

Blockchain:

The Emergence of Distributed Autonomous Institutions

Mariusz Nowostawski

Norwegian University of Science and Technology
Gjøvik, Norway
Email: mariusz.nowostawski@ntnu.no

Christopher K. Frantz

Otago Polytechnic
Dunedin, New Zealand
Email: cf@christopherfrantz.org

Abstract—We present a novel institutional perspective on the distributed consensus and ledger technology known as *blockchain*. We discuss the concept of *Distributed Autonomous Institutions* that are able to facilitate global interactions, contracts, and value transfers, all of which are achieved without the need for the human-based third party trust. We argue that due to its properties and design blockchain technology represents a disruptive change in the modelling paradigms of socio-technical systems. Distributed trust and consensus mechanisms offered by blockchain technology represent a novel, qualitatively different, phenomenon. We present the general design principles, stakeholders, the dynamics between those stakeholders, the incentive models, and the consensus protocols currently used in blockchains, before highlighting the potential of blockchain technology to develop distributed autonomous institutions. We conclude with a discussion of challenges associated with the adoption of blockchain technology.

Keywords—*blockchain; governance; autonomy; distributed autonomous institutions; public ledger; mining; consensus; smart contracts; Bitcoin; DashCoin; Ethereum*

I. INTRODUCTION

One of the important enablers for modern civilization has been the invention of language. Spoken language has enabled the evolution of complex, yet stable communication patterns [1], whereas written language has provided persistence [2] and added the ability to communicate asynchronously, sometimes across centuries or millennia. Communication and persistence have further played a fundamental role in the development of modern computing paradigms. Since the invention of the von Neumann architecture, human development and institutional automation has accelerated, as evidenced in increasingly complex forms of social and economic organisation and associated regulation. Example phenomena include the increasing numbers of digital nomads and flexible organisational boundaries based on procurement of external services. We claim that the accelerated development and growth in complexity in human institutions ultimately relies on the three fundamental elements: (1) *communication*: the ability to communicate and synchronize; (2) *persistence*: the ability to store communication or data; and (3) the *ability to compute*: ie. automatically execute an algorithm, or in other words, a finite set of computational steps. This parallels the characteristics of what we refer to as institutions [3], or “manifestations of social behaviour” [4], which are characterised by (a) *social interaction*, (b) *stability*, i.e., institutions’ ability to survive the constituting behaviour [5], and

TABLE I. INSTITUTIONAL CHARACTERISTICS AND ASSOCIATED ENABLING TECHNOLOGY

Technology	Communication	Persistence	Distributed Computation
Internet	•		
P2P Technology	•	•	
Blockchain	•	•	•

(c) *procedural prescription* of desirable behaviour (or proscriptive of undesirable behaviour) that may or may not be explicitly codified [6] (e.g., as laws vs. social norms). While the internet enabled communication across organisational and national boundaries, laying the foundation for modern virtual organisations, its primary focus was the facilitation of general human communication. The actual state was held within the endpoints, not the network itself. Only the introduction of peer-to-peer technology in the early 2000s (e.g., [7]) moved state into the network itself. Thus state did no longer rely on individual endpoints, but was rather distributed across a collection of participating network nodes. Therefore, the state could be managed in the network itself, the modification required explicit intervention by individual nodes based on externally negotiated semantics. In this context ‘externally negotiated’ implies that the higher-level application-specific semantics (beyond the primitive CRUD operations Create, Read, Update and Delete) are not managed by the system itself. Even though cloud technology reinforced the virtualisation and decentralisation of computation, it did not change the institutional status: the control is retained with a single well-defined entity, generally the owning organisation. The inability to delegate the guaranteed execution of complex instructions, along with assurance of transaction safety to the network itself, limits the adoption for critical services outside the control of organisations such as banks, insurances and governments. We argue that the final missing pillar, the *decentralised execution of procedural prescriptions* makes all the difference in building truly open institutional environments, enabling us to relay critical coordination tasks, such as digital payments, tendering of governmental contracts, or even democratic voting processes to the network itself. Table I summarises the institutional properties of the highlighted technologies.

We believe that *blockchain technology* reflects the natural evolution towards loosely coupled, user-centric, distributed and autonomous institutions, that will fundamentally change

the nature in which humans engage with computers, and, in extension, with other humans. In this context the autonomous nature of institutions reflects the continuous operation without the need for any human intervention.

In Section II we briefly introduce the principles that underlie blockchain technology and highlight the central characteristics that produce the added value that has the potential to redefine the modern economic landscape. We further introduce Bitcoins and DashCoins as example implementations of blockchain technology, before introducing the more advanced blockchain-enabled decentralised computation in Section III. In Section IV, we introduce the concept of *Distributed Autonomous Institutions*, before discussing their impact on socio-technical systems as well as society in the wider sense in Section V, along with an outlook on future work.

II. BLOCKCHAIN

Blockchain technology facilitates the fundamental shift based on automated, yet flexible mechanisms that deal with trust and liability based on adaptive incentive systems. The underlying cornerstone of 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 mathematical consistency of events, or transactions, 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., in principle anyone can participate.

Whereas the distributed nature of state is unproblematic, its synchronised modification is. In an open distributed environment all the nodes need to achieve consensus about whether an individual transaction is accepted or rejected. Accepted transactions must be subsequently integrated into the shared chain of transactions held within the blockchain. Decision-making generally operates based on social choice protocols, such as voting (e.g., majority-based voting). Thus stakeholders cannot modify the distributed shared state or cheat without collaboration by the majority of other stakeholders. The probability of colluding is reduced by network size and anonymity, as well as ensuring that cheating carries a risk of value loss. In this context *value loss* means waste of computational resources or loss of the managed resources, e.g., digital currency. However, any modification puts a computational burden on all members of the network. This aspect could be exploited by injecting large numbers of transactions and reducing the blockchain’s ability to process those, while maintaining global consistency. The associated expectation is that fraudulent transactions (e.g., declaring multiple transactions of the same funds at the same time) will be accepted by a critical number of hosts and eventually be accepted into the global blockchain. 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. This balance of incentive and deterrence is the *proof of work* [8]. 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 [9] the individual participants’ influence is constrained by their commitment to the system, such weighing the influence by the amount 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 the full responsibility or having full authority.

What this means for institutional settings [10] is that trust and liability can now be flexibly shifted on a spectrum ranging between the institution itself and the participating individuals. Let us take an example of a simple asset, e.g., a currency. Let us assume that selfish individuals only trust themselves completely, i.e., one cannot cheat or misuse one’s own trust. Being a custodian of one’s own assets carries liability, e.g., for safekeeping. To relieve oneself from the liability, one can give custody of the asset to a trusted institution (Institution A in Fig. 1), such as a bank. Once an individual deposits an asset, the bank is liable for the safety and security of that deposit. The liability has been transferred from the individual to the institution. However, that transfer also introduces the need for trust. The individual must now put their trust in the bank.




Figure 1. Trust Relationships in the Traditional Institution Concept

To relieve one from liability, and to pass the liability to the bank, one thus has to trust the bank, or, more precisely, institutions that regulate and control the bank’s operation (Institution B in Fig. 1), since the inner workings are inaccessible to trusting individuals and are thus not openly verifiable. But what if one trusts none but oneself, but still wants to pass the reliability to an institution? The solution is a *Decentralised Autonomous Organisation* (DAO) [11] – an algorithm, that codifies the participants, governed resources as well as protocols. The algorithm that is guaranteed to work according to its specification and, if well constructed, never fails. Once instantiated, it would thus never break the trust one puts into it, since the algorithm exhibits verifiable trust. Therefore, with the blockchain it is possible to achieve the liability transfer from individual to institution, without the putting trust into a traditional institution that operates based on human intervention (e.g., a notary). That said, any DAO can only be as good as its implementation. A DAO is governed by verifiable code and reliable execution, but that does not

protect it against bugs introduced at design time. A good example for the importance of thorough development is the recent exploitation of the most prominent DAO and the theft of around one third of all entrusted funds [12].

As another example for a blockchain-enabled application, consider a simple escrow service. Typically, an escrow service is used to assure 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. 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. In the following subsections we thus highlight some examples of blockchain technologies to illustrate the sketched potential.

A. Bitcoin

The first deployment of the blockchain technology and currently the most dominant virtual currency is known as Bitcoin. The creator of the system, known as Satoshi Nakamoto, wrote about the system in a founding white paper [13]. The global network of *miners* and users is one of the largest and most powerful computational resources currently in operation.

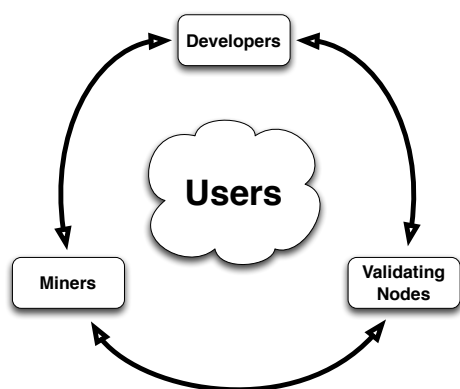


Figure 2. Blockchain Stakeholders

Stakeholders. At its essence the Bitcoin network relies on two operations, a) the *mining* of the currency (i.e., production or minting of the currency tokens), and b) the *validation* of transactions (i.e., facilitating the use of the currency). This is mediated by a set of stakeholders, a schematic overview of which is depicted in Fig. 2. The *developers* provide code for the mining and the consensus library. The *miners* generate new blocks that contain the individual transactions. The *validating nodes* run software to accept or reject transactions. In addition,

validating nodes also accept or reject blocks mined by the miners. To regulate individual influence, the acceptance relies on majority-based voting between the validating nodes. The detailed workings of the system are further explained in [13] and in [14].

Instead of exploring the technical complexity here, we only focus on the circular dependencies between all three stakeholders. Validating nodes are not able to influence the process alone, because they lack the computational power necessary to compute the blocks. Miners, who possess the necessary computational resources, are not capable of influencing the voting process directly, as the network of validating nodes is larger than the mining network. This makes it difficult to obtain 51% of voting power. The developers shape the rules and the consensus protocol, but can neither control the mining nor the network. In principle, all parties thus have a strong incentive to maintain the trust and operational integrity of the network, without the risk of any other group exerting full control, thus giving the system characteristics of a *common pool resource* [15] with distributed governance.

Centralisation. In the early days of the deployment the users of Bitcoin used to be the ones running the mining as well as contributing their computing resources as validating nodes. That was the initial assumption of Satoshi, and mining as well as validating was built into the Bitcoin wallet software. The goal was to keep the network as large and as distributed as possible. However, over time an interesting phenomenon occurred, similar to the development of internet services: centralisation. Due to increased popularity of simplified wallets, increased demands on storage and traffic to maintain fully validating nodes, and the escalation of computational resources needed for mining, most contemporary Bitcoin users are neither miners nor validating nodes. This model has severe limitations, and the community has yet to work out how to address the increase in centralisation of the system. A related phenomenon that exemplifies the complexity of influence factors on the structural characteristics of the network is the fact that majority of mining power for Bitcoin now resides in China [16]. A reason for this lies in the exceptionally cheap access to electricity and the direct access to the mining hardware that is produced in China. Thus micro-economical incentives have tangible impact on the network structure itself. The other property of Bitcoin is that it does not control who the miners or who the validating nodes are. The network can be infiltrated by malicious nodes in an attempt to destabilise the network, or simply to monitor the transactions in order to de-anonymise network participants [17], an aspect we will discuss in the following.

Anonymity and Traceability. Anonymity in Bitcoin network takes a weak form of pseudonymity. That implies that users' identities are hidden behind pseudonyms that can be tracked through the blockchain. Bitcoins are not fungible. Fungibility is the property of a good or a commodity such that its units are completely interchangeable, and can be easily substituted. The Bitcoin protocol allows traceability of transactions between the pseudonyms, and as soon as a given pseudonym is attached to a real person, there is a possibility of de-anonymising other transaction participants. In order to maintain anonymity, specially crafted mixing services need to be used to make tracking harder, or statistically impossible. Those services work in such a way that they generate a large

number of bogus transactions that obfuscate the true coin ownership in the transaction graph.

Governance. From a socio-technical perspective, the most interesting element of the Bitcoin blockchain is its governance model, or, to be precise, the lack of it. The network is fully self-organising, and there is no governance model built in. The decision making and protocol refinement happen through iterative decision-making processes and community adoption. In theory, it means progress can be achieved by the community through majority-based voting. In reality, due to lengthy iterations between the discussions, development, and partitioning of the development efforts, the progress and adoption of ideas is slow. With focus on the reliability and long-term viability of the currency, this can be a desirable property, since it is based on the democratic consensus-based decision-making. On the other hand, consensus-building involves an inherent risk of community partitioning, or even a *hard fork*. A hard fork occurs if the community and the network splits into two chains, out of which one is likely not to persist in the long term. This means assets stored in the eventually discontinued fork are ultimately lost. However, hard forks can occur intentionally: Ethereum's (see Section III) recent funds theft led to precisely that decision based on community consensus [18] in an attempt to revert the fraudulent transactions.

B. DashCoin

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.

Stakeholders. The Dash network is fundamentally similar to the Bitcoin network. However, there are some interesting modifications. The Dash protocol introduces a concept of *second layer* nodes, called *master nodes*. Those are selected nodes that provide a certain proof of stake, or collateral, such that only a limited amount of nodes ever exist in the network. Those nodes are rewarded for participating in the network and they provide certain services, such as governance and voting on new services, allocation of funds, and consensus rules. Those nodes can also provide a distributed oracle service, that is, provide a verifiable ground truth without the need for a trusted third party. Because there is a limited amount of those, and the fact that they can be verifiably trusted (due to the collateral that they deposit), certain operations, such as the confirmation of transactions, can be done much faster than in the Bitcoin model.

Centralisation. The Dash network addresses the issues of centralisation by delegating some of the duties to second-tier nodes (master nodes). The number of those is kept within the range of 3500-4000 nodes, which is sufficient to sustain a robust network. Each of those nodes has deposited 1000 DASH, which means those are core stakeholders in the network, whose incentives are aligned with the group incentives as a whole. Besides the proof of stake, the network also employs a proof of work algorithm that bears no benefits if implemented in hardware – anyone can participate in mining using their graphics processing unit (GPU), and it is not beneficial to implement the algorithm in hardware (e.g., in application-specific integrated circuits (ASICs)).

Anonymity. Coins can be passed in pseudo-anonymous fashion, similar to the Bitcoin. However, there is a built-in transfer mode that mixes the coin in transit, using the master nodes for this purpose. Therefore, the coins can be transferred in statistically anonymous fashion without the need for additional services.

Governance. The most interesting feature is the governance model, consisting of the group of major network stakeholders, i.e., the master nodes. Each master node has the right to vote on resolutions and the majority of voters decide on the structure and future of the network itself. Due to the design, the objectives of the network as a whole, its customers/users, and the master node operators are aligned to promote privacy, consistency, and security. The Dash network does not suffer any of the limitations of the Bitcoin blockchain governance.

C. Example Applications

Fully anonymous, atomic, and reliable peer-to-peer transfer of value is one of the most common examples of the blockchain technology application. It offers the potential to facilitate fully automated micro payments and full remittance automation. Due to built-in mechanisms for delayed payment and multi-party signatures, it is possible to build more complex contractual agreements between parties, and involve multiple participants in the value transfer. Bitcoin and Dash blockchains can be used to issue digital assets, or work as a public registry of ownership (e.g., land title management [19]). Recent developments include prospective adoption of blockchain technology to regulate insurance subsidies based on real-time risk pooling [20]. Alternative use cases involve decentralised identity management [21], or the use of the blockchain to verify and validate the existence of documents based on their hash, without making the actual content public [22]. Despite the novelty of those approaches, all applications share the public ledger concept as the essential operational principle.

III. DECENTRALISED COMPUTATION

Existing public ledgers, such as the Bitcoin blockchain, provide a decentralised, verifiable and mathematically consistent transaction tracking. Each newly created transaction is atomic, that is, it is either fully included into the chain, or it is discarded. This is similar to a distributed database system. The difference being that everyone can participate in maintaining that database, and there is no single central authority that dictates the rules.

The computational expressiveness of such a ledger is limited to several cryptographic operations. This has been a carefully chosen design decision to keep the computational complexity of validating and verifying transactions simple, so as to ensure broad participation. However, this effectively limits the computational capabilities of the ledger itself. Any state transitions or computations that do not use crypto-primitives must be executed by a trusted third party.

Ethereum [11] takes the next incremental step towards automating institutions. It has been designed from ground up to enable execution of arbitrary, Turing-complete code *within* the transaction itself, making it a distributed ledger and distributed execution environment at the same time. This means that the blockchain itself can host a transparent and inspectable process: a sequence of steps that express an algorithm, or state machine transitions that are monitored and executed by

the network itself. The user who wishes to invoke the logic must remunerate the network for the execution of all the operations. No single node or potentially inconsistent client implementations can be held responsible for executing that computation. Though the collective of nodes provides the computational capabilities, the computation itself is distributed across those nodes and cannot be unilaterally modified or prevented. To prevent or circumvent the execution, the entire network would have to be taken down, the prospect of which is unrealistic once a critical adoption level is reached. Going beyond the sovereignty-agnostic currency flow enabled by cryptocurrencies like Bitcoin and DashCoin, this means that specific executions do no longer underlie a single determinable jurisdiction, making the execution truly distributed in the sense of transparency and fungibility.

In practice, the system relies on *ether* (and its subdenomination *wei*) as fundamental unit of exchange that is needed to pay for deploying code. Ether is generated by miners but can be procured via exchanges. The users specify contracts, which can be as simple as modifiable objects, or as complex as long-running decision-making processes, like voting or deploying one's own cryptocurrency inside the Ethereum network. The required payment (*gas*) is estimated based on code complexity and charged to the deploying party. Contracts, or *smart contracts*, are created using the companion Javascript-inspired programming language Solidity [23]. Solidity allows the specification of a contract's stakeholders, permissible modifications, execution conditions (e.g., triggers for voting) as well as termination conditions. Deployed contracts are uniquely identified and publicly visible. The required remuneration for contract deployment deters from excessive use and is deposited during deployment. Unused funds are reimbursed if the initial projection was too high.

The new quality of automated enforcement of codified contracts highlights the importance of thorough development and testing, an aspect that has become evident in the recent first massive hack of an Ethereum DAO [12]. But in this young and dynamic field, solutions are already on the horizon. A proposed solution to this problem is the use of child chains to coordinate asset-based transactions as implemented in the new blockchain alternative Ardor/NXT 2.0 [24], which is under development and to be released for production use in 2017. In contrast to Ethereum's support for general-purpose code, Ardor will concentrate on specific asset-based transactions. The concept of child chains permits the delegation of specified operations onto a given sub chain, and thus increasing the security by limiting the visibility to relevant stakeholders. The security model is further strengthened by supporting complex preconditions for the execution of transactions. In addition, the delegation to child chains increases the scalability of the entire network by reducing the necessary decentralised computations. The concept furthermore includes built-in mechanisms to manage governance and decision-making processes in a reliable and anonymous fashion.

However the blockchain landscape will develop in the future, we see specifically the delegation of code execution into the blockchain itself as the game-changing feature, and the foundation of what we refer to as *Distributed Autonomous Institutions*. Table II provides an overview of essential institutional functions as performed in the discussed instances of blockchain technology.

TABLE II. BLOCKCHAIN TECHNOLOGY INSTANCES AND DISCUSSED CHARACTERISTICS

Technology	Validation	Governance	Managed Capabilities	Artefacts/
Bitcoin	majority-based voting	informal community-based	transactions	
DashCoin	stake-based	representative voting	transactions	
Ethereum	majority-based voting	informal community-based ^a	transactions & stateful autonomous code execution	

^a The community-based governance system is currently undergoing revision in the light of the recent DAO theft, with directions pointing towards the explicit appointment of governing entities based on constitutional principles (see e.g., [25]).

IV. DISTRIBUTED AUTONOMOUS INSTITUTIONS (DAI)

The outlined technological developments suggest that critical cooperative tasks can now be fully automated while retaining oversight, but without the ability to intervene. On first sight, this suggests the complete codification and delegation of cooperative tasks to the blockchain into a DAO. However, this naive conception obscures the reality of useful socio-technical systems. As with conventional socio-technical systems, the value of any system is determined by its usefulness to solve a specific, more or less well-defined task. However, the central determinant of usefulness remains the human stakeholder that interacts with the system, or, in extension, employs an artificial entity to interact with the system on one's behalf. Instead of replacing existing structures, the technological developments allow new formal organisational structures to emerge in such a way that it is the software that is at the centre of explicitly specified objective coordination tasks, freeing external entities from economically inefficient and potentially corruptible third-party oversight. We call those **Distributed Autonomous Institutions (DAI)** (see Fig. 3).




Figure 3. Relationships in Distributed Autonomous Institutions

In DAI the need for trust is eliminated, since the entire workings of the institution (Institution A in Fig. 3) are now transparent. The compliance-enforcing role is taken up by individuals. There is no need for a second institution (Institution B in Fig. 1) that overlooks workings of Institution A. Furthermore, the liability can be partially delegated to individuals.

The software and blockchain technology is capable of providing a transparent, verifiable process to achieve the same effects that traditional organisations achieve with the help of trusted human institutions and governmental services. The essential difference is that DAIs can be made difficult to circumvent and impossible to penetrate. For example, in

economically critical services, such as those offered by banks or governmental agencies, a closed system susceptible to manipulation or fraud can be replaced by a DAI that is more efficient, effective, and which cannot be circumvented by the human element. In contrast to DAOs, the proposed DAI construct includes the consideration of the physical, legal and social environment, as well as contractual relationships residing *outside* the blockchain. This is not meant to reflect a compromise of opportunity and reality, but the merger of the best of two worlds, enabling novel forms of multi-party business-to-business (B2B) operations in which the trust does not need to be mutually negotiated (e.g., relationships between k companies would require $(k(k-1))/2$ contracts) but be attached to an externalised single smart contract accessible and verifiable by either party. This raises enormous potential to construct ad-hoc operations, while providing opportunities to smaller market players that would otherwise not have the capacity to engage in formal negotiations. We specifically want to highlight the explicit formal specification, which, in principle, removes any need for ex-post legal interpretation, since the programmatically encoded agreement is indisputable in legal proceedings, substituting the judicative element necessary for the interpretation of conventional contracts.

The mechanisms discussed above have the potential to fundamentally change the way in which organisations can deal with any form of agreement enforcement, such as individual or collective employment contracts, voting in unions, crowd-funding of startups, or research and development initiatives. However, this notion of verifiable institutions offers novel applications for the revision of the transparent management of funds in governmental organisations, or to facilitate elections. In April 2016, the Minister for the UK Cabinet Office and Paymaster General, Matt Hancock, delivered a speech on Central Government Efficiency, Government Transparency, and Accountability [26], and argued for the use of blockchain technology in the governmental sector:

“We are exploring the use of a blockchain to manage the distribution of grants. Monitoring and controlling the use of grants is incredibly complex. A blockchain, accessible to all the parties involved, might be a better way of solving that problem. [...] Bitcoin proved that distributed ledgers can be used to track currency as it is passed from one entity to another. Where else could we use that? Think about the student loans company tracking money all the way from treasury to a student’s bank account. Or the department for international development tracking money all the way to the aid organisation spending the money in country. [...]”

Currently, we are experiencing the dawn of this technology, and we will experience the rise and demise of various blockchain instances, but we can be certain that the technology core is here to stay. Consequently, we will need to observe how it will change the structure of organisations, how we model socio-technical systems, but also what the ethical implications of concepts such as smart self-enforcing institutions are for our disciplines and society.

Inasmuch as we highlighted the benefits of the technology, we consequently need to be aware of the associated risks that follow suit. Will smart contracts and distributed autonomous institutions mimic the existing brick-and-mortar organisational structures, or will we observe new, qualitatively different loosely-coupled socio-technical systems? Can we

provide mechanisms that control the advent of novel schemes in which users enter contractual agreements they do not fully understand? Is the lack of case-based control, fraud or manipulation always desirable? Will democratic governments or public companies be expected to adopt transparent and verifiable processes based on the blockchain technology? Can blockchain technology be a solution to facilitate effective and efficient electronic voting? An important aspect in this context is to define how to redraw the line between public and private information (and to implement it). Does this technology prevent novel creative accounting practices (based on improved transparency), or will the low adoption threshold in fact stimulate the emergence of new variants of complex services (e.g., mortgage-backed securities) that have caused economic turbulences in the past? What will the accessibility of smart contracts mean for personal privacy in general?

V. SUMMARY, DISCUSSION & OUTLOOK

Distributed ledgers and distributed consensus protocols replace the need for third party trust. We have argued that the new technology enables the formation of private, anonymity-preserving, yet trustworthy automated institutions. This new flavour of institutions will have characteristics not found in current institutional constellations, due to the nature in which trust and liability are managed. This has the potential to fundamentally change the nature of institutions, because the human element can be eliminated. The blockchain technology allows new forms of governance, liability and trust to be shifted from traditional institutions (such as governments, banks, courts) to individuals and delegated to automated distributed autonomous institutions. The old and the new forms of organisations will co-exist by forming complex structures and interdependencies between human-centric and DAIs. We argued that those new forms of organisational structures are qualitatively distinct from existing institutions. Developing such systems will require a change in how we model systems in general, how we interact with them, but most importantly, how to determine and control the authority we delegate to those systems. This will inevitably involve research and analysis into the impact that DAIs will have on society at large.

To realise the benefits of developing transparent open coordination systems, substantial amount of work is required. Beyond the obvious technical challenges, this requires the consideration of social and legal implications. The potential anonymity enabled by the technology requires careful consideration for applications that may afford some public display of identity-related or pseudonymous information, such as in crowd-funding systems, or land title management. An essential aspect here is to prevent potential defamation by anonymous parties, e.g., by leveraging a comparable level of identifiability for all involved parties. Those are important design decisions that lie outside the technical platform provided by Ethereum, or blockchain technology more generally, and precede the implementation of a specific contract. An associated problem is the public nature of the blockchain. This implies the awareness that deployed code is and will be publicly accessible, both for inspection but also potential abuse, which lifts the challenge of developing high-quality non-exploitable code, an aspect we discussed in Section III.

Another important aspect revolves around the handling of conflicts. Whenever operating across system boundaries

– such as conventional private organisational environments and publicly-accessible institutions – conflicts can develop and manifest themselves based on changing local operations or environmental influences. The current state of blockchain technology in Ethereum does not consider a dynamic nature of contracts. Once deployed, contracts have a fixed interaction interface and codified operations. This neither considers the potential to adapt contracts at runtime, nor does include mechanisms to mediate conflicts directly. Instead, an alternative refined contract could be negotiated to replace the original contract (that could continue to coexist or simply be discarded). A central consideration in this context is the management of *ownership* of a given contract, i.e., the party/parties that manage/s the life cycle of a given contract. Per default, the instantiating party gains ownership, an aspect that is important for handling of funds that are allocated to a given contract, etc. Unlike conventional contractual agreements, the technically guaranteed executable contract specification affords the explicitly encoding of infrastructural aspects, such as the redistribution of outstanding funds to individuals, the payment of obligations by individual parties to sponsor the contract execution in the first place (i.e., the necessary *gas*), and the necessary actions for discarding a contract (e.g., multi-party invocation of a specified *discard* function).

These interdisciplinary aspects are grounded in technology, but reach far beyond the purely technical domain into management and the legal discipline. This makes it only more important to ensure the safe specification, deployment, and operation of smart contracts. To make smart contracts truly accessible, future development needs to provide mechanisms that allow non-technical users to write prototypical contracts while maintaining the essential institutional content. A possible approach includes the modelling in a widely accessible specification language and the translation into the corresponding execution language in a (semi-)automated manner. An alternative is to provide domain-specific ‘building blocks’, e.g., for the purpose of ‘voting’ or ‘auctioning’, in order to compose executable contracts that could be specified and reviewed by domain experts. An intermediate step would be the specification of best practices and provision of pattern repositories that contain thoroughly tested contracts ready for immediate instantiation.

Further support for developing smart contracts is complemented by the demand to make *existing* smart contracts easily accessible or interpretable to use blockchain technology for its essential purpose: to coordinate verifiable state in a decentralised manner. This would stimulate the broad adoption of this coordination infrastructure by applications and services in a potentially loosely-coupled manner, and extend the playing field beyond the current currency-centric niche existence of blockchain technology.

Bearing the potential and challenges of this novel technology in mind, one thing is certain: Lawrence Lessig captured the essence of DAI, and blockchain technology more generally, when he stated: “code is law” [27].

REFERENCES

- [1] G. M. Hodgson, “The Evolution of Institutions: An Agenda for Future Theoretical Research,” *Constitutional Political Economy*, vol. 13, no. 2, pp. 111–127, 2002.
- [2] L. Bloomfield, *Language*. New York (NY): Holt, 1933.
- [3] D. C. North, *Institutions, Institutional Change, and Economic Performance*. New York (NY): Cambridge University Press, 1990.
- [4] C. K. Frantz, M. K. Purvis, B. T. R. Savarimuthu, and M. Nowostawski, “Modelling Dynamic Normative Understanding in Agent Societies,” *Scalable Computing: Practice and Experience*, vol. 16, no. 4, pp. 355–378, 2015.
- [5] W. R. Scott, “Approaching Adulthood: The Maturing of Institutional Theory,” *Theory and Society*, vol. 37, no. 5, pp. 427–442, 2008.
- [6] D. C. North, “Institutions,” *Journal of Economic Perspectives*, vol. 5, no. 1, pp. 97–112, 1991.
- [7] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, “Chord: A scalable peer-to-peer lookup service for internet applications,” in *SIGCOMM’01*, 2001, pp. 149–160.
- [8] C. Dwork and M. Naor, “Pricing via processing or combatting junk mail,” in *Proceedings of the 12th Annual International Cryptology Conference on Advances in Cryptology*, ser. CRYPTO ’92. London, UK, UK: Springer-Verlag, 1993, pp. 139–147. [Online]. Available: <http://dl.acm.org/citation.cfm?id=646757.705669>
- [9] QuantumMechanic, “Proof of stake instead of proof of work,” <https://bitcointalk.org/index.php?topic=27787.0>, 2016, accessed on: 1st May 2016.
- [10] M. De Oliveira, M. Purvis, S. Cranefield, and M. Nowostawski, “A distributed model for institutions in open multi-agent systems,” in *Knowledge-Based Intelligent Information and Engineering Systems*. Springer, 2004, pp. 1172–1178.
- [11] G. Wood, “Ethereum: A secure decentralised generalised transaction ledger,” *Ethereum Project Yellow Paper*, 2014.
- [12] K. Finley, “A \$50 million hack just showed that the dao was all too human,” <http://www.wired.com/2016/06/50-million-hack-just-showed-dao-human/>, 2016, accessed on: 1st June 2016.
- [13] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system,” 2008.
- [14] A. M. Antonopoulos, *Mastering Bitcoin: unlocking digital cryptocurrencies*. O’Reilly Media, Inc., 2014.
- [15] E. Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action*. New York (NY): Cambridge University Press, 1990.
- [16] E. Smart, “Does china’s control over bitcoin mining threaten bitcoin?” <http://dcebrief.com/does-chinas-control-over-bitcoin-mining-threaten-bitcoin/>, January 2016, accessed on: 1st June 2016.
- [17] J. Harrington and G. Caffyn, “Chainalysis network monitoring,” http://bit.ly/chainalysis_1, http://bit.ly/chainalysis_2, 2015, accessed on: 1st June 2016.
- [18] A. Quenston, “Ethereum reaches unanimous agreement to hardfork,” <https://www.cryptocoinsnews.com/ethereum-reaches-unanimous-agreement-hardfork/>, 2016, accessed on: 10th July 2016.
- [19] Bitland, “bitland - land title protection ghana,” <http://www.bitland.world/>, 2016, accessed on: 1st June 2016.
- [20] S. Higgins, “Congressional committee hears testimony on blockchain in health care,” <http://www.coindesk.com/us-think-tank-suggests-blockchain-application-insurance-risk-pooling/>, 2016, accessed on: 1st June 2016.
- [21] Blockstack, “What is blockstack?” <https://blockstack.org/docs/what-is-blockstack>, 2016, accessed on: 1st June 2016.
- [22] M. Araoz, “Proof of existence,” <https://proofofexistence.com/>, 2015, accessed on: 1st June 2016.
- [23] Ethereum Team, “Solidity,” <http://solidity.readthedocs.io/en/latest/>, 2016, accessed on: 1st May 2016.
- [24] NXT, “Announcing nxt 2.0!” <http://nxt.org/roadmap/>, 2016, accessed on: 1st June 2016.
- [25] m88888m, “Ethereum with a constitution and legislative initiatives can become a true democracy. while bitcoin still lingers in a plutocratic civil war,” https://www.reddit.com/r/ethereum/comments/4qhp01/ethereum_with_a_constitution_and_legislative/, 2016, accessed on: 10th July 2016.
- [26] M. Hancock, “Digital transformation in government and blockchain technology,” <https://www.gov.uk/government/speeches/digital-transformation-in-government-and-blockchain-technology>, April 2016, accessed on: 1st June 2016.
- [27] L. Lessig, *Code and Other Laws of Cyberspace*. New York (NY): Basic Books, 1999.