Delft University of Technology

Laws for Creating Trust in the Blockchain Age

Pouwelse, Johan; de Kok, André; Fleuren, Joost; Hoogendoorn, Peter; Vliegendorhart, Raynor; de Vos, Martijn

DOI

[10.1515/eplj-2017-0022](https://doi.org/10.1515/eplj-2017-0022)

Publication date

2017

Document Version

Accepted author manuscript

Published in

European Property Law Journal

Citation (APA)

Pouwelse, J., de Kok, A., Fleuren, J., Hoogendoorn, P., Vliegendorhart, R., & de Vos, M. (2017). Laws for Creating Trust in the Blockchain Age. *European Property Law Journal*, 6(3), 321–356.
<https://doi.org/10.1515/eplj-2017-0022>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Research Article

Johan Pouwelse*, André de Kok, Joost Fleuren, Peter Hoogendoorn, Raynor Vliegendorhart, and Martijn de Vos

Laws for Creating Trust in the Blockchain Age

Abstract: Humanity's notion of trust is shaped by new platforms operating in the emerging sharing economy, acting as intermediate matchmaker for ride sharing, housing facilities or freelance labour, effectively creating an environment where strangers trust each other. While millions of people worldwide rely on online sharing activities, such services are often facilitated by a few predatory companies, managing trust relations. This centralization of responsibility raises questions about ethical and political issues like regulatory compliance, data portability and monopolistic behaviour. Recently, blockchain technology has gathered a significant amount of support and adoption, due to its inherent decentralized and tamper-proof structure.

We present a blockchain-powered blueprint for a shared and public *programmable economy*. The focus of our architecture is on four essential primitives: digital identities, blockchain-based trust, programmable money and marketplaces. Trust is established using only historical interactions between strangers to estimate trustworthiness. Every component of our proposed technology stack is designed according to the defining principles of the Internet itself: self-governance, autonomy and shared ownership. Real-world viability of each component is demonstrated with a functional prototype or running code. Our vision is that the highlighted technology stack devises trust, new acts, principles and rules beyond the possibilities in current economic, legal and political systems.

Keywords: Blockchain, Trust, Programmable Economy, Transactions

*Corresponding author: Johan Pouwelse, Department of Software Technology, Delft University of Technology

André de Kok, National Service for Identity Data, Netherlands

Joost Fleuren, Chamber of Commerce, Netherlands

Peter Hoogendoorn, ABN AMRO

Raynor Vliegendorhart, The Netherlands' Cadastre, Land Registry and Mapping Agency

Martijn de Vos, Department of Software Technology, Delft University of Technology

1 Introduction

Humanity has a disposition to trust, meaning the tendency of individuals to be willing to depend on others. Technology is altering our notion of trust. New technology platforms such as Uber and Airbnb disrupt how taxi services are provided and the way in which apartments are rented out.

For these platforms, trust is an essential element: people step into the car of a person they never met before or let strangers sleep in their spare room without worries. Trust is addressed by rating systems, the foundation for their quality and safety measures. For instance, pre-testing for potentially dangerous drivers and pre-screening is reduced and replaced with a continuous quality monitoring system.¹ The Uber mobile application asks both drivers and passengers to provide feedback on their experience. The software averages the one to five-star ratings of the most recent 500 trips.² If this average long-term quality rating of a driver drops below a certain level, it automatically triggers contract termination.³ The traditional dynamics of trust are altered on such platforms. Employment regulations are under pressure: some courts found that Uber violated local taxi regulations.^{4 5 6} For the United Kingdom these platforms are responsible for lowering unemployment levels, reducing fixed employment contracts by a few percent and enforcing the trend towards zero-hour contracts and self-employment. A 2017 United Kingdom government report states that currently 1.1 million people work within this gig economy.⁷ The laws which govern trust, employment

1 Krysia Lenzo, “Understanding Uber’s five-star rating system” (23 February 2016) (<http://www.cnbc.com/2016/02/23/understanding-ubers-five-star-rating-system.html>) visited on 23 July 2017.

2 Uber, “Rate a driver” (<https://help.uber.com/h/7b64dda6-78f5-4575-b7da-3c9e40d2c816>) visited on 23 July 2017.

3 Ellen Huet, “How Uber’s Shady Firing Policy Could Backfire On The Company” (14 October 2014) (<https://www.forbes.com/sites/ellenhuet/2014/10/30/uber-driver-firing-policy/#89f0c6e1527d>) visited on 23 July 2017.

4 Duncan Robinson, “Uber faces regulation in Europe as transport company” (11 May 2017) (<https://www.ft.com/content/6f4ac284-362b-11e7-99bd-13beb0903fa3>) visited on 23 July 2017.

5 Nikolaj Skydsgaard, “Danish prosecutor indicts Uber over driver violations” (2 December 2016) (<http://www.reuters.com/article/us-uber-denmark-idUSKBN13R14G>) visited on 23 July 2017.

6 Hyunjoo Jin, “South Korea court says Uber violated transport law, latest setback for U.S. firm” (26 April 2017) (<http://www.reuters.com/article/us-uber-tech-southkorea-idUSKBN17S09F>) visited on 23 July 2017.

7 RSA, *Good Gigs: A fairer future for the UK’s gig economy* (2017) .

terminations, and access to marketplaces are no longer the exclusive domain of national law. Laws of trust are now partly inside commercial software.

This work aims to help transform the gig economy. We aim to institutionalize the empowerment of workers and small businesses with our software, rules and architectural principles.

First, we design and create a generic, decentralized matchmaking platform for two-sided markets. Our non-profit proof-of-concept application aims to be more transparent, fair, and open. Second, we devise an alternative set of rules and principles for the emergence of trust, using only repeated interactions between strangers. Third, we present a blueprint for a public infrastructure and prototypes for each part of this blueprint for the gig economy. Our proposed public infrastructure in principle should be able to offer a generic non-profit alternative to current matchmaking platforms (e.g. Uber, TaskRabbit, Amazon, Expedia, Booking, etc.). Finally, we hint at how such an infrastructure might be used to realise what has become known as the *programmable economy*.

Blockchain technology is at the core of our proposed Uber alternative and blueprint for a public gig economy. However, cybocurrency-oriented ledgers such as Bitcoin and Ethereum are in our opinion unsuitable for this task or other legal processes such as land registration. Our reason is that all current work based on coin creation, blockchain consensus like proof-of-work or variants thereof lack robust and effective ties to any legal system, or even the real world in general. Additionally, the current generation of cybocurrency-based blockchain technology lacks a durable governance structure. Severe disagreements exist within these communities, as various factions battle for control of ecosystems worth billions. The cybocurrency ecosystem is unsuitable for land registration as it occasionally has periods dubbed civil war by the media.⁸ While there have been promising advances in the field of cybocurrency, we do not consider this research in the remainder of this work. We also do not consider private or consortium blockchains like R3 Corda, as they do not offer any significant advantage over classical distributed databases.⁹ While digital signatures and a block structure increases the security of a private blockchain, it lacks the security, irreversibility

⁸ Lulu Yilun Chen, “Bitcoin Is Having a Civil War Right as It Enters a Critical Month” (10 June 2017) (<https://www.bloomberg.com/news/articles/2017-07-10/bitcoin-risks-splintering-as-civil-war-enters-critical-month>) visited on 23 July 2017.

⁹ Richard Gendal Brown, “Introducing R3 Corda™: A Distributed Ledger Designed for Financial Services” (5 April 2016) (<http://www.r3cev.com/blog/2016/4/4/introducing-r3-corda-a-distributed-ledger-designed-for-financial-services>) visited on 23 July 2017.

and censorship-resistance provided by public blockchains.¹⁰ Instead, we focus on the tamper-proof leaderless database type of solutions. The statements and opinions in this paper are based on our decade-long experience of deploying and gradually improving our own ledger, installed by 1.8 million users. We deployed a very primitive fully distributed ledger in August 2007, pre-dating the launch of Bitcoin.¹¹ In 2017 we have mathematically proven that our blockchain technology, TrustChain, scales linear and surpasses a transaction throughput of 10,000 transactions per second.¹²

2 Problem Description

Contracts, transactions, and the records of them are among the defining structures in our economic, legal, and political systems.¹³ They protect assets and set organizational boundaries. These authentic records create ownership, determine citizenship, and define partnerships between entities. It establishes a governance layer in which states, economic actors, and citizens interact, effectively creating trust.

For thousands of years trusted guardians kept authentic ownership records of land and assets in general. The innovation speed of these centralized bureaucracies which govern these record is slow, while technological innovations like autonomous vehicles and virtual reality are changing the world at a fast pace. This difference in innovation speed is creating increasing levels of friction around data governance. The legal world might be at the beginning of a transition towards full digitization, additional standardization, and further automation. Already 22 years ago the book "Law in a Digital World" was published, however, most land administration systems have only recently moved to digitalized form.¹⁴

Blockchain enthusiasts claim this technology offers an alternative to any trusted guardian and central bank. It offers us another way to organise society,

10 ““Private blockchain” is just a confusing name for a shared database” (15 September 2015) (<https://freedom-to-tinker.com/2015/09/18/private-blockchain-is-just-a-confusing-name-for-a-shared-database/>) visited on 28 August 2017.

11 Colin Barras, “File-sharers forced to play fair” (31 August 2007) (<http://news.bbc.co.uk/1/hi/technology/6971904.stm>) visited on 23 July 2017.

12 Kelong Cong, “multi-chain bottom-up consensus model prototype” (12 June 2016) (<https://github.com/Tribler/tribler/issues/2457>) visited on 23 July 2017.

13 Marco Iansiti and Karim RLakhani, “The Truth About Blockchain” (January 2017) (<https://hbr.org/2017/01/the-truth-about-blockchain>) visited on 23 July 2017.

14 M Ethan Katsh, *Law in a digital world* (Oxford University Press 1995).

with a high degree of decentralization. The resilience of Bitcoin together with the rising gig economy endorses this claim. The Swedish land registration authority Lantmäteriet already explored the potential of blockchain technology and describe the various opportunities in their published report.¹⁵ The breakthrough that made the gig economy possible was the use of reputation systems to build trust between strangers. Reputation systems collect and process information about past interactions, to help people evaluate the trustworthiness of others.¹⁶ An individual's reputation on a platform such as eBay, however, is owned by a profit-driven entity. This leads to the following problems for the social good:

- *Lock-in*: a solid reputation on one platform is locked into that platform: you cannot move your reputation. Like other forms of lock-in, this inhibits competition, encourages monopoly behaviour and reputation manipulation. Ironically, companies in the “sharing economy” do not share reputations. Users are increasingly being protected by European Union law from such data silos.¹⁷ The newly defined right to data portability allows individuals to obtain and reuse their personal data for their own purposes across different services.¹⁸ The ability to move reputation allows for new incentives: for instance, an insurance company can provide discounts for a household insurance if a customer has a high no-claim car insurance bonus with another company. On the other hand, providing a usable, open standard for reputation is challenging.
- *Fragmentation*: each company operates its own closed market, only accessible through their (mobile) application. This leads to lower overall efficiency, compared to a single open market. For example, the ride-hailing market is fragmented into numerous closed markets operated by companies such as Uber, Lyft, BlaBla Car, Didi Kuaidi, GrabTaxi and Karhoo. Each isolated marketplace tries to match drivers and customers in real-time and has to overcome similar challenges.

15 Lantmateriet, Telia Company, ChromaWay and Kairos Future, *The Land Registry in the blockchain* (tech. rep., 2016) .

16 Paul Resnick and Richard Zeckhauser, “Trust among strangers in Internet transactions: Empirical analysis of eBay’s reputation system” in *The Economics of the Internet and E-commerce* (Emerald Group Publishing Limited 2002).

17 “Protection of personal data” (<http://ec.europa.eu/justice/data-protection/>) visited on 21 August 2017.

18 Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC [2016] OJ L119/88.

Centralized approaches are ultimately influenced by the goals and the reliability of the central entity or authority that controls them. One economic model even predicts the rise of predatory, monopolistic platforms within two-sided markets.¹⁹ For this reason, central platforms can never be truly generic and universal. With a decentralized approach, multiple independent individuals cooperate to build a single ecosystem and no one entity has control over the entire environment. This is the model we envision for the future programmable economy: trust relations are not locked in a single profit-driven entity.

The cardinal problem is: *who owns trust*? Creating and maintaining trust within a public, shared infrastructure has proven to be a hard problem. While there is much conducted research on this specific topic, a reliable, secure and real-world deployed trust mechanism stays out.^{20 21 22}

In May 1962 the vision of a time-sharing computer system with many remote stations was presented, known as the Internet today.²³ Nobody owns the Internet. This has been a critical factor for its decades long success story. The Internet consists of numerous Autonomous Systems which are loosely coupled and have a common numbering mechanism. On top of this global communication infrastructure we have built email, video conferencing, entertainment platforms, search engines, marketplaces, countless cloud services, and essentially a digital economy. These examples however often contain a single central point of control and authority. It has proven to be hard to create a decentralized governance layer for such vital public infrastructure.

Joseph Stiglitz co-authored the "Architecture of Economic Systems" in 1985, describing how decision making units can be organized together within a system.²⁴ It presents a basic framework to compare the performance of decentralized and centralized economic systems. Centralized economic systems have given way to decentralized forms. The current challenge is to further refine this form: storage and governance of authoritative answers to ownership questions using a public

19 Simon Loertscher and Andras Niedermayer, *Predatory Platforms* (tech. rep., Working Paper 2016) .

20 Rahim Delaviz, Nazareno Andrade, and Johan A Pouwelse, "Improving accuracy and coverage in an internet-deployed reputation mechanism" (2010).

21 Rahim Delaviz et al., "SybilRes: A Sybil-resilient flow-based decentralized reputation mechanism" (2012).

22 Sepandar D Kamvar, Mario T Schlosser, and Hector Garcia-Molina, "The eigentrust algorithm for reputation management in p2p networks" (2003).

23 Joseph Carl Robnett Licklider and Welden E Clark, "On-line man-computer communication" (1962).

24 Raaj Kumar Sah and Joseph E Stiglitz, "The architecture of economic systems: Hierarchies and polyarchies" [1985] .

and transparent infrastructure. Trust, accountability, and reputation mechanisms are closely tied. Transparency can be used to make an authority accountable in order to establish trust. It promotes integrity of operations by monitoring the correct behaviour of economic actors.²⁵

Creating trustworthy, public, transparent, and decentralized infrastructures is non-trivial. We lack a decentralized, open solution for essential economic primitives.

3 Architecture for Creating Trust

We propose an architecture to create a trustworthy decentralized infrastructure for four economic primitives: digital identities, money, trust and marketplaces. Figure 1 shows the four layers which define our architecture. The novelty of our work is the application of the Internet architecture throughout our architecture, owned by both everybody and nobody. Our work is academically pure: it relies on self-governance, and weak coupling between autonomous entities. Each of the four economic primitives are meticulously designed without relying on any middleman, they are void of any central authority, make traditional intermediaries optional, do not require any central server, remove the need for centralized databases, and even do not depend on Internet connectivity.

Our blueprint builds heavily on the concepts proven within the gig economy and combines it with blockchain technology. We believe our proposal is the first blueprint with real-world viability, as it is the only detailed proposal based on rigorous experimental science: running code. This work is based on a decade of experimentation: for each component within our architecture we crafted various generations of operational prototypes and conducted Internet-deployment tests. Each of the four primitives in our proposed economic architectural blueprint is designed to reinforce the strength, usability and efficiency of the other primitives. Upcoming sections present how we apply and validate our architecture by creating a decentralized alternative for Uber and an open, shared market for mortgage financing.

We believe that blockchain technology enables decentralized economies of higher efficiency and stability than currently known. This work broadens the applications of the organisational principle behind the Internet, offering full sovereignty to all decision making entities, full transparency on their past perfor-

²⁵ Carmela Troncoso et al., “Systematizing Decentralization and Privacy: Lessons from 15 years of research and deployments” [2017] arXiv preprint arXiv:1704.08065.

Fig. 1: Our four-layered architecture for creating trust.

mance, and openness in general. We are applying ideas fundamentally concerned with freedom of individuals, embedding humanities natural disposition to trust, and combining it with a very unforgiving blockchain-based mechanism for rule-breaking economic actors: digital ostracism.²⁶ The threat of being banned forever from an ecosystem may seem to lack compassion and decency, but boosts trust. This threat is also referred to as “shadow of the future” and game theory has shown that it helps to remove cheating incentives.²⁷ In our architecture we apply the shadow of the future to punish dishonesty, lying, cheating, and fraud. In other words, mistakes can be forgiven, but intentional digital manipulations are not.

Our work establishes a blueprint for creating a trustworthy, programmable economy. The term programmable economy was introduced by consultancy firm Gartner Inc. in 2015.²⁸ We expand upon their ideas with various architectural details, propose feasible rules, show operational prototypes, and in general mature this concept. They describe it as “the programmable economy, enabled by meta-coin platforms and smart technologies, will support new forms of value exchange, new kinds of markets (including dynamically defined on-demand markets), and

²⁶ Cristina Bicchieri, John Duffy, and Gil Tolle, “Trust among strangers” (2004) 71(3) *Philosophy of Science* 286.

²⁷ James W Friedman, “A non-cooperative equilibrium for supergames” (1971) 38(1) *The Review of Economic Studies* 1.

²⁸ David Furlonger and Ray Valdes, *Hype Cycle for Blockchain Technologies and the Programmable Economy, 2016* (tech. rep., Gartner 2016) .

new kinds of economies such as the attention economy, the reputation economy, the on-demand economy and the resource optimization economy.”²⁹

First we need to address the digital identity problem, shown as the bottom layer in Figure 1. A wealth of applications require strong authentication and long-lived secure identities. The Internet requires a common continuously evolving strong identity layer. This would make the Internet safer, better and more efficient. A single common identity layer also needs full decentralization and self-governance. This blocks progress, for instance, a single standard for globally legally valid electronic signatures which does not yet exist. This work is motivated by the lack of control over our identity, as formulated within the proposed WebDHT documentation.³⁰

“The Web currently does not have a mechanism where people and organizations can claim identifiers that they have sole ownership over. Identifiers, such as those rooted in domain names like emails addresses and website addresses, are effectively rented by people and organizations rather than owned. Therefore, their use as long-term identifiers is dependent upon parameters outside of their control. One danger is that if the rent is not paid, all data associated with the identifier can be made temporarily or permanently inaccessible. This document specifies a mechanism where people and organizations can cryptographically claim ownership over identifiers such that they control them and the documents that they refer to.”

Self-sovereign identities require that users are the rulers of their own identity. This idea was presented in 2012 by Moxie Marlinspike.³¹ With this concept people and businesses can store their own identity data on their own devices, and present it efficiently to anyone who wants to validate it. The decentralized storage within the concept removes the need for any central database of identity data. This might be increasingly important with the European Union General Data Protection Regulation.³²

29 “Gartner Says the Programmable Economy Has the Potential to Disrupt Every Facet of the Global Economy” (8 October 2015) (<http://www.gartner.com/newsroom/id/3146018>) visited on 24 July 2017.

30 “WebDHT documentation” (<https://opencreds.org/specs/source/webdht/>) visited on 23 July 2017.

31 “What is ‘Sovereign Source Authority’?” (15 February 2012) (<http://www.moxytongue.com/2012/02/what-is-sovereign-source-authority.html>) visited on 23 July 2017.

32 Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC [2016] OJ L119/88 (see n. 18).

This approach to identity is a complete overhaul to the approach used today by governments using central controlled identity administration frameworks. Users become autonomous by using self-sovereign identities and are free to store their digital passport inside their own smartphone device³³. Instead of your government assigning you citizen number "0013" or your telecommunication provider renting you cell phone number "0-013" you can claim ownership of cryptographic key "8A4D48B" as your worldwide identity. Together with novel cryptographic techniques it becomes possible to identify and authenticate people, members of organisations and objects without even requiring Internet access. A shared list of worldwide claimed identities by citizens could replace all central identity administration frameworks, increasing efficiency and data portability.

Being the ruler of your own identity is closely related to the core concept of Bitcoin. Bitcoin creates money without banks, but also provides users with an exceptional level of control. The novelty of Bitcoin is that it provides a system where nobody can stop a specific entity from spending their money. Identity autonomy provides a new level of citizen empowerment and enables financial sovereignty.³⁴

The second layer is our TrustChain fabric, a blockchain design with loose coupling and linear scalability of transaction throughput. Our work differs radically from the current generation of cybocurrency, which does not scale beyond several transactions per second. We specifically avoid the concept of coins. TrustChain exclusively focusses on the emergence of trust through repeated interactions. Instead of using complex mathematical puzzles, our basic building blocks are the interactions between untrusted entities, an idea that naturally extends to the real world. Each actor within our ecosystem operates their own unique database which contains measures against tampering and makes prior agreed transactions irrefutable. Within TrustChain each actor is completely autonomous, owns their own chain, and publishes new electronic business transactions in a tamper-proof append-only manner.

TrustChain and all other layers in our architecture utilise a new model of governance that we call self-governance. It is an ecosystem where ordinary users self-organise into a large-scale Internet-based collective which can be freely joined by anyone, with the special strong property that all authority is temporary. We

33 In this article, we approach self-sovereign identity from a technological perspective, however, this fundamentally different mechanism raises a new sociological question: do people *want* to control their identity themselves? If not, what attitude underlies this objection?

34 Jean Matouk, "ON FINANCIAL SOVEREIGNTY." [2009] *Revue d'Économie Financière*.

define a self-governance system as a distributed system in which autonomous individuals can collectively exercise all of the necessary functions of power without intervention from any authority which they cannot themselves alter. Self-governance implies a mechanism for peaceful transfer of power. Users may directly vote on new laws, changes to rules, routine maintenance updates, and key changes in principles. This approach has proven to be usable, but occasionally leads to fascinating outcomes. On 20 July 2016 the Ethereum community split apart because a majority voted to revert a completely valid transaction, created by faulty user software (known as the DAO hack).³⁵ Another option is to use representatives and appoint professionals responsible for the daily administration of the community through some voting process. However, this goes against the prevailing anti-authority culture within operational blockchain ecosystems. Distributed systems with self-governance have no external overseers and no central controlling servers.

The third layer in Figure 1 transforms our usage of money. Our proposed approach there is again the opposite of the cybercurrency community. We believe building a trustworthy and reliable financial infrastructure from scratch is unwise. By re-engineering and building upon existing bank-based payment systems we provide a realistic, legally compliant, and efficient overlay. We re-use the existing infrastructure and transform it into loosely coupled autonomous entities. We provide a proof-of-principle prototype for making international money transfers with sub-second speed, near-zero commissions, instant clearing, real-time settlement and high availability. Steps towards a programmable economy necessitate a transformation of the expensive transaction network built in 1970s (SWIFT) using a programming language from the 1950s (COBOL).³⁶

Our fourth layer consists of a generic value exchange marketplace and coordination mechanism. This enables efficient, fast, and secure usage of blockchain technology to enable trade at a large scale and coordinate existing businesses. We provide an open alternative for Uber and evaluate the efficiency of our approach using a real-world dataset. This layer is focussed on the open market showcased by the gig economy, but also offers the more traditional alternatives like currency exchange.

35 “The DAO, The Hack, The Soft Fork and The Hard Fork” (5 June 2017) (<https://www.cryptocompare.com/coins/guides/the-dao-the-hack-the-soft-fork-and-the-hard-fork/>) visited on 23 July 2017.

36 Suzan V Scott and Markos Zachariadis, *The Society for Worldwide Interbank Financial Telecommunication (SWIFT): Cooperative governance for network innovation, standards, and community* (Routledge 2017).

In a recent economics publication it is stated that with blockchains, market-places can be bootstrapped without the need of traditional trusted intermediaries, lowering the cost of networking.³⁷ Furthermore they challenge existing revenue models and incumbents's market power, and open opportunities for novel approaches to regulation, auctions and the provision of public goods, software, identity and reputation systems. Little experience, evidence, and knowledge exists for devising this fourth layer. We are currently conducting trials with various businesses to expand our understanding. This is a new, developing area of research, however, we believe too few research teams have the resources for the learn-by-doing methodology required to make serious scientific progress. Building open infrastructures with self-governance is costly (our key funding was the EU FP7 project P2P-Next of 19.500.000 Euro and QLectives of 6.900.000 Euro).^{38 39}

4 Land Registration Application

Recently the World Bank indicated that 70% of the world's population still lacks access to proper land titling or demarcation services.⁴⁰ A digital land registration system using blockchain technology requires a re-design of existing procedures for submission and verification of records and claims.

Obtaining agreement from numerous parties during the several stages of property transactions, detecting errors, and to preserve tamper-proof logging is a significant challenge. Central keepers of records are often entrusted with this critical coordination task. Blockchains clearly have potential here, offering transparency, reducing overhead, and especially providing a single-source-of-truth. Various reports support this opinion and several countries are already exploring

³⁷ Christian Catalini and Joshua S Gans, *Some simple economics of the blockchain* (tech. rep., National Bureau of Economic Research 2016) .

³⁸ "Next Generation Peer-to-Peer Content Delivery Platform" (http://cordis.europa.eu/project/rcn/85326_en.html) visited on 23 July 2017.

³⁹ "Quality Collectives: Socially Intelligent Systems for Quality" (http://cordis.europa.eu/project/rcn/89031_en.html) visited on 23 July 2017.

⁴⁰ Caroline Heider and April Connelly, "Why Land Administration Matters for Development" (28 June 2016) (<http://ieg.worldbankgroup.org/blog/why-land-administration-matters-development>) visited on 23 July 2017.

the possibilities to store land transferral transactions on a global ledger.^{41 42 43} The importance of transparency that blockchains offer is even more apparent when one considers regions in which public's trust has been damaged, e.g., in war-torn Colombia, and land registration requires the involvement of its citizens.⁴⁴

Our presented architecture for the programmable economy is designed to greatly simplify critical public infrastructure such as land ownership. As technology experts we are unable to judge if laws of various countries are ready to transition to a world where analogue printouts hold less authority and the digital world is the authoritative source of truth. To date it seems countries have difficulties recognizing the electronic signatures created using self-sovereign identity systems. While there are various established parties that benefit from uncertainty in land registration, a transparent and searchable infrastructure is key for realising an efficient land allocation mechanism and reducing costs.

The blockchain technology of today is sufficiently mature to start prototyping and understanding the (possible) advantages for land registration. Deeds that currently are stored in a centralized cadastral register, such as land purchases, mortgages and splits, could easily be mapped onto transactions in a blockchain. Things like inheritance rights or easements could be implemented through smart contracts which could rely on, e.g., the existence of attribute attestation transactions (see Section 6.3). In addition to recording the aforementioned transactions, land registration also requires its registers to be searchable and thus to be robust and highly available in order to provide services and certainty to citizens. Support for queries, such as “what are the exact boundaries of my plot?” or “what are the plots in this particular area?”, can be realized through a distributed search index on top of the blockchain. An example of a real-world highly available distributed search index can be found in the Tribler file-sharing software.⁴⁵

⁴¹ Deloitte, *Blockchain applications in the public sector* (2016) .

⁴² Frederick Reese, “Land Registry: A Big Blockchain Use Case Explored” (19 April 2017) (<https://www.coindesk.com/blockchain-land-registry-solution-seeking-problem/>) visited on 26 July 2017.

⁴³ Laura Shin, “The First Government To Secure Land Titles On The Bitcoin Blockchain Expands Project” (7 February 2017) (<https://www.forbes.com/sites/laurashin/2017/02/07/the-first-government-to-secure-land-titles-on-the-bitcoin-blockchain-expands-project/>) visited on 23 August 2017.

⁴⁴ “Landrechten voor vrede in Colombia” (3 March 2017) (<http://ec.europa.eu/justice/data-protection/>) visited on 21 August 2017.

⁴⁵ Johan A Pouwelse et al., “TRIBLER: a social-based peer-to-peer system” (2008) 20(2) *Concurrency and computation: Practice and experience* 127.

While there are possibilities for a blockchain capable of recording land transferral, the adoption of this technology gives raise to complicated regulatory issues, elaborated in the 2016 report of law firm Clyde & Co.⁴⁶ One of the key issues involves the geographical location of data storage, which is under subject of national law. Since the blockchain is an inherent decentralized mechanism, pinpointing the exact location of a specific piece of data (and taking the appropriate cross-border action) might be challenging. Another issue is that computer code often does not contain any notion of law. This questions legal enforceability of smart contracts since a contract involves concepts like acceptance, certainty and consideration, involving human decision making. The digital world does not consider or follow these contractual concepts.

While this technology is not yet sufficiently understood to fully underpin such an essential public service, it will be soon. Similar to self-driving cars, it is likely that legal issues will decelerate adoption of blockchain technology until an adequate legal foundation has been developed.

5 Open Market for Mortgage Finance

We created a minimal prototype to offer mortgage financing on an open, decentralized and blockchain-regulated market. Our application enables financial service providers like banks to offer consumers mortgage products and obtain the required capital investments from the global market. This enables external investors like foreign pension funds to invest in real-estate of financially solid countries with predictable return-on-investment and lower risk. Traditionally, the mortgage market is inaccessible for the public and nontransparent. This prototype aims to explore a viable and open alternative to decrease overhead of mortgage processes and offers a real-estate aftermarket where investors can negotiate agreements and trade. A screenshot of the developed application is presented in Figure 2, showing the user interface from the perspective of a financial institution.

Each agreement reached between individuals in the market are stored as a double-signed contract on a public blockchain. The mortgage financing platform distinguishes between three different types of contracts:

- *Mortgage contract*: this contract represent the initial mortgage agreement between a user and a mortgage provider, containing all relevant informa-

⁴⁶ Clyde & Co, *Blockchain and the Law* (tech. rep., 2016) .

Mortgage Market	Mortgages	Campaigns	Blockchain	Role: financial institution	Logout
-----------------	-----------	-----------	------------	-----------------------------	--------

My campaigns			
Amount needed	Amount invested	Mortgage id	End time
€100,000.00	€25,000.00	1	10/18/2017

Investment offers				
Investor	Amount	Interest rate	Mortgage id	Status
Duncan Sosa	€10,000.00	5%	3	PENDING
Louis Ballard	€15,000.00	2%	1	PENDING

Fig. 2: The user interface of our mortgage market prototype, from the perspective of an investor.

tion about the agreed mortgage (i.e. the house address, mortgage rate and redemption agreements).

- *Investment contract*: this type of contract is created when (a part of) a mortgage is sold to an investor. It holds properties of the resold mortgage.
- *Transaction contracts*: when a mortgage is transferred from one investor to another, a transaction contract is created.

Investment and transaction contracts are expected to depend on another contract. For investment contracts, this should be a mortgage contract whereas transaction contracts should have a investment contract as dependency. Contract dependencies, originating from a single mortgage contract, provide a public overview of all mortgage ownership transferrals since the existence of the mortgage contract.

While our prototype only stores basic attributes of the property like address and price, a future milestone can be to store a material passport of a house. A material passport stores all used materials of a building and acts as a key

component in the *circular economy*. The circular economy is about preventing the use of new resources while reusing materials by recycling them.⁴⁷

To store contractual agreements, a blockchain has been designed and implemented, specifically suitable for transactions that transfer ownership of assets between entities, in this scenario, mortgages. After a contract has been digitally signed by both involved parties, both transaction participants send the contract to all banks (we assume that financial institutions are always available and connected on our platform). Periodically, banks will try to add a new block, containing one or more received contracts, to the blockchain. Contracts can only be appended to a transaction block if all of the contract's dependencies are already available on the blockchain. Invalid contracts will not be added to the blockchain, i.e. a contract that defines a new mortgage on an address already coupled to another mortgage (double-spending). The adopted consensus mechanism here is proof-of-work: while the scalability of this mechanism is limited, it is viable for our prototype since mortgages are usually not negotiated and traded at a high rate. A dynamic difficulty target mechanism assures that on average, one new block is added to the chain every three minutes.

6 Technology Portfolio for Identities and Trust

We now expand on each of the four layers of our envisioned architecture, introduced in Figure 1, for a programmable economy. The detailed architecture of the economic primitives digital identity, trust, programmable money, and marketplaces will be elaborated and is presented in Figure 3. Most of the defined elements of Figure 3 are briefly evaluated by conducting a small experiment with an engineered prototype.

6.1 Physical Unclonable Functions

Some foundations of security are based on the laws of physics. For numerous years scientists have searched for a secure basis for critical infrastructure and we believe physical devices provide the desired basis. We specifically avoid using software and the risk of implementation bugs or weaknesses. Physical devices for

⁴⁷ “The Material Passport As Next Step In Circular Economy” (14 July 2017) (<https://materia.nl/article/material-passport-next-step-circular-economy/>) visited on 28 August 2017.

Fig. 3: Our technology portfolio for trust creation. This figure provides a detailed overview of our technology portfolio as compared to Figure 1.

security purposes are already widely used, mainly as a means to implement a two-factor authentication system.

A Physical Unclonable Function (PUF) is a device that is easy and cheap to produce but practically infeasible to duplicate, due to minor variations in the manufacturing process of the hardware.⁴⁸ Their typical usage can be found in applications that require a high level of security, for instance, self-sovereign identity solutions. The device offers tamper-proof and safe storage of cryptographic keys, representing one's identity. Each PUF device contains a unique fingerprint, determined by randomness of embedded components. A PUF responds to challenges and leads to unique but unpredictable responses, together forming a challenge-response pair. These challenges are often triggered by pushing a physical button attached to the device.

A secure process for identity storage based on a PUF device proceeds in two phases: the enrollment phase (Figure 4a), where a cryptographic key, based on a PUF fingerprint is generated and the reconstruction phase (Figure 4b), which

⁴⁸ Mafalda Cortez et al., "Modeling SRAM start-up behavior for physical unclonable functions" (2012).

Fig. 4: Procedure to store and restore a cryptographic key with a PUF device.

restores the secret cryptographic key that was generated during the enrollment phase. Each phase will now be explained according to Figure 4.

- *Enrollment phase*: first, a so-called reference response of the targeted PUF is measured after initiating a challenge of the device. This reference response is used as input for the Fuzzy Extractor, which consists of a privacy amplification module and Error-Correcting Code (ECC) encoding. The privacy amplification module converts the PUF reference response to a usable cryptographic key. Additionally, some helper data is computed and saved in storage attached to the device. This helper data itself is not sufficient to restore the secret key and is used during the reconstruction phase.
- *Reconstruction phase*: reconstruction of a stored key starts by measuring the PUF response, which is used as input for the Fuzzy Extractor. The helper data, generated during enrollment, is used to perform information reconciliation and generates the PUF reference response. After privacy amplification, the programmed cryptographic key is restored and ready for usage.

PUF devices can also be utilized for identification purposes.⁴⁹ An authentication mechanism based on PUFs works as follows: imagine a bank that needs to identify customers. For each customer, the bank produces a PUF device and stores an initial set of challenge-response pairs securely in their database. Next, the device is given to the customer. When the customer wishes to authenticate himself, he presents the device. The bank, in possession of a set challenge-response pairs unique to this device, sends a random challenge to the hardware. If the device provides a correct response, authentication is successful. Since cloning or mathematical modelling of the device is non-trivial, this is a secure mechanism to deploy for authentication purposes.

⁴⁹ Boris Škoric, Pim Tuyls, and Wil Ophey, “Robust key extraction from physical uncloneable functions” (2005) vol. 3531.

Fig. 5: Extracted minutiae details from a fingerprint.

6.2 Biometric-based Authentication

We present a mobile biometric-based authentication prototype that does not involve any central authority.⁵⁰ A proof-of-principle mobile application has been developed for Android, capable of matching fingerprints using the built-in device camera. By only utilizing device-specific components like the embedded storage and camera, no permissioned or specialized, contractual hardware is needed.

The procedure for fingerprint acquisition is divided in three steps. First, a user opens the application on their smartphone and takes a photo of his or her finger inside an elliptic-shaped area. Next, the captured photo is processed and analysed using various algorithms to extract minutiae fingerprint details. This yields the data presented in Figure 5 where indicating feature points are indicated as small red dots. Finally, the captured fingerprint is matched against known fingerprints available in the embedded database of the mobile device. Smartphones such as the Google Pixel which are equipped with 128GB of storage space, can store up to ten million fingerprints.

We conducted a small experiment with our prototype to estimate how accurately fingerprints are matched. We conclude that the mean fingerprint

⁵⁰ JS Hammudoglu et al., “Portable Trust: biometric-based authentication and blockchain storage for self-sovereign identity systems” [2017] arXiv preprint arXiv:1706.03744.

matching accuracy is 55% with the highest value of 67% for thumb fingers. While this accuracy is not high enough yet for critical authentication purposes, it demonstrates the feasibility of an open-source fingerprint recognition framework, requiring minimal resources. While modern computerized systems are able to accurately match fingerprints more than 99% of the time, these systems are often patented, closed-source and expensive.⁵¹ The ability to capture, process and match fingerprints in mere seconds, renders this concept usable in various scenarios like in remote regions with limited Internet connectivity. This work provides a solution for portable trust and our framework has been specifically designed to serve as a building block for a self-sovereign identity solution.

6.3 Self-sovereign Identity

Most identity services are offered by a single authority, handing out, revoking and managing identities. This trend continued when society adopted digital solutions: issuers of digital identities like Internet Assigned Numbers Authority (IANA) and Internet Corporation for Assigned Names and Numbers (ICANN), the organizations responsible for coordination of digital namespaces on the Internet, were and still are single, large organizations that have a large influence on the layout of the Internet landscape.⁵² Some identity schemes took a small step beyond a centralized structure and offered hierarchical solutions like certificate authorities (CAs). However, such systems rely on root authorities which form a single point of failure.

As Internet usage increased, more and more digital services became available with the need for users to create their digital identity. This resulted in identity fragmentation where a user had to manage multiple identities for different services, sometimes closely related. User-centric Initiatives like OpenID, Facebook Connect and OAuth attempted to address this problem by providing a single identity service. Digital service providers are able to implement such identity services without much effort to allow users to use only one shared identity. Unfortunately, this type of identification is again regulated by mostly centralized authorities like Facebook or Google. In addition, identity service providers are able to track users across different digital services.

⁵¹ Philip Bulman, “NIST Study Shows Computerized Fingerprint Matching Is Highly Accurate” (6 July 2004) (<https://www.nist.gov/news-events/news/2004/07/nist-study-shows-computerized-fingerprint-matching-highly-accurate>) visited on 23 July 2017.

⁵² Christopher Allen, “The Path to Self-Sovereign Identity” (27 April 2016) (<http://www.coindesk.com/path-self-sovereign-identity/>) visited on 23 July 2017.

While some of the described identity solutions are used by millions of users, they do not put the user owning that specific identity in control. Another issue is that companies often want users to reveal more information about their identity than they actually need. Sometimes, this data is used for analysis of user behaviour or data mining. This raises the question whether an identity system can be designed, void of any central authority and where users can decide which attributes they wish to reveal. Entities in need of identity information should be able to verify statements without being dependent on a central authority. In addition, modification of (personal) data must be possible without relying on a trusted third party. Some laws dictate that users are able to control their data and make decisions about availability, like the right to be forgotten.⁵³ Companies that refuse such a request are likely to suffer from reputation damage. This is a prime example of institutional pressure where companies are forced to act in line with the law, without supervision of a central authority.

We present a Self-Sovereign Identity (SSI) mechanism, capable of verifying statements without the requirement for a centralized authority and void of the possibility of data tracking, giving users full control over their identity. Self-sovereign identity serves a purpose in a large range of applications, have the potential to speed up traditional, inefficient verification processes and enables inter-operability between companies, individuals and governments while respecting the privacy of identity owners. In a SSI system, civilians, legal entities and objects are able to prove statements such as "my age is at least 18" to others that legally require verification of this claim, like alcoholic shops or car rental services. A reliable SSI mechanism is an indispensable building block to empower blockchain-based applications and sector-specific processes, operating at higher levels in our technology stack.

The only situation where the input of a centralized identity provider is required is when bootstrapping an identity. Since a cryptographic key on itself does not have much legal basis, the identity should be endorsed by a public authority, like a government. Our proposed SSI mechanism consists of two type of actors: attestors and challengers.

1. *Attestors*: they provide an attestation for a certain attribute, possibly for a fee. Attestations have a weight, indicating the trustworthiness of the specific attestation. For instance, an age attestation performed by a government has a higher value than the same attestation provided by a financial institution

⁵³ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC [2016] OJ L119/88 (see n. 18).

since a government is the authoritative data source of a birth certificate. Optionally, each attestor could be assigned a reputation to prevent abuse of their authority. This reputation should be lowered when the attestor provides a wrong or malicious attestation. Attestors with a low reputation should be excluded from the system.

2. *Challengers*: Challengers wish to prove the validity of a specific statement of the identity owner. It is helpful to make attributes checked by challengers publicly available on a blockchain construction (discussed in Section 7); this allows others to find challengers that have knowledge about a specific attribute.

In our system, attribute values are encrypted and never leave the encrypted domain. Instead, all actions like statement validation and attestations are operations on encrypted data using homomorphic encryption. Statement validation is a prime example of a Zero-Knowledge Proof in which the identity owner proves truth about a statement to a challenger, without actually revealing the exact value of the attribute in the statement. This implies that an alcohol shop is able to verify that a customer reached the eligible age but is not exposed to the actual age of the customer. The described system is in particular useful when trading in the programmable economy since sellers are able to specify trade-specific identity requirements that buyers have to meet before a transaction takes place (i.e. restrictions on age, residence or reputation).

7 Blockchain Layer (TrustChain)

The TrustChain transaction fabric is designed around the notion of entities performing transactions with each other.⁵⁴ Each user in the TrustChain network maintains and grows their own chain of historical transactions. This is in contrast to traditional blockchain constructions like Bitcoin or Ethereum, where a single, global ledger is maintained, containing all transactions since the inception of the platform. TrustChain is significantly faster than Bitcoin or Ethereum and is specifically designed to scale.⁵⁵ TrustChain is simpler, does not have miners, relaxes the need for global replication of transactions, and removes the need

⁵⁴ Pim Otte, Martijn de Vos, and Johan Pouwelse, “TrustChain: A Sybil-resistant scalable blockchain” [2017] Future Generation Computer Systems.

⁵⁵ Cong (see n. 12).

for grouping transactions in blocks. We now discuss the creation, storage and dissemination of transactions recorded on TrustChain.

Fig. 6: A representation of a transaction between two users. When a transaction takes place, both involved parties digitally sign the transaction.

Imagine two users that are performing a transaction with each other. This transaction can be transferral of money, exchanging data, attribute attestation or ownership transferral of a specific asset. Figure 6 shows a representation of a transaction between aforementioned users. Both users digitally sign this transaction, acknowledging that they agree with the content of this record. After the signatures are placed, both users persist the record to their local storage.

Fig. 7: A blockchain of transactions. Each block in the chain points back to the previous block.

A natural way to organize transactions is to chain them together. Now, transactions are stored in a blockchain data structure where each block contains exactly one transaction, both signatures of the interacting participants and a pointer to the prior block in the chain. To be precise, this pointer is a description of the previous block in the form of a digital hash (any secure hashing mechanism can be used for this purpose). This idea is illustrated in Figure 7. Furthermore, each block is accompanied with a sequence number, uniquely identifying a specific block in the chain. Each user creates a genesis block and keeps track of their own

chain of transactions. Every transaction block is present in the chain of both users involved in the transaction.

The data structure shown in Figure 7 is void of any control by other users. As a consequence, a user is able to tamper with their chain of transactions by inserting, removing or reordering records, without being noticed. The integrity of this new transaction chain can be restored by recomputing all prior pointers. In this situation, other users are unable to prove malicious modifications of one's chain. Users might also decide to not append a transaction to their local chain which is tempting if the particular transaction has a negative impact on the standing of the user involved.

Fig. 8: To increase security, each block also references a block in the chain of the other transaction participant. This ensures that each block has exactly two incoming and two outgoing pointers.

In TrustChain, this vulnerability is fixed by adding an additional pointer in each block, see Figure 8. This pointer references the prior block in the chain of the other transaction participant. Now, when two users are interacting, their chains get intertwined or “entangled”. This mechanism strengthens the tamper-proof property of TrustChain. As presented in Figure 8, each block has two incoming and two outgoing pointers. Note that this scheme can easily be extended to support transactions between more than two participants, by increasing the amount of incoming and outgoing pointers of a block.

As entities perform transactions with each other, they become quickly entangled with others. Figure 9 shows a part of the distributed TrustChain ledger where seven blocks, created by seven unique participants, are displayed. Again, note that each block contains exactly two incoming and two outgoing pointers.

Before a new block is added to ones chain, a validation process takes place to verify the consistency and integrity of the local chain and the correctness of the new block. This validation includes a verification of the pointers, transaction data,

Fig. 9: The tamper-proof TrustChain data structure to record transactions.

and signatures. Only if the aforementioned checks pass, the block is appended to the chain, committed to the local storage and optionally shared with other users.

Our approach differs from traditional blockchain fabrics like Bitcoin or Ethereum in various ways. Instead of a global, consistent and distributed ledger, every user maintains a personal history of interactions where in most blockchain-based systems, there exists one ledger which is acknowledged by a majority of the network participants. Consistency of the global ledger is achieved by a consensus mechanism like proof-of-work or proof-of-stake. While network consistency is essential for a cybercurrency system to prevent the double spending attack, it is not a hard requirement for a generic transaction ledger like TrustChain. Some form of consensus is reached however, namely between the involved parties of a transaction. We do not aim to prevent fraudulent operations but rather be able to detect malicious activities afterwards. While one might argue that this is a major limitation, it allows for superior scalability with regard to transaction throughput since parallel transaction processing is inherently possible in TrustChain. In addition, there are many scenarios where the requirement for global consistency leads to an unnecessary layer of complexity and inferior scalability. However, this does not mean that a global consensus mechanism is not beneficial for TrustChain and as such, we present a novel consensus model in Section 7.1. Finally, at a minimum, network participants only have to store the transactions that they are involved in, significantly lowering the storage requirements compared to other blockchains.

TrustChain blocks are designed to be disseminated and replicated throughout the network. In particular, this is important when transaction history is used as input for a reputation mechanism which needs the historical behaviour of others (see Section 7.3). Additionally, replication of blocks makes the system resistant

against network churn where users go on- and offline at a fast rate. Each user operates on their own bulk storage of blocks, resulting in partial storage of the global ledger. Collecting information of other users is challenging due to the vulnerability to various attacks, their limited resources and the burst of their interactions. Prior work investigates this problem and proposes solutions for reliable and secure collection of the interaction history.⁵⁶

We envision TrustChain as an essential building block in our programmable economy, providing the foundation for trustworthiness estimation and transaction recording. The superior scalability and reduced storage requirements make our transaction ledger suitable in the context of identity management and trading.

7.1 Scalable Consensus

Most blockchain fabrics have a requirement for a consistent state which is agreed upon by (a part of) the network. This requirement holds in particular when considering digital money (cybercurrencies) like Bitcoin or Ethereum. Consensus of the global ledger is necessary to prevent the double-spending attack where users transfer the same digital asset twice in multiple transactions. The possibility of such an attack impacts reliability and trust of the overall system, preventing community adoption. Unfortunately, most consensus mechanisms are computationally expensive and limit the global transaction throughput of the system. For instance, the proof-of-work consensus mechanism implemented in the Bitcoin fabric limits the theoretical transaction throughput to around seven transactions per second, which is by far not enough for a medium to large-sized trading platform.⁵⁷ In comparison, Visa processes several hundreds of transactions every second and in April 2017, SWIFT recorded an average of 28.38 million payments per day or around 328 per second.^{58 59} This motivates the need for a scalable blockchain.

We have designed, implemented and evaluated a fault-tolerant, horizontally scalable consensus mechanism on top of TrustChain, capable of detecting mali-

⁵⁶ Dimitra Gkorou et al., “Reducing the history in reputation systems” ICT. open.

⁵⁷ Gay Brandon, “Can the Blockchain Scale?” (13 February 2017) (<https://due.com/blog/can-the-blockchain-scale/>) visited on 23 July 2017.

⁵⁸ Manny Trillo, “Stress Test Prepares VisaNet for the Most Wonderful Time of the Year” (10 October 2013) (<http://www.visa.com/blogarchives/us/2013/10/10/stress-test-prepares-visanet-for-the-most-wonderful-time-of-the-year/index.html>) visited on 23 July 2017.

⁵⁹ SWIFT, “SWIFT FIN Traffic and Figures” (<https://www.swift.com/about-us/swift-fin-traffic-figures>) visited on 23 July 2017.

Fig. 10: The TrustChain structure extended with a checkpoint block. Each checkpoint block contains a signature of the user owning the chain and a hash of the consensus result.

cious activities performed by users like double-spending⁶⁰. As far as the authors are aware, this is the first true horizontally scalable blockchain fabric. We have built our system based on three important objectives:

1. *Reaching global consensus on a global state:* global consensus renders many types of malicious activities useless since the consistent state has consent of honest users in the network.
2. *Resistance against malicious users:* our consensus mechanism should be unaffected by the presence of malicious users, with or without purpose attempting to manipulate the outcome of the global consensus. Usually, these attacks are successful when the number of malicious users reaches a specific threshold.
3. *Horizontal scalability:* in this context, horizontal scalability means that as more users join the TrustChain network, the global transaction throughput increases. Note that the Bitcoin system is not horizontal scalable since the global transaction throughput is not dependent on the number of users in the network.

First, the TrustChain data structure discussed in Section 7 is slightly modified by adding a new type of block, called a checkpoint block. This type of block is displayed in Figure 10, presenting both a transaction block and a checkpoint block. A checkpoint block consists of a pointer to a previous block in the same chain, a number indicating the round of the consensus mechanism after which the specific checkpoint block has been created, a cryptographic description of the result of the consensus and a digital signature, generated by the owner of the chain. We modify the data structure such that the first block in a chain (the genesis block) is always a checkpoint block. The checkpoint blocks in a chain are used during detection of malicious activities.

⁶⁰ Zhijie Ren et al., “Implicit Consensus: Blockchain with Unbounded Throughput” [2017] arXiv preprint arXiv:1705.11046.

Our consensus mechanism proceeds in rounds and each round consists of two phases. The outcome of each round is a set of checkpoint blocks agreed on by the facilitators of that round. Facilitators are special users, elected during the first phase in a round of consensus. These facilitators reach consensus not on the individual transactions like in Bitcoin or Ethereum, but on the state of every chain. If a specific user is not a facilitator, it sends its most recent checkpoint block to all facilitators. When the facilitators received a sufficient number of checkpoint blocks, they start to reach consensus on all received checkpoint blocks using an Asynchronous Subset Consensus (ACS) algorithm.⁶¹ When this algorithm is finished, the facilitators send two messages to all users, first the consensus result and second a signature message where the facilitator signed the consensus result, adding authenticity to the consensus result. When a user receives the consensus result and sufficient valid facilitator signatures, he creates a new checkpoint block and appends it to his chain. Finally, a new set of facilitators is elected, the round number is increased and the process starts over again.

7.2 TrustChain Experiments

We have implemented TrustChain and the scalable consensus mechanism discussed in the previous sections. This section will focus on experimentation to assess the global transaction throughput and consensus duration. It is assumed that every user initiates two transactions per second. We investigate the effect of varying the number of facilitators in the network; due to implementation-specific constraints, our network can host at most 32 facilitators. The size of each transaction is approximately 500 bytes, resembling the average size of Bitcoin transactions. This experiment is conducted on our DAS-5 supercomputer.⁶²

The global throughput of TrustChain is displayed in Figure 11. On the horizontal axis, the network size is shown as the amount of users whereas the vertical axis denotes the transaction throughput in transactions per second. As a first observation, note the linear relationship in the figure between the amount of users in the network and the transaction throughput; this strongly indicates the horizontal scalability property for a network up to 1400 users. The throughput rate lowers when the number of facilitators is increased. This is explained by the fact that more facilitators require additional network communication between users, introducing additional computational effort.

⁶¹ Andrew Miller et al., “The honey badger of BFT protocols” (2016).

⁶² Henri Bal et al., “A medium-scale distributed system for computer science research: Infrastructure for the long term” (2016) 49(5) Computer 54.

Fig. 11: Scalability of TrustChain: the global transaction throughput in TrustChain versus the number of users for different numbers of facilitators.

Figure 11 indicates that we have indeed designed a scalable consensus mechanism, not bounded by a wasteful, expensive consensus mechanism like proof-of-work. With only a few servers, our consensus mechanism is able to reach global throughput rates surpassing that of Visa or the SWIFT payment networks.^{63 64}

7.3 Trustworthiness and Reputation

One way to quantify trustworthiness in an economic environment is by using a reputation mechanism. Reputation systems are widely used within the gig economy but the trust is often centered around a single authority. Fifteen years of research on generic distributed reputation systems has yielded a wide variety of proposed designs and possible solutions.^{65 66 67} Prior work by us challenges the problem of estimating reputation using network flow algorithms. This mechanism is called BarterCast and has been verified by a real-world implementation in

⁶³ Trillo (see n. 58).

⁶⁴ SWIFT (see n. 59).

⁶⁵ Delaviz, Andrade, and Pouwelse (see n. 20).

⁶⁶ Delaviz et al. (see n. 21).

⁶⁷ Kamvar, Schlosser, and Garcia-Molina (see n. 22).

our file-sharing software Tribler.⁶⁸ The most challenging attack on decentralized reputation systems is the Sybil attack, where a malicious user manipulates his or her standing in the network by creating multiple fake identities (sybils) and initiating interactions with them. Our recent research addresses the Sybil attack by proposing a trust mechanism, NetFlow, that bounds the profitability of such an attack. However, this mechanism is computationally expensive and requires significant resources when the size of the network grows. To date, the Sybil attack remains largely unsolved and solutions that prevent manifestation of such an attack are often complex.

We have built and evaluated a distributed reputation mechanism, using irrefutable TrustChain records as foundation. Building a reputation mechanism on top of TrustChain has several advantages. First, TrustChain records are light-weight and designed to be exchanged with other users. More importantly, the inherent tamper-proof property of TrustChain strengthens the reputation system since it becomes infeasible to tamper with past interactions, a major weakness of BarterCast.

Our reputation mechanism is based on PageRank, designed in 1998 by Larry Page, and is called Temporal PageRank since it incorporates the notion of time.^{69 70} PageRank is an accurate model that captures user behaviour when browsing the web and yields a ranking of websites based on relevance for that user. The algorithm models websites and links between websites as a network and explores this network in a structured manner. Temporal PageRank works in the same way, exploring the TrustChain data structure and analysing past transactions, including connections between them. Additionally, it is a relatively simple and cheap technique from a computational perspective, requiring only minimal resources. Temporal PageRank determines trustworthiness scores for other users from the perspective of a user performing the computation and is somewhat resistant against Sybil attacks in a sense that such attacks performed in the past only have minor influence on the outcome.

Temporal PageRank incorporates a weak form of transitive trust. Transitive trust implies that trust is not only established between two entities but is transferred upon further interactions. Specifically this indicates that if a specific entity A trusts another entity B well, he is also tempted to trust an entity C that is introduced to him by B. This is a situation that occurs frequently in the

68 Michel Meulpolder et al., “Bartercast: A practical approach to prevent lazy freeriding in p2p networks” (2009).

69 Lawrence Page et al., *The PageRank citation ranking: Bringing order to the web*. (Tech. rep., Stanford InfoLab 1999) .

70 P Otte, “Sybil-resistant trust mechanisms in distributed systems” [2016] .

Fig. 12: Performance of Temporal PageRank. The horizontal axis shows the number of users in the network and the vertical axis denotes the time it takes to estimate trustworthiness scores for all users in this network.

real world since we are inclined to trust individuals introduced by a trustworthy entity. Note that the property of transitive trust alone does not guarantee a defence against the Sybil Attack: techniques in the trust estimation mechanism should address this attack.

To demonstrate the feasibility of Temporal PageRank, we evaluated the mechanism using a real-world interaction trace, extracted from our file-sharing network Tribler for over a month. This trace includes 917 identities and around 200.000 unique transactions. The result of our experiment is displayed in Figure 12. The horizontal axis displays the number of users in the network and the vertical axis shows the time it takes to compute a reputation score for all these users. For a network with 900 identities, the total duration of Temporal Pagerank is just under 12 seconds. It is often not necessary to determine reputation of all users in the network; in most scenarios, determining reputations of identities you are likely to interact with is sufficient. However, we have demonstrated that Temporal PageRank is a feasible mechanism to use for trustworthiness estimation, even when the amount of users in the network grows.

8 Programmable Money

We now present our vision and work on programmable money. First, a novel overlay capable of performing real-time clearing and settlement will be explained.

Fig. 13: The essence of the Internet-of-Money architecture: the smartphone in the middle acts as a money router using both ING and ABN AMRO bank accounts.

Next, we show how to use primitives in this layer to build powerful business logic to perform conditional payments and build event-based smart contracts.

8.1 Real-time Clearing and Settlement

We present Internet-of-Money, a novel overlay network on top of the Society for Worldwide Interbank Financial Telecommunications (SWIFT) standard, capable of automated, trusted and fast money transfers between financial service providers by using intermediate autonomous entities. This system offers openness, blockchain-based transaction recording, compatibility with traditional banking architectures and does not require a central clearing house. Internet-of-Money has potential to play an important role when transferring value between users, for instance, within the mortgage finance market or in the area of supply finance where fast settlement is key.

Internet-of-Money is motivated by the slow and costly settlement cycles of existing bank infrastructures. While various digital financial solutions like stock exchanges and credit card providers offer sub-second transaction speeds, banks have fallen behind. While intra-banking transactions are often processed within seconds, inter-banking money transfer have a significant longer settlement duration before the money arrives in the account of a receiver, ranging from several hours up to days. International money transfers can have settlement times up to a week and are relatively costly, requiring users to pay high transaction fees. A near real-time irrevocable settlement mechanism is a catalyst for new use-cases centred around transfers in a value network.

As a first step towards Internet-of-Money, we reverse-engineered various mobile banking applications with considerable effort, including the Dutch banks Rabobank, ING and ABN AMRO, HSBC, one of the major financial service

providers in the United Kingdom and PayPal. We created specifications of the private communication protocols for each application and engineered an open client implementation, capable of communicating with the respective banks. This application is planned to be released after extensive closed-source testing using real money. A single unified interface to the services is offered, called *The Money API*, providing primitives for logging in to bank accounts, performing payments and querying for available balance and recent mutations. All operations are performed locally without the need for an additional central authority besides the financial service provider.

The core concept underlying Internet-of-Money is utilizing users, so-called *money switches*, that have accounts at multiple banks. This scenario is displayed in Figure 13. Since intra-banking transactions are often sub-second, transferring money from an ABN AMRO account to ING account proceeds by first sending the money to the ABN AMRO account of the intermediary after which he sends the money from his ING account to the final ING account. Since this only involves intra-banking transactions, settlement is realised within a minimal duration. Figure 13 shows a situation where one switch is used to realise fast relaying of money between different banks, however, the system supports multiple switches, effectively creating a *money circuit*.

The aforementioned mechanism only works when the network contains a sufficient money switches. There are several approaches to incentivize users to act as a money switch. The most basic motivation is realised by introducing a minor transaction fee in the system, rewarding users with a flat or dynamic fee when they provide money routing services for others. A simple switch selection policy is now to select the router charging the lowest transaction fee. This system can be used in conjunction with a reputation mechanism, using the TrustChain history of money routers as foundation for the algorithm (see Section 7.3). Successful money switch operations should increase the reputation of the user, based on the relayed amount of money. A switch that refuses to forward incoming money to an external bank account should be blacklisted and the reputation of this specific relay is lowered. A solid transaction history of a switch could be taken into consideration when selecting switches for a money circuit.

To reduce risk at stake when trading on markets, a technique we named *incremental payments* is proposed. Traditionally, one would exchange some assets in one transaction. However, in some situations, it is possible to split a single transaction into smaller parts (this is often the case when users are trading digital goods or when currency is being converted). Incremental payments can be implemented using the Internet-of-Money architecture since it allows for fast settlement and is built around the notion of programmable money transfers. Incremental payments provide only minimal benefits when using traditional

money transfers due to slow settlement, the requirement for multiple transactions and the lack of automated money transfers. Note that this technique does not completely cancel risk at stake; it only reduces risk to an acceptable level, depending on how the transaction flow proceeds.

Fig. 14: Fast international money transfers and currency conversion using Internet-of-Money. Distinct money circuits are indicated by different arrow styles.

Upcoming experimentation will focus on expanding Internet-of-Money to support additional banks, currency conversion and international bank transfers. For this experiment, we utilize additional type of money switches to send money across the globe in seconds. This structure is presented in Figure 14 and is the first step towards fast and trusted international payments.

8.2 Event-based Smart Contracts

Blockchain platforms like Ethereum allow entities to create and deploy digital, self-executing contracts, *smart contracts*, responsible for exchanging money, property, shares or anything of value in a transparent way without relying on a single authority. Smart contracts are perceived by some to be the silver bullet to a full automation of a digital world. However, the current generation of smart contract technologies lives in a closed world, unable to interact according to events in the real world. For instance, Ethereum smart contracts are unable to adjust their behaviour to legislative changes, respond to bankruptcies of legal entities and grasp the concept of an emerging bank run. Smart contracts on the Ethereum blockchain are defined by a broad set of programming rules, providing programmers a powerful mechanism to define complicated, decentralized applications.

However, the vast capabilities of smart contracts inevitably lead to mistakes made by programmers, resulting in security vulnerabilities. In the past, this has led to theft of large amounts of money and even caused a split in the Ethereum community.^{71 72}

To further automate and expand upon our architecture discussed in Section 8.1, we introduce *event-based smart contracts*, contracts that are limited in functionality as compared to traditional smart contracts, yet providing primitives to fully control the flow of money. An event-based smart contract specifies that a money transfer is triggered when some condition is met or when an event happens. To make these contracts readable and verifiable by lawyers and notaries, we use *Ricardian contracts* to store statements about automated payments.⁷³ The meaning of those contracts can be both interpreted by humans and machines.

Event-based smart contracts are able to reduce invoice payment cycles since traders could specify that a specific money transfer should be initiated when receiving some goods. In particular, the concept of money switches in Internet-of-Money is an example of event-based smart contracts, where a subsequent transaction is initiated as soon as the transaction that transfers money to a switch, is completed. Event-based smart contracting is a simple mechanism to integrate money transfers in a chain of events, similar to the functionality of If This Than That (IFTTT) that allows users to create an event chain out of web services.⁷⁴ To highlight the usage of these contracts in a real-world scenario, consider the scenario of a user buying a house: legal entities are able to define a contract that triggers a money transfer when the buyer has deposited the money, the mortgage conditions are accepted *and* the seller is mandated to sell the property. This additional control over money flows, provided by our open banking API, increases efficiency and allows for automation of tasks that are hard or even impossible to realise when relying on official banking applications.

71 Wolfie Zhao, “Hackers have stolen 32 million Dollar in Ethereum in the second heist this week” (20 July 2017) (<http://www.businessinsider.com/report-hackers-stole-32-million-in-ethereum-after-a-parity-breach-2017-7?international=true&r=US&IR=T>) visited on 23 July 2017.

72 Alyssa Hertig, “Ethereum’s Two Ethereums Explained” (28 July 2016) (<https://www.coindesk.com/ethereum-classic-explained-blockchain/>) visited on 23 July 2017.

73 Ian Grigg, “The ricardian contract” (2004).

74 Steven Ovadia, “Automate the internet with “if this then that”(IFTTT)” (2014) 33(4) Behavioral & social sciences librarian 208.

Fig. 15: The TradeChain data structure, recording a trade between two users.

9 Decentralized Marketplace

The final piece of our trust architecture is an operational prototype for a generic, decentralized two-sided marketplace, capable of trading generic assets like houses, currencies or bonds. The platform facilitates trading without presence of a central clearing house, responsible for matchmaking and trade processing. Every trader keeps track of their own set of buy and sell offers in an orderbook and acts as a matchmaker for others, attempting to match compatible offers. Matchmaking, clearing and settlement proceeds in a completely decentralized way where autonomous entities are directly exchanging buy and sell offers. Assets are stored in unique wallets. A wallet provides functionality to query balances, mutations and allows automated asset transfers to other traders. Our marketplace is secured against malicious behaviour and built to be operational in the presence of a large number of traders. Finally, our market is void of any transaction fee, enabling unrestricted and fair trading unlike most existing blockchain-based exchanges.

All transactions performed in the market are recorded on TradeChain, a distributed ledger that is based on TrustChain (see Section 7). The TradeChain data structure is presented in Figure 15, showing offer creation and a transaction between two users. When a trader creates a new sell or buy offer, a new record with details of the offer is created. To increase security, this record is validated and signed by a random trader, called a *witness*. A witness validates the correctness of the offer and after signing the record, it disseminates the record to other traders. When this offer has been matched with another offer, a transaction is initiated and both traders publicly commit to this trade by signing and storing a block (shown in Figure 15 as a *tx_init* block). Next, transaction participants start exchanging assets. For every transfer operation of assets, a new block

Fig. 16: An orderbook on the decentralized marketplace, implemented in Tribler. Wallet information is presented in the upper-right corner of the screenshot.

(named *payment*) will be created, specifying the volume of assets that have been transferred during a specific payment. Note that when both traders digitally sign this block, it becomes irrefutable that assets have been sent and received by both parties. When a transaction is complete, a block created and signed, finalizing the transaction (named *tx_done*). TradeChain provides a tamper-proof history of transactions and is designed to detect disputes and malicious activities during transactions. External parties can query the chain of specific traders and verify whether historical transactions have been completed successfully (indicated by the presence of a *tx_done* record). This is a powerful method to assign a trustworthiness score to each trader, based on historical encounters and honest behaviour. In the scenario of a dispute, TradeChain provides adequate information to transaction participants and trusted third parties to resolve the issue off-chain.

We implemented our decentralized marketplace and integrated it into Tribler.⁷⁵ The initial release of the marketplace supports trading bandwidth against both cyber- and regulated currencies, using the Internet-of-Money module discussed in Section 8.1. Figure 16 shows the orderbook of a trader in the Tribler software. Information about available wallets of a trader are present in the upper-right corner of the window.

⁷⁵ Pouwelse et al. (see n. 45).

To obtain insight in the real-world efficiency of our system, we created an open alternative to the proprietary Uber taxi marketplace. Uber, one of the largest companies operating in the sharing economy, is a two-sided market where taxi drivers, offering a ride to a location (represented by a sell offer), are matched with users requesting transportation (represented by a buy offer). Since there is no public, reliable dataset provided by the Uber platform, we use historical information of taxi rides published by the government of New York instead.⁷⁶ This dataset provides temporal and geographical information of pick-up and drop-off location of each taxi ride. To explore the limitations of our system, we analysed the dataset and subtracted 1050 offers during the busiest period in 2015, November 1, 00:59:57 to November 1, 01:01:16. The experiment is executed in real-time on a 42 node cluster of the DAS-5 supercomputer where we assume a total of 550 taxi drivers providing rides and 500 passengers requesting transportation.⁷⁷ To guarantee that every passenger can be matched with a taxi driver, we assume an oversupply of 10%, this corresponds to the realistic situation where taxi drivers are waiting for new passengers. During the experiment, each passenger and taxi driver create exactly one buy or sell offer and is connected to at least ten other random taxi drivers (matchmakers) in the network. The experiment starts by taxi drivers placing ride offers with an interval of 100 milliseconds on our decentralized market. After all ride offers have been placed, each passenger creates and disseminates a ride request. In the end, a total of 550 sell offers and 500 buy offers are created. When a new offer is placed, it is disseminated to other matchmakers who inform the creator of this offer about potential matches. In our system, a trader does not accept the first incoming match immediately but instead, waits for and aggregates additional matches for a duration of two seconds after which a trader accepts the optimal match and declines the others. This period should allow a trader to process and store at least 100 incoming matches. It is also assumed that only taxi drivers perform matching since they are likely to be connected to the network for a longer period of time. Matching efficiency is measured by the average distance between a passenger and a taxi driver where we aim to minimize this distance. Taxi drivers and passengers are matched based on their distance as the crow flies.

To study differences in matching efficiency between a market with a central clearinghouse and a fully decentralized network, we varied the percentage of taxi drivers performing matchmaking and the broadcast range of messages. The

⁷⁶ NYC Taxi and Limousine Commission, “TLC Trip Record Data” (http://www.nyc.gov/html/tlc/html/about/trip_record_data.shtml) visited on 24 July 2017.

⁷⁷ Bal et al. (see n. 62).

Fig. 17: The effect of decentralization on matching efficiency when emulating the taxi network of New York, where we vary the message broadcast range.

results are presented in Figure 17, showing matching efficiency as the network becoming increasingly decentralized. The horizontal axis denotes the percentage of taxi drivers performing matchmaking whereas the vertical axis represents the average distance between a matched driver and passenger. The points on each line in Figure 17 correspond to a network where 10%, 20%, 30% etc. of the taxi drivers are performing matchmaking. The only exception are the points near zero percentage on the horizontal axis: they resemble the situation where out of 550 taxis, exactly one taxi conducts all matchmaking activities, like an Uber server (thus 0.18%).

As a first observation, note that matching efficiency decreases when we decentralize the network, however, this performance loss is relatively minor. The differences in average matching distance between a fully centralized and decentralized market are only 30, 124 and 203 meters when a new offer reaches 30, 72 and 110 matchmakers respectively. When we increase the number of matchmakers in our market, the global orderbook is divided amongst more traders, explaining the loss in matching efficiency. Note that in this context, the global orderbook refers to the set of all unfilled offers available on the market. Additionally, increasing the broadcast range increases the matching performance. This can be explained by the fact that when new information is disseminated to more traders, each individual trader has increased knowledge about the global orderbook, resulting in improved matching.

10 Conclusions

Over the past few years, blockchain technology has attracted a significant amount of media attention. The adoption and growth of blockchain-powered platforms like Ethereum and Bitcoin raised questions whether blockchain is able to provide value within economic processes like trading, banking and the value chain. Besides illegal trading and various security weaknesses resulting in compromised digital assets, there are to date few real-world results demonstrating long-term viability of blockchain technology in our legal frameworks.

We propose a novel architecture to create trust. Each component of our technology portfolio is designed and implemented based on the philosophy of the Internet itself: self-governance and loosely coupled autonomous systems, interacting with each other. The essential technology to cultivate trust within our decentralized technology platform is TrustChain, a blockchain built around the notion of transacting real-world entities. Viability of our proposed technology portfolio is demonstrated in the real-world, with running code. Our aim is to adopt the programmable economy as a means to decrease organizational barriers, increase trading efficiency and provide more transparency and openness in existing and new economic systems.

While real-world viability of our blueprint has been proven, there is future work to consider. Some elements should be expanded, like the Internet-of-Money and our self-sovereign identity solution, to unlock their full potential. Additionally, we are unsure how our proposed technology exactly impacts existing legal, political and economic systems. For this, we are reaching out to companies to get insights in the implications of our systems, retrieve feedback on our architectural design and broaden our knowledge in general.

Funding: This work was funded by NWO/TKI grant 439.16.614.