

Towards Application Portability on Blockchains

Kazuyuki Shudo
Tokyo Institute of Technology
Tokyo, Japan

Kenji Saito
Keio University
Kanagawa, Japan

Abstract—We pose a fundamental problem of public blockchain, “incentive mismatch.” It is an open problem, but application portability is a provisional solution to the problem. Portability is also a desirable property for an application on a private blockchain. It is not even clear to be able to define a common API for various blockchain middlewares, but it is possible to improve portability by reducing dependency on a blockchain. We present an example of such middleware designs that provide application portability and especially support migration between blockchains.

Index Terms—blockchain, incentive mismatch, application portability, migration between blockchains

I. INTRODUCTION

Incentive mismatch is a fundamental problem public blockchain has. The incentive of nodes supporting a blockchain is different from the incentives of applications on the blockchain. If possible, it is better for them to match. But it is not so trivial to align them and it is still an open problem.

Portability is another solution to protect an application against collective collapse along with its basing blockchain. Portability is also desirable for an application on a private blockchain. It would be a well-known best practice to minimize dependency of an application on underlying middleware.

However, today’s blockchain middlewares provide their own APIs and an application on a specific blockchain is hardly portable. It is not even clear to be able to form a functioning common API because each blockchain has its own abstraction such as Ethereum’s Solidity language.

In this paper, we pose the incentive mismatch problem and discuss application portability as a provisional solution to the problem. A software architecture and a technique that enable application migration between blockchains follow.

II. MOTIVATION FOR APPLICATION PORTABILITY

Application portability is a desirable property for both public blockchain and private blockchain while the reasons are not identical.

A. Public blockchain

Public blockchain has a fundamental problem we call “incentive mismatch.” Nodes supporting a blockchain and applications on it have no common incentive. Confirmations of transactions in a public blockchain are achieved by Proof of Work [1] and its derivatives such as Proof of Stake [2], that are based on economic incentive of the nodes. The nodes just try to gain coins and have no incentive to support applications themselves.




Fig. 1. Two levels of application portability.

If a public blockchain fails to provide enough economic attraction to its supporting nodes, it loses the ability to confirm transactions in a secure manner. For example, a fall in coin prices incurs defection of the supporting nodes, and the ability of secure confirmation declines. A blockchain with less supporting nodes is more vulnerable to attacks such as majority (51%) attack and eclipse attack [3], [4]. It means that an application on the blockchain have no control of the confirmation ability that the application is based on. Applications on a public blockchain can collapse by uncontrollable economic circumstances.

Is it possible to align the incentives of the nodes supporting a blockchain and applications on the blockchain? It depends on an application. The currency system supporting a public blockchain itself is an application of the blockchain because it utilizes the confirmation ability. Their incentives match by their design. But it is a special case. A public blockchain and its applications generally have no common incentive and it is an open problem how to align their incentives.

Application portability is a provisional solution to the problem. We can migrate an application to another blockchain in case its basing blockchain is collapsing if it is portable, especially migratable (Section III).

B. Private blockchain

Portability is a desirable property also for an application on a private blockchain. If the portability is low, an application suffers so-called vendor lock-in. For example, an application and its users cannot benefit from a better another middleware. At worst, an application dies along with its underlying middleware. It means we cannot execute the middleware due to its staleness.

III. APPLICATION PORTABILITY

There are several levels of application portability. We distinguish the following two levels shown in Figure 1.

- 1) “Runnable” : An application is able to run on different blockchains, but not migratable.
- 2) “Migratable” : An application and its data can migrate from a blockchain to another blockchain.

A common API to blockchains enables “*Runnable*” level of portability. And, it is obvious that an application is portable between separated blockchains but operated with the same middleware. It is uncertain to be able to form a functional common API to different middlewares because each middleware adopts its own abstraction. For example, a middleware adopts directed acyclic graph (DAG) instead of hash chain of blocks and each middleware has its own smart contract mechanism such as Ethereum’s Solidity language. But an application can be portable in case that we can limit blockchain functions the applications utilizes and provide a common API for such limited functions. All the blockchains should support storing a hashed value with its time stamp, and it is possible to provide a common API for the function.

If an application is portable at the “*runnable*” level, the application that migrates to another blockchain can only restart in its initial state. The application has to abandon its accumulated data. Because the proof of existence of data and the verification of state changes are distinctive features of blockchain, migration without logs (Section IV-A) enabling those features is hardly meaningful. The “*migratable*” level of portability enables applications to survive the death of a private blockchain middleware (Section II-B) and the collapse of the underlying public blockchain (Section II-A).

IV. APPLICATION MIGRATION BETWEEN BLOCKCHAINS

This section discusses how we achieve application migration between blockchains and shows a preliminary design of a middleware enabling migration.

A. Data to be migrated

What data should be migratable between blockchains? If an application is not related to blockchain and it migrates between underlying middlewares, it is generally enough to migrate the current state of the application. But one of the strong motivations to adopt blockchain is that it support proof of existence of a data item at a time in the past and verification of state (data) changes. In this case, it is not enough to migrate only the current state. Logs enabling the proof and the verification should also be migratable.

In summary, there are two classes of data to be migrated as follows.

- 1) *Current state* of the application
- 2) *Logs*

For example, in a currency application like Bitcoin, the latter is transactions. A transaction includes state (data) changes and time stamps, that is metadata. The former is account balances though, in Bitcoin, it is not explicitly recorded and can be calculated by adding up related transactions.

B. Middleware design enabling migration

We adopt the following policy to design a middleware enabling the migration.

- To minimize dependency on a specific blockchain middleware.

Only the requirement for blockchain middlewares is to

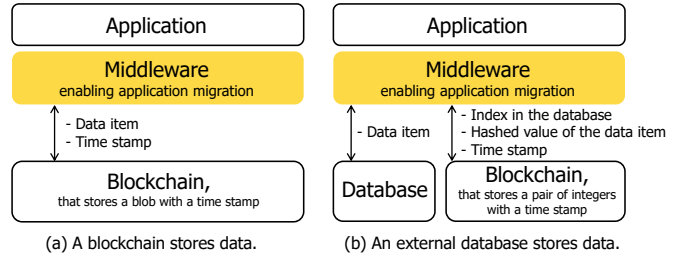


Fig. 2. A software architecture that improves application portability by limiting dependency on a blockchain.

support storing a simple data item, such as numbers or a byte sequence, with its time stamp.

- Not to expect that we can keep trusting the migration source blockchain.

One of motivations for the escape from a blockchain is possible compromise of it. For example, a public blockchain that has lost enough nodes can be compromised (Section II-A). In this case, we cannot keep relying on the blockchain.

Figure 2 shows design candidates. The “middleware” shown in the figure conducts migration between blockchains. To minimize the dependency mentioned above, the preliminary design expects an underlying blockchain to support only storing a set of numbers and a time stamp, at most, a byte sequence with a time stamp. Migration while supporting more features Today’s blockchains provide belongs to future work.

As shown in Figure 2 (b), there is a design choice to store data in an external database, not in a blockchain, though it is necessary for migration. A middleware over blockchains, BBc-1 [5] adopts such a design. This design reduces the size the blockchain. It has an advantage because a blockchain is copied to all the blockchain nodes and occupies storage on them.

In the design shown in Figure 2 (b), fault tolerance of the database has to be considered to achieve enough data availability. We can utilize replication or erasure coding [6] supported by most distributed databases to achieve high availability. We can adjust the extent of fault tolerance by configuring the number of replicas or parameters for erasure coding. The adjustability would be an advantage over simply copying to all the blockchain nodes shown in Figure 2 (a). Today such highly-available databases are provided as public cloud services and even we can record data into multiple databases instances across cloud providers.

In the design shown in Figure 2 (b), if an application does not require verification of state (data) changes, we can overwrite data in the database when they are updated. It reduces the amount of storage occupation further.

C. Migration process

If an application requires access to all the logs (Section IV-A) in the blockchains that it migrated from, the middleware has to provide the access. But we do not expect that we can keep trusting the migration source blockchains (Section IV-B) for the security reasons. Furthermore, it is safer not to expect the migration source blockchain middlewares are still running




Fig. 3. A technique to migrate an application between blockchains.

and accessible online. In any case, we cannot control a public blockchain system. It requires operating work and CPU power to keep running a blockchain middleware in case of a private blockchain.

Figure 3 shows a migration process based on these requirements. The middleware saves a migration source blockchain statically in its storage, truncates the blockchain after the expiration time, and accesses the resulted blockchain. The middleware sets a time just before the migration as the expiration time of the migration source blockchain. The granularity of the expiration time would be millisecond or block height. The static saving obviates the need for a run of blockchain middleware. By the truncation we can stop trusting the blockchain after the expiration time.

The middleware forms a hash chain of the saved blockchains as shown in Figure 3. The hash chain enables us to detect alteration of the saved blockchains and to believe the saved blockchains have not been modified. It also enables us to confirm the expiration time stored in the next blockchain is correct.

The middleware has to have an ability to interpret every migration source blockchains, that are in their own formats. Development of interpreting functions involves much work and the software is prone to be complicated and have a large number of mistakes. It is desirable to simply abandon the migration source blockchains if an application allows it. In case an application requires the proof of existence of data but not the verification of state (data) changes, we can abandon the migration source blockchain after we copy metadata to prove the existence, and data itself in case of Figure 2 (a), to the next blockchain. In this case, how do we believe the copied metadata are correct without the original metadata? Trust provided by blockchain is based on verifiable data. But, in this copying design, the copied metadata are not verifiable because they have no explicit connection to the original metadata. An application originally trust the middleware but such trust by software correctness is weaker than trust by verifiable data, for example, because the middleware can contain human mistakes. The next best solution for this problem would be verification of the copying program. It would be possible to make the copying program verifiable by saving it and its hashed value somewhere.

V. CONCLUSION

While a blockchain application also has motivation for portability as a usual application, public blockchain imposes a

particular motivation “incentive mismatch” on the application. It means that a public blockchain is driven by a different economic incentive from an application and it is not controllable by the application. Because of it, an application also collapses if its underlying blockchain collapses unless it is portable, moreover, migratable.

Such inferior survivability of blockchain applications led us to the discussion on application portability and migration. We demonstrated a middleware design that enables application migration between blockchains and more efficient design alternatives for applications with less requirements.

Standardizing activities in the field of distributed ledgers have started in bodies such as ISO. And, IETF and W3C have already had discussions. It is even challenging to design an effective common API and migration between blockchains is further away. Today a software that have gained much popularity only had an implementation at first and specifications are developed later. For instance, the first Bitcoin Improvement Proposal (BIP) was more than two years and a half later than the start of the Bitcoin network. The effective next step following this paper will be a demonstration of a portable application and a migratable application.

Portability of program code and data for smart contract is part of future work though this paper has hardly discussed it. Even on a single middleware it is difficult to update the data format and program code while keeping existing data.

ACKNOWLEDGMENTS

We thank Toshio Koide for the discussion of portability of smart contracts. This work was supported by Kaula, Inc. This work was supported by SECOM Science and Technology Foundation. This work was supported by New Energy and Industrial Technology Development Organization (NEDO). This work was supported by JSPS KAKENHI Grant Numbers 2570008 and 16K12406.

REFERENCES

- [1] Cynthia Dwork and Moni Naor. Pricing via processing or combatting junk mail. In *Proc. CRYPTO'92*, pages 139–147, August 1992.
- [2] QuantumMechanic. Proof of stake instead of proof of work. <https://bitcointalk.org/?topic=27787.0>.
- [3] Arthur Gervais, Hubert Ritzdorf, Ghassan O. Karame, and Srđjan Čapkun. Tampering with the delivery of blocks and transactions in Bitcoin. In *Proc. ACM CCS 2015*, October 2015.
- [4] Ethan Heilman, Alison Kendler, Aviv Zohar, and Sharon Goldberg. Eclipse attacks on Bitcoin’s peer-to-peer network. In *Proc. USENIX SEC'15*, pages 129–144, August 2015.
- [5] Kenji Saito and Takeshi Kubo. BBc-1 : Beyond Blockchain One – an architecture for promise-fixation device in the air –. Available electronically at <https://beyond-blockchain.org/public/bbc1-design-paper.pdf>.
- [6] Daniel Ford, François Labelle, Florentina I. Popovici, Murray Stokely, Van-Anh Truong, Luiz Barroso, Carrie Grimes, and Sean Quinlan. Availability in globally distributed storage systems. In *Proc. OSDI'10*, October 2010.