Comparing Layer-1 and Layer-2 Blockchain Protocols in Reputation-Based Networks

07-400 Milestone 5 Report

April 15, 2022

1 Major Changes

There have been no major changes since my last milestone report.

2 What I Have Accomplished Since My Last Meeting

Around the time of my last milestone meeting, I had been working on profiling the performance of the 32-byte public keys against the 4-byte unique identifiers for bolt specifications. Since the last milestone meeting, I have extended this analysis to bolt instance accounts. Here are the results for a test set of 10 bolt instance accounts that point to one user and one bolt specification.

Method	Single query (ms)	Double query (ms)
Filtering	116.141	36.241
Deriving PDA	49.723	121.121

As hypothesized, when using the filtering method to query the accounts, we see that as the amount of account data we are storing goes down, the performance improves. Specifically, we see around a three times speedup for a eight times reduction in account data (each account public key is now stored using 4 bytes instead of 32 bytes).
The difference in our results from this week and last week is that the public key of bolt instance accounts can also be derived from the user and bolt specification account public keys. The second row displays the results for using this method with the two setups. Although the cheapest operation overall is still to use double querying with the filtering method, further testing needs to be conducted in order to ensure these results stay the same as we scale the number of accounts in our ledger.

3 Meeting Milestones

Since it is unlikely that the other groups will be finished with their implementations in time for the end of the semester, I have stopped working towards the common interface of the ledger prototype and spent more time trying to profile and improve my existing codebase, as described above.

4 Surprises

One nice surprise was finding out that reducing the amount of pointer data in our accounts not only reduces the cost of keeping the account on the network but also improves our query times!

5 Looking Ahead

It is unlikely that the various teams in our research group will finish working on their respective implementations before the end of the semester in order for me to conduct a comparative analysis against various Layer-1 and Layer-2 solutions for my final report. If this is the case, I will compare performance of my implementation to the existing BoLT ledger implementation, which is on the HyperLedger network.

6 Revisions to Future Milestones

I am revising my next milestone as follows:

- **Milestone 6** (4/18): Test unique identifiers for scalability (i.e. compare performance on 1000+ accounts) and reach a decision about which method the implementation should use to store account public key information.
- **Final report** (5/9): Collect and compare measurements for metrics such as throughput, cost, and latency by running experiments against existing BoLT ledger.

7 Resources Needed

No additional resources are needed at this time.