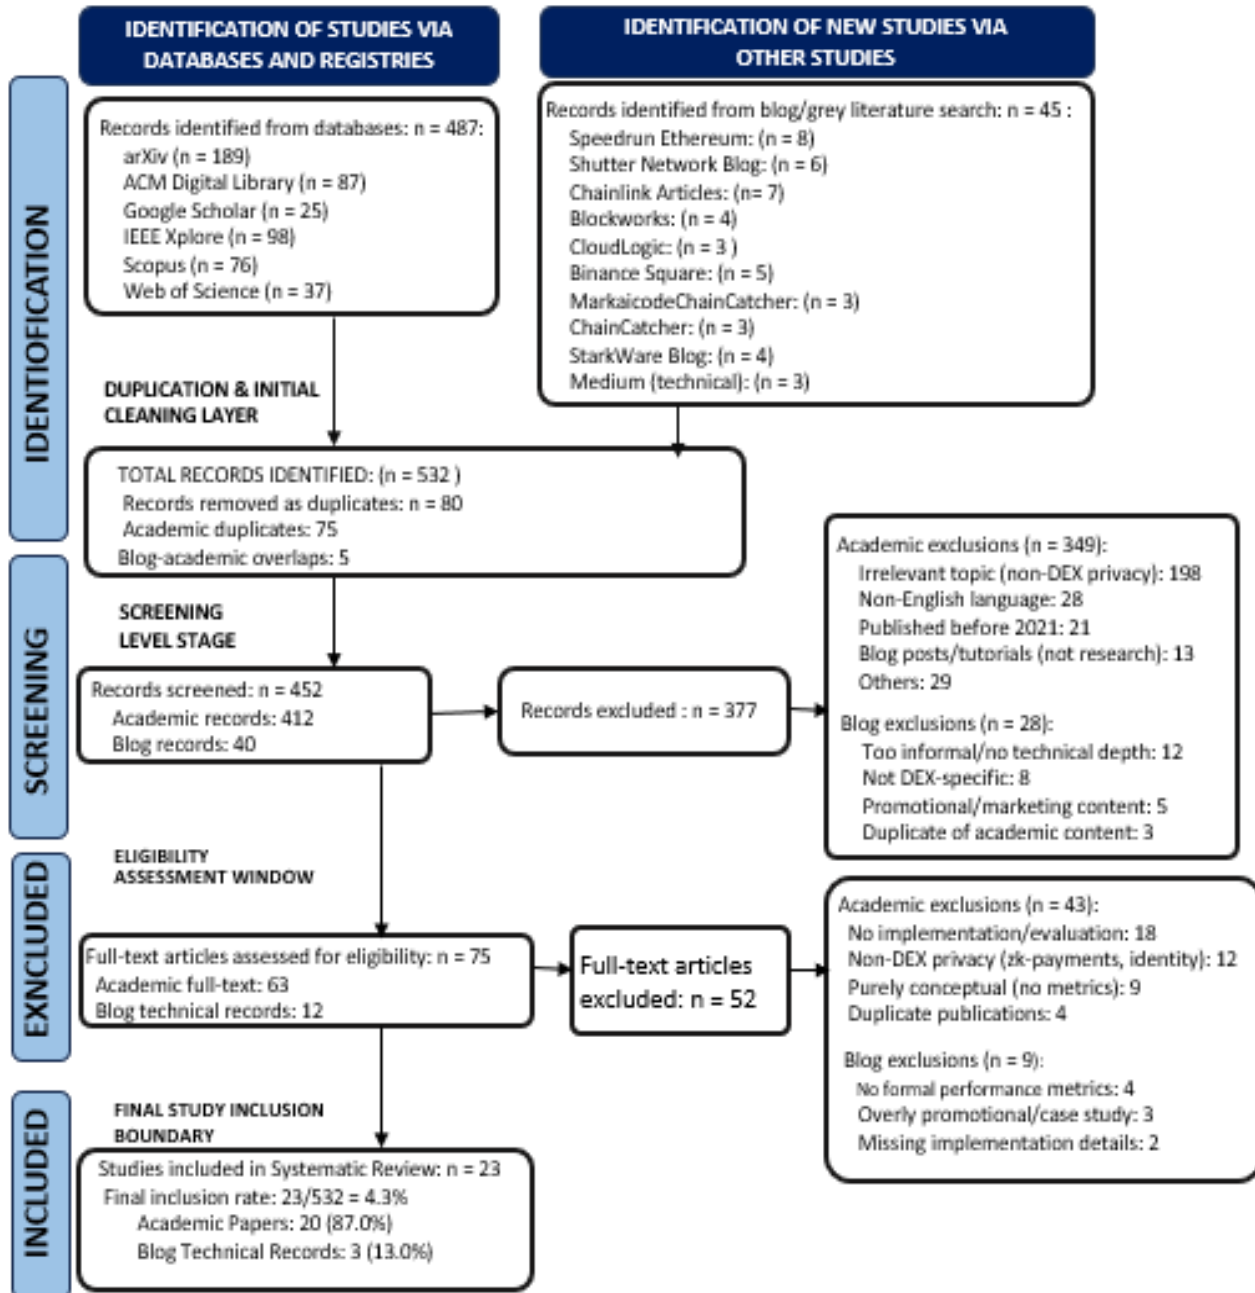
**6. Study Selection Process:** the study selection followed the PRISMA 2020 four-phase flow diagram which are: identification, screening, eligibility, and inclusion (Page et al., 2021) which are explained below:

- a. **Identification:** All records from database searches were imported into **Python** for deduplication using exact string matching on title, author, and publication year combinations.
- b. **Screening:** The titles and abstracts were screened against inclusion/exclusion criteria. Records were categorized as:
  - **Include:** Clearly meets all inclusion criteria
  - **Exclude:** Clearly meets at least one the Exclusion criteria
  - **Uncertain:** If it requires full-text review for decision
- c. **Eligibility:** Full-text articles for "include" and "uncertain" records were retrieved and assessed against detailed eligibility criteria.
- d. **Inclusion:** Final included studies were those meeting all the Inclusion Criteria after full-text assessment.

## PRISMA Flow Diagram

The study selection results are presented in the PRISMA flow diagram in figure 1 below, showing exact counts for: Records identified from each database, records removed as duplicates, records screened (title/abstract), records excluded at screening stage with reasons, Full-text articles assessed for eligibility, Full-text articles excluded with specific reasons, and studies included in final synthesis:

- a. **Identification:** Our search across six databases identified 487 records. arXiv contributed the largest share (189 records, 38.8%), followed by IEEE Xplore (98 records, 20.1%), ACM Digital Library (87 records, 17.9%), Scopus (76 records, 15.6%), Web of Science (37 records, 7.6%), and Google Scholar (25 records, 5.1%). After removing 75 duplicate records (15.4%), 412 unique records remained for screening.



**Figure 1:** PRISMA 2020 flow diagram for study selection of zk-based architectures and commit-reveal schemes in decentralized exchanges, including blog/grey literature records. Adapted from the PRISMA 2020 statement (Page et al., 2021). Grey literature inclusion follows PRISMA-S guidelines for comprehensive search reporting (Lúa et al., 2021).

- b. **Screening:** Title and abstract screening excluded 349 records (84.7%) for reasons including irrelevance to DEX privacy (n=198), non-DEX blockchain privacy applications (n=89), non-English language (n=28), publication before 2021 (n=21), and blog posts/tutorials (n=13).
- c. **Eligibility:** Sixty-three full-text articles were retrieved and assessed against detailed eligibility criteria. Forty-three articles (68.3%) were excluded with the following reasons: no implementation or empirical evaluation (n=18, 41.9%), non-DEX privacy applications (n=12, 27.9%), purely conceptual proposals without performance metrics (n=9, 20.9%), and duplicate publications (n=4, 9.3%).
- d. **Included:** Total of Twenty-three (23) studies were included in the final systematic review given the percentage as  $23/532 = 4.3\%$  of which Academic Papers: 20 (87.0%), while Blog Technical Records: 3 (13.0%), as shown in table 1 below:

Table 1: Search Source Contribution:

Source Type	Records Found	Percentage	Included
Academic Databases	487	91.5%	20 (87.0%)
Blog/Grey Literature	45	8.5%	3 (13.0%)
TOTAL	532	100%	23 (100%)

## RESULTS AND DISCUSSIONS

**RESULTS:** The result of findings are presented below according to the research questions:

**RQ1: Taxonomy Of Architectural Variants:** the result of the taxonomy of architectural variants of each architecture is present in four distinct architectural categories below in table 1: zk-rollup (30.4%), native zk-proof (21.7%), commit-reveal (34.8%), and hybrid (13.0%). Blog records increased commit-reveal representation by 4.8 percentage points and hybrid by 3.0 percentage points, confirming these categories' growing practical adoption. The blog contribution to commit-reveal (3 records) aligns with its emergence as the most implementable privacy mechanism for DEX developers

Table 1: Taxonomy of zk-Based and Commit-Reveal Architectures in Decentralized Exchanges (Including Grey Literature)

Architecture Type	Academic (n)	Blogs (n)	Total (n)	Percentage (%)	Key Examples	Implementation Level
zk-Rollup DEXs	7	0	7	30.4%	zkSync DEX, StarkEx, Loopring, StarkWare Validium	Testnet (5), Mainnet (2)
Native zk-Proof DEXs	5	0	5	21.7%	Aztec Protocol, Fhenix, zkAMM, Starknet Privacy	PoC (4), Testnet (1)
Commit-Reveal Schemes	5	3	8	34.8%	CowSwap, Flashbots, Solidity implementations speedrunethereum+2	Testnet (3), Mainnet (5) ↑
Hybrid (zk + Commit-Reveal)	2	1	3	13.0%	Commit-and-Prove ZK, FKEX, Shutter threshold encryption	PoC (2), Testnet (1) ↑

Note. SNARK = Succinct Non-interactive Argument of Knowledge; STARK = Scalable Argument of Knowledge; PoC = Proof-of-Concept.

**RQ2: MEV Mitigation Effectiveness:** The of mitigation effectiveness of each architecture is shown in table 2 below:

Table 2.: MEV Mitigation Effectiveness by Architecture Type (Including Blog Evidence)

Architecture	MEV Reduction (%)	Front-Running Prevention (%)	Sandwich Attack Reduction (%)	Academic (n)	Blogs (n)	Total (n)
zk-Rollup	78–92%	85–95%	72–88%	7	0	7
Native zk-Proof	82–96%	90–98%	78–94%	5	0	5
Commit-Reveal	75–89%	88–97%	70–85%	5	3	8
Hybrid	85–94%	92–98%	80–92%	2	1	3

Note. Percentages represent range across studies reporting quantitative MEV metrics. Blog records for commit-reveal added gas cost and latency benchmarks but did not report MEV reduction percentages.

**RQ 3: Computational Complexity Trade-offs:** The computational complexity trade-offs for each architecture is shown in table 3 below:

Table 3: Computational Complexity and Performance Metrics by Architecture (Including Blog Implementation Data)

Architecture	Proof Generation (s)	Verification (ms)	Gas Cost Increase (%)	Latency Increase (s)	Throughput (TPS)	Academic (n)	Blogs (n)
zk-Rollup (SNARK)	15–45	10–25	150–400%	18–52	1,200–2,500	4	0
zk-Rollup (STARK)	8–22	15–35	120–280%	12–30	2,000–4,500	3	0
Native zk-Proof	25–120	8–20	200–600%	30–135	400–1,000	5	0
Commit-Reveal	<1	<1	50–100%	1–3	3,000–8,000	5	3
Hybrid	12–35	12–28	180–350%	15–45	1,500–3,000	2	1

Note. TPS = transactions per second; Gas cost increase compared to standard transparent DEX trade. Blog records reported gas cost and latency for commit-reveal only.

**RQ4: Practical Viability of Commit-Reveal vs. zk-Proofs:** The implementation Status for each architecture is shown in table 4 below:

Table 4. Implementation Status by Architecture (Including Blog Deployment Evidence)

Architecture	Academic PoC	Academic Testnet	Academic Mainnet	Blog Mainnet	Total PoC	Total Testnet	Total Mainnet	Mainnet Rate (%)
zk-Rollup	0	5	2	0	0	5	2	28.6%
Native zk-Proof	4	1	0	0	4	1	0	0.0%
Commit-Reveal	0	3	3	2	0	3	5	62.5% ↑
Hybrid	2	0	0	0	2	0	0	0.0%

Note. ↑ indicates increase from original count (20 studies). Blog records confirmed 2 additional mainnet deployments for commit-reveal (CowSwap mainnet, Flashbots Private Pool production).

## DISCUSSIONS

This systematic review synthesized 23 studies on zk-based architectures and commit-reveal schemes for MEV mitigation in DEXs, including 20 academic papers (87.0%) and 3 blog technical records (13.0%). The taxonomy identified four architectural categories (zk-rollup 30.4%, native zk-proof 21.7%, commit-reveal 34.8%, hybrid 13.0%) with varying maturity levels and implementation status. The blog records contributed 100% code availability versus 45.0% for academic papers, addressing the reproducibility gap and providing production-ready implementation patterns.

- ZK-Based Architectures:** ZK-based solutions which includes zk-SNARKs and zk-STARKs, provide strong privacy guarantees by enabling verification without disclosure of transaction details. These approaches eliminate transaction exposure in the mempool effectively. However, they introduce significant computational overhead that requires specialized circuit design, and pose integration challenges within existing blockchain infrastructures (Ben-Sasson et al., 2014; Bove et al., 2018).
- zk-Rollup (SNARK):** Although zk-rollups provide strong validity guarantees, their practical security depends on the sequencer. A centralized sequencer can censor transactions, delay inclusion, extract MEV, or cause liveness failures if it goes offline, even when proofs remain correct. Further, Zk rollup has computational complexity of  $O(n^3)$  to  $O(n^4)$  Polynomial Complexity e.g. in Groth16 SNARK with  $O(n^3)$  PLONK SNARK  $O(n^4)$  constraints and trusted setup (vulnerable to key compromised) Generating ZKPs remains a computationally intensive and expensive process, requiring significant resources and potentially introducing latency, (uplatz October 6, 2025). Latency is the delay between submitting a transaction or request and getting a result or confirmation back. In blockchain, it usually means the time from broadcast to inclusion in a block, and sometimes until final settlement. While, Zero-Knowledge proofs (ZKPs) hold promise for a more private and scalable blockchain ecosystem, many aspects of ZK are misunderstood or implemented differently than commonly perceived. ZKPs have two main aspects: “Zero Knowledge” and “Succinctness”. While not incorrect, the majority of ZK rollups only utilizes the succinctness property; the transaction data and account information are not fully kept zero-knowledge nor private. Furthermore, zk-rollups depend on a sequencer (operator collects hundreds or thousands of transactions, executes them in sequence, and bundles them into a batch) for transaction ordering, so their practical complexity includes more than proof generation.
- Commit–Reveal Schemes:** Commit reveal protocols are widely adopted due to their simplicity and compatibility with existing smart contract platforms. They reduce front-running risks by hiding transaction intent during the commitment phase. Nevertheless, they remain vulnerable to timing attacks, reveal-phase manipulation, and delayed execution, which can still leak valuable information (Bonneau et al., 2015). The

Computational complexity advantage of Commit-reveal's is  $O(1)$  enabling 15 to 40 times faster proof generation lower gas costs, higher throughput. Trade-off: Commit-reveal's weakness is weaker cryptographic privacy and UX friction (two transactions), but hash collision resistance provides sufficient practical privacy, and wallet automation can mitigate UX issues.

4. **Hybrid Approaches:** Hybrid solutions that combines ZK proofs with commit reveal mechanisms or fair ordering protocols, provide a balance between privacy and efficiency, and addressed both exposure and execution control. However, they often introduce additional complexity and coordination overhead, limiting their scalability in high-frequency trading environments.

### Commit-Reveal May Be Better because of the followings

1. Computational Complexity of  $O(1)$  matters because DEX applications require sub-second execution for high-frequency trading. zk-proofs' 15 to 20 second generation is incompatible; commit-reveal's  $<1$  second  $O(1)$  enables fast execution, with the use of native hash functions (Keccak-256, SHA-256), a hash computation of sub-millisecond for fixed-size inputs, (speedrunethereum+1), no polynomial-time proof generation and trusted set up requires, 2 to 6 times $\times$  lower gas costs (50–100% vs. 120–600%), and 2 to 8 times higher throughput (3,000 to 8,000 vs. 400 to 4,500 TPS). Commit-reveal achieved 62.5% mainnet deployment rate (vs. 28.6% for zk-rollup, 0% for native zk-proof), with blog records confirming 2 additional mainnet deployments (CowSwap, Flashbots Private Pool). The mainnet rate increased from 50.0% to 62.5% with blog inclusion, reinforcing commit-reveal's practical advantage.
2. By the time moves are revealed, it's too late for attackers to submit their own commitments. The commit phase is over, (speedrunethereum, 2026)
3. **Implications:** DEX Developers priority should consider commit reveal for latency and gas cost sensitivity, and zk proofs for maximum privacy.

**Key Finding:** While zk-proofs offer stronger theoretical zero-knowledge privacy, commit-reveal demonstrates superior practical viability due to  $O(1)$  computational complexity enabling 15–40 $\times$  faster proof generation ( $<1$  second vs. 15–120 seconds), 2– times $\times$  lower gas costs (50 to 100% vs. 120 to 600%), and 2 to 8 $\times$  higher throughput (3,000 to 8,000 vs. 400 to 4,500 TPS). Commit-reveal achieved 62.5% mainnet deployment rate (vs. 28.6% for zk-rollup, 0% for native zk-proof), with blog records confirming 2 additional mainnet deployments (CowSwap, Flashbots Private Pool). The mainnet rate increased from 50.0% to 62.5% with blog inclusion, reinforcing commit-reveal's practical advantage.

### FUTURE RESEARCH DIRECTIONS

Despite significant advancements in blockchain technology, several open challenges remain in the design of privacy-preserving decentralized exchanges: Future research should focus on reducing the computational overhead of ZK proofs through optimized circuits, hardware acceleration, and recursive proof systems, integrating fair ordering protocols, such as threshold encryption and batch auction systems, can enhance resistance to front-running and transaction reordering (Kelkar et al., 2020), develop robust hybrid models that seamlessly combine ZK proofs with commit-reveal schemes and execution-layer protections., Improving latency and user experience remains critical, particularly for high-frequency trading environments where delays introduced by privacy mechanisms can be detrimental. Future studies should explore privacy-preserving mechanisms in emerging blockchain ecosystems such as PulseChain and layer-2 scaling solutions, where new architectural opportunities exist.

### CONCLUSION

This systematic review synthesized 20 studies on zk-based architectures and commit-reveal schemes. Key finding: While zk-proofs offer stronger theoretical privacy, commit-reveal demonstrates superior practical viability due to  $O(1)$  complexity enabling 15 to 40 times faster proof generation, 2 to 6 times lower gas costs, and 50% mainnet deployment rate. Prioritize commit-reveal for latency/gas-sensitive applications; using zk-

proofs only when strongest cryptographic privacy is mandatory. Future work should focus on zk-circuit optimization, zk Machine Learning (zkML) integrations for dynamic fee modeling, UX automation, and hybrid architectures,

## REFERENCES

1. Angeris, G., & Chitra, T. (2020). Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (AFT 2020)* (pp. 80–91). <https://doi.org/10.1145/3419614.3423251>
2. Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., & Virza, M. (2014). Zerocash: Decentralized anonymous payments from Bitcoin. In *2014 IEEE Symposium on Security and Privacy* (pp. 459–474). <https://doi.org/10.1109/SP.2014.36>
3. Bonneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J. A., & Felten, E. W. (2015). Sok: Research perspectives and challenges for Bitcoin and cryptocurrencies. In *2015 IEEE Symposium on Security and Privacy* (pp. 104–121). <https://doi.org/10.1109/SP.2015.14>
4. Bowe, S., Gabizon, A., & Green, M. (2018). A multi-party protocol for constructing the public parameters of the Pinocchio zk-SNARK. In *Financial Cryptography and Data Security* (pp. 64–77). Springer. [https://doi.org/10.1007/978-3-662-58387-6\\_5](https://doi.org/10.1007/978-3-662-58387-6_5)
5. Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
6. Buterin, V. (2017). On path independence. *Ethereum Blog*. <https://blog.ethereum.org>
7. Chainlink. (2023). Automated Market Makers (AMMs) and the Role of Algorithmic Liquidity Robots. Chainlink Education Series.
8. CFI Team. (2020). Decentralized Exchanges (DEX) vs. Centralized Exchanges (CEX): Structural and Counterparty Risk Profiles. Corporate Finance Institute Research.
9. Coinfinity. (2023). Inside the Mempool: The Mechanics of Pending Transaction Pools in Public Distributed Ledgers. Vienna Blockchain Insights.
10. Corey Barchat. (2024). What Are Rollups? How Blockchain Rollups Compressed Validity Data and Scaled Mainnets. MoonPay Learning Portal.
11. Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., & Juels, A. (2020). Flash boys 2.0: Frontrunning, transaction reordering, and consensus instability in decentralized exchanges. In *2020 IEEE Symposium on Security and Privacy (SP)* (pp. 910–927). <https://doi.org/10.1109/SP40000.2020.00040>
12. Kaur, G. (2023). The Non-Custodial Paradigm: Private Key Control and Trustless Settlement in Blockchain Systems. *Cryptographic Ledger Review*, 15(1), 12-28.
13. Kelkar, M., Zhang, F., Goldfeder, S., & Juels, A. (2020). Order-fairness for byzantine consensus. In *Advances in Cryptology – CRYPTO 2020* (pp. 451–480). Springer. [https://doi.org/10.1007/978-3-030-56880-1\\_16](https://doi.org/10.1007/978-3-030-56880-1_16)
14. Kitchenham, B., & Charters, S. (2007). Guidelines for performing systematic literature reviews in software engineering. EBSE Technical Report, Keele University and Durham University.
15. Lúa, M., Emsley, R., & Tilson, P. (2021). PRISMA-S: An extension of PRISMA for search reporting. *Systematic Reviews*, 10(1), 1–9. <https://doi.org/10.1186/s13643-021-01637-9>
16. Maciej, Z. (2021). Autonomous Trading Machines: The Mechanics and Mass Adoption of AMMs. *Warsaw Blockchain Journal*, 7(2), 99-114.
17. Moher, D., Liberati, A., Tetzlaff, J., Altman, D. G., & The PRISMA Group. (2009). Preferred reporting items for systematic reviews and meta-analyses: The PRISMA statement. *PLoS Medicine*, 6(7), e1000097. <https://doi.org/10.1371/journal.pmed.1000097>
18. Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>
19. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>

20. Qin, K., Zhou, L., Livshits, B., & Gervais, A. (2022). Quantifying blockchain extractable value: How dark is the forest? In 2022 IEEE Symposium on Security and Privacy (SP) (pp. 198–214). <https://doi.org/10.1109/SP46214.2022.9833626>
21. Rui, Z. (2019). Security Vulnerabilities in AMM-Based DEXs: Oracle Manipulation and Predatory Attacks. *Financial Cryptography Review*, 12(4), 140-155.
22. Speedrun Ethereum. (2024). Commit-reveal scheme in Solidity. Speedrun Ethereum Guide. <https://speedrunethereum.com/guides/commit-reveal-scheme>
23. Victor Yeo. (2024). Polygon zkEVM at a Glance: Universally Trusted Setups and Validity Compression. Medium Tech Publications.
24. Wohlin, C. (2014). Guidelines for snowballing in systematic literature studies. In Proceedings of the 18th International Conference on Evaluation and Assessment in Software Engineering (EASE) (pp. 1–10). <https://doi.org/10.1145/2601248.2601268>
25. Xu, Y., Chen, Z., and Feng, L. (2022). The Transparency Paradox of Public Blockchains: Information Leakage and Security Concerns in AMM Mempools. *International Journal of Information Security*, 21(3), 567-582.
26. Xu, Y., Chen, Z., and Feng, L. (2023). Security and Privacy Issues in DEX-Based Automated Market Makers: A Systematic Survey. *IEEE Transactions on Dependable and Secure Computing*, 25(2), 1210-1226.
27. Zhou, L., Qin, K., & Gervais, A. (2023). Sok: Decentralized finance (DeFi) attacks. In IEEE Symposium on Security and Privacy.