



BRNO UNIVERSITY OF TECHNOLOGY

VYSOKÉ UČENÍ TECHNICKÉ V BRNĚ

FACULTY OF INFORMATION TECHNOLOGY

FAKULTA INFORMAČNÍCH TECHNOLOGIÍ

DEPARTMENT OF INFORMATION SYSTEMS

ÚSTAV INFORMAČNÍCH SYSTÉMŮ

**PROCESSING OF THE BLOCKCHAIN EMPLOYING
IPFS**

VYUŽITÍ IPFS PRO ZPRACOVÁNÍ BLOCKCHAINU

MASTER'S THESIS

DIPLOMOVÁ PRÁCE

AUTHOR

AUTOR PRÁCE

Bc. MATÚŠ MÚČKA

SUPERVISOR

VEDOUCÍ PRÁCE

Ing. VLADIMÍR VESELÝ, Ph.D.

BRNO 2020

Master's Thesis Specification



Student: **Můčka Matůš, Bc.**
Programme: Information Technology Field of study: Information Systems
Title: **Processing of the Blockchain Employing IPFS**
Category: Networking

Assignment:

1. Learn about cryptocurrencies and other blockchain technologies (namely, Ethereum, Bitcoin, DigiByte, Decred, Monero).
2. Study the Interplanetary File System, the principles of network communication, distribution and content addressing.
3. Design a prototype of the IPFS connector for generic blockchain access in compliance with the supervisor's recommendations.
4. Implement a prototype that delivers demanded functionality on a particular blockchain of considerable size. Write the API for integration with other applications.
5. Perform validation testing and measure prototype performance characteristics. Discuss possible extensions.

Recommended literature:

- Narayanan, A., Bonneau, J., Felten, E., Miller, A., & Goldfeder, S. (2016). *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton University Press.
- Bitpay, *Guides - Bitcore*, [online] <https://bitcore.io/guides>, [2018-10-19].

Requirements for the semestral defence:

- Tasks 1, 2 and 3.

Detailed formal requirements can be found at <https://www.fit.vut.cz/study/theses/>

Supervisor: **Veselý Vladimír, Ing., Ph.D.**
Head of Department: Kolář Dušan, doc. Dr. Ing.
Beginning of work: November 1, 2019
Submission deadline: June 3, 2020
Approval date: October 22, 2019

Abstract

This work aims to design a platform for processing blockchain data of selected cryptocurrencies for further exploring using IPFS. In this thesis, we created a proprietary decentralized and distributed database system that supports advanced queries. The solution provides a well-arranged graphical user interface for data visualization as well as API, which makes it easy to connect to other applications. The benefit of this work is a new view of blockchain processing which opens up new possibilities in its exploring.

Abstrakt

Cieľom tejto práce je navrhnúť platformu na spracovanie a preskúvanie blockchainu vybraných kryptomien pri použití IPFS. Na riešenie tohoto problému bolo potrebné navrhnúť vlastný decentralizovaný a distribuovaný databázový systém, ktorý podporuje pokročilé dotazy. Vytvorené riešenie poskytuje prehľadné grafické užívateľské rozhranie, ktoré slúži na vizualizáciu dát a taktiež API, vďaka ktorému sa dá systém jednoducho napojiť na iné aplikácie. Prínosom tejto práce je nový pohľad na spracovanie blockchainu čo otvára nové možnosti v jeho prehladávaní.

Keywords

IPFS, decentralization, blockchain, cryptocurrency

Kľúčové slová

IPFS, decentralizácia, blockchain, cryptocurrency

Reference

MÚČKA, Matúš. *Processing of the Blockchain Employing IPFS*. Brno, 2020. Master's thesis. Brno University of Technology, Faculty of Information Technology. Supervisor Ing. Vladimír Veselý, Ph.D.

Rozšírený abstrakt

Kryptomeny sa stali v uplynulých rokoch veľmi populárnymi. Medzi dôvody ich popularity patrí anonymita, prístupnosť, rýchlosť prenosu a možnosť obísť existujúce finančné inštitúcie.

Kryptomeny sú postavené na technológii nazývanej blockchain, ktorá bola predstavená svetu v roku 2008 a vytvorená osobou alebo skupinou ľudí pod pseudonymom Satoshi Nakamoto. Blockchain funguje ako verejná účtovná kniha pre kryptomeny. Sú v ňom zaznamenané všetky transakcie, ktoré v danej mene prebehli. Prechádzaním blockchainu je možné skúmať jednotlivé transakcie, vypočítať aký má adresa zostatok, kedy je adresa najviac aktívna a tým získať dôležité metadáta, ktoré neskôr môžu poslúžiť napríklad k odhalienu totožnosti vlastníka adresy, alebo jeho zaradenie do určitej kategórie.

Pre uľahčenie prechádzania blockchainu vznikli blockchain prieskumníky. Tie poskytujú pre užívateľa grafické užívateľské rozhranie alebo API, ktoré uľahčuje užívateľom preskúvanie blockchainu. Nevýhodou klasických blockchain prehliadačov je, že používajú NoSql databázové systémy, v ktorých nie je možné vykonávať nad dátami zložité dotazy.

Výsledkom tejto diplomovej práce je blockchain prieskumník postavený na decentralizovanom súborovom systéme IPFS, ktorý súboru adresuje na základe ich obsahu. Jeho výhody spočívajú v tom, že obsah databázy, v ktorej je uložený blockchain, je decentralizovaný a zároveň náš vlastný databázový systém podporuje pokročilé databázové dotazy. Celý systém sa skladá z niekoľkých častí.

Hlavná časť je modul Explorer-core, v ktorom sme implementovali decentralizovaný a distribuovaný databázový systém. Databáza používa pre indexovanie dát nad vlastnosťami objektov štruktúru B-stromu, ktorá je optimalizovaná na uloženie do IPFS. Explorer-core taktiež poskytuje bohatý dotazovací jazyk na vytváranie databázových dotazov. Databázový dotaz môže pozostávať z viacerých podmienok, medzi ktorými sú logické spojky konjunkcia a disjunkcia. Databázový systém si následne zvolí ktoré indexy na spracovanie dotazu použije. Databázový dotaz taktiež môže obsahovať informácie koľko výsledkov preskočiť (offset) a koľko vrátiť (limit). Vďaka voľbe programovacieho prostredia (TypeScript), Explorer-core funguje vo webovom prehliadači a tiež ako natívna aplikácia interpretovaná pomocou Node.js. V závislosti v akom prostredí je Explorer-core spustený, komunikuje s ostatnými prvkami systému prostredníctvom TCP socketov, technológie WebSocket alebo WebRTC.

Ďalšia časť systému je terminálová aplikácia Feeder, ktorá je zodpovedná za synchronizáciu dát z blockchainu do IPFS. Môže byť pripojená k viacerým zdrojom dát z blockchainu ako napríklad Blockbook, InsightAPI alebo priamo k blockchainovému uzlu. Feeder dáta spracuje a uloží do nášho databázového systému. Dáta su potom dostupné iným systémovým aplikáciám ako sú ExplorerAPI a ExplorerGUI.

ExplorerAPI je aplikácia ktorá poskytuje jednoduché aplikačné rozhranie na získavanie dát pomocou protokolu HTTP. Aplikácia ExplorerAPI môže byť použitá na integráciu s inými aplikáciami, alebo na tvorbu grafických užívateľských rozhraní pre staršie zariadenia, na ktorých nie je možné rozbehnúť IPFS uzol. ExplorerAPI poskytuje tri adresy, na ktorých je možné získať dáta (pre blok, transakciu a adresu). Obsah žiadostí odoslaných na ExplorerAPI je inšpirovaný dotazovacím jazykom GraphQL.

Prezentačná vrstva s grafickým užívateľským rozhraním sa nazýva ExplorerGUI a je implementovaná ako jednostránková aplikácia (vďaka tomu sa pri každom načítaní stránky neresetuje pripojenie do IPFS siete). ExplorerGUI tvorí niekoľko separátnych obrazoviek (domovská obrazovka s prehľadom pripojených kryptomien, list blokov, detail bloku, transakcia, adresa) a poskytuje moderné a prehľadné grafické rozhranie pre používateľa.

Každá významná komponenta systému bola otestovaná pomocou unit testov. Výkonosť systému bola porovnaná s inými databázovými systémami ako je MySQL a PostgreSQL. Náš databázový systém sa s pribúdajúcim počtom používateľov zrýchľoval, pretože sa obsah databázy lepšie distribuoval, zatiaľ čo rýchlosť klasických databázových systémov sa narastajúcim počtom klientov spomaľovala, pretože server bol preťažený. Pri určitom počte užívateľov dokázal náš decentralizovaný databázový systém spracovať viac dotazov za sekundu ako klasické databázové systémy.

Výsledná platforma dokáže spracovať blockchain kryptomien a poskytuje decentralizované a distribuované riešenie na jeho preskúvanie. Oproti klasickým blockchain prieskumníkom má naša platforma len minimálne náklady na prevádzku a tiež je schopná zotaviť sa po výpadku časti siete.

Processing of the Blockchain Employing IPFS

Declaration

I declare that this master's thesis was prepared as an original author's work under the supervision of Mr. Ing. Vladimír Veselý, PhD. All the relevant information sources, which were used during the preparation of this thesis, are cited and included in the list of references.

.....
Matúš Múčka
June 3, 2020

Acknowledgements

Rád by som poďakoval mojej priateľke za podporu pri písaní tejto práce a za rady spojené typografiou. Ďalej chcem poďakovať môjmu otcovi za odbornú diskusiu pri orientácii sa v databázových systémov. Môjmu bratovi ďakujem za vzájomnú podporu a motiváciu počas štúdia. Tiež ďakujem vedúcemu tejto práce za príležitosti, ktoré sa mi počas štúdia naskytli. V neposlednej rade by som chcel poďakovať celej svojej rodine a priateľom za podporu počas celého štúdia.

Teraz nasleduje recept na výborné pálivé veganské kari. Najpr je potrebné zaobstarat si červenú kari pastu. Zvyčajne býva dostať v obchodoch s orientálnym sortimentom. Na rozohriatej pancivci s olejom opražíme kari pastu spolu s tofu nakrájaným na kocky. Po chvíľe smaženia zmes zalejeme mliekom z dužiny kokosu. Ďalej pridáme rôzne druhy zeleniny, ktoré doma máme (brokolica, karfiol, bambusové výhonky, mrkva a pod.). Varíme kým zelenina nezmäkne. Podávame s ryžou.

Contents

1	Introduction	3
2	Cryptocurrencies	4
2.1	Bitcoin	5
2.2	DigiByte	5
2.3	Ethereum	6
2.4	Decred	6
2.5	Monero	6
2.6	Analysis of Current Blockchain Explorers	7
2.6.1	Blockbook	7
3	IPFS	8
3.1	IPFS Stack	8
3.2	Libp2p	9
3.3	IPLD	9
3.4	IPFS	9
3.4.1	IPFS Core APIs	10
3.5	IPNS	11
3.6	Filecoin	12
3.7	Existing Blockchain Explorers in IPFS	12
4	Design	13
4.1	Database System	13
4.1.1	OrbitDB	13
4.1.2	Textile	13
4.1.3	Database System Design	14
4.2	Platform Components	18
4.2.1	Database Design	18
4.2.2	Feeder	19
4.2.3	Explorer	20
5	Implementation	23
5.1	Explorer-core Implementation	23
5.1.1	IPFS Connector	24
5.1.2	Indexes	24
5.1.3	Query System	25
5.1.4	Database	27
5.2	Feeder	30

5.2.1	ExplorerGUI	31
5.2.2	ExplorerAPI	33
6	Testing and Benchmarking	35
6.1	Testing	35
6.1.1	Unit Testing	35
6.2	Benchmarking	35
6.3	Comparing with Blockchain Explorers	39
7	Conclusion	41
	Bibliography	42
A	Contents of the included storage media	44

Chapter 1

Introduction

HTTP is “good enough” for the most use cases of distributing files over the network. However, when we want to stream lots of data to multiple connected clients at once, we start to hit its limits. When two clients are requesting the same data, there is no mechanism in HTTP that would allow sending the data only once. Sending duplicate data has become a problem in large companies because of bandwidth capacity. Blizzard¹ started to distribute video game content via a distributed solution because it was cheaper for the company and faster for players [13]. Linux’s distributions use BitTorrent to transmit disk images².

The Bitcoin blockchain has now (i.e., May 2020) 242 gigabytes³. When blockchain is processed (i.e., all its data are parsed), the size can grow twice as much. If there are multiple blockchains, then data can have few terabytes. When we are sharing blockchains data from the server for several clients, there is a big chance that multiple clients want the same data. They may be working on the same case and investigating the same wallets. So in standard solution with relational database and some HTTP server, the server has to search in all data (that can have a size of few terabytes) and transmits selected data to the client for every request. This problem happens even if a different client asks for the same data after a few moments/minutes. The behaviour mentioned above dramatically limits the scalability of the server.

Services that are using HTTP, often have client-server architecture, so there is also a problem with a single point of failure. If the server for some reason stops working, the client can not receive data. In a distributed file system such as IPFS, there is no such problem as a single point of failure because all data are duplicated on multiple clients.

This master thesis is divided into eight chapters. Chapter 2 describes the differences between cryptocurrencies used in this project. Chapter 3 describes IPFS and all its layers. The design of the platform created in this thesis is in Chapter 4. The implementation of all applications that are provided with this project is in chapter 5. Finally, the summarization of results achieved in this work is in chapter 7.

¹Computer games company – <https://www.blizzard.com/>

²Image of Debian downloadable by BitTorrent <https://www.debian.org/CD/torrent-cd/>

³A current size of bitcoin blockchain can be seen at <https://www.statista.com/statistics/647523/worldwide-bitcoin-blockchain-size/>

Chapter 2

Cryptocurrencies

There were hundreds of failed attempts of creating cryptographic payment systems before cryptocurrencies like Bitcoin and Ethereum came into existence. Some of these systems are listed in Figure 2.1. All of them were created before Bitcoin. Despite that, only a few of them survived to these days. Some of these attempts were only academic proposals while others were deployed and tested systems. One of the survival is PayPal. It is only because it quickly gave up its original idea of hand-held devices for cryptographic payments. [26]

So there is a question, what makes cryptocurrencies successful nowadays? It may be easy to use principle and no need for external hardware. Another critical component of cryptocurrencies discussed in this work is blockchain. Generally, it is a ledger in which all transactions are securely stored. The idea behind blockchains is pretty old, and it was initially used for timestamping digital documents. [11]

ACC	CyberCents	iKP	MPTP	Proton
Agora	CyberCoin	IMB-MP	Net900	Redi-Charge
AIMP	CyberGold	InterCoin	NetBill	S/PAY
Allopass	DigiGold	Ipin	NetCard	Sandia Lab E-Cash
b-money	Digital Silk Road	Javien	NetCash	Secure Courier
BankNet	e-Comm	Karma	NetCheque	Semopo
Bitbit	E-Gold	LotteryTickets	NetFare	SET
Bitgold	Ecash	Lucre	No3rd	SET2Go
Bitpass	eCharge	MagicMoney	One Click Charge	SubScrip
C-SET	eCoin	Mandate	PayMe	Trivnet
CAFÉ	Edd	MicroMint	PayNet	TUB
CheckFree	eVend	Micromoney	PayPal	Twitpay
ClickandBuy	First Virtual	MilliCent	PaySafeCard	VeriFone
ClickShare	FSTC Electronic Check	Mini-Pay	PayTrust	VisaCash
CommerceNet	Geldkarte	Minitix	PayWord	Wallie
CommercePOINT	Globe Left	MobileMoney	Peppercoin	Way2Pay
CommerceSTAGE	Hashcash	Mojo	PhoneTicks	WorldPay
Cybank	HINDE	Mollie	Playspan	X-Pay
CyberCash	iBill	Mondex	Polling	

Figure 2.1: Electronic payment systems before cryptocurrencies [20]

of data. One block can contain approximately 3,500 transactions given the 2 MB limit. This restriction means that in DigiByte, 230 transactions can be confirmed in one second compared to Bitcoin 4-7 transaction per second. DigiByte also has 1,000:1 DigiByte to Bitcoin ratio, so for every Bitcoin, there is 1,000 DigiByte. [5]

2.3 Ethereum

Ethereum, rather than a cryptocurrency, is a blockchain application platform with Turing-complete programming language. While both the Bitcoin and Ethereum networks are powered by the principle of distributed ledgers and cryptography, the two differ technically in many ways. For example, transactions on the Ethereum network may contain executable code, while data affixed to Bitcoin network transactions are generally only for keeping notes. Other differences include block time (an ether transaction is confirmed in seconds compared to minutes for bitcoin) and the algorithms that they run on (Ethereum uses ethash while Bitcoin uses SHA-256). [27, 2]

2.4 Decred

Decred is cryptocurrency build from Bitcoin. The main difference from Bitcoin is the rewarding system from mining. In Bitcoin, the miner gets full reward for a mined block. Sometimes, the Bitcoin blockchain splits when two or more miners found a block at nearly the same time. The fork is resolved when the subsequent block(s) are added, and one of the chains becomes longer than the alternative(s). In Decred the chance of blockchain forks is minimized by hybrid proof of work and proof of state system. Each time a block is created (by a miner), it is not automatically part of the blockchain. Block needs to be approved by ticket holders. Then miner receives a block reward (newly created DCR). If the block is rejected by ticket holders, the miner does not receive a reward. Tickets holders for a new block are chosen randomly. The ticket validates the previous block. A block needs at least three of the five votes chosen to approve it for it to be validated. This hybrid system has many implications, including making a 51% power attack very difficult and making a minority fork very difficult as well. [12]

2.5 Monero

Three years after Bitcoin, in 2012, the competing Bytecoin cryptocurrency entered the market. The problem with this cryptocurrency, however, was that 80% of all coins were mined in advance by its authors. The chances of mining were, therefore, not balanced. This injustice led to the decision that this cryptocurrency would start again. New cryptocurrency starts on 18 April 2014 and was called BitMonero, a composite of the word coin in Esperanto (Monero) and Bitcoin according to Bitcoin. However, after five days, the community decided to use only Monero for short. Monero's significant advantage is the dynamic size of the mined blocks. Bitcoin has one block size limited to 1 MB, while Monero adapts the block size to the network load. If the number of transactions increases, so does the block size to accommodate all transactions. Thus, unlike Bitcoin, the more transactions users make, the lower the transaction fee. Monero's main benefit is its full anonymity and interchangeability thanks to CryptoNote protocol [25]. Monero hides recipient and sender addresses. [21]

2.6 Analysis of Current Blockchain Explorers

A Blockchain Explorer is a web application that allows us to explore the whole blockchain. Their primary function is to allow everyone with an Internet connection to track in real-time all the transactions or interactions made by each cryptocurrencies holders, regardless of his or her level of knowledge and expertise. [16, 7]

2.6.1 Blockbook

As a representative of traditional explorers, we chose Blockbook. Blockbook⁴ is a blockchain indexer for Trezor Wallet⁵, developed by SatoshiLabs⁶. It currently supports more than 30 coins (and the community implemented some others). For data storage, Blockbook is using RocksDB⁷ developed by Facebook, which is a NoSQL database that stores only key-value pairs. Blockbook is providing fast API for accessing blocks, addresses and transactions. Main limitations of Blockbook are:

- **Not distributed** (client-server architecture) – problem with scaling for more users.
- **Not an SQL database** – it does not have a relational data model, it does not support SQL queries, and it has no support for indexes.
- **Single-Process** – only a single process (possibly multi-threaded) can access a particular database at a time.

⁴<https://github.com/trezor/blockbook>

⁵<https://wallet.trezor.io/>

⁶<https://satoshilabs.com/>

⁷<https://github.com/facebook/rocksdb/wiki>

Chapter 3

IPFS

IPFS¹ stands for Interplanetary File System and is a peer-to-peer distributed file system designed to make the Web faster, safer, and more open. In contrast with a standard file system, objects in IPFS are content-addressed by the cryptographic hash of their contents. In the case of the standard Web, when the user wants some file, he needs to know on which server is a file located and the full path to the file (see Figure 3.1). The principle of the IPFS is that a user needs only to know the hash of the requested file. He does not care about the location of the file (see Figure 3.2). Let us take an MIT licence text, and add it to IPFS. If somebody tries to add this licence as a file to IPFS, it will return `QmWpvK4bYR7k9b1feM48fskt2XsZfMaPfNnFxdbhJHw7QJ` every time. That is from now on the *content address* of that file. Later, when a user tries to get this file by its hash, he can get it from a random person that added it into IPFS in the past.

IPFS can easily represent a file system consisting of files and directories. A small file (less than 256 kB) is represented by an IPFS object with data being the file contents (plus a small header and footer). Note that the file name is not part of the IPFS object, so two files with different names and the same content will have the same IPFS object representation and hence the same hash. A large file (more than 256 kB) is represented by a list of links to file chunks that are less than 256 kB, and only minimal information specifying that this object represents a large file. Currently, there are no known size limitations uploaded file or directory. There are already some big datasets hosted on IPFS such as Geocities archive² (704 TB) or Project Apollo Archives³ (61 GB).

3.1 IPFS Stack

We can split IPFS into layers (see Figure 3.3). *Libp2p*⁴ is at the bottom, which is a peer-to-peer networking module, that handles peer and content discovery, transport, security, identity, peer routing, and messaging. *IPLD* is the data model of the content-addressable web. It is providing linking between objects and multihash computing. On the top is *IPFS*, which allows publishing and share files (or any data). [1]

¹<https://ipfs.io/>

²<https://ipfs.io/ipfs/QmVCjhoEFC9vwvaa8bKyJgwAByp4MXSogcyDGoZ4Lkc3ox>

³<https://ipfs.io/ipfs/QmSnuWmxptJZdLJpKRarxBMS2Ju2oANVrgbr2xWbie9b2D>

⁴<https://libp2p.io/>

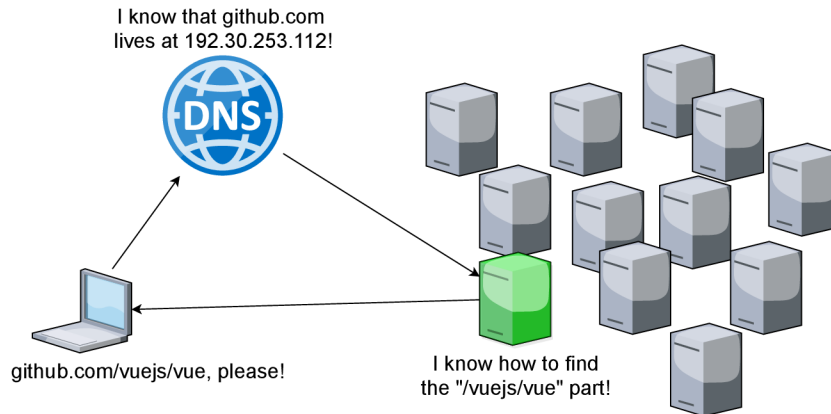


Figure 3.1: Classic web addressing

3.2 Libp2p

Libp2p is a modular system of protocols, specifications and libraries that enable the development of peer-to-peer network applications. It provides NAT Traversal, Peer Discovery, Routing, Stream Multiplexing, Protocol Multiplexing, Encryption, Authentication and more. It grew out of IPFS to solve networking problems in peer-to-peer networks, but now it does not require or depend on IPFS. Today many projects use libp2p as their network transporting layer, which is responsible for the actual transmission and receipt of data from one peer to another. For both content discovery and peer routing, libp2p uses Kademlia-based distributed hash table. With Kademlia, libp2p iteratively routes requests closer to the desired peer or content using Kademlia routing algorithm [18]. In the future, Kademlia might be changed easily to some other solution that implements a simple interface for publishing and requesting data and finding a peer. [8]

3.3 IPLD

IPLD is providing linking and addressing objects with CID (Content ID). CID is hash-based self-describing content identifier (usually encoded to base58⁵ format) which includes codec and multihash. Multihash is then further composed of hash type and hash value. Let us look closer on the MIT licence file, that we add to IPFS at the beginning of this chapter (see Figure 3). Its CID is `QmWpvK4bYR7k9b1feM48fskt2XsZfMaPfnNfxdbhJHw7QJ`. It can be converted to human-readable format as can be seen in Figure 3.1, thanks to multicodec table⁶. We can see that this CID is encoded in base58 format and the file was stored using protobuf⁷ codec (this information is necessary to decode file correctly).

3.4 IPFS

IPFS is the top layer from the IPFS stack. It is used for pinning objects and files, naming system and keys management. File or object is automatically pinned when a user adds it (but other IPFS commands do not include automatic pinning). Pinning a CID tells an

⁵<https://en.wikipedia.org/wiki/Base58>

⁶<https://github.com/multiformats/multicodec/blob/master/table.csv>

⁷https://en.wikipedia.org/wiki/Protocol_Buffers




Figure 3.2: Content-based addressing

Property		Value
Multibase		base58btc
Version		cidv0
Multicodec		dag-pb
Multihash	Hash Type	sha2-256
	Hash Length	256
	Hash	7e1b666c0327...3dc3022f

Table 3.1: Example of a human-readable version of CID

IPFS server that the data is essential and must not be thrown away. When a garbage collector is triggered on a node, any pinned content is automatically exempt from deletion. Non-pinned data may be deleted. The Interplanetary Name System (IPNS) is a system for creating and updating mutable links to IPFS content. Since objects in IPFS are content addressed, an object address changes every time an object's content changes. A name in IPNS is the hash of a public key. It is associated with a record containing information about the hash it links to that is signed by the corresponding private key.

3.4.1 IPFS Core APIs

IPFS provides several APIs that are working on different abstraction levels. The regular top-level API is Files API⁸. The Files API enables users to use the File System abstraction of IPFS. It has `add`, `cat`, `get` and `ls` methods for manipulating with regular files and directories. For our platform are low-level APIs such as DAG⁹ and PUBSUB¹⁰ more interesting. PUBSUB API is used to broadcast messages between peers, and it consists of these methods:

⁸<https://github.com/ipfs/interface-js-ipfs-core/blob/master/SPEC/FILES.md>

⁹<https://github.com/ipfs/interface-js-ipfs-core/blob/master/SPEC/DAG.md>

¹⁰<https://github.com/ipfs/interface-js-ipfs-core/blob/master/SPEC/PUBSUB.md>




Figure 3.3: IPFS stack

- **subscribe** – listen to specific topic. A topic name and handler function are provided as a parameter;
- **unsubscribe** – stop listening on a topic that is provided as parameter;
- **publish** – publish a message to topic. All peers that are subscribed to the topic receive the message;
- **ls** – returns the list of subscriptions the peer is subscribed to;
- **peers** – returns the peers that are subscribed to a topic. A topic name is a parameter of this function.

The DAG (stands for Direct Acyclic Graph¹¹) API provides these methods for manipulation objects in IPFS:

- **put** – stores object in IPFS. Return IPFS hash of a stored object;
- **get** – retrieves an object by its hash from IPFS;
- **tree** – Enumerate all the entries in a graph.

3.5 IPNS

Naming inside IPFS is governed by IPNS¹², the Inter-Planetary Naming System. IPNS takes ideas from SFS¹³ to enable the creation of cryptographically signed mutable pointers, which can be used to the creation of name records inside the network. A *name* in IPNS is the hash of a public key. It is associated with a record containing information about the hash it links to that is signed by the corresponding private key.

¹¹https://en.wikipedia.org/wiki/Directed_acyclic_graph

¹²<https://docs.ipfs.io/guides/concepts/ipns/>

¹³https://en.wikipedia.org/wiki/Self-certifying_File_System

3.6 Filecoin

Another impressive part of the IPFS ecosystem is Filecoin. It is a decentralized storage network that turns cloud storage into an algorithmic market. The market runs on a blockchain with a native protocol token (also called “Filecoin”), which miners earn by providing storage to clients. Conversely, clients spend Filecoin hiring miners to store or distribute data. As with Bitcoin, Filecoin miners compete to mine blocks with sizable rewards, but Filecoin mining power is proportional to active storage, which directly provides a useful service to clients (unlike Bitcoin mining, whose usefulness is limited to maintaining blockchain consensus). This principle creates a powerful incentive for miners to amass as much storage as they can, and rent it out to clients. The protocol weaves these amassed resources into a self-healing storage network that anybody in the world can rely on. The network achieves robustness by replicating and dispersing content, while automatically detecting and repairing replica failures. Clients can select replication parameters to protect against different threat models. The protocol’s cloud storage network also provides security, as the content is encrypted end-to-end at the client, while storage providers do not have access to decryption keys. Filecoin works as an incentive layer on top of IPFS. [14]

3.7 Existing Blockchain Explorers in IPFS

There are already stored a few blockchains of cryptocurrencies in IPFS. For browsing them, we can use dedicated applications¹⁴ ¹⁵ or IPLD explorer¹⁶. Blockchains in IPFS are stored in raw binary format, so custom IPLD codec has to be created for every type of object. Therefore, there needs to be ten different codecs in case of Ethereum (for blocks, transactions, state tries, accounts, contracts etc.). Using custom codecs for cryptocurrencies allows explorers to request blocks and transactions by its hash very fast, but there are also limitations (mainly with filtering or requesting data by different property as a hash).

- Existing IPLD codecs for cryptocurrencies are very limited. Currently, for only a few of cryptocurrencies are codecs available in IPLD (namely Leofcoin, Ethereum, Bitcoin, Zcash, Steller, Decred and Dash)¹⁷.
- Addresses are not stored in IPFS, because they are not part of a blockchain. This means the explorer needs to go through the entire blockchain for computing address balance or to find address transactions.
- No additional information (for example transaction value in US dollars) can be stored with objects because it would change content, thus, the hash of the object.
- There is no sorting or filtering. Explorer can only show the object (block or transaction) by its hash.

¹⁴<https://github.com/arcalineia/IPFS-Zcash-Explorer>

¹⁵<https://github.com/whyrusleeping/zcash-explorer>

¹⁶<https://explore.ipld.io/#/explore/z43AaGEvwdfzjrCZ3Sq7DKxdDhrwoaPQDtqF4jfdkNEVTiqGVFW>

¹⁷<https://github.com/multiformats/multicodec/blob/master/table.csv>

Chapter 4

Design

This chapter describes the proposed design of the whole platform for storing and exploring blockchains in IPFS that is to be created as a result of this thesis. All parts of the platform are described in this chapter.

4.1 Database System

We need a database in our platform to store and index data. Database system needs to be decentralized and distributed. There are several databases build on top of IPFS already implemented.

The two most known are OrbitDB and Textile. However, neither of these solutions fits our use case to store a large amount of data distributed on several nodes and performs queries that can be resolved by downloading only necessary parts of the database. Therefore, we need to create a new database system based on IPFS that would be decentralized and distributed.

4.1.1 OrbitDB

OrbitDB¹ [17] is a serverless, distributed, peer-to-peer database build on top of IPFS, developed by HAJA networks². OrbitDB is a decent solution for small user's databases. However, it is still in the alpha stage of development, and it is not well optimized for storing hundreds of gigabytes of data. The biggest problem is that OrbitDB performs all queries locally. To perform a query that selects transactions that are more valuable than 1BTC, OrbitDB needs to load the whole database locally and then perform a cycle on all transactions to select only those transactions that meet the criteria. So every client ends up with a full copy of the database. This limitation is not usable for our case when we have a database that has hundreds of gigabytes of data.

4.1.2 Textile

Textile³ is a set of open-source tools that provide a distributed peer-to-peer database, remote storage, user management, and more, over the IPFS network. Textile already created applications for storing photos, notes or anything else (Anytype). Textile provides

¹<https://orbitdb.org/>

²<https://haja.io>

³<https://textile.io/>

a high abstraction on top of the IPFS and provides simple API to store and index files securely. Textile provides an open source service called **cafes**, which are nodes that offer offline inboxing and backup. Data are duplicated on several Cafe peers. When a client is performing some query, it contacts one of the Cafe peers to resolve the query for the client. [22]

4.1.3 Database System Design

After some research, we concluded that currently for storing and indexing data in IPFS without large hard disk memory consumption, there is no solution. We created our own indexing system that currently supports three types of indexes. A database system that fits our needs is distributed and decentralized. That brings us lots of synchronization problems to solve. This database system consists of tables that contain records. For faster searching, tables have indexes. Relations between tables are represented via foreign keys. Also, this database system supports fluent query language, used for performing complex queries.

Record

Every record in our database system is stored in an append-only log that contains the whole history of the record. Every update of a record adds a new entry to its log which points to the previous entry.

In centralized systems such as git⁴, conflicts are detected on write (for example, when two git users push changes to the server at the same time, one of them gets an error and needs to pull repository) [3]. This approach is impossible in a decentralized system. When we update some record in our database, we can not know if somebody updates it before us (and we do not receive changes yet). For this reason, we need to solve conflicts while reading. Record with conflict has more than one head in a record log, and users need to solve them in application logic. Look at the example in Figure 4.1. There is a record that is updated by every country in the World when they have new statistics about a pandemic. If two countries updates data at the same time (their updates are pointing to the same previous version of the entry), they create conflict. Luckily, this specific conflict is easy to solve. We just need to look at the previous entry and compute increment for both countries, then update the record with final increment.

Indexes

An index is a B-tree structure optimized for IPFS (with no cycled references and node size less than 256 kB). Each table has at least one primary index. We use the primary index to reference record in foreign keys. A primary key is automatically created when a user does not specify it and has a type of GUID. Value of the primary key for an entity can not be changed, because we would need to scan all tables where the entity is referenced and change the referenced value to a new one. Also, when we execute a query without any condition or sorting, we use a primary key to obtain records. Every index has several components:

- **Comparator** – is a function that has two parameters (two B-tree keys) and outputs a number as a result. An index is using a comparator to correctly traverse B-tree while searching records and inserting records to the right place.

⁴<https://git-scm.com/>

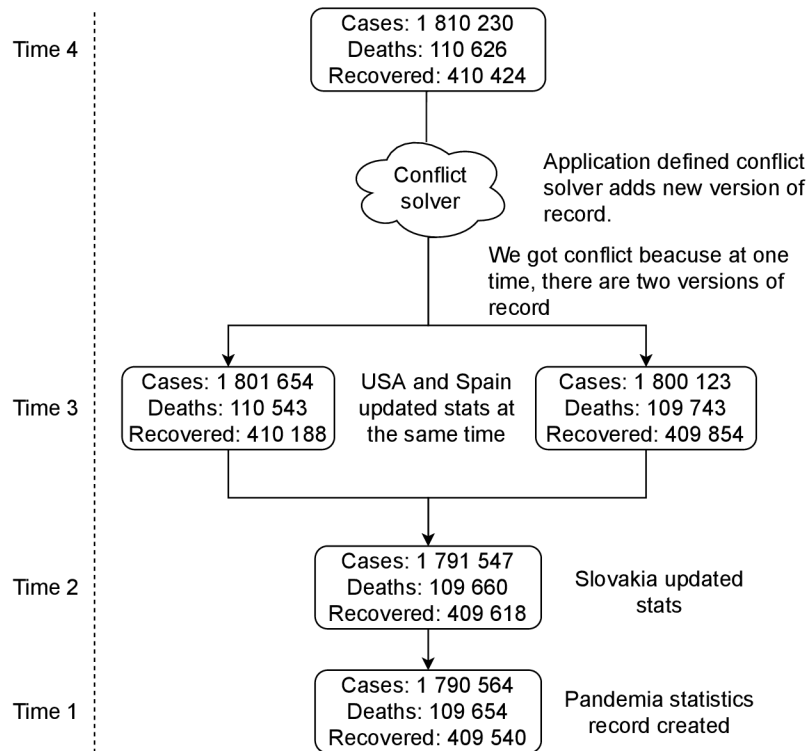


Figure 4.1: Record conflict

- **Key-getter** – is a function that returns a comparable object from the record (a comparable object is an object that can be compared using index Comparator). Every index has Key-getter that is using to obtain keys for objects stored in B-tree.

Table

A table contains indexes and table name. Also, it implements operations that modify its data:

- **Insert** – creates a new record in a table. Insert operation has these steps:
 1. Creates record log with a new entry and store it in IPFS.
 2. Gets record key (via Key-getter) and inserts it to all table indexes with IPFS hash as value.
- **Update** – saves a new version of record. It consists of:
 1. Add new entry to the record log. Save the record log to the IPFS.
 2. Get new record key (via Key-getter) and update all table indexes with new IPFS hash as value.
- **Delete** – removes a record from a table.
 1. Get old record key and remove it from all table indexes.
 2. Add empty entry to the record log. Save the record log to the IPFS.

Foreign Keys

We use foreign keys to represent relationships between tables. A foreign key is simply a table name and a value of a primary key of a referenced record.

Transactions

When we commit a transaction, it is inserted to the transactions queue. If there are more transactions in the queue, a database performs it one by one. We can not execute more transactions in one moment, because it could cause data inconsistency. There are five types of transactions:

- **read** – transaction for reading data from table. It is not logged in a database log (other peers don't have to know what we are reading from database).
- **sync** – when some peer publishes a new version of the database, all other peers have to migrate to it. This transaction has the biggest priority, and therefore it is inserted at the beginning of the transactions queue. After the migration is done, transactions queue can continue working usually.
- **insert, update, delete** – these types of transactions are written to database logs, because they are modifying the database.

Synchronization

Every time a transaction queue is empty (all pending transactions has been applied), a new version of a database is broadcasted to all connected users via IPFS PUBSUB API (see 3.4.1). When we receive information about the new database version, we load its root. Database versions create an append-only log called database log so each database version has information from which version it was created. Database log is an immutable, operation-based conflict-free replicated data structure (CRDT⁵) for storing database versions. Every version in the log is saved in IPFS, and each points to a hash of the previous version forming a graph.

If more than one peer publishes a new version of the database that has been created from the same database version, other peers need to decide which version they would accept. Opposite to records conflicts, we need to solve database versions conflict fast and automatically. There are multiple ways to solve them. For example:

- **The biggest hash wins** – if there is more than one database version at the same (discrete) time, the one with the biggest IPFS hash wins. This strategy is present in Figure 4.2.
- **The longest connecting time wins** – a peer that is connected to the database for the longest time wins.

Lots of more strategies can be implemented, but they need to be deterministic and as fair as possible. If we have access to geolocation, we can implement a strategy based on a distance to the North Pole (the closest node to the North Pole wins). [15]

⁵https://en.wikipedia.org/wiki/Conflict-free_replicated_data_type

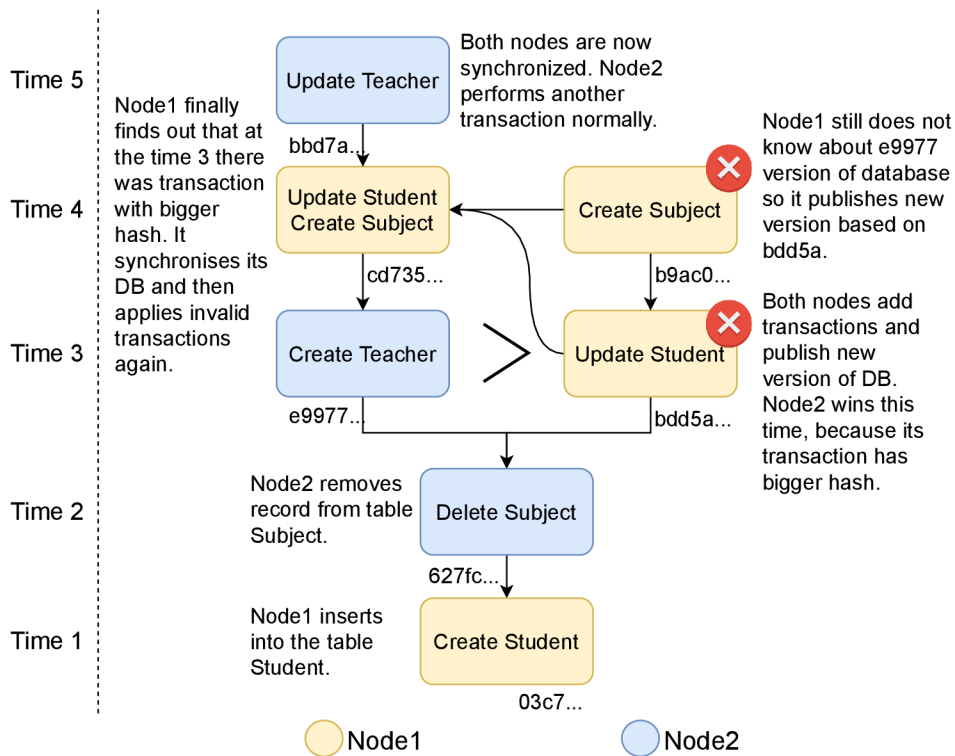


Figure 4.2: Synchronization of a database. Note that in Time 4, node1 adds two transactions in one DB version.

Queries

A database system provides query language for performing selects. A query consists of:

- **conditions** – query may contain multiple conditions. There are logical operators (and/conjunction, or/disjunction) between conditions. When it is possible (i.e., an index is available on condition property), we use indexes for evaluating conditions. If there are more than one conditions, we use an intersection between AND conditions and union between OR conditions. These types of conditions are supported:
 - **equal** – record property equals to specified value;
 - **greater than** – record property is greater than specified value;
 - **less than** – record property is less than specified value;
 - **between** – record property is greater than specified minimal value, and less than specified maximal value.
- **filters** – filters are similar to conditions, but they can be more complicated. They are functions that are being applied to query results. If any filter return **false** for the query result, the result is ignored. A query can contain several filters.
- **offset** – offset is saying how many query results should be ignored from the beginning.
- **evaluator** – we use evaluator for accessing results of the query:
 - **all** – returns array of all results of the query;

- **first** – returns only first result;
- **take** – returns N number of results where N is an argument;
- **paginate** – returns page of results. It accepts two arguments: **perPage** specifies the number of results in one page, **page** is number of requested page;
- **iterate** – returns iterator that can be used in cycles such as **for** or **while**.

4.2 Platform Components

The platform consists of one or more Feeders and Explorers. Feeders are connected to data sources and provide synchronization with cryptocurrency blockchains. Explorer can request data from the platform network and display them to a user (see Figure 4.3).

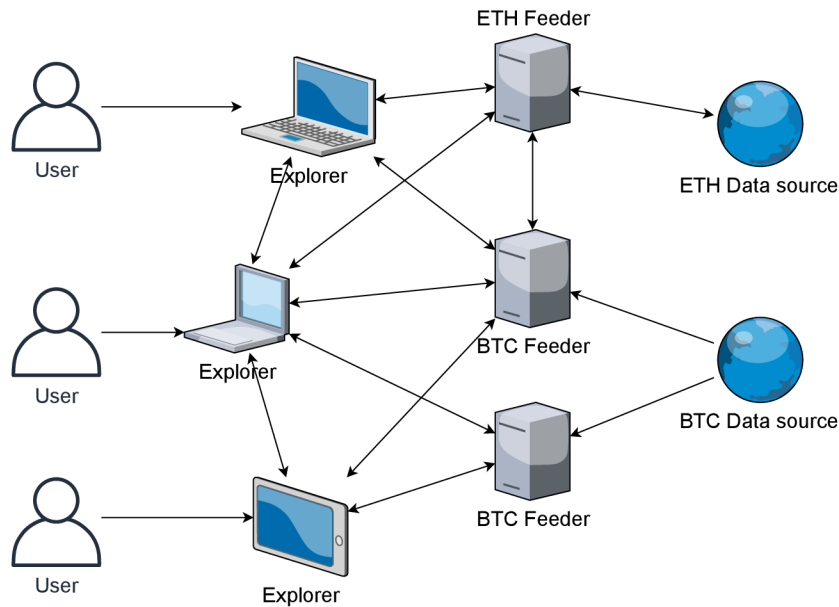


Figure 4.3: Platform architecture

4.2.1 Database Design

Both Explorer and Feeder use the same database schema. The database schema is designed in the way to optimize typical queries. It consists of five tables:

- **Block** – table for storing blockchain blocks. It has a primary key on unique block hash, and indexes on height and time. With these indexes, we can perform queries like search block by its hash or height very efficiently. Also filtering or ordering blocks by time has good performance.
- **Transaction** – a table that contains blockchain transactions has a primary key on transaction hash and foreign key, that references block in which is transaction confirmed, on **blockHash** column. Indexes for efficient filtering are on **blockHeight** and **blockTime** columns.

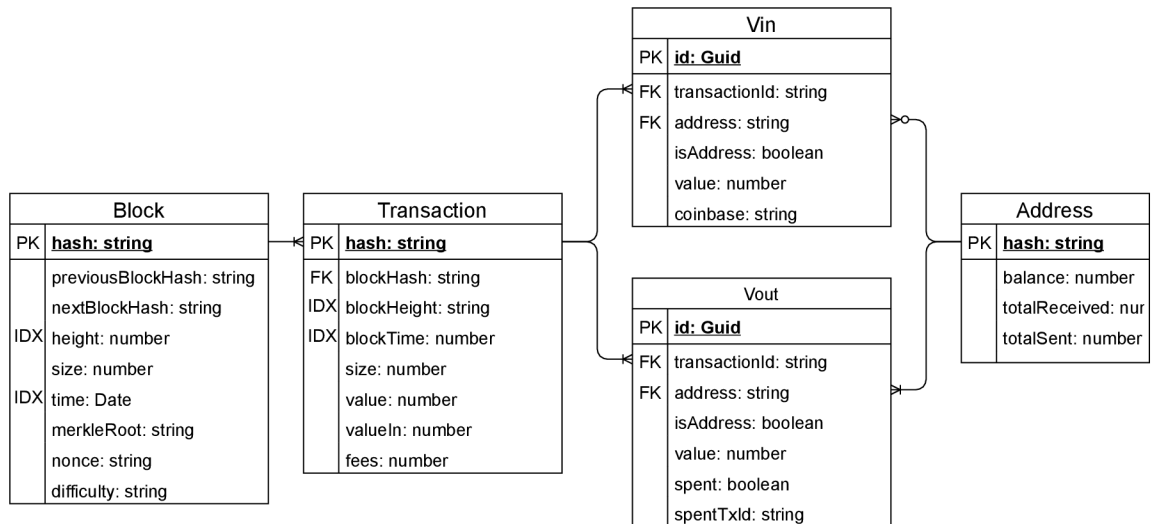


Figure 4.4: Entity-relationships diagram

- **Vin** – every transaction can have multiple inputs. These inputs are stored in Vin table. Vin has not got any unique property, so a custom GUID⁶ needs to be created as a primary key. It has two foreign keys to reference transaction and address. Foreign key **address** can be null, because the input can be mined block reward.
- **Vout** – transactions can have multiple outputs. Every output has a foreign key to transaction and address. If the output is already spent, it has a link to the transaction where this output is used as an input.
- **Address** – Address contains multiple columns. As the primary key, it uses the address hash. Address also contains its **balance**, **totalReceived** and **totalSent**. These values are updated every time new input or output is added to the transaction.

4.2.2 Feeder

A Feeder is a service that stores data in IPFS and indexes them in our database system. Once all blocks are indexed, Feeder waits for new blocks and is periodically publishing most up to date database version. This way, new clients synchronize database more quickly. For optimization purposes, Feeder is performing transactions in bulks. This optimisation means that Feeder is not publishing a new version of the database every time it parsed block or transaction, but rather after it parsed transactions bulk which can consist of hundreds of transactions.

A Feeder can use different sources for obtaining blockchains data. It can be connected directly to the blockchain as a full node or to some other source providing blockchain data as Blockbook⁷ or Insight⁸.

⁶https://en.wikipedia.org/wiki/Universally_unique_identifier

⁷<https://github.com/trezor/blockbook>

⁸<https://insight.is>

4.2.3 Explorer

Explorer can perform basic queries like a search for block by its height or hash and search address and transaction by hash. Nevertheless, Explorer can also make more complex queries (for example, get the first 20 transaction where the sum of inputs is more than some value, or get transactions between some time interval). Explorer is an application that can be used in two environments. In browser with graphical user interface, or in `node.js`⁹ as an API.

ExplorerGUI

ExplorerGUI is browser version of Explorer. It is implemented as an SPA¹⁰ to prevent restarting connection with IPFS after each time a user visits a different page. It has separated views for block, transaction and address details. Also paginating, filtering and sorting objects by all its properties that have index is supported. Every user of ExplorerGUI keeps part of the database that he uses in his local storage. This principle helps balance Feeders load and helps to distribute data for other users. Authentication to the ExplorerGUI is not necessary, because users can not modify data from there. They can only explore multiple blockchain data. All views have the same top toolbar, to allow user performs quick searches. User can see enabled cryptocurrencies at the home screen (see Figure 4.5a). When a user selects cryptocurrency that he wants to explore, he will be redirected to the Blocks view where he can see and filter all parsed blocks from a selected cryptocurrency (see Figure 4.5b). User can select block by clicking on them. This action redirects a user to Block detail view where he can see all blocks transactions (see Figure 4.6a). User can also navigate to the next/previous block from this view. A user gets to the Address view every time he clicks on transaction input or output address. In this view, a user can see the current balance of the address and transactions which this address is part of (see Figure 4.6b). Thanks to the record log, a user can also see the history of the address.

ExplorerAPI

ExplorerAPI is a server-side application that provides a simple HTTP API for obtaining data. Although our platform is decentralized and therefore does not need servers, ExplorerAPI can be used with older devices on which it is not possible to run an IPFS node, or when integrating with other applications. Through RestApi, we could not use the full potential of our database system, so we get inspired by GraphQL¹¹ during the design. ExplorerAPI provides endpoints for these database tables:

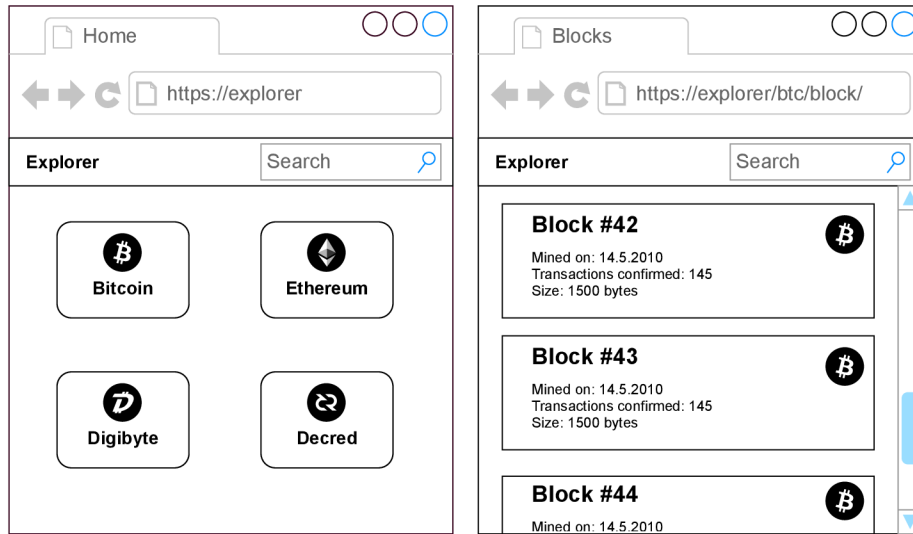
- `/:currency/block` – endpoint for obtaining blocks. Each endpoint hash `currency` path parameter that can be currency unit (for example `btc`) or currency name (for example `bitcoin`).
- `/:currency/transaction` – endpoint for retrieving transactions.
- `/:currency/address` – endpoint for obtaining data from address.

Every endpoint can be accessed with `GET` or `POST` HTTP method. A request contains a query object that consists of conditions, filters, skip and resolver (see Figure 4.7a). This

⁹<https://nodejs.org/>

¹⁰https://en.wikipedia.org/wiki/Single-page_application

¹¹<https://graphql.org/>



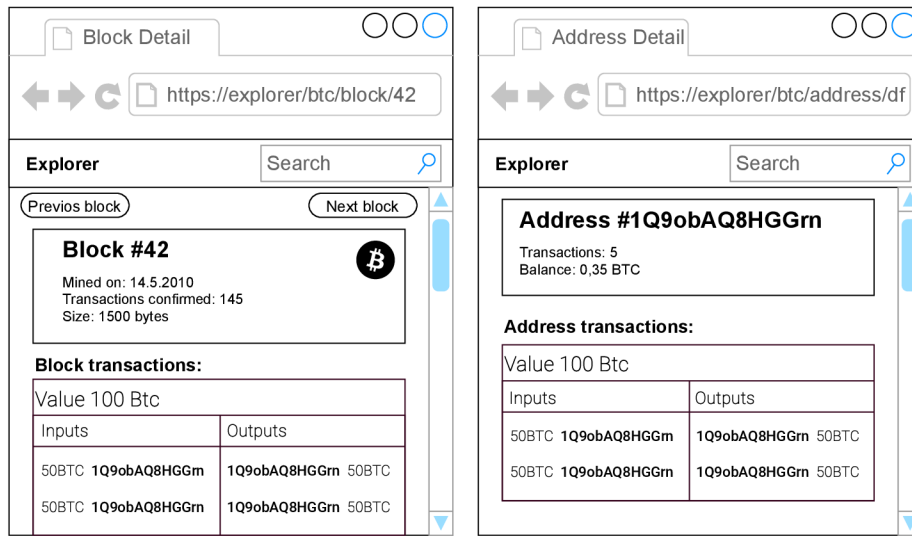
(a) Home view

(b) Blocks list view

Figure 4.5: ExplorerGUI mockups

query object is translated to our Query system and executed. Query results in JSON format are returned (example result is on Figure 4.7b). If a request is sent by GET method, a query object is stringified and urlencoded¹² JSON in query parameters (for example /block?query=%3D%7Bconditions...%7D). If POST method is used, query object is in request body (like on Figure 4.7a). With this design we can use multiple conditions and filter in query.

¹²<https://en.wikipedia.org/wiki/Percent-encoding>



(a) Block detail view

(b) Address view

Figure 4.6: ExplorerGUI mockups

```
{
  "query": {
    "conditions": [
      {
        "type": "and",
        "comparator": "between",
        "property": "height",
        "values": [11990, 12000]
      }
    ],
    "filters": [
      "(b) => b.height % 2"
    ],
    "resolver": {
      "type": "all"
    }
  }
}
```

(a) Request that returns blocks with an odd height that is between 11,990 and 12,000.

```
[
  {
    "hash": "000000001e564...",
    "height": 11991,
    ...
  },
  {
    "hash": "00000000b3637...",
    "height": 11993,
    ...
  },
  {
    "hash": "000000007b4da...",
    "height": 11995,
    ...
  },
  {
    "hash": "00000000e4361...",
    "height": 11997,
    ...
  },
  {
    "hash": "00000000e9935...",
    "height": 11999,
    ...
  }
]
```

(b) Response

Figure 4.7: ExplorerAPI example

Chapter 5

Implementation

This chapter describes the implementation details of the platform and shows the internal architecture. All platform components are implemented in Typescript that is later compiled to JavaScript. This choice of language enables code sharing between individual components and allow us to support multiple platforms (desktop via Node.js, and browser). All platform components are dependent on Explorer-core module where is the whole database system implemented (see Figure 5.1).

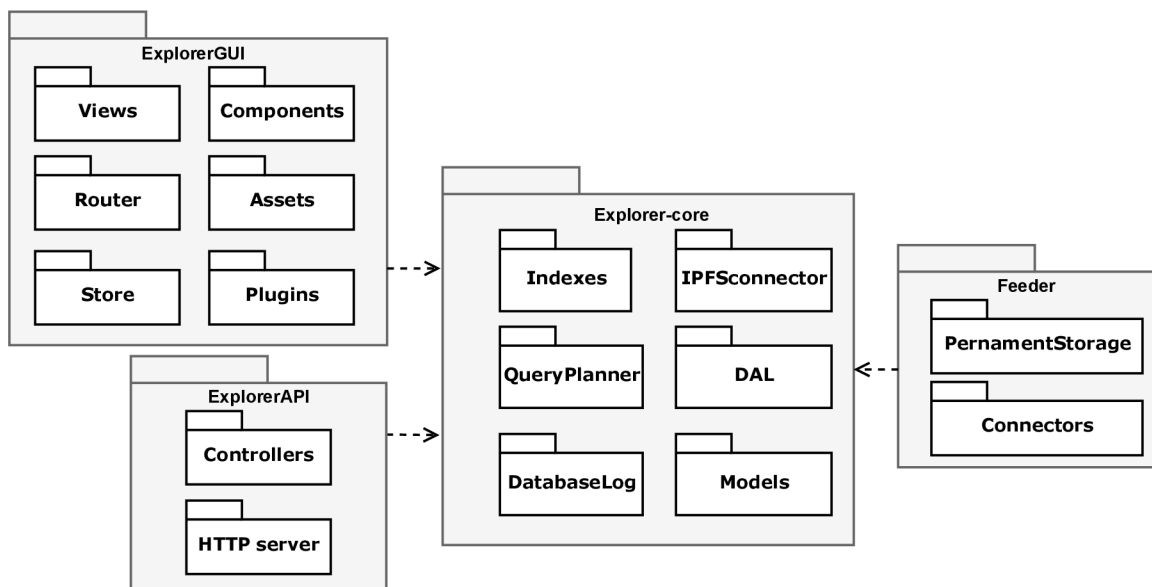


Figure 5.1: Platform architecture

5.1 Explorer-core Implementation

Explorer-core is the most complex module of a whole platform with more than 5,000 lines of code. The database subsystem is a main part of the Explorer-core module. The Explorer-core consists of a Query system, indexes and an abstract database layer. Explorer-core is only out of platform components that communicates with IPFS via [js-ipfs](https://github.com/ipfs/js-ipfs)¹ implementation.

¹<https://github.com/ipfs/js-ipfs>

5.1.1 IPFS Connector

IPFS connector is a singleton class that provides a connection to IPFS. It has a method `getInstanceAsync` which will return promise that resolves into IPFS node instance. This instance is stored in a private static class variable. Next call of function `getInstanceAsync` returns that static class variable. Also, swarm key is present in IPFS connector. This key is used to make private IPFS network with only Explorers and Feeders. Other peers (without swarm key) can not connect to our network. IPFS connector contains settings for node.js and browser IPFS peer. Browser peers use WebRTC and WebSockets for transport. Node.js peers use TCP and WebSockets. That means that Feeders (a node.js application) can communicate with each other with TCP. Explorer communicates with other explorers with WebSockets or WebRTC. Explorer uses WebSockets when communicating with Feeders.

5.1.2 Indexes

For indexing purposes, there is currently implemented only B-tree, but different structures such as tries² can be employed easily. They only need to implement index interface (function such as `insert`, `delete`, `update`, `find`). We can create indexes on database entities with decorators³, which are part of ECMAScript 6 standard. Every database entity has to have `PrimaryKey` decorator on property that is used as a primary index. Index has three parameters:

- `comparator` – is a function that accepts two arguments and returns number that is less than zero if the first argument is greater than a second, zero if arguments are the same, more than zero if a second argument is greater than first. If a user has not set any custom comparator default one (`(a, b) => a < b ? -1 : +(a > b)`) is used. This default comparator works on atomic keys such as string or numbers.
- `keyGetter` – is another function that accepts a whole entity as arguments and returns key, that is used in an index. A default key getter is function that returns value of index property (for example default key getter for property `height` is `(e) => e["height"]`).
- `branching` – the branching factor is the number of children at each node (the outdegree). Default branching factor is 16.

The B-tree structure (see Figure 5.2) is optimized for IPFS. Unlike PostgreSQL B-tree⁴, leaf nodes have not got left and right sibling references due to cyclic reference that is not possible in IPFS, because links are changing CID of an object [23]. We can demonstrate this in the following example:

- If we store object A with a link to the object B, CID of A will change.
- Then if we add backlink from B to A, CID of B will be changed and so A is now linked to an old version of B.
- If we update link of object A to point on the new version of object B, CID of object A will change and therefore object B is now pointing to the old version of object A.

²<https://en.wikipedia.org/wiki/Trie>

³<https://www.typescriptlang.org/docs/handbook/decorators.html>

⁴<https://github.com/postgres/postgres/tree/master/src/backend/access/nbtree>

From this example, it appears that there is no way of making cyclic references in IPFS. Without siblings's references, we do not have to store data only in leaves, but we can store data in nodes itself. This approach leads to better performance when executing queries. Statistically less nodes need to be updated in case of insert queries. Also in search queries, there is a chance that we would find a key in some non-leaf node; therefore, we would need fewer node visits.

All search queries (equal, less, greater, between) has two steps:

- **Find subtree** – First, we need to find minimal subtree that contains all results. We start in a root node of the B-tree. Then, we use an index `comparator` function to determine if we can visit some child node and still have all the result in it. If not, then we have found a minimal subtree for the query.
- **Traverse** – To get results of the query, we can traverse minimal subtree in two directions. In-order for `greater than` and `between`. For `less than` reversed in-order. With `equal`, it does not matter.

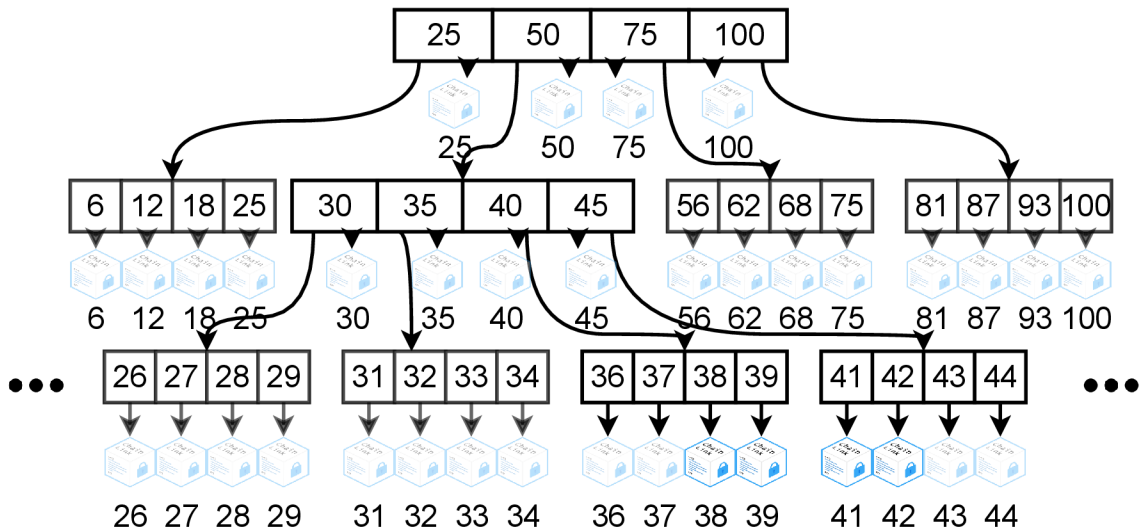


Figure 5.2: B-tree indexing first 100 blocks by their height

5.1.3 Query System

Our database offers a complex Query system. A query can consist of multiple conditions, and the `Query Planner` is responsible for resolving them. It decides which indexes are used for query and choose a strategy. If there is no condition, a primary key is used for query execution. In the case of a single condition, `Query Planner` checks if there is an index on the condition's property. If yes, then this index is used to perform the query. Else condition is transformed to filter, and a primary key is used to obtain results which are then filtered. For multiple conditions connected with logical operators `AND` or `OR`, `Query Planner` creates `OR-hashset` and `AND-hashset`. `AND-hashset` is initialized with the results of a condition which has the smallest number of results and uses `AND` operator. Then we check for each result hash for all other `AND` conditions if `AND-hashset` contains it. If no, a hash is deleted from `AND-hashset`. This creates intersection between all `AND` conditions. The `OR-hashset`

is created empty. Then, we add a hash of every result of all OR conditions to it. This creates union between OR conditions. At the end, we perform union between AND-hashset and OR-hashset. This creates final hashset which is used later to obtain actual data with resolvers such as `all`, `first`, `paginate` etc. If one of the AND conditions has great selectivity (it has less than a hundred results), Query Planner can decide to don't use other indexes, but transform condition to filter and cycle over the results. Example query returning blocks between 38 and 42 is shown in Figure 5.3. A peer needs to download only data that are highlighted in green. After downloading them, they are stored on a peer's filesystem and are offered to other peers.

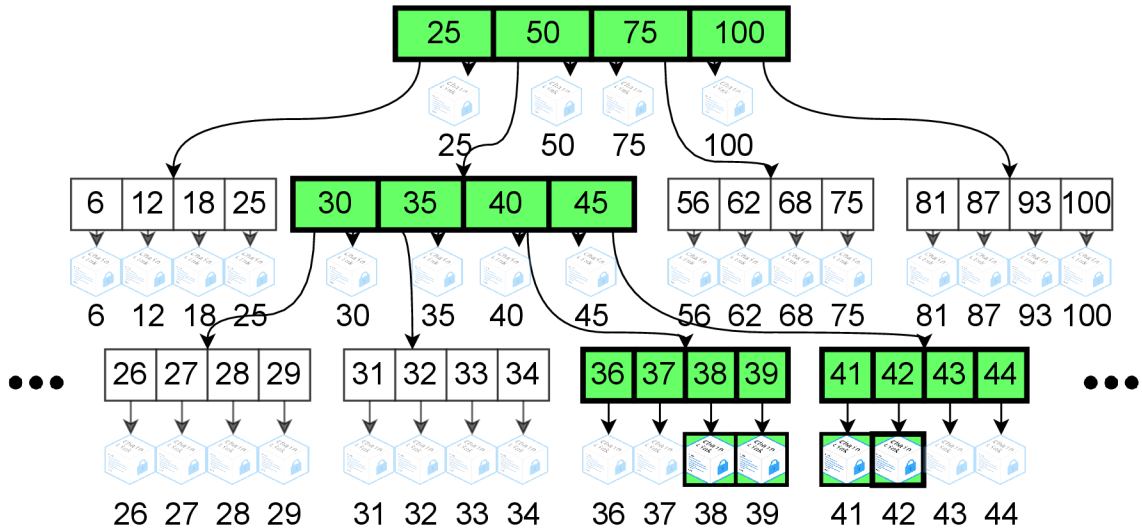


Figure 5.3: Data access for performing query that returns blocks that have a height between 38 and 42.

We can make queries with every database table model (such as model on Figure 5.5). Whole Query system is displayed in the state diagram on the Figure 5.4. The Query system consists of these functions:

- `where(propertyName)` – create a condition on the property. There can be multiple conditions in one query. A condition hash needs to be followed by one of the functions:
 - `gt(value)` – property is greater or equal than `value`. The Query Planner finds in an index first object that has property (set by `propertyName` in `where` function) equal or greater than `value`, and traverse index to the right (to bigger objects);
 - `lt(value)` – property is less or equal than `value`. Similar as in the `gt` function, the QueryPlanner finds the closest object that has property equal or less than `value`. Then, the query traverses index to the smaller objects with smaller index value (to the left);
 - `between(min, max)` – property is greater or equal than `min` and less or equal than `max`;
 - `equal(value)` – returns all objects that `keyGetter` function returns same key as `value`;
- `skip(offsetValue)` – query will skip `offsetValue` number of results;

- `limit(limitValue)` – set maximum number of results. After query has `limitValue` count of matched objects it will stop browsing the index;
- `all()` – return all objects that matched query;
- `first()` – return the first object that matched a query;
- `and(childQuery)` – logical and between two queries. Parent query resolves `childQuery` (call `all()` function) and add its results to `AND-hashset`;
- `and(property)` – Thanks to available generic programming⁵ in Typescript, `and` can be also called with string argument. This register another query condition;
- `or(childQuery)` – logical or between two queries; `childQuery` will be resolved when parent query needs it, and its results stored in `OR-hashset` of parent query;
- `or(property)` – similar as `and` function, `or` can be also called with string parameter. This register another condition to query;
- `iterate()` – returns iterator that can be used in `for (result of query) cycle`.

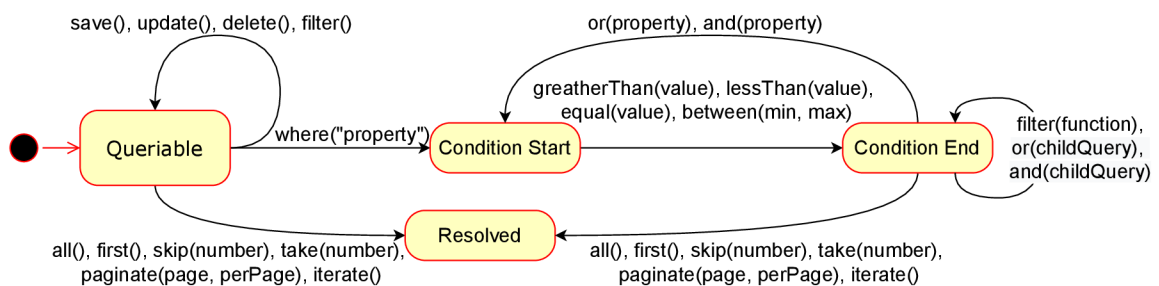


Figure 5.4: State diagram of the Query system

5.1.4 Database

The most important part of the Explorer-core module is the database subsystem which connects the `Query Planner` with indexes. When we make a query, it is precessed by database, and data are obtained with the help of the indexes.

Tables

A database contains tables that consist of indexes. A table has an interface that provides operations to modify table data:

- `create` – creates history log for entity with first entity version. Then adds it to every table index;
- `update` – adds new version of the entity to its history log. Then update every index of the table;
- `delete` – removes entity from every table index. From now, an entity can not be found in this table.

⁵https://en.wikipedia.org/wiki/Generic_programming

There is no `select` operation because a table itself does not perform queries for obtaining data (table has not got necessary logic for choosing optimal index for query). Selects are performed by `Query Planner` after analysing all query conditions.

Example of the class `User` is in Figure 5.5. It has two indexes. One primary on property `name` and one normal index on property `age`. The second index has also specified `comparator` that is a bit faster than default one but works only on numbers, and `keyGetter` that returns a year when a user has born.

```
class User extends Queriable<User> {
    @PrimaryKey()
    name: string;

    @Index(
        (a, b) => a - b,
        u=> new Date().getFullYear() - u.age,
    )
    age: number;
}
```

Figure 5.5: Example of database table abstraction

Transactions

A database can execute only a single transaction at the time to prevent data inconsistency. For that reason, we implemented a transaction queue where transactions are stored before executing in the order in which they came. Executing more transactions in a row is significantly more effective than executing them one by one. If a transactions queue has only one transaction, it waits 50 ms for more transactions to come. Transactions in our platform are represented by classes `Transaction` and `TransactionsBulk`. Both classes implements `ITransaction` interface with `run` function. Transactions bulk are treated the same way as a single transaction thanks to Composite pattern⁶ (see Figure 5.6). After executing, each transaction is appended to the database log. After we execute all transactions in transactions queue (transaction queue is empty), we publish a new database version to all connected peers.

Synchronization

There are two types of transactions: 1) those that changes the database state (`create`, `update`, `delete`); and 2) those that do not (`select`). Every transaction that changes the database state needs to be synchronized with other peers. We created database log for that. It is an append-only log with a discrete-time. Every entry of a database log has these properties:

- **hash** – an IPFS hash of this entry;
- **payload** – payload is object that contains all entry data. Usually, there is a transaction (if there are more transactions published in this entry, there is transactions bulk) and an IPFS hash of the database;

⁶https://en.wikipedia.org/wiki/Composite_pattern


```

interface ITransaction{
    run();
}

public class Transaction implements ITransaction {
    operation: DbOperation;
    data: any;
    run() {
        ...
    }
}

public class TransactionsBulk implements ITransaction {
    public transactions: ITransaction[];
    public run() {
        for (const transaction of this.transactions)
            await transaction.run(database);
    }
}

```

Figure 5.6: Transactions use composite pattern. Body of function `run` of the `Transaction` class is omitted due to complexity.

- **parent** – an IPFS hash to previous entry. We can re-create a whole database log from a head entry by accessing `parent` property until it is `null`;
- **clock** – A Lamport clock ⁷ with entry create time.

There can be only one valid database transaction at every point in time. There are several strategies to choose which transaction is globally accepted and which transactions need to rollbacks (described in design 4.1.3). Merging database logs is a common operation in our database system. Every time a peer publishes new database version, other peers need to merge this version with theirs.

Example of merging database logs can be seen in Figure 5.7. In this example there are two peers (Node A and Node B). First three database version (A1, B1, A2) follows each other and there is no problem with them. But versions A3 and B2 are published in the same time. Node A and B do not know about this issue yet, so they both publish another version (A4 and B3) on top of their head. After nodes get notified about each other versions, their logs are merged by database log class function `merge` which accepts another log as a parameter (see Figure 5.8).

At first, we create a new instance of `TransactionsBulk` where we store transactions that we may need to rollback. Then we check if both logs heads are published at the same time. If one head clock time is greater than another head clock time, we need to get previous database log entry from this log with the same time as another log head to be able to compare these versions. Luckily, both logs head from our example (A4, B3) are published at the same time. We can get a previous entry from accessing `parent` property of database log entry. Finally, after we got both logs heads from the same time, we can start traversing them and comparing their parents. If logs have the same parent (in our example it is version A2), we found a database version where the fork began. Now, peers need to decide which version of the database to accept. The simplest method is to compare hashes

⁷https://en.wikipedia.org/wiki/Lamport_timestamps

of database log entries. In our example A3 entry has bigger hash than B2. This means that Node B needs to rollback transactions published in versions B2 and B3 and apply them after A4. After the merge, database log has a single head again and all published transactions have been applied.

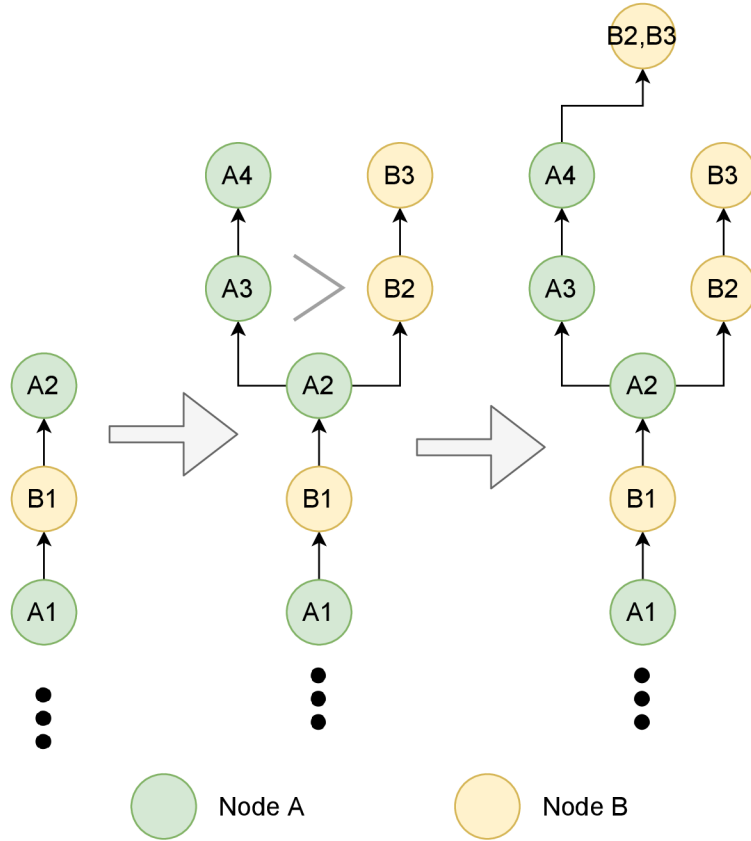


Figure 5.7: Database synchronization

5.2 Feeder

A Feeder is a simple command-line application written in Typescript. Feeder high-level operation is described in Algorithm 1. It strongly depends on the Explorer-core module. Currently, we support only Blockbook connector as a source of blockchain's data. However, Feeder can be simply expanded to support more data source such as InsightAPI or direct connection to a blockchain as a full node. Each Feeder has configuration file (usually called `.env`) with Feeder's settings. Main Feeder settings are URL of the source for the blockchain's data and `DB_NAME` which is a name of the database where Feeder inserts new blocks. A Feeder can be connected to only one blockchain. We can create multiple Feeders for single blockchain, but we need to provide some deterministic algorithm, that ensures that each block is parsed by only one Feeder. This can be done by providing `FeederId` and


```

public async merge(log) {
  let thisHead = this.head;
  let otherHead = log.head;
  const rollbackOperations = new TransactionsBulk();
  while (thisHead.clock.time != otherHead.clock.time) {
    if (thisHead.clock.time > otherHead.clock.time) {
      if (thisHead.identity.isMe())
        rollbackOperations.push(thisHead.transaction);
      thisHead = await this.get(thisHead.parent);
    } else {
      otherHead = await log.get(otherHead.parent);
    }
  }
  while (thisHead.payload.parent != otherHead.parent) {
    if (thisHead.identity.isMe())
      rollbackOperations.push(thisHead.payload.transaction);
    otherHead = await log.get(otherHead.parent);
    thisHead = await this.get(thisHead.parent);
  }
  if (otherHead.compare(thisHead.hash)) {
    await this.migrate(log, rollbackOperations);
  }
}

```

Figure 5.8: Simplified code for merging database logs

`FeedersCount` in `.env` config file. If those values are provided, Feeder parses only blocks that have height modulo `FeedersCount` equals to `FeederId`.

Algorithm 1: Simplified Feeder algorithm

```

load configuration;
while there is new block do
  fetch block;
  for transaction in block do
    | save transaction;
  end
  save block;
end

```

5.2.1 ExplorerGUI

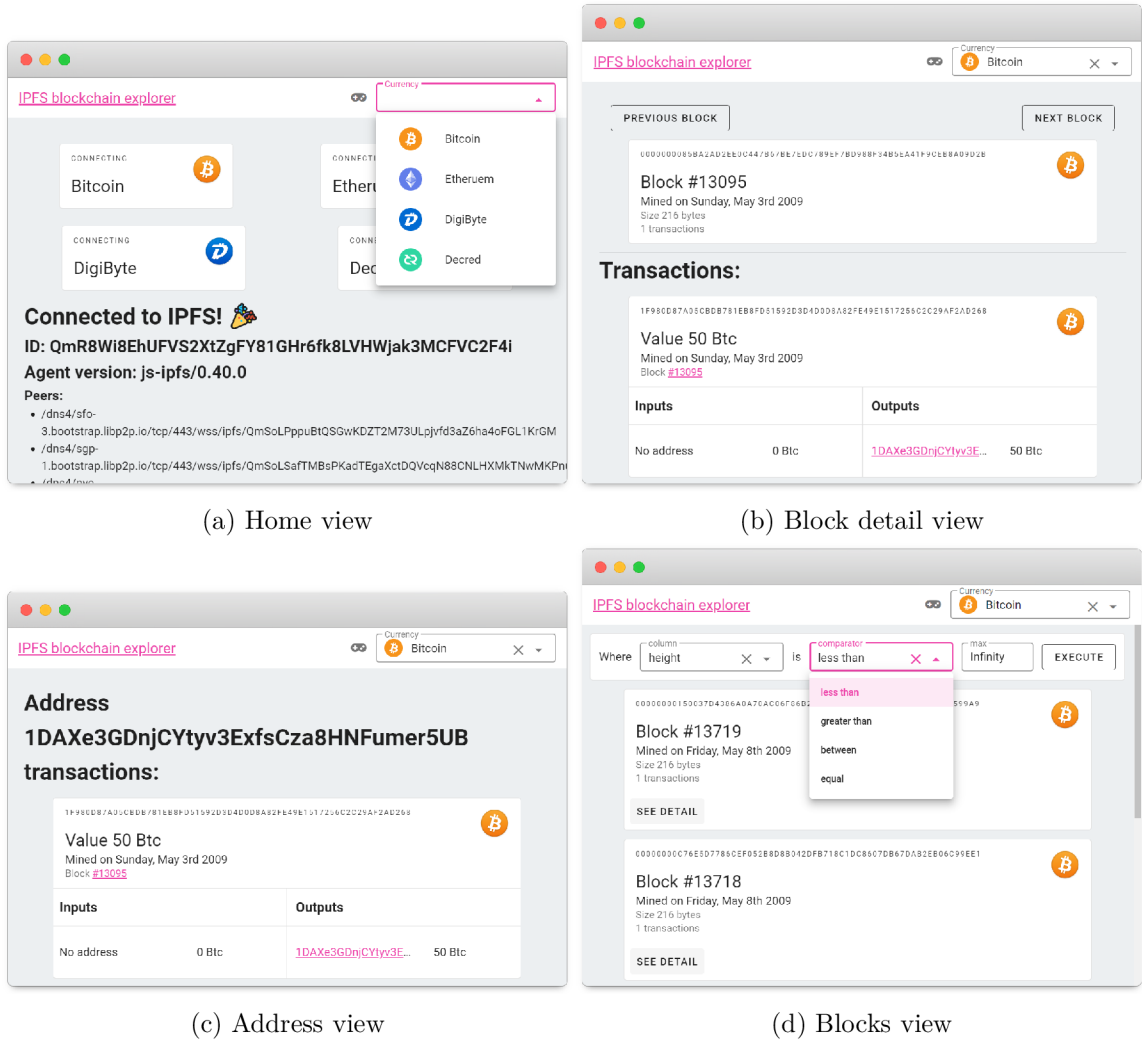
ExplorerGUI is a single page application with a simple user interface implemented with Vue.js⁸ that runs in a browser. We use Vuetify⁹ as a user interface library and browser implementation of IndexedDB¹⁰ as a storage for IPFS. Communication with other peers is provided though WebRTC¹¹ and WebSockets because a web page in a browser can not open TCP socket. Every opened ExplorerGUI browser tab is the same IPFS node instance. Opening a new tab in incognito mode or different browser spawns new IPFS node.

⁸<https://vuejs.org/>

⁹<https://vuetifyjs.com/>

¹⁰https://developer.mozilla.org/en-US/docs/Web/API/IndexedDB_API

¹¹<https://webrtc.org/>



(a) Home view

(b) Block detail view

(c) Address view

(d) Blocks view

Figure 5.9: ExplorerGUI views

After page with ExplorerGUI is loaded, it tries to connect to all supported blockchains. The home screen contains all enabled cryptocurrencies with their statuses (see Figure 5.9a). Blocks list view (Figure 5.9d) is displayed when a user selects specific cryptocurrency. A user can filter blocks by all blocks indexes in this view. A default filter gets all blocks with height less than infinity. This query sorts blocks from the newest (highest) to the oldest (see Figure 5.9d). When a user clicks on the Execute button, ExplorerGUI starts loading blocks that match the selected query by the asynchronous algorithm shown in Figure 5.10. First, we use `iterate` function (implemented in Explorer-core module) to obtain iterator over subtree of all results. This function is asynchronous, so it does not block the main JavaScript thread while it is searching the whole index and looking for the subtree. Then we create an array of promises¹² in which the individual blocks that are displayed on the page are loaded. In for cycle, we push asynchronous tasks to the array. A task contains single parameter that is `pagePosition`. This parameter is the order of the block that tasks loads. Every task for each block then runs in parallel. First, task calls the `next` function on subtree iterator.

¹²https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/Promise

This function returns the next block in the subtree. The `next` function can return left or right sibling of the current block. It depends on the condition. A subtree is traversed to the left if `greaterThan` or `between` condition is used and to the right if there is `lessThan` condition. The `next` function returns a pair of `blockPromise` which is a task that loads blockchain block from IPFS and `done` that is boolean, which signals if there are any more results in the subtree. Finally, task waits until blocks are loaded, and then assigns block in the right index of the blocks array. Thanks to this optimized asynchronous algorithm, all blocks displayed on the page are loaded in parallel.

```
const subtree = await query.iterate()
const tasks = [];
for (let i = 0; i < pageSize; i++) {
  tasks.push(
    (async pagePosition => {
      const { blockPromise, done } = await subtree.iterator.next();
      if (done)
        return;
      blocks[page - 1 * pageSize + pagePosition] = await blockPromise;
    })(i),
  );
}
await Promise.all(tasks);
```

Figure 5.10: Loading blocks from database

When a user clicks on a block from the Block view, he is redirected to the Block detail view (see Figure 5.9b). All transactions that are confirmed in this block are listed here. Transactions have inputs and output that are usually cryptocurrency address (input can also be a coin base). Every address in a transaction in link to the Address View (see Figure 5.9c). The Address view contains all transactions that address is part of as input or output.

The last view that we implemented in ExplorerGUI is Playground (see Figure 5.11). In this view, the user can write a query to the web editor and execute it. Syntax highlighting in editor is done by Prism¹³. Result of a user query is returned in JSON format bellow query.

5.2.2 ExplorerAPI

ExplorerAPI is a server-side application that provides simple HTTP API. After ExplorerAPI starts, it connects to the all enabled cryptocurrency databases (specified in `.env` file). Then it registers all three controllers (for blocks, transactions and addresses) and starts HTTP server on port described in `.env` file or 5000 if there is no specific port in the environment file. ExplorerAPI has only three endpoints (`block`, `address` and `transaction`). All these endpoints accepts query parameter (in case of `GET` method) or body parameter (`POST` method) in which query is described (see Figure 4.7a). This parameter can consist of array of conditions and filters, skip property and resolver. The parameter is translated to the Query system and result is return in JSON format (see Figure 4.7b). ExplorerAPI can accept multiple connections at once, because each request is processed in separated promise (task).

¹³<https://prismjs.com/>

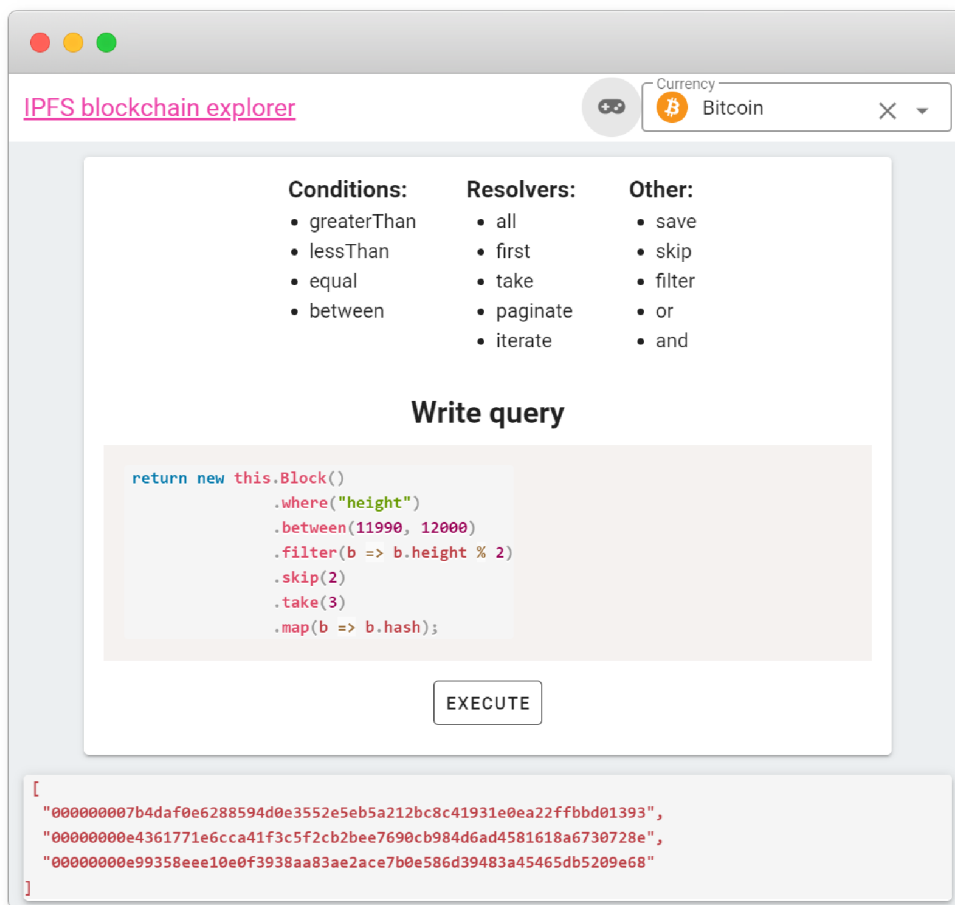


Figure 5.11: Playground view

Chapter 6

Testing and Benchmarking

In this chapter, we describe how our platform is verified and validated. Then, we measured the performance of the database system with different count of clients, and we compared results with other database systems.

6.1 Testing

Testing is an essential part of the development of any project. As part of this master thesis, we performed unit testing to ensure that every platform module works correctly.

6.1.1 Unit Testing

We tested all critical parts of the Explorer-core module through unit tests. We use `jest.js`¹ framework for writing tests. We test every operation on B-Tree index (add, remove, find, update), an append-only log that is used for database log and record log (add, merge, difference). Test for B-Tree insert operation can be seen in Figure 6.1. Tests for every application we have developed in this thesis can be run by command `npm run test`. Tests was automatically performed after each commit to the versioning system, thanks to continuous integration (CI) provided by GitHub². All applications are dockerized (can be run via Docker³), therefore making CI pipeline was easy.

6.2 Benchmarking

To test the performance of our platform, we have focused on the speed of transaction processing in the database system. We compared our database system to other widely used solutions – MySQL and PostgreSQL.

MySQL

MySQL is leading open-source multi-platform relational database system. Originally, MySQL was created by the Swedish company MySQL AB, now owned by Oracle Corporation. It can provide stable and robust database solution for application. Due to its easy

¹<https://jestjs.io/>

²<https://help.github.com/en/actions/building-and-testing-code-with-continuous-integration/about-continuous-integration>

³<https://www.docker.com/>

```

describe("Btree", () =>
{
  it("insert values to btree", async () =>
  {
    const t = new BTree();
    for (let i = 0; i < 100; i++)
    {
      await t.add(i, { name: "test object", value: i });
    }

    for (let i = 0; i < 100; i++)
    {
      expect(await t.find(i)).toStrictEqual({ name: "test object", value: i });
    }
    expect(await t.find(101)).toBe(null);
  });
});

```

Figure 6.1: Unit test for B-Tree insert operation.

deployment, performance and especially due to the fact that it is a freely distributable software, it has a high share in currently used databases. [24, 10]

PostgreSQL

PostgreSQL is object-relational widely used open-source database system. It is available for Linux, FreeBSD, OpenBSD, and Windows. PostgreSQL was formed in 1977. Project started in university of California and later on was developed commercially. To this date PostgreSQL has active development worldwide community. Main advantage of PostgreSQL is approach of object-relational model. This allows developers to store. It is designed to handle a range of workloads, from single machines to data warehouses or Web services with many concurrent users. [9]

Environment

We compare MySQL and PostgreSQL with our peer-to-peer database system. These database systems run virtualized in Docker containers on hardware with 4 cores processor Intel Pentium G4560 3.50GHz and 16 GB of RAM. We start clients that connect to these database systems and performs queries on different hardware. Then we measure average queries per seconds that clients execute. Clients are written in JavaScript with official node.js connectors for MySQL⁴ and PostgreSQL⁵ and are connected in the LAN with database systems over 1Gb Ethernet. Database structure that is used in benchmarks can be seen in Figure 6.2. Table `authors` contains 10,000 records and table `posts` contains 50,000 records (each author has 5 posts). Test query for each database system can be seen in Figure 6.3.

⁴<https://www.npmjs.com/package/mysql>

⁵<https://www.npmjs.com/package/postgres>

<pre>CREATE TABLE 'authors' ('id' int(11) PRIMARY KEY, 'first_name' varchar(50), 'last_name' varchar(50), 'email' varchar(100), 'birthdate' date, 'added' timestamp, PRIMARY KEY ('id'), UNIQUE KEY 'email' ('email')); CREATE TABLE 'posts' ('id' int(11) PRIMARY KEY, 'author_id' int(11), 'title' varchar(255), 'description' varchar(500), 'content' text, 'date' date, PRIMARY KEY ('id'));</pre> <p>(a) MySQL</p>	<pre>CREATE TABLE authors id serial PRIMARY KI first_name varchar(50), last_name varchar(50), email varchar(100), birthdate date, added timestamp); CREATE TABLE posts (id serial PRIMARY KI author_id integer, title varchar(255), description varchar(500) content text, date date);</pre> <p>(b) PostgreSQL</p>	<pre>class Author { @PrimaryKey() id: number; first_name: string; last_name: string; email: string; @Index() birthdate: Date; added: Date; } class Post { @PrimaryKey() id: number; author_id: number; title: string; description: string; content: string; date: Date; }</pre> <p>(c) Our database system</p>
--	--	---

Figure 6.2: Database structures for benchmarking

<pre>SELECT * FROM authors JOIN posts ON posts.author_id=authors.id WHERE authors.id= (SELECT FLOOR(RAND() * 10001 + 1));</pre> <p>(a) MySQL</p>	<pre>SELECT * FROM authors JOIN posts ON posts.author_id=authors.id WHERE authors.id= (SELECT floor(random() * 10001 -</pre> <p>(b) PostgreSQL</p>	<pre>new Author() .where("id") .equal(Math.floor(Math.random() * 10001) + 1) .first();</pre> <p>(c) Our database system</p>
--	--	---

Figure 6.3: Test query

Single Connected Client


With a single connected client, PostgreSQL can make more than 250 queries per second and MySQL 70 queries per second. Our database system is far behind with up to 20 queries per second (see Figure 6.4a). We can see that our database system builds cache, and performance increases by time. With only a single connected client, there is no advantage of a peer-to-peer network. CPU usage with a single connected client is for MySQL and PostgreSQL same – about 20% (see Figure 6.4b). For our database system, it is only 3% due to lots of input-output operations.

20 Connected Clients

With 20 connected clients, PostgreSQL performance dropped to 50 queries per seconds and MySQL to 10 queries per seconds. At the end of benchmarking, our database system is more powerful than PostgreSQL and MySQL, because data gets distributed to clients (see 6.5a). Our database system also requires less processor time thanks to peer-to-peer network. Data are not obtained only from the server (as opposed to MySQL and PostgreSQL) but also from other clients. Thanks to this, our platform only using less than 20% CPU power (see Figure 6.5b).

80 Connected Clients

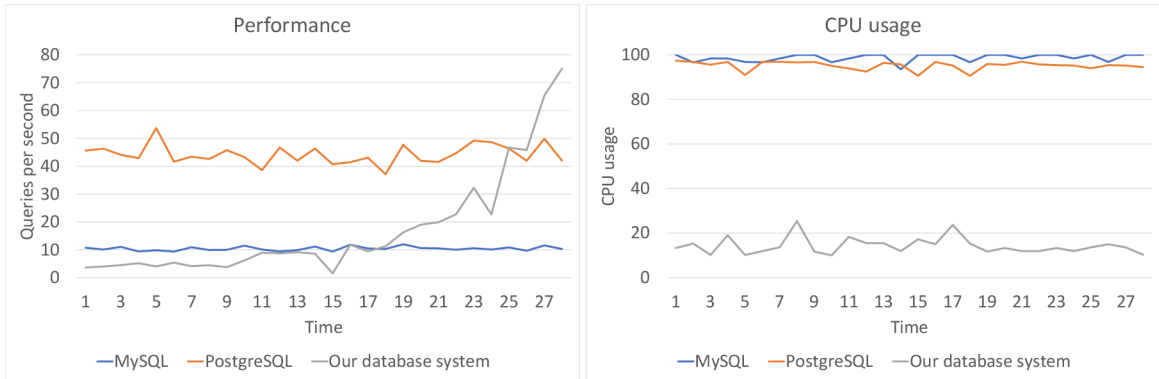
With 80 connected clients, our database system outperforms MySQL and PostgreSQL after few seconds. Every client of our database system performs 100 queries per second. PostgreSQL client performs only 10 queries per second and MySQL only 2 queries per second



(a) Performance

(b) CPU usage with

Figure 6.4: Single connected client benchmark plots



(a) Performance

(b) CPU usage with

Figure 6.5: 20 connected clients benchmark plots


(see 6.6a). CPU usage is in our database system also very low. In the beginning, when content is not distributed on the clients, the central node uses 60% of CPU, but it decreases over time (see 6.6b).

150 Connected Clients

We need to change the default PostgreSQL setting for maximum connections to connects more than 100 clients to the PostgreSQL. Our platform can handle 150 clients very well (opposite to MySQL and PostgreSQL). Each client connected to our database system can make up to 150 queries per second, and it increases by the time (see 6.7a). For the first seconds, CPU usage on the central node is high (near 100%), but it decreases over time (see 6.7b).

Benchmark Conclusion


We can see that with more than 20 connected clients, our database system became more effective than traditional ones. Thanks to IPFS and its peer-to-peer network, a content of the database is distributed on the clients. This reduces the load on the central node.



(a) Performance

(b) CPU usage with

Figure 6.6: 80 connected clients benchmark plots



(a) Performance

(b) CPU usage with

Figure 6.7: 150 connected clients benchmark plots

6.3 Comparing with Blockchain Explorers

Our blockchain exploration platform is unique and differs from traditional blockchain explorers in many ways. See the table 6.1 to compare traditional blockchain explorers with ours.

	Our platform	Traditional blockchain explorer (Blockbook, Insight)
Database system	Proprietary	Typically NoSQL database
Views for blocks, transactions and addresses	Yes	Yes
Complex queries support	Yes	No
Filtering	Yes	No, but filtering can be made by cycling over results
Sorting	Yes	No
Blockchain synchronization	Typically slower but it depends on how many Feeders are working on blockchain	Faster
Get block/transaction/address by hash	Slower	Faster, but may slow down if multiple nodes are connected
When more user is connected	Works faster	Works slower
Multiple coins support	Yes	Yes

Table 6.1: Comparing our platform with typical blockchain explorers

Chapter 7

Conclusion

The goal of this thesis was to create a platform that uses IPFS to explore blockchain. To achieve this, we needed to create our own decentralized and distributed database system that supports advanced queries and indexes on top of the IPFS. Our database system has optimized synchronization. As was described in Chapter 6.2, it is faster with every next connected peer. Explorer-core module in which database is implemented, works in a browser environment and desktop (via node.js) thanks to language choice (Typescript). We created several applications using this decentralized database. Namely Feeder that connects to the blockchain data source and stores it to IPFS. Secondly, ExplorerGUI is used as a presentation layer for users. Lastly, ExplorerAPI provides a simple interface for integrating with other applications.

The research section of this project focuses mainly on the IPFS principles (content-based addressing, object linking, and others). Selected cryptocurrencies for exploring are also briefly discussed. Considerable effort was devoted to designing the whole platform and creating a functional prototype.

Although the task of this work was to create a simple blockchain explorer, this project aroused great interest in the IPFS community, because our database system allows developers to create decentralized serverless applications with a relatively efficient database solution. Thanks to this work, it is possible to host entire information systems in IPFS. There are many benefits to using our database system for developers. The main advantages include the cost of hosting, as in our solution with a higher number of users, the load on the central peer decreases, as well as the availability of data even in the event of a failure of part of the peers. Our Query system is well adapted to today ORM frameworks and uses fluent query language and supports multiple conditions in a single query.

The result of this thesis is a distributed platform, that can parse blockchain data from a data source and then allows user to explore them. Compared to traditional blockchain explorer, our platform has minimal operation cost, and it can recover from a partial network failure.

In the future, it is possible to improve Query system by supporting more operators (`like`, `groupBy`). Other data connectors as Blockbook can be added to the Feeder (for example direct connection to the blockchain via full node). Other index structures can be implemented for better performance in some queries (for example trie¹ for prefix search).

¹<https://en.wikipedia.org/wiki/Trie>

Bibliography

- [1] BENET, J. IPFS - Content Addressed, Versioned, P2P File System. *CoRR*. 2014, abs/1407.3561. Available at: <http://arxiv.org/abs/1407.3561>.
- [2] BUTERIN, V. et al. A next-generation smart contract and decentralized application platform. *White paper*. 2014, vol. 3, no. 37.
- [3] CHACON, S. and STRAUB, B. *Pro git*. Apress, 2014. ISBN 1484200772.
- [4] CHOWDHURY, N. *Inside blockchain, bitcoin, and cryptocurrencies*. Boca Raton: CRC Press, 2020. ISBN 978-1-138-61815-2.
- [5] DASCANO, M. *Digibyte: An Easy Guide to Learning the Essentials*. 1stth ed. North Charleston, SC, USA: CreateSpace Independent Publishing Platform, 2018. ISBN 1725629178.
- [6] DAVIS, J. *The Crypto-Currency*. October 2011. [Online; Accessed 10.5.2020]. Available at: <https://www.newyorker.com/magazine/2011/10/10/the-crypto-currency>.
- [7] DHILLON, V., METCALF, D. and HOOPER, M. *Blockchain Enabled Applications: Understand the Blockchain Ecosystem and How to Make it Work for You*. Springer, 2017.
- [8] DIAS, D. and BENET, J. Distributed Web Applications with IPFS, Tutorial. In: BOZZON, A., CUDRE MAROUX, P. and PAUTASSO, C., ed. *Web Engineering*. Cham: Springer International Publishing, 2016, p. 616–619. ISBN 978-3-319-38791-8.
- [9] DRAKE, J. D. and WORSLEY, J. C. *Practical PostgreSQL*. „ O’Reilly Media, Inc.“, 2002. ISBN 1565928466.
- [10] DUBOIS, P. and FOREWORD BY WIDENIUS, M. *MySQL*. New riders publishing, 1999. ISBN 9780321833877.
- [11] HABER, S. and STORNETTA, W. S. How to time-stamp a digital document. In: Springer. *Conference on the Theory and Application of Cryptography*. 1990, p. 437–455.
- [12] JEPSON, C. DTB001: Decred Technical Brief. Available at <https://coss.io/documents/white-papers/decred.pdf> Additional information available at <https://www.decred.org>. 2015.
- [13] KALLE, K. *Big data in video games*. Lappeenranta, FI, 2017. Bachelor Thesis. Lappeenranta University of Technology, School of Business and Management, Computer Science. Available at: <http://lutpub.lut.fi/handle/10024/147666>.

- [14] LABS, P. Filecoin - A Decentralized Storage Network. 2017. [Online; Accessed 7.5.2020]. Available at: <https://filecoin.io/filecoin.pdf>.
- [15] LAMPORT, L. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*. 1978, vol. 21, no. 7, p. 558–565. Available at: <https://doi.org/10.1145/359545.359563>.
- [16] LAURENCE, T. *Blockchain for dummies*. John Wiley & Sons, 2019. ISBN 1119555019.
- [17] MARK ROBERT HENDERSON, S. P. *The OrbitDB Field Manual*. 2019. [Online; Accessed 13.5.2020]. Available at: <https://github.com/orbitdb/field-manual>.
- [18] MAYMOUNKOV, P. and MAZIERES, D. Kademlia: A peer-to-peer information system based on the xor metric. In: Springer. *International Workshop on Peer-to-Peer Systems*. 2002, p. 53–65.
- [19] NAKAMOTO, S. A peer-to-peer electronic cash system. *Bitcoin*. 2008. [Online; Accessed 21.4.2020]. Available at: <https://nakamotoinstitute.org/bitcoin/>.
- [20] NARAYANAN, A., BONNEAU, J., FELTEN, E., MILLER, A. and GOLDFEDER, S. *Bitcoin and cryptocurrency technologies: a comprehensive introduction*. Princeton University Press, 2016. ISBN 9780691171692.
- [21] NOETHER, S. Ring Signature Confidential Transactions for Monero. *IACR Cryptology ePrint Archive*. 2015, vol. 2015, p. 1098.
- [22] PICK, S. and HAGOPIAN. *A protocol and event-sourced database for decentralized user-siloed data*. 2019. Available at: <https://blog.textile.io/introducing-textiles-threads-protocol/>.
- [23] STONEBRAKER, M. and ROWE, L. A. The design of POSTGRES. *ACM Sigmod Record*. ACM New York, NY, USA. 1986, vol. 15, no. 2, p. 340–355.
- [24] SUEHRING, S. *MySQL Bible*. Wiley Publishing, Inc., 2002. ISBN 0764549324.
- [25] VAN SABERHAGEN, N. *CryptoNote v 2.0*. 2013. [Online; Accessed 20.5.2020]. Available at: <https://static.coinpaprika.com/storage/cdn/whitepapers/1611.pdf>.
- [26] WAYNER, P. *Digital Cash (2nd Ed.): Commerce on the Net*. USA: Academic Press Professional, Inc., 1997. ISBN 0127887725.
- [27] WOOD, G. et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*. 2014, vol. 151, no. 2014, p. 1–32.

Appendix A

Contents of the included storage media

The enclosed CD contains the following files:

- latex/ – source files of this thesis
- code/ – source code of all created applications
 - explorerCore/ – source code of the main module
 - explorerGUI/ – source code of GUI application
 - explorerAPI/ – source code of API application
 - feeder/ – source code of Feeder application
 - benchmark/ – source code of performed benchmarks
- thesis.pdf – PDF of this thesis

All source files are available under GitHub organization IPFSexplorer (<https://github.com/IPFSexplorer>)