

A general framework for blockchain analytics

Massimo Bartoletti, Stefano Lande, Livio Pompianu
University of Cagliari
Italy

Andrea Bracciali
University of Stirling
UK

Abstract

Modern cryptocurrencies exploit decentralised blockchains to record a public and unalterable history of transactions. Besides transactions, further information is stored for different, and often undisclosed, purposes, making the blockchains a rich and increasingly growing source of valuable information, in part of difficult interpretation. Many data analytics have been developed, mostly based on specifically designed and ad-hoc engineered approaches. We propose a general-purpose framework, seamlessly supporting data analytics on both Bitcoin and Ethereum – currently the two most prominent cryptocurrencies. Such a framework allows us to integrate relevant blockchain data with data from other sources, and to organise them in a database, either SQL or NoSQL. Our framework is released as an open-source Scala library. We illustrate the distinguishing features of our approach on a set of significant use cases, which allow us to empirically compare ours to other competing proposals, and evaluate the impact of the database choice on scalability.

CCS Concepts • Applied computing → Electronic commerce; Digital cash;

Keywords Blockchain, Bitcoin, Ethereum, Analytics

ACM Reference Format:

Massimo Bartoletti, Stefano Lande, Livio Pompianu and Andrea Bracciali. 2017. A general framework for blockchain analytics. In *SERIAL'17: Scalable and Resilient Infrastructures for distributed Ledgers, December 11–15, 2017, Las Vegas, NV, USA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3152824.3152831>

1 Introduction

The last few years have witnessed a steady growth in interest on blockchains, driven by the success of Bitcoin and, more recently, of Ethereum. This has fostered the research on several aspects of blockchain technologies, from their

theoretical foundations – both cryptographic [5, 9] and economic [17, 29] – to their security and privacy [1, 6, 10, 13, 20].

Among the research topics emerging from blockchain technologies, one that has received major interest is the analysis of the data stored in blockchains. Indeed, the two main blockchains contain several gigabytes of data (~130GB for Bitcoin, ~300GB for Ethereum), that only in part are related to currency transfers. Developing analytics on these data allows us to obtain several insights, as well as economic indicators that help to predict market trends.

Many works on data analytics have been recently published, addressing anonymity issues, e.g. by de-anonymising users [19, 20, 25, 27], clustering transactions [11, 30], or evaluating anonymising services [22]. Other analyses have addressed criminal activities, e.g. by studying denial-of-service attacks [2, 32], ransomware [15], and various financial frauds [23, 24, 31]. Many statistics on Bitcoin and Ethereum exist, measuring e.g. economic indicators [16, 28], transaction fees [21], the usage of metadata [3], *etc.*

A common trait of these works is that they create *views* of the blockchain which contain all the data needed for the goals of the analysis. In many cases, this requires to combine data *within* the blockchain with data from the *outside*. These data are retrieved from a variety of sources, e.g. blockchain explorers, wikis, discussion forums, and dedicated sites (see Table 1 for a brief survey). Despite such studies share several common operations, e.g., scanning all the blocks and the transactions in the blockchain, converting the value of a transaction from bitcoins to *USD, etc.*, researchers so far tended to implement ad-hoc tools for their analyses, rather than reusing standard libraries. Further, most of the few available tools have limitations, e.g. they feature a fixed set of analytics, or they do not allow to combine blockchain data with external data, or they are not amenable to be updated. The consequence is that the same functionalities have been implemented again and again as new analytics have been developed, as witnessed by Table 1.

In this context, we believe that the introduction of an efficient, modular and general-purpose abstraction layer to manage internal and external information is key for blockchain data analytics, along the lines of the software engineering best practices of *reuse*.

Contributions. The main contribution of this paper is a framework to create general-purpose analytics on the blockchains of Bitcoin and Ethereum. The design of our tool is

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SERIAL'17, December 11–15, 2017, Las Vegas, NV, USA

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-5173-7/17/12...\$15.00
<https://doi.org/10.1145/3152824.3152831>

Analysis goal	Gathered data	Sources
Anonymity	Transactions graph	bitcoind [19, 20, 22, 27, 30], forum.bitcoin.org [27]
	OP_RETURN metadata	bitcoind [22]
Market analytics	IP addresses	bitcoin faucet [27], blockchain.info [22]
	address tags	blockchain.info [19, 20, 30], bitcointalk.org [19, 20, 30]
Cyber-crime	address tags	bitcoin-otc.com [30], bitfunder.org [30]
	Transactions graph	bitcoind [16], blockexplorer.com [28]
Metadata	IP addresses	blockchain.info, ipinfo.io [16]
	address tags	blockchain.info [16]
Transaction fees	trade data	bitcoincharts.com [16]
	Transactions graph	bitcoind [2, 31, 32], blockchain.info [15, 23], Bitcore [4]
Transaction fees	mempool	bitcoind [2]
	unconfirmed transactions	bitcoind [2]
Transaction fees	no longer online services	archive.org [31, 32]
	list of DDoS attacks	bitcointalk.org [32]
Transaction fees	mining pools	blockchain.info, bitcoin wiki [32]
	trades on assets/services	bitcoin wiki [32]
Transaction fees	list of fraudulent services	bitcointalk.org [15, 31], badbitcoin.org [31], cryptohyips.com [31]
	address tags	blockchain.info [31]
Transaction fees	exchange rate	bitcoincharts.com [15, 31], quandl.com [15]
	OP_RETURN transactions	bitcoind [3]
Transaction fees	OP_RETURN identifiers	kaiko.com, oreturn.org, bitcoin wiki [3]
	Transactions graph	bitcoind [21]
Transaction fees	exchange rate	coindesk.com [21]
	mining pools	blockchain.info [21]

Table 1. Data gathered by various blockchain analyses.

based on an exhaustive survey of the literature on the analysis of blockchains. The results of our survey, summarized in Table 1, highlight the need to process external data besides those already present on the blockchain. To this purpose, the workflow supported by our tool consists of two steps: (i) we construct a *view* of the blockchain, also containing the needed external data, and we save it in a database; (ii) we analyse the view by using the query language of the DBMS. The first step is supported by a new Scala library. Distiguishably, we allow views to be organised either as a MySQL database, or a MongoDB collection. Our library supports the most commonly used external data, e.g. exchange rates, address tags, protocol identifiers, and can be easily extended by linking the relevant data sources. We evaluate the effectiveness of our framework by means of a set of paradigmatic use cases, which we distribute, together with the source code of our library, under an open source license¹. We exploit our use cases to evaluate the performance of SQL vs. NoSQL databases for storing and querying blockchain views. As a byproduct of our study, we provide a qualitative comparison of the other tools for general-purpose blockchain analytics.

2 Creating blockchain analytics

We illustrate our framework through some case studies, which, for uniformity, have been developed for the Bitcoin case. We refer to our github repository for some Ethereum examples. Our library APIs provide the following Scala classes to represent the primitive entities of the blockchain:

- **BlockchainLib**: main library class. It provides the `getBlockchain` method, to iterate over `Block` objects.
- **Block**: contains a list of transactions, and some block-related attributes (e.g., block hash and creation time).
- **Transaction**: contains various related attributes (e.g., transaction hash and size).

The library constructs the above-mentioned Scala objects by scanning a local copy of the blockchain. It uses the client, either Bitcoin Core or Parity, to have a direct access to the

¹<https://github.com/bitbart/blockchain-analytics-api>

```

1 object MyBlockchain {
2   def main(args: Array[String]): Unit = {
3
4     val blockchain = BlockchainLib.getBitcoinBlockchain(
5       new BitcoinSettings("user", "password", "8332",
6         MainNet))
7
8     val mongo = new DatabaseSettings("myDatabase",
9       MongoDB, "user", "password")
10    val myBlockchain = new Collection("myBlockchain",
11      mongo)
12
13    blockchain.end(473100).foreach(block => {
14      block.bitcoinTx.foreach(tx => {
15        myBlockchain.append(List(
16          ("txHash", tx.hash),
17          ("blockHash", block.hash),
18          ("date", block.date),
19          ("inputs", tx.inputs),
20          ("outputs", tx.outputs)
21        ))
22      })
23    })
24  }
25 }

```

Figure 1. A basic view of the blockchain.

blocks, exploiting the provided indices. For Bitcoin, it uses the BitcoinJ library as a basis to represent the various kinds of objects, while for Ethereum it uses suitable Scala representations. The APIs allow constructed objects to be exported as MongoDB documents or MySQL records. In MongoDB (a widespread non-relational DBMS) a database is a set of *collections*, each of them containing *documents*. Documents are lists of pairs (k, v) , where k is a string (called *field name*), and v is either a value or a MongoDB document. Conversely, MySQL implements the relational model, and represents an objects as a record in a table. In Sections 2.1 to 2.4 we develop a series of analytics on Bitcoin. Full Scala code which builds the needed blockchain views, queries, and analysis results can be found in the GitHub repository of the project¹.

2.1 A basic view of the Bitcoin blockchain

Since all the analyses shown in Table 1 explore the transaction graph (e.g. they investigate output values, timestamps, metadata, etc.), our first case study focusses on a basic view of the Bitcoin blockchain containing no external data. The documents in the resulting collection represent transactions, and they include: (i) the transaction hash; (ii) the hash of the enclosing block; (iii) the date in which the block was appended to the blockchain; (iv) the list of transaction inputs and outputs.

We show in Figure 1 how to exploit our Bitcoin Analytics APIs to construct this collection. Lines 1–2 are standard Scala instructions to define the main function. The object `blockchain` constructed at line 4 is a handle to the Bitcoin blockchain. At line 5 we setup the connection to Bitcoin Core, by providing the needed parameters (user, password, and port), and by indicating that we want to use the main network (alternatively, the parameter `TestNet` allows to use the test network). At line 6 we setup the connection to MongoDB (alternatively, the parameter `MySQL` allows to use `MySQL`). Since lines 1–6 are similar for all our case studies, for the

```

1 val opReturnOutputs = new Collection("opReturn", mongo)
2
3 blockchain.start(290000).end(473100).foreach(block => {
4   block.bitcoinTxs.foreach(tx => {
5     tx.outputs.foreach(out => {
6       if(out.isOpreturn()) {
7         opReturnOutputs.append(List(
8           ("txHash", tx.hash),
9           ("date", block.date),
10          ("protocol", OpReturn.getApplication(out.
11            outScript.toString)),
12          ("metadata", out.getMetadata())
13        )))))))

```

Figure 2. Exposing OP_RETURN metadata.

sake of brevity we will omit them in the subsequent listings. We declare the target collection `myBlockchain` at line 7.

At this point, we start navigating the blockchain (from the origin block up to block number 473100) to populate the collection. To do that we iterate over the blocks (line 9) (note that `b => { . . . }` is an anonymous function, where `b` is a parameter, and `{ . . . }` is its body), and for each block we iterate over its transactions (at line 10). For each transaction we append a new document to `myBlockchain` (lines 11-16). This document is a set of fields of the form (k, v) , where k is the field name, and v is the associated value. For instance, at line 12 we stipulate that the field `txHash` will contain the hash of the transaction, represented by `tx.hash`. This value is obtained by the API `BitcoinTransaction`.

Running this piece of code results in a view, which we can process to obtain several standard statistics, like e.g. the number of daily transactions, their average value, the largest recent transactions, *etc.*²

2.2 Analysing OP_RETURN metadata

Besides being used as a cryptocurrency, Bitcoin allows for appending a few bytes of metadata to transaction outputs. This is done preeminently through the OP_RETURN operator. Several protocols exploit this feature to implement blockchain-based applications, like e.g. digital assets and notarization services [3].

We now construct a view of the blockchain which exposes the protocol metadata. More specifically, the entries of view represent transaction outputs, and are composed of: (i) the hash of the transaction containing the output; (ii) the date in which the transaction has been appended to the blockchain; (iii) the name of the protocol that produced the transaction; (iv) the metadata contained in the OP_RETURN script. Our API supports the creation of this view by providing the method `OpReturn.getApplication` (used at line 10 in Figure 2), which takes as input a piece of metadata, and returns the name of the associated protocol. This is inferred by the results of the analysis in [3].

²Note that one could also perform these queries during the construction of the view. However, this would not be convenient in general, since — as we will see also in the following case studies — many relevant queries can be performed on the same view.

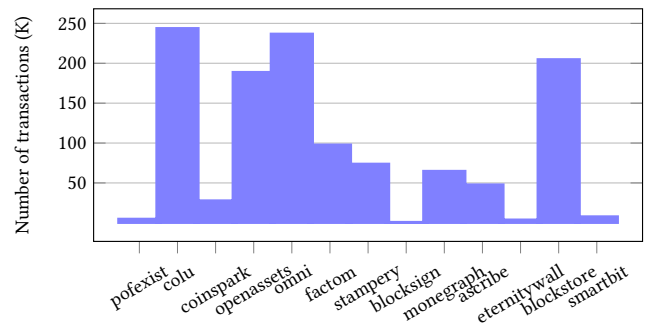
```

1 val blockchain = BlockchainLib.getBitcoinBlockchain(new
2   BitcoinSettings("user", "password", "8332", MainNet, true))
3 val mongo = new DatabaseSettings("myDatabase", MongoDB, "user",
4   "password")
5 val txWithFees = new Collection("txWithFees", mongo)
6
7 blockchain.end(473100).foreach(block => {
8   block.bitcoinTxs.foreach(tx => {
9     txWithFees.append(List(
10      ("blockHash", block.hash),
11      ("txHash", tx.hash),
12      ("fee", tx.getInputsSum() - tx.getOutputsSum()),
13      ("date", block.date),
14      ("rate", Exchange.getRate(block.date))
15    ))))

```

Figure 3. Exposing transaction fees.

The obtained view can be used to perform various analyses. For instance, we show the number of transactions associated with the most used protocols (only those with at least 1000 transactions). The protocol with the highest number of transactions is Colu, which is used to certify and transfer the ownership of physical assets. The second most used protocol is Omni, followed by Blockstore, a key-value store upon which other protocols are based.



2.3 Transaction fees

In this section we study *transaction fees*, which are earned by miners when they append a new block to the blockchain. Each transaction in the block pays a fee, which in Bitcoin is defined as the difference between its input and output values. While the values of outputs are stored explicitly in the transaction, those of inputs are not: to obtain them, one must retrieve from a past block the transaction that is redeemed by the input. This can be obtained through a “deep” scan of the blockchain, which is featured by our library.

We show in Figure 3 how to construct a MongoDB collection which contains, for each transaction: (i) the hash of the enclosing block; (ii) the transaction hash; (iii) the fee; (iv) the date in which the transaction was appended to the blockchain; (v) the exchange rate between *BTC* and *USD* in such date. Since exchange rates are not stored in the Bitcoin blockchain, we resort to an external source (Coindesk) to obtain them.

The extra parameter `true` in the `BitcoinSettings` constructor (missing in the previous example), triggers the “deep” scan. When scanning the blockchain in this way, the library maintains a map which associates transaction outputs to

their values, and inspects this map to obtain the value of inputs³. The methods `getInputsSum` (resp., `getOutputsSum`) at line 10 returns the sum of the values of the inputs (resp., the outputs) of a transaction.

The obtained collection can be used to perform several standard statistics, e.g. the daily total transaction fees, the average fee, the percentage earned by miners from transaction fees, *etc.* We analyse the so-called *whale transactions* [14], which pay a unusually high fee to miners.

To obtain the whale transactions, we first compute the average \bar{x} and standard deviation σ of the fees in all transactions: in *USD*, we have $\bar{x} = 0.41$, $\sigma = 12.09$. Then, we define whale transactions as those which pay a fee greater than $\bar{x} + 2\sigma = 24.58$ *USD*. Overall we collect 242, 839 whale transactions; those with biggest fee are displayed below.

Fee (USD)	Date	Transaction hash
136243.37	2016-04-26 14:15:22	cc455ae816e6cdfb58d54e35d4f46d860047458eacfc7405dc634631c570d
56493.50	2017-01-04 20:01:28	d38bd67153d774a7dab80a055cb52571aa85f6cac8f35f936c4349ca308e6380
39502.15	2017-05-31 14:28:51	cb95ab3aef378c14bc59d0db682d96202b981c1f8fad7d66e23e0be06f2a00c4
25095.71	2017-05-31 14:28:51	8e12a1aba87e4657f5fabec1121ed57f706805ad6d4ffe886cfce78596bd9b75
23518.00	2013-08-28 10:45:17	4ed20e0768124bc67dc684d57941be1482ccdaa45dad64be12afba8c8554537

2.4 Address tags


The webpage `blockchain.info/tags` hosts a list of associations between Bitcoin addresses and *tags* which briefly describe their usage. Table 1 shows that address tags are widely adopted. Cyber-crime studies retrieve addresses tagged as scam or ransomware on forums. Market analyses exploit tags for recognising addresses of business services. Anonymity studies tag the addresses that seem to belong to the same entity. In this section we construct a blockchain view where outputs are associated with the tags of the address which can redeem them (we discard the outputs with untagged addresses). More specifically, we construct a view representing transaction outputs that contains: (i) hash of the enclosing transaction; (ii) the date in which the transaction has been appended to the blockchain; (iii) the output value (in *BTC*); (iv) the address receiving the payment; (v) the tag associated to the address.

To perform the analysis, we query `blockchain.info` to retrieve the tags associated to addresses. Using the obtained view, one can aggregate transactions on different business levels [16] to obtain statistics about the total number of transactions, the amount of *BTC* exchanged, the geographical distributions of tagged service, *etc.* In particular, we aggregate all addresses whose tag starts with *SatoshiDICE*, and then we measure the number of daily transactions which send *BTC* to one of these addresses. The diagram below shows the results of this analysis. The fall in the number of transactions at the start of 2015 may be due to the fact that *SatoshiDICE* is using untagged addresses.

³Since inputs can only redeemed transactions on past blocks, the map always contains the required output. Although coinbase inputs do not have a value in the map, we calculate their value using the total fees of the current block and the block height (reward is halved each 210,000 blocks).

Case	MongoDB			MySQL		
	Create	Query	Size	Create	Query	Size
Basic	9 h	2860 s	300 GB	9 h	3.5 h	266 GB
Metadata	2 h	0.5 s	0.5 GB	1.4 h	2.5 s	0.5 GB
Fees	9 h	448 s	51 GB	8.5 h	614 s	43.5 GB
Tags	4 h	1.8 s	0.8 GB	2.3 h	2.7 s	0.6 GB

Table 2. Performance evaluation of our framework.



3 Implementation and validation

We implement the Ethereum-side of our library by exploiting Parity, queried by means of the `web3j` library. Bitcoin data is provided by both `BitcoinJ` and the RPC interface of Bitcoin Core. While `BitcoinJ` APIs only allow the programmer to retrieve a block by its hash, Bitcoin Core’s interface exposes calls to do so by its height on the chain. Furthermore, `BitcoinJ` block objects do not carry information about block height and the hash of the next block (they only have backward pointers, as defined in the blockchain), which can be fetched by using Bitcoin Core. Our APIs allow to navigate blockchains. Particularly, in the Bitcoin case, we do this by iterating over these steps: (i) get the hash h of the block of height i , by using Bitcoin Core; (ii) get the block with hash h , by using `BitcoinJ`; (iii) increment i . By default, the loop starts from 0 and stops at the last block. The methods `blockchain.start(i)`, and `blockchain.end(j)` allow to scan an interval of blockchains, as shown in Section 2.2. We write the SQL queries exploiting `ScalikeJDBC`, a SQL-based DB access library for Scala. `ScalikeJDBC` provides also a DSL for writing SQL queries.

We carry out our experiments using consumer hardware, i.e. a PC with a quad-core Intel Core i5-4440 CPU @ 3.10GHz, equipped with 32GB of RAM and 2TB of hard disk storage. All the experiments scan the Bitcoin blockchain from the origin block up to block number 473100 (added on 2017/06/27). Table 2 displays a comparison of the size of each view, and the time required to create and query it.

Note that the size of the blockchain view constructed in `Basic` (Section 2.1) is more than twice than the current Bitcoin blockchain. This is because, while Bitcoin stores scripts in binary format, our library writes them as strings, so to allow for constructing indices and performing queries on scripts. Moreover, the SQL query in `Basic` is particularly slow because of the join operations it performs. Indeed, SQL

Tool	Blockchain	Database	Schema	Ext. data	Updated
blockparser	BTC	RAM-only	Custom	Custom	2015-12
rusty-blockparser	BTC	SQL, CSV	Fixed	Custom	2017-09
blockchainsql.io	BTC	SQL	Fixed	None	N/A
BlockSci	BTC	RAM-only	Custom	Custom	2017-09
python-parser	BTC	None	None	Custom	2017-05
Our framework	BTC, ETH	MySQL, MongoDB	Custom	Custom	2017-09

Table 3. General-purpose blockchain analytics frameworks.

and MongoDB query times are quite similar in all the other cases, where no join operation is required.

4 Comparison with related tools

We now compare other general-purpose blockchain analysis tools with ours. Table 3 summarises the comparison, focussing on the target blockchain, the DBMS used, the support for creating custom a schema, and for embedding external data. The rightmost column indicates the date of the most recent commit in the repository. Note that all the tools which support Bitcoin also work on Bitcoin-based *altcoins*.

The projects `blockparser` and `rusty-blockparser` allow one to perform full scans of the blockchain, and to define custom listeners which are called each time a new block or transaction is read. Unlike our library, these tools offer limited built-in support for combining blockchain and external data. The website `blockchainsql.io` has a GUI through which one can write and execute SQL queries on the Bitcoin blockchain. This is the only tool, among those mentioned in Table 3, that does not need to store a local copy of the blockchain. A drawback is that the database schema is fixed, hence it is not possible to use it for analytics which require external data. While the other tools store results on secondary memory, `blockparser` and `BlockSci` keep all the data in RAM. Although this speeds up the execution, it demands for “big memory servers”, since the size of the blockchains of both Bitcoin and Ethereum has largely surpassed the amount of RAM available on consumer hardware. Note instead that the disk-based tools also work on consumer hardware. Some low-level optimizations, combined with an in-memory DBMS, help [12] to overwhelm the performance of the disk-based tools. Unlike the other tools, [12] provides also data about transactions broadcast on the peer-to-peer network.

Remarkably, as far as we know none of the analyses mentioned in Table 1 uses the general-purpose tools in Table 3. Instead, several of them acquire blockchain raw data by using Bitcoin Core⁴ (the reference Bitcoin client), and encapsulate them into Java objects with the BitcoinJ APIs before processing. However, neither Bitcoin Core nor BitcoinJ are natural tools to analyse the blockchain: the intended use of BitcoinJ is to support the development of wallets, and so it only gives direct access to blocks and transactions from their *hash*, but it does not allow to perform forward scans of the blockchain.

⁴<https://bitcoin.org/en/bitcoin-core>. Another popular tool for accessing the blockchain was `Bitcointools` (<https://github.com/gavinandresen/bitcointools>), but it seems no longer available.

On the other hand, Bitcoin Core would provide the means to scan the blockchain, but this requires expertise on its low-level RPC interface, and even doing so would result in raw pieces of JSON data, without any abstraction layer.

A precise comparison of the performance of these tools against ours is beyond the goals of this paper. The performance analysis in Table 2 is a first step towards the definition of a suite of benchmarks for evaluating blockchain parsers.

5 Conclusions and future work

We have presented a framework for developing general-purpose analytics on the blockchains of Bitcoin and Ethereum. Its main component is a Scala library which can be used to construct views of the blockchain, possibly integrating blockchain data with data retrieved from external sources. Blockchain views can be stored as SQL or NoSQL databases, and can be analysed by using their query languages. Our experiments confirmed the effectiveness and generality of our approach, which uniformly comprises in a single framework several use cases addressed by various ad-hoc approaches in literature. Indeed, the expressiveness of our framework overcomes that of the closer proposals in the built-in support for external data, and the support of different kinds of databases and blockchains. Importantly, coming in the form of an open source library for a mainstream language, our framework is amenable of being validated and extended by a community effort, following reuse best practices.

Specifically, on the comparison of SQL vs NoSQL, our experiments did not highlight significant differences in the complexity of writing and executing queries in the two languages. Instead, we observed that the schema-less nature of NoSQL databases simplifies the Scala scripts. From Table 2 we see that both creation and query time are comparable as order of magnitude. As already discussed in Section 3, the difference in the execution time of queries is due to join operations in SQL. A more accurate analysis, carried over a larger benchmark, is scope for future work. Anyway, it is worth recalling that the goal of our proposal is provide to the final user the flexibility to choose the preferred database, rather than ascertain an idea of best-fit-for-all in the choice.

Although our framework is general enough to cover most of the analyses in Table 1, it has some limitations that can be overcome with future extensions. In particular, some analyses addressing e.g. information propagation, forks and attacks [7, 8, 18, 26] require to gather data from the underlying peer-to-peer network. To support this kind of analyses one has to run a customized node (either of Bitcoin or Ethereum). Such an extension would also be helpful to obtain on-the-fly updates of the analyses.

Acknowledgments. This work is partially supported by Aut. Reg. of Sardinia P.I.A. 2013 “NOMAD”, This paper is based upon work from COST Action IC1406 cHiPSET, supported by COST (European Cooperation in Science and Technology).

References

- [1] Elli Androulaki, Ghassan Karame, Marc Roeschlin, Tobias Scherer, and Srdjan Capkun. 2013. Evaluating User Privacy in Bitcoin. In *Financial Cryptography and Data Security (LNCS)*, Vol. 7859. Springer, 34–51. https://doi.org/10.1007/978-3-642-39884-1_4
- [2] Khaled Baqer, Danny Yuxing Huang, Damon McCoy, and Nicholas Weaver. 2016. Stressing Out: Bitcoin “Stress Testing”. In *Financial Cryptography Workshops (LNCS)*, Vol. 9604. Springer, 3–18.
- [3] Massimo Bartoletti and Livio Pompianu. 2017. An analysis of Bitcoin OP_RETURN metadata. In *Financial Cryptography Workshop (LNCS)*, Vol. 10323. Springer.
- [4] Stefano Bistarelli and Francesco Santini. 2017. Go with the -Bitcoin-Flow, with Visual Analytics. In *ARES*. 38:1–38:6. <https://doi.org/10.1145/3098954.3098972>
- [5] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. 2015. SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In *IEEE S & P*. 104–121. <https://doi.org/10.1109/SP.2015.14>
- [6] Joseph Bonneau, Arvind Narayanan, Andrew Miller, Jeremy Clark, Joshua A. Kroll, and Edward W. Felten. 2014. Mixcoin: Anonymity for Bitcoin with Accountable Mixes. In *Financial Cryptography and Data Security (LNCS)*, Vol. 8437. Springer, 486–504. https://doi.org/10.1007/978-3-662-45472-5_31
- [7] Christian Decker and Roger Wattenhofer. 2013. Information propagation in the Bitcoin network. In *P2P*. IEEE, 1–10. <https://doi.org/10.1109/P2P.2013.6688704>
- [8] Joan Antoni Donet Donet, Cristina Pérez-Solà, and Jordi Herrera-Joancomartí. 2014. The Bitcoin P2P Network. In *Financial Cryptography Workshops (LNCS)*, Vol. 8438. Springer, 87–102. https://doi.org/10.1007/978-3-662-44774-1_7
- [9] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. 2015. The Bitcoin Backbone Protocol: Analysis and Applications. In *EUROCRYPT (LNCS)*, Vol. 9057. Springer, 281–310. https://doi.org/10.1007/978-3-662-46803-6_10
- [10] Arthur Gervais, Ghassan O. Karame, Karl Wüst, Vasileios Glykantzis, Hubert Ritzdorf, and Srdjan Capkun. 2016. On the Security and Performance of Proof of Work Blockchains. In *ACM SIGSAC Conference on Computer and Communications Security*. ACM, 3–16. <https://doi.org/10.1145/2976749.2978341>
- [11] Martin Harrigan and Christoph Fretter. 2016. The Unreasonable Effectiveness of Address Clustering. In *UIC/ATC/ScalCom/CBDCom/IoP/SmartWorld*. IEEE, 368–373.
- [12] Harry Kalodner, Steven Goldfeder, Alishah Chator, Malte Möser, and Arvind Narayanan. 2017. BlockSci: Design and applications of a blockchain analysis platform. *arXiv preprint arXiv:1709.02489* (2017).
- [13] Ghassan O. Karame, Elli Androulaki, Marc Roeschlin, Arthur Gervais, and Srdjan Capkun. 2015. Misbehavior in Bitcoin: A Study of Double-Spending and Accountability. *ACM Trans. Inf. Syst. Secur.* 18, 1 (2015), 2. <https://doi.org/10.1145/2732196>
- [14] Kevin Liao and Jonathan Katz. 2017. Incentivizing Blockchain Forks via Whale Transactions. In *Financial Cryptography Workshop (LNCS)*, Vol. 10323. Springer.
- [15] Kevin Liao, Ziming Zhao, Adam Doupe, and Gail-Joon Ahn. 2016. Behind closed doors: measurement and analysis of CryptoLocker ransoms in Bitcoin. In *APWG Symp. on Electronic Crime Research (eCrime)*. IEEE, 1–13. <https://doi.org/10.1109/ECRIME.2016.7487938>
- [16] Matthias Lischke and Benjamin Fabian. 2016. Analyzing the Bitcoin network: The first four years. *Future Internet* 8, 1 (2016), 7.
- [17] Loi Luu, Ratul Saha, Inian Parameshwaran, Prateek Saxena, and Aquinas Hobor. 2015. On Power Splitting Games in Distributed Computation: The Case of Bitcoin Pooled Mining. In *IEEE Computer Security Foundations Symposium*. IEEE, 397–411. <https://doi.org/10.1109/CSF.2015.34>
- [18] Patrick McCorry, Siamak F. Shahandashti, and Feng Hao. 2016. Refund Attacks on Bitcoin’s Payment Protocol. In *Financial Cryptography and Data Security (LNCS)*, Vol. 9603. Springer, 581–599. https://doi.org/10.1007/978-3-662-54970-4_34
- [19] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. 2013. A fistful of bitcoins: characterizing payments among men with no names. In *Internet Measurement Conference*. ACM, 127–140. <https://doi.org/10.1145/2504730.2504747>
- [20] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. 2016. A fistful of Bitcoins: characterizing payments among men with no names. *Commun. ACM* 59, 4 (2016), 86–93. <https://doi.org/10.1145/2896384>
- [21] Malte Möser and Rainer Böhme. 2015. Trends, tips, tolls: A longitudinal study of Bitcoin transaction fees. In *Financial Cryptography Workshops (LNCS)*, Vol. 8976. Springer, 19–33. <https://doi.org/10.1007/978-3-662-48051-9>
- [22] Malte Möser and Rainer Böhme. 2017. Anonymous Alone? Measuring Bitcoin’s Second-Generation Anonymization Techniques. In *EuroS&P Workshops*. 32–41. <https://doi.org/10.1109/EuroSPW.2017.48>
- [23] M. Möser, R. Böhme, and D. Breuker. 2013. An inquiry into money laundering tools in the Bitcoin ecosystem. In *APWG Symp. on Electronic Crime Research (eCrime)*. IEEE, 1–14. <https://doi.org/10.1109/eCRS.2013.6805780>
- [24] Malte Möser, Rainer Böhme, and Dominic Breuker. 2014. Towards Risk Scoring of Bitcoin Transactions. In *Financial Cryptography Workshops (LNCS)*, Vol. 8438. Springer, 16–32. https://doi.org/10.1007/978-3-662-44774-1_2
- [25] Micha Ober, Stefan Katzenbeisser, and Kay Hamacher. 2013. Structure and Anonymity of the Bitcoin Transaction Graph. *Future Internet* 5, 2 (2013), 237–250. <https://doi.org/10.3390/fi5020237>
- [26] Giuseppe Pappalardo, Tiziana di Matteo, Guido Caldarelli, and Tomaso Aste. 2017. Blockchain Inefficiency in the Bitcoin Peers Network. *CoRR* abs/1704.01414 (2017). <http://arxiv.org/abs/1704.01414>
- [27] Fergal Reid and Martin Harrigan. 2013. An analysis of anonymity in the Bitcoin system. In *Security and privacy in social networks*. Springer, 197–223. https://doi.org/10.1007/978-1-4614-4139-7_10
- [28] Dorit Ron and Adi Shamir. 2013. Quantitative analysis of the full Bitcoin transaction graph. In *Financial Cryptography and Data Security (LNCS)*, Vol. 7859. Springer, 6–24. <https://doi.org/10.1007/978-3-642-39884-1>
- [29] Okke Schrijvers, Joseph Bonneau, Dan Boneh, and Tim Roughgarden. 2016. Incentive Compatibility of Bitcoin Mining Pool Reward Functions. In *Financial Cryptography and Data Security (LNCS)*, Vol. 9603. Springer, 477–498. https://doi.org/10.1007/978-3-662-54970-4_28
- [30] Michele Spagnuolo, Federico Maggi, and Stefano Zanero. 2014. Bitiodine: Extracting intelligence from the Bitcoin network. In *Financial Cryptography and Data Security (LNCS)*, Vol. 8437. Springer, 457–468.
- [31] Marie Vasek and Tyler Moore. 2015. There’s No Free Lunch, Even Using Bitcoin: Tracking the Popularity and Profits of Virtual Currency Scams. In *Financial Cryptography and Data Security (LNCS)*, Vol. 8975. Springer, 44–61. https://doi.org/10.1007/978-3-662-47854-7_4
- [32] Marie Vasek, Micah Thornton, and Tyler Moore. 2014. Empirical Analysis of Denial-of-Service Attacks in the Bitcoin Ecosystem. In *Financial Cryptography Workshops (LNCS)*, Vol. 8438. Springer, 57–71. https://doi.org/10.1007/978-3-662-44774-1_5