Evolution of the Bitcoin Address Graph

An Exploratory Longitudinal Study

Erwin Filtz, Axel Polleres
Institute for Information Business
Vienna University of Economics and Business
Vienna, Austria
{firstname.lastname}@wu.ac.at

Roman Karl, Bernhard Haslhofer
Digital Insight Lab
Austrian Institute of Technology
Vienna, Austria
{firstname.lastname}@ait.ac.at

Abstract - **Bitcoin is a decentralized virtual currency, which can be used to execute pseudo-anonymous payments globally within a short period of time and comparably low transaction costs. In this paper, we present initial results of a longitudinal study conducted over the Bitcoin address graph, which contains all addresses and transactions from the beginning of Bitcoin in January 2009 until 31st of August 2016. Our analysis reveals a highly-skewed degree distribution with a small number of outliers and illustrates that the entire graph is expanding rapidly. Furthermore, it demonstrates the power of address clustering heuristics for identifying real-world actors, who prefer to use Bitcoin for transferring rather than storing value. We believe that this paper provides novel insight into virtual currency ecosystems, which can inform the design of future analytics methods and infrastructures.**

Keywords—Bitcoin, network analytics, virtual currencies

I. Introduction

Bitcoin [1] is the most prominent representative of decentralized, unregulated virtual currencies, which are based on cryptographic technologies – also known as "cryptocurrencies". In contrast to other fiat currencies (e.g. EUR, USD), such currencies have no pre-assumed identities, are not controlled by any central authorities, but are organized as a peer-to-peer network. Furthermore, all executed transactions are stored in a public, distributed ledger called the Bitcoin *blockchain*.

While the public ledger provides a high level of transparency on past transactions, it does not explicitly reveal details about real-world actors involved as senders or receivers of financial transactions. A single *transaction* is represented by a list of inputs pointing back to outputs of previous transactions and a list of outputs, each reflecting a certain Bitcoin value that has been transferred to some specific recipient's *address*. A Bitcoin address is an alphanumeric string derived from the public key of an asymmetric key pair generated by some Bitcoin user. Every user can hold multiple key-pairs (and addresses) in a so-called "*wallet*", and is encouraged to use a new address for each transaction to increase the level of anonymity.

The design of Bitcoin implies that the balance of an address is not stored explicitly in the blockchain but must be calculated by summing up all unspent outputs associated with that address. Additionally, a Bitcoin value associated with an address in an output cannot be spent partly and the sum of inputs must be equal to the sum of outputs in each transaction. It is, however, possible to transfer input values exceeding the outputs ("*change*") back to the same address or to another address owned by the same real-world actor.

The goal of this paper is to present initial results of a longitudinal study conducted over the Bitcoin address graph as a Data Science use case. Our contributions can be summarized as follows:

- We provide a comprehensive graph representation of all Bitcoin addresses and transactions from the beginning of its existence (2009-01-03) until the time of this writing (2016-10-31).
- We conduct a structural analysis of the address graph investigating the change in the structure of the graph over time.
- We investigate the fraction Bitcoin addresses that can explicitly or implicitly be assigned to real-world actors and show how it changes over time.
- We examine the transaction behavior of users considering exchange rates between virtual and fiat currencies.
- We analyze the function of virtual currencies from a user perspective by analyzing the activity periods of addresses and address clusters.

A potential practical use of presented methods and results lies in the implementation in real-time virtual currency analytics platforms, which could provide insight into the current state and evolution of virtual currency ecosystems, such as Bitcoin.

The remainder of this paper is organized as follows: Section II provides an overview of related research in the field of virtual currency analytics; in Section III we briefly introduce the core concepts of the Bitcoin system. Section IV introduces the dataset we used and the analyses we conducted, as well as the initial results of our investigations.

II. RELATED WORK

A strong focus of previous research is on the anonymity property of Bitcoin and possible strategies for de-anonymizing addresses. This is strongly motivated by the strong association between virtual currencies and cybercrime (e.g., ransomware, DDoS attacks). Currently, for instance, it is being discussed, whether authorities should force "*wallet providers […] to apply customer due diligence controls, ending the anonymity associated with such exchanges*."[1]

Ron and Shamir [2], for instance, analyze the typical behavior of Bitcoin users and how they act to obfuscate the flow of Bitcoins to remain in anonymity. Meiklejohn et al. [3] try to reveal real-world identities of users by heuristic clustering and re-identification attacks. Linking the IP address of transactions to user pseudonyms even when they are behind NATs or firewalls is described by Biryukov et al. [4]. De-anonymization by identifying users based on their behavior has been shown by Monaco [5], who found that the transaction behavior of users is nonrandom and nonlinear and that users follow the same behavioral patterns in the long run. Web scraping and transaction fingerprinting are applied by Fleder et al. [6] to reveal the identity of real-world actors.

Other previous work investigates the properties and behaviors of real-world actors in the Bitcoin ecosystem. Möser et al. [7] analyze the reliability of *mixing services*, which can be used to camouflage transactions by breaking the connection between a Bitcoin address sending coins and the address(s) they are sent to. They conclude that there are quality differences in existing services. An alleged theft of Bitcoins from a Bitcoin *exchange* is analyzed by Reid and Harrigan [8].

Graph representations extracted from the Bitcoin blockchain have also been studied before: Holtz et al. [9] focused on specific events and investigated the properties of the Bitcoin graph around the announcement of a Bitcoin gaming site by splitting the graph into small parts and comparing certain properties before and after the launch of the gaming site. A deeper insight into the Bitcoin topology, the broadcast method, and the role of influential nodes taking advantage over other nodes is described by Miller et al. [10]. A commonly used address clustering heuristic by allocating all input addresses of a transaction to the same real-world actor, its (dis-) advantages and effectiveness has been addressed by Harrigan and Fretter [11]. Another study by Kondor et al. [12] analyzes the structure of the transaction network and the evolution of wealth in this network. Further analysis of the bitcoin transaction network is conducted by Ober et al. [13] focusing on global properties of the Bitcoin graph with the result that several parameters remained steady over the last 1.5 years. However, a systematic analysis of the Bitcoin graph and its evolution over time has not yet attracted great attention.

Virtual currency analytics tools implementing some of the above methods have been provided by Haslhofer et al. [14] as well as by Spagnuolo et al. [15].

III. BITCOIN

A. Basic Entities

The Bitcoin system and its working principle with *addresses*, *wallets* and *transactions* has first been described by Nakamoto in 2008 [1]. The basic structure and the relations of its elements is illustrated in Figure 1, which shows the relationship between *addresses*, *wallets*, *transactions*, and (real-world) *users* in the Bitcoin ecosystem. Each user U can have zero to multiple addresses A. In this example, we have two users U_1 and U_2, who hold two addresses in their respective wallets: U_1 holds addresses A_1 and A_2 and user U_2 addresses A_3 and A_4. Each (non-coinbase) transaction T then links at least two, but typically three, and up to an arbitrary number of addresses, in which the value of each transaction must be at least 1 Satoshi, which is 10^{-8} Bitcoins (BTC).

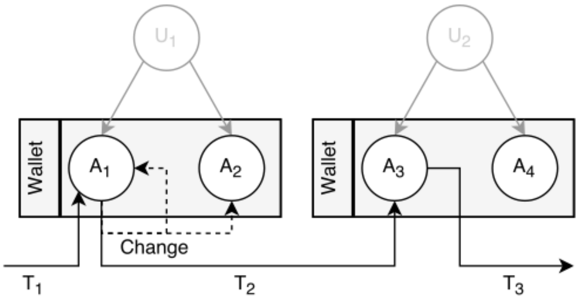

Fig. 1. Schematic view of the Bitcoin system

In Figure 1 let us assume that U_1 received 1 BTC from another user in transaction T_1 and wants to pay U_2 0.75 BTC due to a contractual obligation. As mentioned earlier, U_1 cannot transfer 0.75 BTC to U_2 directly. Instead, U_1 must spend the entire amount of 1 BTC associated with A_1 in T_2, specifying that 0.75 BTC goes to A_3 held by U_2. The remaining 0.25 BTC are change and must be transferred to another address, which can be A_1 or a newly generated address A_2 held by U_1. It is an assumption that the addresses A_1 and A_2 belong to user U_1. The only fact which can be taken from the blockchain is that there were 0.75 BTC sent from address A_1 to address A_3 in transaction T_2 and 0.25 BTC from A_1 to A_1 or A_2 as *change*.

Newly generated transactions are broadcasted to the Bitcoin peer-to-peer network and collected by special-purpose nodes, the so-called *miners*, who try to combine them into a new *block*, which is issued approximately every 10 minutes and added to the blockchain with a timestamp resulting in a monotonously growing, temporarily ordered transaction sequence. Mining is a competitive and highly resource-intensive task (*proof-of-work*) and follows a pre-defined consensus protocol; both are beyond the scope of this paper.

B. Constructing the Bitcoin Address Graph

The Bitcoin address graph can be constructed by extracting all transactions from the blockchain and creating a *property graph*, in which each node represents an address and each edge a transaction that has taken place between a source and a target address.

Each node (address) and edge (transaction) can carry additional descriptive properties: typical properties for

[1] http://www.consilium.europa.eu/en/press/press-releases/2016/12/20-money-laundering-and-terrorist-financing/

addresses are *tags* providing additional contextual information about an address. Such tags might be collected by crawling the Web. Possible properties for edges are the number of transactions or the flow of Bitcoins between two addresses.

In order to quantify the flow of Bitcoins between two addresses, it must be taken into account that unlike in the real-world banking system, a Bitcoin transaction represents an m:n relationship between addresses; thus, a transaction can have multiple input and multiple output addresses, as illustrated in Figure 2: we assume that there is a transaction *T* which has addresses A_1 with a value of 2 BTC and A_2 with a value of 5 BTC as an input. The outputs of *T* are addresses A_3 which receives 3 BTC and A_4 which gets 4 BTC. In Bitcoin, it is impossible to assign a specific value of an input to a specific output address. Even though A_3 receives 3 BTC, the source of the 3 BTC cannot be determined. Therefore, the flow of Bitcoins can only be estimated based on the values of the inputs and outputs as shown in Table I.

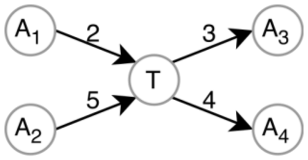

Fig. 2. Bitcoin transaction value assignment

Therefore, we estimate the flow of actual Bitcoins between two addresses using the following formula:

$$Est(I_i, O_j) = O_j * \frac{I_i}{\sum_k I_k}$$

TABLE I. BITCOINFLOW

Transaction	Formula	Estimated BTC
A1 → A3	3 * (2/7)	0.857
A2 → A3	3 * (5/7)	2.143
A1 → A3	4 * (2/7)	1.143
A2 →A4	4 * (5/7)	2.857

IV. ANALYSIS

A. Dataset

For our analysis, we took a dump of the Bitcoin blockchain which includes all transactions from the first block on 3rd of January 2009 until block 430.000 on 15th of September 2016. The analysis is carried out month-wise and considers transaction data until 31st of August 2016. Table II shows the number of addresses, blocks, and transactions included in this dataset.

TABLE II. DATASET STATISTICS

Total number of addresses	176.412.948
Total number of blocks	430.000
Total number of transactions	156.365.848

Bitcoin users are encouraged to use each address only once, which means that ideally each Bitcoin address is involved in at most two - one receiving and one spending - transactions. The difference in the total number of addresses and number of transactions can be explained by *(i)* not all users following this recommendation and *(ii)* addresses that serve as an input for multiple transactions. Addresses are often reused by vendors and organizations that receive Bitcoin donations and refrain from anonymity on purpose by advertising their address publicly on the Web.

The numbers in Table II are also an indicator for the adoption of virtual currencies such as Bitcoin (BTC). While the number of organizations (e.g., Internet Archive[2]) and vendors who accept Bitcoin is still relatively low, the overall transaction volume is steadily increasing. An increase in reuse of addresses over time could in fact indicate a wider adoption by common vendors. An overview of vendors accepting Bitcoin for payment is available online[3] and amounts to around 8,400 worldwide (but also including Bitcoin ATM).

B. Structural analysis

Given the growing number of transactions, we can expect that the structure of the Bitcoin address graph changes over time due to a growing number of participating users and organizations accepting Bitcoin for payment or donations.

For our Bitcoin address graph analysis and for our definition of in- and out-degree of single addresses, we represent addresses as vertices (nodes) and create a labelled directed edge for each transaction *T* that involved two addresses A_i and A_o as in- and output; that is, there may be multiple edges labelled with the same transaction *T*, in case multiple in- and outputs are involved. We associate each transaction *T* with several additional attributes, e.g. the transaction id, in which block the transaction occurred and the minimum, average and maximum transaction value which have been transferred between these addresses.

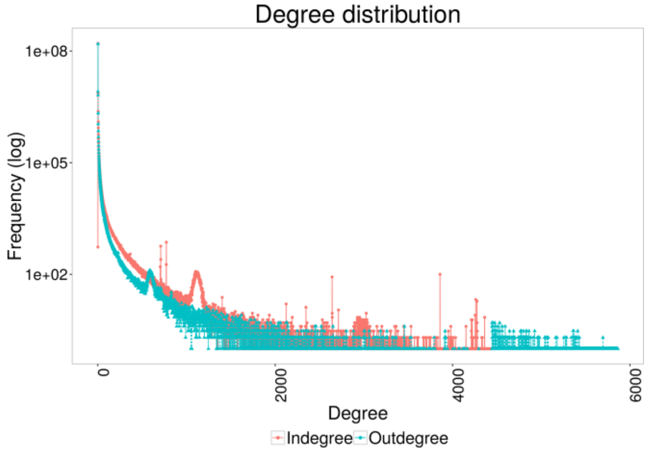

Fig. 3. Degree distribution

The degree of a vertex is the defined as the number of incoming and outgoing edges. Based on the design and

[2] https://archive.org/donate/bitcoin.php (28.10.2016)
[3] http://www.coinmap.org

anonymity of the Bitcoin system, it is expected that most addresses have a low degree. However, since a single transaction can contain an arbitrary number of input and output addresses, the in-degree and out-degree of address nodes varies, as shown in Figure 3.

Especially high- and low-degree address nodes are of interest. Prominent examples for donation addresses are the Internet Archive[4] with an indegree of 1,759 and an outdegree of 105 or WikiLeaks[5] (24,469/125). An address uniting both having the highest indegree (1,595,498) and outdegree (1,600,277) belongs to the online Bitcoin casino satoshiDICE[6] (*1dice8EMZmqKvrGE4Qc9bUFf9PX3xaYDp).*

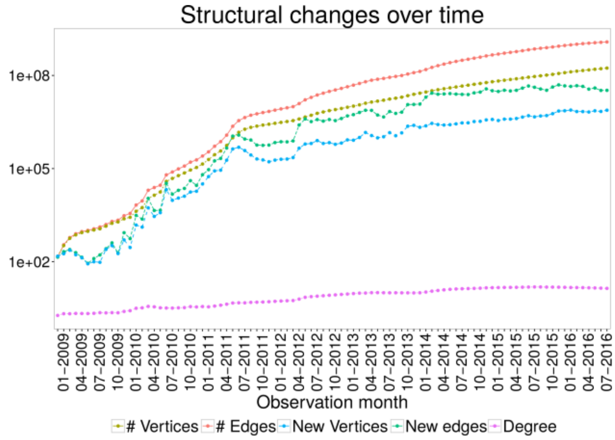

Fig. 4. Structural changes over time

Next, we analyze the evolution of Bitcoin addresses over time. Figure 4 shows the cumulative number of distinct nodes (addresses) and edges in the address graph as well as the number of added nodes and edges and average node degree per month. We can observe that the number of used addresses is increasing which can be interpreted in two ways: on one hand, it shows that the users use new addresses for each transaction, but this does not mean that each new address also leads to a new user. A single user can have lots of transactions within an observation period with many addresses. On the other hand, an increased usage of the virtual currency Bitcoin can be derived. The number of edges (transactions) is increasing in the same way. The degree remains almost steady over the course of time.

C. Real-world actors in the Bitcoin ecosystem

Transactions are anonymous by design and do not reveal the identities of real-world actors who can be individuals, exchanges, payment providers, or any other type of service in the Bitcoin ecosystem. However, as shown by previous research [3,8], it is possible to combine addresses into clusters (wallets), which are likely to be controlled by a certain real-world actor.

A well-known clustering heuristics works under the assumption that multiple addresses, which are used as input of a single transaction must be controlled by the same real-world actor. This assumption holds if transactions are not executed through mixing services, which obfuscate the transaction by breaking the connection between a Bitcoin address sending coins and the addresse(s) they are sent to.

Other heuristics are based on the observation that *change*, is transferred back to the user when the sum of inputs is greater than the sum of outputs. Thus, one of the output addresses of a single transaction T often belongs to the same user or real-world actor; in a typical transaction, this address is even identical to one of the input addresses.

As soon as an address cluster (so called *entities*) is identified, a single address within that cluster carries an explicit tag with contextually relevant information, it is possible to implicitly assign that tag to all other addresses in the cluster and to possibly identify the real-world actor owing that address cluster, which often corresponds to a *wallet*. Therefore, we can group addresses into three different categories:

Unknown addresses: no tag has been assigned to an address and no contextual information is available publicly. From the point of view of the Bitcoin design, this is the desired ideal situation in terms of anonymity and address usage. Unknown addresses have not been used as input with other known addresses.

Explicitly known addresses: additional contextual information can be assigned in the form of a tag. Such tags can be extracted by crawling the Web or gathering data from external information sources such as blockchain.info[7], walletexplorer[8], social media platforms or the Darknet.

Implicitly known addresses: appear in a cluster with at least one other explicitly known address, from which tags can implicitly be derived. In the case shown in Figure 2, it is assumed that the addresses A_1 and A_2 are controlled by the same user and in addition, these addresses appear in a cluster with explicitly known addresses.

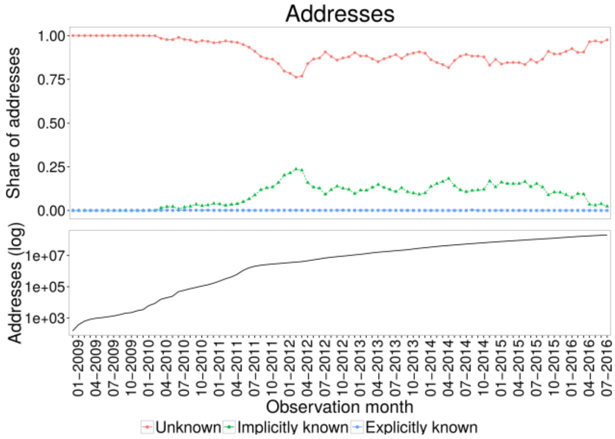

Fig. 5. Share of addresses by category.

Figure 5 shows that the fraction of explicitly known addresses is low throughout the entire Bitcoin's history. However, the fraction of implicitly known addresses starts

[4] https://archive.org/donate/bitcoin.php
[5] https://shop.wikileaks.org/donate
[6] https://www.satoshidice.com/
[7] http://www.blockchain.info
[8] http://www.walletexplorer.com

growing in 2010, reaches its maximum in 2012, remains roughly constant until 2015 and starts decreasing in 2016.

We assume that the decrease towards June 2016 is caused (i) by missing contextual information (tags) for newly generated addresses, (ii) the increasing awareness of end users that reuse of Bitcoin addresses decreases anonymity, and (iii) the increasing usage of Bitcoin mixing and tumbler services.

In general, the need for anonymity depends on the user group and the purpose for which Bitcoin is used. Organizations financing themselves with donations are well advised to publish their Bitcoin address on their homepage or social media to collect donations. They even generate so-called vanity addresses for that purpose, which are personalized addresses which often contain the organizations name in it (e.g. the addresses of organizations or gaming sites). On the other side of the spectrum are organizations conducting illegitimate business transactions such as collecting ransom from cybercrime activities.

D. Transaction behavior and exchange rates

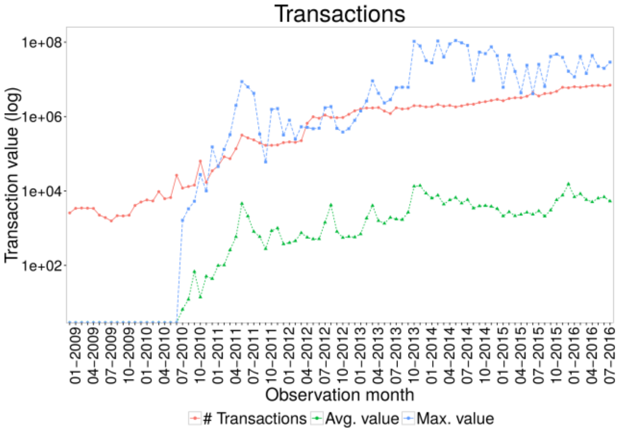

Fig. 6. Number and value of transactions (in USD)

Figure 6 shows the total number of transactions carried out in each month as well as the average and the maximum value of the transactions in USD with the exchange rate at the time of the transaction. The sudden increase of the average and maximum value of the transactions is caused by missing exchange rate data before 2010.

The number of transactions is continuously growing for the first four years and remains steady since then. An explanation could be that a virtual currency like Bitcoin was new and more and more people tried it out. After some years, the interest in Bitcoin seems to have flattened, but the market remains constant and the regular transactions remain. The peaks on the average transaction value in USD go hand in hand with the changes of the exchange rate. The fluctuating maximum real value can also be explained by the changes in the exchange rate, which we will explore in the following subsection.

The average exchange rate BTC/USD over time from January 2009 until December 2016 is shown in Figure 7 together with the minimum and the maximum per month. It is clearly visible, that the exchange rate remains steady for the first two years before it shows volatile behavior afterwards. However, three months are particularly striking: April 2013, December 2013 and May 2016, where the difference between the minimum and the maximum prices for BTC is very high, whereas they are in a certain range for the rest of the months.

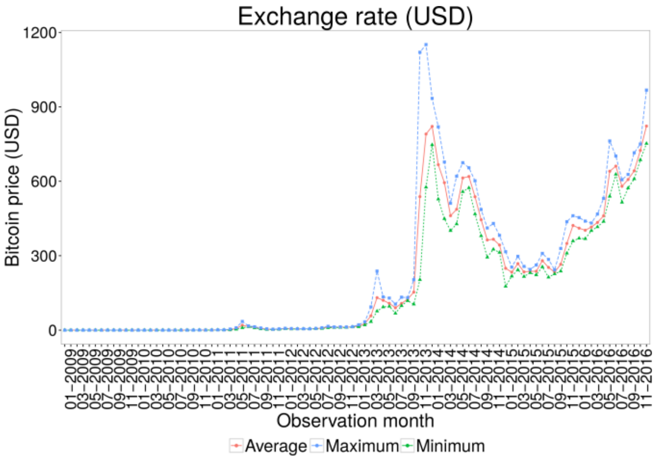

Fig. 7. Exchange rate BTC/USD

Real-world events influence the popularity of virtual currencies. The first peak in the exchange rate is visible in April 2013 shortly after the Cypriote government announced the bailout of Cyprus' banks causing many people to safe their money by switching to Bitcoin[9]. Additionally, the seizure of Mt. Gox, a large Bitcoin exchange, had an impact on the exchange rate[10,11]. The highest price of around 1,250 USD has been reach in late 2013, which is also the last noticeable rise in the number of transactions. The increase in value also caused an increase in the number of transactions. Investors strive to increase their wealth and a strongly increasing exchange rate of a currency is very tempting. The last clearly striking changes in the exchange rate occurred in July 2016. The military coup in Turkey caused a run to the Turkish banks as people wanted to know their money in a safe place[12]. Events in politics which are likely to cause fear among the population which might also be reflected in the virtual currency markets. [13]

E. Activity time

Next, we examine the possible monetary functions (store of value, exchange of value) of Bitcoin by investigating how long users keep their Bitcoins and whether this currency is used as an alternative to long-term savings accounts.

TABLE III. BITCOINFLOW

Metric	Address based	Entity based
Avg. used in transactions	Incoming: 2.25 Outgoing: 1.75	Incoming: 10.5 Outgoing: 3.7
Avg. activity time (days)	12	15
Median activity time (days)	< 1	< 1

[9] http://money.cnn.com/2013/03/28/investing/bitcoin-cyprus/
[10] http://bit.ly/2nrTLGm
[11] http://bit.ly/2mN6J2O
[12] https://news.bitcoin.com/turkish-bitcoin-military-coup/
[13] http://bit.ly/2nmUdbU

Table III shows the activity time of addresses and entities, which is defined as the period between the first and last transaction a single address or an address within a cluster has been involved in. The average activity time of an entity with all accounted addresses is 15 days (median: < 1 day), which strengthens the hypothesis that Bitcoin is not used as a replacement for saving accounts but rather as global payment system. This is in line with our previous observation in Figure 3, showing that the degree for most addresses is rather small.

The number of entities with a positive balance on at least one of their associated addresses amounts to 368,739 with total unspent 1,565,294 BTC on the 15th September 2016. The mean balance is 4.25 BTC (median: 0.000795 BTC).

V. Conclusion

In this paper, we analyzed the Bitcoin address graph from several perspectives based on the publicly available ledger containing all transactions from the beginning of Bitcoin in January 2009 until 31st of August 2016.

We described the procedure we applied for constructing the address graph, presented a possible strategy for estimating currency flows between addresses and described the heuristics we applied for combining addresses in clusters of addresses, which can then be assigned to real-world actors.

Our structural analysis has shown a highly-skewed degree distribution, which implies that the Bitcoin address graph comprises a small number of outliers with high in- and/or outdegree. Manual inspection of those addresses revealed that those addresses are often used by (non-profit) organizations for receiving donations or by online gambling Websites. It has also shown that the address graph is expanding rapidly over time as new addresses and transactions are added to the blockchain. The average node degree, however, remains stable over time.

Investigation of real-world actors has shown that address clustering using well-known heuristics can increase the number of implicitly known addresses in the entire graph. However, most Bitcoin addresses remains anonymous.

Furthermore, our analysis has illustrated the growing transaction volume and the effects of real-world events on the Bitcoin exchange rate. It has also shown that real-world actors use Bitcoin more for transferring value than for storing value. This indicates that Bitcoin is not used as an alternative to savings accounts, probably due to the above-mentioned volatility and instability of the currency.

A clear limitation of our work lies in the selection of tags, which affects the fraction of implicitly and explicitly known addresses. We expected that a more comprehensive tag dataset extracted from various sources would increase the fraction of identifiable addresses but not de-anonymize most addresses.

A possible direction for future work lies in the investigation of effects of external political or economic events (e.g., “Brexit”) on virtual currencies and the prediction of possible micro- and macroscopic reactions within or across virtual currency ecosystems. Furthermore, it would be interesting to extend the analytics methods presented in this paper to other crypto-currencies such as Monero or ZCash.

The dataset for our analysis is a cleansed graph representation of the blockchain and is available to other researchers on request.

VI. Acknowledgement

This work was funded by the Austrian research funding association (FFG) under the scope of the ICT of the Future program (contract # 849906).

References

[1] S. Nakamoto, “Bitcoin: A peer-to-peer electronic cash system.” 2008.

[2] D. Ron and A. Shamir, “Quantitative analysis of the full bitcoin transaction graph,” in *International Conference on Financial Cryptography and Data Security*. Springer, 2013, pp. 6–24.

[3] S. Meiklejohn, M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. M. Voelker, and S. Savage, “A fistful of bitcoins: characterizing payments among men with no names,” in *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 2013, pp. 127–140.

[4] A. Biryukov, D. Khovratovich, and I. Pustogarov, “Deanonymisation of clients in bitcoin p2p network,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 15–29.

[5] J. V. Monaco, “Identifying bitcoin users by transaction behavior,” in *SPIE Defense + Security*. International Society for Optics and Photonics, 2015, pp. 945 704–945 704.

[6] M. Fleder, M. Kester and S. Pillai, “Bitcoin Transaction Graph Analysis”, CoRR, vol. abs/1502.01657, 2015. [Online] Available: http://arxiv.org/abs/1502.01657.

[7] M. Möser, R. Böhme, and D. Breuker, “An inquiry into money laun- ̈ dering tools in the bitcoin ecosystem,” in *eCrime Researchers Summit (eCRS)*, 2013. IEEE, 2013, pp. 1–14.

[8] F. Reid and M. Harrigan, “An analysis of anonymity in the bitcoin system,” in *Security and privacy in social networks.* Springer, 2013, pp. 197–223.

[9] B. Holtz, J. Fortuna and J. Neff, “Evolutionary structural analysis of the bitcoin network”, 2013.

[10] A. Miller, J. Litton, A. Pachulski, N. Gupta, D. Levin, N. Spring and B. Bhattacharjee, “Discovering Bitcoin’s Public Topology and Influential Nodes”, 2015.

[11] M. Harrigan and C. Fretter, “The unreasonable effectiveness of address clustering”, CoRR, vol. abs/1605.06369, 2016. [Online] Available: http://arxiv.org/abs/1605.06369.

[12] D. Kondor, M. Posfai, I. Csabai, and G. Vattay, “Do the rich get richer? An empirical analysis of the bitcoin transaction network,” *PloS one*, vol. 9, no. 2, p. e86197, 2014.

[13] M. Ober, S. Katzenbeisser, and K. Hamacher, “Structure and anonymity of the bitcoin transaction graph,” *Future internet*, vol. 5, no. 2, pp. 237–250, 2013.

[14] B. Haslhofer, R. Karl, and E. Filtz, “O bitcoin where art thou? insight into large-scale transaction graphs,” in Joint Proceedings of the Posters and Demos Track of the 12th International Conference on Semantic Systems - SEMANTiCS2016 and the 1st International Workshop on Semantic Change & Evolving Semantics (SuCCESS’16) co-located with the 12th International Conference on Semantic Systems (SEMANTiCS 2016), Leipzig, Germany, September 12-15, 2016., ser. CEUR Workshop Proceedings, M. Martin, M. Cuquet, and E. Folmer, Eds., vol. 1695. CEUR-WS.org, 2016. [Online]. Available: http://ceur-ws.org/Vol-1695/paper20.pdf

[15] M. Spagnuolo, F. Maggi, and S. Zanero, “Bitiodine: Extracting intelligence from the bitcoin network,” in *International Conference on Financial Cryptography and Data Security*. Springer, 2014, pp. 457–468.