A Systematic Tertiary Review of Privacy in Permissionless Blockchains: Techniques, Limitations, and Future Trends

Saara Saaninkoski, Zheying Zhang and Timo Poranen

Tampere University, Finland

Abstract

Permissionless blockchain operates as a fully decentralized, transparent, and immutable ledger. Preserving privacy in such systems is a complex challenge, as privacy cannot rely on restricting access or deleting data. Bitcoin, the first application of blockchain technology, was initially praised as an anonymous digital currency, but the transactions on the network have been shown to be relatively easy to trace. This realization has led to the development of advanced privacy-enhancing mechanisms with stronger anonymity guarantees.

This article offers a comprehensive overview of privacy-preserving techniques for permissionless blockchain through a systematic tertiary review of existing surveys. It identifies and categorizes key techniques such as zero-knowledge proofs, ring signatures, homomorphic encryption, secure multi-party computation and decentralized mixing protocols. Their capabilities to mitigate risks of linkability and information leakage, as well as limitations like computational overhead, are examined. Furthermore, unresolved challenges and research interests in the field are analyzed. By consolidating fragmented insights into a coherent and accessible resource, this work aims to support the privacy-aware development and adoption of blockchain applications.

The findings highlight a fundamental trade-off between the privacy capabilities, efficiency, and trust assumptions of existing techniques. Privacy in permissionless blockchain often requires computationally complex cryptographic methods, leading to significant delays and increased costs for users. Efficiency can be improved by assuming some level of trust in entities or hardware, but this may conflict with the principle of decentralization. Although many powerful techniques exist, there is no universal solution, and the best results are achieved with combining techniques on a case-by-case basis.

Current research on permissionless blockchain privacy focuses on improving efficiency, interoperability and usability of privacy preserving techniques. Additionally, regulatory compliance and accountability are critical concerns, as the technology must comply with privacy regulations while preventing anonymity in illegal activities such as money laundering.

Keywords

Blockchain, Privacy, Decentralization, Anonymity, Trust

1. Introduction

Blockchain is a decentralized, tamper-resistant digital ledger technology that eliminates the need for a central authority. Its first real-world application, Bitcoin [1], demonstrated how a

TKTP 2025: Annual Doctoral Symposium of Computer Science, 2.—3.6.2025 Helsinki, Finland

saara.saaninkoski@tuni.fi (S. Saaninkoski); zheying.zhang@tuni.fi (Z. Zhang); timo.poranen@tuni.fi (T. Poranen)

0000-0002-6205-4210 (Z. Zhang); 0000-0002-4638-0243 (T. Poranen)

transparent peer-to-peer network could facilitate secure transactions without a trusted intermediary. Over time, additional platforms such as Ethereum [2] introduced smart contracts, expanding the utility of blockchain beyond digital currencies to various use cases, including decentralized finance (DeFi), supply chain management and identity systems [3].

Permissionless blockchains, also known as public blockchains, allow anyone to join, validate, and view the entire ledger, unlike private and consortium blockchains, which rely on a trusted authority to manage access [4]. Because no clear personal information, such as names, is stored on the ledger, cryptocurrencies like Bitcoin initially gained a reputation for anonymity. However, it soon became clear that these systems are not inherently anonymous. Although user identities are hidden behind cryptographic addresses, the open and immutable design of the blockchain presents significant privacy challenges.

A major challenge in blockchain privacy lies in linkability and traceability. Here, linkability refers to the ability to determine that a set of addresses or transactions all belong to the same user. Once linkability is established, it can enable traceability, which we define as the ability to follow the flow of funds or other on-chain events across multiple transactions.

Studies [5, 6] have shown that when multiple addresses sign a single transaction, it indicates that these addresses belong to the same user. These addresses can then be clustered, and the precision of this clustering can be further improved by identifying change addresses [7, 8] and analyzing user behavior patterns [9]. Attackers can also leverage off-chain information such as forum posts, IP logs, or social media data to link pseudonymous blockchain addresses to real-world identities [10, 11].

The public and permanent nature of transactions on permissionless blockchains creates a significant risk of sensitive information disclosure. For example, in Bitcoin, once a transaction is published, it becomes visible to everyone and cannot be deleted. This not only exposes financial details, but also complicates compliance with data protection regulations, such as the GDPR's principles of data minimization and the right to erasure. Therefore, robust privacy mechanisms are essential to meet regulatory requirements and maintain user trust, which are critical for ensuring the future viability of the technology.

However, improving privacy comes with various added costs. Running systems that incorporate advanced privacy mechanisms often requires substantial computational and storage resources, which can increase transaction fees and introduce delays, especially in computationally intensive smart contracts. Introducing trust assumptions can improve efficiency, but compromises decentralization [12].

These challenges highlight the complex trade-offs inherent in designing privacy-preserving permissionless blockchain systems. While research in this area is rapidly advancing, the growing body of literature has led to a fragmented understanding of the capabilities and limitations of various privacy-enhancing techniques. Therefore, to synthesize existing knowledge, guide future research, and promote privacy-aware blockchain development, this article addresses the following research questions:

RQ1. What are the main capabilities and limitations of current privacy-preserving techniques for permissionless blockchain, specifically regarding linkability and traceability, sensitive information disclosure, and the cost of privacy?

RQ2. What are the key research gaps in permissionless blockchain privacy, and what directions does the literature propose to address them?

We address these research questions through a systematic tertiary review of existing literature reviews on privacy in permissionless blockchain, synthesizing their findings on techniques and identifying open research questions. The remainder of this article is organized as follows: Section 2 describes the methodology and data sources, Section 3 provides a detailed discussion of the major privacy-preserving techniques and outlines key research gaps and directions for future work, Section 4 evaluates the limitations and validity of this study, and finally, Section 5 presents our conclusions and summarizes the contributions of this paper.

This paper is based on the first author's master's thesis research [13].

2. Methodology

There are several surveys that review privacy-preserving techniques for blockchain systems. To effectively leverage this existing body of surveys, a systematic tertiary review was conducted to summarize and synthesize the secondary studies addressing a common research question. The review adhered to the guidelines for systematic literature reviews and the tertiary reviews protocol established by Kitchenham and Charters [14].

The search process aimed to identify relevant conference proceedings and journal papers using carefully constructed search strings. These strings were developed iteratively from keywords in the research questions. They comprise three parts: a) "blockchain", b) keywords related to privacy preservation, and c) keywords specifically indicating survey and review studies. We used the logical operator ”AND” to connect these parts. The selected databases were IEEE Xplore, ACM Digital Library and Scopus, selected for their coverage and prominence in the field of computer science and technology. The search and selection process was conducted during week 24 in June 2023.

The inclusion criteria were as follows:

- The study provides an analysis or summary of existing techniques for privacy preservation in blockchain systems.
- The study is written in English.
- The study has undergone a peer-review process.

The exclusion criteria are as follows:

- The study does not focus on blockchain technology.
- The study does not address privacy preservation techniques.
- The study focuses on privacy-preserving techniques deployed with blockchain (as opposed to applied within blockchain)

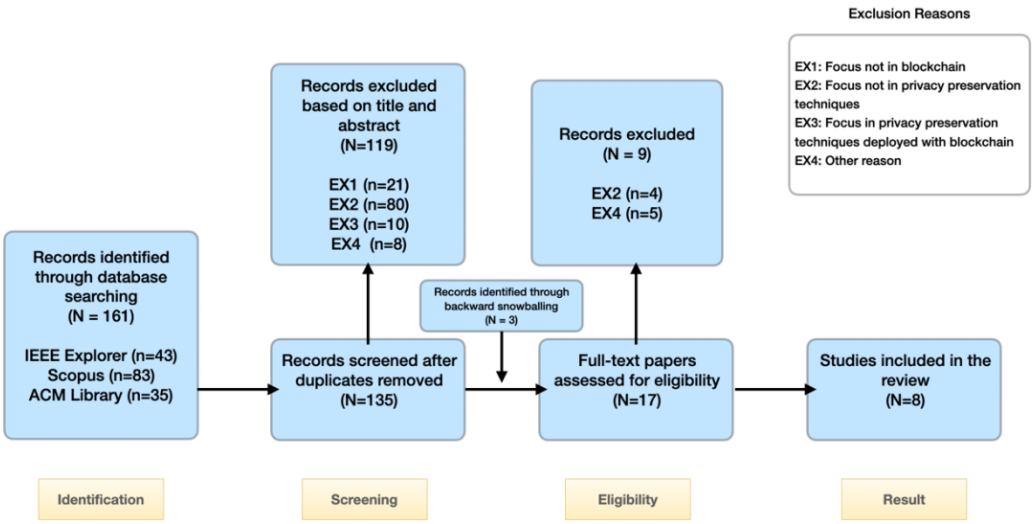

Figure 1: Flow diagram of the study selection process.

To increase the breadth of the review, we included studies encompassing both permissionless and permissioned blockchain, and those that did not differentiate between the two. The first author experimented with several search queries which include the three parts, aiming to identify a comprehensive range of studies. Some queries yielded a large number of results which, upon preliminary review, did not predominantly consist of literature reviews. To manage the workload effectively, only search results with fewer than 60 hits were used for the systematic screening process. The initial database search identified 161 results. From these studies, 119 were excluded based on title and abstract.

The screening process aimed to identify studies that met the inclusion and exclusion criteria, yielding a selection of 14 studies. In addition,the reference lists of the included studies were examined, leading to the identification of three more studies. A thorough review of the 17 selected papers was then conducted to ensure that only studies fulfilling the criteria were included. Furthermore, a set of questions derived from the quality assessment checklist [14] was applied to evaluate the quality and validity of the studies. These questions address aspects such as the adequacy of research process documentation and the presentation of the privacy techniques. The complete selection process is visualized in Figure 1.

In the quality assessment phase, nine papers were excluded from the review. Four of them were excluded due to insufficient coverage of different privacy techniques, as they focused solely on techniques used in specific settings, and four others did not primarily focus on privacy preservation techniques in blockchain.

Based on the evaluation criteria described above, the studies were scored on a scale from 0 to 5, with higher scores indicating greater quality and relevance. One paper was excluded because it scored zero points in the quality assessment. By following this systematic selection and

assessment process, the final set of eight studies were selected as a reliable and comprehensive basis for forming an extensive overview of the privacy preservation techniques in blockchain systems. Table 1 summarizes the selected studies in order of quality score, from highest to lowest. Each study is identified using the prefix SP as an identifier.

During the data extraction phase, information was gathered about the selected studies and the privacy preservation techniques they reviewed. The collected metadata included the title, year and source of publication, and names of authors. Key contributions from these studies were also documented to provide insights into each research paper's perspective to blockchain privacy preservation. Each paper was examined to identify and document the privacy preservation techniques it described. At this stage, all mentioned techniques were recorded without assessing their suitability for fully decentralized environments. For each technique, the following data was extracted:

- The name of the technique
- The high-level description of its operation
- The privacy capabilities of the technique
- The limitations inherent to the technique
- The primary source of the technique the review refers to.

The raw data generated during the tertiary review, including search queries, screened and assessed studies, quality assessment criteria, and extracted data are available at [15].

3. Results

This section presents the results of the systematic tertiary review. The first subsection details the capabilities and limitations of privacy-preserving techniques suitable for permissionless blockchain, addressing RQ1. The second subsection highlights the research gaps and future directions suggested in the reviewed studies, addressing RQ2. Finally, Subsection 3.3 summarizes and discusses the implications of the results.

3.1. Privacy-preserving Techniques for Permissionless Blockchain

Following data extraction as detailed in Section 2, the resulting dataset was used to catalog each technique referenced in the selected studies. Next, each technique was evaluated for its applicability to permissionless blockchain without trust assumptions.

Techniques were excluded from further analysis on the following criteria:

- Dependence on a trusted third party or trusted hardware: techniques requiring a trusted third party or trusted hardware for operation are excluded, except for those relying on the trusted creation of common reference strings (CRS).
- Lack of demonstrated applicability: techniques that have not been shown to be applicable to permissionless blockchain, either in theory or in practice, are excluded.

Table 1
Summary of the selected studies and their quality scores.

Paper ID	Authors	Title	Contributions	Score
[SP1]	Almashaqbeh and Solomon, 2022	SoK: Privacy-Preserving Computing in the Blockchain Era	SoK framework for zero-knowledge proof systems. SoK framework for existing privacy-preserving solutions.	4.5
[SP2]	Peng et al. 2021	Privacy Preservation in Permissionless Blockchain: A Survey	System model of permissionless blockchain Set of requirements for privacy preservation in blockchain. Analysis of existing privacy techniques in permissionless blockchain.	4
[SP3]	Junejo et al. 2021	Empirical Evaluation of Privacy Efficiency in Blockchain Networks: Review and Open Challenges	Classification of privacy solutions with respect to blockchain components. Critical evaluation of current privacy evaluation criteria. Framework for empirical evaluation of privacy solutions.	4
[SP4]	Feng et al. 2019	A Survey on Privacy Protection in Blockchain System	Novel definition of privacy in blockchain. Introduction of existing threats in user and transaction privacy. Introduction of existing privacy techniques in blockchain. Comparison of techniques in practice. Identification of future research directions.	3
[SP5]	Bernabe et al. 2019	Privacy-Preserving Solutions for Blockchain: Review and Challenges	Categorization of privacy challenges in blockchain. Systematic review and categorization of main privacy- preserving techniques for blockchain. Survey of privacy-preserving research proposals for blockchain. Analysis of privacy preserving IdM systems.	3
[SP6]	Zhang et al. 2019	Security and Privacy on Blockchain	Categorization of security and privacy attributes of blockchain. Review of security and privacy techniques for blockchain.	3
[SP7]	Satybaldy and Nowostawski, 2020	Review of techniques for privacy-preserving blockchain systems	Review of privacy-preserving techniques for blockchain. Analysis of techniques based on use cases. Framework for categorization of techniques. Introduction of scenarios for privacy-preserving mechanisms.	2.5
[SP8]	Wang et al. 2020	A Survey on Privacy Protection of Blockchain: The Technology and Application	Survey of existing privacy protection techniques for blockchain. Analysis of applications using these techniques. Exploration of future research.	2

Table 2
Summary of categories of privacy techniques

Name	Description	Capabilities	Limitations	Source	Surveys
Zero-knowledge proofs (ZKP)	Prove a statement without leaking anything about its content	Can prove e.g. that an input satisfies conditions, operations were successful Resistant to transaction graph analysis Prevents content leakage and transaction linkability	Computational and storage overhead	[16]	[SP1]-[SP8]
Ring signatures (RS)	Sign a message with a private key and N public keys to conceal the real signer	External and internal unlinkability Correct and unforgeable signatures Requires no trusted third party or cooperation Users can define custom groups	Limited anonymity set High cost and poor scalability Transaction content not hidden Difficult coordination Identity of signer can never be revealed; supervision is difficult	[17]	[SP1]-[SP8]
Homomorphic encryption (HE)	Encryption that allows computations on ciphertext	Hides transaction content Can enable private on-chain computation Resistant to collision and preimage attacks	Auditing not possible Limited operations Computational and communication overhead	[18]	[SP1], [SP3]-[SP8]
Secure multi-party computation (SMPC)	Jointly compute functions while keeping inputs private, preventing data exposure.	Input confidentiality with decentralized computation	Latency Difficult incentivization	[19]	[SP1]-[SP8]
Confidential Transactions (CT)	Bind a user to a value while keeping the value itself secret	Non-repudiation and non-modifiability while hiding value Efficient	Not alone sufficient for anonymity Unconditional hiding or binding, not both	[20]	[SP1]-[SP5], [SP7], [SP8]
Mixing	Obscure transaction history and unlink sender and recipient addresses	Can work over existing solutions Efficient	No transaction content protection Vulnerable to analysis	[21]	[SP2], [SP4]-[SP8]
Stealth addresses (SA)	Generate one-time address for each tx	Prevents linkability	Weak anonymity Sender not anonymous	[22]	[SP5], [SP8]

From the remaining pool of techniques, seven categories of privacy techniques could be identified. These categories are summarized in Table 2. Although the aggregated data provided a comprehensive overview of various privacy preservation techniques, the descriptions often lacked sufficient detail. For this reason, the primary source for each technique was reviewed to gain deeper insight, and the most frequently cited references (citation count in Google Scholar on 17.11.2024) are summarized in Table 3.

Table 3
Comparison of citations for primary sources of privacy techniques.

Technique	Category	Publication	Year	Citations
zk-SNARK	ZKP	Zerocash [23]	2014	2645
NIZKPoK	ZKP	Zerocoin [24]	2013	1447
Bulletproof	ZKP	Bünz et al. [25]	2013	1444
zk-STARK	ZKP	Ben-Sasson et al. [26]	2018	742
Confidential Transactions	CT	RingCT [27]	2015	669
Linkable ring signatures	RS	RingCT [27]	2015	669
Secure multi-party computations	SMPC	Andrychowicz et al. [28]	2016	592
Decentralized mixing	Mixing	CoinJoin [29]	2013	421
Partially homomorphic encryption	HE	Zether [30]	2020	408
Stealth Addresses	SA	CryptoNote [31]	2013	346
Traceable ring signatures	RS	CryptoNote [31]	2013	346
Fully homomorphic encryption	HE	smartFHE [32]	2023	41

The remainder of this subsection discusses the most prevalent privacy-preserving techniques, detailing their capabilities and limitations in mitigating the risks of linkability, traceability, and sensitive information disclosure, as well as the costs associated with their use. Table 4 provides a summary of this analysis.

Zero-knowledge proofs (ZKPs) [16] are cryptographic protocols that allow a party to prove a statement without revealing the underlying information. ZKPs can mitigate the risks of traceability and sensitive data exposure by concealing transaction details while still enabling the network to verify the transaction. Early blockchain implementations like Zerocoin [24] offered limited anonymity by hiding transaction origins but not destinations or amounts and suffered from poor performance. The introduction of zk-SNARKs in Zerocash [33] enhanced anonymity by hiding all participants and amounts, but this protocol requires a trusted setup phase. zk-STARKs [26] offer similar benefits to zk-SNARKs, without requiring trust assumptions. However, zk-STARKs produce large proofs, and their applications are still maturing. Meanwhile, Bulletproofs [25] both eliminate the need for trust assumptions and provide small proofs, though they suffer from slower execution. Projects like Zether [30] demonstrate that Bulletproofs can enable privacy-preserving and trustless smart contracts.

Ring signatures (RS) [17] conceal the address of a transaction signer within a set of potential signers, reducing the risk of linkability and traceability. The first blockchain implementation of RS was in CryptoNote [31], which uses modified traceable ring signatures to prevent double spending. While traceable ring signatures support auditability by enabling deanonymization of malicious participants, it imposes significant computational and storage overhead as the ring size increases. RingCT [27] builds on CryptoNote by applying a more efficient linkable RS

Table 4
Comparison of techniques based on threat mitigation.

Technique	Linkability and traceability	Sensitive data exposure	Cost
Confidential Transactions	No protection	Hides transaction amounts	Minimal
Stealth Addresses	Weak anonymity: vulnerable to transaction graph analysis	No protection	Minimal
Partially homomorphic encryption	No protection	Hides input data	Slow, not auditable
Fully homomorphic encryption	No protection	Hides input data	Slow, storage overhead, not auditable
NIZKPoK	Hides sender address	Hides transaction amounts	Relatively slow
zk-SNARK	Hides sender and recipient addresses	Hides transaction amounts	Trust assumptions
Bulletproof	Hides sender and recipient addresses	Hides transaction amounts	Slow
zk-STARK	Hides sender and recipient addresses	Hides transaction amounts	Large proofs
Traceable ring signatures	Conditional untraceability	No protection	Not auditable
Linkable ring signatures	Conditional unlinkablility	No protection	Computation and storage overhead
Secure multi-party computation	No protection	Hides input data	Slow, participant trust assumptions
Decentralized mixing protocols	Large anonymity set	No protection	Slow, often requires effort from user

protocol, but does not support auditability. Ring signatures do not inherently hide transaction amounts.

Stealth addresses (SA) [22], also applied in CryptoNote, create temporary intermediate addresses to break direct linkability. They are lightweight but remain vulnerable to transaction graph analysis.

Confidential transactions (CT) [20] use commitment schemes to hide transaction amounts while maintaining verifiability. Although CT alone does not prevent transaction linkability, it can be combined with other techniques, such as ring signatures. RingCT [27], for example, integrates CT into CryptoNote, simultaneously hiding identities and transaction amounts.

Homomorphic encryption (HE) [18] is a class of encryption schemes that allow certain computations to be performed directly on encrypted data, concealing transaction amounts and account balances. HE primarily protects from sensitive data disclosure and does not inherently prevent linkability. HE schemes are classified as partially homomorphic if they support only one type of operation, such as addition or multiplication, and as fully homomorphic (FHE) if they can support multiple types of operations. Partially homomorphic encryption schemes, such as

additive schemes, are suitable for financial contexts. Zether [30], for example, uses an additive HE scheme to allow confidential payments on Ethereum. FHE, on the other hand, enables running a wide range of private smart contract functions, as shown in smartFHE [32]. Despite the benefits, HE – and FHE in particular – tends to be computationally intesive and produce large ciphertexts, causing communication and storage overhead in blockchain. Moreover, it does not inherently support auditing, raising challenges for scenarios that demand accountability or legal traceability.

Secure Multi-Party Computation (SMPC) [19] allows multiple parties to jointly compute on private inputs without disclosing those inputs and only revealing the final outputs. In blockchain, SMPC can support trustless collaboration [28], facilitate private off-chain computations [34], enable private smart contract execution [35] and even replace the trusted third party needed for certain zk-SNARK setups [36]. SMPC effectively mitigates the risk of sensitive information disclosure but does not inherently prevent linkability or traceability. Furthermore, SMPC introduces significant computational and communication overhead, limiting the range of potential applications. Most protocols also assume a majority of honest participants – an assumption that can be challenging to uphold in permissionless blockchain settings.

Mixing is a technique that shuffles transaction inputs and outputs among multiple participants to obscure the original link between senders and receivers, and decentralized mixing protocols [21] achieve this without relying on a trusted third party. While mixing mitigates the risk of linkability and traceability, it does not inherently hide transaction contents. CoinJoin [29], the first decentralized mixing protocol for Bitcoin, allows users to form groups to merge their transactions into a single transaction. However, this approach can only provide external linkability, as transaction details are not hidden from other participants in a mixing group. Due to its self-organizing design, CoinJoin is also vulnerable to Sybil and DoS attacks. Several protocols have since been proposed to mitigate these limitations, including ValueShuffle [37], which extends CoinJoin, and CoinParty [38], which leverages SMPC. Despite these advancements, decentralized mixing protocols typically require a certain level of trust among participants and suffer from computational and communication overhead.

3.2. Research Directions for Permissionless Blockchain Privacy

The surveyed studies identified various research gaps and proposed future directions. To address RQ2, these were recorded separately for each study and tabulated.

The results were then categorized into four thematic groups: (1) performance and scalability, (2) regulatory compliance and accountability, (3) compatibility and interoperability, and (4) emerging technologies and future proofing. Table 5 compiles the research gaps and suggested directions grouped by theme.

Efficiency remains a key challenge for privacy-preserving techniques, with high computational and communication overhead frequently cited as a barrier to adoption and achieving sufficient anonymity. Four studies [SP2, SP3, SP5, SP7] suggested development of new more efficient techniques or improving the performance of existing ones. Two studies [SP1, SP4] emphasized customizing privacy approaches to specific use cases, either by optimizing combinations of techniques and configurations or designing custom solutions. Furthermore, moving heavy computations off-chain [SP1, SP8] was proposed as a practical approach to better per-

Table 5
Thematically grouped research gaps.

Category	Research gap	Future Direction
Performance and scalability	General efficiency	Tailor methods to specific use cases [SP1] Outsource ZKP proof generation [SP1] Develop more efficient decentralized privacy techniques [SP2], [SP5] Improve efficiency of current privacy techniques [SP7] Find optimally efficient and private combinations of techniques and configurations [SP4] Combine blockchain with trusted computing [SP8] Develop privacy protocols that do not create scalability issues in the network [SP3]
	Efficiency of trust-minimized schemes	New notions of trust; the concept of updateable trust [SP1] Optimize trust-minimized privacy techniques for performance [SP4]
	Efficiency of smart contract privacy	Develop efficient decentralized privacy methods for smart contracts [SP2] Conduct research and experiments using Solidity smart contracts [SP3]
Regulatory compliance and accountability	Surveillance against illegal activities	Design conditional privacy preservation methods [SP2], [SP4], [SP8]
	Privacy metrics	Develop common heuristics to enable comparison of privacy techniques [SP3]
	Usability and user control	Enhance privacy usability and control [SP5], [SP7]
	Regulatory compliance	Develop methods that ensure privacy as defined in regulations [SP5]
Compatibility and interoperability	Transaction structure compatibility	Develop privacy preservation methods compatible with diverse transaction structures [SP4]
	Interledger interoperability	Analyze the use of interledger techniques [SP5]
	Alignment with existing standards	Develop blockchain privacy preservation methods complying with current standards [SP5]
Emerging technologies and future proofing	Privacy for multi-user inputs	Employ multi-key FHE [SP1] Combine MPC and ZKP [SP1]
	Built-in privacy	Develop novel privacy-centric consensus mechanisms [SP3]
	Quantum computing resistance	Develop quantum-resistant methods [SP5]

formance. Inefficiencies in trust-minimized protocols [SP1, SP4] and those designed for smart contract applications [SP2, SP3] were also highlighted as areas of improvement.

Balancing privacy with the detection of prevention of illegal activities was another key theme, and three studies [SP2, SP4, SP8] suggested exploration of conditional privacy mechanisms. Improving regulatory compliance was also highlighted, explicitly in one study [SP5] and indirectly in two [SP5, SP7]. Indirect ways to improve regulatory compliance included improving the usability and user control of privacy solutions. Additionally, one study [SP3] suggested developing metrics to facilitate the comparison of privacy techniques.

As blockchain systems diversify, ensuring compatibility across transaction models and platforms has become more important. One study [SP4] noted the dominance of privacy methods designed for Bitcoin's UTXO model and called for research into compatibility with other models. Another [SP5] suggested that techniques supporting interledger interoperability should be analyzed. To further enhance compatibility, it was also suggested [SP5] that privacy solutions be designed to align with existing standards.

To address long-term challenges, one study [SP1] called for research on multi-user input privacy, suggesting FHE and MPC-based techniques as potential solutions. One proposal [SP3] involved incorporating privacy directly into consensus mechanisms, and another one [SP5] stressed the importance of quantum computing resistance.

3.3. Summary

The systematic tertiary review identified several privacy-preserving techniques applicable to permissionless blockchain, which were categorized into seven groups based on their methods of operation and privacy objectives: zero-knowledge proofs (ZKP), ring signatures (RS), stealth addresses (SA), confidential transactions (CT), homomorphic encryption (HE), secure multi-party computation (SMPC) and decentralized mixing protocols.

First, the high-level capabilities and limitations of each category of techniques were inspected and then each technique was individually analyzed. The analysis involved assessing how the technique mitigates the risks of linkability, traceability, and sensitive data exposure, as well as the affiliated performance or storage cost, and trust-assumptions.

Stealth addresses, ring signatures and decentralized mixing protocols all focus on mitigating the risk of linkability and traceability. Stealth addresses entail minimal overhead but provide limited anonymity on their own. Decentralized mixing protocols can support large anonymity sets but often requires user participation and involves significant transaction delays. Ring signatures can provide conditional anonymity, allowing privacy while facilitating some level of auditability of transactions.

Confidential transactions, homomorphic encryption and secure multi-party computation primarily aim to protect sensitive data by concealing transaction amounts or input values. Confidential transactions can efficiently hide transaction amounts and can be combined with methods that target linkability. Homomorphic encryption can be used for more advanced applications, such as smart contracts, but its use limits auditability and introduces delays and storage overhead. SMPC also hides inputs, but relies on trust assumptions among participants.

ZKPs were the only category of techniques that offers both unlinkability and sensitive data protection. Based on citation count, ZKPs are also the most prevalent technique for blockchain

privacy. Despite their versatility, ZKPs face limitations with proof size, verification speed, and trust assumptions affiliated with the setup of some protocols.

Since most privacy-preserving techniques focus either on breaking linkability or hiding sensitive data, combining multiple approaches proves effective for practical use cases. For example, RingCT combines ring signatures with confidential transactions and Zether homomorphic encryption with ZKPs.

This review also recognized several unresolved research gaps outlined in the surveyed studies, which were categorized into four themes: performance and scalability, regulatory compliance and accountability, compatibility and interoperability, and emerging technologies and future proofing.

Potential solutions range from tailoring protocols to specific use cases, outsourcing ZKP proof generation, and combining blockchain with trusted computing, to designing novel consensus mechanisms and ensuring quantum resistance. Among these themes, performance and scalability – particularly the computational overhead introduced by cryptographic privacy solutions – consistently emerge as the most pressing challenge for permissionless blockchain privacy.

4. Evaluation of Methodology and Limitations

The methodology employed in this research was a tertiary review that focused on eight secondary studies published between 2019 and 2022, each discussing privacy-preserving techniques for permissionless blockchain systems. While this approach offered a comprehensive overview of established findings, it also presented certain limitations.

First, relying on a tertiary design meant drawing conclusions from a relatively small set of secondary studies rather than directly from primary research. Although some primary sources were revisited to mitigate bias, it remains possible that significant advancements published after these secondary studies were not captured.

Second, the literature search strategy, which aimed to address privacy in permissionless blockchains broadly, may not have been exhaustive. Specifically excluding keywords such as "Bitcoin" or "cryptocurrency" might have inadvertently omitted relevant works that examine privacy solutions within those contexts, even though such research can be generalizable to permissionless blockchains.

Lastly, the scope was intentionally broad to accommodate a wide array of privacy techniques. Evaluating the level of decentralization and trust assumptions for each technique proved challenging, given the variations in definition and implementation. Consequently, some nuanced aspects of each privacy-enhancing method may not have been analyzed in depth.

5. Conclusions

Preserving privacy in permissionless blockchain systems is challenging due to the technology's inherently transparent and immutable design. Through a review of surveys, this article has shown that privacy solutions span cryptographic methods like zero-knowledge proofs and ring signatures, as well as private computation and encryption protocols. While many techniques

show promising applicability to permissionless blockchain, there is no single comprehensive solution - rather, multiple techniques must be combined as use cases demand.

Efficiency remains a major obstacle, as large anonymity sets often introduce additional latency and cost. Furthermore, compatibility gaps across different blockchain platforms and regulatory concerns around illicit activities underscore the importance of accountability features. Developing scalable, interoperable, and compliant privacy solutions will be crucial for ensuring ethical use and sustained viability of permissionless blockchain applications.

Generative AI Declaration

During the preparation of this manuscript, the authors used Overleaf's Writefull, an AI-assisted editing tool for language improvement, specifically to improve spelling, word choice, readability and clarity. The authors have reviewed any AI-edited content as necessary and assume full responsibility for the accuracy and integrity of the published work.

References

[1] S. Nakamoto, Bitcoin: A peer-to-peer electronic cash system, 2009. Cryptography Mailing list at https://metzdowd.com, Mar. 2009, Accessed: Mar. 14, 2024. Available: http://www.bitcoin.org/bitcoin.pdf.

[2] V. Buterin, Ethereum whitepaper, 2024. Accessed: Apr. 20, 2024. [Online]. Available: https://ethereum.org.

[3] I. Bashir, Mastering blockchain, Packt Publishing Ltd, 2017.

[4] Z. Zheng, S. Xie, H. Dai, X. Chen, H. Wang, An overview of blockchain technology: Architecture, consensus, and future trends, in: 2017 IEEE International Congress on Big Data (BigData Congress), IEEE, 2017, pp. 557–564.

[5] A. Pfitzmann, M. Hansen, A terminology for talking about privacy by data minimization: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management, 2010. Available: https://dud.inf.tu-dresden.de/Anon_Terminology.shtml.

[6] F. Reid, M. Harrigan, An analysis of anonymity in the bitcoin system, in: Security and Privacy in Social Networks, Springer, 2012, pp. 197–223.

[7] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, S. Savage, A fistful of bitcoins: characterizing payments among men with no names, in: Proceedings of the 2013 Conference on Internet Measurement Conference, 2013, pp. 127–140.

[8] M. Spagnuolo, F. Maggi, S. Zanero, Bitiodine: Extracting intelligence from the bitcoin network, in: International Conference on Financial Cryptography and Data Security, Springer, 2014, pp. 457–468.

[9] E. Androulaki, G. O. Karame, M. Roeschlin, T. Scherer, S. Capkun, Evaluating user privacy in bitcoin, in: Financial Cryptography and Data Security: 17th International Conference, FC 2013, Okinawa, Japan, April 1-5, 2013, Revised Selected Papers 17, Springer, 2013, pp. 34–51.

[10] M. Fleder, M. S. Kester, S. Pillai, Bitcoin transaction graph analysis, arXiv preprint arXiv:1502.01657 (2015).

[11] S. Goldfeder, H. Kalodner, D. Reisman, A. Narayanan, When the cookie meets the blockchain: Privacy risks of web payments via cryptocurrencies, Proceedings on Privacy Enhancing Technologies 2018 (2017). doi:10.1515/popets-2018-0038.
[12] M. Kokaras, M. Foti, The cost of privacy on blockchain: A study on sealed-bid auctions, Blockchain: Research and Applications 4 (2023) 100133.
[13] S. Saaninkoski, Enhancing decentralized privacy : A systematic review of techniques for permissionless blockchain, Master's thesis, Tampere University, 47 pages, available at: https://trepo.tuni.fi/handle/10024/162319 (2024).
[14] B. Kitchenham, S. Charters, et al., Guidelines for performing systematic literature reviews in software engineering version 2.3, Engineering 45 (2007) 1051.
[15] S. Saaninkoski, Tertiary Review Dataset for Permissionless Blockchain Privacy, https://1drv.ms/x/s!AvIGqZvmedFbdShkQBCeNQqDrgM?e=rfBQM3, 2024. Accessed: 2025-04-15.
[16] S. Goldwasser, S. Micali, C. Rackoff, The knowledge complexity of interactive proof-systems, in: Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, 2019, pp. 203–225.
[17] R. L. Rivest, A. Shamir, Y. Tauman, How to leak a secret, in: Advances in Cryptology—ASIACRYPT 2001: 7th International Conference on the Theory and Application of Cryptology and Information Security Gold Coast, Australia, December 9–13, 2001 Proceedings 7, Springer, 2001, pp. 552–565.
[18] R. L. Rivest, L. Adleman, M. L. Dertouzos, et al., On data banks and privacy homomorphisms, Foundations of Secure Computation 4 (1978) 169–180.
[19] A. C. Yao, Protocols for secure computations, in: 23rd Annual Symposium on Foundations of Computer Science (SFCS 1982), IEEE, 1982, pp. 160–164.
[20] G. Maxwell, Confidential transactions, 2015. Accessed: Jul. 09, 2024. [Online]. Available: https://nt4tn.net/tech-notes/201505.confidential_values.txt.
[21] D. L. Chaum, Untraceable electronic mail, return addresses, and digital pseudonyms, Communications of the ACM 24 (1981) 84–90.
[22] U. ByteCoin, Untraceable transactions which can contain a secure message are inevitable, Bitcoin Forum. Accessed: Aug. 30, 2024. Available: https://bitcointalk.org/index.php?topic=5965.0, 2011.
[23] E. B. Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, M. Virza, Zerocash: Decentralized anonymous payments from bitcoin, in: 2014 IEEE Symposium on Security and Privacy, IEEE, 2014, pp. 459–474.
[24] I. Miers, C. Garman, M. Green, A. D. Rubin, Zerocoin: Anonymous distributed e-cash from bitcoin, in: 2013 IEEE Symposium on Security and Privacy, IEEE, 2013, pp. 397–411.
[25] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, G. Maxwell, Bulletproofs: Short proofs for confidential transactions and more, in: 2018 IEEE Symposium on Security and Privacy, IEEE, 2018, pp. 315–334.
[26] E. Ben-Sasson, I. Bentov, Y. Horesh, M. Riabzev, Scalable, transparent, and post-quantum secure computational integrity, Cryptology ePrint Archive (2018).
[27] S. Noether, A. Mackenzie, et al., Ring confidential transactions, Ledger 1 (2016) 1–18.
[28] M. Andrychowicz, S. Dziembowski, D. Malinowski, Ł. Mazurek, Secure multiparty computations on bitcoin, in: 2014 IEEE Symposium on Security and Privacy, 2014, pp. 443–458.

[29] G. Maxwell, Coinjoin: Bitcoin privacy for the real world, 2024. URL: https://bitcointalk.org/index.php?topic=279249.0.
[30] B. Bünz, S. Agrawal, M. Zamani, D. Boneh, Zether: Towards privacy in a smart contract world, in: International Conference on Financial Cryptography and Data Security, Springer, 2020, pp. 423–443.
[31] N. Van Saberhagen, Cryptonote v 2.0, 2013.
[32] R. Solomon, R. Weber, G. Almashaqbeh, smartFHE: Privacy-preserving smart contracts from fully homomorphic encryption, in: 2023 IEEE 8th European Symposium on Security and Privacy (euroS&p), IEEE, 2023, pp. 309–331.
[33] N. Bitansky, R. Canetti, A. Chiesa, E. Tromer, From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again, in: Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, 2012, pp. 326–349.
[34] G. Zyskind, A. Pentland, Enigma: Decentralized computation platform with guaranteed privacy, in: New Solutions for Cybersecurity, The MIT Press, 2018, pp. 425–456.
[35] Y. Zhu, X. Song, S. Yang, Y. Qin, Q. Zhou, Secure smart contract system built on smpc over blockchain, in: 2018 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData), IEEE, 2018, pp. 1539–1544.
[36] S. Bowe, A. Gabizon, M. D. Green, A multi-party protocol for constructing the public parameters of the pinocchio zk-snark, in: International Conference on Financial Cryptography and Data Security, Springer, 2019, pp. 64–77.
[37] T. Ruffing, P. Moreno-Sanchez, Valueshuffle: Mixing confidential transactions for comprehensive transaction privacy in bitcoin, Lecture Notes in Computer Science (2017) 133–154.
[38] J. H. Ziegeldorf, F. Grossmann, M. Henze, N. Inden, K. Wehrle, Coinparty: Secure multi-party mixing of bitcoins, in: Proceedings of the 5th ACM Conference on Data and Application Security and Privacy, 2015, pp. 75–86.