DOI:10.1145/3648607

A deep dive into blockchain interoperability: why it is needed, progress that has been made over the past decade, how it is currently deployed and used, and likely paths of future development.

BY RAFAEL BELCHIOR, JAN SÜßENGUTH, QI FENG, THOMAS HARDJONO, ANDRÉ VASCONCELOS, AND MIGUEL CORREIA

A Brief History of Blockchain Interoperability

BLOCKCHAIN INTEROPERABILITY CONFLATES the need for distributed systems to communicate with third-party systems without a canonical chain or orchestration layer. As there is no "chain to rule them all" (for performance, privacy, and market forces), these distributed systems rely on exchanging data and value across network boundaries. Interconnected systems achieve a higher value than the sum of their parts, similar to how the Internet emerged as a set of isolated *local area networks* (LANs)—and, by force of surprising synergies, such networks fundamentally transformed society forever. Concurrently, in the last decade, we have witnessed the astonishing development of blockchain technologies, which seem more connected than ever: via *bridges*,[14] *oracles*,[24] and other *interoperability mechanisms*.[8,26,41] These recent developments have slowly but steadily contributed to the improvement of the scalability of blockchain networks, as well as providing new functionality and use cases,[32] but there is still a long way to go until mass adoption. In this article, we dive into the rabbit hole of blockchain interoperability and explain why it is needed, what has work been done in the last decade (the past), how it is currently deployed and used in practice (the present), and likely paths of development (the future).

Interoperability as a Driver of Evolution

The world is rapidly changing. The current socioeconomic environment, including rapid digitization of information and processes, the rise of machine learning (ML), and ubiquitous access to the Internet, amplifies the need for human-human and human-machine interactions that are transparent, dependable, resilient, and operate at a global scale—without a single point of failure. This might ring a bell; the concept of *distributed ledger technologies* (DLT), or blockchain,

» key insights

- **Blockchain interoperability is an important area of research for developing next-generation decentralized applications and services.**
- **Interoperability research has received increased attention in the last five years. Many challenges have been tackled, such as the security of wallets that support tokens and transfers of value across heterogeneous chains. These advancements often benefit from industry collaborations, where consumer insights guide engineering efforts and research.**
- **Several industry infrastructure providers have contributed to these efforts and key developments.**
- **Despite advancements in security, blockchain interoperability still has a wide attack vector, with multiple malicious actors exploiting popular cross-chain bridges (+3.2B USD). More research is needed to secure cross-chain infrastructure. Other prominent challenges are privacy, improvements in UX, and making bridge aggregators production-ready.**

IMAGE BY INKED PIXELS

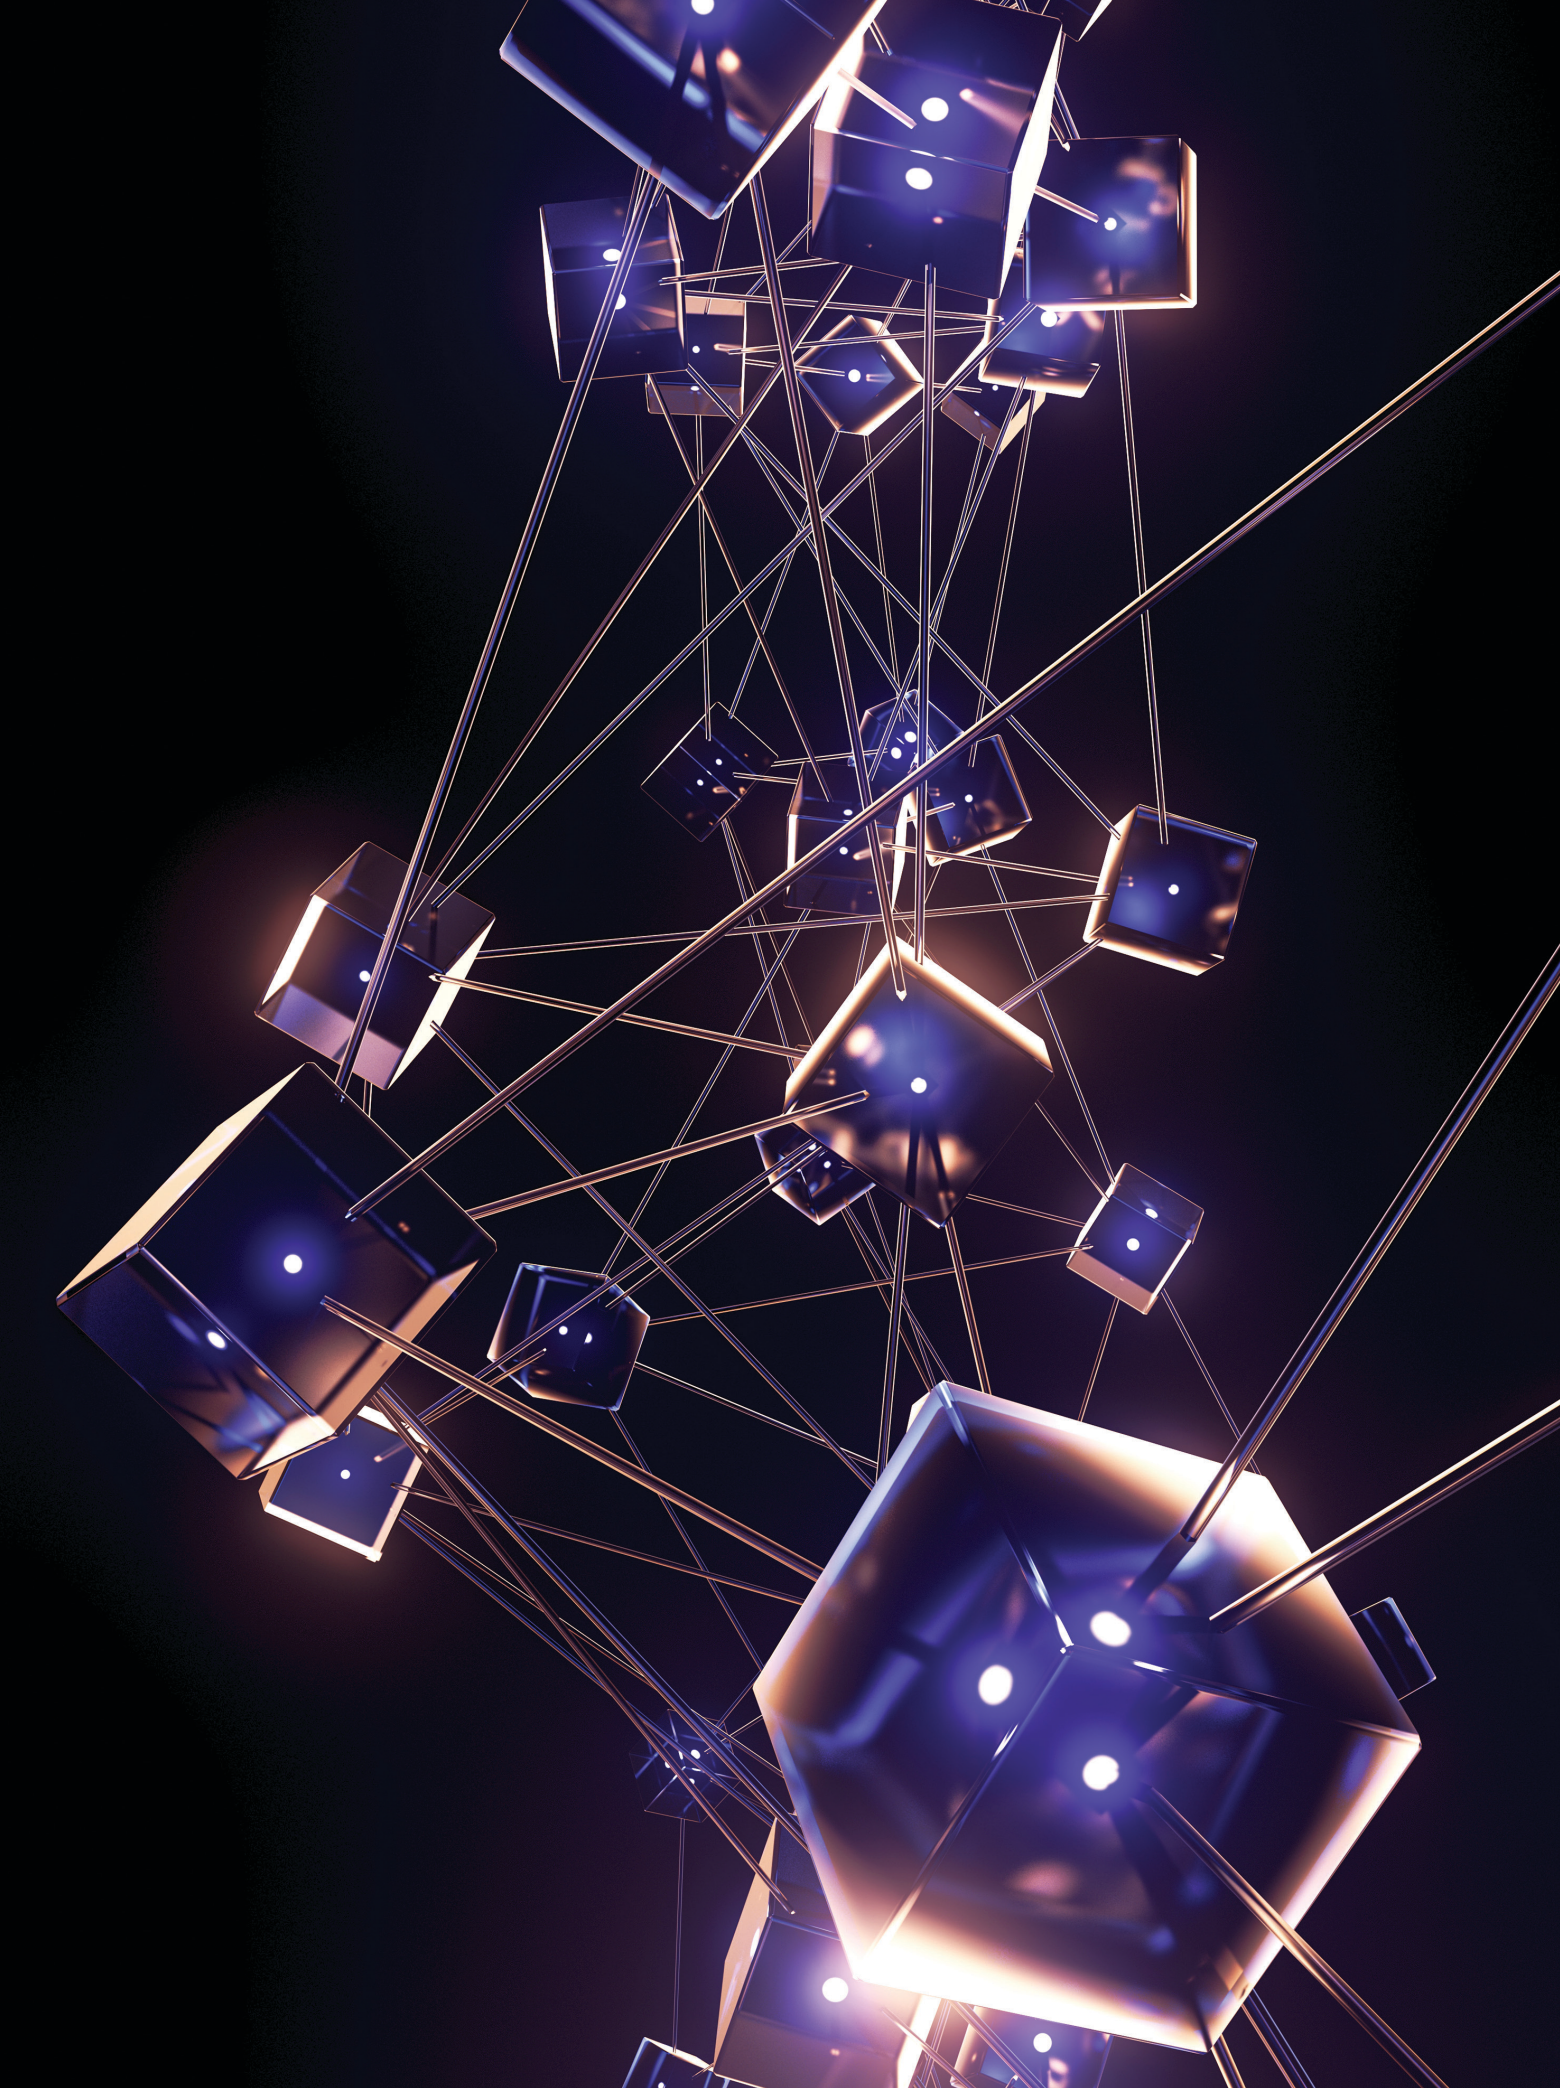

refers to systems implementing these properties. More specifically, DLT refers either to a distributed system of peer nodes that agree on a ledger of records or to a data structure that implements such a ledger. In this design, multiple replicas maintain a global state using a consensus algorithm. The global state is changed via user-submitted transactions, similar to conventional databases. Changing the state is subject to transactions adhering to specific consistency rules.

The innovation that blockchain provides is the ability, for the first time in history, to convey (business) transactions in a decentralized way, allowing the existence of decentralized applications (*dApps*). Many use cases have been either developed as proofs of concept or deployed to production, for instance, in healthcare, supply chain, metaverse, justice, arts/non-fungible tokens (NFTs), decentralized finance (DeFi), and many others. Such systems provide safety and liveness, which in the distributed-system research-area jargon means they do not allow bad behavior from participants (*bad things do not happen*), and desired behavior eventually is processed by the system (*good things happen*).[19] How these properties are realized depends on the desirable decentralization level, the fundamental property of blockchains, and the implementation specifics.

Blockchains have been around since 2008 and come in very different flavors: from the primer blockchain and cryptocurrency *Bitcoin*, a system that revolutionized decentralized peer-to-peer payments without a trusted authority, to *Hyperledger Fabric*, a private blockchain framework that prioritizes privacy and scalability over decentralization,[3] suitable for enterprise-grade use cases. In Bitcoin, safety (that is, security) is realized by the common prefix, chain growth, and chain-quality properties,[22] meaning that, at a high level, honest nodes share a common history of blocks, the chain grows, and the ratio of blocks proposed by malicious nodes is upper-bounded by the ratio of blocks proposed by honest nodes. In Fabric, safety is weaker and realized in terms of accountability. Accountability means that a malicious party can halt the blockchain, but it will be identifiable and, therefore punishable—a sensitive trade-off made in a business network where parties are identified and operate under a certain legal framework. Thus, it is clear that blockchains have evolved in very different directions.

The blockchain trilemma, postulated by one of Ethereum's founders, states that blockchains have an inherent trade-off between security, scalability, and decentralization. Being an equivalent of the CAP theorem[23] for blockchains, the core property chosen is typically security—implemented through consensus algorithms, crypto-economics, formal modeling, and results from distributed systems research (namely crash-fault-tolerant and byzantine-fault-tolerant algorithms[19]). Typically, the more nodes involved in a peer-to-peer network, the harder it is to corrupt it, but the slower the consensus becomes (intuitively, more nodes, more messages exchanged, and therefore, the higher the overall communication latency). Consequently, decentralization and security walk *manus in manu*. Nonetheless, we still have to solve the scalability part of the trilemma. But how? The answer lies within the research area of interoperability, and it will be later apparent to the reader why.

The origins of interoperability—The past. "Interoperability is the ability of two or more software components to cooperate despite differences in language, interface, and execution platform."[38] Counting with a large corpus of research, interoperability has been studied since the 1980s,[30] when engineers started observing the rise of complex software systems that communicated with other systems, heterogeneous in nature. Indeed, interoperability research tends to appear in a later stage of a given technology when modularity, composability, and heterogeneity come into play. As a natural evolution of technological advance, interoperability started gaining more notoriety with the emergence of the Internet.[25] The latter was created in a geo-political context (namely the Cold War) that required the creation of a resilient, dependable, scalable, manageable, and self-healing network that could sustain attacks from a powerful adversary. Effectively, the Internet architecture specified the number of properties that propelled it as a commercial success, enabling considerable economic growth. Those properties are *survivability*, *diversity of services*, and *diversity of networks*.

Not surprisingly, these principles anchored in the Internet architecture are guiding the development of interoperability protocols and standards, with direct application to blockchains.[25] Given the history of the development of the Internet and computer networks in general, it is not surprising that communities are pushing toward cross-chain interoperability. Consequently, the world is settling on several multi-chain blockchains connected by cross-chain solutions (typically bridges, considered major players in DeFi ecosystems) executed by cross-chain transactions. Cross-chain transactions are sets of local transactions that respect a set of business rules or conditions over several domains. Those conditions are called the cross-chain rules.[10] In practice, the rules are restrictions in a sequence of read-and-write operations, orchestrated across different chains. However, unlike traditional databases, a distributed shared ledger lacks a singular or unitary entity that can be relied upon for reading from or writing to it. Instead, the internal consensus protocol assumes the responsibility of ensuring safety and liveness. Typically, cross-chain transactions respect a set of properties equivalent to ACID,[13] but with several fundamental limitations regarding atomicity. While atomicity states that either all the local transactions are executed correctly and committed to the underlying ledger, or none are, they are not guaranteed by default at the cross-chain level. The underlying technical challenge is how to ensure that two or more distributed ledgers mutually agree on a specific ledger state within a defined time limit, unidirectionally or bidirectionally.

One of the first attempts to solve the interoperability problem was to transfer assets between blockchains via atomic swaps,[27] around 2012[31] or 2013.[33] Atomic swaps involve releasing locked assets in one chain upon

a certain time period (that is, using a timelock)—a condition contingent upon the counterparty providing a secret. The first party can use this secret to reclaim tokens on the other blockchain. On the other hand, data transfers and interoperability with non-blockchain infrastructure started with the conceptualization, implementation, and academic study of oracles, around 2011, 2014, and 2020, respectively.[2,24] Although data interoperability was considered first, before asset interoperability, the latter problem was the focus of attention by blockchain communities due to its market interest. Crossing this information with Belchior et al.,[8,14] we can conclude that the area of blockchain interoperability started to get traction around 2016-2017 (when the number of yearly published papers on the topic exceeded 10 documents,[14] and there was enough interest to justify a survey of available solutions[16]).

Interoperability as a requirement of scalability of service. Interoperability was initially studied in the scope of Bitcoin. With the appearance of new blockchains and supporting infrastructure, the scope increased: Interoperability was quickly found to be a sensitive vehicle to offload computation. Practitioners and researchers had to work under the caveat that this new type of interoperability should not sacrifice decentralization and, simultaneously, should achieve a more balanced trade-off set in the referred trilemma. On the one hand, interoperability is a requirement for scalability. On the other, it enables more functionality.

In light of the wide scope of interoperability, we can decompose it into two types: *multi-chain interoperability* and *cross-chain interoperability*. In multi-chain interoperability, instances of a blockchain-of-blockchains framework[14] (for example, Cosmos, Polkadot, Avalanche) communicate with each other through a trust anchor implemented in the protocol. Each instance has a built-in interoperability protocol and data format that other blockchains instantiated by the same framework understand. Consider Polkadot's instantiations called parachains: Each parachain communicates with other parachains via XCMP, a built-in interoperability format.[39] Communications are anchored by the canonical blockchain (the relay chain in Polkadot). In Cosmos, instances are called zones, which communicate via a protocol called Inter-Blockchain Communication (IBC).[29] What anchors the multi-chain communication is a light-client interoperability mechanism that processes cryptographic proofs.[8] Other blockchains that claim to have incredible scalability typically use a sharding system,[37] where each shard is responsible for computing a subset of the overall transactions. However, there is a problem. Polkadot's parachains can communicate with each other, but can they communicate with Cosmos or other blockchain engines? Not natively, because they follow a different protocol and have a different global state (that is, are *heterogeneous*). Those are the boundaries of a blockchain network (otherwise, they would be considered the same system, that is, *homogeneous*). The cross-chain vision connects heterogeneous chains; in the multi-chain vision, a native cross-chain protocol connects homogeneous chains that use the same framework and typically are anchored in a common chain.

Given the history of the development of the Internet and computer networks in general, it is not surprising that communities are pushing toward cross-chain interoperability.

To connect heterogeneous blockchains, we need to use cross-chain communication, a set of techniques allowing us to share data and transfer assets between blockchains by relying on parties external to the involved blockchains. This concept seems prone to security vulnerabilities, and it is indeed—around USD $3B in losses happened only in blockchain bridges, the most popular cross-chain applications,[4,10] (there are more than 110 bridges[a] with a capitalization of almost USD $18B as of December 2023[b]), conquering the rank of having the most devastating attacks in terms of capital lost within DeFi applications. For this reason, at least in part, it has been pointed out by reputable people in the blockchain community that multi-chain is inherently more secure than cross-chain.[17] While the authors tend to agree that multi-chain does seem to lower the

a https://chainspot.io/
b https://bit.ly/4dkzzhZ

Figure 1. Layers of cross-chain communication protocols.[1]

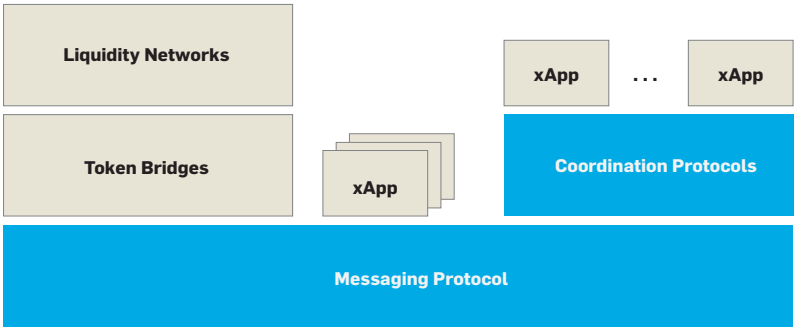

attack vector for interoperable applications, it is also the case that there will not be a blockchain to rule them all: Design decisions need to be made; some give priority to scalability while sacrificing decentralization (namely permissioned blockchains), others focus on privacy,[3] and others are even application-specific[29,39]

Deconstructing Interoperability Mechanisms—The Present

Since 2016, when the interoperability research area started attracting attention, its focus has shifted. Many systematizations of knowledge appeared from 2016 to 2021 (namely 11), highlighting new categories of solutions: sidechains (2015/2016), blockchain-of-blockchains (2016/2017), relays (2019), blockchain-agnostic protocols (2019/2020), solutions for the enterprise (2019/2020), and even preliminary techniques for blockchain migration (2020). Since then, the focus has been on generalization, standardization, and refinement of existing techniques (see Belchior et al[8]). A visible trend is on orchestrating arbitrary logic spanning across centralized and decentralized infrastructure to realize the following interoperability modes acting on the semantic layer. First, the *data transfer* interoperability mode allows arbitrary data transfer to realize general cross-chain business logic.[8] Industry solutions allowing this are called *general message passing* (GMP). Hyperledger Cacti[c] is an example of a cross-chain solution supporting this mode: It connects private to public blockchains and facilitates integration with centralized systems. Such platforms can use multi-chain APIs, such as Blockdaemon's Universal API, as building blocks.[d] The second type, *asset transfer* solutions, are typically implemented through cross-chain bridges. In bridges, an asset is locked in an origin blockchain, and the representation of that asset is created (minted) on a target blockchain (called wrapped or synthetic assets). Bridges have been attacked consistently because the attack surface is very large[4,42] Finally, *asset exchanges* consist of two pairs of transactions, a pair in each blockchain such that: 1) Alice transfers tokens of cryptocurrency A to Bob on blockchain 1; and 2) Bob transfers tokens of cryptocurrency B to Alice on blockchain 2, which are mediated by off-chain processes and smart contracts. Many of these advances were made possible due to the (recent) standardization effort of data formats (for example, view Belchior et al.[12]) and token interfaces (for example, ERC-721, xERC20, ERC-6358), protocols, and blockchain IDs.[e]

A look at the industry. To understand the current interoperability landscape, note that the current market has more than 100 solutions.[f] Out of these, low-level interoperability protocols are more expressive and general than the asset-specific, chain-specific, or application-specific bridges further up the stack, which specialize in one task. We hypothesize that teams are increasing their focus on GMP protocols (for example, the Axelar Team[6]), popularized in 2021/2022, because the expressiveness of the data they can handle allows for developing flexible solutions, by leveraging data transfers as the basis for asset transfers. One can design a GMP protocol that relays messages across blockchains, and expose APIs (on the smart contracts) that can be consumed by coordination protocols (for example, bridges), as Figure 1 illustrates.

Compared to more limited solutions, the development of generalized messaging protocols is more laborious. However, their creators can reduce reliance on individual blockchain networks, applications, and assets. At the same time, they benefit from both the utilization of their own products and those built on their system by partners and customers, for example, through licensing or a pro-rata share of fees. Some examples include Axelar's Satellite, recently extended with cross-chain swaps between the protocol's synthetic and a lot of chains' native assets thanks to the implementation of third-party bridge aggregator Squid Router; liquidity network Stargate and Aptos Bridge, both built on top of LayerZero (see the full version for technical details[11]); and Wormhole's Portal and external Carrier bridge. Even before engaging in a more profound categorization of the systems, it becomes clear that the prevalence of mutually independent solutions is significantly lower than assumed when the underlying messaging protocols are considered.

Asset transfers, the most popular interoperability mode, are typically realized by bridges; there are several types. In recent years, a consensus emerged within the industry regarding the classification of bridges according to the *Interoperability Trilemma—Trustlessness, Extensibility, and Generalizability*. Informally, trustlessness means that the bridge's security is directly pegged to the underlying (source) blockchain. Extensibility means the bridge can support additional blockchains without major refactoring. Generalizability

c https://bit.ly/3AnSSs5

d https://bit.ly/3YDnct4
e https://chainlist.org/
f https://chainspot.io/portal

means the bridge can perform both data and asset transfers. The interoperability trilemma states there is a tradeoff between factors such as latency, cost, and security, implying that different bridge designs exist to accommodate each side of the spectrum. The bridge classification predicts different architectures, systems, and security models. Bridges can be classified into different categories (refer to the full version of this article for a description[11]).

Having already implicitly addressed generalizability (the ability to process arbitrary data) and extensibility (the support of and effort required to expand an interoperability system with new chains), trustlessness undeniably represents the most important dimension, practically speaking, given the number of hacks and amount of damage already suffered by the space.[g,4,10] Trustlessness—a measure for the additional trust required from users of an interoperability system beyond that in the underlying source and destination chains—is closely related to the solution's verification mechanism, potential further trust, and liveness assumptions; and together with these, it constitutes protocol-sided security. However, given the difficulty of reliably assessing highly complex systems with unique architectures, constantly changing maturity, and under permanent threat from a variety of risks and attack vectors, a new approach to trust in interoperability is to look at it as a spectrum.

Current obstacles and challenges. There are ongoing challenges in interoperability, many of which are systematized in Belchior et al.[8,14] and Jin and Xiao[28] and still remain up to date. According to our recent research, the problems we believe to be most prominent as of February 2024 are security monitoring,[4,10] systematic benchmarking of interoperability solutions,[9,36] and privacy.[4] An orthogonal problem in the area is the lack of uniformization of terms and vocabulary: Academia and industry sometimes speak different languages in this research area, in particular on rollups research.

Cross-chain privacy. It is generally agreed upon that anonymity (in terms of unlinkability), confidentiality, and indistinguishability of transactions are beneficial privacy properties in the cross-chain context[4,40] An anonymous asset transfer (or exchange) will hide the identities of the parties involved in the transfer. Confidentiality will hide the number of transferred tokens. Indistinguishability means an external observer cannot say whether or not the transaction is part of a swap. Researchers and practitioners alike have done work in cross-chain, specifically in the areas of asset transfers (namely between privacy-enhanced blockchains as the source and public blockchains as the target,[35] and leveraging promising technologies such as zero-knowledge proofs). Although there is a long road ahead, existing work seems to suggest that preserving the property of "unlinkability" is possible in scenarios where at least one confidential blockchain is involved (by confidential, we mean "permissioned" or privacy-enabled by default like Hyperledger Fabric, ZCash, or Monero, for example, confidential to confidential), therefore achieving some level of anonymity and possibly some confidentiality depending on the blockchain, as ZCash would allow. Privacy on asset exchanges has also been studied.[20] Privacy on asset exchanges appears more straightforward than other interoperability modes: HTLCs share secrets only understandable by the involved parties, making it harder to draw direct associations between transactions. Of course, by analyzing certain heuristics (simpler: amount locked, cryptographic parameters such as the prime field for a private HTLC; more complex: time intervals for swaps, user activity interactions, crossing with off-chain data) one could de-anonymize the actors behind cross-chain transactions. Recent work has revealed interesting insights on cross-chain privacy,[4] namely its deprioritization compared to security, common usage of zero-knowledge proofs, current high-latency and transaction-cost overheads, the need to educate end users, and that full privacy is only attainable if the underlying ledger provides privacy features.

Interoperability solution benchmark. Multiple benchmarking and standardization efforts are in progress. However, there are still considerable challenges since the lack of a uniform API and concrete benchmark datasets hinders a systematic comparison between cross-chain systems (although directions for evaluating interoperability solutions already exist[8]) and a few interoperability solutions are assessed in detail.[18] Methodology and empirical studies to assess components around cross-chain solutions, such as cryptographic primitives, libraries, compilers (especially relevant for SNARK or STARK-based solutions[7]), SDKs, and hardware accelerators, among others, need to be further developed. Studying interoperability solutions in the Web3 world will also give back to traditional interoperability research, as we collect insights on integrating centralized and decentralized systems. A good starting point is directed to evaluate scalability (in terms of the number of blockchains and tokens supported) cross-chain latency, throughput, and transaction costs on popular bridges. There is industry interest in studying this topic.[h]

Security monitoring. Monitoring bridges and the sophisticated and sometimes fragile relationships between ecosystems quickly becomes difficult because the systems to be dealt with are heterogeneous and decentralized, and the systems built on top of them (for example, decentralized applications) may have arbitrarily complex business logic. Imagine a simple case: Your application on blockchain A depends on the consensus of blockchain B. What happens if blockchain B forks, is attacked (for example, 51%), suffers any of the many possible cross-chain attacks, or even collapses?

This last possibility was a reality for the Terra blockchain, with implications for the Cosmos and Ethereum ecosystem, as they were connected by the Osmosis bridge. In the Terra blockchain collapse, exploiters created a destabilization of the stablecoin hosted by Terra. This destabilization caused liquidation

g https://rekt.news/leaderboard/

h https://bit.ly/4cyOap3

cascading, possibly the main cause for a new crypto crash.[21] The collapse of economic security on Luna posed dangers for the Cosmos hub Osmosis, a decentralized exchange bridged to Ethereum. In Osmosis, there was USD$66M of OSMO tokens in the UST/OSMO pool, where UST is the Terra blockchain, that could be stolen over the bridge by an attacker with voting power equal to two-thirds of the staked LUNA. A solution to this problem was for bridge operators to manually shut down bridges, causing impermanent losses. The monitoring of the operations underlying this particular use case could have prevented such a tragic outcome and helped mitigate loss. In a cross-chain setting, automating the discovery of cross-chain models and enabling their monitoring becomes very challenging, as there is a lack of tools to secure and monitor cross-chain applications. Solutions based on modeling by specification[10]could be interesting directions for future work.

The Future of Blockchain Interoperability

What trends will we support in the next few years? To answer this question, there are first some trade-offs to consider, namely the mentioned interoperability trilemma trade-offs: trustlessness, extensibility and generalizability. As the industry seems to have prioritized the last two trade-offs, it is not surprising that the trends reflect an evolution in this sense.

The first trend is the usage of a modular stack design and hence the emergence of cross-chain applications. Instead of having a single interoperability solution to handle all the functions similar to a monolithic Layer 1 network, we observe that blockchain interoperability solutions are increasingly specialized to handle secure arbitrary message passing at a lower level, value transfer, and coordination of remote state-dependent transactions at a higher level.[1] Such a stack framework allows developers to offload the security component to GMPs while focusing on developing applications that coordinate dependent transactions across two or more networks, such as cross-chain *decentralized exchanges* (DEXs), also called DEX aggregators. Sushiwap and Stargate Finance on LayerZero, Squid Router on Axelar, and Osmosis on IBC are examples of cross-chain DEXes enabled by different interoperability solutions. More use cases considered by the IETF are documented here.[34] Those reflect the need of integrating blockchains with centralized systems in the areas of supply chain (transfers of letters of credit, also reported here[12]), currency transfers across central-bank digital currencies (also reported here[5]), delivery vs. payment (DvP) of securities, and transfer of digital art across jurisdictions.

Blockchain is likely to remain an important component for decentralizing our society. However, its full potential needs to be unlocked via synergies with other decentralized and centralized systems, which are not going to be replaced.

The second trend is security-driven model selection. Lower-value transactions typically migrate to Ethereum layer 2 solutions while higher-value ones that demand more security remain on the main chain. Similarly, the selection of particular security models for cross-chain dApps will be largely determined by the use cases and the level of trust and risk users can tolerate. Each model has a clear set of trade-offs in statefulness, security, capital efficiency, speed, and connectivity.[15] For instance, use cases that prioritize speed and cost with lower security requirements can use the external multisig model while those that prioritize security with lower requirements on speed can use the optimistic model[i] or SNARKs.[7] This is related to the emergence of bridge aggregators, software systems that expose several existing bridges in a single interface. Such an interface can provide a better user experience by systematically and explicitly providing details about cross-chain transaction latency, cost, and throughput, and even visualizing the cross-transaction flow.[10] The end user would be able to choose from a range of options depending on their specific needs, availability of liquidity, and connectivity. The trend is analogous to infrastructure providers such as Blockdaemon taking on the complexity of managing the analysis, deployment, and maintenance of hundreds of different blockchain protocols on behalf of their clients.

The third trend is the potential consolidation of GMPs similar to the

i https://bit.ly/46CFCMa

consolidation in layer 1 networks,[j] with most transactions happening on Ethereum, Avalanche, Cosmos, BSC, Solana, and others. There are several contributing factors, such as fragmented liquidity and network effect. On fragmented liquidity, many monolithic solutions utilize different wrap versions of the same asset on the destination chain, resulting in low depth in liquidity pools and hence sub-optimal trading and liquidity provision experience. Such a problem could propel users to migrate to solutions with more adoption across the stack for a better experience and lower capital loss, hence the network effect. From what we have observed, it will be quite likely for different blockchain ecosystems to have canonical interoperability solutions that connect to other ecosystems.

Key Takeaways

Recent developments in blockchain have been incredibly exciting, unveiling a realm of possibilities not possible three years ago. We identified four trends shaping today's interconnected blockchain ecosystems: the adoption of modular stack designs, driven security-model selection, consolidation of GMPs, and usage of bridge aggregators. Indeed, there are few doubts that these technologies will cause fundamental changes in how we interact with each other and how we perceive and exchange knowledge. In spite of its weaknesses, particularly the high computational cost in terms of latency and resources, blockchain is likely to remain an important component for decentralizing our society. However, its full potential needs to be unlocked via synergies with other decentralized and centralized systems, which are not going to be replaced. Among the multiple tasks to be done, the most important ones are enhancing the privacy of cross-chain solutions, creating benchmarks to assess cross-chain systems, and monitoring. We call for a joint endeavor from researchers, engineers, and data and privacy experts as an essential vehicle to unlocking the potential of blockchain for the world at large.

j https://li.fi/; https://bit.ly/4cthu02

Acknowledgments

We warmly thank André Augusto, Jonas Pfannschmidt, Chris Spannos, Freddy Zwanzger, Andie Baker, Dom Martinez, Gabriel Crispino, the colleagues at Hyperledger Cacti, and our colleagues in the IETF's working group Secure Asset Transfer Protocol (SATP) for fruitful discussions. This work was supported by the European Commission under project BIG ERA Chair (grant agreement 952226) and by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UIDB/50021/2020 (INESC-ID) and 2020.06837.BD. Rafael was supported by a Fulbright Scholarship while doing research at MIT Connection Science (MIT Media Lab).

References

1. Abebe, E. et al. Crosschain risk framework. (2023); https://bit.ly/3AjX3VV
2. Al-Breiki, H., Rehman, M.H.U., Salah, K., and Svetinovic, D. Trustworthy blockchain oracles: Review, comparison, and open research challenges. *IEEE Access 8* (2020), 85675–85685.
3. Androulaki, E. et al. Hyperledger fabric: A distributed operating system for permissioned blockchains. In *Proceedings of the 13th EuroSys Conf.* (2018), 1–15.
4. Augusto, A. et al. SoK: Security and privacy of blockchain interoperability. In *Proceedings of the IEEE Symp. on Security and Privacy* (2024); http://tinyurl.com/soksp
5. Augusto, A. et al. CBDC bridging between hyperledger fabric and permissioned EVM-based blockchains. *TechRxiv Preprint* (2023).
6. Axelar. What is general message passing and how can it change web3? (2022); https://bit.ly/3Ai6CEV.
7. Belchior, R. et al. Harmonia: Securing cross-chain applications using zero-knowledge proofs, (2023); https://bit.ly/3yjMOQY
8. Belchior, R. et al. Do you need a distributed ledger technology interoperability solution? *Distributed Ledger Technologies: Research and Practice 2*, 1 (2023), 1–37.
9. Belchior, R. et al. Towards a common standard framework for blockchain interoperability - A position paper. *TechRXIV* (Nov. 2023).
10. Belchior, R. et al. Hephaestus: Modeling, analysis, and performance evaluation of cross-chain transactions. *IEEE Transactions on Reliability* (2023), 1–15.
11. Belchior, R. et al. A brief history of blockchain interoperability. *TechRxiv Preprint*, (2023).
12. Belchior, R. et al. Can we share the same perspective? Blockchain interoperability with views. *TechRxiv Preprint*, 2023.
13. Belchior, R., Vasconcelos, A., Correia, M., and Hardjono, T. Hermes: Fault-tolerant middleware for blockchain interoperability. *Future Generation Computer Systems 129* (2022), 236–251.
14. Belchior, R., Vasconcelos, A., Guerreiro, S., and Correia, M. A survey on blockchain interoperability: Past, present, and future trends. *ACM Computing Surveys (CSUR) 54*, 8 (2021), 1–41.
15. Berenzon, D. Blockchain bridges: Building networks of cryptonetworks. (Sep. 8, 2021); https://bit.ly/3WG2GFi
16. Buterin, V. Chain interoperability. *R3 Research Paper 9* (2016), 1–25.
17. Buterin, V. vitalik.eth on Twitter, arguments for a multi-chain future, Jan. 7, 2022; https://bit.ly/3AdOQm5
18. Chervinski, J.O., Kreutz, D., Xu, X., and Yu, J. Analyzing the performance of the inter-blockchain communication protocol. *arXiv preprint arXiv:2303.10844* (2023).
19. Correia, M. From Byzantine consensus to blockchain consensus. *Essentials of Blockchain Technology* (2019), 41.
20. Deshpande, A. and Herlihy, M. Privacy-preserving cross-chain atomic swaps. In *Proceedings of the Financial Cryptography and Data Security: FC 2020 Intern. Workshops, AsiaUSEC, CoDeFi, VOTING, and WTSC*. Springer, (2020), 540–549.
21. Luna crypto crash wipes out savings of thousands of investors, sparking fears for sector. (May 12, 2022); https://bit.ly/46G46Ep
22. Garay, J., Kiayias, A., and Leonardos, N. The Bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology 9057* (2015), 281–310.
23. Gilbert, S. and Lynch, N. Perspectives on the cap theorem. *Computer 45*, 2 (2012), 30–36.
24. Giulio, C. Before Ethereum. The origin and evolution of blockchain oracles. *IEEE Access* (2023), 1–1.
25. Hardjono, T., Lipton, A., and Pentland, A. Toward an interoperability architecture for blockchain autonomous systems. *IEEE Transactions on Engineering Management 67*, 4 (2019), 1298–1309.
26. Hargreaves, M., Hardjono, T., and Belchior, R. Secure asset transfer protocol (SATP). *Internet Engineering Task Force, Online Draft*(2023); https://bit.ly/4fJTx7k
27. Herlihy, M. Atomic cross-chain swaps. In *Proceedings of the 2018 ACM Symp. on Principles of Distributed Computing* (2018), 245–254.
28. Jin, H. and Xiao, J. Towards trustworthy blockchain systems in the era of "internet of value": Development, challenges, and future trends. *Science China Information Sciences 65* (2022), 1–11.
29. Kwon, J. and Buchman, E. Cosmos whitepaper. A network of distributed ledgers. *Cosmos* (2019), 27.
30. LaVean, G. Interoperability in defense communications. *IEEE Transactions on Communications 28*, 9 (1980), 1445–1455.
31. Lerner, S. P2ptradex: P2p trading between cryptocurrencies. *Bitcoin Forum*, (2012); https://bit.ly/3LWBAVp
32. Narayanam, K., Ramakrishna, V., Vinayagamurthy, D., and Nishad, S. Atomic cross-chain exchanges of shared assets. *arXiv* (Feb. 2022).
33. Nolan, T.T. Re: Alt chains and atomic transfers. *Bitcoin Forum* (2013); https://bit.ly/3YwoO7M
34. Ramakrishna, V. and Hardjono, T. Secure asset transfer (SAT) use cases. *Internet-draft, Internet Engineering Task Force (IETF)*, *7* (2023).
35. Sanchez, A., Stewart, A., and Shirazi, F. Bridging Sapling: Private cross-chain transfers. In *Proceedings of the 2022 IEEE Crosschain Workshop*. IEEE, 1–9.
36. Subramanian, S., Augusto, A., and Belchior, R. Benchmarking bridge aggregators (Jan. 2024). Citation Key: subramanianBenchmarkingBridgeAggregators2024
37. Wang, G., Shi, Z.J., Nixon, M., and Han, S. Sok: Sharding on blockchain. In *Proceedings of the 1st ACM Conf. on Advances in Financial Technologies* (2019), 41–61.
38. Wegner, P. Interoperability. *ACM Computing Surveys (CSUR) 28*, 1 (1996), 285–287.
39. Wood, G. Polkadot: Vision for a heterogeneous multi-chain framework. *White Paper 21*, 2327 (2016), 4662.
40. Yin, R. et al. A survey on privacy preservation techniques for blockchain interoperability. *J. of Systems Architecture* (2023), 102892.
41. Zarick, R., Pellegrino, B., and Banister, C. LayerZero: Trustless omnichain interoperability protocol (2021); https://bit.ly/4dzYPjV

Rafael Belchior is a researcher at INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal and senior engineer at Blockdaemon, Lisbon, Portugal.

Jan Süßenguth is an independent researcher, Frankfurt, Germany.

Qi Feng is an independent researcher, Los Angeles, USA.

Thomas Hardjono is the director of Connection Science & Engineering at Massachusetts Institute of Technology, Cambridge, MA, USA.

André Vasconcelos is a senior researcher at INESC-ID and assistant professor at Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal.

Miguel Correia is a senior researcher at INESC-ID and professor at Instituto Superior Técnico, Universidade de Lisboa, Lisbon, Portugal