Domus Tower Blockchain (DRAFT) March 28, 2016

Rhett Creighton[1]

Domus Tower Inc. San Francisco CA, USA
Patent Pending
everett@domustower.com

Abstract. The purpose of this work is to demonstrate a fast, efficient, highly scalable blockchain. Current blockchain implementations have performance limitations that make them unsuitable for high speed record keeping. For example, the Bitcoin blockchain has a global maximum sustained transaction rate of 7tps [1][2]. Other blockchain implementations have maximum transaction rates in the hundreds or possibly thousands of transactions per second, but do not scale beyond that. In this paper, we describe the Domus Tower Blockchain, which has been benchmarked at ingesting over 1 million tranactions per second on hardware costing less than $50 per hour on Amazon's Web Services with the potential to scale to greater than 10 million transactions per second.

Keywords: blockchain, Merkle graph, big data

1 Background: Real Time Gross Settlement of U.S. Equities

In the current U.S. equities industry, settlement of a trade occurs 3 business days after the execution date of a trade. This is known as T+3. There is an effort to move the industry towards a 2 day settlement period (T+2)[3].

A big part of the reason for this time delay is that trade report errors can and do happen. Trade orders are not cryptographically signed at order creation, so parties are given an opportunity to dispute any trade order claiming that they "don't know" a trade (commonly known in the industry as a DKed trade)[4].

In order to achieve real time gross settlement, we believe that it is necessary to record transactions in an immutable ledger where each transaction is signed by each participant with a cryptographic signature that is impossible to forge. One solution is to build the ledger as a data structure commonly known as a blockchain.

A technical problem with using a blockchain for this application was that, until now, blockchain implementations did not scale well for high transaction rates. Supporting a high rate of transactions is necessary as U.S. stock exchanges intend to support a peak of over 1 Million trades per second now and more in the future.

Therefore, the aim of this research was to design a blockchain capable of high transaction rates which could be used to settle U.S. equities in real time.

2 What is a Blockchain?

Satoshi first described a "block of items" on page 2 of his whitepaper[5]

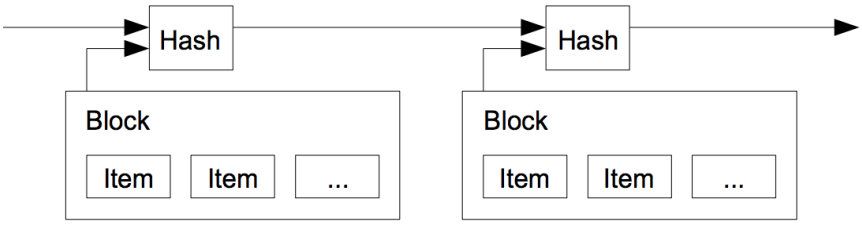

Fig. 1. Satoshi's timestamp server. Items are inserted into blocks. Each block is hashed along with the hash of a previous block. This creates the latest hash, or Merkle root for the blockchain.

This fits our basic definition of a blockchain. Additionally, we require that all transactions in the Domus Blockchain are digitally signed. The basic properties of the Domus Blockchain are:

1. Data storage is contained in a *Merkle directional acyclic graph* (Merkle-DAG), a special case of a Merkle tree, described in: US4309569[6]. We refer to nodes on this graph as "blocks".
2. Data transmitted to the blockchain is signed with a cryptographic digital signature and the signature is verified before the data is written to a block.
3. The latest hash of the Domus Blockchain, known as the "Merkle root", guarantees an immutable history of all signed data stored in all the prior blocks of the entire Merkle graph.

Unlike other blockchains, such as the Bitcoin blockchain, the Domus Blockchain does not include Proof of Work, Mining, Proof of Stake, or any form of Byzantine Fault Tolerance. Bitcoin's block headers are an example of a *dynamic-membership multi-party signature* (or DMMS)[7]. Being a DMMS means that anyone can join the network with no enrollment process.

Conversely, the Domus Blockchain is designed to operate in an environment where participants know each other and independently decide who to trust.

	Max Transaction Rate	Tracks Assets and Liabilities?	Trust Model
Bitcoin	~7 tps	No	Proof of Work
Ripple	~few thousand tps	Yes	Federated Voting
Domus	10M+ (scalable)	Yes	Trusted

Fig. 2. Domus Blockchain compared to other crypto-ledgers.

3 Architecture Overview

We use a set of microservices to work together in an efficient, scalable architecture. Each microservice performs a specific task. This ensures that microservices can be optimized independent of each other. Microservices also operate on simple data-driven interfaces. The major components include a "signature verifier", "transaction batcher", "block maker", and a "client". The components are described in more detail in later sections of this paper.

4 Consensus

Consensus among agents on a network is a fundamental problem in distributed computing. The well-known CAP Theorem states that it is not possible to guarantee consistency and availability across all nodes.

The consensus problem becomes more difficult when actors are anonymous and potentially malicious, incentivized to cheat or lie about transactions. This has been phrased as the well-known Byzantine Generals Problem, for which Bitcoin's Proof of Work provides a clever solution.

The Domus Blockchain achieves high scalability and availability by adopting a permissioned, trusted-node architecture with eventual-consistency. Our model does not use a vote-based approach to writing transaction data. Any agent that has write access to a blockchain has 100% authority to write a transactions to that chain. Authority is centralized under this model.

5 Why Not Use A SQL Database?

By adopting a permissioned, centralized, trusted-node blockchain, one might ask why not use a SQL database, which is also permissioned and centralized. The advantages of using a blockchain are:

1. A blockchain provides an immutable, permanent record of timestamped transactions. All transactions are cryptographically audited in real time with a Merkle root.
2. Every transaction is signed in realtime with a public/private key signature that is impossible to forge.
3. The append-only, timestamped structure of a blockchain actually makes it much easier to scale under a model of eventual-consistency.

In summary, we believe that tracking asset ownership in a corruptible, mutable system that needs to be carefully, manually audited on a frequent basis is a fragile method of record keeping and prone to error.

6 Trusted-Nodes

The Domus Blockchain uses a simple, ad hoc trusted node approach, similar to the GIT version control p2p system of trust. This is not different from how the securities industry works today where parties know each other and trust data feeds.

Disputes can be resolved outside the blockchain. For example, if a customer thinks his account balance is incorrect, he might call his broker on the phone and resolve the issue.

Because the Domus Blockchain is immutable, resolutions to errors are appended to the blockchain. Take, for example, the well-known double-spend attack:

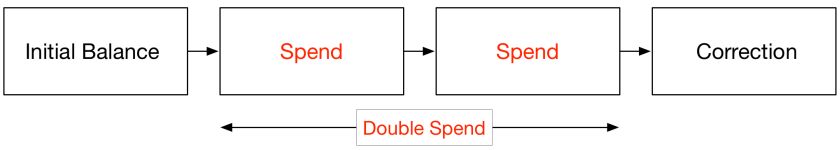

Fig. 3. In the Domus Blockchain, though unlikely to ocur, double spends are theoretically possible for a short time. Once detected, a correction is applied in the form of an appended item to the blockchain.

Because double spend records are not desired, the blockchain application layer can simply detect and filter out potential double spend transactions before

applying them to the blockchain. This will, however, introduce some additional latency in block creation.

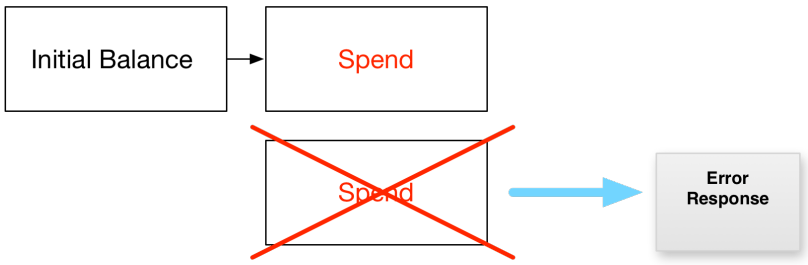

Fig. 4. In this example, the application layer detects a potential double-spend transaction and returns an error before adding it to the blockchain.

7 Signature Verification

Verifying signatures on transactions is a computationally expensive process. However, the problem of verifying many signatures is also Embarrassingly Parallel, meaning that the jobs can easily be computed in parallel across many computers.

Because signature verification is one of the most expensive pieces in this process, we choose to use a high speed public key signature system such as Ed25519[8].

A signature verification process simply takes a message with a signature and public key as an input. If the signature does not match the message and the public key, the process handles the input as an error. If the signature is verified, the process handles the input as verified, and passes it along to be written to the blockchain.

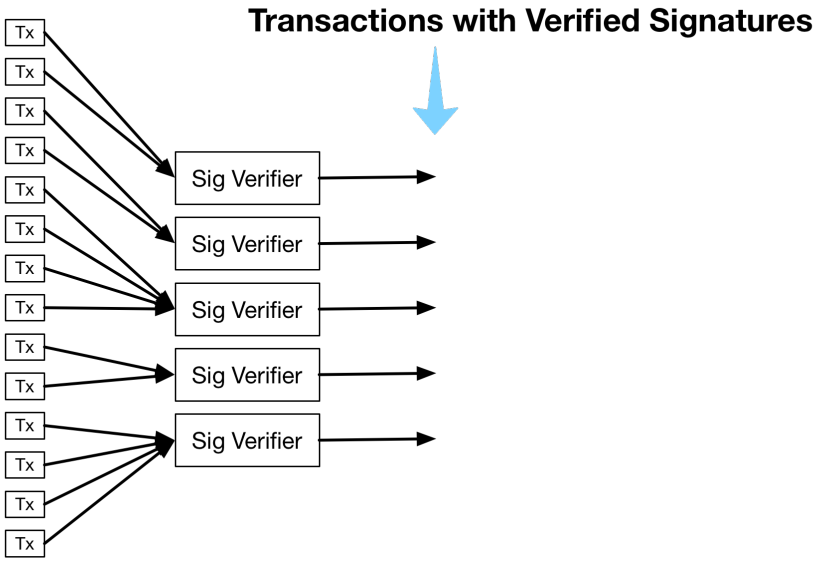

Fig. 5. An overview of the Domus signature verification microservices. If the signature is valid on a transaction, the signature verification service passes the transaction on to be processed into a block.

8 Transaction Batching

Verified signed messages are passed to transaction batchers. Transaction batchers can perform arbitrary operations on batches of input before passing it to a block maker.

One example is that a transaction batcher can compress transactions after every X transactions, or after every Y milliseconds have passed. Our current implementation ingests 10,000 transactions over socket streams and then compresses those transactions using gzip.

A single compressed message is posted to a block making service as an http file upload. The gzip file contains the 10,000 batched and signed transactions.

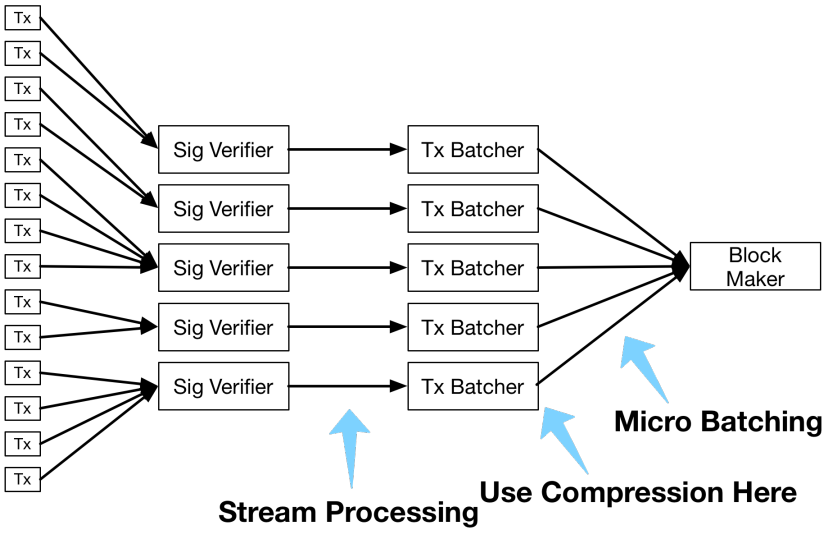

Fig. 6. Transaction batcher services receive verified transactions and batch them before sending them to a block maker service.

9 Block Making

A block maker takes as input compressed batches of transactions with verified signatures. It stores each piece of input in a linked-list that represents the data of the current block being written.

A block represents data written in a specific window of time. It must have been written after the previous block, and before the subsequent block in a blockchain.

In our example, a new block is created for each second. We use a Key-Value Storage system, where the key is used to designate the block time window, and the value is the linked-list of compressed verified transaction data.

The Merkle root of a block is generated by running a cryptographic hash function on the data in the current block concatenated with the previous block's Merkle root. The sole exception is the first block in a blockchain, which has no previous block's Merkle root. In this case, an empty string can be used as the previous Merkle root, or any static value can also be used.

In order to accommodate a high transaction rate, blockchain transactions are recorded using volatile high speed memory, such as RAM. Data persistence is achieved by periodically flushing the in-memory working set of data to a permanent storage device. In our example, fsync is used to perform this operation once per second.

10 Client

A client uses a public/private keypair to digitally sign messages that are sent to a blockchain building service.

11 Benchmarks

Each server is a 32-core c4.8xlarge on Amazon at $1.675 per hour. All servers are in the same availability zone and placement group.

11.1 Setup 1: 8 Servers

- 1x Redis Server (where the blockchain is stored)
- 1X block maker
- 2X Transaction Batchers
- 2X Signature Verifiers
- 2X Client Servers

Max-Transaction-Size is set to 512 Bytes.

Findings

- Total ingest rate 90k-100k tx per second sustained.
- The Redis Server and block maker can handle about 100x the load.
- We are CPU limited on the signature verifiers.
- We are bandwidth limited on transaction batchers.

11.2 Setup 2: 26 Servers

- 1x block maker & Redis Server (where the blockchain is stored)
- 25X Transaction Batcher & Signature Verifier & Client (all 3 processes are run on the same server)

Max-Transaction-Size is set to 256 Bytes.

Findings

- Total ingest rate 1,240,000 tx per second sustained.
- We are CPU limited on the signature verifiers.

12 Double-Entry Accounting

The Domus Blockchain supports double-entry accounting. Under this system, both assets and liabilities are tracked.

This is in contrast to the Bitcoin blockchain where only assets are tracked, even though transactions are recorded in the form of inputs and outputs.

Date	From	To	Amount
March 1	Block Reward	A	25 BTC
March 1	A	B	5 BTC
March 1	B	C	2 BTC

Fig. 7. Single-entry accounting Bitcoin example where only assets are counted and liabilities do not exist.

The advantages of double-entry bookkeeping are numerous and are beyond the scope of this paper. In a simple example, a stock broker might hold stock for the benefit of his clients. The clients' accounts will show a positive balance while the broker's account will show a negative balance. The negative balance reflects the total number of shares that are owed on this balance sheet.

Date	Account	Debits	Credits	Balance
March 1	A		+100 MSFT	+100 MSFT
March 1	C	-100 MSFT		-100 MSFT
March 1	B		+300 MSFT	+300 MSFT
March 1	C	- 300 MSFT		-400 MSFT

Fig. 8. Double-entry accounting on the Domus Blockchain tracking assets and liablities.

In the above example, we are tracking credits, debits, and running balances. In the Bitcoin blockchain, for example, running balances are not tracked. This means that wallets need to either download the entire blockchain, or use an SPV thin wallet[9] just to know their own account balance.

13 Linked Blockchains

Linked blockchains can be created where the assets of an account on one blockchain must match the liabilities on the account of another blockchain. Under this system, a parent chain is used to guarantee the solvency for account holders on a sub-blockchain.

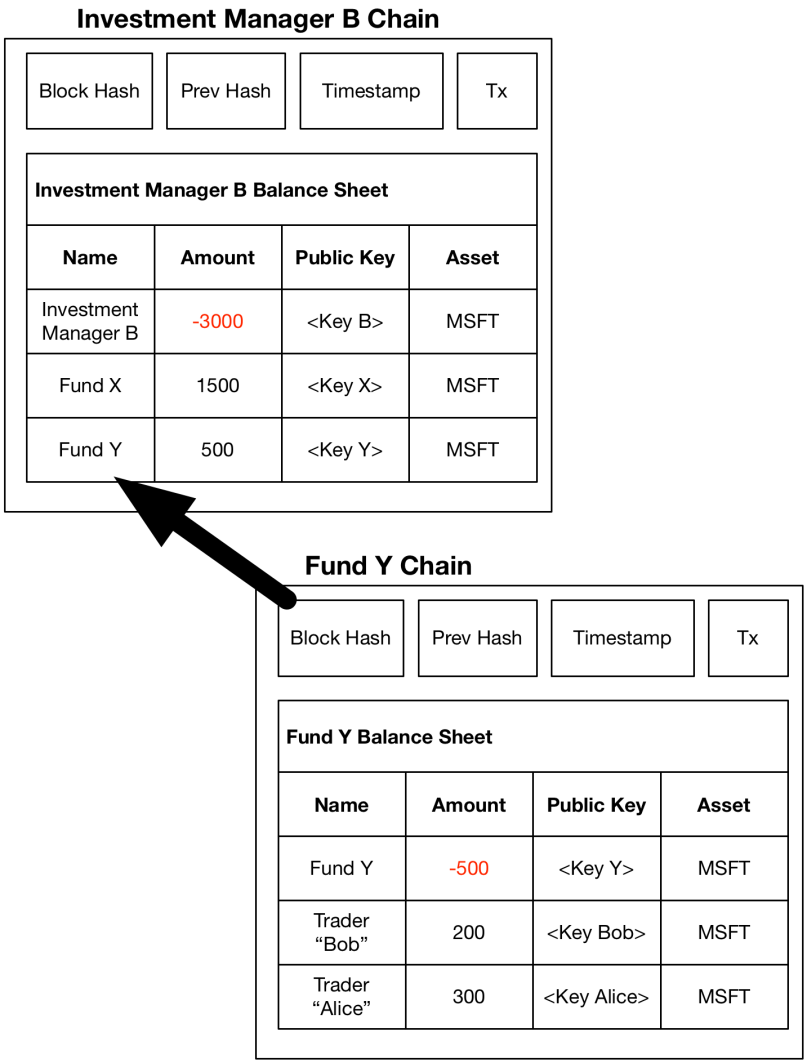

Fig. 9. An example of a linked blockchain. Fund Y has 500 MSFT on Investment Manager B's blockchain. A block on Fund Y's blockchain links to this positive balance. Fund Y is holding stock FBO Alice and Bob on its own blockchain.

14 Scaling Beyond a Block Maker

Because a Block Maker must process a hash function on a set of data, block creation is a difficult problem to parallelize. We recall that blockchains are usually implemented as a single long chain:

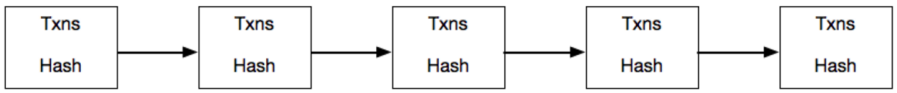

Fig. 10. A basic illustration of a blockchain. Each block contains a set of transactions and a Merkle root hash.

In order to scale, transactions can be sharded into many blocks. After transactions are sharded into blocks, a single block (metadata block) can be created from just the hash of the transaction blocks. Since much less data is contained in the metadata block, metadata blocks can be generated faster than transaction blocks.

This is similar to eDonkey2000 file hash system[10], which splits all files into roughly 9MB parts and hashed each part.

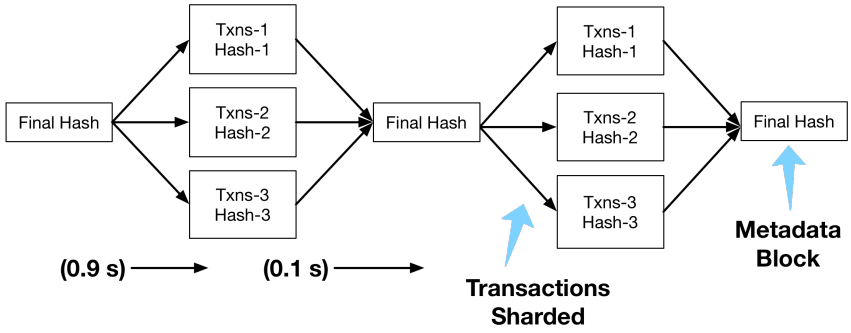

Fig. 11. A scalable Domus Blockchain graph. The chain splits into many branches, and then contracts back to a single branch in a heartbeat manner.

15 Conclusion

With the Domus Blockchain, we have demonstrated a blockchain capable of recording a high rate of transactions in a scalable manner. The Domus Blockchain records a double-entry balance sheet that tracks credits and debits.

Using a blockchain ensures that data is immutable and append-only. Additionally, a Merkle root ensures that the entire blockchain has not been altered or corrupted.

The system uses a hybrid approach in big data architecture. Rather than adopting either a streaming or a batching framework, we designed a custom framework to use each technique where it is most appropriate. This allows Domus to optimize its service for desired low-latency and high throughput.

16 Acknowledgements

Thank you for edits from Kevin Woods, Joe Forster, and Roy Epstein. This work was funded by Domus Tower, Inc.

References

1. *Bitfury Group*. Block Size Increase
   http://bitfury.com/content/5-white-papers-research/block-size-1.1.1.pdf
2. https://en.bitcoin.it/wiki/Scalability
3. http://www.ust2.com/
4. http://www.nasdaq.com/investing/glossary/d/dont-know
5. https://bitcoin.org/bitcoin.pdf
6. Merkle, US 4,309,569, Method of Providing Digital Signatures
7. Enabling Blockchain Innovations with Pegged Sidechains
   https://blockstream.com/sidechains.pdf
8. http://ed25519.cr.yp.to/
9. https://github.com/bitcoin/bips/blob/master/bip-0037.mediawiki
10. https://en.wikipedia.org/wiki/EDonkey2000