UNIVERSITÉ DU LUXEMBOURG

PhD-FSTM-2025-116
The Faculty of Science, Technology and Medicine

DISSERTATION

Defence held on 24 November 2025 in Luxembourg

to obtain the degree of

DOCTEUR DE L'UNIVERSITÉ DU LUXEMBOURG

EN INFORMATIQUE

by

Egor ERMOLAEV
Born on 1 November 1995 in Omsk, Russia

ZERO-KNOWLEDGE PROOFS AND BLOCKCHAIN: APPLYING TECHNOLOGY TO STRIKE A BALANCE BETWEEN PRIVACY AND TRANSPARENCY

Dissertation defence committee

Dr. Gilbert FRIDGEN, Supervisor
Professor, University of Luxembourg

Dr. Marcus VÖLP, Chairman
Professor, University of Luxembourg

Dr. Radu STATE, Member
Professor, University of Luxembourg

Dr. Robert KELLER, Member
Professor, University of Applied Sciences, Kempten, Germany

Dr. Daniel MIEHLE, Member
BMW Group, Munich, Germany

Abstract

A system design built on blockchain technology presents a fundamental challenge: the inherent transparency of the blockchain conflicts with the growing need for user privacy. This dissertation explores how Zero-Knowledge Proofs (ZKPs) can be strategically combined with blockchain to strike a balance between these competing demands. The dissertation analyzes the challenges of privacy and transparency and provides an overview of solutions across the privacy-transparency solution space, drawing from the author's original research and the broader academic landscape.

On the privacy-centric side, it proposes a privacy-preserving design that utilizes off-chain ZKPs. In contrast, on the transparency-centric side, it presents a transparency-enhancing design that leverages on-chain data to build trust. In the middle, it discusses a taxonomy of hybrid applications whose unique combination of on-chain privacy (via ZKPs) and blockchain transparency reveals both a disruptive potential and significant regulatory challenges.

Acknowledgment

First and foremost, I wish to express my sincere gratitude to my supervisor, Prof. Dr. Gilbert Fridgen, for recognizing my potential, guiding me throughout my PhD journey, providing crucial funding, and offering me the freedom to conduct my own research – at times, perhaps more than I thought I could handle. I am also grateful for the incredible opportunities to travel the world; it was a privilege to grow professionally in your interdisciplinary research group.

I extend my heartfelt thanks to my co-authors, whose contributions were instrumental in preparing this cumulative dissertation. I am especially grateful to Prof. Dr. Dr. Tamara Roth, Prof. Dr. Johannes Sedlmeir, Dr. Orestis Papageorgiou, Evgenia Yvonni Tseloni, and Iván Abellán Álvarez. My thanks also go to my CET members, Prof. Dr. Marcus Völp and Dr. Daniel Miehle, and to all the other colleagues from the FINATRAX research group including those who have since graduated, in particular Dr. Sergio Potenciano Menci.

On a personal note, I am deeply grateful to my family: my mom, Svetlana Ermolaeva; my uncle, Valery Ermolaev; my cousin, Igor Ermolaev; and my uncle Alexander Ermolaev and my aunt Galina Ermolaeva. I also wish to honor the memory of my grandparents, Yuriy Ermolaev and Valentina Ermolaeva, who supported my childhood with their pensions.

My heartfelt thanks go to Maryna Chepeleva for her support, which was especially crucial during the final years of my PhD. I wish you a speedy and successful completion of your own.

Finally, I wish to express my gratitude to all my teachers and professors. I am especially grateful to my former supervisor, Prof. Dr. Alexander Shiler, who is sadly no longer with us; to my first legit English teacher, Irina Lavrinenko; and to Prof. Dr. Anton Fedosov, for my first paper. I am also thankful to all my friends around the world, with a special

mention for those in ice-covered Siberia, where the ruthless environment made us strong while leaving its mark.

This research was funded by the Luxembourg National Research Fund (FNR), in the NCER22/IS/16570468/NCER-FT project, and by PayPal, PEARL grant reference 13342933/Gilbert Fridgen.

Declaration

I, Egor Ermolaev, declare that I wrote this thesis by myself and that it has not been previously submitted for any other degree or professional qualification. For the jointly authored research papers, I have clearly distinguished my contributions from those of the other authors. I confirm that I have correctly referenced the work of others whenever I refer to it.

During the preparation of this work I used Gemini in order to improve the language and clarity of parts of my work. After using this tool, I reviewed and edited the content as needed.

I confirm that I have no financial interests to disclose in relation to this research. I commit to conducting my research with transparency, integrity, and adherence to ethical principles, ensuring the reliability of my work.

"Don't believe, don't fear, don't ask."

Varlam Shalamov (1907 – 1982)

Table of Contents

List of Figures

List of Tables

I | Introduction

1 Motivation

The core feature of blockchains is their high degree of transparency, a property that becomes particularly powerful – and problematic (Sedlmeir et al., 2022a) – with the advent of smart contracts on blockchains like Ethereum (Buterin, 2014). While this transparency offers significant advantages, it creates a fundamental tension with the critical need for user privacy (Schmeiss et al., 2019; Sedlmeir et al., 2022a). The public and immutable nature of the ledger poses immediate challenges for compliance with regulations like the General Data Protection Regulation (GDPR) (Ermolaev et al., 2025; Herian, 2020; Politou et al., 2021; Sedlmeir et al., 2022a), as the pseudonymity offered by cryptographic addresses is vulnerable to de-anonymization through blockchain analysis (Kappos et al., 2018; McGinn et al., 2018; Meiklejohn et al., 2013; Werner et al., 2020). This conflict presents a major barrier to adoption in domains where sensitive data is paramount.

A technology particularly well-suited to address this conflict is the Zero-Knowledge Proof, a powerful cryptographic protocol first developed in the last century (Fiege et al., 1987; Goldwasser et al., 1989). Zero-Knowledge Proofs allow one party to prove the truth of a statement without revealing any information beyond the validity of the statement itself, offering a paradigm of "verifiable privacy" (Walfish and Blumberg, 2015). Modern implementations, particularly succinct variants known as non-interactive Zero-Knowledge Proofs (NIZKPs), are efficient enough to be applied directly to blockchains. They can be used at the protocol level for shielded transactions (e.g., Zcash) or at the application level as "Layer-2" scaling solutions like zk-rollups (Hopwood et al., 2016; Thibault et al., 2022).

The true potential for the next generation of decentralized applications lies not in these technologies in isolation, but in their synergistic combination. This powerful fusion is already enabling a new wave of systems that balance privacy with transparency, including compliant Central Bank Digital Currency (CBDC) (Gross et al., 2021), transparent green electricity labeling systems (Sedlmeir et al., 2022b), privacy-preserving stablecoins (Gross et al., 2022), and crypto-mixers that prevent transaction traceability (Barbereau et al., 2023b).

This dissertation argues that the strategic architectural integration of blockchain's transparency with ZKP-driven privacy provides the necessary toolkit to resolve the core tension. While research on blockchain and ZKP applications continues to advance (detailed in Recognition of previous and related work), a focused analysis of these architectural synergies is needed to overcome persistent challenges and find the right balance. Therefore, this dissertation investigates system architectures from the privacy-transparency solution space (Figure III.1) that strategically combine on-chain transparency with ZKP-enforced privacy.

Before outlining the nature and structure of this cumulative dissertation, the following section first situates this research within the existing academic landscape.

2 Recognition of previous and related work

Progress in this field is cumulative. In that spirit, this dissertation situates its contribution within prior research on decentralized architectures, digital identity, market design, and the performance characteristics of distributed ledgers. While the focus is on practical mechanisms such as privacy-preserving designs and auditability, this work also relies on foundational and sector-specific research to present a coherent view of the design space.

Specifically, it draws on foundational work that introduces the core concepts of blockchain and smart contracts (Schlatt et al., 2016; Schütte et al., 2018), as well as analyses of tokenized fundraising and market architectures (Arnold et al., 2019; Bachmann et al., 2022; Fridgen et al., 2018c). Regarding governance, it also leverages research on the diffusion of decentralization technologies in finance (Fridgen et al., 2024a; Fridgen et al., 2024b; Hartwich et al., 2024) and on the power dynamics of token voting in DeFi (Barbereau

et al., 2022b; Barbereau et al., 2023a). Together, these works frame the institutional and technical backdrop against which privacy and transparency mechanisms must operate.

A central thread for this dissertation is the balance between transparency and privacy. This work, therefore, draws on two streams of literature: foundational studies that explore transparency as a mechanism for fostering trust and accountability (Utz et al., 2023), and contemporary research that examines transparency as a significant organizational challenge (Sedlmeir et al., 2022a).

Building on this, the dissertation also relies on studies that operationalize GDPR-compliant blockchain design in cross-organizational settings (Guggenmos et al., 2020; Rieger et al., 2019; Rieger et al., 2021). Complementary work on digital identity and verifiable credentials informs the discussion of selective disclosure and accountability (Glöckler et al., 2024; Höß et al., 03 January 2022; Rieger et al., 2022; Sedlmeir et al., 2021; Weigl et al., 2022; Weigl et al., January 2022). Beyond identity, this research also draws from cryptography-enabled techniques for misinformation countermeasures and privacy-preserving risk analysis (Fridgen and Garizy, 2015; Sedlmeir et al., 2023; Zare-Garizy et al., 2018). Finally, work on tokenization and regulatory compliance highlights the practical need to reconcile the needs of regulators and investors, which is closely related to the privacy-transparency focus of this dissertation (Barbereau et al., 2022a).

Because many privacy and transparency requirements arise in inter-organizational contexts, this work is also informed by public-sector blockchain research covering workflow automation and accountability in government services. This includes the FLORA program and related asylum-process studies (Amend et al., 2022; Amend et al., 2021a; Amend et al., 2021b; Fridgen et al., 2018a; Fridgen et al., 2019; Fridgen et al., 2018b; Guggenmos et al., 2019), and the EBSILUX project (Höß et al., 2022a) – a diploma use case for the European Blockchain Services Infrastructure (EBSI). These cases ground the discussion of verifiability, minimal disclosure, and governance across institutional boundaries.

3 Nature of the Dissertation

This cumulative dissertation synthesizes the findings from five peer-reviewed research articles, which are summarized in Table I.1. A detailed portfolio of the author's individual

contribution to each article, along with the full reprints, is provided in the Appendix, beginning with Chapter VI.

Table I.1: Overview of Research Articles.

RA#	Title	Role[1]
RA1	z-Commerce: Designing a Data-Minimizing One-Click Checkout Solution	[SP]
RA2	Beyond a Fistful of Tumblers: Toward a Taxonomy of Ethereum-based Mixers	[JP]
RA3	Designing a Reporting System for Trust in Environmental Social Governance	[SP]
RA4	What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains	[NP]
RA5	Blockchain-based Reporting System for Trust in Environmental Social Governance*	[JP]

[1]SP = Single Primary Author, JP = Joint Primary Author, NP = Non-Primary Author.

*Under review at a peer-reviewed journal (as of 24 November 2025).

A common methodological thread uniting the core research articles in this dissertation is the Design Science Research (DSR) (Hevner et al., 2004). This approach was particularly well-suited for the interdisciplinary nature of the research, which bridges computer science, information systems, and regulatory compliance. By providing a structured process for designing and evaluating IT artifacts, DSR offered a common framework and terminology that facilitated effective collaboration among co-authors from diverse backgrounds, including software engineering, law, and many others. Ultimately, the DSR approach ensured that the architectural solutions presented were not only technically sound but also grounded in real-world problems and stakeholder needs.

4 Structure of the Dissertation

This dissertation is structured into four main chapters that guide the reader from the foundational challenges to the design of solutions and the conclusion.

- Chapter I, "Introduction," motivates the central topic of balance between privacy and transparency in the technological scope of blockchain and ZKPs, situates the work

within the context of related academic literature, and outlines the cumulative nature and overall structure of the dissertation.

- Chapter II, "The Challenges of Privacy and Transparency," deepens the analysis of the core tension identified in the motivation. It explores the limitations of radical transparency, examines the practical difficulties of achieving Privacy by Design, and provides an overview of related Privacy-Enhancing Technologies (PETs), including Zero-Knowledge Proofs.

- Chapter III, "Designing Solutions to Balance Privacy and Transparency," moves from challenges to detailed system architecture designs. It presents three distinct archetypes – one prioritizing privacy, one prioritizing transparency, and a third navigating the complex hybrid space – to demonstrate how different balances can be struck across the privacy-transparency solution space.

- Chapter IV "Conclusion" summarizes the key findings on how different architectures strike a balance between privacy and transparency using blockchain and ZKPs, discusses the limitations of this research, and outlook.

II The Challenges of Privacy and Transparency

This chapter explores the core dilemma of blockchain technology: the tension between its inherent transparency and the need for user privacy. It first examines the challenges posed by radical transparency in blockchain systems. It then analyzes the complexities of achieving Privacy by Design (PbD) and relevant Privacy-Enhancing Technologies (PETs). Finally, the chapter synthesizes these findings to argue for a nuanced architectural approach to balancing these competing demands.

1 Challenges of Radical Transparency

The definition of transparency is multifaceted. The Cambridge Dictionary offers general definitions such as "the characteristic of being easy to see through" and a business-oriented one: "a situation in which business and financial activities are done in an open way without secrets, so that people can trust that they are fair and honest". Notably, the dictionary even provides a use case that directly highlights its tension with privacy: "we need to strike balance between the need for transparency and respect for individual privacy" (Cambridge University Press, 2025b). Transparency is subtly intertwined with accountability, encourages openness, and heightens concerns for secrecy and privacy (Ball, 2009; Meijer, 2009). Building on these concepts, transparency in the context of this dissertation is an architectural property afforded by the use of blockchain technology, where all transactions are immutably recorded and publicly accessible – creating a verifiable and auditable system designed to foster trust and accountability among participants.

Indeed, the default state of public permissionless blockchains is one of transparency, not privacy (Yli-Huumo et al., 2016). While this transparency can foster trust and accountability (Ermolaev et al., 2025; Ermolaev et al., 2026; Whyte and Macintosh, 2001), it also poses significant risks of surveillance and control when applied without safeguards. The immutable and public nature of the permissionless blockchains creates an excessive degree of transparency that challenges individual autonomy and introduces systemic risks and directly contradicts with regulations like GDPR, such as "right to be forgotten" and beyond (Sedlmeir et al., 2022a). This section delves into these challenges by exploring the "chilling effect" of constant surveillance, the direct conflict with data protection regulations like the GDPR, and the mechanisms through which transparency enables user de-anonymization.

1.1. The Chilling Effect

The constant potential for monitoring on public blockchains can lead to a "chilling effect," where individuals avoid legal yet sensitive activities for fear of being watched (Murray et al., 2023; Nguyen, 2021). This extends into tangible risks in domains like e-commerce, where public transaction histories facilitate behavioral profiling and unfair practices such as price discrimination (Ermolaev et al., 2023), and in healthcare, where the prospect of on-chain health records could deter individuals from seeking care for stigmatized conditions (Liang et al., 2023; Rodriguez-Garcia et al., 2021). This dynamic also creates significant power asymmetries, as radical transparency disproportionately benefits powerful actors – such as corporations and state agencies – that possess the resources to analyze the vast datasets generated by blockchains (Clarke, 2019).

1.2. Conflict with Data Protection Regulations

The architectural properties of public blockchains create a direct conflict with key principles of modern data protection regulations, most notably the European Union (EU)'s GDPR. Three blockchain characteristics are particularly problematic: immutability, public replication, and public readability. The GDPR's "right to erasure" (Article 17), also known as the "right to be forgotten," is fundamentally incompatible with an immutable ledger where data, once written, cannot be deleted (Sedlmeir et al., 2022a). Furthermore, the principle of data minimization (Article 5), which mandates that only necessary data

be processed, is challenged by the public replication of the entire transaction history across all nodes in the network (Cruz, 2019; Herian, 2020).

To mitigate these conflicts, a common architectural pattern has emerged: storing sensitive Personally Identifiable Information off-chain (Guggenmos et al., 2020). In this model, the blockchain is used only to store immutable references – such as cryptographic hashes or ZKPs – while the actual Personally Identifiable Information remains under user control, for instance by being stored as a Verifiable Credential in a digital wallet (Sedlmeir et al., 2021). The "z-Commerce" architecture presented in this dissertation is a practical example of this approach (Ermolaev et al., 2023). The architecture is GDPR-compliant by design. By avoiding the central storage of Personally Identifiable Information (PII) and only making verifiable, data-minimizing proofs public, the system inherently satisfies both the right to erasure and the principle of data minimization.

1.3. From Traceability to De-anonymization

While public blockchains offer pseudonymity, their transparent nature enables a well-defined process of blockchain analysis that can lead to the erosion of privacy (Kappos et al., 2018; Meiklejohn et al., 2013; Werner et al., 2020). The pseudonymity offered by cryptocurrencies like Bitcoin and Ethereum is not true privacy, precisely because their public ledgers enable sophisticated analysis that can de-anonymize users (Bojja Venkatakrishnan et al., 2017; Liu et al., 2022). This process often begins with clustering, where analysts apply graph models and visualization techniques to group different pseudonymous addresses that are likely controlled by a single entity (Fischer et al., 2021; McGinn et al., 2016). The next step is identity linkage, where these clusters are connected to real-world identities, for instance by using statistical methods to associate wallet addresses with network-level identifiers like IP addresses (Juhász et al., 2018). These de-anonymization techniques can be exploited for various purposes, such as mass surveillance or behavioral profiling. The success of these techniques is powerful evidence that on a public blockchain, pseudonymity is not privacy – making the quest for robust privacy-enhancing technologies all the more critical (McGinn et al., 2018).

2 Challenges of Privacy by Design

The challenges of privacy start with its very definition. According to the Cambridge Dictionary, privacy is "someone's right to keep their personal matters and relationships secret" (Cambridge University Press, 2025a). The complexity of the definition deepens when moving from semantics to law, where privacy has been described as "the right to be let alone – the most comprehensive of rights and the right most valued by civilized men" (Solove, 2008). In the digital era, information privacy refers to a user's right to control the access to, use, and dissemination of their personal data (Quach et al., 2022). Therefore, in the context of this dissertation, privacy additionally implies the minimization of a user's information disclosure, particularly through mechanisms that allow for verification without revealing the underlying data.

This is further complicated by the "privacy paradox"; while privacy is a primary concern for citizens (Kokolakis, 2017), there is evidence that individuals are willing to trade their personal data for minimal compensation (Carrascal et al., 2013).

Achieving meaningful privacy is, therefore, a complex endeavor. The next section explores the concept of Privacy by Design and the relevant Privacy-Enhancing Technologies (PETs), including Zero-Knowledge Proofs (ZKPs), which can be used to facilitate Privacy by Design in transparent infrastructures such as public blockchains.

2.1. Privacy by Design

The concept of Privacy by Design (PbD) has been introduced by regulators as a primary solution to information privacy problems, defined as an engineering and management approach that embeds privacy principles directly into a system's architecture from its inception rather than as an afterthought (Cavoukian et al., 2009). However, its practical implementation faced immense challenges. Initially, these included conflicts with data-driven business models, the absence of clear engineering methodologies, and a persistent conceptual confusion between the broader goals of privacy and the narrower scope of security (Spiekermann, 2012). Moreover, the principles of Privacy by Design themselves had drawn criticism for leaving "many open questions about their application when engineering systems" (Gürses et al., 2011). Although nowadays PbD guidelines for developers provide clear and unambiguous instructions with regards to how personal data should be handled,

there is currently no generally accepted PbD standard or best practice (Barth et al., 2022). To overcome these barriers, the suggested path forward is to shift focus from high-level policies to concrete architectural practices, thus requiring developers to embed privacy into systems from the bottom up (Spiekermann, 2012).

2.2. Privacy-Enhancing Technology

In response to these challenges, a diverse array of Privacy-Enhancing Technologies (PETs) has been proposed, operating at different layers of the technology stack. And these PETs can be applied to the landscape of transparency akin to blockchains. The following analysis overviews the relevant PETs that can bring privacy in this context – Fully Homomorphic Encryption (FHE), Trusted Execution Environments (TEEs), Secure Multi-Party Computation (SMPC), ZKPs – examining what each of them offers and where it falls short, having their own challenges.

2.2.1. Fully Homomorphic Encryption

Fully Homomorphic Encryption represents a powerful form of encryption that allows for arbitrary computation directly on encrypted data (Fan and Vercauteren, 2012); however, its practical application is limited by several intertwined challenges. The primary obstacle is its extreme computational cost, with operations being up to 30,000 times slower than on plain text (Sidorov et al., 2022). Most of the papers developed new schemes to enhance efficiency (Yousuf et al., 2021). Widespread adoption is also hindered by a lack of unified standards and the deep cryptographic expertise required for a secure implementation, creating a high barrier to entry for developers (Acar et al., 2017).

2.2.2. Secure Multi-Party Computation

Secure Multi-Party Computation (SMPC) is a generic cryptographic primitive that enables multiple parties to jointly compute a function over their private inputs without revealing those inputs to each other (Evans et al., 2018; Goldreich, 1998; Zhao et al., 2019). However, its practical application is constrained by several significant challenges. Unlike other PETs, the primary bottleneck for SMPC is typically not computation, but communication, especially in the cloud-assisted Secure Multi-Party Computations (Zhao et al., 2019) where lack of trust in the providers is still a major barrier to adoption (Geppert et al., 2022).

Protocols often require a substantial amount of data to be exchanged, and the number of interaction rounds can introduce significant latency, especially over wide-area networks. Furthermore, there is a fundamental trade-off between security and efficiency; protocols that provide security against malicious actors (who can arbitrarily deviate from the protocol) are significantly more expensive than those that only protect against semi-honest (or honest-but-curious) actors (Evans et al., 2018). While researchers are utilizing blockchain to address fairness and scalability in Secure Multi-Party Computation, this approach introduces challenges, such as balancing efficiency and privacy, and necessitates further research (Zhong et al., 2020). Finally, implementing SMPC is complex for developers, as it requires specialized skills to convert standard programs into the boolean or arithmetic circuit representations that the protocols operate on (Goldreich, 1998) which is also a challenge for other PETs.

2.2.3. Trusted Execution Environment

Trusted Execution Environments (TEEs), sometimes called "enclaves", offer a hardware-based approach to privacy by creating a secure, isolated area within a processor for confidential computation, including on mobile devices (Ekberg et al., 2013). However, their practical application is limited by significant challenges, the most critical being their centralized trust model, which relies on a single hardware manufacturer (e.g., ARM's TrustZone, Intel's SGX, etc.) (Sabt et al., 2015). This dependency creates a fundamental trade-off: users gain hardware-enforced security but must in turn trust the centralized vendor. This issue is compounded by the inherent technical tension between domain isolation and resource sharing within the secure environment (Jauernig et al., 2020).

2.2.4. Zero-Knowledge Proof

ZKPs provide a cryptographic protocol for "verifiable privacy" (Walfish and Blumberg, 2015). They allow a party to prove the truth of a statement (e.g., "I am over 18 years old" or "My funds are not from a sanctioned address") without revealing the underlying data supporting the statement. Their strength is providing mathematically verifiable privacy without relying on trusted third parties or hardware.

The foundational concept of Zero-Knowledge Proof originated as interactive protocols, requiring multiple rounds of communication between a prover, who seeks to prove a state-

ment, and a verifier, who verifies such proofs. The security and functionality of any ZKP system are defined by three fundamental properties: completeness, soundness, and zero-knowledge (Sedlmeir et al., 2023). Completeness ensures that an honest prover can always convince the verifier of a true statement. Conversely, soundness guarantees that a dishonest prover cannot deceive the verifier about a false statement, except with a negligible probability that decreases exponentially with each round of interaction. Finally, the zero-knowledge property ensures that the verifier learns nothing beyond the mere validity of the statement itself.

A pivotal evolution in the field was the development of non-interactive ZKPs, largely enabled by the Fiat-Shamir heuristic (Fiat and Shamir, 1987), which replaces the verifier's challenges with cryptographic hash functions. This innovation, which condenses the proof into a single message, was a crucial step that made ZKPs practical for environments with limited interactivity and high verification costs. Thereby, non-interactive Zero-Knowledge Proofs are particularly interesting in the context of blockchain (Sun et al., 2021a) where a variety of use cases emerged not only for privacy but also for scalability (Lavaur et al., 2022). The most notable non-interactive ZKPs are Zero-Knowledge Succinct Non-Interactive Argument of Knowledge (zk-SNARK), Zero-Knowledge Scalable Transparent Argument of Knowledge (zk-STARK) and Bulletproofs.

- zk-SNARK
  - Features: Smallest proof sizes and fastest verification time. This makes zk-SNARKs suitable for resource-constrained environments like blockchains (Barbereau et al., 2023b; Petkus, 2019; Wilcox-O'Hearn, 2018).
  - Main Challenge: The requirement of a "trusted setup." If the secret parameters from this "trusted setup" are compromised (private keys, or so-called "toxic waste"), the security of the entire system can be compromised (Chen et al., 2023).
- zk-STARK
  - Features: Offers zero-knowledge proofs that are both transparent and scalable, making them suitable for verifying computations over large datasets (Ben-Sasson et al., 2018; Oude Roelink et al., 2024a). Unlike zk-SNARK, zk-STARK does not

require a "trusted setup," enhancing their security properties – no "toxic waste" (Oude Roelink et al., 2024a). They are also considered quantum-resistant, making them a more forward-looking solution (Ben-Sasson et al., 2018). Proof time is slightly faster than zk-SNARK.

  - Main Challenges: Proof sizes are two order of magnitude bigger than zk-SNARK, which can be costly for on-chain verification on public blockchains. Verification time is slightly slower than zk-SNARK (Wilcox-O'Hearn, 2018).

- Bulletproofs

  - Features: Enables efficient range proofs, which verify that a number lies within a specific range without disclosing the number itself (Oude Roelink et al., 2024a). Bulletproofs do not require a "trusted setup" – there is no "toxic waste."
  - Main Challenges: Proof sizes and verification time are one order of magnitude larger than zk-SNARK. Verification time is two orders of magnitude slower than for both zk-SNARK and zk-STARK (Wilcox-O'Hearn, 2018).

As we see, aforementioned non-interactive ZKPs have their own challenges in performance and architecture (e.g. zk-SNARK requires a "trusted setup"). It's worth to mention that all of them are complex to implement, requiring specialized expertise to design secure and efficient "circuits" – programs which are written in domain-specific programming languages (Belles-Munoz et al., 2022). And generating proofs can be computationally expensive, demanding significant computational power and time, which can be a barrier for resource-constrained applications (Sun et al., 2021b). Nevertheless, non-interactive ZKPs are recognized PETs, and widely used on public blockchains (Barbereau et al., 2023b; Gross et al., 2022; Lavaur et al., 2022) and beyond (Berentsen et al., 2023; Ermolaev et al., 2023; Lavin et al., 2024; Sedlmeir et al., 2022b), thereby contributing to architectural decision that facilitate practical implementation of Privacy by Design's principles. It is worth mentioning that when compliance is not considered in the initial design of such architectures, the resulting applications can reveal a disruptive potential, posing systemic challenges to the entire blockchain infrastructure (Barbereau et al., 2023b).

3 The Synthesis

The analysis in this chapter highlights a fundamental tension: the radical transparency required for blockchain's trust model often conflicts with the essential need for user privacy. On one hand, excessive transparency can function as a form of surveillance (Bernstein, 2012; Nguyen, 2021) and may conflict with data protection regulations like the GDPR (Sedlmeir et al., 2022a). On the other hand, achieving robust privacy is equally challenging; the principles of Privacy by Design are difficult to implement, and the relevant PETs are themselves complex and present their own trade-offs. Furthermore, systems that achieve near-absolute privacy can introduce their own significant compliance challenges (Barbereau et al., 2023b).

From this landscape, Zero-Knowledge Proofs emerge as a suited solution to balance the requirements of privacy and transparency. Non-interactive ZKPs allow for on-chain verification of private data without its disclosure. This powerful combination enables the creation of systems that can simultaneously leverage the blockchain's strengths – such as public auditability and tamper-resistance – while adhering to privacy principles like data minimization and enabling selective disclosure. It is this unique capability to reconcile transparency with privacy that makes the strategic integration of blockchain and ZKPs a powerful toolkit for building information systems. The following chapter will demonstrate how this integration is instantiated in concrete architectural designs.

III | Designing Solutions to Balance Privacy and Transparency

This chapter presents distinct system architectures from the original research articles, included in this cumulative dissertation (Appendix VI). Each of them occupies a different quadrant of the privacy-transparency solution space, as illustrated on Figure III.1, and demonstrates how the strategic integration of blockchain and ZKPs can create systems tailored to specific requirements, followed by a concluding discussion.

1 An Architecture Prioritizing Privacy

A primary challenge in the digital economy is the excessive collection and centralization of user data. In domains like e-commerce, platforms often gather vast amounts of personal information, creating significant privacy risks and single points of failure. Furthermore, the public nature of many blockchains can lead to the de-anonymization of users through transaction analysis. The research in RA1 (Ermolaev et al., 2023) addresses this challenge by designing a system that puts user privacy and data minimization first.

Grounded in the principles of Privacy by Design and facilitated by Zero-Knowledge Proof, the architecture's primary goals were threefold (1) data minimization at every step of a transaction, (2) sovereign user control over personal information via Digital Identity Wallets, and (3) a seamless "one-click" user experience comparable to incumbent platforms. Simultaneously, the architecture had to satisfy external constraints, such as enabling verifiable regulatory compliance (e.g., for Know Your Customer or age verification) and supporting essential business processes like refund flows. The resulting z-Commerce architecture

achieves this required balance by leveraging off-chain Zero-Knowledge Proofs in combination with Digital Identity Wallets.

In this architecture, users maintain sovereign control over their private data within their personal Digital Identity Wallets, where identity cards and other digital documents are stored as Verifiable Credentials. The architecture, detailed in the original paper, consists of four main components including the merchant's front-end and the user's Digital Identity Wallet, which holds their Verifiable Credentials. When a transaction occurs, the Digital Identity Wallet performs computation to create a Zero-Knowledge Proof. This proof – which confirms conditions like being over 18 without revealing a birthday date – is sent to the checkout page, which verifies the proof. The proof is being generated entirely within the user's Digital Identity Wallet, and only the proof is going to be delivered to the merchant for verification, thus fundamentally shifting the data ownership model.

This architectural approach directly addresses key principles of the GDPR. By design, the system adheres to data minimization, as merchants only receive a cryptographic proof rather than the raw personal data of customers. It also embodies the principles of user consent and control, as the user must actively authorize the generation and sharing of each proof from their personal digital identity wallet. Because Personally Identifiable Information is not stored centrally by the e-commerce platform, the risks associated with data breaches are significantly reduced, and conflicts with the GDPR's "right to be forgotten" are inherently avoided.

2 An Architecture Prioritizing Transparency

While some applications require maximum privacy, others demand high levels of transparency to build trust and ensure accountability. A key example of such a domain is Environmental, Social, and Governance reporting, where stakeholders, including investors and consumers, are increasingly concerned about corporate "greenwashing" – unsubstantiated claims about ESG performance. The core problem is a lack of trustworthy, verifiable data, making it difficult for third parties to audit corporate claims. The research in RA3 (Ermolaev et al., 2025) and RA5 (Ermolaev et al., 2026) address this problem by designing a system that uses blockchain to enhance transparency and trust.

The solution is an ESG KPI reporting system built on a blockchain – an architecture that, in contrast to the privacy-first approach of z-Commerce, prioritizes public verifiability. The process begins with a company selecting a set of industry-specific KPIs relevant to its operations. When ready to report, the company submits the KPI data to the system, along with a metadata identifier corresponding to its methodology document (stored off-chain on InterPlanetary File System). This information is then processed by a smart contract and recorded immutably – creating a permanent, timestamped, and auditable trail on the blockchain.

To further enhance transparency and prevent data ambiguity, the ESG KPI reporting system's design directly addresses the challenge of mutable state within smart contracts. While a blockchain's transaction history is immutable, the state of a smart contract variable can be overwritten, forcing reliance on off-chain services to reconstruct historical data from past events. Therefore, the architecture intentionally implements an explicitly versioned, append-only storage structure directly within the smart contract. Each submitted KPI is stored as a new, timestamped entry, making it impossible to overwrite previous records. This approach transparently exposes the entire history of a KPI, including entries that may have been submitted in error. However, rather than penalizing mistakes, this full auditability serves to prevent a chilling effect – it grants companies the "right to be wrong" by enabling them to correct erroneous data by simply submitting a new, updated version, thus demonstrating accountability and a commitment to accuracy.

This architectural choice makes the data tamper-resistant and transparent. Any stakeholder with access can independently verify the reported information against the permanent record on the blockchain. This design directly counters the problem of data manipulation and lack of trust in corporate reporting. By leveraging the inherent transparency of blockchain, the Environmental, Social, and Governance (ESG) reporting system provides a mechanism for holding organizations accountable for their claims, making it a prime example of a design where public auditability is the primary objective.

3 An Architecture Balancing Privacy and Transparency

The previous sections of this chapter represent architectures designed for the diagonally opposite quadrants of the privacy-transparency solution space. However, many applica-

tions require a complex, multi-layered blend of both privacy and transparency. In these hybrid systems, transparency by default is provided by the public nature of the ledger itself. Simultaneously, Privacy by Design is enforced at the application layer through smart contracts (Buterin, 2014) that use Zero-Knowledge Proofs to break the link between deposits and withdrawals. Crypto mixers are a prime example of such a hybrid architecture.

Our work in RA2 (Barbereau et al., 2023b) provides a taxonomy of such applications. The core mechanism of a ZKP-based mixer relies on a smart contract that functions as a pool of funds. A user initiates the process by depositing a fixed amount of cryptocurrency into the contract, along with a cryptographic hash of a secret known only to them. This hash, called a commitment, is added to a list of deposits stored by the contract. When the user wishes to withdraw their funds, they use a fresh, empty wallet address and submit a ZKP. This proof cryptographically demonstrates that they know a secret corresponding to one of the commitments in the contract's list – without revealing which specific commitment it is. After verifying the proof, the smart contract sends an equivalent amount of cryptocurrency to the new address, effectively breaking the on-chain link between the depositor and the recipient. Crypto mixers introduce a solution for transactional privacy onto public, permissionless blockchains. However, this potent privacy solution proved to be powerfully disruptive, creating immense regulatory and legal challenges, with consequences felt most acutely by its own developers.

To address these regulatory challenges, various technical and policy levers can be integrated into mixer designs, though each comes with a direct trade-off for the system's anonymity set. These levers often create a layer of transparency on demand, where users can provide proofs for compliance without compromising their privacy. For instance, a user could generate a ZKP to demonstrate that their funds do not originate from a sanctioned address, without revealing the address itself. Another example involves modifying the design of so-called "relayers" – services that solve the operational challenge of paying withdrawal fees from a new wallet. By implementing sanctioned-address filters, these relayers can serve as a point for compliance, preventing deposits from known illicit sources (Barbereau et al., 2023b; Burleson et al., 2022).

While the aforementioned approach enhances compliance, it can fragment the anonymity set if different relayers use different blocklists. Furthermore, rate limits can be imposed at the smart contract level to slow down large-scale laundering, but this may also incon-

venience legitimate users and reduce the overall liquidity that forms the anonymity pool. More advanced mechanisms include implementing on-chain deposit and withdrawal censorship through inclusion or exclusion lists (Barbereau et al., 2023b). Each of these controls, while moving the architecture towards regulatory acceptability, systematically reduces the size or uniformity of the anonymity set, highlighting the delicate balance that compliant privacy-preserving systems must strike.

The case of the Ethereum-based crypto mixer Tornado Cash is illustrative. The platform came under intense scrutiny after it was identified as a key tool for laundering funds by North Korea's Lazarus Group (TRM Labs, 2022). This illicit usage prompted the U.S. Department of the Treasury's Office of Foreign Assets Control (OFAC) to formally sanction the service (U.S. Department of the Treasury, 2022), leading to the arrest and sentencing of one of its developers (Nelson, 2024; Schickler, 2025). The original Tornado Cash protocol was not designed for regulatory compliance, demonstrating the severe friction that arises when powerful privacy tools meet real-world legal frameworks.

4 Discussion

The architectural solutions presented in this chapter – the privacy-first z-Commerce architecture, the transparency-first ESG reporting system, and the hybrid crypto mixers – serve to illustrate the diverse solution space for balancing privacy and transparency. These solutions demonstrate how different architectural patterns can be employed to meet specific requirements of privacy and transparency and be mapped onto the four quadrants of the privacy-transparency solution space, as illustrated in Figure III.1. Three of these quadrants are populated by the presented solutions:

- Low Transparency & High Privacy: This quadrant is for applications prioritizing confidentiality, exemplified by the z-Commerce architecture, detailed in RA1 (Ermolaev et al., 2023). This model minimizes on-chain data using off-chain Zero-Knowledge Proofs, but its trustworthiness can be enhanced with a transparent, on-chain registry for revoked Verifiable Credentials.
- High Transparency & High Privacy: This quadrant represents hybrid applications requiring a balance of both properties. The concept of a "compliant crypto mixer"

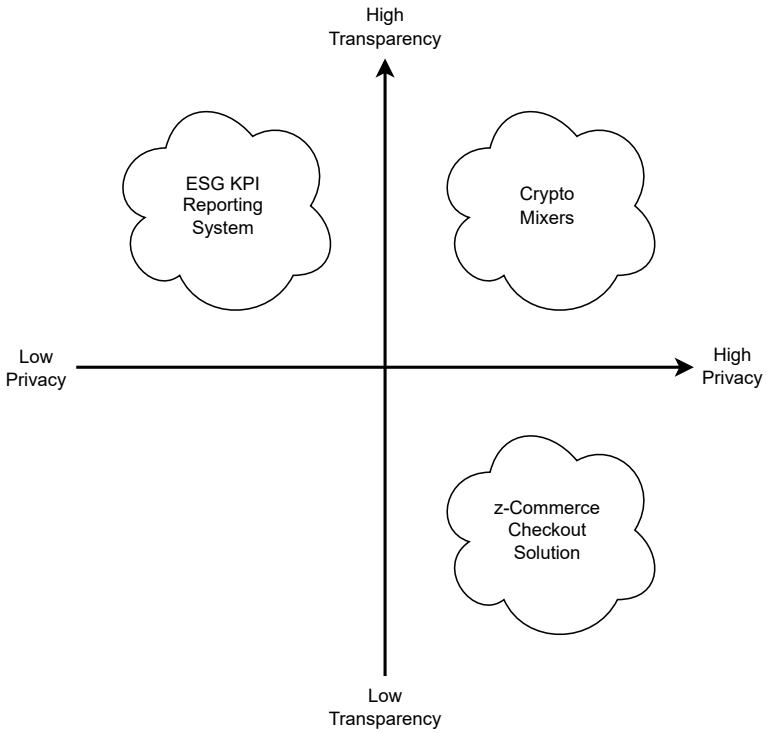

Figure III.1: The privacy-transparency solution space with mapped research articles.

serves as a prime example here, and the taxonomy developed in RA2 (Barbereau et al., 2023b) analyzes the architectural features – such as on-chain ZKP verification and inclusion lists – that enable such a hybrid model.

- High Transparency & Low Privacy: This quadrant is for applications prioritizing auditability, exemplified by the ESG reporting system from RA3, RA5 (Ermolaev et al., 2025; Ermolaev et al., 2026). This architecture leverages on-chain data for an immutable record, yet the design evaluation revealed a demand for layering privacy-preserving techniques (like Zero-Knowledge Proofs) to handle sensitive Key Performance Indicators (KPIs).

Beyond architectural design, the practical viability and performance of these architectures are critical, as blockchain-based systems often face performance challenges such as low throughput and high latency. The research in RA4 (Papageorgiou et al., 2025) addresses this challenge by introducing a generalizable method for benchmarking the performance

of permissioned blockchains, including those that are Ethereum Virtual Machine (EVM)-compatible. Such permissioned systems may also implement ZKPs to enable privacy features. While the fundamental goal of using ZKPs is similar to that in public blockchains – to introduce selective privacy – the context differs. This differs from the fully public nature of permissionless systems; in the permissioned context, ZKPs are often used to manage privacy between known participants rather than providing anonymity to an open network (Politou et al., 2021).

IV | Conclusion

Having demonstrated concrete architectural solutions in the preceding chapter, this final chapter of this dissertation (1) provides a Summary of the overall contributions of this dissertation in the respective section, (2) discusses the Limitations and (3) Outlook in the next sections.

1 Summary

This dissertation is grounded in four years of research that resulted in five original research articles. The work is structured around three core contributions. First, this dissertation analyzes the fundamental challenges of achieving Privacy by Design in the face of the radical transparency inherent to public blockchains. Second, it demonstrates that blockchain and ZKPs are complementary technologies that can be combined in numerous ways to satisfy specific system requirements. Third, it provides an overview of multiple architectural solutions across the privacy-transparency solution space, with each solution embracing a unique set of features to strike a particular balance.

The exploration of these architectural archetypes yielded several key findings. The privacy-first "z-Commerce" model demonstrates how off-chain ZKPs can be used to achieve data minimization and regulatory compliance in user-centric applications. Conversely, the transparency-first ESG reporting system highlights the power of the blockchain as a public ledger for auditability, while also revealing the practical need to layer privacy mechanisms onto transparent systems to protect sensitive data. Finally, the analysis of hybrid systems like crypto mixers underscores the power of on-chain ZKP verification but, more critically, it reveals an inescapable trade-off between the strength of anonymity and the requirements

of regulatory compliance, arguing for a "Compliance by Design" approach alongside the "Privacy by Design" approach.

2 Limitations of the Current Ecosystem

Despite rapid progress, the at-scale deployment of systems combining blockchain and ZKPs still faces interwoven technical, economic, and regulatory challenges that impact adoption. This section outlines these key limitations before discussing the favorable trajectory enabled by maturing standards and technological advancements.

2.1. Adoption and Usability Challenges

Despite growing market valuations, blockchain technology remains in its early stages, with a relatively small user base; for instance, there were just over 600 million identity-verified crypto-asset users globally as of November 2024 (Statista Research Department, 2024). A primary barrier to mainstream adoption is the poor usability of self-custody wallets, which involves complexities with private key management, transaction fees, and limited recourse for errors. Furthermore, public blockchains that use Proof-of-Work face significant environmental concerns regarding their energy consumption (Sedlmeir et al., 2020). While emerging design patterns like Account Abstraction (Buterin et al., 2021) promise safer and more familiar user flows, the limited everyday usability of current systems remains a key hurdle.

2.2. Regulatory and Security Risks

The regulatory landscape for Privacy-Enhancing Technologies (PETs) is uncertain and often adversarial, as exemplified by the sanctions against Tornado Cash. Beyond direct legal action, these systems face significant security risks. Like any cryptographic protocol, Zero-Knowledge Proofs may have yet-undiscovered vulnerabilities (Tang et al., 2024). More concretely, specific implementations can contain flaws or even malicious backdoors, as was reportedly found in an early version of the Tornado Cash protocol (Cimpanu, 2022). Furthermore, the governance mechanisms of these decentralized systems can be manipulated; Tornado Cash, for instance, suffered a governance attack where a malicious actor gained control of the protocol through fraudulent proposals (Halborn, 2023).

2.3. Technical Hurdles for ZKPs

While powerful, ZKPs themselves have practical limitations. Zero-Knowledge Succinct Non-Interactive Argument of Knowledge systems require a trusted setup, a complex parameter-generation ceremony where failure to securely destroy the setup's secret could compromise the system's soundness. Furthermore, proof generation remains substantially more computationally expensive than ordinary computation, though the gap is narrowing thanks to the development of specialized provers. Although, Zero-Knowledge Scalable Transparent Argument of Knowledge systems do not require trusted setups and is generally considered post-quantum-resistant (Oude Roelink et al., 2024b), its proof size is large, making zk-STARKs less applicable in the scope of blockchain.

3 Outlook

3.1. Maturing Standards and Technology

Standardization is a key enabler for adoption. The World Wide Web Consortium (W3C) Verifiable Credential (VC) Data Model 2.0 now officially supports ZK assertions, and the 2024 revision of the electronic Identification, Authentication and Trust Services (eIDAS) regulation that targets selective disclosure in European Digital Identity Wallets. At the protocol layer, the wider availability of cryptographic pre-compiled circuits that accelerate proving by more than 95% (Wang and Gao, 2025; Zhang et al., 2025). This progress in software and standards is complemented by rapid advancements in hardware, with specialized provers making proof generation faster and more energy-efficient.

3.2. Future Directions: Verifiable CPU Execution

Looking further ahead, as computational power and cryptographic techniques continue to advance, it may become feasible to apply ZKPs at an even more fundamental level. Emerging research on "zkCPU" explores the possibility of generating proofs for the correct execution of entire CPU instruction sets (Heath et al., 2021; Yang et al., 2024). This could enable verifiable computation for general-purpose programs, opening up a new frontier for trust-minimized and privacy-preserving systems.

In sum, while challenges persist, the combination of accelerated ZK proving, maturing standards, and groundbreaking new research indicates that a workable balance between privacy and transparency is increasingly attainable.

V | References

Acar, A., H. Aksu, A. S. Uluagac, and M. Conti (2017). A Survey on Homomorphic Encryption Schemes: Theory and Implementation. arXiv: 1704.03578 [cs.CR]. URL: https://arxiv.org/abs/1704.03578.

Amend, J., L. Arnold, L. Fabri, S. Feulner, G. Fridgen, L. Harzer, P. Karnebogen, F. Köhler, P. Ollig, A. Rieger, B. Schellinger, and G. M. Schmidbauer-Wolf (2022). Föderale Blockchain Infrastruktur Asyl (FLORA): Pilotierung und Evaluation des FLORA-Assistenzsystems im Kontext der AnkER-Einrichtung Dresden. de. DOI: 10.24406/publica-941. URL: https://publica.fraunhofer.de/handle/publica/436999.

Amend, J., C. van Dun, G. Fridgen, F. Köhler, A. Rieger, A. Stohr, and A. Wenninger (2021a). "Using Blockchain to Coordinate Federal Processes: The Case of Germany's Federal Office for Migration and Refugees". In: Digitalization Cases Vol. 2: Mastering Digital Transformation for Global Business. Ed. by N. Urbach, M. Röglinger, K. Kautz, R. A. Alias, C. Saunders, and M. Wiener. Cham: Springer International Publishing, pp. 85–100. ISBN: 978-3-030-80003-1. DOI: 10.1007/978-3-030-80003-1_5. URL: https://doi.org/10.1007/978-3-030-80003-1_5.

Amend, J., M. Federbusch, G. Fridgen, F. Köhler, A. Rieger, V. Schlatt, J. Sedlmeir, A. Stohr, and C. van Dun (2021b). Digitization of certification processes in the asylum procedure by means of digital identities: A feasibility study by Germany's Federal Office for Migration and Refugees. English. Tech. rep. URL: https://www.bamf.de/SharedDocs/Anlagen/EN/Digitalisierung/blockchain-whitepaper-2021.pdf?__blob=publicationFile&v=2.

Arnold, L., M. Brennecke, P. Camus, G. Fridgen, T. Guggenberger, S. Radszuwill, A. Rieger, A. Schweizer, and N. Urbach (2019). "Blockchain and Initial Coin Offerings: Blockchain's Implications for Crowdfunding". In: Business Transformation through

Blockchain: Volume I. Ed. by H. Treiblmaier and R. Beck. Cham: Springer International Publishing, pp. 233–272. ISBN: 978-3-319-98911-2. DOI: 10.1007/978-3-319-98911-2_8. URL: https://doi.org/10.1007/978-3-319-98911-2_8.

Bachmann, N. M., B. Drasch, G. Fridgen, M. Miksch, F. Regner, A. Schweizer, and N. Urbach (2022). "Tarzan and chain: exploring the ICO jungle and evaluating design archetypes". In: Electronic Markets 32.3, pp. 1725–1748. ISSN: 1422-8890. DOI: 10.1007/s12525-021-00463-6. URL: https://doi.org/10.1007/s12525-021-00463-6.

Ball, C. (2009). "What Is Transparency?" In: Public Integrity 11.4, pp. 293–308. DOI: 10.2753/PIN1099-9922110400. eprint: https://www.tandfonline.com/doi/pdf/10.2753/PIN1099-9922110400. URL: https://www.tandfonline.com/doi/abs/10.2753/PIN1099-9922110400.

Barbereau, T., J. Sedlmeir, R. Smethurst, G. Fridgen, and A. Rieger (2022a). "Tokenization and Regulatory Compliance for Art and Collectibles Markets: From Regulators' Demands for Transparency to Investors' Demands for Privacy". In: Blockchains and the Token Economy: Theory and Practice. Ed. by M. C. Lacity and H. Treiblmaier. Cham: Springer International Publishing, pp. 213–236. ISBN: 978-3-030-95108-5. DOI: 10.1007/978-3-030-95108-5_8. URL: https://doi.org/10.1007/978-3-030-95108-5_8.

Barbereau, T., R. Smethurst, O. Papageorgiou, A. Rieger, and G. Fridgen (2022b). "DeFi, Not So Decentralized: The Measured Distribution of Voting Rights". In: DOI: 10.24251/HICSS.2022.734.

Barbereau, T., R. Smethurst, O. Papageorgiou, J. Sedlmeir, and G. Fridgen (2023a). "Decentralised Finance's timocratic governance: The distribution and exercise of tokenised voting rights". In: Technology in Society 73, p. 102251. ISSN: 0160-791X. DOI: https://doi.org/10.1016/j.techsoc.2023.102251. URL: https://www.sciencedirect.com/science/article/pii/S0160791X23000568.

Barbereau, T. J., E. Ermolaev, M. Brennecke, E. Hartwich, and J. Sedlmeir (2023b). "Beyond a Fistful of Tumblers: Toward a Taxonomy of Ethereum-based Mixers". English. In: Cybersecurity and Privacy. FNR - Luxembourg National Research Fund. URL: https://hdl.handle.net/10993/57137.

Barth, S., D. Ionita, and P. Hartel (2022). "Understanding Online Privacy—A Systematic Review of Privacy Visualizations and Privacy by Design Guidelines". In: ACM Comput. Surv. 55.3. ISSN: 0360-0300. DOI: 10.1145/3502288. URL: https://doi.org/10.1145/3502288.

Belles-Munoz, M., M. Isabel, J. Munoz-Tapia, A. Rubio, and J. Baylina (2022). "Circom: A Circuit Description Language for Building Zero-Knowledge Applications". In: IEEE Transactions on Dependable and Secure Computing PP, pp. 1–18. DOI: 10.1109/TDSC.2022.3232813.

Ben-Sasson, E., I. Bentov, Y. Horesh, and M. Riabzev (2018). "Scalable, transparent, and post-quantum secure computational integrity". In: Cryptology ePrint Archive.

Berentsen, A., J. Lenzi, and R. Nyffenegger (2023). "An Introduction to Zero-Knowledge Proofs in Blockchains and Economics." In: Review (00149187) 105.4.

Bernstein, E. S. (2012). "The Transparency Paradox". In: Administrative Science Quarterly 57.2, pp. 181–216. ISSN: 0001-8392. DOI: 10.1177/0001839212453028. URL: http://dx.doi.org/10.1177/0001839212453028.

Bojja Venkatakrishnan, S., G. Fanti, and P. Viswanath (2017). "Dandelion". In: Proceedings of the ACM on Measurement and Analysis of Computing Systems 1.1, pp. 1–34. ISSN: 2476-1249. DOI: 10.1145/3084459. URL: http://dx.doi.org/10.1145/3084459.

Burleson, J., M. Korver, and D. Boneh (2022). Privacy-protecting regulatory solutions using zero-knowledge proofs.

Buterin, V. (2014). A Next-Generation Smart Contract and Decentralized Application Platform. Ethereum White Paper. URL: https://ethereum.org/en/whitepaper/.

Buterin, V., Y. Weiss, D. Tirosh, S. Nacson, A. Forshtat, K. Gazso, and T. Hess (2021). ERC-4337: Account abstraction using alt mempool. Tech. rep. Ethereum Improvement Proposals.

Cambridge University Press (2025a). privacy. Cambridge Dictionary. URL: https://dictionary.cambridge.org/dictionary/english/privacy (visited on 08/12/2025).

Cambridge University Press (2025b). transparency. Cambridge Dictionary. URL: https://dictionary.cambridge.org/dictionary/english/transparency (visited on 08/12/2025).

Carrascal, J. P., C. Riederer, V. Erramilli, M. Cherubini, and R. De Oliveira (2013). "Your browsing behavior for a big mac: Economics of personal information online". In: Proceedings of the 22nd international conference on World Wide Web, pp. 189–200.

Cavoukian, A. et al. (2009). "Privacy by design: The 7 foundational principles". In: Information and privacy commissioner of Ontario, Canada 5.2009, p. 12.

Chen, T., H. Lu, T. Kunpittaya, and A. Luo (2023). A Review of zk-SNARKs. arXiv: 2202.06877 [cs.CR]. URL: https://arxiv.org/abs/2202.06877.

Cimpanu, C. (2022). Risky Biz News: Backdoor code found in Tornado Cash. https://news.risky.biz/risky-biz-news-backdoor-code-found-in-tornado-cash/. Accessed: 2025-08-03.

Clarke, R. (2019). "Risks inherent in the digital surveillance economy: A research agenda". In: Journal of Information Technology 34.1, pp. 59–80. ISSN: 0268-3962. DOI: 10.1177/0268396218815559. URL: http://dx.doi.org/10.1177/0268396218815559.

Cruz, R. d. l. (2019). "Privacy Laws in the Blockchain Environment". In: Annals of Emerging Technologies in Computing 3.5, pp. 34–44. ISSN: 2516-029X. DOI: 10.33166/aetic.2019.05.005. URL: http://dx.doi.org/10.33166/aetic.2019.05.005.

Ekberg, J.-E., K. Kostiainen, and N. Asokan (2013). "Trusted execution environments on mobile devices". In: Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security. CCS '13. New York, NY, USA: Association for Computing Machinery, 1497–1498. ISBN: 9781450324779. DOI: 10.1145/2508859.2516758. URL: https://doi.org/10.1145/2508859.2516758.

Ermolaev, E., I. Abellán Álvarez, J. Sedlmeir, and G. Fridgen (2023). "z-Commerce: Designing a Data-Minimizing One-Click Checkout Solution". English. In: Design Science Research for a New Society: Society 5.0. FNR - Fonds National de la Recherche. Springer Nature. DOI: 10.1007/978-3-031-32808-4_1. URL: https://doi.org/10.1007/978-3-031-32808-4_1.

Ermolaev, E., E. Y. Tseloni, T. Roth, and G. Fridgen (2025). "Designing a Reporting System for Trust in Environmental Social Governance". English. In: Proceedings of the Annual Hawaii International Conference on System Sciences. ISSN: 1530-1605. URL: https://scholarspace.manoa.hawaii.edu/items/27db5f81-08f7-4138-a19d-cb371618808e.

Ermolaev, E., E. Y. Tseloni, T. Roth, and G. Fridgen (2026). "Blockchain-based Reporting System for Trust in Environmental Social Governance". English. HDL: https://hdl.handle.net/10993/66729.

Evans, D., V. Kolesnikov, and M. Rosulek (2018). A Pragmatic Introduction to Secure Multi-Party Computation. now publishers Inc. DOI: 10.1561/3300000019. URL: https://securecomputation.org.

Fan, J. and F. Vercauteren (2012). Somewhat Practical Fully Homomorphic Encryption. Cryptology ePrint Archive, Paper 2012/144. URL: https://eprint.iacr.org/2012/144.

Fiat, A. and A. Shamir (1987). "How To Prove Yourself: Practical Solutions to Identification and Signature Problems". In: Advances in Cryptology — CRYPTO' 86. Ed. by A. M.

Odlyzko. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 186–194. ISBN: 978-3-540-47721-1.

Fiege, U., A. Fiat, and A. Shamir (1987). "Zero knowledge proofs of identity". In: Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing. STOC '87. New York, NY, USA: Association for Computing Machinery, 210–217. ISBN: 0897912217. DOI: 10.1145/28395.28419. URL: https://doi.org/10.1145/28395.28419.

Fischer, J. A., A. Palechor, D. Dell'Aglio, A. Bernstein, and C. J. Tessone (2021). "The Complex Community Structure of the Bitcoin Address Correspondence Network". In: Frontiers in Physics 9. ISSN: 2296-424X. DOI: 10.3389/fphy.2021.681798. URL: http://dx.doi.org/10.3389/fphy.2021.681798.

Fridgen, G. and T. Z. Garizy (2015). "Supply Chain Network Risk Analysis : A Privacy Preserving Approach". English. In: Proceedings of the 23rd European Conference on Information Systems. Spring. URL: https://aisel.aisnet.org/ecis2015_cr/49/?utm_source=aisel.aisnet.org%2Fecis2015_cr%2F49&utm_medium=PDF&utm_campaign=PDFCoverPages.

Fridgen, G., T. Guggenberger, J. Sedlmeir, and N. Urbach, eds. (2024a). Decentralization Technologies. Financial Sector in Change. 1st ed. Financial Innovation and Technology. 14 b/w illustrations, 2 illustrations in colour. Published as eBook on 10 December 2024; Hardcover on 11 December 2024; Softcover due 25 December 2025. Springer Cham, p. 263. ISBN: 978-3-031-66047-4. DOI: 10.1007/978-3-031-66047-4. URL: https://doi.org/10.1007/978-3-031-66047-4.

Fridgen, G., T. Guggenberger, J. Sedlmeir, and N. Urbach (2024b). "Introduction: Decentralization Technologies in Finance". In: Decentralization Technologies: Financial Sector in Change. Ed. by G. Fridgen, T. Guggenberger, J. Sedlmeir, and N. Urbach. Cham: Springer Nature Switzerland, pp. 3–17. ISBN: 978-3-031-66047-4. DOI: 10.1007/978-3-031-66047-4_1. URL: https://doi.org/10.1007/978-3-031-66047-4_1.

Fridgen, G., F. Guggenmoos, J. Lockl, A. Rieger, and A. Schweizer (2018a). "Developing an Evaluation Framework for Blockchain in the Public Sector: The Example of the German Asylum Process". In: Proceedings of 1st ERCIM Blockchain Workshop 2018. European Society for Socially Embedded Technologies (EUSSET). DOI: 10.18420/blockchain2018_10.

Fridgen, G., F. Guggenmos, J. Lockl, A. Rieger, and N. Urbach (2019). Supporting communication and cooperation in the asylum procedure with Blockchain technology : A

proof of concept by the Federal Office for Migration and Refugees. English. Tech. rep. URL: https://eref.uni-bayreuth.de/48037/.

Fridgen, G., S. Radszuwill, N. Urbach, and L. Utz (2018b). "Cross-Organizational Workflow Management Using Blockchain Technology : Towards Applicability, Auditability, and Automation". English. In: 51st Annual Hawaii International Conference on System Sciences (HICSS). URL: https://eref.uni-bayreuth.de/39830/.

Fridgen, G., F. Regner, A. Schweizer, and N. Urbach (2018c). "Don't Slip on the Initial Coin Offering (ICO) : A Taxonomy for a Blockchain-enabled Form of Crowdfunding". English. In: 26th European Conference on Information Systems (ECIS). URL: https://eref.uni-bayreuth.de/44524/.

Geppert, T., S. Deml, D. Sturzenegger, and N. Ebert (2022). "Trusted Execution Environments: Applications and Organizational Challenges". In: Frontiers in Computer Science Volume 4 - 2022. ISSN: 2624-9898. DOI: 10.3389/fcomp.2022.930741. URL: https://www.frontiersin.org/journals/computer-science/articles/10.3389/fcomp.2022.930741.

Glöckler, J., J. Sedlmeir, M. Frank, and G. Fridgen (2024). "A Systematic Review of Identity and Access Management Requirements in Enterprises and Potential Contributions of Self-Sovereign Identity". In: Business & Information Systems Engineering 66.4, pp. 421–440. ISSN: 1867-0202. DOI: 10.1007/s12599-023-00830-x. URL: https://doi.org/10.1007/s12599-023-00830-x.

Goldreich, O. (1998). "Secure multi-party computation". In: Manuscript. Preliminary version 78.110, pp. 1–108.

Goldwasser, S., S. Micali, and C. Rackoff (1989). "The Knowledge Complexity of Interactive Proof Systems". In: SIAM Journal on Computing 18.1, pp. 186–208. DOI: 10.1137/0218012. eprint: https://doi.org/10.1137/0218012. URL: https://doi.org/10.1137/0218012.

Gross, J., J. Sedlmeir, M. Babel, A. Bechtel, and B. Schellinger (2021). Designing a Central Bank Digital Currency with Support for Cash-Like Privacy. Tech. rep. SSRN 3891121. Social Science Research Network. DOI: 10.2139/ssrn.3891121. URL: https://ssrn.com/abstract=3891121.

Gross, J., J. Sedlmeir, and S. Seiter (2022). How to Design a Compliant, Privacy-Preserving Fiat Stablecoin Via Zero-Knowledge Proofs. Tech. rep. SSRN 4331465. Social Science Research Network. DOI: 10.2139/ssrn.4331465. URL: https://ssrn.com/abstract=4331465.

Guggenmos, F., J. Lockl, A. Rieger, and G. Fridgen (2019). "Blockchain in der öffentlichen Verwaltung". In: Informatik Spektrum 42.3, pp. 174–181. ISSN: 1432-122X. DOI: 10.1007/s00287-019-01177-y. URL: https://doi.org/10.1007/s00287-019-01177-y.

Guggenmos, F., J. Lockl, A. Rieger, and G. Fridgen (2020). "How to Develop a GDPR-Compliant Blockchain Solution for Cross-Organizational Workflow Management: Evidence from the German Asylum Procedure." English. In: Proceedings of the Hawaii International Conference on System Sciences 2020. FNR - Fonds National de la Recherche. DOI: 10.24251/HICSS.2020.492. URL: http://hdl.handle.net/10125/64234.

Gürses, S., C. Troncoso, and C. Diaz (2011). "Engineering privacy by design". In: Computers, Privacy & Data Protection 14.3, p. 25.

Halborn (2023). Explained: The Tornado Cash Hack (May 2023). https://www.halborn.com/blog/post/explained-the-tornado-cash-hack-may-2023. Accessed: 2025-08-03.

Hartwich, E., T. H. Roth, A. Rieger, L. Zavolokina, and G. Fridgen (2024). "Negotiation and Translation Between Discursive Fields: A Study on the Diffusion of Decentralized Finance". In: ECIS 2024 Proceedings. 4. Available at AIS Electronic Library (AISeL). URL: https://aisel.aisnet.org/ecis2024/track20_adoption/track20_adoption/4.

Heath, D., Y. Yang, D. Devecsery, and V. Kolesnikov (2021). "Zero Knowledge for Everything and Everyone: Fast ZK Processor with Cached ORAM for ANSI C Programs". In: 2021 IEEE Symposium on Security and Privacy (SP), pp. 1538–1556. DOI: 10.1109/SP40001.2021.00089.

Herian, R. (2020). "Blockchain, GDPR, and fantasies of data sovereignty". In: Law, Innovation and Technology 12.1, pp. 156–174. ISSN: 1757-9961. DOI: 10.1080/17579961.2020.1727094. URL: http://dx.doi.org/10.1080/17579961.2020.1727094.

Hevner, A. R., S. T. March, J. Park, and S. Ram (2004). "Design science in information systems research". In: MIS quarterly, pp. 75–105.

Hopwood, D., S. Bowe, T. Hornby, N. Wilcox, et al. (2016). "Zcash protocol specification". In: GitHub: San Francisco, CA, USA 4.220, p. 32.

Höß, A., D. H. Maclennan, A. Rieger, E. Ermolaev, G. Fridgen, and T. Roth (2022a). Issuing and verifying digital diplomas with the European Blockchain Services Infrastructure - Insights from the EBSILUX project. English. Tech. rep. Ministry for Digitalisation.

Höß, A., D. H. Maclennan, A. Rieger, E. Ermolaev, G. Fridgen, and T. Roth (2022b). Issuing and verifying digital diplomas with the European Blockchain Services Infrastructure - Insights from the EBSILUX project. English. Tech. rep. University of Luxembourg.

Höß, A., T. Roth, J. Sedlmeir, G. Fridgen, and A. Rieger (03 January 2022). "With or Without Blockchain? Towards a Decentralized, SSI-based eRoaming Architecture". English. In: Proceedings of the 55th Hawaii International Conference on System Sciences (HICSS). Nominated for the "Internet and the Digital Economy" best paper award. FNR - Fonds National de la Recherche. IEEE Computer Society. URL: https://hdl.handle.net/10125/79899.

Jauernig, P., A.-R. Sadeghi, and E. Stapf (2020). "Trusted Execution Environments: Properties, Applications, and Challenges". In: IEEE Security & Privacy 18.2, pp. 56–60. DOI: 10.1109/MSEC.2019.2947124.

Juhász, P. L., J. Stéger, D. Kondor, and G. Vattay (2018). "A Bayesian approach to identify Bitcoin users". In: PLOS ONE 13.12. Ed. by I. Olier, e0207000. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0207000. URL: http://dx.doi.org/10.1371/journal.pone.0207000.

Kappos, G., H. Yousaf, M. Maller, and S. Meiklejohn (2018). An Empirical Analysis of Anonymity in Zcash. arXiv: 1805.03180 [cs.CR]. URL: https://arxiv.org/abs/1805.03180.

Kokolakis, S. (2017). "Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon". In: Computers & Security 64, pp. 122–134. ISSN: 0167-4048. DOI: https://doi.org/10.1016/j.cose.2015.07.002. URL: https://www.sciencedirect.com/science/article/pii/S0167404815001017.

Lavaur, T., J. Lacan, and C. P. C. Chanel (2022). "Enabling Blockchain Services for IoE with Zk-Rollups". In: Sensors 22.17. ISSN: 1424-8220. DOI: 10.3390/s22176493. URL: https://www.mdpi.com/1424-8220/22/17/6493.

Lavin, R., X. Liu, H. Mohanty, L. Norman, G. Zaarour, and B. Krishnamachari (2024). "A survey on the applications of zero-knowledge proofs". In: arXiv preprint arXiv:2408.00243.

Liang, X., J. Zhao, Y. Chen, E. Bandara, and S. Shetty (2023). "Architectural Design of a Blockchain-Enabled, Federated Learning Platform for Algorithmic Fairness in Predictive Health Care: Design Science Study". In: Journal of Medical Internet Research 25, e46547. ISSN: 1438-8871. DOI: 10.2196/46547. URL: http://dx.doi.org/10.2196/46547.

Liu, X. F., C. G. Akcora, Z.-Y. Zhang, and J.-G. Liu (2022). "Editorial: Cryptocurrency Transaction Analysis From a Network Perspective". In: Frontiers in Physics 10. ISSN: 2296-424X. DOI: 10.3389/fphy.2022.876983. URL: http://dx.doi.org/10.3389/fphy.2022.876983.

McGinn, D., D. McIlwraith, and Y. Guo (2018). "Towards open data blockchain analytics: a Bitcoin perspective". In: Royal Society Open Science 5.8, p. 180298. ISSN: 2054-5703. DOI: 10.1098/rsos.180298. URL: http://dx.doi.org/10.1098/rsos.180298.

McGinn, D., D. Birch, D. Akroyd, M. Molina-Solana, Y. Guo, and W. J. Knottenbelt (2016). "Visualizing Dynamic Bitcoin Transaction Patterns". In: Big Data 4.2, pp. 109–119. ISSN: 2167-6461. DOI: 10.1089/big.2015.0056. URL: http://dx.doi.org/10.1089/big.2015.0056.

Meijer, A. (2009). "Understanding modern transparency". In: International Review of Administrative Sciences 75.2, pp. 255–269. DOI: 10.1177/0020852309104175. eprint: https://doi.org/10.1177/0020852309104175. URL: https://doi.org/10.1177/0020852309104175.

Meiklejohn, S., M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage (2013). "A fistful of bitcoins: characterizing payments among men with no names". In: Proceedings of the 2013 Conference on Internet Measurement Conference. IMC '13. New York, NY, USA: Association for Computing Machinery, 127–140. ISBN: 9781450319539. DOI: 10.1145/2504730.2504747. URL: https://doi.org/10.1145/2504730.2504747.

Murray, D., P. Fussey, K. Hove, W. Wakabi, P. Kimumwe, O. Saki, and A. Stevens (2023). "The Chilling Effects of Surveillance and Human Rights: Insights from Qualitative Research in Uganda and Zimbabwe". In: Journal of Human Rights Practice 16.1, pp. 397–412. ISSN: 1757-9619. DOI: 10.1093/jhuman/huad020. URL: http://dx.doi.org/10.1093/jhuman/huad020.

National Information Standards Organization (2022). ANSI/NISO Z39.104-2022, CRediT, Contributor Roles Taxonomy. https://www.niso.org/publications/z39104-2022-credit. Standard No. Z39.104-2022, Baltimore, MD. DOI: 10.3789/ansi.niso.z39.104-2022.

Nelson, D. (2024). Tornado Cash Dev Alexey Pertsev Sentenced to 64 Months in Prison. https://www.nasdaq.com/articles/tornado-cash-dev-alexey-pertsev-sentenced-to-64-months-in-prison. Accessed: 2025-08-01.

Nguyen, C. T. (2021). "Transparency is Surveillance". In: Philosophy and Phenomenological Research 105.2, pp. 331–361. ISSN: 0031-8205. DOI: 10.1111/phpr.12823. URL: http://dx.doi.org/10.1111/phpr.12823.

Oude Roelink, B., M. El-Hajj, and D. Sarmah (2024a). "Systematic review: Comparing zk-SNARK, zk-STARK, and bulletproof protocols for privacy-preserving authentication". In: Security and Privacy 7.5, e401.

Oude Roelink, B., M. El-Hajj, and D. Sarmah (2024b). "Systematic review: Comparing zk-SNARK, zk-STARK, and bulletproof protocols for privacy-preserving authentication". In: SECURITY AND PRIVACY 7.5, e401. DOI: https://doi.org/10.1002/spy2.401. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/spy2.401. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/spy2.401.

Papageorgiou, O., L. Börtzler, E. Ermolaev, J. Kumari, and J. Schönrich-Sedlmeir (2025). "What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains". English. In: Proceedings of the Annual Hawaii International Conference on System Sciences. ISSN: 1530-1605. URL: https://scholarspace.manoa.hawaii.edu/items/ad447220-c0f5-4004-8877-9130889d12ed.

Petkus, M. (2019). "Why and How zk-SNARK Works". In: CoRR abs/1906.07221. arXiv: 1906.07221. URL: http://arxiv.org/abs/1906.07221.

Politou, E., F. Casino, E. Alepis, and C. Patsakis (2021). "Blockchain Mutability: Challenges and Proposed Solutions". In: IEEE Transactions on Emerging Topics in Computing 9.4, pp. 1972–1986. ISSN: 2168-6750. DOI: 10.1109/tetc.2019.2949510. URL: http://dx.doi.org/10.1109/tetc.2019.2949510.

Quach, S., P. Thaichon, K. D. Martin, S. Weaven, and R. W. Palmatier (2022). "Digital technologies: tensions in privacy and data". In: Journal of the Academy of Marketing Science 50.6, pp. 1299–1323. ISSN: 1552-7824. DOI: 10.1007/s11747-022-00845-y. URL: https://doi.org/10.1007/s11747-022-00845-y.

Rieger, A., F. Guggenmos, J. Lockl, G. Fridgen, and N. Urbach (2019). "Building a Blockchain Application that Complies with the EU General Data Protection Regulation". In: MIS Quarterly Executive 18, pp. 263–279. DOI: 10.17705/2msqe.00020.

Rieger, A., T. Roth, J. Sedlmeir, L. Weigl, and G. Fridgen (2022). "Not yet another digital identity". In: Nature Human Behaviour 6.1, pp. 3–3. ISSN: 2397-3374. DOI: 10.1038/s41562-021-01243-0. URL: https://doi.org/10.1038/s41562-021-01243-0.

Rieger, A., A. Stohr, A. Wenninger, and G. Fridgen (2021). "Reconciling Blockchain with the GDPR: Insights from the German Asylum Procedure". In: Blockchain and the Public Sector: Theories, Reforms, and Case Studies. Ed. by C. G. Reddick, M. P. Rodríguez-Bolívar, and H. J. Scholl. Cham: Springer International Publishing, pp. 73–

95. ISBN: 978-3-030-55746-1. DOI: 10.1007/978-3-030-55746-1_4. URL: https://doi.org/10.1007/978-3-030-55746-1_4.

Rodriguez-Garcia, M., M.-A. Sicilia, and J. M. Dodero (2021). "A privacy-preserving design for sharing demand-driven patient datasets over permissioned blockchains and P2P secure transfer". In: PeerJ Computer Science 7, e568. ISSN: 2376-5992. DOI: 10.7717/peerj-cs.568. URL: http://dx.doi.org/10.7717/peerj-cs.568.

Sabt, M., M. Achemlal, and A. Bouabdallah (2015). "Trusted execution environment: What it is, and what it is not". In: 2015 IEEE Trustcom/BigDataSE/Ispa. Vol. 1. IEEE, pp. 57–64.

Schickler, J. (2025). Tornado Cash Developer Alexey Pertsev Set to Be Released From Prison. https://www.coindesk.com/policy/2025/02/07/tornado-cash-developer-alexey-pertsev-set-to-be-released-from-prison. Accessed: 2025-08-01.

Schlatt, V., A. Schweizer, N. Urbach, and G. Fridgen (2016). Blockchain: Fundamentals, Applications and Potential. German. Tech. rep. URL: https://www.fit.fraunhofer.de/content/dam/fit/de/documents/Blockchain_WhitePaper_Grundlagen-Anwendungen-Potentiale.pdf.

Schmeiss, J., K. Hoelzle, and R. P. G. Tech (2019). "Designing Governance Mechanisms in Platform Ecosystems: Addressing the Paradox of Openness through Blockchain Technology". In: California Management Review 62.1, pp. 121–143. ISSN: 0008-1256. DOI: 10.1177/0008125619883618. URL: http://dx.doi.org/10.1177/0008125619883618.

Schütte, J., G. Fridgen, W. Prinz, T. Rose, N. Urbach, T. Hoeren, N. Guggenberger, C. Welzel, S. Holly, A. Schulte, P. Sprenger, C. Schwede, B. Weimert, B. Otto, M. Dalheimer, M. Wenzel, M. Kreutzer, M. Fritz, U. Leiner, and A. Nouak (2018). Blockchain and smart contracts : Technologies, research issues and applications. s.l. URL: https://www.fim-rc.de/Paperbibliothek/Veroeffentlicht/780/wi-780.pdf.

Sedlmeir, J., H. U. Buhl, G. Fridgen, and R. Keller (2020). "The Energy Consumption of Blockchain Technology: Beyond Myth". In: Business & Information Systems Engineering 62.6, pp. 599–608. ISSN: 1867-0202. DOI: 10.1007/s12599-020-00656-x. URL: https://doi.org/10.1007/s12599-020-00656-x.

Sedlmeir, J., J. Lautenschlager, G. Fridgen, and N. Urbach (2022a). "The transparency challenge of blockchain in organizations". In: Electronic Markets 32.3, pp. 1779–1794. ISSN: 1422-8890. DOI: 10.1007/s12525-022-00536-0. URL: https://doi.org/10.1007/s12525-022-00536-0.

Sedlmeir, J., A. Rieger, T. Roth, and G. Fridgen (2023). "Battling disinformation with cryptography". In: Nature Machine Intelligence 5.10, pp. 1056–1057. ISSN: 2522-5839. DOI: 10.1038/s42256-023-00733-2. URL: https://doi.org/10.1038/s42256-023-00733-2.

Sedlmeir, J., R. Smethurst, A. Rieger, and G. Fridgen (2021). "Digital Identities and Verifiable Credentials". In: Business & Information Systems Engineering 63.5, pp. 603–613. ISSN: 1867-0202. DOI: 10.1007/s12599-021-00722-y. URL: https://doi.org/10.1007/s12599-021-00722-y.

Sedlmeir, J., F. Völter, and J. Strüker (2022b). "The next stage of green electricity labeling: using zero-knowledge proofs for blockchain-based certificates of origin and use". In: SIGENERGY Energy Inform. Rev. 1.1, 20–31. DOI: 10.1145/3508467.3508470. URL: https://doi.org/10.1145/3508467.3508470.

Sidorov, V., E. Y. F. Wei, and W. K. Ng (2022). "Comprehensive performance analysis of homomorphic cryptosystems for practical data processing". In: arXiv preprint arXiv:2202.02960.

Solove, D. J. (2008). Understanding Privacy. Cambridge, MA: Harvard University Press. URL: https://ssrn.com/abstract=1127888.

Spiekermann, S. (2012). "The challenges of privacy by design". In: Commun. ACM 55.7, 38–40. ISSN: 0001-0782. DOI: 10.1145/2209249.2209263. URL: https://doi.org/10.1145/2209249.2209263.

Statista Research Department (2024). Number of crypto users worldwide from 2016 to 2024. https://www.statista.com/statistics/1202503/global-cryptocurrency-user-base/. Accessed: 2025-08-03.

Sun, X., F. R. Yu, P. Zhang, Z. Sun, W. Xie, and X. Peng (2021a). "A Survey on Zero-Knowledge Proof in Blockchain". In: IEEE Network 35.4, pp. 198–205. DOI: 10.1109/MNET.011.2000473.

Sun, X., F. R. Yu, P. Zhang, Z. Sun, W. Xie, and X. Peng (2021b). "A Survey on Zero-Knowledge Proof in Blockchain". In: IEEE Network 35.4, pp. 198–205. DOI: 10.1109/MNET.011.2000473.

Tang, X., L. Shi, X. Wang, K. Charbonnet, S. Tang, and S. Sun (2024). Zero-Knowledge Proof Vulnerability Analysis and Security Auditing. Cryptology ePrint Archive, Paper 2024/514. URL: https://eprint.iacr.org/2024/514.

Thibault, L. T., T. Sarry, and A. S. Hafid (2022). "Blockchain Scaling Using Rollups: A Comprehensive Survey". In: IEEE Access 10, pp. 93039–93054. DOI: 10.1109/ACCESS.2022.3200051.

TRM Labs (2022). North Korea's Lazarus Group Moves Funds Through Tornado Cash. https://www.trmlabs.com/resources/blog/north-koreas-lazarus-group-moves-funds-through-tornado-cash. Accessed: 2025-08-01.

U.S. Department of the Treasury (2022). U.S. Treasury Sanctions Notorious Virtual Currency Mixer Tornado Cash. https://home.treasury.gov/news/press-releases/jy0916. Accessed: 2025-08-01.

Utz, M., S. Johanning, T. Roth, T. Bruckner, and J. Strüker (2023). "From ambivalence to trust: Using blockchain in customer loyalty programs". In: International Journal of Information Management 68, p. 102496. ISSN: 0268-4012. DOI: https://doi.org/10.1016/j.ijinfomgt.2022.102496. URL: https://www.sciencedirect.com/science/article/pii/S0268401222000275.

Walfish, M. and A. J. Blumberg (2015). "Verifying computations without reexecuting them". In: Communications of the ACM 58.2, pp. 74–84. ISSN: 0001-0782. DOI: 10.1145/2641562. URL: http://dx.doi.org/10.1145/2641562.

Wang, C. and M. Gao (2025). "UniZK: Accelerating Zero-Knowledge Proof with Unified Hardware and Flexible Kernel Mapping". In: Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1. ASPLOS '25. New York, NY, USA: Association for Computing Machinery, 1101–1117. ISBN: 9798400706981. DOI: 10.1145/3669940.3707228. URL: https://doi.org/10.1145/3669940.3707228.

Weigl, L., A. Amard, C. Codagnone, and G. Fridgen (2022). "The EU's Digital Identity Policy: Tracing Policy Punctuations". In: Proceedings of the 15th International Conference on Theory and Practice of Electronic Governance. ICEGOV '22. New York, NY, USA: Association for Computing Machinery, 74–81. ISBN: 9781450396356. DOI: 10.1145/3560107.3560121. URL: https://doi.org/10.1145/3560107.3560121.

Weigl, L., T. J. Barbereau, A. Rieger, and G. Fridgen (January 2022). "The Social Construction of Self-Sovereign Identity: An Extended Model of Interpretive Flexibility". English. In: Proceedings of the 55th Hawaii International Conference on System Sciences (HICSS). FNR - Fonds National de la Recherche.

Werner, R., S. Lawrenz, and A. Rausch (2020). "Blockchain Analysis Tool of a Cryptocurrency". In: Proceedings of the 2020 2nd International Conference on Blockchain Technology. ICBCT '20. New York, NY, USA: Association for Computing Machinery, 80–84. ISBN: 9781450377676. DOI: 10.1145/3390566.3391671. URL: https://doi.org/10.1145/3390566.3391671.

Whyte, A. and A. Macintosh (2001). "Transparency and teledemocracy: issues from an 'e-consultation'". In: Journal of Information Science 27.4, pp. 187–198. ISSN: 0165-5515. DOI: 10.1177/016555150102700401. URL: http://dx.doi.org/10.1177/01655515010270 0401.

Wilcox-O'Hearn, Z. (2018). Privacy for Everyone. Accessed on 2025-08-12. URL: https://slideslive.com/38911617/privacy-for-everyone.

Yang, Y., D. Heath, C. Hazay, V. Kolesnikov, and M. Venkitasubramaniam (2024). Tight ZK CPU: Batched ZK Branching with Cost Proportional to Evaluated Instruction. Cryptology ePrint Archive, Paper 2024/456. URL: https://eprint.iacr.org/2024/456.

Yli-Huumo, J., D. Ko, S. Choi, S. Park, and K. Smolander (2016). "Where Is Current Research on Blockchain Technology?—A Systematic Review". In: PLOS ONE 11.10. Ed. by H. Song, e0163477. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0163477. URL: http://dx.doi.org/10.1371/journal.pone.0163477.

Yousuf, H., M. Lahzi, S. A. Salloum, and K. Shaalan (2021). "Systematic Review on Fully Homomorphic Encryption Scheme and Its Application". In: Recent Advances in Intelligent Systems and Smart Applications. Ed. by M. Al-Emran, K. Shaalan, and A. E. Hassanien. Cham: Springer International Publishing, pp. 537–551. ISBN: 978-3-030-47411-9. DOI: 10.1007/978-3-030-47411-9_29. URL: https://doi.org/10.1007/978-3-030-47411-9_29.

Zare-Garizy, T., G. Fridgen, and L. Wederhake (2018). "A Privacy Preserving Approach to Collaborative Systemic Risk Identification: The Use-Case of Supply Chain Networks". In: Security and Communication Networks 2018.1, p. 3858592. DOI: https://doi.org/10.1155/2018/3858592. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1155/2018/3858592. URL: https://onlinelibrary.wiley.com/doi/abs/10.1155/2018/3858592.

Zhang, I., K. Kulkarni, T. Li, D. Wong, T. Kim, J. Guibas, U. Roy, B. Pellegrino, and R. Zarick (2025). "vApps: Verifiable Applications at Internet Scale". In: arXiv preprint arXiv:2504.14809.

Zhao, C., S. Zhao, M. Zhao, Z. Chen, C.-Z. Gao, H. Li, and Y. an Tan (2019). "Secure Multi-Party Computation: Theory, practice and applications". In: Information Sciences 476, pp. 357–372. ISSN: 0020-0255. DOI: https://doi.org/10.1016/j.ins.2018.10.024. URL: https://www.sciencedirect.com/science/article/pii/S0020025518308338.

Zhong, H., Y. Sang, Y. Zhang, and Z. Xi (2020). "Secure Multi-Party Computation on Blockchain: An Overview". In: Parallel Architectures, Algorithms and Programming. Ed. by H. Shen and Y. Sang. Singapore: Springer Singapore, pp. 452–460. ISBN: 978-981-15-2767-8.

VI | Appendix

This appendix chapter is divided into two sections. Section 1 provides an overview of the Research Articles included in this dissertation. Section 2 contains the author's contribution statements for each of the included Research Article. The full-text copies of these Research Articles are appended in the subsequent chapters: Chapter VII, Chapter VIII, Chapter IX, Chapter X, and Chapter XI.

1 Relevant Research Articles

1.1. Articles in this dissertation

- RA1, Ermolaev et al., 2023 – E. Ermolaev, I. Abellán Álvarez, J. Sedlmeir, G. Fridgen. "z-Commerce: Designing a Data-Minimizing One-Click Checkout Solution."
- RA2, Barbereau et al., 2023b – T. Barbereau, E. Ermolaev, M. Brennecke, E. Hartwich, J. Sedlmeir. "Beyond a Fistful of Tumblers: Toward a Taxonomy of Ethereum-based Mixers."
- RA3, Ermolaev et al., 2025 – E. Ermolaev, E. Y. Tseloni, T. Roth, G. Fridgen. "Designing a Reporting System for Trust in Environmental Social Governance."
- RA4, Papageorgiou et al., 2025 – O. Papageorgiou, L. Börtzler, E. Ermolaev, J. Kumari, J. Sedlmeir. "What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains."
- RA5, Ermolaev et al., 2026 – E. Ermolaev, E. Y. Tseloni, T. Roth, G. Fridgen. "Blockchain-based Reporting System for Trust in Environmental Social Governance."

1.2. Other non peer-reviewed publications not included in this dissertation

- Höß et al., 2022b – A. Höß, D. H. MacLennan, A. Rieger, E. Ermolaev, G. Fridgen, T. Roth. Issuing and verifying digital diplomas with the European Blockchain Services Infrastructure - Insights from the EBSILUX project.

2 Author's contribution

Contribution statements of the author (Egor Ermolaev) based on the CRediT system (National Information Standards Organization, 2022).

- RA1 – Single Primary Authorship.

  Conceptualization. Data curation. Investigation. Methodology. Visualization. Writing - original draft, review & editing)

- RA2 – Joint Primary Authorship.

  Conceptualization. Data curation. Investigation. Methodology. Validation. Visualization. Writing - original draft, review & editing.

- RA3 – Single Primary Authorship.

  Conceptualization. Data curation. Investigation. Methodology. Project administration. Software. Supervision. Visualization. Writing - original draft, review & editing.

- RA4 – Non-Primary Authorship.

  Data curation. Software. Writing - original draft, review & editing.

- RA5 – Joint Primary Authorship.

  Conceptualization. Data curation. Investigation. Methodology. Project administration. Software. Visualization. Writing - original draft, review & editing.

VII | Research Paper 1 – z-Commerce Design

Full Title:
z-Commerce: Designing a data-minimizing one-click checkout solution

Publication venue:
18th International Conference on Design Science Research in Information Systems and Technology (DESRIST) 2023

URL:
hdl.handle.net/10993/55406

z-Commerce: Designing a data-minimizing one-click checkout solution*

Egor Ermolaev[0000−0003−3412−5082]**, Iván Abellán Álvarez[0000−0003−4670−433X], Johannes Sedlmeir[0000−0003−2631−8749], and Gilbert Fridgen[0000−0001−7037−4807]

Interdisciplinary Center for Security, Reliability and Trust, University of Luxembourg
{egor.ermolaev,ivan.abellan,johannes.sedlmeir,gilbert.fridgen}@uni.lu
https://wwwen.uni.lu/snt/research/finatrax

Abstract. E-commerce has grown rapidly over the past years, with prevailing e-commerce platforms aggregating large amounts of customer data. This practice has several undesirable side effects, such as facilitating profiling that may lead to price discrimination and data feedback loops that can hamper competition. Moreover, data hoarding carries security risks through data breaches and undermines customers' privacy expectations. On the other hand, convenience aspects and compliance regulation demand the processing and storage of user-related data. To address this tension field, we aim to conceptualize and iteratively refine a data-minimizinig e-commerce platform. Following a design science research approach, we identify design objectives and propose and implement a solution in which stakeholders receive only customer data that is indispensable for their part of the process. Our solution leverages digital identity wallets and general-purpose zero-knowledge proofs (zk-SNARKs). We aim to perform a criteria-based evaluation to assess our artifact's feasibility and fitness from an interdisciplinary perspective. With our results, we hope to illustrate that combining state-of-the-art cryptographic techniques and an emerging digital identity paradigm allows reaching the user experience of incumbent e-commerce platforms while mitigating the undesirable socio-economic side effects of avoidable data disclosure.

Keywords: Compliance, digital wallet, electronic commerce, platform, selective disclosure, privacy, zero-knowledge proof

1 Introduction

Electronic markets facilitate the discovery and coordination of stakeholders such as buyers and sellers, data trading, product matching, and payments through

* This research was funded in part by the Luxembourg National Research Fund (FNR) through the PABLO project (grant reference 16326754) and by PayPal, grant reference "P17/IS/13342933/PayPal-FNR/Chair in DFS/Gilbert Fridgen" (PEARL).

** Corresponding author (e-mail: egor.ermolaev@uni.lu)

digital means [56]. E-commerce is seen as a shift from traditional markets into the digital economy [40]. Corresponding digital platforms use digital technologies to fulfill business and market requirements, yet may also involve physical processes, for instance, for product delivery. E-commerce offers several benefits to merchants and consumers. First, the digital coordination of purchase-related processes, such as business-to-consumer interactions during payment and delivery, makes the corresponding markets ubiquitous [57] and more efficient, especially at large distances. Second, a single integrated platform, among other benefits, enables customers to efficiently discover products or services [65]. Third, e-commerce facilitates delivery based on customer preferences, such as shipment to door, to pick-up points, or just to the closest post office. It thus improves distant goods distribution [43] and benefits customers who would be limited to local sellers with a limited spectrum of available products. Fourth, e-commerce improves user experience by facilitating remote purchase agreements with integrated electronic payments [47]. Studies suggest that in 2022, e-commerce accounted for more than 5 trillion U.S. dollars in retail sales [35] and that e-commerce sales may further increase substantially over the next few years. Particularly large e-commerce platforms such as Amazon have grown enormously. This can be attributed to the presence of indirect network effects, as platforms with the greatest selection of merchants and their goods provide the highest utility to consumers and vice versa [4].

1.1 Benefits of data collection

Data processing is a key aspect of e-commerce due to multiple business requirements that include accounting, legal, recommender systems, and payment processing activities [47]. For instance, a legal requirement would be an age verification for purchasing alcohol, whereas a business requirement could be the disclosure of the residential address for product delivery. In this context, a user first obtains a user account on the e-commerce website by providing personal information. A registration process is not only used by merchants but also underlies payment, as banks or payment service providers (PSPs) need to conduct know-your-customer (KYC) processes based on verifiable personal data retrieved, for instance, from a national ID card [53]. As the registration process is time-consuming for the user and the verification of identity attributes may also be costly for the vendor, e-commerce providers store their customers' identity attributes to have them available for subsequent interactions. At the checkout process, platforms then ensure account ownership via authenticating their users [58]; typically using passwords and potentially additional factors. The inconvenience of registering with multiple service providers by filling out forms not only in e-commerce but also beyond has led to the appearance of identity providers (IdPs) who centrally collect users' identity information to offer single sign-on solutions [54]. To improve user experience, many e-commerce platforms and also independent merchants have integrated this federated identity model to allow users to authenticate with an existing account with their IdP and to transfer corresponding identity attributes without the need for repeated

registration [31]. For instance, Amazon already acts as an IdP for many small vendors to provide their users with a one-click checkout without prior registration, and, thereby, extends its data collection even beyond consumers on its own e-commerce platform.

Data collection also benefits e-commerce platforms beyond improving user experience in the context of identity management. Targeted advertising to promote products is key in the e-commerce business model [3,7]. Recommender systems attract customers by advertising those offers that customers may be most interested in [65]. Personalized product recommendation is known to increase sales [66], which represents a major driver to improve the corresponding recommendation systems, for example, by training machine learning models on collected transaction data [66]. E-commerce platforms that aggregate multiple retail stores and small-sized merchants can, therefore, improve their service and provide more convenience to customers [40]. Data collection also improves the effectiveness of e-commerce platforms toward customer relationships [39]. Customer data gathering helps to tailor relationship management, for example, by making advertising campaigns more compelling, introducing effective customer retention techniques such as loyalty programs, or improving service quality to meet customers' expectations [39].

1.2 Problems in the context of data collection

On the other hand, data collection through e-commerce platforms carries several economic, security, and privacy risks. Large-scale data collection allows for profiling customers, i.e., analyzing their transaction histories and modeling their behavior [34]. This profiling may lead to unfair treatment of customers, such as price discrimination [11]. Moreover, indirect network effects and the corresponding accumulation of market power to large e-commerce platforms are further increased: Direct access to customers' data may make dominant platforms feel tempted to practice anti-competitive measures, for instance, by placing lock-in practices through making data non-portable [15]. Additionally, platforms can also benefit from data feedback loops [29] and data network effects [25], so they could gain market share as advanced data analytics allows them to improve faster. Likewise, these marketplace aggregators, which offer convenience to both retailers and customers, can exploit other revenue streams, such as charging membership, service, or commission fees [40], allowing them to reduce fees and make them more attractive. All these practices significantly increase the market power of large, incumbent e-commerce platforms, ultimately hampering competition [29]. Additionally, collected data is generally stored in centralized silos [12], facing risk of being harvested without users' explicit consent [63]. As such, users have little control over whether their data is being sold to third parties [21]. Moreover, insufficient security measures or targeted attacks on "honey pots" can lead to data leaks or breaches. For instance, a hacker's raid on eBay led to a historic breach that leaked 145M user records in 2014 [48]. It is not only that data breaches pose a risk to privacy, as sensitive information may be exposed to third parties without consent. They also represent a substantial security threat,

as users may be targeted for impersonation or social engineering attacks where publicly available information is exploited for malicious purposes [38]. Opaque data flows, advanced data analytics, and transaction monitoring also raise concerns among users who dislike disclosing personal information [39] or who fear the implementation of surveillance capitalism or an Orwellian state [67].

1.3 Searching for new solutions

We conclude that the handling and processing of customer data need to navigate between convenience and business requirements on the one hand and security, privacy, and socio-economic risks on the other hand. Balancing users' privacy and business and compliance needs that require data collection is already a problem without an easy solution in electronic payments alone [41]. Arguably, balancing out this tension field in e-commerce may be even more challenging as it involves further business and compliance requirements, more complex processes, and additional stakeholders guided by their own interests. Some regulations and initiatives aim to address selected issues; for instance, the general data protection regulation (GDPR) aims to ensure the good and appropriate management of European citizens' personal data [62]. The Payment Services Directive (PSD2) enforces information security of payments conducted by financial institutions and PSPs [18]. Most recently, the Digital Services Act (DSA) introduced a set of measures to protect customer rights and ensure an accountable and fair competitive digital market, targeting "gatekeepers" that include e-commerce platforms such as Amazon [19]. Privacy-enhancing technologies have been proposed also in related areas, such as data markets for the Internet of Things (IoT) [22].

New trends in digital identity and privacy-enhancing technology, such as zero-knowledge proofs (ZKPs), may help pave the way toward convenient and yet privacy-oriented e-commerce solutions. In this paper, we investigate to which extent recent approaches to privacy-oriented digital interactions in the realm of identification [50,54] and payments [17,26,64] can be used for this purpose. Researchers have consistently encouraged the use of cryptography in this context [16,49], and corresponding solutions are increasingly discussed and explored [26,61]. Many approaches rely on ZKPs to facilitate the convenient, selective disclosure of information from users to relying parties, such as digital platforms, service providers, or blockchain-based applications, in a machine-verifiable way. For instance, requirements of an e-commerce transaction from a regulatory or business side should be verified effectively while reducing the processing of sensitive information to a minimum.

However, there has only been limited research on combining such tools in practical applications, particularly in the context of e-commerce. In contrast, most prior research has focused either on exchanging only identity-related data for specific business or administrative processes or on privacy-oriented payments but not on their combination. The work by Schanzenbach et al. [52] poses a notable exception that implements both decentralized identity management and a client-side payment system for e-commerce. However, it has a strong technical perspective, is subject to several limitations, and lacks consideration of how to

coordinate the different stakeholders involved in an e-commerce purchase process. Related work has also not evaluated a data-minimal e-commerce (DMEC) solution with regard to stakeholders' expectations. We believe that the case of e-commerce may indeed illustrate the impact of novel data-minimizing technologies from an individual, economic, and societal perspective, combining both user and business perspectives. Integrating a corresponding wide variety of privacy features in a digital wallet may provide an alternative, "one-click registration and checkout" process that avoids unnecessary data processing and storage while maintaining a high level of user experience [32,51,28]. This direction could also shed light on new directions to preventing the privacy paradox [1].

Our work hence designs and evaluates a proof-of-concept that an alternative, data-minimizing approach to e-commerce is possible. Our system facilitates the selective disclosure of personal and payment information to stakeholders involved in a transaction on an e-commerce platform. We follow the design science research (DSR) methodology by Peffers et al. [45] to develop an artifact that addresses both users' and businesses' needs in the tension field between data use and data minimization. With our artifact, we aim to provide design knowledge on how to build data-minimizing e-commerce solutions. We plan to evaluate the artifact from an interdisciplinary perspective to demonstrate the feasibility and suitability of using digital wallets and ZKPs for achieving minimal information disclosure in e-commerce.

The remainder of this paper is structured as follows: In Section 2, we introduce background knowledge on ZKPs and where they are currently explored in practice in the form of privacy-oriented digital identity management and payments. We then review related work on privacy-oriented e-commerce solutions. Section 3 presents how we aim to conduct our design science research approach in detail and how we believe our research can contribute to the information systems domain.

2 Background and Related Work

2.1 Foundational building blocks

Zero-knowledge proofs: Goldwasser et al. [23] introduced the notion of zero-knowledge in interactive proof systems. It is a property of an interaction between two subjects, a "prover" and a "verifier". The prover probabilistically convinces the verifier of a statement while revealing no additional information about why the statement is true. For instance, the prover could convince the verifier that she knows a solution to a given Sudoku puzzle, without revealing any field of the solution to the verifier. The probability of a malicious prover convincing the verifier of a wrong statement decreases exponentially in the number of interactions between the prover and the verifier. Zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARK) are a family of ZKP that eliminates the need for interactions between the prover and the verifier by obtaining randomness using cryptography. zk-SNARKs also satisfy an additional property that makes them extremely efficient from the perspective of the verifier: They are succinct,

i.e., the resulting proof is very short, and its verification is much quicker than verifying the statement directly through naive re-computation [27]. General-purpose zero-knowledge succinct non-interactive arguments of knowledge (zk-SNARKs) have rapidly evolved during the last few years because of their natural fit for solving technical challenges related to blockchain technology [55,9], with emerging domain-specific languages (DSLs) (e.g., Circom, ZoKrates), libraries (e.g. Circomlib) and tools (e.g. SnarkJS) that allow software engineers to implement ZKPs for a broad class of statements. This development has accelerated research and application of zero-knowledge technology also beyond blockchains, such as in digital identity [6], payment systems [26], and supply chains [42].

Digital identity: Decentralized or self-sovereign identity (SSI) is a concept that involves representing and sharing attributes or authorizations of individuals or organizations, with the aim of providing a more convenient alternative to traditional IdPs while empowering individuals to control data disclosure using digital wallets. SSI's core principles were first described by Allen [2] and later extended with insights from practical experiences [54]. They include not only user-centric requirements such as convenience aspects, control, and data-minimization but also organizational requirements, such as verifiability and authenticity. For instance, a notable application of ZKP is "anonymous credentials" [14] in the realm of digital identity. A prover can ensure a verifier about adulthood without disclosing the date of birth registered in his digital identity by means of a digitally signed ID card issued, for instance, by a national institution [54]. In this context, ZKPs guarantee the verifiability of selectively disclosed attributes, although the underlying digitally signed attestation is never shared. Thereby, ZKPs can be used for selective disclosure – "the ability of an individual to granularly decide what information to share" [60] – and privacy-by-design solutions. A popular implementation is Hyperledger Aries, which presents one of the first efforts for technical specifications and implementation of SSI using blockchain technology [53]. It uses Hyperledger AnonCreds, a library, and specification of anonymous credentials conceptualized by Camenisch & Lysyanskaya [14] that facilitate advanced privacy features such as the selective disclosure of attributes and the avoidance of unique identifiers using special purpose ZKPs. Several public and private sector projects, such as the public-private consortium IDunion, use Hyperledger Aries to implement a SSI-based digital identity management solution that addresses businesses' needs and that is compliant with the EU GDPR [5]. Recently, ZKPs are discussed in the context of the revision of the electronic IDentification, Authentication and trust Services (eIDAS) regulation[13]. These and similar designs support limited functionalities due to the cryptographic primitives they use [6]. However, using these components in the context of e-commerce requires support for adoption at a large scale and more flexible and sophisticated predicates. Therefore, new research focuses on extending the data minimization capabilities of anonymous credentials using general-purpose ZKPs. For instance, several works including [50] and [6] adopt zk-SNARKs to prove complex composable statements about identity attributes without disclosing the signed attestation and to address some scalability shortcomings of the AnonCreds implementation.

Payment systems represent an essential component of e-commerce. Cash is not sufficient for e-commerce due to limitations regarding convenience, such as remote purchasing, and security requirements for a safe transfer. Therefore, electronic payment systems seem more suitable. During the payment process, customers then need to reveal sensitive information to address compliance, technical, and processing requirements to their PSPs. Consequently, PSPs typically have access to transaction data and history. Also, the approach of sharing credit card information with the merchant, including legal name, credit card number, and security code, allows for transaction traceability and is far from data minimizing. To address privacy concerns in electronic payments. Chaum [16] introduced "e-cash" using blind signatures, which provides privacy to the payer. This solution is not flexible enough as the payee is fully transparent. ZeroCash [10] describes a completely anonymous digital currency leveraging ZKP. However, these approaches to anonymous payments do not comply with money laundering and terrorist financing regulation [26]. GNU Taler [17], another implementation of privacy-oriented payment systems based on Chaum's initial e-cash approach, addresses compliance issues by introducing the "auditor" role to which both users and payment processors are accountable. It provides a privacy-friendly payment system solution for the payer, who remains anonymous. Several recent works present alternative centralized or decentralized payment systems to address both privacy demands and compliance by enforcing turnover limits for anonymous transactions [26,64]. These architectures can be considered the next iteration of privacy-focused, ZKPs and blockchain-based cryptocurrencies such as Zcash [10] that in contrast have a solid chance of achieving regulatory compliance.

2.2 Related work

Data markets: DISSENS [52] provides decentralized identity management and a client-side payment system for e-commerce. It aims to be a regulatory-compliant solution built on a decentralized network, specifically a distributed hash table (DHT), that acts as an encrypted file storage for digital identity. Integrating the GNU Taler [17] payment system ensures client-side privacy features and provable regulatory compliance. However, several limitations prevail. Data, although encrypted, is shared on a public network, which may conflict with the GDPR's "right to be forgotten". Its functionality is limited to the disclosed attributes, so extensive functionality is not possible, such that proving general statements concerning one or multiple attributes. Agora [37] proposes a semi-private marketplace for data brokerage. Users generate encrypted data which upon payment brokers decrypt and batch. Brokers act as an unlinkability and aggregator service to interested data buyers, thus protecting users' data in front of end consumers. Due to the nature of the cryptographic algorithms used, their cryptographic protocols are limited in functionality. Agora only supports private data sharing of special mathematical function outputs, such as weighted averages and linear regressions. The paper highlights that functionality could

be improved by leveraging ZKP. For instance, Garrido et al. [22] survey different privacy-enhancing technologies (PETs) in IoT data markets and illustrate their trade-offs. Data brokerage may use various PETs to enable aggregation and obfuscation of the initial private data that makes it irreversible and untraceable. Similarly, a centralized approach is proposed for data collection and service delivery [44]. Data generators use a particular cryptographic primitive to protect and preserve privacy by computing encrypted data. However, the system relies on a third party to create participants' identities, which may pose a risk to data privacy guarantees. If eventually, the registration center becomes corrupted, it could intercept and decrypt the data, thus correlating users to content. Additionally, a tamper-proof device is required, which usually poses an encumbrance to widespread adoption. Bella et al. [8] suggests an e-commerce architecture that balances privacy and trust by using differential privacy. Differential privacy introduces noise to make user data less sensitive, which brings fuzziness to the characteristics of each customer (e.g., adulthood). Hence, differential privacy approaches may not be adequate, particularly when clear boundaries are imposed by regulatory compliance considerations.

Besides some technical limitations of related work that we aim to address, we also note that related work, except for DISSENS [52], so far has not evaluated their solutions with stakeholders, for instance, to assess whether the approach can meet business needs and customers' user experience needs.

3 Method

In this paper, we outline how we aim to create general prescriptive knowledge on reducing the processing and storage of sensitive information in e-commerce following the DSR method [45]. Based on a demonstration of the feasibility of such a design, we want to identify how this approach can help avoid the aggregation of data by incumbent e-commerce platforms and the corresponding security and socio-economic challenges we discussed in Section 1. Our research hence provides an interdisciplinary perspective from an information systems lens by developing a solution based on novel cryptographic tools. Our research also addresses calls for more widespread applications of cryptography for "moral" reasons, for instance, to tackle increasing surveillance threats [49].

Our paper is structured according to the DSR best practices as described by Gregor & Hevner [24], yet considers the limitation of our paper's scope as it represents research in progress. We first comprehensively identify the problem and motivation. A valuable DSR artifact must satisfy a certain societal or business need [30] to determine the relevance and value of the contribution to the IS research field. According to Section 1, the accumulation of customer data and the corresponding economic implications (e.g., opportunities for price discrimination based on customer profiling and negative consequences for competition through data feedback loops), security threats, and moral issues pose such a need [49]. We aim to ensure both the relevance of our research and the fitness of our artifact by first more clearly identifying the research problem, identifying the

roles, tasks, and requirements of the relevant stakeholders in a systematic literature review, and deriving comprehensive objectives of the solution in the future. As we argued in Section 1, the objectives will involve convenience expectations regarding end-user experience, regulatory compliance, and data minimization. We plan to validate and potentially extend these requirements based on expert interviews during the first evaluation cycle of our DSR. We plan to apply the convenience sampling method for interviews. We will form groups of interviewees with respect to their domain of expertise and involve also end-users to have a balanced sample. The selection criteria will correspond to the domains involved in our research, such as cryptographers or IT security researchers, and experts from the business side on e-commerce and adjacent service providers, such as logistics.

We follow the design science research methodology (DSRM) proposed by Peffers et al. [45]. Fig. 1 features the corresponding iterative build-and-evaluation cycles. We aim to design and develop our artifact – a DMEC architecture and corresponding information flow that only shares and stores information that is indispensable for each stakeholder and, therefore, mitigates undesirable side effects such as price discrimination against consumers and accumulation of market power. The design will also define onboarding and end-to-end purchasing processes for customers. The build-and-evaluate process is at the core of DSR and helps IS researchers discover answers to problems that have not been resolved before [30]. The DSR method emerged to combine methods from engineering and social sciences for practically relevant, rigorous research [24]. Thus, one can apply the DSR approach to a broad spectrum of domains. The design of a DMEC platform addresses the tension field between convenience, compliance, and data minimization and, therefore, affects both technical and non-technical considerations that need to be incorporated in the development from an engineering and social sciences perspective. Yet, we believe that the implications of our DSR go far beyond e-commerce and may be applicable to adjacent realms such as e-government [20], healthcare [33], the industrial IoT [59], which also seem to involve tensions between data sharing needs and data protection considerations. Particularly blockchain-based applications require such data minimization by design owing to the inherent transparency of distributed ledgers [55,46].

During the process iterations of the DSRM process model, we also instantiate our artifact in the form of a prototype based on the architecture. We will use this instantiation both for the first set of more technical design iterations and for demonstration in the interview-based evaluation. For preparing the implementation, we will define the components of the high-level system architecture and their functionality. Then we will search for potential open-source solutions which can be used in the components. Particularly, we will analyze which open-source solutions could help us to implement the ZKP stack that provides selective disclosure. There will be multiple criteria for ZKP stack selection, such as the performance of proof generation (especially when considering a prototype for e-commerce on mobile devices). So far, we have closely analyzed several promising

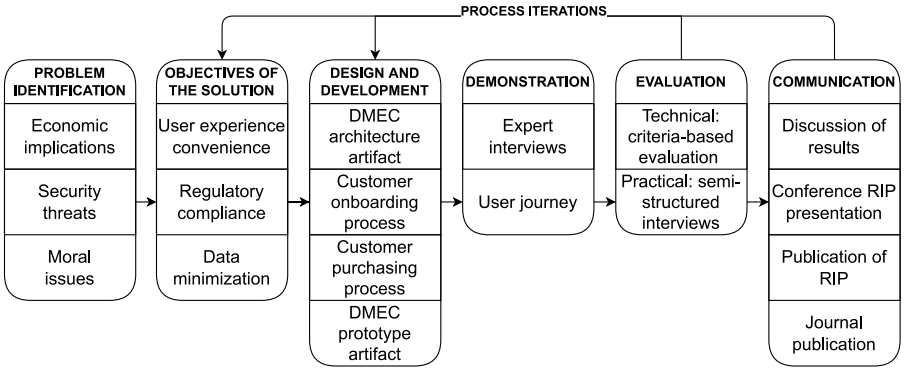

Fig. 1. Our applied DSR approach to design a DMEC architecture following Peffers et al. [45].

open-source solutions as candidates and assessed their potential extension for the needs of our architecture.

The Heimdall framework [6] promises a good fit for our technical requirements. As we highlighted in Section 2, related work on privacy-oriented e-commerce solutions still lacks essential features. The Heimdall framework provides several of these features as it implements anonymous credentials that facilitate seamlessly and verifiably sharing the minimum identity-related data required by the relying party in a scalable way [6]. In particular, Heimdall provides a modular set of tools to create, issue, verify, and manage verifiable credentials. Concretely, the framework supports "prover" and "verifier" functions. Users are able to manage and store issued verifiable credentials. They are also capable of generating data minimizing presentations upon verifiers' request in a private way (e.g., selectively disclosing personal information). They do so by means of ZKP, in particular zk-SNARK. Nonetheless, we envision the need to extend Heimdall's functionality to make it compatible with all the project's needs, which is feasible because zk-SNARKs are general-purpose and can be used to prove any statement. The envisioned artifact is in the form of an e-commerce architecture consisting of four main components, namely a checkout page, a digital identity wallet, a back-end architecture, and an agent. The checkout page implements potential checkout options of an e-commerce platform. The digital identity wallet, in which the customer manages identities, serves as a wallet application to share the minimal necessary information with the merchant. When customers indicate their intention to purchase through the website, this communicates with the e-commerce back-end so that the agent establishes a connection with the digital identity wallet. Customers are then able to generate proofs, which are based on the credentials stored in the wallet, in response to the agent's request. The agent verifies and confirms the user-generated proof and sends the appropriate user data to the checkout page.

During the iterations, we will synchronize the evolved design of the system architecture and development direction. The iterations will go in parallel for all the artifacts: the design, processes, and prototype. At the end of each iteration, the quality of the artifacts will be evaluated both technically (performance benchmarking, as mentioned before) and practically (interviews with relevant stakeholders). For demonstration and continuous improvement, we will conduct multiple interview cycles and evaluate user journeys with groups of people, which will contribute to the design and development as part of the process iterations by disseminating the intermediate progress with the stakeholders and receiving their feedback. Interviews will be essential for collecting feedback regarding the practical side of the prototype to evaluate user experience. There are only a few studies on how users experience the use of digital wallets, particularly data minimization capabilities [51,36], and we believe that presenting users with a familiar flow, such as commerce, may help gain new insights. The collected feedback will also be implemented according to the DSRM process model.

The envisioned artifact is in the form of an e-commerce architecture consisting of four main components, namely a checkout page, a digital identity wallet, a back-end architecture (controller), and an agent. The checkout page implements a checkout option – "checkout with SSI". The digital identity wallet allows customers to share the minimal identity information necessary with the merchant and other stakeholders, such as logistics service providers. We will use the design tools for layouts of both front-end components (checkout page and wallet). When a customer indicates their intention to purchase through the website, the checkout page communicates with the merchant back-end so that the agent establishes a connection with the digital identity wallet and sends a proof request to it. Upon the customer's confirmation, the wallet generates a corresponding cryptographic proof (SNARK) based on the credentials stored in the wallet that addresses the agent request and sends it back to the verifying component of the agent. The agent cryptographically verifies the proof and notifies the merchant's back-end, which transmits the verified identity attributes to the checkout page for a corresponding user.

At the evaluation stage, we will validate the objectives of the solution by means of criteria-based evaluation and semi-structured interviews with stakeholders from both interdisciplinary and specific domains (e.g., business, law, user experience, and software engineering). We also plan to conduct user experience evaluations with our stakeholders. From the business-related interviews, we also aim to investigate the business opportunities of such a solution. At the communications stage, we will focus on the presentation and discussion of the research-in-progress paper at the conference. The second stage would be to publish the full paper. In the end, we will open-source the code of our prototype on GitHub.

4 Conclusion

The increasing market growth of e-commerce has led to the aggregation of large amounts of customer data. With the proliferation of incidents related to sensitive data breaches and abuses, both users and regulators are taking steps to protect privacy rights. In this light, we follow the design science research methodology according to Peffers et al. [45] to identify the feasibility of an e-commerce solution that addresses the tension field between convenience aspects and compliance regulation demand the processing and storage of user-related data. Following DSRM, we aim to derive the system architecture of DMEC and build a corresponding prototype. Our proposed solution utilizes digital identities and zero-knowledge proof as core components for privacy-oriented e-commerce transactions. Using an interdisciplinary, criteria-based evaluation, we aim to demonstrate that our artifact can address the societal and business needs that we previously discussed and serve as a starting point for many relevant studies in information systems on usable privacy.

References

1. Alashoor, T., Keil, M., Smith, H.J., McConnell, A.R.: Too tired and in too good of a mood to worry about privacy: Explaining the privacy paradox through the lens of effort level in information processing. Information Systems Research (2022)
2. Allen, C.: The path to self-sovereign identity (2016), http://www.lifewithalacrity.com/2016/04/the-path-to-self-soverereign-identity.html
3. Alt, R.: Electronic markets on business model development. Electronic Markets **30**(3), 405–411 (2020)
4. Alt, R.: Electronic markets on platform transformation. Electronic Markets **32**(2), 401–409 (2022)
5. Anke, J., Richter, D.: Digitale identitäten. HMD Praxis der Wirtschaftsinformatik (2023)
6. Babel, M., Sedlmeir, J.: Bringing data minimization to digital wallets at scale with general-purpose zero-knowledge proofs (2023), http://arxiv.org/abs/2301.00823
7. Baethge, C., Klier, J., Klier, M.: Social commerce – state-of-the-art and future research directions. Electronic Markets **26**(3), 269–290 (2016)
8. Bella, G., Giustolisi, R., Riccobene, S.: Enforcing privacy in e-commerce by balancing anonymity and trust. Computers & Security **30**(8), 705–718 (2011)
9. Ben-Sasson, E., Bentov, I., Horesh, Y., Riabzev, M.: Scalable, transparent, and post-quantum secure computational integrity (2018), https://eprint.iacr.org/2018/046
10. Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., Virza, M.: Zerocash: Decentralized anonymous payments from Bitcoin. In: Proceedings of the IEEE Symposium on Security and Privacy. pp. 459–474 (2014)
11. Bergemann, D., Brooks, B., Morris, S.: The limits of price discrimination. American Economic Review **105**(3), 921–57 (2015)
12. Braud, A., Fromentoux, G., Radier, B., Le Grand, O.: The road to European digital sovereignty with Gaia-X and IDSA. IEEE Network **35**(2), 4–5 (2021)

13. Busch, C.: eidas 2.0: Digital identity service in platform economy (2022), https://cerre.eu/wp-content/uploads/2022/10/CERRE_Digital-Identity_Issue-Paper_FINAL-2.pdf
14. Camenisch, J., Lysyanskaya, A.: An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In: Proceedings of International Conference on the Theory and Applications of Cryptographic Techniques. pp. 93–118 (2001)
15. Camp, L.J., Osorio, C.A.: Privacy-enhancing technologies for internet commerce (2002), https://papers.ssrn.com/abstract=329282
16. Chaum, D.: Security without identification: Transaction systems to make Big Brother obsolete. Communications of the ACM **28**(10), 1030–1044 (1985)
17. Dold, F.: The GNU Taler system: Practical and provably secure electronic payments (2019), https://syntheses.univ-rennes1.fr/search-theses/notice.html?id=rennes1-ori-wf-1-12183&printable=true
18. European Central Bank: The revised payment services directive (PSD2) (2018), https://www.ecb.europa.eu/paym/intro/mip-online/2018/html/1803_revisedpsd.en.html
19. European Comission: The digital services act: Ensuring a safe and accountable online environment (2022), https://ec.europa.eu/info/strategy/priorities-2019-2024/europe-fit-digital-age/digital-services-act-ensuring-safe-and-accountable-online-environment_en
20. Fedorowicz, J., Gogan, J.L., Culnan, M.J.: Barriers to interorganizational information sharing in e-government: A stakeholder analysis. The Information Society **26**(5), 315–329 (2010)
21. Fienberg, S.E.: Privacy and confidentiality in an e-commerce world: Data mining, data warehousing, matching and disclosure limitation. Statistical Science **21**(2), 143–154 (2006)
22. Garrido, G.M., Sedlmeir, J., Uludağ, Ö., Alaoui, I.S., Luckow, A., Matthes, F.: Revealing the landscape of privacy-enhancing technologies in the context of data markets for the IoT: A systematic literature review. Journal of Network and Computer Applications **207**, 103465 (2022)
23. Goldwasser, S., Micali, S., Rackoff, C.: The knowledge complexity of interactive proof systems. SIAM Journal on Computing **18**(1), 186–208 (1989)
24. Gregor, S., Hevner, A.R.: Positioning and presenting design science research for maximum impact. MIS Quarterly **37**(2), 337–355 (2013)
25. Gregory, R.W., Henfridsson, O., Kaganer, E., Kyriakou, H.: The role of artificial intelligence and data network effects for creating user value. Academy of Management Review **46**(3), 534–551 (2021)
26. Gross, J., Sedlmeir, J., Babel, M., Bechtel, A., Schellinger, B.: Designing a central bank digital currency with support for cash-like privacy (2021), https://papers.ssrn.com/abstract=3891121
27. Groth, J.: On the size of pairing-based non-interactive arguments. In: Advances in Cryptology – EUROCRYPT, pp. 305–326. Springer (2016)
28. Guggenberger, T., Neubauer, L., Stramm, J., Völter, F., Zwede, T.: Accept me as I am or see me go: A qualitative analysis of user acceptance of self-sovereign identity applications. In: Proceedings of the 56th Hawaii International Conference on System Sciences (2023)
29. Hermes, S., Kaufmann-Ludwig, J., Schreieck, M.: A taxonomy of platform envelopment: Revealing patterns and particularities. In: Proceedings of the 26th Americas Conference on Information Systems (2020)

30. Hevner, A., March, S.T., Park, J., Ram, S., et al.: Design science research in information systems. MIS Quarterly **28**(1), 75–105 (2004)
31. Jøsang, A., Fabre, J., Hay, B., Dalziel, J., Pope, S.: Trust requirements in identity management. In: Proceedings of the 44th Australasian Workshop on Grid Computing and e-Research. pp. 99–108 (2005)
32. Jørgensen, K.P., Beck, R.: Universal wallets. Business & Information Systems Engineering **64**(1), 115–125 (2022)
33. Kaye, J.: The tension between data sharing and the protection of privacy in genomics research. Annual Review of Genomics and Human Genetics **13**(1), 415–431 (2012)
34. Kayes, I., Iamnitchi, A.: Privacy and security in online social networks: A survey. Online Social Networks and Media **3–4** (2017)
35. Keenan, M.: Global e-commerce: stats and trends to watch (2022), `https://www.shopify.com/enterprise/global-ecommerce-statistics`
36. Khayretdinova, A., Kubach, M., Sellung, R., Roßnagel, H.: Conducting a Usability Evaluation of Decentralized Identity Management Solutions. In: Selbstbestimmung, Privatheit und Datenschutz : Gestaltungsoptionen für einen europäischen Weg, pp. 389–406. Springer Fachmedien Wiesbaden (2022)
37. Koutsos, V., Papadopoulos, D., Chatzopoulos, D., Tarkoma, S., Hui, P.: Agora: A privacy-aware data marketplace. IEEE Transactions on Dependable and Secure Computing **19**(6), 3728–3740 (2022)
38. Krombholz, K., Hobel, H., Huber, M., Weippl, E.: Advanced social engineering attacks. Journal of Information Security and Applications **22**, 113–122 (2015)
39. Kumar, V., Reinartz, W.: Customer privacy concerns and privacy protective responses. In: Customer Relationship Management, pp. 285–309. Springer (2018)
40. Lee, C.: An analytical framework for evaluating e-commerce business models and strategies. Internet Research **11**(4), 349–359 (2001)
41. Maseeh, H.I., Jebarajakirthy, C., Pentecost, R., Arli, D., Weaven, S., Ashaduzzaman, M.: Privacy concerns in e-commerce: A multilevel meta-analysis. Psychology & Marketing **38**(10), 1779–1798 (2021)
42. Mattke, J., Maier, C., Hund, A.: How an enterprise blockchain application in the U.S. pharmaceuticals supply chain is saving lives. MIS Quarterly Executive **18**(4), 246–261 (2019)
43. Morganti, E., Seidel, S., Blanquart, C., Dablanc, L., Lenz, B.: The impact of e-commerce on final deliveries: Alternative parcel delivery services in France and Germany. Transportation Research Procedia **4**, 178–190 (2014)
44. Niu, C., Zheng, Z., Wu, F., Gao, X., Chen, G.: Achieving data truthfulness and privacy preservation in data markets'. IEEE Transactions on Knowledge and Data Engineering **31**(1), 105–119 (2019)
45. Peffers, K., Tuunanen, T., Rothenberger, M.A., Chatterjee, S.: A design science research methodology for information systems research. Journal of Management Information Systems **24**(3), 45–77 (2007)
46. Platt, M., Bandara, R.J., Drăgnoiu, A.E., Krishnamoorthy, S.: Information privacy in decentralized applications. In: Trust Models for Next-Generation Blockchain Ecosystems, pp. 85–104. Springer (2021)
47. Qin, Z.: Introduction to e-commerce. Springer (2009)
48. Reuters, CNBC: Hackers raid eBay in historic breach, access 145M records (2014), `https://www.cnbc.com/2014/05/22/hackers-raid-ebay-in-historic-breach-access-145-mln-records.html`
49. Rogaway, P.: The moral character of cryptographic work (2015), `https://eprint.iacr.org/2015/1162`

50. Rosenberg, M., White, J., Garman, C., Miers, I.: zk-creds: Flexible anonymous credentials from zkSNARKs and existing identity infrastructure (2022), https://eprint.iacr.org/2022/878
51. Sartor, S., Sedlmeir, J., Rieger, A., Roth, T.: Love at first sight? A user experience study of self-sovereign identity wallets. In: Proceedings of 30th European Conference on Information Systems (2022)
52. Schanzenbach, M., Grothoff, C., Wenger, H., Kaul, M.: Decentralized identities for self-sovereign end-users (DISSENS). In: Proceedings of Open Identity Summit. pp. 47–58 (2021)
53. Schlatt, V., Sedlmeir, J., Feulner, S., Urbach, N.: Designing a framework for digital KYC processes built on blockchain-based self-sovereign identity. Information & Management **59**(7), 103553 (2022)
54. Sedlmeir, J., Huber, J., Barbereau, T., Weigl, L., Roth, T.: Transition pathways towards design principles of self-sovereign identity. In: Proceedings of the 43rd International Conference on Information Systems (2022)
55. Sedlmeir, J., Lautenschlager, J., Fridgen, G., Urbach, N.: The transparency challenge of blockchain in organizations. Electronic Markets **32**, 1779–1794 (2022)
56. Stahl, F., Schomm, F., Vossen, G., Vomfell, L.: A classification framework for data marketplaces. Vietnam Journal of Computer Science **3**(3), 137–143 (2016)
57. Targett, D.: B2B or not B2B? Scenarios for the future of e-commerce. European Business Journal **13**(1) (2001)
58. Trautman, L.J.: E-commerce, cyber, and electronic payment system risks: Lessons from PayPal (2016), https://papers.ssrn.com/abstract=2314119
59. Ukil, A., Bandyopadhyay, S., Pal, A.: IoT-privacy: To be private or not to be private. In: Proceedings of the Conference on Computer Communications Workshops. pp. 123–124 (2014)
60. W3C: Engineering privacy for verified credentials (2022), https://w3c-ccg.github.io/data-minimization/#selective-disclosure
61. Weigl, L., Barbereau, T.J., Rieger, A., Fridgen, G.: The social construction of self-sovereign identity: An extended model of interpretive flexibility. In: Proceedings of the 55th Hawaii International Conference on System Sciences. pp. 2543–2552 (2022)
62. Wolford, B.: What is GDPR, the EU's new data protection law? (2018), https://gdpr.eu/what-is-gdpr/
63. van der Wolk, A., Silva, K.: Insight: A slap on the wrist or show of force – GDPR fines reveal need for EU penalty guidelines (2019), https://news.bloomberglaw.com/privacy-and-data-security/insight-a-slap-on-the-wrist-or-show-of-force-gdpr-fines-reveal-need-for-eu-penalty-guidelines
64. Wüst, K., Kostiainen, K., Delius, N., Capkun, S.: Platypus: A central bank digital currency with unlinkable transactions and privacy-preserving regulation. In: Proceedings of the ACM SIGSAC Conference on Computer and Communications Security. pp. 2947–2960 (2022)
65. Youlong Zhuang, Albert L. Lederer: An instrument for measuring the business benefits of e-commerce retailing. International Journal of Electronic Commerce **7**(3), 65–99 (2003)
66. Zhou, L.: Product advertising recommendation in e-commerce based on deep learning and distributed expression. Electronic Commerce Research **20**(2), 321–342 (2020)
67. Zuboff, S.: Big other: Surveillance capitalism and the prospects of an information civilization. Journal of Information Technology **30**(1), 75–89 (2015)

VIII | Research Paper 2 – Taxonomy of Crypto Mixers

Full Title:

Beyond a Fistful of Tumblers: Toward a Taxonomy of Ethereum-based Mixers

Publication venue:

23rd IEEE/ACIS International Conference on Computer and Information Science (ICIS) 2023

URL:

hdl.handle.net/10993/57137

Beyond a Fistful of Tumblers: Toward a Taxonomy of Ethereum-based Mixers

Completed Research Paper

Tom Barbereau, Egor Ermolaev,
Martin Brennecke, Eduard Hartwich, Johannes Sedlmeir
SnT, University of Luxembourg
Luxembourg, Luxembourg
{tom.barbereau, egor.ermolaev,
martin.brennecke, eduard.hartwich, johannes.sedlmeir}@uni.lu

Abstract

The role played by decentralised services in the obfuscation of crypto-asset transactions performed on transparent blockchains has increasingly captured the attention of regulators. This is exemplified by the headlines about the U.S. Treasury's sanctions on the Ethereum-based mixer Tornado Cash. Yet, despite the existing controversies on the use of mixers, the different functionalities of these information systems with an inherent dark side remain to be explored by the literature. So far, contributions primarily encompass technical works and studies that focus on the Bitcoin ecosystem. This paper puts forward a multi-layer taxonomy of the smart-contract-based – and, therefore, functionally richer – family of mixers on Ethereum. Our proposed taxonomy is grounded on (1) a review of existing literature, (2) an analysis of mixers' project documentation, (3) their corresponding smart contracts, and (4) expert interviews. Our evaluation included the application of the taxonomy to two mixers – RAILGUN and zkBob. The taxonomy represents a valuable tool for law enforcement, regulators, and other stakeholders to explore critical properties affecting compliance and use of Ethereum-based mixers.

Keywords: Anonymisation, classification, compliance, crypto-asset, digital forensics, zero-knowledge proof

Introduction

Despite the transparent nature of public blockchains and the possibility to de-anonymise pseudonymous addresses by linking transactions (Biryukov et al., 2014; Meiklejohn et al., 2013), the extent of criminal activities in the crypto-asset space reached an all-time high in 2022 (Chainalysis, 2023). One of the key tools facilitating these activities are information systems that can be used to obfuscate transaction graphs: so-called crypto-asset "tumblers" or "mixers". The total value of crypto-assets processed by these systems in 2022 alone is estimated to stand at around $10 billion, with at least 24 % originating from illicit sources (Chainalysis, 2023). In August 2022, the public saw the tip of the criminal iceberg when the U.S. Department of the Treasury (2022) sanctioned the Ethereum-based mixer Tornado Cash for "laundering" some $7 billion in crypto-assets; of which $455 million were linked to the cyber-criminal Lazarus Group affiliated with the North Korean regime. In an unprecedented turn, the Treasury's Office of Foreign Assets Control (OFAC) listed smart contract addresses associated with Tornado Cash – making it illegal for U.S. citizens to receive or send assets through this service (Nadler & Schär, 2023). What followed is a cascade of events: the Dutch

authorities arrested one of the mixer's co-founders (FIOD, 2022), the source code on GitHub was deleted and later republished (Shen, 2022), and most recently, the other two co-founders of Tornado Cash were arrested (U.S. Department of the Treasury, 2023).

Typically, scholars differentiate between mixers that are (A) provided by a trusted third-partys (TTPs) acting as custodians; (B) implemented in non-custodial wallets, based on peer-to-peer interactions with other users; or (C) run in a non-custodial and non-interactive way, based on smart contracts that allow users to make the deposit and withdrawal of assets unlinkable with the help of sophisticated cryptographic tools (Béres et al., 2021; Nadler & Schär, 2023; Trozze et al., 2023). Type C is exclusive to smart contract blockchains like Ethereum and arguably most popular given the superior obfuscation provided by zero-knowledge proofs (ZKPs) (Burleson et al., 2022; Nadler & Schär, 2023; Wang et al., 2022).Thus far, mixers' functionalities remain largely understudied beyond research in computer science. Related work has focused on technical improvements to performance (Cirkovic et al., 2022), opportunities to improve anonymity guarantees (Wang et al., 2022), and how to design accountability mechanisms and compliance features (Burleson et al., 2022). To our knowledge, the most comprehensive secondary studies on mixers are those of Feng et al. (2019) and Pakki et al. (2021). Both are classifications that primarily focus on Bitcoin-based mixers. This is a considerable limitation for two reasons: (1) in parallel to an increase in illicit activities on the Ethereum blockchain and superior anonymisation techniques, the corresponding Ethereum-based mixers became those most used (Wang et al., 2022), and (2) the flexibility of smart contracts allows for greater variability in functionalities that impact, for instance, governance-related, economic, and regulatory aspects. We found only two works that address these type C mixers, and both focus solely on Tornado Cash. Nadler and Schär (2023) analysed its design and transaction history. Béres et al. (2021) attempted to de-anonymise transactions processed with it. Hence, to distinguish different aspects related to the functional level of mixers, we pose the following research question: *How to classify characteristics of non-custodial, non-interactive, Ethereum-based mixers?*

This paper systematically derives a taxonomy of Ethereum-based mixers following Nickerson et al. (2013) and Kundisch et al. (2021). Taxonomies are considered a suitable artefact to classify, categorise, and systematise objects and clarify complex fields of interest (Bailey, 1994). Our scope is limited to Ethereum-based mixers implemented via smart contracts. These platforms utilise the Ethereum Virtual Machine (EVM) or EVM-compatible virtual machines (Wang et al., 2022). Our taxonomy development comprises four cycles that respectively consider (1) existing peer-reviewed works, (2) the public documentation of four Ethereum-based mixers, (3) their underlying smart contracts, and (4) expert interviews. Following Kundisch et al. (2021), we further evaluate the final taxonomy by applying it to two mixers, RAILGUN and zkBob.

Our taxonomy contributes to the IS literature on "digital forensics" (Nance et al., 2009) and the "dark side of information technology use" (D'Arcy et al., 2014; Tarafdar et al., 2013). Its development process is inspired by a body of research that combines both on-chain and off-chain data sources related to blockchain-based systems (e.g., Barbereau et al., 2023; Zheng et al., 2021). Our research also responds to calls by the Basel Institute on Governance and Europol (2022) for a greater "understanding of crypto-assets and [associated] services", given that "money laundering schemes can often involve multiple blockchains and mixers located in different jurisdictions". Our research is of a descriptive, explanatory nature (Gregor, 2006): it allows practitioners to understand functionalities and distinguish between mixers more accurately. Classification, including taxonomies, can represent valuable tools as part of the evidence collection process and intelligence cycle in law enforcement and criminal intelligence (see, e.g., Fröwis et al., 2020; Sarre et al., 2018).

Background

In this section, we introduce the technical foundations of Ethereum-based mixers, including a description of the properties and use of ZKPs in this context. Its aim is less to provide a detailed description of each potential design choice of an implementation, but rather to provide a high-level perspective of one foundational and prevalent approach to implementing smart-contract-based mixers. Second, the motivation to provide these foundations lies in the prevalent use of jargon and concepts from computer science and cryptography. By introducing these accordingly, we aim to make our paper accessible to readers that perhaps lack a deep technical background. We restrict our scope to concepts indispensable for the reader to grasp

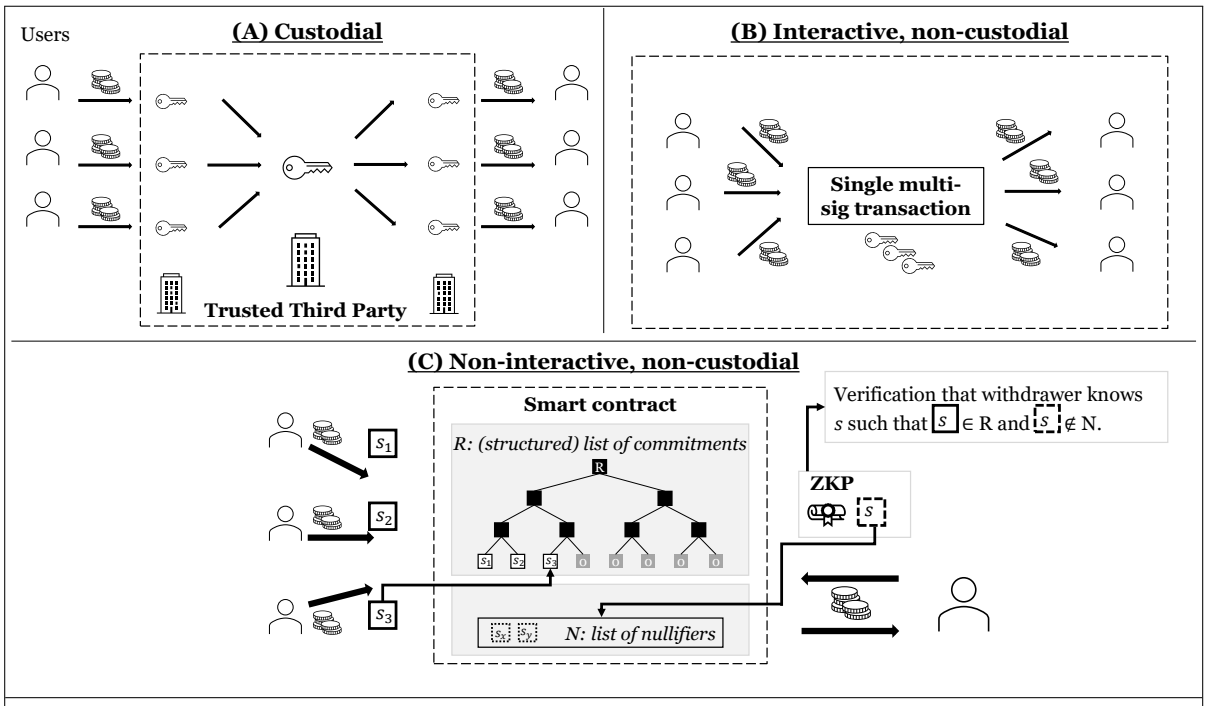

Figure 1. High-Level Perspective on Mixer Implementations

the different dimensions and characteristics later featured as part of the developed taxonomy. Third, establishing these allows for greater generalisability for consideration of projects beyond Ethereum, including other blockchains and decentralised applications that include payment functionalities and navigate privacy-related and regulatory requirements.

Like wallets, mixers can be implemented based on custodial (centralised) and non-custodial (decentralised) designs (Barbereau & Bodó, 2023; Béres et al., 2021). Figure 1 illustrates how mixers can be implemented from a high-level perspective. Custodial mixers (**A**) involve a TTP that offers a service such that entities can send their crypto-assets to an address associated with this TTP as a deposit. The TTP then applies a sequence of joining and splitting transactions with crypto-assets received from other addresses until a satisfactory degree of mixing is achieved. Finally, the TTP transfers the assets back to new addresses as specified by the depositors in bilateral communication (Nadler & Schär, 2023). Users can also engage in interactions to perform non-custodial mixing (**B**), for instance, by engaging in bilateral atomic swaps or CoinJoin-based approaches (Ghesmati et al., 2022). In the latter, each involved user specifies a sending and receiving address and communicates them multi-laterally to participants in the mixing process. The participants then create and jointly authorise (through every participant's digital signature with the cryptographic key corresponding to the sending address) a single transaction that consumes crypto-assets from all sending addresses and sends them back to the specified receiving addresses (Ruffing et al., 2014).

The trust required in a TTP that takes the custody of funds and learns about the connection between deposits and withdrawals (Nadler & Schär, 2023), as well as the complexity and data leakage risk associated with coordinating a substantial number of wallets in interactive approaches, make non-custodial, non-interactive mixers (**C**) a popular alternative. Here, users interact with a smart contract that provides the mixing service directly. This type typically uses non-interactive ZKPs to obfuscate transaction graphs. Deposited crypto-assets are mixed by avoiding the linkability between deposits and withdrawals. ZKPs represent a cryptographic primitive that allows a "prover" to convince a group of "verifiers" (and, therefore, in particular a blockchain-based smart contract) with a single message that a certain computation that the prover performed is correct, without requiring re-execution on the verifier side and in particular the disclosure of inputs, intermediary results, or outputs unless explicitly specified (Goldreich & Oren, 1994).

Non-custodial, non-interactive mixers typically store a structured, append-only list R ("commitments", often in the form of a Merkle tree) and a not necessarily structured, append-only list of nullifiers N. Elements of R and N are computed by users with two different hash functions and then sent to the mixing smart contract. Hash functions represent cryptographic one-way functions that map any input to a fixed-length output (e.g., 256 bits) and that do not allow deriving properties of the input from the output (Nadler & Schär, 2023). When a user wants to deposit crypto-assets, s/he creates a random secret (also called a *note*) s, hashes it with the hash function h_1, and sends crypto-assets and $h_1(s)$ to the mixing smart contract. The mixing smart contract increases its balance ('shielded pool of crypto-assets') by the sent amount and appends $h_1(s)$ to R. When a user seeks to withdraw crypto-assets that s/he previously deposited, s/he can read the current state of R and create a ZKP that they know some secret s such that $h_1(s) \in R$. Notably, neither $h_1(s)$ nor its position in R are revealed in this process. To prevent double-withdrawal despite unlinkability, the user also sends the hash of the secret with another hash function $h_2(s)$ and a ZKP that $h_2(s)$ was correctly computed from the same s underlying $h_1(s)$. The smart contract then appends $h_2(s)$ to the list of nullifiers N and verifies that the value $h_2(s)$ has not been added to it before. The ZKP avoids the creation of an observable link between depositing and withdrawal transaction, as for any observer, all commitments in R are equally likely to correspond to the withdrawal when ignoring potential time-related patterns (Wang et al., 2022).

Related work

Despite cryptocurrencies' decentralised nature and the use of public keys or hashes thereof as pseudonyms, it is possible to de-anonymise users leveraging the public visibility of transaction graphs. Meiklejohn et al. (2013)'s paper *A fistful of Bitcoins* presents one of the first methods to trace and identify crypto-asset users. It was eventually used by the FBI to build a case against the infamous Silk Road. In the following years, other works refined this approach or explored alternative ways to de-anonymise transactions on different blockchains; for instance, by incorporating Bitcoin's network topology (Biryukov et al., 2014) and network traffic (Biryukov & Tikhomirov, 2019) in their analyses. As Bitcoin's unspent transaction output (UTXO) model offers better privacy guarantees than account-based models underlying Ethereum and many other smart contract blockchains, Bitcoin has originally attracted more illicit activities (see also, Möser et al., 2013) and more research on de-anonymisation (Pocher et al., 2023). However, researchers also focused on transactions that went through more sophisticated Ethereum-based mixers like Tornado Cash (Béres et al., 2021) or cryptocurrencies with similar privacy features like Zcash (Biryukov & Tikhomirov, 2019).

Mixers can be seen as a technology whose use is fundamentally dual. On the one hand, these sophisticated anonymisation techniques on public blockchains can be used to address legitimate privacy concerns originating from the replicated transaction processing and storage on blockchains (Sedlmeir et al., 2022), to ensure fungibility (Biryukov et al., 2019), or to uphold (financial) privacy rights, for instance, in countries that suppress political opposition by foreclosing access to financial services (Campbell-Verduyn, 2018). On the other hand, because of their ability to provide a high degree of anonymity, mixers make it difficult to hold their users accountable for activities prior to use. As such, they enable malicious actors to escape regulatory efforts and bypass sanctions in the banking system. Academic research in digital investigations, criminology, and law observed that mixers are frequently used to do so (Trozze et al., 2023). This observation is supported by publications of (inter-)governmental law enforcement agencies (Europol, 2022; Interpol, 2020).

IS research has a standing history of analysing both the potential opportunities given by technological advances and critically reflecting on associated threats. Scholars of dedicated communities focus on topics of "digital forensics" (Nance et al., 2009) and the "dark side of information technology use" (D'Arcy et al., 2014; Tarafdar et al., 2013). In the latter, objects of study included but were not limited to social media (Chan et al., 2019), healthcare systems (Califf et al., 2020), and artificial intelligence (Mikalef et al., 2022). Although the IS literature on mixers is scarce, some works of its communities did consider crypto-assets more generally. For instance, Dhillon (2016) studied how novel technologies and, in particular, cryptocurrencies enable and facilitate money laundering.

Our research into how Ethereum-based crypto-asset mixers may be classified deviates from related work in three fundamental ways. First, several existing studies centre around mixer usage, with some of them focusing primarily on an assessment of the respective mixer's anonymity guarantees (Béres et al., 2021; Nadler & Schär, 2023). Second, previous publications contributed to our understanding of mixers by attempts to de-

anonymise transactions on crypto-asset mixers (Béres et al., 2021; See, 2023). Third, some studies analyse different Bitcoin-based mixing services with limited functionality (Ghesmati et al., 2022; Pakki et al., 2021) or focus on just a singular case, primarily Tornado Cash (Nadler & Schär, 2023). Previous studies thus do not provide an overview of the rich design options surrounding Ethereum-based mixers. We contribute to the latter discourse, as our taxonomy demonstrates that the design space for these mixers is much richer than previously studied, individual examples suggest. A comprehensive understanding of these design options may benefit future studies analysing the (criminal) use and compliance of Ethereum-based mixers.

Research Approach

To lift the fog over non-custodial, non-interactive, Ethereum-based crypto-asset mixers, we develop a taxonomy. Taxonomy building is suitable to classify objects with common characteristics, which, in turn, allows to organise subject areas and knowledge systematically (Bailey, 1994). Since taxonomies can serve as a foundation for future considerations in research and practice, our contribution addresses IS scholars, regulators, law enforcement, and digital forensics professionals. The development of multi-layer taxonomies is common for research on blockchain and crypto-assets given their still nascent status and rapid evolution (see, e.g., Hartwich et al., 2022; Ziegler & Welpe, 2022). Informed and inspired by these works, we follow the taxonomy-building process proposed by Nickerson et al. (2013) and later extended by Kundisch et al. (2021). While Nickerson et al. (2013) define seven steps, we also see relevance in the evaluation proposed by Kundisch et al. (2021). Figure 2 features our 10-step taxonomy development process. We opted for developing a multi-layer taxonomy. The introduction of layers unlocks the potential of including aggregation levels, which aid with structuring and classifying the dimensions and characteristics of the underlying object. We derived these layers inductively from our collected and analysed data and deductively from classifications in previous works.

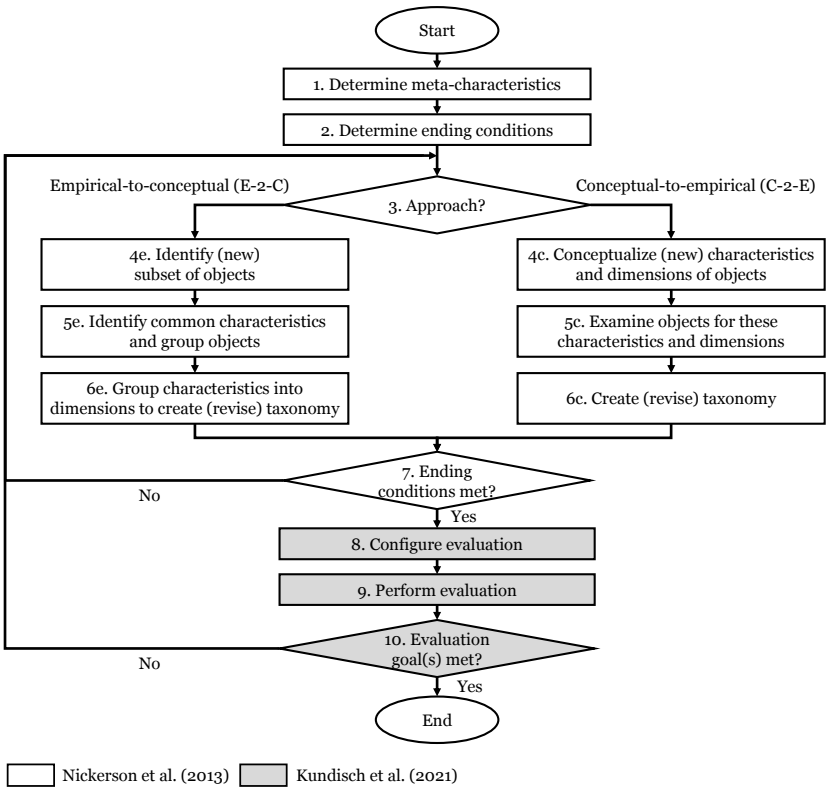

Figure 2. Taxonomy Development Steps

Definition of objectives (1–2): To support our target group of researchers, law enforcement, and regulators in classification and differentiation efforts, we define 'characteristics and functionalities of Ethereum-based mixers' as our meta-characteristic. We draw on the objective and subjective ending conditions proposed by Nickerson et al. (2013). For objective conditions, we summarise the ones suggested by Nickerson et al. (2013) and define that (1) a meaningful sample of mixers is analysed, (2) no layers (Ls), dimensions (Ds), or characteristics (Cs) were added, merged, or split, and (3) all dimensions provide "collectively exhaustive characteristics" (Nickerson et al., 2013, p.8). We also define subjective ending conditions, demanding that our taxonomy be concise, comprehensive, extendable, explanatory, and useful. As such, we check our taxonomy for all conditions defined by Nickerson et al. (2013) and Kundisch et al. (2021).

Design and development (3–7): Given that little information and data on Ethereum-based crypto-asset mixers is available, we follow a mixed approach consisting of four iterations (see Table 1). We started with a conceptual-to-empirical (C-2-E) approach for the first two iterations. We based this first iteration of taxonomy development on the most relevant literature identified via a search of the Scopus and IEEE Xplore databases, using "(mixer OR tumbler) AND (crypto OR Ethereum OR Bitcoin)" as a search string. Our search yielded 48 results on Scopus and 23 on IEEE Xplore. Out of these, we included only publications that are peer-reviewed and available in English. After title screening, abstract screening, and removing duplicates, five relevant publications remained. We hence decided to expand our literature base with two preprints by researchers that had previously published on mixers. For the first draft of our taxonomy, we built particularly on Feng et al. (2019) and Pakki et al. (2021) and supplemented valuable information from Amiram et al. (2022), Barbereau and Bodó (2023), Béres et al. (2021), Nadler and Schär (2023), Rathore et al. (2022), and van Wegberg et al. (2018). Informed by this literature base, we then considered four cases' project documentations (Table 2). Given that no guarantee was provided that project documentation in our specific cases matched the deployed source code of the respective mixer, we classified project documentation under the C-2-E cycle – even though it might also be considered empirical evidence in an empirical-to-conceptual (E-2-C) cycle. For the third iteration, we employed an E-2-C approach which we based on empirical data in the form of the smart contracts corresponding to the mixers in Table 2. Using smart contracts as an empirical data source is well-suited to meet our objectives. It complements project documentation on mixers, which we often found to be incomplete regarding key technical design choices or even differing from the coded reality. Since our goal is to analyse the characteristics and functionalities of crypto-asset mixers, it is paramount to screen the deployed smart contracts. We analysed the smart contracts written in Solidity and dissected their functionalities by mapping code snippets to the taxonomy's layers, dimensions, and characteristics. Table 3 features four examples of this coding scheme for fragments of mixers' Solidity code. For instance, in code fragment 1, the appearance of `_coinDenomination` in the constructor shows that the mixer defines the denomination at the moment of the contract's deployment. Thereby it corresponds to the layer of ***Mixer functionalities*** and the **Denomination** dimension, indicating that the characteristic is *Pre-defined*.

For the fourth iteration, we conducted expert interviews with seven researchers. Interviews are well suited for conducting exploratory and complex studies (Schultze & Avital, 2011). Each interview lasted between

It.	App.	Data Source	L/D/C	Status	Ending conditions
1	C-2-E	Literature review	6/15/47	Defined initial set of Ds and Cs.	Subj. and obj. ending conditions not met: Taxonomy is not explanatory without subject knowledge.
2	C-2-E	Project doc.	7/21/67	Added 6 Ds and 20 Cs. Clustered Ds into 7 Ls.	Obj. ending conditions not met: Cs were added and split, requiring another iteration of the process.
3	E-2-C	Smart contracts	5/20/47	Merged Ls to a total of 5, removed 1 D, merged or removed Cs.	Obj. ending conditions not met: Cs were deleted and added, requiring another iteration of the process.
4	E-2-C	Expert interviews	5/16/45	Removed or merged 4 Ds and 2 Cs. Adapted Ds and Cs.	All ending conditions met.

Table 1. Taxonomy Development Iterations

Case	Documentation	Ethereum smart contract address
Tornado Cash Nova	Tornado Cash Team (2021)	0xca0840578f57fe71599d29375e16783424023357
Cyclone	Cyclone Community (2023)	0xd619c8da0a58b63be7fa69b4cc648916fe95fa1b
Typhoon Cash	TyphoonCash (2021)	0x9cdb933edab885bb767658b9ed5c3800bc1d761b
Aztec	Andrews (2020)	0x737901bea3eeb88459df9ef1be8ff3ae1b42a2ba

Table 2. Cases Considered in Taxonomy Development

30 and 60 minutes (40 minutes average). We conducted the interviews between February and April 2023. To ensure rigour, we recorded and transcribed the interviews for further systematic analysis. We used expert sampling to select interviewees from academia based on recent publications. Although we contacted practitioners from three leading blockchain data analytics firms, they did not respond to our inquiries. Considering recent regulatory measures and increased scrutiny, we decided not to contact developers and users of non-custodial mixers, who may opt to remain anonymous. In each interview, we went through the layers in a semi-structured manner (Myers & Newman, 2007). We asked for the relevance of each dimension in the layer and whether the characteristics are appropriate and exhaustive. After the fourth interview, despite noticing saturation, we conducted three more interviews to ensure that no further layer, dimension, or characteristic had to be added or edited.

Evaluation (8–10): To evaluate our taxonomy, we also used these later interviews to gather feedback on its usefulness. To substantiate the objective validity of the taxonomy, we conducted an additional ex-post evaluation (Kundisch et al., 2021) in which we applied our taxonomy to the mixers RAILGUN and zkBob. This exercise verified that we met our ending conditions and concluded the taxonomy-building process.

In summary, and as indicated in Table 1, the taxonomy underwent various changes over four iterations. The first iteration established initial categories based on the literature review and introduced an "uncategorized" section for miscellaneous features. In the second iteration, the taxonomy shifted towards a more structured approach, while introducing more detailed dimensions. The third iteration refined this structure based on insights from mixers' smart contracts and rearranged certain dimensions. Finally, the fourth iteration further refined terminology based on feedback and integrated and removed various features due to its lack of justification from the mixer code-base.

Taxonomy of Ethereum-based Crypto-asset Mixers

Figure 3 displays the final taxonomy of non-custodial, non-interactive, Ethereum-based mixers. The layers of the taxonomy correspond to smart contracts' lifecycle (Sánchez-Gómez et al., 2020) and the methods they expose for the two core interactions users engage in. ***General features*** describes the choices developers make when designing the smart contract structure and characteristics of the environment in which the mixing service is deployed. The ***Terms of use*** layer describes the prerequisites that users need to fulfil to interact with the mixer (e.g., in terms of fees and compliance). They are reflected by the Solidity modifiers

№	Mixer / Address	Code snippet	Layer / Dimension	Char.
1	Cyclone	`uint256 public coinDenomination;` `constructor(…, uint256 _coinDenomination, …) {…}`	Mixer Func. / Denomination	Pre-defined
2	Cyclone	`function deposit(…) external payable nonReentrant {…` `  uint256 refund = msg.value - coinDenomination; …}`	Deposit Func. / Restitution	Refund
3	Cyclone	`address public govDAO;` `modifier onlyGovDAO { require(msg.sender == govDAO, "..."); _; }`	General Feat. / Governance	DAO

4	Typhoon Cash	`function _processDeposit() internal {`
		`  require(msg.value == 0, "...");`
		`  _safeErc20TransferFrom(…, denomination); }`
			Mixer Func. / Crypto-asset type	ERC20

Table 3. Exemplary Solidity Snippets and Mapped Characteristics

which restrict access to state updates (write methods) exposed by mixers. The remaining layers represent the user flow. We distinguish methods in the mixer smart contract that are uniquely connected to either ***Deposit functionalities*** or ***Withdrawal functionalities***. We associate dimensions that can equally be associated with both deposit and withdrawal functionalities to ***Mixer functionalities***. For instance, the crypto-asset type affects both deposit and withdrawal functionalities because it defines whether Ethereum's native token functionality or an external ERC20 interface will be used within the mixer to handle and transition the assets. In the following, we introduce each layer, dimension, and characteristic of the taxonomy and discuss their meaning, relevance, and implications. We also indicate whether characteristics within a dimension are mutually exclusive (ME), i.e., any mixer is characterised by a unique characteristic in this dimension (Y), or multiple characteristics may apply (N).

General features

Deployment layer refers to the environment in which the mixer is deployed and transactions are executed. *Layer 1 (L1)* refers to the native execution layer of Ethereum, which provides the infrastructure for smart contract-based transactions and the storage of corresponding states. Besides running the EVM, L1 controls the consensus protocol. Layer 2 (L2) refers to solutions that improve scalability and reduce transaction fees. L2 solutions are deployed on top of smart contracts and reduce the complexity of transaction processing on L1. A common family of L2 protocols is referred to as rollups. The two archetypes are optimistic and validity (often called zk-) rollups (Buterin, 2021). The former are based on the assumption that transactions are normally legitimate and perform smart-contract-based verifications of their validity only when they are challenged, while the latter use succinct cryptographic proof systems to ensure the correctness of the processing of individual or batches of transactions in a short ("succinct") proof. Both archetypes reduce the resource requirements in terms of processing, storage, and bandwidth of full nodes by offloading to another, separate execution layer, while the final execution results are still secured by the Ethereum L1. An example of optimistic rollups is Optimism, which allows the integration of arbitrary Solidity smart contracts. Validity rollups are still mostly in use for relatively simple transactions, such as transfers of ETH and ERC20 tokens. Recently, the first validity rollups were deployed on Ethereum test networks that offer support any Solidity smart contract; these projects are generally termed zk-EVM. Rollups are particularly attractive for mixers, as mixers involve relatively complex cryptographic operations. Mixers have been implemented both on top of optimistic rollups (Tornado Cash fork on Optimism, called Privacy Pools) and validity rollups (e.g., Aztec Connect). Both rollup constructions drastically reduce mixer usage fees. For instance, a deposit transaction that costs around $150 on Tornado Cash costs around $3 when deploying Tornado Cash on Optimism (own measurements, with gas prices from 2023-05-04).

Layer	Dimension	ME	Characteristics				
General features	Deployment layer	N	L1		L2 (Optimistic)		L2 (Validity)
	Governance	N	Single key	Multi-sig	DAO		No on-chain governance
	Code verification	Y	Audited		Unaudited		
	Contract architecture	Y	Single contract		Multi contract		
Terms of use	Address registration	Y	Mandatory		Optional		No registration
	Mixing usage fee	Y	Fixed		Variable		No fees
	Usage fee reception	N	(DAO) treasury	Developer(s)	Owner	Liquidity providers	N/A
Mixer functionalities	Denomination	Y	Pre-defined		Variable		
	Crypto-asset type	N	ETH		ERC20		ERC721
	Mixing process	Y	One in, one out		Split and withdraw		Split and use
	Compliance features	Y	Report generation		None		
Deposit functionalities	Restitution	Y	Reversion		Refund		No restitution
	Deposit censorship	N	Inclusion list		Exclusion list		No censorship
Withdrawal functionalities	Execution fee coverage	N	Relayer		Withdrawer		
	Withdrawal censorship	N	Inclusion list		Exclusion list		No censorship

Figure 3. Taxonomy of Ethereum-Based Mixers

Governance refers to how a mixer is controlled after deployment. These include control over mechanisms such as contract upgrades, terms of use, and emergency stops, among others. *Single key* refers to singular, centralised control of any type of smart contracts (Zhang et al., 2019) and, thus, also mixers. The key is held by a singular, distinguished entity that has 'owner' privileges. *Multi-sig* is a way to distribute the risks of a singular controlling entity. For instance, a set of keys and a corresponding majority rule can be pre-defined, or a key can be created in a distributed way such that every entity only receives one or multiple sub-keys. A certain threshold of keys is then needed to authorise updates to the smart contract (Komlo & Goldberg, 2021). However, it is hard to determine if these sub-keys were indeed distributed to more than one entity. Another commonly applied governance structure is *decentralised autonomous organisation (DAO)*-based and uses utility tokens that grant voting rights to their holders. Such governance has limitations as the tokens can be distributed unfairly or accumulated by few holders (Barbereau et al., 2023). Because there is only a limited set of rules to be governed in a mixer, there can also be *No on-chain governance framework*.

Code verification considers whether or not the underlying Solidity code-base of a mixer's smart contracts was *Audited* to identify and remove vulnerabilities. The higher the auditor's reputation, the more users may trust a mixer's code base and deployed contract(s). For instance, Cyclon was audited by two auditors (ChainShield & Slowmist, 2023). Some consider audits an essential process to trust in a smart contract's security (Groce et al., 2020). If no such audit takes place, mixers remain *Unaudited*.

Contract architecture refers to the smart contacts underlying the mixer. Mixers can be implemented within a *Single contract* that comprises all of the functionalities. Many others are composed of multiple smart contracts (*Multi contract*). Deploying multiple contracts is mainly motivated by modularity considerations: Each contract can have a single, thus simpler responsibility (such as complex cryptographic operations) and be governed (updated) individually. Some mixers also use individual smart contracts to manage different asset pools (e.g., one for each denomination, see also the **Mixing process**).

Terms of use

Address registration particularly affects some advanced mixers providing direct interfaces to smart contracts associated with decentralised finance (DeFi) protocols. From a technical perspective, creating a new Ethereum account to interact with mixers is not necessary. Hence, many mixers do not provide internal addresses (*No registration*). Some mixers – particularly those that support internal transfers (see also **Mixing process**) – try to popularise their own wallet or reach a better user experience, similar to other DeFi applications that are compatible with wallets like MetaMask. They do so by requiring the registration of a new account that is specifically associated with the mixer (*Mandatory*). The derived account is 'shielded' because it allows users to engage in transfers inside mixers without any linkability between different transactions associated with the same account. In particular, there is no public visibility of the involved account's address. In some cases, registration is supported but not mandatory (*Optional*).

Mixing usage fee corresponds to the transaction commission of a mixer for each deposit of funds. It is a programmable value on top of the transaction fee (which covers the execution of smart contract's functions on the **Deployment layer**). The entity deploying a mixer can define the recipient of the mixing usage fee. Although the recipient could an externally owned account, the mixing fee recipient in all mixers that we investigated and that implement mixing fees is the mixing smart contract itself or another smart contract (see also **Usage fee reception**.. This address makes a profit whenever users mix their crypto assets. The rules to determine the fee are also specified in the mixer smart contract. The fee is usually either flat (*Fixed*) or a small fraction of each deposit transaction (*Variable*), e.g., 0.25 % in RAILGUN. Many mixing contracts feature *No fees*: Owing to the inherent transparency of the blockchain and the publishing of smart contract code to establish trust in mixer usage, any third party can copy the mixer smart contract code, reduce or remove the usage fees, change the recipient, deploy the mixer with these modifications, and overcoming the chicken-and-egg-problem, i.e., trying to attract enough users through the cheaper usage to reach a sufficient level of anonymity guarantees.

Usage fee reception refers to the entity the usage fee is delivered to, e.g., the *DAO treasury* which governs the mixer. It can also be distributed to *Developer(s)* who designed, developed, deployed, and maintain the mixer, to the *Owner* of the mixer smart contract, to *Liquidity providers* that have deposited crypto-assets in

the mixer, or some combination thereof. It would make more sense to reward deposits based on the duration until subsequent withdrawal ("anonymity mining"), but we found no mixer that implemented this approach. *N/A* corresponds to the case where the mixer does not take usage fees.

Mixer functionalities

Denomination refers to which portions of crypto-assets can be deposited and withdrawn by users. Some mixers have a *Pre-defined* set of denominations such that only previously specified denominations of funds can be deposited. For instance, Tornado Cash Nova supports units of 0.1 ETH, 0.3 ETH, 0.5 ETH, and 1 ETH. The fixed denominations are defined at the deployment of the smart contracts. The corresponding deposits and withdrawals can be managed by a smart contract for each denomination or by a single smart contract that comprises all denominations (see **Contract architecture**). For pre-defined denominations, there is a fundamental trade-off between the number of transactions users need to conduct (and pay) to decompose their custom amounts to match the desired amount, and the level of anonymity guarantees: A small variety of denominations can make it inconvenient for users to mix custom amounts; but on the other hand, bundling all deposits into a single pool with one denomination increases the anonymity set. Due to this trade-off, some implementations of mixers are not restricted to certain denominations. Instead, users can deposit and withdraw *Variable* amounts of funds. However, this approach only provides sufficient degrees of anonymisation if crypto-assets can be split inside the pool (see **Mixing process**).

Crypto-asset type refers to the types of assets that can be mixed. The mixers we investigated all target both the native token of the Ethereum blockchain, *ETH*, and the arguably most popular fungible token standard of the Ethereum blockchain, *ERC20*, albeit in different smart contracts. Technically, there is almost no difference between mixing ETH and mixing ERC20 tokens. Implementations compatible with ERC20 involve an additional interaction with the "approve" method of ERC20 smart contracts. As there is no basic depositing method associated with ERC20s, depositors have to authorise that the address of the mixer may spend a certain amount of tokens ("allowance"). Because of the common ERC20 interface, mixers can easily integrate several ERC20 tokens. For instance, Cyclone supports two ERC20 tokens, USDT and TORN. The compatible tokens should be popular enough to have a sufficient number of deposits, otherwise the anonymity set is too small for effective mixing. Because non-fungible tokens (NFTs) are unique (Hartwich et al., 2022), simple mixing is not useful to improve privacy. However, if the mixer supports internal transfers (see **Mixing process**), supporting NFTs is a useful feature. For instance, Nightfall supports the common non-fungible token standard *ERC721*.

Mixing process relates to the operations supported by the mixer. The archetypal one is denoted *One in, one out* because tokens (typically with a fixed denomination) are deposited, pooled, and withdrawn atomically. Other mixers allow for more advanced actions with the deposited funds. For instance, deposited fungible tokens can be split within the system and then sequentially withdrawn (*Split and withdraw*). Alternatively, the deposited amount can be split (again only for fungible tokens) within the system and used for shielded internal transfers and interaction with, for example, DeFi components (*Split and use*). In this case, ownership relations can be changed without the crypto-assets leaving the mixer, such that the mixer provides similar features as a private cryptocurrency like Zcash (for fungible tokens).

Compliance features refer to mixers that enable users to disclose the links between their deposits and withdrawals verifiably. Most commonly, this is achieved through *Report generation*. While their transaction graphs remain obfuscated for observers on the public blockchain, users can provide detailed information to third parties such as auditors and compliance officers. Doing so in confidential communications permits the verification of the provenance of their funds, even though they were sent through a mixer (Nadler & Schär, 2023). While any ZKP-based mixer could offer this functionality (Garman et al., 2017), not all of them provide this functionality to users (*None*).

Deposit functionalities

Restitution applies to mixers with a fixed denomination. When the amount to be deposited is not equal to that denomination, the mixer smart contract either triggers a *Reversion* (the deposit transaction is not

executed and the user wasted transaction fees) or a *Refund* (the difference between the denomination and the deposit is sent back to the depositor). There is *No restitution* in mixers with a variable denomination.

Deposit censorship refers to functionalities allowing the mixer smart contract to incorporate certain aspects of regulatory compliance. One way to do so is to specify an *Inclusion list* or *Exclusion list* that respectively allows or blocks the interaction of specific depositor addresses with a mixer (Burleson et al., 2022). For instance, using an Oracle, mixers can implement exclusion lists to comply with sanction lists, such as the OFAC's list of sanctioned Ethereum addresses. However, sanctions list updates, e.g., after a major attack on a DeFi project, typically happen too slowly to prevent corresponding deposits to the mixer (Burleson et al., 2022). Inclusion lists, on the other hand, can be useful to restrict access to the mixer to addresses that went through a know-your-customer (KYC) process at certain centralised exchanges. There can also be *No censorship*.

Withdrawal functionalities

Execution fee coverage refers to the mechanism by which transaction fees for the smart contract execution layer are paid upon withdrawals. To withdraw deposited funds, users employ cryptographic keys corresponding to the destination address to sign the transaction for invoking the `withdraw` function of the mixer. To make the mixing useful for anonymisation (and as opposed to the case with deposits), the user will generate a new address with no initial funds. The withdrawal transaction requires a fee and potentially a tip to the block producer to incentivise the inclusion of the transaction in a block. Yet, sending funds to the new address from a user's previous address would spoil the obfuscation of the transaction graph. *Relayers* provide a solution to this issue. Users create a transaction that (1) selects a relayer to cover the transaction fee for the `withdraw` function and to withdraw the funds from the mixer, and that forwards the withdrawn amount less the gas and relayer usage fees to the destination address as specified by the user. As such, the user does not need to trust the relayer in terms of confidentiality or crypto-asset custody. Although funding transactions from user-controlled addresses reduces the degree of anonymisation provided by the mixer, many mixers also support fee payment directly by the *Withdrawer*.

Withdrawal censorship refers to the restriction of access to withdrawals. It has the same purpose as for deposit censorship, and *Inclusion lists* play a similar role (e.g., only withdrawals to KYC-ed addresses are allowed). We did not find exclusion lists for withdrawal addresses – which is plausible because they are typically generated only directly before withdrawal. Yet, *Exclusion lists* can be very useful for compliance if they target the depositor address (in *one-in, one-out* mixers). As mixers only provide a high degree of anonymisation when funds are deposited long enough, illicit actors that deposit their funds (e.g., after exploiting a vulnerability in a DeFi protocol) need to be worried that their depositing addresses will be added to the sanction lists that the mixer refers to, which means that they would not be able to access their funds anymore (Burleson et al., 2022). This places a strong incentive for illicit actors not to use such a mixer and, in turn, makes the mixer more attractive for licit users with legitimate privacy needs. As effective mixing relies on the withdrawal transaction not explicitly referring to the corresponding depositing address, the proof of inclusion or exclusion in corresponding lists must be provided by the withdrawer in the form of a ZKP. There can also be *No censorship*.

Evaluation of the Taxonomy

To demonstrate the usefulness and completeness of our taxonomy, we evaluated it following the guidelines by Kundisch et al. (2021). We employed three distinct methods. First, once we observed theoretical saturation in the interviews (after the fourth interview) during our taxonomy-building process, we started to ask the interviewees for their opinions on the usefulness of the taxonomy. The first three interviewees uniformly described the taxonomy as both "functional" for practitioners and "complete". The fourth interviewee stated that there are "no logical flaws", and the fifth and seventh interviewees described the taxonomy as "thorough". Second, we evaluated the taxonomy by applying it to two distinct mixers not used in the taxonomy design phase, RAILGUN and zkBob. Our taxonomy could capture all relevant characteristics of these two mixers. During the evaluation phase, we made a negative observation related to the characterisation of mixers' user experience. Beyond address registration, although user experience aspects may considerably

Layer	Dimension	ME	Characteristics				
General features	**Deployment layer**	N	L1		L2 (Optimistic)		L2 (Validity)
	Governance	N	Single key	Multi-sig	DAO		No on-chain governance
	Code verification	Y	Audited		Unaudited		
	Contract architecture	Y	Single contract		Multi contract		
Terms of use	**Address registration**	Y	Mandatory		Optional		No registration
	Mixing usage fee	Y	Fixed		Variable		No fees
	Usage fee reception	N	(DAO) treasury	Developer(s)	Owner	Liquidity providers	N/A
Mixer functionalities	**Denomination**	Y	Pre-defined		Variable		
	Crypto-asset type	N	ETH		ERC20		ERC721
	Mixing process	Y	One in, one out		Split and withdraw		Split and use
	Compliance features	Y	Report generation		None		
Deposit functionalities	**Restitution**	Y	Reversion		Refund		No restitution
	Deposit censorship	N	Inclusion list		Exclusion list		No censorship
Withdrawal functionalities	**Execution fee coverage**	N	Relayer		Withdrawer		
	Withdrawal censorship	N	Inclusion list		Exclusion list		No censorship

Figure 4. Taxonomy Applied to RAILGUN

influence the popularity of this mixer, our taxonomy does not account for pronounced differences in user interfaces. Yet, it is difficult to derive discrete and objective characteristics for such a dimension without a large-scale user study – one that is beyond the scope of our research and may also be difficult to realise owing to the legal issues surrounding mixer usage. Another negative observation was related to the composability feature of RAILGUN. During the application of the taxonomy to RAILGUN (see Figure 4), we noted that it enables direct interactions with other smart contracts associated to, for instance, decentralised exchanges, trading and yield farming, prompting users to keep their assets in the mixer for extended periods. Our taxonomy does not include dimensions that can reflect mixers' capabilities of directly interacting with other smart contracts. However, such features do not describe the intrinsic functionality of mixers. Like user experience aspects, we consider it beyond the scope of our taxonomy. Third, we sent the final taxonomy to all of our interviewees via email, asking for additional written feedback on the finished taxonomy. The first three interviewees reaffirmed its value in written form, describing it as both "useful" and "complete". The fifth and seventh interviewees considered the taxonomy "relevant". Interviewee 6, in consideration of the above-mentioned events related to Tornado Cash, additionally described it as "timely". This evaluation suggests our taxonomy's fitness to cover the key characteristics of Ethereum-based mixers.

Discussion

Our interviews revealed that besides the discrete mixer characteristics we identified when designing the taxonomy, there is a need to consider anonymity guarantees as a pivotal aspect in users' choice of mixers. Because anonymity is multi-faceted and represents a continuum in a highly context- (user base), time- (depositing duration), and user- (expertise of use) specific environment, it is difficult to cover it in a "concise" (Nickerson et al., 2013) way and meet the subjective ending conditions of taxonomy building. To account for this, we discuss this aspect in the following. as well as a critical factor for users' choices (Kruisbergen et al., 2019). Anonymity is a central consideration in both the legitimate (privacy needs) and the illegitimate (e.g., money laundering) use of mixers. Here, anonymity represents an abstraction that considers the anonymity set, i.e., the number of depositors (or depositing transaction) a withdrawal could be associated with, as well as potential scenarios that either improve or degrade the feasibility of 'tracing' funds. Tracing refers to the probabilistic linking of the depositor address and the withdrawal address (or multiple withdrawal addresses for split-and-use systems). The anonymity set of a mixer has a profound impact on the traceability of individual addresses (Wang et al., 2022). For mixers with a fixed denomination, anonymity primarily depends on the number of other users who deposited to the mixer without withdrawing their funds. It is further impacted by the delay between deposit and withdrawal. For mixers without a fixed denomination, the extent of anonymity guarantees depends on the extent to which users split funds on withdrawal and whether it

is impossible to restore the deposited amount by selective addition of the withdrawn amounts. Users with strong privacy needs, such as crime offenders, will arguably favour those mixers with a consistent activity of depositors and a large number of depositors who have not conducted a withdrawal, as this directly impacts the degree of anonymity obtained through mixing. They will also prefer mixers that involve relayers to withdraw mixed assets on a new wallet address with no ETH.

Contributions to theory. Our taxonomy provides an initial but comprehensive classification of non-custodial, non-interactive, Ethereum-based mixers. As such, our research contributes to theory and has implications for prac-tice. It further ex-post demonstrates that the properties of being non-custodial and non-interactive make the design space of Ethereum-based mixers much richer than that of Bitcoin-based tumblers. The design choices we identified (e.g., deployment on L1 or L2 relaying, fee mechanisms, and the implementation of block-lists for depositing and withdrawal) have a profound impact on the functionality offered. Some of these choices may contribute to the extent of their use in criminal activities. As such, our taxonomy stands aside other classifications on crypto-assets and the blockchain ecosystem published in the IS discipline (e.g. Hartwich et al., 2022; Ziegler & Welpe, 2022). These taxonomies, including our own, contribute an understanding of 'fuzzy' and dynamic digital phenomena that – given their greater adoption and decentralised nature – one may characterise as *pervasive* in the sense that mixers erect "novel social formations and behaviours not seen before" (Grover & Lyytinen, 2022, p. 4). Inherently, the value of descriptive classifications lies in their capacity to *explain* (Gregor, 2006). Thus, our contribution to theory is primarily explanatory. Our taxonomy helps researchers by elucidating the task environment or, design choices, of Ethereum-based mixers. It provides an understanding of choices certain user groups, such as crime offenders, can make in the digital age (Kruisbergen et al., 2019). Hence, it represents a contribution to IS research in digital forensics (Nance et al., 2009) and IT artefacts with an inherent dark side (D'Arcy et al., 2014; Tarafdar et al., 2013).

Implications for practice. Together with the Basel Institute on Governance, Europol (2022) advocates for a greater "understanding of crypto-assets and services" as "vital to tackle organised crime and money laundering". Similarly, Interpol (2020) advocates for additional "research to determine the proportions in which criminals prefer [...] mixers over privacy coins". Our work contributes to those ends as it provides an understanding of the most prominent tools used in the money laundering process: mixers. Indeed, such classifications represent valuable analytical instruments for practitioners (Sarre et al., 2018). Consequently, our taxonomy has practical implications for law enforcement and investigations. First, it allows to better understand and identify the various types of Ethereum-based mixers used – both in connection to criminals who launder illicit funds and individuals with legitimate goals – and their implemented as well as conceivable functionalities. Second, by categorising mixers based on their features and characteristics, law enforcement agencies can develop targeted strategies for identifying and tracking illicit transactions. These assist in developing digital investigation instruments and techniques for identifying common patterns or signatures associated with specific types of mixers, which could help them trace funds back to their source or develop an understanding of motives of use (Fröwis et al., 2020). Third, our taxonomy can inform the development of regulations and policies to control mixer usage. By understanding the different ways in which mixers are used, regulators can design more effective measures to prevent and deter money laundering and other criminal activities in the crypto-asset space (see also Barbereau & Bodó, 2023). Lastly, consider that many projects in Web3 aim to protect sensitive information (Lacity et al., 2023), with a significant portion of them using public blockchains, their native crypto-asset, and smart contracts as foundation. Therefore, gaining insights into how certain design choices enable improvements in the protection of sensitive information while also addressing compliance requirements for the case of native cryptocurrency assets could prove highly valuable (Barbereau et al., 2022; Sedlmeir et al., 2022).

Conclusion

Following the taxonomy development steps of Kundisch et al. (2021) and Nickerson et al. (2013), we developed a taxonomy of Ethereum-based, non-custodial, non-interactive crypto-asset mixers. Mixers represent a study object suitable for analysis in research on digital forensics and the dark side of IS (D'Arcy et al., 2014; Tarafdar et al., 2013). We considered both qualitative and quantitative data for the taxonomy development and evaluation over four cycles. As such, the taxonomy we present in this paper is based on (1) a

review of existing peer-reviewed research on mixers, (2) project documentation for four mixers in the form of whitepapers and blog posts, (3) an analysis of the corresponding smart contracts, and (4) expert interviews. We evaluate our taxonomy by applying it to two cases, RAILGUN and zkBob. Our taxonomy provides a comprehensive overview of the novel and previously unstructured research field of Ethereum-based mixers.

Limitations. The limitations of our taxonomy are three-fold. First, we considered a sample of four Ethereum-based mixers to develop our taxonomy and a sample of two for evaluation. These samples offered us rich insights and facilitated the ex-post evaluation of characteristics, dimensions, and layers. Nevertheless, future research could use a larger sample to further evaluate, extend, and potentially modify our taxonomy. Second, while we conducted interviews with experts from different fields to iteratively revise and evaluate our taxonomy, all of them were researchers. Future works could include experts from practice, such as law enforcement agents or software engineers. Third, we are limited by the validity of data used in the initial two (C-2-E) iterations, namely by the reliability of mixers' project documentation. This is a common limitation in taxonomy building (Nickerson et al., 2013). We addressed this issue by critically reflecting on the role and capabilities of the involved technical design choices (e.g., ZKPs) and adding analyses of smart contracts in our data collection.

Future research directions. Considering these limitations, in line with recommendations by Kundisch et al. (2021) and recent related taxonomies (e.g., Hartwich et al., 2022), our taxonomy is intentionally extendable. Taxonomies are by design intended to be encompassing though at the same time allowing for a degree of expansion "when new objects appear" (Nickerson et al., 2013). Ethereum-based mixers are a young and evolving field in research and of relevance to practice. Thus, our taxonomy is designed to flexibly adapt to new developments on all three levels of our taxonomy – layer, dimensions, and characteristics (Nickerson et al., 2013). The close alignment of our taxonomy with the lifecycles of smart contract development and mixing transactions may further increase its flexibility to represent technological advances. Future research could focus on the evaluation and modification of our taxonomy as well as the development of ideal types (i.e., archetypes) (Bailey, 1994) by analysing mixers on smart contract blockchains beyond Ethereum as well as dedicated privacy cryptocurrencies such as Zcash. We also believe that linking research on mixers with research on privacy-oriented and compliant payment systems (such as central bank digital currencies) may be a fruitful endeavour. Research in this area has particularly worked on harmonising privacy requirements with regulatory compliance through cryptographic designs that ensure accountability (e.g., Garman et al., 2017) or that enforce certain transaction limits (Groß et al., 2021). Based on such an extended data set, a joint taxonomy for mixers and privacy pools as well as legal studies on the compatibility of these solutions with KYC and AML requirements could be conducted to continue the corresponding research stream (Barbereau & Bodó, 2023; Barbereau et al., 2022; Trozze et al., 2023).

Despite the limitations discussed above, our analysis and results contribute to both theory and practice. We contribute to the IS discipline's understanding of mixers as pervasive digital phenomena by providing insights into why users with legitimate privacy needs and malicious actors alike may use Ethereum-based mixers and which functionalities they may experience. We provided avenues for future research that may enable a better understanding of users interacting with mixers as well as fill knowledge gaps regarding the different ways to integrate functionalities required for compliance. Our detailed description of the taxonomy's layers and dimensions as well as the corresponding characteristics can support law enforcement in developing digital investigation instruments and targeted strategies to combat illicit finance or proactively provide mixers with compliance features as a legitimate alternative.

Acknowledgements

This research was funded by the Luxembourg National Research Fund (FNR) and PayPal PEARL (grant reference 13342933) as well as by the FNR in the FiReSPArX (grant reference 14783405) and PABLO (grant reference 16326754) projects. For the purpose of open access, the authors have applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any author accepted manuscript version arising from this submission.

References

Amiram, D., Jørgensen, B. N., and Rabetti, D. (2022). "Coins for bombs: The predictive ability of on-chain transfers for terrorist attacks," *Journal of Accounting Research* (60:2), pp. 427–466. https://doi.org/10.1111/1475-679X.12430.

Andrews, J. (2020). *Aztec: zkRollup Layer 2 + Privacy*.

Bailey, K. D. (1994). *Typologies and taxonomies: An introduction to classification techniques,* SAGE.

Barbereau, T. and Bodó, B. (2023). "Beyond financial regulation of crypto-asset wallet software: In search of secondary liability," *Computer Security & Law Review* (49). https://doi.org/10.1016/j.clsr.2023.105829.

Barbereau, T., Sedlmeir, J., Smethurst, R., Fridgen, G., and Rieger, A. (2022). "Tokenization and regulatory compliance for art and collectibles markets," in *Blockchains and the Token Economy,* pp. 213–236. https://doi.org/10.1007/978-3-030-95108-5_8.

Barbereau, T., Smethurst, R., Papageorgiou, O., Sedlmeir, J., and Fridgen, G. (2023). "Decentralised finance's timocratic governance: The distribution and exercise of tokenised voting rights," *Technology in Society* (73). https://doi.org/10.1016/j.techsoc.2023.102251.

Béres, F., Seres, I. A., Benczúr, A. A., and Quintyne-Collins, M. (2021). "Blockchain is watching you: Profiling and de-anonymizing Ethereum users," in *IEEE International Conference on Decentralized Applications and Infrastructures,* pp. 69–78. https://doi.org/10.1109/DAPPS52256.2021.00013.

Biryukov, A., Feher, D., and Vitto, G. (2019). "Privacy aspects and subliminal channels in zcash," in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security,* pp. 1813–1830. https://doi.org/10.1145/3319535.3345663.

Biryukov, A., Khovratovich, D., and Pustogarov, I. (2014). "Deanonymisation of clients in Bitcoin P2P network," in *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security,* pp. 15–29. https://doi.org/10.1145/2660267.2660379.

Biryukov, A. and Tikhomirov, S. (2019). "Transaction clustering using network traffic analysis for Bitcoin and derived blockchains," in *IEEE Conference on Computer Communications Workshops,* pp. 204–209. https://doi.org/10.1109/INFCOMW.2019.8845213.

Burleson, J., Korver, M., and Boneh, D. (2022). *Privacy-protecting regulatory solutions using zero-knowledge proofs*. a16z crypto.

Buterin, V. (2021). *An incomplete guide to rollups*.

Califf, C. B., Sarker, S., and Sarker, S. (2020). "The bright and dark sides of technostress: A mixed-methods study involving healthcare IT," *MIS Quarterly* (44:2), pp. 809–856. https://doi.org/10.25300/MISQ/2020/14818.

Campbell-Verduyn, M. (2018). "Bitcoin, crypto-coins, and global anti-money laundering governance," *Crime, Law, and Social Change* (69:2), pp. 283–305. https://doi.org/10.1007/s10611-017-9756-5.

Chainalysis (2023). *The 2023 crypto crime report*.

ChainShield and Slowmist (2023). *Audit report by ChainShield and Slowmist*. Cyclone.xyz.

Chan, T. K., Cheung, C. M., and Wong, R. Y. (2019). "Cyberbullying on social networking sites: The crime opportunity and affordance perspectives," *Journal of Management Information Systems* (36:2), pp. 574–609. https://doi.org/10.1080/07421222.2019.1599500.

Cirkovic, M., Cachin, C., and Le, D. V. (2022). "Cryptographic primitives for on-chain tumbler designs,"

Cyclone Community (2023). *Cyclone Protocol: Development*.

D'Arcy, J., Gupta, A., Tarafdar, M., and Turel, O. (2014). "Reflecting on the "dark side" of information technology use," *Communications of the Association for Information Systems* (35:1), pp. 109–118. https://doi.org/10.17705/1CAIS.03505.

Dhillon, G. (2016). "Money laundering and technology enabled crime: A cultural analysis," in *Proceedings of the 22nd Americas Conference on Information Systems,* AIS.

Europol (2022). *Seizing the opportunity: 5 recommendations for crypto assets-related crime and money laundering*.

Feng, Q., He, D., Zeadally, S., Khan, M. K., and Kumar, N. (2019). "A survey on privacy protection in blockchain system," *Journal of Network and Computer Applications* (126), pp. 45–58. https://doi.org/10.1016/j.jnca.2018.10.020.

FIOD (2022). "Arrest of suspected developer of Tornado Cash," *Nieuws*.

Fröwis, M., Gottschalk, T., Haslhofer, B., Rückert, C., and Pesch, P. (2020). "Safeguarding the evidential value of forensic cryptocurrency investigations," *Forensic Science International: Digital Investigation* (33). https://doi.org/10.1016/j.fsidi.2019.200902.

Garman, C., Green, M., and Miers, I. (2017). "Accountable privacy for decentralized anonymous payments," in *Financial Cryptography and Data Security: 20th International Conference,* Springer, pp. 81–98. https://doi.org/10.1007/978-3-662-54970-4_5.

Ghesmati, S., Fdhila, W., and Weippl, E. (2022). "SoK: How private is Bitcoin? Classification and evaluation of Bitcoin privacy techniques," in *Proceedings of the 17th International Conference on Availability, Reliability and Security,* ACM. https://doi.org/10.1145/3538969.3538971.

Goldreich, O. and Oren, Y. (1994). "Definitions and properties of zero-knowledge proof systems," *Journal of Cryptology* (7:1). https://doi.org/10.1007/BF00195207.

Gregor, S. (2006). "The nature of theory in information systems," *MIS Quarterly*, pp. 611–642. https://doi.org/10.2307/25148742.

Groce, A., Feist, J., Grieco, G., and Colburn, M. (2020). "What are the actual flaws in important smart contracts (and how can we find them)?," in *Financial Cryptography and Data Security: 24th International Conference,* Springer, pp. 634–653. https://doi.org/10.1007/978-3-030-51280-4_34.

Groß, J., Sedlmeir, J., Babel, M., Bechtel, A., and Schellinger, B. (2021). *Designing a central bank digital currency with support for cash-like privacy*. https://doi.org/10.2139/ssrn.3891121.

Grover, V. and Lyytinen, K. (2022). "The pursuit of innovative theory in the digital age," *Journal of Information Technology*, pp. 45–59. https://doi.org/10.1177/02683962221077112.

Hartwich, E., Ollig, P., Fridgen, G., and Rieger, A. (2022). "Probably something: A multi-layer taxonomy of non-fungible tokens," *Internet Research*. https://doi.org/10.1108/INTR-08-2022-0666.

Interpol (2020). "Combatting cyber-enabled financial crimes in the era of virtual asset and darknet service providers," *Global Complex for Innovation*.

Komlo, C. and Goldberg, I. (2021). "FROST: Flexible round-optimized Schnorr threshold signatures," in *Selected Areas in Cryptography,* pp. 34–65. https://doi.org/10.1007/978-3-030-81652-0_2.

Kruisbergen, E. W., Leukfeldt, E. R., Kleemans, E. R., and Roks, R. A. (2019). "Money talks money laundering choices of organized crime offenders in a digital age," *Journal of Crime and Justice* (42:5), pp. 569–581. https://doi.org/10.1080/0735648X.2019.1692420.

Kundisch, D., Muntermann, J., Oberländer, A. M., Rau, D., Röglinger, M., Schoormann, T., and Szopinski, D. (2021). "An update for taxonomy designers: Methodological guidance from information systems research," *Business & Information Systems Engineering* (64), pp. 421–439. https://doi.org/10.1007/s12599-021-00723-x.

Lacity, M., Carmel, E., Young, A. G., and Roth, T. (2023). "The quiet corner of Web3 that means business," *MIT Sloan Management Review* (64:3).

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., and Savage, S. (2013). "A fistful of Bitcoins: Characterizing payments among men with no names," in *Proceedings of the Internet Measurement Conference,* ACM, pp. 127–140. https://doi.org/10.1145/2504730.2504747.

Mikalef, P., Conboy, K., Lundström, J. E., and Popovič, A. (2022). "Thinking responsibly about responsible AI and 'the dark side'of AI," *European Journal of Information Systems* (31:3), pp. 257–268. https://doi.org/10.1080/0960085X.2022.2026621.

Möser, M., Böhme, R., and Breuker, D. (2013). "An inquiry into money laundering tools in the Bitcoin ecosystem," in *APWG eCrime Researchers Summit,* IEEE. https://doi.org/10.1109/eCRS.2013.6805780.

Myers, M. D. and Newman, M. (2007). "The qualitative interview in IS research: Examining the craft," *Information and Organization* (17:1), pp. 2–26. https://doi.org/10.1016/j.infoandorg.2006.11.001.

Nadler, M. and Schär, F. (2023). "Tornado Cash and blockchain privacy: A primer for economists and policymakers," *Federal Reserve Bank of St. Louis Review* (105:2), pp. 122–136. https://doi.org/10.20955/r.105.122-36.

Nance, K., Hay, B., and Bishop, M. (2009). "Digital forensics: Defining a research agenda," in *Proceedings of the 42nd Hawaii International Conference on System Sciences,* ScholarSpace. https://doi.org/10.1109/HICSS.2009.160.

Nickerson, R. C., Varshney, U., and Muntermann, J. (2013). "A method for taxonomy development and its application in information systems," *European Journal of Information Systems* (22:3), pp. 336–359. https://doi.org/10.1057/ejis.2012.26.

Pakki, J., Shoshitaishvili, Y., Wang, R., Bao, T., and Doupé, A. (2021). "Everything you ever wanted to know about Bitcoin mixers (but were afraid to ask)," in *Financial Cryptography and Data Security: 25th International Conference,* Springer, pp. 117–146. https://doi.org/10.1007/978-3-662-64322-8_6.

Pocher, N., Zichichi, M., Merizzi, F., Shafiq, M. Z., and Ferretti, S. (2023). "Detecting anomalous cryptocurrency transactions: An AML/CFT application of machine learning-based forensics," *Electronic Markets* (33:1). https://doi.org/10.1007/s12525-023-00654-3.

Rathore, M. M., Chaurasia, S., and Shukla, D. (2022). "Mixers detection in Bitcoin network: A step towards detecting money laundering in crypto-currencies," in *IEEE International Conference on Big Data,* pp. 5775–5782. https://doi.org/10.1109/BigData55660.2022.10020982.

Ruffing, T., Moreno-Sanchez, P., and Kate, A. (2014). "Coinshuffle: Practical decentralized coin mixing for Bitcoin," in *Proceedings of the 19th European Symposium on Research in Computer Security,* Springer, pp. 345–364. https://doi.org/10.1007/978-3-319-11212-1_20.

Sánchez-Gómez, N., Torres-Valderrama, J., García-García, J. A., Gutiérrez, J. J., and Escalona, M. (2020). "Model-based software design and testing in blockchain smart contracts: A systematic literature review," *IEEE Access* (8), pp. 164556–164569. https://doi.org/10.1109/ACCESS.2020.3021502.

Sarre, R., Lau, L. Y.-C., and Chang, L. Y. (2018). "Responding to cybercrime: Current trends," *Police Practice and Research* (19:6), pp. 515–518. https://doi.org/10.1080/15614263.2018.1507888.

Schultze, U. and Avital, M. (2011). "Designing interviews to generate rich data for information systems research," *Information and Organization* (21:1). https://doi.org/10.1016/j.infoandorg.2010.11.001.

Sedlmeir, J., Lautenschlager, J., Fridgen, G., and Urbach, N. (2022). "The transparency challenge of blockchain in organizations," *Electronic Markets* (32), pp. 1779–1794. https://doi.org/10.1007/s12525-022-00536-0.

See, K. (2023). "The Satoshi laundromat: A review on the money laundering open door of Bitcoin mixers," *Journal of Financial Crime*. https://doi.org/10.1108/JFC-11-2022-0269.

Shen, M. (2022). "Crypto mixer Tornado Cash says sanctions can't apply to smart contracts," *Bloomberg*.

Tarafdar, M., Gupta, A., and Turel, O. (2013). "The dark side of information technology use," *Information Systems Journal* (23:3), pp. 269–275. https://doi.org/10.1111/isj.12015.

Tornado Cash Team (2021). *Tornado Cash introduces arbitrary amounts & shielded transfers*.

Trozze, A., Davies, T., and Kleinberg, B. (2023). "Of degens and defrauders: Using open-source investigative tools to investigate decentralized finance frauds and money laundering," *Forensic Science International: Digital Investigation* (46). https://doi.org/10.1016/j.fsidi.2023.301575.

TyphoonCash (2021). *Introducing Typhoon Cash – A new protocol for yield-capable private transactions*.

U.S. Department of the Treasury (2022). *U.S. Treasury sanctions notorious virtual currency mixer Tornado Cash*.

U.S. Department of the Treasury (2023). *Treasury designates Roman Semenov, co-founder of sanctioned virtual currency mixer Tornado Cash*.

van Wegberg, R., Oerlemans, J.-J., and van Deventer, O. (2018). "Bitcoin money laundering: mixed results? An explorative study on money laundering of cybercrime proceeds using Bitcoin," *Journal of Financial Crime* (45:2), pp. 419–435. https://doi.org/10.1108/JFC-11-2016-0067.

Wang, Z., Chaliasos, S., Qin, K., Zhou, L., Gao, L., Berrang, P., Livshits, B., and Gervais, A. (2022). "On how zero-knowledge proof blockchain mixers improve, and worsen user privacy," in *Proceedings of the ACM Web Conference,* pp. 2022–2032. https://doi.org/10.1145/3543507.3583217.

Zhang, R., Xue, R., and Liu, L. (2019). "Security and privacy on blockchain," *ACM Computing Surveys* (52:3). https://doi.org/10.1145/3316481.

Zheng, P., Zheng, Z., Wu, J., and Dai, H.-N. (2021). "On-chain and off-chain blockchain data collection," in. *Blockchain Intelligence: Methods, Applications and Challenges,* pp. 15–39. https://doi.org/10.1007/978-981-16-0127-9_2.

Ziegler, C. and Welpe, I. M. (2022). "A taxonomy of decentralized autonomous organizations," in *Proceedings of the 43rd International Conference on Information Systems,* AIS.

IX | Research Paper 3 – ESG KPI Reporting System I

Full Title:
Designing a Reporting System for Trust in Environmental Social Governance

Publication venue:
58th Hawaii International Conference on System Sciences (HICSS) 2025

URL:
hdl.handle.net/10993/63596

Designing a Reporting System for Trust in Environmental Social Governance

Egor Ermolaev
University of Luxembourg
egor.ermolaev@uni.lu

Evgenia Yvonni Tseloni
University of Luxembourg
evgenia.tseloni@uni.lu

Tamara Roth
University of Arkansas
troth@walton.uark.edu

Gilbert Fridgen
University of Luxembourg
gilbert.fridgen@uni.lu

Abstract

In an effort to make companies adhere to green practices, they are increasingly required to publicly disclose environmental, social, and governance (ESG) key performance indicators (KPIs). Typically published within annual reports, the current format may not deliver all legally mandated information, and also does not attract investors who have become more focused on sustainable investing. Thus, various industries aim at improving their ESG KPIs reporting practices, including life insurance companies. Life insurance companies need to integrate ESG-focused strategies into their client interactions to obviate greenwashing allegations and appeal to their environmentally-friendly target group. This helps them remain competitive and build long-lasting institution-based trust. To help these companies address the complexities of ESG KPIs reporting, this study proposes the development of a reporting system architecture supported by Distributed Ledger Technology (DLT) for compliant, transparent, reliable, and standardized ESG reporting.

Keywords: Reporting, Architecture, Greenwashing, Green Information Systems, Blockchain

1. Introduction

The call for more sustainable practices across all industries has made the integration and transparent reporting of ESG practices in companies indispensable (Kim et al., 2021). Both consumers and business partners demand transparency that goes beyond the occasional mention of ESG KPI in annual reports (KPIs for ESG, 2010; Nugroho et al., 2024). More specifically, the corporate image, perceived price fairness, and purchasing behavior can be influenced by how accessible and transparent a company's ESG practices are (Park & Han, 2021). AbuRaya (2017) and Zhan (2023) emphasize the importance of governance to improve the frequency and quality of corporate environmental sustainability, while also calculating ESG risk (Awan, 2011; Vasiu & Bratu, 2022; Ziolo et al., 2023), i.e., potential financial repercussions due to the implementation of ESG practices (Bax et al., 2023), and their influence on corporate performance. However, these ESG risks appear small given the potential cost savings, increased operational efficiency, long-term resilience (Sousa et al., 2023), and green productivity afforded by business models that are well aligned with a company's ESG strategies (Gaur et al., 2011).

Moreover, regulators have tightened their legal frameworks regarding reporting and disclosure requirements (Ho, 2023) to improve the transparency and quality of ESG-related data, e.g. in the European Union (EU). Their frameworks focus particularly on the financial industry – based on the FinTech Action Plan (2018) and the Sustainable Finance Strategy (2021) – to help market participants compare the sustainability of different financial products (Duran & Tierney, 2023). Anticipating that such frameworks will soon also be issued for other industries, companies are well-advised to systematically improve their ESG reporting (Jean & Grant, 2022). Additionally, ESG strategies and goals have evolved into an investment strategy, with sustainable investing being projected to surpass the $40 trillion mark by 2030 (Bloomberg, 2024).

Despite the evident need for companies to demonstrate excellence in reaching ESG goals (Hang et al., 2023), ESG implementation and reporting comes with challenges. Specifically, companies lack comprehensive and reliable metrics (Vasiu & Bratu, 2022). Although Cruz and Matos (2023)

URI: https://hdl.handle.net/10125/109520
978-0-9981331-8-8

have envisioned a solution in the form of a software framework, it still needs comprehensive real-world testing and validation, as it is still in development and the data on its effectiveness and user satisfaction are limited. Moreover, the software framework by Cruz and Matos (2023) does not anticipate the potential institution-based trust issues related to potential greenwashing allegations. Greenwashing describes marketing tactics (Chueca Vergara & Ferruz Agudo, 2021) aimed at deceiving stakeholders about the environmental impact of an organization or promoting an unsustainable product or service as sustainable by highlighting desirable and masking undesirable features (Delmas & Burbano, 2011; Duran & Tierney, 2023). Building on studies such as Lindgren et al. (2021), Prabawani and Hadi (2022), we thus not only highlight the need for better measurement and certification of green business models and products but also ask how we can build an information system that supports ESG integration and facilitates ESG reporting while avoiding greenwashing and providing trustful and reliable results.

We aim to tackle this need using a design science research (DSR) approach to design an architecture for an ESG KPI reporting system. Our paper is structured according to the framework suggested by Gregor and Hevner (2013): it begins with a *Literature Review* in the theoretical background which sheds light on the convergence of ESG reporting and the burgeoning discussion of blockchain in IS within the EU legal framework and institution-based trust theory. This is followed by the *Method* section that details our approach. We then describe and evaluate our artifact in the *Artifact Description* and *Evaluation* sections. The paper continues with the *Discussion* of our findings and implications, culminating in the *Conclusions* section that summarizes key insights and potential areas for future research in FinTech.

2. Theoretical background

2.1. Fundamentals of ESG (reporting)

Good ESG reporting follows two purposes (Ho, 2023): (1) to showcase the positive impact of investments in sustainability, which ideally, encourages more such investments (Chopra et al., 2024) and (2) to mitigate investment risks resulting from reputational damage due to non-compliance with promoted sustainability strategies (EBA Report, 2024). Encouraged to leverage the benefits of good ESG reporting, many businesses proactively focus on establishing or refining standards and accountability frameworks. While current and imminent regulations provide guidance, they often need to be translated into the organizational context with the help of ESG KPIs.

ESG KPIs serve as quality indicators for products and aim to improve the presentation, guidelines, structure, and content of ESG reporting. As part of company performance reports they ensure that relevant data is complete and provided in an accessible format (KPIs for ESG, 2010). This way, the overall report provides different stakeholders and potential investors with a transparent and comprehensive overview of a company's ESG strategies and practices, allowing for the comparison of products (KPIs for ESG, 2010). Dependent on a company's focus and goals, ESG KPIs can be versatile. They can, for instance, encompass energy consumption, carbon footprint, and the circular economy, as well as social topics, such as issues of inequality, labor relations, and human rights (European Commission, 2024). Regarding the actual ESG reporting, KPIs such as transparency on management structures or accountability frameworks are crucial.

2.2. EU ESG-related Legal Frameworks

These accountability frameworks can, for instance, be inspired by EU legislation. The EU has already declared its intention for the financial industry to accelerate its transition to a net-zero economy by 2050 (EU Taxonomy, 2020). To support this transition, the EU has established a comprehensive package of measures, including policy papers, the delegation and implementation of acts, directives, and regulations. The package should also make corporate sustainability reporting more common and standardized, reducing greenwashing and preventing market and regulatory fragmentation among the member states. Moreover, it can provide transparency to investors so that they can adequately assess and value investments risk related to, for instance, EU climate. The most significant EU legislative documents on ESG sustainability reporting are the EU Taxonomy Regulation (EU Taxonomy) 852/2020 (2020), Sustainable Finance Disclosure Regulation (SFDR) 2088/2019 (2019), and Corporate Sustainability Reporting Directive (CSRD) 2022/2464 (2022).

The EU Taxonomy is a classification system that sets six main environmental objectives (article 9 EU Taxonomy) and respective criteria to ensure economic activities are sustainable, functioning as a basis for standards and labels for sustainable financial products. The SFDR regulation supplements the EU Taxonomy by enhancing investors' protection and is directly related to the ESG KPI reporting(s). It specifies sustainability disclosure requirements for

financial market participants, such as the inclusion of product-related and entity-related disclosure in pre-contractual documents (articles 8-9) or the publication of concrete sustainability strategies on websites (article 10) and periodic reports (article 11) (Macchiavello & Siri, 2020). The CSRD introduces the concept of double materiality according to which companies report their impact from two mutually constitutive views. That is, it assesses the effects of environmental changes and sustainability issues due to, for instance, ESG, on a company's (economic) performance and probes into different social sustainability issues, such as human rights.

2.3. Trust Reduction in ESG Reporting due to Greenwashing

While ESG-related legal frameworks already try to tackle greenwashing in ESG reporting, more needs to be done to avoid the impression of fraudulent practices. Greenwashing typically aims to mislead consumers regarding the environmental practices of a company (firm-level) or the environmental benefits and impact of a product or service (product/service level), both of which appear sustainable but are not (Delmas & Burbano, 2011; Duran & Tierney, 2023). In ESG reporting, greenwashing particularly manifests in the form of vague descriptions of concrete sustainable practices, omitting or concealing undesirable information, or overstating the effects of the few ESG strategies companies actually implement (Chueca Vergara & Ferruz Agudo, 2021).

This practice may not only violate compliance with established legal frameworks but can also negatively influence institution-based trust. Trust is a fundamental factor in any kind of (business) relationship when risk or uncertainty are involved (McKnight et al., 2009). Institution-based trust, in particular, describes an individual's feelings of relative security towards an impersonal structure despite potential risks (Goo & Nam, 2007; McKnight et al., 1998). Structural assurance and situational normality are the two subconstructs of institution-based trust (Goo & Huang, 2008). They can be directly impacted by, for instance, greenwashing practices since the legislative normative efforts currently only constitute an intention to implement sustainability strategies but are not binding, thus, threatening Structural Assurance (Foley et al., 2024). Greenwashing also unsettles the belief in sustainability promises and the associated Situational Normality, which causes trust to crumble (Foley et al., 2024; McKnight & Chervany, 2001) and violates the qualities of a trustworthy trustee (McKnight & Chervany, 2001). That is, greenwashing not only undermines trust in the entire ESG mechanism but also the trustee's integrity, as it provides a false willingness to help, thereby breaching the principle of good faith (McKnight & Chervany, 2001). To prevent not only the practice but also potential allegations of greenwashing that could negatively impact a company's reputation, technical solutions can be used to ensure ESG data reliability and integrity, improve data accuracy, and facilitate regulatory compliance.

2.4. Trust-Building through Blockchain Technology

Blockchain technology could prove to be a valuable building block in creating such a solution. Since its inception in 2008 (Satoshi, 2008), blockchain has been hailed as a technology that would replace interpersonal trust with trust in technology (Casey & Vigna, 2018; De Filippi et al., 2020; Utz et al., 2023; Ziolkowski et al., 2020). Technically speaking, blockchains are distributed databases that record transactional data in a chronological order on several blockchain nodes in a blockchain network (Jones, 2019; Swan, 2015). Each basic ordering element, called 'block', is cryptographically linked via hash functions with the previous block, creating a chain (Scholl et al., 2020). The proposition of the next block is typically tied to a scarce resource such as the amount of energy required to solve computational puzzles or a certain cryptocurrency balance put at stake (Rieger et al., 2022). This not only makes blockchains difficult to manipulate but each attempt would also be transparently recorded (Pincheira et al., 2020). Additionally, the increase in energy consumption of a specific blockchain should always be balanced with the savings and benefits it provides (Sedlmeir et al., 2020).

Most blockchains also support the deployment of deterministic programming logic, so-called Smart Contracts (Szabo, 1994). They are characterized by self-enforceability (Rozas et al., 2021). From the moment they are created, no human interference is required since the pre-specified programming logic underlying smart contracts ensures that all requirements are met and transactions executed as well as recorded on the blockchain (Raskin, 2016). In an environment characterized by uncertainty and risks, the assurance that contractually agreed parameters will be met, given that involved parties cannot meddle with the functionality of the smart contracts after deployment, contribute another valuable trust-building component of blockchain technology (Mendling et al., 2018; Savelyev, 2017) and decrease ex-ante smart contract

specification costs changing the predominance of transaction governance mode (Halaburda et al., 2024).

3. Research Method

3.1. DSR approach

We used DSR to explore how blockchain technology can be used to improve institution-based trust in ESG reporting through the development of a blockchain-based ESG KPI reporting system. This system is designed to improve compliance, reliability, transparency, and standardization in ESG reporting. DSR is a well-established research method widely used in the design and creation of various Information Technology (IT)-based artifacts, including constructs, frameworks, architectures, models, methods, and algorithms (Peffers et al., 2007). Additionally, DSR addresses more abstract artifacts such as social innovations and design propositions, as well as technical and social properties, design principles (DPs), and theories (Utz et al., 2023). This makes DSR a suitable method to develop a technical tool that improves ESG reporting.

We started our DSR project with an extensive literature review to establish our initial design requirements and features. These were further refined during a 48-hour long and intensive hackathon, which involved collaboration with a life insurance company and technical representatives from a blockchain lab. The project was then presented and defended before an interdisciplinary jury of seven expert jurors. Additionally, we developed a proof-of-concept and conducted on-site testing during the hackathon. The resulting artifact convinced the jurors to select our project for the final, international competition.

During the design iterations, we could gain generalizable knowledge in the form of DPs that can offer guidance for the incorporation of trust-building components in ESG KPI reporting systems and contribute to theories on trust-building in institutions through technology. Moreover, we incorporated and evaluated our artifact against current legal frameworks for ESG KPI reporting. This makes our artifact relevant from both a theoretical and practical perspective (Gregor & Hevner, 2013; Hevner & Chatterjee, 2012). Our DPs make a knowledge contribution of the exaptation type. Exaptation requires the extension of a known solution to new problems (Gregor & Hevner, 2013).

3.2. Problem and objectives

We started with a comprehensive literature review across various databases, including Google Scholar, Scopus, and Web of Science. Additionally, we employed AI-powered tools such as Elicit. Our keywords were “reporting” AND “ESG” OR “greenwashing”, “ESG” AND “responsible investment”, “greenwashing” AND “trust” OR “institutional trust”, as well as “ESG” AND “reporting” AND “trust” OR “institutional trust”, “ESG” AND “blockchain”, and “blockchain” AND “trust” OR “institutional trust”. Our initial search yielded a plethora of different articles and book chapters. In line with (Webster & Watson, 2002), we excluded articles that were not written in English, book chapters, and articles published in journals with a percentile below 85 % on Scopus. We then reviewed titles, and abstracts of a high-quality subset relevant to our study. While reading some of the selected articles, we included additional papers that proved relevant to our topic. From the combined set, we developed a preliminary problem statement and identified a first set of design requirements. We further refined these initial requirements during the hackathon.

3.3. Demonstration and evaluation

We created and refined our high-level architecture (HLA) for a blockchain-based ESG KPI reporting system in three design and development iterations during the hackathon. This included multiple iterations involving senior management and the technical experts from various sectors (Tab. 1). These interactions helped us contextualize and refine our project concept. It also helped us substantiate our design requirements (DRs), translate them into design features (DFs), and continuously improve our conceptual architecture through iterative build-and-evaluate cycles.

Discussions with stakeholders from the life insurance company helped us better understand the challenges associated with assessing ESG KPIs in practice. This reinforced our decision to use blockchain technology as an essential building block. Consultations with technical representatives from the blockchain lab supported the technical feasibility of our resulting proof-of-concept. We then presented a preliminary version of our conceptual architecture for a blockchain-based ESG KPI reporting system to an interdisciplinary panel of jury members – the domain experts from various business sectors – to collect feedback for further improvement. Additionally, we conducted an evaluation of the proof-of-concept against established ESG frameworks, such as EU Taxonomy, SFDR, through conceptual analysis to ensure its feasibility and relevance.

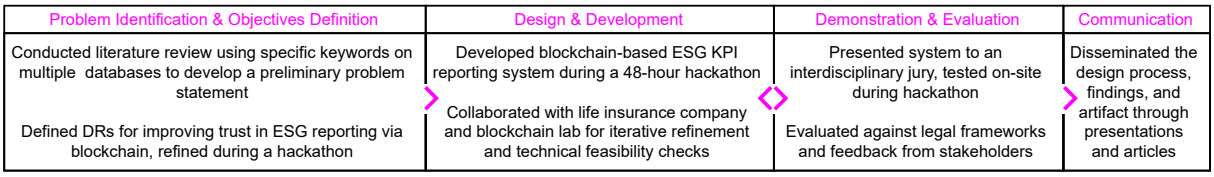

Figure 1. Adapted DSR Model based on Peffers et al., 2007

4. A blockchain-based ESG KPI reporting system

4.1. Design requirement

After our discussion with stakeholders from the life insurance company, we first refined our DRs and translated them into concrete DFs relevant for the development of our artifact.

DR1 – Compliance. The implementation of the ESG KPI reporting system should adhere to the aforementioned regulatory framework and *allow companies to submit ESG KPIs (DF1)* to the blockchain. Dependent on a company's business domain, *the system should allow companies to select a set of ESG KPIs that best fit their context (DF2)*.

DR2 - Reliability. Submitted ESG KPIs should be *immutable (DF3)* after the moment of submission and must not be changed or deleted.

DR3 - Transparency. Submitted ESG KPIs should be *publicly accessible (DF4)* for anyone online at any point in time from the public ledger.

DR4 - Standardization. All ESG KPIs and sets thereof are customized for different business domains and predefined. So are *types and lengths of each ESG KPI's value (DF5)* that allows for operating the values of submitted ESG KPIs in a standardised way.

4.2. Description of the artifact

Our artifact intends to provide a blockchain-based tool for compliant, reliable, and transparent ESG KPI reporting. The aim is to enable companies to submit their ESG KPIs (DF1) in a standardized manner (DF5), which also supports investors in assessing the sustainability of companies. Based on blockchain technology, companies will submit ESG KPIs on the public ledger, instead of publishing them in their annual reports, where they are transparently and immutably stored (DF3). The artifact also helps companies become ESG-compliant by default, since it is tailored to their specific business domain needs. Eccles et al. (2012) sustain the need for reporting standards that reflect the need of different industries so that only the most essential ESG dimensions are identified and

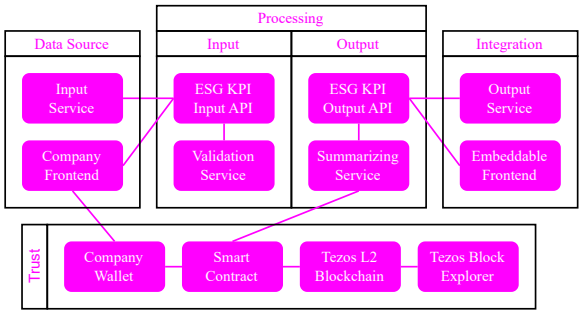

Figure 2. High-level architecture

reported. We envision that companies select their business domain and industry-specific ESG KPIs as a part of an on-boarding process (DF2). This not only increases the relevance of ESG KPIs for companies but also fosters transparency for investors, who can better assess where to allocate their sustainable investments based on available and recorded information (DF4).

The proposed artifact has four layers: the Data Source Layer, the Input/Output Processing Layers, the Trust Layer, and the Integration Layer (Fig. 2).

The **Data Source Layer** allows companies to submit ESG KPI data to the blockchain in a standardised manner (DF5) and in line with EU Taxonomy (2020) framework. The Company Frontend enables the manual submission of ESG KPI data. The Input Service supports non-manual submission of the ESG KPI data via application programming interface (API) of the Input Processing Layer. The submitted ESG KPI data is pipelined into the Input Processing Layer.

The **Input Processing Layer** operates via the ESG KPI Input API and validates ESG KPI input data within its Validation Service. The Validation Service cross-checks the submitted ESG KPI data according to the company's business domain under the EU Taxonomy (2020) itself (and its alignment), CSRD and SFDR. Different business domains also have different sets of ESG KPIs that need to be differentiated and submitted (DF2). The Input Processing Layer does not write a company's ESG KPI data to the blockchain but prepares raw transactions (TXs), which are returned to the Company Frontend (or in a response to the API call

within the ESG KPI Input API) for employees of the company to include data. All returned raw TXs are signed by the employees in the Company Wallet and submitted on the Tezos Layer 2 (L2) on their own. The choice of Tezos L2 was dictated by its Ethereum virtual machine (EVM)-compatibility. The TX is essentially a call of the smart contract's function to write a company's ESG KPI data that has been pre-validated within the Input Processing Layer. All functionalities of the Input Processing Layer can be transferred to the Company Frontend in the next version of the proof-of-concept.

The **Trust Layer** is built around the existing Tezos ecosystem on which testnet we deployed the smart contract. Companies use the Company Wallet to submit pre-validated ESG KPI data to the single smart contract. The smart contract, written in Solidity, manages both public ESG KPIs. It provides functions to write and read public KPI values directly using a unique key generated from a company's wallet address and KPI identifier. This smart contract allows access to KPI data where they are openly accessible. It is possible to read the ESG KPI data directly from the Trust Ledger, using, for instance, Tezos Block Explorer (DF4). Once written to the smart contract, ESG KPIs are stored there permanently in an append-only manner and cannot be changed retrospectively (DF3).

The **Output Processing Layer** enables investors, for instance, to request a summary of a company's KPIs by reading ESG KPI data from the Trust Layer and summarizing however many ESG KPIs into a sustainability metric (0–100 %) – including explanations on how this one-dimensional metric has been calculated – that reflects the company's environmental impact. The Summarizing Service processes the ESG KPI data and outputs in a convenient format for the end users.

The **Integration Layer** uses the ESG KPIs Output API that provides access to the Output Processing Layer and retrieves the summary of company-specific ESG KPI data for each asset in a company's portfolio. Our life-insurance company, for instance, provides clients with an investment app showcasing their portfolio and explaining how they invest the client's premiums.

4.3. Evaluation of the artifact

To evaluate and refine our ESG KPI reporting tool, we used the hackathon and its various feedback rounds as a workshop-like evaluation method (Thoring et al., 2020). Hackathons provide opportunities for gathering concrete and constructive feedback in a multi-stage workshop setting with a final expert jury evaluation. Although we did not expect such a setting to produce a final, refined architecture for our envisioned ESG KPI reporting tool, the multi-stage workshop setting demonstrates its practical viability and potential (Thoring et al., 2020). Moreover, hackathons typically select jury members according to their specific domain knowledge so that they can evaluate the design according to defined parameters. Thus, different representatives of both the business and the technical side provide feedback to develop and refine the artifact. In our case, business representatives from the life insurance sector and technical experts were present. These stakeholders provided vital input on business needs and technical feasibility, ensuring the artifact met both business and technical requirements within existing infrastructures (Thoring et al., 2020).

Throughout the various group discussions – our primary source for feedback (Thoring et al., 2020) – we extensively documented observations and insights, ensuring that we captured the nuances and specific feedback from the stakeholders involved. These notes provided a rich source of qualitative data that helped us iteratively improve the design and functionality of our ESG KPI reporting system. Also, the qualitative data supported our reflection on the overall artifact, which helped us better understand its practical implications and supported our development of generalizable DPs.

For the **first design iteration** we drafted the initial HLA design. During this initial iteration we evaluated the functionality of each component with representatives of the life-insurance company. They pointed out that ESG KPI were currently difficult to process and should instead be summarized per company so that clients can "build emotional repertoire" (Expert 7) with a company's ESG goals. To address this feedback, we proposed the Summarising Service. Our proposal was well received, however, the representatives suggested to "mitigate input flaws by validation" (Expert 6), and requested the inclusion of only the most relevant ESG KPIs for different business domains (EU Taxonomy, 2020). Thus, we added the Input Validation Service.

In the **second design iteration** we refined the HLA based on the prior feedback. During this iteration, we evaluated the HLA with technical experts. They gave us feedback on our design's feasibility, specifically its smart contract functionality, based on a digital schematic HLA on Draw.io. The technical team recommended using the test network of Tezos for our proof-of-concept. They also suggested utilizing its L2 should we be familiar with EVM tools and Solidity so that our solution could be interoperable with all the EVM-compatible ecosystems. Based on their feedback, we included the suggested tools (Solidity, Hardhat, Wallet Connect) and parts of the proposed ecosystem (testnet of Tezos L2) in our proof-of-concept and HLA.

During the **third design iteration**, the proposed solution was comprehensively evaluated by the expert jury (Tab. 1) based on a detailed criteria framework. The jury was chosen according to their experience in, for instance, software development, IT services, insurance, aviation, and education. They had been in their respective professions from four to over twenty years and occupied professional roles such as Adoption Manager (AM), Business Advisor (BA), Chief Executive Officer (CEO), Chief Information Officer (CIO), Chief Communications Officer (CCO), Marketing Director (MD), and Head of corporate social responsibility (CSR). Combined, they could provide a well-rounded perspective on the industry-specific challenges and practices. Their evaluation focused particularly on the project (team), its feasibility, and the viability of the design. Additionally, the quality of the pitch, the clarity and conciseness of the project summary, the business model, the presentation deck, and responses to the jury's inquiries were part of the rating. Blockchain experts were responsible for the technical evaluation and focused particularly on the integration of the Tezos blockchain, the innovativeness of the design's features, the user experience, and interface design.

Initially, all participating teams competed at the hackathon stage. From there, the semi-final included all teams, divided into two groups, with each group evaluated by a dedicated jury board. The top six teams from the semi-finals advanced to the final round, where a single jury board evaluated their projects to determine the winners. Our project won a price and was selected for the finale in Paris, scheduled for autumn 2024.

After reflection on the evaluation during the hackathon, we decided to add a criterion-based analysis with legal frameworks such as SFDR, EU Taxonomy, and CSRD to test our **final design**. We compiled materials from the hackathon and performed the analysis to ensure that our ESG KPI reporting tool complies with regulation. The analysis revealed that we currently do not yet deliver on the double materiality of CSRD since our ESG KPI have primarily focused on environmental sustainability. However, the domain-specific selection of KPI could contain social sustainability goals if they are a particular priority for one company or industry. Regarding SFDR and EU Taxonomy, our artifact delivered on the relevant transparency and standardization of ESG KPI for each domain while accounting for domain-specific differences.

№	Industry	Size, EE	Position	Exp., y.
1	Blockchain	51-200	AM	6
2	IT Services	2-10	BA	7
3	Insurance	51-200	CEO	20+
4	Airlines	1001-5000	CIO	20+
5	Education	11-50	CCO	4
6	Insurance	51-200	MD	20+
7	Insurance	51-200	CSR	15

Table 1. Experts

5. Discussion

The evaluation of our conceptual architecture produced valuable insights for the design of blockchain-based ESG KPI reporting systems. Following in the footsteps of Gregor and Hevner (2013), we have identified three design principles that can be useful to practitioners who wish to design and successfully implement such a system. Using blockchain technology as a key building block in our artifact, we also contribute to theory on trust-building in institutions through technology, since blockchain appears to prevent greenwashing allegations by enhancing reliability, transparency, standardisation, and compliance for companies dealing with ESG reporting (Pöll, 2024; Smits & Hulstijn, 2020; Utz et al., 2023).

5.1. Practical implications

We began the design of this artifact with an observation that current ESG KPI reporting often lacks rigor and standardization, which can confuse customers, stakeholders, and potential investors (KPIs for ESG, 2010). This confusion can also result in the impression of greenwashing practices, since ESG goals are often ambitious and their implementation strategies and accountability frameworks vague (EBA Report, 2024). To address this problem, we provided not only an innovative tool for ESG KPI reporting but also illustrated in the development of our design how practitioners could improve their current ESG reporting.

Specifically, we demonstrate how regulation could be used as a guide in the innovation of reliable and trustworthy ESG KPI reporting tools. Many companies have settled for infeasible ESG goals (EBA Report, 2024) and unfocused strategies, which made it difficult to discern relevant raw data to assess the outcomes of a strategy. Inspired by EU legislation for the financial industry EU Taxonomy (2020), SFDR and CSRD, we used their comprehensive package of measures to revise ESG KPI for other industries. Our artifact makes companies compliant by default by entering domain- and company-specific ESG KPI. That is, the **Data Source Layer** enforces the submission of ESG KPI data in a standardised manner – either automatically or manually – via the Input Service. Once submitted, the **Input Processing Layer** validates ESG KPI input data, returns the raw TX, and **Data Source Layer** writes it to

Trust Layer through the Company Wallet. This results in a first design principle **DP1 – Compliant submission of raw data**.

Moreover, our project highlights the use of the Tezos blockchain as an essential building block for delivering transparency. In the form of a **Trust Layer**, the submitted pre-validated ESG KPI data entries for companies, managed by the smart contract, are publicly accessible. We also made sure that people around the world can access the public ESG KPI data via the Tezos Block Explorer to prevent greenwashing allegations due to information accessibility issues (Utz et al., 2023). Thus, our second DP is **DP2 – Public access to immutable raw data**.

Finally, we ensured that ESG KPI reporting encourages real change and prevents accidental greenwashing by distinguishing ESG KPI not only at a company but also an industry level. This way, companies have to comply with industry standards and are less likely to formulate vague ESG goals that sound more like mission statements instead of binding commitments. Implementing the **Output Processing Layer** with a Summarizing Service that gives companies a clear indication of what is required and how their ESG will be assessed. The resulting report, which derives its data directly from the **Trust Layer**, also contains a sustainability metric that provides explanations on how it has been calculated. This results in our third design principle, namely **DP3 – Industry-specific reporting goals and transparent assessment metrices**.

5.2. Theoretical implications on trust in technology

The design of our artifact contributes to literature on trust-building in institutions through technology. Fostering such trust requires an individual's willingness and belief in the positive attributes of functionality, helpfulness, and reliability of a technology even in situations where negative consequences are possible (McKnight et al., 2009). However, once individuals perceive a technology to be a driver of trust, it can be used across different domains to mend or establish trust where trust-levels have been low (Ying et al., 2018).

This is particularly relevant in the (re-)establishment of institution-based trust (Ying et al., 2018). Our choice of blockchain technology for the development of our artifact also lies in its reputation as a motor of trust in an otherwise trust-less environment (McKnight et al., 2020). As our first (DP1) and second design principle (DP2) demonstrate, we find that the use of blockchain enhances structural assurance because it mandates devotion to the duality of legal and algorithmic rule frameworks (Ziolkowski et al., 2020). By complying with the EU's regulatory framework for sustainable finance and governance of financial and corporate actors, our artifact provides a high level of transparency regarding ESG KPI and reduces risks related to ESG by employing standards and easy-to-understand metrics. Moreover, using blockchain in our artifact also satisfies situational normality because it has the reputation of being tamper-proof and reduces human interference through smart contracts (De Filippi et al., 2020; McKnight et al., 2009). Giving users access to ESG KPI data and transparent as well as reliable reports further delivers on trust dimensions that knowledge and the ability to verify data in order to build trust in an institution (McKnight et al., 2017).

5.3. Limitations of this study and potential for further research

The current study contains the first design of an ESG KPI reporting tool and has not yet been implemented in a real-world setting. Thus, our current DPs may either be extended or slightly adapted once we reach the implementation stage. In particular, we need to further investigate the use of privacy enhancing technologies for submission and verification of private data entries. Research in this area is already underway but requires a more in-depth exploration (Sedlmeir et al., 2022). Moreover, the incorporation of regulation in the design of artifacts and simultaneous evaluation may warrant a separate investigation.

6. Conclusion

Our artifact for general-purpose ESG KPI reporting can make a valuable contribution for companies interested in improving their ESG KPI reporting. The use of blockchain technology as an essential building block may prevent greenwashing allegations by enhancing reliability, transparency, standardisation, and compliance with EU regulation. Moreover, it shows potential to enhance institution-based trust by providing an immutable and transparent ledger of a company's ESG KPI.

7. Acknowledgements

This research was funded by the Luxembourg National Research Fund (FNR) and PayPal, PEARL grant ref. 13342933/Gilbert Fridgen and NCER22/IS/16570468/NCER-FT. For the purpose of open access, and in fulfillment of the obligations arising from the grant agreement, the author has applied a

References

2022/2464. (2022). *Directive (eu) 2022/2464 of the european parliament and of the council of 14 december 2022 amending regulation (eu) no 537/2014, directive 2004/109/ec, directive 2006/43/ec and directive 2013/34/eu, as regards corporate sustainability reporting*. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022L2464

2088/2019. (2019). *Regulation (eu) 2019/2088 of the european parliament and of the council of 27 november 2019 on sustainability-related disclosures in the financial services sector*. https://eur-lex.europa.eu/eli/reg/2019/2088/oj

852/2020. (2020). *Regulation (eu) 2020/852 of the european parliament and of the council of 18 june 2020 on the establishment of a framework to facilitate sustainable investment, and amending regulation (eu) 2019/2088*. https://eur-lex.europa.eu/eli/reg/2020/852/oj

AbuRaya, R. (2017). Corporate environmental disclosure and corporate governance: A critical review. *Journal of Empirical Research in Accounting; Auditing An International Journal*, *04*(01), 24–53.

Awan, U. (2011). Green marketing: Marketing strategies for the swedish energy companies. *International Journal of Industrial Marketing*, *1*(2), 1.

Bax, K., Sahin, Ö., Czado, C., & Paterlini, S. (2023). Esg, risk, and (tail) dependence. *International Review of Financial Analysis*, *87*, 102513.

Bloomberg. (2024). *Global esg assets predicted to hit $40 trillion by 2030, despite challenging environment, forecasts bloomberg intelligence*. https://www.bloomberg.com/company/press/global-esg-assets-predicted-to-hit-40-trillion-by-2030-despite-challenging-environment-forecasts-bloomberg-intelligence

Casey, M. J., & Vigna, P. (2018). *In blockchain we trust*. https://www.technologyreview.com/2018/04/09/3066/in-blockchain-we-trust/

Chopra, S., Senadheera, S., Dissanayake, P., Withana, P., Chib, R., Rhee, J., & Ok, Y. S. (2024). Navigating the challenges of environmental, social, and governance (esg) reporting: The path to broader sustainable development. *Sustainability*, *16*, 606.

Chueca Vergara, C., & Ferruz Agudo, L. (2021). Fintech and sustainability: Do they affect each other? *Sustainability*, *13*(13), 7012.

Cruz, C. A., & Matos, F. (2023). Esg maturity: A software framework for the challenges of esg data in investment. *Sustainability*, *15*(3), 2610.

De Filippi, P., Mannan, M., & Reijers, W. (2020). Blockchain as a confidence machine: The problem of trust & challenges of governance. *Technology in Society*, *62*, 101284.

Delmas, M. A., & Burbano, V. C. (2011). The drivers of greenwashing. *California Management Review*, *54*(1), 64–87.

Duran, R. E., & Tierney, P. (2023). Fintech data infrastructure for esg disclosure compliance. *Journal of Risk and Financial Management*, *16*(8), 378.

EBA Report. (2024). *Greenwashing monitoring and supervision*. https://www.eba.europa.eu/publications-and-media/press-releases/esas-call-enhanced-supervision-and-improved-market-practice-sustainability-related-claims

Eccles, R. G., Krzus, M. P., Rogers, J., & Serafeim, G. (2012). The need for sector-specific materiality and sustainability reporting standards. *ERN: Regulation (IO) (Topic)*.

EU Taxonomy. (2020). *Eu taxonomy for sustainable activities*. https://finance.ec.europa.eu/sustainable-finance/tools-and-standards/eu-taxonomy-sustainable-activities_en

European Commission. (2024). *Overview of sustainable finance*. https://finance.ec.europa.eu/sustainable-finance/overview-sustainable-finance_en

FinTech Action Plan. (2018). *Fintech action plan: For a more competitive and innovative european financial sector*. https://finance.ec.europa.eu/publications/fintech-action-plan-more-competitive-and-innovative-european-financial-sector_en

Foley, A. M., Heffron, R. J., Al Kez, D., Furszyfer Del Rio, D. D., McInerney, C., & Welfle, A. (2024). Restoring trust in esg investing through the adoption of just transition ethics. *Renewable and Sustainable Energy Reviews*, *199*, 114557.

Gaur, S., Matta, G., & Singh, V. (2011). Importance and role of green productivity in industries: A review. *Environment Conservation Journal*, *12*(1 and 2), 129–133.

Goo, J., & Huang, C. D. (2008). Facilitating relational governance through service level agreements in it outsourcing: An application of the commitment–trust theory. *Decision Support Systems*, *46*(1), 216–232. https://doi.org/https://doi.org/10.1016/j.dss.2008.06.005

Goo, J., & Nam, K. (2007). Contract as a source of trust–commitment in successful it outsourcing relationship: An empirical study. *Proceedings of Hawaii International Conference on System Sciences (HICSS)*.

Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, *37*(2), 337–355.

Halaburda, H., Levina, N., & Semi, M. (2024). *Digitization of transaction terms within tce: Strong smart contract as a new mode of transaction governance*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4501318

Hang, L., Wang, L., & Xu, P. (2023). The impact of corporate esg performance on corporate green innovation.

Hevner, A., & Chatterjee, S. (2012). *Information systems theory*. Springer New York.

Ho, J. (2023). Banking, blockchain, and esg. In *The role of distributed ledger technology in banking: From theory to practice* (pp. 235–266). Cambridge University Press.

Jean, M. S., & Grant, E. (2022). Management system enabled esg performance. *Proceedings of International Pipeline Conference (IPF)*.

Jones, E. (2019). Rethinking greenwashing: Corporate discourse, unethical practice, and the unmet potential of ethical consumerism. *Sociological Perspectives*, *62*(5), 728–754.

Kim, J., Kim, M., Im, S., & Choi, D. (2021). Competitiveness of e commerce firms through esg logistics. *Sustainability*, *13*(20), 11548.

KPIs for ESG. (2010). *Kpis for esg*. https://effas.com/wp-content/uploads/2021/09/KPIs_for_ESG_3_0_Final.pdf

Lindgren, P., Knoth, N. S. H., Sureshkumar, S., Friedrich, M. F., & Adomaityte, R. (2021). "green multi business models" how to measure green business models and green business model innovation?

Macchiavello, E., & Siri, M. (2020). Sustainable finance and fintech: Can technology contribute to achieving environmental goals? a preliminary assessment of 'green fintech"'. *SSRN Electronic Journal*.

McKnight, D., Carter, M., & Clay, P. F. (2009). Trust in technology: Development of a set of constructs and measures. *Proceedings of Diffusion Interest Group In Information Technology (DIGIT)*.

McKnight, D., & Chervany, N. (2001). Trust and distrust definitions: One bite at a time. https://doi.org/10.1007/3-540-45547-7_3

McKnight, D., Cummings, L., & Chervany, N. (1998). Initial trust formation in new organizational relationships. *The Academy of Management Review*, *23*(3), 473.

McKnight, D., Lankton, N. K., Nicolaou, A., & Price, J. (2017). Distinguishing the effects of b2b information quality, system quality, and service outcome quality on trust and distrust. *The Journal of Strategic Information Systems*, *26*(2), 118–141.

McKnight, D., Liu, P., & Pentland, B. T. (2020). Trust change in information technology products. *Journal of Management Information Systems*, *37*(4), 1015–1046.

Mendling, J., Weber, I., Aalst, W. V. D., Brocke, J. V., Cabanillas, C., Daniel, F., Debois, S., Ciccio, C. D., Dumas, M., Dustdar, S., Gal, A., García-Bañuelos, L., Governatori, G., Hull, R., Rosa, M. L., Leopold, H., Leymann, F., Recker, J., Reichert, M., … Zhu, L. (2018). Blockchains for business process management - challenges and opportunities. *ACM Trans. Manage. Inf. Syst.*, *9*(1). https://doi.org/10.1145/3183367

Nugroho, D. P. D., Hsu, Y., Hartauer, C., & Hartauer, A. (2024). Investigating the interconnection between environmental, social, and governance (esg), and corporate social responsibility (csr) strategies: An examination of the influence on consumer behavior. *Sustainability*, *16*(2), 614.

Park, Y.-n., & Han, S.-L. (2021). The effect of esg activities on corporate image, perceived price fairness, and consumer responses. *korean management review*, *50*(3), 643–664.

Peffers, K., Tuunanen, T., Rothenberger, M., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, *24*(3), 45–77.

Pincheira, M., Donini, E., Giaffreda, R., & Vecchio, M. (2020). A blockchain-based approach to enable remote sensing trusted data. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, *IV-3/W2-2020*, 35–40.

Pöll, E. (2024). Engineering the trust machine. aligning the concept of trust in the context of blockchain applications. *Ethics and Information Technology*, *26*, 1–16. https://doi.org/10.1007/s10676-024-09774-6

Prabawani, B., & Hadi, S. P. (2022). Sustainability indicator: An initial parameter for convenience product. *Jurnal Presipitasi : Media Komunikasi dan Pengembangan Teknik Lingkungan*, *19*(1), 179–189.

Raskin, M. (2016). *The law and legality of smart contracts*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2959166

Rieger, A., Roth, T., Sedlmeir, J., & Fridgen, G. (2022). We need a broader debate on the sustainability of blockchain. *Joule*, *6*(6), 1137–1141. https://doi.org/https://doi.org/10.1016/j.joule.2022.04.013

Rozas, D., Tenorio-Fornés, A., Díaz-Molina, S., & Hassan, S. (2021). When ostrom meets blockchain: Exploring the potentials of blockchain for commons governance. *Sage Open*, *11*(1), 21582440211002526.

Satoshi, N. (2008). *Bitcoin: A peer-to-peer electronic cash system*. https://bitcoin.org/bitcoin.pdf

Savelyev, A. (2017). Contract law 2.0: 'smart' contracts as the beginning of the end of classic contract law. *Information & Communications Technology Law*, *26*(2), 116–134.

Scholl, H. J., Pomeshchikov, R., & Rodríguez Bolívar, M. P. (2020). Early regulations of distributed ledger technology/blockchain providers: A comparative case study. *Proceedings of Hawaii International Conference on System Sciences (HICSS)*.

Sedlmeir, J., Buhl, H., Fridgen, G., & Keller, R. (2020). The energy consumption of blockchain technology: Beyond myth. *Business & Information Systems Engineering*, *62*.

Sedlmeir, J., Lautenschlager, J., Fridgen, G., & Urbach, N. (2022). The transparency challenge of blockchain in organizations. *Electronic Markets*, *32*. https://doi.org/10.1007/s12525-022-00536-0

Smits, M., & Hulstijn, J. (2020). Blockchain applications and institutional trust. *Frontiers in Blockchain*.

Sousa, L. M., Viana, D. C., Neto, A. P. d. L., Castro, Z. R., Aguiar, G. Q. M. d., & Silva, I. R. D. (2023). The evolutions achieved in companies with the implementation of environmental, social and governance: Integrative review. *International Journal of Business, Economics and Management*, *10*(4), 44–53.

Sustainable Finance Strategy. (2021). *Strategy for financing the transition to a sustainable economy*. https://finance.ec.europa.eu/publications/strategy-financing-transition-sustainable-economy_en

Swan, M. (2015). *Blockchain: Blueprint for a new economy*. O'Reilly Media.

Szabo, N. (1994). *Smart contracts*. https://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart.contracts.html

Thoring, K., Mueller, R., & Badke-Schaub, P. (2020). Workshops as a research method: Guidelines for designing and evaluating artifacts through workshops. *Proceedings of Hawaii International Conference on System Sciences (HICSS)*.

Utz, M., Johanning, S., Roth, T., Bruckner, T., & Strüker, J. (2023). From ambivalence to trust: Using blockchain in customer loyalty programs. *International Journal of Information Management*, *68*, 102496.

Vasiu, D. E., & Bratu, R. (2022). An overview on environmental social and governance – esg-topics from the financial markets' perspective. *Management of Sustainable Development*, *14*(2), 76–82.

Webster, J., & Watson, R. (2002). Analyzing the past to prepare for the future: Writing a literature review. *MIS Quarterly*, *26*.

Ying, W., Jia, S., & Du, W. (2018). Digital enablement of blockchain: Evidence from hna group. *International Journal of Information Management*, *39*, 1–4.

Zhan, S. (2023). Esg and corporate performance: A review (M. Yin, P. Wang, & T. Kuang, Eds.). *SHS Web of Conferences*, *169*, 01064.

Ziolkowski, R., Miscione, G., & Schwabe, G. (2020). Decision problems in blockchain governance: Old wine in new bottles or walking in someone else's shoes? *MIS Quarterly*, *37*(2), 316–348.

Ziolo, M., Bak, I., & Spoz, A. (2023). Incorporating esg risk in companies business models: State of research and energy sector case studies. *Energies*, *16*(4), 1809.

X | Research Paper 4 – Blockchain Benchmarking

Full Title:

What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains

Publication venue:

58th Hawaii International Conference on System Sciences (HICSS) 2025

URL:

hdl.handle.net/10993/63520

What Blocks My Blockchain's Throughput? Developing a Generalizable Approach for Identifying Bottlenecks in Permissioned Blockchains

Orestis Papageorgiou
SnT – University of Luxembourg
orestis.papageorgiou@uni.lu

Lasse Börtzler
Karlsruhe Institute of Technology
lasse.boertzler@student.kit.edu

Egor Ermolaev
SnT – University of Luxembourg
egor.ermolaev@uni.lu

Jyoti Kumari
SnT – University of Luxembourg
jyoti.kumari@uni.lu

Johannes Sedlmeir
SnT – University of Luxembourg
johannes.sedlmeir@uni.lu

Abstract

Permissioned blockchains have been proposed for various use cases where a certain degree of decentralization is necessary yet enterprise IT requirements must be met. However, their throughput remains considerably lower than that of established centralized systems. Previous studies that address permissioned blockchains' performance remain blockchain-specific, lacking a generalizable approach for locating and understanding bottlenecks. This paper presents a unified, graphical method for identifying bottlenecks in permissioned blockchains. We augment the DLPS – an open-source benchmarking tool – with graphical evaluation functionalities and use them to identify performance bottlenecks of Hyperledger Fabric and Quorum, two widely used permissioned blockchains with distinct architectural designs. Our work provides researchers and practitioners with a toolkit, guidelines on blockchain performance data analytics, and insights that assist with the bottleneck identification and improvement of permissioned blockchains.

Keywords: Blockchain, Distributed Ledger, Performance Evaluation, Hyperledger Fabric, Quorum

1. Introduction

Since its inception by Nakamoto (2008), blockchain technology has been explored across various industries far beyond its use in the cryptocurrency Bitcoin. Researchers and practitioners have analyzed its potential in a variety of applications where a neutral platform is desirable (Sedlmeir et al., 2022b), such as supply chain management (Queiroz et al., 2020) and the streamlining of cross-organizational workflows (Fridgen et al., 2018). Organizations looking to implement blockchain-based information systems often opt for permissioned blockchains as they restrict participation in consensus, reduce data visibility and latency, and allow for better throughput. However, transitioning projects based on permissioned blockchains from pilot stages to business applications still presents many technical challenges (Toufaily et al., 2021). One major obstacle is the technology's inferior performance compared to established centralized systems, stemming from the resource-intensive nature of replication and consensus (Sedlmeir et al., 2022a). As a result, a substantial part of research focuses on examining the performance characteristics of permissioned blockchains (Fan et al., 2020). Blockchain benchmarking research has primarily examined high-level performance indicators such as throughput, with less emphasis on identifying performance-limiting factors. While assessing throughput is useful for comparing different blockchains and assessing deployment parameters (Guggenberger et al., 2022), it provides limited insights into crucial aspects like node resource utilization, which are usually reported only as aggregated data (Fan et al., 2020), offering a limited understanding of the inner workings of blockchain performance.

This paper addresses this shortcoming by analyzing resource-related metrics of blockchain nodes and their impact on throughput in a systematic and graphical way, providing a general approach that can be used to detect bottlenecks. We survey related work to ground our method and then conduct *exploratory data analysis* (EDA) to determine the performance bottlenecks of Hyperledger Fabric (Fabric) and Quorum, using an extended version of the distributed ledger performance scan (DLPS) (Sedlmeir et al., 2021).

URI: https://hdl.handle.net/10125/109740
978-0-9981331-8-8

2. Background

2.1. Hyperledger Fabric

Fabric has become one of the industry's leading permissioned blockchains (Guggenberger et al., 2022), known for its unique architecture that provides ample opportunities for finetuning performance (Androulaki et al., 2018). In Fabric, nodes are grouped into organizations, and a node can take at least one of the roles of a peer node (*peer*) or an orderer node (*orderer*) (Androulaki et al., 2018). Peers maintain an append-only ledger and a corresponding running aggregate (state), whereas orderers create and broadcast blocks. Fabric relies on the *execute-order-validate* paradigm that involves three phases. In the *Execution Phase*, a client sends a signed transaction proposal to the peers. Peers *simulate* the transaction, i.e., they run the required smart contract ("chaincode") on their current version of the state without updating it. Peers then respond to the client with a signed *endorsement* containing the peer's digital certificate, the transaction's read-write set, and the simulation outcome (Androulaki et al., 2019). During the *Ordering Phase*, and once a client has collected sufficient endorsements according to a chaincode's policy, it packs them into a transaction and forwards it to the ordering service. Orderers use a consensus protocol, such as RAFT (Ongaro et al., 2014), to sort transactions, group them into a batch ("block"), and sign this block without evaluating the transactions' validities. Subsequently, they broadcast the block to a subset of peers (to one *anchor peer* per organization) for validation. In the *Validation Phase*, upon receiving a block – either directly from an orderer or via a fellow peer – a peer validates the included transactions in three steps (Thakkar et al., 2018). First, through parallel verification, *validation system chaincode* (VSCC) ensures transactions have the required endorsements and consistent execution results. Next, valid transactions undergo *multi version concurrency control* (MVCC) – a sequential check of whether the simulations were conducted on compatible ledger versions by comparing the respective read-write sets. Finally, each peer commits the transaction to their local ledger (including a flag for its validity) and – if both checks have passed – updates its state accordingly.

2.2. Quorum

Quorum is another blockchain that has emerged as a significant player among the industry's permissioned blockchains because of its similarity to the public Ethereum blockchain. Unlike Fabric, Quorum relies on the *order-execute* paradigm. It supports four different consensus mechanisms, including RAFT. In the *Ordering Phase*, clients send signed transactions to the nodes, which perform preliminary validations (e.g., correct nonce, syntax, and signatures). Verified transactions are shared with other nodes via a gossip protocol and added to their unconfirmed transaction pools ("mempool"). The RAFT leader sorts and batches valid transactions from its mempool, compiles them into a block, and disseminates the block to follower nodes. Followers attach the received block to their ledger and send acceptance messages to the leader. After receiving acceptance messages from the majority of nodes, the block becomes the new head of the blockchain ("2-phase commit"). During the *Execution Phase*, each node then updates its state deterministically according to each included transaction.

2.3. Distributed Ledger Performance Scan

The DLPS is an open-source, end-to-end benchmarking framework where users can define specifications for various blockchain and client network configurations using a single configuration file (Sedlmeir et al., 2021). We simplified its complex deployment process by dockerizing both the deployment and experiment handlers and extended its graphical capabilities for analyzing performance and resource utilization data.[1] The benchmarking follows a recursive localization of maximum throughput by gradually increasing the request rate to determine the network's maximum throughput. The first run starts by sending (asynchronous) requests at a base rate. The DLPS measures each component's (e.g., nodes, clients) resource utilization, as well as the average request frequency (f_{req}) and response frequency (f_{resp}), of successful transactions by collecting request and response timestamps from all clients. To achieve this, the DLPS leverages system monitoring tools available by the operating system, such as `vmstat`, `mpstat`, `sar`, and `ping`, to gather comprehensive performance metrics on CPU usage, memory, network activity, and I/O operations. If f_{resp} does not deviate from f_{req} by more than a given threshold, the next run increases f_{req}. Otherwise, a certain number of retries is performed. If f_{resp} fails to get close to f_{req} also in the retries, the ramp-up sequence is terminated and the maximum throughput is set to the highest average f_{resp} in any of the runs of the previous ramp-up sequence. After the experiment, the DLPS stores fine-granular data in CSV format for further analysis and generates summary figures for illustrative purposes.

[1] The source code and higher resolution figures are available at: `https://github.com/orepapas/What_Blocks_My_Blockchains_Throughput_Data`.

3. Related work

To gain an overview of the academic literature on bottleneck identification in permissioned blockchains, we conducted a systematic literature review. We used the broad search string *(blockchain OR "distributed ledger technology") AND (performance OR throughput OR latency) AND (benchmarking OR measurement OR evaluation OR analysis)* on Google Scholar, ACM Digital Library, IEEE Xplore, and arXiv. Our search yielded 4,248 results. After manually reviewing titles and abstracts and removing duplicates, 57 publications remained. For these, we performed a full-text screening and excluded articles that did not provide empirical results on blockchain performance. Additionally, we excluded papers that lacked insights into potential bottlenecks, such as papers focusing on comparisons of different configurations or comparisons between different blockchains. For Fabric, we excluded research on v0.6 since it still used the order-execute architecture. We ended up with six relevant publications.

In Fabric v1.0, Androulaki et al. (2018) identify the validation phase and, in particular, the VSCC as a major bottleneck. Thakkar et al. (2018) find three major bottlenecks of v1.0, which are related to the validation phase and were addressed in subsequent Fabric versions. Ruan et al. (2020), using Fabric v1.3, also point to the validation phase as the bottleneck, especially when many unserializable transactions are included in the ledger. Wang et al. (2020) find that in v1.4, the VSCC remains the bottleneck due to limited parallelization. For the same version, Chacko et al. (2021) trace transaction failures to systemic issues, with the validation phase being the bottleneck. Specifically, MVCC read conflicts result in transaction failure, necessitating a return to the execution phase for a new round of endorsements and endorsement policy failures, which significantly slow down the VSCC and transaction processing. We could not find research focusing on identifying bottlenecks in Fabric v2.0 or higher. For Quorum, Mazzoni et al. (2021) posit that a potential bottleneck lies with the node's remote procedure call (RPC) server buffers being capped at 128KB. While this limitation suggests a maximum transaction size of 128KB, it is improbable to be the main bottleneck, as average transaction sizes are only a few hundred bytes.

4. Method and Results

We developed our approach by analyzing the experiments of the papers mentioned in Section 3, using the DLPS to collect data, and EDA (Chatfield, 1986) to identify bottlenecks. We collected extensive data on factors influencing node performance, ensuring a comprehensive overview of resource utilization. After cleaning and validating the data, we used EDA to detect performance irregularities. We then examined the relevant metrics to determine potential root causes. We divided our analysis into two major parts. The first part identifies potential bottlenecks by analyzing the different node resources that can impact blockchain performance. In the second part, we examine the relationship between these candidates and the blockchain's throughput in varying degrees of resolution, such as different time windows and component selections.

Since many enterprise applications have focused on Fabric and Quorum, and Section 3 indicates that their bottleneck analysis is intricate, we detail our bottleneck identification method for both blockchains. We selected an experiment for Fabric v2.0, based on the findings of (Guggenberger et al., 2022) to determine a configuration that seems robust under modifications and extended it to Quorum v23.4. For Fabric, the network configuration comprised 16 clients, 8 peers, and 4 orderers, with four organizations comprised of four clients, two peers and one orderer each. The Quorum configuration consisted of 16 clients and 8 nodes. In both configurations, we selected RAFT as consensus mechanism due to its minimal overhead, as previous publications suggest that consensus is not the bottleneck in this case (Guggenberger et al., 2022; Mazzoni et al., 2021). We conducted the experiment on the AWS cloud platform (Amazon EC2), where each node was configured in an independent EC2 instance allocated with 16 vCPUs, 64 GB of RAM, and 1 Gbps of bandwidth running Ubuntu Server 18.04 LTS (HVM) – a typical configuration for enterprise blockchain nodes.

4.1. Fabric: Resource utilization

The initial analysis focuses on the impact of incremental increases in request rate on the resource utilization of each node. We plot data points for each node every second during a 14-second period within a 20-second experiment, excluding the first and last three seconds to avoid distortions related to the discontinuity of f_{req} (Figure 1). This approach aims to uncover trends and correlations between resource usage and increased f_{req}. We directly see that Fabric peers' central processing unit (CPU) and network utilization exhibit clear plateaus, indicating they could be bottlenecks.

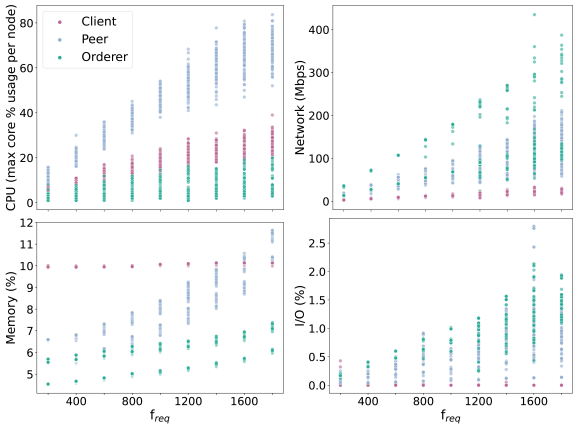

Figure 1: Fabric – Nodes' key resource utilizations.

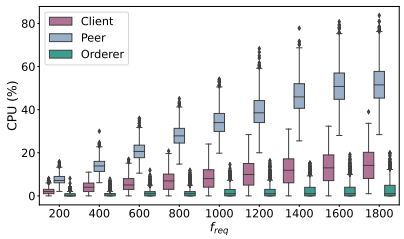

Figure 2: Fabric – CPU utilization of all cores.

4.1.1. CPU The analysis begins by examining the average CPU utilization across a 14-second period. We observe that the linear relationship between the request rate and CPU utilization of peers breaks down at $f_{req} \approx 1600\,s^{-1}$, indicating a potential saturation point or limitation in CPU capacity. Looking into the CPU usage across individual cores (Figure 2) reveals a limited utilization for orderers and clients that is not plateauing at higher request rates. Focusing on peers, we observe a significant variance in CPU utilization, ranging from 25 % to 80 % at higher request rates. This variability suggests an uneven distribution of resources, with some peers bearing a heavier workload than others or an imbalanced allocation of tasks within some peers' cores. First, we investigate the CPU utilization per peer in Figure 3a, which indicates an equitable distribution of computational resources among the peers, with peer 5 falling behind slightly. Analyzing the mean CPU utilization of individual cores for a single peer (peer 0 in this case) in Figure 3b reveals that this also is not the cause of the high fluctuations in CPU usage since all cores show similar utilization. Because of the inconclusive results of both possible explanations, our analysis progresses to evaluate the temporal evolution of CPU usage across individual cores at the highest

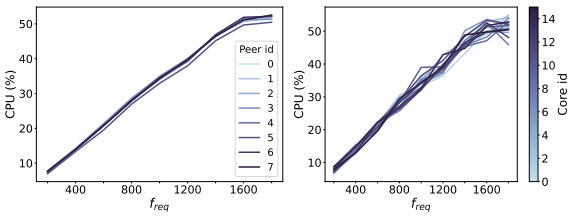

(a) Mean utilization per peer. (b) Mean utilization of peer 0 per core.

Figure 3: Fabric – peer CPU utilizations.

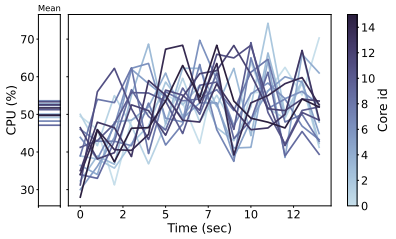

Figure 4: Fabric – Peer 0 CPU utilization for f_{req}=1600 s^{-1}.

f_{req} before the utilization plateaus (Figure 4). This f_{req} represents the peak stress on the network before any performance degradation. Here, we observe that individual cores' CPU utilization can fluctuate as much as 30 % within a single run. We are unable to deduce the reasons behind these fluctuations from CPU utilization data alone. However, as these fluctuations balance out over time according to the narrowly distributed mean CPU utilization, they are likely not the reason for the plateau. It is worth noting that Figure 3 indicates that average CPU usage plateaus at around 50 % across all cores on all peers, an improvement over previous Fabric versions. However, in scenarios utilizing a higher number of vCPUs (16 in our case), Fabric v2.0 still demonstrates a mediocre mean utilization, leaving room for further improvement (see also Thakkar et al. (2018)).

4.1.2. Network The network-related analysis begins by looking into the mean network utilization, distinguishing between inbound and outbound traffic for the different types of nodes. Orderers' traffic does not plateau at high request rates, suggesting the ordering service is not the bottleneck (Figure 5a). Notably, orderer 0 broadcasts a disproportionate amount of traffic compared to the rest, which indicates that orderer 0 is the RAFT leader and, as a result, has the additional task of broadcasting new blocks to each following orderer. Detailed traffic analysis allows its decomposition into individual components, such as the traffic generated by block propagation. According to

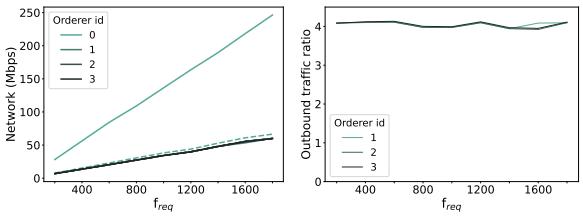

(a) Mean utilization per orderer . (b) Ratio of leader outbound traffic over follower outbound traffic.

Figure 5: Fabric – orderer mean network utilization (outbound traffic as continuous and inbound traffic as dashed lines).

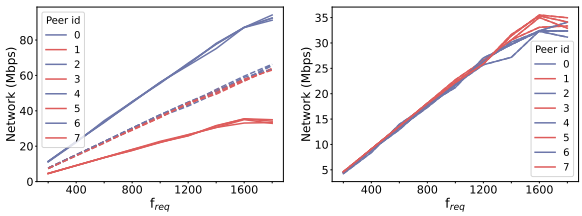

(a) Mean utilization per peer. (b) Outbound traffic per peer .

Figure 6: Fabric – peer mean network utilization (outbound traffic as continuous and inbound traffic as dashed lines).

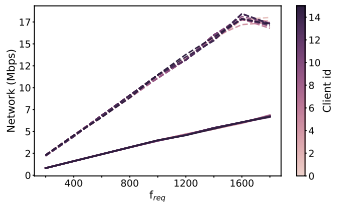

Figure 7: Fabric – Client mean network utilization (outbound traffic as continuous and inbound traffic as dashed lines).

Fabric's architecture, the outbound traffic of follower orderers predominantly consists of sending blocks to the peers. From Figure 5a, and in line with expectations, consistency of inbound and outbound traffic of the followers is observed: followers receive each block once from the leader, and each orderer forwards a received block only once to an anchor peer. On the other hand, when comparing followers' outbound traffic with the leader's outbound traffic (Figure 5b), the ratio consistently stands around four, as the RAFT leader sends each block to each follower (three in our case) and to one peer. Our analysis excludes traffic generated by consensus-related messages, such as appended entries and heartbeat, because it is difficult to distinguish them from block propagation traffic. Nonetheless, the consensus-related messages generate significantly less traffic compared to block dissemination, making the outbound traffic of follower orderers a viable approximation for traffic related to block propagation. Regarding peers (Figure 6), we see that inbound traffic scales linearly with f_{req} for all of them, indicating it is not a bottleneck. Concerning outbound traffic, we observe that it plateaus for some peers while remaining unaffected for others (Figure 6a). Thus, we classify the peers into two main clusters, color-coded as blue and orange. We infer that blue peers act as the gossip leaders of their respective organizations, each with one follower.

The primary distinction in outbound traffic among the two types of peers stems from block propagation between them. To confirm that blue peers are gossip leaders, we deduct the traffic associated with block propagation (as obtained from the ordering service analysis) from the outbound traffic of blue peers (Figure 6b). The resulting traffic is relatively uniform, which is in line with our hypothesis that the blue peers are the gossip leaders. Furthermore, we observe a significant overlap in outbound traffic among peers, particularly at lower request frequencies, with some discrepancies at higher rates. This is expected, as traffic from consensus-related messages increases at higher request rates, making our outbound traffic approximation less accurate. Peer outbound traffic consists of sending requested endorsements and transaction confirmations back to clients. Since the peer outbound traffic plateaus at high f_{req}, it suggests that these components are potential bottlenecks.

We focus on client traffic to pinpoint the specific bottleneck (Figure 7). The inbound traffic of clients, which drops at higher request rates, is generated by the same components as peers' outbound traffic, leaving us with the same potential bottlenecks. Client outbound traffic includes transaction proposals submitted to peers and endorsed transactions sent to the orderers. Since the traffic does not plateau, it suggests that these are not the bottleneck. Overall, the traffic from endorsed transactions sent to the ordering service and blocks sent to peers never plateaus, indicating that the execution and ordering phases are not the bottleneck. Since the endorsements that peers send back to the clients come in between the execution and ordering phases, they are also not the bottleneck. This leaves transaction confirmations from peers to clients, sent after the validation phase, as the likely candidate for a bottleneck. In other words, the validation phase seems to remain the bottleneck even in Fabric v2.0.

4.1.3. Memory & Hard Drive Starting with memory usage, from Figure 1, we observe that all node types exhibit non-plateauing usage levels, suggesting

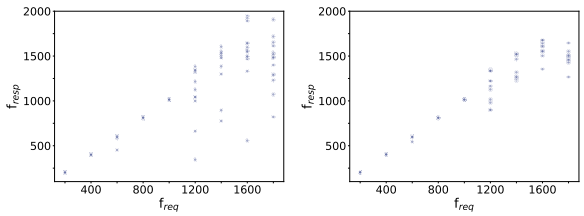

(a) 1s window (b) 3s window

Figure 8: Fabric – Rolling mean of throughput with different window sizes.

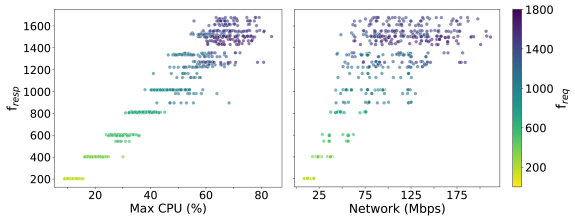

Figure 9: Fabric – Throughput against key resources.

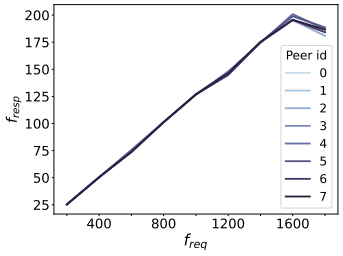

Figure 10: Fabric – Throughput per peer.

that memory constraints are not the bottleneck. For hard drive utilization, only peers' usage appears to plateau. However, peers' I/O operations (such as ledger updates or block processing) occur during or after the validation phase, which does not provide new insights. Additionally, with the peak utilization at approximately 2 %, it is evident that hard drive usage is far from reaching capacity. This indicates that constraints lie within a different component, which in turn limits hard drive utilization. As mentioned in Section 2, the validation phase is comprised of three steps: VSCC, MVCC, and each peer updating its state. As the peers' hard drive utilization is minimal, this leaves only VSCC and MVCC as the potential bottlenecks.

4.2. Fabric: Throughput

We start by gaining an overview of how the request rate affects throughput. This is achieved by plotting the throughput as a rolling mean across two window sizes (Figure 8). Using a one-second window, each data point is plotted individually, revealing significant fluctuations in throughput beyond, f_{req}=1200 s^{-1}, with fluctuations in f_{resp} reaching up to 1000 s^{-1}. We increase the window size to three seconds to observe the overall network performance trend. We selected this interval as it matches the average time it takes for a transaction to be committed to the blockchain under high request rates (see also Guggenberger et al. (2022)). This is significant because queuing effects become prominent at elevated f_{req}, and opting for a shorter time window could underestimate throughput. We see that, on average, the system keeps up with the request rate until reaching approximately f_{req}=1600 s^{-1}.

Next, we explore the correlation between Fabric's throughput and the components and resources highlighted as potential bottlenecks in the first part of the analysis, namely peer CPU utilization and peer network traffic. Figure 9 plots the two resources against the network's throughput using the three-second window. For CPU usage, we observe an initial linear increase with throughput until the characteristic plateau. Additionally, we observe that the instability in CPU utilization starts at around f_{resp}=1200 s^{-1}, which is also the point at which the correlation between f_{req} and f_{resp} starts breaking down (Figure 8a). Consequently, CPU usage correlates more closely with f_{resp} than f_{req}. Similarly, although network traffic increases with throughput, it does so at a lower rate and begins to exhibit instability at f_{resp}=600 s^{-1} already, where the throughput still manages to keep up with f_{req}.

Next, we examine the throughput of individual peers (Figure 10). We observe that every peer contributes similarly to throughput, with minor variations at high f_{req}. Given the significant differences in network traffic between gossip leaders and followers, network traffic is likely not a significant factor in determining throughput. If it were, we would expect noticeable differences in throughput between gossip leaders and followers. Thus, peer CPU utilization appears to be the primary factor behind the leveling off of throughput. The crucial role of peer CPU utilization is in line with our conclusion in Section 4.1 that VSCC and MVCC are the only candidates for bottlenecks in Fabric as both depend on peer CPU. However, given that the checks in MVCC are sequential and Figure 4 shows similar mean core utilization for peers, it is unlikely that MVCC is the bottleneck. This leaves VSCC as the only bottleneck candidate. This hypothesis is further supported by the fact that the validations executed in VSCC are parallelized, and we have noted that the parallelization capacity of Fabric is limited.

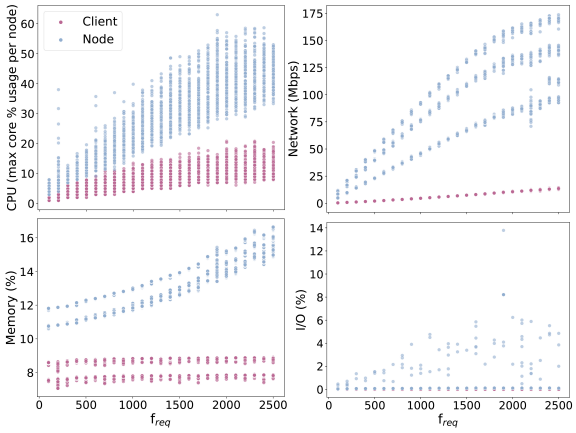

Figure 11: Quorum – Key resource utilization for different request rates f_{req}.

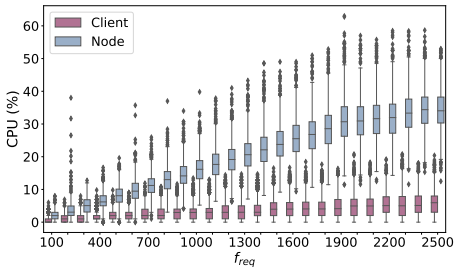

Figure 12: Quorum – CPU utilization of all cores.

4.3. Quorum: Resource utilization

Our benchmarking experiment for Quorum increases the request rate in increments of 100 requests per second. The Quorum bottleneck analysis based on this series of measurements also starts by gaining an overview of the four resources in relation to f_{req} (Figure 11). Similar to the Fabric case, it is apparent that CPU and network utilization are closely correlated with the request rate. Memory and hard drive usage are limited, with I/O operations exhibiting high fluctuations but generally showing less than 1 % usage. Across all resources, client utilization appears to be minimal, and it is either unaffected or grows linearly with the request rate. Therefore, we again focus on the resource utilization of nodes.

4.3.1. CPU We begin the node CPU analysis by examining the utilization across all cores, where – as for the case of Fabric – we note significant fluctuations among the cores of the nodes as request rates increase (Figure 12). Examining the mean node CPU utilization (Figure 13a), we see that nodes can be grouped into three distinct categories. Node 0 (green) exhibits the highest utilization, indicating it is the RAFT leader, which is responsible for additional operations in consensus, such as transaction ordering and block composition. Nodes 1, 2, and 3 (blue) display slightly lower utilization levels, as their role in consensus is receiving and pre-validating transactions. The remaining nodes (orange) exhibit significantly lower CPU usage as they do not have additional responsibilities beyond appending blocks to their local ledgers and applying them to their state.

Examining the utilization per core of node 0 (Figure 13b), we observe that one core (core 15) bears a higher workload across all request rates. This is because, except pre-validation, all other tasks of the leader are executed sequentially, leading to a disproportionate strain on one core. Following this observation, we examine the temporal evolution of CPU usage across individual cores at f_{req}=2300 s^{-1} (Figure 14). Here, we see fluctuations by as much as 20 % for the RAFT leader and followers, with disparities of up to 10 % between individual cores when excluding core 15 for node 0. The main difference between the mean utilization of node 0 and nodes 1, 2, and 3 comes only from core 15. These results suggest that parallel processing in Quorum is even more limited than in Fabric, with average core usage not exceeding 40 % and specific leader tasks overburdening one core. The CPU utilization of nodes 0, 1, 2, and 3 reaches a plateau at f_{req}=2400 s^{-1}. Additionally, we notice the first break in the linear relationship between CPU usage and the request rate at f_{req}=1800 s^{-1}. While this does not necessarily indicate a bottleneck, it could provide clues for identifying factors behind the decline in CPU utilization. From Figure 13a, we infer that node 0 and nodes 1, 2, and 3 have similar behavior after reaching a plateau. Therefore, transaction ordering and block building, the main unique operations performed by the leader, are likely not the bottlenecks. Examining Figure 13b, we see a rapid decline in the utilization of core 15 at high request rates. Considering that the remaining sequential operations, such as transaction ordering and block propagation, require minimal CPU resources, this sharp drop cannot be justified, and it appears that another component limits CPU utilization.

4.3.2. Network Focusing on mean network utilization (Figure 15), we can categorize the nodes into three groups. Due to the complexity of Quorum network traffic, we cannot accurately decompose it into individual components, so we rely on the architectural design to identify bottlenecks. Starting with the blue nodes, outbound traffic levels off at f_{req}=2400 s^{-1} while

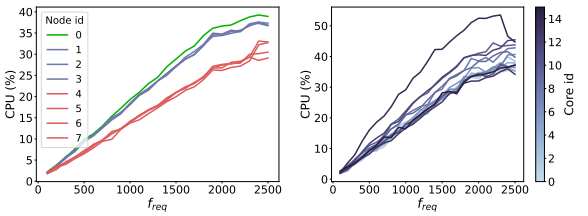

(a) Mean utilization per Node. (b) Mean utilization of Node 0 per core.

Figure 13: Quorum – Mean CPU utilization.

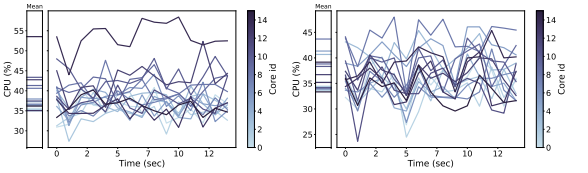

(a) Node 0. (b) Node 1.

Figure 14: Quorum – CPU utilization for f_{req}=2300 s^{-1}.

signs of plateauing in their inbound traffic appear at f_{req}=1800 s^{-1}. The inbound traffic primarily consists of transactions received either directly from clients or through gossip and blocks from the leader, while outbound traffic originates from the dissemination of pre-validated transactions. This suggests that these operations are the limiting factors at their respective request rates. While we see similar patterns for the orange nodes, the patterns of the leader are essentially the opposite of those of other nodes. The inbound traffic is comprised of all the transactions that are broadcasted to the network and reach the leader through gossip or directly from the clients, and plateaus at f_{req}=2400 s^{-1}. The outbound traffic involves mainly block dissemination to the other nodes and plateaus at f_{req}=1800 s^{-1}. Since their inbound traffic plateaus at f_{req}=2400 s^{-1}, the leader likely keeps receiving the transactions from the blue nodes normally up until that point, leaving us only with block dissemination as the main bottleneck at f_{req}=1800 s^{-1} and transaction propagation as the main issue for f_{req}=2400 s^{-1}.

Considering our CPU utilization findings, we posit that the rapid drop in CPU utilization for core 15 at f_{req}=2400 s^{-1} (Figure 13b) is caused by the leader not receiving enough transactions from the other nodes. At f_{req}=1800 s^{-1}, the decline could be attributed to either the block propagation or one of the processes preceding it, such as transaction pre-validation and adding the block to the chain. Considering that appending the block to a node's local ledger is not CPU intensive,

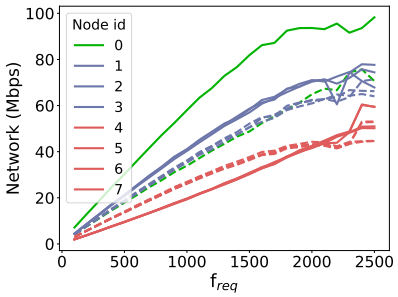

Figure 15: Quorum – Node mean network utilization (outbound traffic as continuous and inbound traffic as dashed lines).

the primary issues likely lie with block propagation or pre-validation. Using the same line of reasoning for f_{req}=2400 s^{-1}, the bottleneck appears to be either transaction propagation or pre-validation, as it is the only operation that precedes propagation.

4.3.3. Memory & Hard Drive Starting with memory utilization, we see consistent behavior across all nodes, with no signs of plateauing (Figure 11). As such, memory does not seem to contribute to performance degradation. Despite observing significant peaks in hard drive utilization (Figure 11), the mean utilization remains mostly below 1 %, indicating it is not a constraining factor. The large fluctuations are probably related to the writing of the block into each node's database, but since it is improbable that it leads to a bottleneck, we do not examine it further.

4.4. Quorum: Throughput

Analyzing the correlation between f_{resp} and f_{req}, we see that even for the one-second window size, throughput remains stable but starts to exhibit higher fluctuations at f_{req}=1800 s^{-1}. Examining the three-second window, throughput plateaus at around f_{resp}=2100 s^{-1}, significantly lower than the maximum request rate. This suggests that the overall performance of the blockchain started to decline before reaching the highest request rate. This indicates that the performance degradation noted at f_{req}=1800 s^{-1} for both CPU and network utilization may be more critical in identifying the bottleneck.

Examining throughput against the CPU and network utilization (Figure 17), we see similar patterns. Both resources keep up with throughput initially and experience higher fluctuations around f_{resp}=1800 s^{-1}. Beyond this point, differences in throughput between request rates become less pronounced (even close

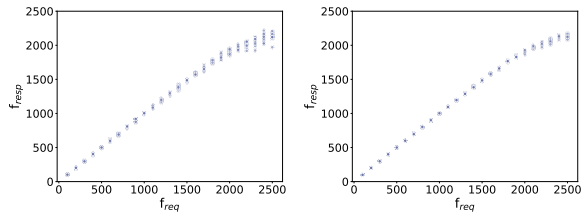

(a) 1s window (b) 3s window

Figure 16: Quorum – Rolling mean of throughput with different window sizes.

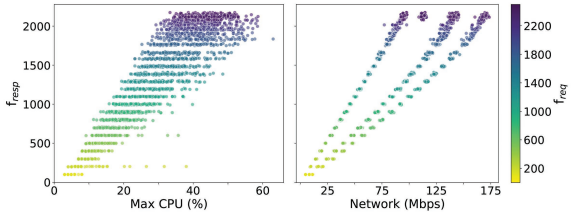

Figure 17: Quorum – Throughput against key resources.

to indistinguishable near f_{resp}=2100 s^{-1}). Given these similar patterns, we search for a potential common source behind the blockchain's performance degradation. Since block and validated transaction propagation are network-related, if they were the bottleneck, they would mainly impact network traffic, and, as a result, they are less likely to be the bottleneck.

This observation leaves only transaction pre-validation as the possible constraining factor, which impacts both CPU and network traffic when nodes are not receiving enough transactions. According to Figure 11, network utilization of clients increases linearly with the request rate, which suggests they send the proper number of transactions to the nodes. Examining further the interaction of nodes with the incoming transactions, we look into the number of rejected transactions (Figure 18a). Rejected transactions are those that nodes decline to propagate, leading to clients receiving nearly immediate (within 50 ms) notifications of transaction failure. Initially, the count of rejected transactions is minimal but begins to surge at f_{req}=1800 s^{-1}, culminating in approximately 7000 rejections by f_{req}=2500 s^{-1}. This corresponds to a rejection rate of 20 %, prompting further investigation into the underlying causes.

In Quorum, transactions are categorized as either executable or non-executable. Executable transactions can be immediately included in a block, while non-executable transactions are out of nonce order and must wait for preceding transactions with lower nonce to execute first. Clients are limited to keeping

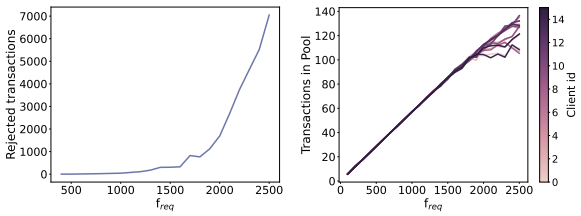

(a) Number of rejected transactions. (b) Number of transactions in the pool per client.

Figure 18: Quorum – Transaction metrics.

16 executable and 500 non-executable transactions in the pool at any given time. Figure 18b illustrates that at f_{req}=1800 s^{-1}, the number of transactions in the pool from six clients begins to exhibit instability, diverging from the previous linear relationship with the request rate, and most clients diverge at f_{req}=2400 s^{-1}. This pattern, along with the fact that the number of client transactions in the pool never exceeds 500, suggests that clients are reaching the limit of executable transactions in the pool, leading to more rejections. This hypothesis is further supported by the number of rejected transactions, which increases sharply at f_{req}=1800,s^{-1}, coinciding with the initial reduction in CPU and network utilization and the subsequent plateau at f_{req}=2400,s^{-1}. This indicates that the main bottleneck for Quorum is the clients' limit of executable transactions in the pool.

5. Conclusion

This paper introduces a general illustrative method for blockchain bottleneck identification, demonstrated through an analysis of a 12-node Fabric network and an 8-node Quorum network. Our method leverages EDA to analyze blockchain performance metrics, highlighting their specific characteristics and bottlenecks. By employing a combination of proportional analysis and the study of plateau-shaped trends in resource utilization versus transaction metrics, we uncover performance anomalies. This approach allows us to narrow down the reasons for bottlenecks by comparing the correlation between data trends, the request rates (f_{req}), and response rates (f_{resp}).

For Fabric, we identify the validation phase as the main bottleneck, even for v2.0, with VSCC being the most likely component behind the bottleneck. We were also able to showcase the average moderate degree of parallelization within Fabric, which leaves ample room for improvement. In this sense, our findings align with previous studies (see Section 3). For Quorum, we posit that the bottleneck stems from the restriction

on the number of executable transactions a client can have in the transaction pool, leading to many rejections. Additionally, our findings illustrate Quorum's relatively limited capacity for parallel processing.

Acknowledgements

Funded by the Luxembourg National Research Fund (FNR), grant reference 16326754 and NCER22/IS/16570468/NCER-FT, and by PayPal, PEARL grant reference 13342933/Gilbert Fridgen. To meet grant obligations, the author has applied a Creative Commons Attribution 4.0 International (CC BY 4.0) license to any Author Accepted Manuscript version arising from this submission.

References

Androulaki, Elli et al. (2018). "Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains". In: *Proceedings of the 13th EuroSys Conference*. ACM. DOI: 10.1145/3190508.3190538.

Androulaki, Elli et al. (2019). "Endorsement in Hyperledger Fabric". In: *Proceedings of the International Conference on Blockchain*. IEEE. DOI: 10.1109/Blockchain.2019.00077.

Chacko, Jeeta Ann et al. (2021). "Why Do My Blockchain Transactions Fail? A Study of Hyperledger Fabric". In: *Proceedings of the International Conference on Management of Data*. ACM. DOI: 10.1145/3448016.3452823.

Chatfield, Chris (1986). "Exploratory Data Analysis". In: *European Journal of Operational Research*. DOI: 10.1016/0377-2217(86)90209-2.

Fan, Caixiang et al. (2020). "Performance Evaluation of Blockchain Systems: A Systematic Survey". In: *IEEE Access* 8. DOI: 10.1109/ACCESS.2020.3006078.

Fridgen, Gilbert et al. (2018). "Cross-Organizational Workflow Management Using Blockchain Technology – Towards Applicability, Auditability, and Automation". In: *Proceedings of the 51st Hawaii International Conference on System Sciences*. DOI: 10.24251/hicss.2018.444.

Guggenberger, Tobias et al. (2022). "An In-Depth Investigation of the Performance Characteristics of Hyperledger Fabric". In: *Computers & Industrial Engineering* 173. DOI: 10.1016/j.cie.2022.108716.

Mazzoni, Marco et al. (2021). "Performance Evaluation of Permissioned Blockchains for Financial Applications: The ConsenSys Quorum Case Study". en. In: *Blockchain: Research and Applications*. DOI: 10.1016/j.bcra.2021.100026.

Nakamoto, Satoshi (2008). *Bitcoin: A Peer-to-Peer Electronic Cash System*. URL: bitcoin.org/bitcoin.pdf.

Ongaro, Diego et al. (2014). "In Search of an Understandable Consensus Algorithm". In: *USENIX Annual Technical Conference*. URL: raft.github.io/raft.pdf.

Queiroz, Maciel M et al. (2020). "Blockchain and Supply Chain Management Integration: A Systematic Review of the Literature". In: *Supply Chain Management: An International Journal* 25. DOI: 10.1108/SCM-03-2018-0143.

Ruan, Pingcheng et al. (2020). "A Transactional Perspective on Execute-Order-Validate Blockchains". In: *Proceedings of the ACM SIGMOD International Conference on Management of Data*. DOI: 10.1145/3318464.3389693.

Sedlmeir, Johannes et al. (2021). "The DLPS: A New Framework for Benchmarking Blockchains". In: *54th Hawaii International Conference on System Sciences*. URL: hdl.handle.net/10993/45620.

Sedlmeir, Johannes et al. (2022a). "A Serverless Distributed Ledger for Enterprises". In: *55th Hawaii International Conference on System Sciences*. URL: http://hdl.handle.net/10125/80228.

Sedlmeir, Johannes et al. (2022b). "The Transparency Challenge of Blockchain in Organizations". In: *Electronic Markets* 32 (3). DOI: 10.1007/s12525-022-00536-0.

Thakkar, Parth et al. (2018). "Performance Benchmarking and Optimizing Hyperledger Fabric Blockchain Platform". In: *Proceedings of the 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*. IEEE. DOI: 10.1109/MASCOTS.2018.00034.

Toufaily, Elissar et al. (2021). "A Framework of Blockchain Technology Adoption: An Investigation of Challenges and Expected Value". In: *Information & Management* 58.3. DOI: 10.1016/j.im.2021.103444.

Wang, Canhui et al. (2020). "Performance Characterization and Bottleneck Analysis of Hyperledger Fabric". en. In: *Proceedings of the 40th International Conference on Distributed Computing Systems*. DOI: 10.1109/ICDCS47774.2020.00165.

XI | Research Paper 5 – ESG KPI Reporting System II

Full Title:

Blockchain-based Reporting System for Trust in Environmental Social Governance

Publication venue:

Under review at a peer-reviewed journal (as of 24 November 2025).

URL:

hdl.handle.net/10993/66729

Blockchain-based Reporting System for Trust in Environmental Social Governance

Egor Ermolaev[1*], Evgenia Yvonni Tseloni[1], Tamara Roth[2], Gilbert Fridgen[1]

[1*]Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, 29 Avenue J.F Kennedy, Luxembourg, 1855, Luxembourg.

[2]Sam M. Walton College of Business, University of Arkansas, 1 University of Arkansas, Fayetteville, 72701, Arkansas, United States of America.

Abstract

Towards growing efforts to promote sustainable business practices, companies are required to publicly disclose their environmental, social, and governance (ESG) key performance indicators (KPIs). However, current reporting formats often lack the transparency and verifiability needed to meet legal requirements and build trust with responsible investors. To address these issues, we leverage blockchain technology, whose inherent properties of immutability and decentralization can foster institution-based trust in the context of ESG reporting and mitigate accusations of greenwashing. Following the design science research method, we conducted two design iterations to develop a system architecture and prototype. These iterations were guided by six design objectives & fifteen design requirements. The blockchain-based ESG KPI reporting system artefact was then evaluated through twelve semi-structured interviews with ESG experts from industry and academia. As a contribution to the theory of trust, we distilled a set of three design principles.

Keywords: ESG Reporting, Greenwashing, Green Information Systems, Blockchain, Trust

1 Introduction

In recent years, industries have shown increasing interest in adopting governance models to improve the quality of corporate environmental sustainability (AbuRaya 2017; Zhan 2023). This growing focus on sustainability governance has far-reaching implications, both for internal practices and external perception. Specifically, corporate image, perceived price fairness, and purchasing behavior can be shaped by how accessible and transparent a company's environmental, social, and governance (ESG) practices are (Park and Han 2021). Moreover, the formalization of ESG strategies and goals has given rise to a new investment paradigm: responsible investment, with global ESG assets that is projected to surpass the $40 trillion mark by 2030 (Bloomberg 2024). Reflecting this broader shift, companies are increasingly highlighting their annual reports by disclosing ESG-related data (Ho 2023), indicating sustainable and responsible practices and products to stakeholders such as consumers and business partners (Kim et al. 2021; European Federation of Financial Analyst Societies 2010; Nugroho et al. 2024).

Notably, this shift has been especially pronounced within the financial industry, where the integration of ESG considerations has moved beyond reputational concerns to become a core component of investment strategy, risk assessment and regulatory compliance (Cantero-Saiz et al. 2024). More specifically, institutional investors, asset managers, insurers and pension funds play a critical role in advancing sustainability objectives through their capital allocation and shareholder engagement practices, redefining the environmental, social, and governance (ESG) performance as a material factor in financial decision-making to include long-term environmental and social risks (Mandas et al. 2023; Amel-Zadeh and Serafeim 2018; Ströher et al. 2025). Following this growing trend of ESG disclosure and transparency, legislators around the world are moving toward more standardized ESG reporting frameworks to ensure consistency, comparability, and accountability in ESG reporting (Department for Business and Trade, Government of the United Kingdom 2024; Legislative State Bureau, State of California 2023; Australian Accounting Standards Board, Australian Government 2024). In particular, the European Union (EU) has established a comprehensive package of measures aimed at preventing regulatory fragmentation across member states and promoting a unified internal market (European Parliament 2024b). Recent strategies include the FinTech Action Plan (European Commission 2018) and the Strategy for Financing the Transition to a Sustainable Economy (European Commission 2021), extending the scope of the new frameworks to both financial (Sustainable Finance Disclosure Regulation 2019/2088 (2019)) and non-financial market participants (Corporate Sustainability Reporting Directive 2022/2464 (2022)), with the aim of encompassing a broad spectrum of businesses. Financial institutions are now required to disclose how they integrate ESG risks into investment decisions and advisory practices, promoting comparability of financial products and reducing greenwashing (Driessen 2021). At the same time, non-financial companies must report on the sustainability of their operations, using the double materiality approach to capture both financial and societal impacts. In light of these developments, companies are well-advised to systematically improve their ESG reporting practices (Jean and Grant 2022) since the disclosure of ESG-related data serves a dual purpose: it not only helps

market participants compare the sustainability of different financial products (Duran and Tierney 2023), but also addresses issues of greenwashing, including marketing tactics (Chueca Vergara and Ferruz Agudo 2021) aimed at deceiving stakeholders about the environmental impact of an organization or promoting unsustainable products or services as sustainable (Duran and Tierney 2023; Delmas and Burbano 2011).

Despite growing regulatory pressure and corporate efforts to enhance ESG transparency, the practical implementation and reporting of ESG strategies remain fraught with challenges (Hang et al. 2023). A key concern is the lack of comprehensive and reliable metrics, which makes it difficult for companies to effectively assess and communicate their ESG performance (Vasiu and Bratu 2022). This challenge is particularly critical as ESG metrics are becoming more central to strategic decision-making (Awan 2011; Zhan 2023). Although the need for better measurement and certification of green business models is well recognized (Lindgren et al. 2021; Prabawani and Hadi 2022), progress toward sustainable corporate governance in finance remains slow, as aligning investors' sustainability preferences with organization' practices remains a critical challenge, and companies often lack the tools needed to translate ESG into practice (Driessen 2021; Dye et al. 2021).

Moreover, trust, particularly institution-based trust, plays a critical role in establishing the credibility of ESG reporting. While third-party assurance of ESG reports can help mitigate stakeholder distrust (Shen et al. 2017), concerns about greenwashing persist and continue to undermine confidence in reported ESG (Cruz and Matos 2023). Existing digital tools for ESG reporting often fail to address these deeper concerns about institutional trust, as they rarely provide mechanisms to verify the disclosed information. This is further complicated by literature remaining inconclusive on whether specific governance mechanisms effectively promote transparency and accountability (García-Sánchez et al. 2022). Traditional ESG reporting approaches – often static PDFs combined with centralized audit processes – offer limited verifiability and are susceptible to manipulation (Yu 2024). Despite growing efforts to improve assurance practices (Liu et al. 2024; Pizzi et al. 2022; Gu et al. 2023), current systems remain opaque, making it difficult to build the level of transparency required for institutional trust. These limitations underscore the urgent need for new reporting systems that embed trust by design, ideally through verifiability and tamper-resistance of data that supports transparent and auditable disclosures.

In response to these limitations, there is growing interest in developing ESG reporting systems that move beyond static toward more dynamic, trustworthy, and auditable disclosures (Aggarwal et al. 2023; Liu et al. 2024). Such systems aim to integrate verifiability (Yu 2024) and auditability (Miranda et al. 2023) directly into the reporting process, providing the foundation for improved institutional trust. Emerging approaches explore the use of distributed ledger technologies, such as blockchain, to create tamper-resistant records and enable traceability and transparency of relevant ESG data (Gramlich et al. 2024; Liu et al. 2021, 2023). While these approaches make progress in addressing greenwashing and institutional trust, they still struggle with a lack of real-time transparency, poor data quality, limited customization for stakeholder needs, low interoperability across systems, and inadequate comparability despite existing standards (Gramlich et al. 2024). Specifically, issues of data integrity,

system-wide functionality, and compliance by design are often insufficiently addressed. Our study thus aims to integrate institutional trust-building elements throughout the design process by focusing on the following research questions.

RQ 1: How can trust in blockchain technology improve institution-based trust in ESG reporting while avoiding greenwashing?

RQ 2: How can an ESG reporting system be designed on blockchain?

To answer these questions, we follow the Design Science Research (DSR) approach as proposed by Peffers et al. (2007). This approach enables us to identify the design objectives (DOs) for developing an ESG key performance indicator (KPI) reporting system artefact, as well as the design requirements (DRs) needed to enhance trust in technology and, by extension, institutional trust. Our paper is structured according to the framework suggested by Gregor and Hevner (2013). We began with a literature review, presented in the *Theoretical Background* section, which explores the convergence of ESG reporting and the burgeoning discussion of blockchain in IS, particularly within the context of the EU legal framework and institutional-based trust theory. This is followed by a *Research Approach* section that details our methodological process. Specifically, we conducted a systematic literature review of current ESG reporting systems based on Webster and Watson (2002), which was complemented by our participation in a hackathon. Both activities helped us identify an initial set of DOs & DRs. During the hackathon, we also designed a system architecture and developed a first prototype of the ESG reporting artefact. To further refine our artefact and gain deeper insight into industry needs, we engaged in multiple evaluation and development loops with a diverse set of experts across. These included interdisciplinary jurors at the hackathon, subject-matter reviewers following the submission of a conference paper, participants at United Nations (UN) Forum on Sustainability and Digital Transformation workshop, and twelve academic and industry professionals through semi-structured interviews. The evolution of DOs & DRs across these iterative cycles is presented in the *ESG KPI Reporting System section*, along with the architecture, implementation and the evaluation. The paper continues with the *Discussion* of our findings, where we analyze practical and theoretical implications by formalizing the prescriptive design knowledge in the form of design principles (DPs), thereby contributing to a nascent design theory (Gregor and Hevner 2013) on using blockchain for facilitating and enhancing trust in ESG reporting, acknowledges limitations, and suggests potential areas for future research. Finally, the paper ends with a *Conclusion* section that summarizes key insights.

2 Theoretical Background

2.1 An industry perspective on ESG reporting

The concept of ESG has gained increasing importance in the business world, as companies and investors aim to prioritize environmental, social, and governance factors to drive meaningful change and enhance performance (Vasiu and Bratu 2022). The term ESG encompasses three dimensions of corporate responsibility. First, the environmental dimension focuses on the reduction of a company's negative impact on the environment, such as minimizing carbon emissions and resource consumption.

Second, the social dimension emphasizes a company's influence on its stakeholders, including employees, customers, suppliers, and the broader community. Finally, the governance dimension relates to a company's leadership, organizational policies, and internal structures that ensure accountability and ethical decision-making (Li et al. 2023; Cabaleiro-Cerviño and Mendi 2024; Liu et al. 2024). An ESG report captures a company's performance across these three dimensions. In the financial industry, the integration of ESG not only fosters sustainable and responsible investment practices, products, and services, but it also opens opportunities for innovation and long-term value creation (Mandas et al. 2023). That is, ESG has the capacity to shape financial decision-making and risk assessment, requiring coordinated action across the entire FinTech ecosystem. This includes integrating ESG into investment and lending decisions, systematically tracking progress toward ESG objectives through both voluntary and regulatory reporting frameworks, and dynamically adjusting financial portfolios accordingly. Moreover, assessing climate, environmental, and transition risks at both portfolio and organizational levels has become critical for managing the financial and reputational risks associated with sustainable finance (Galeone et al. 2024).

Despite growing interest in sustainability and the increasing integration of ESG into financial decision-making (Gharpure 2025), research focusing specifically on ESG reporting within banking, financial services, and FinTech remains limited (Galeone et al. 2024; Buallay et al. 2023). This, however, is required as financial institutions play a pivotal role in ensuring global financial stability. Unsustainable practices, such as weak governance, poor environmental management, or social inequality, can lead to systemic risks, asset devaluation, and market volatility, threatening financial market resilience (Gharpure 2025). As key actors in capital allocation and risk assessment, financial institutions are under growing pressure to disclose their sustainability performance and improve their governance frameworks (Buallay et al. 2023). Investors, in turn, increasingly factor ESG into their decisions, seeking alignment with both financial (i.e., risk and return) and ethical (i.e., own values and social values) objectives (Gharpure 2025; Galeone et al. 2024). To meet these needs and support both financial system sustainability and stability, we require transparent and verifiable ESG reporting.

ESG reports integrate both qualitative and quantitative evaluations to present relevant information in an accessible and comprehensive format (AbuRaya 2017). These reports offer a transparent overview of ESG strategies and practices, enabling a broad range of stakeholders, including internal and external decision-makers, and shareholders (Keynes 1936), to make more informed assessments and comparisons across firms and industries (European Federation of Financial Analyst Societies 2010). Besides enhancing transparency and comparability, ESG reporting also serves two fundamental functions as emphasized by Ho (2023): a promotional and a protective function. The promotional function highlights the positive societal impact of sustainable investments and encourages further investments in sustainable initiatives (Chopra et al. 2024). Empirical studies suggests that ESG integration can yield substantial benefits such as cost savings, improved operational efficiency, and long-term resilience (Sousa et al. 2023). However, these findings are not consistent. For example, Buallay et al. (2023) report a significant negative relationship between ESG scores and a company's

operational, financial, and market performance across their full sample, underscoring the need for more nuanced insights into ESG's effects in the financial industry. The protective function of ESG reporting, in turn, focuses on mitigating investment risks, particularly reputational risks, arising from non-compliance with publicly promoted sustainability commitments (European Supervisory Authorities 2024c). To accomplish both of these functions, companies must more proactively establish or refine internal standards, accountability frameworks, and processes that drive systematic improvements in ESG performance (Jean and Grant 2022). A key instrument in this endeavor is the use of ESG KPIs, which offer measurable benchmarks within ESG reporting systems (European Federation of Financial Analyst Societies 2010). These ESG KPIs vary according to strategic focus and company goals. For instance, environmental KPIs may encompass metrics such as energy use, carbon emissions, or circular economy initiatives. Social KPIs often address equity, labor relations, or human rights (European Commission 2024a), while governance KPIs typically focus on transparency in management structures, and risk management practices (Cabaleiro-Cerviño and Mendi 2024).

2.2 Reduced trust in ESG reporting due to greenwashing practices

Despite regulatory progress and growing efforts to standardize ESG performance reporting well-defined KPIs, many companies continue to engage in misleading or deceptive representations of their sustainable activities (Jones 2019; Aggarwal and Kadyan 2011). This practice, widely known as "greenwashing", was first coined in 1986 by activist Jay Westerveld and has since evolved into a complex, multifaceted phenomenon (Lyon and Montgomery 2015) demanding closer scrutiny and more effective countermeasures (De Freitas Netto et al. 2020). At the organizational level, greenwashing typically involves attempts to mislead consumers, investors, or other stakeholders about a company's environmental or social practices (Delmas and Burbano 2011; Duran and Tierney 2023). From a financial regulatory perspective, the EU Taxonomy defines the "greenwashing" as "the practice of gaining an unfair competitive advantage by marketing a financial product as environmentally friendly, when in fact basic environmental standards have not been met" (Recital 11 of the EU Taxonomy Regulation (EU Taxonomy), Directorate-General for Financial Stability, Financial Services and Capital Markets Union, European Comission (2020)). De Freitas Netto et al. (2020) identify five common organizational-level forms of greenwashing: dirty business, ad bluster, political spin, "it is the law, stupid", and fuzzy reporting. Of these, fuzzy reporting is particularly relevant to ESG reporting, as it exploits the largely one-directional and often unverified nature of sustainability reports to obscure or distort the truth. At the product or service level, greenwashing may involve overstating environmental benefits without credible third-party certification or offering vague, misleading, or irrelevant claims. Tactics include concealing trade-offs, omitting critical information, using deceptive labels, or applying the "lesser of evils" strategy to divert attention from more harmful impacts (Delmas and Burbano 2011; Duran and Tierney 2023; Terrachoice 2011). In practice, greenwashing of ESG often takes the form of ambiguous language, selective disclosure, or exaggerated claims about the scope or

success of a company's sustainability strategies (Chueca Vergara and Ferruz Agudo 2021).

These deceptive practices have far-reaching implications for trust. Trust is fundamental in relationship involving risk or uncertainty, especially in financial and investment contexts, where stakeholders must rely on the integrity and good intentions of companies (McKnight et al. 2009). Greenwashing undermines trust in ESG reporting and erodes the perceived integrity and benevolence of the reporting company, violating the principle of good faith and damaging institution-based trust (McKnight and Chervany 2001). Institution-based trust refers to an individual's sense of security in impersonal systems, such as companies and governments, where stakes are high and potential risk is involved (McKnight et al. 1998; Goo and Nam 2007). It is grounded in two sub-constructs: structural assurance, or belief in the protective capacity of formal systems (e.g., laws, contracts, and institutional norms), and situational normality, the assumption that the context is orderly, predictable, and favorable (McKnight et al. 2009; Goo and Huang 2008).

Greenwashing erodes institution-based trust by creating perceptions that the organization may intend harm (malevolence) and is likely to be dishonest or provide false information (deceit) (Moody et al. 2013, 2017). First, it weakens structural assurance by revealing the shortcomings of current normative frameworks, which are often non-binding or poorly enforced (Foley et al. 2024). Second, it disrupts situational normality by undermining the expectation that companies will follow through on their stated ESG commitments. Consequently, both trust in ESG reporting and the broader institutions that govern it begins to crumble (Foley et al. 2024; McKnight and Chervany 2001). The result is not only reputational damage for companies engaged in (McKnight et al. 2017). To address these challenges and restore trust, technological solutions such as blockchain, cloud computing, Internet of things (IoT), Machine Learning (ML), and decentralized digital identities (Kshetri 2021; Di Vaio and Varriale 2020; Asif et al. 2023; Truant et al. 2023; Gramlich et al. 2024) are gaining traction. These technologies offer capabilities for ensuring the verifiability, integrity, and traceability of ESG data (Chen et al. 2023). By embedding trust into the reporting structure itself, they can improve data and reporting accuracy, prevent tampering, and enhancing compliance with evolving regulatory standards.

2.3 Building trust with blockchain technology

Blockchain technology provides a transparent and tamper-resistant infrastructure for recording, verifying, and automating transactions (Macchiavello and Siri 2020). Since its introduction in 2008 (Nakamoto 2008), it has been widely recognized for its ability to reconfigure trust, shifting reliance from centralized authorities and interpersonal assurances to technologically enforced guarantees (Casey and Vigna 2018; De Filippi et al. 2020; Utz et al. 2023; Ziolkowski et al. 2020). Technically, blockchains are distributed databases that store transactional data chronologically across synchronized nodes in a blockchain network (Swan 2015; Jones 2019). Each basic ordering element, commonly referred to as "block", is cryptographically linked to its predecessor using hash functions, creating a chain that ensures the immutability and traceability of

recorded data (Scholl et al. 2020; Kshetri 2021; Pincheira et al. 2020). To ensure consistency between nodes and include transactions in the next block, blockchains rely consensus protocols, most notably Proof of Work proof-of-work (PoW) and Proof of Stake proof-of-stake (PoS). In PoW, miners compete to solve a cryptographic puzzle – secured by spent energy. In PoS validators are chosen to create to create blocks based on the amount of coins they have locked up – secured by staked capital. In both cases, decision-making power is tied to a scarce resource, thereby securing the network (Sedlmeir et al. 2020). Although data replication introduces some inefficiencies across all blockchains, the energy demands of PoW systems remain a particular concern. Still, any evaluation of blockchain's energy use must be context specific, considering the specific context and the type of blockchain required (Rieger et al. 2022).

Beyond data security and transparency, many blockchains also support the execution of smart contracts, i.e., custom deterministic programs that automatically enforce predefined rules and conditions (Szabo 1994; Rozas et al. 2021; Gramlich et al. 2024). These smart contracts are not autonomous; rather they must be activated by a transaction initiated by an external actor, using externally owned accounts externally owned account (EOA) (such as a human user or an automated bot), or another smart contract. That is, the transaction calls a specific function within the contract through the contract's application binary interface (ABI), and if conditions are met, the execution and results are recorded immutably on the blockchain (Raskin 2016; Rozas et al. 2021). In the context of ESG reporting, this capability can provide structural assurance by preventing the retroactive alteration of sustainability data and ensuring consistent rule application (Mendling et al. 2018; Savelyev 2017). Through the integration of data integrity, automated enforcement of predefined conditions, and transparency, blockchain presents a promising approach for combating greenwashing and enhancing accountability, even in the absence of robust regulatory frameworks.

Existing literature has explored blockchain-based trust mechanisms in many industries and contexts, often addressing only parts of the broader challenge. For example, Balzani and Corsi (2022) examine how Hyperledger Fabric can improve the traceability of renewable energy and greenhouse gas (GHG) emissions by disintermediating information flows. While their focus is on post-recording immutability, they do not address the authenticity of initial data entries. Similarly, Mugurusi and Ahishakiye (2022) investigate a private blockchain-based platform for tracing responsibly sourced cobalt in the Democratic Republic of Congo, noting the "Garbage-In-Garbage-Out" problem and advocating for trusted on-site data verifiers. Liu et al. (2021) propose a framework using smart contracts and "fact index tokens" (cryptographically signed and verified facts) to enhance ESG data authenticity and transparency, partially addressing the oracle problem. While these efforts demonstrate technical solutions to establishing trust in sustainability reporting, they often lack attention to regulatory integration, standardization, and usability.

As a result, current solutions remain fragmented, variously focusing on traceability Mugurusi and Ahishakiye (2022), information flow Balzani and Corsi (2022), or data authenticity Liu et al. (2021), without offering a unified, holistic approach. There is a significant need for developing an artefact that integrates these technical functions with regulatory compliance, a flexible audit model, and a user-friendly design for, for

instance, investors and regulators. Such an approach could improve institutional trust and reduce greenwashing in corporate ESG reporting.

2.4 Supporting trust-building with machines by harmonizing regulation regarding ESG in the EU

While blockchain technology provides a robust technical foundation for enhancing trust, its effectiveness for ESG reporting is realized only when integrated with clear and enforceable regulatory frameworks. Trust in ESG data is not solely a function of technological safeguards; it also hinges on institutional factors that define reporting requirements, verification mechanisms, and enforcement processes across jurisdictions. Recognizing this, the European Union EU has taken a leading role in formalizing ESG-related obligations. As part of its broader commitment to achieving a net-zero economy by 2050 (Directorate-General for Financial Stability, Financial Services and Capital Markets Union, European Comission 2020), the EU has issued a comprehensive set of legislative measures aimed at aligning financial markets with sustainability goals. These initiatives are designed to redirect financial flows toward sustainable activities and provide investors with clearer insight into non-financial risks, the societal and environmental impacts of business practices, and long-term corporate sustainability performance. To this end, the EU has developed an integrated package of policies, , directives, delegated acts, and regulations intended to standardize and harmonize sustainability across member states (European Parliament 2024b). With this harmonization, the EU seeks to reduce the risk of greenwashing, and prevent market and regulatory fragmentation (Randazzo and Perozzi 2023). Moreover, it can provide transparent data to help investors evaluate how companies contribute to EU climate and energy target (Truant et al. 2023).

Key legislation in this context includes the EU Taxonomy Regulation (EU Taxonomy, 2020/852 (2020)), the Sustainable Finance Disclosure Regulation (SFDR, 2019/2088 (2019)), and the Corporate Sustainability Reporting Directive (CSRD 2022/2464 (2022)). The EU Taxonomy establishes a classification system and outlines six core environmental objectives (Article 9) to define sustainable activities and prevent greenwashing. Complementing this, the Sustainable Finance Disclosure Regulation (SFDR) imposes specific sustainability disclosure requirements for financial market participants, such as asset managers and institutional investors, mandating transparency at both the entity and product levels through pre-contractual documents (Articles 8-9), website disclosures (article 10), and periodic sustainability reports (Article 11).

The CSRD, a central pillar of the European Green Deal and the Sustainable Finance Action Plan (Secretariat-General, European Commission 2025) significantly expand the scope of reporting obligations. It requires a wider range of companies to disclose information using the European Sustainability Reporting Standards (ESRS, 2023/2772 (2023)), (Article 1) of the delegated regulation, and introduces the principle of double materiality. The Corporate Sustainability Reporting Directive (CSRD) additionally requires companies to report on both how environmental and social issues affect their business and how their operations impact people and their environment. Finally, the Corporate Sustainability Due Diligence Directive (CSDDD, 2024/1760

(2024)) compels companies to identify and address human rights violations and environmental harms within their operations and value chains. To balance regulatory ambition with economic competitiveness, the European Commission introduced the Omnibus proposal package earlier this year (Directorate-General for Financial Stability, Financial Services and Capital Markets Union, European Comission 2025b). This package aims to ease the compliance burden for small and medium-sized enterprises small and medium-sized enterprises (SMEs), particularly those indirectly subject to the CSRD and the CSDDD via contractual relationships with larger companies or financial institutions. The package also introduces voluntary reporting standards for companies outside the scope of these directives, reflecting a broader EU goal to streamline sustainability regulation and enhance economic vitality (Directorate-General for Financial Stability, Financial Services and Capital Markets Union, European Comission 2025c; Directorate-General for Communication, European Commission 2025a).

Although these regulatory measures lay a strong institutional foundation for defining what must be reported, how data is verified and accountability enforced across jurisdictions, regulation alone is insufficient to ensure trustworthy ESG reporting. Differences in interpretation, variation in implementation, and delayed enforcement leave room for inconsistencies and manipulation, diminishing the overall credibility of ESG reporting. This underscores the complementary role of technologically embedded mechanisms - such as blockchain - to reinforce trust. The design artefact and design principles proposed in this study seek to address this gap by providing a secure, verifiable infrastructure that supports accurate, auditable, and tamper-resistant ESG reporting aligned with evolving regulatory requirements.

3 Research Approach

3.1 Design Science Research Approach

We applied the DSR methodology, a well-established approach for developing information technology (IT)-based artefacts that address real-world problems through purposeful design (Peffers et al. 2007). This makes DSR particularly well-suited for our goal: designing a blockchain-based ESG reporting system that integrates regulatory with institutional requirements to derive more generalizable design knowledge on how foster institution-based trust in corporate responsibility practices. DSR has been successfully used to develop a wide range of artefacts, from technical architectures to more abstract constructs, such as design propositions and design theories (Utz et al. 2023). This methodological flexibility enables us to bridge concrete system development with theoretical abstraction. To guide our development process, we structured the Design Science Research Methodology (DSRM) Process Model around two sequential iterations (Figure 1). Each iteration followed the nominal DSRM phases: (1) problem identification & objectives definition, (2) design & development, (3) demonstration, (4) evaluation, and (5) communication (Peffers et al. 2007). Across these interactions, we incrementally developed and refined a blockchain-based ESG KPI reporting system. The design was informed by a set of emerging DOs, continuously shaped by empirical insights and stakeholder feedback. The final set includes six core DOs - compliance,

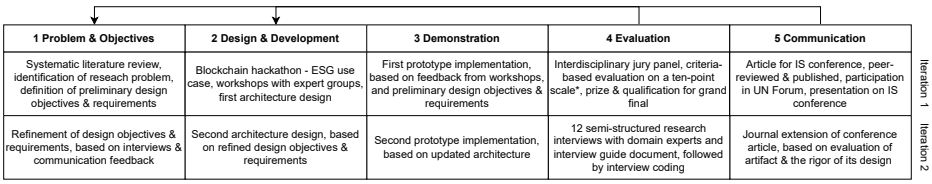

Fig. 1 Adapted DSRM Process Model, based on Peffers et al. (2007)

reliability, transparency, standardization, comparability, and simplicity - supported by fifteen detailed DRs.

Our research followed the two iterations of the DSRM Process Model (Peffers et al. 2007). The initial set of DOs and DRs was primarily shaped by insights from a systematic literature review and expert workshops during a blockchain hackathon. The evolving artefact was then demonstrated and evaluated through a multi-phase process, including peer feedback at an academic conference, stakeholder consultation at a UN forum, and a final, rigorous round of expert interviews.

This iterative process not only shaped the artefact but also enabled the extraction of generalizable design knowledge in the form of DPs. These principles aim to guide the integration of trust-building technical and regulatory components into ESG KPI reporting systems and to contribute to broader theories on technology-mediated trust-building in institutional contexts. At each stage, we continuously assessed the evolving artefact against current legal frameworks for ESG reporting to ensure both practical applicability and theoretical relevance (Gregor and Hevner 2013; Hevner and Chatterjee 2010). The resulting set of DPs constitutes an exaptation-type knowledge contribution, extending a known solution – blockchain-based systems for trust – to address the emerging challenge of declining trust in ESG reporting (Gregor and Hevner 2013).

3.2 Identifying the Problem and Defining the Objectives

3.2.1 Literature Review

We began with a *systematic literature review* to *identify the research problem*. Using a broad keyword search - **"Blockchain" AND ("ESG reporting" OR "ESG Framework" OR "ESG software" OR "ESG tool")** - we screened seven major databases (Springer, AISeL ArXiv, Emerald, Elsevier, IEEE, ACM Digital Library), supplemented by Google Scholar and two AI-powered databases, Dimensions and Elicit. The initial search produced 373 articles and book chapters. Subsequently, we excluded items that were not written in English, book chapters, and articles either not accessible or from journals that fall below the 85th percentile in Scopus (Webster and Watson 2002; Brocke et al. 2009, 2015). We then reviewed titles and abstracts of the remaining subset to manually filter irrelevant papers. We retained 9 papers for full-text review and included additional relevant papers through snowballing. From this refined corpus, we developed a preliminary problem statement and *derived a preliminary set of DOs and their corresponding DRs*, using two rounds of coding following Saldaña

(2022): the first employed initial and in vivo coding to identify design objective (DO) & design requirement (DR), followed by axial coding to specify their dimensions and interrelations.

3.2.2 Hackathon

To initiate the design & development phase, we participated in a 48-hour *blockchain hackathon* in March 2024, focused on sustainability goals and ESG applications in life insurance. During the hackathon, we participated in *workshops, conducted by groups of experts* from life insurance, who provided critical insights into ESG KPIs assessment challenges, and from a blockchain lab, who validated the technical feasibility of our proposed artefact and approved the choice of development tools for implementation. Since workshops are a recognized method for artefact development and evaluation in DSR (Thoring et al. 2020), the hackathon provided an ideal setting to *design the first architecture* of our ESG KPI reporting artefact grounded in the preliminary DOs & DRs.

3.3 Demonstration and Evaluation

The artefact's *first prototype* was implemented and demonstrated during the hackathon (Peffers et al. 2007). An *interdisciplinary jury panel evaluated the projects using a ten-point scale based on criteria* such as project viability, technical soundness, and pitch quality. The blockchain lab experts were specifically valuable in giving feedback on the technical soundness. Although scores were confidential, our team was awarded a token grant, qualified for the international final, and became eligible to apply for a post-hackathon grant.

3.3.1 Conference

Following the hackathon, we conducted a conceptual and criterion-based analysis of the artefact using established ESG regulatory frameworks, including EU Taxonomy, SFDR, and CSRD. This helped us to better align the artefact with regulatory requirements and expand its domain-specific capabilities. We summarised the development of the ESG KPI reporting system in the form of a *research article, submitted to a conference in Information System*. The article was *peer-reviewed and accepted*. We also received feedback from the reviewers and were invited to present our research on site. *We presented the research on the conference*, and *the article was published in the conference proceedings*. At this stage, our artefact was grounded in four DO, and five corresponding DRs.

3.3.2 UN Forum

In December 2024, we participated in the UN's Forum on Energy for Sustainable Development (Economic and Social Commission for Asia and the Pacific, United Nations 2024b) (hereafter: UN Forum), engaging in multiple workshops and group sessions where we demonstrated our artefact to diverse stakeholders including country delegates, diplomats, ambassadors, transnational corporate leaders, and UN officials.

During the Capacity Building Workshop (Economic and Social Commission for Asia and the Pacific, United Nations 2024a), we explored the role of digital technologies in the transition from centralized to distributed energy systems to enhance sustainability, resilience, accessibility, and fairness in energy services. This engagement provided critical insights and helped identify potential interviewees for the evaluation phase of the second iteration.

3.3.3 Interviews

Feedback from both the UN Forum and IS Conference informed the second iteration of our design & development process (Figure 1). To rigorously evaluate the artefact, we conducted*twelve semi-structured interviews* with domain experts , including ESG specialists, industry practitioners, academics, representatives of intergovernmental organizations (IGOs) and stakeholders familiar with UN sustainability initiatives (Table 1). This qualitative approach allowed us to examine perceptions of the artefact's utility, usability, and trust-building potential. Throughout the interviews, we used *an interview guide* (C) to structure the conversation while allowing participants to explore topics they found particularly relevant in this context. In this final evaluation phase, we challenged and refined our DOs and DRs, while also surfacing relevant DPs, resulting in overall six DOs, their corresponding fifteen DRs, and three DPs.

Table 1 Interview Participants

Industry	Org. Size	Position	Exp. (Y)
Insurance	100	Senior Vice President	17
Renewable Energy	11-50	Management Consultant	30
ESG Consulting	11-50	Project Manager	16
Higher Education	2k	Full Professor	47
ESG Consulting	2.5k	Senior Auditor	12
Financial Services	3.8k	Partner, Advisor	15
Higher Education	4k	Full Professor	17
Higher Education	10k	Postdoctoral Researcher	15
ESG Consulting	201-500	Reporting Manager	9
IGO	37k	Director Global Strategy	26
Higher Education	1.2k	Postdoctoral Researcher	8
Natural Resources	4.2k	Reporting Manager	19

3.3.4 Coding

The transcripts were analyzed by the same authors who conducted the interviews using higher-order conceptual categories. Following the recommendations of Saldaña (2022); Miles et al. (2018), we adopted a structured, iterative, and pragmatic approach in our two coding cycles to systematically analyze our multi-source data and refine and expand our design knowledge.

Our goal in the first coding cycle was to develop an initial set of DOs and DRs that could inform the development of blockchain-based ESG reporting systems. We began with a provisional list of codes derived from a review of academic and practitioner

literature on ESG reporting challenges in the financial industry and beyond. These initial codes reflected key constructs such as transparency, verifiability, and regulatory ambiguity. We supplemented these with inductive codes as new insights emerged during our analysis of data, which included expert interviews, hackathon/workshop notes, and written evaluation feedback. This first coding cycle was characterized by a combination of descriptive, process, and in vivo coding (Saldaña 2022) to capture participants' terminology and the underlying dynamics of institution-based trust, as well as characteristics of blockchain that appeared to have a mediating effect, and regulation. In the second coding cycle, we employed pattern coding to consolidate and abstract first-level codes into broader themes and categories (Saldaña 2022). Then these themes were mapped to the evolving sets of DO and DR. In line with Miles et al. (2018)'s recommendations, we focused on clarifying interrelationships between codes, including tensions and complementarity. As part of this process, we validated and refined our DOs and DRs via a criterion-based analysis against major ESG regulatory frameworks. We also used stakeholder feedback from hackathon participants and domain experts, and peer-review feedback to ensure practical relevance and derive generalizable learnings. These learnings were then consolidated from the refined DOs and DRs as well as the architecture, resulting in preliminary DPs. Additional interviews and iterations with the theory of institutional trust further helped to enhance their theoretical coherence, generalizability, and applicability. Our refined DOs & DRs were also instrumental in revising the *system architecture and prototype implementation*, which will be explained more in detail in the architecture section.

4 ESG KPI Reporting System

This section outlines the design and initial implementation of the proposed ESG KPI reporting system, translating the previously established DOs & DRs into a tangible artefact. This section is structured in four parts; the first subsection details the final set of DOs & DRs (Table 2). The second subsection presents the system's Unified Modeling Language (UML) component diagram (Figure 2), illustrating the structural relationships between the system's modules and components (Bell 2004; Briand et al. 1999). The third subsection describes the current stage of the prototype and its key technical features. And finally, the fourth subsection describes the two design iterations that led to the final artefact.

4.1 Design Objectives & Design Requirements

Our set of six DOs was systematically formulated to address critical shortcomings in current ESG reporting practices identified in our literature review, such as a lack of transparency, reliability, and comparability, which contribute to "greenwashing" and erode institutional trust. Each DO provides a direct response to these challenges. For instance, DO2 (Reliability) and DO3 (Transparency) are introduced to counter data manipulation, while DO5 (Comparability) directly addresses the problem of inconsistent benchmarking.

In this subsection, we explicitly transformed the problem (Peffers et al. 2007) into six DOs supported by fifteen DRs (Table 2). The initial set of objectives (DO1 –

DO4) and requirements (DR1 – DR5) was grounded in the literature and validated during the hackathon, while the final framework was extended and derived directly from our analysis of expert interviews, with anonymized references to the interview partners (e.g., I1, I5), and our criterion-based analysis against major ESG regulatory frameworks.

Table 2 DO interdependencies and supporting DRs

DOs	**Supported by DRs**	**Interdependent DOs**
DO1: Compliance	DR1: Scope: KPI classification (mandatory/optional, public/private) DR2: Context: Context-specific KPI selection DR6: Regulatory: Up-to-date alignment with EU legislation	DO4: Standardization DO5: Comparability
DO2: Reliability	DR3: Technical: Immutability of KPIs DR7: Data Quality: Robustness of data DR8: Metadata: Source of data	DO3: Transparency DO4: Standardization
DO3: Transparency	DR4 Technical: Publicly accessible KPIs DR9: Procedural: Clarity of methodologies, system boundaries, data process, scope and target audience DR10: Assurance: Verifiability of data	DO1: Compliance DO2: Reliability DO5: Comparability DO6: Simplicity
DO4: Standardization	DR5: Data Structure: predefined types, lengths, equivalents, and conversions for KPI values	DO1: Compliance DO2: Reliability DO5: Comparability
DO5: Comparability	DR11: Analytical: Internal historical cross-comparison DR12: Analytical: Comparison with other companies in the sector DR13: Analytical: Comparison with external standards and benchmarks	DO1: Compliance DO3: Transparency DO4: Standardization DO6: Simplicity
DO6: Simplicity	DR14: Design: Readily accessible user interface DR15: Functionality: Streamlined data input process	DO3: Transparency DO5: Comparability

DO1 – Compliance. Achieving the design objective "compliance" in ESG reporting requires a multidimensional approach to ensure that reported KPIs are contextually relevant and are aligned with legal, regulatory, accounting, and technical standards (I6) (Yu 2024). This entails distiguishing between KPIs that are selected based on organizational context and those that require continuous alignment with evolving legislation.

First, sustainability disclosure - whether mandatory or voluntary - is essential for organizations, encompassing both financial and non-financial data to meet stakeholder information needs (Balzani and Corsi 2022; Liu et al. 2023). The system must therefore enable companies to clearly *distinguish between mandatory (public) and optional (public or private) ESG KPIs (DR1)* to ensure comprehensive reporting and facilitate

appropriate data sharing, including relevant ESG requirements (Duran and Tierney 2023). This distinction guides how companies allocate their reporting resources (I5). Mandatory KPIs are typically publicly accessible and reflect generic performance indicators (I1, I5 I3, I9), while optional KPIs can offer more detailed, stakeholder-specific insights and may only be shared selectively with investors and banks (I1, I10). The system must also provide transparent selection criteria for both types of KPIs (I3).

Second, compliance requires simplification and contextualizing of data reporting based on a company's size, operations and financial turnover (Markopoulos et al. 2023; Miranda et al. 2023; Gong et al. 2024). Interviewees further emphasized the importance of a company's geography, jurisdiction, business activities, and strategic goals (I2, I4, I5, I6, I10). This *context-specific KPI selection (DR2)* adds complexity: for instance, emissions KPIs depend on country-specific factors such as energy use, verified through smart meters or invoices (I7); meanwhile, water or waste KPIs require detailed information on operational boundaries and a deep understanding of value chains (I7). To mitigate the risk of greenwashing (I4), companies must provide clear justification and traceable evidence for KPI selection (Balzani and Corsi 2022; Mugurusi and Ahishakiye 2022; Gong et al. 2023; Liu et al. 2023; Pizzi et al. 2022), aligning disclosures with both internal goals (I5, I6), and external frameworks like the Sustainable Development Goals (SDGs) (I5, I11).

Third, the system must enable tracking progress with relevant legislation (I3). Specifically, it must support continuous *up-to-date alignment with EU legislative frameworks (DR6)*, which involves balancing data protection with data disclosure. This includes not only sustainability-specific regulations such as the EU Taxonomy, CSRD and SFDR, but also data-centric laws such as eIDAS and the General Data Protection Regulation (GDPR) (I1, I5, I10, I11, I12).

DO2 – Reliability. The design objective of "reliability" centres on the structural quality of original ESG data – its consistency, provenance, and resistance to manipulation (I7). While transparency addresses the visibility of facts (I8), reliability ensures those facts are accurate, traceable, and credible, particularly for stakeholders (Balzani and Corsi 2022) like banks and investment funds (I10). This is achieved through three DRs that ensure data immutability, robustness, and a clear and verifiable origin (provenance).

First, the system must guarantee that ESG information is immutable and tamper-resistant, representing a "single version of truth" (Balzani and Corsi 2022; Yu 2024; Wu et al. 2022; Liu et al. 2023; Miranda et al. 2023). Blockchain technology supports this by making data manipulation and distortion practically impossible (Balzani and Corsi 2022; Pizzi et al. 2022). The system must therefore ensure *immutability of ESG KPIs (DR3)* while also preserving authenticity (Chen et al. 2024; Luo et al. 2024) and credibility (Yu 2024). However, once data is stored, errors become permanent and cannot be corrected (I7), which presents a risk (Luo et al. 2024). As one interviewee put it: "if the data that flows in the first place is incorrect, then you will not be able to change it. So, maybe placing a certain number of cheques before confirming the data with the appropriate [data]" and "a preview of the future output prior to sending the report" (I3) could allow companies to detect and correct errors prior to final upload, avoiding the permanence of inaccurate data on the blockchain.

Second, ensuring *robustness of data (DR7)* requires multi-level controls and validations to verify data truthfulness (I11) and thereby satisfy stakeholder expectations (I10) (Liu et al. 2024). While historical data supports benchmarking and trend analysis, it does not guarantee robustness or prevent manipulation (I11). Techniques like "backfilling", i.e., the retroactive adjusting of data to align with desired outcomes (I12), remain a risk. According to one interviewee, "the source does matter, [and] is one of the requirements for data robustness" (I11). Thus, techniques such as "digital tagging" can help trace the "entire history about the data point" (I11), but ultimately, robustness hinges on clear provenance and strong protection against tampering.

Third, reliability depends on the credibility of the data source, which is also a core component of robust reporting (I11), and requires *metadata about the source of data (DR8)*. Metadata should clarify whether data is self-reported, estimated, or independently assured, adding essential context for interpretation and comparability of ESG information (I8). In cases of self-reporting, any use of estimates, assumptions, or specific calculation methods must be explicitly declared (I8).

DO3 – Transparency. The DO of transparency is to balance public disclosure with business confidentiality (I5), particularly for quantitative ESG data – the core focus of our blockchain-based ESG KPI reporting system (I4, I5, I6, I7). This is achieved through three DRs: data accessibility, procedural clarity, and data verifiability.

First, ESG KPI data must be *publicly accessible (DR4)* via the public ledger. Availability in real time is important for stakeholders engagement (Yu 2024; Balzani and Corsi 2022; Markopoulos et al. 2023; Luo et al. 2024) and allows them and others, for instance investors, to quickly assess a company's ESG performance before allocating funds, enabling more informed decision-making (Mugurusi and Ahishakiye 2022; Yu 2024). Similarly, regulatory bodies can monitor compliance with ESG standards in real time, reducing the burden of manual audits and ensuring accountability (Duran and Tierney 2023).

Second, transparency requires procedural *clarity of methodologies, system boundaries, data process, scope, and target audience (DR9)*. Specifically, transparency must extend beyond final figures to include every step of the reporting process, from data generation and methodological choices to the final published report, so that others can accurately interpret what the data truly represents (I4, I6, I11). Without this context, transparency may "backfire", allowing greenwashing through misrepresentation of or masking of the actual performance (I10, I11). Providing full transparency of data workflows – including how data is gathered from dashboards or Excel sheets, and avoiding "copy-pasting" (I5, I7) – allows stakeholders to determine whether deviations are justified or constitute "red flags" (I10). Additionally, the criteria used in filtering mechanisms, such as the distinction between listed and unlisted companies, must be easily accessible (I8). However, explaining calculations, such as for Scope 3 emissions, which rely heavily on assumptions and complex upstream data methodologies, remains a challenge (I7, I11).

Third, immutability alone does not guarantee initial data quality. Even a secure blockchain record – that technically does not require a validator (Balzani and Corsi 2022) – needs verification of the raw data's authenticity (Yu 2024; Liu et al. 2021;

Wu et al. 2022) so that "the original data entered into blockchain is factual from the start" (I7). As trust cannot rely on "goodwill" (I11), ESG KPIs – after being written on blockchain – should be stamped by an accredited auditor, adding an external layer of verification (I1, I6, I11, I12). That is, to ensure the system's integrity, *verifiability of data is essential (DR10)* and can be reinforced through a robust regulatory framework (I2, I8, I10, I11) and strong internal system control (I11). In practice, ESG data may be verified against reliable sources such as invoices or smart meters – preferred due to their reliability, precision, and auditability (I7).

DO4 – Standardization. Standardization aims to address the diversity in which ESG data – often non-standardized and unstructured – is reported by companies, third parties, and independent bodies. By translating this data into more comparable formats, standardization enhances both transparency and reliability of ESG data (Macchiavello and Siri 2020; Gong et al. 2023; Liu et al. 2021; Yu 2024; Liu et al. 2024). In this context, interviewees emphasized the need for a *"solid set of KPIs"*, warning that without such standardization, ESG reporting *"would be a mess"* (I1). They stressed the importance of a shared taxonomy that enables consistent classification across industries and companies within the same industry – taking into account different requirements based on company size, industry, and data presentation methods (I1, I5, I6, I9, I7, I10, I12). To ensure consistency and comparability, the system should enforce *predefined types, lengths, equivalents, and conversions for ESG KPI values (DR5)*. This is particularly important in emerging and frontier markets, where ESG reporting practices are still maturing. Standardized input formats, data types, and conversions, support fairer and more reliable comparisons across companies (I1, I5, I6, I12).

Importantly, standardization is not only a compliance requirement – it is a prerequisite for meaningful long-term comparability (I8, I10, I12). Ensuring that KPIs remain stabile over time while allowing necessary adaptation or fluctuation is essential for historical comparisons to be valid (I12). This structured approach must be supported by clear governance mechanisms, such as the appointment of accredited auditors, to ensure the integrity of the reported data (Liu et al. 2024).

DO5 – Comparability. The design objective "comparability" is essential building investors trust and enabling market-driven sustainability performance. Transparency and standardization lay the foundation for meaningful comparisons, which reduce ambiguity and allow stakeholders to assess performance effectively (I10, I11). This objective is supported through three distinct comparison types.

First, the system must support *internal historical cross-comparison (DR11)*, allowing companies to track their own ESG performance over time. Having "historical data storage that is difficult to tamper with or that is difficult to backfill (I12)" is key to demonstrating real progress, while avoiding bias, minimizing data manipulation risk (I4), and holding companies accountable to past commitments (I11). Consistent year-over-year reporting, with companies submitting data from the current and prior year (I7, I11), enables the detection of restatements, updates, anomalies, and methodology changes – especially in cases of rebaselining (I7). To ensure fairness, reports should include disclaimers for unusual data and explain whether deviations are due to sustainability measures or incidental events like machine shutdowns or infrastructure changes,

reflecting intentional progress rather than accidental outcomes (I4). Thus, core figures should be contextualized with base years, material events, system boundaries, and organizational or operational scope (I4, I7).

Second, the system must enable *comparison with other companies in the sector (DR12)*. Peer benchmarking is valuable for investors (I3, I4, I6, I11, I12), but only meaningful if supported by transparent disclosures of system boundaries, calculation methods, and assumptions (I4, I7). Inconsistent calculation models – even with identical scopes – can distort the assessment outcomes (I4, I7, I11, I12). Knowing "how clean the data is or how complete the data is for any sort of investor (I7)", is critical for interpretation. Without this transparency, comparative assessments risk being misleading and may open the door to greenwashing (I4). To this end, quantitative metrics must remain stable over time to allow for tracking progress with integrity (I5, I6, I7).

Third, the system must facilitate *comparison with external standards and benchmarks (DR13)*. This functionality directly supports financial decision-making, such as loan evaluations by banks or risk assessments by investors (I6, I10). Although some companies may be hesitant to share detailed data due to competitive concerns (I8), comparability can also serve as a performance driver, helping companies identify gaps and motivating improvement (I10). A comparable reporting environment further enables companies to identify manipulative behavior or unsubstantiated claims, thus uncovering potential greenwashing (I10).

DO6 – Simplicity. Simplicity is a key requirement for widespread system adoption, particularly among non-technical users and mainstream audiences (I3, I7). Thus, the system must minimize complexity, maximize ease of use, and provide clear cost-benefit value, avoiding a "technology push", that overwhelms users with blockchain-specific technicalities (I8, I9, I12). This is achieved through two core requirements.

First, "[the system] has to be readily accessible [for users not] the designers of this tool, [...] but for mainstream people [...] it should be easy, just like Googling " (I3). Indeed, clear and intuitive as well as *readily accessible user interface (DR14)* not only helps retain users and attract new ones, but can also enhance public perception and market value by making credible sustainability practices of a company more visible (I1, I9, I12). For partners like banks, simplified access to well-structured ESG data via an intuitive user interface (UI) is critical for efficient assessments (I6, I7).

Second, "auditing nowadays is something that is still very labour-intensive, where you have to talk to people and understand what was done, [...] going towards more and more documentation" (I11). Thus, a *streamlined data input process (DR15)* eases the burden on auditors (I11). It allows companies to easily report the most relevant ESG data while simultaneously verifying its accuracy (I12). It lowers barriers for stakeholders like asset managers and insurers by replacing repetitive questionnaires (I8). Ultimately, reducing complexity lowers audit cost, motivating tool adoption and encouraging voluntary reporting (I9, I12).

4.2 Architecture

The architecture of the ESG KPI reporting system is illustrated as a UML component diagram (Figure 2) and comprises four primary modules: *Frontend*, *Backend*,

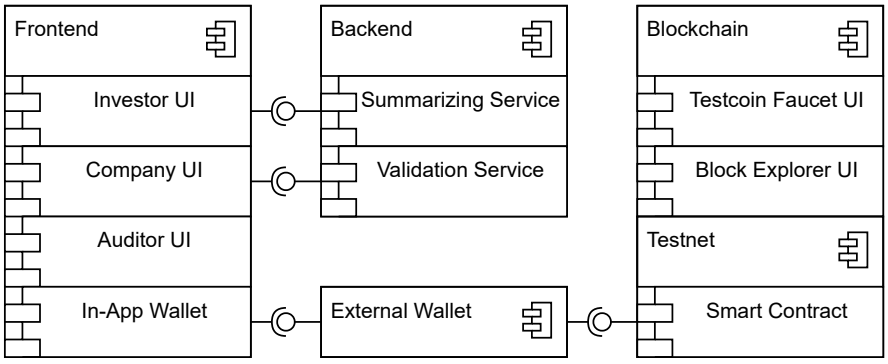

Fig. 2 UML component diagram of the ESG KPI reporting system

External Wallet, and *Blockchain*. Each module includes its own components and other embedded submodules, and they interact with each other via Application Programming Interfaces (APIs). The *Frontend* module is composed of the UI components (*Company/Investor/Auditor UIs*) and an integrated *In-App Wallet*. In line with simplicity (DO6), all UIs, all UIs are designed to be intuitive, simple, and easily accessible (DR14), enabling users to interact with the *Blockchain* module directly through the integrated *In-App Wallet*. The *Blockchain* module includes an embedded *Testnet* module containing a *Smart Contract* component, along with *Testnet Fuacet/Explorer UI* components for development and testing. To accommodate experienced users, an *External Wallet* module acts as an intermediary between the *Frontend* and the *Blockchain*. In production scenarios, an organizational wallet with four-eyes approval is anticipated to meet (DO1) compliance requirements (Abriani and Catania 2022; Milosavljevic 2023). The *Backend* module includes *Validation/Summarizer Services*, which pre- and post-process ESG KPI data before it is submitted to and retrieved from the blockchain.

The *Blockchain* module is designed to enhance trust and reduces greenwashing by ensuring transparency (DO3) through public access to ESG KPIs (DR4). These ESG KPIs can be viewed both from the application's *Frontend* and directly through the *Testnet Explorer UI*. For development and testing, the system uses standard blockchain tools such as a *Testcoin Faucet UI*, which provide developers with test coins to deploy and interact with smart contracts (Franzoni et al. 2020).

At the core of this module is the *Smart Contract*, deployed on the *Testnet* of a public Proof-of-Stake *Blockchain* chosen for its energy-efficient consensus protocol, which pre-emptively addresses environmental concerns. Once an ESG KPIs is recorded in the *Smart Contract*, it is reliably (DO2) stored with guaranteed immutability (DR3). This *Smart Contract* manages all on-chain logic for submitting, auditing, and reading ESG KPIs. A key feature of the design is the ability to "patch" individual KPIs in an append-only manner, typically in response to an auditor's request. This eases the tension between the need for corrections and the principle of immutability. Although a KPI is patched, all its previous values remain on the blockchain. This complete,

unalterable history is accessible through both the *Investor UI* and the *Testnet Explorer UI*, ensuring a fully transparent (DO3) and auditable trail of all data modifications. This mechanism directly enables internal historical cross-comparison (DR11), as it preserves the full lineage of every reported ESG KPI.

The *Company UI* facilitates the submission of ESG KPIs (DR1) and plays a central role in achieving compliance (DO1) and standardization (DO4). To promote standardization the UI enforces predefined data types, lengths, and formats for ESG KPI values (DR5), ensuring all data is compatible with the input requirements of the *Smart Contract*. Furthermore, it supports the selection of context-specific ESG KPIs based on geography, jurisdiction, business activities, goals, and size (DR2). To address compliance (DO1), the UI embeds reporting schemas that are kept in up-to-date alignment with EU legislative frameworks (DR6), including the EU Taxonomy, CSRD, and SFDR. Finally, the UI enhances data integrity by mandating the inclusion of metadata with every KPI submission. The metadata serves two purposes: it ensures reliability (DO2) by documenting the source of the data (DR8), and it provides procedural transparency (DO3) by requiring clarity of methodologies, system boundaries, and scope (DR9).

The *Validation Service* functions as a data quality gate, directly contributing to the overall reliability (DO2) of the system. It also pre-processes all ESG KPIs before they are committed to the *Blockchain* and conducts automated checks, cross-referencing new ESG KPI values against historical data. This helps facilitate both internal historical cross-comparison (DR11) and sector-based benchmarking (DR12), directly supporting comparability (DO5). It also ensures data verifiability (DR10) by flagging missing, invalid, or inconsistent data entries (significantly deviate from historical trends or expected formats) and suggesting corrections, which simplifies (DO6) the reporting process for users. Only after this pre-validation is complete does the company proceed to the final step: confirming the ESG KPI submission by signing the transaction (TX) in the *In-App Wallet*.

The *In-App Wallet* component serves as the primary interface for all on-chain interactions, bridging the UIs with the *Blockchain*. By abstracting the complexities of wallet management, the *In-App Wallet* is a key enabler of simplicity (DO6). It accommodates both mainstream users through its integrated functionality and experienced crypto users via an interface to their *External Wallet*. After a reporting TX is signed, the *In-App Wallet* submits it to the *Blockchain* by invoking a dedicated function in the *Smart Contract* to record the ESG KPI value and ensure reliability (DO2) and transparency (DO3). Following a successful submission, the system allows the company to initiate an ESG audit request for the newly recorded data.

The *Auditor UI* supports external verification of submitted ESG KPI data and serves as a digital "stamp of approval" from an ESG auditor. This human-led verification process represents a critical layer of data verifiability (DR10), significantly strengthening the overall system reliability (DO2). Through the *Auditor UI*, auditors can access a queue of "pending" reports submitted by companies that are awaiting verification. This design enables direct peer-to-peer interactions between individual auditors and companies and can potentially reduce the cost of ESG assurance through disintermediation. While the specific incentive models, accountability frameworks, and

operational shifts for this new generation of auditors are a subject for future research, our artefact primarily provides the technical foundation for such a model to emerge. Ultimately, by providing a mechanism for independent assurance, the *Auditor UI* enhances transparency (DO3). The final act of verification is performed when the auditor uses their *In-App Wallet* to cryptographically sign an attestation for the verified ESG KPI. This function is also essential for compliance (DO1), as it provides the necessary tooling for the external assurance mandated by frameworks (DR6) like the CSRD.

The *Investor UI* presents submitted and verified ESG KPI data pulled from the blockchain. The interface facilitates comparability (DO5) by enabling two analytical features: First, it facilitates historical comparison of ESG data within a single company (DR11), allowing users to measure the company's progress over time. Second, it supports benchmarking between companies within a given business sector (DR12), offering insights into how the company performs relative to its competitors.

To enrich the user experience (UX), we propose the *Summarizing Service* component, which processes complex ESG KPIs into more intuitive sustainability metrics. This makes the information more readily accessible (DR14) and directly supports the goal of simplicity (DO6), though the design of specific methodologies for this summarization remains a topic for future research.

4.3 Implementation

This subsection outlines the technical implementation of the ESG KPI reporting system as a prototype, based on the previously described architecture. The prototype is developed using Next.js (Vercel 2024), a full-stack React (Meta Open Source 2024) framework, and includes both the *Company UI*, and *Investor UI*. Interaction with the on-chain *Smart Contract* is managed via thirdweb Connect, which provides the integrated *In-App Wallet*. The *Smart Contract* itself is deployed on the Alfajores *Testnet*, a public testing environment for the Celo blockchain. Key technical refinements in this version include the implementation of an append-only "patching" mechanism for ESG KPIs and an enhancement of the data displayed via the public *Block Explorer UI*. An overview of the implementation workflow is presented as a UML activity diagram (Figure 3), with UI screenshots available in the appendix (App. A, B).

Blockchain: To support sustainability, we deployed the prototype on the Celo blockchain, which is recognized for its low carbon footprint (J 2024). As an Ethereum virtual machine (EVM)-compatible blockchain, Celo allows for the direct deployment of *Smart Contracts* compiled from Solidity. The migration involved configuring our Hardhat development environment (Nomic Foundation 2024) to connect to the Alfajores *Testnet* by specifying its RPC endpoint and chain ID (44787). To cover the gas fees required for deployment and end-to-end testing, funds were obtained from the public Alfajores *Faucet* (Celo 2024b). This step is crucial, as the reporting functionality fundamentally relies on writing data to the *Blockchain*, which consumes gas. For testing purposes, a corresponding custom network was also configured in the MetaMask browser extension (*External Wallet*) (Celo 2024a).

Our analysis of the TX costs revealed that a single ESG KPI submission currently costs approximately 0.0039 CELO on the *Testnet*. Based on an exchange rate of 1

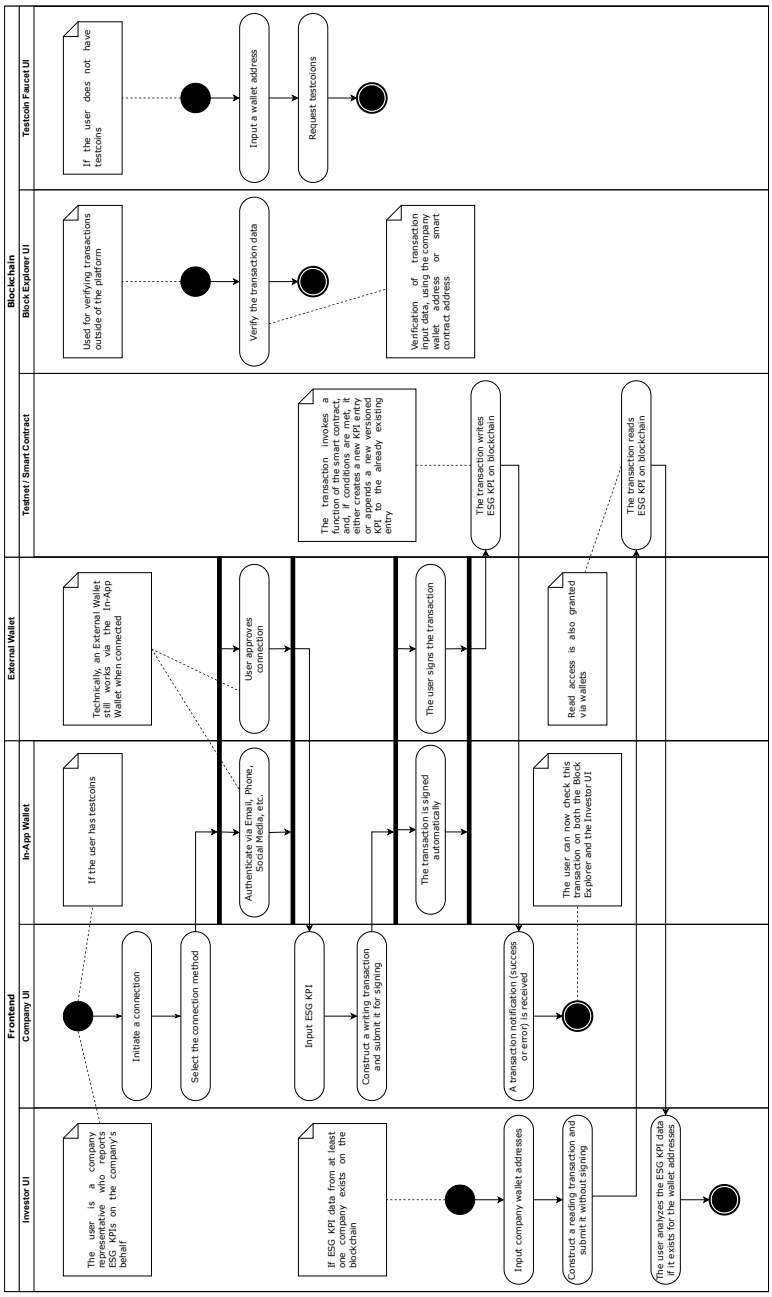

Fig. 3 UML activity diagram of the ESG KPI reporting system

CELO = €0.219 in June 2025 (Revolut 2025), this translates to a mainnet cost of less than one-tenth of a cent per transaction. To put this in perspective, reporting all 1100+ individual data points required by the European Sustainability Reporting Standards (ESRS) (Envoria 2025) would cost a total of approximately €0.94. While the current contract is only partially optimized, we plan to introduce batch submissions in future iterations to further reduce costs.

Explorer: Beyond a sustainable deployment, the public verifiability of the *Smart Contract* is critical for trust and auditability. An unverified contract on a block explorer appears as opaque bytecode, hindering analysis and interaction. To address this, we automated the contract verification process on Celoscan, the official block explorer for Celo. This involved incorporating a Celoscan application programming interface (API) key into our Hardhat configuration, which executes a command-line task to submit the source code for verification immediately after deployment to the Alfajores testnet. This step ensures the deployed contract is not only hosted on a sustainable platform but is also fully transparent and accessible for public review.

This transparency extends beyond the verified source code. The block explorer's "Events" tab provides a structured and human-readable log of key on-chain activities. In our implementation, the *Smart Contract* emits a Solidity event (an abstraction on top of the EVM's logging functionality) for each ESG KPI submission. Consequently, critical data attributes are not merely embedded within raw TX data but are indexed and readily available for manual verification. This feature directly enhances the transparency (DO3) of the reporting process and builds trust among stakeholders.

Smart Contract: The smart contract is the core on-chain component responsible for the management and historical tracking of ESG KPIs. It enables companies to submit their KPIs, ensuring an unalterable historical record through a versioning feature that implements an "append-only" patching mechanism.

The current implementation defines a `KpiVersion` data structure that stores the numerical `value` of the KPI, a `submissionTimestamp`, and a `metadataCid` used as content identifier for the methodology (DR9), which itself can be stored on InterPlanetary File System (IPFS). The `metadataCid` is stored as a `bytes32` type to efficiently reference this off-chain metadata (DR8). To optimize for gas efficiency, KPIs are stored on-chain using a nested mapping structure: (`kpiOwner` → `kpiTypeId` → `reportingYear` → `versionNumber` → `KpiVersion` struct), which avoids the high costs of dynamic arrays.

The contract's primary functions include `submitKpiVersion`, which creates new KPI versions by incrementing the version number and recording the `block.timestamp`. Upon each submission, a `KpiVersionSubmitted` event is emitted with indexed parameters to facilitate efficient off-chain querying. For data retrieval, the contract provides two key functions: `getLatestKpiVersion` to access the most recent submission, and `getKpiVersion` to retrieve any specific historical version by its version number. This architecture inherently supports immutability (DR3), traceability, and transparent (DO3) audit of verifiable data (DR10), as all submissions are permanently recorded and verifiable on-chain.

In-App Wallet: Interaction with Web3 applications, including our ESG KPI reporting prototype, requires users to have a blockchain-specific identity, commonly

referred to as a Web3 wallet (Murray et al. 2023). Traditionally, these wallets are external applications, such as browser extensions or mobile apps, which are well-suited for experienced crypto users. However, for the non-native crypto users who are the primary target audience for this system, the requirement to install and manage an external wallet creates a significant barrier to entry and hinders mass adoption (Krause 2025). This directly challenges our objective of creating a readily accessible user interface (DR14).

To address this challenge, we utilized thirdweb Connect, a toolkit from the full-stack Web3 development framework thirdweb (thirdweb 2024, 2025). This toolkit offers a seamless solution by providing an *In-App Wallet* that supports out-of-the-box integration with EVM-compatible blockchains like Celo. Crucially, it allows users to authenticate using familiar Web2 methods, including email, phone, passkeys, or social logins (e.g., Google, Apple). This design choice is a key enabler of simplicity (DO6), as it abstracts away the complexities of wallet creation and management. The toolkit acts as a wallet aggregator and UI layer and relies on Multi-Party Computation (MPC) for non-custodial key management in a decentralized way, so one can infer the supposed blockchain effect on trust. Furthermore, when a user is authenticated, thirdweb Connect can automatically handle the signing of transactions, creating a much smoother user experience.

Company UI: The reporting workflow, designed as a streamlined data input process (DR15), is initiated when a company representative accesses the prototype's homepage and connects their wallet. To facilitate frictionless onboarding during the prototype phase, the user is guided to acquire testnet tokens from the public Celo *Faucet* (Celo 2024b) if their balance is insufficient to cover transaction fees, ensuring cost-free testing.

Once connected and funded, the user can proceed with KPI submission. The current implementation focuses on a single mandatory ESG KPI: "GHG Emissions Scope 1 & 2 (tCO2e)". This metric is quantitative, and its unit of measurement is fixed to ensure standardisation (DO4) and compatibility with the *Smart Contract*. After the user signs and submits the TX, the UI provides immediate feedback by displaying a "Success! Transaction confirmed." message and a link to view the TX on the *Block Explorer*. Further details and screenshots of this process are available in appendices (App. B).

Investor UI: The *Investor UI* is the primary interface for ESG KPI data retrieved from the blockchain. To ensure broad and readily accessible (DR14) use, it operates in a read-only mode that does not require investors to connect a wallet. In the current prototype, users can retrieve a company's reporting history by entering its wallet address.

The retrieved data is visualised on a chronological chart, which is a key component for enabling comparability (DO5). Specifically, this view facilitates both internal historical cross-comparison (DR11), by displaying a company's performance over time, and creates the foundation for future features to support comparison with other companies (DR12). To make the presented information more intuitive, the *Summarizing Service*, as described in the architecture, will process raw KPI data into accessible metrics, further supporting simplicity (DO6).

To enhance usability in future iterations, we plan to implement an onboarding process for companies, allowing them to link their wallet addresses to human-readable names, potentially via the Ethereum Name Service (ENS). This will enable investors to search for companies by name rather than by cryptographic addresses.

4.4 Evaluation

The artefact was evaluated in two consecutive design iterations. The first focused on conceptual development and initial validation through a systematic literature review, a hackathon, and a criterion-based analysis, resulting in a preliminary set of four DOs (DO1 – DO4) and five DRs (DR1 – DR5). The second iteration involved a more rigorous expert-based evaluation via interviews with domain experts, expanding and refining the artefact to a final set of six DOs and fifteen DRs.

4.4.1 First design iteration

The first evaluation of our artefact followed a workshop format of a blockchain hackathon (Thoring et al. 2020), enabling rapid artefact development and feedback collection from a diverse group of stakeholders. This process involved three feedback loops, iterative refinement of the architecture, and a final assessment by an interdisciplinary expert jury.

Assessment of business needs. In the initial stage, the concept was presented to representatives of a life insurance company. They highlighted that ESG KPIs are currently difficult for clients to interpret and recommended summarizing them to help users "build emotional repertoire" around a company's ESG goals and impact. They also suggested "mitigating input flaws through validation" and prioritizing the most relevant ESG KPIs by business domain. In response, two key components – the Summarizing Service and the Validation Service – were added to the architecture.

Assessment of technical feasibility. In the second stage, the updated architecture was reviewed by technical experts from a blockchain lab. They confirmed its feasibility and recommended using an L2 EVM-compatible test network and Solidity to ensure interoperability with other ecosystems. These recommendations were incorporated into the first versions of our architecture and prototype.

Final jury evaluation. During the final hackathon stage, the artefact was evaluated by an interdisciplinary jury panel based on its technical design, buisness model, and implementation maturity. The project received an award and was selected for the international finals, validating its practical relevance and innovation potential.

Regulatory compliance analysis. Following the hackathon, we conducted a criterion-based analysis of the artefact against key EU legal frameworks (SFDR, EU Taxonomy, and CSRD). The analysis revealed that while the artefact meets the transparency and standardization requirements of SFDR and the EU Taxonomy, it does not yet fully address the "double materiality" principle of the CSRD, as the initial focus was on environmental KPI.

4.4.2 Second design iteration

The second iteration focused on systematically evaluating and refining the artefact through twelve semi-structured interviews with domain experts in ESG, audit, finance, and academia. These interviews confirmed the initial DOs and refined existing and surfaced new DRs, contributing to an improved artefact design.

Refinement of DO1 Compliance. We confirmed the validity of DR1 (KPI classification) and DR2 (context-specific KPI selection). To address specific regulatory mandates like the CSRD, we introduced DR6 (up-to-date alignment with EU legislation), which requires tooling for external assurance.

Refinement of DO2 Reliability. The interviews confirmed the importance of DR3 (immutability of KPIs). Experts also emphasized the need for verifiable trust, leading to two additions: DR7 (robustness of data), and DR8 (source of data), reinforcing provenance and immutable audit trails.

Refinement of DO3 Transparency. We confirmed DR4 (publicly accessible KPI) and extended the objective to include DR9 (clarity of methodologies, system boundaries, data process, scope and target audience) and DR10 (verifiability of data), ensuring that both content and the verification process are transparent to stakeholders.

Emergence of DO5: Comparability. Several interviewees (I3, I4, I6, I10, I11, I12) noted that reliable data has limited value without comparability. This led to the articulation of a new *DO5: Comparability*, and three new DRs: DR11 (internal historical cross-comparison), DR12 (comparison with other companies), and DR13 (comparison with external benchmarks).

Emergence of DO6: Simplicity. Non-technical stakeholders (I3, I7, I8, I9, I12) repeatedly pointed out that complexity in UI & UX is a key barrier to future adoption. This resulted in the introduction of *DO6: Simplicity*, and two DRs: DR14 (readily accessible user interface) and DR15 (streamlined data input process).

5 Discussion

Drawing on the framework of Gregor and Hevner (2013), we present three DPs aimed at guiding the development of blockchain-based ESG KPI reporting systems. These principles emerged through iterative design and evaluation cycles and are informed by regulatory frameworks, technical features of blockchain technology, and mechanisms for institutional trust-building. At a theoretical level, our findings contribute to research on technology-mediated institutional trust. Specifically, we demonstrate how blockchain – through characteristics such as reliability, standardization, transparency, comparability, and compliance – can serve as an infrastructural mechanism to mitigate greenwashing and support the development of institutional trust (Utz et al. 2023; Smits and Hulstijn 2020).

5.1 Practical implications

Practitioners engaged in ESG KPI reporting are currently challenged by a lack of standardization, limited comparability, and growing scepticism about the credibility of disclosures due to greenwashing concerns (European Federation of Financial Analyst

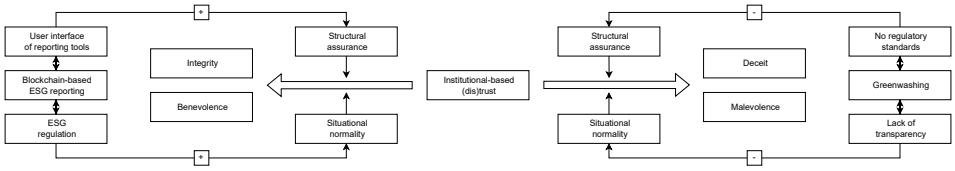

Fig. 4 Institutional-based (dis)trust

Societies 2010; European Supervisory Authorities 2024c). Our research responds to these challenges by presenting both a the blockchain-based system prototype for ESG KPI reporting and a generalizable set of design principles. The practical relevance of our work lies not only in its technical implementation, but also in the systematic design process, which can be adapted and replicated across other contexts.

The resulting DPs provide a concrete roadmap for practitioners aiming to enhance transparency and institutional trust in ESG reporting. They illustrate how regulatory frameworks can guide and constrain the design and development of reliable and trustworthy reporting systems. By abstracting these insights, we offer actionable guidance to support a broad range of practitioners in improving the credibility, comparability, and utility of ESG reporting across diverse contexts.

5.1.1 DP1: Compliant submission of raw data

A persistent challenge in ESG reporting is the ambiguity around what constitutes relevant and reliable data for performance assessment (European Supervisory Authorities 2024c). Our first DP addresses this issue by advocating a shift from outcome-based to verifiable input-based reporting, enforcing the compliant submission of raw data. Leveraging blockchain, regulatory compliance with ESG-related legal frameworks can be programmatically enforced, reducing manual intervention and potential errors (Yu 2024; Amend et al. 2024). At submission time, the smart contract enforces standardized KPI schemas and immutably records each value, timestamp, and metadata CID. Each submission triggers an indexed event, allowing auditors/regulators to automatically cross-check entries, turning the act of submission into a traceable and verifiable control point rather than a passive upload.

Context-specific legal frameworks provide criteria for KPIs selection (DR2) and support ongoing alignment with legal requirements (DR6). "Submission of raw data" emphasizes the need to capture unaggregated operational inputs instead of post-processed performance metrics. In our artefact, this is operationalized through standardized input formats (DR5), a real-time *Validation Service* that pre-validates these inputs against internal historical (DR11) and peer benchmarks (DR12), and immutable *Smart Contract*-based (DR3) recording. While Liu et al. (2021), Gong et al. (2023) and the interviewees frequently raised concerns about disclosing sensitive data, they perceived that a tiered access model – granting full access to auditors and aggregated summaries to the public, without any sensitive operational data that could influence competitiveness (I5, I6) – could balance transparency with confidentiality, and enhance compliance (DO1) as well as reliability (DO2). Another key insight from

the interviews was that ESG reporting systems should be compliant, aligning with frameworks like the CSRD, the SFDR, the EU Taxonomy (I6, I10, I11) as "if I am interested in a company, for me it's much more difficult to identify the data that I'm interested in, so particular ESG metrics as well. Everybody encounters the same problem. [But] if I have a database which already picks up all the information, it's really great in the moment" (I10). Contrary to expectations, compliance was not viewed as a constraint, but as a structural DP that provides a legitimate basis for defining what data is material and must be collected (I6). This approach necessitates a clear definition of organizational and operational boundaries (DR2), as different scopes significantly impact how performance is interpreted (I7). Furthermore, it requires full methodological transparency, including disclosure of standards, assumptions, and variables, which is a prerequisite for valid interpretation and comparability.

Together, these interconnected requirements and objectives coalesce into our first DP **compliant submission of raw data**. This principle reframes ESG reporting from a narrative-driven task to a structured, traceable, compliance-driven, and verifiable process.

5.1.2 DP2: Public access to an immutable yet append-only correctable raw data

A recurring concern in our evaluation was the tension between data immutability (DR3), necessary to transparency (DO3), and the flexibility needed to accommodate errors in ESG reporting. While public accessibility (DR4) and immutability (DR3) enhance trust and reduce greenwashing, some interviewees cautioned that honest mistakes in data submission are inevitable (I4, I10). Without a mechanism for post-submission rectification, companies may feel trapped by unchangeable errors, leading to reluctance to adopt blockchain-based systems.

First, they may be reluctant because strict immutability could be misinterpreted as permanent endorsement of flawed data, even when the error was unintentional (I4). Moreover, public visibility of ESG data – while fostering comparability (DO5) – can create pressure encouraging adverse behavior. In fact, "companies have become quite good in spinning their results and showing off the good stuff, [...] concealing the not so good stuff (I5)", manipulating future reports to appear competitive rather than genuinely improving performance (I10). To address this, we implemented in our prototype an "append-only patching" mechanism: companies can sumbit a new, correct version of an ESG KPI, which is recorded as a new entry, while the full and auditable history of changes remains transparently on-chain. This preserves data integrity and enables justified updates without erasing past records or impacting immutable traces (DR3).

This approach leads to our second DP: **public access to an immutable yet append-only correctable raw data**. This principle reconciles transparency and flexibility by making corrections themselves part of the auditable record (DO3), fulfilling data verifiability (DR10), while ensuring that the history is immutable (DR3).

5.1.3 DP3: Industry-specific and materiality-driven contextualization

One of the most significant obstacles to effective ESG reporting is the prevalence of vague, mission-statement-style objectives lacking accountability (I6, I10). A key insight from our interviews is the need to ground ESG reporting in the principle of double materiality (I1, I6, I7). This principle requires companies to consider two perspectives: internal materiality – sustainability-related risks to financial performance – and external materiality – a company's impact on society and the environment. This reframes compliance not as a mere constraint, but as a strategic exercise in identifying, managing, and mitigating material risks (I6).

Furthermore, interviewees highlighted the need to align with emerging EU legislation on "cross-cutting and sector-specific" ESG reporting standards (I1, I10). This industry-specific approach is important as material topics vary substantially across industries. For example, beverage companies must report plastic consumption alongside general emissions to avoid the omission of key ESG dimensions that can lead to "passive" greenwashing (I7). Such standards also enable the creation of meaningful benchmarks that distinguish genuine performance improvements from superficial ones (I7).

These insights lead to our third DP, namely **industry-specific and materiality-driven contextualization**. This principle ensures that ESG KPI reporting is not genetic but tailored to industry-specific risks and opportunities, guided by the logic of double materiality. We acknowledge that operationalizing this principle across diverse may pose challenges (I7), and that a holistic approach would also require companies to consider the ethical and environmental implications of the reporting technology – such as blockchain's energy use – as part of their broader double materiality assessments (I7, I11).

5.2 Theoretical implications

5.2.1 Impact on institution-based trust

The three design principles we identify go beyond offering actionable guidance for developing a blockchain-based ESG reporting system. They also contribute to the theoretical understanding of how institutional trust can be built through technology. Specifically, fostering such trust is based on the user's belief in the technology's key attributes, such as functionality, helpfulness, and reliability (McKnight et al. 2009, 2011). Once a technology is perceived as a driver of trust, it can become a catalyst for establishing *institutional trust* (Müller et al. 2024), particularly in areas like sustainability reporting, where such trust is often lacking (Ying et al. 2018). In the following we demonstrate how blockchain-based architecture and its underlying principles enhance *institution-based trust* (McKnight et al. 2017; Moody et al. 2013) and mitigate *institution-based distrust* (McKnight and Chervany 2001; McKnight et al. 2017). However, our evaluation also reveals that a purely technology-centric view of trust is insufficient. Trust was described not merely as a feature of the system but as a reciprocal socio-technical relationship, requiring continuous contribution and engagement from all involved actors (I10). This aligns with prior research that views

trust in blockchain as emerging from socio-technical configurations, not technological characteristics alone (Utz et al. 2023).

As illustrated in our figure (Figure 4), the blockchain-based artefact builds *structural assurance* and *situational normality* by aligning with EU legal frameworks and offering a user-friendly interface – two factors widely recognized as foundations of institution-based trust (McKnight et al. 1998; McKnight and Chervany 2001; McKnight et al. 2009, 2017). In contrast, greenwashing practices erode these foundations, triggering institution-based distrust (McKnight and Chervany 2001; McKnight et al. 1998; Moody et al. 2013), particularly when stakeholders perceive inconsistencies or ambiguity in ESG disclosures (Liu et al. 2024). While prior research has acknowledges blockchain's potential in ESG contexts (Balzani and Corsi 2022; Mugurusi and Ahishakiye 2022; Liu et al. 2021), few studies have unpacked how blockchain characteristics interact with trust-building mechanisms to meaningfully support ESG reporting.

We address this gap through our three DPs. DP1 directly responds to one of the core challenges in ESG reporting, that is, the ambiguity over what constitutes reliable and relevant performance data (European Supervisory Authorities 2024c). DP1 operationalizes trust by embedding regulatory compliance directly into the technical logic of the system. It ensures that ESG reporting is grounded in input-based, standardized, verifiable disclosures, rather than outcome-based narratives prone to manipulation. More specifically, through smart contracts and real-time validation services, the system automates compliance with legal frameworks like the CSRD and SFDR (Pizzi et al. 2022), reducing reliance on human interpretation and reinforcing perceptions of technological integrity and helpfulness. Thus, by binding regulatory legitimacy with technological enforcement, DP1 increases the stakeholders' willingness to depend on and engage with the system, contributing quality data (I10), which constitute a core dimension of both *trust in technology* and *institution-based trust* (McKnight et al. 2009, 1998). Importantly, our findings show that different stakeholder groups define trust differently (Liu et al. 2023; Luo et al. 2024): while investors prioritize traceability and data reliability (I10) (Yu 2024; Mugurusi and Ahishakiye 2022), companies highlight the importance of data privacy and control over access (I2, I6). To balance firms' privacy requirements with societal demands for verifiability, future iterations of our design could integrate privacy-preserving verification mechanisms (e.g., Zero-Knowledge Proofs (ZKPs)) that allow outsiders to confirm rule compliance without accessing raw data.

In defining our second design principle DP2, we address a fundamental tension between transparency and flexibility in ESG reporting. System gaps could be exploited to present misleading data as truthful, potentially enabling greenwashing rather than preventing it (I10, I12) (Chen et al. 2024). However, while blockchain's immutability and public verifiability foster trust by making tampering virtually impossible (Liu et al. 2024), there are concerns that rigid data permanence could backfire. In fact, in the current industry practice, ESG data is often locked in static, untraceable PDF files, lacking both versioning and historical context (I7, I9) and increasing the risk for counterfeit or manipulated disclosures (Mugurusi and Ahishakiye 2022). At the same time, concerns were raised about companies self-reporting erroneous or overly optimistic

data, sometimes due to a lack of expertise rather than intention (I4, I12). Definitely, simple "trust [on the blockchain technology] does not replace the quality of the input of the data [...]. It only goes as far as the data is correct in the first place" (I7). However, while blockchain reduces the likelihood of errors (Liu et al. 2024), errors in data submission are inevitable, and without a means of correction, the system could appear unforgiving or even misleading, or risking the perception that outdated or incorrect data has institutional endorsement (I4, I10). To resolve this, our system introduces an append-only correction mechanism. This allows versioned updates to be layered on the top of original values (I7), providing a clear indication of patches without requiring analysis of blockchain events. This way, it reinforces *structural assurance* through embedded safeguards that allow for methodological refinements, data updates, and error correction without compromising data integrity but retaining full transparency (I7) (Luo et al. 2024), and accountability, which positively impacts the trust of stakeholders (Chen et al. 2024). Moreover, by aligning with familiar practices like audit trails (Balzani and Corsi 2022) and revision tracking, it fosters *situational normality*, helping users feel at ease with its use and at the same time "creating a safe space (I11)". This design principle enhances *trust in technology* by demonstrating that the system is not only stable and transparent (Luo et al. 2024), but also capable of accommodating justified changes (Müller et al. 2024). Blockchain's tamper-proof nature, traceability (Pizzi et al. 2022) from raw ESG data to ESG scores in ESG assessment activities (Liu et al. 2023) and the use of smart contracts create user confidence in the stability and predictability of the reporting process (De Filippi et al. 2020), while its transparent handling of mistakes signals *benevolence* and *integrity*, thereby extending the perception of trustworthiness to the system itself (McKnight et al. 2009).

In defining our third design principle DP3, we address a critical gap in ESG reporting: the lack of relevance and foundations of accountability in reported metrics. Definitely, willingness is heavily influenced by system usability and pragmatic integration (Liu et al. 2023) since trust in any new solution depends on its ability to seamlessly fit into existing ESG reporting workflows, thereby avoiding duplication of effort and understanding industry-specific needs (I5). This alignment fosters *situational normality* and reinforces *institution-based trust*. Our evaluation revealed that vague, generic disclosures often undermine trust, as they fail to reflect the actual risks and impacts that matter within a given sector. To counter this, DP3 emphasizes the importance of aligning acESG reporting with the principle of double materiality and shows that the system is not only technically robust, but also attuned to contextual relevance. More specifically, rather than enforcing one-size-fits-all rules, it supports sector-specific guidance informed by emerging EU regulations, enabling reporting systems to reflect meaningful, industry-relevant priorities (Liu et al. 2021, 2023; Luo et al. 2024) and thereby reinforcing *structural assurance*, demonstrating to users that the system is grounded in legitimate and evolving regulatory standards. It also strengthens *situational normality*, as companies are more likely to trust and engage with a system that mirrors their operational reality and risk landscape. Moreover, by guiding companies toward disclosing material issues that truly reflect their impact, the design principle helps distinguish genuine sustainability efforts from superficial compliance. This responsiveness to context signals *integrity* in how the system defines and

enforces relevance. Yet, interviewees cautioned that trust in the system is conditional upon the quality of input data (I2, I5, I7). They noted that trust cannot be assumed based on technological design alone; it must be earned through demonstrable validity and verifiability (Liu et al. 2023, 2024). This finding supports the need for robust data triangulation mechanisms, such as aligning ESG claims with financial or operational data (I12). Moreover, by framing compliance not as a constraint but as a strategic means of identifying and mitigating risk, the system fosters *benevolence*. By pairing technological guarantees with regulatory legitimacy, all the relevant users obtain a sense that the system supports rather than polices them, building *institutional trust*, offering a principled foundation for ESG reporting.

Ultimately, even a perfectly designed system faces external challenges that moderate trust. The interviewees acknowledged the symbolic and market effects of trust, where a platform with committed users can generate a network effect, drawing in more stakeholders seeking credible sustainability indicators (I10, I12). As I12 mentioned "the more user use your platform, the more even investment manager fund managers. Come on your platform to fish out some sustainability and matrices to use for their own investments decisions". However, they emphasised that this network effect is conditional on the demonstrable truthfulness and integrity of the reported data, not just its availability (I11, I12). This conditionality is further complicated by the lack of specific regulatory oversight for the technology, which were mentioned as potential barriers to adoption (I6, I7, I11). Lastly, "educate people about the benefits of blockchain. Because if you take for instance, old school people [...] to them, blockchain is a means for tax evasion, which is not the main purpose. So, [...] there is a huge deal of education that needs to be done in order to make it reliable in the eyes of the public (I3)".

5.2.2 Foundation of trust in the blockchain-based ESG reporting system

As illustrated through our blockchain-based ESG reporting system, trust is not an intrinsic feature of the technology itself but emerges from the broader sociotechnical configuration that governs how data is verified, contextualized, and safeguarded. Building on our design principles (DP1–DP3), we demonstrate how regulatory alignment, sector-specific relevance, and embedded correction mechanisms collectively enhance perceived system integrity and benevolence, thereby enhancing trust in technology and institution-based trust and mitigating institution-based distrust by minimizing uncertainty and risk. For instance while automated legal compliance (DP1) reduces ambiguity, the capacity for error correction (DP2) reinforces system forgiveness, and relevance through contextual materiality (DP3) signals that the system acts in users' best interests. Moreover, our findings suggest that trust must be continuously earned through system transparency, usability, and responsiveness to diverse stakeholder needs, rather than assumed through technical features alone.

5.3 Limitations and further research

The current study describes the second iteration of an ESG KPI reporting system. We conducted several evaluations, updated the DOs & DRs, updated the architecture, and developed a prototype. While our prototype demonstrates the feasibility of ESG KPI reporting via blockchain, certain functionalities, such as tailoring KPI reporting schemas to specific business domains for automatic ESG compliance, remain conceptual at this stage, as the current implementation is limited to a single numeric GHG Emissions Scope 1 & 2 ESG KPI. Furthermore, key areas for future research have been identified, including: (1) the technical optimization of the smart contract through TX batching; (2) the architectural design of cross-organizational data links to prevent greenwashing; and (3) the socio-economic models for incentivizing and governing a disintermediated audit process. These functionalities are critical for achieving the full vision of the system and are planned for future iterations of development. Despite this, the prototype lays a solid foundation for further enhancements by establishing the core processes of data submission, validation, and immutability on the blockchain.

Defining reporting schemas to specify optional and mandatory KPIs across different clusters of business domains warrants a separate investigation. A fair comparability framework for companies that report based on different reporting schemas is also an avenue for future research. It will involve a rigorous analysis of regulatory frameworks and specificities of business categories. We hypothesise that a combination of Machine-Readable Legislation (MRL) and AI could help to define such reporting schemas. Elaborated reporting schemas will drive development of the *Backend* – the *Validation/Summarizing Services*: on the UN Forum we confirmed suitability and feasibility of Explainable Artificial Intelligence (XAI) technology for (1) validation of ESG KPIs against historical data, and (2) summarization of ESG KPI data on the *Investor UI* with explanations. This step will require training of XAI models for two distinct purposes. First, for the *Validation Services*, models will be trained on publicly available data for all the reporting schemas to ensure quality of input data. Second, for the *Summarizing Services*, models will be trained on investor data to enable the system to identify responsible investment, based on ESG KPIs.

6 Conclusion

This paper addressed the critical challenge of trust in ESG reporting by financial institutions, a problem fueled by a lack of transparency, comparability, and verifiability. Following two DSR iterations, we designed, developed, and evaluated a blockchain-based system aimed at restoring integrity to the ESG reporting process. Blockchain's core features – decentralization, traceability, immutability, and smart contracts – offer a strong technical basis for enhancing ESG transparency. Leveraging these can address limitations in traditional disclosure systems and support more reliable ESG implementation.

More specifically, our research yields several key contributions. The resulting artefact demonstrates the technical feasibility of leveraging blockchain's inherent features – such as immutability and transparency – to create a more reliable reporting infrastructure. More importantly, our primary theoretical contribution is a socio-technical

model of trust. We found that technology alone is insufficient to build institutional trust; rather, trust emerges from a symbiotic relationship between Technological Guarantees provided by the artefact and Regulatory Legitimacy derived from established frameworks. This model is instantiated in a set of three generalizable DPs: (1) Compliant submission of raw data, (2) Public access to an immutable yet correctable record, and (3) Industry-specific and materiality-driven contextualization.

For practitioners, our work provides not only a prototype architecture but also a clear guide of more credible and effective ESG reporting solutions. While acknowledging the limitations of our prototype's current scope, this research lays the foundation for future work, particularly in expanding reporting schemas and integrating advanced AI technologies for automated validation. Ultimately, this research demonstrates that rebuilding trust in sustainability reporting is not merely a technical challenge, but a holistic design that must harmoniously weave together technology, regulation, and the complex, reciprocal needs of all stakeholders.

Appendix A Auth & top up In-App Wallet

A series of screenshots illustrates the process of connecting the In-App Wallet on the Frontend and acquiring Testnet tokens – testcoins. The journey begins with the user landing on the Frontend's home page, which prompts them to connect the In-App Wallet to get started. Upon attempting to connect, a "Sign in" modal appears, offering various authentication options including Google, Apple, Facebook, email, phone number, Passkey, or connecting a wallet (an External Wallet). Choosing to connect an External Wallet leads to a selection of popular wallet providers, including MetaMask. Once an External Wallet (demonstrated with MetaMask) is successfully connected, the user's wallet details, including their address and asset balance (e.g., Celo Alfajores), are displayed within the application, along with options to send, receive, buy, view transactions, and manage or disconnect the In-App Wallet. Subsequently, the user navigates to a "Fund Your Testnet Account" page, specifically the "Alfajores Token Faucet", where they can claim testcoins. Finally, returning to the "ESG Reporting" application, the increased balance of CELO testcoins is reflected in the connected wallet display, indicating a successful top-up of the In-App Wallet's testcoin balance.

Fig. A1 In-App Wallet: the Web3 conneciton button shows that the Frontend is not currently connected.

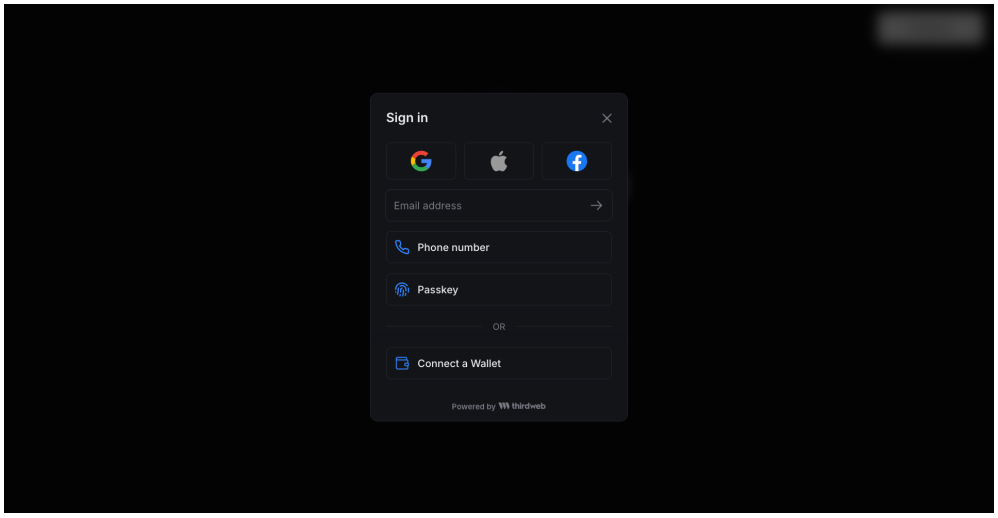

Fig. A2 In-App Wallet: multiple authentication options.

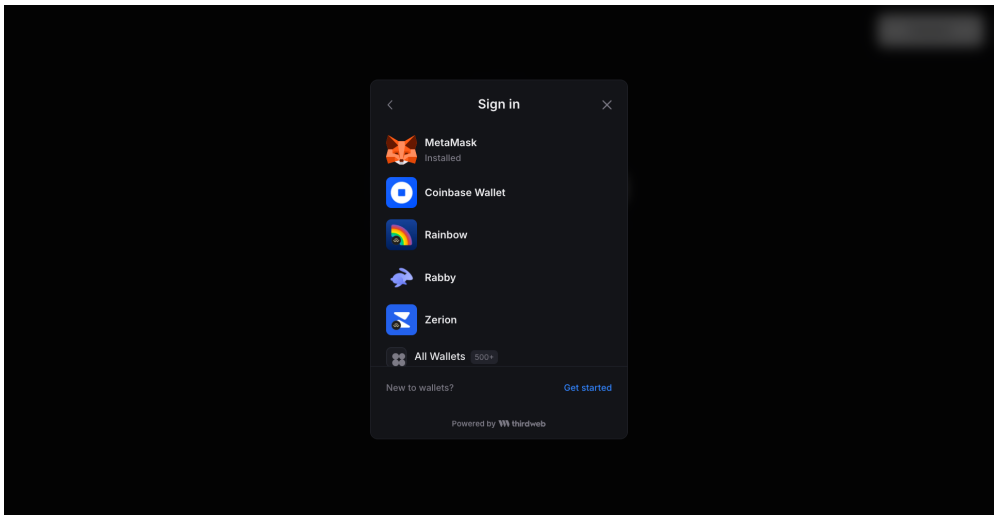

Fig. A3 In-App Wallets: authentication options via External Wallets.

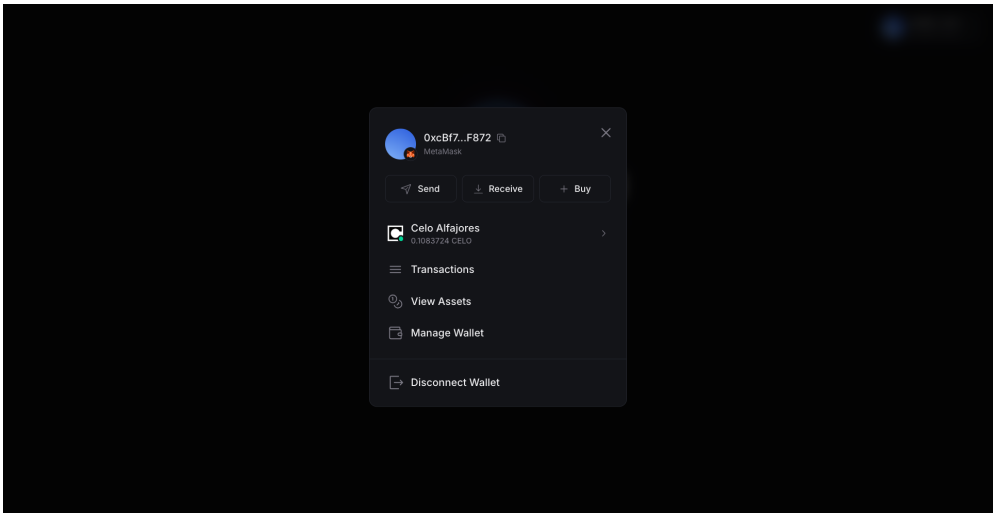

Fig. A4 In-App Wallet: authenticated via the External Wallet (MetaMask). The balance is 0.1083724 CELO testcoins.

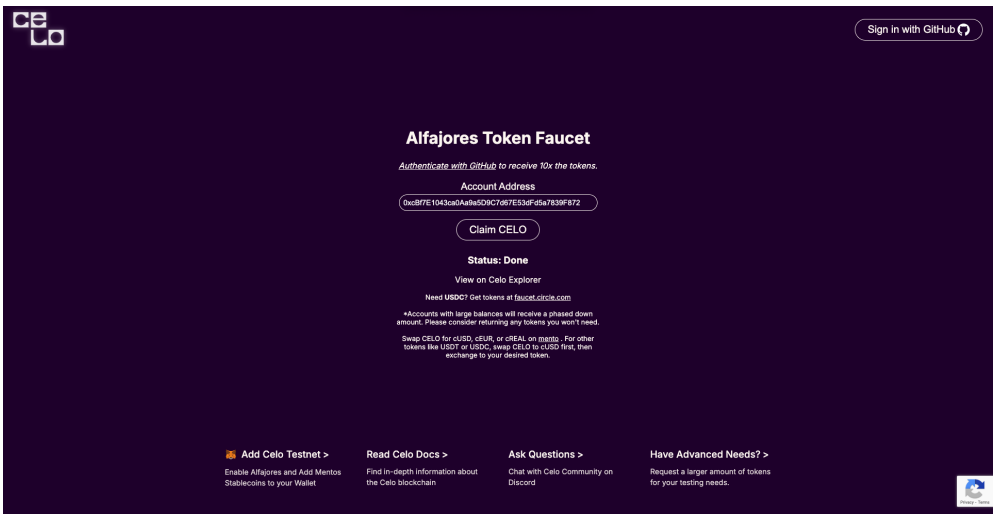

Fig. A5 Testcoin Faucet UI: claim of 0.1 CELO testcoins.

Fig. A6 In-App Wallet: the Web3 conneciton button shows that the Frontend is currently connected. The balance is 0.20837 CELO testcoins.

Appendix B ESG KPI submission & analysis

This series of screenshots illustrates the full lifecycle of ESG KPI data on a blockchain-enabled platform, from company reporting to investor analysis. It begins with a Block Explorer UI which shows the verified Smart Contract, indicating the transparent and auditable foundation for data management. Next, the Company UI is shown, where companies can input their ESG KPI data. Following this, a confirmation screen appears, confirming that the ESG KPI data has been successfully submitted by the company. The process then returns to the Block Explorer UI, showcasing the processed submission of ESG KPI data with decoded input, which demonstrates that the data is immutably recorded on-chain and its contents are publicly verifiable in a human-readable format. Finally, the perspective shifts to an Investor UI, first displaying a comparison of ESG KPIs across different companies, leveraging the on-chain data for transparency and benchmarking. This concludes with a detailed view within the investor interface, offering comparison details for specific ESG metrics.

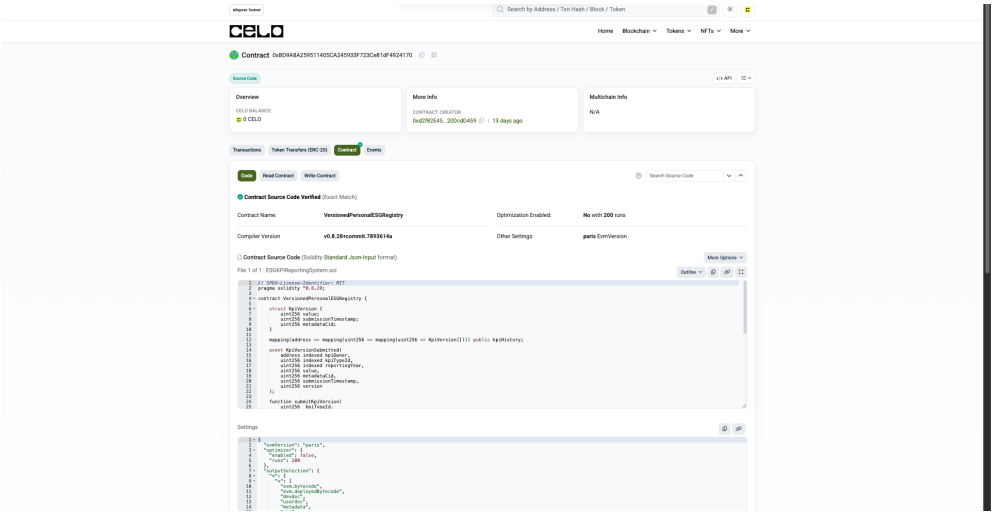

Fig. B7 Block Explorer UI: the "Code" tab of the "Contract" tab shows the verified Smart Contract, deployed on the Testnet of the Celo Blockchain.

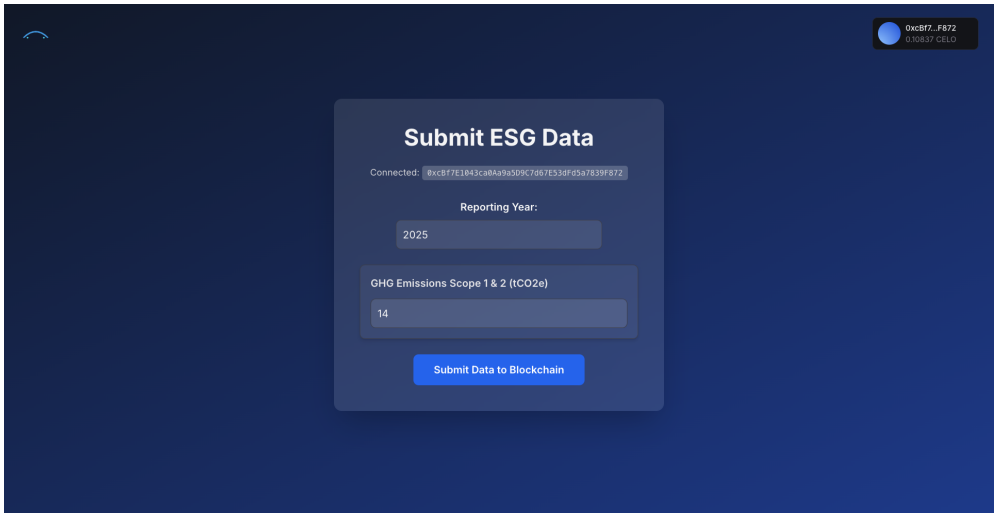

Fig. B8 Company UI: input of Carbon Dioxide Equivalent (tCO2e) ESG KPI in tons for 2025 year.

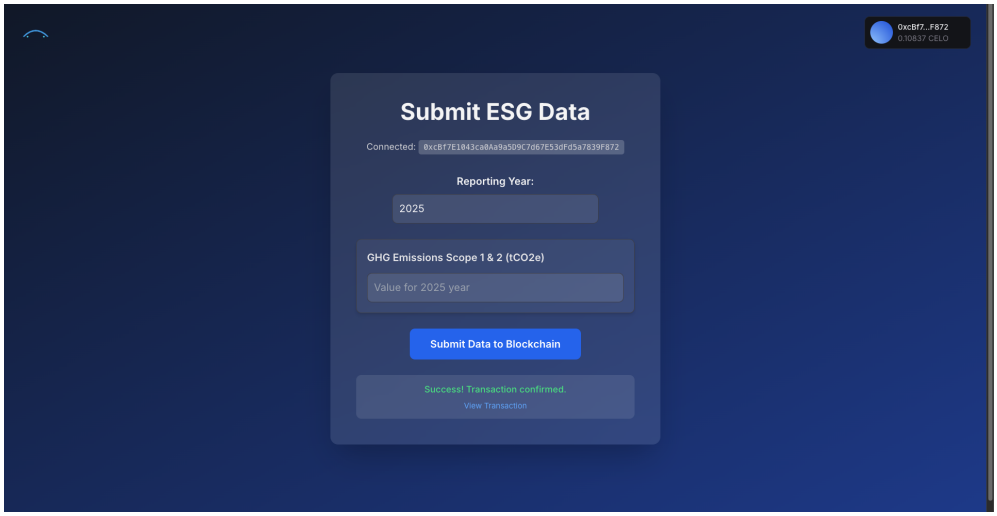

Fig. B9 Company UI: successful submission of 14 tCO2e – the TX has been already confirmed on the Blockchain.

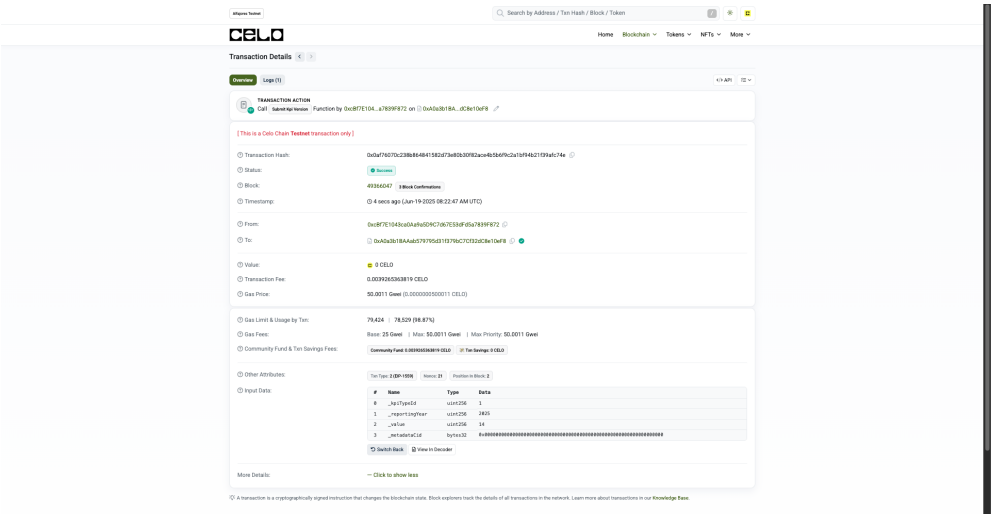

Fig. B10 Block Explorer UI: the decoded ESG KPI data entry.

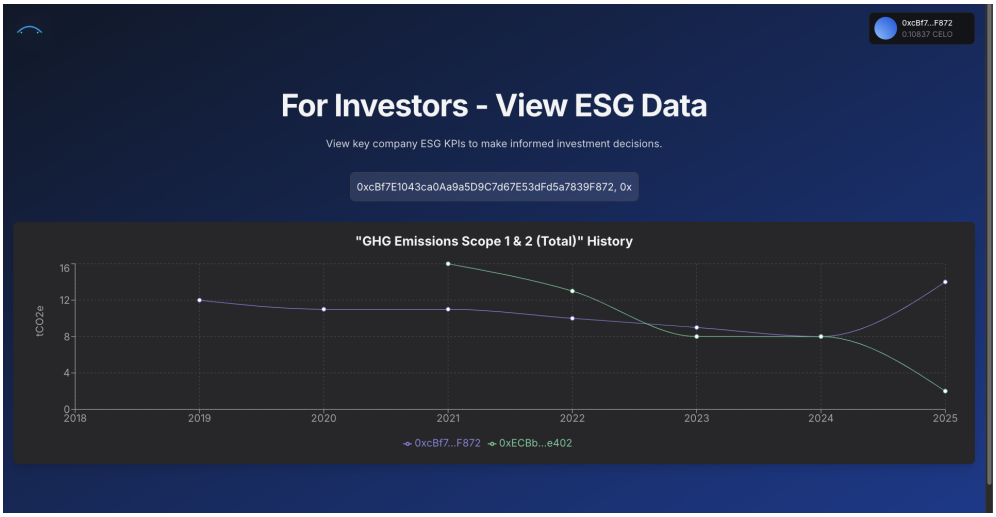

Fig. B11 Investor UI: a historical comparison of the ESG KPIs between two companies.

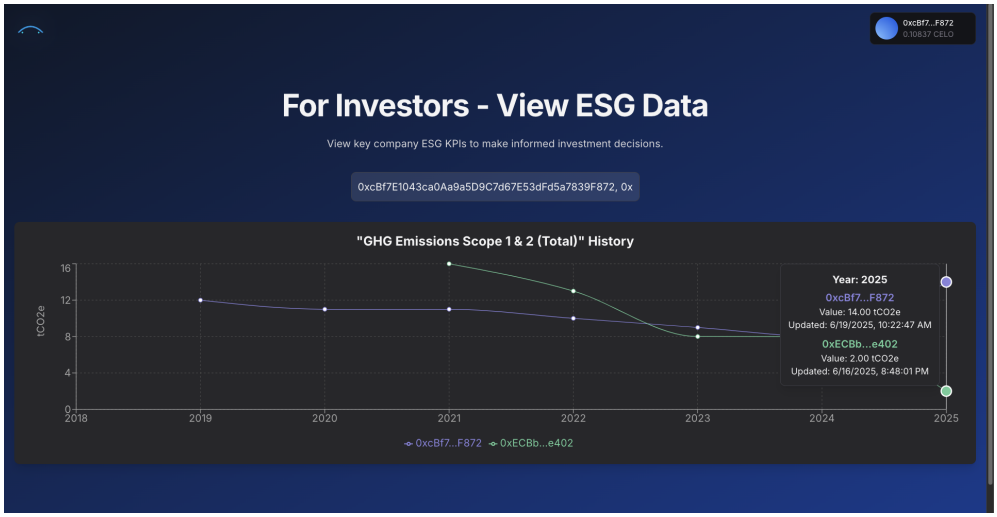

Fig. B12 Investor UI: Comparison between two companies based on two data points, including the recently submitted 14 tCO2e.

Appendix C Interview Guide

The interview guide which was prepared for the semi-structured interviews in April 2025. Attached as a PDF across 3 pages.

Introduction

Time – 30 minutes

Purpose of the Interview

- To gather insights, feedback and evaluation on the design and functionality of the blockchain-based ESG KPI reporting system (ESG KPIs are sustainability metrics).
- To understand user needs and challenges in ESG KPI reporting and compliance.
- To explore industry-specific requirements for ESG KPI selection and assessment.
- Gather insights on potential improvements and alignment with other stakeholder needs.

Context

The project was initiated during a Hackathon in Luxembourg. Throughout the iterative development process, informed by feedback from industry stakeholders, the initial prototype was developed. The insights gained from this process were summarized in a conference contribution, and we were subsequently invited to present the prototype at the United Nations' 13th International Forum on Energy for Sustainable Development. Feedback from both the academic community and a broader range of stakeholders facilitated further refinement of the artefact. We now seek to assess whether our prototype meets the requirements of financial institutions in and around Luxembourg that aim to enhance their ESG reporting.

Confidentiality

- Inform the interviewee that their responses will be anonymized.
- Consent – this interview will be recorded, transcribed, and used in the research paper.

Interview Structure

Part 1: Background Information

1. Personal and Professional Background
    a. Could you tell us about your professional background and your current role?
    b. How familiar are you with ESG reporting and blockchain technologies?

2. Organizational Context
    a. What is your organization's role in ESG reporting (e.g., reporting entity, investor, regulator)?
    b. Could you describe any challenges your organization faces in ESG reporting or compliance?

Part 2: Current Practices and Challenges

1. ESG KPI Reporting Practices
    a. How does your organization currently collect and report ESG KPIs?
    b. Are there specific frameworks or standards your organization follows (e.g., EU Taxonomy, CSRD, SFDR)?
    c. Do you use tools or platform for ESG data management and if so, which are these?
2. Challenges in ESG Reporting
    a. What challenges do you encounter in ensuring compliance with ESG reporting standards?
    b. How do you address issues like greenwashing, data validation, and standardization?

Part 3: Evaluation

1. Design Objectives
    a. Do you see ESG KPI compliance (DO1) as a key challenge for organizations?
    b. How important is reliability (DO2) in ESG reporting for stakeholders?
    c. Submitted ESG KPIs should be *publicly accessible (DO4) (transparency)* for anyone online at any point in time from the public ledger.
    d. Would standardized ESG KPI formats (DO4) improve reporting efficiency among industries?
2. Design Requirements
    a. How useful is the ability to select relevant ESG KPIs (DR2) for a company?
    b. What are your thoughts on immutability of submitted ESG data (DR3)?
    c. Does public accessibility (DR4) of ESG KPI data create value or concerns for your organization?
    d. Would predefined ESG KPI types and values (DR5) simplify ESG reporting process?
3. Design Principles

    a. How do you perceive the automatic compliance enforcement (DP1) when submitting ESG data?
    b. Does public access to immutable ESG data (DP2) increase trust or introduce risks for companies?
    c. How beneficial is industry-specific KPI reporting (DP3) for ensuring relevant sustainability goals?
4. Specific Features
    a. To what extent do you think the use of blockchain enhances trust in ESG KPI reporting?
    b. Do you see any risks or limitations in relying on blockchain for ensuring trust in ESG disclosures?
    c. How important is institutional trust in ESG reporting, and do you think our system supports or challenges existing trust structures?
5. User Experience
    d. How easy or difficult do you think it would be for your organization to adopt this system?
    e. Are there specific features or functionalities that you would like to see added?
6. Broader Applications
    f. Would a system like this influence trust in ESG disclosures among different stakeholders, such as investors, regulators, and the public?
    g. How do you think regulatory bodies or industry associations might perceive the trustworthiness of this system?
    h. What factors, beyond technological transparency, contribute to institutional trust in ESG reporting systems?

Closing

Wrap-Up

1. Summary of key points discussed during the interview.
2. Inquiry whether there is anything the interviewee would like to add.

Next Steps

1. Inform the interviewee about the next steps in the project and how their feedback will be incorporated.
2. Thank them for their time and insights.

Appendix D Acronyms

ABI	application binary interface
API	application programming interface
CSRD	Corporate Sustainability Reporting Directive
DO	design objective
DP	design principle
DR	design requirement
DSR	Design Science Research
DSRM	Design Science Research Methodology
EOA	externally owned account
ESG	environmental, social, and governance
EU	European Union
EU Taxonomy	EU Taxonomy Regulation
EVM	Ethereum virtual machine
GDPR	General Data Protection Regulation
GHG	greenhouse gas
IGO	intergovernmental organization
IoT	Internet of things
IPFS	InterPlanetary File System
IT	information technology
KPI	key performance indicator
ML	Machine Learning
MPC	Multi-Party Computation
MRL	Machine-Readable Legislation
PoS	proof-of-stake
PoW	proof-of-work
SDG	Sustainable Development Goal
SFDR	Sustainable Finance Disclosure Regulation
SME	small and medium-sized enterprise
TX	transaction
UI	user interface
UML	Unified Modeling Language
UN	United Nations
UX	user experience
XAI	Explainable Artificial Intelligence
ZKP	Zero-Knowledge Proof

Funding Declaration

This study was funded by the Luxembourg National Research Fund.

References

2019/2088: Regulation (EU) 2019/2088 of the European Parliament and of the Council of 27 November 2019 on sustainability-related disclosures in the financial services sector (Text with EEA relevance) (2019). https://eur-lex.europa.eu/eli/reg/2019/2088/oj Accessed July 18, 2025

2020/852: Regulation (EU) 2020/852 of the European Parliament and of the Council of 18 June 2020 on the establishment of a framework to facilitate sustainable investment, and amending Regulation (EU) 2019/2088 (2020). https://eur-lex.europa.eu/eli/reg/2020/852/oj Accessed July 18, 2025

2022/2464: Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 amending Regulation (EU) No 537/2014, Directive 2004/109/EC, Directive 2006/43/EC and Directive 2013/34/EU, as regards corporate sustainability reporting (2022). https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022L2464 Accessed July 18, 2025

2023/2772: Commission Delegated Regulation (EU) 2023/2772 of 31 July 2023 supplementing Directive 2013/34/EU of the European Parliament and of the Council as regards sustainability reporting standards (2023). https://eur-lex.europa.eu/eli/reg_del/2023/2772/oj/eng Accessed July 18, 2025

2024/1760: Directive (EU) 2024/1760 of the European Parliament and of the Council of 13 June 2024 on corporate sustainability due diligence and amending Directive (EU) 2019/1937 and Regulation (EU) 2023/2859 (2024). https://eur-lex.europa.eu/eli/dir/2024/1760/oj/eng Accessed July 18, 2025

AbuRaya, R.: Corporate Environmental Disclosure and Corporate Governance: A Critical Review. Journal of Empirical Research in Accounting; Auditing An International Journal **04**, 24–53 (2017)

Abriani, N., Catania, A.: In: Marano, P., Noussia, K. (eds.) Corporate Governance and the So-Called 'Four-Eyes Principle', pp. 3–24. Springer, ??? (2022). https://doi.org/10.1007/978-3-030-85817-9

Aggarwal, P., Kadyan, A.: Greenwashing: The Darker Side Of CSr. Indian Journal of Applied Research **4**, 61–66 (2011)

Asif, M., Searcy, C., Castka, P.: ESG and Industry 5.0: The role of technologies in enhancing ESG disclosure. Technological Forecasting and Social Change **195**, 122806 (2023)

Aggarwal, P., Singh, A., Malhotra, D.: Role of Blockchain in Strengthening ESG Reporting: A Systematic Review and Directions for Future Research. CPJ Law Journal **14**, 486–499 (2023)

Amend, J., Troglauer, P., Guggenberger, T., Urbach, N., Weibelzahl, M.: Facilitating cooperation of smallholders in developing countries: design principles for a cooperative-oriented decentralized autonomous organization. Information Systems and e-Business Management **22**(1), 1–31 (2024)

Australian Accounting Standards Board, Australian Government: Australian Sustainability Reporting Standard (2024). https://standards.aasb.gov.au/aasb-s1-sep-2024 Accessed July 18, 2025

Awan, U.: Green Marketing: Marketing Strategies for the Swedish Energy Companies. International Journal of Industrial Marketing **1**, 1 (2011)

Amel-Zadeh, A., Serafeim, G.: Why and how investors use esg information: Evidence from a global survey. Financial Analysts Journal **74**(3), 87–103 (2018)

Balzani, L., Corsi, K.: Blockchain and sustainability disclosure: reliable information on renewable energies. In: In Proceedings of 19th Annual Conference of the Italian Chapter of AIS (ITAIS) (2022). https://aisel.aisnet.org/itais2022/14

Bell, D.: Uml basics: The component diagram. IBM Global Services, 1–10 (2004)

Bloomberg: Global ESG assets predicted to hit $40 trillion by 2030, despite challenging environment, forecasts Bloomberg Intelligence (2024). https://www.bloomberg.com/company/press/global-esg-assets-predicted-to-hit-40-trillion-by-2030-despite-challenging-environment-forecasts-bloomberg-inte Accessed July 18, 2025

Briand, L.C., Morasca, S., Basili, V.R.: Defining and validating measures for object-based high-level design. IEEE Transactions on Software Engineering **25**(5), 722–743 (1999)

Buallay, A.M., Marri, M.A., Nasrallah, N., Hamdan, A., Barone, E., Zureigat, Q.: Sustainability reporting in banking and financial services sector: a regional analysis. Journal of Sustainable Finance & Investment **13**(1), 776–801 (2023)

Brocke, J.v., Simons, A., Niehaves, B., Riemer, K., Plattfaut, R., Cleven, A.: Reconstructing the giant: On the importance of rigour in documenting the literature search process. (2009)

Brocke, J.v., Simons, A., Riemer, K., Niehaves, B., Plattfaut, R., Cleven, A.: Standing on the shoulders of giants: Challenges and recommendations of literature search in information systems research. Communications of the Association for Information Systems **37** (2015) https://doi.org/10.17705/1CAIS.03709

Cabaleiro-Cerviño, G., Mendi, P.: ESG-driven innovation strategy and firm performance. Eurasian Business Review **14**(1), 137–185 (2024)

Chen, R.R., Chen, K., Ou, C.X.J.: Facilitating interorganizational trust in strategic alliances by leveraging blockchain-based systems: Case studies of two eastern banks. International Journal of Information Management **68**, 102521 (2023) https://doi.org/10.1016/j.ijinfomgt.2022.102521

Celo: Add Celo to MetaMask (2024). https://docs.celo.org/wallet/metamask/add-celo-testnet-to-metamask Accessed July 18, 2025

Celo: Alfajores Token Faucet (2024). https://faucet.celo.org/alfajores Accessed July 18, 2025

Cruz, C.A., Matos, F.: ESG Maturity: A Software Framework for the Challenges of ESG Data in Investment. Sustainability **15**(3), 2610 (2023)

Chopra, S., Senadheera, S., Dissanayake, P., Withana, P., Chib, R., Rhee, J., Ok, Y.S.: Navigating the Challenges of Environmental, Social, and Governance (ESG) Reporting: The Path to Broader Sustainable Development. Sustainability **16**, 606 (2024)

Cantero-Saiz, M., Polizzi, S., Scannella, E.: ESG and asset quality in the banking industry: The moderating role of financial performance. Research in International Business and Finance **69**, 102221 (2024)

Casey, M.J., Vigna, P.: In blockchain we trust (2018). https://www.technologyreview.com/2018/04/09/3066/in-blockchain-we-trust Accessed July 18, 2025

Chueca Vergara, C., Ferruz Agudo, L.: Fintech and Sustainability: Do They Affect Each Other? Sustainability **13**(13), 7012 (2021)

Chen, W., Wu, W., Ouyang, Z., Fu, Y., Li, M., Huang, G.Q.: Event-based data authenticity analytics for IoT and blockchain-enabled ESG disclosure. Computers & Industrial Engineering **190**, 109992 (2024)

Delmas, M.A., Burbano, V.C.: The Drivers of Greenwashing. California Management Review **54**(1), 64–87 (2011)

Department for Business and Trade, Government of the United Kingdom: UK Sustainability Reporting Standards (2024). https://www.gov.uk/guidance/uk-sustainability-reporting-standards Accessed July 18, 2025

De Freitas Netto, S.V., Sobral, M.F.F., Ribeiro, A.R.B., Soares, G.R.d.L.: Concepts and forms of greenwashing: a systematic review. Environmental Sciences Europe **32**(1) (2020)

Directorate-General for Financial Stability, Financial Services and Capital Markets Union, European Comission: EU taxonomy for sustainable activities (2020). https://finance.ec.europa.eu/sustainable-finance/tools-and-standards/

eu-taxonomy-sustainable-activities_en Accessed July 18, 2025

Directorate-General for Communication, European Commission: Commission proposes to cut red tape and simplify business environment (2025). https://commission.europa.eu/news-and-media/news/commission-proposes-cut-red-tape-and-simplify-business-environment-2025-02-26_en Accessed July 18, 2025

Directorate-General for Financial Stability, Financial Services and Capital Markets Union, European Comission: Commission simplifies rules on sustainability and EU investments, delivering over €6 billion in administrative relief (2025). https://finance.ec.europa.eu/publications/commission-simplifies-rules-sustainability-and-eu-investments-delivering-over-eu6-billion_en Accessed July 18, 2025

Directorate-General for Financial Stability, Financial Services and Capital Markets Union, European Comission: Omnibus package (2025). https://finance.ec.europa.eu/news/omnibus-package-2025-04-01_en Accessed July 18, 2025

De Filippi, P., Mannan, M., Reijers, W.: Blockchain as a confidence machine: The problem of trust & challenges of governance. Technology in Society **62**, 101284 (2020)

Dye, J., McKinnon, M., Byl, C.: Green Gaps: Firm ESG Disclosure and Financial Institutions' Reporting Requirements. Journal of Sustainability Research **3**(1) (2021)

Driessen, M.: Sustainable Finance: An Overview of ESG in the Financial Markets, pp. 329–350. Springer, ??? (2021). https://doi.org/10.1007/978-3-030-71834-3_10

Duran, R.E., Tierney, P.: Fintech Data Infrastructure for ESG Disclosure Compliance. Journal of Risk and Financial Management **16**(8), 378 (2023)

Di Vaio, A., Varriale, L.: Blockchain technology in supply chain management for sustainable performance: Evidence from the airport industry. International Journal of Information Management **52**, 102014 (2020) https://doi.org/10.1016/j.ijinfomgt.2019.09.010

Economic and Social Commission for Asia and the Pacific, United Nations: Capacity-building workshop on digitalization in energy (2024). https://unescap.org/sites/default/d8files/2024-12/Concept%20note_Digitalization_IFESD%202024_v.3_0.pdf Accessed July 18, 2025

Economic and Social Commission for Asia and the Pacific, United Nations: Thirteenth International Forum on Energy for Sustainable Development (2024). https://www.unescap.org/events/2024/thirteenth-international-forum-energy-sustainable-development Accessed July 18,

2025

Envoria: ESRS Data Points: Guide for Successful CSRD Reporting (2025). https://envoria.com/insights-news/esrs-data-points-guide-for-successful-csrd-reporting Accessed July 18, 2025

European Federation of Financial Analyst Societies: KPIs for ESG (2010). https://effas.com/wp-content/uploads/2021/09/KPIs_for_ESG_3_0_Final.pdf Accessed July 18, 2025

European Commission: FinTech action plan: For a more competitive and innovative European financial sector (2018). https://finance.ec.europa.eu/publications/fintech-action-plan-more-competitive-and-innovative-european-financial-sector_en Accessed July 18, 2025

European Commission: Strategy for financing the transition to a sustainable economy (2021). https://finance.ec.europa.eu/publications/strategy-financing-transition-sustainable-economy_en Accessed July 18, 2025

European Commission: Overview of sustainable finance (2024). https://finance.ec.europa.eu/sustainable-finance/overview-sustainable-finance_en Accessed July 18, 2025

European Parliament: The Internal Market: General Principles (2024). https://www.europarl.europa.eu/factsheets/en/sheet/33/the-internal-market-general-principles Accessed July 18, 2025

European Supervisory Authorities: ESAs call for enhanced supervision and improved market practice on sustainability-related claims (2024). https://www.eba.europa.eu/publications-and-media/press-releases/esas-call-enhanced-supervision-and-improved-market-practice-sustainability-related-claims Accessed July 18, 2025

Franzoni, F., Abellan, I., Daza, V.: In: Bonneau, J., Heninger, N. (eds.) Leveraging Bitcoin Testnet for Bidirectional Botnet Command and Control Systems, pp. 3–19. Springer, ??? (2020). https://doi.org/10.1007/978-3-030-51280-4_1

Foley, A.M., Heffron, R.J., Al Kez, D., Furszyfer Del Rio, D.D., McInerney, C., Welfle, A.: Restoring trust in ESG investing through the adoption of just transition ethics. Renewable and Sustainable Energy Reviews **199**, 114557 (2024)

Gu, Y., Dai, J., Vasarhelyi, M.A.: Audit 4.0-based ESG assurance: An example of using satellite images on GHG emissions. International Journal of Accounting Information Systems **50**, 100625 (2023)

Goo, J., Huang, C.D.: Facilitating relational governance through service level agreements in IT outsourcing: An application of the commitment–trust theory. Decision

Support Systems **46**(1), 216–232 (2008)

Gregor, S., Hevner, A.: Positioning and Presenting Design Science Research for Maximum Impact. MIS Quarterly **37**(2), 337–355 (2013)

Gharpure, A.: ESG Reporting in Enterprise Financial Systems: Challenges and Innovations. Journal of Computer Science and Technology Studies **7**, 194–201 (2025)

Gramlich, V., Jelito, D., Sedlmeir, J.: Maximal extractable value: Current understanding, categorization, and open research questions. Electronic Markets **34**(1), 49 (2024)

Gramlich, V., Körner, M.-F., Ströher, T., Strüker, J., Volland, M.: In: Fridgen, G., Guggenberger, T., Sedlmeir, J., Urbach, N. (eds.) Decentralization Technologies in the Context of ESG Accounting and Reporting, pp. 177–194. Springer, ??? (2024). https://doi.org/10.1007/978-3-031-66047-4_10

Goo, J., Nam, K.: Contract as a Source of Trust–Commitment in Successful IT Outsourcing Relationship: An Empirical Study. In: In Proceedings of 40th Annual Hawaii International Conference on System Sciences (HICSS) (2007). https://doi.org/10.1109/HICSS.2007.148

Galeone, G., Ranaldo, S., Fusco, A.: ESG and FinTech: Are they connected? Research in International Business and Finance **69**, 102225 (2024)

García-Sánchez, I.-M., Hussain, N., Khan, S.-A., Martínez-Ferrero, J.: Assurance of corporate social responsibility reports: Examining the role of internal and external corporate governance mechanisms. Corporate Social Responsibility and Environmental Management **29**(1), 89–106 (2022)

Gong, X., Tao, X., Das, M., Kwok, H.H.L., Cheng, J.C.P.: In: Capone, P., Vito, G., Rahimian, F.P., Dawood, N., Bruttini, A., Sorbi, T. (eds.) Integrating ESG Factors into Construction Projects: A Blockchain-Based Data Management Approach, pp. 327–334. Firenze University Press, ??? (2023). http://dx.doi.org/10.36253/979-12-215-0289-3.31

Gong, X., Tao, X., Zhang, M., Xu, Y., Kwok, H.H.L., Dai, J., Cheng, J.C.P.: Secure environmental, social, and governance (ESG) data management for construction projects using blockchain. Sustainable Cities and Society **114**, 105582 (2024)

Hevner, A., Chatterjee, S.: Design Science Research in Information Systems. Springer, ??? (2010). https://doi.org/10.1007/978-1-4419-5653-8_2

Ho, J.: Banking, Blockchain, and ESG, pp. 235–266. Cambridge University Press, ??? (2023). https://doi.org/10.1017/9781009411783

Hang, L., Wang, L., Xu, P.: The Impact of Corporate ESG Performance on Corporate Green Innovation. Highlights in Business, Economics and Management **22**, 399–404 (2023)

J, K.: What Is Celo? A Mobile-First Carbon Negative Layer 1 Turned Layer 2 (2024). https://www.coingecko.com/learn/celo Accessed July 18, 2025

Jones, E.: Rethinking Greenwashing: Corporate Discourse, Unethical Practice, and the Unmet Potential of Ethical Consumerism. Sociological Perspectives **62**(5), 728–754 (2019)

Keynes, J.M.: The General Theory of Employment, Interest and Money. Macmillan, ??? (1936). https://www.bibsonomy.org/bibtex/2c27c51d717b88bdf405c38cbfdfa65ac/butz

Kim, J., Kim, M., Im, S., Choi, D.: Competitiveness of E Commerce Firms through ESG Logistics. Sustainability **13**(20), 11548 (2021)

Krause, D.: In: Zarifis, A., Cheng, X. (eds.) The Rise of Web3: Opportunities and Challenges, pp. 255–266. Springer, ??? (2025). https://doi.org/10.1007/978-3-031-83402-8_9

Kshetri, N.: Blockchain and sustainable supply chain management in developing countries. International Journal of Information Management **60**, 102376 (2021) https://doi.org/10.1016/j.ijinfomgt.2021.102376

Legislative State Bureau, State of California: Climate Corporate Data Accountability Act (2023). https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=202320240SB253 Accessed July 18, 2025

Lindgren, P., Knoth, N.S.H., Sureshkumar, S., Friedrich, M.F., Adomaityte, R.: "Green Multi Business Models" How to measure Green Business Models and Green Business Model Innovation? Wireless Personal Communications **121**, 1303–1323 (2021)

Lyon, T.P., Montgomery, A.W.: The Means and End of Greenwash. Organization & Environment **28**(2), 223–249 (2015)

Liu, L., Ma, Z., Zhou, Y., Fan, M., Han, M.: Trust in ESG reporting: The intelligent Veri-Green solution for incentivized verification. Blockchain: Research and Applications **5**(2), 100189 (2024)

Luo, Y., Shen, J., Liang, H., Sun, L., Dong, L.: Supporting building life cycle carbon monitoring, reporting and verification: A traceable and immutable blockchain-empowered information management system and application in Hong Kong. Resources, Conservation and Recycling **208**, 107736 (2024)

Liu, X., Wu, H., Wu, W., Fu, Y., Huang, G.Q.: Blockchain-Enabled ESG Reporting Framework for Sustainable Supply Chain. In: In Proceedings of 7th International Conference on Sustainable Design and Manufacturing (SDM)) (2021). https://doi.org/10.1007/978-981-15-8131-1_36

Li, L., Wang, M., Zhou, X.: Creating value beyond commercial outcomes: The esg practices of online marketplaces for sustainable development. Electronic Markets **33**(1), 62 (2023)

Liu, X., Yang, Y., Jiang, Y., Fu, Y., Zhong, R.Y., Li, M., Huang, G.Q.: Data-driven ESG assessment for blockchain services: A comparative study in textiles and apparel industry. Resources, Conservation and Recycling **190**, 106837 (2023)

Mugurusi, G., Ahishakiye, E.: Blockchain technology needs for sustainable mineral supply chains: A framework for responsible sourcing of Cobalt. Procedia Computer Science **200**, 638–647 (2022)

Markopoulos, E., Al Katheeri, H., Al Qayed, H.: A decision support system architecture for the development and implementation of ESG strategies at SMEs. In: In Proceedings of the 6th International Conference on Intelligent Human Systems Integration (IHSI) (2023). https://doi.org/10.54941/ahfe1002916

Miranda, Y., Alves, P., Paskin, R., Nasser, R., Robichez, G., Faria, L., Trindade, R., Silva, J., Peixoto, L., Miranda, F.: Enhancing Corporate Social Responsibility with Blockchain-based Trackable ESG Tokens. In: In Proceedings of 6th Blockchain Workshop: Theory, Technology and Applications (WBlockchain) (2023). https://doi.org/10.5753/wblockchain.2023.777

McKnight, D.H., Chervany, N.L.: In: Falcone, R., Singh, M., Tan, Y.-H. (eds.) Trust and Distrust Definitions: One Bite at a Time, pp. 27–54. Springer, ??? (2001). https://doi.org/10.1007/3-540-45547-7_3

McKnight, D.H., Cummings, L.L., Chervany, N.L.: Initial Trust Formation in New Organizational Relationships. The Academy of Management Review **23**(3), 473 (1998)

McKnight, D.H., Carter, M., Clay, P.: Trust in technology: development of a set of constructs and measures. In: In Proceedings of 14th Diffusion Interest Group In Information Technology Workshop (DIGIT) (2009). https://aisel.aisnet.org/digit2009/10

McKnight, D.H., Carter, M., Thatcher, J., Clay, P.: Trust in a specific technology: An Investigation of its Components and Measures. ACM Transactions on Management Information Systems **2**(2), 12–32 (2011)

Meta Open Source: React – A JavaScript library for building user interfaces (2024). https://react.dev Accessed July 18, 2025

Moody, G.D., Galletta, D., Lowry, P.B.: Unifying Conflicting Models of Trust and Distrust for Enhanced Understanding and Predictive Power in Organizational Relationships: Proposing the Unified Trust-Distrust Model (UTDM). SSRN Electronic Journal (2013)

Miles, M.B., Huberman, A.M., Saldana, J.: Qualitative Data Analysis: A Methods Sourcebook. SAGE Publications, ??? (2018). https://books.google.lu/books?id=fjh2DwAAQBAJ

Milosavljevic, A.: 4 Eyes Principle: Effective Governance Management (2023). https://get-newton.com/the-4-eyes-principle Accessed July 23, 2025

Murray, A., Kim, D., Combs, J.: The promise of a decentralized internet: What is Web3 and how can firms prepare? Business Horizons **66**(2), 191–202 (2023)

Moody, G.D., Lowry, P.B., Galletta, D.F.: It's complicated: explaining the relationship between trust, distrust, and ambivalence in online transaction relationships using polynomial regression analysis and response surface analysis. European Journal of Information Systems **26**(4), 379–413 (2017)

McKnight, D.H., Lankton, N.K., Nicolaou, A., Price, J.: Distinguishing the effects of B2B information quality, system quality, and service outcome quality on trust and distrust. The Journal of Strategic Information Systems **26**(2), 118–141 (2017)

Mandas, M., Lahmar, O., Piras, L., De Lisa, R.: ESG in the financial industry: What matters for rating analysts? Research in International Business and Finance **66**, 102045 (2023)

Müller, L.S., Nohe, C., Reiners, S., Becker, J., Hertel, G.: Adopting information systems at work: a longitudinal examination of trust dynamics, antecedents, and outcomes. Behaviour & Information Technology **43**(6), 1096–1128 (2024) https://doi.org/10.1080/0144929X.2023.2196598 https://doi.org/10.1080/0144929X.2023.2196598

Macchiavello, E., Siri, M.: Sustainable Finance and Fintech: Can Technology Contribute to Achieving Environmental Goals? A Preliminary Assessment of 'Green FinTech'. SSRN Electronic Journal (2020)

Mendling, J., Weber, I., Aalst, W.V.D., Brocke, J.V., Cabanillas, C., Daniel, F., Debois, S., Ciccio, C.D., Dumas, M., Dustdar, S., Gal, A., García-Bañuelos, L., Governatori, G., Hull, R., Rosa, M.L., Leopold, H., Leymann, F., Recker, J., Reichert, M., Reijers, H.A., Rinderle-Ma, S., Solti, A., Rosemann, M., Schulte, S., Singh, M.P., Slaats, T., Staples, M., Weber, B., Weidlich, M., Weske, M., Xu, X., Zhu, L.: Blockchains for Business Process Management - Challenges and Opportunities. ACM Transactions on Management Information Systems **9**(1), 1–16 (2018)

Nakamoto, S.: Bitcoin: A Peer-to-Peer Electronic Cash System (2008). https://bitcoin.

org/bitcoin.pdf Accessed July 18, 2025

Nugroho, D.P.D., Hsu, Y., Hartauer, C., Hartauer, A.: Investigating the Interconnection between Environmental, Social, and Governance (ESG), and Corporate Social Responsibility (CSR) Strategies: An Examination of the Influence on Consumer Behavior. Sustainability **16**(2), 614 (2024)

Nomic Foundation: Hardhat Documentation (2024). https://hardhat.org/docs Accessed July 18, 2025

Pizzi, S., Caputo, A., Venturelli, A., Caputo, F.: Embedding and managing blockchain in sustainability reporting: A practical framework. Sustainability Accounting, Management and Policy Journal **13**(3), 545–567 (2022)

Pincheira, M., Donini, E., Giaffreda, R., Vecchio, M.: A Blockchain-Based Approach To Enable Remote Sensing Trusted Data. In: In Proceedings of IEEE Latin American GRSS & ISPRS Remote Sensing Conference (LAGIRS) (2020). https://doi.org/10.1109/LAGIRS48042.2020.9165589

Park, Y.-n., Han, S.-L.: The Effect of ESG Activities on Corporate Image, Perceived Price Fairness, and Consumer Responses. Korean Management Review **50**, 643–664 (2021)

Prabawani, B., Hadi, S.P.: Sustainability Indicator: An Initial Parameter for Convenience Product. Jurnal Presipitasi : Media Komunikasi dan Pengembangan Teknik Lingkungan **19**, 179–189 (2022)

Peffers, K., Tuunanen, T., Rothenberger, M., Chatterjee, S.: A Design Science Research Methodology for Information Systems Research. Journal of Management Information Systems **24**(3), 45–77 (2007)

Raskin, M.: The law and legality of smart contracts. Georgetown Law Technology Review **1**(2), 305–341 (2016)

Revolut: CELO to EUR: Convert Celo (CELO) to Euros (EUR) – Revolut (2025). https://www.revolut.com/crypto/price/celo/eur Accessed July 18, 2025

Randazzo, R., Perozzi, F.G.: EU sustainable finance and greenwashing: where are we and what Lies ahead? Business Law International **24**, 37–50 (2023)

Rieger, A., Roth, T., Sedlmeir, J., Fridgen, G.: We need a broader debate on the sustainability of blockchain. Joule **6**(6), 1137–1141 (2022)

Rozas, D., Tenorio-Fornés, A., Díaz-Molina, S., Hassan, S.: When Ostrom Meets Blockchain: Exploring the Potentials of Blockchain for Commons Governance. Sage Open **11**(1), 21582440211002526 (2021)

Saldaña, J.: The Coding Manual for Qualitative Researchers. SAGE Publications Ltd,

??? (2022). https://psycnet.apa.org/record/2009-06064-000

Savelyev, A.: Contract Law 2.0: 'Smart' Contracts As the Beginning of the End of Classic Contract Law. Information & Communications Technology Law **26**, 116–134 (2017)

Sedlmeir, J., Buhl, H., Fridgen, G., Keller, R.: The Energy Consumption of Blockchain Technology: Beyond Myth. Business & Information Systems Engineering **62**, 599–608 (2020)

Secretariat-General, European Commission: Omnibus I - simplification package of legislative frameworks and sustainability rules (2025). https://commission.europa.eu/publications/omnibus-i_en Accessed July 18, 2025

Smits, M., Hulstijn, J.: Blockchain Applications and Institutional Trust. Frontiers in Blockchain **3**, 5 (2020)

Ströher, T., Körner, M.-F., Paetzold, F., Strüker, J.: Bridging carbon data's organizational boundaries: toward automated data sharing in sustainable supply chains. Electronic Markets **35**(1), 1–22 (2025)

Scholl, H., Pomeshchikov, R., Rodríguez Bolívar, M.P.: Early Regulations of Distributed Ledger Technology/Blockchain Providers: A Comparative Case Study. In: In Proceedings of 53rd Hawaii International Conference on System Sciences (HICSS) (2020). https://doi.org/10.24251/hicss.2020.218

Sousa, L.M., Viana, D.C., Neto, A.P.d.L., Castro, Z.R., Aguiar, G.Q.M.d., Silva, I.R.D.: The evolutions achieved in companies with the implementation of Environmental, Social and Governance: Integrative review. International Journal of Business, Economics and Management **10**, 44–53 (2023)

Swan, M.: Blockchain: Blueprint for a New Economy. O'Reilly Media, Inc., ??? (2015). https://dl.acm.org/doi/book/10.5555/3006358

Shen, H., Wu, H., Chand, P.: The impact of corporate social responsibility assurance on investor decisions: Chinese evidence. International Journal of Auditing **21**(3), 271–287 (2017)

Szabo, N.: Smart Contracts (1994). https://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart.contracts.html Accessed July 18, 2025

Truant, E., Borlatto, E., Crocco, E., Bhatia, M.: ESG performance and technological change: Current state-of-the-art, development and future directions. Journal of Cleaner Production **429**, 139493 (2023)

Terrachoice: The Sins of Greenwashing (2011). https://twosides.info/UK/

terrachoice-the-sins-of-greenwashing-home-and-family-edition-2010/ Accessed July 18, 2025

thirdweb: Thirdweb Documentation (2024). https://portal.thirdweb.com Accessed July 18, 2025

thirdweb: Connect – The complete toolkit for web3 onboarding, wallets and authentication (2025). https://portal.thirdweb.com/connect Accessed July 18, 2025

Thoring, K., Mueller, R., Badke-Schaub, P.: Workshops as a research method: Guidelines for designing and evaluating artifacts through workshops. In: In Proceedings of 53rd Annual Hawaii International Conference on System Sciences (HICSS) (2020). https://dx.doi.org/10.24251/HICSS.2020.620

Utz, M., Johanning, S., Roth, T., Bruckner, T., Strüker, J.: From ambivalence to trust: Using blockchain in customer loyalty programs. International Journal of Information Management **68**, 102496 (2023)

Vasiu, D.E., Bratu, R.: An Overview On Environmental Social And Governance – Esg-topics From The Financial Markets' Perspective. Management of Sustainable Development **14**(2), 76–82 (2022)

Vercel: Next.js Documentation (2024). https://nextjs.org Accessed July 18, 2025

Wu, W., Fu, Y., Wang, Z., Liu, X., Niu, Y., Li, B., Huang, G.Q.: Consortium blockchain-enabled smart ESG reporting platform with token-based incentives for corporate crowdsensing. Comput. Ind. Eng. **172**, 108456 (2022)

Webster, J., Watson, R.T.: Analyzing the Past to Prepare for the Future: Writing a Literature Review. MIS Quarterly **26**(2), (2002)

Ying, W., Jia, S., Du, W.: Digital enablement of blockchain: Evidence from HNA group. International Journal of Information Management **39**, 1–4 (2018)

Yu, T.: In: Patnaik, S. (ed.) Blockchain Technology and the Improvement of ESG Information Transparency, vol. 56, pp. 211–219. IOS Press, ??? (2024). http://dx.doi.org/10.3233/ATDE240431

Zhan, S.: ESG and Corporate Performance: A Review. SHS Web of Conferences **169**, 01064 (2023)

Ziolkowski, R., Miscione, G., Schwabe, G.: Decision Problems in Blockchain Governance: Old Wine in New Bottles or Walking in Someone Else's Shoes? MIS Quarterly **37**(2), 316–348 (2020)

Jean, M.S., Grant, E.: Management System Enabled ESG Performance. International Pipeline Conference, vol. Volume 1: Pipeline Safety Management Systems;

Project Management, Design, Construction, and Environmental Issues; Strain-Based Design and Assessment; Risk and Reliability; Emerging Fuels and Greenhouse Gas Emissions, pp. 001–01004 (2022). https://doi.org/10.1115/IPC2022-86870 . https://doi.org/10.1115/IPC2022-86870