Effective governance of and by the blockchain

Anwaar Ali

University of Cambridge
Computer Laboratory
Wolfson College

Principal supervisor: Professor Jonathan Crowcroft

January 2019

This report is submitted as
the second year's report for PhD

Declaration

This report is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text.

This report does not exceed the regulation length of 4000 words, excluding tables, figures and bibliography.

Effective governance of and by the blockchain

Anwaar Ali

Executive Summary

As stipulated in this footnote[1] by the Laboratory, this document is primarily a *dissertation schedule* (in blockchain's vernacular, it is akin to writing a smart contract's interface much like this one[2]) with cross references to my first year report [1] to gauge my progress as per the original proposed plan (see Table 4.1 on Pg. 36 in [1] for the summary of the said plan). Further, in this document I present a tentative outline for my final doctoral thesis (subject to change) with a few placeholders (such as for prospective images and tables etc) deliberately left blank which indicate what sort of experiments, processes, pending tasks are inline and will be completed alongside populating this document with the intention that it eventually morphs into my final doctoral thesis.

Summary of progress so far: At a glance it can be seen that in Table 4.1 (Pg. 36 in [1]) I divided my second year into four quarters with specific tasks assigned to each quarter. As per the current progress, the task of submitting a survey article to IEEE COMST in the first quarter is complete. Further, an extensive boilerplate has been implemented in the form of an overarching smart contract (OSC) (the task specified in the second quarter). The code is mainly based on this Github repository[3] (*Note:* I am planning to make my own code online later during my PhD's third year). However, the experimentation process (as specified in the third quarter) is currently still under progress. Although this task seems to be carrying over to my third year, however, I am quite confident it will not be long before it is completed as well. As the task in the fourth quarter was conditional on the completion of the task in third quarter, completing the task in third quarter will then imply an immediate completion of the task in the fourth quarter.

The boilerplate setup, as mentioned above, will primarily be used to study the problem of resource consumption and automatic pricing against it for community networks. The very nature of such networks demands a trusted and efficient *trustless environment* for record keeping and to facilitate trusted inter-peer interactions. Blockchain with its trusted record keeping and automation capabilities in the form of smart contracts just happen to provide such provisions. There are, however, quite a few caveats while doing so with blockchain such as upgradability of smart contract logic, fetching off-chain data for a

[1] `https://www.cst.cam.ac.uk/local/phd/year2-report`

[2] `https://solidity.readthedocs.io/en/v0.4.20/contracts.html#interfaces`

[3] `https://github.com/StephenGrider/EthereumCasts/tree/master/kickstart`

particular application, compliance with data protection regulation, and scalability. These are the main issues around which my research revolves and, ideally, an efficient solution to these problems will be the distinguishing feature of my doctoral research. Throughout this report I highlight how I am planing to tackle these issues along with details of what sort of experiments I already have or I am currently conducting. In my opinion, the use case of considering the resource consumption and automatic and fair pricing provides a good substrate to study the overarching theme of my PhD which is to study an effective integration of policies and laws with blockchain technology. I also intend to study the problem of trusted data provenance collection for trusted data audits. This will involve interfacing my boilerplate with CamFlow[4] [5].

I would also like to mention a few of my other activities during the second year of my PhD. I was fortunate enough to do an internship at Universitat Politcnica de Catalunya (UPC) in Barcelona Spain during the summer of 2018. It was there when I discovered the use of automating resource consumption in community mesh networks with blockchain technology [2]. As a result of my this collaboration I managed to take part in two publications [7, 4] that are quite relevant to my own doctoral research. These two publications are mainly about our experience of deploying two blockchain platforms (Hyperledger Fabric and Ethreum) in community mesh networks and measuring their performance. Given my this experience, I will be conducting further such evaluation experiments in my own research as well mainly for benchmarking purposes.

[4] `http://camflow.org/`

Acknowledgements

Thanks Jon for being an understanding and accommodating supervisor.

Contents

Chapter 1

Outline of the dissertation

This chapter presents a tentative outline for my final PhD dissertation. I provide a brief chapter-by-chapter summary of my prospective PhD dissertation with checkpoints to let the readers gauge the current state of progress in each chapter.

Note: I prepend the following headings in this chapter with (Chapter x) to highlight that I intend this to be Chapter x of my PhD dissertation.

1.1 (Chapter I) Motivation

This chapter is intended to motivate the usage of blockchain for a few use cases particularly those where a trustless environment and distributed record-keeping makes more sense. In my case, the use cases of resource consumption in community mesh networks [3, 2] and data provenance-based trusted data audit [6] will be motivated. The subject matter of this chapter will be quite similar to my Introduction Chapter 1 in [1]. It will also contain a literature review surveying the current state of blockchain deployments in different areas (this chapter will be highly influenced by our recent survey article [**?**]).

This chapter will address the questions such as why consider blockchains at all? What is trustless environment and why it would make sense in a few use cases (as described in the last paragraph). Why I am considering the above mentioned use cases? What do I mean by compliance, compliance by design, making blockchain-based system in compliance with a given set of policies and rules (what kind of sets of policies and rules I am considering in particular)? The important question of *extent* to which a blockchain-based system can be made compliant with a policy set. I will introduce the idea of *compliance by design* much like the idea of *privacy by design* as put forward by the GDPR. The question of extent will, in general, be discussed throughout the subject matter of the dissertation with a special focus with lessons learnt and concluding remarks treated extensively in its own chapter (as shown by Section 1.6 in this report).

Current progress: The actual chapter hasn't been written yet but I believe I have reasonable material to populate this chapter as I will be borrowing quite a lot from my first year report [1] and a recent extensive literature survey [].

1.2 (Chapter II) Use case: Community networks

The study of economic sustainability of community mesh networks can be considered as one of the significant changes to the original text of my first year report [1]. I visited Prof. Leandro Navrarro[1] in the summer of 2018 and it was there when I was highly motivated to study this problem through blockchain's perspective.

This study can be justified by observing that community mesh networks are, in general, heterougenous in nature. Further, such networks are usually a result of a community's volunteer efforts where different entities come together to pool their networking infrastructure resources to provide internet access to the still disconneted masses of a region. The main ingredients of blockchain which are distributed record keeping, decentralized consensus, and trusted automation of interactions among the peers of a network seem to fit quite well with the idea of community networks which inherently require a decentralized and trustless environment in order for them to function efficeintly and sustainably [3, 2].

Throughout this chapter, I will extensively discuss the background related to community mesh networks (such as Guifi.net[2]), how make such networks economically sustainable [3], how to automate resource consumption inventory keeping and automatic and fair pricing using blockchain and smart contracts.

Current progress: The actual chapter hasn't been written yet but the research is well underway. So far these publications have been worked upon [4, 7], a boilerplate project is implemented (based on this footnote[3], see [2] for futher details).

[1] `http://people.ac.upc.edu/leandro/`

[2] `https://guifi.net/en/what_is_guifinet`

[3] `https://github.com/StephenGrider/EthereumCasts/tree/master/kickstart`

1.2.1 Background

1.2.2 Economic sustainability

1.2.3 Resource consumption and pricing

1.2.4 Blockchain feasibility

1.2.5 Blockchain automation

1.3 (Chapter III) Use case: Data provenance and system audit

I will mainly be adopting the approach as described in [1].

Current progress: Much of the work related to this theme still needs to be done. However, the boilerplate setup and the set of experiments as mentioned in Section 1.2 will make is relatively easier to conduct a similar evualation study with this use case.

1.4 (Chapter IV) Experimental components and setup

Most of what I will be doing in this chapter will be highly influenced by what I say in this internship report [2]. Basically I will be talking about: Why adopt a kickstart mockup as a toy substrate example? Why oracles? Why delegateCall, need for logic update? Which platform? Which consensus mechanism? How is the boilerplate setup suitable to study the above mentioned use cases.

Current progress: The actual chapter hasn't been written yet but the implementation is well underway. The immediate next step is to harmonise the still isolated components (oracles, logic update, and the overall set of smart contracts) and do an initial system's evaluation.

1.4.1 Choice of platform

1.4.2 Choice of consensus

1.4.3 Fetching outside data: Oracle

The main questions that will be discussed are: Why do it? Why oracles? How do they work? What are the security caveats? Do we still trust the third parties eventually?

1.4.4 Logic update: Call delegation

The main questions that will be discussed are: Why do it in the first place? Why deletagate call? Caveats e.g., state maintenance considerations.

1.4.5 Lessons learned and smart contract best practices

Inspiration comes directly from how Truffle suite works[4]. It encourages one to think of implementing all the functionality (control/meta or otherwise) in the form of a set/system of smart contracts. This source can be considered for inspiration regarding the best practices related to smart contracts[5]. Under the light of these mentioned resources I will summarise and document my own best practices pertaining to implementing the set/system of smart contracts for the use cases of resource consumption in community networks and trusted data provenance. An exhaustive transaction flow graph/network analysis (somewhat akin to a state machine) of the overall system will be highly desirable. It will then provide the opportunity to model my overall system mathematically which will pave the way for further analytical study on it. This analysis will directly influence what I will have to say in the immediate next Section 1.4.6.

1.4.6 A step towards self-contained self-governing and trusted system

Room for ML and AI if so then how (this is aimed to study the possibility of the feedback loops as shown in Figure 3.1 on Page 29 in [1])? Would you have resources, expertise and time to do something with it? Given data you might just be able to do that much like how data is available on BigQuery[6],[7] and similar data analysis and feeding that as a feedback for informed decision making.

1.5 (Chapter V) System evaluation

For motivations I will be conducting similar experients as done in my recent publications here [7, 4]. Outline as many evaluation metrics as you can and then as many relations among them as you can. Highlight different areas for evaluation such as scalability and then discuss each metric individually such as Tx rate, consensus latency, choice of platform, and choice of consensus mechanism.

[4] https://truffleframework.com/tutorials/pet-shop

[5] https://github.com/ConsenSys/smart-contract-best-practices

[6] https://cloud.google.com/blog/products/gcp/bitcoin-in-bigquery-blockchain-analytics-on-public-data

[7] https://cloud.google.com/blog/products/data-analytics/ethereum-bigquery-public-dataset-smart-contra

Current progress: The actual chapter hasn't been written but I have two relevant publications [4, 7]. The actual subject matter of this chapter will be influenced by what I did in these publications.

1.5.1 Scalability

1.5.2 Sustainability

Talk about adoption.

1.6 (Chapter VI) Compliance

Current set of policies. How the concepts translate from the legal domain to blockchain's domain? *Most important: Extent.* I will be following an approach closer to how Dave describes how to make blockchain coexist with GDPR[8].

Current progress: The actual chapter hasn't been written but I have done a similar study in our recent survey article []. The text in this section will be highly influence by the subject matter of this survey.

1.7 (Chapter VII) System verification (tentative)

An ambitious step.

Current progress: The actual chapter hasn't been written. This chapter is tentative will only be dealt with if time permits. However, I will be extensively conducting tests similar to the ones described in this tutorial[9]. This will be to provide certain concrete assertions related to the logic that I program in my set of smart contracts..

[8] `https://thenextweb.com/syndication/2018/07/26/gdpr-blockchain-cryptocurrency/`
[9] `https://truffleframework.com/tutorials/pet-shop#testing-the-smart-contract`

Chapter 2

Timetable

Table 2.1: Summary of Research plan and timetable

Plan	**Output and milestones**
Third year: First quarter (Feb 19-Apr 19)	
Automation of a self-governing and self-contained resource consumption and economic sustainability model for community networks with blockchain.	Completion of Chapters I (Section 1.1) II (Section 1.2), and (some parts of) IV (Section 1.4). This will potentially result in at least a publication at a reputable venue.
Third year: Second quarter (May 19-Jul 19)	
A system's evaluation of the above automation.	Completion of Chapter V (Section 1.5) with respect to Chapter II (Section 1.2). This will also have the potential for a publication.
Second year: Third quarter (Aug 19-Oct 19)	
Implementation of trusted data audit using data provenance with blockchain.	Completion of Chapter III (Section 1.3) and Chapter V (Section 1.5) with respect to it. This will also have the potential for a publication.
End of third year: Fourth quarter (Nov 19-Jan 20)	
Study *compliance by design* extensively in the light of my research thus far.	Dissertation wrap up with special focus given to the lessons learned and exhaustive concluding remarks documented in Chapter VI (Section 1.6). Time permitting, I will also, at least in part during my PhD, work on Chapter VII (Section 1.7).

Chapter 3

Research activity and output so far

3.1 Publications and reports

1. Blockchain And The Future of the Internet: A Comprehensive Review [In submission at COMST]
2. Towards Blockchain-enabled Wireless Mesh Networks [7].
3. Blockchain for Economically Sustainable Wireless Mesh Networks [4].
4. Summer internship report at Universitat Politcnica de Catalunya (UPC) [2].

3.2 Presentations and talks

1. Presentation at Trusted System Design at Computer Lab, Cambridge[1].
2. Panel discussion at Legal Things One event in London[2].

[1] https://techroose.com/trusted-system-design.html
[2] https://bit.ly/2BOhJQC

Bibliography

[1] Ali, A. Effective governance of and by the blockchain. https://www.cl.cam.ac.uk/~aa980/papers/fyr_final_v1.pdf, January 2018. (Accessed on 16/01/2019).

[2] Ali, A. Spanish summer internship report. https://www.cl.cam.ac.uk/~aa980/papers/spanishCaseStudy.pdf, May 2018. (Accessed on 23/01/2019).

[3] Baig, R., Dalmau, L., Roca, R., Navarro, L., Freitag, F., and Sathiaseelan, A. Making community networks economically sustainable, the guifi. net experience. In *Proceedings of the 2016 workshop on Global Access to the Internet for All* (2016), ACM, pp. 31–36.

[4] Kabbinale, A. R., Dimogerontakis, E., Selimi, M., Ali, A., Navarro, L., and Sathiaseelan, A. Blockchain for economically sustainable wireless mesh networks. *arXiv preprint arXiv:1811.04078* (2018).

[5] Pasquier, T., Han, X., Goldstein, M., Moyer, T., Eyers, D., Seltzer, M., and Bacon, J. Practical whole-system provenance capture. In *Proceedings of the 2017 Symposium on Cloud Computing* (2017), ACM, pp. 405–418.

[6] Pasquier, T., Singh, J., Powles, J., Eyers, D., Seltzer, M., and Bacon, J. Data provenance to audit compliance with privacy policy in the internet of things. *Personal and Ubiquitous Computing* (2017), 1–12.

[7] Selimi, M., Kabbinale, A. R., Ali, A., Navarro, L., and Sathiaseelan, A. Towards blockchain-enabled wireless mesh networks. *arXiv preprint arXiv:1804.00561* (2018).