Quasi-Social
Software as the 'Social' in Socio-Technical Design

Mariusz Nowostawski[0000−0002−2809−8615]
Christopher Frantz[0000−0002−6105−8738]

Norwegian University of Science and Technology, Norway
{mariusz.nowostawski,christopher.frantz}@ntnu.no

Abstract. In traditional socio-technical system design, we typically discuss three core layers: the *social*, the *technical*, and the *socio-technical* layer. The social layer represents human aspects, the technology represents the advancements in software and technology development, and the socio-technical layer captures the interplay between the social systems and the technology-enabled or technology-mediated interactions. The socio-technical research programme responds to this pattern by integrating these layers and focusing on the interplay between social and technical. However, modern peer-to-peer technology, cryptography and encryption protocols together with decentralised technology enable mixing and the interplay of various layers, and the emergence of other, novel intermediate layers that are not well-captured, or at least not systematically identified, by the traditional methods of socio-technical design. There is a shift of software, with code and automation penetrating places that have, historically, been moderated by social aspects. This shift reflects a continuous pattern that reflects the ability of software to take over increasingly sophisticated tasks previous occupied by human actors. The resulting secondary order of complexity is insufficiently captured by traditional methods and blueprints for socio-technical systems. As a result, modern technology-mediated social systems form complex dependencies, in which the social gradually shifts towards technical, and software is replacing other components. When software starts crowding out traditionally human-mediated (social) tasks, what can be used to model and design these next generation systems? The purpose of this paper is to highlight the need for new models and metaphors that could be used to help in design of complex contemporary socio-technical systems.

Keywords: quasi-social · decentralised · blockchain · governance · sovereign · anonymity · autonomy

1 Introduction

A socio-technical system can be thought of as a social system operating on a technical base. It is a system that is both: social and technical. Traditional socio-technical system design [2] focuses on the social aspects, and how the social constructs can be designed, or built-in into the technology fabric. The traditional 4-layer model, explained by Whitworth and Ahmad [15], depicts mechanical/technological aspects as the base, followed

by the information layer, on top of which an HCI/Personal layer is constructed, to culminate in the social layer on top. The approach claims: "Whether electronically or physically mediated, a social system is always people interacting with people. Electronic communication may be virtual, but the people involved are real." [15], section 24.1.7. This might have been true for a broad range of systems, however, there is a growing category of systems for which this is no longer the case. The interaction patterns are no longer as clear cut: people are interacting directly with people, other autonomous systems and institutions. With institutions we refer to any organisation enacting centralised regulatory functions. Those can be also done in software [6]. Institutions are interacting with other institutions, too. With the advent of AI and autonomous systems is gets increasingly complex.

The relevant components of modern socio-technical design are: humans that are the actual social (human) element of the system; the technology, that provides the means of interactions and the infrastructural support; and finally the software artefacts (code [7]) that act autonomously on behalf of users, individually, or on behalf of a group. With the growth of AI and autonomous systems, the software artefact can also, potentially, act on behalf of itself. We refer to those layers respectively as socio-, technical-, and *quasi-social* layers. Those layers interact with one another, influence one another, and co-evolve. The modern information systems that provide means for value transfer, social organisation and engagement utilise decentralised democratic governance mechanisms, run without a single point of control, provide consistency under various distributed trust models and act autonomously, not requiring human control or external interventions, beyond the initial design of the system. For this reason, such systems are sometimes referred to as *allegal* as they operate in the zone of unregulated social systems and those traditionally regulated by legal systems. The systems that mediate human interactions, organisation and value transfer allow human to interact and engage in contractual agreements anonymously or pseudo-anonymously, both, with other humans as well as with software artefacts and computational systems acting on behalf of other humans, other institutions, or on behalf of itself. The existing socio-technical design methodologies are not well-equipped to model and represent the new challenges related to this technological progression. We lack proper terminology, abstractions, and foremost, methodologies to research and analyse those new social systems that are intertwined within this quasi-social layer. Some of the early experiments with the autonomous anonymous systems have highlighted the importance of the systematic approach to addressing governance challenges, such as decision-making, information signalling, and consensus protocols as well as the actual evolution and maintenance of the system itself. Those experiments have also highlighted the need to conduct experiments in the real world, since the behaviour patterns observable in the real world are a response to the interaction *with* the system – driving the need for action research [11]. In addition, there are a number of other challenges such as the contextualisation of legal jurisdictions, the concept of sovereignty, identity and enforcement, both, within the system itself as well as in the wider, social or legal contexts. In this work, we focus on the new developments in the anonymous peer-to-peer space and discuss their influence on the designs and modelling of socio-technical and quasi-social systems. For illustrative purposes, we will apply those ideas to advanced decentralised systems in the financial sector.

2 Trust

Traditional financial, insurance and governance institutions work through a trusted, centralised, controlled and monitored organisations established within particular national boundaries and bound by the national and international rule of law. These centrally-managed institutions consist of increasingly complex, monolithic computational systems that are difficult to adapt to new, growing and changing requirements. In fact, the implicit lack of transparency regarding operational details relies on sustained trust into their functioning – trust in humans that constitute the fundamental elements of those institutions. This gets increasingly complex in the context of multi-national jurisdictions, often conflicting legal requirements, cross-border taxation, political instability, privacy-preserving and data protection directives such as the new European Union GDPR provisions. Similar problems are also present in the context of large-scale cross-national social systems in which software algorithms that are primarily used to determine and enforce operations and constraints on the actual workings of the system itself. The complexity of the social systems mediated by technology becomes impossible to be handled by traditional means and the detailed workings of those complex systems become extremely difficult, or impossible to trace. This is particularly true for situations in which decisions within the system are mediated through machine learning black-boxes.

Good examples of that complexity are modern financial systems. On the one hand, such systems are regulated and monitored by the complex legal frameworks that, in large, is agreed upon internationally. The regulations safeguard, to some extent, the actual workings of the system, but they also limit the innovation and customer choice. On the other hand, the complexity of financial instruments offered by the incumbents makes it actually impossible for full control and predictability of the resulting emerging properties of the system itself. The history of financial crashes provides us with a plethora of examples. The next financial crises cannot be predicted or prevented, partially due to the human elements that influence it, but partially, due to inherent self-dependencies, built-in feedback loops and internal mechanisms that give raise to untraceable emergent properties. Therefore, traditional systems lack transparency and exhibit strong resistance to analysis and predictability. The regulations have limited control over the potential misuse/abuse. Nevertheless, the systems require high levels of data disclosure from the participating individuals, an aspect that can be in conflict with personal freedoms and private data protection directives.

Due to the problems above and with the help of technological innovations a new class of systems is emerging, which exhibit qualitatively novel properties. The ideas behind distributed consensus and ledger technology have gained significant traction since the launch of the first public experiment known as Bitcoin in 2008. Multiple examples of decentralised technology have been deployed and successfully used for value transfer, digital assets, document validation, and crypto-currencies. Despite these experiments computer and social scientists are in the early stages of understanding the impact this technology has, and is bound to have, on social institutions and organisations more generally. Emerging distributed institutions that utilise the blockchain technology are able to facilitate international interactions, contracts, and value transfer, all of which can be achieved *without* the need for the human-based third-party trust, central authority, or externally managed audits. Moreover, those interactions can be au-

tomated and conducted autonomously and anonymously by the distributed peer-to-peer networks. In other words, on the highest level, layer 7, or application layer, it is not communities interacting with communities. It is software interacting with software, on behalf of users that provide only implicit regulatory mechanisms and policies within the self-governing computational peer-to-peer system. Those innovations have a significant impact on the future structure of our social and economical environment. We argue, that due to the properties and the design, anonymous, privacy-conscious decentralised technologies represent a qualitatively new and disruptive change in the construction and structure of traditional socio-technical systems. It also provides novel ways to management, governance, conflict resolution and the societal organisation in general, which are directly relevant to the class of systems traditionally investigated by socio-technical research methodologies. To argue our point, we will explore both the technological and organisational foundations of blockchain technology. We will do so by first highlighting historical examples for information and communication systems in Section 3, before discussing the corresponding characteristics in blockchain technology in Section 4. This is followed by a discussion of specific applications of blockchain technology that are dissected along specific layers in Section 5.1 to provide a basis for a comparative analysis of different blockchain-based applications in Section 5.2. We conclude the discussion on the refined perspective in Section 8.

3 Peer-to-peer and decentralised systems

The Internet is one of the largest multi-national projects that we, as humanity, have been engaged with. It has enabled new and innovative ways for the social organisation as well as communication. It has provided an unmatched repository of knowledge, training, and various information-centric or knowledge-centric business models. The Internet itself represents "technology" in traditional socio-technical design, yet, one feels there is something missing. A simple Word Processor can be thought of as a technology, as a tool, that facilitates a class of interactions that are build on top of it. Yet, systems such as Facebook, which are part of the Internet fabric are in themselves socio-technical systems, in fact require a recursive application of socio-technical design principles as part of the system analysis, let alone such systems' further constituents. This exemplifies a break of the idealised hierarchical design of 4-layered socio-technical systems as argued by Whitworth and Ahmad [15]. We need something new, to capture both, the recursive nature of the inter-dependencies of the system itself, as well, as the fundamental inter-dependencies of the social, technological, and software artefacts that are the fundamental building blocks of the complex systems we deal with today.

The fundamental design of the Internet protocols is inherently distributed and peer-to-peer, with centralisation occurring only in places where it is aligned with the geopolitical organisation of the actual physical world. For example, Domain Name Services are hierarchically organised with country root domains managed by the central services of a given country. Nevertheless, ownership and rights are distributed among multiple entities, even on the top-level of this hierarchical structure. Similarly, all core protocols such as SMTP (for mail delivery), or HTTP (for web content delivery) are inherently peer-to-peer-like, due to the nature of the underlying TCP/IP layer. This peer-to-peer

environment and design gave rise to a large number of innovation and a plethora of new services, that were not possible before. The standardised way of communicating enabled innovation and exploration of possible interaction patterns and organisations that human groups could form freely. The early days of the Internet could also be referred to as *alegall*, as the groups organised spontaneously and followed their own codes of conduct, often across borders (e.g., topical forums, social networks). This dynamic process continues today too, even though many of the core services are dominated by few large stakeholders, such as Amazon, Apple, Facebook, Google and Microsoft.

Some of the most complex computer systems are designed in a peer-to-peer, evolutionary fashion, through something that is referred to as the *bazaar* model [13]. Linux or FreeBSD kernels, for example, are the most advanced and widespread operating systems kernels. They have not occurred within the traditional, top-down management but rather in an open, dynamic, and constantly changing and adapting *workplace*, that is increasingly often virtual. Peer-to-peer and decentralised systems are essential to facilitate unconstrained innovation and exploration on a large scale. This is one of the main reasons why decentralised technology is re-shaping the communication paradigm. Building on this idea and encroaching functions of the social-coordinative realm, permissionless open-source blockchain developments and blockchain-based systems change the way liability, trust and ownership are handled, aspects of which we discuss in the following.

4 Absence of central authority

Decentralised technology relies on an agreement of a system state achieved in a situation without central authority and with potentially hostile and fraudulent actors. The technology, through an interesting interplay of incentive system, automation and distribution of power, allows achievement of an agreement (consensus) on the system state, and system state record. In essence, a public blockchain technology is solving the consistency problem, that is, ensuring a consistent indisputable representation of state and transitions outside of the control of either single stakeholder. The consistency of the events log is assured by aligning the incentive model with the goals of the distributed network of peers. In this context 'public' implies that blockchain applications operate in the open public sphere and coordinate interaction between unknown participants in a permissionless fashion, i.e. anyone can participate.

In the absence of a central sanctioning authority, blockchain modifications (i.e., transactions) need to be cheap enough not to discourage the system's use, yet expensive enough to prevent opportunistic abuse (e.g., by submitting fraudulent transactions). Mechanisms that facilitate this trade-off include the consumption of high amounts of processing power or per-transaction payments or mechanisms involving stake that can be lost when abuse is detected. This balance of incentive and deterrence is the hard socio-technical challenge. The best mechanisms to-date rely on so-called *proof of work* [5]. In that model, a cryptographic riddle is posed and requires a provable amount of time spent on computation, to be resolved. Validating the riddle result is easy, solving the riddle can be made arbitrary hard.

An alternative approach that avoids the inefficiencies associated with the proof of work, such as wasted power and processing time, as well as to limit the computational 'arms race' for computing power, is the *proof of stake*. In the proof of stake [1] the individual participants' influence is constrained by their commitment to the system, such weighing the influence of the number of resources individual participants hold. Naturally, this introduces hierarchical characteristics into the system, but increases the efficiency of the system without unproductive use of computing resources. Whatever the specific protocol employed by a given blockchain implementation, the proof of work, proof of stake, and the voting model used for validation work in unison; the stable long-term strategy is not to cheat. Decentralised blockchain technology offers third-party trust without any single entity taking full responsibility or having full authority.

To discuss novel artefacts of blockchain technology, we need to first understand some of the architectural underpinnings – which we systematically highlight in the following section.

5 Architectural Layers of Blockchain Systems

5.1 Layer 1: Base layer

The decentralised technology can be seen as consisting of two fundamental layers. Layer 1 facilitates the consensus and transactions sub-system. They represent the core functionality. Layer 2, ie. any protocols build on top of Layer 2, provide additional facilities and can provide application layer logic.

One of the core Layer 1 technology today is Bitcoin. The creator of the system, known as Satoshi Nakamoto, wrote about the system in a founding white paper [10]. The global network of *miners* and users is one of the largest and most powerful computational resources currently in operation. Bitcoin is a good example that highlights the main components of a blockchain-based ecosystem. It comprises of software developers, that are either paid by the users owning the virtual currency, or own the currency themselves, the miners, and the node operators. Bitcoin operates with pseudo-anonymous identities. There is no reputation subsystem, and the transactions are safeguarded by proof-of-work mechanism. Participation is encouraged through mechanisms of incentives, including e.g. mining rewords.

To address some of the shortcomings of the original Bitcoin structure, alternative currencies have emerged. One example for this development are DashCoin, whose structural characteristics we will compare to Bitcoin, in order to disambiguate blockchain technology from specific applications built on its principles. DashCoin uses the original codebase for Bitcoin node, however, it changes some of the core fundamental mechanisms. DashCoin introduced the notion of masternodes, that provide a different, hierarchical structure to the consensus mechanism. It also provide built-in mechanism for so called Private-Send, which is equivalent to the CoinJoin protocol in the Bitcoin blockchain [8]. Those mechanisms obfuscate the source and destination of transactions to form long complex chains of ownership that is difficult to analyse. The mechanism also provides provable deniability.

5.2 Layer 2: Application Layer

Fully anonymous, atomic, and reliable peer-to-peer transfer of value is one of the most common examples of the blockchain technology application. That can be built-in into Layer 1, and all of the existing blockchain systems have a native built-in currency. Bitcoin, Ether, Dash, and many others crypto-currencies operate directly on Layer 1. However, new assets or digital currencies can be developed and provided on Layer 2. In fact, Ethereum provides a formal specification through ERC-20, and as of April 2019, over 180,000 ERC-20-compliant tokens are found on the Ethereum network. Most of the stable coins (crypto-currency with the value pegged by one of the existing FIAT currencies, such as EURO or US dollars) are also using Ethereum as Layer 1. As another example of a blockchain-enabled application, consider a simple escrow service. Typically, an escrow service is used to ensure atomicity of a transaction between two non-trusted entities, and to have the ability to roll back a partially fulfilled transaction. An escrow service, a trusted third party is used to work as a trusted intermediary to facilitate the transaction. With the blockchain, such transactions are atomic by design, without the need for a trusted third party. Escrow services, or decentralised exchanges are now possible through the mechanisms that allow to conduct Atomic Swaps. Those are available on many blockchains and sit somewhere between Layer 1 and Layer 2. Atomic swap allows to automate the process of creating escrow services based on simple smart contracts.

What those examples demonstrate is that many centrally-managed services, in particular those provided by insurance companies, banks, or governments, can be made more secure and more transparent with the use of blockchain technology. This means that the human element can be eliminated from selected institutions or contractual agreements, especially in areas in which the ability to maintain accountability is challenging. This has a fundamental impact on how we will perceive and deal with fraud, data leaks or power abuse. This potential and the associated challenges become clearer when exploring examples of blockchain technology with respect to structural and governance characteristics. We thus use the following subsections to highlight some examples of blockchain technologies, so-called blockchain applications, in an attempt to illustrate the sketched potential.

6 Blockchain Applications

6.1 Lightning Network

The Bitcoin Lightning network is a mechanism to overcome the inherent scalability limitations, in terms of transaction throughput, of traditional Bitcoin blockchain systems. Due to the required consensus mechanisms and decentralised nature of open blockchain systems, new entries in the ledger (transactions) can only be recorded with a predefined step-like mechanism, that limits the per-second transaction throughput. One of the solutions to this inherent limitation is a technique based on off-chain transactions. It is based on the payment channel concept, and it relies on the idea that not all transactions need to be recorded in the ledger. Only those that establish payment channels between participants and those that resolve, or close the payment channels (channel closure).

The actual payments within the channel can be done without leaving any trace in the blockchain proper, and instead happen off-chain (hence the name, off-chain transactions). For this to work, however, additional guarantees need to hold for the participants to limit the possibility of abuse or misuse of such a system. The Lightning Network Protocol addresses this by the existence of additional services, so-called *Watchtowers* [4]. Watchtowers will provide additional guarantees and triggers necessary for the system to function properly. There is, however, limited research of how they would operate, how they would be conceptualised and implemented, and how information they reveal will be made available to the general public as well as to the system participants – and last but not least, how we can devise mechanisms to react to irregularities identified by watchtowers. This collection of open issues offer a fertile ground for research, and our study of quasi-social systems as discussed in this paper targets precisely this unexplored niche.

However, this approach is not limited to this precise example, but also existing applications with tangible economic impact, such as prediction markets.

6.2 Prediction markets

Let us consider a system in which traders can trade virtual asset that represents the outcome of the future event. Consider it as information or decision markets, idea futures, or event derivatives. A market is created for the purpose of trading the outcome of an event. The outcome is binary, therefore, the virtual option will expire at the price of 0% or 100%. A prediction market contract trades between 0 and 100%. Prediction markets can be treated as crowdsourcing information, with the main purposes of eliciting aggregating beliefs over an unknown future outcome. Traders with different beliefs will trade on contracts whose payoffs are related to the unknown future outcome. Then, the market price of the contract is considered as the aggregated belief. Markets like that can be used for risk assessment or risk hedging in order to establish a likelihood of a future event, market value, future market value, and so on. It is stipulated that three necessary conditions need to hold for such markets to function well: *diversity of information*, *independence of decisions*, and *decentralisation of organisation* [14]. It is exactly this third property that makes blockchain-based, smart-contracts driven prediction markets an appealing value proposition. There is one ongoing experiment started in 2014, Augur [12], that provides an open prediction market. However, one of the most severely limiting factors is that those markets can be subject to manipulation [3]. Similar to the challenges outlined for the lightning network and its trust-based operation on sidechains, such markets suffer the same problem of intermediaries detecting and addressing patterns of illegitimate behaviour – without being confined to monitoring merely technical aspects of the system, but clearly operating at the intersection of social and technical activity, thus assuming the proposed quasi-social role. Coming back to the specific example, this real-time monitoring of the events and manipulation detection is something that the aforementioned Watchtowers could provide by facilitating the necessary feedback loop and ensuring the system's self-regulating properties. To the best of our knowledge, there is no active research on Watchtowers for this particular type of systems.

6.3 Autonomous Loan Dispenser Systems

Let us consider an autonomous automated secured loan system. In this scenario, the borrower needs to borrow money, that she promises to pay back at a certain time in the future. The security of the loan is based on the collateral, that is a property or assets, conveniently represented electronically, that a lender accepts as security for a loan. The borrower must first obtain an appraisal with the estimated fair market collateral value of the property to be considered in the loan - this process itself requires expert opinion and delegating it to an open market, such as the above example of the prediction market, provides certain benefits. While the collateral value is an important component of the loan, the loan system should also use the information on a borrowers credit profile and credit history. In a traditional system, the value of the collateral, collection of the credit information and assessment are left to the institution offering a loan. Typically, it is a bank. The market value of the collateral needs to be checked on a regular basis against the market value and the loan value, such that the loan never exceeds a certain threshold (50-90%) of the current market value of the collateral. In our case, we could consider a blockchain-based system based on Maker tokens (MakerDAO, http://makerdao.com/en/). It is a smart-contract-driven system based on Ethereum network, that provides loans based on electronic collateral. The value of the loan has been established to be at a maximum of 50% of the collateral value.

If the value of the collateral drops below a certain safety margin, the lender would not be secured anymore against the failure of the loan repayment. In such a case, a safety trigger needs to force the borrower to repay the loan, or, the collateral would need to be re-sold on an open market as to provide the necessary security for the lender. In this simple example, the lender actually is a set of autonomous smart contracts, that execute the predefined logic of securing (freezing) the collateral (e.g. cryptocurrency or some forms of digital assets, like house ownership) such that the loan system does not need a human institution to facilitate the financial service. It is self-regulating and autonomous. However, for its operation there needs to be an autonomous, real-time monitoring entity that checks which of the borrowers agreements trigger the safety event. This form of monitoring ensures and replace the role of financial auditors that verify and validate that the system operates as prescribed. In the system, the same as in the previous two examples, this is left as non-colluding 3rd party service, such as the quasi-social Watchtower concept highlighted before, that provides meta-layer on top of the base-layer protocol.

7 Discussion & Proposal for Integration

7.1 Discussion

As we can see from the provided examples, all monitoring functionality necessary for developing open self-organising real-world applications – as currently advocated in the form of blockchain technology – rely on technology that is, of course, technology, but can no longer be clearly associated with a technical perspective, but rather sits in a layer that operates in concert with the social affordances a system provides. However, those can often only be observed at runtime, and worse, can change over time – preventing

these intermediary entities from being considered at design time only, but rather considering their change and evolution over time.

To make a simple real-world analogy, imagine that a certain institution is formed, with the rules of how it operates, e.g., a bank. Then, a set of meta-layer institutions (such as auditors, regulators and law-enforcement) need to be built on top of the base layer to facilitate the proper operation of the institution itself. This is exactly what Watchtowers are for. However, until today it is unclear how such systems of quasi-social nature are embedded into (or within) socio-technical systems. The nuanced characteristics are not explicitly reflected in the traditional socio-technical design perspective.

7.2 Proposed Integration in the Socio-technical Design

To explore pathways towards integrating these perspectives, let us provide an overview of the aspects discussed in the context of explicit examples. Figure 1 depicts the traditional stratification into the social layer at the top, capturing all the characteristics of human participants. At the bottom, we can see the technological layer that involves all the structural aspects discussed in Section 5.1.

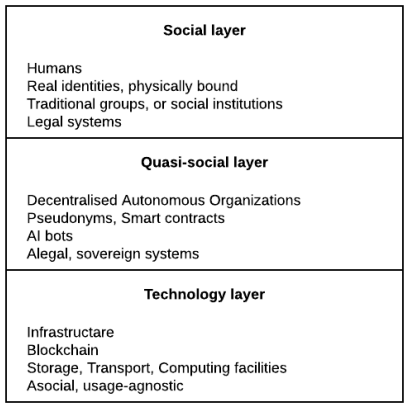

Fig. 1. Three-layer Quasi-Social model

The novel characteristics with coordinative function on both the social and technical layer, we propose, needs to be conceptualised on a novel *quasi-social* layer that can both involve the observation and moderation of social interaction by autonomous and human entities, but primarily adopting functions that in traditional socio-technical systems would clearly be associated with the social side of the system. This involves aspects such as Decentralised Autonomous Organisations based on smart contracts, AI bots, Watchtowers, and more generally, the regulative functions that guarantee, for example, compliance with legal regulation and associated liability – activities traditionally associated with human actors.

Existing systems for audit and monitoring are focused on transaction tracking and transaction monitoring. There are existing commercial systems such as Chainalysis, Elliptic, CipherTrace, as well as a large number of block explorers that provide real-time or close to real-time transaction visualisation of records from a given blockchain ledger. There has been also a number of research projects related to transaction analysis [9], and others. Those focus on transactions, that is, individual recorded ledger entries in the given blockchain. This is an important and fundamental source of information about the underlying system, however, this is not sufficient. ontologies to deal with the complex interactions in the Quasi-Social layer of the complex software systems. Given that the field is mostly practitioner-driven, it faces unprecedented challenges in terms of social security and legal compliance. Similarly, the challenges faced by law enforcement agencies, courts and legal experts rely on auditability, forensic readiness, ability to obtain evidence and transaction and value transfer tracking. Those technologies and requirements are new and rely on expertise and tools that effectively are being on the cutting edge of research and development. Even though the core crypto technologies have been known for over 25 years and the Bitcoin has been operating for over 10 years, the socio-technical implications of the various decentralised autonomous system deployments are not yet well understood. Given this it is even more important to structure and analyse such systems in a way that faciliates the differentiation of the technological, social and coordinative quasi-social functionality. The Watchtowers are an example of this. They provide the necessary feedback loop and mechanisms to validate, verify and audit if the underlying systems actually behave within the regimes that they have been designed for.

Sticking with the illustrative examples of *Watchtowers*, there are significant scientific implications of the higher-order feedback mechanisms for decentralised systems that operate across all three levels – with respect to technological impact, quasi-social and societal impact.

Technological. Watchtowers will have a significant impact on how the decentralised, autonomous and smart-contract driven systems are designed, and how they operate. The higher-order indicators that will be measured by Watchtowers can be re-integrated into the fundamental lower-layers such as to offer enhanced self-regulating properties. The necessary feedback will make the systems more robust, resilient, more autonomous and self-regulating.

Quasi-social. The Watchtowers will make the decentralised systems more transparent and robust. Harder to misuse. The Watchtowers will enable better, richer and more complex services that can be designed in the future with the use of the automated feedback mechanisms. This is similar to how Lightning requires and builds on top of the required Lightning watchtowers.

Societal. Regulators, auditors, traditional law enforcement agencies as well as the general population can use the results of this project, indirectly, through the feedback mechanisms that will be established between the underlying technological layer and the indicators and metrics that are provided by the Watchtower services. This will provide better and more transparent information spread, ease the operation of the system as well as ensures that the systems are not misused or abused.

8 Summary

We have argued that traditional characterisations of socio-technical design are not sufficient to capture the full spectrum of complexity arising in modern systems. We have argued, that a three-level modelling approach, based on social-, technical- and quasi-social conceptualisations will lead to a way forward in understanding and subsequently analysing the new category of emerging systems that involve anonymous, peer-to-peer networks and complex social-machine interactions that span individual, collective and institutional layers. We discussed this using the concept of watchtowers as an example and explored their potential impact across all three layers of the socio-technical design, as well as all layers from individual to institutional. For both the socio-technical research community, it will be important to consider the aspects thrown into this debate by the emergence of such novel socio-technical system components, since, one way or another (i.e., explicitly or implicitly), decentralised coordination technology, such as evidenced in blockchain technology will be pervasive in any future information systems. This leaves us with two options: observe the space and risk being sidelined by the complexity such systems pose, or to actively engage in such discussion in order to influence how such systems will look like in the future.

References

1. Proof of stake instead of proof of work. https://bitcointalk.org/index.php?topic=27787.0, accessed on: 1st May 2016
2. Baxter, G., Sommerville, I.: Socio-technical systems: From design methods to systems engineering. Interacting with computers **23**(1), 4–17 (2011)
3. Chakraborty, M., Das, S.: Trading on a rigged game: Outcome manipulation in prediction markets. In: IJCAI. pp. 158–164 (2016)
4. Decker, C., Russell, R., Osuntokun, O.: eltoo: A simple layer2 protocol for bitcoin. White paper: https://blockstream. com/eltoo. pdf (2018)
5. Dwork, C., Naor, M.: Pricing via processing or combatting junk mail. In: Proceedings of the 12th Annual International Cryptology Conference on Advances in Cryptology. pp. 139–147. CRYPTO '92, Springer-Verlag, London, UK, UK (1993), http://dl.acm.org/citation.cfm?id=646757.705669
6. Frantz, C.K., Nowostawski, M.: From institutions to code: Towards automated generation of smart contracts. In: Foundations and Applications of Self* Systems, IEEE International Workshops on. pp. 210–215. IEEE (2016)
7. Lessig, L.: Code is law. The Industry Standard **18** (1999)
8. Maxwell, G.: Coinjoin: Bitcoin privacy for the real world (2013)
9. Moser, M.: Anonymity of bitcoin transactions (2013)
10. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system (2008)
11. Pasmore, W.: Action research in the workplace: The socio-technical perspective. Handbook of action research **2**, 38–48 (2006)
12. Peterson, J., Krug, J.: Augur: a decentralized, open-source platform for prediction markets. arXiv preprint arXiv:1501.01042 (2015)
13. Raymond, E.: The cathedral and the bazaar. Knowledge, Technology & Policy **12**(3), 23–49 (1999)
14. Surowiecki, J.: The wisdom of crowds. Anchor (2005)
15. Whitworth, B., Ahmad, A.: Socio-technical system design. In: Stephanidis, C. (ed.) The encyclopedia of human-computer interaction, chap. 24. Interaction Design Foundation (2012)