Managing Knowledge Provenance using Blockchain

Utkarsh Mittal, Tom Gedeon, Sabrina Caldwell
Research School of Computer Science
The Australian National University
Canberra, Australia
{Utkarsh.Mittal|Tom.Gedeon|Sabrina.Caldwell}anu.edu.au

***Abstract*—The emergence of a “post-fact” world has seen knowledge being misinterpreted and manipulated to suit diverse purposes. Information from credible sources is often drowned out by ‘facts’ with dubious provenance. It is not enough to simply establish knowledge provenance if this meta-data can be changed afterwards. This paper presents a possible solution using blockchains to manage knowledge provenance. The solution is modelled using a concrete example of knowledge institutions. Variations and extensions to suit other domains are suggested. Particularly, a graph based variation of the blockchain is outlined. We also discuss the impact of our model beyond technology.**

Keywords—knowledge provenance; knowledge management; blockchain; permissioned blockchain; evidenced knowledge statements; cognitive infocommunications; post-fact; post-truth

I. Introduction

The events and politics of the 2016 U.S. elections and UK European Union referendum made it incontrovertibly obvious that we are now in a “post-fact” world. A “post-fact” or “post-truth” world is one where information from questionable sources is bandied about as fact. Appeals to emotion are often used to convince people of the validity of knowledge. This phrase and the events surrounding it were significant enough that Oxford Dictionaries named “post-truth” their 2016 Word of the Year [1]. Given the nature and origin of this problem, there has been a lot of discussion surrounding it in the public sphere. Purveyors of knowledge such as newspapers have taken steps to be more transparent, especially online. In late 2017, BBC added a link at the bottom of all of its web pages titled “Why you can trust BBC News” [2].

Notably, multiple authors have commented on what this phenomenon means for academics and scientists. Coughlan mentions relativism as a philosophical parallel to “post-truth” [3]. If truth is relative, are facts also? Higgins in her article calls relativism itself relative. She argues that even staunch relativist philosophers such as Nietzsche “held intellectual honesty at a premium” [1]. Coughlan suggests that the way to combat this erosion of trust is for scientists to speak out when scientific facts are ignored [3]. But what makes a scientist more deserving of trust? Is it the implicit assumption that scientists seek only the truth? Other authors note that simply broadcasting good ideas to combat bad ideas is no longer enough [4, 5].

There is another facet of the “post-fact” era, which often gets ignored. Not only are lies common, it is becoming harder to hold someone accountable and answerable to their earlier statements, or the knowledge bases for their decisions. This problem is compounded by the fact that in the digital age, information can be falsified or skewed with little effort. This makes it easy for an individual to claim that they “never said that” and that all proof to the contrary is lies or post-hoc manipulated evidence.

The issues surrounding trust and fact communication will play a part in the future of cognitive infocommunications (CogInfoCom). Existing work in the field focusses on interaction between humans [6-8] and between human and artificial systems [9-11], among other things. Both of these paradigms (intra-cognitive and inter-cognitive communication [12]) have an implicit assumption of trust. While multiple authors [11, 13] in the field have written about the representation of knowledge in the digital space and the trustworthiness of knowledge on the web [14], we would like to focus on the origin or provenance of that knowledge statement. As knowledge systems become increasingly autonomous, provenance and accountability will become an important layer such frameworks.

We believe knowledge provenance using blockchains could solve these issues. Specifically, we focus on two aspects of the problem. One, providing a secure platform for knowledge institutions to publish facts and provide provenance for knowledge. Two, extending this model to individuals and organizations and showing how the public can use it to hold social actors such as governments or corporations accountable. Our model leverages the decentralized and update-only properties of the blockchain to achieve these goals.

To limit discussion within the scope of this paper, it is necessary to define what we mean by knowledge. In the field of Information Technology, knowledge is often considered the third rung in a ladder comprising of data, information, knowledge and wisdom [15]. However, McInerney, takes a broader view and considers knowledge dynamic and not something that can be placed so neatly in a hierarchy [16]. Based on Davenport and Prusak’s definition [17], knowledge can be considered to be the subjective experience, values and insights of a knower. However, this definition does not capture the notion of data and information being a part of knowledge. In our paper, we consider a knowledge statement to be

978-1-5386-7094-1/18/$31.00 ©2018 IEEE

composed of statements of fact ("Average global temperature has increased by 0.8 C in the past 40 years") and/or statements of evidenced belief ("Global warming is caused by humans").

In the following section, we provide a general overview of blockchains and the variant used in our model. The method section illustrates our model with the help of a use case. Following that, we discuss some modifications to the base model. We finish by discussing the impact of the model on society.

II. Blockchains

Blockchains were first described in a paper by Nakamoto [18] about Bitcoin, the digital currency. While Nakamoto discusses the use of blockchains only for digital currencies, the blockchain distributed ledger data structure on its own has properties that make it useful for many ventures.

The original paper proposes a formulation which has a distributed timestamp server. Each node accepts a block of items which need to be timestamped and then hashes them with the time value and broadcasts the hash. Each timestamp includes the previous timestamp in the hash. This leads to a chain of blocks. Each block is essentially a collection of transactions. To ensure the blockchain is trustworthy, the hashed values in the blockchain should be difficult to replicate. Nakamoto achieves this by using a proof-of-work formulation which requires time intensive computations [18]. Thus, mining a block is time intensive, and so if the majority of the nodes in the network are honest, the longest chain in the network will be honest as well. Therefore, the longest chain is accepted as the version depicting the true ordering of transactions. This process is known as consensus.

Since the invention of the blockchain, much research has gone into developing improvements and variations which make it more suitable for different domains. For our model, we use a private/permissioned blockchain. The major difference from the original public blockchain is that the list of users who can mine nodes or make transactions on the blockchain is restricted. The access control mechanism can differ depending on the implementation and use case [19]. There is some confusion over terminology in this area as private blockchains are also referred to as permissioned or consortium blockchains [19-22]. Some authors consider private and permissioned blockchains to have different meanings [21, 22]. Specifically, private blockchains allow only one authority while consortium or permissioned blockchains have multiple controlling members.

In this paper, we use the term permissioned blockchain to signify a blockchain where only some members can add and verify blocks. All members are semi-trusted and their identities are known. Given the semi-trusted nature of the contributors, we can relax the trustless constraint that made proof-of-work consensus necessary in the original formulation and in the process, overcome drawbacks such as energy intensiveness [23]. Bano et al provide a comprehensive overview of different consensus algorithms and their classifications [24]. For our model we use a committee voting based consensus approach [24, 25]. In this approach, a majority vote of the committee is required to verify a block. Since all members in our committee will be known, this consensus algorithm is valid for our domain. We will later show that scalability should also not be an issue for our model.

III. Method

To elucidate our model, we outline a specific use case.

A. Use Case

Consider a fictitious group of seven Universities: Uni A, Uni B, Uni C, Uni D, Uni E, Uni F and Uni G. All the Universities exist in the same country but are otherwise independent of each other. They have formed a coalition together called the 'Group of 7" or Go7. Together, the Go7 wants to provide a stronger platform for scientific evidence and exchange. To this end, each University has decided to publish knowledge statements about their beliefs based on current research. These statements are not traditional research papers but are rather statements of evidenced knowledge agreed upon within a University. Due to the possible contrasting views between Universities, each statement will be authored by only one University. But together they want to provide a strong platform for this service.

B. Model

The Go7 use a permissioned blockchain with a committee voting consensus algorithm. Only the members of the committee (the Go7) are allowed to write blocks to the blockchain. Each block will consist of one or more statements by a single University. If a majority of the committee (4 of out 7 members) verify the block as being authentic, it is written to the blockchain. Voting here indicates that a committee member believes the block originated from the author and does not indicate agreement with the views in the statement itself. This setup necessitates that all committee members are connected in the network. This is a given if the network is assumed to be the internet. We call this blockchain the "Knowledge Blockchain".

A University can choose to update a statement by writing a block which also includes the hash of the previous block which contained the earlier statement. For example, Uni A issues a statement recording its current views on global warming every year. During this time, other blocks have been added to the knowledge blockchain. In 2018, when Uni A wishes to issue a new statement, it will include a hash of the block from 2017. This forms a link between the updates which makes searching easier. The use of blockchain also establishes an ordering between the knowledge statements and their contents.

The public can access the knowledge blockchain in two modes. The first mode is the "Headers-only" mode [26], which was described as "Simple Payment Verification" [18] in the original Nakamoto paper. For our model, we call this a "thin client". A thin client stores only the headers of the blocks from the blockchain and the data which the user specifically requests. This approach allows the user to access the data with minimal resources while still providing verification. The second mode is the standard "Full Node" [26] mode, where the client maintains a full copy of the blockchain. In our model, we call this a "thick client". A thick client allows a user to peruse the data without being connected all the time. We refer to both

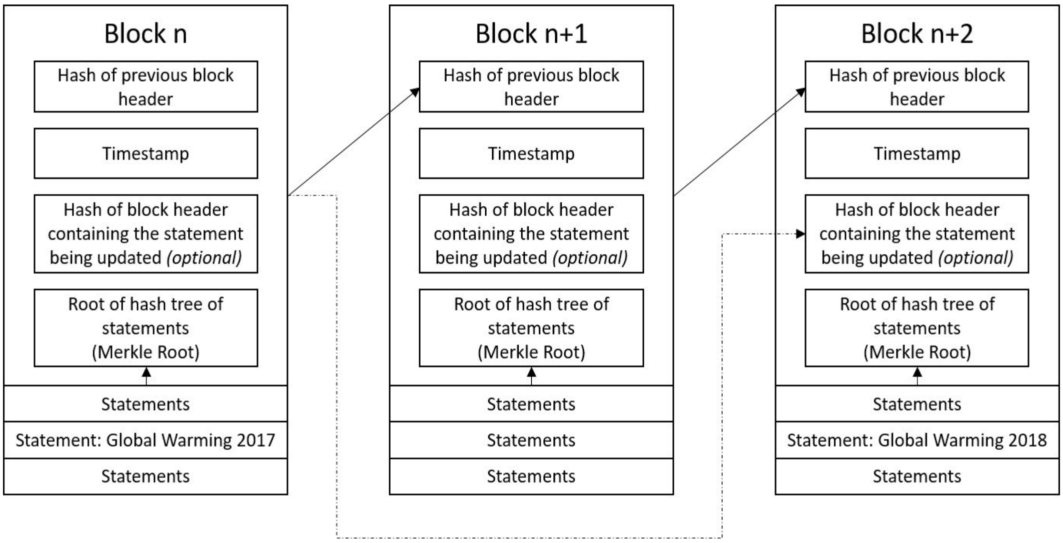
Block n
Hash of previous block header
Timestamp
Hash of block header containing the statement being updated *(optional)*
Root of hash tree of statements (Merkle Root)
Statements
Statement: Global Warming 2017
Statements
Block n+1
Hash of previous block header
Timestamp
Hash of block header containing the statement being updated *(optional)*
Root of hash tree of statements (Merkle Root)
Statements
Statements
Statements
Block n+2
Hash of previous block header
Timestamp
Hash of block header containing the statement being updated *(optional)*
Root of hash tree of statements (Merkle Root)
Statements
Statement: Global Warming 2018
Statements

Fig. 1. Block n+2 contains a statement updating the statement from Block n

types of clients together as “Hosting Nodes”. Hosting nodes have read-only capabilities and do not have write permissions.

To use the statements in a formal setting, for example journalism, we want the users to be able to identify themselves and record when they have accessed the information. This is similar to the citation system common in research. It serves as a powerful accountability tool on both sides, for the Universities and more importantly, for the users.

To avoid clutter on the knowledge blockchain, this access information is stored on a separate blockchain that we term “Access Record Blockchain”. The access record blockchain is maintained by the committee (Go7) in a similar fashion as the knowledge blockchain. Each block consists of one access request by a user which will contain a timestamp, a hash of the last block on the knowledge blockchain and optionally a field denoting which information was accessed. The access itself is provided through a smart contract [27], where a user has to provide valid identification to register their access on the blockchain. The identity validation method can be chosen as per the domain the model is being used in. For example, the user can be required to provide valid banking details and the smart contract will execute a $0.01 transaction to verify it and hence prove their identity.

C. *Example*

These two blockchains in tandem provide a mechanism for knowledge dispersal, storage and accountability. Let us consider an example. Uni B publishes a statement saying global warming is untrue on the knowledge blockchain at time T1. This information is now accessible from any hosting node or committee member node. Say User X wishes to write about Uni B’s statement. If they did so without registering access, their article is not credible as they cannot prove where or when they got the information. Say User X registers access on the access record blockchain at time T2. They can now formally say that all views in the article are based on Uni B’s statements of time T2. Say Uni B publishes an updated and contradictory statement at time T3 declaring global warming is true. Uni B cannot now call out User X for publishing false information as the article was true as of time T2 only (as proved by the access record). Now say another user, User Y, registers access on the access record blockchain at time T4 but publishes an article stating Uni B does not believe in global warming. Uni B can call them out and prove that as of time T4, their belief is that global warming is true.

D. *Limitations*

This model might not scale well as the voting consensus algorithm can be quite slow if there are a large number of blocks to be written per second. But since the purpose of the knowledge blockchain is to store statements, it will not see such huge traffic. Practically, no University is expected to issue more than a dozen statements a day. A similar problem exists for the access record blockchain. This limitation is not as easy to overcome but we discuss a possible solution in the next section.

Another potential problem with scalability is if the number of committee members gets too large. If the example was extended to include all the Universities in the world, it might become infeasible to do majority voting. In that case, instead of a majority, something like 10% of all members could be required. Damaging the knowledge blockchain would still require an attacker to control 10% of all the trusted Universities in the world.

IV. Variations and Extensions

In this section, we present some modifications of the base model which are useful in different domains or have other properties.

A. *Publicly maintained access record blockchain*

The access record blockchain can face problems with scalability if it is maintained by the committee. If the number of access registrations is too high, the relatively smaller number of verifying authorities can cause a bottleneck. The solution is to increase the number of nodes which can add blocks to this access record blockchain. Since the identity verification is handled by a smart contract, the real purpose of signing and verifying a block is to establish ordering within this blockchain. This task can be farmed out to the public and the access record blockchain as a whole can be publicly maintained. The problem then becomes to provide incentives for the public to maintain the blockchain. The traditional practice is to award digital currency coins for mining. If this is implemented, the blockchain must be converted to a proof-of-work implementation and a valid use must be found for the currency itself. Considering the state of cryptocurrencies, we believe the coins may be useful simply to trade for other currencies. This is somewhat similar to stocks of a company. We would like to note that we are not comfortable with proof-of-work implementations due to the massive environmental effect associated with them [23], but the notion of a tradeable currency based on knowledge is otherwise appealing.

B. DAG based blockchain structure

A key feature of our knowledge provenance model is to provide the ability to issue updated statements and link them together. Since the blockchain is sequential, the only way to maintain links across blocks is to include the hash of some older block in the new block. For our Go7 use case, we consider sequencing between statements in the same field to be very important. To facilitate easy traversal, we propose a directed acyclic graph (DAG) version of the blockchain similar to the one described by Popov [28]. Unlike Popov, who presented the DAG based "tangle" to increase scalability, we propose using a DAG for reasons of coherency.

The knowledge blockchain would start with a root node authored by no one. It is simply a genesis node to build the graph on. For this model, we consider all blocks to be nodes. Each University will link a separate node to this genesis node. All statements published by a University will now be within their own chain. For example, Uni A will publish a statement by attaching a node to their own. In fact, Uni A can create different branches within its chain. They might prefer to maintain a separate branch for all their statements regarding global warming. Each update statement will belong to this chain. The DAG structure also allows for a single node linking to multiple predecessor nodes. Perhaps Uni A wants to issue a statement which builds on their existing statements about global warming and also their existing statements about coral bleaching events. Moreover, the structure can be used by Universities to issue joint statements. Uni A and Uni B can insert a node linking to both their chains containing a joint statement about global warming.

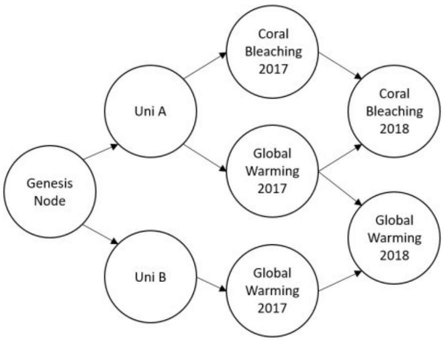

Fig. 2. DAG blockchain with joint statements and multiple predecessors

Another advantage of this structure is that it allows us to merge the access control blockchain into the knowledge blockchain. Each access registration can be maintained in a separate chain stemming from the genesis node or even attached to the node that it actually accesses. If a user wants to access Uni A's 2018 statement on global warming, the access registration is attached to that node.

The DAG structure makes the flow of knowledge easier to understand and navigate and is feasible for a slow moving domain with semi-trusted entities (Go7).

C. Blockchains for individuals or organizations

The model can be extended to commercial organizations or individuals. A group of companies from a specific sector can come together to host a blockchain for their press releases or evidenced belief statements. Powerful individuals who are in the public eye can do the same thing. In cases like these, the knowledge blockchain acts as a powerful accountability tool for the public. An organization cannot change published information on the blockchain and deny it altogether. Hence, the public (or the media) can call out the organization when it does not follow through or act as per the statements they have issued on the knowledge blockchain. The presence of the statements on a decentralized secure platform bolsters accountability.

D. Knowledge blockchains as a tool for journalism

Newspaper and news agencies already maintain vast archives of their published articles and pieces. If this was maintained on the knowledge blockchain, it would increase transparency for the public. Journalists will be capable of proving what they said and when they said it. The public will be able to hold journalists responsible for contradictory statements, as incontrovertible proof will exist outside of the control of the company. Currently, equivalent verifiability is provided by the paper versions of past publications. As publishing moves more and more online, and even the evidence of the paper versions may only be available in digital form, the need for proof will only mount.

V. Discussion

Using emerging technologies to solve problems such as knowledge provenance has ramifications which extend beyond computer science. It is necessary to understand and discuss the effects of our model on society.

One of the main tenets of the internet is that everyone has an equal voice. It provided a platform for people to bypass the traditional gatekeepers of knowledge and balanced the power equation by giving everyone equal opportunity. Our model implies that statements on the knowledge blockchain are inherently more trustworthy since they cannot be changed or altered. Anyone using this platform could have a stronger voice. This creates a potential power imbalance. The obvious solution is for all parties to use a blockchain based solution. But currently blockchains are expensive and early adopters of the model are likely to be the rich and powerful. We consider this to be a smaller problem than the accountability issues our model solves and hence our work still represents progress.

A corollary of having a platform maintained by a specific group is that the views on the platform are likely to be similar and only cover a small range. The minority voice may not be on the blockchain. As discussed earlier, this causes a potential power imbalance. We envision knowledge statements in blockchains to be initially be used in areas considered incontrovertible by scientists yet regularly attacked and distorted by vested interests. Our oft quoted case of global warming is a prime example. Moreover, we consider our model to be a data structure for maintaining knowledge statements and not a platform for debate (academic or otherwise).

The knowledge blockchain model also allows accountability beyond our example with Universities. In the previous section, we discussed how the model can be extended to individuals or organizations participating in public discourse. For example, a politician who chooses to use the blockchain is now accountable to the public. Journalists can quote their statements when they do not follow through. The individual cannot deny the statement as they do not individually have the power to change the blockchain record of the statement. They can also not claim that their views were not reported correctly as the statements are their very own commits to the blockchain.

Widespread use of this model can also cause pressures on individuals. Choosing not to make a statement on the blockchain about an issue can be seen as a statement in itself. Individuals may be forced to make statements simply because others in the public eye have made statements along similar lines. This problem exists outside the scope of our paper however and is an issue faced in all existing spheres of public discourse as well. Our model neither exacerbates nor ameliorates the situation.

Having a small group in charge of the model (as in the Go7 use case) increases the chances of collusion. This undermines the very security of a blockchain and is not really within the scope of any blockchain based solution but we would like to present some solutions which are applicable in this domain. Let us say that all 7 Universities have at some point published a statement about global warming being caused by humans. In the future, if there is conclusive proof that global warming is not caused by humans, the Universities may want to collude and change their past statements on the blockchain to save face. We can place a check on this by adding a rule to our design, that updates to any blocks older than a day are invalid. The hosting nodes are essential to this process. If they receive an update from the committee nodes where blocks in the knowledge blockchain older than a day have been changed, they can refuse to accept them and invalidate the trustworthiness of the whole knowledge blockchain. Another variation is to add the latest block from the access record blockchain to the knowledge blockchain at regular intervals. This creates a cross reference between the chains and checkpoints both of them. Essentially, corrupting the knowledge blockchain requires corruption of both chains, which is a substantially more difficult task as there are likely to be many more thin clients. Adding knowledge blockchain blocks to the access control blockchain provides the same benefit in the other direction. The whole process can be made "neater" by adding only the hashes instead of the actual contents since the hash is all we need for validity checks.

VI. Conclusion

Our model for managing knowledge provenance leverages the benefits of a blockchain to provide a solution which is decentralized, immutable and easy to update. By presenting knowledge statements in this form, we provide a data structure which overcomes many of the problems presented by the "post-fact" world without fundamentally changing the nature of academic discourse itself. Extensions such as the DAG based blockchain model showcase a better way to organize the information in the same paradigm. With our model as the base, we would like to try and find technical solutions to further problems in the area of trust and accountability. Implementing the model in a usable domain would also help to uncover further challenges and areas of research.

Acknowledgment

Utkarsh Mittal thanks Sneha Vaidya for her help in understanding the impacts of this model on society.

References

[1] K. Higgins, "Post-truth: a guide for the perplexed," Nature News, vol. 540, no. 7631, p. 9, 2016.

[2] BBC News Services. (2017, 25 May, 2018). Learn how the BBC is working to strengthen trust and transparency in online news.

[3] S. Coughlan, "What does post-truth mean for a philosopher?," BBC, 12 January, 2017. Accessed on: 25 May, 2018.

[4] N. Enfield, "We're in a post-truth world with eroding trust and accountability. It can't end well," The Guardian, 17 November, 2017. Accessed on: 27 May, 2018. Available: https://www.theguardian.com/commentisfree/2017/nov/17/were-in-a-post-truth-world-with-eroding-trust-and-accountability-it-cant-end-well

[5] F. Fukuyama, "The Emergence of a Post-Fact World," Project Syndicate, 12 January, 2017. Accessed on: 27 May, 2018. Available: https://www.project-syndicate.org/onpoint/the-emergence-of-a-post-fact-world-by-francis-fukuyama-2017-01

[6] G. Csapó, "Sprego virtual collaboration space," in Cognitive Infocommunications (CogInfoCom), 2017 8th IEEE International Conference on, 2017, pp. 000137-000142: IEEE.

[7] I. Horváth, "Innovative engineering education in the cooperative VR environment," in Cognitive Infocommunications (CogInfoCom), 2016 7th IEEE International Conference on, 2016, pp. 000359-000364: IEEE.

[8] V. Kövecses-Gősi, "Cooperative Learning in VR Environment," Acta Polytechnica Hungarica, vol. 15, no. 3, 2018.

[9] J. Irastorza and M. I. Torres, "Analyzing the expression of annoyance during phone calls to complaint services," in Cognitive Infocommunications (CogInfoCom), 2016 7th IEEE International Conference on, 2016, pp. 000103-000106: IEEE.

[10] V. Sarathy, M. Scheutz, and B. F. Malle, "Learning behavioral norms in uncertain and changing contexts," in Cognitive Infocommunications (CogInfoCom), 2017 8th IEEE International Conference on, 2017, pp. 000301-000306: IEEE.

[11] S. Savić, M. Gnjatović, D. Mišković, J. Tasevski, and N. Maček, "Cognitively-inspired symbolic framework for knowledge representation," in Cognitive Infocommunications (CogInfoCom), 2017 8th IEEE International Conference on, 2017, pp. 000315-000320: IEEE.

[12] P. Baranyi and A. Csapo, "Definition and synergies of cognitive infocommunications," Acta Polytechnica Hungarica, vol. 9, no. 1, pp. 67-83, 2012.

[13] A. Adamkó, B. Balázs, E. Krisztián, F. Attila, H. N. Kristóf, and K.-F. Norbert, "Smart campus service link: Adaptation and interaction planes for campus and university environments," in Cognitive Infocommunications (CogInfoCom), 2017 8th IEEE International Conference on, 2017, pp. 000271-000276: IEEE.

[14] E. Jozsa, A. Komlodi, R. Ahmad, and K. Hercegfi, "Trust and credibility on the web: The relationship of web experience levels and user judgments," in Cognitive Infocommunications (CogInfoCom), 2012 IEEE 3rd International Conference on, 2012, pp. 605-610: IEEE.

[15] M. Alavi and D. E. Leidner, "Knowledge management and knowledge management systems: Conceptual foundations and research issues," MIS quarterly, pp. 107-136, 2001.

[16] C. McInerney, "Knowledge management and the dynamic nature of knowledge," Journal of the Association for Information Science and Technology, vol. 53, no. 12, pp. 1009-1018, 2002.

[17] T. H. Davenport and L. Prusak, Working knowledge: How organizations manage what they know. Harvard Business Press, 1998.

[18] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," ed, 2008.

[19] P. Jayachandran, "The difference between public and private blockchain," in Blockchain Unleashed: IBM Blockchain Blog vol. 2017, ed, 2017.

[20] V. Buterin, "On Public and Private Blockchains," in Ethereum Blog vol. 2017, ed, 2015.

[21] J. Garzik, "Public versus Private Blockchains Part 1: Permissioned Blockchains," ed: BitFury, 2015.

[22] A. Varshney, "Types of Blockchain—Public, Private and Permissioned," 1 April 2017. Accessed on: 21 November 2017.

[23] "The great chain of being sure about things," The Economist,, 31 October 2015. Accessed on: 20 November 2017. Available: https://www.economist.com/news/briefing/21677228-technology-behind-bitcoin-lets-people-who-do-not-know-or-trust-each-other-build-dependable

[24] S. Bano et al., "Consensus in the Age of Blockchains," arXiv preprint arXiv:1711.03936, 2017.

[25] Hyperledger Architecture, Volume 1 (White Paper). [Online]. Available: https://www.hyperledger.org/wp-content/uploads/2017/08/Hyperledger_Arch_WG_Paper_1_Consensus.pdf. Accessed on 25 May, 2018.

[26] R. Skudnov, "Bitcoin clients," Degree Program in Information Technology, Turku University of Applied Sciences, 2012.

[27] V. Buterin, "A next-generation smart contract and decentralized application platform (White Paper)," Accessed on: 27 May, 2018. Available: https://github.com/ethereum/wiki/wiki/White-Paper

[28] S. Popov, "The Tangle (White Paper)," Accessed on: 22 November 2017. Available: https://iota.org/IOTA_Whitepaper.pdf