Towards the Science of Essential Decentralised Infrastructures

Johan Pouwelse
Delft University of Technology

Abstract

Dependence of our society on digital infrastructures is growing daily, confronting us with an urgent task of building ethical and democratic alternatives to monopolistic big-tech platforms. We call upon the scientific community to put our talents to this challenge by creating decentralised infrastructures for trust-based economic and social cooperation. We empirically demonstrate that a public infrastructure to establish trust between peers in decentralized networks is possible at significant scale. Our work is based on over 15 years of improving our distributed systems which were used by more than a million people. We present six stringent criteria for designing trustworthy infrastructure, called *zero-server architecture*. Adhering to these principles, we designed a novel trustworthy networking infrastructure, called P2P-Apps. It enables smartphone apps to communicate without *any* servers, by forming a scalable overlay that uses our generic mechanism to build trust between peers, Trustchain. P2P-Apps are generic and can be expanded to serve as an alternative to centralized infrastructure owned by Big Tech.

CCS Concepts • **Networks** → **Peer-to-peer networks**;
• **Security and privacy** → **Trust frameworks**;

Keywords Creating trust, zero-server architecture, Big Tech

ACM Reference Format:

Johan Pouwelse Delft University of Technology. 2020. Towards the Science of Essential Decentralised Infrastructures. In *1st International Workshop on Distributed Infrastructure for Common Good (DICG '20)*, December 7–11, 2020, Delft, Netherlands. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3428662.3429744>

1 Introduction

We call for more scientific effort to solve fundamental problems for decentralized infrastructures such as peer-to-peer trust. We demonstrate that a public infrastructure to establish trust between strangers is possible without Big Tech intermediaries. We present empirical work towards this solution, one of the first at significant scale, based on a re-usable method we call Trustchain.

This work is licensed under a Creative Commons Attribution International 4.0 License.

DICG '20, December 7–11, 2020, Delft, Netherlands

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8197-0/20/12.

<https://doi.org/10.1145/3428662.3429744>

For over 15 years we have designed, implemented, deployed, operated, and improved 3 generations of experimental distributed social networks, seen over 1 million unique users, and developed a generic decentralised technology stack [15]. Our work contains a wealth of practical engineering with open-source code contributions from 173 developers [9]. This engineering-rich path-finding has provided us with unique insights on how to design trustworthy infrastructures and develop alternatives to monopolies of Big Tech. Our scientific contributions are:

- *Zero-server architecture* – a design principle for fair architectures capable of underpinning the global online economy. This design principle states that participants must retain their sovereignty to run code and choose roles in the system, form a collective based on self-governance, and use democratic decision-making. These principles provide a desirable property for commons infrastructure, that it is owned by both nobody and everybody.
- *Zero-server technology stack* – we have created an entire technology stack for smartphones based on the zero-server principle. Our Internet-deployed stack does not use *any* server or cloud provider. This stack consists of three layers. First, the zero-server communication layer at the lowest level. Unlike smartphone apps of today, that exclusively communicate with (web) servers, we pioneer direct authenticated communication between smartphone apps themselves. Communication is based on *carrier-grade NAT puncturing* for 3G/4G and local Bluetooth/WiFi when no Internet is available. Second, our stack creates a trustworthy overlay. The overlay consists exclusively of smartphones and may grow to any size. Third, our stack provides a generic mechanism to create trust at the top layer. Transactions between smartphone users are recorded with tamper-resilience measures to form Trustchain, a distributed ledger that scales [14]. Trustchain records are spread among users and used to calculate trustworthiness [20] (similar to star-ratings on Big Tech platforms).
- *Internet-deployment* – we expanded the above work into fully functional prototypes, called P2P-Apps. Our deployed P2P-Apps include 1) decentralised social networking, 2) decentralised music distribution, and 3) leaderless organisation featuring shared ownership of money with democratic decision making.

2 Problem Description

Our goal is to design a new class of trustworthy middleware – to facilitate *any* economic activity and *any* social activity, without centralized intermediaries and trust brokers. Our approach is to remove *all* central elements from the ecosystem to avoid points-of-failure of both technical (e.g. hardware malfunction) and socio-economic (e.g. freeriding) nature. We are inspired by the ecosystems based on permissionless innovation and "working anarchy" (Wikipedia, Linux kernel, Bitcoin).

Bitcoin represents a new level of financial sovereignty. A user remains in full control, if the cryptographic key remains secret. Bitcoin also solves the trust problem around coin minting and double spending, at the cost of scalability.

However, the grand challenge of building a cooperative decentralised network online between billions of strangers, is still a cardinal unsolved scientific problem. One of the key underlying issues is a facilitation of trust within open, permissionless, free-to-join and self-organising platforms. Enabled by the breakthrough in software-based reputation systems to facilitate trust, the sharing economy emerged around 2008. It brought strangers together on central platforms to share resources - spare accommodations (Airbnb) and taxi rides (Uber). However, users get locked in these platforms, where their reputation and data is owned by a profit-driven entity, creating conflicts between users' sovereignty and business interests of platforms.

We believe a *public infrastructure for creating peer-to-peer trust* is feasible. As a proof, we present the first empirical work that shows decentralized trust-building infrastructure working at a significant scale.

3 Related work

Economists have studied trust-based cooperation required for transactions for decades.

In his "The Market for Lemons" paper, George Akerlof (Nobel prize winner in 2001) describes the failing of markets, as buyers can only guess the quality of goods. Cooperation is also extensively studied at community level versus the micro-level of transactions. Overfishing, overpopulation, and pollution are all instances of social dilemma known as "the tragedy of the commons" [7].

Elinor Ostrom (Nobel prize 2009) studied groups that managed natural resources successfully and autonomously, without top-down government. She discovered eight core design principles for any group whose members must work together to achieve a common purpose. The essential principles are: self-governance, democratic decision-making, prevention of free-riding, and trust. Her field work showed, contrary to conventional thinking, that communities are able to self-organize in ways that punish free-riders who use common resources without contributing.

Figure 1. P2P-App user interface with network overlay screen and key exchange

Another research direction studies trust required for societal cooperation, such as indirect reciprocity. Indirect reciprocity can be described as "I will help you if you have helped someone" [17]. Game theory (Nobel prize 1994) incorporates reciprocity into rigorous models and provides fertile theoretical grounding. Yet, the standard economic approach cannot account for altruism and our need for fairness ("dictator" games, Nobel prize 2017) [5].

The multi-agent systems community has a long tradition of investigating formal trust models. Ideas have been numerous, with little advancement around labor-intensive empirical grounding. No academic design has yet been demonstrated to *work at scale* in the real world after many years past since their publication [2, 4, 8, 11, 13, 18]. Recent work from 2020 also does not address the core problem, leaving data outside user control [16]. A hard open problem is how to spread information on who is trustworthy and who is not through the network in an attack-resilient manner.

There is also the problem of enabling pseudonymous or anonymous interactions in such infrastructures. The Tor network represents the state-of-the-art in privacy-enhancing technology, composed exclusively of voluntarily-operated servers. However, managing this common resource has proven to be difficult. Various designs were proposed to prevent the tragedy of the commons, such as, GoldStar, PAR, BRAIDS, LIRA, TEARS, and TorCoin [10]. Finally, initiatives such as OpenBazaar, Steem.io, GNU Social, Mastodon, Diaspora offer an alternative to Big Tech services. They do not, however, directly address the trust problem, lack efficient mechanisms of democratic decision making, and often rely on server federation.

4 Zero-server architecture and stack

The zero-server concept is the key insight which enables public trustworthy cooperation infrastructure, suitable for any socio-economic activity. We present both our architecture and fully implemented technology stack (secure communication, overlay, and Trustchain).

Zero-server concept – Bitcoin, BitTorrent file sharing, and the Scuttlebutt gossip protocol are rare examples where no central server or trusted third party controls the ecosystem. Other online systems are almost without exception (in)directly controlled by a single central entity or multiple federated servers. Our zero-server concept provides users with participatory freedom and data sovereignty. We define a *zero-server architecture* as a network architecture without hierarchy (i.e. same responsibilities and capabilities for peers), no intermediaries exist, no single point-of-failure exists, participants have full data sovereignty, a democratic decision-making process is used, and the ecosystem itself is based on self-governance.

These criteria are stringent and considerable rigor is required to meet them. It goes beyond classical *logical* decentralisation, by providing decentralised ownership.

Zero-server secure communication – Based on this zero-server concept we designed and implemented a complete technology stack for smartphones with three layers. Our secure communication primitive between any pair of smartphone users forms the lowest layer [1, 19]. Due to pervasive deployment of carrier-grade NAT hardware, no such direct communication is possible normally [12]. Based on extensive experimental studies of NAT behavior we pioneered a NAT puncturing method for 3G/4G networks [6, 19]. After establishing a direct communication path we provide mutual authenticated communication between users, identified by their public key. It supports both offline, online, and mixed connectivity seamlessly. Networking details such as IPv4 or Bluetooth addresses are abstracted away. We use a standard challenge/response protocol with protection against spoofing, man-in-the-middle, and replay attacks.

Zero-server trustworthy overlay – On top of this communication primitive we designed the first overlay network consisting exclusively of smartphones [1]. Each smartphone connects directly to an ample amount of neighbors within the trustworthy overlay, see Figure 1 for two screenshots.

The "Peers" screen shows a real-time view into the connected peers of the overlay. For each peer it shows their public key, Internet addresses (IPv4), number of seconds since receiving data, seconds since sending data, and network latency. The cryptographic public keys are shown only for those peers which already authenticated. To avoid using a bootstrap server, we use a list of smartphones with well known fixed publicly connectable IPv4 addresses. Mobile Internet providers such as T-Mobile and Vodafone do

Figure 2. Measurements of Trustchain deployment

not offer connectable Internet addresses. We manage to bypass these restrictions through assistance from other smartphones, without using a coordinating (STUN) server [19].

The screen titled "add nearby contact" depicts the manual key exchange using QR codes. Full details of our protocol, carrier-grade NAT traversal, open source code, and our developers community are on Github [1].

Trustchain – The creation of trust is based on our integrated storage mechanism dubbed Trustchain. It operates on top of our network overlay. Figure 2 shows several years of Trustchain data, collected by crawling our public network. Our distributed ledger was first deployed in 2007 and received numerous upgrades which continue today. It shows 139 million Trustchain records collected from 87.700 users in total, records from earlier primitive ledgers are omitted. On most days around 150.000 and 250.000 new records are created and exchanged. The peaks are due to press attention and new software releases. The two visible dips are due to network failure at our crawlers. *Trustchain records* are a surprisingly simple data structure. One record merely includes the public key of the peer providing help, the amount of help, and the receiving peer identity. We do not restrict the nature of "helping", it may be storing a file or answering a query. All the records of one particular peer together show freeriding, neutral, or altruistic behavior.

Together these records contain a global ledger on provided and offered service within the entire ecosystem. Standard cryptographic measures are in place to prevent tampering, such as sequence numbering, hashing, cryptographic signatures, and append-only logs. Full specifications may be found in our IETF Internet Standard draft [14].

The storage layer is merely responsible for tamper-resilience of the interaction records (e.g. duplicate detection, preventing flooding). The application itself is responsible for: the believability of these self-signed records, indirect reciprocity policies (freeriding), and punishment of fraudsters. Our effective strategy is to assume no trustworthiness for new peers and enforcement of a continuous positive balance with

Figure 3. P2P-App: an Internet-deployed social network

the community. In prior work we provide the mathematical proofs of our Sybil-resilient peer-ranking policies [20].

5 Decentralised social networking

We created a proof-of-principle prototype of a social network on top of our zero-server technology stack. Our Internet-deployed prototype features the ability to connect to friends, exchange messages and view your friend feed (Figure 3). Exchange of animations with friends is also supported through a TFTP implementation. Obviously our open source academic work only represent a small subset of typical social network features.

Our primary interest is the scalability of our zero-server architecture work. The results of an initial experiment with real hardware are shown in Figure 4. We test the impact of maintaining a direct live communication connection with numerous friends. During our experiment we increase our connected friends in steps and capture the resulting messages per minute of our overlay protocol. Results show a roughly linear increase in message workload when increasing the number of direct connected friends. We also aim to quantify the performance impact of not using any server. From another hardware device we continuously probe the responsiveness of the hardware under testing. We send ordinary ICMP network pings in order to roughly estimate responsiveness. Network latency of usually around 50 ms is fairly typical for a smartphone device, without any signs of device overload. Our efforts are the first to provide empirical results in this direction with the Trustchain layer in place to pioneer new collaborative moderation processes. These preliminary finding support the notion that – in principle –

Figure 4. Scalability experiment with up to 500 connections

it is possible to offer a social media experience without any server and without any company.

6 Music Industry in a P2P-App

We identified the music industry as an area where a single open platform could replace *all* intermediary roles. We have designed and implemented a fully operational alternative for music distribution from the artist to their fans, without any middleman. Our functional P2P-App enables musicians to self-publish music and receive payments without any of this money going to record labels, producers, bankers, credit card companies, or tech companies. We are expanding our user community by offering Creative Commons licensed music. Real-world usage enables our ongoing deployment of distributed machine learning algorithms and trails of music recommendations [3].

Our work differs from related initiatives (Mycelia creative passport, Musicooin, OPUS audio). Due to our zero-server approach the platform is a public good, enabling 100% of money to go to artists.

Figure 5 shows the ability to directly rewards artists for their work using Bitcoin. Audio playback is implemented using Bittorrent streaming. By using Trustchain each artist can self-publish their recordings along with a Bitcoin address, and other metadata. A Decentralized Autonomous Organization (DAO) is an entity which only exists in software. Our P2P-App for music services is called the *MusicDAO*. In the next section we explore generic DAO primitives we experimented with. Our findings support the notion that – in principle – it is possible to disrupt all intermediaries within a value chain by replacing them with software.

Figure 5. P2P-App: Our music distribution platform using Bitcoin and Bittorrent

Figure 6. DAO voting for new users and sending EuroTokens

7 DAO voting and shared ownership of money

We now present experiments to understand the principle economic building blocks of the upcoming "robot economy". Our goal is to design an international software-only entity with self-governance to offer leaderless alternatives to ICANN, Wikipedia and the Linux Foundation.

Based on our technology stack and Bitcoin we devised a full end-to-end proof-of-principle of a DAO which is capable of 1) controlling money 2) democratic decision making and 3) continuous sustained self-evolution. Figure 6 shows two screens of our experimental work. The "Proposals" screen shows the details of four new users who all issued a *join request* to a new DAO. The existing owner(s) of the DAO are

Figure 7. DAO onboarding process using Bitcoin multisignature wallet

able to vote on these requests. Each joining member pays the fixed entrance fees and thereby grows the DAO fund.

The "Alice" chat history screen within our social network shows the simplicity of money transfers. Tokenised e-Euro and Bitcoin transactions are supported. We have successfully integrated our EuroToken into the existing permissioned financial system using IBAN accounts. We are in discussion with our financial regulator and central bank for conducting joint closed trails. For legal liability reasons the automated gateway for real-time token exchange with legacy bank accounts is still firewalled from our DAO.

A quorum-based voting mechanism is used as the underpinning mechanism for joining a DAO, spending funds, and expanding the DAO with new functionality. Each DAO member may submit a proposal for voting. This may be a proposal to transfer funds, investment of DAO funds (not implemented), or approve software for an additional DAO capability. Our proof-of-principle supports the whole workflow of self-evolution. In detail these are: creating a DAO, joining a DAO, initiating a democratic decision making process, broadcasting the request to create new DAO software functionality, any developer globally may seed new software using Bittorrent, voting process to determine if the new feature meets expectations, payment of developer from DAO funds, dissemination of software in .apk form using Bittorrent, each user may upgrade their software, and finally, code injection within the running Android P2P-App with a bypass of the Android security model. Detailed documentation and source code can be found on Github [1].

Figure 7 shows the onboarding process of new DAO member Bob. Our work builds upon the standard Bitcoin m-n multi-signature script. The first step is that joining DAO member Bob creates an empty multi-signature wallet. Second, Bob gives partial control of this wallet to existing DAO members Alice and Dave. Third, Bob creates a single partial signed transaction which includes transfer all funds from the old DAO into the new wallet and Bob his payment of the entrance fee. Fourth, Alice and Dave verify they have received their stake of the new DAO. Finally, Alice and Dave vote on the joining of Bob in the DAO by either withholding their signature or signing the proposed DAO transfer transaction. Our experiments uncovered an unsolved security flaw for any DAO design with such an open membership structure.

Dilution of voting power is possible with a non-unanimous voting quorum and sufficient funds to create *DAO Sybils*. This shows the need for further research into, for instance, strong digital identities [21].

Our findings demonstrate that – in principle – it is possible to democratically control money without any physical presence or legal entity and we show that possibly significant complexity may emerge through self-evolution.

8 Conclusions

The first "decentralisation winter" set in several years after Napster popularised the peer-to-peer paradigm. Today we witness an unsustainable level of "decentralisation hype", inspired by Bitcoin. Developing decentralised infrastructure remains a very difficult engineering challenge which easily requires a decade of effort, while specialised engineering talent is scarce. Many decentralisation projects lack long-term sustainability due to a lack of practical guidance and theoretical grounding. The engineering of self-organising systems, compilers, encoders, and kernels is beyond the resources of most. The second "decentralisation winter" will soon arrive if the community remains fragmented and resources are spread thin.

Our work aims to provide proof that Big Tech alternatives are possible. Our empirical Trustchain research suggests that a generic critical infrastructure to build peer-to-peer trust is possible. A total of 173 software developers contributed to our work in the past 15 years. We demonstrate that our zero-server infrastructure can host potentially billions of smartphones, without any corporate entity, middleman, or even central server to bootstrap.

We consolidate various scientific fields within a single decentralised infrastructure: secure communication, money, trust, democratic control, and leaderless organisation. A "loss-free revenue channel" is now available for artists using our MusicDAO P2P-App. We welcome others to build and expand upon our work.

With sufficient efforts and resources we believe such self-organising international collectives could grow in complexity, sophistication, and scope beyond what is possible for traditional top-down megacorporations.

References

- [1] 2020. TrustChain Super App. <https://github.com/Tribler/trustchain-superapp/>. (2020).
- [2] K. S. Barber and J. Kim. 2001. Belief Revision Process Based on Trust: Agents Evaluating Reputation of Information Sources. In *Proceedings of the workshop on Deception, Fraud, and Trust in Agent Societies held during the Autonomous Agents Conference*. Springer-Verlag, London, UK, 73–82.
- [3] Kornél Csernai and Márk Jelasity. 2012. Distributed Machine Learning Using the Tribler Platform. http://www.kl.csko.hu/projektek/msc_thesis.pdf. (2012).
- [4] Z. Despotovic. 2005. *Building trust-aware P2P systems*. Ph.D. Dissertation. EPFL, Swiss Federal Institute of Technology Lausanne, Switzerland. <http://library.epfl.ch/theses/?nr=3313>
- [5] Andreas Diekmann. 2004. The power of reciprocity: Fairness, reciprocity, and stakes in variants of the dictator game. *Journal of conflict resolution* 48, 4 (2004), 487–505.
- [6] Gertjan Halkes and Johan A Pouwelse. 2011. UDP NAT and Firewall Puncturing in the Wild. In *Networking 2011, 10th International Conferences on Networking (IFIP'11)*. <http://pds.twi.tudelft.nl/reports/2010/PDS-2010-007.pdf>
- [7] G. Harding. 1968. The tragedy of the commons. *Science* 162, 3859 (1968), 1243–1248.
- [8] K. Hoffman, D. Zage, and C. Nita-Rotaru. 2009. A survey of attack and defense techniques for reputation systems. *ACM Computing Surveys (CSUR)* 42, 1 (2009), 1–31.
- [9] Synopsys Inc. 2020. Github statistics by Black Duck Open Hub. <https://www.openhub.net/p/tribler/>. (2020).
- [10] Rob Jansen. 2014. <https://blog.torproject.org/blog/tor-incentives-research-roundup-goldstar-par-braids-lira-tears-and-torcoin>. (2014).
- [11] Audun Jøsang, Roslan Ismail, and Colin Boyd. 2007. A survey of trust and reputation systems for online service provision. *Decision support systems* 43, 2 (2007), 618–644.
- [12] I. Livadariu, K. Benson, A. Elmokashfi, A. Dhamdhere, and A. Dainotti. 2018. Inferring Carrier-Grade NAT Deployment in the Wild. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*. 2249–2257.
- [13] T. G. Papaioannou and G. D. Stamoulis. 2008. Achieving Honest Ratings with Reputation-based Fines in Electronic Markets. In *Proceedings of IEEE Infocom 2008*.
- [14] J.A. Pouwelse. 2018. IETF - TrustChain protocol. <https://www.ietf.org/archive/id/draft-pouwelse-trustchain-01.txt>. (2018).
- [15] J.A. Pouwelse, P. Garbacki, J. Wang, A. Bakker, J. Yang, A. Iosup, D.H.J. Epema, M. Reinders, M.R. Van Steen, and H.J. Sips. 2008. TRIBLER: a social-based peer-to-peer system. *Concurrency and Computation: Practice and Experience* 20, 2 (2008), 127–138.
- [16] Mohsen Rezvani and Mojtaba Rezvani. 2020. A Randomized Reputation System in the Presence of Unfair Ratings. *ACM Trans. Manage. Inf. Syst.* 11, 1, Article 2 (March 2020), 16 pages. <https://doi.org/10.1145/3384472>
- [17] Tatsuya Sasaki, Isamu Okada, and Yutaka Nakai. 2017. The evolution of conditional moral assessment in indirect reciprocity. *Scientific Reports* 7 (2017), 41870.
- [18] Michael Schillo, Petra Funk, and Michael Rovatsos. 2000. Using Trust for Detecting Deceitful Agents in Artificial Societies. *Applied Artificial Intelligence* 14, 8 (2000), 825–848.
- [19] Matouš Skála. 2020. Technology Stack for Decentralized Mobile Services. <https://repository.tudelft.nl/islandora/object/uuid%3AAbd3a5fbd-430b-4af6-bc33-eab436f4f7db>. (2020).
- [20] Alexander Stannat. 2020. On the Sybil-Proofness of Accounting Mechanisms in P2P Networks. <https://repository.tudelft.nl/islandora/object/uuid%3A6b4011c6-1668-4a1c-a49e-f86d226063b1>. (2020).
- [21] Quinten Stokkink, Dick Epema, and Johan Pouwelse. 2020. A Truly Self-Sovereign Identity System. (2020). arXiv:2007.00415