Title
Understanding the Spatial, Platial, and Temporal Properties of Cryptocurrency Ecosystems

Authors
Janowicz, Krzysztof
Regalia, Blake
Zhu, Rui
et al.

Publication Date
2022-06-22

Understanding the Spatial, Platial, and Temporal Properties of Cryptocurrency Ecosystems*

Krzysztof Janowicz
UC, Santa Barbara

Blake Regalia
UC, Santa Barbara

Rui Zhu
UC, Santa Barbara

Bo Yan
UC, Santa Barbara

Gengchen Mai
UC, Santa Barbara

August 29, 2018

Abstract

Cryptocurrencies and their underlying technologies such as blockchains and smart contracts are rapidly gaining traction in sectors such as banking, identity management, supply chain management, cloud-computing, voting, forecasting, and so forth. With this change in visibility and first signs of mainstream adoption, there is a growing interest in understanding the cryptocurrency ecosystem, e.g., regarding market trends or inherent risks. Interestingly, however, spatial and platial aspects have not yet received much attention. One possible reason for this lack of analysis may be due to the perception of cryptocurrencies being global and living outside of legal frameworks. We will show that this is a misconception and that understanding the cryptocurrency ecosystem requires looking at the spaces and places involved in their creation, consumption, and regulation.

1 Introduction

A digital distributed ledger is essentially a database of shared, replicated, and synchronized records that are collaboratively managed by peers without the need for central administration. Typically, the ledger is open in the sense

*This is a working draft of a paper to be submitted for review and made available here for discussion; use on your own risk as content may change.**

that anybody can contribute, e.g., by having their transactions included, as well as in the sense that anyone can act as a node, e.g., to validate the state of the ledger. The contributing nodes are usually dispersed across the globe, spanning countries as well as institutions, and it is often said that the networks created by these nodes lack geographic boundaries as they operate the same way throughout the world regardless of local jurisdiction.

The technologies underlying distributed ledgers including blockchains and smart contracts have the potential to transform many sectors including banking, identity management, supply chain management, cloud-computing and storage, voting, and forecasting, and in many cases, they are already being tested or deployed in these sectors. Coins and tokens – often summarized as *cryptocurrencies* – constitute a significant building block of distributed ledgers and serve a variety of functions including providing the incentive for mining blocks, i.e., for processing transactions, bundling them into blocks of a given size, and linking these blocks to the chain of the previous transaction blocks. Hence, with distributed ledger technologies rapidly gaining traction and cryptocoins increasing (or decreasing) in value by often hundreds of percent within months, understanding the crypto-economy and the ecosystems it forms becomes a pressing issue.

Interestingly, most work so far has focused on either studying the underlying technologies from a privacy perspective such as by focusing on methods to (de)anonymize users [3], by studying the properties of the peer-to-peer networks of contributing nodes, the utilized cryptographic technologies, different economic measures such as proof-of-work against various kinds of attacks, or from a trading perspective by applying well-known technical chart indicators used in stock market analysis. To the best of our knowledge, spatial, platial (i.e., place-based), and temporal aspects have been largely ignored or only received attention in news articles discussing power consumption and the mining industry more broadly. One of the few exception is the work by Zyskind et al. [3] on estimating the location of users.

We believe that the impression of a *borderless*, global crytpo-ecosystem is naive at best and potentially misguided. In the following, we will outline how space, place, and time impact the entire ecosystem in various ways and discuss why we believe that studying the crypto-economy will yield valuable new insights and pose interesting new research questions to the GIScience community. To give an intuitive example, proof-of-work based systems such as Bitcoin (BTC)[1] are greatly affected by energy costs which vary geograph-

[1] While we will use Bitcoin as main example, our results generalize to many other coins as well.

ically and are directly linked to factors from physical geography. The wide availability of geothermal energy on Iceland, for instance, puts the otherwise small country on the forefront of the rapidly growing mining industry. Human factors play an equally important role to physical aspects as miners have to operate within the legal frameworks of their countries. For example, China has recently taken a more conservative stance and is thereby driving out substantial parts of the mining industry with direct consequences for neighboring states. Simply put, what is often overlooked is that even abstract digital goods have a physical and cultural grounding [1]. In the following sections we will present a series of small experiments to substantiate our claims.

2 Variation of Profitability

The profit of mining proof-of-work based cryptocurrencies depends on at least three aspects: the value of the coin, the difficulty level and hash rate, i.e., the amount of work required to mine a block, and the costs associated with performing this work, i.e., electricity, required hardware, and so on. As we will discuss later, the value of a cryptocoin such as Bitcoin varies geographically. However, we will treat it as invariant for our first experiment and highlight the cost of electricity instead. Electricity costs vary between and within countries, even by season, thereby having substantial consequences for mining profitability. This alone is a trivial insight, but what makes it interesting from a spatial analysis perspective is the fact that electricity is paid in local currency while most cryptocoins are valued against the US Dollar (USD) either directly, via Bitcoin (and its relation to USD), or through a so-called stable coin (such as USDT) that has a fixed 1:1 rate against USD. Consequently, as the cryptocurrency ecosystem is built on the promise of reduced financial frictions, e.g., wrt transaction costs and market access, geographic regions and countries with access to cheap energy and/or currencies that are very weak, as compared to USD, provide a strong competitive advantage for large-scale mining farms. To give an extreme example, the potential cost of mining one Bitcoin in February of 2018 in South Korea was estimated at above $26k, while it was just $513 in Venezuela, given that Bitcoin fluctuated between $6.6k and $11k during this period, it is easy to see why Venezuela is currently experiencing a rapid rise in mining activity and an increasingly favorable political/legal climate.

In our experiment, we use the ASICminer 8 Nano Pro as an example of state-of-the-art mining hardware. The system was released in May 2018

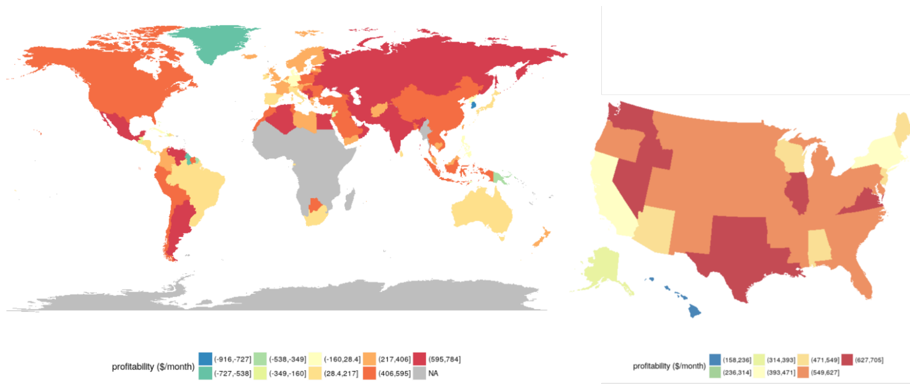

Figure 1: Est. monthly profit worldwide and for the US.

and its profitability will diminish quickly with an increasing difficulty level and with potentially declining bitcoin prices, while increasing prices will provide earlier return of investment. Hence, miners have to amortize the cost of the hardware early on and then make enough profit to afford new hardware within the months to come. Put differently, successful Bitcoin miners[2] have to be good at estimating risk along the following dimensions: costs of electricity, Bitcoin prices over time despite large fluctuations, and relative political stability with respect to the acceptance of mining and cryptocurrencies. We will show that all of these factors show strong place-based and space-based variation.

Figure 1 shows the estimated profitability of Bitcoin mining in USD per month in June 2018 using the aforementioned hardware. US electricity costs per state and their monthly variation are well-documented, while global data is less reliable due to a lack of up-to-date sources and exchange rates. As one can clearly see, variance across countries and even within the US is substantial. US-based miners will not be able to tolerate decreasing prices for much longer, while many countries in Asia and South America will remain profitable even in the event of very large fluctuations. Electricity prices are determined by physical factors, thus making this a spatial problem, as well as political and economical factors such as taxes, exchange rates, and so forth, thus showing a platial dimension as well. The spatial aspect will translate into strong autocorrelation while the platial aspect would not necessarily. Computing Moran's I using common boundaries confirms this balancing effect and yields 0.25 for the US and 0.68 globally. If the political climate

[2] As well as miners of many other coins with a similar setup such as Ethereum, which will switch to a proof-of-stake model in the future.

within one administrative unit changes, miners can move (and have moved) to neighboring units. Finally, it is worth mentioning that large-scale miners often have direct access to power plants or even reactivated hydro power stations across Europe, effectively gaining independence from energy prices.

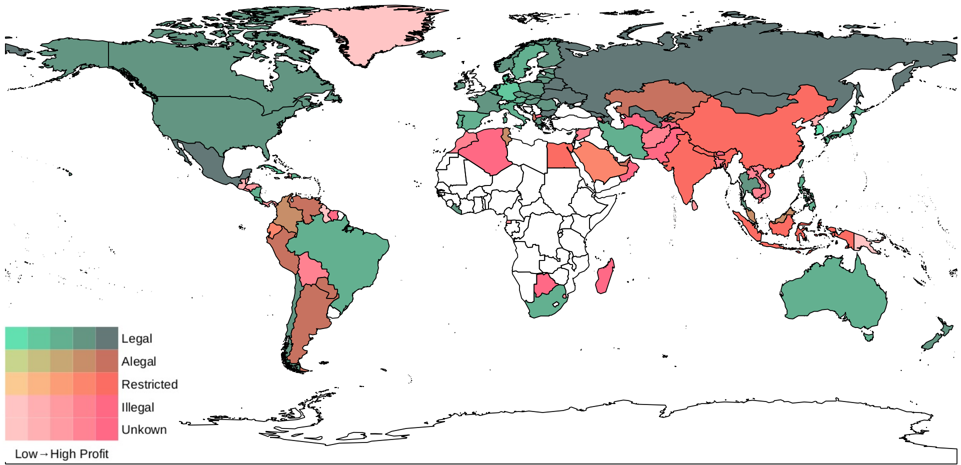

Figure 2: Bivariate map: political climate and mining profit.

Summing up, the bivariate map in Figure 2 illustrates the interplay of spatial and platial effects by showing the combinations of political climate and electricity costs. [3]

3 Exchanges and Local Prices

Taking on a consumer-centric view, we conduct two experiments on the exchange markets between different fiat currencies and Bitcoin. We analyze the exchange volumes using Jensen-Shannon divergence to understand the general trend and the similarity between fiats and Bitcoin. In order to gain more insight about the fine-grained patterns, we use dynamic time warping to analyze the sequential patterns of these exchanges. Finally, we discuss local Bitcoin prices as a function of a place's ease of access to, and utility of, the cryptocurrency.

[3]Data on the political climate classification was taken from `https://coin.dance/poli/legality`.

3.1 Exchange Volume Analysis

First, we analyze the general trend of exchange volumes over a one year period (2017). The exchange data was collected from 79 online Bitcoin exchanges, such as GDAX, Bithumb, Coinone, and so on. In order to get an holistic view of exchange volumes, we selected 8 fiat currencies that are actively traded against Bitcoin (BTC): Chinese Yuan Renminbi (CNY), European Euro (EUR), Indian Rupee (INR), Japanese Yen (JPY), South Korean Won (KRW), Russian Ruble (RUB), Singapore Dollar (SGD), and USD.

Hourly volume data from each exchange is then aggregated by fiat currency. To better understand exchange volume trends, we compute Jensen-Shannon divergence to measure their similarity (see Figure 3), where a smaller Jensen-Shannon divergence value signifies a higher similarity. As can be seen by the top row of the matrix, the CNY market tends to exhibit similarity with other fiat markets. The exchange volume patterns for INR, RUB, and SGD are more similar to each other than to the other 5 markets. USD and JPY show the most distinct patterns. This aligns with our expectations since USD is the largest fiat currency traded against cryptocurrencies (and the most common stable coin, USDT, is coupled to the USD), while Japan has a very advanced crypto-ecosystem such as ATM machines, stores supporting Bitcoin, and a favorable political situation. Another implication of the results is that sudden changes in volume for some fiat pairs correlate strongly, while others do not. For instance, before becoming illegal, China had a very active and turbulent market for so-called Initial Coin Offerings (ICOs) to such a degree that international traders were trying to anticipate the effect of news and technical analyses in order to buy coins during China's early morning hours.

3.2 Exchange Time Series Analysis

We also examine the temporal characteristics of exchanges by collecting time series data for each market pair from 2017. A normalized subset of this time series data is illustrated in Figure 4. As can be seen, the trends for different market pairs have distinguishing features. For example, the trend for BTC-CNY exchange stands out from the others where it peaks in January 2017, it witnesses a dramatic decrease from the middle of the month into February, when its trade volumes virtually dissappear. The likely explanation for this change is due to the regulations issued by the Chinese government around this time. Similarly, trends for INR and SGD to BTC exchanges are significantly different from others as well, probably due to the different economical,

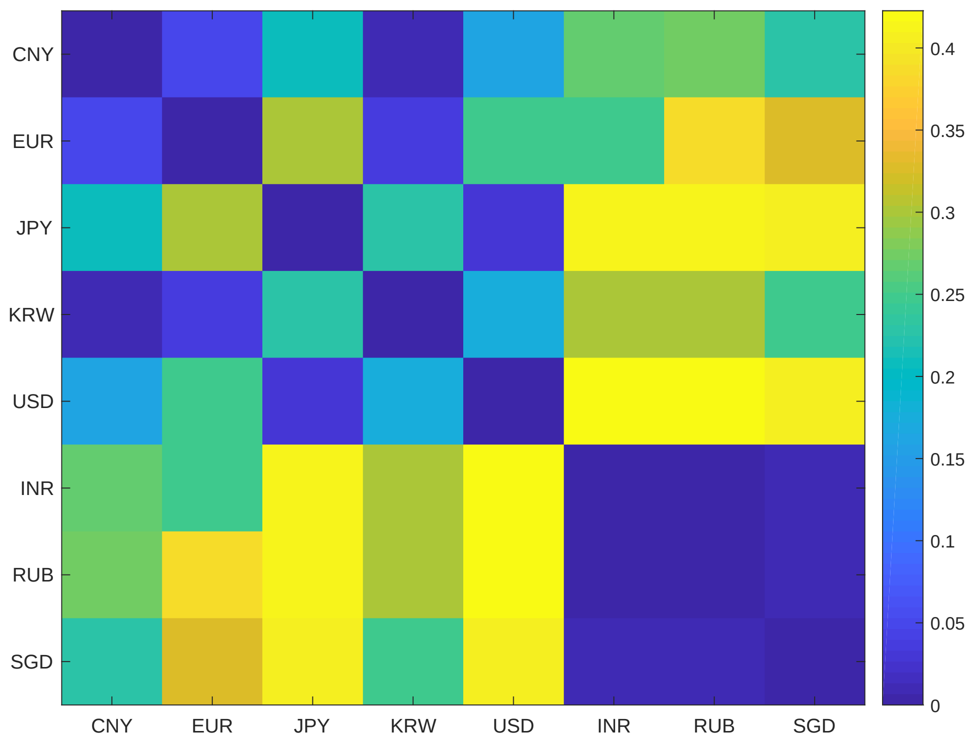

Figure 3: Jensen-Shannon Divergence Matrix

political, and cultural settings.

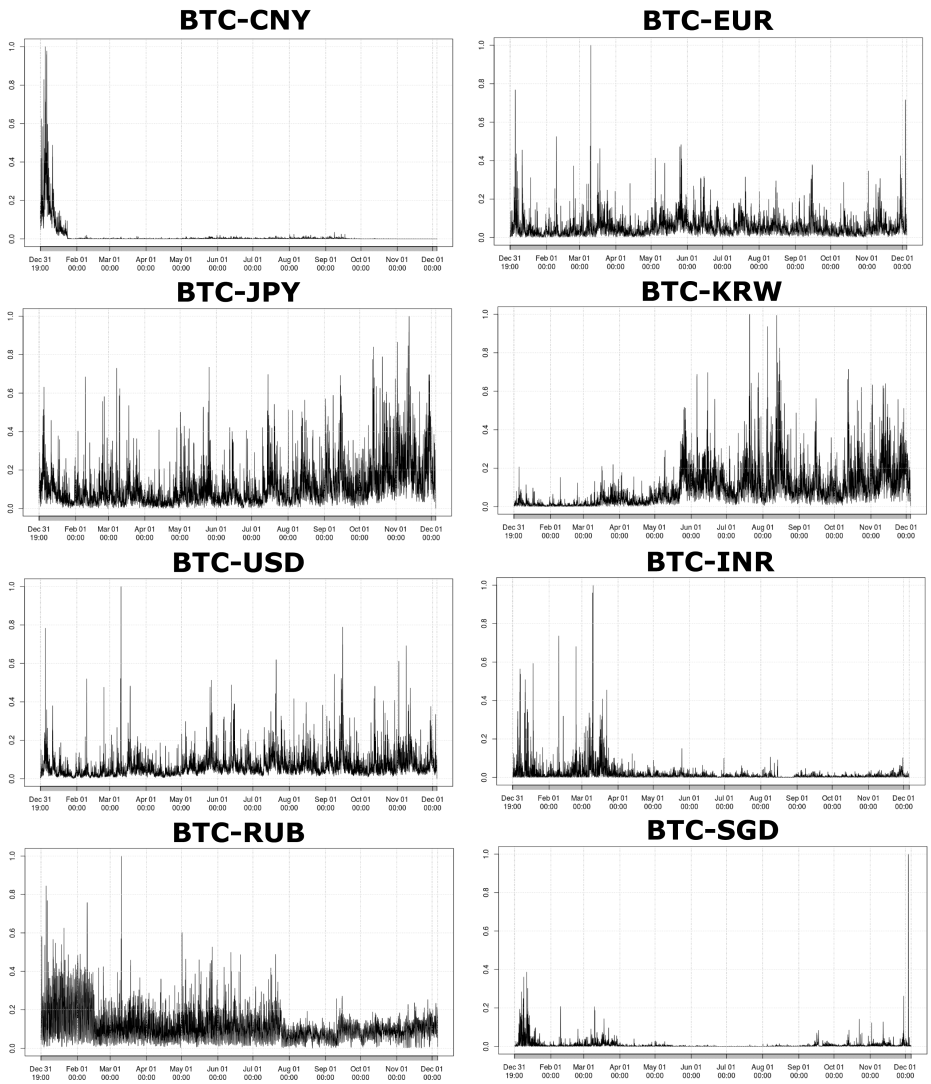

Figure 4: Time series for main exchange markets

Considering that there are temporal lags across the time series of different market pairs, we use dynamic time warping (DTW) to compare their trends (see Figure 5). DTW is a technique to calculate the similarity between two time series by taking into account variation in speed. An important event, e.g., new regulations wrt. one specific exchange market, might not immediately influence another market. Instead, the second market will be affected following some delay. From Figure 5, it is clear that the trends of BTC-CNY exchange markets are relatively similar to the ones of BTC-INR and BTC-SGD, which means that these three markets may have influenced each other more directly. In contrast, the trends of BTC-JPY, BTC-KRW and BTC-RUB are rather distinct from the other exchanges which could imply that they are relatively independent.

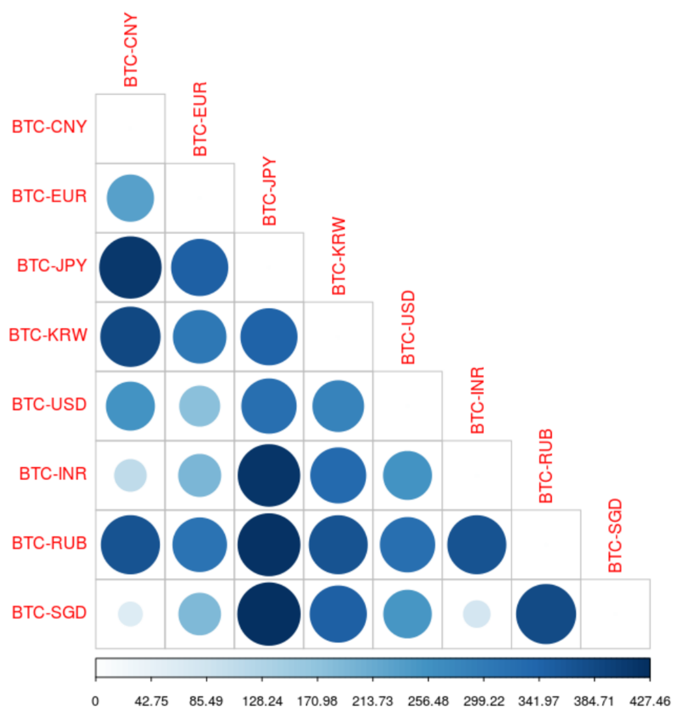

Figure 5: Dissimilarity among different market pairs

3.3 Local Prices

The two previous experiments focus on temporal aspects. In this last example, we want to briefly address a place-based effect as well. In several countries, such as Zimbabwe, Bitcoin and other coins are broadly used to address liquidity gaps or hyperinflation. Nonetheless, local exchanges are not always available in these countries and some of them have imposed restrictions or even bans on crypto-coins in fear of losing control over the financial sector. As major exchanges increasingly require user authentication, many people in such countries have to turn to services that facilitate over-the-counter trading of local currency for Bitcoins. Localbitcoins.com is such a service that enables users to meet in person to exchange coins and currencies in their home city. We believe that it will be a great source for insights into the spatial and platial dynamics of the cryptocurrency ecosystem. Due to the limited space and scope of this paper, we only showcase one example on how place-based, here country-specific, regulations impact pricing. The more difficult and risky trading is, the higher the prices will be. This gives us a unique opportunity to *quantitatively* estimate trading risk and trust in the local currency. E.g., as of 6/11/18 Bitcoin trades at about \$6.6k globally but at \$15.5k in Venezuelan whose VEF is plagued by hyperinflation and at \$9.1k in Afghan Aghani.

4 Network Influence

A blockchain's network hashrate refers to the sum total computational power of all contributing nodes with respect to how many attempts to solve the next block they can make per unit time. *Mining power* is an individual's hashrate (e.g., that of a single node or mining pool) relative to the network hashrate. It is often thought that mining power is the most effective means to influence the evolution of a network since *clients* depend on miners to commit their transactions to the blockchain. However, mining power does not immediately translate to *voting power*, which is rather based on the sheer number of nodes that adopt a certain fork of the blockchain, or switch to running a certain protocol defined by some code branch of the underlying software. While these decisions are often driven by economic consensus, smaller sects within the community can pressure the majority to follow their lead by controlling a large percentage of nodes that disseminate blocks quickest. As we demonstrate in this section, the effect that a node can have on its network is deeply tied to its position within the routing infrastructure, which is largely determined by *geographic location*.

Starting with a snapshot of the Bitcoin nodes[4], we find that 7% of them are Amazon EC2 instances. Even given the cheapest *spot instance pricing*, none of these instances lead to profitable mining operations. Instead, these instances are acting as *full nodes* or *pruning nodes*, verifying the work of miners and propagating the latest blocks, a crucial component of the ecosystem as other miners also depend on these nodes to receive the latest blocks. It's in the best interest of miners to peer with nodes that can consistently broad-

[4] `http://bitnodes.earn.com/`

cast the latest blocks to them as soon as possible, giving them a competitive advantage to mine the next block.

Out of Amazon's 16 availability zones, we positioned a probe at the top 6 zones populated by Bitcoin nodes and include the top 3 here (18% in Virginia, U.S., 16% in Oregon, U.S., and 14% in Singapore) as well as Germany to illustrate the vastly different Bitcoin network connectivity available at each of these locations; see Figure 6. Inspection of these graphs reveals that Virginia is most strongly connected to the core Bitcoin network, an observation that aligns with the distribution of Bitcoin nodes across all Amazon availability zones. In contrast, Oregon and Germany exhibit a spread-out pattern of dense clusters, possibly serving the larger subnetworks at the further reaches of the hub seen from Virginia. However, Singapore exhibits the most unique pattern, showing a highly diffused graph of connectivity. The discrepancy between this zone's observed EC2 node count and its limited connectivity could be explained if one considers the likely proximity to concentrated mining power in Asia, whereby a node in this zone stands to gain favor with the continent's mining giants.

5 Conclusions

While the crypto-ecosystem may resemble regular stock trading and its indicators, this view is equally naive as the idea of a borderless system uncoupled from its physical and cultural grounding. In this paper, we have provided a series of experiments to highlight the spatial, platial, and temporal characteristics of the crypto-ecosystem. We believe that spatial thinking and analysis have an important role to play in understanding this new sector and that novel methods such as spatially-explicit indicators will be required to develop better risk assessment models.

References

[1] Helen Couclelis. Pizza over the internet: e-commerce, the fragmentation of activity and the tyranny of the region. *Entrepreneurship & Regional Development*, 16(1):41–54, 2004.

[2] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *International conference on financial cryptography and data security*, pages 436–454. Springer, 2014.

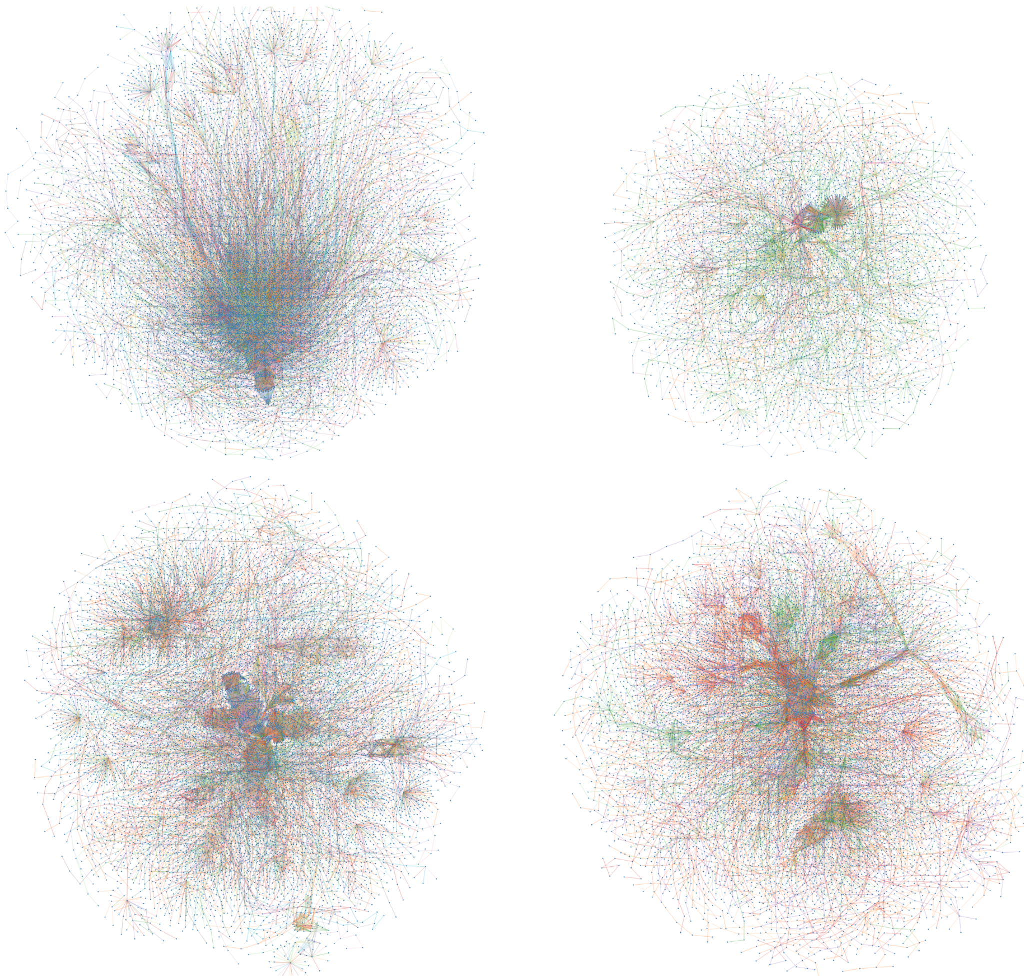

Figure 6: Force-based graph simulations of *traceroutes* to all ipv4 BTC nodes based on packet round-trip-time to intermediate routing nodes from an EC2 node in Virginia, U.S. (tl), Singapore (tr), Oregon, U.S., (bl), and Germany (br).

[3] Guy Zyskind, Oz Nathan, et al. Decentralizing privacy: Using blockchain to protect personal data. In *Security and Privacy Workshops (SPW), 2015 IEEE*, pages 180–184. IEEE, 2015.