Review of Temtum Blockchain

Dr Ioannis Psaras
University College London
i.psaras@ucl.ac.uk

June 26, 2019

Abstract

This document is a technical review of "temtum", the Temporal Blockchain. The present review document discusses technical details and features of the temtum whitepaper, as well as suggestions for improvement. The views and opinions included herein DO NOT constitute investment advice and the author IS NOT to be held responsible for any decision taken by any party to invest (or not) in the project, or collaborate (or not) with the team.

Summary

Overall, temtum is a very interesting and exceptionally novel approach to a fast and scalable, blockchain transaction system. The system design is borrowing concepts from well-known operational systems, such as the Tor Network, but is advancing them in several interesting ways. The whitepaper is providing a lot of details on the structure of the system and the architecture design, as well as on the application design and the team. There are, however, a few issues that need further clarification and/or restructuring, on which I'm focusing on below.

It has to be said that it is clear that the temtum team has gone a long way to design the system, build the required components, identify applicability areas, but also implement and test the temtum blockchain system. Most of the issues brought up below can be worked on to strengthen the design and presentation of the system. It also has to be noted that the version of the whitepaper that has been reviewed is slightly outdated as the team has done substantial progress and has advanced several parts of the system design. These will be included in the updated whitepaper that the temtum team will release in the immediate future. The technical lead has clarified several issues, during a conversation we had that are not detailed in the whitepaper which was public during the review period in early June 2019.

1 Technical Issues & Improvement Suggestions

1.1 System and Architecture

For the purposes of this review, I am assuming that there is a number of "archive nodes" that store the entire blockchain and then there are normal nodes that

only confirm transactions. For structural purposes, it would be nice to have a section to describe the overall architecture in the beginning of the document.

The overall design of the system is very novel and is borrowning some fundamental concepts from the Tor Network, which is proven to be a successful and operational network for many years. That said, *it is interesting to see that the temtum team has extended those concepts to build a blockchain-based network that is already live.*

It would be interesting to show how many archive nodes will the system need as the network grows. It would also be interesting to provide some estimates for the size of the headers kept by normal nodes. I would not expect the size of the headers to be excessive, but it would be good to show this expectation graphically in the document.

1.2 NPD Document

I assume that all normal nodes need to have a copy of the NPD document (process described in page 25). *The overall design is simple, but very effective and seems to be efficient from the operational network that the temtum team has built.* I have also assumed that the NPD document needs to be synchronised among all the nodes and furthermore, that this needs to be done in very short time intervals - a requirement that might be difficult to achieve in practice and at scale. After talking with the technical leader of temtum, it turns out that the document does not need to be synchronised across the temtum network nodes. The details and justification of why this is so will be included in the next version of the whitepaper.

I would strongly suggest that the temtum team considers a slightly more complicated structure, but more efficient and secure, such as a "key-value store" structure of some form. Distributed Hash Tables (DHTs) are well-known such systems. Depending on the use-case, DHTs might face scalability issues from a point on, but it seems that this can be a good approach for the temtum case. I understand that the team had to focus on other priorities up until this point.

1.3 Transaction Fees and Incentives

It is great to see that temtum is developing a solution where transactions can be confirmed with minimal energy consumption. Such designs are certainly needed at this point in the development of blockchain networks. It is worth noting that transaction fees have been used for a variety of purposes in the past, one of them being to avoid DoS and DDoS attacks. I suggest that a note is added in the document to showcase how such behaviour is avoided in the temtum network.

In some systems, transaction fees are used as a form of incentivisation to participate and contribute in the network and are therefore collected by mining/minting nodes. *The temtum project follows a novel approach in this regard* and is borrowing concepts from the Tor Network, where users are not rewarded for contributing to the network. The interesting point here is that *temtum can be tuned to integrate transaction fees, if the community so demands.*

1.4 Distributed vs Centralised Trust

The temtum consensus algorithm assumes that there is a number of "semi-trusted authority nodes", whose task is to carry out measurements and approve nodes that can act as leaders in the temtum network. This is increasing the centralisation of the system (as temtum itself is "pre-approving" nodes).

After a discussion with the technical leader of temtum, it turns out that pre-elected nodes are only going to exist during the bootstraping of the system (which is a common strategy in deploying distributed systems) in order to avoid attacks, but will not be needed when the fully-fledged version of the system is released. More details on this will be provided in the updated version of the whitepaper.

1.5 Sharding & PIP:

The whitepaper is referring very briefly in a sharding solution under development by the temtum team. This is an important and very central part of the overall system. The reason why the team has not included more information on these issues in the early version of the whitepaper is due to IPR protection. The team has now submitted their patent applications, hence, will be able to include more technical details in the upcoming version of the whitepaper.

1.6 Competitor Analysis:

It is generally very nice to have the "Competitor Analysis" section. I would suggest that the authors provide a more quantitative comparison between temtum and the competition. I would also suggest that projects such as IOTA and Algorand can be discussed in this section. This will only improve the position of temtum against the competition.

1.7 NIST Randomness Beacon

The idea of using the external NIST Randomness Beacon as a source of randomness is very interesting and truly novel. The NIST Randomness Beacon is random, can provide unpredictability, autonomy and consistency, but at the same time any output is *public knowledge*. That is, as far as I understand, *anyone can get access to the output produced by the beacon*. It would be great to clarify whether this has any implication to the temtum system. For instance, a malicious node that knows the beacon, but also has access to the NPD can become aware of the next leader's IP address and carry out a DDoS attack. However, given that initially the system works based on "pre-approved" nodes", this is successfully prevented.

After discussing with the technical lead of the temtum project, he clarified that those issues will be addressed in the next version of the whitepaper.

Author Bio: Dr Ioannis Psaras is an EPSRC Fellow and University Lecturer (Assistant Professor) at University College London (UCL, UK). He is interested in resource management techniques for current and future networking architectures with particular focus on routing, caching and congestion control. He has recently been focusing on function-centric networks to realise distributed and decentralised edge computing, also referred to as "computing in the network". He holds a prestigious EPSRC Early Career Fellowship in the area of *"decentralised content-oriented and service-centric edge-computing architectures"*. He has been heavily involved in the effort to shift the Internet towards an Information or Content-Centric Networking environment and lately has been focusing on the application of blockchain technologies to future, decentralised Internet architectures.

Dr Psaras has extensive experience in leading cutting-edge research and managing the development and implementation of proof-of-concept prototypes. He has been collaborating and consulting for some of the biggest industries in the area of future networks. He has received *five Best Paper Awards* for his work in content-centric and mobile edge-computing networks and has attracted more than £2.5M in research funding to date from the Engineering and Physical Sciences Research Council (EPSRC, UK), the EU FP7/H2020 framework programmes and from Innovate UK. More details can be found here: http://www.ee.ucl.ac.uk/∼ipsaras/.