Dutch Blockchain Research Agenda

dutch digital delta

connect and create

www.dutchdigitaldelta.nl/blockchain

Summary

Blockchain research is best directed at **identifying and creating the conditions** to steer the development of blockchain technology toward maximizing **its potential for societal good**; **and to the exclusion or remediation of undesirable developments**.

These conditions arise from ethical, technological, economical, legal, and societal perspectives, all in close interrelationship. Blockchain research, therefore, must take a systems' point of view.

Blockchain Characteristics

Specific characteristics of blockchain technology may generate either great opportunities or great causes for concern. We identified some distinguishing blockchain characteristics, that all require a better understanding:

- Decentralization for a distributed, immutable ledger
  (consensus and immutability in both open and permissioned blockchains)
- Automation and standardization of transactions
  (smart contracts, interoperability, legal compliance)
- Digital Scarcity
  (value versus information, by means of cryptographic tokens and incentives)
- Disintermediation
  (trust in institutions and technology, interfacing with off-chain world)

Overarching Research Concerns

Specifically, three overarching concerns need to be addressed, in order to align analytic and design challenges for creating and adopting blockchain technology that realizes a positive societal potential:

- Trustworthiness
  - Trust in social and legal institutions, that could either govern the transition to blockchain technology, cohabitate peacefully with it, or might be replaced by it
  - Personal trust in the veracity, accuracy and security of information on the blockchain, including transparent user interfaces for domain engineers and end-consumers
  - Technological reliability and security of consensus and immutability of the ledger; correctness of smart contracts; scalability and performance of blockchain technology

- Sustainability
  - Energy consumption versus the inherent cost of reaching consensus
  - Scalability, both in number of transactions and in number of participants
  - Resilience against disruption, power concentration, hostile takeover
  - Economic viability of blockchain technology versus alternative technologies, including a techno-economic analysis of its use cases

- Governance
  - Legal compliance, in particular on privacy (including the right to be forgotten) and (selfsovereign) identity management
  - Governance of a blockchain, rule- and decision-making, life-cycle management
  - Management of technology transition and evolution; including interoperability and migration between blockchains; emergence of blockchain infrastructure and services

Contents

1. Preface

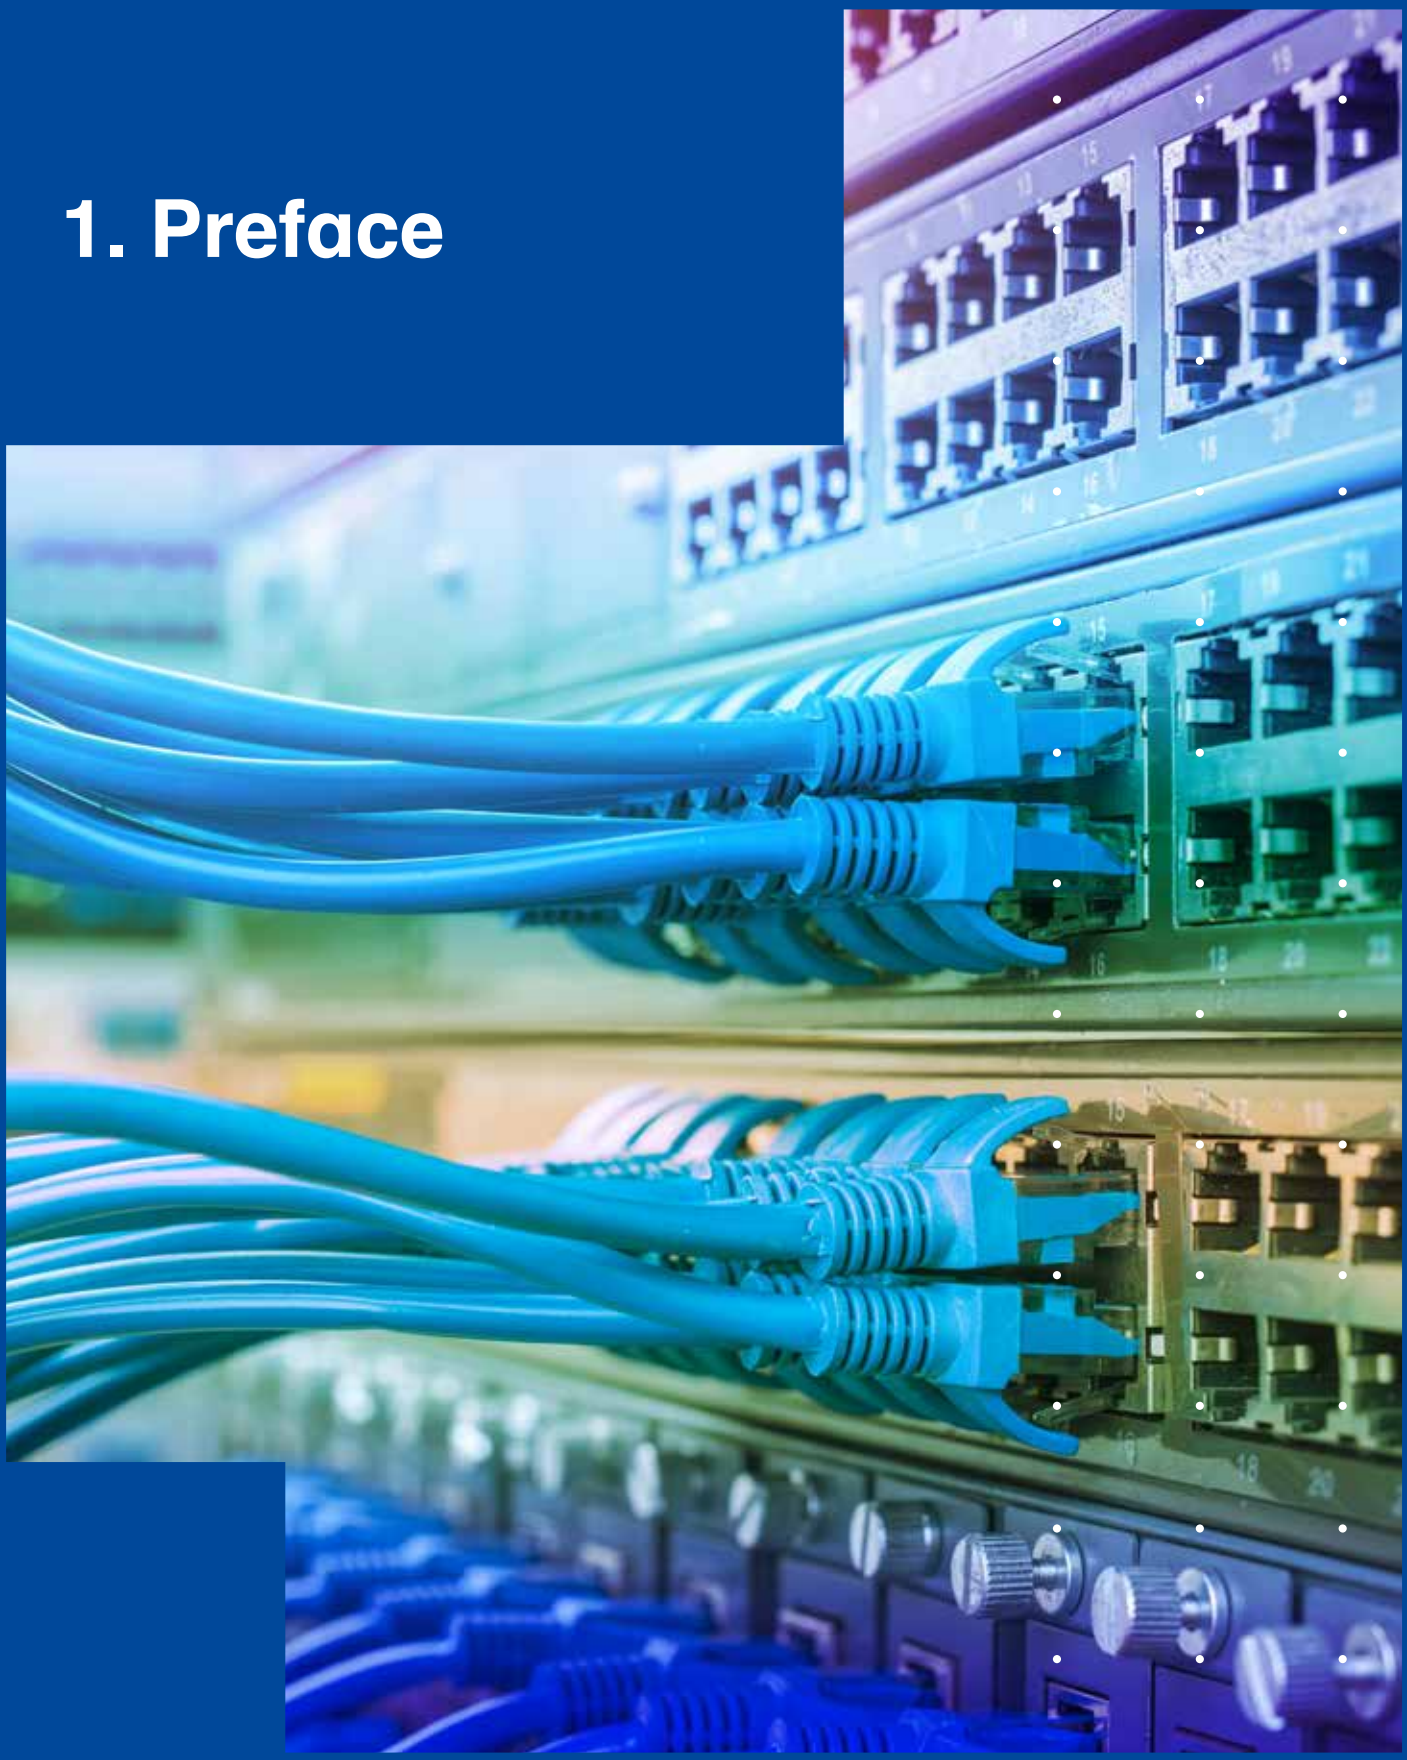

Preface

This research agenda on blockchain has been produced by the Dutch Advisory Committee on Blockchain Research, commissioned by NWO on request of the ICT Top-team. The assignment identified *blockchains* as an upcoming and potentially impactful technology, and asked for a vision on blockchain research, addressing both long-term and mid-term research questions, put in an international context. The assignment also called to organize broad support for this research agenda, with input from multiple disciplines. As a consequence, NWO composed a committee, with representatives from companies, research institutes, government, and scientists from various universities and disciplines (see appendix).

After identifying the most relevant aspects for blockchain research, the committee organised an expert working conference to collect the views from experts from various disciplines and sectors. This resulted in a draft agenda, which was presented and commented in an open consultation, co-organised with the Dutch Blockchain Coalition, leading to this final version.

This blockchain research agenda takes its starting point in the perspective of *responsible science*: any technology should be designed to maximize the good for society and avert negative consequences. Since blockchain technology intermediates in transactions, collaboration and trust between people (individuals or organisations), it must implement our shared public values. Since blockchains form complicated socio-technical systems, a multi-disciplinary systemic view is required, addressing ethical, technical, economic, legal, and social perspectives.

2. Introduction

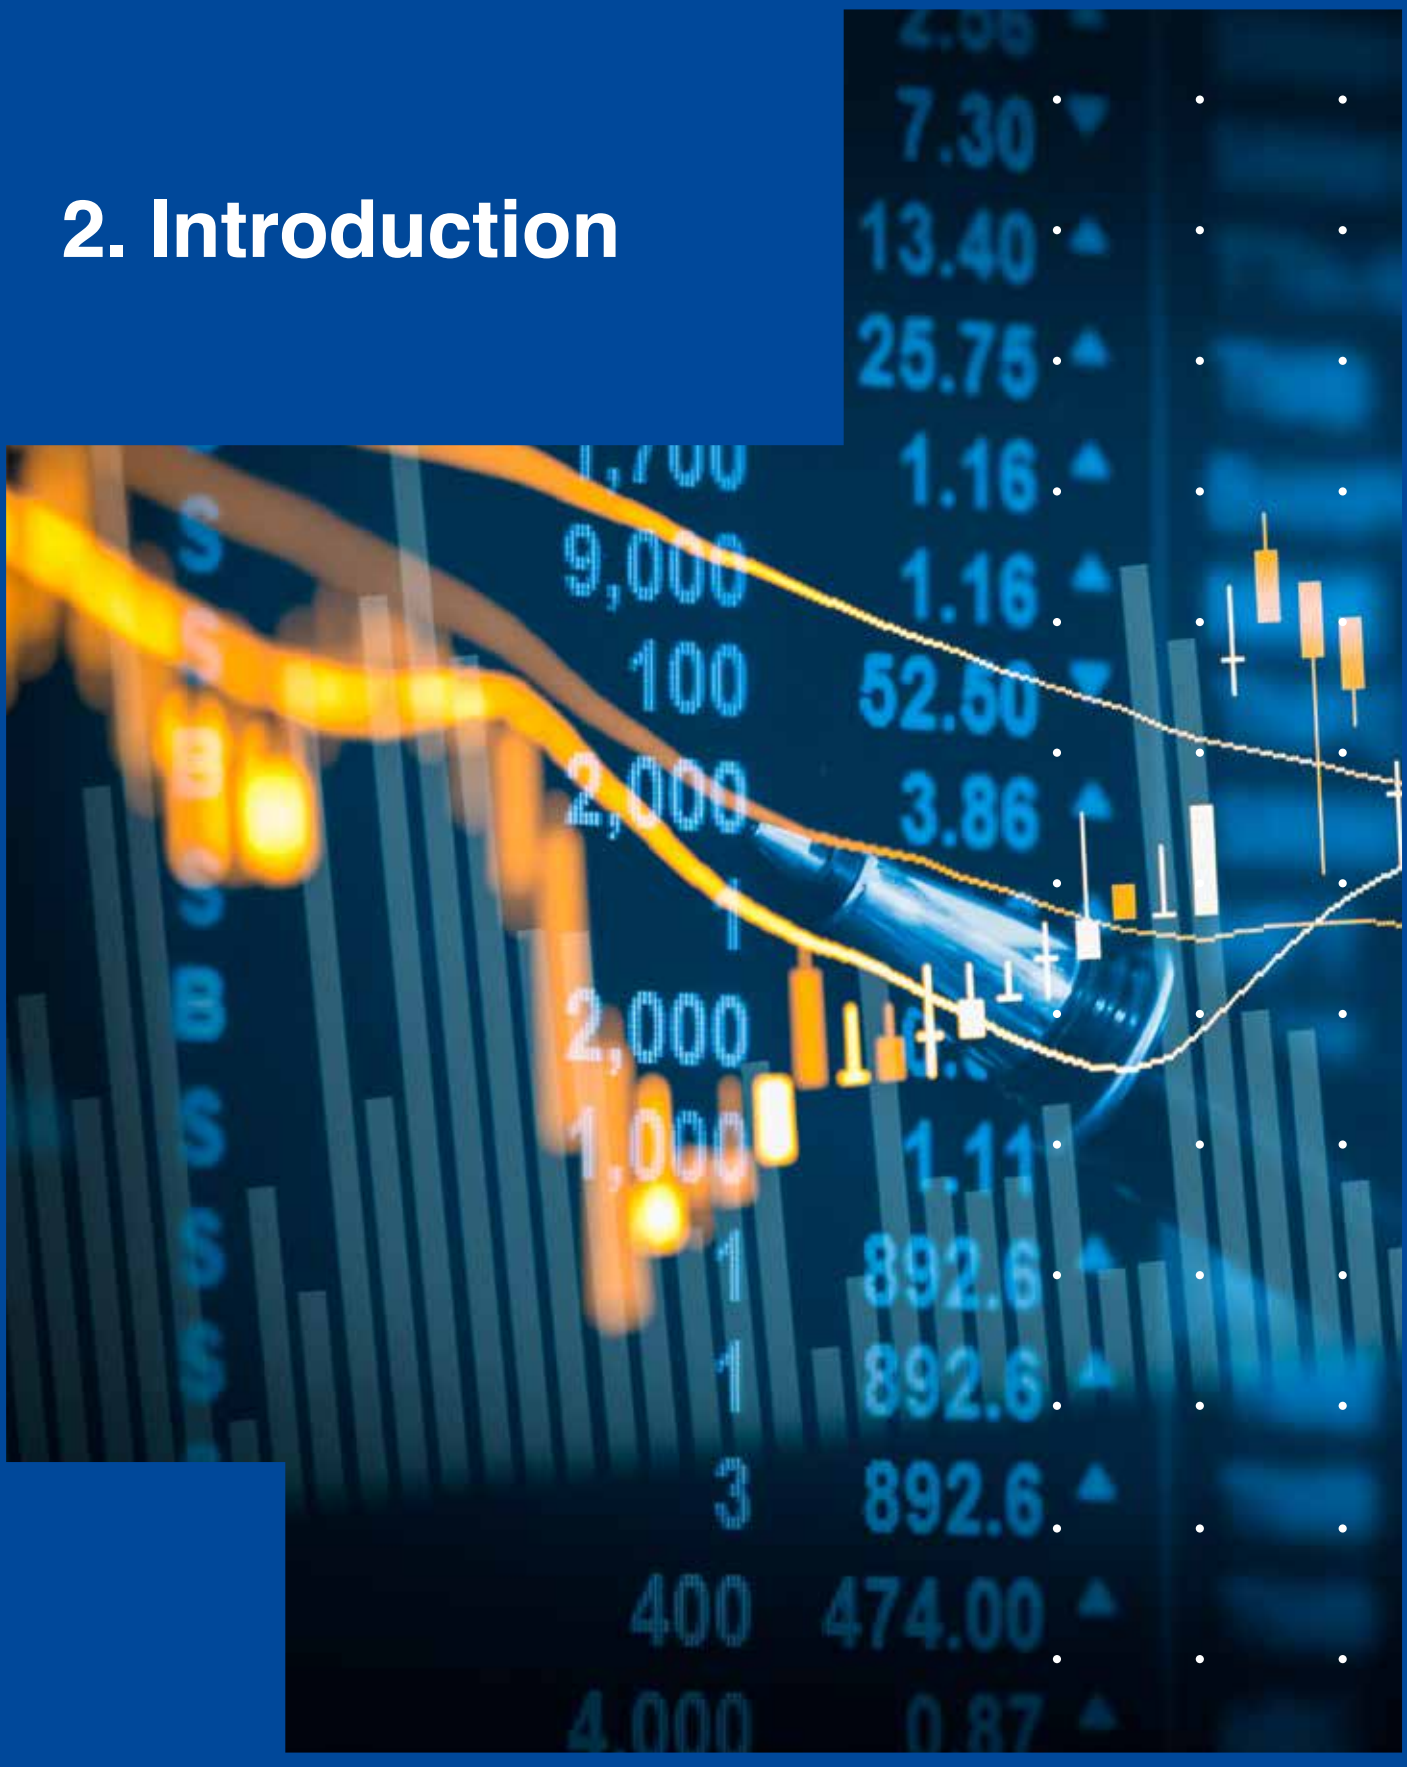

Introduction

Blockchain Technology has received an enormous amount of attention recently and has led to numerous initiatives in research and industry, especially in the financial sector. A number of claims have been made with regards to their large innovation potential. Blockchains, i.e., distributed ledgers based on collaborative, peer to peer, distributed computing architectures, would allow us to redesign traditional institutions and even put society and organizations on a completely new footing. Many believe the promise by technologists and entrepreneurs, that we can start anew and do better this time; yet others are highly skeptical.

Blockchain is welcomed as a technology that allows us to organize our businesses, government and society in a radically different and decentralized way, by implementing radical subsidiarity and local control. Blockchain is also seen as offering solutions to nasty problems that all societies have, like collaboration on an ever-larger scale, by keeping track of economic transactions, and by managing trust and reputation of people and organisations.

We recognize similar claims about the revolutionary potential from a range of digital technologies that emerged in the last decades: Artificial Intelligence (AI) in the sixties of the previous century, the internet in the eighties and nineties, big data and Internet of Things (IoT) at the beginning of the 21st century. By now we must be aware that the impact of new technology is almost impossible to predict. Positive social, economic, political or cultural changes induced by technology rarely come automatically, while unforeseen and often undesirable outcomes happen more often than we tend to expect in advance. Will Blockchain and related distributed computing solutions truly deliver on the expectations of radically innovating society? Our goals should be to design blockchain technology as a desirable socio-technological system, and to develop the capacities to detect, prevent and mend societally undesirable developments. This all boils down to asking the right questions at the earliest possible moment.

Blockchain technology therefore needs to be designed, deployed and harnessed - to obtain a good society and implement shared public values. This opens up exciting research questions at various scales. At the technical level, future-proof algorithms and protocols are required that achieve communication, consensus and immutability in a reliable and secure, fair and resource-efficient manner. Ensuring that these protocols indeed provide the required functionality, needs thorough analysis and experiments.
Design and analysis of blockchain infrastructures, blockchain access services, and blockchain business models are needed as well. Blockchain technology needs to be built, managed, standardised and monetised. We may well see the emergence of blockchain infrastructure and service providers, that specialise in jointly operating a multitude of blockchains efficiently, providing blockchain access both collaboratively and competitively, similar to the internet services. However, it is important to take an even much wider systemic point of view. The functioning of blockchain infrastructures and their applications rely on the decisions of people and organisations, each with their specific core businesses, missions, psychology, values and cultures. This requires the design and justification of the right mix of economic and legal incentive systems. Finally, one must take into account a still bigger picture, of how blockchain will interact with society through organisational contexts, institutional environments, and local and international legal structures. We should understand the effect on institutions that adopt blockchain, or that are disrupted by it.

3. Characteristics of Blockchain Technology

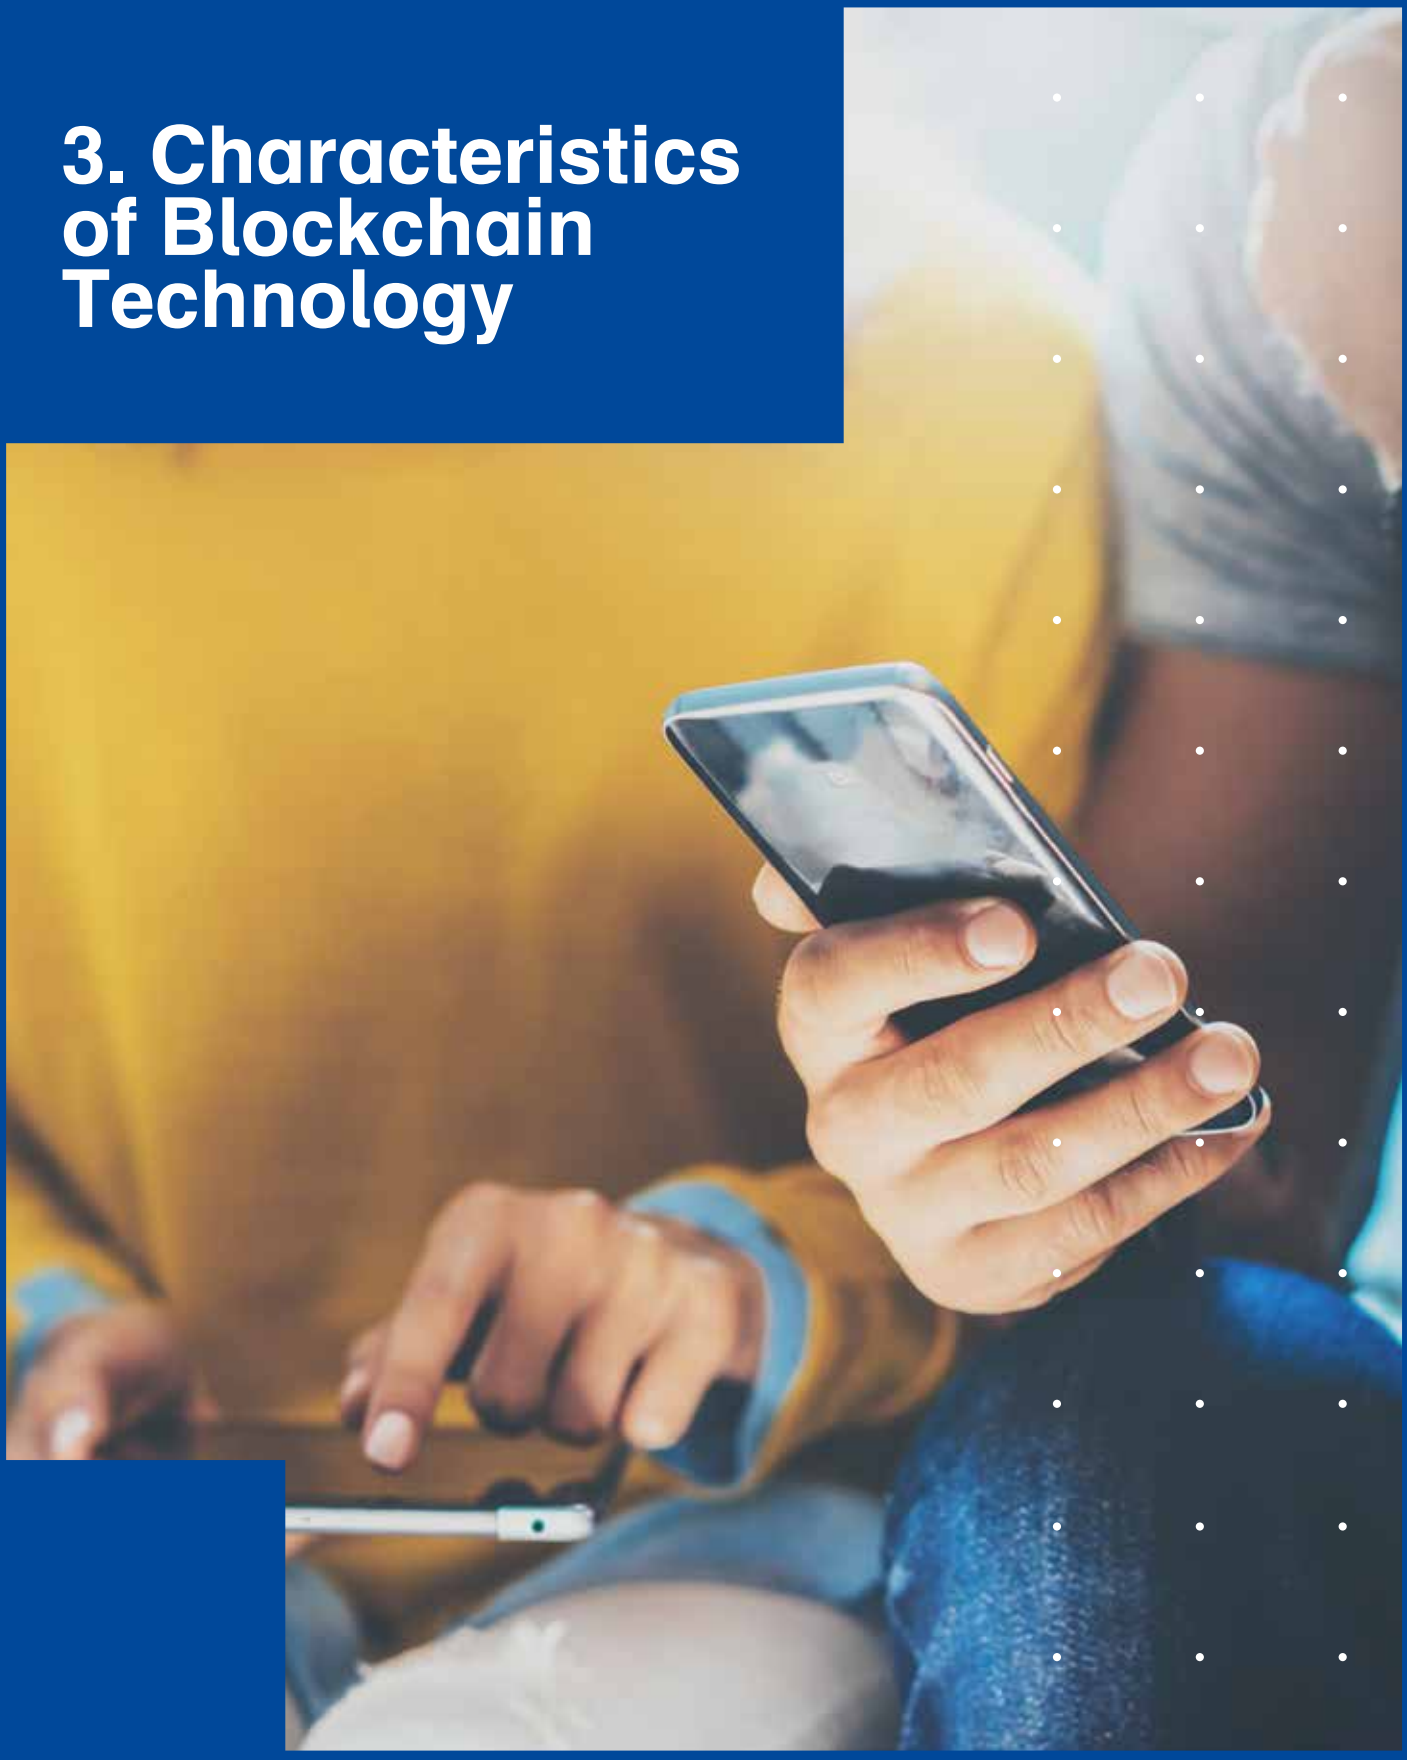

Characteristics of Blockchain Technology

Distributed ledger technologies (DLTs) come in many flavours. We do not replicate their technical definitions here but point out some characteristics that are shared across most, if not all designs, and taken together they define the unique elements of blockchain technology. These shared characteristics include (1) the distribution of infrastructure and redundancy of data storage; (2) the automation and standardization of transactions; (3) the re-introduction of digital scarcity; and (4) the disintermediation of transactions, markets and exchanges. We briefly touch upon each of these below, and subsequently discuss the differences between various blockchain technologies.

(1) Distributed Ledgers are defined by **distributed protocols**, providing consensus on and immutability of the transactions in the ledger. The security of consensus and immutability are technically based upon redundant data storage and cryptographic properties. Actually, there are many competing blockchain proposals, from which this agenda tries to abstract, but one important distinction is between *open blockchains* (i.e., public, permissionless blockchains, in which anyone may enter, even under a pseudonym, like Bitcoin and Ethereum) and *permissioned blockchains* (which regulate the access of participants to provide a limited trust level, like Hyperledger and Quorum). Open blockchains typically use variants of proof-of-work, performed by miners, while permissioned blockchains use different consensus algorithms, like Byzantine Agreement. In both cases, the validity of transactions must be checked by "validators" before they enter the ledger.
In open blockchains, security depends on the relatively equal distribution of power and chances of participation. Yet it has become clear that external factors, such as skewed incentives and energy costs can lead to undesired centralisation of control over the consensus. But permissioned blockchains also require a careful design of business models and governance, especially when these are baked into their distributed protocols. The emergence of oligopolistic control of key network resources, such as mining pools on the bitcoin network or the emergence of a dominant infrastructure vendor in permissioned blockchains, could ultimately affect the fundamental security and trustworthiness of the infrastructure. Hence, the issue of reliable, secure, fair and efficient distribution should be addressed as a technical, economic, political and regulatory challenge.

(2) Smart contracts are programmable transactions, that run on and are part of the immutable blockchain. They may **automate** increasingly complex transactions between individuals, businesses governments, and machines. There is an intense debate about the conditions of, and the extent to which, it is possible to re-implement our legal systems in computer code. To provide maximal expressivity of the complex rules and tasks involved, Turing-complete smart contract languages have been proposed, which raises fundamental limitations to their algorithmic analysability. Meanwhile, as is the case with any other software ecosystem, the development of smart contract languages, solutions and applications already triggered a considerable level of **standardization** at the level of contract fundamentals, leading to the emergence of global standards of smart contract-based interactions. In the past, the emergence of such standards resulted in substantial decreases in transaction costs and the opening of new markets, but the impact of this process in the smart contract domain is largely unknown.

(3) DLTs re-introduce **scarcity** into digital information markets, forcing us to rethink the abundance-based logic of organizing digital information. Cryptographic tokens are (practically) unique, and they exist only in a single copy at a time. In the last decades, information markets used to be organized around digital abundance: the unlimited ability to copy digital information. While digital scarcity allows us to represent value and physically scarce resources digitally, it forces

us to rethink the rules by which digital markets operate.

(4) The key promise of blockchain is **disintermediation**. Since (dis)trust is regulated through technology, transactions between unknown participants can happen at a global scale directly, and the "middle-man" can be cut out. However, the suggestion that trust in institutions can be replaced with trust in a protocol turns out to be a more complex issue than anticipated, opening up a new debate about the conditions, costs and benefits, feasibility and desirability of disintermediation. Moreover, also the function that is being disintermediated needs to be designed, built, incentivised and governed. The technical ability to store a time-stamped, immutable ledger across a large number of nodes forces us to reconsider the type of information that actually needs and merits such investment into preservation and replication.

Differences

Beyond the shared characteristics, the *differences* between blockchain technologies raise their own important questions. The various DLTs are fundamentally different at the level of infrastructure and at the application layer. Some are designed to operate at a planetary scale, while others are based on a core technology with a limited local scope. Some are open for anonymous participation, while others restrict access permissions to known participants. Some are based on open source code run by a community, while others are based on proprietary code operated by a single organisation. Also, their incentive schemes differ widely.

So, there is no such thing as 'the' blockchain technology, and we must deconstruct and nuance the monolithic, undifferentiated understanding of what different blockchain technologies are in theory and in practice. In particular, we must differentiate their consensus models, and specify and analyse the exact meaning of the claims on their reliability and security. We should also be aware of the large gap between the ideal distributed ledger (that promises decentralization, immutability and consensus) and their actual implementations (that fall short).

Finally, special attention must be paid to the non-technical components of the complex socio-technical assemblage that is blockchain technology. The focus on the technical opportunities offered by purely on-chain applications has hidden the importance of the interface between on-chain and off-chain systems, and the fact that this interface must offer two-way communication. Having so called *oracles* to reflect a state in the off-chain world for on-chain transactions is not sufficient. On-chain transactions with off-chain consequences cannot do without entities that are able to enforce such consequences. The task of maintaining synchronicity between the off-chain and the on-chain state of the world is a challenge that needs proper attention, as it may ultimately make or break the real-world applicability of blockchain technologies.

4. Overarching concerns and research challenges

Overarching concerns and research challenges

There are a number of fundamental concerns which apply to almost any context in which blockchain technology may be used. We have bundled these concerns into three clusters: trustworthiness, sustainability, and governance of blockchain technology. Those concerns give rise to new research challenges, both in design and analysis.

4.1. Trustworthiness of Blockchain Technology

Blockchain is usually referred to as a *trustless* infrastructure. However, the actual trust fabric appears to be very complicated. We discuss the trust-related concerns in three circles: the social, personal and technical. We also mention the related threats and vulnerabilities.

Institutional trust (social circle).

It is a challenge to track down what and who we are actually supposed to trust besides the technical infrastructure. Do we need to trust individual miners or validators, who enforce the rules over the network? How are these rules established anyway? What if the overall validation power tends to concentrate in a few hands, like miners or blockchain service providers? Finally, trustable interfaces to the real-world are required, for instance conflict resolution regimes that can provide justice, and institutions that can maintain the synchronicity of the on-chain and off-chain state of the world.

We need a better understanding how the often millennia-old social systems of trust accommodate this technological newcomer. At the same time, it is interesting to investigate how blockchain technology could overcome limitations of traditional social institutions, and what the fate of the currently trusted institutions will be.

User confidence (personal circle).

To what extent can an individual user be certain that a blockchain system or component does what it promises, and that the information provided by it is accurate and veridical? This question pertains to the epistemic quality of the infrastructure and its environment. This information must be communicated through transparent user interfaces. In the first place, this concerns end-consumers, who need to understand the status of all configurations of and interactions with blockchain applications. This requires a smooth user experience when it comes to identity management. A particular challenge is to make the technology inclusive for all kinds of users. This also applies to blockchain and domain engineers, who should oversee the consequences of a particular blockchain technology selection, (open) software design choices, blockchain service infrastructure choices, and the expression of smart contracts, which requires expressive but simple domain-specific programming languages and programming interfaces (API).

Infrastructure reliability (technical circle).

All levels of the blockchain technology stack come with claims about correctness and well-functioning, e.g. on eventually reaching consensus, immutability of the ledger, scalability and performance of infrastructure, cost of operation, confidentiality of encrypted information, correct functionality of smart contracts, integrity of software and applications that access and process information on a blockchain, etc. Alternative blockchain technologies realize slightly different claims by completely different means, for instance Proof of Work versus Byzantine agreement. An important question is how such technical claims can be specified precisely, how they can be substantiated and certified, and how the realised infrastructures can be audited on this. In particular, practical verification methods for distributed protocols with an unknown and dynamic number of participants are not known. Verification of smart contracts expressed by Turing-complete languages is formally undecidable. Studies on performance and scalability require extensive models for simulation and analysis, experiments and piloting in testbeds and field-labs. Besides analytic methods, synthetic methods are required

to design alternative solutions, with demonstrably better properties of correctness, security, performance and resource-efficiency.

Security aspects.

To assess the security of blockchain-based applications, the specific vulnerabilities and threats must be known. In open blockchain networks, other participants can have different intentions, be it unforeseen business models, or even criminal intents. Defining "bad" behaviour in this setting is highly non-trivial. Besides attacker models, blockchain security requires the design and analysis of countermeasures to detect or prevent misbehaviour, like encryption, security protocols and system monitoring. Security should not be restricted to the core DLT, but focus on the edge with the surrounding environment, e.g. integrity of wallet software, theft of wallets, off-chain rule enforcement, and interaction with IoT devices.

4.2. Sustainability of the Blockchain System

Blockchains cannot become the backbone in the global economy of social interactions, without developing itself as a sustainable system. We give a short overview of the obstacles that must be overcome, in order to achieve system that is energy-efficient, scalable, and resilient. Last but not least, a condition for the widespread adoption of blockchain technology is to prove their economic viability and the added value over alternative technologies. This requires insight in the demands and incentives for all stakeholders.

Energy.

It is well known that the *Proof-of-Work* principle underlying the immutability of the Bitcoin ledger currently wastes the same amount of energy as a medium country. Although several competing Proof-of-X schemes have been proposed, the long-term behaviour emerging from such schemes is currently unclear. A number of fundamental questions arise: What levels of global distributed consensus are actually required for certain use cases? What are the inherent costs for reaching those levels of consensus in a reliable and secure manner, under various trust assumptions? And can alternative technologies be designed, like P2P computing and Byzantine Agreement, to reach the required levels of consensus in a more secure and efficient manner?

Scalability.

Another limiting factor to embrace blockchain technology at a large scale is the current lack of scalability, both in the number of transactions (for open blockchains) and the number of participants (for permissioned blockchains). These limitations are inherent in the current design of blockchain solutions. Usually, fundamental improvements in performance require a breakthrough in algorithms and system design. Also, a differentiation of the precise requirements for particular applications would increase scalability and efficiency for particular application domains, but this requires a good understanding of the trade-offs involved. The rapidly developing field of Internet governance warns us that the scalability of the governance of any planetary scale technology is a complex challenge, and it remains to be seen whether blockchain technologies would follow a similar governance development path as the internet did in the last few decades.

Resilience.

If blockchain systems would become a prominent ubiquitous technology in society, new safety hazards arise. If we were to build our organizations, institutions in finance, healthcare and transport around them, what would happen if the systems would fail or no longer be available? Given the complex technological and social construct, resilience considerations go beyond the mere observation that a distributed ledger has no single point of failure. For instance, potential re-centralisation of a blockchain system, through mining power or a hostile takeover, should be taken into account as well. Resilience has multiple non-technical aspects as well, such as data portability, or the interoperability of different blockchain systems, which, may or may not happen organically, and should, but maybe cannot be mandated.

Economic viability.

It is important to understand for which use cases blockchain systems are economically viable, either because their costs are justified, or because they provide cheaper solutions than other

technologies. One should understand for which demands P2P distributed systems, classical web services, or centralised solutions might provide better alternatives.
Besides a fundamental complexity analysis, this requires a technoeconomic analysis of how blockchain based solutions perform against more traditional or alternative solutions, in terms of cost, scalability, and acceptability. To understand the public value chains, a sector-specific and application-oriented research approach is required. The business case can be quite different for different sectors in industry, government and services. Also, the business case and scale of economy for service providers of generic blockchain infrastructures requires attention. They could specialise in jointly operating a multitude of blockchains efficiently.

Incentives.

Another angle is to study the *incentives* that reward "good" behaviour. An important aspect of blockchain systems are the built-in incentives for participants to maintain the system. The long-term effects of artificial incentive systems are still poorly understood. Blockchain infrastructures may turn out to become two-sided markets with complex interactions between the retail and wholesale sides. The game-theoretic challenge of blockchain economics is to design incentive systems that ensure that groups of selfish individuals or organisations work together towards collaborative goals. The interplay between incentive-based and reputation-based systems provides another topic for further investigation.

4.3. Governance and regulation

The last cluster of concerns is related to blockchain governance. We start pointing out open problems in regulation and legal compliance to sometimes competing regimes, in particular for issues in identity and privacy. Next, we discuss governance and organisational structures around blockchain technology, blockchain infrastructures and blockchain applications. Finally, we discuss issues related to technology evolution and migration.

Legal compliance, identity and privacy.

Within the European Union a number of legal frameworks may apply to blockchain infrastructures and applications. These frameworks include the General Data Protection Regulation, the Payment Services Directive, the E-commerce directive, and rules on law enforcement, especially on money laundering and cybercrime. Note that these rules are often contradictory to other legal frameworks and jurisdictions under which a planetary scale blockchain technology would resort.
On the one hand the original, open blockchain technologies were developed especially to be immune to, and even bypass such regulatory regimes. On the other hand, a certain level of legal recognition of blockchain technologies is an important condition for their widespread adoption by states, businesses, and institutions. Some more recent, permissioned blockchain technologies are being developed with regulatory compliance in mind.
In particular, questions arise on the relationship between blockchain transactions and smart contracts, versus legal concepts of property and contracts. It is yet unclear whether limitations of legal compliance can be amicably resolved, or will lead to long term antagonism between this technology and its legal environment.
Questions on identity, privacy and confidentiality need immediate attention as well. The launch of the European General Data Protection Regulation seems to be in stark conflict with the fundamental design of most open blockchain technologies, which store data on distributed, publicly available, immutable ledgers (albeit possibly in encrypted or hashed form). Certain proposed applications of blockchain in health-care, e-government and banking have led to the hope that this technology can return control over personal data back to citizens, while some blockchain infrastructures (like Sovrin) address privacy and self-control explicitly. This would enable users to be responsible ("self-sovereign") for their own credential and identity management.
However, this raises serious questions on the nature of anonymity, privacy, and the protection of personal data. All these well-established legal concepts and their corresponding systems need to be better understood and adopted. At the same

time, new privacy-enhancing technology should be designed to implement the requirements in a compliant manner. Also, methods to demonstrate and certify compliance to privacy regulations are called for.

Governance of blockchain technology.

The (good) governance of technology requires agreed and clear rules and structures defining the boundaries of the technology, assigning responsibilities of management, decision making, execution, describing rules of rulemaking, structures of monitoring, accountability, liability, sanctioning of bad behaviour, regimes of conflict resolution, etc. At this moment quite a few blockchain based projects can demonstrate sophisticated governance structures, for example the various Hyperledger flavours. A lack of governance is especially pronounced in open blockchain technologies, such as Bitcoin or Ethereum, which develop their open-source code base in loosely defined anarchistic, or charismatic meritocratic communities.

Since it is exactly these internal governance structures that need to interface with existing legal and institutional regimes, it is imperative to develop a better understanding of the evolution of these governance structures. Good governance is required at all levels of blockchain technology, infrastructures and applications. It is important to evaluate which blockchain technologies would be applicable to Dutch business sectors and government, and to design the governance of the infrastructures based on these technologies, in order to provide the standards of governance that qualify blockchain application for certain uses by public institutions.

Also, in the case of permissioned blockchains, explicit governance is required to settle questions like which participants are admitted, who is responsible to keep the infrastructure running, what happens when main stakeholders would leave the project, how technical change management should be orchestrated, or how a blockchain infrastructure does reach its end-of-life. It is an interesting challenge to investigate how much of this management can be performed automatically on-chain, and where human decisions remain necessary.

Management of technology transition and technology evolution.

Blockchain technology promises to significantly reduce bureaucracy and in the view of some to upset the institutional status quo. The latter is usually discussed in terms of disruption, yet radical disruption is usually not the most frequent scenario to happen. Any transition certainly benefits from being properly managed as opposed to it happening in an ad-hoc, uncontrolled manner. Transition management for the introduction of blockchain, including the mitigation of the negative effects of disruption, is currently utterly lacking. Also, a view is lacking on how a post-blockchain-transition world might look like.

We need to develop capacities to obtain reassuring answers to questions such as: what are the different types of (legacy) institutions, and institutional functions that can be replaced, phased out, and which ones need special attention to be preserved shall not be replaced? What functions and institutions can be gradually improved? Where should we prepare for the long-term co-existence of blockchain based novel, and legacy institutions and processes?

At the level of the technology itself, it must be avoided that current blockchain software will become the legacy of the future. This calls for sound software evolution principles, including well understood mechanisms to improve standards, e.g. for smart contract languages. Also, this calls for whole life-cycle management for blockchain infrastructure and applications. Practical questions, like when and how to start introducing blockchain technology, and how to bridge it with the surrounding ICT and organisational context, are largely open.

Technical mechanisms to achieve interconnection and interoperability between dissimilar blockchain infrastructures under joint blockchain umbrellas need to be designed. Methods to migrate blockchain applications to an alternative blockchain infrastructure will be essential to provide sustainable solutions on top of this fast-moving technology. However, such trust-preserving interconnection, interoperability, and migration mechanisms are currently not existing and hard to imagine.

5. Notes on the research methods and programmes:

Notes on the research methods and programmes:

Given the complex fabric of technological and societal questions around blockchain, future research seems to require at least the awareness of this multi-disciplinarity, or even seek collaboration across the boundaries of disciplines. Blockchain research carries many challenges on the level of research design and methodology. As is the case with systems focused research, the proper demarcation of scope of future research projects and programmes is essential. This scope also sets the disciplinary mix that needs to be involved. At the same time, it should be ensured that the required disciplinary progress can happen, especially since different disciplines require research at different time scales.

Since blockchain technology is a moving target, in terms of research methodology one must also consider more exploratory, theory generating, high risk and open-ended approaches, including tools such as mathematical modelling and analysis, business modelling, techno-economic analysis, functional and non-functional design and testing, action research, simulations and experiments in research labs and living labs, horizon scanning, etc. As this research agenda includes both fundamental and applied research, it requires active involvement from non-academic stakeholders from public bodies, industry, market sectors and the general public.

Another methodological challenge is the future-proofing of research. In such a volatile field, it is often difficult to distinguish issues relevant only in the short term, versus long term blockchain specific problems, versus fundamental research questions that cut across multiple digital technologies and have been and will be with us for decades.

There are several streams of investment that fuel research in the blockchain technology domain. Private investment through venture capital and ICOs (crowdsourcing) as well as public investment by governments, universities, and research funding bodies should be aligned in a smart way. In that context it seems inevitable to identify the fields that Dutch academia, research institutes and research departments of Dutch organisations are best positioned to answer, either because they already excel in certain domains, or because they want to build skills and research capacity through strategic investment.

6. International Context and State-of-the-Art

International Context and State-of-the-Art

Blockchain was invented as the underpinning of the cryptocurrency Bitcoin (Nakamoto, 2008). It cleverly combined several existing techniques (Narayanan & Clark, 2017) to maintain a decentralized yet global immutable ledger of transactions among mutually distrustful peers (Tschorsch & Scheuermann, 2016; Narayanan et.al., 2016). Soon alternative blockchain systems emerged. Ethereum (Buterin, 2013; Wood, 2017) extended Bitcoins limited scripting capabilities into Turing complete smart contracts. Alternative consensus mechanisms, for permissioned ledgers, based on Byzantine fault tolerant consensus protocols (Lamport et.al., 1982; Castro & Liskov, 2002) were proposed. In recent years we have witnessed a surge in academic research into the foundations of distribute ledger technology (DLT). Among the areas of study are the incentive structure of Bitcoin (Sompolinsky & Zohar, 2018), the different consensus mechanisms and their properties (Bano et.al., 2017; Cachin & Vukolic, 2017), critiques on the actual amount of decentralisation and security offered (Eyal & Sirer, 2014; Efe Gencer et.al., 2018) and a start of a formal analysis of distributed ledgers (Garay et.al., 2017).

An interesting attempt to identify gaps in research on blockchain (Yli-Huumo 2016) lists as prominent system-level challenges for blockchain technology: its scalability (latency, throughput, versioning, etc.), its usability (human interfaces, APIs), and the investigation of the wide diversity of alternative DLT technologies. In the humanities, recent surveys on the philosophy of blockchain have appeared as well (Swan & de Filippi, 2017). For a survey of research on legal and socio-technical aspects we refer to (Atzori, 2015; Campbell-Verduyn, 2017; Finck, 2018; De Fillippi & Wright, 2018).

Similarly, we see incredible business activity around DLT. Currently, financial institutes including banks are developing financial products on top of blockchain technology, among others to simplify international trade. Other applications in logistics focus on provenance: tracking the origin of components or food on an immutable public ledger. There are proposals to manage immaterial goods on the blockchain, like energy, emissions, externalities, digital contents, etc. Finally, we mention envisaged applications centered around identity and credential management. At the same time, people start studying the question whether all these applications truly benefit from applying DLT (Peck, 2017; Wüst and Gervais, 2017). Several international blockchain research centers have already emerged. In the Dutch Blockchain Coalition (DBC), governments, companies and research institutes combine forces in blockchain technology, often organised in field labs. As prototypical example of concentrated blockchain research initiatives in EU, we mention the UCL Center for Blockchain Technology (CBT). UCL-CBT's blockchain research connects over hundred researchers from eight departments in three research pillars: science & technology, economics & finance, and regulation & law. Another initiative, Infrachain (Luxemburg, EU) aims to establish European blockchain infrastructures. We also mention the Swiss ecosystem, with the Crypto Valley in Zug and IBM Research, with a clear financial focus. Many competing Distributed Ledger Technologies, like Ethereum and HyperLedger, are rooted in this area, probably encouraged by the liberal legal context that facilitates experiments with financial technologies. Fraunhofer (Germany) recently launched a technical roadmap on Blockchain and Smart Contracts (Prinz and Schulte, 2017).

7. Literature

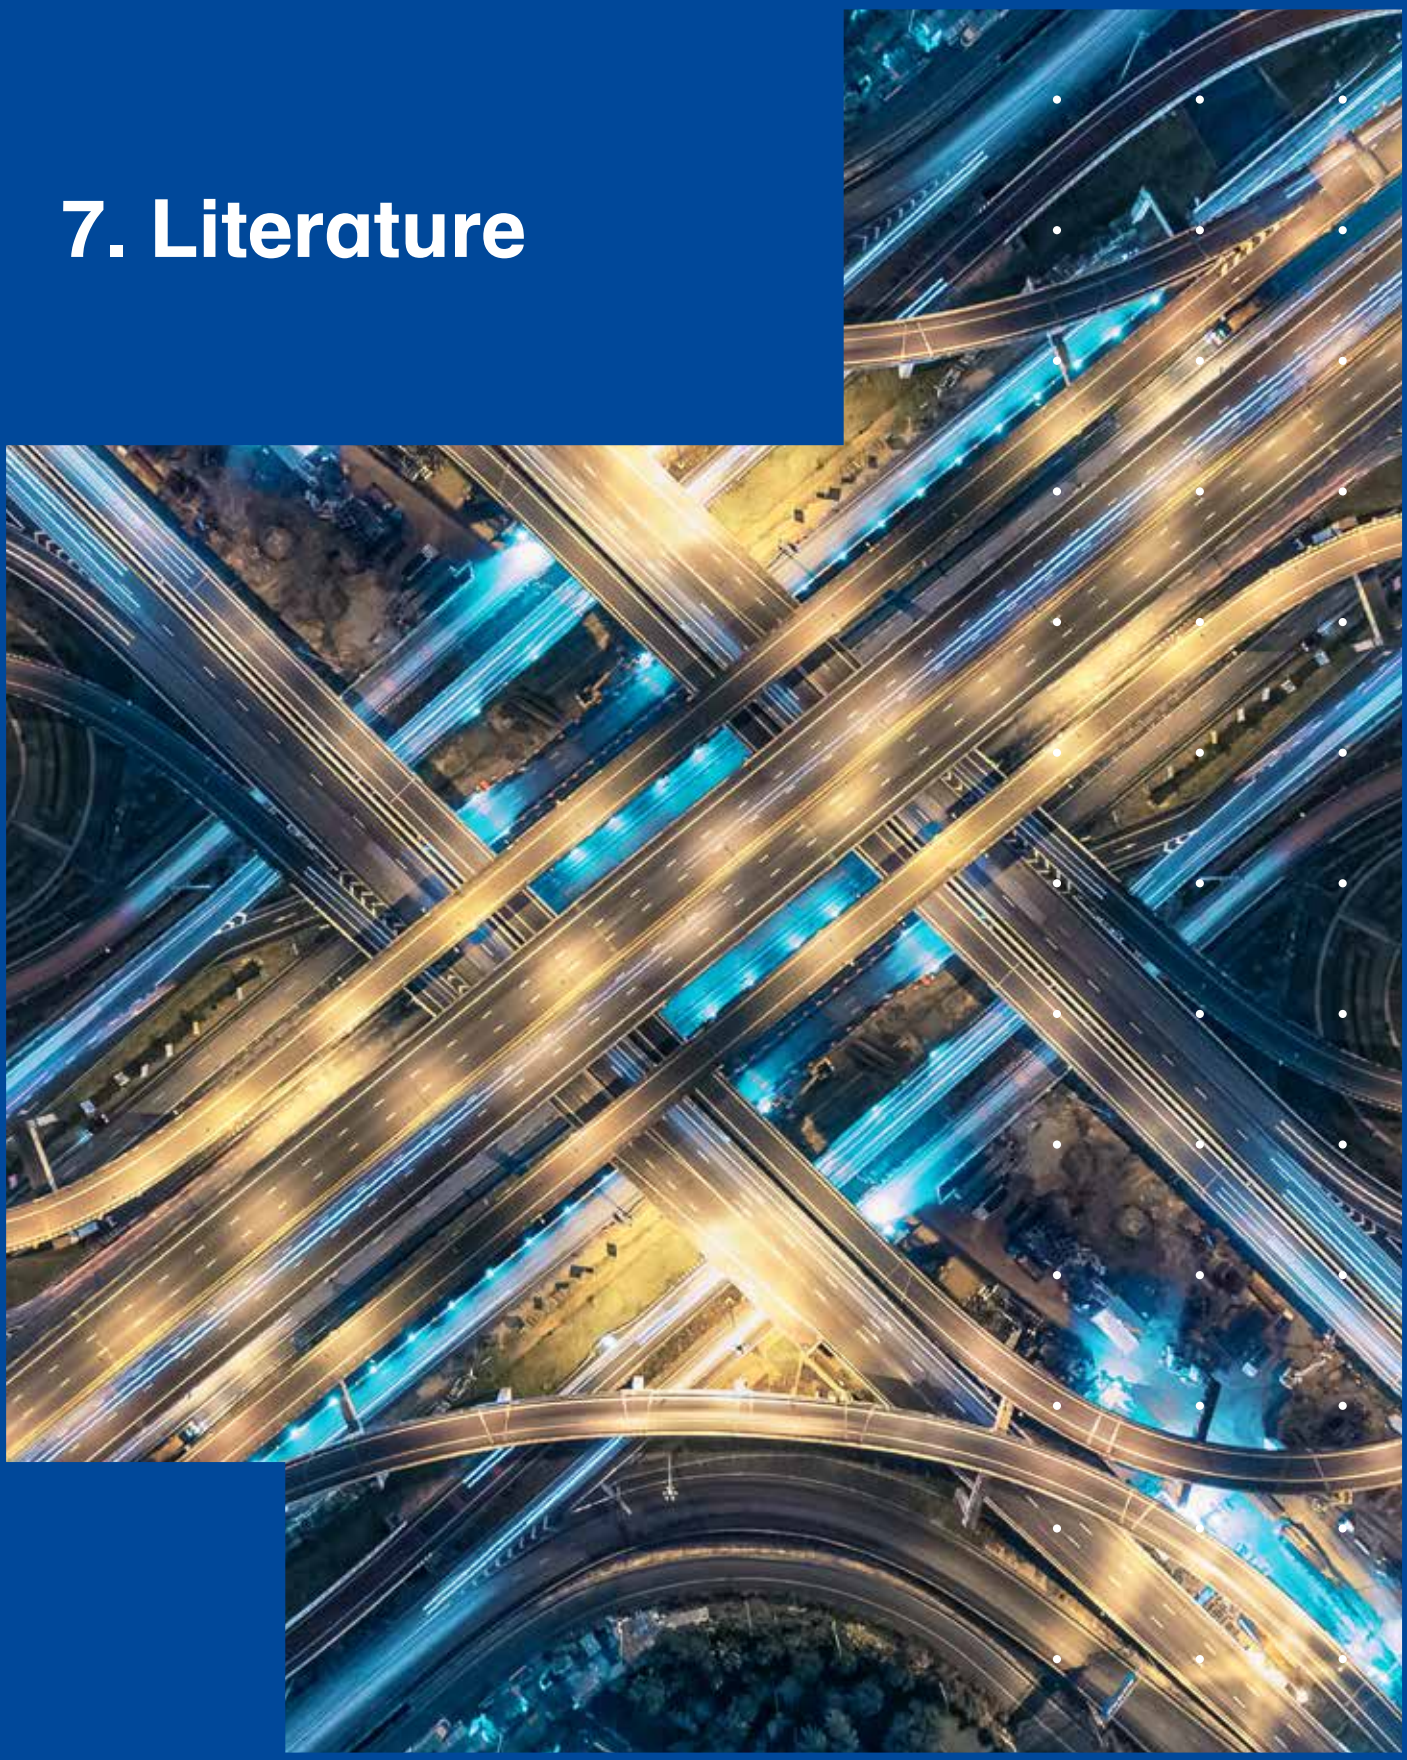

Literature

- M. Atzori, Blockchain Technology and Decentralized Governance: Is the State Still Necessary? (December 1, 2015). Available at SSRN: https://ssrn.com/abstract=2709713 or http://dx.doi.org/10.2139/ssrn.2709713
- S. Bano, A. Sonnino, M. Al-Bassam, S. Azouvi, P. McCorry, S. Meiklejohn, and G. Danezis. Consensus in the age of blockchains. CoRR, abs/1711.03936, 2017.
- V. Buterin. Ethereum: A next-generation smart contract and decentralized application platform. https://github.com/ethereum/wiki/wiki/White-Paper, 2013.
- C. Cachin and M. Vukolic. Blockchain consensus protocols in the wild. CoRR, abs/1707.01873, 2017.
- M. Campbell-Verduyn (Ed.). (2017). Bitcoin and Beyond: Cryptocurrencies, Blockchains, and Global Governance. Routledge.
- M. Castro and B. Liskov. Practical byzantine fault tolerance and proactive recovery. ACM Trans. Comput. Syst., 20(4):398–461, 2002.
- A. Efe Gencer, S. Basu, I. Eyal, R. van Renesse, and E. Gün Sirer. Decen- tralization in Bitcoin and Ethereum Networks. ArXiv e-prints, January 2018.
- I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In: 18th Int.. Conf on Financial Cryptography and Data Security (FC 2014), Christ Church, Barbados, March 3-7, 2014, Revised Selected Papers, pages 436–454, 2014.
- M. Finck, Blockchain Regulation (August 7, 2017). German Law Journal, 2018, Forthcoming; Max Planck Institute for Innovation & Competition Research Paper No.1713 Available at SSRN: https://ssrn.com/abstract=3014641 or http://dx.doi.org/10.2139/ssrn.3014641
- P. de Filippi and A. Wright. Blockchain and the Law: The Rule of Code. Harvard University Press, 2018.
- J. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol: Analysis and applications. Cryptology ePrint Archive, Report 2014/765, June 2017. https://eprint.iacr.org/2014/765.
- L. Lamport, R. E. Shostak, and M. C. Pease. The byzantine generals problem. ACM Trans. Program. Lang. Syst., 4(3):382–401, 1982.
- S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Oc- tober 31 2008.
- A. Narayanan, J. Bonneau, E. Felten, A. Miller and S. Goldfeder, Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction, Princeton University Press, 2016.
- A. Narayanan, J. Clark, Bitcoin's Academic Pedigree, In: Communications of the ACM, Vol. 60 No. 12, Pages 36-45. Dec 2017. DOI 10.1145/3132259
- M. E. Peck. Do you need a blockchain? IEEE Spectrum, 54(10):38– 39,60, oct 2017.
- W. Prinz, A.T. Schulte (eds), Blockchain und Smart Contracts, Technologien, Forschungsfragen und Anwendingen, Fraunhofer Gesellschaft, Nov. 2017.
- Y. Sompolinsky, A. Zohar, Bitcoin's Underlying Incentives, Communications of the ACM, Vol. 61 No. 3, Pages 46-53. March 2018. DOI 10.1145/3152481
- M. Swan, P. de Filippi, Toward a Philosophy of Blockchain: A Symposium: Introduction. In: Metaphilosophy 48(5), Oct 2017.
- F. Tschorsch and B. Scheuermann, Bitcoin and Beyond: A Technical Survey on Decentralized Digital Currencies. IEEE Communications Surveys & Tutorials 18(3), 2016.
- G. Wood. Ethereum: A secure decentralised generalised transaction ledger. EIP-150 Revision. (The Ethereum Yellow Paper), 2017.
- Karl Wüst and Arthur Gervais. Do you need a blockchain? IACR Cryptology ePrint Archive, 2017:375, 2017.
- J. Yli-Huumo, D. Ko, S. Choi, S. Park and K. Smolander, Where Is Current Research on Blockchain Technology?—A Systematic Review. In: PLOS ONE, Oct 2016. DOI:10.1371/journal.pone.0163477.

Appendix: Composition of the Committee

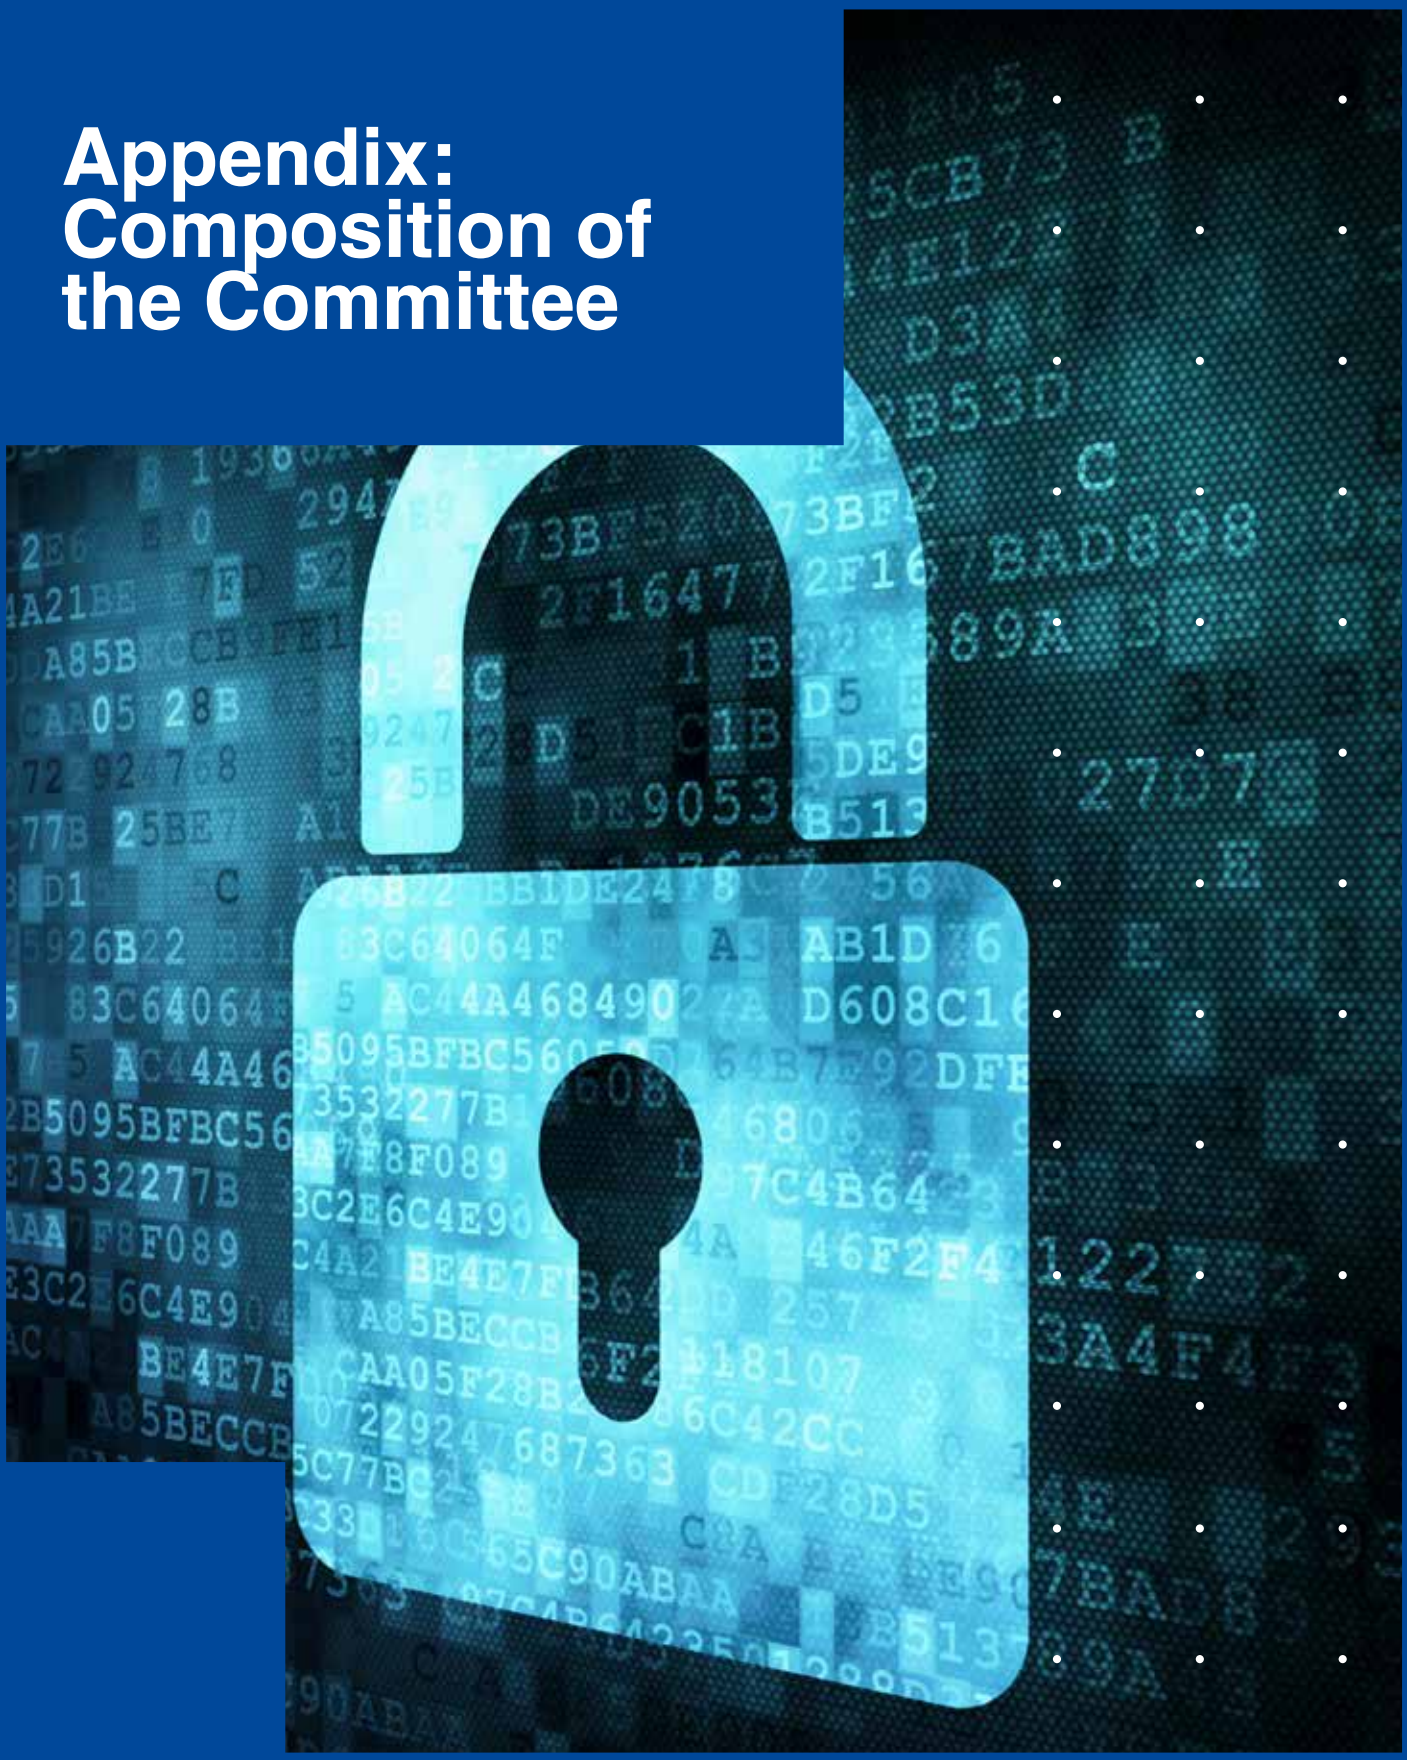

Composition of the Committee

Prof. Jaco van de Pol (chairman), U of Twente, Computer Science

Dr. Balázs Bodó, U of Amsterdam, Fac. of Law, Institute Information Law
Dr. Ir. Oskar van Deventer, TNO, senior scientist, director TNO Blockchain Lab
Dr. Jaap-Henk Hoepman, Radboud U Nijmegen (CS) & U of Groningen (Law)
Prof. Jeroen van den Hoven, TU Delft, Philosophy
André de Kok, Ministry of the Interior, Rijksdienst voor Identiteitsgegevens
Marjan van der Plas, ABN Amro, Innovation Manager Blockchain & Smart Contracts
Dr. Ir. Johan Pouwelse, TU Delft, Computer Science
Prof. Marcel Thaens, Erasmus U Rotterdam & PBLQ
Dr. Marc Stevens, CWI Amsterdam, Cryptology

Organizational support by NWO:
Dr. Ruben Sharpe, senior policy officer

Acknowledgments.

The committee acknowledges the numerous contributions received from the community, in the form of reports, inspiration from around 40 experts in the expert panel, and feedback from around 25 professionals from government, industry, and research institutes in the open consultation meeting.

NWO | Netherlands Organisation for Scientific Research

dutch
digital
delta

www.dutchdigitaldelta.nl/blockchain