

# Blockchain Design for an Embedded System: Open Review

Authors: Sara Falcone,<sup>†</sup> John Zhang,<sup>‡</sup> Agnes Cameron,<sup>§</sup> Amira Abdel-Rahman<sup>††</sup>

Reviewers: Reviewer A, Reviewer B

**Abstract.** The final version of the paper “Blockchain Design for an Embedded System” can be found in Ledger Vol. 4, S1 (2019) 7-16, DOI 10.5915/LEDGER.2019.172. There were two reviewers involved in the review process, none of whom have requested to waive their anonymity at present, and are thus listed as A and B. After initial review by Reviewers A and B (1A), it was determined that the submission required revisions. The authors responded to their feedback and revised the manuscript (1B). The changes were accepted, thus completing the peer-review process. Authors’ responses are bulleted for clarity.

## 1A. Review

### Reviewer A:

This paper presents a preliminary design of a blockchain protocol for embedded systems, such as robot swarms.

The topic of the paper is relevant to the symposium, and the proposal is sufficiently well argued to raise interesting discussion.

The paper is sufficiently well written and clear. The text is longer than the limit of 6-8 pages, and the formatting is incorrect (the paper seems to have been written as a Word document rather than in Latex). In the title of Section 1.2, “lit review” is too colloquial.

This point in the paper is not well motivated:

“Many critics of blockchain target the redundancy of this data storage as an unnecessary inefficiency. However, for our application it is advantageous as all robots can use a complete map of the surface they are exploring to decide where to move and explore. Engineers are naturally interested in the full history of each agent’s locomotion as this information helps in the development of future systems.”

---

<sup>†</sup> S. Falcone (sfalcone@mit.edu) is a researcher in the Center for Bits and Atoms at the MIT Media Lab.

<sup>‡</sup> J. Zhang (johnz@mit.edu) is a student at the Massachusetts Institute of Technology.

<sup>§</sup> A. F. Cameron (agnescam@mit.edu) is a researcher at the Viral Communications Group at the MIT Media Lab.

<sup>††</sup> A. Abdel-Rahman (amira.abdel-rahman@cba.mit.edu) is a researcher in the Center for Bits and Atoms at the MIT Media Lab.

Neither the point about “full history”, neither the point about “complete map” are true in the way they are argued.

- About full history: While it is true that in the development phase of a system there is interest in maintaining a log of the past operations of a robot, it is not required that this data should be kept by the robots. It is instead better to offload such data to external storage, for later analysis on more powerful hardware.

Keeping a complete map: This depends on what “map” means, as there is an obvious trade-off between map resolution and available storage. In a space application, it is unlikely to require that every robot maintains a cm-precision map of an entire planet. Again, more realistically this could (and should) be offloaded to an external system.

Another aspect that is left untouched is the definition of Byzantine robot. Every robot will experience a certain degree of noise in the data that it senses, and noise filtering techniques might not be sufficient to readily and conclusively decide whether a certain part of a map looks in a certain way. This is common practice - maps are built over time, minimizing the noise covariance through loop closures. This blockchain concept seems to neglect completely this aspect of mapping, and it would be useful to see a discussion of it.

#### **Reviewer B:**

Well written paper with an experimental validation of the blockchain method for multiple robots. The paper would benefit from a mathematical formulation of the method and a discussion on the time outs for verifying a “transaction”.

### **1B. Authors’ Response**

#### **Reviewer A:**

This paper presents a preliminary design of a blockchain protocol for embedded systems, such as robot swarms.

The topic of the paper is relevant to the symposium, and the proposal is sufficiently well argued to raise interesting discussion.

The paper is sufficiently well written and clear. The text is longer than the limit of 6-8 pages, and the formatting is incorrect (the paper seems to have been written as a Word document rather than in Latex). In the title of Section 1.2, “lit review” is too colloquial.

- Changed titles.

This point in the paper is not well motivated:

“Many critics of blockchain target the redundancy of this data storage as an unnecessary inefficiency. However, for our application it is advantageous as all robots can use a complete map of the surface they are exploring to decide where to move and explore. Engineers are

naturally interested in the full history of each agent's locomotion as this information helps in the development of future systems.”

- Aimed to change language to adjust and focus motivation.

Neither the point about “full history”, neither the point about “complete map” are true in the way they are argued.

- About full history: While it is true that in the development phase of a system there is interest in maintaining a log of the past operations of a robot, it is not required that this data should be kept by the robots. It is instead better to offload such data to external storage, for later analysis on more powerful hardware.

- Reduced the emphasis on these points in the introduction, though they are still there.

Keeping a complete map: This depends on what “map” means, as there is an obvious trade-off between map resolution and available storage. In a space application, it is unlikely to require that every robot maintains a cm-precision map of an entire planet. Again, more realistically this could (and should) be offloaded to an external system.

- Clarified what “map” means for this experimental demonstration. Emphasized and described in better detail the lattice locomotion. Added Fig. 3. To help visualize better this system.

Another aspect that is left untouched is the definition of Byzantine robot. Every robot will experience a certain degree of noise in the data that it senses, and noise filtering techniques might not be sufficient to readily and conclusively decide whether a certain part of a map looks in a certain way. This fact is common practice - maps are built over time, minimizing the noise covariance through loop closures. This blockchain concept seems to neglect completely this aspect of mapping, and it would be useful to see a discussion of it.

- Language was changed to make our implementation clearer. As this is specifically an implementation for mapping a lattice structure, we do not run into resolution issues, though the comments sparked an interesting addition to the paper where we included the following:
- “Sample size is an important consideration for experimentation. For our proof of concept exploration, we were mainly interested in implementing a useful form of blockchain that would work with BILL-E, but the developed protocol can be utilized towards other systems. If the robot were not constricted to the lattice structure one could imagine them traversing the surface of a planet. In that case a map is much more extensive than an array of Boolean values describing if a voxel is present or not. It would be nearly impossible to record redundant topological

readings with meaningful resolution given noise and actual changes that can occur in exposed environments, though average readings with a threshold deviation could be considered PoV.

- “With PoV difficulty is not determined by incrementing a nonce, instead the amount of time robots need to map determines the cadence between block addition. Incorporating statistical metrics on the sensor readings and adjusting the threshold could also be used to control the mining difficulty.”

**Reviewer B:**

Well written paper with an experimental validation of the blockchain method for multiple robots. The paper would benefit from a mathematical formulation of the method and a discussion on the time outs for verifying a “transaction”.

- We developed the simulations used for Fig 3 and ran it hundreds of times on different geometries. We found that the state machine was sufficient to map the surface in every test performed with 100 voxels. Thus, all transactions were verified in the simulation, but we were not able to develop a mathematical formula to prove this. Though the simulations are useful to show the robot system, we wanted to focus on the embedded system, its performance and development.



Articles in this journal are licensed under a Creative Commons Attribution 4.0 License.



Ledger is published by the University Library System of the University of Pittsburgh as part of its D-Scribe Digital Publishing Program and is cosponsored by the University of Pittsburgh Press.